Search Duration*

Raluca M. Ursu[†] Qingliang Wang[‡] Pradeep K. Chintagunta[§]

Current version: August 21, 2018
First version: June 20, 2017

PRELIMINARY AND INCOMPLETE
PLEASE DO NOT CITE OR CIRCULATE WITHOUT THE AUTHORS' PERMISSION

Abstract

In studying consumer search behavior, researchers typically focus on which products consumers add to their consideration sets (the extensive margin of search). In this article, we attempt to additionally study how much consumers search individual products (the intensive margin of search), by analyzing the time they spend searching (search duration). We develop a sequential search model where consumers who are uncertain (and have prior beliefs) about their match value for a product search to reveal (noisy) signals about it that they then use to update their beliefs in a Bayesian fashion. Search duration, in this context, is an outcome of the decision by a consumer to seek information on the same product multiple times; with a unit of time corresponding to one signal, the more the number of signals sought greater is the search duration. We also show how the model can be used to study revisits - a feature not easily accommodated in Weitzman's (1979) sequential search model. We build on the framework by Chick and Frazier [Chick, S., P. Frazier. 2012. Sequential Sampling with Economics of Selection Procedures. *Management Sci.* **58**(3), 1-16] for describing the optimal search rules for the full set of decisions consumers make (which products to search, for how long, and whether to purchase), and develop the model's empirical counterpart. We estimate the proposed model using data on consumers searching for restaurants online. We document that search duration is considerable, even when consumers search few restaurants, and that restaurants that are searched longer are more likely to be purchased. Using our model, we quantify consumer preferences, search costs, and prior uncertainty parameters. We find that consumers start their search with high initial uncertainty, which rationalizes their choice to search few options intensively in the presence of search costs. In counterfactual simulations, we analyze the effect of prior uncertainty and search environment on consumer choices.

Keywords: online consumer search, sequential sampling, search duration, search with learning, optimal search rules, revisits, online reviews.

---

*We are thankful for comments from Paulo Albuquerque, Simon Anderson, Ron Berman, Eric Bradlow, Bart Bronnenberg, Stephen Chick, Babur De los Santos, JP Dube, Daria Dzyabura, Tulin Erdem, Pedro Gardete, Wesley Hartmann, Elisabeth Honka, Ali Hortacsu, Xiao Liu, Xueming Luo, Carl Mela, Eitan Muller, Harikesh Nair, Sridhar Narayanan, Thomas Otter, Navdeep Sahni, Stephan Seiler, Andrey Simonov, Miguel Villas-Boas, Russell Winer, Pinar Yildirim, and seminar participants at Stanford University, Wharton Marketing Camp, Virginia University, Temple University, and attendees of the 2017 Marketing Science/INFORMS, 2017 QME, and 2018 SICS conferences for their thoughtful comments. The usual disclaimer applies.

[†]Stern School of Business, New York University, E-mail: rursu@stern.nyu.edu.
[‡]Northwestern Politechnical University, E-mail: qingliang.wang@nwpu.edu.cn. Corresponding author.
[§]Booth School of Business, University of Chicago, E-mail: pradeep.chintagunta@chicagobooth.edu

1 Introduction

Nowadays, consumers have access to a plethora of information, especially online. This additional information allows consumers to make better or more informed choices. At the same time, paying attention to this information is costly. To understand how consumers make choices in such an environment, previous work has devoted a considerable amount of attention to the question of which products consumers add to their consideration sets before making a purchase decision. We refer to this type of consumer search decision as the extensive margin of search. In contrast, relatively little is known about how much consumers choose to search individual products, or what we refer to as the intensive margin of search. However, examples of such search decisions abound including the decision to spend time searching a product, the decision to revisit a previously searched option to resolve further uncertainty about it, etc.

In this article, we develop a sequential search model that endogenizes not only the decisions to search and purchase a product, but also the consumer's decision to spend time searching. In this model, consumers who are uncertain (and have prior beliefs) about their match value for a product, search to reveal (noisy) signals about it that they then use to update their beliefs in a Bayesian fashion. We model search duration as the consumer's decision to search the same product multiple times. This approach allows us to use the sequential sampling theory developed by Chick and Frazier (2012) to characterize consumers' optimal search rules. This model captures richer patterns of consumer search behavior, relating not only search costs and consumer preferences, but also their prior awareness of the product. We build on the work of Chick and Frazier (2012) and develop the model's empirical counterpart, which we then estimate with data from an online restaurant review website.

Our approach differs from related models such as the sequential search model (Weitzman, 1979), the multiarmed bandit model (Gittins 1979; Brezzi and Lai 2002), and the model of consumer search with learning (Rothschild, 1974; Koulayev, 2013; De los Santos et al., 2016). Since it allows consumers to search the same option multiple times before a single purchase decision, it differs from the sequential search and the multiarmed bandit models. In addition, the fact that consumers learn about their match values rather than the distribution of rewards in the market distinguishes it from existing models of search with learning following Rothschild (1974). Our approach provides a general framework to study consumer engagement with a product through search, being able to also capture decisions such as revisits to a previously searched product to resolve further uncertainty.

The empirical setting in which we estimate the proposed model is consumer search for restaurants on an Asian review website. On this website, consumers start their search by visiting the homepage and specifying a query, that is either searching for a restaurant by typing in a keyword or using one of the menu options, such as cuisine type or location. In response to this query, consumers see an ordered list of restaurants, which contains some information about each option, such as average reviews and price. The consumer can then search a restaurant to obtain additional information by clicking on it. In this case, consumers navigate to a second screen dedicated to that restaurant. Finally, the consumer can decide to visit the restaurant (offline) when her search ceases. We have data on all these restaurant characteristics and the choices consumers make in response to the information observed.

These data provides insights on consumers searching on both the intensive and the extensive margins. More precisely, we have data on which products consumers clicked/searched (extensive margin), how much time they spent searching each product (intensive margin), and which products they purchased, if any. Using these data, we find that search duration is considerable, even when consumers search few restaurants. In addition, we find that search duration is related to purchases, so that restaurants that consumers spend more time searching are also the ones they purchase. For a given restaurant, we find that the more information is displayed upon search, the longer is the search duration. This suggests that duration imposes a cost by preventing consumers from searching additional options, and at the same time provides a benefit in the form of information discovered through search that might influence purchase decisions. That is, search duration is a choice made in addition to search and purchase decisions. This finding makes the empirical setting favorable for estimating a model that endogenizes the consumer's search duration decision.

To estimate the model, we first adapt it to the empirical setting. More precisely, we assume consumers search to decide whether to visit (purchase) a restaurant. The utility from visiting a restaurant is influenced by consumers' preferences for restaurant characteristics observed on the list page before search, and their beliefs about product match values given information observed through search up to that point. Consumer beliefs are influenced by their prior uncertainty, as well as by the restaurant page characteristics observed through search. Search is costly but provides (noisy) signals about the restaurant's match value that are a function of the restaurant page information observed. Consumers then use these signals to update their beliefs about restaurant match values in a Bayesian fashion. At each stage, consumers decide whether to continue spending time on that restaurant's page; return to the list page and click a different restaurant; or make a purchase decision. The more time spent on the restaurant's page, the more information consumers can gather.

We estimate our model and quantify consumer preferences, search costs, and prior uncertainty parameters. The model allows estimation not only of consumers' expected utility before search (preferences for restaurant characteristics observed on the list page), but also of their preferences for information discovered through search (preferences for restaurant characteristics observed on the restaurant page), better recovering consumer preferences from the data. In addition, adopting a framework where consumers learn their match values by searching to update their prior beliefs allows us to estimate consumers' prior uncertainty and how it relates to consumer preferences and search costs. We find that consumers start their search with high initial uncertainty, which rationalizes their choice to search few options intensively in the presence of search costs. This result sheds new light on previous results that estimated large search costs: without information on search duration, the fact that consumers search few options suggests that consumers have large search costs. However, with information on search duration, we show that the reason consumers search few options is because they start with high initial uncertainty and spend a considerable amount of time searching those options. Finally, we use a latent class model to study consumer heterogeneity and uncover an additional layer of the relation between consumers' prior uncertainty and search costs. We find that one segment of consumers searches many options for a short amount of time since it has lower search costs and begins search with lower uncertainty than a second segment that searches few options intensively. Without information on search duration, the impact of prior uncertainty on search would be misattributed mainly to differences in search costs.

In counterfactual simulations, we investigate the importance of consumer prior knowledge and search environment in affecting consumer search and purchase decisions. As such, this analysis is relevant to managers of platforms or intermediaries that determine the format in which consumers interact with information on third party sellers. We consider two questions. First, we ask what is the impact on searches and purchases of the fact that as the website matures, consumers gain greater knowledge of the products displayed, and thus start search with lower uncertainty. We find that lower prior uncertainty leads to fewer searches, since consumers know more about the options available, and thus to fewer transactions. Second, an important decision for managers is to determine what information to make available before and after consumers click. To answer this question, we consider the effect of shifting restaurant page information to the list page which we find increases both searches and transactions.

The rest of the paper is organized as follows. The next section discusses relevant prior work. Section 3 introduces our model and describes consumers' optimal search rules. In section 4, we discuss the specific empirical context of our analysis, the data we employ, and provide descriptive statistics on search duration. Section 5 describes our empirical model and estimation procedure, as well as its identification. In section 6 we describe our estimation results, while in section 7 we provide managerial implications using a counterfactual. The last section concludes.

2 Literature

This paper connects two literature strands: consumer search and Bayesian learning. The consumer search literature generally follows either Stigler's (1961) theoretical model of simultaneous search or Weitzman's (1979) sequential search theory. In these models, consumers know the distribution of rewards (e.g. price, match values, etc), but search to reveal the reward of a specific product. Search reveals all uncertainty about the searched product. Both papers derive optimal search rules for consumers. In the case of simultaneous search, consumers search the set of products providing the maximum expected reward net of search costs. In the case of sequential search, the optimal search rules are characterized by a reservation policy: consumers search options in order of their reservation values (an index describing the hypothetical observed reward that would make the consumer indifferent between searching and stopping), stop when no alternative has a reservation value greater than the realized reward of the searched alternatives, and buy if the product with the highest realized reward is greater than the value of not purchasing.

The empirical search literature quantifies the impact of search frictions on search and purchase decisions (Hong and Shum, 2006; Moraga-Gonzalez and Wildenbeest, 2008; Kim et al., 2010, 2016; Ghose et al., 2012; Seiler, 2013; Honka, 2014; Koulayev, 2014; Moraga-Gonzalez et al., 2015; Chen and Yao, 2016; Honka and Chintagunta, 2016) by estimating consumer preferences and their search cost parameters. Ours is also an empirical search paper. However, it differs from this literature because in our model consumer information is not revealed fully in one search action, but rather consumers update their prior beliefs on product match values in a Bayesian fashion through search. As such, our paper is related to the literature on Bayesian learning in dynamic brand choice models. Papers such as Erdem and Keane (1996), Ackerberg (2003), and Erdem, Keane, and Sun (2008), consider the signaling effect on product quality of a host of factors, such advertising content and frequency, experience and prices. Researchers estimate consumer preferences jointly with their prior uncertainty about products. Our model also differs from this work in that we consider learning through search prior to a purchase decision, rather than through purchases. By combining the two strands of the literature on consumer search and Bayesian learning, we develop a model that allows consumers to learn their match value while searching sequentially for product information. With this model, we are able to quantify a richer set of factors influencing consumer choices that combine those of these two literature strands. More precisely, we estimate consumer preferences, search costs and prior uncertainty, which provides new insights relating three stages of the consumer decision making journey: choice, search and awareness.

Since we model search duration as the consumer's decision to search the same option multiple times, naturally our model is related to work on multiarmed bandit problems (Gittins 1979; Brezzi and Lai 2002). In these models, consumers look to maximize the sum of rewards from sequentially sampling options, including sampling the same option multiple times. Such models are suitable to study repeat purchase occasions, as it is done in Lin, Zhang and Hauser (2014). In contrast, in our model consumers maximize the rewards from a single option chosen after sequentially searching for information about available options. As a result, we model repeated search decisions (e.g. time spent searching, revisits) and a single purchase decision.

Most closely related to our work are papers that model search with learning. In Table 1 below we provide a quick overview of papers in this literature and show how our paper fits in. Papers are ordered chronologically and by the type of search learning assumed: gathering information on product attributes, learning about the market distribution of the product characteristic consumers search for (e.g. price), or learning about consumers' match values. We also differentiate papers by theoretical versus empirical work, by whether they model search as simultaneous or sequential, by whether they derive or use optimal search rules to describe consumer behavior, by the type of distribution assumed for the beliefs of consumers, and finally by the number of products considered.

The first group of papers models search for product attributes. In these models, consumers have some basic product information and decide sequentially whether to obtain additional information on other product features. Branco et al. (2012, 2016) consider the case of one product and study how the optimal stopping rule (there is no selection rule in this case, since the model contains just one product) depends on model parameters, such as search costs. Ke et al. (2016) focus

Table 1: Literature on search with learning

	Learning	Theory/ Empirical	Simultaneous/ Sequential	Optimal search rules	Distribution of beliefs	Number of products
Branco et al. (2012)	Attributes	Theory	Sequential	Yes*	Symmetric	1
Branco et al. (2016)	Attributes	Theory	Sequential	Yes*	Symmetric	1
Ke et al. (2016)	Attributes	Theory	Sequential	Yes	Symmetric	≥ 2
Gardete and Hunter (2018)	Attributes	Both	Sequential	No	Empirical	N
Rothschild (1974)	Market	Theory	Sequential	Yes	General	N
Koulayev (2013)	Market	Empirical	Sequential	Yes	Dirichlet	N
De los Santos et al. (2016)	Market	Empirical	Sequential	Yes	Dirichlet	N
Hu et al. (2017)	Market	Empirical	NA	Yes	Dirichlet	N
Chick and Frazier (2012)	Match	Theory	Sequential	Yes	Normal	N
Dukes and Liu (2015)	Match	Theory	Simultaneous	Yes	Extreme value	N
Ma (2016)	Match	Empirical	Sequential	No	Normal	N
Ke and Villas-Boas (2017)	Match	Theory	Sequential	Yes	Two point	≥ 2
Our paper	**Match**	**Empirical**	**Sequential**	**Yes**	**Normal**	**N**

*=selection rule NA

on two products that they later extend to more products, and derive optimal search rules when attributes are independent and search informativeness is constant. Most recently, Gardete and Hunter (2018) propose a model where consumers can search to learn about products with possibly correlated characteristics. Consumers can search the same option multiple times by choosing what characteristic to learn about. The authors then estimate this model on online used car dealer data, assuming consumers are myopic, that is consumers focus only on the highest expected immediate benefit when deciding whether and what to search. In contrast to this set of papers, we do not observe what characteristics consumers reveal through search, but only observe the time they spend searching, which we model as search to learn consumers' match value for the product. Also, we use Chick and Frazier (2012) to characterize the optimal search rules for forward looking consumers and any number of alternatives. Similar to Gardete and Hunter (2018), in the counterfactual we consider the effect of the search environment on search decisions.

In the second group of papers, Koulayev (2013) and De los Santos et al. (2016) follow the theoretical framework of Rothschild (1974) and assume consumers search for the highest reward (for example, the lowest price), while at the same time learning about the market distribution of rewards. They form beliefs about this distribution (both assume Dirichlet priors), search to reveal information about one company's product and update their beliefs about the distribution using Bayes' rule, which they use to decide whether to search another product. In this setting, the optimal search rule is also an index/reservation policy as in the case of search without learning (e.g. in Weitzman, 1979), except that here reservation utilities are non increasing over time, so consumers are more likely to accept an offer over time in the model with learning than in the one without. In contrast to our setting, here consumers will not search the same product more than once (except in order to return and accept a previous offer). Also in our model consumers learn about their individual match values (instead of the distribution of rewards across all products in the market), allowing them to decide when it is optimal to switch from learning about one product to learning about another. Recently, Hu et al. (2017) model the decisions of consumers who observe Groupon deals daily, decide whether to click, learn about the distribution of deals if they click, and make purchase decisions. This model also differs from ours in several respects: search is assumed passive (deals arrive every day, rather than consumers seeking out options to search), and consumers learn about the market distribution of rewards rather than their match value. As a result, the problem that consumers are solving in these settings is fundamentally different from ours, leading to different optimal search rules and different consumer behavior.

Most closely related to our work, papers in the third group model learning for match values. Dukes and Liu (2015) investigate the strategic interplay between search intermediaries, companies and consumers, when consumers decide optimally both the extensive and the intensive margins of their search. Under simultaneous search and assuming consumers choose the same intensive margin of search for all searched products, they show how search costs, the intermediary's search engine design and firm pricing decisions affect the amount of search in equilibrium. In our paper, we focus on the demand side of the model and estimate a sequential search model using consumer optimal search rules. Ma (2016) embeds an Erdem and Keane (1996) like framework into a sequential search model, where consumers can learn about product quality by choosing whether to observe signals from different types of product reviews (e.g. star rating versus photos). However, the paper does not use/derive the optimal search rules for the full set of decisions made by consumers. Finally, Ke and Villas-Boas (2017) focus on the case of two products (which they later extend to three products) and derive the optimal sequential search strategy when rewards are drawn from a two point distribution. In contrast, we follow Chick and Frazier (2012) who derive optimal search rules when consumers search sequentially among any number of alternatives and hold normally distributed beliefs. In addition, we take this model to data and propose an estimation strategy for it.

Two recent papers provide data patterns on search duration. More precisely, De los Santos (2017) finds that, in searching for books, duration is affected both by previous consumer choices (e.g. past bookstores visited) and their demographics, as older consumers with lower education or income levels tend to spend more time searching. This shows that consumers' opportunity cost of time is an important factor in the decision to search on the intensive margin. Seiler and Pinna (2017), without modeling search on the extensive margin, measure the change in price paid from spending an additional minute searching in a super market setting and find a benefit of $2.10 per minute. Both of these papers provide insights on the importance of search duration hinting at its dual effects: providing a benefit but at a measurable cost. This suggests the need to quantify these two effects using a model that endogenizes the search duration decision, which is the focus of this paper.

3 Model

3.1 Consumer problem

Consider a consumer $i \in \{1, \ldots, N\}$ who seeks to purchase an alternative $j \in \{1, \ldots, J\}$ or choose the outside option (denoted by $j = 0$). For notational simplicity, we suppress the subscript i in what follows and present the model from the perspective of one consumer. However, it is understood that there are N consumers in the market and we will account for possible heterogeneity in their preferences and search costs in the estimation of our model (see section 5). The expected utility of the outside option is known, but the consumer faces uncertainty about the J options. To (partially) resolve this uncertainty, the consumer can search for information before making a purchase decision, which involves paying a cost per search, $c_j > 0$. The consumer's goal is to maximize her expected utility net of total search costs from the best option she will choose to purchase when search ceases.

Searching an alternative once does not resolve all the consumer's uncertainty. To obtain further information, the consumer can search the same option in multiple time periods, where time is discrete and indexed by $t = 1, \ldots, T$. To model search duration, we will interpret time spent searching an option as the consumer's choice to search the same option multiple times. Although we focus on search duration, this same approach can be used to model revisits of previously searched options, as we demonstrate below (see Figure 2).

Each time period, the consumer decides whether to continue searching, in which case she chooses a product to search, or whether to stop, in which case, she decides which product to purchase, if any. We model the consumer's utility from purchasing product j at time t as

$$u_{jt} = \mu_{jt} + \epsilon_j, \tag{1}$$

where μ_{jt} is the consumer's perceived match value with product j in period t, and $\epsilon_j \sim N(0, \sigma_\epsilon^2)$ is an idiosyncratic shock, unobserved by the researcher, but known to the consumer before search.[1]

The consumer is uncertain about the true match value of each of the J alternatives, which is normally distributed with unknown mean μ_j and known variance σ_j^2. She holds beliefs about her match value, which she updates in a Bayesian fashion using information gained through search. More precisely, the consumer can search to learn the unknown mean by obtaining (unbiased) signals on j at t given by

$$s_{jt} \sim N(\mu_j, \sigma_j^2), \tag{2}$$

at a cost of c_j per search. Because draws are independent, the consumer does not learn about the match value of one product by searching another. Initially, the consumer's prior beliefs are summarized by

$$N(\mu_{j0}, \sigma_j^2 / n_{j0}), \tag{3}$$

[1] Consumers search to learn about μ_{jt} and know ϵ_j before search. Thus, the search rules derived by Chick and Frazier (2012) that we present below (section 3.2) are optimal in our case.

where μ_{j0} is the prior mean and n_{j0} gives the implied number of samples drawn to form the prior belief. We follow Chick and Frazier (2012) to write the prior variance as a ratio of the true variance σ_j^2 and n_{j0}. This allows us to use their results to characterize the optimal search rules of the consumer, as we do in the next section. Also, this approach simplifies the exposition of the model in light of the fact that only the ratio of prior and signal variances can be recovered through estimation (for more details on the identification of the empirical model, see section 5.3).

After searching j at t, the consumer's posterior belief about her match value is formed using Bayes' rule and equals

$$N(\mu_{jt+1}, \sigma_j^2/n_{jt+1}), \tag{4}$$

where

$$\begin{aligned} \mu_{jt+1} &= \frac{n_{jt}\mu_{jt} + s_{jt}}{n_{jt+1}} \\ n_{jt+1} &= n_{jt} + 1, \end{aligned} \tag{5}$$

while for $k \neq j$, $\mu_{kt+1} = \mu_{kt}$ and $n_{kt+1} = n_{kt}$, that is, options that are not searched are not updated.

In contrast to the J alternatives' unknown mean match value, the outside option of not purchasing has a known expected utility, which we normalize to zero, that is $u_0 = \epsilon_0$. The interpretation of this assumption is that the consumer chooses between one of the J alternatives or rejects all of them, obtaining zero mean utility.

Figure 1: Model illustration: Sequential search

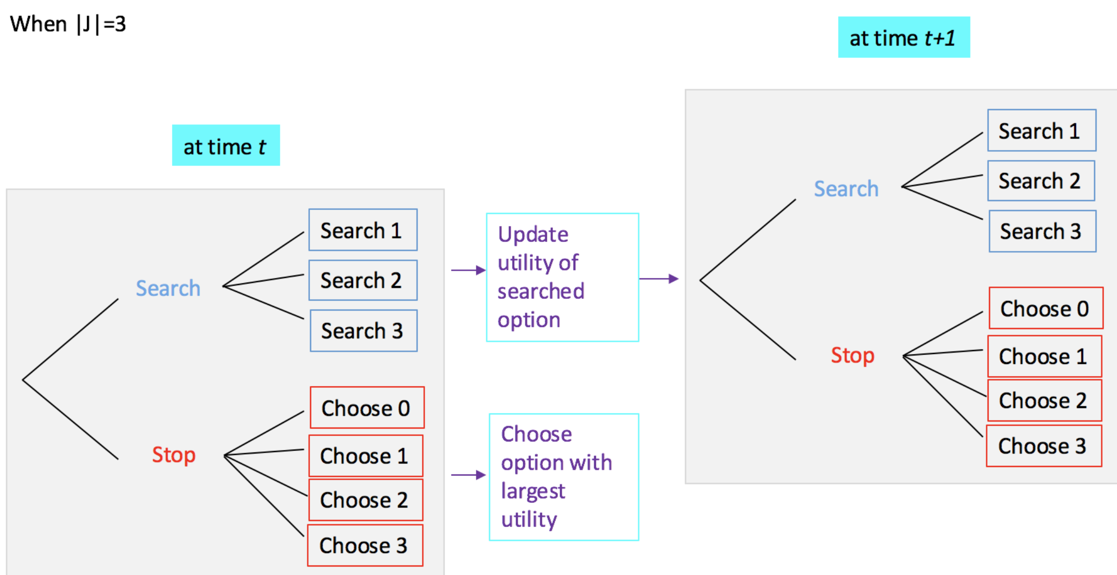

To illustrate the consumer's sequential search procedure, consider Figure 1. Suppose the consumer has three options to choose from. In a given time period t, the consumer can either continue or stop searching. If she continues searching, then she must choose the option to search next. If instead she decides to stop searching, then she chooses whether to purchase one of the three options or choose the outside option (option 0) of not purchasing. At any stage in the search process, the consumer uses all data observed thus far to make a decision. For example, at time t, if she chooses to search option 2, then she observes a signal about this product, which she uses to update her belief about her posterior utility of 2. Her beliefs about the other two options stay the same since the consumer did not observe any new information about these. Then in the next period, she again can choose to continue searching or to stop. Importantly, she can choose to search any of the three options available, including the previously searched option 2. The possibility of searching the same option as before is what distinguishes this model from previous search models in the literature that assume all the uncertainty of the consumer about a product is resolved in a single search action.

As mentioned previously, this same model can also be used to model revisits of previously searched options. To see why this is the case, consider Figure 2. Suppose, for example that the consumer searches option 2 in period t and switches to searching option 1 in period $t+1$. The decision to search option 2 again in the next period, which defines a revisit, can be accommodated by our model, since searching an option does not reveal all uncertainty about it, so previously searched options can be searched again.

Figure 2: Model illustration: Revisits

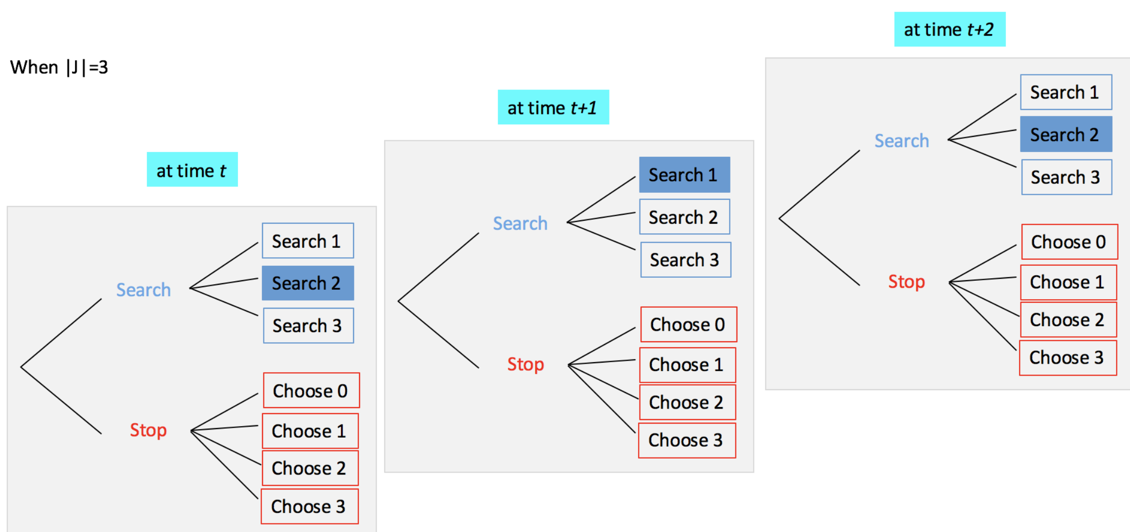

3.2 Optimal search

At a given point during search, the consumer must decide whether to continue searching, and if so, which alternative to sample next. Upon stopping her search, she must decide whether to purchase. To model consumer's optimal search decision, we follow Chick and Frazier (2012). They provide an optimal policy of choosing at each time period whether to continue searching, if so which alternative, and upon termination, which alternative to choose. To characterize such a policy, let the state of information about option j at t be given by $\Theta_{jt} = (\mu_{jt}, n_{jt})$ and the state of the system at t be

$\vec{\Theta}_{jt} = (\Theta_{0t}, \Theta_{1t}, \ldots, \Theta_{Jt})$. The optimal policy is one that determines which j to search/purchase at each t given $\vec{\Theta}_{jt}$ in order to maximize the expected utility from the outcome chosen once search terminates net of total search costs. Chick and Frazier (2012) frame this problem using dynamic programming and show that the optimal policy is one that attains the maximum of the following Bellman recursion problem

$$V(\vec{\Theta}_t) = \max\left\{\max_{j=1,\ldots,J} E(-c_j + V(\vec{\Theta}_{t+1}|\vec{\Theta}_t), \max_{j=0,1,\ldots,J} E(u_{jt}|\vec{\Theta}_t)\right\}. \tag{6}$$

Chick and Frazier (2012) solve for the optimal policy in two steps. First, after proving the existence of an upper bound on the total number of searches a consumer will make, they consider the case with one alternative ($J = 1$), one outside option, and normally distributed rewards with unknown means and known variances. This problem can be solved using the Bellman recursion above given the upper bound on the total number of searches. However, the solution depends on parameters of the problem and it would have to be recomputed when these change. Thus, instead of this approach, they choose to transform the discrete-time problem to continuous-time and use diffusion approximation to describe the solution (similar to approaches used for multiarmed bandit problems, e.g. Chernoff and Ray 1965, Lai 1987, Brezzi and Lai 2002, Chick and Gans 2009). This approach leads to a solution that is independent of parameters when $J = 1$. Second, they use results from the case of one alternative to provide approximations to the solution for the case of $J > 1$. We note that the search rules may not be optimal when $J > 1$, because they are derived from an approximation to the dynamic programming problem. However, Chick and Frazier (2012) show that these perform very well when using numerical results and they are easier to implement than solving the dynamic programming problem using Bellman recursion.

The optimal policy for the case of $J > 1$ is characterized by three search rules. We follow the search rules based on the stopping boundary that Chick and Frazier (2012) derive:[2]

1. **Stopping Rule**: Continue to search at t if and only if $\exists\ j \in (1, \ldots, J)$ such that its posterior mean utility u_{jt} lies within the *continuation set*, that is, $u_{jt} \in (\max_{k \neq j} u_{kt} \pm M_{jt}(c_j, \sigma_j, n_{jt}))$, for $k \in (0, 1, \ldots, J)$, where M_{jt} is the boundary of search, a function of search costs and product uncertainty that we define in equation (9) below. This condition can be rewritten as follows: search will continue at t if and only if $\exists\ j \in (1, \ldots, J)$ such that

$$M_{jt}(c_j, \sigma_j, n_{jt}) > \Delta_{jt}, \tag{7}$$

where $\Delta_{jt} = |u_{jt} - \max_{k \neq j} u_{kt}|$ for $k \in (0, 1, \ldots, J)$.

2. **Selection Rule**: While the stopping rule is not satisfied, choose to sample the alternative $j \in (1, \ldots, J)$ such that

$$\arg\max_j \frac{M_{jt}(c_j, \sigma_j, n_{jt}) - \Delta_{jt}}{c_j^{1/3} \sigma_j^{2/3}}. \tag{8}$$

[2]This is the approach recommended by Chick and Frazier (2012), because of its ease of implementation.

3. **Choice Rule**: Conditional on stopping, choose the alternative $j \in (0, 1, \ldots, J)$ with the largest posterior expected utility.

The optimal search rules can be understood as follows. The stopping rule dictates that the consumer will continue searching if at least one alternative falls within the continuation set, that is if comparing its posterior mean utility and the maximum posterior mean utility of all other alternatives (and the outside option) is smaller than the boundary of search. Thus, all alternatives that fall within the continuation set are potential candidates for further search, while those outside the continuation set will not be searched at t (although they might be searched at a different time). The selection rule says that if the consumer finds it optimal to search at t, then she should search the option that is furthest inside the continuation set as measured in standardized coordinates. Finally, the choice rule, a familiar decision rule in marketing, says it is optimal for the consumer to pick the product with the largest posterior utility once search stops.

To complete the description of the optimal search rules, it remains to describe the boundary of search, $M_{jt}(\cdot)$. The specific functional form for the boundary of search comes from solving the dynamic programming problem in equation (6) using diffusion approximation. Chick and Frazier (2012) show that it is given by

$$M_{jt}(c_j, \sigma_j, n_{jt}) = c_j^{1/3} \sigma_j^{2/3} b(\sigma_j^{2/3} / (c_j^{2/3} n_{jt})) \tag{9}$$

where $b(h)$ can be approximated by[3]

$$\hat{b}(h) = \begin{cases} 0.233h^2, & \text{if } h \leq 1 \\ 0.00537h^4 - 0.06906h^3 + 0.3167h^2 - 0.02326h, & \text{if } 1 < h \leq 3 \\ 0.705h^{1/2} ln(h), & \text{if } 3 < h \leq 40 \\ (3h log(h) - ln(8\pi) - 2/ln(h) - 170)^{1/2}, & \text{if } h > 40 \end{cases}$$

This approximation for $b(\cdot)$ is similar to results in the related multiarmed bandit problem (Gittins, 1989; Brezzi and Lai, 2002).

The boundary of search is non-negative since $\hat{b}(\cdot)$ is non-negative. It is straightforward to derive the relation between $M_{jt}(\cdot)$ and its arguments. We illustrate these relations in Figure 3 below. As can be seen in this figure or derived more formally from equation (9), the boundary of search is decreasing in search costs c_j and the number of samples n_{jt} drawn, and increasing in the signal variance, σ_j^2. As a result, $M_{jt}(\cdot)$ has intuitive properties: higher search costs and lower prior uncertainty both lead to a lower boundary of search and thus a lower likelihood of search.

[3]Note that the expression for $\hat{b}(h)$ when $h > 40$ has been updated relative to Chick and Frazier (2012) using results in Chernoff (1965). This update has been mentioned in Chick et al. (2018) and the code is available at `https://github.com/sechick/pdestop`.

Figure 3: The boundary of search $M_{jt}(\cdot)$

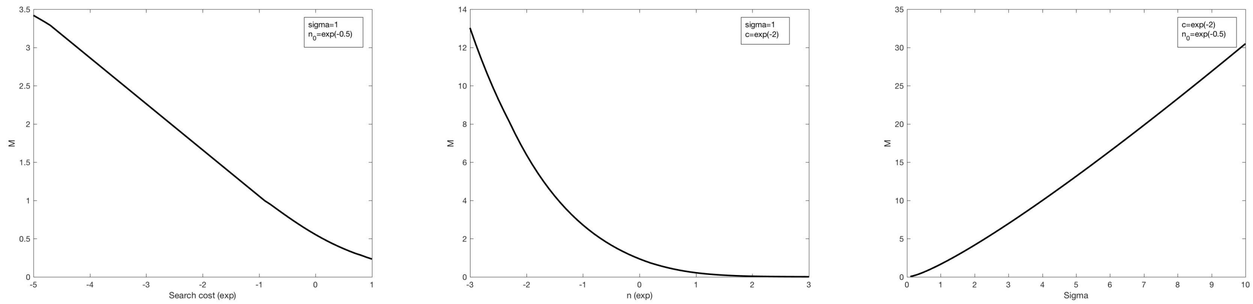

3.3 Relation to other approaches

In this section, we compare the model presented in section 3 to the sequential search model of Weitzman (1979) which is commonly used in marketing and to the more general framework of multiarmed bandits. Although these models are related, there are important differences that distinguish them. In particular, we focus on two distinguishing features of these approaches. First, we describe what problem consumers are assumed to solve in each model, and thus what consumer behavior each model can capture in empirical settings. Second, we contrast the optimal search rules that arise from these different problem formulations and show which features are shared across approaches.

As mentioned before, the model presented in this paper assumes consumers want to purchase a product and are initially unsure of their match with the product. However, they can search before deciding whether to purchase, which is costly, but reveals informative signals about their match value. The goal of a consumer is to maximize her expected utility net of total search costs from the best option she will choose when search ceases.

In Weitzman's (1979) search problem, the consumer faces a set of options and can sequentially sample each. Searching an option reveals all uncertainty about it and the consumer focuses on deciding whether to continue searching any of the unsearched options or stop and choose one of the searched options. Thus, this model cannot be used to study the consumer's decision to search the same option multiple times (e.g. spend time searching, revisit a previously searched option). The optimal policy involves three search rules. The stopping rule dictates that the consumer will terminate search when the maximum utility observed u_j exceeds the reservation utility z_j of any unsearched option, where the reservation utility is defined by the solution to $c_j = \int_{z_j}^{\infty}(u_j - z_j)f(u_j)du_j$, with c_j giving the search cost and $f(\cdot)$ the utility distribution. This equation shows that the reservation utility is a function of both expected utility and search costs. More precisely, a higher expected utility or lower search costs lead to a higher reservation utility, that is a product that is more likely to be searched. The selection rule says that if a search is to be made (if the stopping rule is not satisfied), the option with the highest reservation utility z_j should be searched next. Finally, the choice rule says that once the consumer stops searching, she will choose to purchase the option with the highest utility u_j among those searched (including the outside option).

The stopping and selection rules in Weitzman's (1979) sequential search model are fundamentally in contrast to those in our paper. These rules use an index (the reservation utility) to describe the consumer's decision to continue searching. In the Weitzman (1979) model, the consumer always searches the option with largest reservation utility, which is a function of the expected utility of the product. In our model, search continues if at least one option's posterior utility lies within the continuation set. This implies that both options with very large and those with very small posterior utility might not be candidates for search in a given period. As a result, the consumer might prefer to search an option with a lower posterior utility than the largest one, if the former is within the continuation set, while the latter is not. This is so since consumers are uncertain about the true match value of a product and are only willing to pay to search if the expected increase in utility together with the reduction in uncertainty is sufficient. This difference arises from the fact that in the Weitzman (1979) model consumers reveal all uncertainty after one search, while our framework allows for learning.

The assumed consumer search behavior is another distinguishing feature of our approach. More precisely, in our model the consumer begins search with a prior on her match value and an unobserved (from the researcher's perspective) term that influence her utility. By searching, the consumer observes a signal which she uses to update her beliefs. Thus, in our framework, the consumer searches to reveal uncertainty about the match value, but knows the part of utility that is unobserved by the researcher. This is in contrast to most empirical instantiations of the Weitzman (1979) model (see for example, Kim et al 2010; Honka and Chintagunta, 2016; Chen and Yao, 2016). In these models, consumers are certain about the expected utility from each product, and search in order to reveal an unobserved (from the researcher perspective and from the consumer's perspective before search) set of search attributes, which are revealed in one step.

The fundamental differences between the two approaches will make it hard to compare the estimation results from the two models. One way to see this is to consider what modifications of our approach would be needed to make it more comparable to the Weitzman (1979) approach. Since in the Weitzman (1979) model all uncertainty about an option is revealed after paying one search cost, we would not be able to model learning or search duration. This would mean setting $\mu_{ij0} = 0$, $\mu_{ij} = 0$, and $\sigma_j = 0$. We could also set $n_{ij0} = 1$, so that it does not affect the boundary of search. In this setup, the consumer searches, reveals all uncertainty about an option in one search action, and decides whether to continue searching another alternative or whether to stop and purchase one of the searched alternatives, if any. Note that in this case, the stopping and the selection rules in equations 7 and 8 would need to be modified, so that only unsearched options are candidates for search at every point in time. To recover the effect of product characteristics on consumer choices, one can include these observables in the consumer's expected utility before search, as is done typically in the literature (see for example, Kim et al 2010; Honka and Chintagunta, 2016; Chen and Yao, 2016). However, this approach does not allow estimating consumer preferences for product characteristics observed through search, as we do using in our approach (see section 5). Thus, even with these modifications, it is still difficult to compare our approach to empirical models based on the Weitzman (1979) approach.

In contrast to Weitzman's (1979) sequential search problem which describes the case where consumers cannot search the same option multiple times, more general multiarmed bandit problems deal with the case where the same option is sampled multiple times. In such problems, the consumer has the option to observe rewards from a number of alternatives, described by different reward distributions. By sampling an option, she learns about the distribution of that option and has the option to continue sampling from all options. There is an implicit tradeoff the consumer is facing between exploiting her knowledge from the sampled options or exploring potentially less appealing options currently in order to learn about their reward distribution and make better choices in the future. The goal of the consumer is to maximize the (discounted) sum of rewards. The optimal policy is characterized in Gittins and Jones (1974) and Gittins (1979) in terms of an index rule that dictates the consumer should choose in each time period the option with the largest index. The Gittins index resembles the role of the reservation utility in the Weitzman's (1979) sequential search model. This model is well suited to study repeated purchase decisions (as is done in Lin, Zhang and Hauser, 2014). However, since in this model the consumer accumulates rewards after each period, it is not well suited for the problem in this paper where the consumer's goal is to maximize the single utility net of search costs from the option chosen after search ceases. In addition, adding a cost to sample, leads in most cases to the breakdown of the optimality of the index (see Bank and Sundaram, 1994), making the use of the multiarmed bandit framework in our case less desirable. A natural extension of the multiarmed bandit model to account for multiple searches before a single purchase and for search costs is the model by Chick and Frazier (2012), which is the framework we adopt here.

3.4 Relation to the consumer's decision making journey

Our model describes richer patterns of consumer search behavior than other search models in the literature. More precisely, the model captures the consumer's decision of not only which products to search and which (if any) to purchase, but also the number of times to search each individual product, that is the consumer's search duration. Observing the duration of search can provide information on the consumer's awareness of products prior to search, a feature which is not captured by previous models.

To see this, consider the following example. Suppose there are two consumers, A and B, with similar preferences, and suppose both search the same number of products, but A spends twice as much time searching each product than B. A search model that ignores search duration, such as the Weitzman (1979) search model, would infer that the two consumers must have similar search costs as well. However, observing duration introduces the possibility that the two consumers differ in terms of initial awareness of the products, even though they might have similar preferences and search costs. That is, in the example, consumer A may have higher initial uncertainty of the products than consumer B, making it rational for her to spend more time searching products to resolve her uncertainty. In other words, modeling duration allows us to analyze (indirectly) an

Figure 4: Prior uncertainty and consumer choices

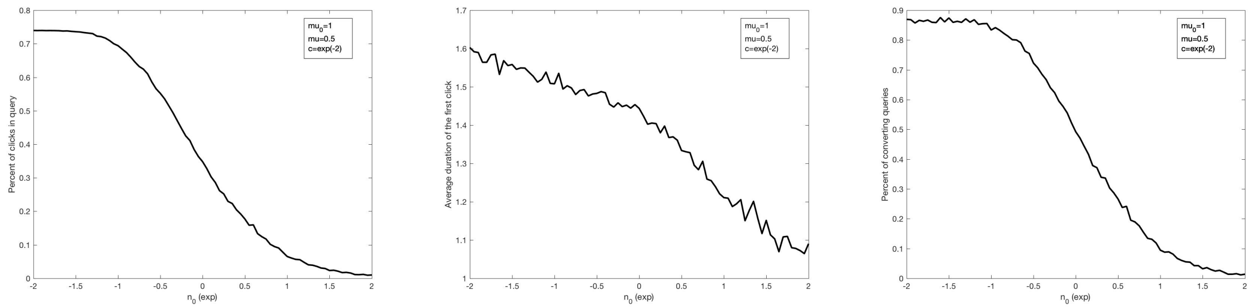

Searches | Duration of first search | Transactions

additional stage of the consumer's decision making journey (beyond the information search and the purchase stages typically captured by all search models), that is the awareness stage.[4] Taking this model to data (as we do in section 6), will allow us to estimate not only consumer preferences and search costs, but also their level of prior uncertainty for products, a feature that is not shared with other search models.

Using our model, we can study the relation between prior uncertainty, that is σ_j^2/n_{j0}, and consumer choices. To illustrate this relation, we use a simulation exercise where we generate a data set of 2,000 consumers, each performing two queries in which they observe three to five products they can search. We then use our model to simulate their choices. More precisely, we use the stopping, selection and choice rules presented above to determine: which products consumers will search, for how long, and whether or not they will purchase.

In Figure 4 we illustrate the effect of increasing n_{j0} or decreasing prior uncertainty (assuming all products are affected equally) on three choices of interest: duration of first search, percent of searches in a query and percent of converting queries.[5] From equation (9), we know that increasing n_{j0} decreases the boundary of search M_{j0} before the first search, making it less likely that consumers will search any products, as can be seen in the left panel in Figure 4. At the same time, even conditional on a search, a higher n_{j0} means that the boundary of search M_{jt} for observing an additional signal from the same product (spend time searching it) would be lower, so thus we expect search duration to also decrease with n_{j0}. To illustrate this most clearly we focus on the time spent on the first product searched in a query as a function of n_{j0} and this negative relation is shown in the middle panel of Figure 4. Together, the fact that consumers search fewer options and spend less time searching means that consumers are less likely to purchase.[6] These relations show new types of data patterns that we can accommodate in our model.

In sum, in order to model the time spent searching an option (or to revisit a previously searched

[4]We thank JP Dube for this comment.

[5]In the figure, we vary the value of r from -3 to 3, where $n_{j0} = exp(r)$.

[6]Note that in the model, consumers can purchase an option based solely on their prior knowledge, that is without searching it. However, to construct Figure 4 we assume consumers can only purchase an option if it was searched, as is typically the case in empirical applications in the literature, including our own.

option) a model that allows consumers to search the same option multiple times with each search being costly is needed. The model presented in this paper provides exactly such a framework. In the next section, we discuss the data we will use to estimate this model.

4 Data

4.1 Search Process

The specific empirical context of our analysis is consumer search for restaurants on an Asian review website. At the time of our data collection, this website provided review information for many products and services, but mainly focused on restaurants (similar to Yelp). We start by describing consumers' three step search process on this website to introduce the observables in our data set. Figure 5 provides an illustration of these steps.

Consumers start their search by visiting the homepage of the website. Here they can find a restaurant either by typing in a keyword in the search bar at the top, or by searching by cuisine type, location, dish tags, or another menu option (step 1). We refer to these actions as "specifying a query". In response to a query, consumers see an ordered list of restaurants,[7] typically divided into pages with 10 results per page. The consumer can make multiple queries. If these queries are less than one hour apart, we will interpret them as belonging to the same session, consistent with previous work (Wu et al. 2015). The list page contains some information about the displayed restaurants, such as the name, location, average rating information on three dimensions (taste, ambience and service) and average price. The consumer can then search a restaurant to obtain additional information by clicking on it (step 2). In this case, they navigate to a second screen reserved for that restaurant, which we refer to as the restaurant page. Here they can see photos, restaurant and dish tags, a brief description of the restaurant, as well as previous consumer reviews ordered by posting date. Given the amount of information on this page, consumers decide how much time to spend on a restaurant page, whether to return to the list and make another search, or whether to purchase (step 3). We interpret clicking on the list page as search on the extensive margin, while spending time on the restaurant page as search on the intensive margin.

4.2 Data Sources

There are two main data sources we use in this paper. One data source is obtained from the Asian review website. This has three components. First is a click stream data set containing searches consumers made on the site from December 2007 to March 2008. Importantly, these data contain information on the date and time of the click, which allows us to compute the duration of a click, using differences in time stamps. One concern with using time stamps to measure duration is measurement error. More precisely, we observe when consumers clicked on the restaurant page and

[7] Restaurants are ordered by default according to the proprietary ranking algorithm used by the website. However, consumers can further sort or filter search results. Modeling such decisions is beyond the scope of this paper. The interested reader should refer to Chen and Yao (2016).

Figure 5: Search process illustration

Step 1: Visit homepage; Specify query

Step 2: See list page information; Click

Step 3: See restaurant page information; Duration; Purchase

when they clicked to go back to the list page or another page. However, we do not observe exactly what they did in this time interval, that is whether they spent time reading about the restaurant or whether they were engaged in another activity. Although we cannot fully alleviate this concern, we do two things to partially address it. First, we collapse duration above 10 minutes since this is more likely to include activities not related to restaurant viewing. Second, we use the duration variable measured by comScore to cross check the time spent on a click on a similar website (Yelp). As we show in the section below, we find very similar duration measures in the comScore data as in our own.

In general, having time stamp information would allow us to obtain duration information for all clicks but the last click made by the consumer (duration would be truncated). However, the data include not only clicks made on restaurants, but also clicks to the homepage of the website, clicks on the consumer's profile on the site, other member's profiles clicks, clicks to chat pages, etc. Thus, we are able to directly observe duration information for 79% of clicks and 40% of last clicks.

The second data component describes restaurant page characteristics of the clicked restaurants for the period April 2003 to March 2008.

Third, we have individual level transactions for the period May 2005 to March 2008 of consumers who have a loyalty card distributed by the website from restaurants which collaborate with it. By using the loyalty card at the restaurant, consumers obtain a 10-30% discount at collaborating restaurants. Note that consumers' use of this loyalty card allows us to link online queries to offline transactions for (possibly) only a subset of purchases, and thus our transaction data is truncated. However, given the significant discount provided by the loyalty program, we anticipate this truncation to only have a minor impact on our data collection efforts. To further minimize the impact of truncation, we will focus our analysis on consumers who make a purchase. Although we limit the analysis to converting consumers, we observe both converting and non-converting sessions, where we call a non-converting session one in which more than 75% of the clicked restaurants participate in the loyalty program.

Since we are interested in modeling consumer search, we need to observe not only which restaurants consumers clicked, but also those they did not search, information which is not included in the first data source. Thus, to augment the data on the restaurants clicked, we use a second data source, which comes from an Internet archiving website called "Wayback Machine" (WBM).[8] Using the keywords that consumers searched and the time of search, we retrieve from the WBM the list of restaurants that consumers likely saw as a response to their query. We require that the keywords consumers searched should be exactly matched with the ones save on WBM. However, because the time of search usually cannot be exactly matched, we retrieve the closest time that the keyword search was saved. Given that data on WBM becomes more sparse going further back in time, we are able to match 68% of queries, which we will use in the analysis.

[8] The Wayback Machine website can be found at `https://archive.org/web/`.

4.3 Final Data Sample

In our final data sample, there are 343,270 observations, where an observation is a restaurant displayed to consumers in response to their query. We observe 5,465 consumers searching across a total of 17,852 sessions and 34,912 queries and making 50,439 clicks and 7,538 transactions (21.59% of queries end in a transaction). There are 10,632 restaurants in the data, and in Table 2, we summarize the average restaurant characteristics we observe.

Table 2: Restaurant characteristics

	All		Clicked		Duration>median		Purchased	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
List page information								
Rating	2.71	0.32	2.72	0.31	2.74	0.29	2.78	0.25
Price (100 RMB)	0.91	0.92	0.88	0.76	0.87	0.68	0.81	0.37
Promotion	0.20	0.40	0.24	0.43	0.27	0.45	0.30	0.46
Number of reviews (1000)	0.47	0.62	0.49	0.61	0.53	0.60	0.61	0.55
Position	5.59	3.03	5.10	3.04	5.15	3.07	4.98	3.08
Restaurant page information								
Average review length (1000 words)			0.37	0.18	0.38	0.19	0.41	0.19
Number of photos			85.01	100.37	94.23	101.37	104.18	97.69
Length of introduction (words)			98.54	49.44	103.26	45.33	112.54	32.25
Observations	343270		50439		25226		7538	

The list page includes information on the rating of the restaurant (weighted sum of taste, ambience and service measures), the average price, a promotion indicator, the number of reviews of the restaurant, as well as its position in the list.[9] As can be seen, clicked, purchased or restaurants on which consumers spent more time generally have a higher rating and lower prices. If the consumer clicked on a restaurant on the list page, then we observe additional information as contained on the restaurant page, such as the number of photos, the length of the introduction posted by the owner of the restaurant, as well as the individual consumer reviews posted, which we summarize using the average review length. In general, the more information is displayed on the restaurant page (e.g. in terms of the number of photos, length of the introduction or review length), the longer consumers spend on the restaurant page.

In addition, we have information on several query observables, which we summarize in Table 3. For instance, we observe that on average queries are made by consumers who registered with the website approximately two years in advance and that on average a transaction happens less than one week after the query. Note that we can also compute the time that consumers spent on the list page before clicking any restaurant, since we have both the time stamp of the query and of the first click. Consumers spend less than one minute on the list page before making the first restaurant click, and, as we show next, this is small relative to the time consumers spend on a click on the restaurant page, which is the focus of this paper.

[9]We follow the policy of the website and weight taste by 0.6, ambience by 0.25, and service by 0.15 to construct the rating of a restaurant.

Table 3: Query characteristics

	Mean	SD
Days since registered on website	681.62	418.24
Days between session and transaction	6.21	12.79
Time before first click (minutes)	0.72	2.73
Weekend	0.21	0.41
Office hour	0.66	0.47
Observations	34912	

4.4 Data Patterns on Search Duration

In this section, we show how consumers search on the intensive margin using data on the time they spend searching restaurants. More precisely, we provide evidence on how much time consumers search, on what affects search duration and on the effect of duration on purchase decisions.

4.4.1 How much time do consumers spend searching?

We find that consumer search on the intensive margin is considerable: the average (median) consumer spends 3.47 (2.45) minutes on a click, with a standard deviation of 3.07 minutes.[10] [11] In Figure 6, we display the distribution of search duration for clicked restaurants, showing a large variation in search duration, with many clicks lasting less than one minute and a large right tail. To provide external validation for this result, we use comScore to check the time spent on a similar website, and we find that click duration on Yelp is 3.55 minutes (January, 2013), which is similar to duration in our data. Also, this finding is generally in line with estimates from the literature. More precisely, Fradkin (2017) reports that on Airbnb consumers spend 58 minutes before sending contact information to sellers and browse on average 31 listings, implying an average time spent on any listing of 1.87 minutes. In contrast, although consumers spend a relatively long time searching each restaurant, their search on the extensive margin (that is, the number of restaurants clicked) is small. More precisely, we find that 60% (71%) of sessions (queries) have only one click, with an average (median) click number by session of 2.83 (1). Correspondingly, at the query level, the average (median) click number is 1.44 (1). Thus, even when consumers search very few restaurants, they search each option intensively.

4.4.2 What influences search duration?

We showed that search duration is considerable. Next, we ask what influences the consumer's search duration decision, by considering three related questions. First, is duration considerable because certain consumers spend a long time searching or because certain restaurants are searched longer by all consumers? To answer this question, we divide consumers into types by experience,

[10] Considering only the observations for which we observe full duration information (79% of clicks), the average (median) consumer spends 2.88 (1.45) minutes on a click, with a standard deviation of 3.21 minutes.

[11] We express duration here as follows: 1.50 minutes means one minute and 30 seconds.

Figure 6: Extent of search on the intensive margin: Search duration

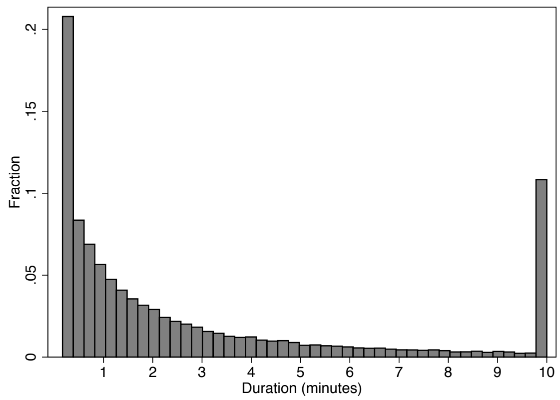

Notes: Histogram of duration (minutes) for observations with full duration information (no imputed values). The spike at the right tail is due to truncation and collapse of duration larger than 10 minutes (the 90th percentile).

that is the days since they registered with the website. Consumers who joined the website further in advance should be more experienced with it and thus might differ in terms of how much time they require to process restaurant information. A high (low) type consumer is one with experience higher (lower) than the 75th (25th) percentile of that distribution. We also divide restaurants into types such that a high (low) quality restaurant is one with rating and price above (below) the 75th (25th) percentile of the respective distributions.

In Table 4, we compute the average search duration by consumer and restaurant types. A $*$ identifies significant differences (at the 5% level) between two entries by means of a t-test. We find that low quality restaurants are searched less than high quality ones by all consumers. This is to be expected since low quality restaurants present relatively less interest to most consumers. In addition, experienced consumers spend less time on both types of restaurants than those less experienced. In sum, we find that both consumer and restaurant characteristics affect search duration.[12]

Table 4: Duration by consumer and restaurant types

		Restaurant type		
		High quality		**Low quality**
Consumer type	**Experienced**	3.55	*	2.93
				*
	Inexperienced	3.72	*	3.19

Second, which restaurant and consumer characteristics affect search duration? To answer this question, we have considered using survival analysis to model the consumer's decision to continue searching the same option after a given time period (for example, one minute), rather than switching to searching another restaurant or terminating search. This model would be preferable to a linear

[12]Our results hold also when considering the median duration by type. The analysis is available upon request.

regression if censoring of the data were of significant concern. However, in our data, we observe duration larger than 10 minutes in only 10% of the clicks, leading us to model the relation between duration and observables using a linear regression.

Table 5: Effect of observables on duration (OLS)

	Estimates	Std. Err.	Estimates	Std. Err.
	(1)		(2)	
List page information				
Rating	0.1668**	(0.0533)	0.1309*	(0.0536)
Price	−0.0820***	(0.0195)		
Price fixed effects				
(20-50RMB)			0.3672***	(0.0876)
(51-79RMB)			0.4351***	(0.0851)
(80-119RMB)			0.4974***	(0.0866)
(120-200RMB)			0.1534	(0.0929)
(>200RMB)			0.2272*	(0.1080)
Promotion	0.3732***	(0.0346)	0.3425***	(0.0350)
Number of reviews	0.2794***	(0.0523)	0.2365***	(0.0533)
Position	0.0363***	(0.0045)	0.0360***	(0.0045)
Restaurant page information				
Average review length	1.5938***	(0.1248)	1.4646***	(0.1271)
Number of photos	0.1010***	(0.0111)	0.0915***	(0.0114)
Length of introduction	0.0930***	(0.0084)	0.0964***	(0.0085)
Query information				
Days since registered on website	−0.0001***	(0.0000)	−0.0001***	(0.0000)
Weekend	−0.2036***	(0.0328)	−0.2001***	(0.0328)
Office hour	0.3616***	(0.0286)	0.3576***	(0.0286)
Constant	1.4013***	(0.1309)	1.1134***	(0.1553)
R-squared	0.0317		0.0327	
Observations	50439		50439	

Standard errors in parentheses
Notes: OLS regression with dependent variable search duration. In column (1), price is measured in 100RMB. In both columns, number of reviews, average review length, number of photos, and length of introduction enter the regression in logarithmic form. In column (2), price effects are relative to the left out category: prices <20RMB. All estimates are conditional on a click.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

In Table 5 we consider the effect of a large set of factors on the consumers' decision to spend time on a restaurant using a linear regression. In particular, we divide restaurant characteristics by whether they are displayed on the list or the restaurant pages and we also consider the effect of query characteristics on search duration.[13] Focusing on the list page information, we find that cheaper restaurants with a higher rating, that are promoted and that are displayed lower on the list page lead to a higher search duration. Also, the more information is displayed on the restaurant page, for example in terms of a larger number of photos, longer description of the restaurant or

[13] For 9.58% of searched restaurants, we do not observe review information. Therefore, to be able to measure the effect of restaurant characteristics on consumer choices, we add one and take the log of the affected variables (number of reviews, number of photos, average review length and the length of the introduction).

Figure 7: Click order and duration

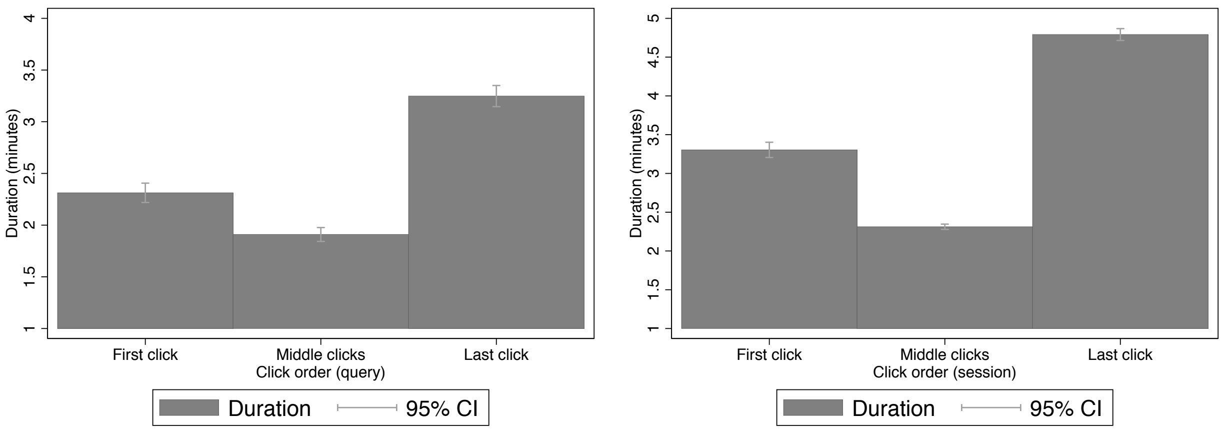

Query level

Session level

longer reviews, the more time consumers spend reading about the restaurant. In terms of query characteristics, we find that consumers are more likely to spend a longer time searching on weekdays between 9am and 5pm. Finally, less experienced consumers (measured by the number of days since website registration) also spend a longer time on a restaurant page, consistent with our results in Table 4.

In sum, we find that search duration is increased by restaurant rating, consumer inexperience with the website, and by the amount of information displayed on the restaurant page.

Third, we investigate the relation between search duration and clicks. In particular, we ask whether click order influences time spent searching. To this end, we restrict our attention to sessions (or queries) with at least three clicks, and compute the average click duration for the first, last, and middle clicks. Figure 7 shows our results. We find that consumers spend more time on the first and last clicked restaurant in a session (or query) than on middle clicks.[14] Consistent with a sequential model of search, click order captures consumers' expected utility from the considered alternatives (net of search costs). Our finding suggests that search duration captures consumers' revealed preference for restaurant characteristics, which is not captured completely by click order.

4.4.3 Relation between search duration and purchases

Finally, we consider the relation between the time that consumers spend on a click and the probability of purchasing the searched restaurant. A t-test reveals that clicked restaurants that were purchased have a higher search duration (1.52 minutes difference, $t = -40.31$) than those that were not purchased. To further decompose this effect, we condition on clicks, and model the purchase decision in a session as being influenced by duration, restaurant, and query characteristics. We include an outside option in the model, meant to represent the option of not purchasing, which has

[14]The same analysis by total click number shows a similar pattern and is available upon request.

an expected utility normalized to zero.[15]

Table 6: The effect of search duration on transactions (Conditional logit at session level)

	Estimates (1)	Std. Err.	Estimates (2)	Std. Err.	Estimates (3)	Std. Err.
Duration (minutes)			0.0698^{***}	(0.0045)	0.0690^{***}	(0.0045)
List page information						
Rating	0.4214^{***}	(0.0630)	0.4088^{***}	(0.0633)	0.3023^{***}	(0.0639)
Price	-0.4443^{***}	(0.0354)	-0.4453^{***}	(0.0356)		
Price fixed effects						
(20-50RMB)					4.7787^{***}	(1.0021)
(51-79RMB)					4.5221^{***}	(1.0019)
(80-119RMB)					4.7473^{***}	(1.0020)
(120-200RMB)					3.9765^{***}	(1.0033)
(>200RMB)					3.5757^{***}	(1.0059)
Promotion	0.0462	(0.0329)	0.0374	(0.0330)	0.0285	(0.0335)
Number of reviews	0.3505^{***}	(0.0478)	0.3559^{***}	(0.0481)	0.2339^{***}	(0.0496)
Position	-0.0045	(0.0043)	-0.0069	(0.0043)	-0.0077	(0.0044)
Restaurant page information						
Average review length	1.1898^{***}	(0.1135)	1.1454^{***}	(0.1142)	0.8545^{***}	(0.1173)
Number of photos	0.0223	(0.0116)	0.0166	(0.0116)	0.0075	(0.0118)
Length of introduction	0.2395^{***}	(0.0135)	0.2376^{***}	(0.0135)	0.2584^{***}	(0.0137)
Query information						
Days since registered on website	0.0001^{*}	(0.0000)	0.0001^{*}	(0.0000)	0.0001^{**}	(0.0000)
Weekend	0.2209^{***}	(0.0413)	0.2275^{***}	(0.0411)	0.2229^{***}	(0.0411)
Office hour	0.2724^{***}	(0.0361)	0.2547^{***}	(0.0359)	0.2551^{***}	(0.0360)
Outside option	3.6347^{***}	(0.1648)	3.8652^{***}	(0.1661)	8.4218^{***}	(1.0158)
Log-likelihood	-17,068		-16,950		-16,796	
AIC	34,159		33,926		33,626	
BIC	34,269		34,044		33,782	
Observations	68,291		68,291		68,291	

Standard errors in parentheses
Notes: Conditional logit model with dependent variable transactions within a session. In columns (1) and (2), price is measured in 100RMB. In all columns, number of reviews, average review length, number of photos, and length of introduction enter the regression in logarithmic form. In column (3), price effects are relative to the left out category: prices <20RMB. All estimates are conditional on a click.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Our results can be found in Table 6. In column (1), we omit the effect of duration and show that restaurants with a higher rating that are cheaper and that have more reviews are more likely to be purchased. Similarly, restaurants with more information displayed on the restaurant page are also more likely to be purchased. In column (2), we include duration as an explanatory variable in the model and find that even after accounting for restaurant and query characteristics, observing consumers' search duration decision helps predict purchases, as restaurants with higher search duration are more likely to be purchased. The last column shows the same effect when we decompose price into price level fixed effects. It shows that consumers prefer to avoid extreme prices, making the magnitude of the linear price effect shown in the first column hard to interpret.

[15]Including an outside option in the model increases the number of observations by the number of sessions.

In sum, in this section we have shown that consumers spend a considerable amount of time searching, even when they search few products. Search duration increases with consumer inexperience, higher rating of the restaurants and the amount of information displayed on the restaurant pages. Finally, we have shown that search duration is higher for purchased restaurants. These results demonstrate that search duration represents both a benefit to consumers (in terms of the amount of information gained through time spent) and a cost for accumulating this information, as consumers search few restaurants. In other words, our results imply that search duration is a choice, made separately from the two other choices consumers have: which restaurants to click and whether or not to purchase. This implies the need to incorporate the intensive margin decision into a consumer search model, which is the focus of this paper.

5 Estimation and identification

5.1 Empirical model

In section 3, we developed a sequential search model that accounts not only for the consumer's decision of which products to search and whether or not to purchase, but also how much time to spend gathering information about each product searched. We now consider an empirical application of this model that uses the data introduced in the previous section. To this end, we show how to adjust the model to the empirical setting considered and describe the estimation procedure we use to recover consumer preference, search costs and uncertainty parameters.

The empirical context of our paper is consumer search for restaurant information on an Asian review website. After visiting the homepage of the website, a consumer $i \in \{1, \ldots, N\}$ types in a query and observes a list of restaurants, with an individual restaurant on this list denoted by $j \in \{1, \ldots, J\}$. On the list page, the consumer observes some information about each restaurant, but can search by navigating to the restaurant page to obtain additional information about it. Each time period $t = 1, \ldots, T$, the consumer decides whether to continue searching, in which case she chooses a restaurant to search (including a previously searched option) or whether to stop, in which case she decides which restaurant to visit (purchase), if any. Deciding not to purchase is interpreted as choosing the outside option (denoted by $j = 0$).

The consumer's utility from visiting restaurant j in period t is equal to

$$u_{ijt} = X_j^{\text{list}} \beta_i + \mu_{ijt} + \epsilon_{ij}. \tag{10}$$

We model a consumer's utility from visiting restaurant j in period t as having three components. First, the consumer values the characteristics of each restaurant revealed on the list page, X_j^{list}. We assume consumers observe this information without paying a search cost, consistent with prior literature (Kim et al. 2010, 2016, Chen and Yao, 2016, De los Santos and Koulayev, 2016). Second, by searching j, the consumer obtains additional information hidden on the restaurant page. The longer she spends searching j, the more information she obtains. However, search does not reveal all uncertainty about j. Instead, μ_{ijt} represents the consumer's perceived match value at t with restaurant j. Third, also on the list page the consumer observes information that is hidden from the researcher modelled as an idiosyncratic shock, $\epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$. We normalize the expected utility of not purchasing, i.e. choosing the outside option, to zero so that $u_{i0} = \epsilon_{i0}$.

Consumer learning is modeled as in section 3. More precisely, the consumer is uncertain about her match value with a restaurant j, which we assume is drawn independently from $N(\mu_{ij}, \sigma_j^2)$, with unknown mean μ_{ij} and known variance σ_j^2. To resolve this uncertainty, the consumer spends time searching restaurants. To model search duration, we will interpret time spent searching an option as searching the same option multiple times (e.g. spending 5 minutes on a restaurant will be equivalent to searching it 5 times). The consumer begins her search with a prior belief about the match value, summarized by $N(\mu_{ij0}, \sigma_j^2/n_{ij0})$. Since n_{ij0} gives the number of samples assumed by the prior belief distribution, it is positive, so we model it using an exponential function.

By searching j in period t, the consumer observes a signal of the restaurant's match value, given by $s_{ijt} \sim N(\mu_{ij}, \sigma_j^2)$. We model signals as a function of restaurant page information, as follows

$$\mu_{ij} = \mu_i + X_j^{\text{rest}} \alpha_i. \tag{11}$$

Using Bayes Rule, the consumer updates her belief about a restaurant's match value given this additional information. More precisely, searching j at t implies a posterior belief $N(\mu_{ijt+1}, \sigma_j^2/n_{ijt+1})$, with $\mu_{ijt+1} = \frac{n_{ijt}\mu_{ijt} + s_{ijt}}{n_{ijt+1}}$ and $n_{ijt+1} = n_{ijt} + 1$, while for $k \neq j$, $\mu_{ikt+1} = \mu_{ikt}$ and $n_{ikt+1} = n_{ikt}$, that is, options that are not searched are not updated. The posterior mean is what affects the consumer's purchase decision.

To obtain restaurant page information, the consumer pays a cost of c_{ij} per search. To ensure that search costs are positive, we model c_{ij} as an exponential function,

$$c_{ij} = exp(k_i + \rho_j \gamma_i), \tag{12}$$

where k_i is the consumer specific mean search costs and ρ_j denotes factors that affect the consumer's search cost for j (e.g., the restaurant's position, Ursu 2018).

The empirical model differs in two respects from the one presented in section 3. First, motivated by the empirical application, consumers have access to list page information before search. This information is available without paying a search cost and influences the consumer's utility. Thus, consumers search to obtain information on μ_{ijt}, but are certain about the list page information observed. Second, restaurant page information influences the signals that the consumer observes through search. This assumption is motivated by the reduced form results in the previous section (such as those in Table 5), showing that search duration is affected by the information revealed on the restaurant page. For example, restaurants with longer reviews or more photos are searched longer. In our model, restaurants with more favorable restaurant page information imply more favorable signals, making the consumer more likely to search those longer.

5.2 Likelihood function

To estimate the model, we use the optimal search rules in section 3.2 to construct the likelihood of consumer search, search duration, and purchase decisions. The optimal search rules translate into the following restrictions on the parameters of interest. Suppose the consumer searched an option in each of the $t \leq T$ periods, with period T denoting the final period in which search occurs. Then, in period $T+1$ we observe the consumer making a purchase decision (visit a restaurant or choose the outside option of not purchasing). In this case, the stopping rule imposes two types of restrictions on parameters. First, in all periods $t \leq T$ when the consumer $i \in \{1, \dots, N\}$ searched a restaurant, for the searched option $j \in \{1, \dots, J\}$ it must be that

$$M_{ijt}(c_{ij}, \sigma_j, n_{ijt}) - \Delta_{ijt} > 0, \tag{13}$$

where $\Delta_{ijt} = |u_{ijt} - \max_{k \neq j} u_{ikt}|$ for $k \in (0, 1, \dots, J)$.

Second, in period $T+1$ when the consumer does not search, it must be that

$$M_{ijT+1}(c_{ij}, \sigma_j, n_{ijT+1}) - \Delta_{ijT+1} < 0, \;\; \forall j \in \{1, \dots, J\}. \tag{14}$$

The selection rule requires that, if the consumer searched $j \in \{1, \dots, J\}$ at $t \leq T$, then

$$\frac{M_{ijt}(c_{ij}, \sigma_j, n_{ijt}) - \Delta_{ijt}}{c_{ij}^{1/3} \sigma_j^{2/3}} > \max_{k \neq j} \frac{M_{ikt}(c_{ik}, \sigma_k, n_{ikt}) - \Delta_{ikt}}{c_{ik}^{1/3} \sigma_k^{2/3}}, \;\; \forall k \in \{1, \dots, J\}. \tag{15}$$

Finally, consistent with the choice rule, if the consumer chooses j (including the outside option) after terminating search, her utility from this choice must exceed the utilities of all other options. Formally,

$$u_{ijT+1} \geq \max_{k \neq j} u_{ikT+1}, \;\; \forall j, k \in \{0, 1, \dots, J\}. \tag{16}$$

If consumers search sequentially, they make search, search duration and purchase decisions jointly. Thus, the probability of observing a certain outcome in the data in period t for consumer i is characterized by the joint probability of the stopping, selection and choice rules holding in that period, as given by

$$L_{it} = Pr(\text{Stopping rule}_{it}, \text{Selection rule}_{it}, \text{Choice rule}_{it}). \tag{17}$$

Because consumers make these three decisions jointly, the likelihood function does not have a closed form solution. As a result, we use a simulated maximum likelihood (SMLE) approach to estimate the parameters of the model. In choosing the simulation method, we follow McFadden (1989), Honka (2014), Honka and Chintagunta (2017) and use the logit-smoothed AR simulator.

Simulation using the logit-smoothed AR simulator involves the following steps:

1. For each consumer, determine the last period of search T.

2. Make $d = \{1, \dots, D\}$ draws of ϵ_{ij} for each consumer-product-time period.

3. For each ϵ_{ij} draw, make $f = \{1, \dots, F\}$ draws of s_{ijt}.

4. Compute search costs c_{ij}, use draws in steps 2 and 3 and Bayesian updating formulas in equations 5 to form u_{ijt}, M_{ijt}, and Δ_{ijt}.

5. Use the stopping, selection and choice rules to construct the following expressions:

   (a) $\nu^1_{it} = M_{ijt} - \Delta_{ijt}$ for the searched option $j \in \{1, \ldots, J\}$ and $t \leq T$.

   (b) $\nu^2_{iT+1} = \Delta_{ijT+1} - M_{ijT+1}$ $, \forall j \in \{1, \ldots, J\}$.

   (c) $\nu^3_{it} = \frac{M_{ijt} - \Delta_{ijt}}{c_{ij}^{1/3} \sigma_j^{2/3}} - \max_{k \neq j} \frac{M_{ikt} - \Delta_{ikt}}{c_{ik}^{1/3} \sigma_k^{2/3}}$ for the searched option $j \in \{1, \ldots, J\}$, any other option $k \in \{1, \ldots, J\}$, and $t \leq T$.

   (d) $\nu^4_{iT+1} = u_{ijT+1} - \max_{k \neq j} u_{ikT+1}, \quad \forall j, k \in \{0, 1, \ldots, J\}$.

6. Compute expression

$$R_i = exp(-\nu^2_{iT+1}/\lambda) + exp(-\nu^4_{iT+1}/\lambda) + \sum_t [exp(-\nu^1_{it}/\lambda) + exp(-\nu^3_{it}/\lambda)] \qquad (18)$$

   where $\lambda > 0$ is a scaling parameter.[16]

7. Obtain $R = \sum_i R_i$ by summing over consumers and compute S

$$S = \frac{1}{1 + R}.$$

8. The average of S over the D and F draws of the error terms gives the simulated likelihood function.

We estimate two versions of the model. One that assumes consumers are homogenous, and another were we capture consumer heterogeneity using a latent class approach with two segments of consumers. To construct the likelihood function in the former case, we consider consumer choices at the query level. Thus, the set of options that is available for search is the set of restaurants displayed in the query. To construct the likelihood function in the latter case, we assume sessions can be grouped into segments, making all queries that correspond to a certain session belong to the same segment. Thus, we need to compute the likelihood function across all queries in a session, conditional on belonging to a segment, and take the weighted sum of the conditional likelihoods (with the segment sizes serving as weights).

5.3 Identification

In this section, we describe how utility, uncertainty, and search costs parameters are identified. Identification is similar for the two versions of the model (homogenous parameters or latent class model), so for exposition simplicity we suppress the consumer subscript below. The set of parameters to estimate is composed of the following: list and restaurant page utility parameters

[16] Little guidance is available on choosing the scaling parameter λ. We will determine the appropriate scaling parameter using Monte Carlo simulations, which are described in section 5.4.

(β, α), prior belief (μ_{j0}, n_{j0}), signal mean constant and variance (μ_j, σ_j), search costs (k, γ), and the variance of the unobserved idiosyncratic shock (σ_ϵ^2).

The specific empirical setting which we consider to estimate the model leads to certain normalizations that are needed to fully identify it. First, as is typical in Bayesian learning models, we can only recover the ratio of the prior and the signal variance, which implies that we will not be able to estimate σ_j^2, but can estimate n_{j0}. For this reason, we fix $\sigma_j^2 = 1$, for all restaurants. Second, we fix $\sigma_\epsilon^2 = 1$, as is common in the literature (see for example, Kim et al., 2010, 2016, Honka and Chintagunta, 2016, Chen and Yao, 2016). The variance of the unobserved idiosyncratic shock affects the benefits from search, but unfortunately our data do not allow us to separately identify it from search costs except through functional form (see Dong et al, 2018 for a more detailed description of this issue). Our data provide a short three month glimpse into consumers' behavior on the Asian review website. As a result, we are not able to observe differences in consumer prior restaurant information. Thus, we assume consumers start with uniform priors across all restaurants, that is $\mu_{j0} = \mu_0$ and $n_{j0} = n_0$, for all j. Given these considerations, the set of parameters we seek to identify for each segment becomes $(\mu_0, \mu, \beta, \alpha, n_0, k, \gamma)$.

The data provides information on three types of choices that consumers make. First, conditional on a query, we observe which restaurants consumers searched. Second, we observe search duration, that is the number of minutes a consumer spent on each searched restaurant page. Finally, we observe whether the consumer chooses to visit one of the searched restaurants or whether she chooses the outside option of not purchasing. These choices together with the optimal search rules we presented in section 3.2 constitute the necessary components of our identification argument.

Utility parameters $(\mu_0, \mu, \beta, \alpha)$ are identified from the choice rule conditional on search. More precisely, utility parameters for characteristics that vary by restaurant (β, α) are identified from the correlation in restaurant characteristics displayed on the list and the restaurant pages and the frequency with which restaurants are purchased, given information prior to a purchase. For example, the odds of a consumer visiting a restaurant with a higher rating or with more favorable restaurant page information reveals consumers' weights for these restaurant characteristics. The prior and the signal mean (μ_0, μ) are identified from the prevalence of restaurants purchased given different search duration. In other words, the evolution of consumer purchase decisions given information accumulated through search identifies these parameters.

Stopping and selection rules identify prior uncertainty (n_0) by observing both consumers who continue searching the same restaurant and those who switch, conditional on product characteristics revealed upon search. For example, observing a consumer search the same restaurant again reveals her high prior uncertainty (low n_0) relative to the scenario in which she decides to search another restaurant, ceteris paribus.

The mean search cost (k) is identified from the stopping and the selection rules. These rules impose an upper and a lower bound on mean search costs, respectively, that must have made it optimal for the consumer to stop after searching a certain number of restaurants. The level of search costs is pinned down by the functional form of the boundary of search and the distribution of the utility function error terms. The product specific search cost (γ) is identified from differences in search and purchase odds. For example, restaurants that are searched frequently (infrequently), but not purchased (purchased), have low (high) search cost.

The main challenge in identifying the parameters of the model lies in disentangling why a consumer stops searching. This might happen for three reasons: high search cost, high utility, or low prior uncertainty. There are two factors that allow us to separately identify these three effects. First, we can separately identify search costs from utility parameters since search costs do not affect purchase decisions. Thus, utility parameters are identified from the choice rule conditional on search, while search costs are then identified from the stopping and the selection rule. Second, search costs and prior uncertainty are separately identified since they affect the stopping and selection rules differently. More precisely, as the consumer searches, observing an additional match value signal affects prior uncertainty, but has no effect on search costs, that is the search cost remains constant, while prior uncertainty decreases.

5.4 Monte Carlo Simulation

In what follows, we show that Simulated Maximum Likelihood using the logit-smoothed AR simulator can recover utility, prior uncertainty, and search cost parameters in this model. We do so using Monte Carlo simulation. More precisely, we generate a data set of 2,000 consumers. Since in our empirical application consumers make several queries and queries are of various lengths (that is various numbers of restaurants are displayed on the list page in a query), we assume in the simulation that each consumer makes two queries and observes three to five options (one outside option and the rest restaurants). Restaurants have both list and restaurant page characteristics, which we assume are drawn from a normal distribution with mean and standard deviation equal to those found in the data. After search ceases, the consumer chooses whether to purchase from a restaurant on the list (with varying search duration) or whether to choose the outside option, which has expected utility normalized to zero. The true values of the parameters are chosen to be consistent with those from a preliminary estimation of the model. For estimation, we follow the steps described in Section 5 and use 100 draws from the distribution of the utility error terms and 30 draws from the signals distribution for each consumer-restaurant-time period combination. We repeat the estimation 50 times and report mean results.[17]

We perform the estimation on two different data sets. One where consumers are homogenous and another where there are two equally sized segments of heterogenous consumers. In the latter case, we use a latent class approach to estimate the model, where we aggregate queries by consumer type. Our results can be found in Tables 7 and 8, respectively.[18] In Table 7, the first column shows the true parameters and the second column shows the estimated parameters. Similarly, in Table 8, we show the true parameters and the estimated parameters by segment. Generally, we find that our

[17]We also perform robustness checks and estimate the model from 50 different starting points. Results are similar and are available upon request.

[18]The results are obtained with (inverse) scaling factor $1/\lambda$ equal to 10. However, upon request, we can share results for simulations with $1/\lambda$ ranging from 1 to 20. The results are similar.

Table 7: Monte Carlo Simulation Results: One Segment of Consumers

	True values	Estimates	Std. Err.
List page information			
Rating	1.00	0.8708***	(0.0028)
Price	−4.00	−3.2900***	(0.0015)
Restaurant page information			
Average review length	0.50	0.4147***	(0.0041)
Prior			
μ_0	1.50	1.1835***	(0.0041)
n_0 (exp)	−0.50	−0.5470***	(0.0054)
Signal			
μ	0.50	0.2816***	(0.0036)
Search cost (exp)			
Constant	−4.00	−2.7226***	(0.0306)
Log-likelihood		-10,363	
Observations		16,046	

Standard errors in parentheses

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 8: Monte Carlo Simulation Results: Two Segments of Consumers

	Segment 1			Segment 2		
	True values	Estimates	Std. Err.	True values	Estimates	Std. Err.
List page information						
Rating	1.00	0.9833***	(0.0030)	1.00	0.9069***	(0.0047)
Price	−4.00	−3.6037***	(0.0037)	−2.00	−1.7669***	(0.0025)
Restaurant page information						
Average review length	0.50	0.4064***	(0.0052)	−0.50	−0.3884***	(0.0047)
Prior						
μ_0	1.50	1.1790***	(0.0028)	1.00	1.0750***	(0.0075)
n_0 (exp)	−0.50	−0.5491***	(0.0096)	−1.00	−1.0727***	(0.0032)
Signal						
μ	0.50	0.3058***	(0.0052)	1.00	0.7749***	(0.0041)
Search cost (exp)						
Constant	−4.00	−2.9469***	(0.0582)	−1.50	−0.7164***	(0.0327)
Probability Segment 1						
π	0.00	0.1462*	(0.0655)			
Log-likelihood		-10,735				
Observations		16,046				

Standard errors in parentheses

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

method recovers the parameters of the model well. The exception is that we seem to underestimate the value of search costs under both methods. However, additional simulation results with more consumers and utility and signal draws improve this estimate.[19]

In sum, in this section we showed how the model we presented in section 3 can be applied to estimate consumer preference and search costs using the data we described in the previous section. Next, we provide estimation results of our model.

5.5 Estimation sample

Before discussing our results, we clarify a few data choices we make to help with the interpretation and estimation of the model. First, since restaurant characteristics vary in their scale, we choose to demean the following restaurant characteristics to make them comparable: rating, number of reviews, number of photos, average review length and the length of the introduction. Second, since for 9.58% of searched restaurants we do not observe review information, we add one and take the log of those variables related to review information (number of reviews, number of photos, average review length and the length of the introduction) to be able to measure their effect on consumer choices. The logarithmic transformation also helps emphasize the effect of marginal changes in these variables when the base value is relatively small. Third, we create an outside option to represent not purchasing. As in the model we presented in section 3, the outside option has an expected utility of zero. Note that a feature of the data is that consumers cannot choose to visit a restaurant that they did not previously search. Fourth, we focus on estimating price fixed effects for the six price levels determined by the company, instead of estimating a linear price effect. This approach is preferable given our reduced form results presented in Tables 5 and 6, where we show that consumers prefer to avoid extreme prices, thereby making the linear effect hard to interpret. Fifth, we discretize search duration and round it up to one minute increments. This assumption will only affect the units in which we express the mean search cost estimate. Finally, to make the estimation feasible, we restrict the data sample as follows. We only consider queries for which we observe search duration information for all searches made. In 18.14% of searches, consumers revisit a previously searched restaurant. However, we only consider for estimation those queries in which consumers did not revisit a previously searched restaurant. Although our model can capture revisits (as we demonstrated in section 3), modeling this introduces the question of forgetting, which is beyond the scope of our analysis. We estimate the model at the session level, where each session can have one or more queries. We restrict our attention to sessions with at most four queries, which represent 90 percent of sessions in the data. Finally, we randomly select a subsample of 1,000 sessions for estimation. This leads to a data set with 16,136 observations and 1486 queries, of which 21% lead to a transaction (similar to the conversion rate in the full data set). We find that 1955 restaurants are searched, with an average click duration of 2.77 minutes and a median of 2 minutes.

[19]This estimation is time intensive, but our results are available upon request.

6 Results

In this section, we present estimation results using the sequential search model and the estimation procedure presented in previous sections. For the estimation, we use 50 draws from the distribution of the utility error terms and 30 draws from the distribution of signals for each consumer-restaurant-time period combination. We repeat the estimation 50 times and present mean results across these 50 separate estimations.[20,21] We estimate two versions of the model. One where consumers are homogenous and another where there are two heterogenous consumer segments. In the latter case, we use a latent class approach to estimate the parameters of the model.

Table 9 presents our main estimation results for the case of homogenous consumers. A distinguishing feature of our model is that it allows us to estimate not only consumer preferences for characteristics observed before search (list page information), but also their preferences for information observed after search (restaurant page information). In the first column, we restrict the number of list page and restaurant page characteristics we include in the estimation, while in the second column, we include a more complete set of characteristics we observe.

Generally, we find utility and search cost estimates that are economically meaningful and significant. In particular, a higher rating and the restaurant's promotional efforts increase utility. To measure the impact of prices, we include price fixed effects in the model and find that consumers tend to avoid extreme prices, preferring the middle range of prices between $80 - 119RMB$, which correspond to \$12 − 18. In addition, the more information is revealed on the restaurant page, in terms of the length of the reviews and introduction, and the number of photos, the higher the utility. Since most sessions do not lead to a transaction regardless of the time spent searching, we find that mean prior beliefs (μ_0) and the signal constant (μ) are negative. However, the difference between these two parameters shows that consumers start their search with relatively poor expectations about product matches, but these improve through search. Consistent with our previous results in section 4.4, this shows that consumers are more likely to purchase a restaurant they spend more time searching.

In the data, on average consumers click few restaurants in a session, but spend a considerable amount of time searching each of them. The model rationalizes this with relatively high search costs and high initial product uncertainty ($1/n_0$), as can be seen in Table 9. Without information on search duration, we would not have been able to separate consumers' prior uncertainty from their search costs. Given such data, we determine the empirical relation between initial uncertainty and search costs: in the presence of high initial uncertainty and search costs, searching a restaurant requires a larger time investment, leading to only few restaurants that are searched.

Our results using a latent class approach with two segments can be found in Tables 10 and 11, where the latter includes a larger set of factors that might influence consumer choices. The data favor the model with two segments of consumers, as demonstrated by the larger differences in log-likelihood, AIC and BIC. The first segment is relatively smaller, corresponding to approximately

[20] The results are obtained with (inverse) scaling factor $1/\lambda$ equal to 10.

[21] We estimate the model from 50 different starting points. Results are similar and are available upon request.

Table 9: Estimation Results: One Segment of Consumers

	Estimates (1)	Std. Err.	Estimates (2)	Std. Err.
List page information				
Rating	0.0230^{***}	(0.0018)	0.0186^{**}	(0.0060)
Price fixed effects				
Price ($20-50$RMB)	0.2466^{***}	(0.0148)	0.3187^{***}	(0.0145)
Price ($51-79$RMB)	0.4381^{***}	(0.0116)	0.4799^{***}	(0.0111)
Price ($80-119$RMB)	0.5014^{***}	(0.0115)	0.5570^{***}	(0.0112)
Price ($120-200$RMB)	0.2312^{***}	(0.0265)	0.3137^{***}	(0.0234)
Price (>200RMB)	0.2975^{***}	(0.0348)	0.3595^{***}	(0.0284)
Promotion			0.1585^{***}	(0.0159)
Number of reviews			0.0108	(0.0071)
Restaurant page information				
Average review length	0.0277^{+}	(0.0158)	0.0171^{+}	(0.0100)
Number of photos			0.0295^{+}	(0.0123)
Length of introduction			0.0323^{***}	(0.0097)
Prior				
μ_0	-1.7069^{***}	(0.0028)	-1.8130^{***}	(0.0107)
n_0 (exp)	-1.1083^{***}	(0.0611)	-1.1476^{***}	(0.0479)
Signal				
μ	-1.4709^{***}	(0.0130)	-1.5782^{***}	(0.0094)
Search cost (exp)				
Constant	1.5957^{***}	(0.1022)	1.4756^{***}	(0.0903)
Position			0.0314	(0.0054)
Log-likelihood	-11,464		-11,436	
AIC	22,950		22,904	
BIC	23,035		23,027	
Observations	16,136		16,136	

Standard errors in parentheses

Notes: In both columns, the average review length enters the model in logarithmic form. Price effects are relative to the left out category: prices <20RMB.

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 10: Estimation Results: Two Segments of Consumers

	Segment 1		Segment 2	
	Estimates	Std. Err.	Estimates	Std. Err.
List page information				
Rating	0.0502^{***}	(0.0083)	0.0211^{***}	(0.0046)
Price fixed effects				
Price (20 – 50RMB)	0.5556^{***}	(0.0317)	0.4141^{***}	(0.0139)
Price (51 – 79RMB)	0.8673^{***}	(0.0068)	0.5822^{***}	(0.0127)
Price (80 – 119RMB)	0.8659^{***}	(0.0089)	0.6882^{***}	(0.0114)
Price (120 – 200RMB)	0.4405^{***}	(0.0387)	0.4378^{***}	(0.0242)
Price (> 200RMB)	0.5891^{***}	(0.0428)	0.4456^{***}	(0.0423)
Restaurant page information				
Average review length	0.0278	(0.0456)	0.0441	(0.1850)
Prior				
μ_0	-2.3071^{***}	(0.0073)	-1.7107^{***}	(0.0080)
n_0 (exp)	1.1473^{***}	(0.0684)	0.3286^{*}	(0.1308)
Signal				
μ	-2.1479^{***}	(0.0477)	-1.2997^{***}	(0.1632)
Search cost (exp)				
Constant	-2.6200^{***}	(0.0835)	2.0034^{***}	(0.1564)
Probability segment 1				
π	-0.1376	(0.0881)		
Log-likelihood	-10,959			
AIC	21,963			
BIC	22,140			
Observations	16,136			

Standard errors in parentheses
Notes: In both columns, the average review length enters the model in logarithmic form. Price effects are relative to the left out category: prices <20RMB.
$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 11: Estimation Results: Two Segments of Consumers and Additional Characteristics

	Segment 1		Segment 2	
	Estimates	Std. Err.	Estimates	Std. Err.
Rating	0.0433^{***}	(0.0111)	0.0318^{+}	(0.0193)
Price fixed effects				
Price $(20 - 50\text{RMB})$	0.2462^{***}	(0.0231)	0.3484^{***}	(0.0461)
Price $(51 - 79\text{RMB})$	0.4052^{***}	(0.0199)	0.5076^{***}	(0.0153)
Price $(80 - 119\text{RMB})$	0.5182^{***}	(0.0195)	0.6158^{***}	(0.0319)
Price $(120 - 200\text{RMB})$	0.4336^{***}	(0.0339)	0.2171^{***}	(0.0586)
Price $(> 200\text{RMB})$	0.3100^{***}	(0.0556)	0.3063^{***}	(0.1238)
Promotion	0.1730^{***}	(0.0274)	0.1185^{***}	(0.0311)
Number of reviews	0.0168	(0.0204)	0.0119	(0.0265)
Restaurant page information				
Average review length	0.0050	(0.0665)	0.0488	(0.3551)
Number of photos	0.0780	(0.0668)	0.1037	(0.3637)
Length of introduction	0.1351^{*}	(0.0669)	0.1480	(0.4842)
Prior				
μ_0	-2.0531^{***}	(0.0119)	-1.7033^{***}	(0.0189)
n_0 (exp)	1.2695^{***}	(0.0087)	0.5385^{**}	(0.1951)
Signal				
μ	-1.8472^{***}	(0.0548)	-1.1555^{***}	(0.3048)
Search cost (exp)				
Constant	-3.2865^{***}	(0.0043)	2.3410^{***}	(0.1769))
Position	0.0003	(0.0059)	0.0242	(0.0257)
Probability segment 1				
π	-0.1792^{*}	(0.0835)		
Log-likelihood	-10,831			
AIC	21,728			
BIC	21,982			
Observations	16,136			

Standard errors in parentheses

Notes: In both columns, number of reviews, average review length, number of photos, and length of introduction enter the model in logarithmic form. Price effects are relative to the left out category: prices <20RMB.

$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

46% of consumers.[22] Preferences for list and restaurant page characteristics differ slightly between the two segments, and as in the case of one segment, we find that consumers begin their search with less favorable prior expectations, but these improve through search.

Most interesting for our study, by allowing parameters to vary with consumer segments, we uncover an additional layer of the relation between consumers' prior uncertainty and their search costs. More precisely, our results show that the two segments differ starkly in their search primitives and thus their search strategies. The first segment begins search with lower uncertainty (higher n_0) and has lower search costs than the second segment. This means that the first segment of consumers searches many options but for a short amount of time, while the second segment searches few options but spends a relatively larger amount of time searching. Without information on search duration, the impact of prior uncertainty on search decisions would be missed and attributed manly to differences in search costs.

7 Counterfactual analysis

In this section, we study the role of website design in affecting consumer search and purchase decisions. This analysis is relevant to managers of platforms and intermediaries that determine the format in which consumers interact with information on third party sellers. To this end, we use the parameter estimates from the previous section to simulate consumer choices in relevant scenarios. We focus on the case of one segment of consumers and use the estimates that include additional restaurant characteristics (column 2 in Table 9) to simulate consumer choices using the sequential sampling model presented in section 3. To integrate over the unobserved components in consumers' utility function, we repeat the simulation 100 times and report the mean result.

Our results can be found in Table 12. The table provides results on the percent change in the number of clicks and transactions, as well as the percent change in the price of the restaurant purchased across two scenarios. These results represent short run effects of changing the structural parameters.

First, we ask what is the impact on searches and purchases of the fact that as the website matures, consumers gain greater knowledge of the products displayed, and thus start their search with lower prior uncertainty. To answer this question, we increase the number of samples assumed to form the prior (n_0) by 10%, which lowers prior uncertainty, and present our results in the first column of Table 12. We compare consumer choices in this scenario with their simulated choices under the current design of the website, assuming the options available for search stays the same. We find that lower prior uncertainty leads to fewer searches as well as transactions. Since consumers know more about the options available before starting their search, they do not need to search as many options, saving on search costs. However, fewer searches leads them to find poorer matches, and thus make fewer transactions. The restaurants that benefit from this change are those that are

[22]To obtain the probability of being in segment 1, we use the value for π in Tables 10 and 11 and compute $exp(\pi)/(1 + exp(\pi))$.

priced in the preferred middle range.[23] Note that our counterfactual only considers the symmetric case where all restaurants are impacted in the same way by lower prior uncertainty. With data revealing consumers' asymmetric prior uncertainty (beyond the scope of this paper), these results might differ.

Table 12: Counterfactual Results

Percent change	Lower prior uncertainty (1)	Shift restaurant page information to list page (2)
Choices		
Restaurants clicked	−15.05	1.97
Transactions	−11.70	1.33
Price of restaurant purchased		
Price (<20RMB)	−0.94	−1.74
Price (20-50RMB)	0.40	−0.43
Price (51-79RMB)	0.04	−0.25
Price (80-119RMB)	0.02	0.61
Price (120-200RMB)	−0.41	0.22
Price (>200RMB)	−0.59	−0.34

Second, an important question in designing e-commerce websites is how much and which piece of information to display to consumers before they need to click (i.e. on the list page), and how much to relegate to the page they access through a click (i.e. on the restaurant page).[24] Including too little information on the list page means consumers need to click for additional information frequently, which we know is typically very costly, since consumers only search few options. At the same time, including too much information on the list page may make it difficult for the consumer to compare alternatives on relevant attributes in deciding what to search.

To determine how much information to include on the list page, we simulate consumer choices when the restaurant page information, is moved to the list page. Our results can be found in the second column in Table 12. We show that shifting restaurant page information to the list page encourages more clicks, as well as increases the number of transactions. With more information available on the list page, consumers' prior expected utility is higher, encouraging more search. In addition, conditional on search, since this information is observed with certainty (rather than influencing their signal of the match value of the restaurant), their posterior mean utility is also higher, leading to a higher conversion rate.

8 Conclusion

In this article, we study consumers' decision to spend time searching, in addition to the decision of which products to search and whether to purchase. To this end, we develop a sequential search

[23] In this section, we use estimation results to study the impact of prior uncertainty on consumer choices. These results mirror those from an exercise we performed on simulated data in Figure 4 in section 3.2.

[24] Using different methods, this question has also been investigated by Gu and Wang (2017) and Gardete and Hunter (2018).

model in which consumers are uncertain about their match value for a product and search to reveal (noisy) signals about it. Consumers then use these signals to update their beliefs about products searched in a Bayesian fashion. We model search duration as the consumer's decision to search the same product multiple times. We build on the framework by Chick and Frazier (2012) to describe the optimal search rules for the consumer and develop the model's empirical counterpart, which we then estimate using data on consumers searching for restaurants on an Asian review website. We document that search duration is considerable, that consumers search few options, and that restaurants that are searched longer are more likely to be purchased. Estimating our model on this data allows us to quantify consumer preference, search cost, and prior uncertainty parameters. We find that consumers start their search with high initial uncertainty, which leads them to search few options for a long time, even in the presence of search costs. In counterfactual simulations, we investigate the effect of lower prior uncertainty and of shifting information revealed after to before search. This analysis is of interest to managers of web intermediaries and it would not be possible without a model of search on the intensive margin. Our approach provides a general framework to study consumer engagement with a product through search, and can also be used to capture a consumer's decision to revisit a previously searched product to resolve further uncertainty.

While our approach provides one way in which the duration of search can be endogenized, there could be other approaches that need to be explored given the empirical importance of the duration decision. Two other potentially useful extensions of our study are the following. First, looking at search within and across sites for a product would give a more complete picture of the search process for a product or service. Second, a validation of our counterfactual predictions in the context of a field study would enhance our understanding of the benefits of considering duration data when analyzing search behavior. We leave these and other related topics to future research.

References

[1] Ackerberg, D. 2003. Advertising, Learning and Consumer Choice in Experience Good Markets: A Structural Empirical Examination. *Internat. Econom. Rev.* **44** 1007-40.

[2] Banks, J., R. Sundaram. 1994. Switching Costs and the Gittins index. *Econometrica.* **62** 687-694.

[3] Branco, F., M. Sun, J. M. Villas-Boas. 2016. Too Much Information? Information Provision and Search Costs. *Marketing Sci.* **35**(4) 605-618.

[4] Branco, F., M. Sun, J. M. Villas-Boas. 2012. Optimal Search for Product Information. *Management Sci.* **58**(11) 2037-2056.

[5] Brezzi, M., T. Lai. 2002. Optimal Learning and Experimentation in Bandit Problems. *J. Economic Dynam. Control.* **27**(1) 87-108.

[6] Chen, Y., S. Yao. 2016. Sequential Search with Refinement: Model and Application with Click-Stream Data. *Management Sci.* **34**(4) 606-623.

[7] Chernoff, H. 1961. Sequential Tests for the Mean of a Normal Distribution. *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability.* 79-91.

[8] Chernoff, H., S. N. Ray. 1965. A Bayes Sequential Sampling Inspection Plan. *Ann. Math. Statist.* **36**(5) 1387-1407.

[9] Chick, S. E., N. Gans. 2009. An Economic Analysis of Simulation Selection Problems. *Management Sci.* **55**(3) 421-437.

[10] Chick, S. E., N. Gans, O. Yapar. 2018. Bayesian Sequential Learning for Clinical Trials of Multiple Correlated Medical Interventions. *INSEAD Working Paper No. 2018/20/TOM/ACGRE.* Available at SSRN: `https://ssrn.com/abstract=3184758`

[11] Chick, S., P. I. Frazier. 2012. Sequential Sampling with Economics of Selection Procedures. *Management Sci.* **58**(3), 1-16.

[12] De los Santos, B. 2017. Consumer Search on the Internet. Available at SSRN 1285773.

[13] De los Santos, B., A. Hortaçsu, M. R. Wildenbeest. 2017. Search With Learning for Differentiated Products: Evidence from E-Commerce. *Journal of Business and Economic Statistics.* **35**(4), 626-641.

[14] De los Santos, B., S. Koulayev. 2016. Optimizing Click-through in Online Rankings for Partially Anonymous Consumers. Working paper.

[15] Dong, X., I. Morozov, S. Seiler, L. Hou. 2018. Estimation of Preference Heterogeneity in Markets with Costly Search. Working paper.

[16] Dukes, A., L. Liu. 2015. Online Shopping Intermediaries: The Strategic Design of Search Environments. *Management Sci.* **62**(4), 1064-1077.

[17] Erdem, T., M. P. Keane. 1996. Decision-making under Uncertainty: Capturing Choice Dynamics in Turbulent Consumer Goods Markets. *Marketing Sci.* **15** 1-21.

[18] Erdem, T., M. P. Keane, B. Sun. 2008. A Dynamic Model of Brand Choice When Price and Advertising Signal Product Quality. *Marketing Sci.* **27**(6) 1111-1125.

[19] Fradkin, A. 2017. Search, Matching, and the Role of Digital Marketplace Design in Enabling

Trade: Evidence from Airbnb. Available at SSRN 2939084.

[20] Gardete,P., M. Hunter. 2018. Avoiding Lemons in Search of Peaches: Designing Information Provision. Working paper.

[21] Ghose, A., P. Ipeirotis, B. Li. 2012. Surviving Social Media Overload: Predicting Consumer Footprints on Product Search Engines. Working paper.

[22] Gittins, J., D. Jones. 1974. A Dynamic Allocation Index for the Sequential Design of Experiments. Gani J., Sarkadi K., Vince I., eds. *Progress in Statistics* (North-Holland, Amsterdam) 241-266.

[23] Gittins, J. 1979. Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society.* Series B 41, 148-177.

[24] Gittins, J. 1989. Multiarmed Bandit Allocation Indices. *John Wiley and Sons, New York*.

[25] Gu, C., Y. Wang. 2017. Consumer Online Search with Partially Revealed Information. Working paper.

[26] Hong, H., M. Shum. 2006. Using Price Distributions to Estimate Search Cost. *RAND J. Econom.* **37**(2) 257-275.

[27] Honka, E. 2014. Quantifying Search and Switching Costs in the U.S. Auto Insurance Industry. *RAND J. Econom.* **45**(4) 847-884.

[28] Honka, E., P. Chintagunta. 2017. Simultaneous or Sequential? Search Strategies in the U.S. Auto Insurance Industry. *Marketing Sci.* **36**(1) 21-42.

[29] Hu, M., C. Dang, P. Chintagunta. 2017. Groupon Fatigue: Search and Learning in a Daily Deals Site. Working paper.

[30] Ke, T., J.M. Villas-Boas. 2017. Optimal Learning Before Choice. Working paper.

[31] Ke, T., Z. M. Shen, J.M. Villas-Boas. 2016. Search for Information on Multiple Products. *Management Sci.* **62**(12) 3576-3603.

[32] Kim, J. B., P. Albuquerque, B. J. Bronnenberg. 2010. Online Demand under Limited Consumer Search. *Marketing Sci.* **29**(6) 1001-1023.

[33] Kim, J. B., P. Albuquerque, B. J. Bronnenberg. 2016. The Probit Choice Model Under Sequential Search with an Application to Online Retailing. *Management Sci.* Forthcoming.

[34] Koulayev, S. 2013. Search With Dirichlet Priors: Estimation and Implications for Consumer Demand. *Journal of Business and Economic Statistics.* **31** 226-239.

[35] Koulayev, S. 2014. Search for Differentiated Products: Identification and Estimation. *RAND J. Econom.* **45**(3) 553-575.

[36] Ma, L. 2016. Only the Interested Learn: A Model of Proactive Learning of Product Reviews. Working paper.

[37] Moraga-Gonzalez, J., and M. Wildenbeest. 2008. Maximum Likelihood Estimation of Search Cost. *European Economic Review.* **52** 820-848.

[38] Moraga-Gonzalez, J.L., Z. Sandor, and M. Wildenbeest. 2015. Consumer Search and Prices in the Automobile Market. Working paper.

[39] Lai, T. 1987. Adaptive Treatment Allocation and the Multiarmed Bandit Problem. *Ann.*

Statist. **15**(3) 1091-1114.

[40] Lin, S., J. Zhang, J. R. Hauser. 2014. Learning from Experience, Simply. *Marketing Sci.* **34**(1) 1-19.

[41] McFadden, D. 1989. A Method of Simulating Moments for Estimation of Discrete Response Models Without Numerical Integration. *Econometrica.* **57**(5) 995-1026.

[42] Rothschild, M. 1974. Searching for the Lowest Price When the Distribution of Prices Is Unknown. *Journal of Political Economy.* **82**(4) 689-711.

[43] Seiler, S. 2013. The Impact of Search Costs on Consumer Behavior: A Dynamic Approach. *Quant. Marketing Econom.* **11**(2) 155-203.

[44] Seiler, S., F. Pinna. 2017. Estimating Search Benefits from Path-Tracking Data: Measurement and Determinants. *Marketing Sci.* (forthcoming).

[45] Stigler, G.J. 1961. The Economics of Information. *Journal of Political Economy.* **65**(3) 213-225.

[46] Ursu, R. 2018. The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions. *Marketing Sci.* **37**(4) 530-552.

[47] Weitzman, M. 1979. Optimal Search for the Best Alternative. *Econometrica.* **47**(3) 641-654.

[48] Wu, C., H. Che, T. Chan, X. Lu. 2015. The Economic Value of Online Reviews. *Marketing Sci.* **34**(5) 739-754.