

ONLINE APPENDIX

Consumer Price Search and Platform Design in Internet Commerce

by Dinerstein, Einav, Levin, and Sundaresan

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_s^L which may differ by search option s . Loosely, the “Product Page” will offer

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

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

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 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” ($s = 1$) became the default search option and the cost to deviation was d_1 . The size of i ’s consideration set, L_i , also changed, 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 we assume that the platform samples L_{i3} listings from J_t without replacement, where each listing j has sampling weight ω_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:

$$u_{ij} = v_{ik} + w_j + \epsilon_{ij}, \tag{1}$$

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

$$v_{ik} = \theta_k + \sum_{l=1}^k \lambda_{ik} \text{ for } k = 1, 2, 3, \quad (2)$$

with $\lambda_{in} \sim^{iid} N(0, \sigma_n^2)$. 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

$$w_j = \alpha_1 p_j + \alpha_2 TRS_j + \alpha_3 p_j TRS_j, \quad (3)$$

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 , ϵ_{ij} , with noise. She may then incur a costly click to learn ϵ_{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 ϵ_{ij} for listings of the targeted product ($k = 3$). For other listings, ϵ_{ij} is observed without noise. Denote the signal by η_{ij} and the correlation between ϵ_{ij} and η_{ij} by ρ . To observe ϵ_{ij} , the user must click on the listing at cost κ . Clicking has a dual role: it reveals ϵ_{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 ψ the user makes a clicking mistake – clicking when a click was not intended or not clicking when a

³We 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.

click was intended. Including the click noise is not strictly necessary, as extreme draws of ϵ_{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 $\bar{u}_i = \max_{k < 2} u_{ik}$ 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

$$E(u \text{ if click on } 1) = E(u_{i1} | u_{i1} > \bar{u}_i) \text{Prob}(u_{i1} > \bar{u}_i) + \bar{u}_i \text{Prob}(u_{i1} \leq \bar{u}_i) - \kappa \quad (4)$$

The expected utility from not clicking is \bar{u}_i . The user clicks if $E(u \text{ if click on } 1) > \bar{u}_i$, or equivalently if the expected gain from the click ($E(u_{i1} | u_{i1} > \bar{u}_i) - \bar{u}_i$) $\text{Prob}(u_{i1} > \bar{u}_i)$ is greater than the click cost κ . 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, λ_{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 μ_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 $D_j(p_j, TRS_j, \omega_j) = \int_i D_{ij}(p_j, TRS_j, \omega_j) dG_i$, where the arguments are the listing’s price, whether the seller is a top-rated seller, and the listing’s sampling weight (ω_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. Facing the platform’s transaction fee T and ad valorem fee t , each seller sets its price to solve

$$\max_{p_j} ((1-t)p_j - c_j - T) D_j(p_j, TRS_j, \omega_j). \quad (5)$$

We set $t = 0.1, T = 0.3$ and recover each seller’s marginal cost c_j by inverting the first-order condition:⁴

$$c_j = (1-t)p_j - T + (1-t) \left(\frac{\partial D_j}{\partial p_j} \right)^{-1} D_j(p_j). \quad (6)$$

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, ω_j (except for the “Product Page” option after the redesign, for which the platform samples according to price). We estimate the weights ω_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 click 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_{ij}^C be the model-predicted probability that i clicks on j , and Q_{ij}^Y 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

⁴We exclude the platform fees, but because they can be subsumed into the cost estimates, our demand and search estimates are identical.

j_i be the listing (or outside option) that i bought. Then the likelihood function is

$$\mathcal{L} = \prod_i P_i\{s = 0\}^{1-s_i} (1 - P_i\{s = 0\})^{s_i} (Q_{ij_i}^Y)^{Y_{ij_i}} \prod_j (Q_{ij}^C)^{C_{ij}} (1 - Q_{ij}^C)^{1-C_{ij}} \quad (7)$$

The 14 parameters to estimate are the product type fixed effects $(\theta_1, \theta_2, \theta_3)$, the variances of the random effects $(\sigma_1^2, \sigma_2^2, \sigma_3^2)$, the coefficients on targeted listings' characteristics $(\alpha_1, \alpha_2, \alpha_3)$, the correlation of the signal with the idiosyncratic listing-specific preference (ρ) , the cost of an intentional click (κ) , the probability of a click mistake (ψ) , 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 (λ_{ik}) , the signal (η_{ij}) , and the listing-specific preference term (ϵ_{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 variation that is useful in identifying 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 $(\sigma_1^2, \sigma_2^2, \sigma_3^2)$ 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 probit 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.

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 a \$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 α 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.”

Parameter	Description	Estimate
θ_1	product 1 FE	-3.7807
θ_2	product 2 FE	-3.2929
θ_3	product 3 FE	4.2683
σ_1^2	nest 1 RE Variance	0.9329
σ_2^2	nest 2 RE Variance	0.0164
σ_3^2	nest 3 RE Variance	1.6000
α_1	coefficient on price	-0.2000
α_2	coefficient on TRS	2.0764
α_3	coefficient on price*TRS	-0.0442
ρ	correlation(signal, ϵ)	0.9998
κ	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

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, κ , can thus be subsumed into the product 3 fixed effect, θ_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 (λ_i) or idiosyncratic preferences over search strategies (μ_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 γ parameter. This reduced form representation could distort substitution patterns if it ignores that consumers with different preferences over product types (λ_i) will select into pages with different sampling processes. If the selection is instead driven by the idiosyncratic preference over search strategies (μ_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.⁵ We also kept a smaller group of products – all iPhone 4 products – as a separate subcategory for comparison.

Experiment. For the experiment results presented in Section 3.3, we collected data on all products active during the experiment (7/25/12 – 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.

⁵The 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.

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 $D_i = 1 + \exp(\delta + \lambda \ln |J_i^M|) + \sum_{k \in J_i^J} \exp(\alpha_0 + \alpha_1 p_k + \alpha_2 TRS_k + \alpha_3 p_k TRS_k + \alpha_4 q_k)$. The

likelihood function is:

$$\mathcal{L} = \prod_i \left(\frac{1}{D_i} \right)^{1-Q_i^M - \sum_{k \in J_i^J} Q_{ik}^J} \left(\frac{\exp(\delta + \lambda \ln |J_i^M|)}{D_i} \right)^{Q_i^M} \prod_j \left(\exp \left(\begin{array}{c} \alpha_0 + \alpha_1 p_j + \alpha_2 TRS_j + \\ \alpha_3 p_j TRS_j + \alpha_4 q_j \end{array} \right) / D_i \right)^{Q_{ij}^J} \quad (8)$$

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, ω_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:

$$a_l(\omega_j) = \binom{1}{1} \binom{|J_i^J| - 1}{l - 1} \int_0^1 (1 - x^{\omega_j/D}) (1 - x^{1/D})^{l-1} dx \quad (9)$$

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, ω_j , such that the model-predicted fraction of appearances matches the data:⁶

$$v_j = \sum_{l=1}^{l^{max}} a_l(\omega_j) v_j^l. \quad (10)$$

⁶There are a few listings with $v_j = 0$ or $v_j = 1$. Our model is unable to find a unique positive ω_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 $\omega_j = \max_{k:0 < v_k < 1} \omega_k$.

We follow the same procedure in the before and after periods. In the after period, we project these weights onto listing prices:

$$\omega_j = \exp \left[-\gamma \left(\frac{p_j - \min_{k \in \mathcal{J}_i^J}(p_k)}{\text{std}_{k \in \mathcal{J}_i^J}(p_k)} \right) + \nu_j \right]. \quad (11)$$

We estimate γ with the following OLS regression:

$$\ln \omega_j = \tau - \gamma \frac{p_j}{\text{std}_{k \in \mathcal{J}_i^J}(p_k)} + \nu_j \quad (12)$$

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, $\tilde{\omega}_j = \exp \left[-\hat{\gamma} \left(\frac{p_j - \min_{k \in \mathcal{J}_i^J}(p_k)}{\text{std}_{k \in \mathcal{J}_i^J}(p_k)} \right) \right]$. 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 ω_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 ω_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 (η_{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 η_{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:

$$\tilde{\omega}_j = \exp \left[-\hat{\gamma} \left(\frac{p_j - \min_{k \in \mathcal{J}_i^J}(p_k) - |\hat{\alpha}_4/\hat{\alpha}_1|q_j}{std_{k \in \mathcal{J}_i^J}(p_k)} \right) \right] \quad (13)$$


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 $\tilde{\omega}_j = q_j + 15$.

References


- Chen, Yuxin, and Song Yao (2016). ‘‘Sequential Search with Refinement: Model and Application with Click-Stream Data.’’ *Management Science* 63(12): 4345-4365.
- Kim, Jun P., Paulo Albuquerque, and Bart J. Bronnenberg (2010). ‘‘Online Demand under Limited Consumer Search.’’ *Marketing Science* 29(6): 1001-1023.
- Ursu, Raluca M. (forthcoming). ‘‘The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions.’’ *Marketing Science*, forthcoming.

Appendix Figure A1: Profit-Maximizing Platform Choice

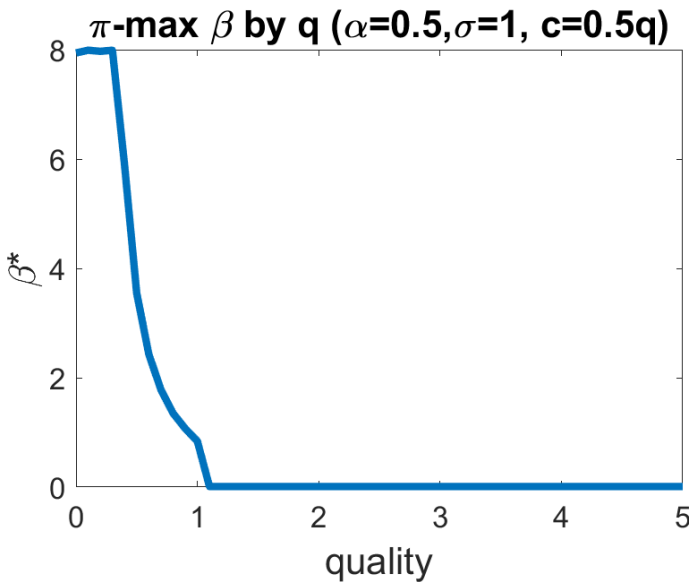
Panel A: β^* as a Function of α



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 platform profits, 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 $\alpha=0.5$, $q=1$, $c=0.5q$, and $\sigma=1$ unless they are the parameter being varied for the comparative static. The platform charges 0.3 per transaction and keeps 0.1 of the transacted price.

Appendix Figure A2: Timeline of Platform Changes

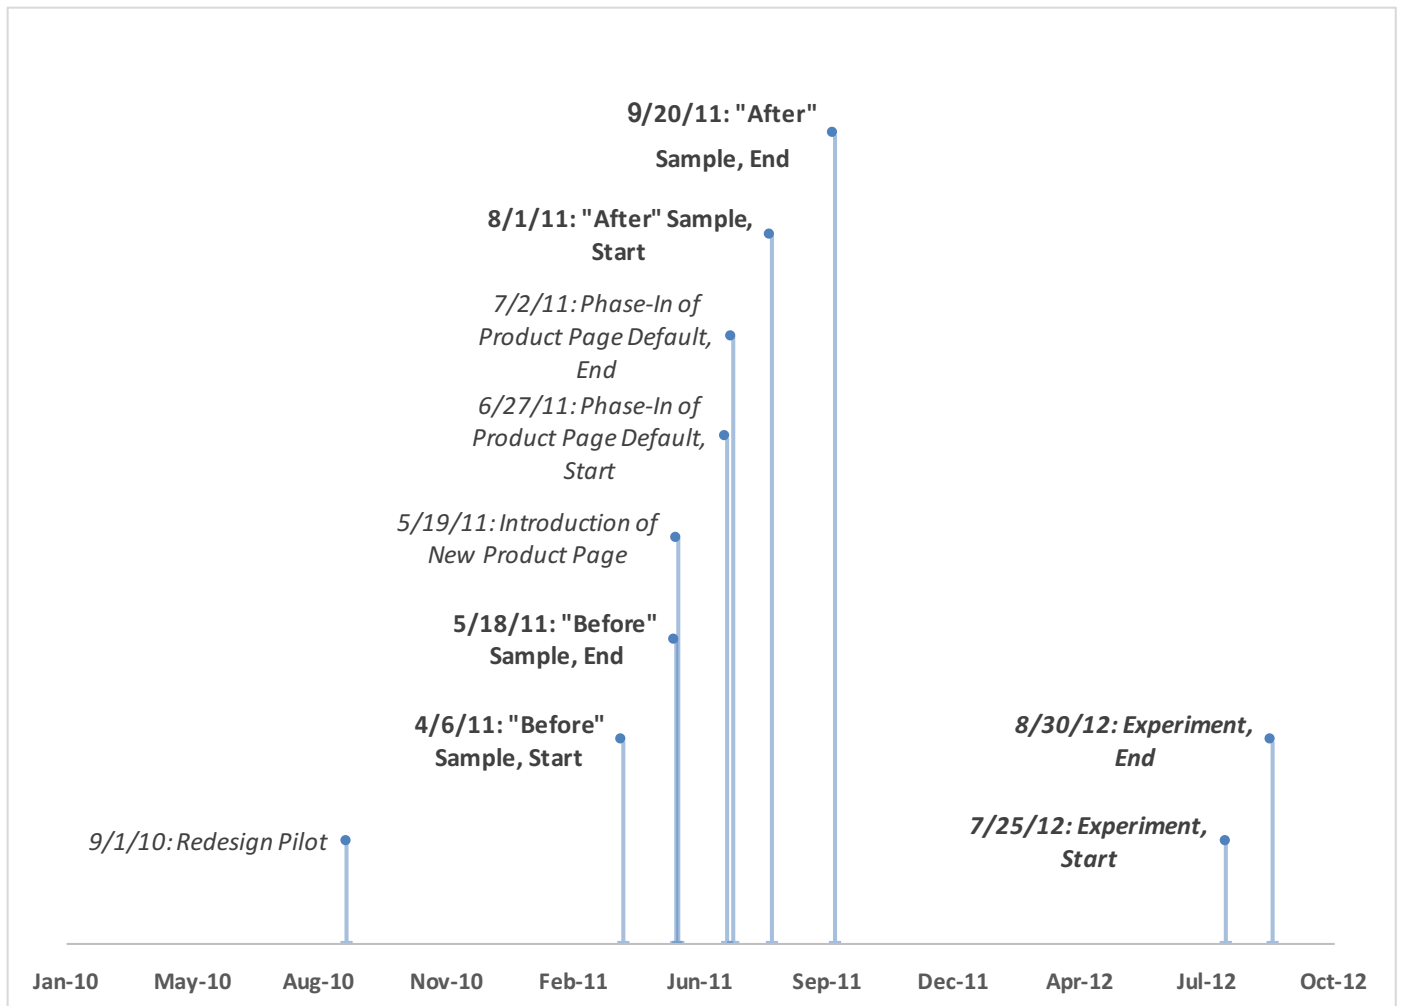


Figure shows the timeline of the platform changes and experiment, as well as the periods from which we pull our data for the empirical model. Bolded labels indicate the event corresponds to the timing of our data sample. Italic labels indicate the event corresponds to a change of eBay's platform.