You Can Lead a Horse to Water: Spatial Learning and Path Dependence in Consumer Search*

Charles Hodgson
Yale University and NBER

Gregory Lewis
Amazon

September 8, 2023

Abstract

We develop and estimate a model of consumer search with spatial learning. Consumers make inferences from previously searched objects to unsearched objects that are nearby in attribute space, generating path dependence in search sequences. The estimated model rationalizes patterns in data on online consumer search paths: search tends to converge to the chosen product in attribute space, and consumers take larger steps away from rarely purchased products. Eliminating spatial learning reduces consumer welfare by 13%: cross-product inferences allow consumers to locate better products in a shorter time. Spatial learning has important implications for product recommendations on retail platforms. We show that consumer welfare can be reduced by unrepresentative product recommendations and that consumer-optimal product recommendations depend both on consumer learning and competition between platforms.

*The authors would like to thank Bart Bronnenberg, Jun Kim, and Carl Mela for sharing their data. We would also like to thank Vasilis Syrgkanis and Eliot Abrams for discussions that inspired this paper, and Liran Einav, Matthew Gentzkow, Tim Bresnahan, Marc Rysman, Elizabeth Honka, and participants at numerous seminars.

1 Introduction

An ever-increasing share of consumer activity takes place online. The dominance of online platforms means that, in many markets, consumers have access to choice sets containing hundreds or thousands of alternatives.[1] Consumers may have limited prior knowledge of a product category, and are unlikely to consider every available alternative before making a purchase decision. Search-mediating platforms such as Amazon, Netflix, and AirBnB therefore play a significant role in guiding consumers' search paths through product recommendations and other information provision. Understanding the process of search - how consumers choose their path through alternatives and how this path influences purchase decisions - is therefore increasingly important to understanding consumer markets and the role of platforms.

In most classic models of sequential search, an agent wants to choose one item from a set of heterogeneous objects (products, jobs, etc.) that appear identical (perhaps up to some observable characteristics) prior to search (McCall 1970, Rothschild 1974, Weitzman 1979). Sampling an alternative allows the searcher to learn the payoff from that option, resulting in an optimal stopping problem. Crucially, these models impose independence of the ex-ante unobserved part of utility across alternatives (conditional on observables). What a searcher learns from one alternative does not diferentially affect the expected payoffs of other alternatives.

This paper starts with the observation that, in many real life settings including online consumer search, it is possible that learning about the payoff from one alternative may change the consumer's beliefs about the payoff from other, similar alternatives. We introduce the idea of *spatial learning*: when a searcher samples an option and observes an unexpectedly high or low payoff from that option, they update on the payoffs to other options *that are close in the space of observables*.[2] For example, a job seeker receiving an attractive offer at Microsoft might infer that a potential Google offer would be better than they had expected, but not update on the value of an offer from McKinsey; a student deciding which colleges to apply to may cancel their campus visits to liberal arts colleges after a bad experience with one of them; a consumer looking for a camera who reads negative reviews for a model with low resolution will probably update her beliefs about all low resolution cameras.

This paper makes three contributions to understanding spatial learning and the

[1] At the time of writing, there were 577 different microwave ovens available on Amazon.com.

[2] Note that spatial here refers to product attribute space, and not geographical space. A similar model has been applied to search in geographical space by Hodgson (2019).

role it plays in online consumer markets. First, we develop a model of search with spatial learning and argue that it is identified by data on search paths. Second, we estimate the model using data from online consumer search and show that it can rationalize patterns in consumer search sequences. Third, we show how spatial learning changes the effect of information provision on search paths and affects the design of consumer-optimal recommendations.

The building blocks of the model are a *characteristic space* consisting of ex-ante observable characteristics of the options, and utility functions modeled as a *Gaussian process* over that characteristic space, specified by a mean function (giving the expected payoff to any unsearched option) and a kernel function (giving the covariance between pairs of options). The kernel function takes as inputs the locations of any two options in characteristic space, and outputs a covariance between them. Searchers will update more about close-by options than far-away options. The kernel specifies the distance metric, and encodes the mental model that searchers use to extrapolate. We show that this model of learning leads to *path dependence* in search — a consumer who has a bad experience when sampling some part of the product space will tend to focus their search elsewhere in the future.

We apply our model to data which records the search paths of consumers shopping online for digital cameras, originally collected by Bronnenberg, Kim and Mela (2016). We document a series of stylized facts that are consistent with spatial learning. Consumers tend to take significantly larger steps in attribute space after viewing rarely purchased products, the products searched by consumers converge in attribute space to the product ultimately purchased, and step size in attribute space and the variance of product attributes searched declines as search progresses.

We argue that these search path patterns identify the parameters that control spatial learning - the variance and spatial correlation of the ex-ante unobserved part of utility. The use of search sequences to identify cross-product covariance in utility is novel. With the increasing availability of clickstream data, we expect that this type of identification strategy will become increasingly feasible.

We estimate the model using an approximate dynamic programming approach similar to Keane and Wolpin (1997) and Bertsekas and Tsitsiklis (1996). The estimated model suggests that consumers are spatial learners and make inferences about the utility of unsearched objects that guide their search paths. Search paths simulated using the estimated model fit the data well, and in particular replicate the patterns we highlight as being suggestive of spatial learning - convergence to the chosen attribute levels over the search path and jumps away from rarely purchased products. These patterns cannot be replicated by a constrained version of the model estimated under the assumption of no learning. Simulated search paths also show that learning is quantitatively important to consumer welfare. Expected consumption utility is about 12% lower for simulated consumers who do not make inferences across products than for consumers with correct beliefs. Non-learning consumers have to extend their search length by about 25% to obtain the same utility as consumers with correct beliefs.

The path dependence generated by spatial learning implies that product recommendations can affect consumers' search paths and purchases by changing their beliefs about unsearched alternatives. For example, by highlighting worse-than-expected products in some parts of the product space a search intermediary can steer consumers away from those areas and towards a desired purchase. We investigate this novel mechanism through simulations of the estimated model under different information provision scenarios. We show that recommending products with idiosyncratically high or low utility reduces consumer welfare by providing misleading information about the utility of nearby options, shifting search paths and purchases toward or away from the recommended product in attribute space.

Finally, we demonstrate the importance of spatial learning for platforms' information provision problem by computing consumer-optimal recommendations. The platform faces an explore-exploit trade-off in choosing which products to recommend. When the variance of the spatially correlated part of utility is high, learning is important, and optimal recommendations will push consumers towards products in unexplored parts of the product space that are informative about other unsampled products. Conversely, when there is no spatial learning, optimal recommendations will direct consumers to high expected utility products. The platform's incentive to help consumers explore is moderated by competition. When consumers are likely to switch to another platform, the incentive to recommend high expected utility products dominates, even if spatial learning is important.

Related literature. Search is a well-studied topic in microeconomic theory, empirical industrial organization and marketing. Theory papers have examined how consumers learn product payoffs through search, and how this affects the resulting equilibrium on the supply side (Branco, Sun and Villas-Boas 2012, Branco, Sun and Villas-Boas 2016, Ke and Villas-Boas 2019). Generally speaking, the ex-ante unobservable payoffs are assumed independent across products. An important exception is Adam (2001), who analyzes a model which allows for payoffs to be sampled from a discrete set of nests so that searchers who sample an option from one nest will update their posterior on the distribution for all other items on this nest.

Empirical work in this area has proceeded in several directions. Some of this work has studied the identification and estimation of some of the classic search models (Koulayev 2014, De Los Santos, Hortaçsu and Wildenbeest 2012). A second strand has taken the Weitzman (1979) model to data, including Kim, Albuquerque and Bronnenberg (2010), Honka and Chintagunta (2017), and Ursu, Zhang and Honka (2023) Ursu (2018). A third area of research has followed Rothschild (1974) in allowing for learning. In these models, consumers update their beliefs about the distribution from which searched objects are drawn (De Los Santos et al. 2012, Koulayev 2013, Dickstein 2018, De Los Santos, Hortaçsu and Wildenbeest 2017, Ursu, Wang and Chintagunta 2020).

Among these papers, Gardete and Hunter (2018) is most closely related but focuses on within-product rather than across-product learning. Jindal and Aribarg (2021) provide experimental evidence that consumers update their beliefs about product attributes as they search.

This paper is also related to the literature on platform design and optimal information provision, including Dinerstein, Einav, Levin and Sundaresan (2018), De Los Santos and Koulayev (2017), Ellison and Ellison (2009), Hagiu and Jullien (2011), and Fradkin (2018). Our findings on product recommendations identify cross-product learning as an additional channel through which a platform can influence search. Our results on the optimal design of recommendations complement the theoretical and simulation based results of Dzyabura and Hauser (2019) who show formally that it may not be optimal to recommend the products with the highest ex ante expected utility when consumers learn about their weighting of product attributes.

The Gaussian process model of beliefs builds on the literature on Gaussian processes in machine learning, as summarized by Rasmussen and Williams (2005). Gaussian processes have been widely used as *predictive* models in fields such as geostatistics (Cressie 1992), real estate (Xu, Zhang and Li 2021), and financial time series (Andersen, Davis, Kreiß and Mikosch 2009). In the economics literature, Gaussian processes (including Brownian motion) have been used to model Bayesian priors in theoretical and empirical studies of study of optimal experimentation in product development (Callander 2011) policymaking (Callander and Hummel 2014), and oil and gas exploration (Covert 2015, Hodgson 2019). In each of these papers, the agent faces an exploration problem in similar to the spatial search problem in this paper.

Finally, Korganbekova and Zuber (2023) builds on our model, applying it to data from an online furniture retailer and documenting similar evidence of spatial learning. They combine our model with experimental variation in recommendations to design an optimal personalization algorithm.

Paper outline. The remainder of the paper proceeds as follows. Section 2 provides an illustrative example of spatial learning and path dependence. Section 3 outlines a general model and derives implications for consumer search behavior. Section 4 describes the data on consumer search paths we use to test our model, and presents stylized facts from this data that match model predictions. Section 5 describes the estimation of the model using data on search paths. Sections 6 and 7 present the results of the estimation and counterfactuals, and Section 8 concludes.

2 An Illustrative Example

We begin with an example that illustrates the main forces present in our model. Consider a world with 3 products, A, B and C. A consumer has to buy one of the three (we add an outside option in the main model, but omit it here for simplicity). Their payoff from consumption depends on price and quality according to:

$$u_j = q_j - p_j$$

Quality is unknown to the consumer ex-ante; all they observe are the prices, which are ordered as $p_A < p_B < p_C$. By searching a product, they learn the payoff u_j. Each search costs $c > 0$, and products must be searched before purchase.

Assume that consumers know that $\boldsymbol{q} \sim N(\boldsymbol{p}\mu, \Sigma)$ where $\boldsymbol{q} = (q_A, q_B, q_C)$, $\boldsymbol{p} = (p_A, p_B, p_C)$. $\mu > 1$ is a known scalar. Because μ is positive, price acts as a signal of quality, and because it is greater than one, consumers believe that increasing price implies higher expected utility. The variance-covariance matrix Σ is also known ex-ante. Consistent with the spatial logic offered in the introduction, we assume that it takes the form $\Sigma_{ij} = \lambda \exp(\frac{-(p_i - p_j)^2}{\rho})$. This means that, for example, $cov(u_A, u_B) > cov(u_A, u_C)$. The ex-ante unobserved part of utility (quality in this example) is more highly correlated between products that are closer in terms of ex-ante observable

Figure 1: Optimal Search Strategies

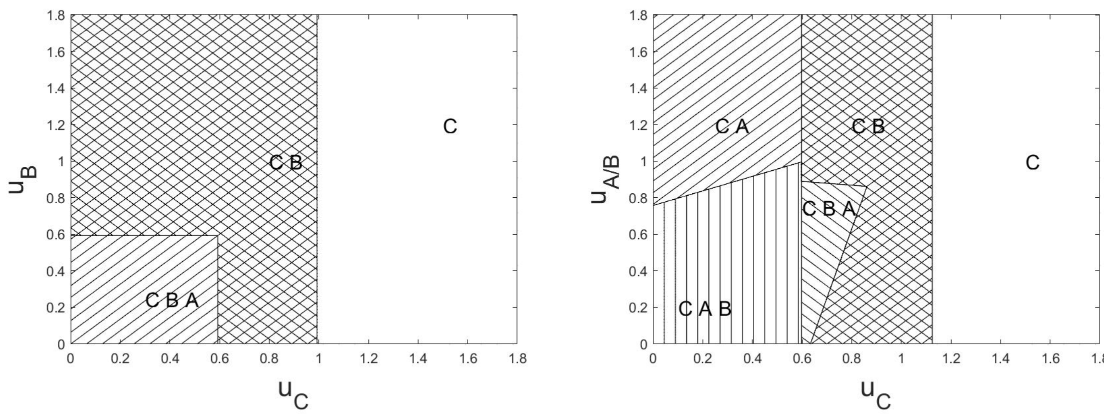

Notes: The left panel shows the optimal search strategies when there is no correlation in quality across products, and observing u_j only provides information about product j. The right panel illustrates how search strategies change when consumers believe that there is positive cross-product covariance in quality. The x-axis is the realized utility of the first product searched, and the y-axis is the realized utility of the second product searched. Each region records the order in which products are searched before the consumer stops searching. In this example, $p_A = 2$, $p_B = 3$, and $p_C = 4$. $\mu = 1.3$, $c = 0.4$, $\Sigma_{ii} = 1.4$, and $\Sigma_{ij} = 1.4 \exp(\frac{-(p_i - p_j)^2}{\rho})$. In the left panel, $\rho \approx 0$ and in the right panel $\rho = 0.8$.

attributes (price in this example).

As an initial baseline, consider a model where $\rho \approx 0$, so that all the off-diagonal elements of Σ are zero and there is no spatial correlation in payoffs. The consumer's optimal policy is illustrated graphically in the left panel of Figure 1 for a specific numerical example.[3] After searching product C, consumers will stop if the observed value of u_C is above the reservation utility z_B, and otherwise will search product B. If the observed utilities u_C and u_B are both below z_A, then the consumer will then search product A. Notice that there is no path dependence; regardless of the utility realizations, consumers will search products in the order C, B, A.

Next consider a model in which $\rho > 0$, so that payoffs are spatially positively correlated. Since $|p_A - p_C| > |p_A - p_B| = |p_B - p_C|$, consumers will update more about B than A the after sampling C. There is no straightforward characterization of the optimal search strategy, and we solve for it numerically by backward induction. The right panel of Figure 1 illustrates the results of this exercise. As before, the

[3]This is a special case of the Weitzman (1979) model. The optimal search algorithm assigns each option a score z_j — which in our example satisfies $z_A < z_B < z_C$ — and requires searching those in decreasing order of score, stopping if the maximum payoff found thus far exceeds the search index of the next option to be searched.

consumer starts by searching product C. But the next product they search depends on the observed value of u_C. If u_C is sufficiently high, they stop and buy it. If u_C is intermediate, they move on to product B, buying either B or C if B is good enough, and only searching A if the max of B and C is low.[4] If u_C is low, they infer that μ is also low, and instead target product A next, moving onto product B if the maximum payoff of A and C is sufficiently low, and otherwise stopping.

This example exhibits the basic logic of spatial learning in consumer search. The differential correlation of utility between products, which is a function of the distance between products in the ex-ante observable attribute space, induces path dependence: each successive outcome determines not only *whether* to stop but *where* to go next. This example is a special case of the general model of search with spatial learning which we develop in the next section.

3 Model

3.1 Environment

A consumer i with unit demand faces a finite set $\boldsymbol{J}$ of available products. Each product has a set of characteristics $X_j \in \boldsymbol{X} \subseteq \mathbb{R}^K$ that are observable to consumers before search. Each product also has an associated search cost c_j. In online search, search costs may differ across products j because of, for example, product rankings of results pages. By paying the search cost, the consumer may learn the payoff u_j from buying product j. Consumers may search as many products at they like. After terminating search, they may consume any product they have searched (they may not purchase a product without searching it first) or choose to consume the outside option instead, with payoff $u_0 = 0$. Their final utility is the payoff from the product consumed, less the sum of the search costs.

We assume that the payoffs have the following structure:

$$u_{ij} = m_i(X_j) + \xi_j + \epsilon_{ij} \tag{1}$$

where $m_i : \boldsymbol{X} \to \mathbb{R}$ is a function that maps a vector of characteristics to aver-

[4]The values of u_B and u_C matter individually too. The downward sloping line at the top of the blue region indicates that for a fixed u_B just above 0.8, the decision to search A depends on whether the news about C was good. If it was good, then the posterior μ is higher and price is a stronger signal of quality, so it is optimal not to search A; whereas if it was bad the converse applies.

age payoffs, ξ_j is a product-level random effect drawn iid across products from a distribution $N(0, \sigma_\xi)$ common to all consumers, and ϵ_{ij} is an idiosyncratic shock sampled iid across consumers and products from a distribution $N(0, \sigma_\epsilon)$. The function $m_i(X)$ is sampled from a Gaussian process with prior mean function $\mu_i(X)$ and covariance function $\kappa_i(X, X')$. We assume that μ is a continuous function, and that $\kappa(X, X') \equiv f_\kappa\left(\frac{\|X-X'\|}{\rho}\right)$ for some weakly positive, continuous and decreasing function f_κ, where $\|\cdot\|$ is the Euclidean norm and ρ is a parameter that controls how covariance declines with distance.[5] Draws of $m_i(X)$ are therefore continuous functions of X centered around $\mu_i(X)$. After searching a product j, consumers observe u_j but do not observe values of ξ_j or ϵ_{ij}, or the function $m_i(X)$. We assume that the consumers know the observables $X_j \in \boldsymbol{X}$, the prior mean and covariance functions, $\mu_i(X)$ and $\kappa_i(X, X')$, and the distributions of ξ_j and ϵ_{ij} prior to search.[6]

3.2 Beliefs and Learning

As consumers search through alternatives, they update their beliefs about the joint distribution of utility for the remaining alternatives. When the consumer searches an alternative j, they observe u_j and updates their beliefs according to Bayes' rule. Their posterior beliefs about $m(X)$ are a Gaussian process with mean and covariance functions:

$$\mu'(X) = \mu(X) + \frac{\kappa(X, X_j)(u_j - \mu(X_j))}{\kappa(X_j, X_j) + \sigma_\xi^2 + \sigma_\epsilon^2} \tag{2}$$

$$\kappa'(X, X') = \kappa(X, X') - \frac{\kappa(X, X_j)\kappa(X_j, X')}{\kappa(X_j, X_j) + \sigma_\xi^2 + \sigma_\epsilon^2} \tag{3}$$

Notice that $\kappa(X_j, X_j)$ is the variance of the "signal" in observed utilities, the part of utility that comes from $m(X)$, and $\sigma_\xi^2 + \sigma_\epsilon^2$ is the variance of the "noise", the part of the observed utility that comes form product-level and idiosyncratic shocks. Figure 2 illustrates an example of the consumer's learning process. Panel A represents a consumer's prior beliefs their preferences over products on a a one-dimensional characteristic space $X \in [0, 100]$. The consumer's prior mean function, $\mu(X) = 0$,

[5]Where here we use X and X' to denote two different locations in product space.

[6]One could imagine a model in which the common utility shock, ξ_j, is also spatially correlated. For example, $u_{ij} = m_i(X) + \xi(X) + \tilde{\xi}_j + \epsilon_{ij}$, where $\xi_j = \xi(X) + \tilde{\xi}_j$ and $\xi(X)$ is drawn from a Gaussian process. Although such a model would have interesting implications, it is not likely that the correlation structure of $m_i(X)$ and $\xi(X)$ could be separately identified following the identification arguments in Section 5 below.

Figure 2: Gaussian Process Learning

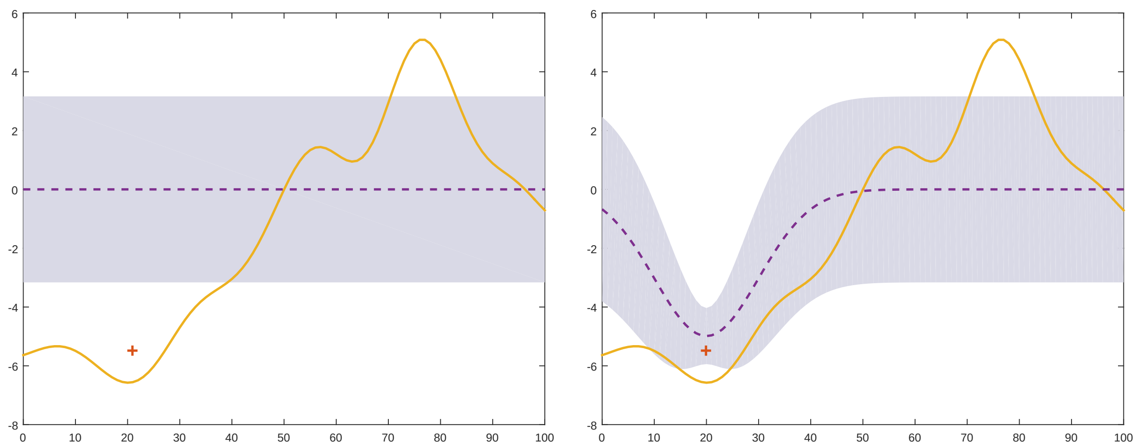

(A) Prior Beliefs (B) Posterior Beliefs

Notes: This figure illustrates Bayesian updating in a single dimensional Gaussian process with mean 0. In Panel A, the dashed line is the prior mean, and the shaded area is a one standard deviation interval around the mean. The solid line is the "true" function which is drawn from the Gaussian process, and the cross is the value observed by an agent, which is equal to the value of the Gaussian process draw plus noise. In Panel B, the dashed line reflects the mean of the agent's posterior beliefs. The shaded area is a one standard deviation interval of the posterior beliefs.

is indicated by a dashed line. The shaded area is a one standard deviation band of consumer's prior. The solid line is the consumer's utility function $m(X)$ which is drawn from the Gaussian process. The consumer searches a product j at location $X_j = 20$ and observes the utility u_j, indicated by the by the point in Panel A. Panel B shows the consumer's posterior beliefs. Notice that the observation has revised down the consumer's expected utility and reduced the consumer's uncertainty, especially for products close to X_j in product attribute space.[7]

There are at least two ways to interpret this model. Under one interpretation, consumers do not know their preferences over characteristic space, $m(X)$. As they search, they get noisy signals of the function. Under the other interpretation, consumers know their preferences over the observable characteristics X, $\mu(X)$, but there are other unobservable product characteristics whose values are unknown without

[7]Notice that the type of learning in this model is different from that modeled by Koulayev (2013) and De Los Santos et al. (2017). In those papers, the ex-ante unobservable part of utility is distributed iid across products, and consumers learn about the common distribution of utility. In our model, products are ex-ante differentiated and utility is differentially correlated across products, so inference from one product's observed utility to beliefs about other products depends on the covariance structure. Our model does not formally nest these models, which use Dirichlet priors. However, when $\rho \to \infty$, the spatial covariance in our model becomes uniform, reducing to model to one in which utilities are drawn iid from a normal distribution with an unknown mean.

search, $m(X) - \mu(X)$. As they search, consumers refine their model of the mapping between the observable and the unobservable characteristics. Under either interpretation, consumers are aware of some set of observable characteristics X which they use to direct their search. In applications, the X can be thought of as those product features that are most salient and likely to be observable to consumers ex-ante.

Two special cases are worth noting. As $\rho \to 0$, the correlation in average payoffs between any two points goes to zero, so that each product has independent and unknown payoffs prior to search. This is the model of Weitzman (1979), specialized to the case of normally distributed payoffs (see Appendix A.2). As $\rho \to \infty$, the correlation in average payoffs goes to one, so that learning the payoffs at any one point is equally informative for all other points.[8]

3.3 Consumer Behavior

The search process is a Markov Decision process. We model the state as a tuple $S = (\mu(X), \kappa(X, X'), \hat{u}, J)$, where $\mu(X)$ are the current mean beliefs, $\kappa(X, X')$ is the current covariance, $\widehat{j}$ is the best product found so far, $\widehat{u}$ is the payoff to the best product found so far and J are the available products remaining to be searched. The transitions on the state variables $(\hat{u}, J)$ are straightforward, and the transitions on $(\mu(X), \kappa(X, X'))$ are given by equations 2 and 3 above.

The consumer's problem is described by the Bellman equation:

$$V(S) = \max \left\{ \hat{u} - c_0, \max_j \left(E\left[V(S') | S \right] - c_j \right) \right\} \tag{4}$$

Given state S, the consumer chooses whether to stop searching and obtain consumption utility $\hat{u}$ or continue searching, in which case they chooses the alternative j that maximizes the expected continuation value, less the cost of search, c_j. The expectation is over the realization of u_j with respect to the consumer's current beliefs, which together with the current state, S, determines next period's state, S', according to the Bayesian process described above.

Unlike in the case of Weitzman (1979), it is not possible to solve the consumer's

[8]Under independence, the beliefs never update and consequently the order of search is pre-determined. The optimal order of search and stopping rule is given by Weitzman (1979). Under perfect correlation, the mean beliefs update everywhere symmetrically so that each update $\mu' - \mu$ is constant in X. In this case, consumers are learning about the average *level* of utility as they search. This can be thought of as learning the common distribution from which utilities are drawn iid across products, as in Rothschild (1974).

problem analytically. It is, however, possible to obtain some intuition about optimal consumer behavior under the assumption that conusmers are myopic, and only look one search ahead. Under this assumption, we can derive a closed form solution for the optimal search rule, which allows us to illustrate some of the forces in the model.[9]

The myopic policy scores the available options based on their expected marginal contribution over the current best option $\hat{u}$. Define $s_j = \sqrt{\kappa(X_j, X_j) + \sigma^2}$, the standard deviation of the payoff of product j (which includes the idiosyncratic shock). Define $a_j = (\hat{u} - \mu(x_j))/s_j$, the current best option normalized by the mean and variance in payoffs for item j. Then we score option j according to:

$$z_j = \Phi(a_j)\hat{u} + (1 - \Phi(a_j))\mu_j + \phi(a_j)s_j - c_j \tag{5}$$

where the first term captures the chance that product j is worse than the current best, the second two are the expected value of product j conditional on being better times the probability of that event, and the last term subtracts the product-specific search cost. That is, the score for product j is the expected value of $\hat{u}$, the best utility searched so far, *after* searching product j, less the search cost. The optimal myopic policy is to search the option with the highest score z_j, so long as it exceeds $\hat{u}$; otherwise to stop and buy the current best option because the expected increase in $\hat{u}$ is less than the search cost. We provide a derivation of z_j in Appendix A.1.

It is straightforward to prove some useful comparative static properties using the analytical characterization of the score in (5).

Proposition 1 (**Comparative statics**).

$$\frac{\partial z_j}{\partial \hat{u}} = \Phi(a_j) > 0 \ , \ \frac{\partial z_j}{\partial \mu_j} = 1 - \Phi(a_j) > 0 \ , \ \frac{\partial z_j}{\partial c_j} = -1 < 0 \ ,$$

Moreover the impact of the payoff to the last search u_k on current scores is given by:

$$\frac{\partial z_j}{\partial u_k} = \frac{\partial z_j}{\partial \mu_j}\frac{\partial \mu_j}{\partial u_k} + 1(u_k = \hat{u})\frac{\partial z_j}{\partial \hat{u}} = (1 - \Phi(a_j))(\kappa(X_j, X_k)/s_k^2) + 1(u_k = \hat{u})\Phi(a_j)$$

These properties are intuitive, but have some interesting implications. First, an improved current best option affects the score of a product at a rate that depends on whether its payoff may fall below the best option - i.e. based on the tail risk

[9] This myopic or "one-period look ahead" assumption is common in the literature on Gaussian processes, which typically employ n-period look ahead assumptions (Osborne, Garnett and Roberts 2009).

of an option. It follows that consumers score risky options more highly when they have better existing options. Second, the comparative static on search cost implies an important role for product rankings and visibility in driving search paths.

Finally, the main beneficiaries of a higher payoff for the last search are options that have high covariance, $\kappa(X_j, X_k)$, with the last search location. Since the consumer's prior $\kappa(X_j, X_k)$ is decreasing in the distance between X_j and X_k, this means that observing a high (low) utility draw from a product k will increase (decrease) the search index z_j of products j that are close to k in attribute space more than products that are far from k in attribute space.[10] Thus, differential covariance across products induces path dependence in search - a low draw of u_k will make a consumer less likely to search similar products in future.

While the myopic policy is not generally optimal, it can provide a close approximation to the optimal policy in certain cases. Frazier, Powell and Dayanik (2009) provides explicit bounds on the suboptimality of the myopic, or "knowledge gradient" policy in the case of Gaussian process beliefs. They show that the this policy is close to optimal when $\kappa(X, X')$ varies little across pairs of products and is exactly optimal when the mean payoffs are perfectly correlated ($\rho = \infty$) or independent ($\rho = 0$). Intuitively, the myopic policy does not allow consumers to search products *because* they are differentially informative about other products that may be searched in future. If consumers are forward-looking, it may be optimal to search some product j with a low value of z_j if the expected continuation value after searching this product is high. Our empirical application in Section 6 below will use approximate dynamic programming techniques to allow for such forward looking incentives.

4 Empirical Evidence from Online Search

4.1 Data

We apply our model of consumer search with spatial learning to data which records the search paths of consumers shopping online for digital cameras. The data comes from ComScore, who track the online browsing behavior of panelists who have installed

[10]This is precisely true when k is the first product searched. For later searches, $\kappa(X_j, X_k)$ is not only a function of distance but also of past searches. Intuitively, variation in $\kappa(X_j, X_k)$ should largely be a function of distances between products in areas of the search space that are less well explored. This is not the case in models without differential correlation across products, including Weitzman (1979) and Rothschild (1974).

ComScore's tracking software. The sample we use was constructed by Bronnenberg et al. (2016) (henceforth BKM), and comprises the browsing activity of 967 ComScore panelists who were searching for digital cameras between August and December 2010. Although this is a selected sample and not necessarily representative of the population of consumers, it covers search across all domains and covers a product category that is a good candidate for the application of our model. Digital cameras are infrequent purchases and heterogeneous is quality and features. There are hundreds of models of digital camera available, so consumers are unlikely to know their mapping from models to utility before searching. However, there are several salient attributes such as zoom and pixel that many consumers are likely to understand and use to direct their search. In this section we illustrate patterns in this data that are suggesting of spatial learning, before discussing how these patterns identify learning in the structural model in the next section. [11]

For an individual panelist, we observe the sequence of products viewed, the product eventually purchased (if any), and the date and time of each observation. Product views were detected by scraping the sequence of URLs visited by consumers for product information. The data covers all browsing behavior and therefore is not limited to one retailer. A product "view" or "search" (we use the terms interchangeably) in the data is recorded when a product listing page on a retail site (e.g. Amazon) is loaded. Searches are therefore recorded at the product-retailer level. Purchases are identified using a second ComScore dataset that tracks online transactions carried out by panelists. For each product search, the data records the product make, model, and four continuous product attributes - price, zoom, display size, and pixels. The conversion of the raw ComScore browsing data and the matching of this data to product attributes was performed by BKM, and extensive details on the preparation of the data are provided in that paper.

Defining a product as a unique combination of brand, pixel, zoom, and display, and taking the average price recorded for that combination results in 357 products and 1022 product-retailers. The left panel of Table 1 records summary statistics on the distribution of four continuous attributes (price, zoom, display size, and pixels)

[11] Similar search patterns to those documented here have been found in work by Korganbekova and Zuber (2023), who use a data set covering all consumers on Wayfair. One could apply our model to other setting where we do *not* expect consumers to have limited information or to make cross product inferences, for example frequently purchased household staples. As discussed in Section 5 below, we would expect to estimate learning parameters (λ, ρ) close to 0 in a setting where we do not observe these search patterns.

Table 1: Summary Statistics

Products	Mean	SD
Price	299.52	431.00
Zoom	6.47	5.95
Pixel	11.13	2.79
Display	2.73	0.37
Availability	2.88	1.44

Searches	Mean	SD	Min	Max
Length	5.86	7.46	1	69
Purchase Discovered	0.84	0.26	0.03	1
Price Searched	275.30	365.18	16.99	5250.09
Zoom Searched	6.40	5.95	1	35
Pixel Searched	12.04	2.29	1	21
Display Searched	2.79	0.29	1.1	3.5
Retailers Searched	1.87	1.03	1	5

Notes: Left panel records statistics on products from the digital camera data. An observation in the first four rows is a product-retailer. Availability is the number of retailers at which we observe the product being searched. An observation in the fifth row is a product. Right panel records statistics on search paths from the digital camera data. Search path length is the number of product-retailer combinations viewed. Product discovered is recorded in terms of search percentile, as defined in the text. Product attributes searched record the distribution over all consumer-product observations.

attributes across products. Products may be sold by multiple retailers (i.e. domains) at different prices. The fifth row of this panel indicates that the average product is available at 2.88 retailers, where we treat the top four retailers separately and combine all other retailers into a composite "other". Appendix Table A.2 records the share of products for each retailer and brand.

The right panel of Table 1 records summary statistics on the 967 consumer search paths. Recall that search paths are a sequence of product-retailers. In our discussion of search paths we will refer to product-retailers as "product:s for brevity, where retailer can be thought of as one of the product attributes. The first row of the records path length - the number of unique product combinations searched.[12] The average consumer views about 5.9 products. There is a tail of consumers with very long search paths, the longest of which is 69 products. The second row documents the search percentile at which the ultimately purchased product is first discovered. If a consumer searches T products in total, then the search percentile of the tth product is $\frac{t}{T}$. Note that the Tth product is not necessarily the product purchased. The chosen product is typically discovered towards the end of search.[13] The remaining

[12]Note that we include only a consumer's first visit to a product URL in the analysis data, and we therefore drop any revisits to the same product after that product has been searched once. These revisits represent around 24% of the URL queries recorded in the raw data. The model described in Section 3 cannot rationalize revisits, since the consumer learns their utility for a product after searching it once.

[13]9.6% of recorded search paths end in no purchase. These paths are omitted from the statistic "purchase discovered" in Table 1. When we apply our model to the data, we treat these consumers

rows documents the distribution of attributes among products *searches*. Comparing these distributions to the distributions of product attributes in the top panel indicates that products which are less expensive are searched more. Similarly, products have higher resolution and a larger display are searched more often.

4.2 Convergence in Product Space

In this section we present several stylized facts that describe how consumers move through the product attribute space as they search. We argue that these descriptive statistics suggest that consumers begin search with some uncertainty about their preferences over these four attributes, and that they update their beliefs about their preferences for un-searched items after viewing each product in their search path.

Figure 3 replicates one of the main findings of BKM - that the attributes of products searched get closer to the attributes of the product eventually purchased as search progresses. The left panel plots search percentile on the x-axis against the distrance in log price between the product searched at that search percentile and the product eventually purchased. This Figure shows that the attributes of the product being viewed get closer to those of the product eventually purchased over time. Products considered, but not purchased, in late search are more similar in price to the purchased product than products considered in early search. The right panel shows that the same is true of log zoom. The same pattern can be observed in other product attributes (pixels, and display size), as documented in Appendix Figure A.6. Note that this result is not driven by the fact the the purchased product tends to be first discovered towards the end of the search path, since the purchased product is excluded from the data.[14]

In addition to getting closer to the purchased product on average, consumers search a wider variety of products and take larger "steps" through attribute space early in the search path than later in the search path. We document these additional facts in Appendix C.1. Taken together, these patterns suggest that consumers explore a wider variety of products early in their search before narrowing in on close substitutes to the product that is ultimately purchased. This behavior is not predicted by standard models of sequential search. In contrast, correlated Gaussian process

as choosing an outside option.

[14]In Appendix Table A.3 we present these results as linear regressions including consumer fixed effects, showing that distance to the chosen product decreases significantly with search percentile. These regressions also show that consumers do not get further from the chosen product after switching domains.

Figure 3: Convergence to Chosen Attribute Level

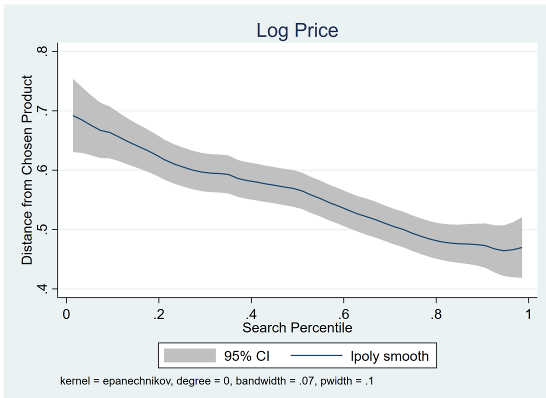

Notes: The y-axis for each panel records, for the relevant product attribute, the absolute difference in standard deviations of the attribute between the searched product and the product ultimately purchased. The x-axis reports the search percentile, as defined in the text. The product ultimately purchased is excluded from the data for each consumer. The solid line is a kernel regression using an Epanechnikov kernel, and the shaded area is 95% confidence interval. The estimation sample includes all search paths from the ComScore data on search for digital cameras, with revisits to the same product dropped.

learning has been shown to exhibit this type of convergence behaviour. Frazier et al. (2009) show that agents following a myopic policy searching over alternatives with payoffs drawn from a multivariate normal will tend to explore the search space early on, and then concentrate later search in high-payoff regions. These findings are difficult to rationalize without a model in which there is a spillover of information between searched and un-searched objects.

4.3 Step Size and Path Dependence

In this subsection we test a direct implication of the model of search with spatial learning developed in Section 3. Proposition 1 implies that when an object is observed to have a higher than expected utility, other objects that are nearby in attribute space move up the search ranking more than objects that are distant in attribute space. Likewise, when a searched product had lower than expected utility, objects that are closer in attribute space move down the search ranking more than distant objects.

These implications of the model are difficult to test directly, since we do not observe consumer preferences. An ideal experiment would randomly expose consumers to one of two objects, j and k, with $X_j = X_k$, but $\xi_j > 0 > \xi_k$. That is, two objects at the same location in the ex-ante observable product space, but with different unobservable product effects. After viewing object j, consumers should, on average, make the inference that similar objects also yield higher utility than expected, and should be more likely to subsequently search nearby products. Consumers that view object k should, on the other hand, be less likely to subsequently search nearby products.

To approximate this experiment we rely on the observation that different values of ξ_j not only generate different search path patterns, but also generate different purchase patterns. In particular, products with high values of ξ_j should be purchased more frequently than similar products, conditional on being searched. We test whether this is true: do products that are purchased less (more) often, relative to observably similar products, also induce larger (smaller) "jumps" in attribute space? To do this we construct a product level index $\hat{\theta}_j$ which measures how much more or less likely a product is to be purchased than other products with similar attributes X_j. High values of $\hat{\theta}_j$ mean that a product is purchased more, conditional on being searched, than similar products. Vice versa for low $\hat{\theta}_j$. In the context of our model, variation in $\hat{\theta}_j$ across products is explained by variation in product effects, ξ_j. Details on the construction of this measure are in Appendix C.2. We then regress a measure of the "step size" of search after a consumer observes product j on this index.[15]

Let $j(i,t)$ be the product searched by consumer i on the tth search (we will sometimes write this j_{it} to make expressions easier to read). To test for consumer learning, we regress measures of step size, for example $\Delta price_{it} = |price_{it} - price_{it-1}|$ on the estimated index of the last product viewed, $\hat{\theta}_{j(i,t-1)}$. If consumers are spatial learners, Proposition 1 implies that the size of the consumer's tth search step should be negatively correlated with $\hat{\theta}_{j(i,t-1)}$. We run this regression for four observable attribute dimensions - log price, log pixels, log display size, and log zoom - and record coefficients in Table 2. All regressions include controls for search percentile, an indicator for whether product $j(i,t-1)$ is the product ultimately purchased, product density controls, and consumer fixed effects.[16] $\hat{\theta}_{j(i,t-1)}$ is standardized so that the first row reports the effects of one standard deviation changes of $\hat{\theta}_{j(i,t-1)}$.

$\hat{\theta}_{j(i,t-1)}$ has a significant, negative effect on step size for each of the four attribute

[15]Note that the index $\hat{\theta}_j$ does not have a structural interpretation. It is a descriptive statistic that captures how frequently a product is purchased compared to other products with similar observables. The results in Table 2 should therefore be thought of as descriptive regressions that are consistent with the behavior predicted y the model.

[16]Product density is the average distance between $j(i,t-1)$ and all other products in the relevant observable attribute dimension. If "surprisingly bad" products tend to be located in regions of the attribute space that are sparsely populated by other products, then step size after searching one of these products will mechanically be larger.

Table 2: Effect of Product Residuals on Step Size

	$\Delta price_{it}$	$\Delta pixel_{it}$	$\Delta zoom_{it}$	$\Delta display_{it}$	$\Delta domain_{it}$	$\Delta brand_{it}$
$\hat{\theta}_{j(i,t-1)}$	-0.197***	-0.255***	-0.118**	-0.244***	0.037	0.036
	(0.042)	(0.045)	(0.060)	(0.048)	(0.036)	(0.038)
$SearchPercentile_{it}$	-0.130***	-0.020	-0.026	-0.015	0.178***	-0.091***
	(0.036)	(0.038)	(0.051	(0.041)	(0.031)	(0.032)
$Purchased_{it-1}$	-0.076***	0.031	-0.085*	-0.011	-0.096***	-0.080***
	(0.034)	(0.037)	(0.049)	(0.039)	(0.030)	(0.031)
N	3976	3976	3976	3976	3976	3976
Consumer FE	Yes	Yes	Yes	Yes	Yes	Yes
Density Controls	Yes	Yes	Yes	Yes	Yes	Yes
Mean of Dep. Var.	0.569	0.550	0.667	0.598	0.365	0.491

Notes: Table presents regressions of search step size on the product residual index $\hat{\theta}_{j(i,t-1)}$. Values of $\hat{\theta}_{j(i,t-1)}$ are standardized so that estimated coefficients are the effect of one standard deviation. Any product observations where j_{it-1} is never purchased, and hence a value $\hat{\theta}_{j(i,t-1)}$ is not computed, are omitted from the regression. All regressions include consumer fixed effects. The data includes all search paths in which at least two products are searched, with revisits to the same product dropped. *** indicates significance at the 99% level. ** indicates significance at the 95% level. * indicates significance at the 90% level.

dimensions. A one standard deviation decrease in $\hat{\theta}_{j(i,t-1)}$ increases step size in log price by 0.093, which is 18% of the average step size in log price recorded in the final row of Table 2. Similarly, a one standard deviation decrease $\hat{\theta}_{j(i,t-1)}$ increases step size in in log pixels by 30% of the average, in log zoom by 15% of the average, and in log display by 18% of the average. The results indicate that consumers take larger than average steps in attribute space after viewing products that are rarely purchased (those with low values of $\hat{\theta}_{j(i,t-1)}$). That is, *purchase* behavior associated with a specific product predicts *search* behavior after consumers have viewed that product. This finding is strongly suggestive of learning, and is in line with what we would expect to observe if consumers made inferences about nearby products after each search, per Proposition 1. When consumers view products with "surprisingly low" utility (those with low values of ξ_j), they jump further away in attribute space.[17]

The fifth and sixth columns of Table 2 report the results of analogous regressions of $\Delta domain_{it}$ and $\Delta brand_{it}$ on $\hat{\theta}_{j(i,t-1)}$. $\Delta domain_{it}$ is an indicator that is equal to 1 when the consumer switches domain, or retailer, between search $t-1$ and search t, for example from eBay.com to Amazon.com. $\Delta brand_{it}$ is an indicator that is equal to

[17] In Appendix Table A.6 we run examine the robustness of these results to the definition of $\hat{\theta}_j$ using alternative binary classification of products as "frequently" or "infrequently" purchased. The results are consistent with the pattern in Table 2.

1 when the tth and $t-1$th products searched are of different brands. In both cases, the coefficients on $\hat{\theta}_{j(i,t-1)}$ are not statistically significan. These results are consistent with a model in which consumers learn about the value of product attributes, and not about the value of brands or retailers. This is perhaps reasonable if brand reputation is established through advertising before search begins, and the value of retailers is ex-ante known from prior shopping experience.

These effects suggest that the information consumers obtain from search affects not only their purchase decisions but also the direction of their search paths. If the effects recorded in Table 2 persist, then they induce path dependence in search. Viewing a product with a low value of ξ_j rather than an otherwise identical product with a high value of ξ_j could permanently divert the consumer's search path by pushing search to another area of the attribute space. On the other hand it could be that the effects in Table 2 are transient, and any change in the step size is undone by subsequent search.

To determine the extent to which jumps in step size are persistent, we regress two and three step differences in product attributes, for example $|price_{it} - price_{it-2}|$, on two and three step lags of $\hat{\theta}_j$. The results of these regressions are recorded in Appendix Table A.7. The estimated coefficients indicate that the correlation between $\hat{\theta}_j$ and step size persist. The coefficients are mostly negative and slightly lower in magnitude than the one-step coefficients in Table 2.[18]

Together, the results discussed in this subsection indicate that consumers jump away from from low-$\hat{\theta}_j$ products and tend to *stay away* in subsequent search, although this effect fades with subseqeust steps as consumers obtain more information. This pattern is consistent with a persistent effect of observing low-ξ_j products on consumers' beliefs generating path dependence in search. To quantify the importance of these effects to consumer welfare, and to further investigate the implications of path dependence in search for platform power we next turn to estimating the structural parameters of the model.

[18] In Appendix Table A.8 we report further regressions of *forward* one-step differences, for example $|price_{it+1} - price_{it}|$ and $|price_{it+2} - price_{it+1}|$, on lags of $\hat{\theta}_j$. We find no significant effects of $\hat{\theta}_{j(i,t-1)}$ on any one-step difference size except the tth. We also find no effect of $\hat{\theta}_{j(i,t-1)}$ on *past* step sizes.

5 Structural Estimation

In order to take the model developed in Section 3 to the data on consumer search paths, we make several additional assumptions.

5.1 Multiple Retailers

The empirical setting deviates from the model described in Section 3 because search take place over multiple retailers that may sell overlapping sets of products. To deal with this, we let each alternative j denote a combination of product and retailer. Let $\tilde{j}_j$ denote the product identity (e.g. Nikon P100) corresponding to alternative j and r_j denote the retailer (e.g. Amazon.com). This has two implications for the model.

First, we assume that the alternative-specific random effects, ξ_j, are identical at the product level, so $\xi_j = \xi_{j'}$ if $\tilde{j}_j = \tilde{j}_{j'}$. Because of this, learning across alternatives takes place not only because of the correlation in $m(X)$ but also because multiple alternatives share the same ξ_j. This is accommodated through an appropriate modification of the Bayesian updating rules in equations 2 and 3. See Appendix D for details.

Second, we introduce a retailer switching cost. Let $\tilde{r}$ indicate the retailer of the last viewed alternative. If the consumer chooses to search an alternative with $r_j \neq \tilde{r}$ they pay c_{switch} in addition to the search cost c_{ij}. Because this introduces a dependency between $\tilde{r}$ and the value of searching each alternative, j, $\tilde{r}$ must be included in the consumer's state variables.

5.2 Empirical Specification

We assume that consumers' prior means are linear in product characteristics:

$$\mu(X_j) = \alpha + X_j\beta_i + \gamma_{r(j)} + \delta_{b(j)} \tag{6}$$

Where $\gamma_{r(j)}$ are retailer fixed effects and $\delta_{b(j)}$ are brand fixed effects. Notice that we allow for consumer heterogeneity in the prior mean functions through consumer-specific coefficients β_i. The model therefore nests the random coefficients discrete choice model of Berry, Levinsohn and Pakes (1995).[19] We assume that the coefficients

[19] In particular, when $\gamma = 0$ and $c_{ijt} = 0$ the model collapses to a probit choice model with linear utility and random coefficients.

β_i, are normally distributed according to equation 7, where we restrict Ω to be a diagonal matrix and denote the kth diagonal element ω_{kk}.[20]

$$\beta_i \sim N(\beta, \Omega) \tag{7}$$

We assume that that consumers' prior covariance function $\kappa_i(X_j, X_l)$ is of the form given by equation (8). This is similar to the square exponential covariance function introduced earlier in the text but allows the covariance between $m_i(X_j)$ and $m_i(X_l)$ to decay with distance at different rates along different dimensions of the product characteristic space. In particular, there are K parameters ρ_k that control spatial correlation in utility along the K dimensions. The parameter λ controls the overall variance level of the prior Gaussian process.

$$\kappa(X_j, X_l) = \lambda^2 exp\left(\sum_{k=1}^{K} \frac{-\left(X_{jk} - X_{lk}\right)^2}{2\rho_k^2}\right) \tag{8}$$

Let $\boldsymbol{\rho}$ be the vector with kth entry ρ_k.[21] To further simplify the consumer's problem, we suppose that consumer i's cost of searching product j at period t, c_{ijt}, is given by equation 9, where c is a parameter, and ζ_{ijt} is a logit error term that is drawn independently across t, i, and j. The logit assumption simplifies subsequent computation, and captures variation in search costs across consumers and products due to, for example, variation in page rankings. We do not observe page ranking in our data, but in applications where such data is available costs could be conditioned on page ranking or other variables capturing product salience.[22]

$$\begin{aligned} c_{ijt} &= c + \zeta_{ijt} \\ c_{i0t} &= \zeta_{i0t} \end{aligned} \tag{9}$$

[20] Due to the standard price endogeneity concern, the parameters β should be interpreted as the net effect on *expected* utility of changes in characteristics, $\beta_{price} = \frac{\partial E(u)}{\partial price}$, fixing consumer beliefs. This is not a problem for the counterfactual exercises we perform, which focus on information provision rather than price changes. We provide a discussion of this issue in Appendix F.1.

[21] Note that $\kappa(X_j, X_l)$ is not a function of retailer, $r(j)$, or brand, $b(j)$. This amounts to assuming that consumers know the brand's and retailer's contribution to utility ex-ante. This assumption is made for computational tractability, but is consistent with the results in Table 2.

[22] Costs could also be conditioned on demographic variables. Prior work has shown that search costs are related to consumer demographics (De Los Santos et al. 2017). For simplicity, we have chosen not to project search costs onto demographics here, though this could be accommodated at some computational cost.

Finally, we normalize the level of utility by giving consumers an outside option with utility zero, setting $\hat{u}_{i0} = 0$ for all i. Note that in our application to digital cameras we only observe an individual if they make at least one search. To deal with this, we assume that consumers must make at least one search, and afterwards can choose to stop searching without purchasing a product and obtain outside option utility $\hat{u}_{i0} = 0$.

Thus the parameters to be estimated comprise those determining the prior mean, $\{\beta, \alpha, \Omega, \gamma_{r_j}, \delta_{r_j}\}$, those determining the prior covariance function, $\{\lambda, \boldsymbol{\rho}\}$, the search cost parameters c and c_{switch}, and the parameters the control the "noise" in consumers' learning process - the variances $\{\sigma_\xi, \sigma_\epsilon\}$ and the values of the product effects, ξ_j. Let ψ be the set of parameters to be estimated. Given ψ and a K dimensional vector of product attributes for each of the J products, the model generates a distribution of search paths and purchase decisions.

5.3 Estimation

We estimate the model by constructing the likelihood of the observed consumer search paths and choices. Under the assumption that search costs are given by equation (9) with logit errors, we can write the expectation of the value function as,

$$V(\mathcal{S}, \psi) = \log \left(exp(\hat{u}) + \sum_j exp\left(E\left[V(\mathcal{S}')|\mathcal{S}\right] - c\right) \right). \tag{10}$$

Which is equal to the expectation of Bellman equation 4 with respect to c_j.[23] Note that here we make explicit the dependence of the value function on the parameters ψ. The probability of a consumer choosing to search product $j \in J$ conditional on being at state $\mathcal{S}$, but unconditional on the realizations of the logit cost shocks is then given by:

$$P_i(j|\mathcal{S}, \psi) = \frac{exp\left(E[V(\mathcal{S}', \psi)|\mathcal{S}] - c\right)}{exp\left(\hat{u}\right) + \sum_{l \in J} exp\left(E[V(\mathcal{S}', \psi)|\mathcal{S}] - c\right)} \tag{11}$$

The expectations are over realizations of the sampled utility, u_j, which determines the evolution of the state variable. In particular, after sampling u_j, $\hat{u}' = \max\{\hat{u}, u_j\}$.

Suppose consumer i searches T_i times before stopping. Let j_{it} be the tth product

[23] The is the expected value up to the additive Euler-Mascheroni constant. Because $V(\mathcal{S}', \psi)$ and c enter linearly in the lkelihood below, we can think of this constant as part of c.

searched. Let $j_{it} = 0$ indicate stopping and purchasing the highest utility sampled product (or the outside option). Finally, let $\hat{j}_i$ indicate the product purchased. If the consumer's state variable, $\mathcal{S}$, was fully observable to the econometrician, the likelihood of the consumer's search path would then be given by equation 12.

$$L_i(\{j_{it}\}_{t=0}^{T_i}, \hat{j}_i | \{\mathcal{S}_t\}_{t=0}^{T_i}, \psi, \beta_i) = \left(\prod_{t=0}^{T_i - 1} P_i(j_{it}|\mathcal{S}_t) \right) P_i(0|\mathcal{S}_{T_i}) 1 \left(u_{\hat{j}_i} = \hat{u}_{j_{iT_i}} \right) \tag{12}$$

Since the econometrician does not observe the utility draws that enter $\mathcal{S}$, it is necessary to integrate them out of the likelihood function. Conditional on ψ, the vector of utilities observed by consumer i, $u_i = \left(u_{i,j(i,t=1)}, ..., u_{i,j(i,t=T_i)}\right)$, is distributed according to a multivariate normal distribution, $G(u_i) = N(\bar{u}_i, \Sigma_i)$. The vector of mean utilities, $\bar{u}_i$, has a τth entry given by $\mu(X_j) + \xi_{j(i,\tau)}$. The covariance matrix Σ_i has diagonal elements $\kappa(X_{j(i,\tau)}, X_{j(i,\tau)}) + \sigma_\epsilon^2$ and off-diagonal elements $\kappa(X_{j(i,\tau)}, X_{j(i,\tau')})$ for $\tau \neq \tau'$.The likelihood function for consumer i unconditional on utility draws is given by equation 13.

$$L_i(\{j_{it}\}_{t=0}^{T_i}, \hat{j}_i | \psi) = \int \int L_i(\{j_{it}\}_{t=0}^{T_i}, \hat{j}_i | \{\mathcal{S}_t\}_{t=0}^{T_i}, \psi) dG(u_i) dF(\beta_i) \tag{13}$$

The inner integral is taken over the distribution of u_i given β_i. In practice, we approximate these integral by averaging over draws from $G(u_i)$. The outer integral is taken over the distribution of β_i, given by equation 7. We search over parameter vectors ψ to maximize the objective function,

$$L(\psi) = \prod_{i=1}^{N} L_i(\{j_{it}\}_{t=0}^{T_i}, \hat{j}_i | \psi) \prod_{j=1}^{J} \phi(\xi_j, 0, \sigma_\xi) \tag{14}$$

Which is the product of the likelihood of N conusmer search paths, and the prior $\xi_j \sim N(0, \sigma_\xi)$. We multiply the likelihood with the prior on ξ_j to obtain a maximum a posteirori objective function. This helps discipline the estimates of ξ_j in finite sample, for example for products j that are rarely viewed, and can be thought of as a regularization or shrinkage term. Notice that as $N \to \infty$ the prior drops out of the log likelihood. The esimates are therefore asymptotically equivalent to maximum likelihood.

5.4 Approximating the Value Function

To compute the likelihood for a candidate parameter value, ψ, we need to solve for the consumer's continuation value, $V(\mathcal{S}', \psi)$, in equation 11. Observe that the consumer's state variable $\mathcal{S}$ includes current beliefs, which are described by a vector of J means with elements $\mu(X_j)$ and a $J \times J$ covariance matrix with (j, j') element $\kappa(X_j, X_{j'})$. The state variable therefore has dimension at least $J + \frac{1}{2}J(J-1)$,with $J \approx 1000$ alternatives in the data. Rather than solving for the value function at every possible state point for every candidate ψ, we instead adopt an approximation to the value function that allows us to interpolate between states as in Keane and Wolpin (1997) and Crawford and Shum (2005).

We perform value function iteration on equation 15. On the kth iteration, we use this equation to compute the value function at a set $\mathcal{S}_W$ of W simulated states and parameter vectors. We then estimate a neural network regression of $V_{k+1}(\mathcal{S}, \psi))$ on $(\mathcal{S}, \psi)$, generating an approximated value function, $\hat{V}_{k+1}(\mathcal{S}, \psi)$. The approximated value function is then used as the continuation value in the next iteration of the Bellman Equation. This procedure, which alternates between iterating the Bellman equation and approximating the value function using a neural net, is based on the "approximate value iteration" procedure described by Bertsekas and Tsitsiklis (1996), who provide results on the conditions required for convergence.

$$V_{k+1}(\mathcal{S}_{it}, \psi) = \log\left(exp(\hat{u}) + \sum_j exp\left(E\left[\hat{V}_k(\mathcal{S}_{it}, \psi)|\tilde{S}_{i\tau}\right] - c\right)\right). \tag{15}$$

The function approximation in is necessary because computing the expectation in equation 15 requires evaluating the value function at states outside $\mathcal{S}_W$. The neural network regression thus serves as a means of extrapolating the value function from a finite number of sampled states, allowing for forward looking incentives without forcing the researcher to choose an arbitrary functional form for the value function. For large enough W and a sufficiently rich neural network specification, the approximation $\hat{V}_k$ can be made arbitrarily accurate (Park, Yun, Lee and Shin 2020). Details of the implementation of this method, including a discussion of further state space reduction, are provided in Appendix E.

We substitute $\hat{V}(\mathcal{S}', \psi)$ into equation 11 in place of $V(\mathcal{S}', \psi)$ and find ψ that maximizes the likelihood in equation 14. Notice that since all the parameters that are relevant for the consumer's problem enter $\hat{V}(\mathcal{S}', \psi)$, it is not necessary to repeat the value function iteration for each candidate value of ψ.[24]

5.5 Identification

Our model differs from standard models of sequential search because of the presence of the spatially correlated beliefs controlled by the parameters $\{\lambda, \rho, \sigma_\xi, \sigma_\epsilon\}$. The identification argument we outline below relies on patterns in the search *sequences* to identify the parameters $\{\lambda, \rho, \sigma_\xi, \sigma_\epsilon\}$. Note that we do not normalize any of the three variance parameters $\{\lambda, \sigma_\xi, \sigma_\epsilon\}$ that control the variance of the spatially correlated, within product, and iid components of utility respectively. We argue that our model puts sufficient structure on the covariance of utility across products that each of these variance parameters is identified from data on sequential search[25]

The heuristic argument is as follows. The probability of each possible search and purchase sequence is identified directly from the data as the number of consumers grows large. The probability that each product is searched first identifies the parameters of the prior mean, β. The variance, Ω, of the random coefficients is identified by the relative variation in search product attributes across and within individuals. If there is more variation across individuals than within individual search paths in the attributes of searched products, this suggest greater heterogeneity in β_i. This is similar to the standard argument for identification of preference heterogeneity in discrete choice panel data as in Keane (1997). The intercept α is identified by the share of consumers choosing the outside good, and the search cost parameter, c, is identified by the by the distribution of search lengths. The probability that product j is purchased, conditional on being searched, identifies ξ_j. Note that the argument so far makes no use of the *sequence* of search.

The identification of the variance and covariance parameters, $\{\lambda, \rho, \sigma_\xi, \sigma_\epsilon\}$, is novel and merits a more detailed argument. These parameters are identified by cross-product variation in ξ_j and *second search* probabilities. To see this, suppose for simplicity that $\beta_i = 0$ and $\alpha_i = 0$ and that the observable product attribute space is one-dimensional. Consider, for example $P(j_{i2} = B | j_{i1} = A)$, the probability that a consumer searches product B second conditional on searching product A first. This

[24]One concern is that the approximation may be poor at values of ψ that are far from any that appear in $\mathcal{S}_W$. We discuss this issue in Appendix E.

[25]To provide additional supporting evidence for this heuristic argument, we present additional Monte Carlo exercises in Appendix F.3 in which we show that attempting to "normalize" one of the variance parameters reduces the fit of the model. Appendix F.2 also includes discussion of the reliance of identification on functional form.

probability depends on consumers' posterior beliefs about the distribution of utility from product B. Following equations 2 and 8, the average (Across consumers) expected utility of product B after searching product A is

$$\begin{aligned} E(u_B|j_{i1}=A)=&\frac{\kappa(X_A,X_B)(\xi_A)}{\lambda^2+\sigma_\xi^2+\sigma_\epsilon^2} \\ \kappa(X_A,X_B)=&\lambda^2 exp\left(\frac{-\left(X_A-X_B\right)^2}{2\rho^2}\right). \end{aligned} \tag{16}$$

Which depends on ξ_A and $(X_A-X_B)^2$. Thus, variation in distance between pairs of products, $(X_j-X_k)^2$ and variation in ξ_j across products will generate variation in the expected utility of product k after searching product j, and therefore variation in the second search probabilities probabilities $P(j_{i2}=k|j_{i1}=j)$.[26] This allows us to identify ρ from differential second search probabilities for products at different distances from the first searched product, and $\frac{\lambda^2}{\lambda^2+\sigma_\xi^2+\sigma_\epsilon^2}$ from variation in second search probabilities after first searching products with different values of ξ_j. Intuitively, if there is no spatial correlation in beliefs, $\frac{\lambda^2}{\lambda^2+\sigma_\xi^2+\sigma_\epsilon^2}=0$, and variation in ξ_j should not affect beliefs after the first search. If there is spatial correlation, then the extent to which consumers "jump" away from nearby products after searching a product with $\xi_j<0$ depends on ρ.

Recall that the data includes searches in which consumers view the same product across multiple retailers. Product A yields the same utility across all retailers (up to retailer fixed effects). Therefore, the average expected utility of product A on retailer 2 after the consumer first searches product A on retailer 1 is

$$E(u_{\tilde{j_{i2}}=A,r_j=2}|\tilde{j}_{i1}=A,r_{j_{i2}}=1)=\frac{\left(\sigma_\xi^2+\lambda^2\right)(\xi_A)}{\lambda^2+\sigma_\xi^2+\sigma_\epsilon^2}. \tag{17}$$

Following a similar argument to that above, the extent to which the probability of second searches for the same product at different retailers varies with ξ_j identifies $\frac{\lambda^2+\sigma_\xi^2}{\lambda^2+\sigma_\xi^2+\sigma_\epsilon^2}$. Intuitively, the lower $\lambda^2+\sigma_\xi^2$, then the less correlated is utility within product across retailers, and therefore the less likely to view products with high ξ_j at multiple retailers.

Consider now the *variance* across consumers of the expected utility of product B

[26] Proposition 1 formalizes this in the case of the myopic model. With forward looking consumers, the link between posterior beliefs and search probabilities is more complex and does not have an analytic expression.

after first searching product A. Cross consumer variation in utility realizations comes through ϵ_{ij}. This variance is given by

$$Var(u_B|j_{i1} = A) = \sigma_\epsilon^2 \left(\frac{\kappa(X^A, X^B)}{\lambda^2 + \sigma_\xi^2 + \sigma_\epsilon^2} \right)^2. \tag{18}$$

Since ρ and $\frac{\lambda^2}{\lambda^2+\sigma_\xi^2+\sigma_\epsilon^2}$ are already identified, the variance in second search probabilities across consumers identifies σ_ϵ^2. Consumers with the same prior make different second search choices because of their idiosyncratic draws ϵ_{ij} and search cost shocks ζ_{ijt}, the variance of which is normalized. Thus, σ_ϵ^2 rationalizes the overall variance in second search probabilities. Putting equations 16, 17, and 18 together, identification of ρ, $\frac{\lambda^2}{\lambda^2+\sigma_\xi^2+\sigma_\epsilon^2}$, $\frac{\lambda^2+\sigma_\xi^2}{\lambda^2+\sigma_\xi^2+\sigma_\epsilon^2}$, and σ_ϵ can be combined to identify the parameters $\{\lambda, \rho, \sigma_\xi, \sigma_\epsilon\}$.[27]

The path dependence patterns in search path data that this model seeks to explain are therefore the source of variation in the data that helps identify $\{\lambda, \rho, \sigma_\xi, \sigma_\epsilon\}$. Crucially, a model without learning could be identified without data on the *sequence* (i.e. the order) of searches, and it is this sequential variation that we lean on to identify the novel components of our model.[28]

To provide additional evidence of identification we run a Monte Carlo exercise. Appendix Table A.4 reports the mean and standard deviation of the estimated parameters for 150 simulated data sets with $N = 1000$ search paths each. Estimated parameters are close to the true values and covered by 95% confidence intervals.

[27]The argument above relies on *second* search probabilities only, but the data includes the full search sequence, providing additional identifying variation. Intuitively, consider the population of consumers who search a sequence of products $(1, 2, ..., J)$ in that order. Suppose there are two observably identical products A and B. The distribution of utilities among consumers who search $(1, 2, ..., J, A)$ and $(1, 2, ..., J, B)$ are identical. However, if $\xi_A \neq \xi_B$ then the distribution of subsequent searches will differ between the two groups, providing variation that helps to identify the learning parameters. With sufficiently large data it may also be possible to identify heterogeneity in consumer uncertainty, λ. For instance, if some consumers consistently take larger steps away from bad products than other consumers, they may have more prior uncertainty. We do not pursue this due to computational limitations and the small size of the data.

[28]Note that without the learning component ($\lambda = 0$), our model is based on a standard linear random coefficients utility specification (e.g. Berry, Levinsohn and Pakes (1995)). The identification argument offered here demonstrates how a random utility specification with a *richer* cross-product covaraince structure can be identified using sequential search data under assumptions about consumer learning.

6 Results

6.1 Parameter Estimates

We estimate the model on the digital camera search path data from BKM using the approach discussed above. Observable characteristics known to the consumer before searching are log price, pixels, display size, and log zoom. All characteristics are standardized to have mean 0 and standard deviation 1 across products.

The estimated parameters are presented in Table 3. The coefficient on price is negative and statistically significant and the coefficients on pixels and display are positive and statistically significant. The coefficient on zoom is negative but small is magnitude. In particular, the estimated variance of the prior expected marginal utility of zoom, ω_2,is high relative to the mean, β_2, suggesting significant heterogeneity in prior preference for zoom across consumers. This is consistent with zoom being an important dimension of horizontal differentiation - with some consumers searching for larger professional cameras with a higher zoom and others compact, low zoom cameras. There is also significant heterogeneity in prior divisibility of price.

The standard deviation of the Gaussian process $m(X)$ from which consumers' preferences are drawn, λ, and the covariance parameters ρ_k for all four attribute dimensions are positive and significant. Recall that as $\rho_k \to 0$, the model converges to a standard sequential search model without learning. The data on search paths therefore provides evidence that consumers update their beliefs about un-searched objects as they search.

The estimated value λ is of the same order of magnitude as the the standard deviations of the product effects, σ_ξ, and the idiosyncratic error, σ_ϵ. That is about one third of the ex-ante unobservable variation in utility is attributable to the spatially correlated component, $m(X)$, and consumers therefore make meaningful inferences about the utility of unsearched products from observed utilities.[29]

Table 4 illustrates the fit of the model to the data. The first two columns record the mean and standard deviation of various statistics across search paths in the data. The third and fourth column record these same statistics across 5,000 search paths simulated using the estimated parameters. For each simulation, we draw a new value of $m(X)$ from the Gaussian process and new values of the idiosyncratic errors ϵ_{ij}. We hold ξ_j fixed across simulations at their estimated values. The results in the first

[29] The "signal to noise ratio" is approximately $1:3$.

Table 3: Estimated Parameters

	Estimate	SE		Estimate	SE
β_1 (log price)	-2.452	0.086	c_{switch}	1.988	0.040
β_2 (log zoom)	-0.224	0.098	c	7.860	0.025
β_3 (pixels)	2.602	0.090	ρ_1 (log price)	1.066	0.029
β_4 (display)	0.444	0.090	ρ_2 (log zoom)	0.608	0.015
ω_1 (log price)	2.386	0.081	ρ_3 (pixels)	1.920	0.036
ω_2 (log zoom)	1.960	0.097	ρ_4 (display)	2.213	0.033
ω_3 (pixels)	0.925	0.072	λ	22.321	0.005
ω_4 (display)	0.966	0.100	σ_ξ	25.840	0.496
α	-44.559	0.121	σ_ϵ	22.820	0.007
γ_1(Amazon)	-1.486	0.097	δ_1(Canon)	3.372	0.130
γ_2(BestBuy)	-1.347	0.106	δ_2(Kodak)	2.880	0.135
γ_3(eBay)	-3.983	0.132	δ_3(Nikon)	4.545	0.142
γ_4(Other)	-2.163	0.190	δ_4(Other)	2.221	0.205

Notes: Table reports estimated parameters and standard errors computed using the observed Fisher information. For more details on the estmiation procedure, see Appendix D.

two rows indicate that the average search percentile at which the purchased product is first discovered and the share of consumers choosing the outside option in the simulated paths match the data reasonably well. The next four row rows record the average observable characteristics of the products searched. For each attribute, the average characteristic searched in the data is close to the simulated value. The final two rows show that the simulated paths are somewhat longer than the paths in the data, and consumers sample products from more platforms in the simulations. The imperfect fit to search length is likely the result of numerical imprecision in computing the integrals in equation 13.[30]

6.2 Search Path Patterns

As discussed in Section 4, the model of search and learning is motivated by descriptive patterns from the search data. First, as recorded by Table 2, consumers take systematically larger steps in attribute space after viewing products that are rarely

[30] The integrals are computed using Monte Carlo draws of $m_i(X)$ and β_i. Simulated search length is further from the data for estimation results that use fewerMonte Carlo draws. We suspect that, for any finite number of draws, c will tend to be underestimated because reducing the search costs increases the likelihood of any particular search sequence up to the stopping decision.

Table 4: Model Fit

	Data		Simulations	
	Mean	SD	Mean	SD
Outside Option Share	0.041	0.199	0.047	0.212
Chosen Product Discovered	0.835	0.262	0.784	0.252
Average Price Searched	275.299	365.179	277.521	421.358
Average Zoom Searched	6.404	5.954	6.177	5.706
Average Pixel Searched	12.045	2.292	11.743	2.608
Average Display Searched	2.792	0.289	2.774	0.341
Search Length	5.862	7.459	7.917	5.787
Outlets Searched	1.873	1.035	2.516	1.171

Notes: The first two columns report statistics on search paths from the data used in estimation, as in Table 1. The second two columns records analogous statistics for 5,000 simulated search paths, holding all parameters at their estimated level and redrawing $m_i(X)$ and ϵ_{ij} for each simulated consumer.

purchased. Second, as recorded by Figures 3 and A.2, consumers get closer to the purchased product as they search. We now show that our estimated model can replicate these patterns, while a restricted version of our model without spatial learning cannot. We illustrate this by replicating some of these descriptive exercises with simulated search paths. Two sets of search paths are simulated: one uses the baseline parameter estimates, and the other uses "no learning" parameters estimates. The no learning estimates impose the restriction $\lambda = 0$.[31] They are recorded in Appendix Table A.9.

Table 5 replicates the step size regressions recorded in Table 2. At the baseline parameters, the model matches these step size patterns well. As with the real data, the coefficient on $\hat{\theta}_{j(i,t-1)}$ for the simulated data is negative and statistically significant for each of the four dimensions. Data simulated from the model generates these patterns because products with large or small residuals $\hat{\theta}_j$ correspond to products with large or small product effects, ξ_j. Products have large estimated residuals in the simulated data *because* they have large product effects, and product effects ξ_j affect step size through consumer beliefs. As discussed in Section 5, these patterns are an important source of identification for the parameters λ and ρ of the Gaussian process beliefs.

Under the no learning restriction, the estimated model cannot replicate these step size patterns. Indeed, the estimated parameters on $\hat{\theta}_{j(i,t-1)}$ in the lower panel of

[31] We also impose the restriction that consumers do not update their beliefs across listings of same product at multiple retailers. Estimation of no-learning parameters requires normalizing σ_ϵ, as the identification argument above does not apply when $\lambda = 0$.

Table 5: Simulations: Effect of Product Residuals on Step Size

Baseline Parameters

	$\Delta price_{it}$	$\Delta pixel_{it}$	$\Delta zoom_{it}$	$\Delta display_{it}$
$\hat{\theta}_{j(i,t-1)}$	-0.0252***	-0.0129***	-0.0352***	-0.0113***
	(0.0035)	(0.0034)	(0.0035)	(0.0035)
N	64403	64403	64403	64403

$\lambda = 0$ Parameters

	$\Delta price_{it}$	$\Delta pixel_{it}$	$\Delta zoom_{it}$	$\Delta display_{it}$
$\hat{\theta}_{j(i,t-1)}$	0.0067	-0.0145***	0.0042	-0.0033
	(0.0041)	(0.0040)	(0.0041)	(0.0041)
N	53782	53782	53782	53782

Notes: Table presents regressions of search step size on the product residual index $\hat{\theta}_{j(i,t-1)}$. Sample is 5,000 simulated search paths at the estimated parameter values. The top panel uses simulations at the baseline parameter estimate. The bottom panel uses simulations at parameters estimated under the restriction $\lambda = 0$. Specifications are otherwise identical to those described in Table 2. *** indicates significance at the 99% level. ** indicates significance at the 95% level. * indicates significance at the 90% level.

Table 5 are not statistically different from zero for three of the four product attributes. Without spatial learning, there is no mechanism through which product effects ξ_j can affect beliefs about other products.

Figure 4 replicates the exercise recorded in Figure 3, which records the relationship between search percentile and distance of the searched product from the purchased product. The left panel of Figure 4 reports this relationship for log price in simulations using the baseline parameter estimates. As in the real data, simulated consumers get significantly closer to the purchased product as they search, along both product attribute dimensions.[32] This pattern is generated by the dynamics of spatial learning in the model, and is not an artifact of the data. The right panel records the same relationships in search paths simulated using the $\lambda = 0$ parameters. The convergence in attribute space is eliminated in these simulations. Indeed, when there is no learning the searched product moves *away* from the chosen product as search progresses. When we shut down cross-product learning, all product utilities have the same variance and so, apart from the effects of search cost shocks, consumers will search products in order of ex-ante expected utility, from high to low.[33] Because prior expected utility is linear is product attributes, consumers will tend to move away from the ex-ante

[32]The same exercise for price, display and pixels is recorded in Appendix Figure A.8. Price and display exhibit similar patterns.

[33]This is a direct implication of Weitzman (1979).

Figure 4: Simulations: Convergence in Attribute Levels

Baseline Parameters $\qquad$ $\lambda = 0$ Parameters

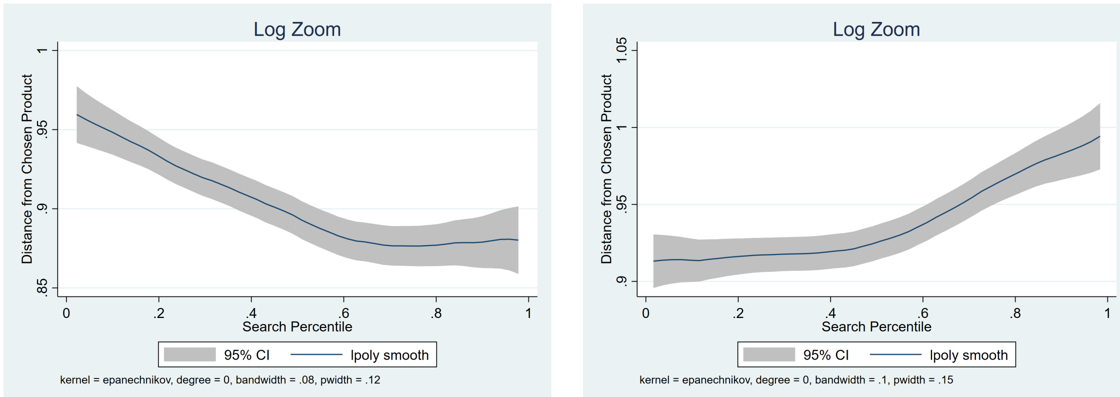

Notes: Figures are constructed using 5,000 search paths simulated at the estimated parameters. The left panels uses the baseline estimates, and the right panel uses the estimates under the restriction the $\lambda = 0$. Specifications are otherwise identical to those used in Figure 3.

most preferred product as they search. Our model generates convergence rather than divergence in attribute space because the variance of the spatially correlated part of utility, λ, is sufficiently large that consumer's refinement of beliefs as they search outweighs the importance of the prior expected utility. The ability of the model to rationalize these patterns and the step size patterns in Table 5 therefore makes the presence spatial learning a plausible explanation for several aspects of search behavior documented by BKM that cannot be rationalized by standard models.

6.3 The Value of Learning

How important is spatial learning to consumer welfare? To answer this, we use the estimated model to ask how consumer search paths would be different under different assumptions about consumer beliefs and learning. The model is estimated under the assumption that consumers know the distribution of the ex-ante unobserved part of utility, $m(X)$, and use the utilities they observe for searched products to make correct Bayesian inferences about unsearched products. In particular, the model assumes that consumers know the true spatial covariance parameters, $\boldsymbol{\rho}$, that govern the correlation of the unobserved part of utility along observed attribute dimensions. To quantify the value of learning to consumers we simulate consumer search paths assuming consumer utilities are distributed according the to estimated parameters but consumers have incorrect beliefs about this distribution. In particular, we assume consumers believe

Figure 5: The Value of Learning

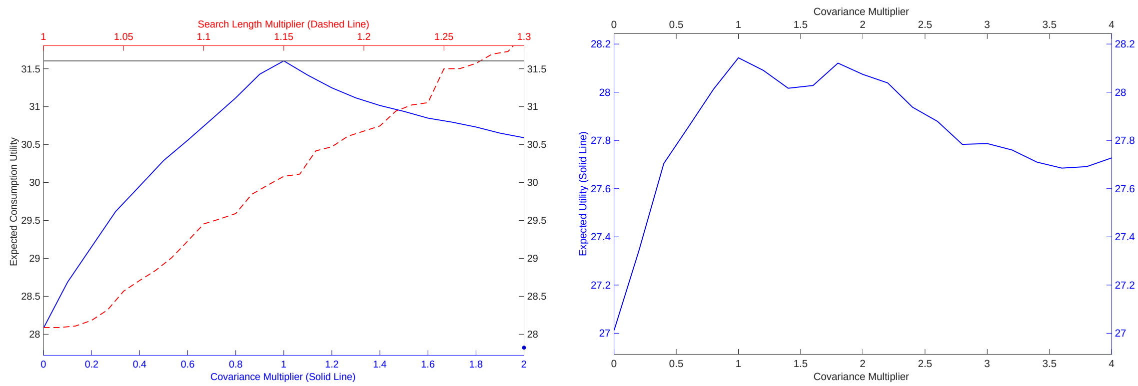

Fixed Search Length | Optimal Search Length

Notes: In the left panel, the blue solid line records, for values δ along the lower x-axis, the average consumption utility of 20,000 simulations when consumer beliefs have covariance parameters equal to $\delta\hat{\boldsymbol{\rho}}$,with all other parameters are at their estimated value. Search length is held fixed for each simulated consumer its length in the $\delta = 1$ simulation. The blue point is the limit of the blue line as $\delta \to \infty$. The dashed red line records, for values of γ along the upper x-axis, the average consumption utility for analogous simulations where search length for consumer i is set to γl_i and the covariance multiplier is set to $\delta = 0$. The right panel records the average total utility (consumption utility less search costs) for analogous simulations without fixed search length.

the spatial covariance parameters to be $\delta\hat{\boldsymbol{\rho}}$. For example, if $\delta = 0$, then although consumers have correct beliefs about the total variance of unobserved utility, they do not make inferences across products because they believe the covariance of $m(X)$ along all dimensions to be 0. We fix search path lengths to isolate the effect of different learning assumptions on the *consumption utility* of the best product located in a fixed number of searches. This allows us to benchmark the effect of different beliefs to changes in search length and ask how much more consumers with incorrect beliefs would have to search to achieve the same level of consumption utility.

The left panel of Figure 5 records the results of these simulations. The solid blue line plots the mean consumption utility across simulations for different values of of the covariance multiplier, δ, indicated by the lower x axis. Consumers obtain the best match to a product in a fixed number of searches when $\delta = 1$. Consumption utility is highest when consumers have correct beliefs about the covariance parameters, ρ. When $\delta < 1$, consumers *under-extrapolate* from observed products to unobserved products, such that if a consumer obtains a particularly high utility draw or a given product, they do not update their beliefs about surrounding products as much as a consumer with correct beliefs, and is therefore more likely to move away from that region of the product space. this under-extrapolation leads to a monotonic reduction in consumption utility as $\delta \to 0$. At $\delta = 0$, expected consumption utility is about 13% lower than at $\delta = 1$.

When $\delta > 1$, consumers *over-extrapolate*, and will, for example, move too far away from a region of the product space based on a low utility draw. This also results in a decrease in consumption utility. As $\delta \to \infty$, the perceived correlation in $m(X)$ across products tends to 1. At this limit, consumers update beliefs equally at all distances from an observed product. There is therefore no spatial learning (only learning about the overall level of utility), and, for fixed search lengths, the expected consumption utility is about the same as if $\delta = 0$. This is illustrated by the blue dot at the far right of Figure 5, which simulates a counterfactual in which $\kappa(X, X') = \lambda^2$, the limit of the function given by equation 8 as $\rho \to \infty$.

To benchmark the value of learning, we ask how much longer a consumer who does not update their beliefs ($\delta = 0$) would have to search to obtain the same level of utility as a consumer with correct beliefs ($\delta = 1$). To do this, we run simulations where $\delta = 0$ and each consumer's search length is set to γl_i for values of γ between 1 and 2. When search length is extended, consumers obtain better matches even though they do not update their beliefs as they search. The results of these simulations are recorded by the red dashed line in Figure 5. Consumers that do not learn as they searches have to sample about 27.5% more products than consumers who learn optimally to obtain the same level of utility in expectation.

Similar patterns obtain when search length is not fixed in simulations. The right panel of Figure 5 repeats the simulations in which consumers believe the spatial covariance parameters to be $\delta\hat{\boldsymbol{\rho}}$, but does not fix search length. The blue line records average total utility - consumption utility minus total search costs. As in the fixed length simulations, utility is maximized when $\delta = 1$ due to over-and under-extrapolation when $\delta \neq 0$.

7 Path Dependence and Product Recommendations

These findings suggest that consumer learning and in particular cross-product inference plays an economically significant part in determining consumers' search paths. Affecting consumer beliefs and search paths through information provision is therefore a potentially important channel through which online retail platforms can influence purchase decisions. It is well established that (for example, see Ursu (2018)) that highly ranked, recommended, or salient products on online platforms are more likely to be searched first. In addition to the direct effect of recommendation on search costs, spatial learning introduces a secondary "path dependence" effect on consumer beliefs. In particular, consider an experiment in which all consumers are forced to view a particular product before beginning their search through the remaining products. In the model with spatial learning, changing this "recommended product" will change the beliefs consumers have at the beginning of their search, and therefore change their subsequent paths.

7.1 Search Diversion

To illustrate the path dependence effect of information provision on search and welfare, we use the estimated model to simulate search paths under different information provision scenarios. We draw 50,000 values of $m(X)$ and simulate search paths. For each search path, we draw a "focal product", $F \in J$ from the set of existing product locations, and "show" the consumer the utility they would obtain from this product before they begin their search.[34]

The effect of this type of information provision on consumers' search paths depends on the values of the unobserved product effect, ξ_F, for the focal product F. Consumers do not observe ξ_F separately from total utility. Particularly large (positive or negative values) of ξ_F can therefore divert search away from or towards different areas of the product space. For example, if consumers view a product with a particularly large negative value of ξ_F, they will attribute this partly to the Gaussian process draw $m(X_F)$ and infer that nearby products will also yield low utility and will divert their subsequent search path. In this sense, products with large values of ξ_F are *misleading* and not *representative* of the spatially correlated part of preferences, $m(X)$.

To illustrate this effect, we run the information provision simulation for a range of values of ξ_F. The left panel of Figure 6 illustrates the effects of information provision on utility. The blue solid line plots the mean utility (consumption utility minus search costs) across simulations for different values of the focal product effect. Expected utility is maximized locally when ξ_F is close to 0 because in this case product F is *representative* of the part of utility that is correlated across observable dimensions.

When ξ_F is increased or decreased from 0, average consumption utility falls. This efffect is driven by misleading information diverting search paths away from or towards

[34] To isolate the effect of this change in consumers' beliefs on the search path, we require consumers to pay a search cost and view the focal product again before buying it. This means that we are only providing information, not reducing the search cost of obtaining a particular product.

Figure 6: Search Diversion

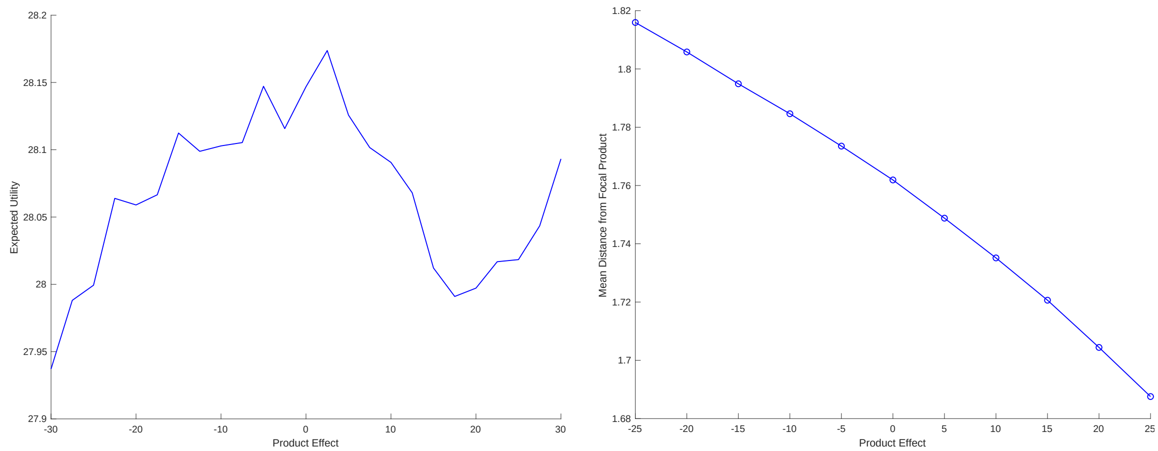

Notes: The left panel records the average utility (consumption utility minus search cost) over 50,000 simulated paths for differen values of ξ_F.The right panel records the mean of $|logprice_j - logprice_{iF}|$ where j is searched by consumer i for simulations using different values of ξ_F.

the focal product, as illustrated in the right panel of Figure 6 Past some threshold, expected utility is increasing in ξ_F. The effect of misleading information reducing on horizontal match quality, $m(X)$ is eventually offset by the vertical quality component of F, ξ_F - in the limit as $\xi_F \to \infty$, consumers always buy F and obtain infinitely high utility. The significant search diversion recorded in Figure 6 imply that if a retail platform wants to direct consumers towards or away from certain products, the platform's information provision design should take account of these effects.

7.2 Optimal Recommendations and Platform Competition

The results discussed so far have illustrated how information provision can divert search paths. These effects could be exploited by a platform to increase revenue. For example, if different products are deferentially profitable to an online retail platform, the platform may want to choose the set of products which are displayed most prominently on the page to direct consumers towards high margin products.[35] However, a forward looking platform may also have an interest in maximizing consumer utility to encourage consumers to return to the platform in future.[36] We now ask what the

[35]In Appendix A, we show that it can be optimal for a firm to recommend a "bad product" in order to divert search towards a chosen high-margin product using the numerical example of Section 2.

[36]Indeed, a recent report in the Wall Street Journal described an internal debate at Amazon.com over the extent to which the search algorithm should highlight more "relevant results" or more

model tells us about the characteristics of consumer-optimal product recommendations when consumers are spatial learners.

Let the set of products on platform (retailer) r be $J_r = \{j \in J : r(j) = r\}$. The platform has information I_{it} about consumer i after their tth search. For each consumer i after each search t, the platform scores products $j \in J_r$ according to a recommendation function $f(I_{it}, j)$. The platform then reduces the search cost of the alternative with the highest score by a factor $\Delta \in [0, 1]$. That is,

$$c_{ijt} = c\left(1 - \Delta 1\left(j = \arg\max_{j \in J_r}\{f(I_{it}, j)\}\right)\right) + \zeta_{ijt}. \tag{19}$$

Consumers are thus more likely to search the product that is given the highest score, $f(I_{it}, j)$, by the platform.[37] We suppose that the platform is interested in maximizing total consumption utility of purchases *on the platform*. The platform's problem is therefore,

$$\max_f E_i\left[P\left(\hat{j}_i \in J_r\right) E\left(u_{i\hat{j}_i} | \hat{j}_i \in J_r\right)\right]. \tag{20}$$

Where the outer expectation is taken over consumer types. The expression inside the square brackets is the probability of a consumer i making a purchase on platform r, multiplied by their expected utility conditional on making a purchase on platform i. Notice that the platform is interested in both increasing on-platform sales, and increasing the quality of the consumer-product match conditional on a sale.

To simplify the problem assume that $f(I_{it}, j)$ is linear in: product quality, ξ_j, the average distance from product j to the set of products that have already been *searched* by consumer i, $DistS_{ijt}$, and the average distance from product j to the remaining *unsearched* products, $DistU_{ijt}$.[38] The platform can therefore choose to direct consumers towards high-ξ_j products, towards products in unexplored parts of the product space, or towards products that are similar to what they have already

"profitable results" (Mattioli 2019), with a spokesperson for the company emphasizing that the algorithm's historical focus on relevant results was in the interest of "long-term profitability".

[37] We assume that consumers are unsophisticated in that they do not use the identity of the recommended product to update their beliefs. That is, they behave as if they believe $c_j = c + \zeta_{ijt}$. This means, for example, they cannot use the fact that product j was recommended to infer something about ξ_j, nor can they choose to search alternatives in order to manipulate future recomemendations.We also omit the potential reoptimization of prices after a change in recommendations, which would require an additional model of price setting.

[38] $DistS_{ijt}$ is defined as $DistS_{ijt} = \frac{1}{|J_S|}\sum_{\ell \in J_S}\sqrt{\sum_k (X_{jk} - X_{\ell k})^2}$ where k indexes the four observable product attributes, and J_S is the set of searched products. $DistU_{ijt}$ is defined analogously.

Table 6: Optimal Sequential Recommendation: Amazon.com

	Specification				
Coefficient on:	Baseline	No Learning	$\beta_i = \beta$	$c_{switch} = 10\hat{c}_{switch}$	$c_{switch} = 0$
ξ_j	1	1	1	1	1
$DistS_{ijt}$	-0.182	-1.085	-0.322	-0.246	-0.983
$DistU_{ijt}$	-0.477	0.222	-1.042	-1.064	0.394
At Baseline Model:					
On-Platform Purchases	1306	1302	1301	1302	1312
On-Platform Utility	37.825	37.104	37.832	37.765	36.997
Utility	33.723	33.326	33.763	33.763	33.390

Notes: Top panel records the coefficnets on ξ_j, $DistS_{ijt}$, and $DistU_{ijt}$ is the linear recommendatoion function $f(I_{it}, j)$ that maximized the platform's objective for each specification. Bottom panel records statistics from simulations of the baseline model using the recommendation coefficneis in each column. Each simulatio uses 3,000 search paths.

searched. These variables are representative of the type of data typically available to online platforms: ξ_j could be measured by aggregate sales data, and $DistS_{ijt}$ and $DistU_{ijt}$ are functions of a conumer's search history.

We find the coefficints on these variables that solve the problem in equation 20 under various model specifications. The results for Amazon.com, the largest platform in the data, are reported in Table 6. Since the function $f(I_{it}, j)$ is used to rank products, we normalize the magnitude of the coefficient on ξ_j to 1 (allowing it to be either positive or negative). Each of the three variables is standardized so that the magnitudes of the coefficients are comparable. For these simulations, we set $\Delta = 0.5$, so recommendations reduce the search cost by half, and we restrict recommendations to each consumer's first 10 searches. We simulate 3,000 search paths for each candidate recommendation function $f(I_{it}, j)$.

The first column records the optimal recommendation coefficients for the baseline model. The results illustrate the "explore-exploit" incentives faced by the platform in choosing which product to recommend. The coefficient on product quality, ξ_j, is positive. By pushing consumers towards high-ξ_j products, the platform generates higher expected utility. The coefficient on $DistS_{ijt}$ is negative. This means that alternatives that are *closer* in product space to previously searched alternatives are more likely to be recommended. Although the platform does not know consumers' idiosyncratic prior coefficients, β_i, they do know what each consumer has searched so far. Since consumers are likely to search products with high prior expected utility, initial search provides a signal of the consumer's prior preferences, and pushing consumers towards similar products raises expected utility.[39]

The platform therefore *exploits* its information about ξ_j and consumers prior searches to push consumers to high ex-ante expected utility products. However, when consumers are spatial learners, the platform also has an incentive to help consumers *explore* the product space to learn their preferences more quickly, in order to generate better matches on the ex-ante unknown part of utility, $m_i(X)$. This is reflected by the negative coefficient on $DistU_{ijt}$. This means that products that are close to more *unsearched* alternatives are more likely to be recommended. Because utility is spatially correlated, products that are near more unsearched alternatives are more informative.

Optimal recommendations therefore balance two incentives, to direct consumers to high expected utility products and to help them explore. Which of these dominates depends on the importance of the different components of utility. When $\lambda = 0$, there is no cross-product learning, and there is no informational advantage to searching products that are close to unsearched alternatives. In this case, as illustrated in the second column, the coefficient on $DistU_{ijt}$ is positive. On the other hand, when $\beta_i = \beta$, and there is no heterogeneity in consumers' priors, then variation in past searches is less informative about consumer utility. In this case, illustrated in the third column, exploration is more important, and the coefficient on $DistU_{ijt}$ is negative and larger in magnitude.

Ignoring learning incentives in optimizing recommendations lowers consumer utility. In particular, when the no-learning recommendations are applied in the baseline model (with learning), average consumption utility for on-platform purchases is about 2% lower than is obtained under the baseline recommendations. Applying the recommendations from the $\beta_i = \beta$ specification to the baseline model actually increases consumption utility, since consumers are encouraged to explore more than under the baseline recommendations. These recommendations are sub-optimal from the platform's perspective because they reduce the number of on platform purchases (and thus the objective given by equation 20).[40]

[39] This is similar to the intuition behind the familiar "you might also like" suggestions on Amazon.com, and more generally the collaborative filtering algorithms that are widely used in online platforms(Schafer, Frankowski, Herlocker and Sen 2007).

[40] Although these utility differences appear small, it is not possible to quantify how important they are for the platform, as we do not model the platform's dynamic profit maximization problem directly but rather assume an objective function. If small utility differences lead to significant long run platform switching, then ignoring learning in recommendations could have significant effects on the present discounted value of platform profit.

The fact that recommendations which encourage more exploration also lead to fewer on-platform purchases highlights the explore-exploit trade-off faced by the platform. Pushing consumers towards high ex-ante expected utility products is more likely to generate an immediate sale at the cost of allowing the consumer to explore and learn about their preferences, while helping consumers explore by putting more weight on $DistU_{ijt}$ can generate better matches but will lead to more consumers leaving the platform in expectation. This suggests that a platform's optimal recommendations will also depend on the level of competition between platforms.

We illustrate this effect in the fourth and fifth columns in Table 6. The fourth column presents results from a model in which we multiply the estimated platform switching cost by 10, limiting platform switching so that each platform is effectively a monopolist. In this case, optimal recommendations put more weight on $DistU_{ijt}$ compared to the baseline. The platform encourages consumers to explore because they are unlikely to switch to another platform. The fifth column presents results with $c_{switch} = 0$, so that platform switching is free. In this specification the coefficients are similar to those in the no learning specification: the coefficient on $DistU_{ijt}$ is positive and the coefficient on $DistS_{ijt}$ is negative and large. In this case it is optimal for the platform to direct consumers towards high expected utility products and away from unexplored parts of the space. The high probability of platform switching means that the platform is incentivised to prioritize immediate sales rather than encourage exploration. This has the surprising implication that increasing competition between platforms can result in recommendations that *reduce* consumer welfare. More competition generates optimal recommendations that discourage exploration, resulting in lower consumption utility: applying the $c_{switch} = 0$ recommendations to the baseline model lowers on-platform utility by 2%. This result depends on the model of spatial learning: when $\lambda = 0$ the platform faces no explore-exploit trade off.

8 Conclusion

In this paper, we develop a model of search with spatial learning and investigate its implications for information design in online platforms. Consumers are initially uncertain of the utility yielded by the set of available products, which they learn about through search. Searching a particular product not only provides information about that product, but provides a signal about how much the consumer is likely to value similar products - those that are "nearby" in product attribute space. Learning induces path dependence: the decision of *which* product to search next depends on past observations. We establish some simple comparative statics on the consumer's "search ranking" of products under a one period look ahead assumption that formalize this intuition.

We argue that our model is identified by data on sequential search paths. The use of search *sequences* to identify covariance of utility across products is novel, and contributes to the broader literature on demand estimation. In particular, our model is based on a random utility specification that allows for a richer cross-product covariance structure than standard random coefficients models. This suggests that incorporating search sequence data, together with variation in prices induced by instrumental variables, could result in the estimation of more flexible cross-product elasticities. This is beyond the scope of this article, but a potentially valuable direction for future research as sequential search data in online retail becomes more widespread.

The path dependence induced by spatial learning has important implications for the role of search intermediaries such as online retail platforms. We use simulations to show that platforms can exploit spatial learning using product recommendations of idiosyncratically high or low payoff products to divert search towards or away from regions of the search space. We then show that consumer-optimal recommendations depend on the extent of learning and cross-platform competition. Our findings suggest that studies of online search and the design and analysis of recommendation systems should account for cross-product learning, and our methodology provides a framework for future research. One possible direction is combining our model with a supply model to investigate how spatial learning affects price setting and product positioning in equilibrium.

References

Adam, Klaus, "Learning While Searching for the Best Alternative," *Journal of Economic Theory*, 2001, *101* (1), 252–280.

Andersen, T.G., R.A. Davis, J.P. Kreiß, and T.V. Mikosch, *Handbook of Financial Time Series* Mathematics and Statistics, Springer Berlin Heidelberg, 2009.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, jul 1995, *63* (4), 841.

Bertsekas, Dimitri P. and John N. Tsitsiklis, *Neuro-Dynamic Programming*, Belmont, MA: Athena Scientific, 1996.

Branco, Fernando, Monic Sun, and J. Miguel Villas-Boas, "Optimal search for product information," *Management Science*, nov 2012, *58* (11), 2037–2056.

____, ____, and ____, "Too much information? Information provision and search costs," *Marketing Science*, jul 2016, *35* (4), 605–618.

Bronnenberg, Bart J, Jun B Kim, and Carl F Mela, "Zooming In on Choice: How Do Consumers Search for Cameras Online?," *Marketing Science*, 2016, *35* (5), 693–712.

Callander, Steven, "Searching and Learning by Trial and Error," *American Economic Review*, October 2011, *101* (6), 2277–2308.

____ and Patrick Hummel, "Preemptive Policy Experimentation," *Econometrica*, 2014, *82* (4), 1509–1528.

Covert, Thomas R., "Experiential and Social Learning in Firms: the Case of Hydraulic Fracturing in the Bakken Shale," Working Paper 2015.

Crawford, Gregory and Matthew Shum, "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, 2005, *73* (4), 1137–1173.

Cressie, Noel, "Statistics for Spatial Data," *Terra Nova*, 1992, *4* (5), 613–617.

De Los Santos, Babur, Ali Hortaçsu, and Matthijs R Wildenbeest, "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior," *American Economic Review*, may 2012, *102* (6), 2955–2980.

____, ____, and Matthijs R. Wildenbeest, "Search With Learning for Differentiated Products: Evidence from E-Commerce," *Journal of Business Economic Statistics*, oct 2017, *35* (4), 626–641.

____ and Sergei Koulayev, "Optimizing click-through in online rankings with endogenous search refinement," *Marketing Science*, aug 2017, *36* (4), 542–564.

Dickstein, Michael, "Efficient Provision of Experience Goods: Evidence from Antidepressant Choice," Working Paper, Harvard University 2018.

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan, "Consumer price search and platform design in internet commerce," *American Economic Review*, jul 2018, *108* (7), 1820–1859.

Dzyabura, Daria and John R. Hauser, "Recommending Products When Consumers Learn Their Preference Weights," *Marketing Science*, 2019, *38* (3), 417–441.

Ellison, Glenn and Sara Fisher Ellison, "Search, Obfuscation, and Price Elasticities on the Internet," *Econometrica*, mar 2009, *77* (2), 427–452.

Fradkin, Andrey, "Search, Matching, and the Role of Digital Marketplace Design in Enabling Trade: Evidence from Airbnb," *SSRN Electronic Journal*, apr 2018.

Frazier, Peter, Warren Powell, and Savas Dayanik, "The knowledge-gradient policy for correlated normal beliefs," *INFORMS Journal on Computing*, sep 2009, *21* (4), 599–613.

Gardete, Pedro and Megan Hunter, "Avoiding Lemons in Search of Peaches: Designing Information Provision," *SSRN Electronic Journal*, jun 2018.

Hagiu, Andrei and Bruno Jullien, "Why do intermediaries divert search?," Technical Report 2 2011.

Hodgson, Charles, "Information Externalities, Free Riding, and Optimal Exploration in the UK Oil Industry," Working Paper 2019.

Honka, Elisabeth and Pradeep Chintagunta, "Simultaneous or Sequential? Search Strategies in the U.S. Auto Insurance Industry," *Marketing Science*, 2017, *36* (1), 21–42.

Jindal, Pranav and Anocha Aribarg, "The Importance of Price Beliefs in Consumer Search," *Journal of Marketing Research*, 2021, *58* (2), 321–342.

Ke, T. Tony and J. Miguel Villas-Boas, "Optimal learning before choice," *Journal of Economic Theory*, mar 2019, *180*, 383–437.

Keane, Michael P., "Modeling heterogeneity and state dependence in consumer choice behavior," *Journal of Business and Economic Statistics*, 1997, *15* (3), 310–327.

—— **and Kenneth I. Wolpin**, "The Career Decisions of Young Men," *Journal of Political Economy*, 1997, *105* (3), 473–522.

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

Korganbekova, Malika and Cole Zuber, "Balancing User Privacy and Personalization," Working Paper 2023.

Koulayev, Sergei, "Search With Dirichlet Priors: Estimation and Implications for Consumer Demand," *Journal of Business Economic Statistics*, apr 2013, *31* (2), 226–239.

——, "Search for differentiated products: identification and estimation," *RAND Journal of Economics*, sep 2014, *45* (3), 553–575.

Mattioli, Dana, "Amazon Changed Search Algorithm in Ways That Boost Its Own Products," sep 2019.

McCall, J J, "Economics of Information and Job Search," *The Quarterly Journal of Economics*, 1970, *84* (1), 113–126.

Osborne, Michael A., Roman Garnett, and Stephen J. Roberts, "Gaussian processes for global optimization," *3rd International Conference on Learning and Intelligent Optimization (LION3)*, 2009, pp. 1–15.

Park, Sejun, Chulhee Yun, Jaeho Lee, and Jinwoo Shin, "Minimum Width for Universal Approximation," *CoRR*, 2020, *abs/2006.08859.*

Rasmussen, Carl Edward and Christopher K I Williams, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, The MIT Press, 2005.

Rothschild, Michael, "Searching for the Lowest Price When the Distribution of Prices Is Unknown," *Journal of Political Economy*, 1974, *82* (4), 689–711.

Schafer, J. Ben, Dan Frankowski, Jon Herlocker, and Shilad Sen, *Collaborative Filtering Recommender Systems*, Berlin, Heidelberg: Springer Berlin Heidelberg,

Schmidt-Hieber, Johannes, "Nonparametric regression using deep neural networks with ReLU activation function," *The Annals of Statistics*, 2020, *48* (4), 1875 – 1897.

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

____, Qianyun Zhang, and Elisabeth Honka, "Search Gaps and Consumer Fatigue," *Marketing Science*, 2023, *42* (1), 110–136.

____, Qingliang Wang, and Pradeep K. Chintagunta, "Search duration," *Marketing Science*, sep 2020, *39* (5), 849–871.

Weitzman, Martin, "Optimal Search for the Best Alternative," *Econometrica*, 1979, *47* (3), 641–654.

Xu, Yanran, Haoyang Zhang, and Haoxuan Li, "Gaussian Process Regression for Forecasting of Boston Housing Prices," 2021, pp. 1191–1200.

Appendix - For Online Publication

A. Proofs

A.1 Optimality of Myopic Search Rule

The myopic consumer's expected value at state S is,

$$V(S) = \max\left\{\hat{u}, \max_j \left(E\left[\max(\hat{u}, u_j)|S\right] - c_j\right)\right\}$$

Note that there is no continuation value on the right hand side. The myopic consumer compares the vlue of stopping, $\hat{u}$, to the maximum over all remaining products of the expeted value of searching once more and then stopping, $\max_j \left(E\left[\max(\hat{u}, u_j)|S\right] - c_j\right)$. Let $x_j \sim N(\mu_j, s_j^2)$ be the posterior distribution of $x_j$at state S. Let $a_j = \frac{\hat{u}-\mu_j}{s_j}$. Define

$$\begin{aligned}
z_j(\hat{u}) &= E\left[\max(\hat{u}, u_j)|S\right] - c_j \\
&= \int_{-\infty}^{\hat{u}} \frac{\hat{u}}{s_j}\phi\left(\frac{x-\mu_j}{s_j}\right)dx + \int_{\hat{u}}^{\infty} \frac{x}{s_j}\phi\left(\frac{x-\mu_j}{s_j}\right)dx - c_j \\
&\int_{-\infty}^{\hat{u}} \frac{\hat{u}}{s_j}\phi\left(a_j\right)dx + \int_{\hat{u}}^{\infty} \frac{x}{s_j}\phi\left(a_j\right)dx - c_j \\
&= \hat{u}\Phi\left(a_j\right) + \left(1-\Phi\left(a_j\right)\right)\int_{\hat{u}}^{\infty} \frac{\frac{x}{s_j}\phi\left(a_j\right)}{\left(1-\Phi\left(a_j\right)\right)}dx - c_j \\
&= \hat{u}\Phi\left(a_j\right) + \left(1-\Phi\left(a_j\right)\right)\mu_j + \phi\left(a_j\right)s_j - c_j
\end{aligned}$$

Where $\Phi(\cdot)$ and $\phi(\cdot)$ are the CDF and PDF of the standard normal distributionand the final line uses the expected value of the truncated normal.

It therefore follows that, given state S, the myopic consumer's optimal policy is to stop if $\hat{u} > \max_j\left(z_j - c_j\right)$, and otherwise search the product with the highest value of z_j.

A.2 Equivalence to Weitzman

Suppose that $\kappa(X, X') = 0$ for $X \neq X'$, which is the case in the limit as $\rho \to 0$. We can then write utility as,

$$\begin{aligned}
u_{ij} &= m_i(X_j) + \xi_j + \epsilon_{ij} \\
&= \mu_i(X) + \delta_{ij}
\end{aligned}$$

Where $\delta_{ij} = m_i(X_j) - \mu(X) + \xi_j + \epsilon_{ij}$ is iid normal over products for a speific consumer, i. This is a special case of the model in Weitzman (1979), in which $u_j \sim F_j$ *independently* across alternatives j.

A.3 Proof of Proposition 1

We take the derivatives of the score z_j in turn:

$$\begin{aligned}
\frac{\partial z_j}{\partial \hat{u}} &= (\phi(a_j)/s_j)\hat{u} + \Phi(a_j) - (\phi(a_j)/s_j)\mu_j + \phi'(a_j)s_j/s_j \\
&= \phi(a_j)(\hat{u} - \mu_j)/s_j + \Phi(a_j) + \phi'(a_j) \\
&= \phi(a_j)a_j + \Phi(a_j) - a_j\phi(a_j) \\
&= \Phi(a_j)
\end{aligned}$$

where on the third line we use the fact that $\phi'(x) = -x\phi(x)$. Similarly:

$$\begin{aligned}
\frac{\partial z_j}{\partial \mu_j} &= -(\phi(a_j)/s_j)\hat{u} + (1 - \Phi(a_j)) + (\phi(a_j)/s_j)\mu_j - \phi'(a_j) \\
&= 1 - \Phi(a_j)
\end{aligned}$$

The partial on costs is immediate. Finally, notice that from the transition equations (2) and (3) the last observation's payoff only influences mean beliefs and potentially the current highest payoff $\hat{u}$ (if it was better than the prior best option). Applying the chain rule, we can use the partial derivatives derived above, along with the derivative $\frac{\partial \mu_j}{\partial u_k}$ from (2) to get the result.

B. Search Rankings and Manipulation of Beliefs in the Illustrative Example

In the example in Section 2, we assumed equal costs of searching all products. However, it is well documented (for example, see Ursu (2018)) that the ranking or salience of products on online platforms affects the order in which consumers search through those products. We thus allow for different search costs, where higher-ranked products have lower search costs. The direct effect of this is to ensure that higher-ranked products are more attractive to search. But under spatial learning, there are also spillover effects: what consumers learn from searching a highly-ranked product can affect consumers' beliefs about other products. *Product rankings can therefore be used*

Figure A.1: Belief Updating

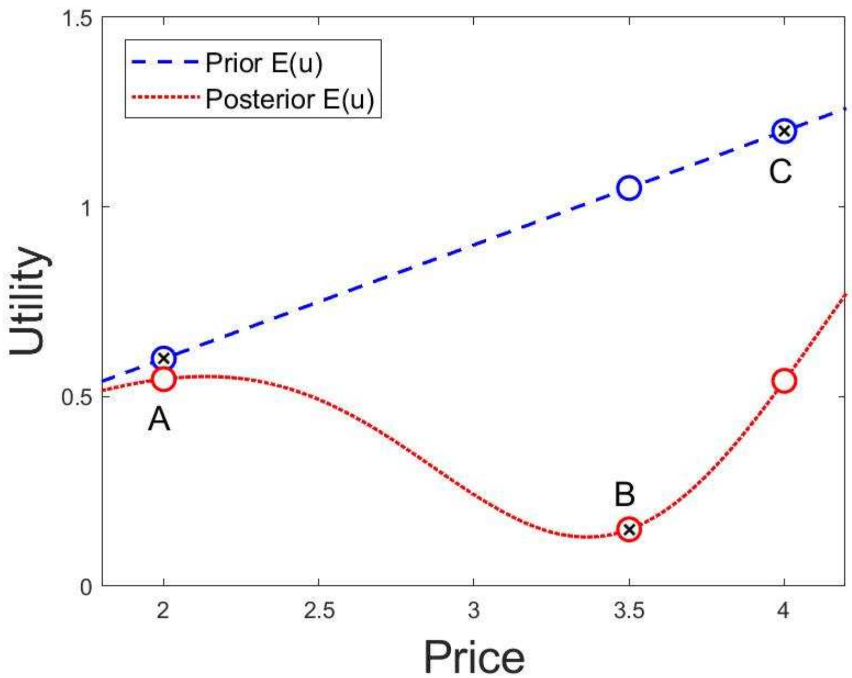

Notes: Black crosses indicate the location of the three products, A, B,and C, in price-utility space. The blue dashed line is the consumer's prior expected utility of hypothetical products at different price levels. This is given by $E(u_j) = (\mu - 1)p_j$. The red dashed line indicates the consumer's posterior expected utility at different price levels after searching product B. The posterior is computed using the Bayesian updating rule described in the text. Parameters are as described in the notes to Figure 1.

to manipulate both search costs and beliefs.

To show the ways in which rankings can be used to change purchase behavior, we modify our example from before by setting $p_B = 3.5$ so that product B is closer in price space to product C than product A. We set the search cost for product B to zero, so that searching it is free — and therefore it is optimal to search B first (this is an attempt to model it being heavily promoted by the platform). Last we assume that the latent payoff for B is $u_B \approx 0.2$, much worse that expected.

Figure A.1 illustrates how the consumer's beliefs about u_A and u_C are updated after she searches product B. This bad initial experience drags down the posterior mean beliefs about C more than product A, so that after the "free" search of B, the consumer believes that A is a better option.

This "belief manipulation" can be effective in driving consumers towards a desired option. Suppose for example that the search intermediary wants the consumer to buy product A, perhaps because it earns the highest commission on sales from that seller or because it is a "house brand". Intuitively, one might expect that the best the intermediary can do is to promote product A, driving its search cost to zero and ensuring it enters the consumer's consideration set. Yet it turns out the answer is more subtle and depends on the search costs.

Table A.1 records the consumer's purchased product as a function of the product they are shown first, and the search cost, c. For low search costs ($c < 0.05$), the consumer will search every product and ultimately purchase C, the highest utility product. In this search cost regime, the platform cannot control the purchase outcome. On the other hand, for very high search costs ($c > 0.91$), the consumer will not search beyond the product initially shown to them by the platform. The platform has complete control over the purchase decision, and therefore should show product A first so that it is purchased. The surprise is that in intermediate cases ($c \in (0.05, 0.78]$), the platform can achieve its aim of getting product A purchased only by showing *product* B first. If the consumer views either A or C first, the observed utility will be equal to the prior expected utility, and the consumer will not update their expectation about the other products. Thus, if the consumer is shown product A first, she will search product C second, since $E(u_C|u_A) = (\mu - 1)p_C > (\mu - 1)p_B = E(u_B|u_A)$. After viewing C she will purchase it. However, if she is shown the inferior product B first she will infer that product C is likely to be low quality since it is close to product B in price space, and will therefore search product A second. With intermediate search costs, it is then optimal to stop and purchase product A.

Notice that this "intermediate case" is likely to be the most prevalent in practice, since we think of platforms as having some but not perfect control over what is purchased on their sites. They also often have considerable prior data on purchasing decisions which may allow them to predict with high accuracy which products are "surprisingly bad" and may therefore be used to steer consumers in this way (we ourselves do such prediction using a relatively small Comscore dataset later in the paper). So belief manipulation is a realistic possibility, depending on the motivation and sophistication of the search intermediary.

Table A.1: Purchase as a Function of Starting Product and Search Cost

Starting Product	c$\in [0, 0.05]$	c$\in (0.05, 0.78]$	c$\in (0.78, 0.91]$	c$\in (0.91, \infty]$
A	C	C	C	A
B	C	A	B	B
C	C	C	C	C

Notes: Each cell records the product purchased by a consumer with search cost c given by the column headers who is shown the starting product indicated by the first column before starting to search. Parameters are as described in the notes to Figure 1.

C. Additional Descriptive Statistics

C.1 Narrowing of Search

Figure 3 shows that consumers search a wider variety of products early in the search path than later in the search path. The left panel of Appendix Figure A.2, which also replicates a finding from BKM, shows that consumers are not only getting closer to the purchased product in attribute space, but are focusing on smaller areas of the attribute space as search progresses. This narrowing of search is illustrated by plotting the distribution of prices searched in each decile of the search path, where the tth search of a search path of length T is in search decile d if $\frac{d-1}{T} < \frac{t}{T} \leq \frac{d}{T}$. Prices are normalized by taking the difference in log price from the price of the product eventually purchased. The figure shows that the distributions of prices searched in the first search deciles are more spread out than in later deciles. For example, the interquartile range in normalized log price is 2.62 for the 1st decile and 1.83 for the 10th decile.

The right panel of Appendix Figure A.2 supports the finding that consumers gradually narrow the scope of their search. The y-axis records the average "step size" in log price. For example, a consumer's nth search has a step size in price of $\Delta price_t = |price_t - price_{t-1}|$ where $price_t$ is the price of the consumer's tth searched product. The x-axis records search percentile, as in Figure 3. The results indicate that step size is declining. For example, in early search the average step size in price is around 60% of the cross product standard deviation in log price, falling to less than 50% by the end of the search path.[41]

[41] This pattern is documented for other product attributes in Appendix Figure A.7. These step size patterns are not documented by BKM.

Figure A.2: Narrowing of Search

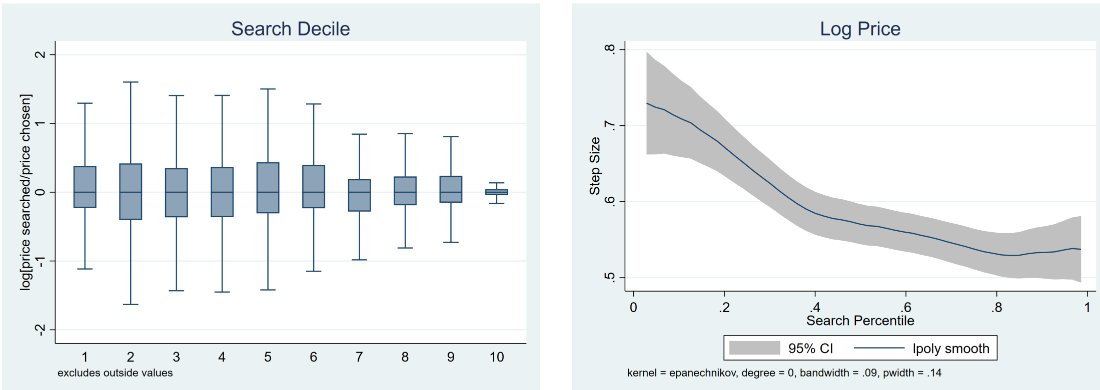

Notes: The left panel displays box plots that record the distribution of the log difference in searched price from the price of the product ultimately purchased, for each search decile as defined in the paper. The Box records the 25th, 50th, and 75th percentiles of the distribution and the whiskers record the upper and lower adjacent values. The y-axis of the right panel records the absolute distance in standard deviations of log price between the product searched and the previous product searched. The x-axis reports the search percentile, as defined in the text. The solid line is a kernel regression using an Epanechnikov kernel, and the shaded area is 95% confidence interval. For both panels, the estimation sample includes all search paths from the ComScore data on search for digital cameras, dropping revisits to the same camera and excluding consumers who do not make a purchase.

C.2 Estimation of $\hat{\theta}$

The index $\hat{\theta}_j$ for each product j is constructed as follows. Let J_i be the set of products that are searched by consumer i. We find the values $\tilde{\theta}_j$ that maximize the likelihood of observed purchases when the probability that consumer i purchases product $j \in J_i$ is given by:

$$P_{ij} = \frac{exp(\tilde{\theta}_j)}{1 + \sum_{k \in J_i} exp(\tilde{\theta}_k)} \tag{21}$$

$\tilde{\theta}_j$ is an index that measures the probability of purchase conditional on search. Note that this is *not* a structural object but a convenient statistical device for classifying products, and that equation 21 is not derived from the model.[42] We use OLS to decompose $\tilde{\theta}_j$ into part that can be explained by product attributes and a residual:

$$\tilde{\theta}_j = X_j \gamma + \theta_j \tag{22}$$

The estimated residuals, $\hat{\theta}_j$, are our measure of how much more or less likely product j is to be purchased relative to products with similar attributes X_j. High

[42]Some objects are never purchased, we omit these objects from J_i and do not construct an index $\hat{\theta}_j$ for them. They are omitted from the regressions in Table 2.

values of $\hat{\theta}_j$ mean that a product is purchased more, conditional on being searched, than similar products. Vice versa for low $\hat{\theta}_j$. In the context of our model, variation in $\hat{\theta}_j$ across products is explained by variation in product effects, ξ_j. Products that are purchased more frequently that others with similar observable attributes must have higher unobservable utility across consumers.

D. Learning with Multiple Retailers

The Bayesian updating rules in equations 2 and 3 apply when correlation across alternatives is determined only by product chatacteristics X_j. In the empirical application, there is additional correlation across alternatives j because two alternatives (j, j') may represent the same product at two different retailers. That is $\tilde{j}_j = \tilde{j}_{j'}$ and $r_j \neq r_{j'}$. In this case, it is useful to define

$$\tilde{m}_i(X_j) = m_i(X_j) + \xi_j, \tag{23}$$

which is distributed as a Gaussian process with mean $\tilde{\mu}(X) = \mu(X)$ and covariance function

$$\tilde{\kappa}(X_j, X_{j'}) = \kappa(X, X') + \sigma_\xi^2 1(\tilde{j}_j = \tilde{j}_{j'}). \tag{24}$$

Consumers update their beliefs according to,

$$\tilde{\mu}'(X) = \tilde{\mu}(X) + \frac{\tilde{\kappa}(X, X_j)(u_j - \tilde{\mu}(X_j))}{\tilde{\kappa}(X_j, X_j) + \sigma_\epsilon^2} \tag{25}$$

$$\tilde{\kappa}'(X, X') = \tilde{\kappa}(X, X') - \frac{\tilde{\kappa}(X, X_j)\tilde{\kappa}(X_j, X')}{\tilde{\kappa}(X_j, X_j) + \sigma_\epsilon^2}. \tag{26}$$

Notice that the σ_ξ terms in the denominator of equations 2 and 3 do not appear in equations 25 and 26 , becuase the variance of the product effects, ξ_j, are incorporated in $\tilde{\kappa}(X_j, X_{j'})$.

E. Estimation Algorithm

E.1 Overview of Value Function Approximation

In this appendix we describe in detail the estimation algorithm outlined in Sections 5.3 and 5.4 above. First, we provide an overview of the value function approximation procedure.

To restate the problem, observe that the consumer's state variable $\mathcal{S}$ includes current beliefs, which are described by a vector of J means with elements $\mu(X_j)$ and a $J \times J$ covariance matrix with $(j, j^{'})$ element $\kappa(X_j, X_{j'})$. The state variable therefore has dimension at least $J + \frac{1}{2}J(J-1)$,with $J \approx 1000$ alternatives in the data.

To make progress, we represent the consumer's state more efficiently by defining an augmented state variable,

$$\tilde{\mathcal{S}}_{i\tau} = \big((u_{j(i,t)})_{t<\tau}, (X_{j(i,t)})_{t<\tau}, (r_{j(i,t)})_{t<\tau}, \psi_i, \tilde{r}, c, c_{switch}\big) \tag{27}$$
$$\psi_i = (\beta_i, \alpha, \gamma_{o(j)}, \delta_{b(j)}, \lambda, \boldsymbol{\rho}, \sigma_\xi, \sigma_\epsilon).$$

For a consumer i who has searched τ alternatives so far. The first three elements of the state variable are a list of τ utilities, τ locations X_j in attribute space, and τ retailers of the products that have already been searched. The fourth element, ψ_i, is a list of parameters that control the distribution of utilities for consumer i, including the consumer-specific random effects β_i. Together, these three elements of the state variable uniquely define consumer beliefs $\mu(X_j)$ and $\kappa(X_j, X_{j'})$ through the prior mean and covariance functions given by equations 6 and 8 and the Bayesian updating rules in equations 2 and 3. Note however, that $\tilde{\mathcal{S}}_{i\tau}$ has dimensionality of only $6\tau + 19$, which is several orders of magnitude smaller than $\mathcal{S}$ even for the longest observed search path.

$\tilde{S}_{i\tau}$ contains all the information relevant for the consumer's problem, so we can write we can write the value function as $V(\tilde{S}_{i\tau})$ rather than $V(\mathcal{S}, \psi)$. Notice in particular that, conditional on $\tilde{S}_{i\tau}$ the consumer's problem does not depend on the distribution of β_i, defined by parameters (β, Ω), or the draws of the product effects, ξ_j. (β, Ω) enter the likelihood through the outer integral in equation 13. ξ_j affect the means $\bar{u}_i$ of the inner integral in equation 13.

We solve the consumer's problem using this lower dimensional representation of the state using a value function iteration and approximation routine. The steps are as follows:

1. Initialize the value function to $\hat{V}_0(\tilde{S}_{i\tau}) = 0 \, \forall \tilde{S}_{i\tau}$

2. Draw W values of $\tilde{\mathcal{S}}_{i\tau}$ from a proposal distribution. Call the set of sampled states $\mathcal{S}_W$

3. Iterate the Bellman equation once at each state $\tilde{\mathcal{S}}_{i\tau} \in \mathcal{S}_W$. On the $k+1$th

iteration,

$$V_{k+1}(\tilde{S}_{i\tau}) = \log\left(exp(\hat{u}) + \sum_j exp\left(E\left[\hat{V}_k(\tilde{S}_{i\tau+1})|\tilde{S}_{i\tau}\right] - c\right)\right). \tag{28}$$

4. Estimate a neural network regression of $V_{k+1}(\tilde{S}_{i\tau})$ on $\tilde{S}_{i\tau}$. Denote the predicted values of the neural network at a generic state $\hat{V}_{k+1}(\tilde{S})$.

5. Return to step 2 and repeat to convergence. Let $\hat{V}(\tilde{S})$ be the final approximation.

The following subsections provide further details.

E.2 Obtain myopic estimates

To obtain initial parameter values, we define myopic choice probabilities,

$$P_i(j|\mathcal{S},\psi) = \frac{exp\left(E[\max\{\hat{u}, u_j|\mathcal{S}] - c\right)}{exp\left(\hat{u}\right) + \sum_{l\in J} exp\left(E[\max\{\hat{u}, u_l|\mathcal{S}] - c\right)}, \tag{29}$$

which correspond to the optimal choice probabilities of consumers who always behave as if they can only search once more before stopping, as in theoretical results in Section 3. The expected value of an alternative j is just the expectation of the max of u_j and $\hat{u}$, with continuation value set to 0. Using these choice probabilities, we find the value of the parameters ψ that maximizes the likelihood 14. Call these first step estimates ψ^1.

E.3 Draw augmented states

Recall the augmented state variable, $\tilde{\mathcal{S}}_{i\tau}$, defined in equation 27. $\tilde{\mathcal{S}}_{i\tau}$ comprises data on the set of products a consumer has viewed up to time τ, $\big((u_{j(i,t)})_{t<\tau}, (X_{j(i,t)})_{t<\tau}, (o_{j(i,t)})_{t<\tau}, \tilde{o}\big)$, and a vector of parameters, (ψ_i, c, c_{switch}). Recall that this augmented state variable contains all the information relevant to the consumer's problem, so we can write the value function as $V(\tilde{S}_{i\tau})$. Because the state space is continuous, it is impossible to evaluate the value function at each possible value of $\tilde{S}_{i\tau}$. Therefore we must generate a sample set of state points, $\mathcal{S}_W$, at which to evaluate $V(\tilde{S}_{i\tau})$.

The first step in generating $\mathcal{S}_W$ is to draw values of the parameters, (ψ_i, c, c_{switch}). Let $\tilde{\psi^1} = \psi^1\backslash\{\boldsymbol{\xi}^1\}$ be the first step estimates of all the parameters except the product effects, $\boldsymbol{\xi} = \{\xi_j\}_{j=1}^J$. We draw n parameter vectors indexed by w, $\tilde{\psi}_w$ from a distribution $\tilde{\psi}_w \sim N\left(\tilde{\psi}^1, \Sigma_\psi\right)$, where Σ_ψ is diagonal. Then, for each sampled parameter vector, $\tilde{\psi}_w$, we draw a value of the individual random coefficients $\beta_{iw} \sim N(\beta_w, \Omega_w)$. Thus, we generate n individual-specific parameter draws indexed by w, $(\psi_{iw}, c_w, c_{switch,w})$, centered around the first step estimates, ψ^1. In practice we simulate $n = 500$ such parameter vectors.

We then simulate search paths of length 100 for each parameter vector. Given some initial value of the function $V(\tilde{S}_{i\tau})$, we can generate search paths using the choice probabilities in equation 11, where we draw the realizations of utilities according to the model. In practice, we add noise to the choice probabilities so that the probability of simulated consumer w searching alternative j is:

$$\tilde{P}_w(j|\mathcal{S}, \psi_w) = \mathit{\Delta}_0 P_w(j|\mathcal{S}, \psi_w) + \Delta_1 \frac{1}{J}, \tag{30}$$

where $\Delta_0 + \Delta_1 = 1$ and we shut down the outside option. This generates search paths with a greater coverage of the state space than the optimal search paths, which follow choice probabilities $P_w(j|\mathcal{S}, \psi_w)$.

These simulated search paths are then recorded as a sequence of observed utilities and product locations, $\left((u_{j(w,t)})_{t<\tau}, (X_{j(w,t)})_{t<\tau}, (r_{j(w,t)})_{t<\tau}, \tilde{r_{wt}}\right)$, for each simulated consumer w, corresponding to the consumer's state after their t'th search.

$\mathcal{S}_W$ is then the collection of these simulated states with the corresponding simulated parameter vectors. In practice, $\mathcal{S}_W$ contains 50,000 simulated states.

E.4 Compute approximation to the value function

To compute the approximated value function, we start by initializing the value function to $\hat{V}_0(\tilde{S}_{i\tau}) = 0 \, \forall \tilde{S}_{i\tau}$. Using this initial value, we draw the set of augmented states, $\mathcal{S}_W$, as described above. For every state point $\tilde{\mathcal{S}}_{w\tau} \in \mathcal{S}_W$, we compute $V_1(\tilde{S}_{i\tau})$ using Bellman equation 28.

Notice that we have iterated the Bellman equation once at the points in $\mathcal{S}_W$. To solve for the value function, we would like to continue iterating the Bellman equation until we achieve convergence. However, on the $k+1$th iteration, the evaluation of the Bellman equation at a state in $\mathcal{S}$ requires that we know $V_k(\mathcal{S}')$ at all states $\mathcal{S}'$ that the agent could reach after one more search, including states that are outside $\mathcal{S}_W$. In order to extrapolate to states outside states outside $\mathcal{S}_W$, we run a neural network regression of $V_1(\tilde{S}_{i\tau})$ on $\tilde{S}_{i\tau}$ for the points in $\mathcal{S}_W$.

Figure A.3: State Representation for Neural Network

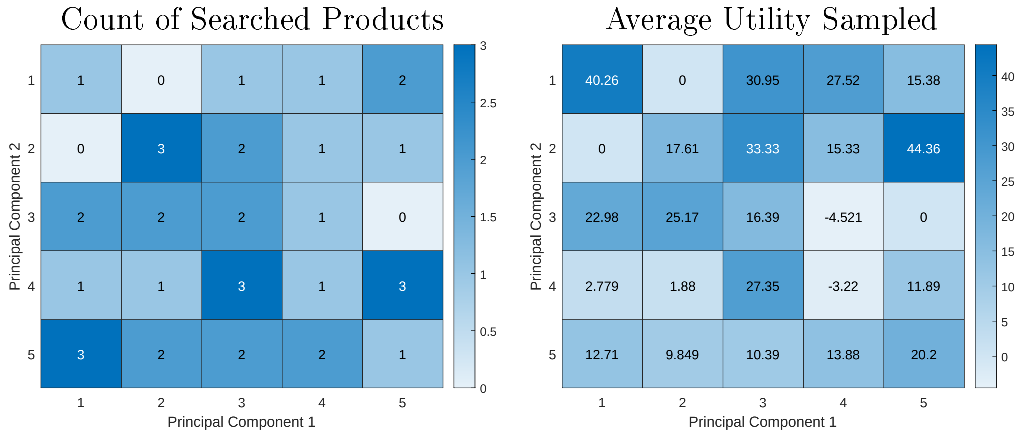

The neural network regression procedure is as follows. First, we further reduce the dimensionality of $\tilde{\mathcal{S}}_{w\tau}$. We reduce the four dimensions (price, zoom, pixels, display) of the location of products in attribute space, $(X_{j(i,t)})_{t<\tau}$, to two dimensions using principal component analysis. We then construct a 5×5 grid over the dimension-reduced product space and for each $\tilde{\mathcal{S}}_{w\tau}$ record the number of products searched in each grid square and the average utility sampled in each grid square. This procedure generates 50 features describing the location and utilities of the products already searched. An example of this representation is illustrated in Figure A.3.

We append to these features the highest utility observed so far, $\hat{u}$, the number of unsearched alternatives at the current outlet and at all other outlets, and the parameter vector $(\psi_{iw}, c_w, c_{switch,w})$. The result is a representation of the state with a total of 221 dimensions (or "features" in the language of neural networks).

We regress $V_1(\tilde{S}_{i\tau})$ on these features of the state using a single layer neural network with a single fully connected layer with 20 neurons and a "rectified linear unit" activation function (see e.g. Schmidt-Hieber (2020)). The estimated function $\hat{V}_1(\tilde{S})$ is an approximation to the value function after one iteration that can be evaluated at any state point. Note that because of the dimension reduction described above, we are are averaging over different states before we estimate the neural network.

We iterate the sampling of $\mathcal{S}_W$, application of the Bellman equation 28, and the neural network approximation until we achieve convergence. We say that convergence has been achieved when $\sum_{\mathcal{S}_W} \left(\hat{V}_k(\tilde{S}_{i\tau}) - \hat{V}_{k-1}(\tilde{S}_{i\tau})\right)^2$ falls below a critical value.

In Figure A.4 below, we plot $\hat{V}_k(\tilde{S}_{i\tau})$ against $\hat{V}_{k-1}(\tilde{S}_{i\tau})$ for the final iteration of this algorithm at the set of sampled state points.

Figure A.4: Convergence of Value Functions

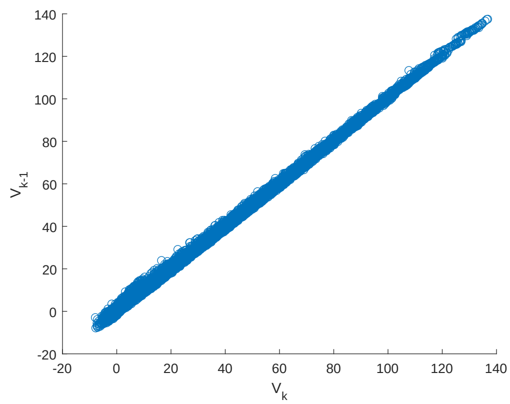

Notes: Scatter plot of $\hat{V}_k(\tilde{S}_{i\tau})$ against $\hat{V}_{k-1}(\tilde{S}_{i\tau})$ over 500,000 sampled state points for the final iteration of the Bellman fucntion iteration.

E.5 Obtain dynamic estimates

The final step is to obtain parameter estimated using the dynamic model. We plug the approximated value function, $\hat{V}_k(\tilde{S}_{i\tau})$ into the choice probabilities given by equation 11, and use these to construct the likelihood of the data. We then find the value of ψ that maximizes the likelihood. Notice that we do not need to recompute $\hat{V}_k(\tilde{S}_{i\tau})$ for each candidate parameter vector, since the consumer's parameter vector is an argument of the approximated value function.

One might be concerned about candidate values of ψ drifting too far away from values in $\mathcal{S}_W$, reducing the accuracy of the approximation. This is why we center the sampling of parameter vectors on ψ^1.We also start the parameter search for the dynamic estimates at ψ^1. In practice, the dynamic estimates do not differ too much from the myopic first step estimates (except for the search cost, c, which is higher when consumers are forward looking). In principle, one could iterate this procedure, repeatedly obtaining dynamic estimates and the re-estimating the approximation to the value function at states centered around the new estimates.

F. Identification Details

F.1 Parameter Interpretation and Price Endogeneity

The standard price endogeneity concern applies here. Prices may be positively correlated with product quality that is unobserved by the econometrician. However, we can still meaningfully interpret the estimated coefficients, β, in light of our model. We assume that consumers have rational beliefs about the distribution of utility, and we *explicitly model* this distribution. β_{price} therefore measures the *net effect on expected utility* of price and any positive correlation between price and unobserved quality, fixing beliefs. That is, $\beta_{price} = \frac{\partial E(u)}{\partial price}$, where the expectation is taken with respect to consumers' prior beliefs.

This interpretation limits the counterfactual exercises we can perform. For instance, we cannot think about price changes. Under counterfactual prices, the estimated consumer beliefs about the relationship between price an expected utility would no longer be correct. To recompute counterfactual rational beliefs we would need to decompose $\beta_{price} = \frac{\partial E(u)}{\partial price}$ into the direct effect of price on utility and the correlation of price with unobserved quality. In order to separately estimate the direct defect of price on utility and the correlation of price with expected quality we would need exogenous variation in prices over which consumers' beliefs about the relationship between price and expected utility can be credibly argued to be held fixed. This would be an useful exercise but is outside the scope of this paper.

This is not an issue for the exercises we perform using the estimated model, since we are interested primarily in the effect of information provision about products on search paths and consumption, *fixing* product locations in attribute space.

F.2 Rational prior restriction

Notice that the identification arguments do not use the assumption that $\xi_j \sim N(0, \sigma_\xi)$. That is, ξ_j and σ_ξ are separately identified without imposing rational beleifs about the distribution of ξ_j on conusmers. In estimation, we impose this additional restriction (see equation 14) which implies that the values of ξ_j contain information about σ_{xi}.

F.3 Alternative covariance function and distance metrics

The identification argument for the variance and covariance parameters, $\{\lambda, \rho, \sigma_\xi, \sigma_\epsilon\}$, relies on the functional form of the covariance function (see equation **??**). In particular, the parameters are identified by varying the distance $(X_A - X_B)^2$ between products, fixing the product effect ξ_A, and varying the product effect, fixing the distance between products.

Note that the structure of the covariance matrix allows this type of identification. If we instead freely parameterized the covariance as $cov(u_A, u_B) = \sigma_{AB}$, the model would likely not be identified. What is crucial is the assumption that $cov(u_A, u_B)$ depends *only* on some notion measure of distance between products, $d(A, B)$, so that $cov(u_A, u_B) = \lambda^2 g(d(A, B))$. It is not essential that $d(A, B)$ is the euclidean norm in attribute space. For example, one could make a similar identification argument using

$$\kappa(X_A, X_B) = \lambda^2 exp\left(\frac{-\min_k \left\{(X_{Ak} - X_{Bk})^2\right\}}{2\rho^2}\right).$$

For a general covariance function $cov(u_A, u_B) = \lambda^2 g(d(A, B))$ the identification argument would require fixing ξ_A and examining the probability of second search at different distances $d(A, B)$ to identify $g(\cdot)$. Note that to obtain non-parameteric identification of $g(\cdot)$ would likely require asymptotics in the number of products. In our case, $g(\cdot)$ depends on a single parameter and thus can be identified from searches over a finite number of products.

F.4 Variances cannot be normalized

In Section 5 we argue that the parameters of the model are identified by data on search sequences. In particular, the variance parameters, σ_ξ, σ_ϵ, and λ are separately identified and cannot be normalized. To provide some additional evidence, we simulate search paths at fixed parameter values and then estimate the model using the simulated data, with and without imposing an additional "normalization" restriction. The data is simulated at the parameters listed in Table A.4, with $\lambda = 100$. We estimate the model under the additional restriction $\lambda = 50$. If one of the variance parameters can be normalized, then the maximum likelihood with this restriction should be equal to the maximum likelihood without it. In Figure A.5 we plot the constrained and unconstrained maximum likelihood for 200 simulated data sets. The points almost all lie below the 45 degree line, suggesting that the restriction reduces

Figure A.5: Comparison of Likelihood with Normalized Variance

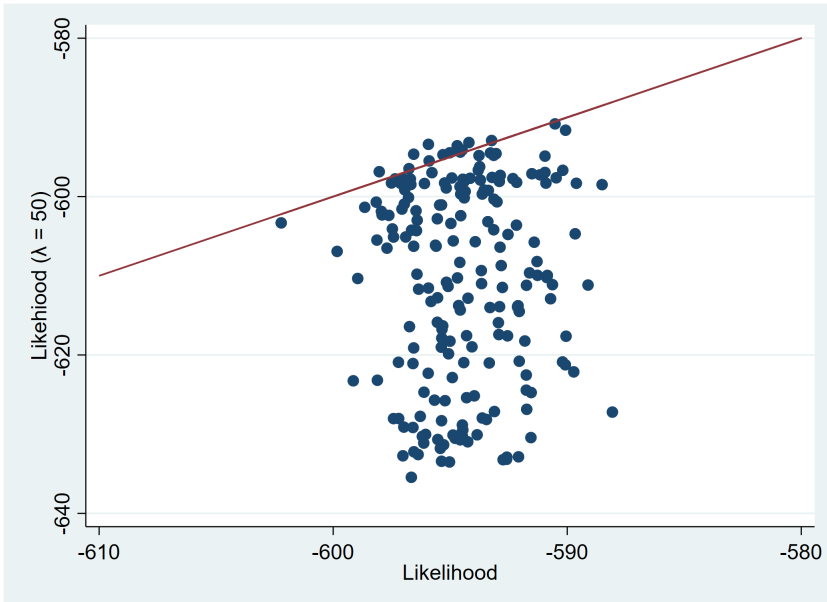

Notes: Scatter plot shows the maximum likelihood obtained applying the estimation procedure to 200 simulated datasets. Data was simulated at the parameters listed in Table A.4. Each simulation contains 1000 search sequences. The x-axis records the unconstrained maximum likelihood. The y-axis records the maximum likelihood under the restriction that $\lambda = 50$. A 45 degree line is plotted in red.

the fit of the model and is therefore not a normalization.

G. Definition of KL Divergence

We compute the *expected* KL divergence for each (j_1, j_2) according to the following equation:

$$KL(j_1, j_2) = \frac{1}{2} \int \left(\left(\Sigma_0^{-1} \Sigma_1 \right) + (\boldsymbol{\mu_0} - \boldsymbol{\mu}_1)' \, \Sigma_0^{-1} \, (\boldsymbol{\mu_0} - \boldsymbol{\mu}_1) + \ln \left(\frac{\det \Sigma_0}{\det \Sigma_1} \right) \right) dF_0(u_{j_1}, u_{j_2}) \tag{31}$$

Where Σ_0 is the prior covariance of all product utilities, $\boldsymbol{\mu_0}$ is the prior mean vector, Σ_1 is the posterior covariance after observing u_{j_1} and u_{j_2}, $\boldsymbol{\mu}_1$ is the posterior mean vector, and $F_0(u_{j_1}, u_{j_2})$ is the prior distribution of the utilities of j_1 and j_2.

Table A.2: Brand and Outlet Shares

Brand	Product Share	Domain	Product Share
Canon	13.41	Amazon.com	30.23
Kodak	11.45	BestBuy.com	12.82
Nikon	14.68	eBay.com	21.04
Sony	11.64	Walmart.com	12.62
Other	48.83	Other	23.29

Notes: Table records the share of the top four brands and retailers among all product-retailer combinations observed in the data.

I. Additional Tables and Figures

Figure A.6: Covergence to Chosen Attribute Level

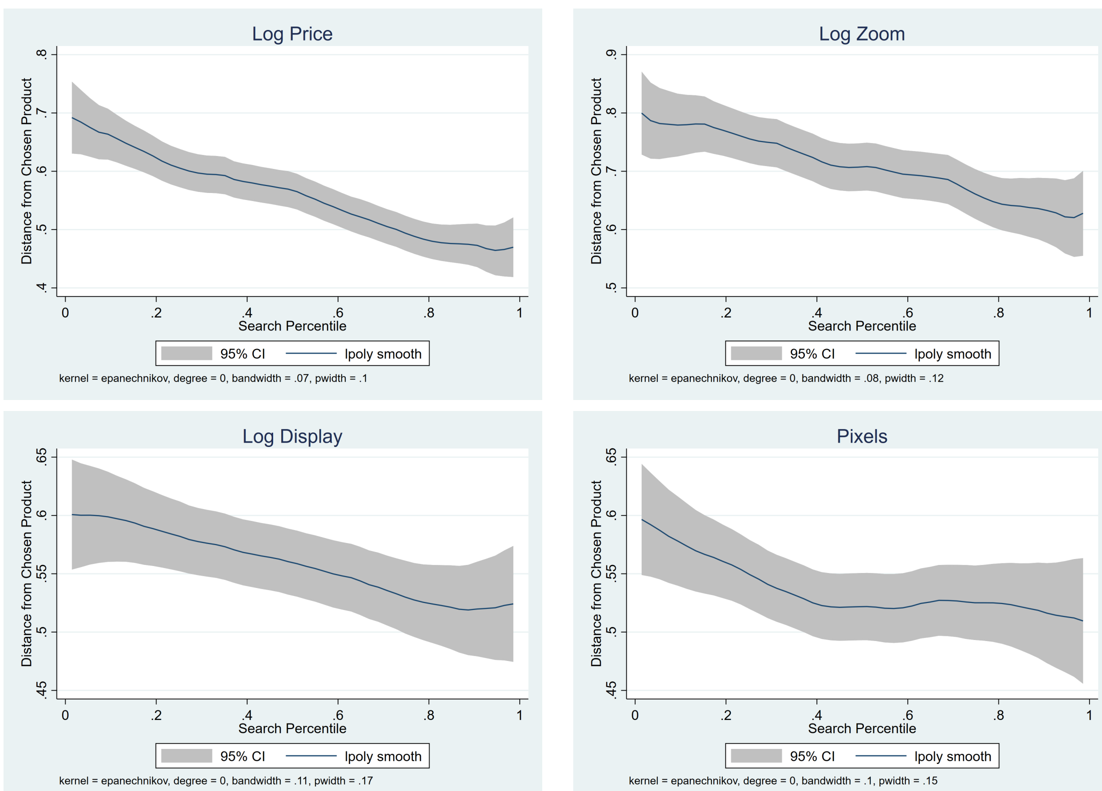

Notes: The y-axis for each panel records, for the relevant product attribute, the absolute difference in standard deviations of the attribute between the searched product and the product ultimately purchased. The x-axis reports the search percentile, as defined in the text. The product ultimately purchased is excluded from the data for each consumer. The solid line is a kernel regression using an Epanechnikov kernel, and the shaded area is 95% confidence interval. The estimation sample includes all search paths from the ComScore data on search for digital cameras, with revisits to the same product dropped.

Figure A.7: Covergence in Step Size

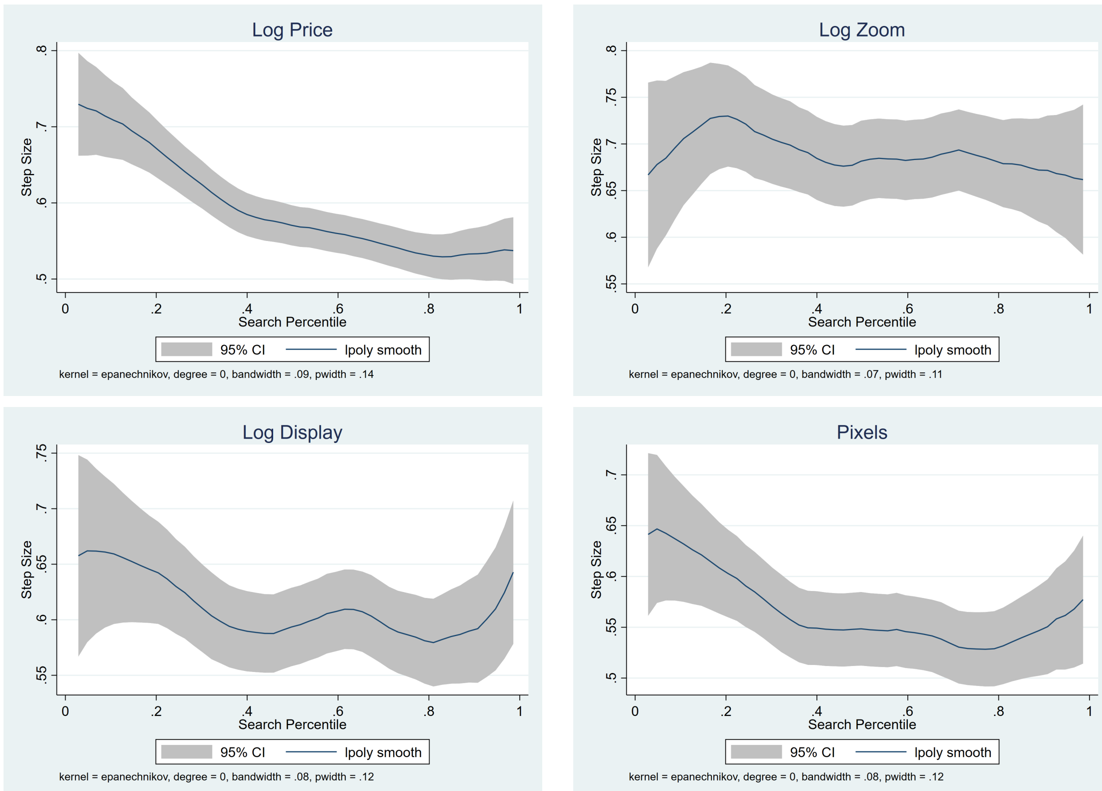

Notes:The y-axis of the each panel records the absolute distance in standard deviations of relevant attribute between the product searched and the previous product searched. The x-axis reports the search percentile, as defined in the text. The solid line is a kernel regression using an Epanechnikov kernel, and the shaded area is 95% confidence interval. For both panels, the estimation sample includes all search paths from the ComScore data on search for digital cameras, including revisits to the same camera and excluding consumers who do not make a purchase.

Table A.3: Effect of Search Percentile and Retailer Switching on Distance to Chosen Product

Distance to chosen product in:	Log Price	Pixel	Log Zoom	Long Display
$\Delta domain_{it}$	-0.052***	-0.034**	-0.068***	-0.025
	(0.014)	(0.016)	(0.020)	(0.0170
$SearchPercentile_{it}$	-0.487***	-0.374***	-0.543***	-0.409***
	(0.022)	(0.025)	(0.031)	(0.027)
N	6407	6407	6407	6407
Consumer FE	Yes	Yes	Yes	Yes

Notes: The dependent variable in each regression is the absolute distance in standard deviations of relevant attribute between the tth product searched and the product ultimately purchased. $\Delta domain_{it}$ is an idicator that is equal to 1 if the consumer switched domain between the $t-1$th and the tth search. The data includes all search paths in which at least two products are searched. *** indicates significance at the 99% level. ** indicates significance at the 95% level. * indicates significance at the 90% level.

Table A.4: Monte Carlo Exercise

	True Parameter	$N = 1000$		True Parameter	$N = 1000$
β_1	-5	-4.997	c_{switch}	1	0.842
		(0.256)			(0.140)
β_2	5	5.023	c	7	7.070
		(0.253)			(0.451)
ω_1	1	1.081	λ	100	99.926
		(0.204)			(3.494)
ω_2	2	1.962	ρ_1	0.5	0.490
		(0.237)			(0.085)
α	-5	-6.095	ρ_2	0.5	0.505
		(2.0)			(0.065)
γ_1	5	4.842	σ_ξ	5	4.459
		(0.643)			(0.863)
			σ_ϵ^2	20	22.341
					(5.388)

Notes: Table reports the mean and standard deviation of the estimated parameters across 150 Monte Carlo replications. For each replication, N search paths are simulated, fixing the parameters are the values reported in the "True Parameter" column, and fixing X_j and ξ_j for $J = 20$ products at values as described in the text. We draw the locations of 20 products in a two dimensional attribute space where attribute k is distributed $X_j^k \sim N(0,1)$. For each product we then draw product effects according to $\xi_j \sim N(0, \sigma_\xi)$. We repeat the search path simulation and parameter estimation 150 times, fixing the product characteristics and parameters over these iterations.

Figure A.8: Covergence to Chosen Attribute Level: Simulations

Baseline Parameter Estimates

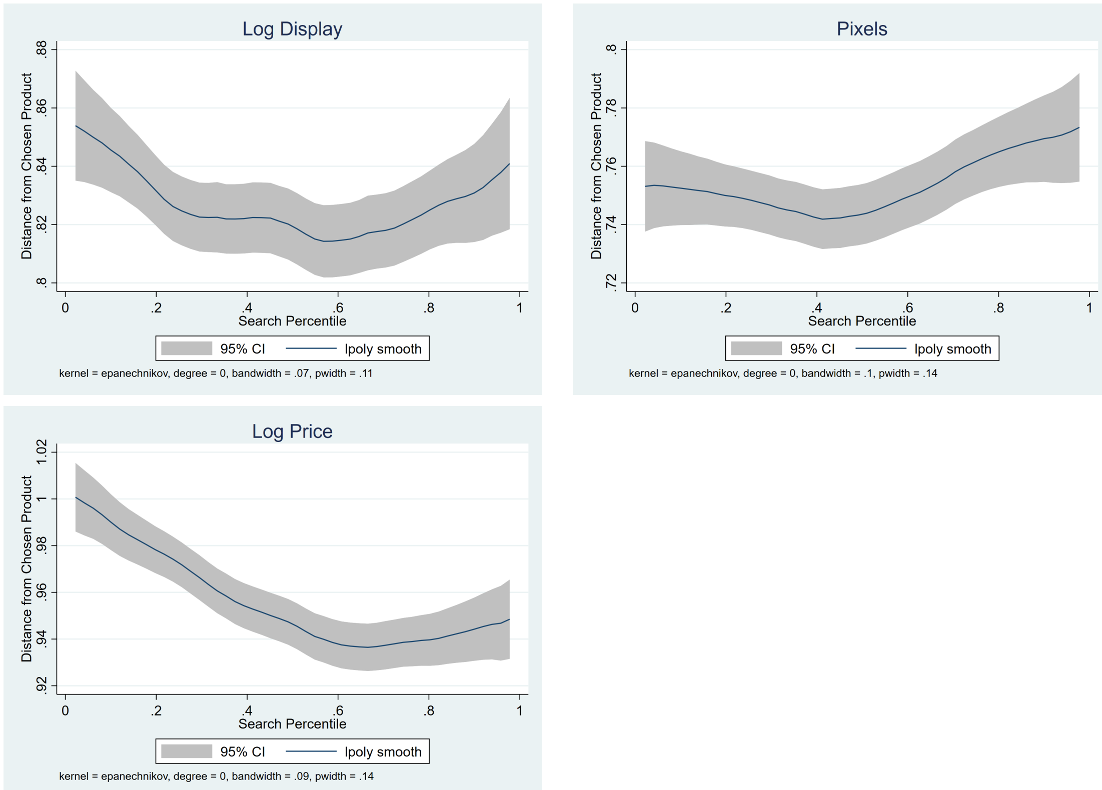

Notes: Figures are constructed using 5,000 search paths simulated at the estimated parameters. The y-axis records, for the relevant product attribute, the absolute difference in standard deviations between the searched product and the product ultimately purchased. The product ultimately purchased is excluded from the data for each consumer. The x-axis reports the search percentile, as defined in the text. The solid line is a kernel regression using an Epanechnikov kernel, and the shaded area is 95% confidence interval.

Figure A.9: Covergence to Chosen Attribute Level: Simulations

$\lambda = 0$ Parameter Estimates

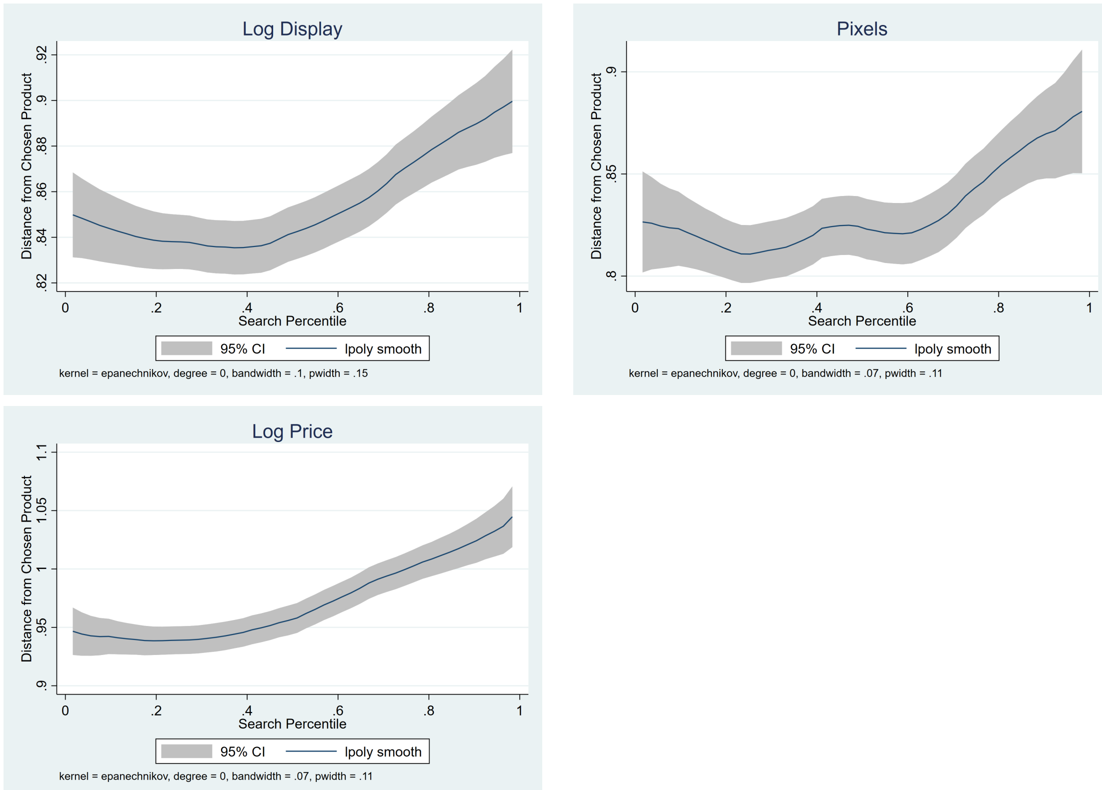

Notes: Figures are constructed using 5,000 search paths simulated at the estimated parameters under the restriction the $\lambda = 0$. The y-axis records, for the relevant product attribute, the absolute difference in standard deviations between the searched product and the product ultimately purchased. The product ultimately purchased is excluded from the data for each consumer. The x-axis reports the search percentile, as defined in the text. The solid line is a kernel regression using an Epanechnikov kernel, and the shaded area is 95% confidence interval.

Table A.5: Effect of Product Residuals on Step Size: Fewer Controls

	$\Delta price_{it}$	$\Delta price_{it}$	$\Delta price_{it}$	$\Delta price_{it}$
$\hat{\theta}_{j(i,t-1)}$	-0.037	-0.010	-0.003	-0.148***
	(0.022)	(0.0253)	(0.025)	(0.035)
$SearchPercentile_{it}$		-0.181***	-0.176***	-0.156***
		(.037)	(0.037)	(0.036)
$Purchased_{it-1}$			-0.094***	-0.122***
			(0.035)	(0.031)
N	3976	3976	3976	3976
Consumer FE	No	No	No	No
Density Controls	No	No	No	Yes

Notes: Table presents regressions of search step size on the product residual index $\hat{\theta}_{j(i,t-1)}$. Step sizes are measured using the absolute difference in standardized log product attributes between the tth and the $t-1$th search. $\hat{\theta}_{j(i,t-1)}$ is constructed as described in the text. Values of $\hat{\theta}_{j(i,t-1)}$ are standardized so that estimated coefficients are the effect of one standard deviation. Any product observations where j_{it-1} is never purchased, and hence a value $\hat{\theta}_{j(i,t-1)}$ is not computed, are omitted form the regression. Other covariates are described in the text. The data includes all search paths in which at least two products are searched. *** indicates significance at the 99% level. ** indicates significance at the 95% level. * indicates significance at the 90% level.

Table A.6: Effect of Rarely Purchased Product on Step Size

	$\Delta price_{it}$	$\Delta pixel_{it}$	$\Delta zoom_{it}$	$\Delta display_{it}$
$BadProduct_{it-1}$	0.052**	0.057**	0.067**	0.046*
	(0.022)	(0.024)	(0.031)	(0.025)
$SearchPercentile_{it}$	-0.140***	-0.021	-0.012	-0.006
	(0.034)	(0.036)	(0.047)	(0.039)
$Purchased_{it-1}$	-0.092***	0.010	-0.101**	-0.026
	(0.035)	(0.038)	(0.049)	(0.040)
N	4697	4697	4697	4697
Density Controls	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	Yes	Yes

Notes: Table presents regressions of search step size on an indicator, $BadProduct_{it-1}$, for whether the last product searched is rarely purchased. $BadProduct_{it-1} = 1$ if product j_{it-1} is searched by at least 5 consumers in the data and purchased with probability less than 10% conditional on being searched. . Step sizes are measured using the absolute difference in standardized log product attributes between the tth and the $t-1$th search. All regressions include controls for search percentile, product density, and consumer fixed effects. The data includes all search paths in which at least two products are searched. *** indicates significance at the 99% level. ** indicates significance at the 95% level. * indicates significance at the 90% level.

Table A.7: Path Dependence: Multi-Step Differences

	Two Steps			
	$\Delta^2 price_{it}$	$\Delta^2 pixel_{it}$	$\Delta^2 zoom_{it}$	$\Delta^2 display_{it}$
$\hat{\theta}_{j(i,t-2)}$	-0.110**	-0.218***	-0.120	-0.097
	(0.053)	(0.055)	(0.075)	(0.060)
N	3348	3348	3348	3348
	Three Steps			
	$\Delta^3 price_{it}$	$\Delta^3 pixel_{it}$	$\Delta^3 zoom_{it}$	$\Delta^3 display_{it}$
$\hat{\theta}_{j(i,t-3)}$	-0.114*	-0.110*	0.080	-0.051
	(0.059)	(0.061)	(0.082)	(0.067)
N	2833	2833	2833	2833

Notes: $\Delta^2 price_{it}$=$|price_{it} - price_{it-2}|$ and $\Delta^3 price_{it}$=$|price_{it} - price_{it-3}|$ Table presents regressions of multi-step differences in product attributes on the product residual index $\hat{\theta}_{j(i,t-2)}$ or $\hat{\theta}_{j(i,t-3)}$. Step sizes are measured using the absolute difference in product attributes between the tth and the $t-1$th search. All product attribute are in logs and standardized.$\hat{\theta}_{j(i,t-1)}$ is constructed as described in the text. Values of $\hat{\theta}_{j(i,t-1)}$ are standardized so that estimated coefficients are the effect of one standard deviation. Any product observations where j_{it-1} is never purchased, and hence a value $\hat{\theta}_{j(i,t-1)}$ is not computed, are omitted form the regression. All regressions include the same covariates as in Table 2, including consumer fixed effects. Two step regressions in the top panel include all search paths wwith at least three products searched. Three step regressions in the lower panel include all search paths with at least four products searched. *** indicates significance at the 99% level. ** indicates significance at the 95% level. * indicates significance at the 90% level.

Table A.8: Placebo Tests: Leads and Lags of Product Residuals

	$\hat{\theta}_{j(i,t-3)}$	$\hat{\theta}_{j(i,t-2)}$	$\hat{\theta}_{j(i,t-1)}$	$\hat{\theta}_{j(i,t)}$	$\theta_{j(i,t+1)}$
$\lvert price_{it} - price_{it-1}\rvert$	-0.012	-0.027	-0.197***	0.019	0.021
	(0.059)	(0.055)	(0.042)	(0.038)	(0.043)
$\lvert pixel_{it} - pixel_{it-1}\rvert$	0.027	0.080	-0.255***	0.049	-0.079
	(0.068)	(0.061)	(0.045)	(0.044)	(0.049)
$\lvert zoom_{it} - zoom_{it-1}\rvert$	0.053	0.141*	-0.118**	0.010	0.026
	(0.084)	(0.078)	(0.060)	(0.052)	(0.059)
$\lvert display_{it} - display_{it-1}\rvert$	-0.058	0.040	-0.244***	0.040	-0.084*
	(0.072)	(0.066)	(0.048)	(0.045)	(0.051)

Notes: Each cell in this table is the coefficient from a regression of step sizes indicated by the row titles on lagged product residuals indicated by the column headers. Regression specifications are otherwise as rcorded in the notes to Table 2. *** indicates significance at the 99% level. ** indicates significance at the 95% level. * indicates significance at the 90% level.

Table A.9: Estimated Parameters with $\lambda = 0$

	Estimate	SE		Estimate	SE
β_1 (log price)	-2.734	0.124	c_{switch}	1.867	0.027
β_2 (log zoom)	-0.280	0.108	γ_1	-1.764	0.220
β_3 (pixels)	2.202	0.095	γ_2	-1.158	0.180
β_4 (display)	0.648	0.103	γ_3	-0.604	0.225
ω_1 (log price)	3.076	0.105	γ_4	-3.170	0.216
ω_1 (log zoom)	3.515	0.113	δ_1	3.567	0.153
ω_1 (pixels)	0.911	0.117	δ_2	2.091	0.142
ω_1 (display)	0.689	0.120	δ_3	4.384	0.138
α	-21.681	0.753	δ_4	2.481	0.155
c	6.376	0.265	σ_ϵ	17.489	0.038

Notes: Table reports estimated parameters under the restriction $\lambda = 0$ and standard errors. Estimation uses the procedure described in Section 5. Note that, as discussed in Section 5, we adopt the normalization $\sigma_\xi = 20$ when $\lambda = 0$. The reported estimates are the mean and standard deviations computed using the observed Fisher information. For more details on the estmiation procedure, see Appendix D.