# Consumer Price Search and Platform Design in Internet Commerce*

Michael Dinerstein, Liran Einav,
Jonathan Levin, and Neel Sundaresan†

August 2017

**Abstract.** Despite low physical search costs, online consumers still face potentially large search frictions due to the proliferation and high churn of products and sellers. Consequently, the platform design – the process that helps potential buyers navigate toward a product they may purchase – plays a critical role in reducing these search frictions and determining market outcomes. In this paper we study a key trade-off associated with two important roles of efficient platform design – guiding consumers to their most desired product while also strengthening seller incentives to offer low prices. We begin by illustrating this trade-off in a simple theoretical framework, and then combine detailed browsing data from eBay and an equilibrium model of consumer search and price competition to quantitatively assess this trade-off in the particular context of a change in eBay's marketplace design. We find that retail margins are on the order of 10%, and use the model to explore how pricing and purchase rates vary with the platform redesign. Our model explains most of the effects of the redesign, and allows us to identify conditions where narrowing consumer choice sets can be pro-competitive. The counterfactual exercises also point to a very different resolution of the platform design trade-off when products are more heterogeneous, a result that is also qualitatively supported by a subsequent A/B experiment run by eBay.

*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 and many seminar participants for helpful comments. Data access for this study was obtained under a contract between Dinerstein, Einav, and Levin and eBay Research. Neel Sundaresan was an employee of eBay at the time this research began.

†Department of Economics, University of Chicago (Dinerstein), Department of Economics, Stanford University (Einav and Levin), NBER (Dinerstein, Einav, and Levin), and Microsoft (Sundaresan). Email: mdinerstein@uchicago.edu, leinav@stanford.edu, jdlevin@stanford.edu, and nsundare@gmail.com.

# 1 Introduction

Search frictions play an important role in retail markets. They help explain how retailers maintain positive markups even when they compete to sell near-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., 2015). One explanation for remaining search frictions in online markets is that the set of competing products is often very large and changes regularly such that consumers cannot be expected to consider, or even be aware of, all available products.

To deal with this proliferation of options, consumers shopping online can use either price search engines or (more often) compare prices at e-commerce marketplaces, or internet platforms, 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). Sellers on these platforms 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). These often conflicting incentives highlight the important role of the platform design, which structures online search in a way that affects consumer search and seller incentives at the same time. In markets where the set of potential offers is large, the platform's design may have first-order implications for price levels and the volume of trade.

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; Honka et al., 2014). 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, and use a simple theoretical framework to illustrate the associated trade-offs. In particular, we assess the trade-off between guiding consumers to their most desired products and strengthening seller incentives to offer low prices.

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 also present results from a randomized A/B experiment that eBay ran subsequent to the search redesign. The experiment randomized the default search results presented to consumers. The experiment results highlight that the impact of the search redesign varies considerably across product categories that are more homogenous or less so. It also points to the limitation of an A/B experiment in testing equilibrium predictions, which may require longer time and greater scale to materialize and cannot capture equilibrium responses that occur at a level higher than the randomization.

Motivated by these limitations, the primary empirical exercise of the paper proposes a model of consumer demand and price competition in Section 4, and estimates 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, and 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 into three sources: variation in seller costs, perceived seller vertical and horizontal 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 buyers tends to increase, rather than decrease, price competition. In contrast, when we apply the same exercise to a product category that exhibits much greater heterogeneity across items, prioritizing prices in the search design has negative consequences, and appears less efficient than search designs that prioritize product quality.

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., 2015), 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. Lewis and Wang (2013) examine the theoretical conditions under which reducing search frictions benefits all market participants. 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

## 2.1 Conceptual Framework

We begin by describing the simple economics of platform design. Consider $J$ sellers, each listing one product for sale on a single platform. Each product $j$ (offered by seller $j$) is associated with a fixed vector of product attributes $x_j$ and is offered for sale at a posted price $p_j$, which is determined by the seller. Each consumer $i$ who arrives at the platform is defined by a vector of characteristics $\zeta_i$, drawn from a population distribution $F$. Each consumer has a unit demand and decides which product to purchase, or not to purchase at all. Conditional on purchasing product $j$, consumer $i$'s utility is given by $u(x_j, p_j; \zeta_i)$.

So far we described a standard, traditional setting of demand and supply of differentiated goods. The distinction, which is the focus of this paper, is the existence of a platform as a market intermediary, whose main role is in allocating consumers' attention and/or awareness to different products. This role is less essential in more traditional markets, where the number of products is limited and consumers are likely to be reasonably familiar with most of the products. But in online markets, where there are hundreds or sometimes thousands of different competing products available for sale at a given time, and product churn is high, consumers cannot be expected to consider, or even be aware of, all these products. This is the context in which the platform has an important role in deciding which products to make visible to a given consumer.

A simple generic way to model the platform is by assuming that the platform sets an awareness/visibility function $a_{ij} \in [0, 1]$, where $a_{ij}$ is the probability that product $j$ is being considered by consumer $i$. For example, the platform can decide not to show product $j$ to anyone, in which case $a_{ij} = 0$ for all $i$, or can decide to rank order certain products when it presents search results, which would imply $a_{ij} > a_{ik}$ for all $i$ iff product $j$ is ranked higher than product $k$ for all searches. We will consider below the trade-offs associated with different platform designs, and while we will not explicitly model the optimality of the platform design, it would be implicit in the discussion that technological or consumer attention generates a constraint of the form $\sum_j a_{ij} \leq K_i$. To keep things simple, and consistent with the empirical setting presented below, we further assume that $a_{ij} = a_j = a(p_j, x_j; p_{-j}, x_{-j})$ for all $i$. That is, the platform presents products to consumers based on their prices and attributes, but does not discriminate presentation across consumers.[1]

Given this setting, platform design implies (possibly stochastic) choice sets, $L$, for consumers, so that overall demand for product $j$ is given by

$$D_j(p_j, p_{-j}) = \sum_{L \in 2^J} a_L D_j(p_j, p_{-j}; L), \tag{1}$$

where

$$D_j(p_j, p_{-j}; L) = \int 1(u(x_j, p_j; \zeta_i) \geq u(x_k, p_k; \zeta_i) \forall k \in L) dF(\zeta_i) \tag{2}$$

and

$$a_L = \left(\prod_{j \in L} a_j\right) \left(\prod_{j \notin L} (1 - a_j)\right). \tag{3}$$

This consideration set approach to modeling demand is not new (see, e.g., Goeree 2008; Honka et al. 2014), but our focus is different. While earlier papers mostly took the consideration sets as given, our focus is on the platform's decision as to how to affect it.[2] Note also that we make the assumption that the platform design affects choices, but does not enter the consumer's utility directly; this can be motivated by the fact that conditional on engaging in a search process, the consumer exerts a fixed amount of effort regardless of the outcome.

Consider now the seller's pricing decisions. Seller $j$ sets $p_j$ to maximize profits

$$\pi_j = \max_{p_j} D_j(p_j, p_{-j})(p_j - c_j), \tag{4}$$

[1] From an ex ante perspective, this still allows for setting $0 < a_j < 1$, which would be implemented by randomizing across consumers, and thus generates discrimination ex post.

[2] The literature sometimes draws a distinction between consumer actively "considering" a product and the consumer seeing a product but ultimately disregarding it. We will treat a product as part of a consumer's consideration set if she is shown the offer, regardless of how seriously she considers it when deciding whether to purchase. We discuss in Appendix A why our data do not allow us to make such a distinction.

leading to the familiar first order condition

$$p_j = c_j - \left( \frac{\partial D_j(p_j, p_{-j})}{\partial p_j} \right)^{-1} D_j(p_j, p_{-j}). \tag{5}$$

Note that we can write

$$\frac{\partial D_j(p_j, p_{-j})}{\partial p_j} = \sum_L a_L \frac{\partial D_j(p_j, p_{-j}; L)}{\partial p_j} + \sum_L \frac{\partial a_L}{\partial p_j} D_j(p_j, p_{-j}; L), \tag{6}$$

so the price has two distinct effects. One is the usual effect on demand: conditional on considering product $j$, consumers are more likely to buy it if its price is lower. The second effect of price depends on the platform design. If the platform is more likely to show the product when its price is lower – that is, if $\frac{\partial a_j}{\partial p_j} < 0$ – it provides yet another incentive for sellers to reduce prices.

This will be a key point that we will focus on throughout the paper. The platform has two distinct roles in choosing its search design. One is the familiar role of generating more efficient sorting: trying to help imperfectly informed (or imperfectly attentive) consumers find their desired product within a large assortment of different products. The second role of the platform design is to exert stronger pricing incentives on sellers. It seems natural that if the platform tries to maximize consumer surplus (which we assume is the case),[3] the platform should tilt its optimal design from trying to predict demand for product $j$ towards a design that assigns greater weight to price, as a way to increase demand elasticities faced by the seller.

## 2.2 A Toy Example

We now use this framework to present a highly stylized example, which illustrates some key elements that will be the focus of the empirical exercise. Consider two products ($J = 2$), which are associated with differentiated qualities $q_1 > q_2$ such that $q_1 = q$ and we normalize $q_2 = 0$. Corresponding marginal costs are $c_1 = c$ and $c_2 = 0$. Consumers have unit demand,

[3] In the context of most e-commerce platforms, including eBay, it seems reasonable to approximate platform revenues as a fixed share of transaction volume. To the extent that the platform maximizes long-run (rather than short-run) volume, and driving consumers to the platform (rather than sellers) is the main challenge, short-run consumer surplus would be highly correlated with long-run platform revenues.

and consumer $i$'s utility from product $j = 1, 2$ is given by $u_{ij} = \zeta_i + q_j - \alpha p_j$ where $\zeta_i$ is distributed uniformly on $[0, 1]$, and utility from the outside option for all consumers is normalized to $u_{i0} = 0$. We further assume that the platform can only show to consumers a single product and (as before) cannot discriminate what it shows across consumers. Within this context, the platform design is reduced to the probability it would show each product, $a_1$ and $a_2 = 1 - a_1$, as a function of qualities ($q_1$ and $q_2$) and prices ($p_1$ and $p_2$).

From a seller's perspective, demand is driven by consumer demand and the platform strategy:

$$D_j(p_j, p_{-j}) = \begin{cases} a_j(p_j, p_{-j}) & if \quad p_j < q_j/\alpha \\ a_j(p_j, p_{-j})(1 + q_j - \alpha p_j) & if \quad p_j \in [q_j/\alpha, (1 + q_j)/\alpha] \\ 0 & if \quad p_j > (1 + q_j)/\alpha \end{cases}, \tag{7}$$

and sellers set prices to maximize profits.

Finally, for illustration, it will also be convenient to assume that the platform cannot perfectly implement its design strategy (e.g., because there are thousands of products and quality is estimated/measured, by the platform, with noise). Specifically, we assume that product 1 is shown to consumers with probability

$$a_1 = \frac{[\exp(q - \beta p_1)]^{1/\sigma}}{[\exp(q - \beta p_1)]^{1/\sigma} + [\exp(-\beta p_2)]^{1/\sigma}} \tag{8}$$

and product 2 is shown with probability $a_2 = 1 - a_1$. The platform's design depends on its choice of the parameter $\beta$; that is, on the extent to which lower prices are more likely to be shown to consumers.

Figure 1 illustrates the trade-off associated with different platform strategies, that is with different choices of $\beta$. We do so by solving for equilibrium pricing for a given set of parameters – $\alpha = 0.5$, $\sigma = 1$, $q = 1$, and $c = 0.5$ – but the basic insights apply more generally. When $\beta = 0$ both sellers set the monopolistic price, $p_j^M = (1 + q_j + \alpha c_j)/2\alpha$, so that $p_1 = 2.25$ and $p_2 = 1$. The figure then illustrates the two offsetting forces that are in play as $\beta$ increases and the platform assigns greater weight to prices. On one hand, as $\beta$ increases, sellers' effective demand becomes more price sensitive, and in equilibrium both sellers set lower prices, benefiting consumers. On the other hand, as $\beta$ increases, the cheaper (and lower quality) product obtains "preferential" treatment by the platform, and is being shown more often. The inefficiency is easy to see at the extreme, as $\beta$ approaches infinity; then, the cheaper product is always shown, and (given the cost differences) the higher cost (and higher quality) product (product 1) is never shown, which is inefficient. As the bottom panel of Figure 1 shows, the trade-off is then resolved with an intermediate value of $\beta$ ($\beta^* = 4.55$ at the given values of the parameters), which maximizes consumer surplus. It is important to note that this optimal value of $\beta$ is still significantly greater than the corresponding weight assigned to price by consumers (recall $\alpha = 0.5$).

In Figure 2 we use the same setting to illustrate some comparative statics, which are useful in thinking about the optimal platform design across a range of different product categories. The top left panel shows how the optimal platform design, $\beta^*$, varies with the price sensitivity of consumer demand. Naturally, all else equal, as consumers are more price sensitive (higher $\alpha$), it is more efficient to increase the importance of price in the platform design, thus leading to higher $\beta^*$. In the top right panel, we show how the platform design changes with the cost $c$ of the higher quality product. As the cost increases, the seller of the higher quality product has less ability to mark up its price, so the value to emphasizing price in the platform design is lower, and $\beta^*$ is lower. Similarly, the bottom left panel shows that as the products are more vertically differentiated, again the optimal platform design should apply lower weight to price as distorting demand toward the cheaper product leads to greater inefficiency. In the bottom right panel, we show that as the noise in measuring quality ($\sigma$) increases, the platform design applies a higher weight to price, which is the product characteristic that can be targeted without error.

## 2.3 Existing Approaches to Platform Design

The above framework captures what we view as the two key dimensions of consumer search in online markets. The first is to try to "predict" consumers' demand, and guide them 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 desires, and by doing so amplify the effective price elasticity faced by sellers. Empirically, due to different consumer mix and different product offering, online platforms adopt heterogeneous approaches to the search problem by emphasizing one of the above dimensions, or both.

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 (in the context of the model, this can be thought of as a case with relatively low $\sigma$). 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 3 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 her 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 unproductive tactics that avoid head-to-head price competition.

# 3 Setting and Motivating Evidence

## 3.1 Background: Changes in Platform Design on eBay

With this general framework in mind, the rest of the paper will use detailed data from eBay, taking advantage of an interesting episode of platform design changes to eBay's marketplace, which allows us to compare the different approaches. Appendix B provides more details about the data construction.

The top panel of Figure 4 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.[4] 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.[5] 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 spring 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 4.[6] The product page has a prominent "Buy Box" that displays the seller with the lowest posted price (plus shipping)

[4] When 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.

[5] At various times, 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.

[6] The 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.

among those reputed sellers who are classified as "top rated" by eBay. 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. The two designs correspond closely to the cases we considered in our stylized example of Section 2.2. The Best Match algorithm incorporated only non-price product characteristics into the ordering of search results, which is analogous to setting $\beta = 0$ in our example, while the product page ordered fixed price listings based only on price, which is analogous to setting $\beta$ to be quite high.

About a year later, however, in summer 2012, eBay evaluated the redesign with 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).[7] The experiment, which we were not involved in, was run on 20 percent of the site's traffic. After being shown initial results using the randomized type of results page as a default, users could choose to browse using the other type of results page. 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. Subsequent to the experiment, eBay made the original, Best Match results the default view for searchers.[8]

While much of our analysis below will focus on the initial, 2011 changes, we also report the main patterns that emerge from the subsequent, 2012 A/B experiment.

[7] The 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. Users with customized search preferences, such as preferring results sorted by shipping distance, were not affected by the experiment.

[8] The search design has continued to evolve, but the default search results continue to be a Best Match relevance ranking, albeit one that it likely to be correlated with price for well-defined products.

## 3.2 The Impact of The Product Page: Descriptive Evidence

The new product page was introduced on May 19, 2011.[9] However, the traditional search results 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 (and defaulted into) 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 example, in the digital camera category, consumers pay on average around 18% more than if they had selected the 10th percentile price.

[9]eBay 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.

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%, demonstrating the redesign's long-run effect on seller pricing.

Several statistics are suggestive of changes in which listings consumers considered. 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 (TRS).[10]

Figure 5 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 will use to estimate our model below. The top panel shows the distribution of new, fixed price 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 a new, fixed price Halo Reach video game listing. After the search redesign, consumers generally clicked fewer times on their way to a purchase, consistent with a more streamlined process.

[10]As 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.

## 3.3 Moving Back to Traditional Best Match: Results from eBay's A/B Experiment

As the stylized example in Section 2.2 highlights, the effects of platform design likely depend on the product's degree of quality differentiation, $q$. The A/B experiment provides a clean comparison of demand behavior under the different platform designs. Note, however, that because the two designs were simultaneously active and sellers set a single price per listing, the experiment will not induce any differential changes to pricing incentives. Therefore, the experimental results will only capture the platform's ability to efficiently sort consumers to listings and not its effect on pricing. We will return to this shortcoming in the next section.

We first examine the experiment's average results, aggregating across all product categories. A starting point is that the experiment did succeed in steering users toward particular results. For users randomly assigned to the product page by default, 3.45% of all sessions included a product page visit, compared to 1.87% for users who were randomly assigned to 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). As mentioned, this comparison made eBay make the traditional Best Match results the default view for searchers.

The higher purchase rate for the Best Match group (despite slightly higher average prices paid) suggests that non-price characteristics play an important role. To explore further, 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 the 200 products with product pages that were visited at least 1,000 times and had at least 20 purchases during the experiment, and to fixed price listings for these products.

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. 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.[11]

Figure 6 reports statistics based on this cut of the experimental data. The top panel shows, by product, the price effect of making the Best Match the default search results page relative to making the Product Page the default. The bottom panel shows the same for quantity. The effect is highly heterogeneous across products, presumably reflecting a combination of sampling variation and idiosyncrasies across products in how much residual heterogeneity across listings exists. Overall, while the average price and quantity effects are both small, but positive, there is a remarkable variation across products, with products that are more heterogeneous having the greatest (and non-trivial) positive quantity and price effect when Best Match is used, while more homogeneous products are associated with essentially no quantity effect and a slight negative price effect due to Best Match.

### 3.4 Discussion

The results in Section 3.2 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 to have resulted from both a change in purchasing patterns and a fall in the distribution of posted prices. The A/B experiment results (reported in Section 3.3) highlight the important heterogeneity in this response across products, even within a fairly narrow product category. The estimated heterogeneous effects confirm that the platform's trade-off between prioritizing price or non-price characteristics depends closely on the product's level of differentiation. Taken together, the collection of descriptive results reported in this section suggests that the platform design is an important feature of the eBay market, and that platform changes could make a non-trivial difference for market outcomes. At the same time, while the patterns are suggestive regarding some of the channels that are in play, the analysis also highlights the difficulties in interpreting the empirical patterns without a model.

[11] Implicitly 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 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. For our empirical model, we will construct a more direct measure of listing quality. The measure will rely on extensive search results data and thus is not practical for analysis across many products.

Consider the results from the A/B experiment first in light of the empirical framework presented in Section 2. The premise of the framework is that pricing on the platform responds to the platform design, yet the A/B experiment, while useful in highlighting the importance of heterogeneity across products, cannot capture this pricing response for two reasons. First, sellers respond to their expected demand, and the experiment affected only a small share of users. Furthermore, expected demand is integrated over users reaching both types of search results, so we cannot compare across experimental groups. Second, sellers' pricing decision and strategies are unlikely to respond immediately, so although the short run response (captured in the results presented earlier) might be indicative of the longer run effects, quantitatively it could be quite different. On the other hand, the A/B results suggest that heterogeneity across products appears to be quite important, and this may make it difficult to interpret the category-average patterns we presented in Section 3.2.

Therefore, in the next section we develop and estimate a more complete model of the underlying economic primitives. The model allows us to explain the price levels 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 and product types not present in the data.

# 4 An Empirical Model

In this section, we describe a model of consumer search and price competition. In the next section, we estimate the model's parameters using data from a single product market, and 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 redesign affected all of these.

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 either the targeted or a non-targeted product. The targeted product is the product that is the focus of the market (e.g., the product corresponding to the search terms the user specifies) while non-targeted products are other, possibly-related products. We allow listings to vary only by their price $p$, vertical quality $q$,[12] and by whether they are listed by a top-rated seller (denoted $TRS$). We attribute any additional differentiation to a logit error. We assume that consumer $i$'s utility from listing $j$ of the targeted product is given by

$$u_{ij} = \alpha_0 + \alpha_1 p_j + \alpha_2 TRS_j + \alpha_3 p_j TRS_j + \alpha_4 q_j + \varepsilon_{ij}, \tag{9}$$

where $\epsilon_{ij}$ is distributed Type I extreme value and is independent of the listing's price, quality, and $TRS$ status.

We assume that consumer $i$'s utility from listing $m$ of the non-targeted product is given by

$$u_{im} = \delta + \lambda \varepsilon_{im}, \tag{10}$$

where $\epsilon_{im}$ is independently distributed Type I extreme value. We parameterize the degree of horizontal differentiation of non-targeted products by $\lambda$ to allow non-targeted products to be more or less differentiated than targeted products.[13]

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. Let $J_i^J$ and $J_i^M$ denote the set of targeted and non-targeted listings, respectively, in the consideration set. 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 $u_{i0} = \varepsilon_{i0}$, where $\varepsilon_{i0}$ is also an independent Type

[12] We describe the way we measure quality in the next section.

[13] This parameterization yields the same substitution patterns across the non-targeted and targeted products as a nested logit.

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 redesign, 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 $L_i = (L_i^J, L_i^M)$ 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 targeted and non-targeted listings. We assume that $L_i$, the number of the items in the consideration set, is independent of any particular buyer characteristics, or the distribution of prices.

Which listings of the targeted product make it into the consideration set? Prior to the redesign, we noted that price did not factor directly into search ranking, but that after the redesign, 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 listings are selected onto the displayed page. Specifically, we assume that products are sampled from the set of available targeted products $\mathcal{J}_i$, such that each product $j \in \mathcal{J}_i$ is associated with a sampling weight of $\omega_j$. Before the redesign, the sampling weight is the listing's quality, $q_j$. While a listing's quality may be correlated with its price, it is fixed and thus price changes do not affect the listing's sampling weight. This reflects eBay's use of the Best Match algorithm, which attempts to rank listings based on a single-dimensional measure of a listing's quality. This model therefore allows us to use browsing data from before the redesign to infer listings' quality, which we then use in our demand estimation.

After the redesign, listing $j$'s sampling weight is

$$\omega_j = \exp\left[-\gamma\left(\frac{p_j - \min_{k\in\mathcal{J}_i^J}(p_k)}{std_{k\in\mathcal{J}_i^J}(p_k)}\right)\right]. \tag{11}$$

Consumer $i$'s consideration set is then constructed by sampling $L_i^J$ products from $\mathcal{J}_i^J$, without replacement. This implies that the consideration set of targeted listings is drawn from a Wallenius' non-central hypergeometric distribution. We expect $\gamma > 0$ so that lower price items are disproportionately selected into the consideration set after the platform redesign. We further modify the sampling process after the redesign to incorporate a Buy Box by reserving one position in the consideration set for a TRS product. Specifically, we draw the first product in the consideration set from the set of available targeted products from TRS sellers, $\mathcal{J}_i^{TRS,J}$. Denote this product $j_i^0$. We then draw the remaining $L_i^J - 1$ products from $\mathcal{J}_i^J \backslash j_i^0$, without replacement. Below we estimate $q$ and $\gamma$ using the browsing data that records the listings that appeared on pages buyers actually visited.

### 4.3 Pricing Behavior

We model seller as pricing using a standard Nash Equilibrium assumption.[14] Seller $j$ of a targeted product with marginal cost $c_j$ sets its price to solve

$$\max_{p_j}(p_j - c_j)D_j(p_j). \tag{12}$$

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

[14] In our data, the modal seller is associated with a single listing. For simplicity, even for sellers who sell multiple items, we assume that prices are set for each listing independently. This assumption is unlikely to affect the results much given that the large number of sellers and products make substitution across listinjgs of the same seller minimal.

choice probabilities, we have

$$D_j(p_j) = \sum_{J:\, j \in J \subseteq \mathcal{J}} \left[ \frac{\exp\left(\alpha_0 + \alpha_1 p_j + \alpha_2 TRS_j + \alpha_3 p_j TRS_j + \alpha_4 q_j\right)}{1 + \exp\left(\delta + \lambda \ln |J^M|\right) + \sum_{k \in J^J} \exp\begin{pmatrix} \alpha_0 + \alpha_1 p_k + \alpha_2 TRS_k + \\ + \alpha_3 p_k TRS_k + \alpha_4 q_k \end{pmatrix}} \right] \Pr\left(J | \mathcal{J}\right). \tag{13}$$

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 (stochastic) set of competing products over the entire lifetime of the listing. The competing items are drawn from the approximately one month (either "before" or "after") period considered. When a competing listing is simulated to be purchased, it is replaced on the site by another listing from the period.[15] See Backus and Lewis (2016) for related work on stochastic sets of competing products.

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 consideration set. With this notation, the optimal price $p_j$ satisfies:

$$\frac{p_j}{c_j} = \left(1 + \frac{1}{\eta_D}\right)^{-1} = \left(1 + \frac{1}{\eta_A + \eta_Q}\right)^{-1}, \tag{14}$$

where $\eta_D$, $\eta_A$, $\eta_Q$ are respective price elasticities.[16] 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 $\gamma$ intensifies the first effect. In addition, increasing $\gamma$ effectively faces each seller with tougher competition conditional on making it

[15] The new listing that replaces the purchased one is sampled according to the length of time the listings were actually active on eBay during our estimation period. Thus, sellers are more likely to face competitors who are selling many units at once or competitors with relatively unattractive products as they remain on the simulated site for longer. Every 100 searches we exogenously reset the set of competing products to account for the feature that some eBay listings expire without being purchased.

[16] To see this, note that $\eta_D = D'(p/D)$, $D = AQ$, and $D' = Q'A + A'Q$. This implies that $\eta_D = D'(p/D) = Q'(p/Q) + A'(p/A) = \eta_Q + \eta_A$.

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. In Appendix A we report results from a much richer consumer search model, which more explicitly models the decision of how to search, and which item to click, before a final purchase decision is made. As we discuss in Appendix A, the estimated model described above can be viewed as a more general demand framework, which captures some of the key elasticities that affect the platform design using free parameters, while summarizing many other components of the consumer search process in a reduced form.

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 quality and none are top-rated. As $\gamma \to \infty$, 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 markups.

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 consumers. It might be 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.[17]

# 5 Estimation and Results

## 5.1 Estimation Sample

To estimate the model, 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 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. In fact, 51% of Halo Reach listings share the same title, our proxy for degree of product homogeneity when we compared products from the A/B experiments in Section 3. This places Halo Reach at the 82nd percentile for product homogeneity.[18]

The data for the analysis come directly from eBay and is described in more detail in Appendix B. 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.[19] The

[17] One 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 non-targeted 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 targeted listings, or some mixture of the two.

[18] The Halo Reach viedo game is a farily homogenous product. It would have been interesting to compare and contrast results from this product against a less homogeneous product, but once a product becomes heterogeneous the challenge faced by the platform immediately translates to a challenge faced by the researcher: to identify the set of listings who would be classified as such a product. Instead, we therefore use the counterfactual exercise as way to qunatitatively assess the tradeoff.

[19] As 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.

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." This results in $14,753$ visits to the search results page ($9,409$ of them in the pre-period) and $6,733$ visits to the product page (18 in the pre-period).[20]

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). We define "targeted listings" as new Halo Reach items, listed either with a posted price, or as an auction but with a Buy-It-Now price.[21] The non-targeted fixed price listings are those that appear in search results but do not meet our definition of targeted because they are used items or are not the Halo Reach video game itself.

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 if 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.[22]

## 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

[20] The "product page" in the pre-period was more rudimentary than one introduced on May 19 (see footnote 6), and relatively few people navigated to it.

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

[22] For 89% of the targeted 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.

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.

Table 2 also reports our measure of item quality, which relies on eBay's internal rankings of listings that enters the Best Match algorithm. eBay assigns each listing a score, which does not depend on price, that is intended to reflect how attractive the listing is to consumers. While we do not observe the score directly, we infer it from the frequency with which a listing appears in Best Match search results when it is active on the site. As the Best Match algorithm samples listings onto the page without replacement, we use Wallenius' non-central hypergeometric distribution, just as we specified in our model, to estimate a listing's Best Match score.[23] We are only able to estimate quality for listings appearing in the before period,[24] which (as discussed below) further motivates our use of only before period data in estimating the demand parameters. The scores are only identified up to a scalar, so we normalize them to be between 0 and 1. As reported in Table 2, the median listing has extremely low quality while the mean quality is an order of magnitude larger, though still small relative to the best listing's quality. The quality distribution is highly right skewed with a small number of listings at much higher quality levels than the rest.

The bottom panel in Table 2 shows statistics on searches. In the before period, listings appearing in search results were positively selected on quality. Consumers saw lower prices in the after period, and a larger fraction of searches resulted in purchases of targeted listings (1.2% compared to 1.0%) and non-targeted listings (2.6% compared to 2.0%). Recall that in Figure 5, displayed earlier, we already showed that there was a significant increase in the number of targeted listings consumers saw after a search. We also showed in Figure 5 that eventual purchasers seem to have had an easier time getting to the point of sale: eventual

[23] When estimating the listing's quality, we account for variation in the number of search results on a page. For instance, if listings A and B each appeared in half of their eligible searches, but searches when listing A was active led to many more search results on average, we would infer a higher score for listing B. Additionally, solving for all of the listings' scores simultaneously would be computationally infeasible. We therefore make the simplification that each listing is competing for page space with other listings all of average quality. Simulations suggest that this simplification has minimal effect on our estimates.

[24] While there are many searches in the after period that use the Search Results Page (see Table 2), only few are sorted by Best Match (compared to, say, time ending soonest).

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 sets to estimate the demand parameters, and then impose an assumption of optimal pricing to back out the implied marginal costs of each listing. Appendix B provides more details.

The first step is to estimate the consideration set model. We obtain the empirical distribution of $L_i$ (the number of targeted and non-targeted items sampled by a consumer) directly from the browsing data, and separately for the before and after periods (see Figure 5). We use the browsing data to estimate the sampling weight of each listing in the before period.[25] For the sampling process in the after period, we estimate the sampling parameter $\gamma$ in equation (11) that determines the extent to which cheaper listings are more likely to enter the results page. We estimate $\gamma$ using ordinary least squares (see Appendix B) and obtain an estimate of 0.80 (with a standard error of 0.14). This implies that a ten percent reduction in the posted price would, on average, make the listing 27% more likely to be part of a consumer's consideration set.

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. A top-rated seller pricing at \$37 has an equal probability of transacting as a non-top-rated seller of similar quality pricing at \$35.21. 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 redesign, so this effect is large. Price also has a very large effect. The price elasticity implied by the estimates is about -11. It is even higher (closer to -14) for TRS sellers. The profit margin implied by

[25] For the sampling weight, we use all searches in the before period that reached the search results page. A small number of these searches were made with customized search preferences (e.g., ordering by time ending soonest) that meant the results were not displayed according to Best Match.

these estimates is about 10%: \$2.94 for TRS sellers and \$4.20 for other sellers. The degree of quality differentiation is limited and right-skewed. The difference between the lowest- and highest-quality listings is equivalent to just \$2.67. Finally, we estimate that the average non-targeted listing is less desirable than the average targeted listing, and non-targeted listings have a higher degree of horizontal differentiation. This is consistent with the non-targeted listings including a diversity of products.

The last step is to estimate seller costs. From the seller's optimization problem,[26] we have:

$$c_j = p_j + D_{jt}(p_j) / D'_{jt}(p_j), \tag{15}$$

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 listing in the "before" period. Then we use the first order condition above to back out the cost $c_j$ that rationalizes each listing's price as optimal.

The implied cost distribution is presented in Figure 7, 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 just slightly under \$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.[27] We discuss the high price sellers in more detail in Section 6.3.

[26] As in the example of Section 2, we assume that sellers set prices simultaneously to maximize expected profits, where the expectations are taken over the all consumers and consideration sets the seller's item could be part of, taking the platform design as given and assuming (as in a Nash Equilibrium) that competing sellers set prices in the same way.

[27] We 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 redesign, the observed consideration sets are quite representative, in terms of listed prices, of the full set of listings.

# 6 The Effect of Search Design

## 6.1 Changing the Search Design

We first use our estimates to assess the introduction of the product page and compare the model predictions to the data. 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 listing cost-quality distribution remains unchanged. The results from this exercise are reported in Table 4, and Figures 8 and 9. 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 8 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. The implication is that seller margins should fall. Comparing the top and bottom rows of the top panel of Table 4 shows that the median optimal margin fell from \$2.94 (or 8% of price) to \$2.46 for TRS sellers, and from \$4.20 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 5, there was a noticeable increase in the size of consideration sets, and buyers had a much smaller chance of seeing just a single targeted listing. In addition, price became an important factor in entering the consideration set. With our estimate of $\gamma = 0.80$ for the after period, a ten percent price reduction increases the odds of appearing in the consideration set from 0.08 to 0.11, providing sellers with a new incentive to reduce prices. The new platform also included a "Buy Box" that guaranteed at least one listing from a top-rated seller would appear in the consideration set. 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 consideration set size, the increase in listings, the Buy Box, and the increase in $\gamma$. 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 $\gamma$) has by far the largest effect on seller incentives. The increased size of consideration sets, the increase in the number of sellers, and the introduction of the Buy Box have minimal effects on equilibrium margins. The increase in purchase rates is driven by making price a factor in forming consideration sets in combination with the other redesign elements.

In the bottom panel, we evaluate the importance of the supply response in explaining the increased purchase rates. We implement the four components of the redesign – the larger consideration sets, the increase in the number of sellers, the Buy Box, and making price a factor in forming consideration sets – but fix listing prices. We find that 62% of the total effect on purchase rates is driven by the redesign without a price response. Thus, the remaining 38% comes from the supply response.

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 9 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.[28] Our model's predictions outperform other reasonable benchmarks. Between the before and after periods, Amazon prices were remarkably similar on average, dropping only from \$37.59 to \$37.39, though prices for third-party sellers listing used versions of the game on Amazon fell from \$22.93 to \$19.51. We therefore consider three alternative predictions: no price change, price change equal to the change in Amazon's list price, and price change equal

[28] While the estimation of $\gamma$ projects platform outcomes onto prices from after the redesign, the correlation between prices and probability of appearing in a choice set is a much larger driver of the results than a small shift in the price distribution. We thus view this use of data from after the redesign as having little effect on how well we fit the price distribution.

to the change in Amazon's third-party used list prices. In all three cases, a Kolmogorov-Smirnov test of the null hypothesis that the distribution of the predicted prices is the same as the actual distribution is rejected at the 5% level. On the other hand, when we test our model's predicted price distribution, we fail to reject (p-value of 0.13).

It is also possible to compare how well our model predicts other data moments. In the bottom rows of Table 3 we show the consumer purchase rates of targeted listings predicted by the model, and those that we observe after the redesign. We predict a sizable increase (the model predicts 1.47%, which is a bit higher than the 1.23% observed in the data). Despite modeling the Buy Box without adding parameters to the model, our estimates of the percentage of targeted purchases coming from the Buy Box listing are very close (the model predicts 65% compared to the 64% observed in the data). We match some additional unreported moments fairly well. We predict a large increase in TRS purchase share (44% to 83%) which matches the trend in the data (38% to 65%). We also predict a drop in average transacted prices ($34.30 to $33.54) similar to the one in the data ($34.56 to $33.30) and confirm that the decrease is larger for listings not from top-rated sellers.

## 6.2 Platform Design and Product Differentiation

Recall that in the simple framework of Section 2, and motivated by the A/B experiment results reported in Section 3, we discussed how the optimal platform design varies with product type, especially with the degree of horizontal and vertical differentiation. The estimated model provides a way for us to obtain a quantitative assessment of these effects by analyzing how different platform designs perform empirically across a variety of product types. Therefore, in this section we use the estimates more broadly (and, consequently, more out of sample) to consider various ways of reducing search frictions across different types of products and to identify the sources of online price markups.

We analyze these factors in Table 5. The Table compares equilibrium outcomes for variations of the model that differ along three dimensions. Across the columns, we vary the consideration set design. In the first and second columns, we consider using quality rank and price rank to form consideration sets, analogous to the before and after search regimes. In the third column, we introduce demand weights rank where listings are sampled based on purchase probability. In the last column, we assume that all Halo Reach listings on the platform enter each consumer's consideration set. Across the two rows in each panel, we vary the degree of horizontal 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 and non-targeted listings form one nest and all sellers of targeted listings are part of a second nest. Specifically, the $\varepsilon_{ij}$ in our logit demand model (9) becomes $\zeta_{iJ} + (1-\sigma)\,\varepsilon_{ij}$, where all sellers of targeted listings share the same $\zeta_{iJ}$, whose distribution depends on $\sigma$ (see Berry, 1994). The "limited differentiation" model assumes $\sigma = 0.8$, which reduces the weight on the listing-specific error and makes the products much less differentiated than the baseline logit "differentiation" case, which corresponds to $\sigma = 0$. Across the panels, we vary the degree of vertical product differentiation. Panel A uses our estimated qualities for the Halo Reach listings, a relatively homogenous product. Panel B draws qualities from a uniform distribution with a range 30 times larger than the estimated range for Halo Reach.[29]

In the eight scenarios in Panel A, we fix the joint distribution of seller costs and quality, and draw costs and quality for each seller on the platform (assuming 28 sellers, which is the mean from the after period).[30] For the eight scenarios in Panel B, we fix the distribution of seller costs and draw quality from a uniform distribution over (-15,15). 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 and qualities. To solve for equilibrium prices and markups, 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 the sources of seller margins and differential purchase rates. Consider first the case with no search costs and limited horizontal and

[29] We choose to increase the range by a factor of 30 because this matches the degree of vertical differentiation we estimate for a less homogenous product, the Canon Mark II digital camera, active at the time of the redesign. Most of this product's listings include additional accessories (an extra lens, a battery pack, etc.), consistent with a large degree of vertical differentiation.

[30] Across all scenarios, we also fix the consideration set size distribution to match the after period and exclude the Buy Box.

vertical differentiation (top right of Panel A). In this scenario, sellers sustain positive margins 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 markup is less than \$1, and the average transaction price is \$25. Purchase rates are very high, as markups are low and users' consideration sets include all available listings.

As we incorporate search frictions through smaller consideration sets (moving from right to left on the top row), we see that search frictions lead to substantially increased markups and transaction prices and decreased purchase rates. Listing horizontal differentiation, however, is an even more potent force for pricing power. For any assumption about search design, increased seller differentiation leads to higher markups and higher prices.[31] Moreover, even with no search frictions prices are higher than in any of the limited differentiation cases. Interestingly, once listing differentiation is present, the "price rank" search design actually leads to more intense price competition 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 $\varepsilon_{ij}$) as well as price.

For a product with considerable vertical differentiation (Panel B), we reach similar conclusions when comparing outcomes between the no search frictions design and the designs with limited consideration set sizes, albeit we find somewhat higher markups for this product, supported by the increased differentiation. The main difference we see is that whereas price rank led to a higher purchase rate for the less vertically differentiated product, quality and demand rank outperform price rank here. This insight – which is obtained from a complete model that uses parameter estimated from actual data – is qualitatively similar to the key insight we obtained in Section 2 from our toy example, where increasing the degree of vertical differentiation called for a lower weighting of price in determining which product enters the consideration set.

[31] Purchase rates are not directly comparable across the rows as they come from different demand systems.

### 6.3 Discussion and Extensions

We considered a number of other permutations of the model. In one exercise, we investigated how the platform's choice of how many targeted versus non-targeted listings to include in the consideration set affects pricing incentives and purchase rates. The platform may want to have several listings of both types if some consumers are engaging in product search. The return to replacing a targeted listing with a non-targeted listing depends on how different the new listing is compared to the other listings already in the consideration set. In other words, if the non-targeted listings are relatively homogenous ($\lambda \approx 0$), adding another will not increase purchase rates much. The discussion also potentially relates to the importance of obfuscation and the ability of the platform to filter less relevant listings. If the non-targeted listings show up even for users who specify clear search terms, then better filtering might replace some extraneous listings (e.g., iPhone covers or chargers when the search terms clearly specify the device itself) with more relevant ones.

We first assessed the effect of replacing a single targeted listing with a non-targeted listing. We find the purchase rate of targeted listings decreases significantly and the purchase rate of non-targeted listings increases slightly. The effect on targeted purchase rates is stronger because most consideration sets already include many non-targeted listings while many include just one or two targeted listings. We also see minimal effect on seller margins. We then examined the consequences of a less marginal change by replacing all non-targeted listings with targeted listings 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 targeted listing 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.

As a second exercise, 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) and ignoring differences in quality and TRS status, we estimated the probability of sale as a function of the listing's posted price.[32] We did this flexibly using a local polynomial regression to obtain the demand estimate shown in Figure 10. The Figure shows that listings priced above \$41 — which constitute 35 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.[33] 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" (offering free shipping and using posted prices) also are negatively correlated with high prices. But the relationships are rather noisy, and other measures such as being top-rated and 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 offer low price listings, nor do they change their prices frequently. 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

[32] As demonstrated with our parametric model, quality differences across Halo Reach listings are minimal. TRS status, however, may strongly affect probability of sale. If we estimate the probability of sale for non-TRS listings only, we reach the same conclusions.

[33] Recall 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 $\frac{pD(p)-p'D(p')}{D(p)-D(p')} > c$.

reasonable solution for our current purposes.

# 7 Conclusion

This paper has explored the role of platform design in online markets, emphasizing the trade-off between reducing search frictions by matching buyers to their most desired products and intensifying price competition among sellers. We began with a stylized theoretical framework that illustrated the trade-off, and then used a particular episode of a platform redesign at eBay to examine this trade-off empirically. We presented descriptive evidence pointing to the impact of the platform design on both consumer and seller pricing behavior, and used results from an experiment run by eBay to show that the impact of the platform design varies quite considerably across product categories that cover more versus less homogeneous sets of products.

The descriptive evidence also highlighted the distinction between short-run and longer-run effects of platform design changes, and the potential for equilibrium effects that a smaller-scale A/B experiment may miss. In the last part of the paper, we therefore developed a complete equilibrium model and estimated it on a narrow, yet well defined product category, where we can quantitatively assess the platform design trade-off using counterfactual exercises.

Of course, our analysis is narrow, in the sense that we focused on specific product markets, where products vary only in price and quality. Yet, the broader lesson that we draw from our analysis regards the importance of the platform design in affecting not only consumer behavior and reducing their search frictions, but also in affecting sellers' decisions. Our analysis focused on price, but similar forces would be at play for other product attributes that can be changed in the short run, such as service quality or information disclosure. We view our work as an initial step. With the increasing importance of internet platforms, such as eBay, Amazon, Uber, and Airbnb, to the overall economy, we think that further studies that would assess the efficiency of different platform designs in a variety of contexts is a promising direction for further work.

# References


Backus, Matthew, and Gregory Lewis (2016). "Dynamic Demand Estimation in Auction Markets." NBER Working Paper No. 22375.

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.

Chen, Yuxin, and Song Yao (forthcoming). "Sequential Search with Refinement: Model and Application with Click-Stream Data." *Management Science*, forthcoming.

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 (2015). "Assessing Sales Strategies in Online Markets Using Matched Listings." *American Economic Journals: Microeconomics* 7(2), 215-247

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.

Honka, Elisabeth, Ali Hortacsu, and Maria Ana Vitorino (2014). "Advertising, Consumer Awareness and Choice: Evidence from the U.S. Banking Industry." Mimeo, University of Chicago.

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.

Kim, Jun P., Paulo Albuquerque, and Bart J. Bronnenberg (2010). "Online Demand under Limited Consumer Search." *Marketing Science* 29(6), 1001-1023.

Lewis, Gregory and Albert Wang (2013). "Who Benefits from Improved Search in Platform Markets?" Harvard Working Paper.

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.

Ursu, Raluca M. (2016). "The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions." New York Univesity Working Paper.

# Figure 1: Comparative Statics in Platform Choice

## Panel A: Seller Prices as a Function of Platform Choice (β)

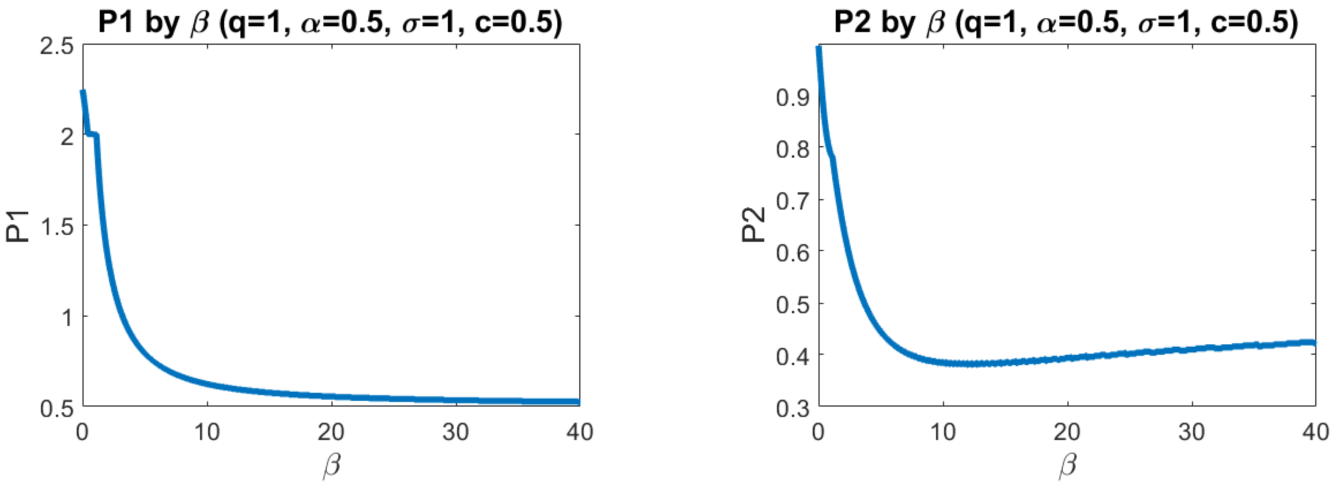


## Panel B: Seller Probabilities of Appearing as a Function of Platform Choice (β)

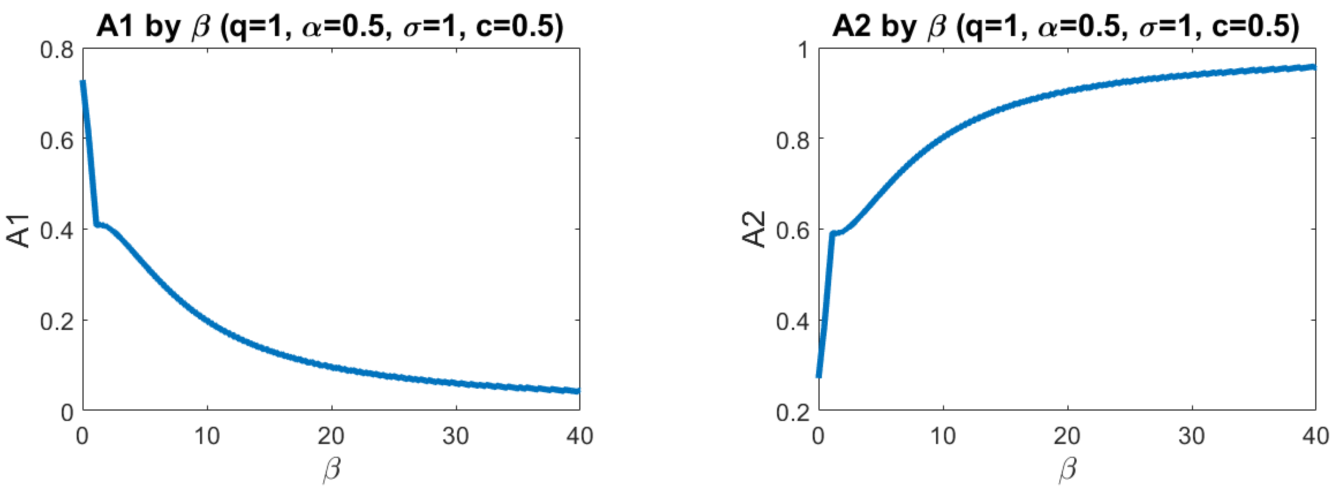


## Panel C: Consumer Surplus as a Function of Platform Choice (β)

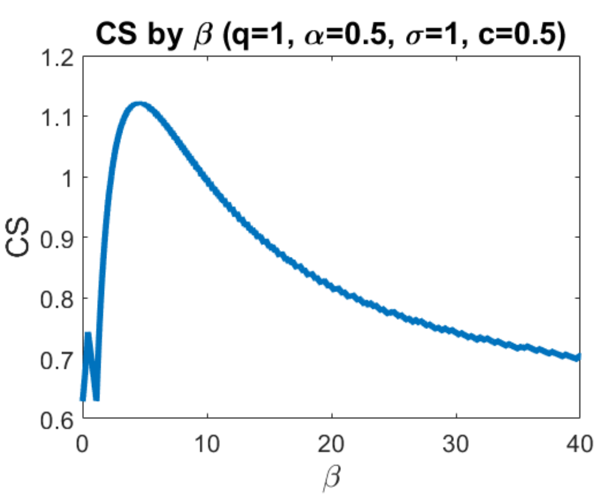


Figure shows seller prices (Panel A), seller ex ante probabilities of being shown (Panel B), and consumer surplus (Panel C) as a function of the platform's choice of the relative weight (β) on price when determining which product to show users.

# Figure 2: Optimal Platform Choice

**Panel A: β* as a Function of α**

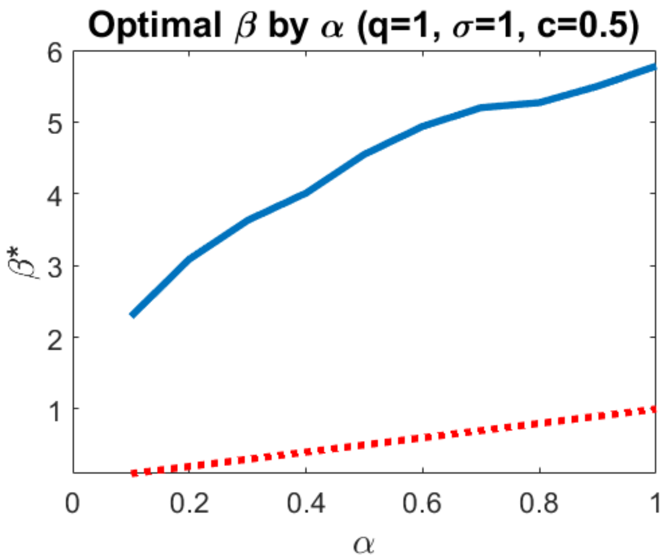


**Panel B: β* as a Function of c**

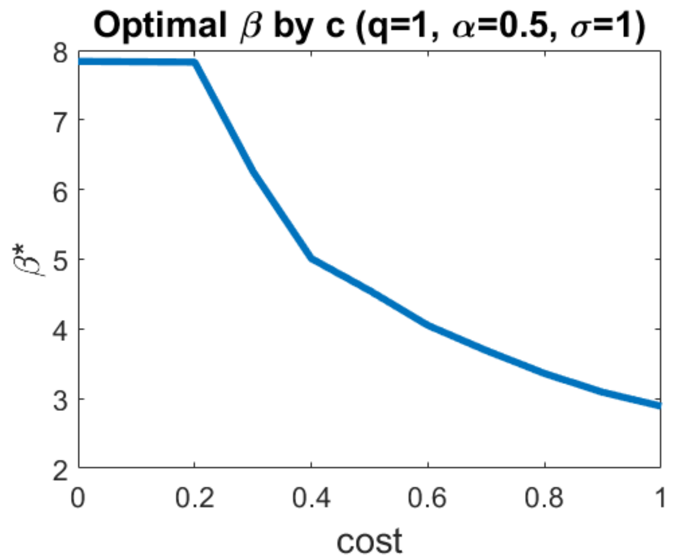


**Panel C: β* as a Function of q**

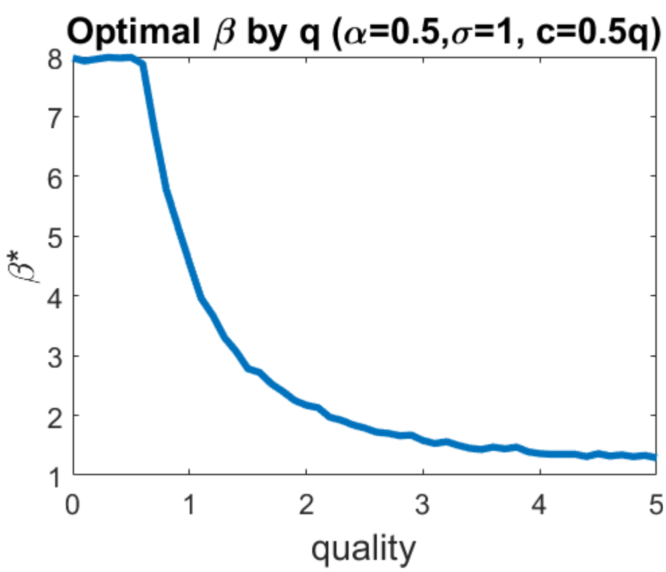


**Panel D: β* as a Function of σ**

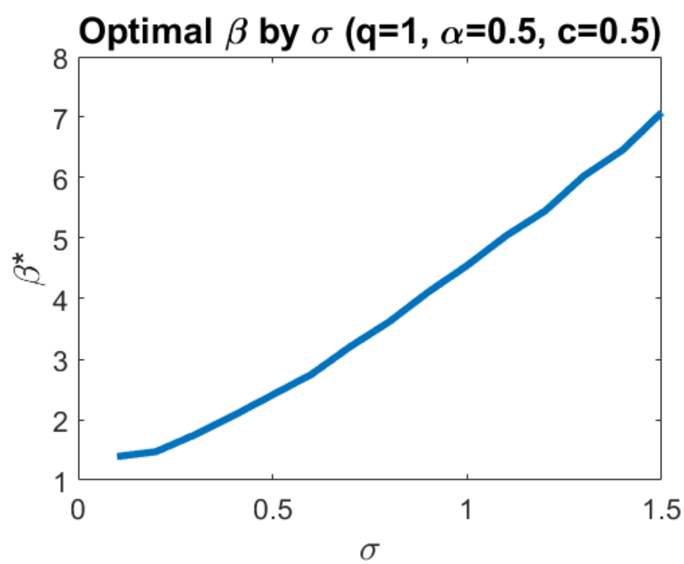


Figure shows β*, the platform's choice of the relative weight to put on price that maximizes consumer surplus, as a function of consumers' price sensitivity in demand (Panel A), the cost of the higher-quality product (Panel B), the quality difference between the products (Panel C), and the platform's noise in observing quality (Panel D). Parameters are fixed at α=0.5, q=1, c=0.5q, and σ=1 unless they are the parameter being varied for the comparative static.

# Figure 3: Different approaches to platform design

Craigslist

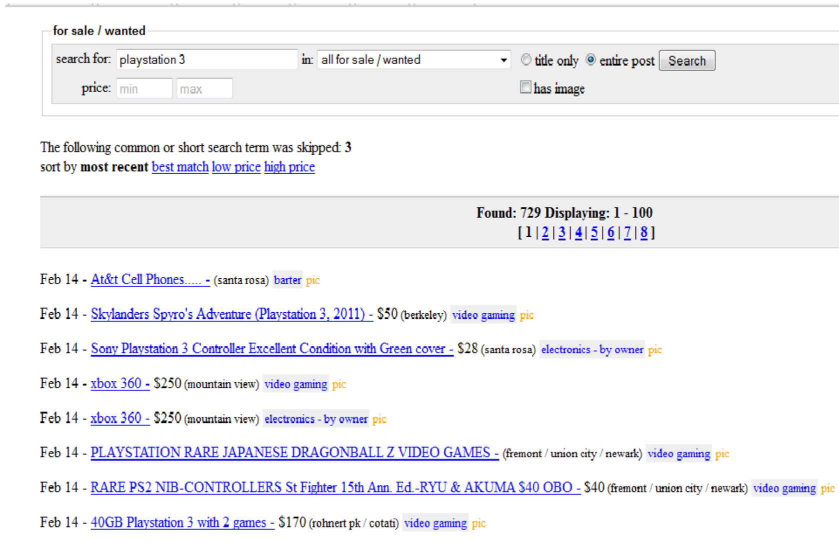


Google Shopping

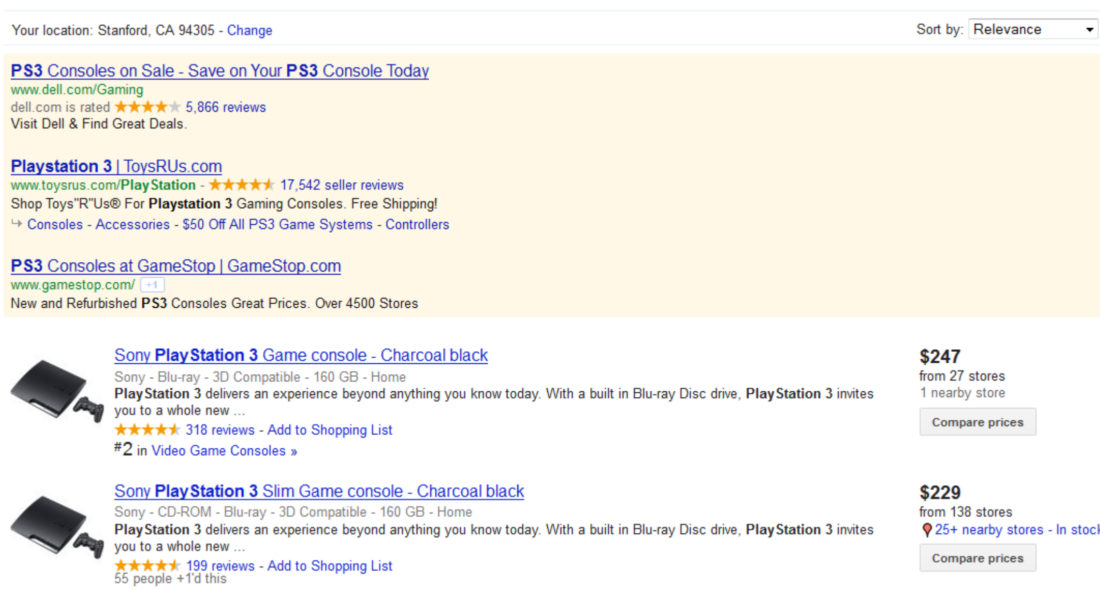


Amazon


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

# Figure 4: eBay's platform re-design

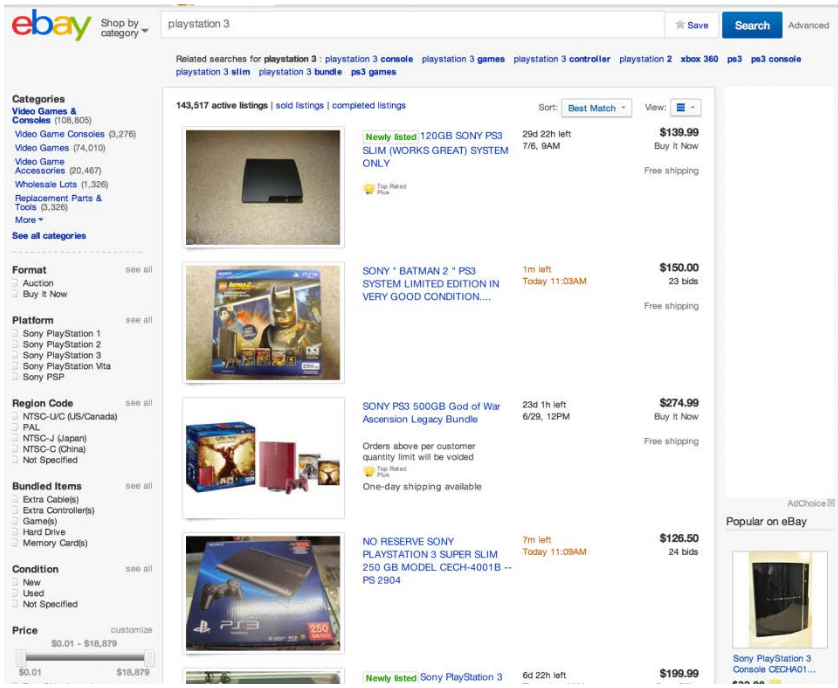


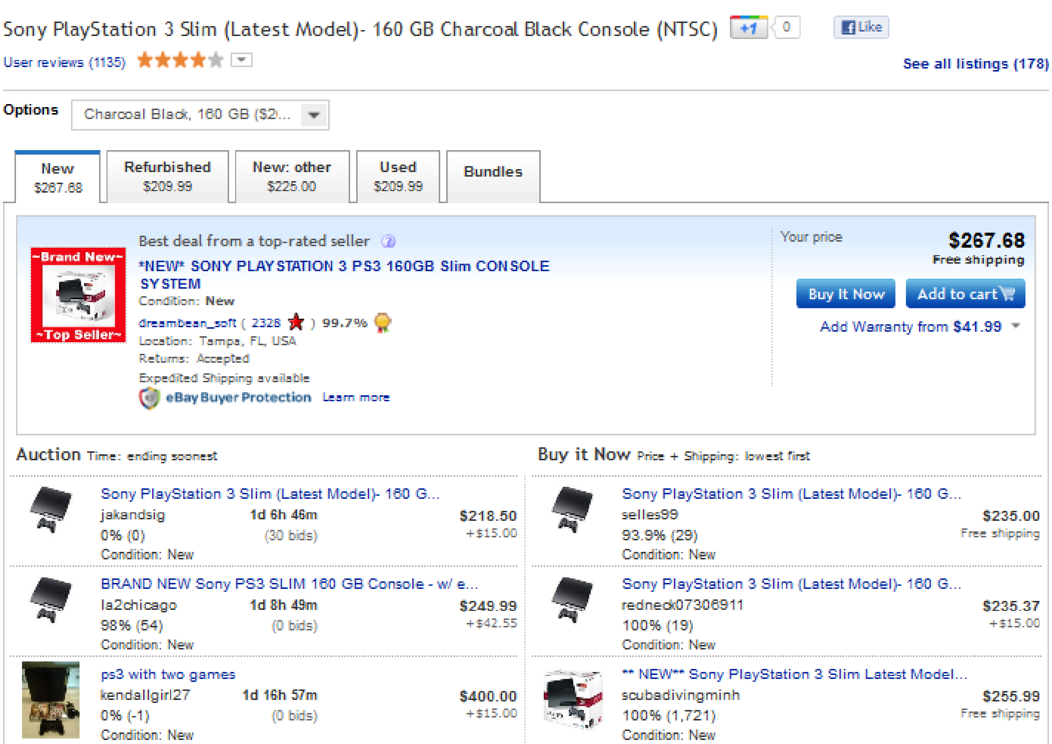


Figure shows the change in eBay's presentation of search results. The top panel shows eBay's Best Match results. The bottom panel shows a product page, with listings ordered by sales format and price.

# Figure 5: Change in Size of Consideration Set

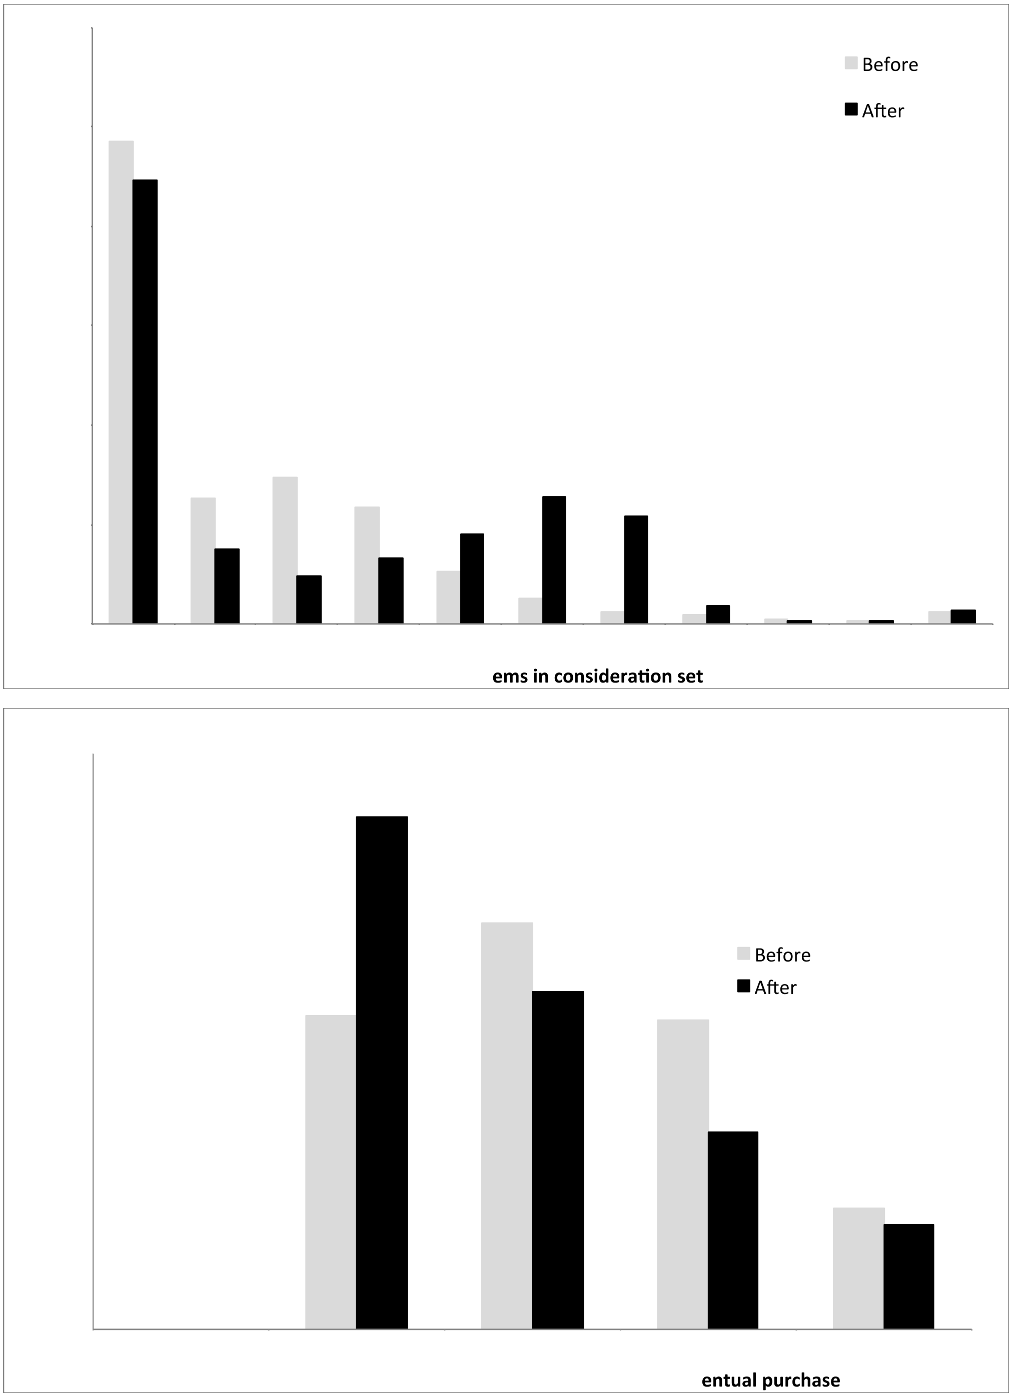


Figure shows changes in browsing experience between the Before (4/6/11-5/18/11) and After (8/1/11-9/20/11) periods. Top panel shows distributions of the size of the targeted consideration set, $L^J$ - that is, the number of targeted items shown on the search results page (the default in the “Before” period) or the product page (the default in the “After” period) – for Halo Reach listings. For users that visited multiple pages, the consideration set includes all listings on the pages. Bottom panel plots the distribution of clicks per search session prior to eventual purchase of a targeted (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 6: A/B Experiment

## Panel A: Effect on Price (%) by Share of Most Common Title

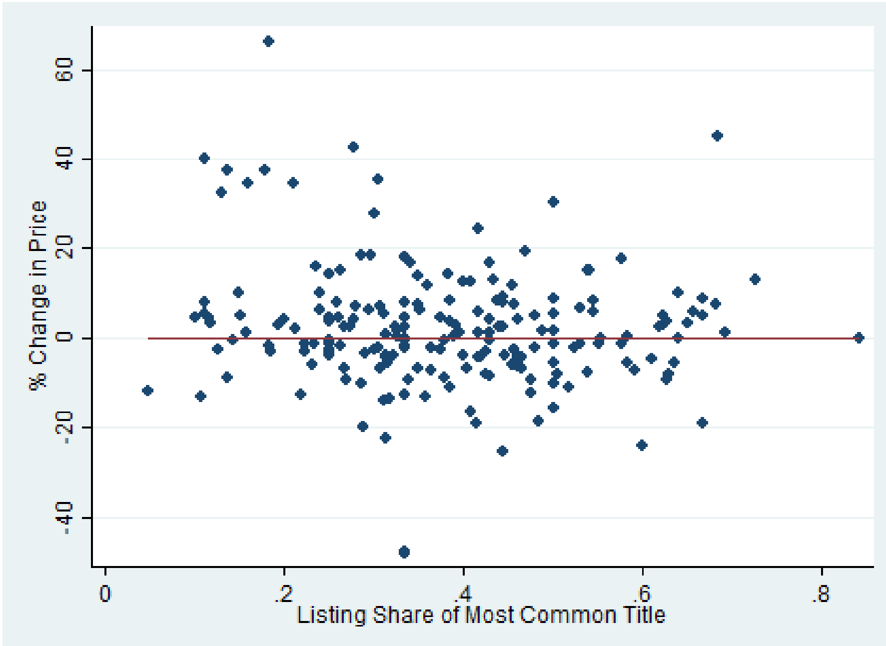


## Panel B: Effect on Quantity Sold (%) by Share of Most Common Title

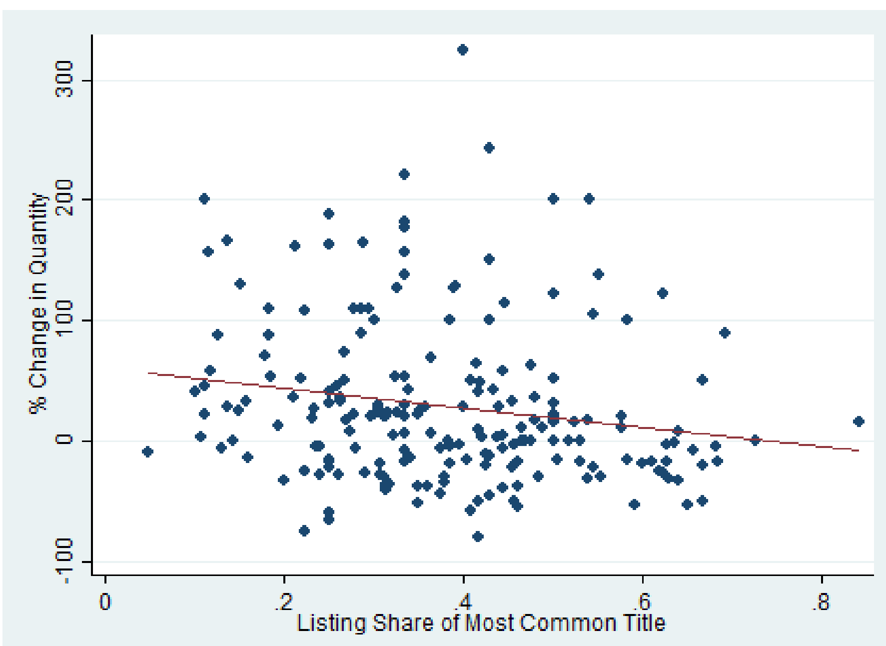


Figure shows results of the A/B experiment on transacted prices (Panel A) and transacted quantities (Panel B). Each point is an eBay product. The sample is restricted to products with at least 1,000 visits to its product page and at least 20 total purchases in the experiment. The y-axis is the percentage change in prices/quantities comparing users given the Best Match default to users given the Product Pages default. The x-axis is the share of listings of the product that use the most common (i.e., suggested) listing title.

# Figure 7: Implied Cost Distribution

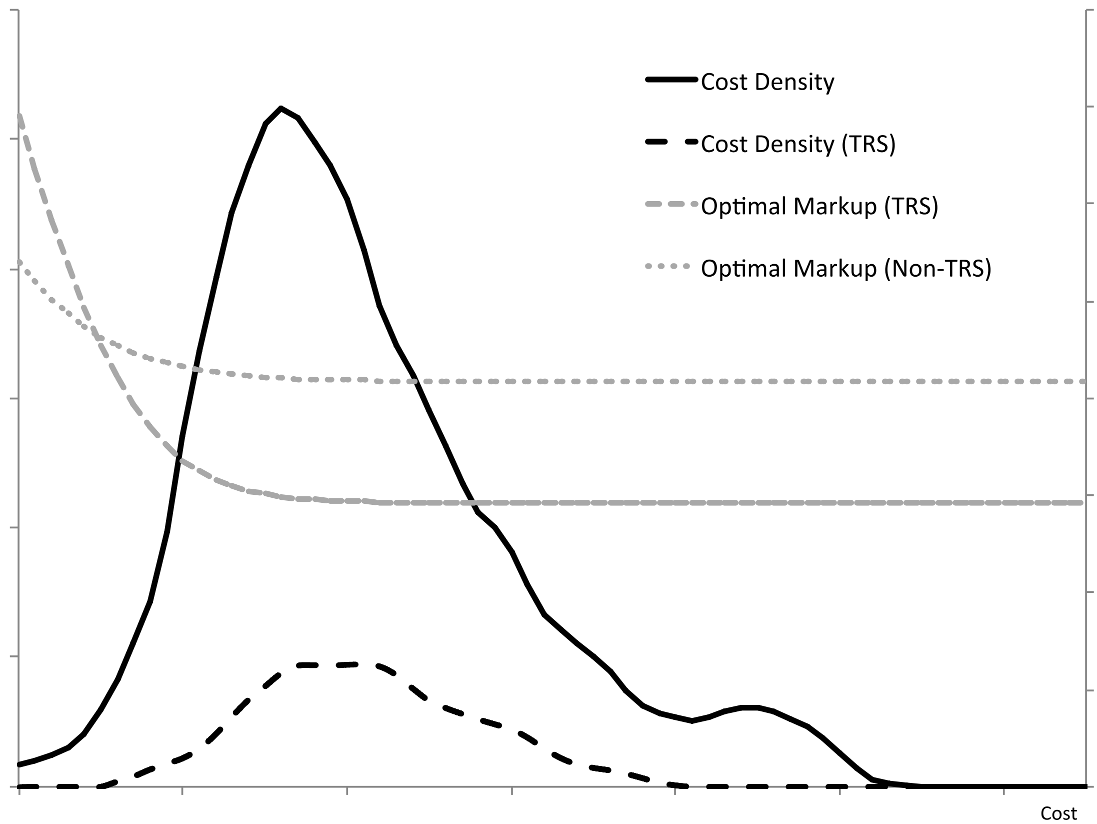


Figure shows the distribution of seller costs imputed from the observed prices and the sellers' first order condition. This cost distribution is assumed to remain the same after the platform re-design, and is held fixed in the counterfactual exercises. The dashed black line shows the cost distribution 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 8: Estimated Demand Curves

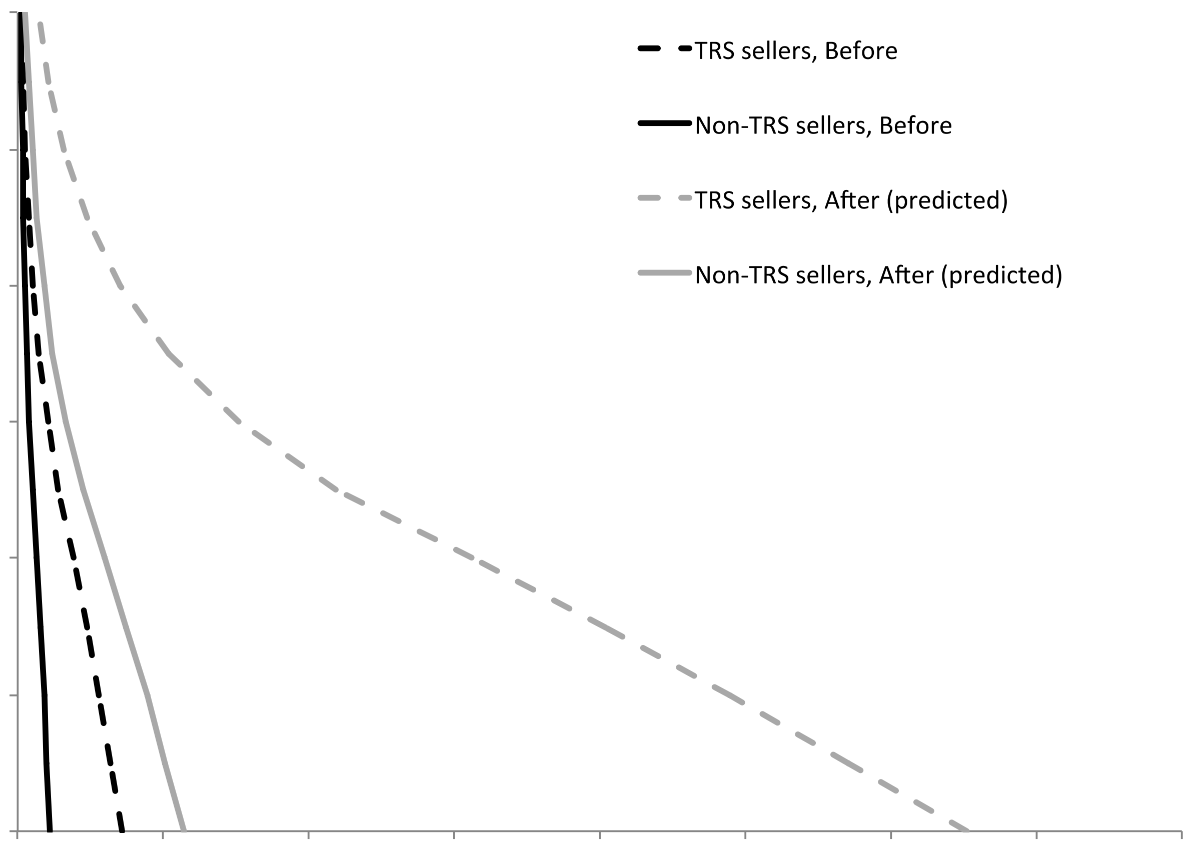


Figure plots demand curves based on our model estimates. The x-axis is the per-search probability of being transacted, which is the probability of appearing in the consideration set multiplied by the probability of being transacted conditional on being in the consideration set.

# Figure 9: Observed and Predicted Price Distributions

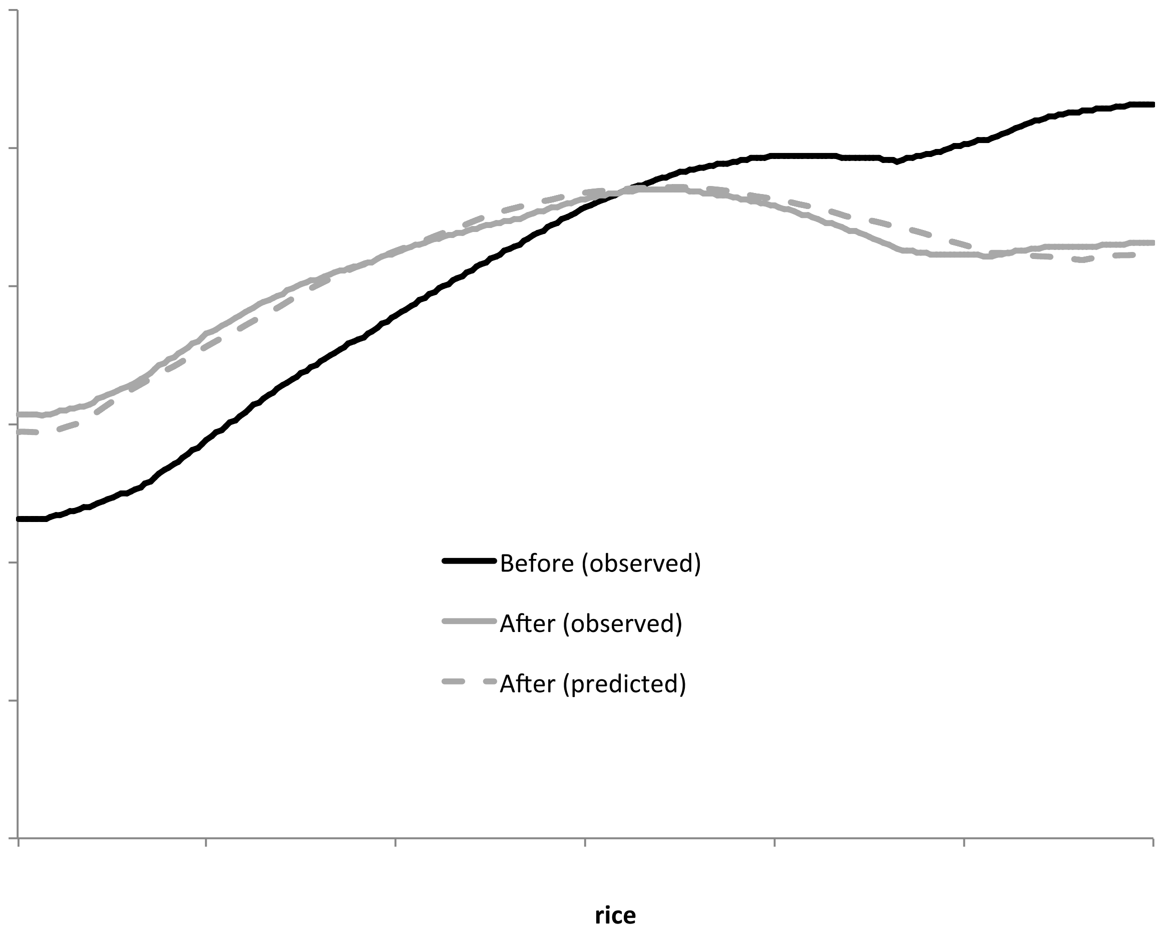


Figure shows distributions of posted prices from the Before (4/6/11-5/18/11) and After (8/1/11-9/20/11) periods, and the predicted price distribution for the after period based on our estimated model. Note that the model is estimated using only before data (except for the use of the after data to estimate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

# Figure 10: Non-Parametric Plot of Listing Demand

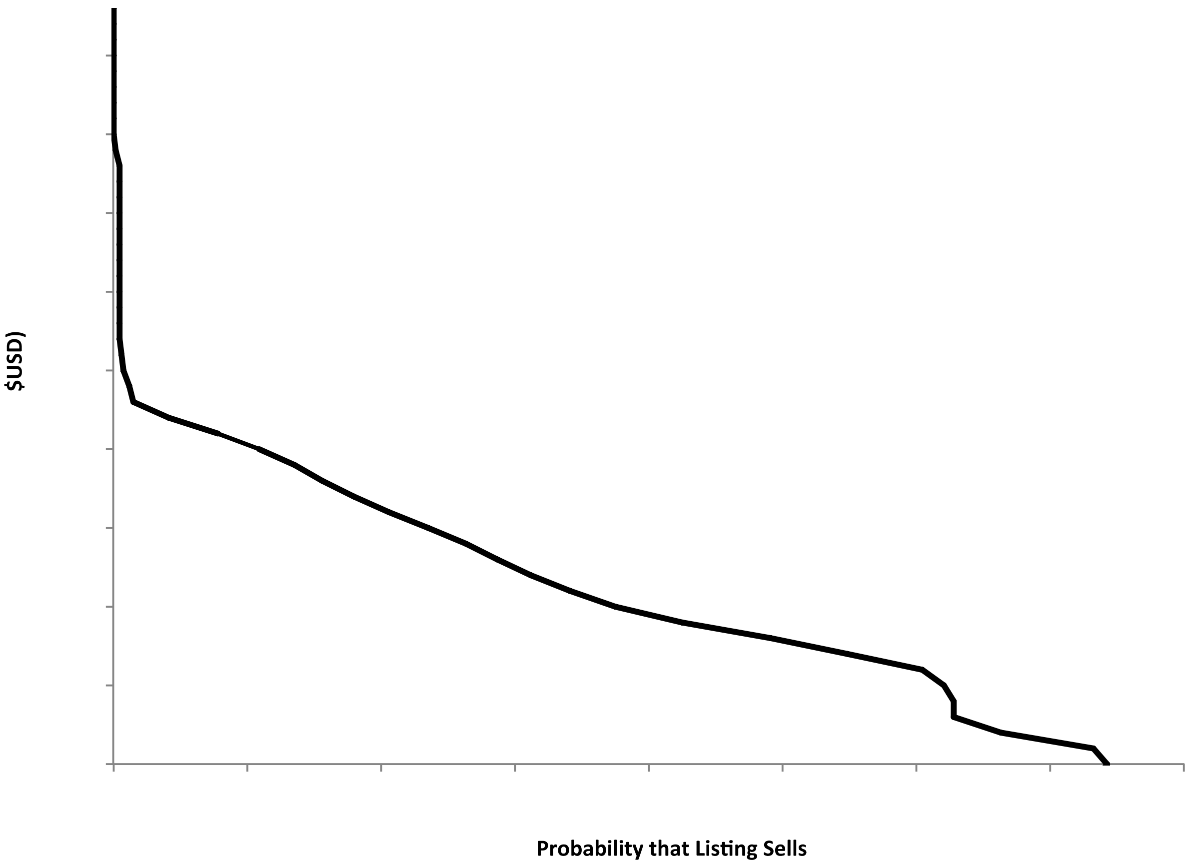


Figure shows the probability of sale for listings in the before period, estimated using a local polynomial regression plotted against listing price. The sample size is N=270 listings.

# 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 |
| | ***Change*** | -2.32 | -6.90 | -9.07 | 2.38 | -8.61 | -4.33 |
| | | (0.98) | (1.21) | (1.76) | (0.84) | (0.61) | (0.62) |
| **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 |
| | ***Change*** | -$99.57 | -$247.50 | -$4.86 | -$4.92 | -$0.45 | -$181.50 |
| | | ($11.47) | ($13.14) | ($2.61) | ($5.32) | ($0.84) | ($12.04) |
| **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 |
| | ***Change*** | -$8.55 | -$182.46 | -$8.14 | $34.53 | -$3.53 | -$121.90 |
| | | ($11.26) | ($10.78) | ($0.59) | ($3.47) | ($0.40) | ($6.83) |
| **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 |
| | ***Change*** | 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** | 39.12% | 78.93% | 45.27% | 42.62% | 44.42% | 40.52% |
| | ***Change*** | -4.75% | 8.00% | 17.88% | -1.31% | 17.31% | 3.90% |
| | | (1.44%) | (1.93%) | (2.18%) | (1.96%) | (1.22%) | (1.20%) |

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. We note that we use eBay's classification of listings to product identifiers. In some cases (most commonly in the context of cell phones, accessories get classified as the product itself, leading to what may appear as large price dispersion, which more likely reflects product misclassification).

# 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% |
| Mean Quality | 0.037 | |
| Median Quality | 0.002 | |
| Search Data | | |
| Number of "Search Results Page" Searches | 9,409 | 5,344 |
| Number of Product Page Searches | 18 | 6,715 |
| Total Number of Searches | 9,427 | 12,059 |
| Mean Transacted Price (+Shipping) | $34.56 | $33.30 |
| Median Transacted Price (+Shipping) | $34.99 | $34.00 |
| Standard Deviation of Transacted Price (+Shipping) | $2.59 | $3.89 |
| Mean Quality | 0.111 | |
| Median Quality | 0.021 | |
| Number of Halo Reach Fixed Price New Transactions | 97 | 148 |
| Number of Other Transactions | 185 | 317 |

The first panel uses listing-level data. The second panel uses search-level data. Targeted 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. "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 | 21 | 28 |
| Average Number of TRS Listings on the Site | 3 | 10 |
| Prob. of a Single-Targeted Item Consideration Set | 0.13 | 0.08 |
| Estimated Gamma | | 0.80 (0.14) |
| Median Prob. of Appearing in a Search | | 0.08 |
| Prob. of Appearing if Lower Price by 10% | | 0.11 |
| **Demand** | | |
| Constant (Halo Reach New Fixed Price) | 3.72 (1.07) | |
| Price | -0.24 (0.03) | |
| Top-Rated Seller (TRS) | 4.13 (2.75) | |
| Price*TRS | -0.10 (0.08) | |
| Quality (0 to 1) | 0.64 (0.36) | |
| Constant (Other Listings) | -8.37 (0.41) | |
| Size of Epsilon (Other Listings) | 1.70 (0.14) | |
| **Implied Price Elasticities** | | |
| Average Own-Price Elasticity | -10.64 | -13.53 |
| Average Own-Price TRS Elasticity (TRS) | -14.16 | -17.09 |
| Average Own-Price TRS Elasticity (Non-TRS) | -9.95 | -12.84 |
| **Supply** | | |
| Median Price - Cost (TRS) | $2.94 | $2.46 |
| Median Margin (% of P) (TRS) | 0.08 | 0.07 |
| Median Price - Cost (Non-TRS) | $4.20 | $3.23 |
| Median Margin (% of P) (Non-TRS) | 0.11 | 0.09 |
| **Purchase Rates** | | |
| (Halo Reach) Observed | 1.03% | 1.23% |
| (Halo Reach) Predicted | 0.93% | 1.47% |
| (Other) Observed | 1.96% | 2.63% |
| (Other) Predicted | 1.27% | 1.85% |
| (Buy Box, Conditional on HR Purchase) Observed | | 64% |
| (Buy Box, Conditional on HR Purchase) Predicted | | 65% |

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 Halo Reach and Other purchase rates are defined as the shares of relevant search queries that end up transacting in a targeted (new, fixed price) Halo Reach listing or other listing, respectively. The Buy Box purchase rates are the percentage of Halo Reach purchases that came from the Buy Box listing.

# 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 | $2.94 | $4.20 | 8% | 11% | 0.9% |
| Larger Consideration Set | $2.94 | $4.20 | 8% | 11% | 1.0% |
| Additional Sellers | $2.94 | $4.20 | 8% | 11% | 1.1% |
| Buy Box | $2.94 | $4.20 | 8% | 11% | 1.0% |
| Increase in Gamma | $2.45 | $3.29 | 7% | 9% | 1.1% |
| Predicted After | $2.46 | $3.23 | 7% | 9% | 1.5% |
| Implementing the Platform Change without Price Changes | | | | | |
| Before Prices, After Design | $2.94 | $4.20 | 8% | 11% | 1.3% |

The top and bottom rows of the “Implementing the Platform Change” panel report the margins and purchase rate of targeted listings 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, adding additional sellers, introducing a Buy Box that samples a TRS listing, and increasing price-dependence in the search. The “Implementing the Platform Change without Price Changes” panel estimates the effect of the platform change on purchase rates while keeping the prices at the before levels.

# Table 5: The Impact of Search Frictions

## Panel A: Using Estimated Quality

| Using Estimated Quality | | Search Design: Quality Rank ("Before") | Price Rank ("After") | Demand Weights Rank | No Frictions (all items visible) |
|---|---|---|---|---|---|
| Limited Seller Differentiation | Median Markup | $2.50 | $2.23 | $2.23 | $0.83 |
| | Mean Transacted Price | $33.37 | $32.03 | $32.04 | $24.61 |
| | Purchase Rate (HR) | 1.1% | 1.4% | 1.4% | 2.6% |
| | Purchase Rate (Other) | 2.4% | 2.2% | 2.2% | 2.1% |
| Seller Differentiation | Median Markup | $4.18 | $3.20 | $3.21 | $4.18 |
| | Mean Transacted Price | $34.70 | $33.02 | $33.05 | $33.56 |
| | Purchase Rate (HR) | 1.2% | 1.4% | 1.4% | 2.2% |
| | Purchase Rate (Other) | 2.1% | 2.0% | 2.0% | 1.7% |

## Panel B: Using Quality with 30x the Dispersion of Estimated Quality

| Scaling up Quality Differentiation x30 | | Search Design: Quality Rank ("Before") | Price Rank ("After") | Demand Weights Rank | No Frictions (all items visible) |
|---|---|---|---|---|---|
| Limited Seller Differentiation | Median Markup | $3.56 | $2.82 | $3.50 | $1.34 |
| | Mean Transacted Price | $37.84 | $36.33 | $37.79 | $40.61 |
| | Purchase Rate (HR) | 2.7% | 2.6% | 2.7% | 9.8% |
| | Purchase Rate (Other) | 2.3% | 2.5% | 2.2% | 1.6% |
| Seller Differentiation | Median Markup | $4.18 | $3.21 | $4.17 | $4.17 |
| | Mean Transacted Price | $38.24 | $36.55 | $38.20 | $40.26 |
| | Purchase Rate (HR) | 2.7% | 2.6% | 2.8% | 7.0% |
| | Purchase Rate (Other) | 2.2% | 2.4% | 2.1% | 1.6% |

Panel A uses the estimated quality distribution from the Halo Reach product used in estimation. Panel B considers a hypothetical product with 30 times the dispersion in quality of Halo Reach. The labels "Seller Differentiation" and “Limited Seller Differentiation" refer to whether we include a seller-specific logit error for targeted listings. The version with differentiation keeps the error, while the “Limited Differentiation” specification assumes a nested logit model in which all “targeted listings” are in the same nest and the nested logit σ parameter is set to 0.8. Each column refers to a different platform design: Quality Rank (the “Before” regime), Price Rank (the “After” regime), a counterfactual regime with Demand Weights Rank, and a counterfactual regime in which consumers are shown the entire set of targeted listings available on the platform.

# Table 6: Explanation of High Seller Costs

| | Univariate Regressions<br>Dep Var: Dummy for Seller's Max Cost > $40 | |
|---|---|---|
| Top-Rated Seller Dummy | 0.041 | (0.093) |
| Years on eBay (Truncated at 5) | -0.004 | (0.041) |
| % Listings Fixed Price | -0.221 | (0.081) |
| % Listings Free Shipping | -0.094 | (0.084) |
| Log(Total Halo Reach Listings) | 0.066 | (0.023) |
| Log(Total Videogame Listings) | 0.014 | (0.011) |
| Log(Total eBay Listings) | 0.015 | (0.010) |
| Log(Q Available in All Listings) | 0.017 | (0.009) |
| Log(Q Sold in All Listings) | 0.006 | (0.011) |
| Log(Num Categories Listing in) | 0.024 | (0.025) |
| Any Halo Reach Price Change | 0.067 | (0.081) |

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.

# Appendix: not for publication

# Appendix A: Complete Search and Platform Model

In this appendix, we develop and estimate a more complete model than the version presented in the main text. This more complete version more explicitly models the sequential search process where the user specifies a search strategy as a function of her preferences, decides which listings to click on to acquire additional information, and decides which of the clicked items to purchase, if any.[1] We conclude this appendix by discussing how the model and its estimates relate to the baseline model we use in the main text, and why this simplified approach should capture the most important aspects of the platform design in our empirical context.

## A.1 Setting

**The role of the eBay platform.** When a user $i$ makes a search on eBay, she receives search results that typically include multiple, often-related products. As in the main text, we focus on searches for Halo Reach, but consider all possible search results We group listings into three product types, indexed by $k$: listings unrelated to Halo Reach ($k = 1$), Halo Reach accessories or used games ($k = 2$), and new, fixed price listings for the Halo Reach (HR) video game ($k = 3$). As usual, we also index by $k = 0$ the outside option of not purchasing a fixed price listing.[2] We refer to product type $k = 3$ as the targeted product. Let $L_i = (L_{i0}, L_{i1}, L_{i2}, L_{i3})'$ be a vector of the number of listings of each product type considered by user $i$. We normalize $L_{i0} = 1\ \forall i$. In the next section we specify a nested demand structure that allows for correlated preferences across these product types.

The platform offers user $i$ two search options, indexed by $s$: a "Search Results Page" ($s = 0$) or a "Product Page" ($s = 1$). The options correspond to eBay's most common ways of presenting search results and we borrow eBay's labels. The search options differ along two dimensions. First, user $i$ draws the size of consideration set by product type, $L_i$, from a distribution $F^L_s$ which may differ by search option $s$. Loosely, the "Product Page" will offer more targeted products while the "Search Results Page" will offer more listings from the other product types. Second, in filling the $L_{i3}$ positions on the page, the platform samples from all listings of the targeted product that are available at the time of the search. The rules for sampling may differ across the search options.

User $i$ may choose which search option to pursue, depending on her preferences across product types. The platform sets a default search option and a cost to deviating from the default option. As we have maintained throughout the paper, the platform treats all users identically from an ex ante perspective. Users, however, may experience different search processes ex post due to stochastic draws from a common distribution of listings or because users select different search strategies.

The platform thus affects user $i$'s search process in three ways: (i) by choosing the default search option and the cost to deviating; (ii) by choosing the distribution of consideration set

[1] Kim, Albuquerque, and Bronnenberg (2010), Chen and Yao (forthcoming), and Ursu (2016) estimate sequential search models that share many similar features.

[2] As in the simple model, auction listings are bundled with not purchasing any listing as the outside option.

size across different product types; and (iii) by choosing how the targeted product listings are sampled into the consideration set.

**The platform redesign.** Before the redesign, the "Search Results Page" ($s = 0$) was the default search option and the cost to deviation was $d_0$. After the redesign, the "Product Page" ($s = 1$) blame the default search option and the cost to deviation was $d_1$. The size of $i$'s consideration set, $L_i$, is has changed as well, and we assume that it is drawn from a distribution that is specific to the period – before or after – and the search option.

For the "Search Results Page," the targeted product listings are sampled according to the same process in both periods. The process is identical to how we specify the sampling process in the before period in our model in the main text: let $J_t$ be the set of targeted product listings that are active at time $t$, and the we assume that the platform samples $L_{i3}$ listings from $J_t$ without replacement, where each listing $j$ has sampling weight $\omega_j$. This weight may correlate with price but it does not change if a seller decides to change the price of a given listing.

For the "Product Page," we specify separate sampling processes in line with the empirical setting. In the before period, listings are sampled according to the same process as for the "Search Results Page." In the after period, the "Product Page" is now sorted by price and includes a buy box. The platform places the lowest-priced listings on the page and reserves one spot for the lowest-priced listing from a top-rated seller. Specifically, let $J_t^{TRS}$ be the targeted product listings from top-rated sellers that are active at time $t$. The platform identifies the lowest-priced listing in this set and places it in the Buy Box. Then the remaining $L_{i3} - 1$ spots on the page are filled with the lowest $L_{i3} - 1$ remaining prices in the set $J_t$.

## A.2 Demand

**Utility Specification.** User $i$'s utility from purchasing listing $j$ of type $k$ is given by:

$$u_{ij} = v_{ik} + w_j + \epsilon_{ij}, \tag{16}$$

where $v_{ik}$ is a product type component that may vary across users, $w_j$ is a listing component, and $\epsilon_{ij}$ is an idiosyncratic preference distributed i.i.d. from a standard normal distribution $N(0, 1)$. We normalize $v_{i0} = 0$, and parameterize

$$v_{ik} = \theta_k + \textstyle\sum_{l=1}^{k} \lambda_{ik} \text{ for } k = 1, 2, 3, \tag{17}$$

with $\lambda_{in} \sim N(0, \sigma_n^2)$ This parameterization imposes a particular correlation structure in the random effects across product types, which seems (to us) natural; for example, it makes product type 2 and 3 closer substitutes than product types 1 and 3.

For the listing component, we set $w_j = 0$ for all products except those that belong to the targeted product type ($k = 3$). For targeted product listings, we parameterize $w_j$ as

$$w_j = \alpha_1 p_j + \alpha_2 TRS_j + \alpha_3 p_j TRS_j, \tag{18}$$

where $p_j$ is the item's posted price and $TRS_j$ is an indicator equal to one if listing $j$ is listed by a top-rated seller.[3]

**Search Option, Information, and Clicks.** User $i$ chooses a search option and the platform gives her a consideration set with size drawn from the distribution corresponding to the chosen search option. The user observes her entire consideration set with one exception: she observes her idiosyncratic valuation of a targeted listing $j$, $\epsilon_{ij}$, with noise. She may then incur a costly click to learn $\epsilon_{ij}$ without noise. Conditional on the listings she observes without noise, the user chooses the one that provides her with the highest level of utility (including the outside option). We start by describing the click process and then work backward to the choice of search options.

When the user is given her consideration set by the platform, she only observes a signal of $\epsilon_{ij}$ for listings of the targeted product ($k = 3$). For other listings, $\epsilon_{ij}$ is observed without noise. Denote the signal by $\eta_{ij}$ and the correlation between $\epsilon_{ij}$ and $\eta_{ij}$ by $\rho$. To observe $\epsilon_{ij}$, the user must click on the listing at cost $\kappa$. Clicking has a dual role: it reveals $\epsilon_{ij}$ and is also a necessary action before purchasing listing $j$. Making clicking a necessary action before purchase parallels the actual process of buying an item on eBay. We treat the Buy Box separately from other listings to account for its more detailed presentation in search results and the fact that the user can proceed to purchase the Buy Box listing immediately from the search results. We model this difference by making the click cost to the Buy Box equal zero. Finally, because some of the clicks in our data are difficult to rationalize with the above model, we also introduce click noise by assuming that with probability $\psi$ the user makes a clicking mistake – clicking when a click was not intended or not clicking when a click was intended. Including the click noise is not strictly necessary, as extreme draws of $\epsilon_{ij}$ can rationalize any click. But we still include the noise because we think some eBay users are likely to click unintentionally and because it speeds up our computation considerably by requiring fewer draws for our maximum simulated likelihood estimation.

The user employs an optimal sequential click strategy for up to three costly, intentional clicks. This restriction on the maximum number of clicks is made for computational tractability and is consistent with the data, where very few users (0.12%) click on more than 3 targeted listings during a single browsing session. Let $u_{ij}$ be the utility user $i$ derives from purchasing listing $j$. Let $\hat{u_{ij}}$ be the utility user $i$ expects to derive from owning listing $j$ prior to clicking on it. For listings of non-targeted products, $u_{ij} = \hat{u_{ij}}$, and let $\bar{u}_i = \max_{k<2} u_{ij}$ be the utility of the most preferred non-targeted listing. The user employs the following optimal sequential click strategy. First, she finds the listing of the targeted product with the highest expected utility. Without loss of generality, denote this listing by $j = 1$, so the expected utility is $\hat{u_{i1}}$. The expected utility from clicking on item 1 is

$$E(\text{u if click on 1}) = E(u_{i1}|u_{i1} > \bar{u}_i)Prob(u_{i1} > \bar{u}_i) + \bar{u}_i Prob(u_{i1} \leq \bar{u}_i) - \kappa \qquad (19)$$

The expected utility from not clicking is $\bar{u}_i$. The user clicks if $E(\text{u if click on 1}) > \bar{u}_i$, or equivalently if the expected gain from the click $(E(u_{i1}|u_{i1} > \bar{u}_i) - \bar{u}_i)\, Prob(u_{i1} > \bar{u}_i)$ are

[3] We note that $w_j$ differs from our simple model as it does not include a listing's quality. This omission is a necessary shortcoming of the more complete model as it requires data from both the before and after periods for identification. Our measure of quality is derived from "Best Match" search results, but the platform redesign lowered the number of "Best Match" searches to the point where we are unable to estimate quality consistently in the after period.

greater than the click cost $\kappa$. If the click is made, the user observes $u_{i1}$ and listing 1 is eligible to be purchased. We use the normality of the errors for a closed-form representation of $E(u_{i1}|u_{i1} > \bar{u}_i)$, which depends only on the parameters and $\bar{u}_i - \hat{u_{i1}}$. The user then considers whether to make a second click, and the process is similar. We redefine $\bar{u}_i = \max\{\bar{u}_i, u_{i1}\}$, and the user finds the listing of the targeted product with the highest expected utility, excluding all listings that have already received clicks. The user then follows the same rule as above in deciding whether to click again. The user repeats this process for the third click.

With the click process described, we return to the user's choice of search option: Search Result Pages ($s = 0$) or Product Page ($s = 1$). Let $I_{is}$ be the expected continuation utility from choosing search option $s$, which excludes the cost to deviating from the default. Note that the continuation utility will depend on $i$'s idiosyncratic preferences across product nests, $\lambda_{ik}$, such that users with strong preferences for the targeted product will systematically prefer the search option that on average includes more targeted listings in the consideration set (empirically, $s = 1$).

Let $\mu_i$ be an i.i.d. unobserved (to the econometrician) preference for search option 0. We parameterize $\mu_i \sim N(0, 1)$. Then in the before period, $i$ chooses $s = 0$ if and only if $I_{i0} - I_{i1} + d_0 > \mu_i$. In the after period, $i$ chooses $s = 0$ if and only if $I_{i0} - I_{i1} - d_1 > \mu_i$.

## A.3 Supply

Our model of supply mirrors the supply side of the model in the main text, except that the sellers now faces consumer demand determined by the more complete model search process. Let $G$ be the cumulative density function over user types $(\lambda_{i1}, \lambda_{i2}, \lambda_{i3}, \mu_i)$. Then total demand for listing $j$ is $D_j(p_j, TRS_j, w_j) = \int_i D_{ij}(p_j, TRS_j, \omega_j) dG_i$, where the arguments are the listing's price, whether the seller is a top-rated seller, and the listing's sampling weight ($\omega_j$) and $D_{ij}(p_j, TRS_j, \omega_j)$ is $i$'s expected demand for listing $j$, integrated over the distribution of competitor listings and the distributions of consideration sets.

For each targeted listing $j$, we model seller pricing using a standard Nash in prices assumption. Each seller sets its price to solve

$$\max_{p_j} (p_j - c_j) D_j(p_j, TRS_j, \omega_j). \tag{20}$$

We recover each seller's marginal cost $c_j$ by inverting the first-order condition:

$$c_j = p_j + \left(\frac{\partial D_j}{\partial p_j}\right)^{-1} D_j(p_j). \tag{21}$$

## A.4 Estimation and Identification

**Estimation.** We start by estimating the process that forms consideration sets directly from data on search results. We estimate the distribution of consideration set sizes, $F_s^L$, separately for each search option and period as the empirical distribution. To determine which targeted listings enter the consideration set, the platform relies on a set of sampling weights, $\omega_j$ (except for the "Product Page" option after the redesign, for which the platform samples according to price). We estimate the weights $\omega_j$ using the same procedure as in the simple model, which relies only on search results data. We estimate separate weights for "Search Results Page" searches ($s = 0$) before and after the redesign. For the "Product Page" searches ($s = 1$) before the redesign, we have insufficient observations so we use the weights from the $s = 0$ searches.

We then estimate the demand parameters plus the cost to deviating from the platform's default search option. There are three user decisions within our model: (i) the choice of search strategy; (ii) which listings in the consideration to clicks on; and (iii) which listings in the consideration set to purchase. We observe each of these decisions in our data. Let $K_i$ be $i$'s observed consideration set. Then for user $i$ we label the observed strategy with $s_i$, whether $i$ clicked on listing $j \in K_i$ with $C_{ij}$, and whether $i$ purchased listing $j \in K_i$ with $Y_{ij}$.

For a given set of demand parameters, we use our model to generate the predicted probabilities of these decisions. Let $P_i\{s = 0\} = \Phi(I_{i0} - I_{i1} + d_0)$ be the model-predicted probability that $i$ chooses search option 0, $Q^C_{ij}$ be the model-predicted probability that $i$ clicks on $j$, and $Q^Y_{ij}$ be the model-predicted probability that $i$ chooses $j$. We estimate $I_{is}$, $i$'s expected continuation utility from choosing search option $s$, by simulation. We draw a series of consideration sets for each search option and predict $i$'s click and purchase decisions. Let $j_i$ be the listing (or outside option) that $i$ bought. Then the likelihood function is

$$\mathcal{L} = \prod_i P_i\{s = 0\}^{1-s_i}(1 - P_i\{s = 0\})^{s_i}(Q^Y_{ij_i})^{Y_{ij_i}} \prod_j (Q^C_{ij})^{C_{ij}}(1 - Q^C_{ij})^{1-C_{ij}} \tag{22}$$

The 14 parameters to estimate are the product type fixed effects $(\theta_1, \theta_2, \theta_3)$, the variances of the random effects $(\sigma^2_1, \sigma^2_2, \sigma^2_3)$, the coefficients on targeted listings' characteristics $(\alpha_1, \alpha_2, \alpha_3)$, the correlation of the signal with the idiosyncratic listing-specific preference $(\rho)$, the cost of an intentional click $(\kappa)$, the probability of a click mistake $(\psi)$, and the costs to deviating from the default search strategy $(d_0, d_1)$. We estimate using maximum simulated likelihood where we simulate from the distributions of the random effects $(\lambda_{ik})$, the signal $(\eta_{ij})$, and the listing-specific preference term $(\epsilon_{ij})$.

Once we have the demand estimates, we estimate a seller's expected demand by simulating sets of competitors (both set size and members of the set), consideration sets, and users. We draw the set of competitors from a smoothed empirical distribution of prices. We then recover seller costs by inverting the estimated first-order conditions.

**Identification.** The seller costs are identified from an optimal price setting assumption. We thus focus on identification of the demand and platform parameters. The product type fixed effects $(\theta_1, \theta_2, \theta_3)$ are identified by the relative purchase probabilities of listings of different product types. The coefficients on the targeted product characteristics $(\alpha_1, \alpha_2, \alpha_3)$ are identified by how purchase probabilities vary across new, fixed price Halo Reach listings depending on their prices and whether they are sold by TRS. The variances of the random effects $(\sigma^2_1, \sigma^2_2, \sigma^2_3)$ are identified by how purchase probabilities of listings in different nests vary as the number of listings in each nest changes in different searches. Users with a strong preference for listings in a specific nest may select into different search strategies, but this selection depends on an expectation about which search results she will receive, not the consideration set that is actually realized. We therefore have residual exogenous variation in the number of listings from each nest, even conditional on selection into a search strategy, and this variation allows us to identify the variances of the random effects.

Identifying parameters related to the click process is more complicated. Because we specified a sequential search process where the choice of making intentional clicks depends on the same characteristics that affect utility from purchase, differences between observed click patterns and observed purchase patterns conditional on click are rationalized as click mistakes. For instance, in the data we see no purchases of very high-priced listings even though they occasionally appear in search results. Thus, the extent to which users click on these listings identifies the probability of making a click mistake.

The number of clicks we see is informative about the click cost and the correlation of the signal with the true idiosyncratic preference. Few clicks could be indicative of a high click cost or that clicking reveals little new information about a listing. To separate these components, we rely on two sources of variation: the order of clicks and the introduction of the Buy Box in the after period. If we observe that users' first clicks are more likely to go to listings based on their observable characteristics (price, TRS) than the logit probabilities would imply, then we would infer that the signal is not very informative about the true idiosyncratic preference.

The Buy Box does not require a costly click and already has all of its information displayed. Conditional on the user making a purchase, if we see that the lowest-priced TRS listing on the page has a higher market share in the after period than the before period, our model attributes this to not having to incur a click cost or better information about the listing. But even in the before period when users only see a signal of the lowest-priced TRS listing's value, in expectation this signal is the same as the true value. Thus, conditional on making a purchase, the lowest-priced TRS listing should not capture more market share in the after period if the value to the click is learning about the listing's value. Therefore, the observed change in market share for this listing type identifies the click cost. We note that because having the Buy Box is so helpful for identifying this click process, our ability to identify the complete model using only data from the before period, as we do for the simpler model in the main text, is limited.

Finally, the costs to deviating from the default search strategies relate to the fraction of users who end up on the default search page, combined with the exogenous change to which page is the default. For instance, in the before period nearly all users see "Search Results Pages," implying a very large cost to deviating. In the after period, users are split more evenly across the two page types; hence, we estimate a smaller cost to deviating from the default.

## A.5 Estimates

We present the estimates of the demand and platform parameters from the more complex model.

| Parameter | Description | Estimate |
|---|---|---|
| $\theta_1$ | product 1 FE | -3.7807 |
| $\theta_2$ | product 2 FE | -3.2929 |
| $\theta_3$ | product 3 FE | 4.2683 |
| $\sigma_1^2$ | nest 1 RE Variance | 0.9329 |
| $\sigma_2^2$ | nest 2 RE Variance | 0.0164 |
| $\sigma_3^2$ | nest 3 RE Variance | 1.6000 |
| $\alpha_1$ | coefficient on price | -0.2000 |
| $\alpha_2$ | coefficient on TRS | 2.0764 |
| $\alpha_3$ | coefficient on price*TRS | -0.0442 |
| $\rho$ | correlation(signal,$\epsilon$) | 0.9998 |
| $\kappa$ | click cost | 1.5454 |
| $p$ | probability of a click mistake | 0.0586 |
| $d_0$ | search strategy deviation cost, before | 3.2461 |
| $d_1$ | search strategy deviation cost, after | -0.1475 |

As in the simple model, we find that price and whether the listing is from a top-rated seller are important in predicting which targeted listing a user purchases. Compared to $35 listing from a non-TRS seller, a TRS seller could price at $37.17 and be purchased at the same rate. The large price coefficient implies very elastic demand. The implied median price elasticities are -8 and -11 for non-TRS and TRS sellers.

The estimates also imply considerable preference heterogeneity across the different product types. The estimated standard deviations of the random effects on all listings and the targeted listings are 0.93 and 1.60, respectively (or $4.65 and $8.00 when divided by the price coefficient). The random effect on Halo Reach-related listings is estimated to be very small.

As for the click process, we estimate a very large click cost, equivalent to $7.73, but that clicking reveals essentially no new information about the listing. The estimated correlation between the signal and the idiosyncratic valuation is 1. The large click cost is driven by a small number of total clicks and the Buy Box's high market share. The lowest-priced TRS listing captured 27% of purchases of targeted listings in the before period and 64% in the after period. The very informative signal is consistent with most (84%) users who purchased a targeted listing clicking only on the purchased item. The additional clicks can be explained with minimal noise, and indeed we estimate that click mistakes are rare, occurring for under 6% of listings. This is unsurprising as click patterns are highly predictable given the coefficients on listings characteristics. If we use our estimates of $\alpha$ to rank targeted listings according to their non-idiosyncratic valuations, 45% of users who click once clicked on the highest ranked listing and 17% clicked on the second highest ranked listing. Just 1% clicked on the lowest ranked listing (conditional on seeing at least four targeted listings).

Finally, the estimated cost to deviating from the default search strategy is very large in the before period and small in the after period. These estimates are consistent with the "Product Pages" being very hard to find prior to the platform change while both types of pages were readily accessible after the platform change. In the before period, just 18 of 9,427 searches led to "Product Pages." In the after period, users were more evenly split between search strategies in the data, as 56% reached "Product Pages."

## A.6 Discussion and Relationship to the Baseline Model

There are three main differences between the more complete model and the simpler model we use in the main text: the product type definition, the click process, and the choice of search strategy. With the above estimated parameters, we find that it reduces to a version similar to our simpler model.

The first difference between the models is that the more complete version allows for more product types than the simpler model. Product types $k = 1$ (non-Halo Reach related listings) and $k = 2$ (Halo Reach accessories) are separate to allow consumers to have different preferences for listings closer to the new Halo Reach video game than completely unrelated listings. The simpler model, on the other hand, pools these product types into one category. While the more complete model has the flexibility to allow for heterogeneous preferences across these product types, the estimates indicate that consumers' purchase patterns do not vary across these product types. We find that $\hat{\theta}_1 \approx \hat{\theta}_2$ and $\hat{\sigma}_2^2 \approx 0$, and thus the simpler model captures the relevant product type distinctions.

The more complete model also allows for consumers to click on listings to learn more information about them and as a prerequisite for purchase. Our estimates, however, find that the information acquisition motive does not empirically drive clicking patterns. We estimate that the signal, observed prior to clicking, already contains all of the information about the product: $\hat{\rho} \approx 1$. Therefore, the only reason to click on an item is because the platform requires a click before purchase. The click cost, $\kappa$, can thus be subsumed into the product 3 fixed effect, $\theta_3$, which is included in the baseline model. The exception is that the complex model specifies the Buy Box click as costless. For the simpler model to capture this, we would need to include a separate indicator for the Buy Box in the utility specification. Estimating such a parameter in the simpler model would require using purchase data from after the platform change. But as described in the text, our inability to estimate listing quality after the platform change, plus the ability to validate our model predictions using actual data, lead us to leave out a Buy Box-specific parameter. Even with this omission, we still predict a large shift in purchases to the listing in the Buy Box. In our data, the lowest-priced TRS listings accounts for 27% of purchases of the targeted product in the before period and 64% in the after period. Our simple model predicts a similar shift, from 40% to 65%.

Finally, the complex model allows consumers to select between two possible search strategies in each period while the simple model imposes a single search strategy per period. In the before period, we estimate that the cost to deviating from the default strategy, $\hat{d}_0$, is so large that nearly all users choose the "Search Results Page." The simple model thus approximates the search environment well in the before period. In the after period, we estimate a fairly small cost to deviating from the default strategy, $\hat{d}_1$. This implies that different users will select into different search strategies. Users might choose a specific search strategy due to idiosyncratic preferences over product nests ($\lambda_i$) or idiosyncratic preferences over search strategies ($\mu_i$).

Our simple model is a reduced form version that summarizes the sampling process of the two search strategies with a single sampling process, where a listing's sampling weight is summarized as a function of the reduced form $\gamma$ parameter. This reduced form representation could distort substitution patterns if it ignores that consumers with different preferences over product types ($\lambda_i$) will select into pages with different sampling processes. If the selection is instead driven by the idiosyncratic preference over search strategies ($\mu_i$), which is independent from preferences over product types, then treating consideration sets as exogenous does not yield inconsistent estimates of the preference parameters.

To evaluate whether such a distortion is large, we assess how much a consumer $i$ might gain, in expectation, from choosing a specific search strategy. The expected gain depends on consumer $i$'s preferences over the different product types but not $i$'s preference over search strategies. For example, a consumer with a strong idiosyncratic preference for the targeted product is more likely to have large gains from searching via the "Product Page."

We find that consumers choose their search strategy largely based on their idiosyncratic search strategy preference rather than their product type taste heterogeneity. We simulate 1,000 consumer types $i$, where the consumer type is a draw of the distribution of nest random effects, and estimate the expected gain from searching via "Search Results Pages" versus "Product Pages" in the after period, $\hat{I_{i0}} - \hat{I_{i1}}$. We find that the standard deviation across consumer types of these estimated gains is 0.10, which is small relative to the standard deviation of search strategy preferences, set to 1. While the product type preference heterogeneity is empirically relevant in affecting purchase decisions, both search strategies include multiple listings from all product types in most of their consideration sets, and thus the gain from using one search strategy over another is small. This lack of targeted search is perhaps not surprising based on the search patterns we see in the data. Under 5% of sessions, and about 3% of sessions ending in a purchase, included multiple search queries for Halo Reach. Users also rarely clicked beyond the first page of search results, with users averaging just 1.1 results pages per session. We thus conclude that our reduced form representation, which treats consideration sets as exogenous to the consumer, is a reasonable approximation and captures the relevant substitution patterns.

# Appendix B: Data and Estimation

In this appendix, we provide further details of how we constructed our data samples and estimated our empirical model.

## B.1 Data Samples

**Product Category Analysis.** For the product category analysis presented in Section 3.2, we gathered data from products in the five categories affected by the platform redesign in Summer, 2011. For each product as defined by eBay's catalog, we counted the number of visits to its product page from 6/27/11 – 7/2/11, the week during which the platform redesign became fully implemented. Within each product category, we chose the 10 products that had the highest number of product page visits.[4] We also kept a smaller group of products – all iPhone 4 products – as a separate subcategory for comparison.

**A/B Experiment.** For the A/B experiment results presented in Section 3.3, we col-

[4] The textbooks category had 10 products, but one of them did not have transactions in the before and after periods so it is dropped from the analysis.

lected data on all products active during the A/B experiment (6/25/12 – 8/30/12). We restrict our sample to products with at least 1,000 visits to its product page and at least 20 total purchases in the experiment. This left us with 200 different products.

**Estimation Sample.** For our empirical model (Section 4, 5, and 6), we focus on a single product, the Halo Reach video game for Xbox 360. The data for the analysis come directly from eBay and cover 4/6/11 – 5/18/11 and 8/1/11 – 9/20/11. 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 keep searches that lead to at least one click or transaction on any listing following the query. We keep all search results (listings shown to the user) derived from the user's last search query. This results in $14,753$ visits to the search results page ($9,409$ of them in the pre-period) and $6,733$ visits to the product page (18 in the pre-period).

We further drop two types of search results: auctions and listings with missing prices. Some auction listings have a Buy-It-Now price that lets the user purchase the listing at a posted price. After the first auction bid, the posted price is no longer available. We only drop the listing after its posted price is no longer available. As mentioned in the text, there is a special case when we may not observe a listing's price during a portion of its active time on the site. If the listing will subsequently have a price change but prior to the change the listing never receives a click nor is transacted, then we sometimes do not observe its price. In these cases, we drop the listing from the search results during the period when we do not observe its price.

We define the user's consideration set as all listings that eBay included on the search results or product page in the user's search. As discussed in the paper and Appendix A, we do not distinguish whether listings received clicks or by their placement on the page.

For listings that appear in users' consideration sets, we divide them into "targeted" and "non-targeted" products. Targeted products are new, fixed price (or auction, while a posted price is available) listings of the Halo Reach video game. We identify listings as the Halo Reach video game if eBay catalogues them as such. We further 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). The non-targeted products therefore include listings of used goods, listings catalogued as products other than Halo Reach, or listings catalogued as Halo Reach but whose titles indicate they are not the video game itself.

For the supply model, our sample consists of all listings classified as the targeted product.

## B.2 Estimation

**Demand.** We estimate the consumer demand parameters using maximum likelihood. For user $i$ and targeted listing $j$ in $i$'s consideration set, let $Q_{ij}^J$ be an indicator that equals 1 if $i$ purchased $j$. Let $Q_i^M$ be an indicator that equals 1 if $i$ purchased a non-targeted product. Let $D_i = 1+\exp(\delta+\lambda \ln |J_i^M|)+\sum_{k\in J_i^J} \exp\left(\alpha_0 + \alpha_1 p_k + \alpha_2 TRS_k + \alpha_3 p_k TRS_k + \alpha_4 q_k\right)$. The likelihood function is:

$$\mathcal{L} = \prod_i \left(\frac{1}{D_i}\right)^{1-Q_i^M-\sum_{k\in J_i^J} Q_{ik}^J} \left(\frac{\exp(\delta+\lambda\ln|J_i^M|)}{D_i}\right)^{Q_i^M} \prod_j \left(\exp\left(\begin{array}{c}\alpha_0+\alpha_1 p_j+\alpha_2 TRS_j+\\ \alpha_3 p_j TRS_j+\alpha_4 q_j\end{array}\right)/D_i\right)^{Q_{ij}^J} \tag{23}$$

The likelihood only depends on observables and parameters, with one exception: a listing's quality, $q_j$. We describe at the end of the next section how we recover quality.

**Platform.** From the search estimation sample, we recover the joint empirical distribution of the number of targeted and non-targeted listings in a consideration set, $L = (L^J, L^M)$. We estimate separate distributions for the before and after periods.

From the search estimation sample, we construct the empirical sampling probability, $v_j$, for each targeted listing $j$, separately for the before and after periods. From the eBay data, we calculate $v_j$ as the percentage of searches made while $j$ was active on the eBay site in which $j$ appeared in the consideration set (in the search results). For each listing $j$ we also calculate the percentage of searches made while $j$ was active on the eBay site that had consideration set size $l$: $v_j^l$.

The platform forms consideration sets by sampling $L_i^J$ products from $\mathcal{J}_i^J$, without replacement. Listings are sampled according to their heterogeneous sampling weights, $\omega_j$. This implies that the consideration set of targeted listings is drawn from a Wallenius' non-central hypergeometric distribution. The probability any given listing is drawn into the consideration set depends on the sampling weights of all competing listings. Estimating the full vector of sampling weights is computationally intractable, so we make the simplification that all competing listings are of a normalized sampling weight, 1. With this simplification, the probability that listing $j$ is drawn into a consideration set of size $l$, with $|J_i^J|-1$ competing targeted listings is:

$$a_l(\omega_j) = \binom{1}{1}\binom{|J_i^J|-1}{l-1}\int_0^1 (1-t^{\omega_j/D})(1-t^{1/D})^{l-1}dt \tag{24}$$

where $D = |J_i^J| - l$. In the before period, we set $|J_i^J| = 21$, and in the after period $|J_i^J| = 28$.

Using the model-predicted probability that listing $j$ is drawn into a consideration set of size $l$, $a_l(\omega_j)$, we can construct the model-predicted fraction of searches that listing $j$ appears in: $\sum_{l=1}^{l^{max}} a_l(\omega_j)v_j^l$. We then solve to find the sampling weight, $\omega_j$, such that the model-predicted fraction of appearances matches the data:[5]

$$v_j = \sum_{l=1}^{l^{max}} a_l(\omega_j)v_j^l. \tag{25}$$

We follow the same procedure in the before and after periods. In the after period, we

[5] There are a few listings with $v_j = 0$ or $v_j = 1$. Our model is unable to find a unique positive $\omega_j$ to rationalize the data. Therefore, for listings with $v_j = 0$, we set $\omega_j = \min_{k:0<v_k<1}\omega_k$, and for listings with $v_j = 1$, we set $\omega_j = \max_{k:0<v_k<1}\omega_k$.

project these weights onto listing prices:

$$\omega_j = \exp\left[-\gamma\left(\frac{p_j - \min_{k\in\mathcal{J}_i^J}(p_k)}{std_{k\in\mathcal{J}_i^J}(p_k)}\right) + \nu_j\right]. \tag{26}$$

We estimate $\gamma$ with the following OLS regression:

$$\ln\omega_j = \tau - \gamma\frac{p_j}{std_{k\in\mathcal{J}_i^J}(p_k)} + \nu_j \tag{27}$$

We simulate new consideration sets with the following procedure. First, we determine the $\mathcal{J}_i^J$ targeted listings that are on the simulated site for user $i$. We form a queue of listings where we sample listings from the full set of listings available in the before or after period. We sample each listing with equal probability except we duplicate multi-unit listings according to their listed quantities. Thus, a listing with two units for sale will appear twice as frequently in the queue than a single-unit listing, on average. The first 21 (before period) or 28 (after period) listings in the queue are active on the simulated site.

Second, we draw the consideration set size from the empirical distribution. Third, we fill the targeted product positions in the consideration set by sampling from the listings active on the simulated site. We sample according to heterogeneous sampling weights, $\tilde{\omega_j}$. In the before period, $\tilde{\omega_j} = \omega_j$. In the after period, $\tilde{\omega_j} = \exp\left[-\hat{\gamma}\left(\frac{p_j - \min_{k\in\mathcal{J}_i^J}(p_k)}{std_{k\in\mathcal{J}_i^J}(p_k)}\right)\right]$. In some versions of the model, we also include a Buy Box. We model the Buy Box by reserving the first position in the consideration set for a listing from a TRS seller. This seller is drawn according to the same process, but the set of competing listings is comprised only of other TRS listings.

Once the consideration set is formed for user $i$, we simulate a purchase decision. We then reconstruct the simulated site for the next user, $i'$. If user $i$ purchases one of the targeted listings, we replace that listing on the simulated site with the next one in the queue. Otherwise, if user $i$ does not purchase a targeted listing, the set of active listings on the simulated site is unchanged. Note that unpopular listings are likely to last longer on the site. We repeat this process for 100 users and then reset the site by drawing an entirely new queue. This resetting of the site accounts for the feature that some eBay listings expire without being purchased.

While unrelated to the model of the platform, the estimation of listing $j$'s quality, $q_j$, follows a similar procedure. We repeat the process of estimating $\omega_j$ except we use only searches from the before period that led to Best Match results (i.e., we exclude results from time-ending soonest searches, etc.). Let $\omega_j^{BM}$ be this estimated listing weight. We then set $q_j = \omega_j^{BM}$.

**Supply.** As detailed in the text, estimating marginal cost $c_j$ amounts to estimating the elasticity of demand ($\eta_{D_j}$). It is useful to write $D_j(p_j) = \sum_l^{l^{max}} a_j^l(p_j) Q_j^l(p_j) Pr_l$, where $a_j^l(p_j)$ is the probability listing $j$ appears in a consideration set that includes $l$ targeted listings, $Q_j^l(p_j)$ is the expected probability of transacting given a consideration set size $l$ (where the expectation is taken over different sets of competitors and different numbers of non-targeted listings in the consideration set), and $Pr_l$ is the probability the consideration set will consist of $l$ targeted listings. We estimate $Q_j^l(p_j)$ by simulating 1000 searchers per listing and forming their consideration sets according to the model of the platform.

We estimate $\partial a_j^l(p_j)/\partial p_j$ using the platform model. With the chain rule, we have $\partial a_j^l(p_j)/\partial p_j = (\partial a_j^l/\partial \tilde{w}_j)(\partial \tilde{w}_j/\partial p_j)$. In the before period, we have $(\partial \tilde{w}_j/\partial p_j) = 0$. In the after period, $(\partial \tilde{w}_j/\partial p_j) = -\gamma \tilde{w}_j$. We use the probability mass function for Wallenius' non-central hypergeometric distribution to numerically estimate $\partial a_j^l/\partial \tilde{w}_j$.

We use the logit formula to get $\partial Q_j^l(p_j)/\partial p_j = (\alpha_1 + \alpha_3 TRS_j)Q_j^l(p_j)(1 - Q_j^l(p_j))$. With these components, we can then estimate $\eta_{D_j}$ and back out $c_j$.

**Counterfactuals.** The counterfactuals alter components of the platform design, the distribution of listing quality, or the substitution patterns across targeted listings. The counterfactuals are largely self-explanatory with two exceptions. For the third column of Table 5, we construct consideration sets with "Demand Weight Rank." We seek to include both price and quality as determinants of a listing's sampling weight. We construct the sampling probability as:

$$\tilde{\omega_j} = \exp\left[-\hat{\gamma}\left(\frac{p_j - \min_{k\in\mathcal{J}_i^J}(p_k) - |\hat{\alpha_4}/\hat{\alpha_1}|q_j}{std_{k\in\mathcal{J}_i^J}(p_k)}\right)\right] \tag{28}$$

For Table 5, Panel B, we increase the degree of quality differentiation. To do so, we draw new listing quality, $q_j$, from a Uniform[-15,15] distribution and set sampling weights for the before period to $\tilde{\omega_j} = q_j + 15$.