

Returns to Consumer Search: Evidence from eBay

THOMAS BLAKE, eBay Research

CHRIS NOSKO, University of Chicago

STEVEN TADELIS, University of California, Berkeley, NBER and CEPR

A growing body of empirical literature finds that consumers are relatively limited in how much they search over product characteristics. We assemble a dataset of search and purchase behavior from eBay to quantify the returns, and thus implied costs, to consumer search on the internet. The extensive nature of the eBay data allows us to examine a rich and detailed set of questions related to search in a way that previous structural models cannot. In contrast to the literature, we find that consumers search a lot: on average 36 times per purchase over 3 (distinct) days, with most sessions ending in no purchase. We find that search costs are relatively low, in the region of 25 cents per search page. We pursue the analysis further by, i) examining how users refine their search, ii) how search behavior spans multiple search sessions, and iii) how the amount of search relates to finding lower prices.

Additional Key Words and Phrases: Search, Economics

1. INTRODUCTION

Across a wide range of markets, from online retail to consumer packaged goods to mutual funds, consumers must actively search to obtain information about available products and their prices. Following Stigler [1961], models of costly search helped explain imperfectly competitive behavior and price dispersion in seemingly competitive markets with homogeneous goods. A recent wave of empirical papers argue that consumers do not search extensively, implying that search costs must be high. We argue that these conclusions are based on limited data, and as such, were forced to rely on theoretical models and structural estimation to infer search costs. In essence, this literature substitutes models and structure for data, leading to search costs estimates that seem unrealistically high. We, instead, analyze both search behavior and the implied search costs with panel data containing comprehensive search and purchase behavior. We find that users search extensively, which implies that search costs are quite modest.

The theory underpinning the structural literature generally follows one of the two canonical models of search in economics. The first, developed by Stigler [1961], assumes that consumers sample a *fixed* number of stores, or websites, and choose to buy the lowest priced item. The second, more widely adopted theoretical model was developed first by McCall [1970] and Mortensen [1970], who posit that a model of *sequential* search is a better description of consumer search behavior. Both modeling approaches make assumptions and develop structure that can reduce the need for data in an attempt to measure search costs. For example, Hong and Shum [2006] infer search costs from a structural model that only uses online price data with no consumer behavior data at all. They estimate a median search cost for a book using a Stigler-like fixed sample size model to be \$2.32, whereas a sequential-search model delivers an estimate of \$29.40 per search. De los Santos et al. [2012] use Comscore data that includes online book purchases with some limited search

Author's addresses: T. Blake, thblake@ebay.com, C. Nosko, cnosko@chicagobooth.edu and S. Tadelis, stadelis@berkeley.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'16, July 24–28, 2016, Maastricht, The Netherlands. Copyright © 2016 ACM 978-1-4503-3936-0/16/07 ...\$15.00.

<http://dx.doi.org/10.1145/2940716.2940754>

behavior data. Their data is more consistent with the fixed sample size search model, and they estimate search costs that are on average around \$1.35 or \$4.14 per search, depending on some assumptions.

In a more recent paper, Koulayev [2014] develops a model motivated by the features of a hotel search website from which he obtained data. He estimates the median search cost of an extra online search for a hotel to be around \$10 per page of results, though it can be as high as \$30 for some consumers. This is a direct consequence of the fact that in his data, consumers search very little, many only engaging in one or two searches. This observation is, as we suggest, a consequence of limited data; because users in his sample are anonymous, Koulayev [2014] cannot connect searches made by the same person more than 24 hours apart.

Unlike the empirical studies described above, we use comprehensive data from eBay to shed light on the search process with minimal modeling assumptions. The data show that consumers actually search significantly more than other studies have suggested – on average 36 times per purchase. Consumer search is a protracted process, which can span 3.5 distinct days over a period of several weeks. Furthermore, there is a large tail of heavy searchers; we find that 5 percent of users are still searching for the same product 30 days after starting a search. Unlike previous studies like those mentioned above, we estimate the average costs per search to be about \$0.25. To put this in perspective, if a user spends about a minute per search then the estimates of Hong and Shum [2006], De los Santos et al. [2012] and Koulayev [2014] suggest an hourly opportunity cost of search in the range of \$81 to \$1,800, while our estimate is about \$15, which we believe is a much more realistic estimate.

Our rich data allows us not only to use minimal assumptions to estimate search costs, but it also offers some insights on the way people search. Consider a consumer's book purchase decision: You may have a vague intuition about the type of book you want to read (fiction vs. non-fiction), so you go to Amazon and search for "non-fiction", trusting Amazon's search engine to return best-selling results to you. You then click on a few different titles and read some reviews. The next day you return to Amazon and search for "non-fiction WW2" having decided you wanted to read a book in that category. After a few more sessions like this, you settle on a book and then check Barnes and Noble's website, and maybe a couple more websites, to see if they have a cheaper price. When Johnson et al. [2004] document that "On average, households visit only 1.2 book sites" they are referring to the very last piece of this search process. Similarly, when Kim et al. [2010] estimate that for consumers searching for a digital camera, "The mean of the search set size distribution is 14," they are inferring from decisions that consumers made within a single search session, ignoring previous (or later) search sessions that the same consumer engaged in.

In our data we are able to track individuals over time, and are therefore able to link their search sessions over time all the way through to either a purchase or abandonment of the search. The richness of the data allows us to shed light on the search process by using simple econometric techniques that do not rely on complex models or structure. We therefore contribute to the existing literature in three ways.

First, we start by documenting behavior within a given search session using a cross section of search efforts at a single point in time. We examine the clickstream data for a particular day and find that users refine searches as they engage with the site, very much like the example of choosing a book that we described above. The data suggests that the average number of terms in the query increases as the session unfolds. We also find that the use of the default "best match" ranking declines as users move to more "deterministic" searches that sort over prices, and the average price of search results declines over time. Thus, it seems that people search in ways that are consistent neither with the fixed-set nor with the sequential search models.

Second, we construct a cohort of 500,000 users who searched on a single day and track their activity over long periods of time. We then tabulate all search and purchase activity for these users for the following 30 days. We find that users search frequently and over a protracted period, sometimes spanning weeks. Users perform about 144 searches per user for an average of 4 transactions, suggesting an average of 36 searches per transaction. These searches also span an average of 11 distinct (i.e., non-consecutive, active) days. This is in stark contrast to previous studies who used data that is significantly more constrained. For example, De los Santos et al. [2012] analyze (limited) browsing data up to 7 days before a transaction and state that “one week is long enough to capture all search behavior related to a transaction.” Our data show that this is too restrictive an assumption, which will understate the amount of search performed by consumers.

Third, we put together a third data set based on a “reverse cohort” analysis to quantify the returns to searching, and through this analysis, propose a simple estimate of the cost of searching. Like much of the previous literature we create the cohort by conditioning on a purchase and then observing the purchaser’s prior search levels. However, eBay’s rich data allows us to go back over many days to fully capture the search process and to compare purchase outcomes to comparable purchases to see how the purchase price compares to the expected price for the item. Then, using a simple revealed preference approach, we back out the implied search costs. We find that gains from search are modest compared to the prior literature but still demonstrably positive. Consumers save, on average, 25 cents per search page and about 75 cents for each day spent searching. We then identify a proxy for patience and show that users who are more patient actually search more, consistent with intuition from economic theory.

The richness of our data, and our simple revealed preference approach puts our paper in stark contrast to other recent empirical studies. Previous studies either use data from Comscore [De los Santos et al. 2012; Johnson et al. 2004] or infer search costs from purchase or scraped “view-item” behavior [Kim et al. 2010; Seiler 2013]. In both instances, the data constraints are severe. With the Comscore data, researchers can observe the purchased product and the sites that were visited, but don’t observe the products that were searched for, the results returned to the user, or the number of searches within a site. When search behavior is inferred from purchases, a whole host of assumptions go into a model that substitutes for the lack of actual search data. Papers that work with actual primitive search data do not link users across sessions and, perhaps consequently, find that users search very little and have high search costs [Ghose et al. 2014; Koulayev 2014].


It is important to note that a limitation of our analysis is that we only observe what users do on the ebay.com website. This implies, however, that our already low estimates of search costs can be considered upper bounds. Because many users probably search on more than just eBay’s marketplace, actual search costs are likely to be even lower than what we estimate. Furthermore, if consumers may have sampled both online competitors and traditional brick-and-mortar retailers before buying on eBay then their implied search costs are even lower.

2. BACKGROUND & DATA

2.1. Background: Search at eBay

As with many online retail or marketplace websites, search is the main way that consumers find products on eBay. A common pattern is shown in Figure 1, where a consumer arrives at the ebay.com homepage and is confronted with a large search box on the top of the page. After entering a search term (or “query”), the user is taken to a “search results page” (SRP). As figure 1a documents, a list of available products, together with some information about each product, are available directly on that page. The information typically includes a picture, whether or not it is an auction or fixed price listing, the item’s price (or current

auction price), and when the listing ends. Fifty items are listed per page by default, with the user needing to click “next” in order to advance to the next page of results. If the user sees a product that interests him or her, he or she clicks on the title and is taken to a “view-item page” (VIP) with more information about the item, including detailed information about the seller, the product’s condition, and any other notes that the seller has entered about the product, as shown in Figure 1b. From there, if the user is interested in purchasing the product, he or she can bid in the auction (if it is an auction listing) or purchase the item from the seller if the item is a fixed price (“buy it now”) listing. If the user is not interested in the product or wishes to do more exploration, he or she can return to the SRP page (not counted as a separate search) and click on other items or refine the search query to change or refine the focus of the search.



(a) eBay Search Results Page

(b) eBay View Item Page

Fig. 1: eBay Search Experience

From the perspective of a marketplace platform like eBay (or any retail website more generally), the search process provides many levers for influencing a consumer’s decisions. Perhaps the most important is the order in which search results are displayed. Figure 1a illustrates that a search for the term “watch” returns over 1.4 million listings. With so many options to choose from, finding a product match without a good ranking algorithm would be a herculean task to say the least. It also indicates that the orderings of search results potentially plays a large role in determining which products are purchased or whether a product is purchased at all. With a platform like eBay, these results are a mix of different products (as illustrated by all of the different watch types available) and different sellers listing the same product for different prices or in different conditions. This provides us with opportunities, such as observing the same user searching through different sellers of the same product, but also with challenges, such as the difficulty in determining whether two listings are actually the same product because product characteristics are potentially amorphous.

By default, eBay displays search results using a ranking algorithm called “Best Match.”¹ The best match algorithm was created to display items in the order that best predicts expected eBay revenue, maximized by increasing the probability that a product is purchased times its sale price.² Behind the scenes is a machine learning algorithm where the target is eBay’s revenue, which is trained on data that is associated with both product and seller characteristics. The results of this machine learning process are fit to the current set of products available that match any given search term.

2.2. Available Data

Data for this project come from internal eBay records. For years eBay has done an extremely thorough job recording data from the search and purchase process. There are two sources for this data. First, all transaction relevant information such as bids, purchases, price paid, buyer and seller IDs, etc., are recorded in a structured database used as a record of all transactional activity on eBay. These records tend to be precise with very few errors or leakage. Second, eBay logs “clickstream” data that tracks how users navigate through the site. This is a much messier process given the amount of data and its semi-structured nature. These data are divided into sessions, defined by 30 minutes of inactivity for a given user. Within the session, eBay records all clicks that occur, and, for search results, an extremely rich set of information about what was displayed to the user, including the ordering of items displayed and their properties. Essentially, as far as search and economic choice is considered, eBay captures and records the whole “consideration set”, what a user typed to obtain that set, and how that translated into a user’s click behavior, all the way up to a purchase, or to abandonment of the search process.

One large benefit of eBay’s internal records data is that a user can be tracked across sessions and purchases, which allows us to link a user’s behavior across sessions over time. Correctly recording this is easy if a user is signed in across multiple sessions (which they would need to be in order to purchase or participate in activities like bidding or making offers, which show up in the transactional logs) but trickier given that most users are often not signed in during exploratory searches even if they have an eBay account. Fortunately, eBay does a substantial amount of work to unpack who these users are. “Cookies”, or little bits of information that are stored on a user’s computer and then transmitted to sites every time a page is requested, are key to this process. Whenever a new browser is seen by the eBay servers, a cookie with a unique ID is “dropped” on the browser/computer. This ID is then tracked through all of the clickstream data. If a user ever logs in on that browser/computer, the system automatically backfills all clickstream information to reflect the fact that the system has learned who the user was who generated all of that clickstream data. There is noise in the process, for instance, if multiple people sign in from the same browser/computer or if a user never signs in from a computer but browses on it. For the most part, however, the process works well – internal eBay audits indicated that somewhere around 70% of search behavior can be tracked back to an eBay user account.

2.3. Data Selection

In theory, transactional records are available for all users going back to 2005 and clickstream records going back to 2010. The volume of this data, however, is way too large

¹Users have the option of sorting according to other ranking schemes, including by highest price, by lowest price, and time ending soonest (for auctions). Interestingly, most users do not “unsort” best match, but we are cognizant of the potential concerns that these options give rise to and it will be discussed in the context of selecting a sample for our study.

²Historically, eBay has generated revenues that are comprised of some fees for listing an item on the site (listing fees) and a percentage of the sales prices (final-value fees, typically around 9%) when an item successfully sells. Most of eBay’s revenue is generated by the final value fees.

for any meaningful analysis without careful selection rules. We believe that three styles of analysis make the most sense in the context of analyzing search behavior: 1) **Cross-sectional**. For any given day or relatively short window, construct a dataset that includes all search behavior and associated product/seller characteristics. This sort of analysis is most suited for getting a sense of complete search behavior for a given search term. 2) **Cohort analysis**. Track all search and purchase behavior for a cohort of users, selected with some larger goal in mind. This style of analysis has the advantage of tying search behavior to users at a very detailed level and gives a complete picture for any given individual, but does not allow for working through models of equilibrium behavior. 3) **Condition on purchase**. Find all buyers of items and examine *pre*-purchase behavior, that is, search behavior that occurred before the purchase was recorded. This data, which effectively would be constructed looking backward (hence the term “reverse cohort”), would allow for the comparison of search strategies and effort across different types of purchases. We explore each of these approaches in the following sections.

3. CROSS SECTION - WITHIN SESSION BEHAVIOR

We start by documenting behavior within a given search session. This constitutes a cross section of search efforts at a single point in time. We examined the clickstream data for a particular day and summarized select measures of search as they evolve with time within a given session. For this purpose we used data generated by searches that started on July 27th, 2014. We computed the time since the users’ first site activity (on that day) for each search event and then examined how different measures of search behavior evolved as users refined their search.

As expected, many users drop off in the first few minutes of sessions suggesting that many searchers abandon the site if they do not quickly find what they are looking for. This, for example, might be a consequence of a consumer checking across several websites, eBay being just one, in order to get an idea of selection and prices. The rate of quick termination is not very high, and the majority of sessions last many minutes.


We then considered measures that quantify the specificity of each search. For example, early in a search a user may use the query “watch”, then refine it to “men’s watch” and later add further qualifying words such as color, shape, strap type, and more. Hence, the number of words used can proxy for how refined a search is. Similarly, early in the exploratory phase of a search a user may be happy to use eBay’s Best Match algorithm for sorting items, but once a search has been refined and the user has a clear item in mind for purchase, the user may want to use price as a ranking in order to find a good deal.

The data suggest that users do indeed refine searches as they engage with the site over time within a session. The results are summarized in Figure 2. The average number of terms in the query rises over time. The propensity to use the default ranking declines over time as users move to more ‘deterministic’ searches like price sorts. The average price of search results also declines over time.

In summary, the simple cross-section search analysis suggests that search follows an intuitive pattern. Users seem to focus and refine their searches as the search session progresses, with indications that the process starts with more of a “what would I like” phase, followed by more of a “find a good deal” phase. This behavior fits neither the fixed-sample search models nor the sequential search models described earlier.³

³Though not reported in our paper, similarly to De los Santos et al. [2012] we find that users return to view items that they already viewed earlier, which is inconsistent with the sequential search model. And the way in which people refine their searches is not captured by the fixed-sample models either.

Fig. 2: Evidence of Refinement from Within Session Behavior



This plots several metrics within a single search session. Each plot shows mean values for the indicated value on the vertical axis for the minute from session start on horizontal axis. Plots show (starting in the upper left) the i) number of users still actively searching , ii) the average number of words in each query searched, iii) the percent of performed searches that use the default algorithmic ranking, iv) the average price of all items returned on the search page, v) the number of purchase events, and vi) a normalized plot of the log of transaction prices which represents the percentage deviation from average purchase price.

4. COHORT ANALYSIS

For the purposes of trying to understand user behavior across time, we constructed a cohort of searchers by identifying a pool of searchers from a single seed day, July 27, 2014.⁴ We identified the full list of unique, logged in, users that performed at least one search on that day and then took a sample of 500,000 users to make the analyses manageable. We focused on logged-in users because their activity can be tracked over long periods of time. Any site activity that occurs on a browser or device on which the user has ever logged in can then be matched by mapping cookies to a common user. We then tabulated all search and purchase activity for these users for the following 30 days after the July 27 search. We believe that this period of time should be long enough to capture all of at least one purchase intent.⁵

⁴This date is a Sunday, which we chose deliberately to increase the sample as the broadest set of users search on Sundays. When instead sampled a Friday and performed our analyses on that data set, the results were virtually unchanged.

⁵We note that this sample construction is left truncated; this is likely to be searching activity related to the same purchase prior to the seed date for some users. Given that search activity is continuous and ongoing for many

We found that users search frequently and over a protracted period, sometimes spanning several weeks. Table I presents summary statistics of the resulting panel. Users search, on average, a great deal. There are 144 searches per user for an average of 4 transaction. This suggest an average of 36 searches per transaction. These searches also span an average of 11 distinct (i.e. non-consecutive, active) days. Users tend to search within a narrow product range, spanning 2.4 categories of products.⁶

Table I: Panel Summary Statistics

Variable	Mean	Std. Dev.	N
Transactions	4.074	14.388	500000
Searches	144.551	269.593	500000
Categories	2.39	1.5	499997
Number of Days Searching	11.353	8.965	500000
Clicked Items	12.553	2.036	500000
Days Repeating a Search	3.516	7.182	500000

Much of this search activity is undoubtedly for many overlapping search efforts because users may be searching eBay for many items at the same time. That said, there is evidence of substantial repetition of searches across separate days. An obvious way to measure repeat search behavior is to track the individual search query strings across time in our panel. That is, we identified all of the search queries on our seed date, and then identified which users repeated one of those searches for every subsequent day in the panel.

We found that the average user repeats a specific search query on 3.5 separate days during this 30 day panel window. Figure 3 shows that these repeat searches taper off over time but over 5 percent of the panel is still searching 30 days after the seed date. For reference, Figure 3 also plots the fraction of users that purchase on each day in the panel. The panel was selected based on actively searching on the seed date, so there is naturally a greater purchase volume in the beginning of the window. The purchase rate appears to reach a stable weekly cycle (peaks are Sundays) about half way through the panel, which is about 2 weeks. Interestingly, the search efforts continue past that time, which is indicative that unsuccessful search efforts last longer than successful ones.

In summary, the cohort analyses suggest that users search a lot more than one might have concluded from previous studies. As mentioned earlier, De los Santos et al. [2012] assume that “one week is long enough to capture all search behavior related to a transaction.” Our analyses shows that this is too restrictive an assumption, which will understate the amount of search performed by consumers. Moreover, we are restricted to search activity recorded only on eBay. Since multi-homing is common (i.e., users will search sites like Amazon, Etsy, and others), actual search activity is likely to be even more intense than we are able to show with our data.


5. REVERSE COHORT ANALYSIS

We now turn to an alternative approach to quantify the returns to searching, and through this analysis, propose a simple estimate of the cost of searching. We turn back to the tradition in the literature of conditioning on a purchase and then connecting price to the purchaser’s prior search levels. However, eBay’s rich data allows us to add two important

users, identifying the exact beginning of any given search intent would require additional assumptions about the definition of search intents. If anything, the fact that we may be truncating prior search behavior means that users may search even more than we infer they do from our data.

⁶We are using a slightly broad definition of category, such that there are 110 unique categories in our panel.

Fig. 3: Evidence of Long Horizon Search from Panel Data



features to the analysis that are critical: 1) we collect data going back over many days to try and fully capture the search process and 2) we compare purchase outcomes to comparable purchases to see how the purchase price compares to the expected price for the item.

This approach allows us to use a simple method that does not rely on the structure of a more complex theoretical model to tease out search costs from the data. That said, we still follow the most basic of economic theories: that by revealed preference on search intensity and the resulting price savings, we can identify the search costs. One disadvantage of our approach is that we do not attempt to explicitly identify search costs separately from other underlying structural parameters. The reverse cohort strategy uses variation from two different consumers searching different amounts for the same product. If other fundamentals differ between these consumers then the interpretation of our reduced form numbers as search costs might be incorrect. For instance, if an individual's price coefficient is correlated with their search costs, then inelastic customers might search less not because they have high search costs, but because the mapping of the gains from searching into utility (the price coefficient) is different relative to someone that searches a lot. Nevertheless, with this caveat in mind, we believe that the gains to searching that we document can be thought of approximately as search costs. If in order to obtain a reduction in price of \$2 a user must engage in 5 searches, then we infer the cost of a search to be 40 cents. Of course, to perform this kind of analysis we must set other things equal, in particular, the product under consideration.

To do this we identified all purchasers on an arbitrary date, July 27th, 2014. We then limited the sample to purchases of *common* and *well defined* goods which have defined product identifications (declared by sellers or flagged by eBay). These 'well defined' items are generally those with Universal Product Codes (UPCs). We defined a product as 'common' if we found at least 10 purchases of that same product in the 6 weeks prior to our selection

date. This allowed us to construct a distribution of prices for each of the goods in our sample.

Next we identified all search behavior of the buyer in the 6 weeks prior to the purchase. A challenge is to identify searches related to the product purchased, knowing that the queries over time may have changed due to refinements of all sorts. To do this, we first we counted the number of searches that *returned* items which are identified as being the exact same product that was eventually purchased. That is, we relied on eBay’s search engine to infer that the user is searching for the kind of items that were eventually purchased. We then identified the length of search as the time between the first search and purchase as another measure of search intensity. Finally, we counted the number of distinct days on which the user searched for the product.

We then computed the expected product price by taking the mean of all of the purchases of a given product in the 6 weeks prior to the selection date. We treat this as the expected price one would pay for a product in lieu of search activity. We proceeded to derive the discount relative to the expected product price as the percentage difference between expected price and the buyer’s realized purchase price.

Fig. 4: Reverse Cohort Search Distributions





Figure 4 displays the distribution of the deviation from mean price for the product purchased during the sample period, as well as the distribution of our three measures of search intensity described above. There is a general left skewness in the price data; most users realize a slightly below average price but some pay very high prices. The search metrics all exhibit long right tails; most users search very little yet some search quite intensively, consistent with the cohort analysis described in Section 4. We will use the search intensity measures to impute returns to search.

Using the data we collected we proceeded to explore the relationship between measures of prices paid and of search intensity, which are displayed in Figure 5. The first row of Figure 5 shows the mean price paid for the different levels of the indicated search intensity (days searching, days since first search, and the number of searches). There is generally a positive relationship between price and search, which at first glance may be surprising. However, this does not control for the product purchased. Users presumably spend more time searching for costlier purchases because they expect to get a larger absolute value of savings from additional searches. Hence, this should not be interpreted as a causal relationship but rather one driven by selection.

Fig. 5: Returns to Searching



We controlled for this selection by controlling for the products being purchased. That is, rather than use price paid on the *y*-axis we plot the percent difference from the *expected* price that would have been paid for the purchased product. The second row of Figure 5 shows a clear negative relationship between our three measures of search intensity and the price paid. That is, the more a consumer searches for a given product, the lower the price paid for that product.

The results shown in Figure 5 sit well with search theory: the more a user can benefit from search, the more they search, and conditional on a given product, the more they search the less they pay. We can quantify this return on searching using a set of simple regressions to infer the actual costs of a search. That is, using a revealed preference approach we use the actual searches and the resulting price discount to infer the willingness to search, and hence, the implied search costs. Table II shows the results.

Columns 1 through 3 show regressions of price on search with product fixed effects. Each additional search is associated with a 26 cent reduction in the price. Columns 4 through 7 show results from a regression using percent discount and log price as dependent variables. The coefficients in these columns can be interpreted as percentage gains to searching. An additional search is associated with a 0.2% to 0.3% gain. For the mean sample purchase price in this sample, that is also about 25 cents. Each additional day spent searching yields a 0.8% or 75 cents savings.

These magnitudes are much more sensible than those described earlier from the existing literature and seem like reasonable returns to such modest efforts. It is hard to exactly assess how long a search takes. As described in the introduction, if a user spends about a minute per search then the estimates of Hong and Shum [2006], De los Santos et al. [2012] and Koulayev [2014] suggest an hourly opportunity cost of search in the range of \$81 to \$1,800, while our estimate is about \$15 and hour. If a search takes half the time then these estimates should be doubled. It is, of course, the fact that we observe a lot of search behavior that results in our lower and more believable measures of search costs.

Table II: Quantifying Returns to Search

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Price Paid	Price Paid	Price Paid	% Diff from Expected Price	Ln(Price Paid)	Ln(Price Paid)	Ln(Price Paid)
Searches Returning Product ID	-0.264*** (0.0308)	-0.0882*** (0.0341)	0.0588 (0.0541)	-0.00204*** (0.000208)	-0.00333*** (0.000323)	-0.00118*** (0.000354)	0.000418 (0.000561)
Days Since First Search		-0.317*** (0.0268)	-0.272*** (0.0297)			-0.00399*** (0.000279)	-0.00350*** (0.000309)
Days Searching			-0.759*** (0.217)				-0.00824*** (0.00225)
Product Expected Price	0.884*** (0.00247)	0.886*** (0.00246)	0.886*** (0.00246)				
Ln(Product Expected Price)					1.015*** (0.00270)	1.020*** (0.00270)	1.020*** (0.00270)
Constant	0.492 (0.469)	2.040*** (0.484)	2.447*** (0.498)	-0.127*** (0.00266)	-0.260*** (0.0111)	-0.258*** (0.0110)	-0.254*** (0.0110)
Observations	14331	14331	14331	14331	14331	14331	14331

Standard errors in parentheses
* $p < .1$, ** $p < .05$, *** $p < .01$

6. SEARCH COST HETEROGENEITY

One might reasonably expect there to be a large amount of heterogeneity in search behavior across consumers. Indeed, there are surely a myriad of factors that distinguish intensive searching consumers from more passive consumers. An obvious one would be heterogeneity in search costs. Any economic theory of search would imply that consumers with higher search costs will search less. A complete exploration of the mechanisms underlying search intensity is beyond the scope of this paper, but we can show that comprehensive data like ours can unlock insights that narrow data and modeling assumptions cannot.

To shed some light on the issue of heterogeneity, we explore one consumer characteristic that would explain heterogeneity in search intensity: patience. More patient consumers would be willing to spend more time delaying a purchase in order to achieve a better deal. Hence, if we could measure patience, then we could test whether more patient users tend to engage in more search.

Clearly, patience is not observable to us. However, we can use another revealed preference choice of consumers to rank them as more or less patient. As a proxy for patience we use a consumer's choice of shipping methods. Many sellers on ebay offer potential buyers a variety of shipping methods from slowest to fastest, and the faster shipping methods obviously cost more. Hence, users who choose faster shipping methods reveal less patience compared to those who choose those that are cheaper, yet come with more delay.

We constructed an indicator for each consumer recording whether or not they chose the fastest option when they were faced with multiple shipping options, hence being a proxy for impatience. The assumption is that the more impatient a consumer is, the more willing they are to pay extra for a faster shipping option. We then regressed this indicator of impatience on our measures of search intensity in the reverse cohort dataset.

Table III shows that choosing expedited shipping is generally negatively correlated with search intensity. This confirms our intuition from search models: the more impatient a

consumer is, the less they should engage in search behavior that will delay their purchase. Columns (1) through (3) show that each of our search intensity measures are significantly negatively correlated with our proxy of impatience. When all three measures are put in the same regression as shown in column (4), only the number of searches returning the same product ID remains significant. This should not be surprising because the three measures are positively correlated.

Table III: Search and Patience: Shipping

	(1)	(2)	(3)	(4)
	Pr(Expedite)	Pr(Expedite)	Pr(Expedite)	Pr(Expedite)
Days Searching	-0.00225*** (0.000707)			0.000132 (0.00136)
Days Since First Search		-0.000429** (0.000168)		-0.000217 (0.000206)
Searches Returning Product ID			-0.000695*** (0.000192)	-0.000618* (0.000335)
Constant	0.0735*** (0.00242)	0.0724*** (0.00239)	0.0726*** (0.00216)	0.0738*** (0.00254)
Product FE	Yes	Yes	Yes	Yes
N	14509	14509	14509	14509

Standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

7. DISCUSSION AND CONCLUSION

Putting all of the analysis together leads to a rather coherent story of search behavior. Consumers search a lot on eBay both within and across sessions. This intensive amount of search behavior should imply that search costs are relatively modest online. Indeed, our measure of the returns to search is on the order of 25 cents per search. This is in contrast to the existing literature findings, which indicate very little product search and consequently estimate high search costs as a justification of the low amount of search.

The data suggest not only that search costs are low, but that search proceeds as a kind of “funnel” where initially search is along broad categories, and then search becomes more refined to obtain a good at the lowest cost given a consumer’s cost of search. In thinking with this heuristic model, a few things become apparent. Firm behavior, and specifically the ordering of search results, should cater to these types of search patterns. Search engines might be better served if it built holistically around this search funnel and helped consumers learn about the attributes of products in a way that ultimately led them step by step down the process instead of assuming that they are at the end of it. Search queries that indicate browsing behavior might be met with results designed to encourage learning. In a way, this idea of endogenous firm behavior relates closely to the literature on informative advertising and consideration sets. In that literature [Draganska and Klapper 2011; Goeree 2008], products enter a consumer’s consideration set through costly advertising by firms. Here, products enter a consumer’s consideration set through search ranking by firms. Either way, one begins to question whether firm behavior coincides with actions a social planner might take and whether oligopolistic markets create inefficient outcomes.

This search behavior has implications both for optimal firm actions and for what market equilibria might look like. Consider a multi-stage model where consumers are unaware of

either the distribution of product characteristics or of the individual draws of characteristics at different web sites. The consumer must decide, in essence, whether to walk down a search funnel at eBay or at another site. The user could of course do both, but the narrowing down process is costly and does not necessarily translate across sites. For instance, the user might observe the current set of results that are returned from a search query for “watch” on eBay, which include used and old-fashioned watches, narrowing down the search criteria over time, and homing in on a self-winding antique. The user might then search for this same or similar watches on other sites, which unlike eBay, would probably have a very limited selection of antique used watches. Thus, because it is costly to do every iteration of every search across sites, the initial decision of which site to search on may have large consequences despite a large number of searches.

This heuristic model of consumer behavior also sheds light on an empirical regularity that has been a bit puzzling – evidence that consumers have strong brand preferences across websites that essentially offer very similar services. For instance Chevalier and Goolsbee [2003] argue that despite similarities between Amazon and Barnes and Noble across a wide variety of dimensions, consumers still have strong preferences for one or the other (see also Ellison and Ellison [2005] for a summary of this literature). Plenty of explanations, such as the fixed cost of registering for an account, have been put forward to explain this, though these do not seem terribly convincing to us. The heuristic search funnel model might be another explanation. If consumers start to learn about the types of products that each site carries through search, then brand preferences will develop endogenously and be in part controlled by firm behavior, in particular the search experience.

REFERENCES

- Judith Chevalier and Austan Goolsbee. 2003. Measuring prices and price competition online: Amazon.com and BarnesandNoble.com. *Quantitative marketing and Economics* 1, 2 (2003), 203–222.
- Babur De los Santos, Ali Hortaçsu, and Matthijs R Wildenbeest. 2012. Testing models of consumer search using data on web browsing and purchasing behavior. *The American Economic Review* 102, 6 (2012), 2955–2980.
- Michaela Draganska and Daniel Klapper. 2011. Choice set heterogeneity and the role of advertising: An analysis with micro and macro data. *Journal of Marketing Research* 48, 4 (2011), 653–669.
- Glenn Ellison and Sara Fisher Ellison. 2005. Lessons about Markets from the Internet. *Journal of Economic Perspectives* (2005), 139–158.
- Anindya Ghose, Panagiotis G Ipeirotis, and Beibei Li. 2014. Examining the impact of ranking on consumer behavior and search engine revenue. *Management Science* 60, 7 (2014), 1632–1654.
- Michelle Sovinsky Goeree. 2008. Limited information and advertising in the US personal computer industry. *Econometrica* 76, 5 (2008), 1017–1074.
- Han Hong and Matthew Shum. 2006. Using price distributions to estimate search costs. *RAND Journal of Economics* 37, 2 (2006), 257–275.
- Eric J Johnson, Wendy W Moe, Peter S Fader, Steven Bellman, and Gerald L Lohse. 2004. On the depth and dynamics of online search behavior. *Management Science* 50, 3 (2004), 299–308.
- Jun B Kim, Paulo Albuquerque, and Bart J Bronnenberg. 2010. Online demand under limited consumer search. *Marketing science* 29, 6 (2010), 1001–1023.
- Sergei Koulayev. 2014. Search for differentiated products: identification and estimation. *RAND Journal of Economics* 45, 3 (2014), 553–575.
- John J. McCall. 1970. Economics of information and job search. *Quarterly Journal of Economics* 84, 1 (1970), 113–126.

- Dale T. Mortensen. 1970. Job search, the duration of unemployment and the phillips curve. *American Economic Review* 60, 5 (1970), 847–862.
- Stephan Seiler. 2013. The impact of search costs on consumer behavior: A dynamic approach. *Quantitative Marketing and Economics* 11, 2 (2013), 155–203.
- George J. Stigler. 1961. The economics of information. *Journal of Political Economy* 69, 3 (1961), 213–225.