Consumer Online Search with Partially Revealed Information

Chris Gu*

Yike Wang‡

Abstract

Modern-day search platforms generally have two layers of information presentation. The outer layer displays the collection of search results with attributes selected by platforms, and consumers click on a product to reveal all its attributes in the inner layer. The information revealed in the outer layer affects the search costs and the probability of finding a match. To address the managerial question of optimal information layout, we create an information complexity measure of the outer layer, and study the consumer search process for information at the expense of time and cognitive costs. We first conduct online random experiments to show that consumers respond to and actively reduce cognitive cost for which our information complexity measure provides a representation. Then, using a unique and rich panel tracking consumer search behaviors at a large online travel agency (OTA), we find cognitive cost is a major component of search cost, while loading time cost has a much smaller share. By varying the information revealed in the outer layer, we find price revelation shifts search behavior most dramatically compared to the other product attributes. We propose information layouts that Pareto-improve both revenue and consumer welfare for our OTA.

Keywords: online consumer search, cognitive modeling, information complexity, search intermediaries, platform design.

JEL Classifications: D83, L81, L86

*Scheller College of Business, Georgia Institute of Technology, chris.gu@scheller.gatech.edu.

†Department of Economics, London School of Economics and Political Science, y.wang379@lse.ac.uk

‡We are grateful to Jean-Pierre Dubé, Pradeep Chintagunta, Günter Hitsch, Michael Wichura, Coley Dale, Zhiying Hu, Elisabeth Honka, Sumon Datta, Axel Stock, Bradley Shapiro, Rong Song, Mingjie Zhu, Hua Zhao. We also thank the feedback from the audience of the Marketing Science Conference 2016, Georgia Research Symposium 2017, and the University of Chicago Marketing working group. Last but not least, we would like to thank the generous support of INFORMS Society for Marketing Science Doctoral Dissertation committee and Mary Kay Foundation Doctoral Dissertation committee. The usual disclaimer applies.

1 Introduction

The Internet has opened an expansive ocean of information and product choices to consumers. Search platforms are valuable compasses in this ocean helping consumers acquire information and make choices effectively and efficiently. Modern-day search platforms can generally be considered as having two layers of information presentation. The outer layer displays the collection of search results, and consumers may click on a search-result entry to learn about all the product details in the inner layer. To assist consumers during their search process, platforms usually summarize the products by selecting a subset of product attributes to display in the outer layer. If too much information is presented in the outer layer, the need to click beyond the outer layer is reduced; yet the search-result pages may become too complex and require too much cognitive effort for consumers to understand. If too little information is presented, the search-result pages may become trivial to understand, but consumers need to incur high clicking cost to find what they need.

We use the historical layout change in Google's search-result pages as a motivating example. Figure 1 shows the layout of search-result pages for the keyword "minimalistic" in 2000 on the left and in 2016 on the right.[1] The left panel shows the early design, which only revealed the title of the search-result items. The right panel shows the modern design, under which Google reveals additional article information to consumers, including host names, URLs, and the tagline of the articles. Consumers gain more certainty about the link content by reading the additional information, and hence click fewer times.

Similar design decisions need to be made for shopping search platforms, such as Expedia, where the platform needs to decide which hotel attributes to display on the search-result pages, and how the search-result pages have to be revised for smaller screens, for example, mobile phones. Understanding how much information to reveal in the outer layer to help consumers search more effectively is in the interest of search platforms. This issue is a central platform design decision that balances simplicity and usefulness (Wong (2017)).

Before we delve into further details, we lay out the key elements of the consumer search process and the terms used. A more thorough introduction will be provided in Section 3. During their search process on modern search platforms, consumers navigate the outer layer by visiting different

[1]Photo courtesy of http://sebastians-pamphlets.com/rediscover-googles-free-ranking-checker/.

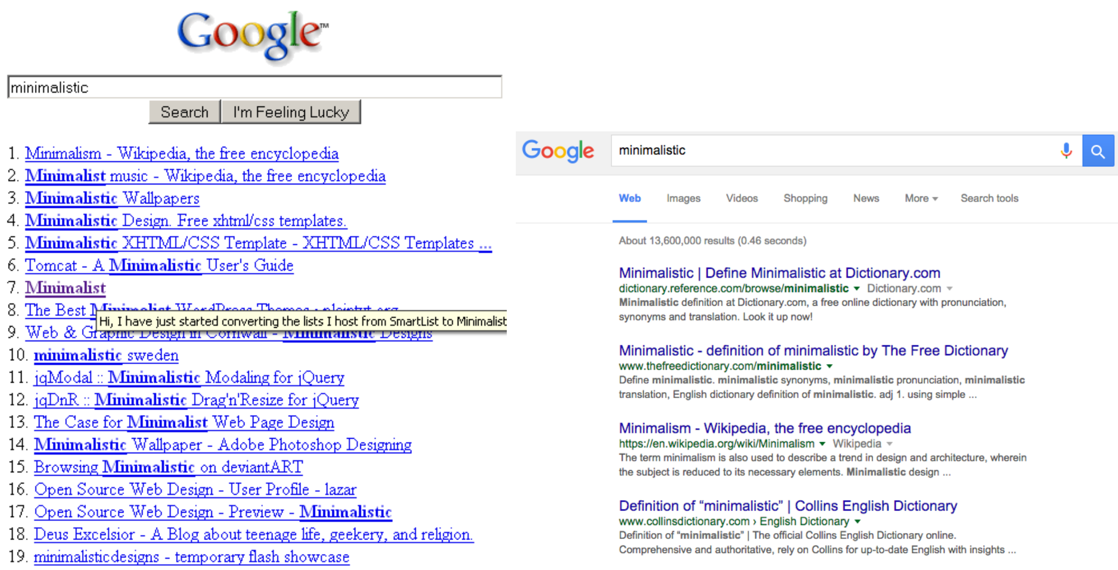

Figure 1: Changes in Google's search interface. The left panel shows the design of the search-result page in 2000, and the right panel shows the design in 2016.

"refinements" that are collections of search results. Each refinement is defined by some filter criteria that narrow the search-result set, and by some sorting criteria that order the search results. If the consumers specify no sorting criteria, a default platform-suggested ranking is applied to order the search results. Once they observe the product availability and the partially revealed product information on the visited refinements, the consumers decide which products to click on for more information. While the consumers acquire product information in this process, they also incur various search costs, including time and cognitive cost. In particular, consumers need to wait for webpages to load, and to exert cognitive effort to understand and compare the revealed product information.

In this paper, we answer the research question of how consumers acquire information at the expense of time and cognitive costs. Answering this question helps us address the managerial question of the optimal amount of information to reveal in the outer layer. To answer these questions, we collect measures of time cost and innovate measures of cognitive costs. We then find model-free evidence that supports the impact of these costs on the consumer search process. We conceptualize the consumer search process described above by proposing a new structural model. Finally, given the above knowledge, we try to improve the platform design that balances consumer welfare and platform revenue.

We further expound on our research and illustrate its connection to the most related literature in terms of the research question, cost measures, data, and model.

In terms of the research question, the consumer search literature grew out of the economic theory of information search (Stigler (1961)), in which consumers are no longer assumed to be all-knowing agents as in the traditional demand theory. Instead, consumers form their consideration sets for final purchase decisions by engaging in information search (Hauser and Wernerfelt (1990) and Roberts and Lattin (1991)). Search cost estimation is the focus of the empirical search literature because it lies at the heart of search models. Depending on the managerial questions being answered, past literature has looked at the overall search cost per consumer search, such as in Honka (2014) for customer retention study, and at positional search cost for optimal default platform-suggested ranking, such as in Ghose et al. (2012, 2014) and Ursu (2018).

Both the refinement visiting and product clicking decisions are important for effective search. Scant research has studied refinement choice, because of the high demands on data. The closest papers to ours are Chen and Yao (2016) and De los Santos and Koulayev (2017), which highlight the importance of studying the consumer refinement choice in that they reflect consumer preferences. For example, De los Santos and Koulayev (2017) assert that “consumers who sort by price in search are more price sensitive in choice.” Our insight is that, besides the relationship with consumer preferences, different refinements also have different “information complexity” (to be defined next) and hence require different cognitive effort to process. When consumers switch refinements, they not only change the potential information they will find and hence reflect consumer preference, but also alter the information complexity and hence reveal information on the cognitive search cost.

In terms of the cost measures, we propose to use webpage-loading time for the time cost, and we innovate “orderedness entropy” as an information complexity measure, which, together with consumer responsiveness to it, accounts for cognitive cost. The concept of entropy comes from information theory, first proposed by Shannon (1948). The earliest appearance of an application of entropy in the marketing literature is Herniter (1976), who builds a probabilistic model of market share based on an underlying consumer taste distribution, and estimates this distribution by maximizing the entropy of the probabilistic model. For a survey of how information theory is used in marketing as a unified modeling approach, see Brockett et al. (1995). Our research also relates to the recent stream of economics research on rational inattention in e.g., Caplin (2016) and Matějka and McKay (2015). The rational inattention literature has also applied the concept of entropy, but mostly in its original sense to quantify the randomness of and the distances between distributions.[2]

The essence of our measure lies in the insight that filtering and sorting not only bring the consumers' favored products to display, but also reduce the consumers' cognitive effort. As a result, an effective information complexity measure needs to capture, at its bottom line, the impact of the well-sortedness of a string of values on consumer perception. Once we have an orderedness measure, we use entropy to quantify the well-sortedness. Computing entropy for a string of values directly would miss the target, because permuting that string will generate the same entropy value.

In terms of the data, we have collected a rich data set that allows us to accomplish our task. As we mentioned, the study of refinement choice puts a high bar on the data for the following reasons. Each refinement is composed of many products, such as 25 items typical in the online hotel agency context, and each product has many attributes, some of which are dynamic. For a particular refinement on different days, the search platform may display different inventory and product attribute values, such as price levels, due to their dynamic nature. Therefore, consumers may observe different information on different days even for the same refinement. To carefully model the consumer search process, we need to reconstruct the exact information consumers have observed. The search platform we investigate managed to store a large amount of refinements over a long period of time. We managed the big data challenge to retrieve the dynamic product attributes and rebuild the decision context during the time window in which consumers arrive. Finally, we organized the data amenable for efficient search model estimation. The panel data researchers have used in previous consumer search literature mainly contains the information for the products consumers have clicked on, but not the refinements consumers have visited, such as in Chen and Yao (2016), Kim et al. (2017) and Ursu

[2] The rational inattention (RI) literature studies a general decision context where the decision makers try to acquire information optimally. To characterize the cost of processing information, the RI framework proxies a decision maker's learning effort using the reduction of Shannon entropy from her prior belief distribution to the posterior belief distribution. Although the entropy concept plays an important role in our research as well, there are the following differences. Our empirical context involves a consumer making a sequence of actions to resolve uncertainty of product availability and attributes dynamically. The standard RI framework models a single static decision to reduce uncertainty and make a choice. In addition, there is a natural notion of ordering (for table-formatted data) that search platforms utilize by providing sorting and filtering functionality on the platforms. In particular, contemplating a refinement page with different ordering of products due to sorting and filtering would require different levels of cognitive effort. The original RI framework models a generic decision context without such notion of ordering and cannot capture ordering-induced cognitive costs. Our model tailors to the specific dynamic search context and captures the notion of ordering using the new orderedness entropy measure.

(2018). The data that are closest to ours are used in Ghose et al. (2014) and De los Santos and Koulayev (2017).

In terms of the model, our central principle is to characterize the evolution of consumers' information and cost structure during their search process, accounting for model tractability. Our model is rooted in the sequential search model proposed by Weitzman (1979) and follows the same tradition as in Kim et al. (2010, 2017), Chen and Yao (2016), De los Santos and Koulayev (2017) and Ursu (2018). The two papers that study refinement choices have the following defining characteristics. Chen and Yao (2016) model each consumer search action as a combined refinement and click choice. The authors assume consumers are unaware of the attribute values of all unclicked products at any stage of the search process. As a result, after a consumer arrives at a particular refinement and clicks on a product, if she decides to click on another product on the same refinement, her information about and cost of the second product is assumed to be the same as if she had not previously arrived on the refinement, even though the refinement reveals some information about the second product. De los Santos and Koulayev (2017) model refinement choice using sequential search model, and adopt a discrete choice model of product clicks in a reduced form fashion by assuming consumers click on at most one product on any refinement. The use of a discrete choice model also assumes that the attributes and utility shocks of all products are known to consumers after they visit a refinement. We cannot adopt these models to address our research and managerial questions, because our model needs to have the flexibility to account for the following scenario. When more information is revealed on refinement pages, for example, our model should allow the consumers to utilize the additional information to help their product clicking decisions. They may click less because less uncertainty surrounds the products they have not sampled, even though the search cost for clicking has not changed. On the other hand, the additional revealed information should come at the expense of higher cognitive cost for refinements. As a result, our model needs to capture both the refinement visiting and product clicking decisions, to distinguish between what consumers know before and after these two types of search decisions, and to decompose search costs further into time and cognitive components. Our model uses sequential search framework to allow both refinement visiting and product clicking as search options, while adopting the same one-step-ahead simplification as in De los Santos and Koulayev (2017) for model tractability.

In our empirical work, we find that consumers tend to visit more refinements than to click on individual products within a search session (a session is the set of search activities within a start and end timestamp provided by the online travel agency (OTA), similar to the definition in the previous literature), which means consumers actively use search-result pages as an important source of information. The increase in the loading time of refinement and product-detail pages decreases the number of refinement visits and product clicks. Furthermore, we use model-free analysis and online random experiments to demonstrate that consumers respond to the type of information complexity represented by orderedness entropy, and they are more likely to sort the information on a refinement when sorting induces a higher reduction in cognitive cost. Based on our structural model estimates, we find cognitive cost is a major component of all search costs. For a typical refinement visited, consumers incur the cognitive cost equivalent of \$21.56 on average. Time cost is estimated to be \$20.28 per hour on average, and its share in the search costs is much smaller compared to cognitive cost. We find the position cost to be \$0.07, which is equivalent to a cost of \$1.69 if consumers scroll to the end of a refinement page with 25 products on it. This position cost is similar to that of \$0.05 estimated by Chen and Yao (2016) for an average consumer in their context.

Supported with the information complexity measure and the structural model, we study a new managerial question of optimal information layout that is difficult to address otherwise. We vary the information revealed in the outer layer to characterize the various changes in the consumer search process. We find price is an important product attribute whose presence in the outer layer will dramatically change consumer search behavior. We also study the relationship between consumer welfare and platform revenue. On the Pareto frontier of this relationship, we find a general tradeoff in which consumer welfare increases and revenue decreases as more information is revealed. For our OTA under study, we find room for Pareto improvement in both revenue and consumer welfare by changing the composition of the revealed product attributes. In the consumer welfare-maximizing layout, price is revealed (among other changes in attribute revelation), and the platform may increase consumer welfare by 95% without losing any revenue. In the revenue-maximizing layout, price is not revealed (among other changes in attribute revelation), and the platform may increase revenue by about 60% without hurting consumer welfare. We also look at win-win configurations in which both consumer welfare and platform revenue may be improved nontrivially.

To summarize our contributions in this paper, we extend the search literature to study cognitive and time costs that are crucial in the consumer online search process. We then make a first attempt at answering an important managerial question of how much information should be revealed in the outer layer, which could be relevant for all modern search platforms. We innovate an information complexity measure for cognitive cost, and propose a new structural model to model consumers' endogenous refinement visiting and product clicking decisions. We believe our modeling principle is in full accordance with the essence of search model literature, which is to capture what the decision makers do and do not know at each decision point, and to advance our understanding of the associated finer-level costs involved in the search process.

The rest of the paper is organized as follows. Section 2 describes the empirical context of our study. Section 3 illustrates the conceptualization of the consumer search process. Section 4 details our various search cost measures, and in particular, cognitive cost. Section 5 summarizes the data. Section 6 presents model-free evidence and online random experiment results aimed at exploring whether consumers respond to the cost measures. In Section 7, we develop our structural model. We discuss model identification and our estimation approach in Section 8. Section 9 presents and discusses our estimation results. Section 10 presents two managerial studies of optimal information layout and Pareto improving local sorting. Section 11 concludes.

2 Empirical Context

Our empirical context is a major online travel agency (OTA) in China, which is the counterpart to Expedia in the United States. The site design can be seen in Figure 2. This search platform provides search functionalities in terms of filters and sorting variables. The filters allow consumers to reduce search-result sets by retaining results that fit the exact criteria specified, in terms of city, check-in and check-out dates, price, hotel star and so on. The sorting variables, including default ranking, consumer rating, price, and hotel star, allow consumers to rearrange the order of display, and only one active sorting variable can be applied. Search results are displayed according to default ranking if no sorting variable is specified; price and hotel star have both ascending and descending sorting capability, whereas the other two variables only have one sorting direction. A search query is defined in terms of both the filters and the sorting variable chosen by a consumer. When a consumer submits a search query, the search platform retrieves the search-result set that satisfies the filters, and generates a refinement page ordered by the sorting variable for the consumer to

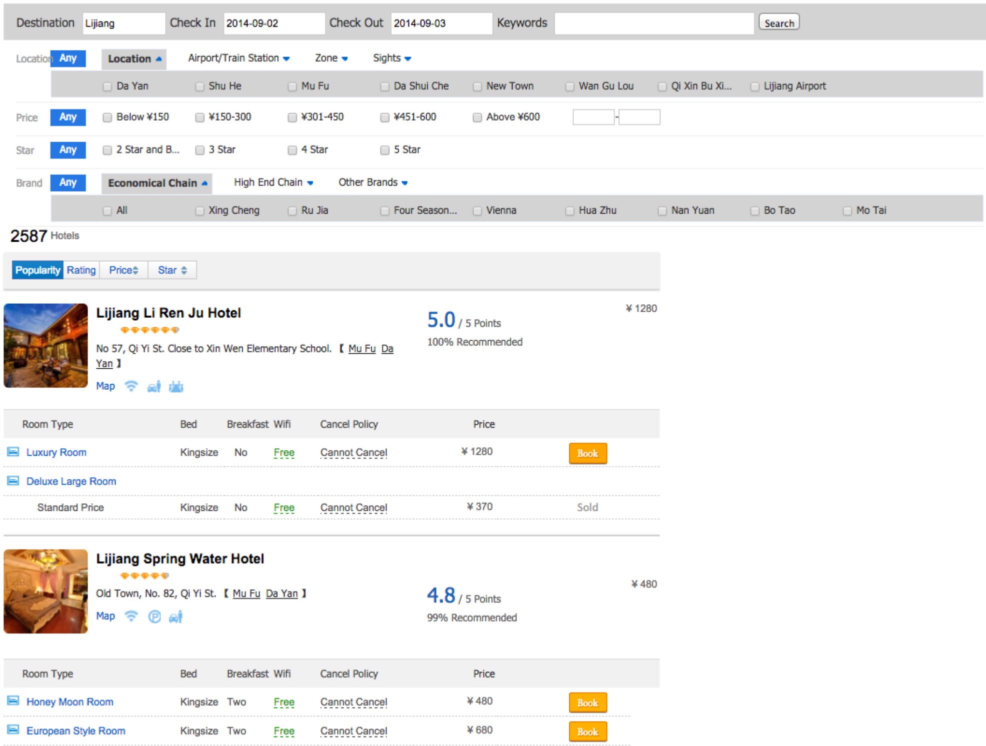

Figure 2: OTA in our empirical study.

interact with further. A given search query may return different content at different times, because product availability and some product attribute values are dynamic, such as price and cancellation policy.

Beneath the search functionalities, the site displays the search-result set of up to 25 hotels on a single page, with more results to display in latter pages. These search-result pages are what we call "outer layers," where the platform selects a subset of product attributes to display, including price, consumer rating, hotel star, breakfast, cancellation policy, and so on. If the consumer clicks on a hotel, she lands at the "inner layer" of the product-detail page, where she can learn about all the remaining attributes of the clicked hotel, including free newspaper, free guide service, free afternoon tea, and so on. All the products on the OTA under study have the same inner layer structure; they only differ in the value of the product attributes.

3 Conceptualization of Consumer Search Process

As we mentioned in the introduction, our central modeling principle is to characterize the evolution of consumers' information and cost structure during their search process, accounting for model tractability. As a result, our model needs to capture both the refinement visiting and product clicking decisions, to distinguish between what consumers know before and after visiting refinements and clicking products, and to decompose search costs further, particularly into time and cognitive components.

We first recognize that refinement pages in the outer layer and product-detail pages in the inner layer are different information sources. Consumers observe product availability and the partially revealed attributes of all the products on a given refinement, whereas they learn about the hidden product attributes of a single product on a product-detail page. The search costs associated with these two information sources are also different. Instead of having separate mean search costs for these two sources, we further recognize that information collection takes time and cognitive effort, and variation of the factors that influence these costs may influence the search process. For example, slow platform loading time or complex website design may induce consumers to search less and make do with less-than-ideal products they have found. In our conceptualization, we will parse out the time and cognitive components of the search costs.

Furthermore, in the platform design of the OTA we are studying, the inner layer of product-detail pages have the same page layout, except that the attribute values are different for different products. By contrast, the outer layer of refinement pages will, by consumer choice, have different compositions of products in the variation of product attribute values and well-sortedness. For example, a search-result page that has a filter of 5-star hotels will have no variation in the hotel star attribute, enabling consumers to easily understand the star attribute. If the consumer further sorts the result set on price, the search-result page will be well sorted on price attributes, enabling consumers to easily compare products in the price dimension. This example calls for our information complexity measure, which is meant to capture the cost of *comprehending* and *comparing* a list of options with different product attribute values. When many products are available to choose from, comparing a randomly juxtaposed list of products is more costly than comparing a well-sorted list. Note that our information complexity measure will not capture the cognitive effort required to comprehend a single product, such as when consumers are reading product-detail pages, and we leave that component in the residual search cost for product clicking.

Finally, we assume consumers rely on the expected search costs to guide their search decisions, and they have rational expectations about time and cognitive costs. In terms of loading time, because consumers can fairly easily get a sense of the platform response time after the first one or few pages they visit, we assume each consumer forms her own session-specific expectations about the loading time of refinement and product-detail pages. Because the computation required to construct different refinements on the server is the same, the expected loading time for different refinements within a session will be assumed the same. The same holds true for product-detail pages, and the expected loading time for different product-detail pages will be assumed the same. Furthermore, the consumers expect to incur the same loading time cost when they revisit a refinement due to the platform design that prohibits caching of refinement content. In terms of cognitive cost, since a given search query may return different refinement content over time due to the variation of product availability and attribute values, we assume consumers' rational expectations for each search query reflect the average degree of information complexity of the returned refinement pages over time.

In summary, we have the following types of search costs. To decide whether to visit a particular refinement, consumers will account for the session-specific expected loading time cost, the expected cognitive cost to understand a given refinement, and the residual search cost for refinement visiting. To decide whether to click on a particular product, consumers will account for the session-specific expected loading time cost of the product-detail page, the product position cost of going down the search-result page to click on the link (as studied in the previous literature), and the residual search cost for product clicking.

Instead of describing the conceptualization of the consumer search process by writing down mathematical models at the first encounter, we illustrate the conceptualization by walking through a hypothetical search scenario. During this hypothetical search, we lay out the information consumers acquire and the costs they incur at each step, using the diagrams in Figure 3. We will formalize the conceptualization as a model in Section 7. We also introduce some terminology we will use throughout the paper.

3.1 Hypothetical Search Scenario

Figure 3 shows 10 steps, each of which is depicted in a row. Suppose three refinements, Ref1, Ref2 and Ref3, are available, denoted by the rectangles at the top of each row of a step. From the consumers' perspective, a refinement is defined by the associated search queries, which include search filters and sorting variables (together with pagination for empirical work as we explain below). We use "search a query," "visit a refinement," and "visit a search-result page" interchangeably. Each refinement displays a certain number of search results for users to further click on, which, in our context, is 25 hotels. These hotels are denoted by the circles at the bottom and are linked to the refinement they belong to, using an edge. We use "click on a link" and "sample a product" interchangeably.

First the consumer arrives at the default refinement Ref1 in step (1). She observes the available products in Ref1 and the partially revealed product information x_i^I, but does not know the unrevealed information x_i^N for product $i = 1, \ldots, 25$. The search costs involved in step (1) include the expected time cost of waiting for the search-result page to load, the expected cognitive cost of understanding the page, and any related residual cost. In step (2), the consumer can choose to search either one of the 25 partially revealed hotels on Ref1, or the unvisited refinements Ref2 and Ref3. These search options are denoted in yellow in Figure 3. Because she has not visited Ref2 and Ref3, she *does not* know the product availability of these refinements and *cannot* sample any product in those refinements directly.

In step (3), suppose the consumer decides to sample hotel 1, which reveals the hidden information x_1^N. The costs for step (3) include the expected time cost of waiting for the product-detail page to load, the position cost, and any related residual cost. The consumer ponders her search options again in step (4), which include the unsampled products on Ref1 and the unvisited refinements Ref2 and Ref3, all denoted in yellow in Figure 3. Suppose she decides to switch to Ref3 in step (5). At that point, she observes the product availability of Ref3 and the partially revealed product information. She incurs the same cost of visiting a refinement in step (5) as in step (1). Now all the information contained in Ref3 is known to the consumer, and Ref3 is no longer a search option; instead, all its related products become search options, because they still have uncertainty to be resolved. After consideration in step (6), she samples hotel 25 of Ref3 in step (7) to learn the unrevealed product attributes, and incurs the usual clicking associated costs as in step (3). When she considers her search options in step (8), even though she is at Ref3, she still has unsampled products in Ref1 as her search options. She decides to switch back to Ref1 and samples hotel 2 in step (9). She learns the unrevealed product attributes and incurs the usual clicking cost as in step (3). Moreover, she incurs the time cost of loading Ref1, but she does not need to incur the other associated costs of visiting a refinement, including the cognitive and residual cost, because she already understands the refinement. Finally, she decides not to search anymore and to purchase hotel 25 of Ref3 among all the sampled products.

In addition, in our empirical context, one hotel could satisfy more than one search queries and thus appear in different refinement pages. In Section A of the online appendix, we present additional diagrams to illustrate the consumer search process under that scenario.

Moreover, in our empirical work, we define refinements by search queries that include not only search filters and sorting variables, but also pagination. The rationale for this level of granularity is as follows. Typical search queries return many more results than the maximum number of 25 hotels that can be displayed on a search-result page. As a result, all the search results will be split into different pages, where the product attribute distributions are different on different pages due to sorting. For example, if a consumer is price sensitive, she may sort the search results by price in ascending order. The product attribute distribution on different pagination will be different, and this difference has behavioral implications whereby the consumer is more likely to look at earlier than later pages. By defining refinements with consideration of pagination as well, we can control for different product attribute distributions as pagination increases. Furthermore, our empirical model can account for refinements that have fewer results than the maximum number of 25 hotels that can be displayed on a search-result page. We just use 25 hotels in our illustration for clarity.

This search process is the basis of our search model and we will detail the model specification in Section 7.

4 Cost Measures

In Section 3, we described the different search costs consumers incur during the search process. Four basic types of search costs exist: time, cognitive, position, and residual costs. We can measure the time cost of waiting for the search-result and product-detail pages to load, using their respective

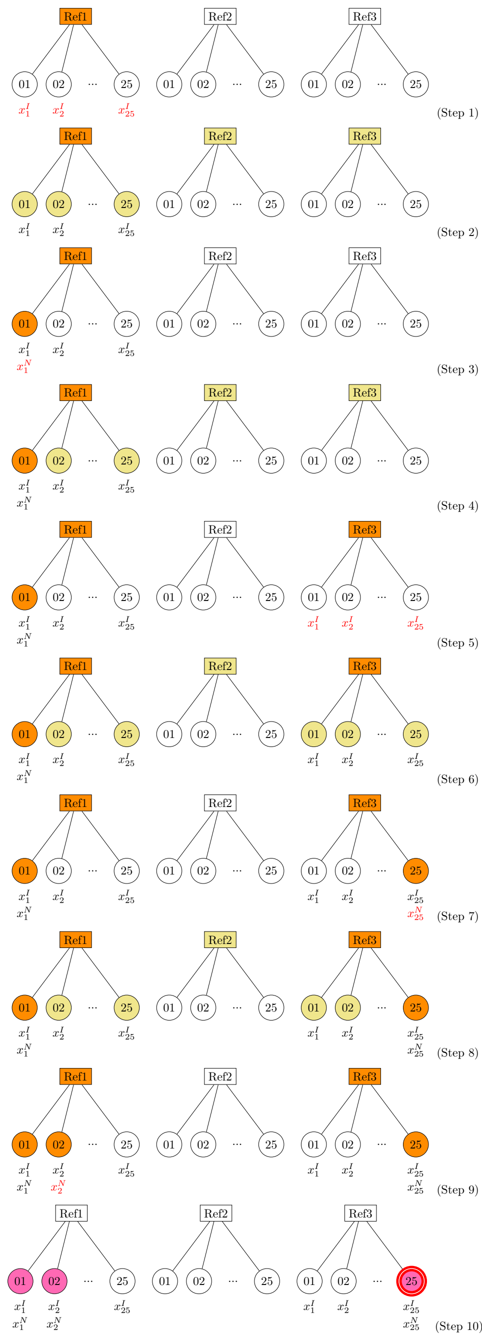

Figure 3: Illustration of the consumer search process in the context of three available refinements, each with 25 products to sample. The consumer arrives at the default refinement Ref1 in step (1), and then decides the search option to proceed with among the 25 partially revealed products on Ref1 and the unvisited refinements Ref2 and Ref3 in step (2). Subsequently, she samples product 1 in step (3), considers again in step (4), switches to Ref3 in step (5), considers in step (6), samples product 25 of Ref3 in step (7), and considers again in step (8). At this point, even though she is at Ref3, she still has unsampled products in Ref1 as search options. She decides to switch back to Ref1 and samples product 2 in step (9). Finally, she decides to stop sampling and to purchase product 25 of Ref3.

loading time. We use the observed position of the products to measure their position costs. The residual costs capture all the remaining costs not captured by the observed covariates. Because no quantification of the cognitive costs exists, we create an information complexity measure of a search-result page to help gauge the cognitive cost.

4.1 Information Complexity Measure

The information complexity measure is based on the entropy of orderedness of the product attributes on the refinement. Recall that information complexity measure, which, together with consumer responsiveness to it, accounts for cognitive cost, is meant to capture the cost of *comprehending* and *comparing* different product attributes.

Information entropy, first introduced by Shannon (1948) and has fundamental importance in communication theory, captures the amount of uncertainty of a random variable. The more random and less structured a variable is, the higher entropy it has.[3] Specifically, the entropy for a discrete random variable Y with probability mass function $P(y)$ on the support $\{y_1, ..., y_n\}$ is defined as $H(Y) = -\sum_{i=1}^{n} P(y_i)\log_2 P(y_i)$ with units in bits. To get a sense of this mathematical construct, consider the scenario of flipping a biased coin with a 90% chance being head, another biased coin with a 70% probability being tail, and finally a fair coin. It is quite obvious that the outcome for the first coin flip is more certain than the second and the third, and the fair coin provides the largest uncertainty among the three coins. If we apply the Shannon Entropy formula, the entropies for the three coin flips are respectively 0.47, 0.88 and 1.0 bits, which aligns with our intuition of uncertainty. In the empirical context, we compute the entropy for a sequence of attribute values by using the empirical probability distribution, $\hat{P}(y_1), ..., \hat{P}(y_n)$. For example, the sequence, $1, 1, 1, 1, 2, 2, 2, 2, 2, 2,$ has empirical probability mass $\hat{P}(1) = 0.4$ and $\hat{P}(2) = 0.6$, and thus its entropy is 0.97 bits. For continuous random variables in the empirical context, we first discretize them and then apply the same computation.

It is helpful to know that Shannon Entropy has the following desirable properties:

- First, Shannon Entropy is non-negative for any random variable. Thus, it is proper to use

[3]Entropy can also be interpreted as a measure of the amount of information. In particular, entropy can be considered as the average minimum description length of a random variable. Intuitively, entropy measures how many binary inquiries one needs to ask the person who knows the realization from a random variable to learn its value, and the number of such binary inquiries is the description length. In this paper, we instead use another interpretation of entropy as a measure of uncertainty, which fits the goal of quantifying cognitive efforts more closely.

entropy to measure uncertainty and cognitive cost, since both are non-negative quantities by nature.

- Second, Shannon Entropy is largest for random events distributed uniformly, and is zero for deterministic events. This corresponds to our intuition that a uniform random variable has the highest uncertainty and a constant random variable involves no uncertainty.

- Third, uniform distributions with larger support have higher Shannon Entropy. Intuitively, the outcome of rolling a fair die is less certain than the outcome of flipping a fair coin, since rolling a die has more potential outcomes. It is easy to compute that the entropy is 2.58 bits for the former and 1 bit for the latter.

Another prominent feature of Shannon Entropy is that it is invariant to flipping the arguments. For example, in the setting of coin flip, the entropy is the same no matter the head/tail has either 30/70 or 70/30 split of the probability. Thus, in our empirical context, the entropy is the same for the following two value sequences of a product attribute:

$$4, 2, 5, 1, 3, 2, 3, 4, 1, 5 \text{ vs } 1, 1, 2, 2, 3, 3, 4, 4, 5, 5.$$

This is because the empirical probability mass function is the same for the two sequences, namely $\hat{P}(1) = \hat{P}(2) = ... = \hat{P}(5) = 0.2$. By definition, entropy ignores the order of the realizations of a random variable. Therefore, entropy cannot directly quantify the level of complexity involved in understanding a numeric value sequence of a product attribute: although the two sequences above have the same entropy level, understanding the second sequence takes less effort than the first. Computing summary statistics, such as the mean and median, and understanding the frequency distribution are much easier to do for the second sequence than for the first.

Instead, we argue that *orderedness* is an important quantity that matters for the *comprehension* and *comparison* of a product attribute: the more ordered the value sequence of a given product attribute, the less *cognitive* costs necessary to understand the product attribute. Therefore, we first define a new sequence, called an orderedness sequence, based on a given single product attribute. Then, we define information complexity for a single product attribute by combining the concept of entropy together with the orderedness sequence.

Specifically, the orderedness sequence of the numeric sequence x is defined to be

$$ord(x) = (x_{i+1} - x_i)_{i=1}^{m-1}, \tag{1}$$

where m is the length of x. It constructs the first differences of the consecutive values in the original value sequence. This definition can first capture whether a numeric sequence is well sorted or not, which is the *qualitative* aspect of *orderedness*. Consider again the two numeric sequences above, where the second sequence is sorted but otherwise equivalent. The orderedness sequences are computed as

$$-2, 3, -4, 2, -1, 1, 1, -3, 4 \text{ vs } 0, 1, 0, 1, 0, 1, 0, 1, 0.$$

The first orderedness sequence has more uncertainty than the second, which corresponds to the lack of sortedness of the first sequence.

The orderedness sequence can also capture how regular the successive changes are in the original sequence values, that is, the *quantitative* aspect of *orderedness*. For instance, consider the following numeric sequences:

$$1, 4, 5, 7, 11, 23, 27, 31, 32, 36 \text{ vs } 1, 2, 3, 4, 5, 6, 7, 8, 9, 10.$$

Comprehending the second sequence is easier than comprehending the first, even though both are sorted. The rationale for this phenomenon is that a sequence of ordered and evenly spaced values is easier to understand than another ordered yet jumpy sequence. The derived orderedness sequences for this example are

$$3, 1, 2, 4, 12, 4, 4, 1, 4 \text{ vs } 1, 1, 1, 1, 1, 1, 1, 1, 1.$$

We again observe that the first orderedness sequence has more uncertainty than the second.

Given these observations and the definition of orderedness sequence, we complete the definition for the information complexity measure of a value sequence, x, as the Shannon Entropy of the derived orderedness sequence. In particular, the orderedness entropy measure is defined as $H(ord(x))$. This definition fundamentally assumes that the more uncertainty exists in the orderedness sequence, the harder it is to understand the original sequence. Applying this definition, the orderedness entropies for the first pair of sequences are

$$H(ord(4,2,5,1,3,2,3,4,1,5)) = 2.75 \text{ vs } H(ord(1,1,2,2,3,3,4,4,5,5)) = 0.99,$$

and for the second pair of sequences are

$$H(ord(1,4,5,7,11,23,27,31,32,36)) = 2.06 \text{ vs } H(ord(1,2,3,4,5,6,7,8,9,10)) = 0.$$

These patterns align with our intuitions for the information complexity. Furthermore, we define the information complexity measure of a search-result page by summing over the information complexity of its constituent revealed product attributes, all of which are in the unit of bits and hence summable.

A desirable property of orderedness entropy is that this measure is invariant to affine transformation. Consider two sequences below, where the second sequence is two times the first sequence minus one:

$$1,2,3,4,5,6,7,8,9,10 \text{ vs } 1,3,5,7,9,11,13,15,17,19.$$

These two sequences are similarly complex, and thus it is desirable if they have the same amount of orderedness entropy. Indeed, the orderedness sequences are respectively $1,1,1,1,1,1,1,1,1$ and $2,2,2,2,2,2,2,2,2$, and thus their orderedness entropies are both zero. It is appealing that orderedness entropy is not affected by translation and scaling. Other uncertainty measures, such as standard deviation, do not possess this property.

In terms of our empirical context, when consumers rearrange search results by applying the sorting variable to, for example, hotel star, they drastically reduce the complexity of the search-result page in that dimension. Similarly, applying filtering variables will reduce the complexity in those dimensions. For example, filtering hotels such that only 4-star hotels remain in the search result will reduce the information complexity of that dimension to 0. Furthermore, due to the correlation among product attributes, sorting or filtering a subset of the product attributes will change the information complexity of the remaining product attributes, and thus the overall information complexity of the search-result page. We will show empirically in Section 6 that applying sorting and filtering variables, in general, reduces information complexity. More importantly, we show that people actively apply sorting variables as a response to information complexity using both reduced-form and experimental evidence.

Although we have discussed the desirable properties of orderedness entropy as a measure of information complexity, we would like to remind the cautious readers to not consider this measure as a universal truth grounded in information theory. We have created this measure to proxy the cost of comprehending and comparing product attributes on refinement pages, but certainly there are scenarios where the consumer's behaviors are inconsistent with the specified cost functions, and other proxies may perform better. We discuss the various limitations of the orderedness entropy measure in Section B of the online appendix, while further enhancements await future research.

5 Data

In addition to the OTA description we provide in Section 2, we observe a detailed panel of consumer interactions with the site, including the search queries used to refine the search results, the content and loading time of the search-result and product-detail pages, the refinements visited, the clicks made, and the rooms purchased. The data set we use covers 29,065 randomly selected consumers who visited the site searching for hotels in six cities in China (Beijing, Shanghai, Guangzhou, Hangzhou, Chengdu, and Lijiang) from Jan 27, 2014, to May 29, 2014.

5.1 Data Description

Some key summary statistics of the data are presented in Table 1. The median number of hotels in the typical search-result set is 171, which need to be split into several pages to be displayed. We have removed search queries that provide no results. The number of hotel bookings made by consumers in the observed window is slightly more than 1 on average, and the maximum number made by a single consumer is 42 times. A consumer may have several search sessions, any of which may result in no hotel booking. Our data do not have consumers who have not made any purchase in any search session.

In Table 2, we show the number of search-result pages visited and the number of product-detail pages clicked for each consumer session in different cities. We can see that the distribution for the number of search-result pages visited and the number of product-detail pages clicked are quite similar across cities.

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.	SD
Search-Result Size	1	66	171	908	1550	4896	1271.5
Number of Hotel Bookings per Consumer	1	1	1	1.25	1	42	0.922

Table 1: Descriptive Statistics. Search-result size counts the number of matching hotels for a given search query, which may need to be split into several pages to display. Number of hotel bookings per consumer counts the total number of bookings a consumer makes over multiple search sessions in our observation window.

# Search-Result Pages Visited	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.	SD
Beijing	1	1	1	2.15	2	52	2.16
Shanghai	1	1	1	2.21	2	62	2.31
Guangzhou	1	1	2	2.34	3	43	2.35
Hangzhou	1	1	2	2.38	3	61	2.21
Chengdu	1	1	2	2.44	3	42	2.3
Lijiang	1	1	2	2.43	3	30	2.09
# Product-Detail Pages Clicked	**Min.**	**1st Qu.**	**Median**	**Mean**	**3rd Qu.**	**Max.**	**SD**
Beijing	0	0	0	0.75	1	55	1.64
Shanghai	0	0	0	0.77	1	78	1.79
Guangzhou	0	0	0	0.85	1	33	1.63
Hangzhou	0	0	0	1.06	1	186	2.37
Chengdu	0	0	0	1.08	1	55	2.05
Lijiang	0	0	1	1.5	2	60	2.76

Table 2: Distribution of the number of search-result pages and product-detail pages visited for each consumer session in different cities.

A noticeable and consistent phenomenon across different cities is that consumers tend to visit more search-result pages than to click on individual products. The implication is that consumers actively use search-result pages as an important source of information, and they only click on a product if they are sufficiently interested in it to learn its hidden product attributes. This finding motivates us to treat refinement choice as an integral endogenous choice in the consumer search process.

5.1.1 Loading Time

Table 3 shows the loading time distribution of product-detail pages and search-result pages in seconds. We observe that the loading time of search-result pages is longer and has more variation than that of product-detail pages, because more substantial data extraction, computation, and rendering are performed at the platform side. For the same reason, the two sources of loading time are not

Loading Time (sec)	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.	SD
Product-Detail Page	0.056	0.931	1.364	2.273	2.182	99.500	4.004
Search-Result Page	0.032	1.547	2.249	3.369	3.504	99.070	4.571

Table 3: Distribution of the loading time of product-detail pages and search-result pages.

highly correlated, with a correlation coefficient of 0.163. The variation in the loading time is also large, with some outlying sessions with very long loading time. We have removed sessions with loading time longer than 100 seconds, which usually correspond to failed connections due to either consumer or platform side errors.

5.1.2 Orderedness Entropy

In our empirical context, we have both numeric and binary product attributes. We convert binary product attributes to numeric values by setting one level to 1 and the other to 0. This conversion is without loss of generality due to the invariance property of orderedness entropy to translation and scaling. All the product attributes are discrete, except that price and consumer rating of the hotels are continuous product attributes. We discretize price by rounding up in bins of 5 RMB and consumer rating in bins of 0.1 points. The chosen bin sizes are the smallest amount that people discern a difference of these variables in daily life.

For all the discrete revealed attributes, we sum up their orderedness entropies, because they have the common natural unit, namely, bits. To be extra cautious in the discretization step, we keep the orderedness entropies of the continuous attributes separate, allowing their coefficients of the cognitive search cost to be different. We have more discussions on the pros and cons of summing up the discrete-features orderedness entropies in Section 7.2. The unit of orderedness entropy does not have an immediate economic equivalent unit. However, our structural model will transform bits into equivalent units of time or monetary value on par with other variables that enter consumer decisions.

The distributions of orderedness entropies of price, consumer rating, and discrete revealed attributes are shown in Table 4. The minimum orderedness entropy being 0 comes from very small result sets that have no variation in the corresponding attributes. The interquartile range of the distributions is more reflective of the typical range of orderedness entropies.

(bits)	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.	SD
Price	0.000	3.255	3.509	3.384	3.689	4.189	0.512
Consumer Rating	0.000	1.041	1.142	1.128	1.241	2.925	0.248
Discrete Revealed Attributes	0.000	4.338	4.913	4.857	5.462	11.330	1.020

Table 4: Distribution of orderedness entropies of price, consumer rating, and discrete revealed attributes.

To better understand orderedness entropy, we compare the consequences of adding search filters as well as applying sorting variables to any given search-result page. We show, based on the following analysis, that information complexity decreases on average as search filters are added or sorting variables other than the default ranking are applied. Furthermore, we show, in the next section, that consumers proactively reduce information complexity during their search process. These findings suggest that cognitive cost resulting from information complexity is an important component for studying consumer's refinement choice.

First, within each consumer search session, we identify all the pairs of search-result pages with the addition of search filters, while keeping sorting variables fixed. We find that, after the addition, the orderedness entropies of price, consumer rating, and discrete attributes combined decrease, on average, by 0.289, 0.029, and 0.363 bits, respectively, with p-value < 0.001. These values correspond to 8.5%, 2.6%, and 5.3% decrease, respectively, when compared to the mean levels.

Second, within each consumer search session, we identify all the pairs of search-result pages that switch from sorting based on the default ranking to either sorting on price, consumer rating, or hotel star, while keeping the search filters fixed. We compare the differences in the orderedness entropy before and after the switching of sorting variables in Table 5. First, all the changes are significant with p-value < 0.001. Second, the largest percentage decreases in orderedness entropies are on the diagonal, which correspond to the variables being sorted. This finding shows that sorting is effective in simplifying the search-result page in the dimension being sorted. Third, due to the correlation structure in the product attributes, when some variables become sorted, some other variables may become unsorted. For example, when sorting search-result pages based on price, the consumer rating variable becomes more complex than sorting based on the default ranking. Fourth, the overall orderedness entropies decrease when we apply any of the sorting variables on price, consumer rating, or hotel star. This finding shows that sorting is also effective in simplifying the

(bits)	Add a Filter		Sort on Price		Sort on Consumer Rating		Sort on Hotel Star	
Price	-0.289	(-8.5%)	-0.461	(-13.6%)	0.033	(1.0%)	-0.247	(-7.3%)
Consumer Rating	-0.029	(-2.6%)	0.147	(13.0%)	-0.331	(-29.3%)	0.143	(12.7%)
Discrete Revealed Attributes	-0.363	(-5.3%)	-0.713	(-10.4%)	0.032	(0.5%)	-0.697	(-10.2%)

Table 5: Within each consumer search session, we first identify all the pairs of search-result pages with the addition of search filters, while keeping sorting variables fixed. Then, we identify all the pairs of search-result pages that switch from sorting based on the default ranking to either sorting on price, consumer rating, or hotel star, while keeping the search filters fixed. This table shows the differences in orderedness entropy before and after the adding of search filters or switching of sorting variables, all significant with probability < 0.001. The percentage change is compared to the mean levels of orderedness entropy of each attribute.

overall search-result page.

6 Model-Free Evidence

In this section, we first empirically show that people respond to longer loading time by reducing their search and clicking. This finding establishes the basic empirical observation that clicking and searching is beneficial to the consumer as a way of information acquisition, and loading time is one of the relevant costs of search. Second, we confirm that orderedness entropy is a good measure of information complexity and a good proxy for the cognitive cost of search in our context, by showing consumers are more likely to sort when sorting induces a higher reduction in orderedness entropy. In Section C of the online appendix, we design a website and implement two online random experiments to further reinforce the reduced-form findings. We find supportive evidence that consumers consider the type of information complexity that orderedness entropy represents, and their sorting decisions incorporate the orderedness entropy reduction. The experimental findings can also be of independent interests to the researchers who study consumer choice in other empirical contexts beyond consumer search.

6.1 The Effects of Loading Time on Number of Searches and Clicks

We run a Poisson regression of the number of result pages searched and the number of product-detail pages clicked on the loading time of the result and detail pages, respectively, controlling for individual, city, and time (day, morning, and afternoon) fixed effects. Table 6 shows significant consumer responses to search and click loading time. For search-result pages, a 1-second increase in

	Number of Searches	Number of Clicks
Search Loading Time	-0.094∗∗∗ (0.008)	
Click Loading Time		-0.047∗∗∗ (0.007)
Individual Fixed Effect	Yes	Yes
Dest. City Fixed Effect	Yes	Yes
Day Fixed Effect	Yes	Yes
Morning Fixed Effect	Yes	Yes
Afternoon Fixed Effect	Yes	Yes
Number of Observations	35,545	35,545

Table 6: Poisson regression of the number of search-result pages searched and the number of product-detail pages clicked on the loading time of the search-result and product-detail pages, respectively, controlling for individual, city, and time (day, morning, and afternoon) fixed effects. Standard errors in parentheses. $*p < 0.1$; $**p < 0.05$; $***p < 0.01$.

the search-result page loading time leads to a 9.4% decrease in the number of searches conducted. At the mean number of 2.77 searches, that percentage amounts to a reduction of 0.26 searches. For product-detail pages, a 1-second increase in the product-detail page loading time leads to 4.7% decrease in the number of product-detail pages clicked, which amounts to a reduction of 0.05 clicks at the mean number of 1.02 clicks.

6.2 The Effects of Orderedness Entropy on Consumer Choice of Sorting Variable

In the previous section, we showed that adding search filters and changing sorting variables change the orderedness entropy. In this section, we further show that consumers respond to the cognitive cost of search.

The purpose for consumers to use refinements is possibly two-fold: one is to change the set of hotels on the refinement page to better fit consumers' preferences, and the other is to reduce cognitive costs by making the search results more orderly. To nail the channel of reducing cognitive costs using reduced-form evidence, we look at decision occasions in which consumers use refinements to only change information complexity but not the content of the search-result page. As a result, we look at whether consumers respond to information complexity by looking at their choices of sorting variables but not filtering variables, because filtering by definition changes the content of the search-result page. We also need to be careful with the choice of sorting variables, because sorting a search-result set larger than 25 will also change the content on a particular page. Therefore, we

Intercept (Price)	-3.997***	(0.000)
Intercept (Consumer Rating)	-3.567***	(0.000)
Intercept (Hotel Star)	-5.477***	(0.000)
Reduction of Orderedness Entropy of Price	0.965**	(0.029)
Reduction of Orderedness Entropy of Consumer Rating	0.493*	(0.065)
Reduction of Orderedness Entropy of Hotel Star	0.696***	(0.000)
Number of Observations	8,184	

Table 7: Multinomial logit regression of the choice probability of sorting on default ranking, price, consumer rating, or hotel star as a response to the change in the orderedness entropies. The positive coefficients in front of the reduction of orderedness entropies show that a sorting variable is more likely to be chosen if that leads to a larger reduction in the orderedness entropy of the product attributes. Standard errors in parentheses. $*p < 0.1$; $**p < 0.05$; $***p < 0.01$.

look at searches in which the result set can be fit on a single page. In this case, changing the sorting variable will not change the content on the refinement but the orderedness of the search results.

In particular, we identify 8,184 consumer searches that use the default ranking and result in fewer than 25 products. Consumers can choose to switch the sorting variable to price, consumer rating, or hotel star, or they can choose not to switch. We observe switches to sorting on price 165 times (2.0%), sorting on consumer rating 253 times (3.1%), and sorting on hotel star 43 times (0.5%).

Given the above empirical setup, we study how much more likely consumers are to switch the sorting variable away from the default ranking given how much the switching reduces information complexity. In particular, we use a multinomial logit regression to study the choice probability of sorting on $j \in SV = \{\text{Default Popularity, Price, Consumer Rating, Hotel Star}\}$ as a response to the change in the orderedness entropies, computed by $\left(\text{Entropy}_{Default,v} - \text{Entropy}_{j,v}\right)$ for $v \in KV = \{\text{Price, Consumer Rating, Discrete Revealed Attributes}\}$:

$$P(\text{sort on } j) = \frac{\exp\left(\text{Intercept}_j + \sum_{v \in KV} \beta_v \left(\text{Entropy}_{Default,v} - \text{Entropy}_{j,v}\right)\right)}{\sum_{k \in SV} \exp\left(\text{Intercept}_k + \sum_{v \in KV} \beta_v \left(\text{Entropy}_{Default,v} - \text{Entropy}_{k,v}\right)\right)}. \tag{2}$$

Table 7 shows that the coefficients β_v in front of the reduction of orderedness entropy $\left(\text{Entropy}_{Default,v} - \text{Entropy}_{j,v}\right)$ are positive for all $v \in KV$. Thus when the entropy reduction along one attribute dimension of the sorting option j increases, ceteris paribus, the consumer is more likely to choose that sorting option. The choice-specific intercepts are all negative, absorbing the loading time cost of switching away from the default ranking.

7 Model

In this section, we formalize the conceptualization of the consumer search process specified in Section 3. We will first describe the components of the consumer product utility and search cost. Then we will lay out the consumer information and cost structure during the search process. Finally, we will specify the decision model for the search process.

7.1 Components of the Consumer Utility and Search Cost

The product attribute is composed of a revealed part displayed on the refinement $x^I \in \mathbb{R}^n$ and an unrevealed part $x^N \in \mathbb{R}^m$ displayed after clicking on the product-detail page. Because the attribute values for a given product may change over search sessions due to their dynamic nature, we use se to denote a search session.

Consumer h's utility for product j at search session se is given by

$$\begin{aligned} u_{hj,se} &= V^I_{hj,se} + V^N_{hj,se} + \epsilon_{hj,se} \\ &= \beta^I_h x^I_{j,se} + \beta^N_h x^N_{j,se} + \epsilon_{hj,se}, \end{aligned} \tag{3}$$

where β^I_h and β^N_h are the consumer-specific preference coefficients for the revealed and unrevealed product features, respectively, and $\epsilon_{hj,se}$ is the consumer's session-specific utility shock. We use the superscript I to denote the quantities relating to the revealed product attributes (i.e., $V^I_{hj,se}$, β^I_h and $x^I_{j,se}$), and the superscript N to denote those relating to the unrevealed ones (i.e., $V^N_{hj,se}$, β^N_h and $x^N_{j,se}$). The session-specific utility of the outside option $j = 0$ is assumed to be known by the consumer from the beginning. $\epsilon_{hj,se}$ is unknown to the researcher.

Following our assumptions on consumers having rational expectations in Section 3, consumer h has rational expectations about the current search session's refinement loading time $SearchLoad_{h,se}$ and the product-detail page loading time $DetailLoad_{h,se}$, both computed as the sample average of the loading time of the visited refinement pages and product-detail pages during her search session se. For refinement query rf, the consumer also has rational expectations about the orderedness entropy of price $EntrPrice_{rf}$, of consumer rating $EntrRating_{rf}$, and of discrete revealed attributes $EntrDiscrete_{rf}$. Since one search query (composed of the filtering variables, sorting variables, and pagination) may return different refinement pages over time due to the dynamics of product availability and attribute values, we assume each of these quantities is computed as the sample average of the respective orderedness entropy of the realized refinement pages generated by query rf over time. Moreover, the consumer also observes the product position $R_{j,se}$ on the page.

7.2 Consumer Information and Cost Structure

Our model will capture the consumer information and cost structure as follows:

- **Before visiting a refinement** rf, a consumer h forms a belief on the utility distribution of *visiting* refinement rf, which is denoted as $U_{h,rf} \sim F_{h,rf}$, while she does not know the product availability on this refinement yet. As a result, the unvisited refinement rf is a search option for the consumer, but products on this refinement are not. Also, the consumer has not observed any information within the refinement, such as the partially revealed product information, to form a belief conditional on this information. At this point, we need to make a modeling decision regarding the utility distribution of visiting a refinement for the consumer. We follow the same distribution assumption adopted by De los Santos and Koulayev (2017), which assumes that $F_{h,rf}$ is a normal distribution and its first and second moments reflect the variation of the maximum product utility on the refinement. In particular, since the product attribute values and product availabilities on the platform change over time, the realized refinement content for a given refinement query also changes over time. This is the source of the uncertainty of visiting a refinement. Following De los Santos and Koulayev (2017), the uncertainty could be reasonably proxied by the distribution of the first order statistic, since, when visiting a refinement, the consumer only cares about the product with the highest utility on the refinement, not the rest. As a result, we collect all the realized refinement pages for a given query rf (which is composed of the filtering variables, sorting variables, and pagination), and we compute the sample mean and variance of the maximum product utilities on the refinement pages, i.e., $\hat{E}\left(\max_{j\in rf}\left\{\beta_h^I x_{j,se}^I + \beta_h^N x_{j,se}^N\right\}\right)$ and $\hat{Var}\left(\max_{j\in rf}\left\{\beta_h^I x_{j,se}^I + \beta_h^N x_{j,se}^N\right\}\right)$, respectively. They are the mean and variance of $F_{h,rf}$, the belief of the utility distribution of visiting refinement rf for consumer h. We will discuss this assumption in more details when we present the decision model.

In terms of the cost to search the unvisited refinement page, the consumer pays the expected refinement loading time cost, the expected cognitive cost, and the refinement residual cost:

$$\begin{aligned} c_{R,h,se}(rf) &= \gamma_{h,Search} + \gamma_{h1} SearchLoad_{h,se} + \\ &\quad \gamma_{h2,1} EntrPrice_{rf} + \gamma_{h2,2} EntrRating_{rf} + \gamma_{h2,3} EntrDiscrete_{rf} + \\ &\quad \gamma_{h2,4} EntrPrice^2_{rf} + \gamma_{h2,5} EntrRating^2_{rf} + \gamma_{h2,6} EntrDiscrete^2_{rf}. \end{aligned} \tag{4}$$

We have included the second-order terms of the orderedness entropy measures. In the literature of neoclassical economics, a wide range of empirical evidence has been found to support the hypothesis that the cost functions are generally convex with increasing marginal costs. In practice, the consumers may choose to not digest so much information when the information complexity is very high, which suggests that the refinement cost could be a concave (or less convex) function in orderedness entropy. Thus, the second order terms are useful and flexible enough to allow the incremental increase in cognitive cost to be decreasing (or less increasing) in orderedness entropy.

In addition, we have separate coefficients for the orderedness entropies of the continuous attributes (namely price and consumer rating), as their coefficients depend on the discretization intervals. Another reason for having separate coefficients for price and rating is that consumers may want to spend more efforts on understanding the product features that they care more about. Needless to say, price is particularly important as is found in many empirical studies. Consumer rating is also discovered to be very important in recent literature, as discussed in, for example, Chen and Yao (2016) that "consumer ratings have on average the highest impact on consumer utility." Thus, we use separate coefficients for price and rating to allow for this flexibility.[4]

Finally, the refinement residual cost, $\gamma_{h,Search}$ (also known as the baseline cost), captures the

[4]In principle, we would like to have separate cognitive cost coefficients for each revealed discrete attributes to account for their different importance. However, this is too demanding for estimation. In the model, an orderedness entropy measure is associated with 4 unknown coefficients, namely the mean and standard deviation of the random coefficients for the first and second order terms of the measure. As there are 12 revealed discrete attributes in the empirical context, there are 48 ($= 12 \times 4$) unknown coefficients to estimate if separate coefficients are allowed. The orderedness entropies for the discrete attributes have the same units of bits and thus are summable. We decide to pool them together to not lose too much statistical power in the estimation.

remaining psycho-physiological costs related to refinement visiting, which are not captured by the existing observed covariates.

- **After visiting a given refinement**, the consumer knows the product availability on the refinement and the partially revealed product attributes on the refinement. As a result, the visited refinement page is no longer a search option for the consumer, but the unsampled products on this page are. Furthermore, the unsampled products on all previously visited refinement pages are also search options for the consumer. More specifically, the consumer observes the partially revealed product attributes $x^I_{j,se}$ for product j in search session se, and the consumer forms conditional belief on the utility distribution for product j, that is $\left\{u_{hj,se}|x^I_{j,se}\right\} \sim G_h$. Recall that utility $u_{hj,se}$ contains both the revealed component, $\beta^I_h x^I_{j,se}$, and the unrevealed component, $\beta^N_h x^N_{j,se}+\epsilon_{hj,se}$. We assume that $\left\{u_{hj,se}|x^I_{j,se}\right\}$ follows a normal distribution, whose mean and variance depend on how the consumer predicts the unrevealed product attribute values. In particular, we assume that the consumer forms predictions of the unrevealed product attributes through a multivariate linear regression which regresses the unrevealed attributes on the revealed ones:

$$x^N_{j,se} = \phi + \Gamma x^I_{j,se} + \eta_{j,se}.$$

The prediction rule is assumed linear for simplicity. The multivariate regression estimation provides the estimated intercept vector $\hat{\phi}$, the estimated coefficients matrix $\hat{\Gamma}$, and the estimated residuals $\hat{\eta}_{j,se}$. Therefore, it is natural to have the mean of the conditional utility distribution, $G_h\left(u_{hj,se}|x^I_{j,se}\right)$, as $\beta^I_h x^I_{j,se} + \beta^N_h\left(\hat{\phi} + \hat{\Gamma}x^I_{j,se}\right)$, and the variance as $\beta^N_h \hat{Cov}\left(\hat{\eta}_{j,se}\right)\beta^{N'}_h + \sigma^2_\epsilon$, where $\hat{Cov}\left(\hat{\eta}_{j,se}\right)$ is the sample variance-covariance matrix of the regression residuals $\hat{\eta}_{j,se}$, and σ^2_ϵ is the consumer's belief on the variance of the idiosyncratic utility shock, $\epsilon_{hj,se}$, which is assumed to be the same as the true variance of $\epsilon_{hj,se}$. To estimate the multivariate regression, we collect all the products on all the refinements and search sessions.[5]

[5]The simplification assumption for the above prediction rule is that it is homogeneous across refinements. We assume that a single prediction rule captures the main underlying mental reasoning for consumers, since otherwise it would be too demanding to assume that the consumer forms a large number of prediction rules over the large number of possible refinements. Another practical concern behind this simplification is that the unknown coefficients, ϕ and Γ, are large-dimensional due to the high dimensionalities of the revealed and unrevealed product attributes. We would lose lots of precisions if we separately estimate the prediction coefficients for each refinement.

Also, although the prediction rule for the unrevealed product attributes is assumed to be homogeneous across

In terms of the cost to click the unsampled products,

– To click on an unsampled product on the *current* refinement page where the consumer is at, the consumer pays the expected product-detail page loading time cost, the observed position cost, and the product-detail page residual cost:

$$c_{P,h,se}(rf, j) = \gamma_{h,Detail} + \gamma_{h1} DetailLoad_{h,se} + \gamma_{h3} R_{j,se}. \tag{5}$$

The position cost captures the lower tendency for the lower-ranked products to be clicked due to the additional cost of scrolling down. The product-detail page residual cost, $\gamma_{h,Detail}$ (also known as the baseline cost), captures the remaining psycho-physiological costs associated with product clicking, which are not captured by the existing observed covariates.

– To click on an unsampled product on a *previously visited* refinement page rf', the consumer pays the additional expected refinement page loading time cost to switch refinement:

$$c_{P,h,se}(rf, j) = \gamma_{h,Detail} + \gamma_{h1} \left(SearchLoad_{h,se} + DetailLoad_{h,se}\right) + \gamma_{h3} R_{j,se}. \tag{6}$$

In the clicking cost function $c_{P,h,se}(rf, j)$, the argument rf indicates the current refinement, and the argument j is an index for product j, which contains all the relevant information for product j, including the refinement $rf_{j,se}$ of product j and the position $R_{j,se}$ of product j in session se. When the refinement of product j is different from the current refinement ($rf_{j,se} \neq rf$), the consumer pays the additional cost for switching refinement ($SearchLoad_{h,se}$).

- For product j on a refinement page the consumer has visited and clicked, the consumer knows all the product attributes $x^I_{j,se}$ and $x^N_{j,se}$ and their idiosyncratic utility shocks $\epsilon_{hj,se}$; that is, the consumer knows $u_{hj,se}$. The sampled products are no longer search options, but are candidates

consumers, the prediction for the unrevealed product utility distribution depends on consumer-specific preference coefficients and hence is heterogeneous. Note that the preference coefficients β^N_h affect the option value of clicking on a product through both the mean and variance of the conditional utility distribution, and hence will explain the different clicking behaviors across consumers.

for purchase decision. The costs associated with the purchase decisions are the observed prices of the sampled products.

7.3 Multiple Sequential Search Model

We model the consumer search process using multiple search problems. Each search problem is anchored to the current refinement the consumer is visiting. The consumer starts her search from the default refinement on the search platform.

Within a given refinement, the consumer is assumed to engage in the traditional sequential search in which both the *unvisited refinements* and *unsampled products on the visited refinements* are search options. Furthermore, the model shares the same cost-benefit tradeoff as in Weitzman (1979), where the consumer will evaluate each search option based on the expected gain of search and search cost. This evaluation can be equivalently characterized by the reservation utility of each search option. In particular, for each unsampled product j, the consumer computes its reservation utility $z_{h,rf,j,se}$ as the utility level that makes her indifferent between sampling product j versus choosing the current best offer precisely at the utility level $z_{h,rf,j,se}$:

$$c_{P,h,se}(rf,j) = \int_{z_{h,rf,j,se}}^{\infty} \left(u_{hj,se} - z_{h,rf,j,se}\right) dG_h\left(u_{hj,se}|x^I_{j,se}\right). \tag{7}$$

For each unvisited refinement rf, the consumer computes its reservation utility $z_{h,rf,se}$ similarly as

$$c_{R,h,se}(rf) = \int_{z_{h,rf,se}}^{\infty} \left(U_{h,rf} - z_{h,rf,se}\right) dF_{h,rf}\left(U_{h,rf}\right). \tag{8}$$

During the search process, the consumer will update her current best offer to be the maximum of the realized utility of the sampled products as

$$u^*_{h,se} = \max_{j\in S}\{u_{hj,se}\},$$

where S is the set of sampled products for the consumer. The optimal strategy contains two steps, a selection rule and a stopping rule:

- Selection rule: If sampling a product or visiting a refinement, the consumer picks the search option with the highest reservation utility.

- Stopping rule: The consumer stops the search process if the current best offer exceeds the reservation utilities of all the unvisited refinements and unsampled products on the visited refinements.

When the consumer decides to visit a different refinement from the one she is currently at, we assume the consumer enters a new search problem whereby she needs to evaluate the search options again, because she may observe a new set of products as search options with partially revealed product information, and the search costs of sampling products on previously visited refinements may change. Hence, in our model, switching refinements means switching search problems. When the consumer switches refinements, she faces the new search problem in which she performs the following four transition steps:

1. Carries over the current best offer;

2. Updates her search costs for the unsampled products on the current and last-visited refinements, as described in Section 7.2;

3. Carries over all the search options she has discovered so far, if sampling a product from a *previously visited refinement*;

4. Merges all the search options she has discovered so far with a new set of partially revealed products, if searching an *unvisited refinement*.

Finally, when the consumer stops the search process, she chooses either a sampled product or the outside option, depending on which one has the highest realized utility.

7.3.1 Discussion

The *multiple* sequential search model is a simplification of a *unified* sequential search model that contains all the search problems within. The difference is that, in the unified model, no simple search strategy for this general dynamic programming problem exists, and the consumer needs to mentally compute the payoff for all possible search paths. We can have a glimpse into this computation challenge by looking at some necessary computation steps needed. To evaluate the period utility of visiting a particular refinement along any search path, the consumer needs to account for the high-dimensional *joint* distribution of all the partially revealed attribute values for an ordered list of products on that refinement. Estimating and integrating over this distribution is very difficult. The consumer needs to repeat this computation for all refinements, along all paths. The curse of dimensionality in the number of search paths and the size of the probability space of the product attributes on refinement pages render this modeling approach infeasible. In the online appendix Section D.1, we present the unified sequential search model with more details of its computation difficulty.

The simplification we adopt for our model is the same one-step-ahead approach as in De los Santos and Koulayev (2017), where the consumer does not mentally hypothesize what happens beyond reaching the new refinement, because that computation depends on too much information unobserved to her, and accounting for such uncertainty is too costly to perform. Instead, the consumer uses a heuristic yet intuitive utility distribution to characterize the uncertain benefit of searching a refinement. When she arrives at a new refinement, the consumer observes new information and updates her search costs, based on which she re-evaluates the search options in that refinement. As a result, we split the unified sequential search problem into multiple related search problems, each anchored to the current refinement the consumer is situated at during her search process. These multiple search problems are separated by the different amount of information the consumer has acquired before and after switching a refinement. When the consumer switches refinements, she switches search problems.

Within each problem, we retain the sequential search structure, because the curse of dimensionality of the form mentioned above does not exist: when the consumer samples products on the current refinement she is situated at, she has already observed product availability and the partially revealed product attributes. Furthermore, Weitzman (1979) has characterized a simple search strategy that only relies on the reservation utility of each product using the conditional belief of the hidden product utility for that product, *independent of other* products. As a result, with the assumption that the consumer uses an intuitive utility distribution for searching a refinement, she treats unvisited refinements in the same way as unsampled products, and Weitzman (1979) gives us a complete characterization of the consumer search process within each search problem and hence overall. In the online appendix Section D.2, we present the multiple sequential search model and discuss its deviations from the unified model in details.

In summary, we only simplify our model structure for the choices the consumer does not have much information about, while still maintaining the overall sequential search framework and model expressiveness. We hope this selective simplification is consistent with our central modeling principle to capture the consumer information and cost structure more realistically, while accounting for model tractability. Our model extends that of De los Santos and Koulayev (2017). It grants us the flexibility to have distinct search costs for refinement visiting and product clicking, and allows the change in the platform design to affect refinement visiting through cognitive costs and product sampling through partially revealed information. This flexibility is crucial to answering our managerial question.

8 Estimation

Based on the model specification in Section 7 and the characterization of the optimal search strategy, we derive the model likelihood and describe the estimation strategy in this section. We defer the discussion on model identification, Monte Carlo simulations, and model fit to Section E of the online appendix.

8.1 Likelihood and Estimation Strategy

For a given consumer h and search session se, suppose she engages $B_{h,se} \geq 1$ search problems and performs $K_{h,se,b} \geq 1$ searches within each bth search problem. For the bth problem and kth search, the model has the following state variables (suppressing subscripts h, se, b, k): the set of unsampled products on the current refinement $\bar{S}$ (S for the sampled counter part), the set of unsampled products on previously visited refinements $\bar{T}$ (T for the sampled counter part), the set of unvisited refinements $\bar{R}$, the current refinement rf, and the current best utility u^*. Note that $u^*_{h,se,b,k} = \max\{\max_{j \in S \cup T} u_{hj,se}, u_{h0,se}\}$, where $u_{h0,se}$ is the outside option utility.

For the bth problem and kth search (suppressing subscripts h, se, b, k), the search options include $\bar{S}$, $\bar{T}$, and $\bar{R}$. Let $J = |\bar{S} \cup \bar{T} \cup \bar{R}|$ denote the total number of search options, let $R(n)$ denote the index of the search options with the nth largest reservation utility, and recall that $z_{R(n)}$ is the reservation utility for the search option that ranks n.

We incorporate the observed consumer search process into the indicator function $1_{h,se}(\beta_h, \gamma_h, \epsilon_{h,se}; data)$ that simultaneously satisfies all the following conditions, specified by Equations (9) - (14), for a given utility shock realization vector of the sampled products $\epsilon_{h,se} = (\epsilon_{h,j,se})_j$. Then the likelihood for the observed consumer search process is

$$L_{h,se}(\beta_h, \gamma_h; data) = \int 1_{h,se}(\beta_h, \gamma_h, \epsilon_{h,se}; data) dF(\epsilon_{h,se}).$$

We suppress subscripts h, se, b, k below, because each quantity pertains to the relevant index in its own context. The conditions for the indicator function $1_{h,se}(\beta_h, \gamma_h, \epsilon_{h,se}; data)$ are specified as follows:

- For the bth search problem before the last one in the search session, that is, $1 \leq b < B$, the consumer does not stop the search process. By definition of our search problem, the last search in the bth search problem must be an option among $\bar{T} \cup \bar{R}$ in order for the consumer to switch search problem, and all the other searches (if $K > 1$) must be among $\bar{S}$. Therefore,

  - Each kth search action out of K searches removes the search option from the unsearched set and implies

$$z_{R(k)} \geq \max_{n=k+1}^{J} z_{R(n)} \tag{9}$$

$$z_{R(k)} \geq u^*. \tag{10}$$

- If the consumer searches an *unvisited refinement* or clicks on a product from a *previously visited refinement*, she switches refinements and starts a new search problem. She performs the four transition steps as described in Section 7.3.

- For the last search problem, namely, $b = B$, the consumer searches in the first K steps among $\bar{S}$, stops at the $(K+1)$-th step, and makes the purchase decision:

  - Each kth search action out of K searches removes the search option from the unsearched set and implies

$$z_{R(k)} \geq \max_{n=k+1}^{J} z_{R(n)} \tag{11}$$

$$z_{R(k)} \geq u^*. \tag{12}$$

  - Stopping at the $(K+1)$-th step implies

$$u^* \geq \max_{m=K+1}^{J} z_{R(m)}. \tag{13}$$

– Purchasing the jth option ($j = 0$ for choosing the outside option) implies

$$u_j = u^*. \tag{14}$$

To estimate the consumers' preference and search cost heterogeneity, we introduce a distribution assumption for β_h and γ_h following the current literature (see, e.g., Chen and Yao (2016) and De los Santos and Koulayev (2017)). In particular, we assume that β_h and γ_h follow $\beta_h \sim N(\beta, \Omega_\beta)$ and $\gamma_h \sim N(\gamma, \Omega_\gamma)$,where Ω_β and Ω_γ are diagonal matrices.

The likelihood of the model, integrated over consumers' heterogeneity and aggregated over all individuals and search sessions, is given by

$$L(\beta, \gamma, \Omega_\beta, \Omega_\gamma; data) = \prod_h \int_{\beta_h} \int_{\gamma_h} \prod_{se} \left[L_{h,se}(\beta_h, \gamma_h; data)\right] dF(\gamma_h) dF(\beta_h). \tag{15}$$

We use simulated maximum likelihood to estimate the model. In particular, for each consumer, h, we draw 50 pairs of (β_h, γ_h) from their normal distributions, and for each pair of (β_h, γ_h), we draw 100 sets of $\epsilon_{h,j,se}$ from the standard normal distribution for the clicked products. We compute the individual likelihood using the simulated draws, and we smooth over the likelihood using a logit-smoothed AR simulator (Hajivassiliou and Ruud (1994)). Finally, we repeat the process for each individual.

9 Results

The estimated preference coefficients can be found in Table 8. We normalize the coefficient estimates to RMB by dividing them by the magnitude of the price coefficient. We find that consumer rating is very influential in the evaluation of products, consistent with the findings of Chen and Yao (2016). A hotel with 5-point rating on average is USD 46.04 more expensive than a hotel with 1-point rating. The estimates of the other preference coefficients also align closely with our intuition. Many of the preference coefficients have considerable heterogeneities, including price, consumer rating, hotel star, breakfast, cancellation allowed, Wi-Fi in room, parking, and outdoor swimming pool.

The estimated search cost coefficients can be found in Table 9. To make the cognitive search cost more interpretable, we convert the cognitive cost consumers encounter in a typical search-result page into a monetary value as follows:

$$\text{Search Cost Coeff (util/bits)} \times \text{Entropy Search Cost (bits)/Price Coeff (util/RMB)}.$$

For an average search-result page, the orderedness entropy of price, consumer rating, and discrete variables are 3.38 bits/page, 1.13 bits/page, and 4.88 bits/page, respectively. Thus, consumers on average pay the cognitive cost of 45.96 RMB, 3.96 RMB, and 94.51 RMB, respectively, per refinement visit, or 144.43 RMB (21.56 USD) in total. We find significant heterogeneity in consumers' reactions to discrete variable entropy. The estimated coefficients to the second order terms suggest that the refinement search cost is convex in the orderedness entropies.

Furthermore, the share of loading time cost in the overall search cost is much smaller. An average search-result page takes 3.37 sec to load, and consumers on average pay the time cost of 0.13 RMB (0.019 USD) per refinement visit. An average product-detail page takes 2.27 sec to load, and consumers on average pay the time cost of 0.09 RMB (0.013 USD) per click. Thus, time cost is estimated to be 135.85 RMB, or 20.28 USD, per hour on average.

Our estimated position effect is 0.45 RMB (0.07 USD), which is equivalent to the cost of 11.32 RMB (1.69 USD) if consumers scroll to the end of a refinement page with 25 products on it. This estimated position effect is close to that discovered by Chen and Yao (2016), who find a multiplicative position effect that inflates the overall search cost from 4.85 USD for the first position to 5.31 USD for the 10th position, totaling approximately 0.05 USD for each position. Both the position effect and clicking baseline cost have large heterogeneities across consumers.

10 Managerial Implications

In this section, we study our opening managerial question of how much information to reveal in the outer layer. To approach this question, we vary the set of product attributes revealed in the outer layer and leave the remaining product attributes to the inner layer. We only perturb whether a product attribute is displayed in the outer or inner layer, without removing any product attribute from the platform. Each such configuration corresponds to one information layout, which is applied

Revealed Attribute	Mean (std. err.)	Heterogeneity (std. err.)	RMB ¥	USD $
Price	**-0.00053 (0.00002)**	**0.00013 (0.00001)**		
Consumer Rating	**0.04122 (0.00656)**	**0.01594 (0.00750)**	77.11	11.51
Hotel Star	**0.01049 (0.00069)**	**0.00268 (0.00064)**	19.63	2.93
Breakfast	**0.02804 (0.00334)**	**0.03864 (0.00978)**	52.47	7.83
Cancel Allowed	**0.00822 (0.00076)**	**0.00456 (0.00114)**	15.38	2.30
Wi-Fi in room	**0.04779 (0.00418)**	**0.02372 (0.00576)**	89.42	13.35
Kingsize bed	**0.03077 (0.01035)**	0.01548 (0.01119)	57.58	8.59
Wi-Fi in public areas	**0.07113 (0.03477)**	0.04611 (0.03137)	133.08	19.86
Parking	**0.02933 (0.00165)**	**0.00782 (0.00200)**	54.87	8.19
Meeting rooms	**0.00898 (0.00287)**	0.00381 (0.00637)	16.80	2.51
Gym	**0.07076 (0.00528)**	0.01014 (0.00598)	132.39	19.76
Restaurant	**0.02277 (0.00460)**	0.01557 (0.00856)	42.60	6.36
Airport pickup service	0.02575 (0.03018)	0.23038 (0.53489)	48.17	7.19
Outdoor swimming pool	**0.15488 (0.00860)**	**0.03797 (0.01214)**	289.78	43.25
Unrevealed Attribute	Mean (std. err.)	Heterogeneity (std. err.)	RMB ¥	USD $
Free toiletries (less than six items)	0.00013 (0.00035)	0.00012 (0.00079)	0.25	0.04
Free toiletries (more than six items)	**0.04935 (0.00812)**	0.02472 (0.01902)	92.33	13.78
Free domestic long-distance calls	**0.05771 (0.00512)**	0.01437 (0.00796)	107.97	16.12
Free bottled water	0.00158 (0.00377)	0.02488 (0.10956)	2.95	0.44
Free map	0.00230 (0.00687)	0.05312 (0.29399)	4.31	0.64
Free newspaper	0.00403 (0.01184)	0.16200 (0.80627)	7.53	1.12
Pets Allowed	0.01830 (0.01155)	0.11224 (0.13049)	34.24	5.11
Free Afternoon Tea	**0.01123 (0.00326)**	0.00595 (0.00460)	21.01	3.14
Free Guide Service	**0.03398 (0.00445)**	0.00832 (0.00570)	63.57	9.49

Table 8: Preference Estimates. The RMB ¥ column is the estimates divided by the absolute value of the price coefficient, and the USD $ column is further converted from RMB to USD by dividing by the currency exchange rate of 6.7. Bold fonts indicate estimates significant at the 95% level.

	Coefficients	Mean (std. err.)	Heterogeneity (std. err.)	RMB ¥	USD $
Refinement	Base	**0.00147 (0.00035)**	0.00037 (0.00116)	2.75	0.41
	Price Entropy	**0.00048 (0.00016)**	0.00017 (0.00043)	0.89	0.13
	Rating Entropy	0.00048 (0.00028)	0.00028 (0.00062)	0.90	0.13
	Discrete Entropy	**0.00026 (0.00005)**	**0.00013 (0.00006)**	0.48	0.07
	Price Entropy (Sq)	**0.00199 (0.00018)**	**0.00039 (0.00013)**	3.72	0.56
	Rating Entropy (Sq)	**0.00122 (0.00044)**	0.00048 (0.00049)	2.28	0.34
	Discrete Entropy (Sq)	**0.00205 (0.00019)**	**0.00071 (0.00009)**	3.84	0.57
Time	Time (Sec)	**0.00002 (0.00001)**	0.00001 (0.00001)	0.04	0.01
Click	Base	**0.04547 (0.00170)**	**0.02186 (0.00229)**	85.08	12.70
	Position	**0.00024 (0.00001)**	**0.00013 (0.00002)**	0.45	0.07

Table 9: Search Cost Estimates. The RMB ¥ column is the estimates divided by the absolute value of the price coefficient, and the USD $ column is further converted from RMB to USD by dividing by the currency exchange rate of 6.7. Bold fonts indicate estimates significant at the 95% level.

to all refinement pages. Then for each information layout, we simulate the various aspects of the consumer search process using our estimated model, including the number of clicks, the number of searches, purchase probability, purchase session utility, purchased product price, revenue, and consumer welfare. We define revenue to be the purchased product price accounting for purchase probability. We also define consumer welfare to be the purchase session utility, which is the utility of the purchased product or the outside option conditional on purchase decision, net of the total accumulated search cost. These behavioral aspects provide us another opportunity to further understand the properties of the model. Our final aim is to study the relationship between consumer welfare and platform revenue, which are often in conflict, and which platforms hope to balance. We will try to characterize the Pareto frontier of this relationship, which will guide us in answering the managerial question. We also include an alternative counterfactual in Section G of the online appendix, where we analyze the impact on consumer search behaviors by making the refinements with default ranking more orderly. We find that, a special kind of sorting, called local sorting, provides a Pareto improving strategy that platforms may consider adopting to increase consumer welfare without losing on platform revenue.

10.1 Counterfactual Setup

Due to the large space of possible configurations (about 8.3 million), we randomly sample information layouts. We first generate layouts that cover the extreme cases, from revealing all attributes to revealing only one attribute. Then we randomly sample 12,000 information layouts. Once we have the general configuration space covered through randomization, we improve the resolution near the Pareto frontier between consumer welfare and platform revenue by adding and removing up to four attributes from the outer layer of the layouts on the sampled Pareto frontier. This procedure allows us to better construct the Pareto frontier.

Fixing a particular information layout and using the estimated model, we simulate 500 runs of the consumer search process by generating random utility shocks. Due to the dynamic nature of product attributes, a given refinement will have different product availability and attribute values at different times of visits. As a result, each simulation run will be based on the randomly sampled refinement realizations. We then average the behavioral aspects mentioned above over the 500 simulation runs as the typical consumer response to the given information layout.

10.2 Counterfactual Results

Figure F.1 in the online appendix shows the various behavioral responses of the consumer search process as information layouts vary.

Here we collate the behavioral responses of revenue and consumer welfare, and focus on their relationship as information complexity varies in Figure 4. Each point corresponds to one information layout, and the color of the point corresponds to the average information complexity of the layout. The red end corresponds to layouts with extremely low information provision, and the blue end to layouts with extremely high provision. Two point clouds exist: the top cloud of round-shaped points corresponds to the layouts when product price is not revealed, and the bottom cloud of diamond-shaped points corresponds to when price is revealed. When product price is not revealed, revenue decreases and consumer welfare increases as information complexity increases in Figure F.1 in the online appendix, and hence we see a downward-sloping trend in the top point cloud. Because revenue changes less when product price is revealed, the same pattern holds with a smaller slope in the bottom cloud. The grey line traces out the Pareto frontier of revenue versus consumer welfare. We find a general tradeoff whereby, as information complexity increases in the outer layer, revenue decreases and consumer welfare *reaches the maximum at an intermediate level.*

Finally, the current information layout of the search platform under study is marked by the circle in the figure, which shows room for improvement in both revenue and consumer welfare. This room for improvement motivates us to characterize the win-win situations that benefit both the consumers and the search platform in Section F.2 of the online appendix.

11 Conclusion

In this paper, we study the research question of how consumers acquire information at the expense of time and cognitive costs, in addition to the search costs we already understand from the search literature. Answering this question also helps us address new managerial questions of the optimal amount of information to reveal in the search-result pages of search platforms and Pareto improving local sorting. To answer these questions, we collect a rich data set from a large OTA search platform. To quantify cognitive cost, we innovate an information complexity measure called orderedness entropy. We find through model-free analysis and online random experiments that (1) consumers

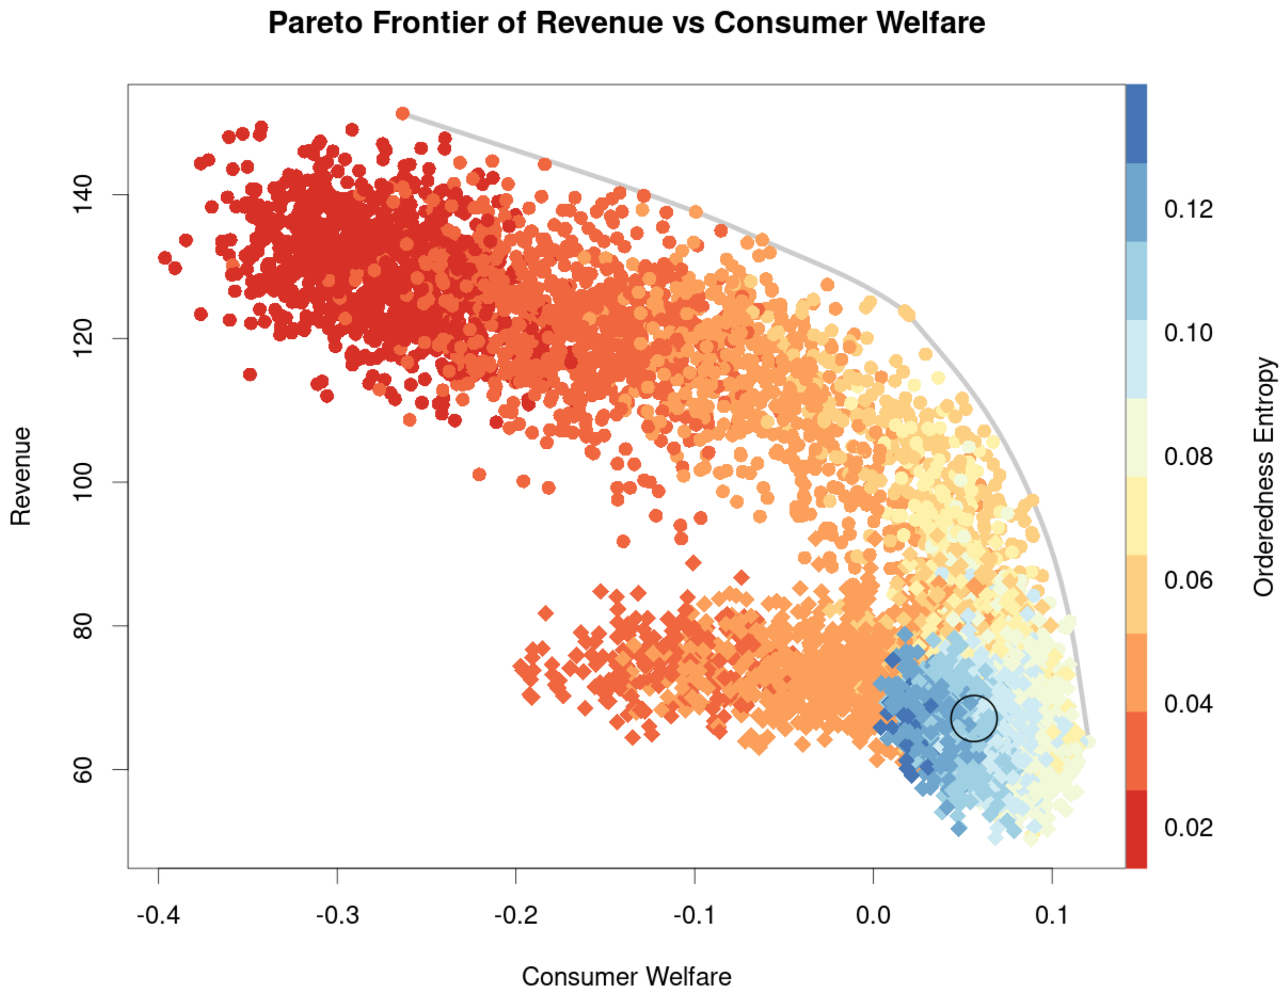

Figure 4: This figure shows the relationship between revenue and consumer welfare as information complexity varies. Each point corresponds to one information layout, and the color of the point corresponds to the average information complexity of the layout. The round-shaped points correspond to the layouts when product price is not revealed on the outer layer, and the diamond-shaped points correspond to the layouts when price is revealed. The grey line traces out the Pareto frontier of revenue versus consumer welfare. The current information layout of the search platform under study is marked by the circle in the figure.

actively use search-result pages as an important source of information; (2) consumers are responsive to refinement and click loading time; and (3) orderedness entropy is a good information complexity measure for cognitive cost, and consumers actively apply sorting variables in order to reduce orderedness entropy in their search process.

We then conceptualize the consumer search process by proposing a new structural model to capture both the refinement visiting and product clicking decisions, to distinguish between what consumers know before and after visiting refinements and clicking products, and to decompose search costs further into time and cognitive components. We find through model estimation that cognitive cost is a major component of the total search cost, while loading time cost has a much smaller share.

Finally, we vary the information revealed in the outer layer to characterize the various changes in the consumer search process, and study the relationship between consumer welfare and platform revenue. We find price is an important product attribute whose presence in the outer layer will dramatically change consumer search behavior. On the Pareto frontier of this relationship, we also find a general tradeoff whereby consumer welfare increases and platform revenue decreases as more information is revealed. For our OTA under study, we find room for Pareto improvement in both revenue and consumer welfare by changing which product attributes are revealed. We also find that local sorting of the default-ranking refinements is a Pareto improving strategy that platforms may consider adopting to increase consumer welfare without losing on platform revenue.

Our research is a first attempt in understanding and modeling cognitive cost in consumer search process. The information complexity measure we create has the potential to be adapted in other information-rich decision contexts. We hope our paper will inspire future research along this direction.

References

Brockett, P. L., Charnes, A., Cooper, W. W., Learner, D. and Phillips, F. Y.: 1995, Information theory as a unifying statistical approach for use in marketing research, *European Journal of Operational Research* **84**(2), 310–329.

Caplin, A.: 2016, Measuring and modeling attention, *Annual Review of Economics* **8**(1), 379–403.

Chen, Y. and Yao, S.: 2016, Sequential search with refinement: Model and application with click-stream data, *Management Science* .

Ching, A. T., Erdem, T. and Keane, M. P.: 2014, A simple method to estimate the roles of learning, inventories and category consideration in consumer choice, *Journal of choice modelling* **13**, 60–72.

De los Santos, B. and Koulayev, S.: 2017, Optimizing click-through in online rankings with endogenous search refinement, *Marketing Science* **36**(4), 542–564.

Ghose, A., Ipeirotis, P. G. and Li, B.: 2012, Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content, *Marketing Science* **31**(3), 493–520.

Ghose, A., Ipeirotis, P. G. and Li, B.: 2014, Examining the impact of ranking on consumer behavior and search engine revenue, *Management Science* **60**(7), 1632–1654.

Hajivassiliou, V. A. and Ruud, P. A.: 1994, Chapter 40 classical estimation methods for LDV models using simulation, *Handbook of Econometrics*, Elsevier BV, pp. 2383–2441.

Hauser, J. R. and Wernerfelt, B.: 1990, An evaluation cost model of consideration sets, *Journal of Consumer Research* **16**(4), 393–408.

Herniter, J. D.: 1976, An entropy model of brand purchase behavior, *Lecture Notes in Economics and Mathematical Systems*, Springer Nature, pp. 76–81.

Honka, E.: 2014, Quantifying search and switching costs in the US auto insurance industry, *The RAND Journal of Economics* **45**(4), 847–884.

Kim, J. B., Albuquerque, P. and Bronnenberg, B. J.: 2010, Online demand under limited consumer search, *Marketing Science* **29**(6), 1001–1023.

Kim, J. B., Albuquerque, P. and Bronnenberg, B. J.: 2017, The probit choice model under sequential search with an application to online retailing, *Management Science* **63**(11), 3911–3929.

Matějka, F. and McKay, A.: 2015, Rational inattention to discrete choices: A new foundation for the multinomial logit model, *American Economic Review* **105**(1), 272–98.

Roberts, J. H. and Lattin, J. M.: 1991, Development and testing of a model of consideration set composition, *Journal of Marketing Research* **28**(4), 429–440.

Shannon, C. E.: 1948, A mathematical theory of communication, *Bell System Technical Journal* **27**(3), 379–423.

Stigler, G. J.: 1961, The economics of information, *Journal of Political Economy* **69**(3), 213–225.

Tehrani, S. S. and Ching, A. T.: 2019, A heuristic approach to explore: The value of perfect information, *Johns Hopkins Carey Business School Research Paper* (19-05).

Ursu, R. M.: 2018, The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions, *Marketing Science* **37**(4), 530–552.

Weitzman, M. L.: 1979, Optimal search for the best alternative, *Econometrica* **47**(3), 641–654.

Wong, E.: 2017, Simplicity in design: 4 ways to achieve simplicity in your designs.
URL: *https://www.interaction-design.org/literature/article/simplicity-in-design-4-ways-to-achieve-simplicity-in-your-designs*

Online Appendix of "Consumer Online Search with Partially Revealed Information"

A Alternative Hypothetical Search Scenario

In addition to the conceptualization of the consumer search process described in Section 3.1 of the paper, we pay further attention to some details in our empirical work. The hotels displayed under different refinements are generally different, but some hotels could appear in multiple refinements if they satisfy the search query. For example, suppose hotel 1 appears in both Ref1 and Ref3 because it satisfies the criteria imposed by both search queries. Figure A.1 shows such a scenario. Before visiting Ref3, the consumer does not know the product availability. In step (5), she visits Ref3 and observes its product availability. Then the product information for hotel 1 is carried over from Ref1, and she does not need to sample hotel 1 again.

B Limitations of Current Information Complexity Measure

Although the orderedness entropy measure is desirable in many perspectives, there are certain scenarios where this measure is not ideal and new measures for cognitive costs are required. First, some product attributes might be categorical, such as the color of products, which cannot be transformed naturally into numeric values. For example, consider an attribute sequence of colors: *Red, Blue, Yellow, Green, Red, Blue, Yellow, Green*. It is challenging to quantify the differences between two consecutive colors. One possibility is to assign numeric values arbitrarily, say Red to 1, Blue to 2, Yellow to 3, and Green to 4. Then, this sequence becomes 1,2,3,4,1,2,3,4 and its orderedness sequence becomes 1,1,1,-3,1,1,1, which has relatively low orderedness entropy (0.59) since the orderedness sequence is quite concentrated at 1. Another similarly complicated value sequence can be obtained by cycling the sub-sequence *Blue, Yellow, Green* to obtain *Red, Yellow, Green, Blue, Red, Yellow, Green, Blue*. Using the same numeric coding, we obtain 1,3,4,2,1,3,4,2. Its orderedness sequence is 2,1,-2,-1,2,1,-2, which has much higher orderedness entropy (1.95), since the orderedness sequence is more uniformly distributed. Therefore, it is not valid to directly apply the orderedness entropy measure of the current form to categorical variables through arbitrary numeric coding.

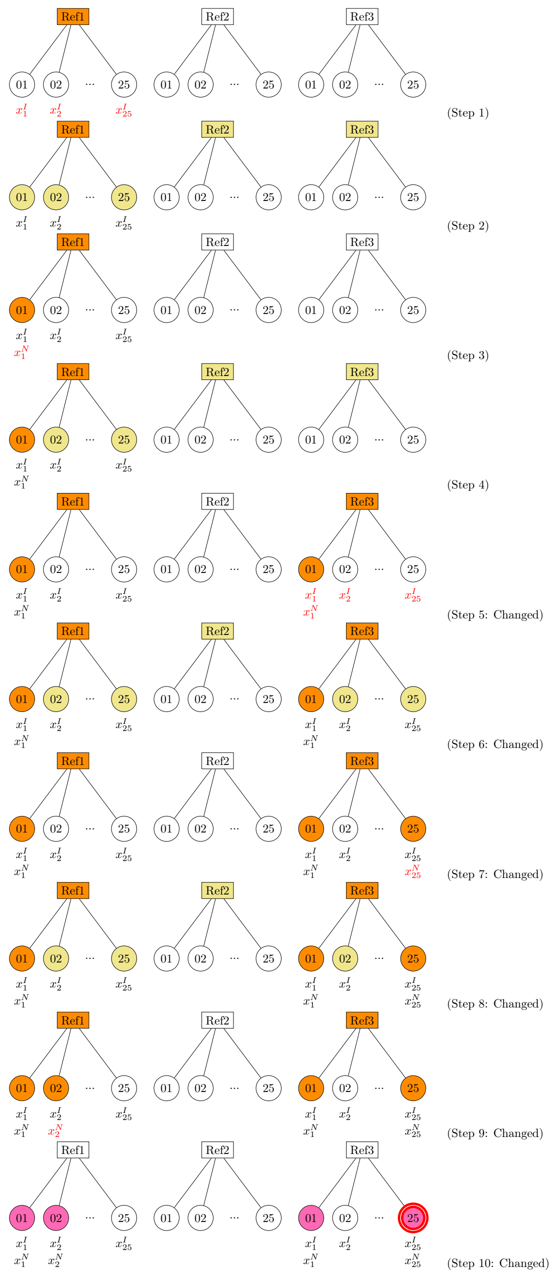

Figure A.1: Alternative scenario in which hotel 1 of Ref1 also appears as hotel 1 in Ref3. It could also appear as a hotel at some other position in Ref3 depending on the ranking. All the steps are the same as in Figure 3 in Section 3.1 of the paper, except after step (5). Before visiting Ref3, the consumer does not know the product availability of Ref3. In step (5), the consumer visits Ref3 and observes its product availability. The product information for hotel 1 is carried over from Ref1, and hotel 1 in Ref3 is not a search option for the consumer for the rest of the search process, which differs from the previous scenario.

Another possibility might be to transform the categorical sequence into one that indicates whether two consecutive values are different. After such transformation, both sequences of colors become 1, 1, 1, 1, 1, 1, 1 and thus have zero orderedness entropy. However, both sequences are harder to comprehend than, for instance, *Red, Blue, Red, Blue, Red, Blue, Red, Blue,* which also has zero orderedness entropy. As a result, a new metric is desired to capture the cognitive costs of understanding categorical variables.

Second, the orderedness entropy measure captures the well-orderedness of a sequence. However, sometimes an attribute sequence is hard to understand, not only because of its disorderedness, but also because of the difficulty to grasp the *intrinsic* information in the attribute values. For instance, consider the following two value sequences: 1038660.24, 1038910.27, 1039160.30, 1039410.33, 1039660.36, 1039910.39, 1040160.42 versus 1, 3, 5, 7, 9, 11, 13. Clearly, the first sequence is harder to understand than the second, whereas from the perspective of orderedness entropy, both sequences are equally complex since they have constant incremental values between consecutive numbers. Thus, orderedness entropy by itself does not capture the levels of cognitive efforts necessary to understand the *intrinsic* values of different product attributes.

Third, orderedness entropy as a measure for quantifying cognitive costs does not take into account the possibility that consumers may not comprehend and compare all the products on refinements. In practice, the consumers may endogenously form consideration sets as a reaction to information complexity. When a refinement page has high information complexity, the consumers may only consider a subset of product information on it. In order to seriously study the formation of consideration sets, ideally we would like to have consumers' eye-tracking data to accurately know which products they have considered on each refinement page. Unfortunately, this detailed information is not available in our data. Thus, if we extend the current framework to explicitly model how consumers form their consideration sets, we need to integrate out all the possible consideration sets in our likelihood computation.

There are different ways to form possible consideration sets on one refinement. The most comprehensive way is to consider the power set of all products on the refinement, whose size is 2^{25} in our context with 25 products. This renders the estimation infeasible. The most simplistic way is to assume that the consumers only consider the top-k products listed on each refinement page, and thus the number of possible consideration sets is 25 for each refinement. However, we need to account for all combinations of possible consideration sets over the visited refinements by the consumer, which is still computationally intractable. Nevertheless, adjusting the cognitive cost measure to incorporate the possible formation of consideration sets together with collecting detailed eye-tracking data from consumers is an important research topic, which we leave for future research.

C Evidence From Online Random Experiments

In this section, we provide findings from online random experiments, which complement our reduced-form evidence. In particular, in the reduced-form analysis, we effectively rule out the possibility that consumers apply sorting to change the content of the search results, since we focus on those refinements with small-result sets only. However, the reduced-form setting cannot guarantee that the only purpose for consumers to apply sorting is to enhance well-orderedness. Instead, consumers may also sort using a highly relevant product attribute that is very *predictive* for the overall product utilities, and thus they could easily anchor the highest utility product after sorting on this product attribute. However, without knowing consumers' utility coefficients, we cannot examine or control for this channel using the observational data and reduced-form analysis. Hence, we introduce an experiment to strengthen the evidence.

C.1 Experimental Design

Experiment 1 on Response Time: To enhance our understanding on whether consumers respond to the orderedness of product attributes, we investigate the simplest scenario where the products have only one attribute. In particular, we first generate a list of 20 products, each with its only attribute and hence its utility being an even number from 0 to 20, where 0 and 2 appear once and the rest integers appear twice in the list. Therefore, the fully-sorted list of the product utility is $v = [0, 2, 4, 4, 6, 6, 8, 8, ..., 20, 20]$. Then, we introduce two more lists. One is a 3-block-sorted list, which is generated by cutting the fully-sorted list v into 3 blocks and randomly shuffling the three blocks. Another one is a fully-random list, which is generated by randomly sampling the elements in v without replacement. Finally, for each generated list, we draw a random integer from 1 to 20 and add it to all the elements in the list.

Among the fully-sorted, 3-block-sorted, and fully-random lists, we randomly assign two of the lists to each experiment subject. We have recruited 87 subjects on Amazon Mechanical Turk (MTurk) for this experiment. Each subject earns a bonus when she selects the highest-utility product in a list.

Clearly, the level of orderedness as measured by orderedness entropy varies significantly across lists. In particular, for the fully-sorted, 3-block-sorted, and fully-random product lists, the orderedness entropy on average is 0.69, 0.99 and 2.25, respectively. Then, we examine if the experiment subjects need to spend more time to make a choice when they encounter higher orderedness entropy. A screenshot of one example of this experiment can be found in Figure C.1.

Experiment 2 on Sorting Behaviors: We introduce an experiment similar to the reduced-form setting, in which a product has multiple attributes and the consumers decide whether to sort the list of options and which attribute to sort on. The difference between the experiment and the reduced-form setting is that we observe consumers' utilities in the experiment. As indicated in the reduced-form analysis, consumers tend to sort on the product attribute which induces a larger amount of reduction in orderedness entropy. However, one potential concern with this interpretation is that a product attribute corresponding to more orderedness entropy reduction could also be more relevant for consumers; as a result, consumers sort on that attribute not only to reduce orderedness entropy, but also to better predict the ordering of the overall product utilities. Without knowing consumers' utilities, we cannot disentangle these two hypotheses. With the help of experimental design, we not only know consumers' utilities, but we can also make sure that all the product attributes are equally predictive of the overall product utilities. Based on this foundation, we can test if consumers apply sorting to reduce orderedness entropy.

In this experiment, the total utility of a product is set to be the sum of its three attribute values. Different product attributes may have different numbers of support points. More specifically, each decision context includes 20 products, and each of the three product attributes is a list of 20 values that could have 2, 3, 5 or 11 support points with values evenly spread out between 0 and 10. In particular, a fully-sorted list of a 2-support-points product attribute is set to be $v_2 = [0, 0, ..., 0, 10, 10, ..., 10]$, a 3-support-points one is $v_3 = [0, ..., 0, 5, ..., 5, 10, ..., 10]$, a 5-support-points one is $v_5 = [2, ..., 2, 4, ..., 4, 6, ..., 6, 8, ..., 8, 10, ..., 10]$, and finally an 11-support-points one is

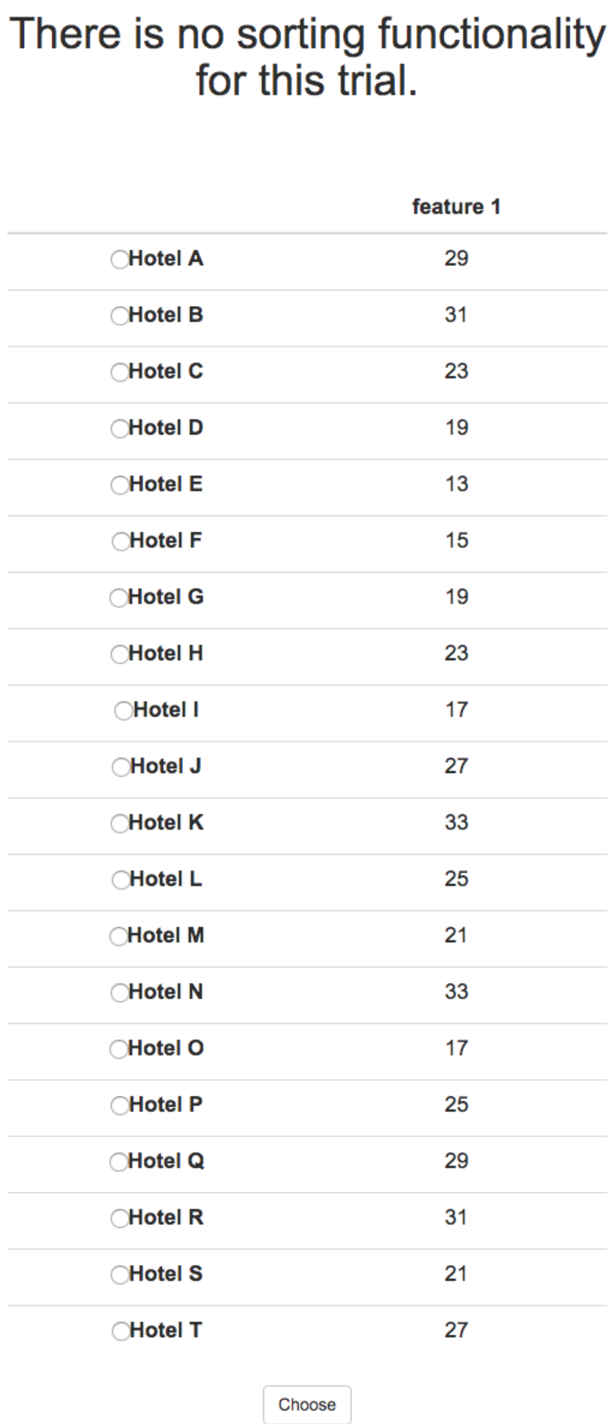
There is no sorting functionality for this trial.

	feature 1
Hotel A	29
Hotel B	31
Hotel C	23
Hotel D	19
Hotel E	13
Hotel F	15
Hotel G	19
Hotel H	23
Hotel I	17
Hotel J	27
Hotel K	33
Hotel L	25
Hotel M	21
Hotel N	33
Hotel O	17
Hotel P	25
Hotel Q	29
Hotel R	31
Hotel S	21
Hotel T	27

Choose

Figure C.1: Screenshot of Experiment 1: we examine the subjects' response time for understanding a list of one-attribute products.

$v_{11} = [0, 1, 2, 2, ..., 9, 9, 10, 10]$.

Furthermore, a decision context could be either a combination of the (2,3,11)-support-points product attributes, or the (2,5,11)-support-points product attributes. Also, to allow different ranges of orderedness entropy reduction, each product attribute could be fully-sorted, 3-block-sorted or fully-random, as described in Experiment 1. Thus, for the decision context with either the (2,3,11)-support-points or the (2,5,11)-support-points product attributes, there are seven designs: (1) one design where all the attributes are fully-random, (2) two designs where the first attribute is either 3-block sorted or fully-sorted, (3) two designs where the second attribute is either 3-block sorted or fully-sorted, and (4) two designs where the third attribute is either 3-block sorted or fully-sorted. Overall, we have 14 designs, and the orderedness entropy varies significantly across designs. We also randomly permute the columns of the product attributes.

Last but not least, we define the predictability of a product attribute for the overall product utilities as the Spearman's Rank-Order Correlation between the ranking of a given product attribute value and that of the overall product utilities. Then, we adjust the range of each product attribute so that different product attributes have similar amount of predictability for the overall product utilities.

We have 85 subjects for the second experiment. Each subject is assigned four out of the seven designs described above for the (2,3,11)-support-points product attributes as follows: design (1) and one random design from (2), (3), and (4). Furthermore, each subject is assigned another four designs out of the seven for the (2,5,11)-support-points product attributes. The screenshot of one example of the second experiment can be found in Figure C.2, which presents a (2,5,11)-design with fully random attributes.

C.2 Experimental Findings

Experiment 1: The findings of the first experiment indicate that the more sorted the product utilities are, the less time spent on making a choice decision. As is reported in Table C.1, the experimental results show that when the orderedness entropy of the product attribute list increases by one unit, the subjects on average spend 2 more seconds on comprehending the list. Hence, the

Sort By: feature 1 feature 2 feature 3 Default

	feature 1	feature 2	feature 3
Hotel A	8	7	17
Hotel B	2	9	17
Hotel C	3	12	7
Hotel D	10	12	7
Hotel E	8	17	7
Hotel F	6	12	7
Hotel G	9	12	17
Hotel H	4	9	17
Hotel I	7	7	17
Hotel J	5	15	7
Hotel K	12	9	7
Hotel L	11	15	7
Hotel M	5	15	7
Hotel N	6	17	17
Hotel O	10	15	7
Hotel P	12	9	17
Hotel Q	4	17	17
Hotel R	7	7	17
Hotel S	9	7	17
Hotel T	11	17	7

Choose

Figure C.2: Screenshot of Experiment 2: we examine the subjects' sorting behaviors for understanding a list of three-attributes products.

	Response Time (std. err.)
Orderedness Entropy	**2.01 (0.59)**
Intercept	**8.55 (0.91)**
Number of Observations	174

Table C.1: Experimental evidence: consumers' response time as a function of orderedness entropy. This experiment contains a list of products with one product attribute. The positive coefficient of orderedness entropy indicates that, on average, consumers spend more time to understand, compare, and choose from a product list when the list is more disordered. Bold fonts indicate estimates significant at the 95% level.

subjects on average spend 3.13 additional seconds to understand a fully-random attribute list of 20 products than to understand a fully-sorted one. Therefore, the experimental findings support the idea that the orderedness entropy measure reasonably proxies the amount of cognitive efforts required for understanding a list of product attribute.

Experiment 2: The outcomes of the second experiment indicate that the subjects tend to sort on a product attribute which generates a page with lower orderedness entropy. Like the reduced-form analysis, we use a multinomial logit model to study the choice probability of sorting on $j \in SV = \{\text{No Sorting, Feature 1, Feature 2, Feature 3}\}$ as follows:

$$P(\text{sort on } j) = \frac{\exp\left(\text{Intercept}_j + \eta\text{Entropy}_j + \xi\text{RankCorr}_j\right)}{\sum_{k \in SV} \exp\left(\text{Intercept}_k + \eta\text{Entropy}_k + \xi\text{RankCorr}_k\right)}.$$

As can be seen from Table C.2, the estimated coefficient of orderedness entropy is negative and significant at the 95% level. Once again, this finding suggests that orderedness entropy appropriately proxies the amount of cognitive cost required for comprehending and comparing a list of products, and consumers actively apply sorting to reduce that cost. Moreover, the insignificance of the estimated rank correlation coefficient in the second specification confirms that we have successfully matched the degree of predictability across product attributes in the experimental design. The estimate of the orderedness entropy coefficient is robust to the introduction of the rank correlation variable, which supports our conclusion that the reduction of orderedness entropy is indeed one purpose (which may not be the only purpose) for the consumers to apply sorting to the search-result pages.

	Specification 1	Specification 2
Orderedness Entropy Coefficient η	**-0.36 (0.11)**	**-0.35 (0.11)**
Rank Correlation Coefficient ξ		-0.14 (0.21)
Intercept Sorting Attribute 1	**0.39 (0.11)**	**0.36 (0.12)**
Intercept Sorting Attribute 2	-0.06 (0.12)	-0.10 (0.13)
Intercept Sorting Attribute 3	**-0.36 (0.13)**	**-0.40 (0.14)**
Number of Observations	680	680

Table C.2: Experimental evidence: consumers' sorting behaviors as a reaction to orderedness entropy. This experiment contains a list of products with three product attributes. The negative coefficient of orderedness entropy indicates that consumers tend to sort on a product attribute which generates a page with lower orderedness entropy. The insignificance of the rank correlation coefficient in the second specification confirms that consumers indeed sort a list of options for the purpose of reducing orderedness entropy, rather than for the purpose of locating the highest utility product. Bold fonts indicate estimates significant at the 95% level.

D Unified and Multiple Sequential Search Models

D.1 Unified Sequential Search Model

The unified sequential search model has the following state variables: the set of unsampled products $\bar{O}$ (which is the union of the set of unsampled products on the current refinement $\bar{S}$ and the set of unsampled products on the previously visited refinements $\bar{T}$), the set of unvisited refinements $\bar{R}$, the current refinement rf, and the current best utility y. We use O, S, T, R for the respective sampled counterparts. We suppress the subscripts for consumer h and search session se in this section to avoid notational clutter, because the search model concerns the search process for a given consumer within a search session. When the context calls for the distinction of different consumers and search sessions, such as when we specify the likelihood across consumers and search sessions, we restore the subscripts.

Equations (D.1) - (D.4) specify the value functions of the fully forward looking search process.

$$\Phi\left(\bar{O}, \bar{R}, rf, y\right) = \max\{y, \tag{D.1}$$

$$\max_{j\in\bar{O}}\left\{-c_P(rf, j) + E_j\Phi\left(\bar{O} - \{j\}, \bar{R}, rf_j, \max\{y, u_j\}\right)\right\}, \tag{D.2}$$

$$\left.\max_{rf'\in\bar{R}}\left\{-c_R(rf') + E_{rf'}\Phi\left(\bar{O}\cup\bar{S}_{rf'}, \bar{R} - \{rf'\}, rf', y\right)\right\}\right\} \tag{D.3}$$

$$\Phi\left(\emptyset, \emptyset, rf, y\right) = y \tag{D.4}$$

Lines (D.1) - (D.3) specify the choices that a consumer has at any stage:

1. The consumer can choose to exit the search with the current best utility y. (Line (D.1))

2. The consumer can choose to sample a product j that generates the highest option value, computed as the tradeoff between the product clicking cost $c_P(rf, j)$ and the continuation value to sample product j. (Line (D.2)). Note that when the consumer is evaluating product j as a search option, the product could be on the current refinement or on a previously visited refinement. Recall that for a product on the current refinement, the consumer incurs clicking baseline cost, position cost, and the expected detail page loading time cost; for a product on a previously visited refinement, the consumer incurs these three costs as well as the expected refinement loading time cost. In order to differentiate the product clicking cost for either case, $c_P(rf, j)$ is a function of rf and j, where rf indicates the current refinement, and the argument j is an index for product j, which contains all the relevant information for product j including the refinement rf_j of product j. When the refinement of product j is different from the current refinement ($rf_j \neq rf$), the consumer pays the additional cost for switching refinement. The value function to sample a product will have the following as its updated state variables: a new set of unsampled products $\bar{O} - \{j\}$ with product j being no longer a search option, the same set of unvisited refinements $\bar{R}$ as before, the refinement rf_j that product j is on, and the updated current best utility $\max\{y, u_j\}$.

3. The consumer can choose to visit a refinement rf' that generates the highest option value, computed as the tradeoff between the refinement search cost $c_R(rf')$ and the continuation value to visit refinement rf' (Line (D.3)). The refinement search cost contains the expected refinement loading time cost, the expected cognitive cost, and the refinement baseline cost. The value function to visit a refinement will have the following as its updated state variables: a new set of unsampled products $\bar{O} \cup \bar{S}_{rf'}$ with new product availabilities $\bar{S}_{rf'}$ being discovered and added as search options, a new set of unvisited refinements $\bar{R} - \{rf'\}$ with refinement rf' being no longer a search option, the new current refinement rf', and the same current best utility y as before.

Equation (D.4) specifies the base case that the value function equals the current best utility when there are no more searches to be performed.

The new model tries to capture the important information and cost structures during the consumer search process on modern platforms, as described in the conceptualization. It also accommodates search behaviors that are commonplace on modern platforms. In particular,

- Consumers discover new product availabilities, $\bar{S}_{rf'}$, and expand their search options $\bar{O}$ after they visit a new refinement rf', rather than knowing them all before commencing any search.
- Product information is separated into the revealed portion that is discovered as the consumer visits a new refinement, and the hidden portion to be searched as the consumer visits a product detail page. The model assumes that the consumer accounts for such uncertain information through the following integrations:
    1. For the value function of visiting a new refinement, the integration ($E_{rf'}$) is taken over the joint distribution of the partially revealed product attribute values for an ordered list of products on refinement rf';
    2. For the value function of sampling a product, the integration (E_j) is taken over the conditional distribution of the hidden component of the product utility given the revealed product attribute values.
- Consumers may re-visit old refinements to sample products, because search options $\bar{O}$ include unsampled products on all the visited refinements.
- The search cost is flexible enough to distinguish sampling products on the current refinement versus on the previously visited refinements.

These model capabilities are important for researchers to study the preference and cost parameters of the consumers. The additional model flexibility also allows the researchers to answer new managerial questions, including the determination of the optimal amount of information to reveal in the outer layer.

However, the capabilities and flexibility of the new model come at a cost. Since the consumers could expand their search options when they visit new refinements, there is no simple optimal solution as specified in Weitzman (1979) for the dynamic programming problem. Rather, solving the fully forward looking problem poses the following enormous computation challenges:

1. The curse of dimensionality arises due to the exponential growth of decision paths of varying length along which the decisions to sample products and to visit refinements interleave;

2. The integration, $E_{rf'}$, for each refinement is computationally difficult, and even estimating the relevant distributions involved in the integration is a formidable task itself.

Therefore, regrettably, using the fully forward looking model in its original form is not viable for our empirical analysis.

D.2 Multiple Sequential Search Model

The multiple sequential search model simplifies the unified sequential search model by balancing model flexibility and computation tractability. In the following, we first present our multiple sequential search model, and then we compare it with the unified model.

In the multiple sequential search model, each search problem is anchored to the current refinement the consumer is visiting. When the consumer decides to visit a different refinement from the one she is currently at, we assume the consumer starts a new search problem. Suppose a consumer in a search session has B search problems. Within the bth search problem ($b \in \{1, ..., B\}$), the consumer is assumed to engage in the traditional sequential search where both the *unvisited refinements* and *unsampled products on the visited refinements* are search options. Equations (D.5) - (D.8) specify the value function of the bth problem in the multiple sequential search model. The notations are the same as the ones used in the unified model, except that we use subscript b to indicate that the quantities are for the bth search problem.

$$\Phi_b\left(\bar{O}_b, \bar{R}_b, y_b\right) = \max\{y_b, \tag{D.5}$$

$$\max_{j\in\bar{O}_b}\left\{-c_P(rf_b, j) + E_j\Phi_b\left(\bar{O}_b - \{j\}, \bar{R}_b, \max\{y_b, u_j\}\right)\right\}, \tag{D.6}$$

$$\left.\max_{rf'\in\bar{R}_b}\left\{-c_R(rf') + E_{rf'}\Phi_b\left(\bar{O}_b, \bar{R}_b - \{rf'\}, \max\{y_b, U_{rf'}\}\right)\right\}\right\} \tag{D.7}$$

$$\Phi_b(\emptyset, \emptyset, y_b) = y_b \tag{D.8}$$

The multiple sequential search model has a very similar structure with the unified model. In particular, the option values for stopping (Line (D.1) vs (D.5)) and for sampling a product on the visited refinements (Line (D.2) vs (D.6)) are computed in the same way. The base case (Equation (D.4) vs (D.8)) is also the same. The most prominent deviation from the unified model is the computation of the continuation value of searching a refinement (Line (D.3) vs (D.7)). Below we discuss in details the differences between the multiple and unified sequential search models.

First, the multiple sequential search model simplifies the computation of the exact continuation value of searching a refinement. Since the exact continuation value takes into account the benefit of expanding the search options, no simple optimal solution as specified in Weitzman (1979) is available for the dynamic programming problem. Also, from the consumer's perspective, computing the exact continuation value needs to mentally hypothesize what happens beyond reaching a new refinement, which depends on too much information that is unobserved, including the would-be partially revealed product attribute values of an ordered list of products on the refinement. Provided that accounting for such uncertainty is quite difficult to perform, both mentally and computationally, we assume that the consumer uses a heuristic yet intuitive utility distribution to characterize the benefit of searching a refinement. Following the same one-step-ahead approximation adopted in De los Santos and Koulayev (2017), we assume that the consumer believes that the benefit of searching a refinement is captured by the distribution of the highest product utility on the refinement. In other words, as De los Santos and Koulayev (2017) put it, the consumer mentally hypothesizes that she will select the highest utility option on the refinement as is in a standard multinomial discrete choice setting.

Accordingly, instead of integrating over $\bar{S}_{rf'}$ as in Line (D.3), where the uncertainty comes from the joint distribution of the partially revealed product attribute values of an ordered list of products on refinement rf', we integrate over the highest product utility on the refinement, $U_{rf'}$, which is much simpler as the uncertainty is based on a scalar.[6] With this simplification, the curse of dimensionality of the dynamic programming problem is addressed. The bth search problem has the same structure as the problem in Weitzman (1979), and therefore Equations (9) - (14) in the paper characterize the optimal solutions of the dynamic programming problem, which utilize the reservation utilities of searching a refinement and sampling a product as defined in Equations (7) and (8) of the paper.

Second, when the consumer decides to visit a new refinement, the consumer updates the state variables in the same way as she does in the unified model, whereas the deviation from the unified model is that the consumer now enters a new search problem. In particular, when the consumer visits a new refinement rf', the consumer starts the $(b+1)$th search problem and rf_b is updated

[6]For our empirical analysis, we assume that the distribution of U_{rf}, namely $F_{h,rf}$, is a normal distribution for which its first and second moments are unbiased estimates of the mean and variance of the maximum product utility on refinement rf, as discussed in details in Section 7.2 of the paper.

to $rf_{b+1} = rf'$. Meanwhile, the consumer updates her state variables as follows: $\bar{O}_b$ becomes $\bar{O}_{b+1} = \bar{O}_b \cup \bar{S}_{rf'}$, where the available products $\bar{S}_{rf'}$ on refinement rf', are added as the search options; $\bar{R}_b$ is updated to $\bar{R}_{b+1} = \bar{R}_b - \{rf'\}$ with refinement rf' removed from the search options; and the current best utility, y_b, is updated to be the best product utility discovered so far. With the updated state variables $\bar{O}_{b+1}$, $\bar{R}_{b+1}$ and y_{b+1}, the consumer starts the $(b+1)$th search problem as characterized in Lines (D.5) - (D.8).

Note that, due to the one-step-ahead approximation, the consumer *necessarily* starts a new search problem when she switches to a new refinement, because the product availabilities on the new refinement are *approximately* considered before she switches the refinement, and the product availabilities (i.e. the partially revealed product attributes of an ordered list of products on the new refinement) are fully updated after she switches the refinement. Overall, the multiple sequential search model contains a sequence of the Weitzman's problems, which are separated by the different amount of information the consumer has acquired before and after switching a refinement.

Third, we note that rf is a state variable in the unified sequential search model, but it is not a state variable in the multiple sequential search model. This is because in the multiple sequential search model, each search problem is uniquely indexed by subscript b; since the subscript b corresponds one-to-one to the current refinement rf_b, we no longer need rf_b as a state variable provided that we index each search problem using b. In the unified sequential search model, we need rf as a state variable to keep track of the refinement the consumer is looking at.

Finally, we compare our multiple sequential search model with the model studied in De los Santos and Koulayev (2017). We adopt the same one-step-ahead approach by using an intuitive utility distribution to characterize the benefit of searching a refinement. We differ, however, in how the current best offer is updated and in the range of search options being considered during the search process. In terms of the update of the current best offer, their model assumes that, after switching refinements, the consumer decides whether to click on at most one product on the refinement, captured in a reduced-form fashion by a discrete choice model. As a result, all product attributes and random utility shocks are assumed known to consumers after they visit a refinement. In our model, we assume the consumer still engages in sequential search, because she only observes the partially revealed product attributes and does not know the hidden attributes and utility shocks. The update of the current best utility is realized through selective sampling of products, which informs the researchers about the product clicking costs, separately from the refinement search costs. In terms of the refinement choice, both models assume the consumer computes and compares the reservation utilities of the unvisited refinements for search decision. The difference is that, in our model, the consumer additionally considers the reservation utilities of the unsampled products, and the set of the unsampled products may increase as the consumer uncovers new partially revealed products as she visits new refinements.

In summary, we only simplify the component of the unified sequential search model (i.e. the computation of the refinement option value) for which the consumer does not have much information about before they proceed further in the search process, while still maintaining the overall sequential search framework and model capabilities. We hope this simplification is consistent with our central modeling principle to capture the consumer information and cost structure more realistically, while accounting for model tractability. Our model extends that of De los Santos and Koulayev (2017).[7] It grants us the flexibility to have distinct search costs for refinement visiting and product clicking, and allows the change in the platform design to affect refinement visiting and product sampling through changes in cognitive costs and partially revealed information. This flexibility is crucial to answering our managerial question.

[7] Our paper also relates to recent papers in the context of dynamic consumer learning through modeling consumers' bounded cognitive ability via heuristic approaches. For example, Ching et al. (2014) points out that it is cognitively taxing for consumers to solve a learning problem in each purchase occasion. Given a product category, they propose to have a heuristic reduced-form first stage of decision to consider. The consumers only solve the learning problem on the purchase occasions when they decide to consider. As another example, Tehrani and Ching (2019) starts from the concern of cognitive burden placed on consumers for solving dynamic problems. They introduce a heuristic valuation of exploration, which additively augments the myopic valuation of a choice. This improved model, myopic-VPI, does not solve the dynamic learning problem, but, due to its computational tractability, more plausibly approximates the consumer learning process under bounded cognition. In our paper, we too believe bounded cognitive ability is an important component in the consumer search process, as manifested by platforms providing sorting and filtering facilities and consumers actively using such facilities to reduce cognitive cost. This cost also has important managerial implications. As a result, we come up with a heuristic measure of cognitive cost to enrich the consumer search model.

E Identification, Monte Carlo Simulations, and Model Fit

E.1 Identification

The foundation of our empirical study is the identification of the parameters of consumer preference and search costs. We adopt the same identification strategy as the one in Chen and Yao (2016), while our study emphasizes the distinctions between the refinement visiting costs and product clicking costs, and our identification relies on different exclusion restriction variables. We also use Monte Carlo simulations to demonstrate the identification and estimation performance of our empirical framework.

Preference Parameters: Various sources of the observational data variations contribute to the identification of the mean preference parameters. First, the consumers' decisions on which refinements to visit reveal their preferences as follows. At each decision stage, the consumer evaluates the expected returns of visiting different refinements to determine whether to search a refinement and which refinement to search. The expected return depends on the distribution of the best product utility on each refinement, which in turn depends on the joint distribution of the product attributes on the refinement and consumer's preferences for the product attributes. The sequence of consumer refinement visiting reveals the ordering of the expected returns of the refinements, and hence the consumer preferences. For instance, suppose a consumer sorts on "prices from low to high" and visits the first few pages of the returned results. Given that the same sorting variable is applied, intuitively the first few pages and the remaining pages have similar levels of orderedness entropies (and thus similar amount of refinement costs), while consumers' decisions on visiting the first few pages rather than the later pages reveal the consumers' preferences for lower prices. Therefore, controlling for the variations of search costs across refinements, the average orderings of consumers' refinement visiting reveal the mean preference parameters.

Second, in a similar manner as the identification based on refinement visiting, the aggregate orderings of the consumers' product clicking also reveal their mean preferences. In particular, the utility of clicking a product depends on the revealed product attributes and the conditional distribution of the hidden product attributes, as well as consumers' mean preferences for these product attributes. Consumers' decision on clicking a product suggests that the product has the highest expected returns among the unclicked products, which in turn is informative about consumers' preferences (after controlling for the different clicking costs across products).

Finally, provided that the attribute values vary across products, the purchase frequencies over products among the set of sampled products provide identification for the mean preference coefficients. This source of identification is similar as in the classical multinomial discrete choice setting, where we need to normalize the outside option utility to be $u_{h0,se} = \epsilon_{h0,se}$, and $\epsilon_{hj,se}$ follows the standard normal distribution.

The heterogeneity of the preference coefficients across consumers can be recovered since we observe a sequence of decisions for each individual consumer, including refinement visiting, product clicking and purchasing. The repeated observations for each consumer are informative for consumer-specific preference coefficients, and the variations of consumer-specific preference coefficients provide identification for the heterogeneity of the preference coefficients across consumers.

Search Cost Parameters: The consumers encounter both the costs of visiting refinements and the costs of clicking products. More specifically, the refinement visiting cost parameters include the consumers' responsiveness to the expected information complexity and to the expected loading time, and the refinement baseline cost. Also, the product clicking cost parameters include the consumers' reactions to the product positions and to the expected loading time, and the clicking baseline cost. For these parameters, we first discuss the identification of the mean parameter values and then the heterogeneity across consumers.

The information complexity level as measured by orderedness entropy varies significantly across refinement options. Controlling for the effects of applying various sorting and filtering variables on the product utility distributions of refinements, the correlation between the frequency with which refinements are visited and the corresponding orderedness entropies informs about consumers' mean responsiveness to orderedness entropy.

Similarly, products have different positions on refinements. Controlling for the differences in the conditional utility distributions across products given the partially revealed product attributes, the correlation between the frequency with which products are clicked and the specific product positions provides identification for consumers' mean reactions to positions.

Moreover, consumers consider the unvisited refinements and unclicked products on the visited refinements together. The expected loading time is different for visiting refinements and clicking products, and is different across consumers. Thus, the effects of the variations in loading time on consumers' search decisions reveal consumers' reactions to loading time.

Consumers make purchase decisions based on the product utilities as functions of product attributes. The exclusion restrictions are defined as the variables that only affect consumers' search costs but not the product utilities, and thus they do not influence purchase decisions directly. Following the literature of selection models and the identification strategy in Chen and Yao (2016), the availability of the exclusion restriction variables enables the separate identification of the baseline utility parameter and the baseline search cost coefficients. In our empirical setting, we have different exclusion restriction variables for the refinement visiting costs and product clicking costs. In particular, the expected orderedness entropies and refinement loading time only affect the refinement visiting costs, and the positions and expected product-detail page loading time only determine the product clicking costs, while all of them do not appear in the consumer utility functions. Therefore, the baseline costs for both the refinement visiting and product clicking, as well as the baseline utility coefficient are separately identifiable.

Finally, we have repeated observations for consumers' decisions on searching and clicking, which are informative for consumer-specific search cost coefficients, and thus enable the identification of the search cost heterogeneities.

E.2 Monte Carlo Simulations

We use Monte Carlo simulations to demonstrate the identifiability and estimation performance of the model.

We first simulate the platform data, and then we use the search model described in Section 7 of the paper to simulate consumer behaviors. For the platform data, we simulate attribute values for 28 products, each having two revealed attributes, i.e., price and hotel star, and one hidden attribute, hotel rating. Prices and hotel stars are distributed independently from the continuous uniform distribution between 100 and 300 and from the discrete uniform distribution between 1 and 5, respectively. Hotel rating is related to price and hotel star linearly as

$$\begin{aligned}\text{Rating} &= 1 + 2.0 \times \log(\text{Price}) + 0.4 \times \text{Star} + \epsilon \\ \epsilon &\sim N(0, 0.25).\end{aligned}$$

Each refinement can display up to 7 products. For the default-ranking refinement, the 28 products are randomly distributed on 4 pages. In terms of the filtering functionality, for simplicity, we allow consumers to filter price to the range of [100, 200] for lower priced options, which reduces the search-result size to 14 products. In terms of the sorting functionality, consumers may sort on price or hotel star. Consumers may also simultaneously apply filtering and sorting to the products. In total, there are 18 refinements for consumers to consider (i.e., the unfiltered search results span 4 pages for using default ranking, sorting on price, and sorting on hotel star; and the search results filtered for lower priced products span 2 pages for using default ranking, sorting on price, and sorting on hotel star). The refinement loading time is an i.i.d random draw for each consumer per search session from the continuous uniform distribution between 2 and 5, and similarly the product-detail page loading time is an i.i.d random draw uniformly distributed between 1 and 3. We simulate 200 sets of platform data, based on which rational expectations of orderedness entropy and best utility distribution are derived for each refinement, and the parameters in the belief distribution for the unrevealed product utilities conditional on the revealed product attributes are estimated.

For the consumer search model, product utility depends on price, hotel star, hotel rating, and utility shock. Consumer refinement search cost depends on the refinement baseline cost, price entropy cost, hotel star entropy cost, and refinement loading time cost. Consumer clicking cost depends on the clicking baseline cost, position cost and product-detail page loading time cost. In each bootstrap sample, we simulate the behaviors for 100 heterogeneous consumers, each having 2 purchase sessions. We estimate the model parameters using the simulated consumer behavior. Due to the computation time constraint, we set the number of heterogeneity draws to be 50, and the number of random utility draws to be 50 for each consumer in the simulated maximum likelihood estimation. Finally, we repeat the process 30 times, and we report the bootstrap means and standard errors for the estimated preference and cost coefficients below in Table E.1.

Table E.1 shows the true model parameters and the estimated parameters. The estimates are overall quite close to the true values and are all significant. Therefore, the Monte Carlo simulations provide supportive evidence that our model parameters can be identified and well estimated.

			True Parameter	Est. Parameter (SE)
Preference Mean	Revealed	Price	-0.120	**-0.106 (0.014)**
		Star	1.730	**1.508 (0.228)**
	Unrevealed	Rating	1.920	**1.901 (0.175)**
Preference Heterogeneity	Revealed	Price SD	0.030	**0.031 (0.002)**
		Star SD	0.468	**0.468 (0.027)**
	Unrevealed	Rating SD	0.538	**0.539 (0.064)**
Search Cost Mean	Refinement	Base	0.309	**0.357 (0.056)**
		Price Entropy	0.106	**0.114 (0.010)**
		Star Entropy	0.145	**0.153 (0.012)**
	Time	Time (Sec)	0.008	**0.008 (0.001)**
	Click	Click Base	0.054	**0.047 (0.005)**
		Click Position	0.012	**0.011 (0.001)**
Search Cost Heterogeneity	Refinement	Base SD	0.070	**0.081 (0.013)**
		Price Entropy SD	0.036	**0.039 (0.004)**
		Star Entropy SD	0.055	**0.058 (0.006)**
	Time	Time (Sec) SD	0.004	**0.004 (4.3e-04)**
	Click	Click Base SD	0.023	**0.020 (0.003)**
		Click Position SD	0.006	**0.005 (0.001)**

Table E.1: Monte Carlo estimation results. Bold fonts indicate estimates significant at the 95% level.

E.3 Model Fit

To examine how well the model fits the data, we compare the simulated consumer behaviors with the observed ones in terms of the following moments that are free of model parameters: number of searches, number of clicks, purchase probability, and platform revenue. We consider both the in-sample and out-of-sample settings. In the in-sample setting, both the training and test data are the same as the full data. In the out-of-sample setting, training data is composed of a random sample of 80% of consumers, and test data is composed of the remaining 20%. In either setting, we first estimate the heterogeneous model on the training data, using both the platform's decision context data and the consumer behavior data. Then we sample 400 sets of coefficients from the heterogeneity distribution of the model parameters. For each set of coefficients, we simulate consumer behavior on the respective test data using the platform's decision context data only. We then compare the simulated against the observed behaviors in the test data using the aforementioned moments.

We find that the observed and simulated moments are very close. In terms of the in-sample fit, the average number of searches, number of clicks, purchase probability and revenue for the observed data are 3.769, 3.015, 0.198 and 87.001, respectively, and for the simulated ones are 3.990, 3.260, 0.308, 89.734, respectively. The p-values for testing the differences between the observed and simulated means are 0.478, 0.482, 0.394 and 0.496, respectively, which suggest that we cannot reject the null hypotheses that the observed and simulated means are the same for these behavior measures. Also, in terms of the out-of-sample fit, the average number of searches, number of clicks, purchase probability and revenue for the observed data are 3.718, 2.679, 0.244 and 102.930, respectively, and for the simulated ones are 3.999, 3.313, 0.312 and 88.764, respectively. The p-values for testing the differences between the observed and simulated means are 0.473, 0.455, 0.439 and 0.520, respectively, which again suggest that the observed and simulated means are very close.

F Main Counterfactual

F.1 Behavioral Responses

Within any given plot of Figure F.1, the horizontal coordinate of each point corresponds to the average orderedness entropy over all refinements on the search platform under a particular information layout, and the vertical coordinate corresponds to the average behavioral response simulated under that information layout.

A clear separation exists between two point clouds. The blue point cloud corresponds to the information layouts that do not reveal price, and the green point cloud corresponds to those that do reveal price. Because price is the most complex product attribute to understand, the information layouts in the blue point cloud have lower information complexity than those in the green point cloud, and hence are situated to the left in the horizontal coordinates. Putting the price attribute aside, the information layouts that have the same set of attributes (excluding price) revealed will have the same relative position within each of the two clouds. In other words, the leftmost points in each cloud correspond to the layouts that have few product attributes revealed, whereas the rightmost points correspond to those that have the most attributes revealed. We see that any two such identical information layouts (excluding price) differ in the level of the behavioral responses. The solid lines are the nonparametric regression fit to the point clouds.

The separation of the two point clouds shows that price is an important product attribute whose presence in the outer layer will dramatically change the consumer search behavior. This phenomenon is not only because price is the most complex attribute to understand, but also because it has considerable uncertainty and the presence of this attribute in refinements will change the benefit of the product clicking behavior. We will further explain the impact of the presence of the price attribute when we discuss each behavioral response below.

To understand the product clicking behavior in the “Number of Product Clicks” panel in Figure F.1, we note an inverse relationship exists between the information complexity and the conditional standard deviation of the unrevealed product utility: as the outer layer reveals more attributes, each refinement page has greater information complexity, and each product has less uncertainty for the consumers; furthermore, the conditional standard deviation of the unrevealed product utility becomes lower. Because the conditional standard deviation of the unrevealed utility captures the marginal benefit of clicking, this inverse relationship implies that as the information complexity increases, the benefit of clicking decreases, and hence the number of clicks decreases. The difference between the two point clouds comes from the fact that revealing price in the information layouts reduces the conditional standard deviation of the unrevealed utility, hence lowering the number of product clicks.

The “Number of Refinement Searches” panel in Figure F.1 is a direct illustration of how cognitive cost affects the consumer search process: as information complexity increases, the marginal cost of visiting additional refinement increases, and hence the number of refinement searches decreases. The two point clouds are different because revealing price adds to the cognitive cost and hence reduces the number of refinement searches in the green point cloud. This finding shows that, when considering platform design, search platforms need to take into account the cognitive cost embedded in information layouts.

The previous two results explain the pattern in the “Purchase Probability” panel in Figure F.1: as the consumer searches less, she is less likely to find what she would like to purchase. As a result, we see purchase probability decreases as information complexity increases. Also, purchase probability is lower, at the same relative position, in the information layouts that reveal price versus those that do not reveal price.

The “Purchase Session Utility” panel in Figure F.1 shows how consumer utility at the end of the search process responds to information layouts in a purchase session. Here, consumer utility is defined to be the purchased product or the outside option utility, depending on the consumer purchase decision. Recall that this purchase session utility, like all the other behavioral statistics, is an average quantity over simulation runs under a given information layout, so it is the expected consumer utility as she goes through the search process. On average, consumer utility is higher when the consumer searches and clicks more. This finding is true both within and across the two point clouds.

The “Purchased Product Price” panel in Figure F.1 shows the response of the chosen product price conditional on purchase with respect to the changes in information layouts. When consumers search within information layouts that reveal price, they may directly select on price. As information complexity increases, consumers search fewer refinements, and hence reduce their price selection ability. This leads to higher transaction prices as information complexity increases.

On the other hand, when consumers search within information layouts that do not reveal price, they need to form rational predictions about the product price given the revealed product attributes. This lack of price information limits consumers’ ability to select on price, and the positive search cost leads to general higher purchase prices when price is not revealed than when it is revealed. Furthermore, as more attributes are revealed, consumers may use that information to predict price more accurately, and hence select on expected price, leading to the downward sloping trend of the transaction price.

The "Revenue" panel in Figure F.1 shows the response of revenue to the change in information layouts. Recall that we define revenue to be the purchased product price accounting for purchase probability, so the observed pattern is a natural outcome of the respective panels we have seen above. In general, we see revenue decreases as more information is revealed, and not revealing price generates more revenue than revealing price does. Also, not revealing price makes revenue much more responsive to the change in information complexity than revealing price does.

Search platforms care not only about revenue, but also about consumer welfare, because arguably consumer welfare creates customer loyalty, which makes consumers more likely to return to the search platform even when their outside options increase. Recall that we define consumer welfare to be the purchase session utility (defined above) net of the total accumulated search costs. The "Consumer Welfare" panel in Figure F.1 shows the response of consumer welfare to the variation in information complexity. First, we observe that, in general, revealing price is more beneficial to the consumers than not revealing price. Second, providing consumers with more product information will improve consumer welfare, but too much information provision is not beneficial either. We see that very low information provision makes consumers incur a large amount of product clicks and refinement visits on the extensive margin, which corresponds to high total search cost and welfare loss. At the other extreme, too much information causes consumers to stop searching prematurely, leading to utility loss, and eventually welfare loss. Hence, the optimal information layouts for consumer welfare are in the intermediate range of information complexity.

Finally, we need to be cautious about the implication of the bounded rationality assumption on the counterfactual results. In the multiple sequential search model, we assume that the refinement utility distribution is approximated by the best product utility distribution on each refinement. Under this assumption, when the platform moves product attributes from the outer to the inner layer, the refinement utility distribution is not affected. However, for a fully-forward-looking consumer, when some important product features are moved to the inner layer, she may click more products and pay more clicking costs to find out the attribute information. As a result, not only the consumer's cognitive cost of understanding the outer layer is reduced, but also the benefit of visiting a refinement is likely to be reduced as well. If we were able to solve the fully-forward-looking dynamic programming problem, these effects would be correctly reflected in the value function of visiting a refinement. However, it is neither easy nor feasible to relax the bounded rationality assumption which utilizes the first-order-statistic approximation. This approximation could lead to an *over-estimation* of the inverse relationship between the information complexity and the number of refinement searches. Hence, we need to be aware of the implication and restriction of the approximation when interpreting the counterfactual result.

F.2 Pareto Improving Information Layouts

To understand how the search platform may increase revenue and consumer welfare compared to the current information layout, we select several Pareto improving information layouts, and compare the associated information complexity, consumer welfare, and platform revenue in Table F.1.[8] In particular, we compare layouts that are minimalist style, welfare maximizing, and revenue maximizing, as well as two win-win scenario layouts to the current layout. We define minimalist style layout to be the one that reveals the fewest product attributes, given that it is a Pareto improving layout.

We see the minimalist style layout simplifies the current layout by removing three attributes from display. This simplification increases consumer welfare by 15.28% while keeping the current revenue level intact. In the consumer welfare maximizing layout, the platform may in addition display two more attributes that consumers care about to reduce search cost. This addition will increase consumer welfare by 95% without losing revenue. In the revenue maximizing layout, some important product attributes, including price, are replaced with less important ones. The main driver of the revenue hike is that price is not revealed, which will reduce consumers' direct selection on price. The

[8] Similar to the counterfactual analysis in most structural model papers, our observational data set does not include the information on how the subjects respond to the policy change. As a result, we need to rely on the model to extrapolate beyond the data variations we observe. In particular, as we perturb the information layout in the consumer's search context, we need to assume that the consumer search mechanisms, namely the benefit and cost trade-offs behind the consumer's decisions, including refinement searching, product clicking and purchasing, remain the same.

As there is *no data to verify* the predicted consumer's response given by the model, taking the counterfactual results too literally could be misleading. Intuitively, the farther away the alternative layout from the current layout, the more room there is for the deviation from the model-predicted behavior. We include the whole spectrum of the counterfactual settings in order to gain a more complete understanding of the model behavior. As the Pareto improving layouts we consider are relatively close to the current layout, they are also relatively more robust. For practitioners, the Pareto improving layouts suggested by the model could be potentially interesting directions to further explore in the vast strategy search space.

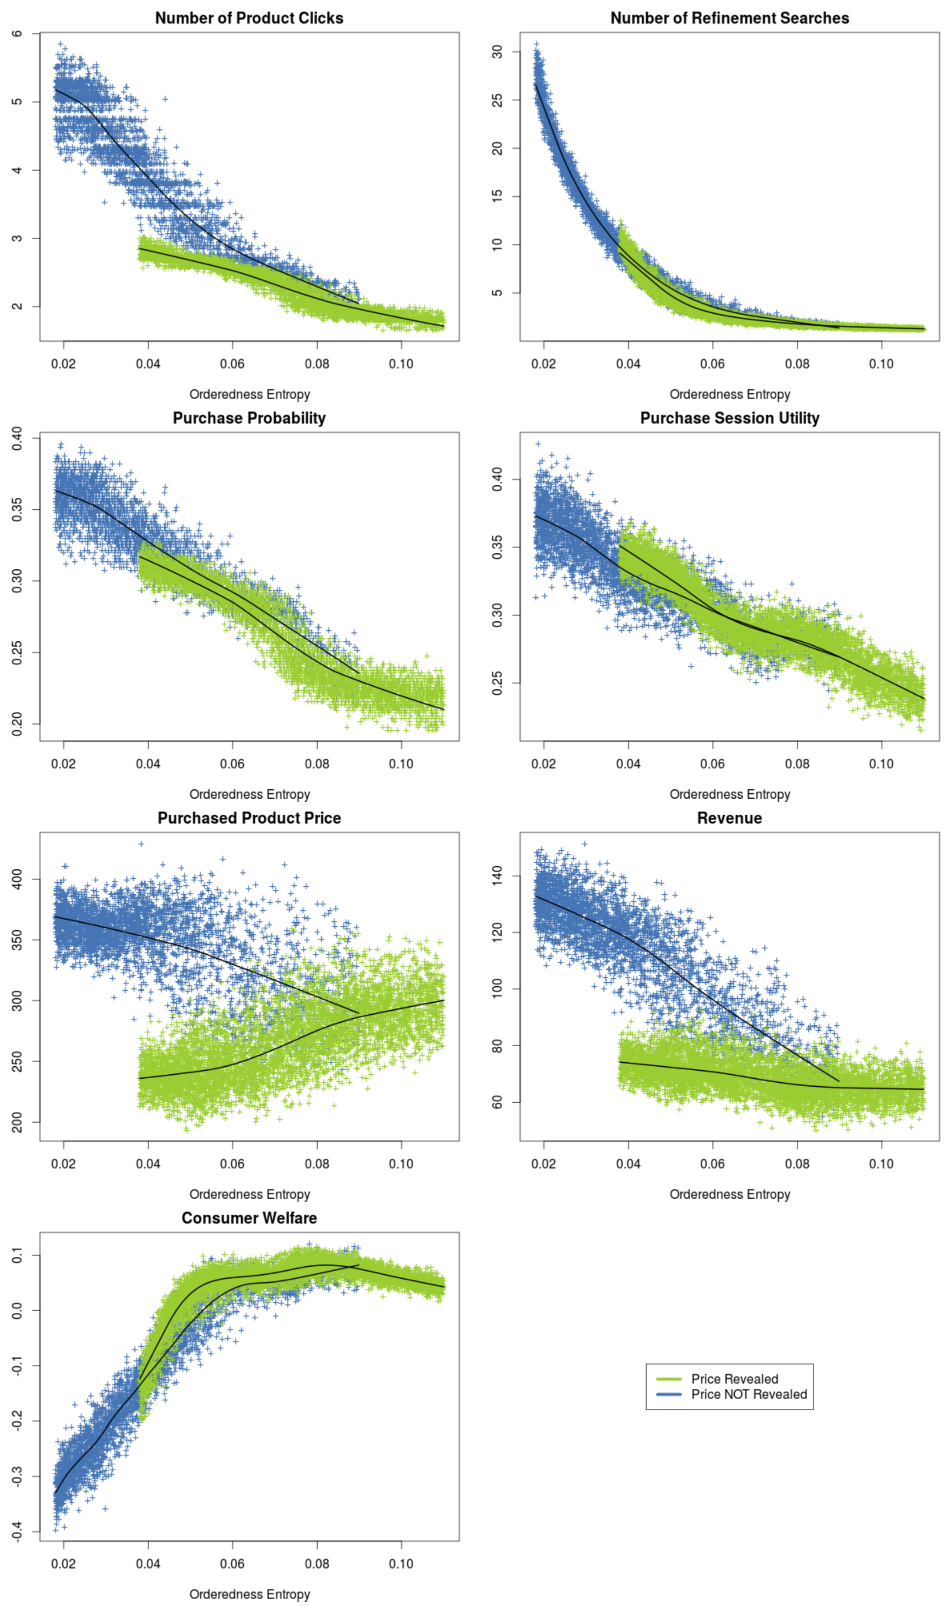

Figure F.1: This figure shows the various behavioral responses of the consumer search process as information layouts vary. Within any given plot, the horizontal coordinate of each point corresponds to the average orderedness entropy over all refinements on the search platform under a particular information layout, and the vertical coordinate corresponds to the average behavioral response simulated under that information layout.

addition of the less important attributes will deter consumers from searching excessively, avoiding too much search cost. This change in layout results in about a 60% increase in revenue without hurting consumer welfare.

Many win-win information layouts exist in which both consumer welfare and platform revenue may be improved. We have selected two such layouts, one revealing product price and the other not. When product price is revealed, we may remove some attributes that consumers care about. This reduces the cognitive cost of search refinements, which results in a 46% increase in consumer welfare. At the same time, consumers are more likely to choose the higher priced product, using price as predictors for the important unrevealed attributes. This results in a 25% increase in revenue. When price is not revealed, we may reveal some additional attributes to consumers to compensate for price and prevent consumers from engaging in excessive search. Overall consumer search cost is reduced. This change in layout leads to a 19% increase in consumer welfare and a 53% increase in platform revenue.

G Alternative Counterfactual

Online platforms often display products according to aggregate popularity ranking as the default sorting method, measured by, for example, sales volume, and this is the case for the platform in our study. However, such ranking does not account for the cognitive cost of comprehending and comparing products displayed. In this counterfactual study, we examine if we could Pareto improve both platform revenue and consumer welfare by having a more ordered presentation for the refinements using the default ranking.[9]

In particular, for a given refinement using the default ranking, we apply "local sorting" as opposed to "global sorting." Global sorting is the same as applying regular sorting variable. For a search result that spans multiple pages, global sorting will completely reorder the products over different paginations. In contrast, local sorting retains the set of products on a given page and reorders them without extending beyond the pagination boundary. Local sorting is relevant to platforms, because platforms may have special sponsorship agreements with product vendors by placing products on certain pages of particular search results. Local sorting could retain such agreements, while poten-

[9]We thank an anonymous referee for suggesting this counterfactual study.

	Current	Minimalist	Max Welfare	Max Revenue	Win-Win 1	Win-Win 2
Information Complexity	0.0663	0.0484	0.0787	0.0629	0.0574	0.0566
Consumer Welfare	0.0563	0.0649	0.1097	0.0620	0.0821	0.0670
Revenue	67.11	68.56	67.48	107.12	83.73	102.47
Price	x	x	x		x	
Consumer Rating	x	x	x		x	x
Hotel Star	x	x	x	x	x	x
Breakfast	x		x		x	x
Cancel Allowed	x	x	x		x	x
Wifi in room	x	x	x	x	x	x
Kingsize bed	x	x	x	x	x	x
Wi-Fi in public areas	x	x	x	x	x	x
Parking	x	x	x	x	x	
Meeting rooms	x	x	x	x		x
Gym	x	x	x	x		x
Restaurant	x		x	x	x	x
Airport pickup service	x		x	x	x	
Outdoor swimming pool	x	x	x		x	x
Free toiletries (less than six items)				x		
Free toiletries (more than six items)				x		x
Free domestic long-distance calls				x		x
Free bottled water				x		
Free map			x	x		x
Free newspaper						
Pets Allowed						
Free Afternoon Tea						x
Free Guide Service			x	x		

Table F.1: We select several Pareto improving information layouts, and compare the associated levels of information complexity, consumer welfare, and platform revenue in order to understand how the search platform may increase revenue and consumer welfare compared to the current information layout. An "x" in the layout column represents the revelation of the respective product attribute.

tially help both platforms and consumers. We apply local sorting to a given refinement page using a product attribute that reduces entropy cost most. Then we simulate 500 runs of the consumer search process using the mean estimation result in Section 9 of the paper and the randomly sampled refinement realizations. We then average the consumer behavior moments over the simulation runs as the typical consumer response to this platform modification, detailed in Table G.1.

We see that information complexity, which is the average refinement entropy cost, has reduced by 8.5% due to the local sorting of the default-ranking refinements. Also, the number of refinement searches has decreased by about 22.4%, but the fraction of refinement searches using the locally sorted default ranking has increased by 19.8%. At the same time, the number of product clicks has increased by 7.4%.

The counterfactual findings are very intuitive. As the information complexity of the default-ranking refinements decreases, consumers visit these refinements more, which can be seen from the increasing fraction of the refinement searches using the locally sorted default ranking. Provided that the default-ranking refinements are based on aggregate popularity ranking, products that better fit the average consumer tastes are ranked higher, and the local sorting mostly maintains the overall popularity ranking. In contrast, global sorting completely reshuffles products along one dimension of product attributes over different paginations, which may not be ideal in other product attributes. Consumers need to search over more refinements using global sorting in order to find what they like compared to using locally sorted default ranking. This explains why in the counterfactual setting, the number of refinement searches decreases as the consumers use the locally sorted default ranking more, and the number of product clicks increases as the default-ranking refinements cater more preferable product options. As a result, purchase session utility has increased by 7.3%, and total search cost has reduced by 5.7%, which leads to an increase of consumer welfare by 59.3%. The revenue remains about the same.

In summary, by accounting for the cognitive cost that consumers may incur during their search process, local sorting of refinements using the default ranking is a Pareto improving strategy that platforms may consider adopting to increase consumer welfare without losing on platform revenue.

	Current	Locally Sorted
Information Complexity	0.0663	0.0607
Number of Refinement Searches	2.2580	1.7520
Fraction of Refinement Searches using Locally Sorted Default Ranking	0.7307	0.8756
Number of Product Clicks	2.4980	2.6820
Purchase Session Utility	0.2799	0.3004
Total Search Cost	0.2236	0.2108
Revenue	67.11	67.61
Consumer Welfare	0.0563	0.0897

Table G.1: The comparison of consumer behaviors under the current platform design versus the modification where the refinements using the default ranking are made more orderly by locally sorting the attribute that reduces the cognitive cost most. We see that local sorting is a Pareto improving strategy that platforms may consider adopting to increase consumer welfare without losing on platform revenue.