7-29-2020

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

Charles Hodgson

Gregory Lewis

YOU CAN LEAD A HORSE TO WATER: SPATIAL LEARNING AND PATH DEPENDENCE IN CONSUMER SEARCH

By

Charles Hodgson and Gregory Lewis

July 2020

COWLES FOUNDATION DISCUSSION PAPER NO. 2246

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.yale.edu/

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

Charles Hodgson
Yale University

Gregory Lewis
Microsoft Research

July 29, 2020

Abstract

We develop a model of consumer search with spatial learning in which sampling the payoff of one product causes consumers to update their beliefs about the payoffs of other products that are nearby in attribute space. Spatial learning gives rise to path dependence, as each new search decision depends on past experiences through the updating process. We present evidence of spatial learning in data that records online search for digital cameras. Consumers' search paths tend to converge to the chosen product in attribute space, and consumers take larger steps away from rarely purchased products. We estimate the structural parameters of the model and show that these patterns can be rationalized by our model, but not by a model without spatial learning. Eliminating spatial learning reduces consumer welfare by 12%: cross-product inferences allow consumers to locate better products in a shorter time. Spatial learning has important implications for the power of search intermediaries. We use simulations to show that consumer-optimal product recommendations are that are most informative about other products.

*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, and seminar participants at IIOC 2017, EC 2019, MIT Sloan, Duke/UNC, the Simons Institute at UC Berkeley, and the Yale IO workshop for helpful comments.

1 Introduction

We live in a world where the supply of information is substantial and increasing; it is more widely-shared (through the internet) and cataloged (by search engines) than ever before. As Herbert Simon presciently argued (Simon 1971), this has substantially increased the time spent acquiring information: searching online is now a significant part of the day for many people.[1] In online retail markets, consumers are presented with ever-growing sets of alternatives that cannot all be costlessly considered before a purchase decision is made. Search-mediating platforms such as Amazon, Netflix, and AirBnB play a significant role in guiding consumers' search paths through judicious information provision. Search algorithms and recommendation systems that determine how product information is presented to consumers are central to the economic value generated by these firms (Evans and Schmalensee 2016). 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.

Economists have long considered information frictions important, and have frequently relied on models of costly search to rationalize phenomena from price dispersion to unemployment (Varian 1980, Pissarides 1976). Less attention has been paid to the *process* of search itself - the ways in which people learn as they search, how accumulated information and changing beliefs direct their subsequent inquiries, and how this process affects economic outcomes.

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. Because of this, these models imply that options are sampled at random (if payoffs are iid) or in a pre-specified order (if payoffs

[1]For example, Boik, Greenstein and Prince (2016) show that the average US household participating in the Comscore survey spent around 2 hours a day online in 2013 (although much of this is content consumption, rather than information acquisition).

depend on ex-ante observable characteristics). In particular, what a searcher learns from searching an alternative only determines *whether* to continue searching, not *what* to search next.

This paper starts with the observation that, in many real life settings, learning about the payoff from one alternative should 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.* For example, a job seeker receiving a very attractive offer at Microsoft might reasonably infer that a potential Google offer would be better than they had previously 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 online reviews for a model with low resolution and decides it is not for her will probably negatively update her beliefs about all low resolution cameras.

We offer a framework for modeling spatial learning. The building blocks are a *characteristic space* consisting of ex-ante observable characteristics of the options (in the context of online retail, these could include price and star rating), and utility functions modeled as a *Gaussian process* over that characteristic space. Gaussian processes have found wide application in machine learning (Rasmussen and Williams 2005). They are very flexible and yet are fully 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. First, consistent with the model's prediction of path dependence, we show that consumers tend to take significantly larger steps in attribute space after viewing rarely purchased products. Second, replicating results documented by Bronnenberg et al. (2016), we show that the products searched by consumers converge in attribute space to the product ultimately purchased and that step size in attribute space and the variance of product attributes searched declines as search progresses. These patterns are suggestive of exploration of the characteristic space giving way to concentrated search in the neighborhood of previously explored high-payoff options.

We argue that these search path patterns can be used to identify the learning parameters of our model. For instance, the extent to which consumers jump away from rarely purchased products is informative of the spatial correlation in beliefs. In this way, the model opens the black box of the search path and uses variation that has not previously been exploited in the consumer search literature to learn about the forces that determine search sequences. With the increasing availability of online search data, we expect that this type of identification strategy will become increasingly feasible.

We estimate the model by Markov Chain Monte Carlo under a one-period look ahead assumption (similar assumptions have been made by Gabaix, Laibson, Moloche and Weinberg (2006), Ursu and Zhang (2020), and Yang, Toubia and De Jong (2015)). 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. As we hypothesized, allowing consumers to make inferences across products is essential to matching these search path patterns.

Simulated search paths show that learning is quantitatively important to consumer welfare. Expected consumption utility is about 112% lower for simulated consumers who do not extrapolate across products than for consumers with correct beliefs. Utility is similarly reduced if consumers over-extrapolate and update their beliefs about unsearched objects more than is implied by the estimated model. Spatial learning with correct beliefs about the covariance of utility across the product attribute space allows consumers to locate high-utility products in a shorter time. To benchmark the value of spatial learning, we show that non-learning consumers have to extend their search length by about 25% to obtain the same utility as consumers with correct beliefs.

These results imply that spatial learning plays an important part in directing consumers' search paths and purchase decisions, and that incorrect beliefs can lead to welfare losses. This suggests that changing consumer beliefs and search paths through information provision is a potentially important mechanism which online search platforms can use to influence purchase decisions. 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 the extent to which this type of search diversion is possible given the estimated model of learning. We simulate product recommendations by providing consumers information about the payoff of selected products before they begin searching. 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. This mechanism leads to the surprising finding that recommending a product that generates *higher* utility than similar products can *reduce* final consumer utility by up to 2.5%.

Finally, we ask what the model implies about the characteristics of consumer-optimal recommendations. We show that consumer utility is maximized when the set of recommended products are maximally *informative* about other products. Informative product recommendations are located in dense regions of the attribute space and, when multiple products are recommended, represent diverse areas of the attribute space. We show that informative recommendations allow consumers to locate higher utility alternatives in fewer searches.

This set of counterfactuals sheds some light on the power held by search engines and online platforms. We do not take a stand on the platform's objective - as highlighted by Hagiu and Jullien (2011), a profit maximizing platform may have competing incentives to maximize consumer welfare, to direct consumers towards (or away from) particular products, or to keep the consumer on the platform for as long as possible (thereby maximizing advertising revenue). Instead, we highlight the importance of consumer learning as a channel through which platforms can shape search and consumption, with significant potential for both improving and harming consumer welfare. These findings point to consumer learning and the potential for manipulation of beliefs as an important consideration in debates about the regulation of the online platforms and recommendation systems that mediate an ever-growing share of our consumption choices.

Related literature. Search is a well-studied topic in microeconomic theory, empirical industrial organization and marketing. Recent empirical work has studied the identification and estimation of some of the classic search models (Koulayev 2014) and testing the alternative of sequential versus non-sequential search using Comscore data (De Los Santos, Hortaçsu and Wildenbeest 2012). In the marketing literature there have been a number of papers that have taken the Weitzman (1979) model to data, including Kim, Albuquerque and Bronnenberg (2010), Honka and Chintagunta (2017) and Ursu (2018).

The strand of this literature most closely related to this paper is the work following Rothschild (1974) on search with learning. Typically, these models involve consumers updating their beliefs about the distribution from which searched objects are drawn. Recent papers in this literature include De Los Santos et al. (2012), Koulayev (2013), and Anghel (2020). A small number of papers study consumer learning where ex-ante unobservable payoffs are correlated across products. Adam (2001) 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. Dickstein (2018) applies this type of model to data on prescription drug demand. Our work extends this literature by estimated a flexible model of cross-product correlation of beliefs that does not use pre-specified nests.

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), and Fradkin (2018). Some papers in this literature, such as Ellison and Ellison (2009) and Hagiu and Jullien (2011), have considered the incentives for platforms to mislead consumers or divert search. Of particular interest is Gardete and Hunter (2018), who use a model of consumer search over multiple attributes to study optimal information design when the platform can choose which attributes to make salient, finding that there are limited incentives for platforms to obfuscate product attributes and mislead consumers. Although Gardete and Hunter (2018) consider consumers learning through search, the learning modeled is across attributes *within* a product, not across products as in this paper. Our findings show that cross-product learning is an additional channel through which a platform can influence search.

The Gaussian process model of beliefs builds on the literature on Gaussian processes in machine learning, as summarized by Rasmussen and Williams (2005). In this literature, optimal Bayesian learning based on Gaussian processes is used to construct algorithms for maximizing unknown functions. The observation that assuming a Gaussian process is a tractable way to design machine learning algorithms applies similarly to the modeling of learning by rational agents. The one period look ahead policy we adopt to aid computation is widely used in this literature, where it is known as the knowledge gradient policy (Powell and Ryzhov 2012, Frazier, Powell and Dayanik 2009). The Gaussian processes model of spatial beliefs has also been used in recent studies in the economics of oil and gas exploration (Covert 2015, Hodgson 2019).

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. Section 6 presents the results of the estimation, and Section 7 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 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.[2] 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 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 consumer starts by searching product C. But the next product they

[2]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.

Figure 1: Optimal Search Strategies

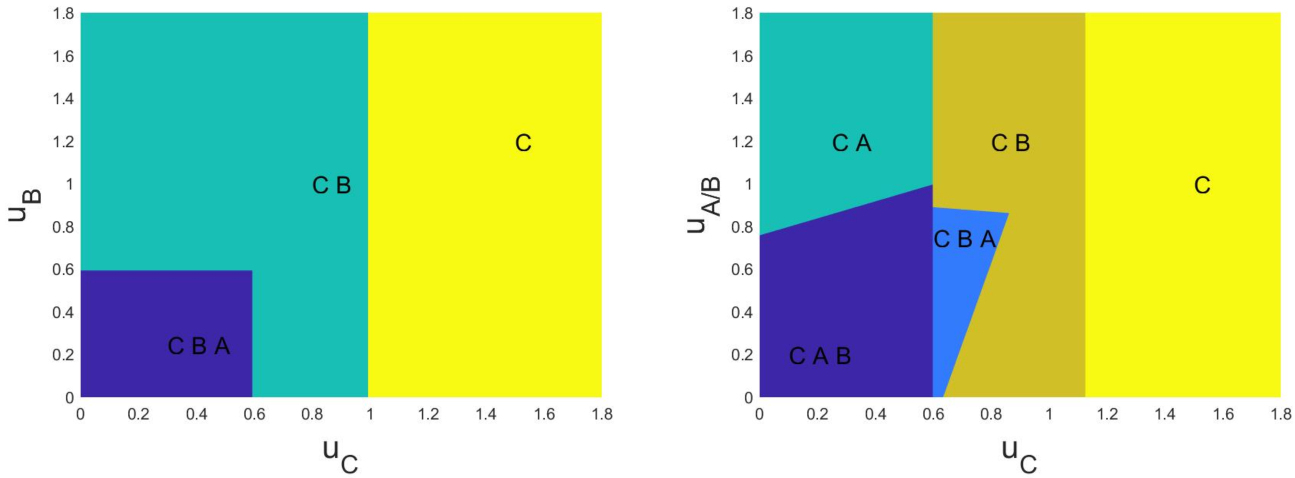

Notes: The left panel shows the optimal search strategies when there is no correlation in quality across products, and observing u_j only provides information about product j. The right panel illustrates how search strategies change when consumers believe that there is positive cross-product covariance 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 $\rho = 0.8$.

search depends on the observed value of u_C. If u_C is sufficiently high (the yellow region), 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 (brown region), and only searching A if the max of B and C is low (blue region).[3] If u_C is low, they infer that μ is also low, and instead target product A next, moving onto product B (purple region) if the maximum payoff of A and C is sufficiently low, and otherwise stopping (green region).

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]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.

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. 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 $m : \boldsymbol{X} \to \mathbb{R}$ is a function that maps a vector of characteristics to average payoffs, ξ_j is a product-level random effect drawn iid across products from a distribution $N(0, \sigma_\xi)$ and common to all consumers, and ϵ_{ij} is an idiosyncratic shock sampled iid across consumers and products from a distribution $N(0, \sigma_\epsilon)$.[4] The function $m(X)$ is unknown to consumers, and sampled from a Gaussian process with prior mean function $\mu(X)$ and covariance function $\kappa(X, X')$. We assume that μ is a continuous function, and that $\kappa(X, X') \equiv \tilde{\kappa}\left(\frac{\|X-X'\|}{\rho}\right)$ for some weakly positive, continuous and decreasing function $\tilde{\kappa}$, where $\|\cdot\|$ is the Euclidean norm and ρ is a parameter that controls how covariance declines with distance. We assume that the consumers know the prior. As consumers search, they update their beliefs about $m(X)$ according to Bayes rule (see the section on beliefs and learning below).

An interpretation of the model is that initially consumers don't know their preferences over characteristic space, which are summarized by $m(X)$. As they search, they get noisy signals of the

[4] Note that in discussing the model, we sometimes drom the i subscript on $m(X)$.

function (noisy because they observe u_j, and not $m(X_j)$). Because $\tilde{\kappa}$ is decreasing, the covariance in payoffs declines with distance in characteristic space, and learning is spatial: sampled payoffs are more informative about the payoffs of close-by products than those far away.

Another interpretation consistent with the model is that consumers know their preferences over the observable characteristics X, but there are other unobservable product characteristics whose values are unknown without search. As they search, consumers refine their model of the mapping between the observable and the unobservable characteristics, updating the model $m(X)$.

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. 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.

3.2 Beliefs and Learning

The search process is a non-stationary Markov Decision process. We model the state as a tuple $S_t = (\mu_t(X), \kappa_t(X, X'), \hat{j}, \hat{u}, J)$, where $\mu_t(X)$ are the current mean beliefs, $\kappa_t(X, X')$ 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 $\hat{j}, \hat{u}, J$ are straightforward. The mean and covariance functions update according to:

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

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

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 the consumer's learning process. Panel A represents a consumer's prior beliefs and ex-ante unknown preferences over a one-

Figure 2: Gaussian Process Learning

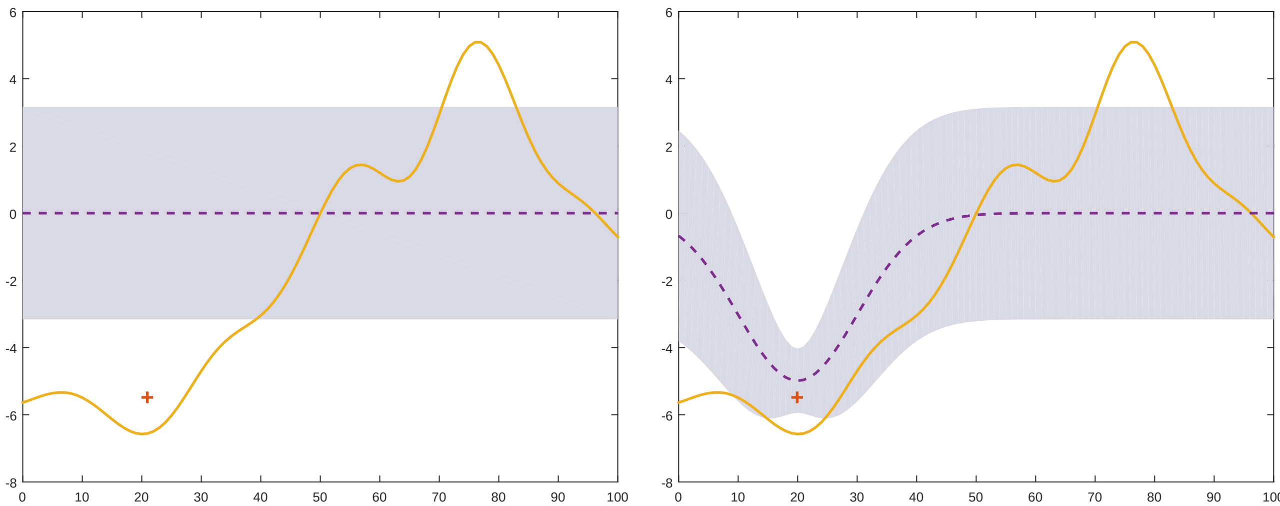

(A) Prior Beliefs (B) Posterior Beliefs

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.

dimensional characteristic space $X \in [0, 100]$. The consumer's prior mean, $\mu(X) = 0$ is indicated by a dashed line. The shaded area is a one standard deviation band of the prior Gaussian process around the mean. The solid line is the consumer's utility function $m(X)$ which is drawn from the Gaussian process. The consumer searches a product j 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 reduced the consumer's uncertainty about her utility function $m(X)$, especially for products close to X_j in parameter space.

3.3 Consumer Behavior

Because there are a finite set of products that can be searched, the consumer's decision problem can be solved by backward induction. Doing so will generally be computationally intractable with a reasonable number of products, since the state variables are continuous functions.[5]

[5]One could instead take the set of realized utilities for searched objects as the state variable — since these are sufficient for the beliefs — but with a large number of products this remains intractable. For example, Crawford and Shum (2005) interpolate the value function between a discrete set of states in a setting with 5 products. In section 5 we apply the model to a setting with around 300 products.

We would like to get a closed form solution so that we can illustrate some of the forces in the model. Accordingly, we analyze a heuristic solution to the problem: one period look ahead search. This is a common solution method in the literature on Gaussian processes, which typically employ n-period look ahead assumptions (Osborne, Garnett and Roberts 2009).

The one-period look ahead 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:

$$z_j = \Phi(a_j)\hat{u} + (1 - \Phi(a_j))\mu_j + \phi(a_j)s_j - c_j \tag{4}$$

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. The one-period look ahead 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.

While one-period look ahead is not generally optimal, it can often provide a close approximation to the optimal policy. In Appendix Figure A.2, we show that the optimal and one-period look ahead policies nearly coincide in the example of Section 2 above. Frazier et al. (2009) provides explicit bounds on the suboptimality of the one period look ahead, 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 ($\rho = 0$).[6]

It may also be the case that a myopic policy provides a good empirical approximation to the behavior of boundedly rational consumers. Gabaix et al. (2006) provide evidence that a one-period look ahead policy matches search behavior better than a fully rational policy in an experimental

[6] 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. The updates thus don't change the underlying decision problem — they are just affine transformations of the utility — and thus the problem is essentially iid.

setting. Our empirical application in Section 6 below will also show that consumer behavior can be matched well under such a policy.

Under the one-period look ahead assumption it is straightforward to prove some useful comparative static properties using the analytical characterization of the score in (4).

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 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 utility draw from a product k will incease 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.[7] Likewise, a low utility draw from product k will reduce z_j more for products close to k in attribute space. 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.

[7] 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.

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 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.

For an individual panelist, we observe the sequence of products viewed, the product eventually purchased, 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 webpage providing information about a single product is loaded. This could include product pages on retail sites such as Amazon.com, manufacturer websites, and review sites. Purchases are identified using a second ComScore dataset that tracks online transactions carried out by panelists. For each product view, 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. Note that this is a selected sample and not representative of the population of consumers. We use this data to illustrate broad patterns that motivate our modeling approach and to test our model.

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 described by four continuous attributes (price, zoom, display size, and pixels). The top panel of Table 1 records summary statistics on the distribution of these attributes across products.

Table 1: Summary Statistics

	Products			
	Mean	SD	Min	Max
Price	302.13	478.47	16.99	5250.00
Zoom	6.04	5.60	0	35
Pixel	10.54	3.13	1	21
Display	2.67	0.41	1.1	3.5
N	357			
	Searches			
	Mean	SD	Min	Max
Search Length	5.59	4.00	6.52	1.00
Purchase Discovered	0.79	1.00	0.29	0.03
Price Searched	285.45	406.45	16.99	5250.00
Zoom Searched	6.43	5.97	0.00	35.00
Pixel Searched	11.96	2.37	1.00	21.00
Display Searched	2.78	0.30	1.10	3.50

Notes: top panel records statistics on products from the digital camera data are defined by unique values of brand, zoom, pixel, and display. If there are multiple prices recorded for the same product, this table uses the average price recorded over all searches. Bottom panel records statistics on search paths from the digital camera data. Search path length is the number of prducts viewed. Chosen 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.

An observation in the data is a sequence of products viewed and the identity of the product purchased. The bottom panel of Table 1 records summary statistics on consumer search paths. The first row of the records path length - the number of products searched before purchase. The average consumer views about 5.6 products. There is a tail of consumers with very long search paths, the longest of which is 58 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. The remaining rows documents the distribution of attributes among products *searched*. For example, the mean price of products searched in \$285.45. This is the average over all searches by all consumers - including multiple counts of the same product if multiple consumers search that product. Compating 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 zoom, 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 - products considered late in search are more similar to the purchased product in terms of zoom than products searched early. The same pattern can be observed in other product attributes (pixels, and display size), as documented in Appendix Figure A.3. Note that purchased products are excluded from the data used to construct these figures, so the results are not artifacts of the fact that purchased products are discovered late in the search path.

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 Figure 4, 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. Note that this finding of a “narrowing” of search is

Figure 3: Convergence to Chosen Attribute Level

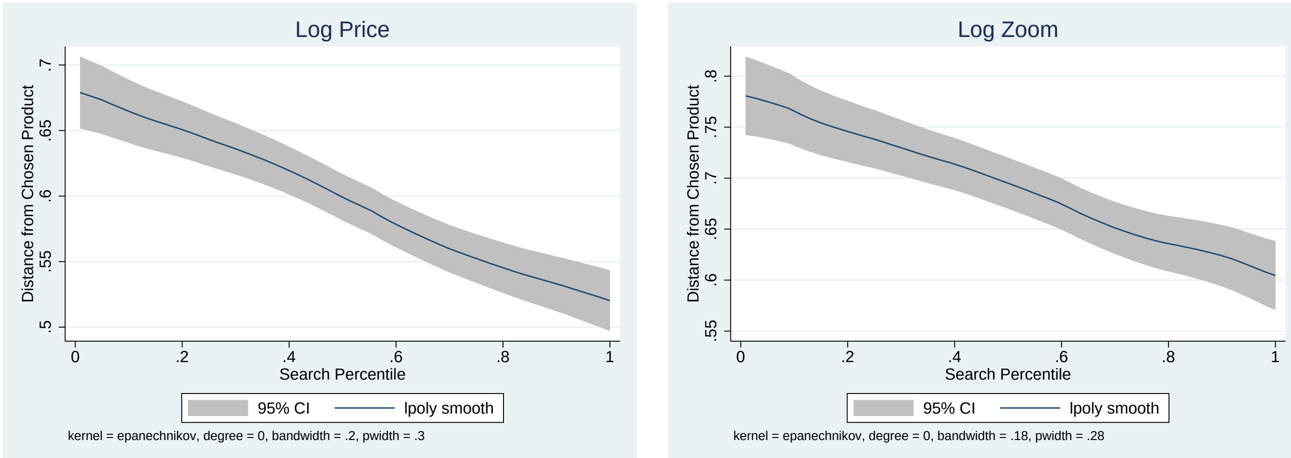

Notes: The y-axis for each panel records, for the relevant product attribute, the absolute difference in standard deviations of the attribute between the searched product and the product ultimately purchased. The x-axis reports the search percentile, as defined in the text. The product ultimately purchased 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.

not necessarily implied by the convergence of search to the chosen attribute levels illustrated by Figure 4: it could be that consumers always search a narrow area of the attribute space, but move their focus towards the chosen product over time.

The right panel of Figure 3 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 $|price_{t-1} - price_t|$ where $price_t$ is the price of the consumer's tth searched product. The x-axis records search percentile, as in Figure 4. 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.[8] The search paths used in Figures 3 and 4 include revisits - cases in which the consumer views a product more than once, perhaps on different websites. The patterns described here persist, but are less statistically significant, when revisits are excluded.[9]

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.

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

[9]See Appendix Figure A.4.

Figure 4: Narrowing of Search

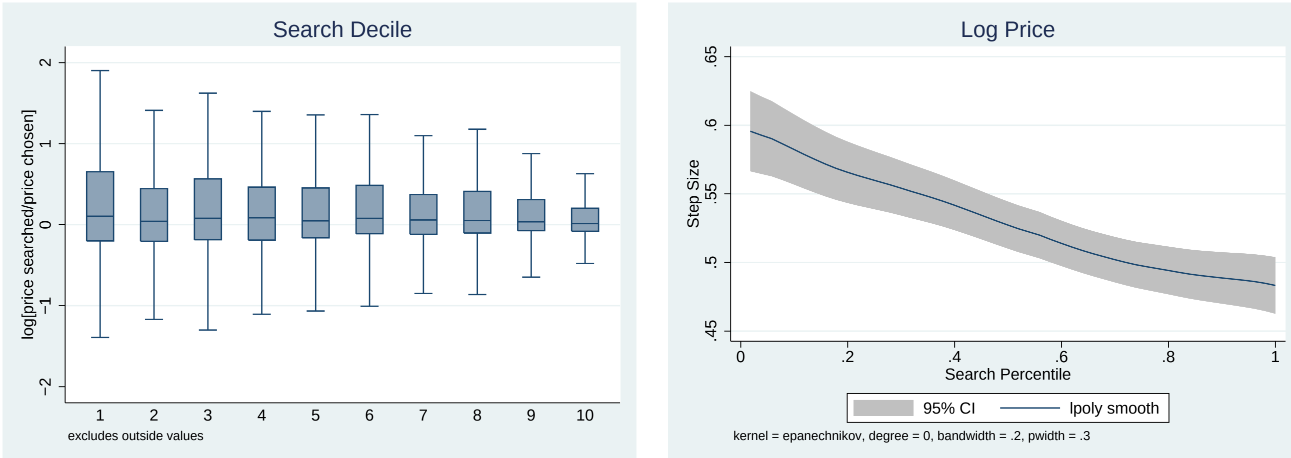

Notes: The left panel displays box plots that record the distribution of the log difference in searched price from the price of the product ultimately purchased, for each search decile as defined in the paper. The Box records the 25th, 50th, and 75th percentiles of the distribution and the whiskers record 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, including revisits to the same camera and excluding consumers who do not make a purchase.

This behavior is not predicted by standard models of sequential search. In contrast, correlated Gaussian process learning has been shown to exhibit this type of convergence behaviour. Frazier et al. (2009) show that agents following a one-period look ahead rule 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. The findings documented in Figures 3 and 4 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

Together these findings describe non-stationary search paths that are inconsistent with standard sequential search models and are suggestive of consumer learning. In this subsection we test a direct implication of the model of search with spatial learning developed in Section 3. The model assumes that beliefs are correlated in attribute space, so that after viewing an object, consumers learn their utility for that object and update their beliefs about other objects. 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.[10] 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, and hence we do not know what a particular consumer learns when she views a particular item, nor what her beliefs are before searching. 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. We then regress a measure of the "step size" of search after a consumer observes product j on this index.

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

[10] More precisely, Proposition 1 implies that $\frac{\partial z_j}{\partial u_k}$ is increasing in $\kappa(X_j, X_k)$. Consumers' prior $\kappa(X_j, X_k)$ is decreasing in the distance between k and j before consumers have started searching. As consumers search, their posterior beliefs do not necessarily follow this pattern for every pair of products. for example, if posterior variance in very low in a region of the atrreibute space, then $\kappa(X_j, X_k)$ and $\frac{\partial z_j}{\partial u_k}$ will also be low if j is in that region.

when the probability that consumer i purchases product $j \in J_i$ is given by:

$$P_{ij} = \frac{exp(\tilde{\theta}_j)}{1 + \sum_{k \in J_i} exp(\tilde{\theta}_k)} \tag{5}$$

$\tilde{\theta}_j$ is an index that measures the probability of purchase conditional on search.[11] We use OLS to decompose $\tilde{\theta}_j$ into a part that can be explained by product attributes and a residual:

$$\tilde{\theta}_j = X_j \gamma + \theta_j \tag{6}$$

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 values of $\hat{\theta}_j$ mean that a product is purchased more, conditional on being searched, than similar products and 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.

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 $|price_{it} - price_{it-1}|$ 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 a number of controls: search percentile, an indicator for whether product $j(i,t-1)$ is the product ultimately purchased, product density controls, and consumer fixed effects.

The definition and reasoning behind each of these controls is as follows. We include search percentile to control for the trends recorded in Figure 3. The indicator for purchase allows us to treat these observations differently (notice that in order for there to be a subsequent step, the consumer

[11]Note that this is *not* a structural object but a convenient statistical device for classifying products, and that equation 5 is not derived from the model. Some objects are never purchased, and so 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.

Table 2: Effect of Product Residuals on Step Size

	$\|price_{it} - price_{it-1}\|$	$\|pixel_{it} - pixel_{it-1}\|$	$\|zoom_{it} - zoom_{it-1}\|$	$\|display_{it} - display_{it-1}\|$
$\hat{\xi}_{j(i,t-1)}$	-.064***	-.274***	-.076***	-.280***
	(.019)	(.029)	(.026)	(.030)
$SearchPercentile_{it}$	-.130***	-.105**	-.087**	-.101**
	(.028)	(.041)	(.038)	(.042)
$Purchased_{it}$	-.104***	.002	-.083*	.000
	(.034)	(.050)	(.046)	(.051)
$ProductDensity_{it}$	.153***	2.159***	.311***	22.299***
	(.010)	(.160)	(.022)	(.973)
N	5590	5590	5590	5590
Consumer FE	Yes	Yes	Yes	Yes
Mean of Dep. Var.	.523	.707	.609	.732

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 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. Other covariates are described in the text. All regressions include 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.

must have returned to this product and purchased it later on). 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. $\hat{\theta}_{j(i,t-1)}$ is standardized so that the first row reports the effects of one standard deviation changes of $\hat{\theta}_{j(i,t-1)}$. Last, consumer fixed effects eliminate the possibility that these results are driven by a correlation between taking large or small steps and a idiosyncratic taste for products that are often or rarely purchased.

$\hat{\theta}_{j(i,t-1)}$ has a significant, negative effect on step size for each of the four attribute 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 2. When consumers view products with "surprisingly low" utility (those with low values of ξ_j), they jump further away in attribute space.

The data used in these regressions includes revisits to the same product (perhaps on other websites). In Appendix Table A.2 we present a version of these regressions that omits revisits and find that the results do not change dramatically. In Appendix Table A.3 we 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.

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 Table 3. All specifications include the same controls as those in Table 2. The estimated coefficients indicate that the correlation between $\hat{\theta}_j$ and step size persist. The coefficients are significant and most are slightly lower in magnitude than the one-step coefficients in Table 2.[12]

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 subsequest steps as consumers obtain more information. This pattern is consistent with a persistent effect of

[12]In Appendix Table A.4 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.

Table 3: Path Dependence: Multi-Step Differences

	Two Steps			
	$\lvert price_{it} - price_{it-2}\rvert$	$\lvert pixel_{it} - pixel_{it-2}\rvert$	$\lvert zoom_{it} - zoom_{it-2}\rvert$	$\lvert display_{it} - display_{it-2}\rvert$
$\hat{\xi}_{j(i,t-2)}$	-.041***	-.263***	-.083***	-.241***
	(.021)	(.030)	(.029)	(.107)
N	4939	4939	4939	4939
	Three Steps			
	$\lvert price_{it} - price_{it-3}\rvert$	$\lvert pixel_{it} - pixel_{it-3}\rvert$	$\lvert zoom_{it} - zoom_{it-3}\rvert$	$\lvert display_{it} - display_{it-3}\rvert$
$\hat{\xi}_{j(i,t-3)}$	-.050***	-.243***	-.045**	-.220***
	(.022)	(.033)	(.030)	(.035)
N	4398	4398	4398	4398

Notes: 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.

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.

5 Structural Estimation

5.1 Econometric Specification

In order to take the model developed in Section 3 to the data on consumer search paths, we make additional assumptions on the forms of the consumers' prior mean and covariance functions. We assume that consumers' prior means are linear in product characteristics:

$$\mu(X_j) = \alpha + X_j\beta \tag{7}$$

Note that in this application we restrict all consumers to having the same prior mean, $\mu_i(X_j) = \mu(X_j) \, \forall i$. This could be relaxed by allowing consumer-specific coefficients β_i, in which case the model would nest the random coefficients discrete choice model of Berry, Levinsohn and Pakes (1995). [13]

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. 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.

$$c_{ijt} = c + \zeta_{ijt} \tag{9}$$

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 (i.e there is no initial outside option), 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, $\{\beta, \alpha\}$, those determining the prior covariance function, $\{\lambda, \boldsymbol{\rho}\}$, the search cost parameter c, and the parameters

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

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 vector of parameters. 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.

5.2 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. 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.

5.3 Estimation

We estimate the model by constructing a likelihood function on the observed consumer search paths and choices. Under the assumption that search costs are given by equation (9) with logit errors, the probability of a consumer choosing to search product $j \in \tilde{\boldsymbol{J}}$ conditional on being at state $\mathcal{S}$, but unconditional on the realizations of the logit cost shocks is given by:

$$\tilde{P}_i(j|\mathcal{S}) = \frac{exp\left(E[\max\{\hat{u}, u_j\}|\mathcal{S}] - c\right)}{exp(\hat{u}) + \sum_{l \in \tilde{\boldsymbol{J}}} exp\left(E[\max\{\hat{u}, u_l\}|\mathcal{S}] - c\right)} \tag{10}$$

Suppose consumer i searches T_i times before stopping. Let j_{it} be the tth product 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 11.

$$L_i(\{j_{it}\}_{t=0}^{T_i}, \hat{j}_i|\{\mathcal{S}_t\}_{t=0}^{T_i}, \psi) = \left(\prod_{t=0}^{T_i-1} \tilde{P}_i(j_{it}|\mathcal{S}_t)\right) \tilde{P}_i(0|\mathcal{S}_{T_i}) 1\left(u_{\hat{j}_i} = \hat{u}_{j_{iT_i}}\right) \tag{11}$$

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 $\alpha + X_{j(i,\tau)}\beta + \xi_{j(i,\tau)}$. 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 unconditional on utility draws is given by equation 12.

$$L_i(\{j_{it}\}_{t=0}^{T_i}, \hat{j}_i|\psi) = \int L_i(\{j_{it}\}_{t=0}^{T_i}, \hat{j}_i|\{\mathcal{S}_t\}_{t=0}^{T_i}, \psi) dG(u_i) \tag{12}$$

In practice, we approximate this integral by averaging over draws from $G(u_i)$. The likelihood of the data for all N consumers is then given by equation 13.

$$L(\psi) = \prod_{i=1}^{N} L_i(\{j(i,t)\}_{t=0}^{T_i}, \hat{j}(i)|\psi) \tag{13}$$

Our estimation procedure uses this likelihood to form a Monte Carlo Markov Chain with flat priors. Let $\tilde{\psi} = \psi \backslash \{\sigma_\xi, \boldsymbol{\xi}\}$ be the set of all parameters except the product effects and the variance of the product effects. Using the distributional assumption $\xi_j \sim N(0, \sigma_\xi)$, and assuming an inverse gamma conjugate prior distribution of $\sigma_\xi \sim IG(1,1)$ on the variance of the product effects, the posterior distributions are given by equation 14 (Train, 2009).

$$
\begin{aligned}
&P(\tilde{\psi}|\sigma_\xi, \boldsymbol{\xi}) \propto L(\tilde{\psi}, \sigma_\xi, \boldsymbol{\xi}) \\
&P(\sigma_\xi|\tilde{\psi}, \boldsymbol{\xi}) \propto L(\tilde{\psi}, \sigma_\xi, \boldsymbol{\xi}) k \left(\sigma_\xi; 1+J, \frac{1+\sum \xi_j^2}{1+J} \right) \\
&P(\xi_j|\tilde{\psi}, \sigma_\xi, \{\xi_k\}_{k\neq j}) \propto L(\tilde{\psi}, \sigma_\xi, \boldsymbol{\xi}) \phi \left(\frac{\xi_j}{\sigma_\xi} \right)
\end{aligned} \tag{14}
$$

Where $k(\sigma_\xi; a, b)$ is the density of the inverse gamma distribution at σ_ξ with degrees of freedom a and scale parameter b. The assumption $\xi_j \sim N(0, \sigma_\xi)$ is discussed in the description of the model in Section 3, but not used in construction the likelihood, since the likelihood is conditional on ξ_j. The inverse gamma assumption on σ_ξ is a frequently used diffuse conjugate prior for the standard deviation of normal distributions (Train, 2009). In Appendix B we provide details on the Markov Chain procedure that generates draws from these posteriors. For each parameter, we take the mean of the draws as our parameter estimate, and the standard deviation of these draws to be the standard error of the estimate.

5.4 Identification

To investigate whether the model is identified by data on search paths, we run a Monte Carlo exercise. 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 then simulate N search paths and estimate the parameters of the model on the search path data by maximizing the likelihood given by equation 13. We repeat the search path simulation and parameter estimation 500 times, fixing the product characteristics and parameters over these iterations.

Table 4 reports the mean and standard deviation of the estimated parameters over the 500 iterations for $N = 500$ and $N = 1000$. In both cases, estimated parameters are close to the true values,

Table 4: Monte Carlo Exercise

	True Parameter	$N = 500$	$N = 1000$		True Parameter	$N = 500$	$N = 1000$
β_1	-0.5	-0.490	-0.492	λ	10	10.028	10.008
		(0.073)	(0.064)			(1.426)	(0.931)
β_2	0.5	0.499	0.493	ρ_1	1	1.022	1.002
		(0.094)	(0.090)			(0.278)	(0.227)
α	-5	-4.983	-5.005	ρ_2	2	2.074	2.015
		(0.319)	(0.207)			(0.763)	(0.282)
c	4	3.983	3.985	σ_ξ	10	9.951	9.967
		(0.193)	(0.119)			(0.730)	(0.415)
				σ_ϵ^2	10	9.970	9.987
						(0.238)	(0.154)

Notes: Table reports the mean and standard deviation of the estimated parameters across 500 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. The $N = 500$ and $N = 1000$ columns report the results for two separate exercises, one that uses 500 search paths, and one that uses 1000 search paths.

and standard errors are small relative to parameter magnitudes. For all parameters, the standard deviation of the estimates is lower for the 1000 search path case than for the 500 search path case, suggesting convergence to the truth as $N \to \infty$.

The Monte Carlo exercise provides some reassurance that the model is identified. To see *how* it is identified, notice that our model is different from standard models of sequential search, whose identification has been studied by Koulayev (2014) among others, because of the presence of the spatially correlated beliefs controlled by the parameters $\{\lambda, \rho\}$. As argued above, spatial correlation in beliefs is consistent with certain search path patterns that cannot be rationalized by a model without learning. For instance, the patterns recorded by Table 2 that show that consumers take larger jumps in attribute space after searching rarely purchased products. These patterns in the search sequences identify the parameters $\{\lambda, \boldsymbol{\rho}\}$.

The intuition 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 indentifies the parameters of the prior mean, β and α, and the total variance of prior beliefs. The probability that product j is purchased, conditional on being searched, identifies ξ_j. For instance, products that are rarely purchased relative to other products with similar attributes must have $\xi_j < 0$.

Cross-product variation in ξ_j and the distance between pairs of products j and k then identifies the spatial covariance parameters, $\{\lambda, \boldsymbol{\rho}\}$. For example, if $\xi_j < 0$ then $P(k|j)$, the probability of searching each product k after searching j, should be lower for k close to j in attribute space. Likewise, if $\xi_j > 0$ then $P(k|j)$ should be higher for k close to j. That is, $\frac{\partial P(k|j)}{\partial \xi_j \partial |X_j - X_k|} > 0$, which is an implication of Proposition 1. The size of this cross-derivative depends on the variance of the spatially correlated part of utility, λ, and the spatial covariance parameters, $\boldsymbol{\rho}$.[14] The empirical analogues of these cross derivatives are the patterns recorded by Table 2.

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 the learning parameters. Appendix C presents a more detailed argument along these lines. In the next Section, we will present the results of the estimation and show that the estimated model rationalizes these patterns, while a no-learning model does not.

6 Results

6.1 Parameter Estimates

We estimate the model on the digital camera search path data from BKM using the MCMC approach discussed above. We drop revisits from the data used in estimation, since the model does not rationalize multiple visits to the same product. 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 5. As we might expect, 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 positive but not significant. This does *not* mean that, for an individual consumer, the marginal utility of zoom is 0. In particular, $\frac{\partial u_{ij}}{\partial X} = \beta + \frac{\partial m_i(X)}{\partial X}$. In expectation, the second term is equal to zero, so the *expected* effect of zoom on utility is zero.

[14] For instance if $\lambda = 0$ then $\frac{\partial P(k|j)}{\partial \xi_j \partial |X_j - X_k|} = 0$ and consumers do not "jump" away from low-ξ_j products.

The actual effect of zoom on utility includes the Gaussian process draw, $m(X)$, which consumers do not learn until they start searching. This finding suggests that preference for cameras with different zoom is heterogeneous across consumers and that consumers use search to learn about their preference for this attribute.[15]

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. Since we can reject the hypothesis that $\rho_k = 0$ in favor of the alternative $\rho_k > 0$, the data on search paths 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 half 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.[16]

As discussed in the previous Section, the parameters β are identified by the probability that each product is searched first, since they reflect consumers beliefs about expected utility *before* search begins. If consumers search paths tended follow this ex-ante ordering of products, the estimated value of λ would be small. The extent to which consumers' search paths deviate from these prior beliefs - for example by taking larger jumps after sampling products with negative ξ_j - is rationalized by more uncertainty about $m(X)$ through a higher value of λ.[17]

Finally, the search cost parameter, c, rationalizes the observed search lengths. Note that the estimated coefficient on price cannot be used to give a dollar interpretation to c since the coefficient on price includes both the direct effect of price on utility and the indirect effect of price on consumers' prior beliefs about quality, as discussed above.

[15] Note that a richer model might include a consumer-specific coefficient, β_i and a Gaussian process component, $m_i(X)$. This would separate consumer heterogeneity in beliefs about the expected relationship between attributes and utility and uncertainty about $m(X)$.

[16] The "signal to noise ratio" is approximately $2:3$.

[17] This result could also be driven by cross-consumer heterogeneity in expected utility, $X\beta$. See footnote 15 above.

Table 5: Estimated Parameters

	Estimate	SE		Estimate	SE
β_1 (log price)	-0.704	0.084	ρ_1 (log price)	0.321	0.011
β_2 (log zoom)	0.082	0.082	ρ_2 (log zoom)	1.292	0.039
β_3 (log pixels)	2.404	0.088	ρ_3 (log pixels)	1.442	0.060
β_4 (log display)	0.541	0.107	ρ_4 (log display)	2.462	0.068
α	-24.373	0.370	λ	20.171	0.121
c	7.065	0.120	σ_ξ	24.850	0.628
			σ_ϵ	3.854	0.080

Notes: Table reports estimated parameters and standard errors. Estimation uses the MCMC procedure described in Section 5. 5,000 draws from the chain are dropped for burn in, and the reported estimates are the mean and standard deviations of the 15,000 draws from three parallel chains. For more details on the estmiation procedure, see Appendix X.

Table 6 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 10,000 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 two rows indicate that the distribution of search path lengths, and the search percentile at which the purchased product is first discovered in the simulated paths match the data reasonably well. The remaining rows record the average observable characteristics of the products searched. For instance, the third column records the average value of price across all searches by all consumers (or simulated paths). For each attribute, the average characteristic searched in the data is close to the simulated value. Notice also that the average *simulated* searched product characteristics are different from the average product characteristics recorded in Table 1, because — just as in the real search data — simulated consumers do not sample products uniformly at random.

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

Table 6: Model Fit

	Data		Simulations	
	Mean	SD	Mean	SD
Search Length	5.390	6.429	6.054	4.702
Chosen Product Discovered	0.806	0.280	0.796	0.247
Average Price Searched	284.156	403.989	312.705	469.277
Average Zoom Searched	6.437	5.968	6.469	5.826
Average Pixel Searched	11.961	2.372	11.839	2.705
Average Display Searched	2.784	0.304	2.782	0.329

Notes: The first two columns report statistics on search paths from the data used in estimation. Sample includes all search paths in the data, dropping revisits to the same product. Search length and chosen product discovered are consumer averages. Chosen product discovered records the average search percentile at which the prodcut that was ultimately purchased was first searched. Average product attributes searched are averages over all searches. That is, a consumer that views 5 products will enter the average 5 times, once for each product. the second two columns recod alalogous statistics for 10,000 simulated search paths, holding all parameters at their estimated level and redrawing $m_i(X)$ and ϵ_{ij} for each simulated consumer.

attribute space after viewing products that are rarely purchased. Second, as recorded by Figures 3 and 4, consumers get closer to the purchased product and take smaller steps in attribute space 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 a set of "no learning" parameter estimates. The no learning estimates are constrained: we impose the restriction $\lambda = 0$ and then otherwise proceed as described in Section 5 above. They are recorded in Appendix Table A.5.

Table 7 replicates the step size regressions recorded in Table 2 using 10,000 simulated search paths at the estimated parameter values and the $\lambda = 0$ parameters. After generating the search paths, we regress purchase probabilities on observable product attributes and obtain product-level residuals, $\hat{\theta}_j$ as described above in Section 4. We then regress the tth search step size in each of the four observable dimensions on the product residual of the $t-1$th product searched. At the baseline parameters, the model matches these step size patterns closely. 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

Table 7: Simulations: Effect of Product Residuals on Step Size

	Baseline Parameters			
	$\|price_{it} - price_{it-1}\|$	$\|pixel_{it} - pixel_{it-1}\|$	$\|zoom_{it} - zoom_{it-1}\|$	$\|display_{it} - display_{it-1}\|$
$\hat{\theta}_{j(i,t-1)}$	-.039***	-.024***	-.039***	-.022***
	(.004)	(.004)	(.004)	(.005)
N	47825	47825	47825	47825

	$\lambda = 0$ Parameters			
	$\|price_{it} - price_{it-1}\|$	$\|pixel_{it} - pixel_{it-1}\|$	$\|zoom_{it} - zoom_{it-1}\|$	$\|display_{it} - display_{it-1}\|$
$\hat{\theta}_{j(i,t-1)}$	-.002	-.001	.005	.001
	(.004)	(.004)	(.004)	(.005)
N	48547	48547	48547	48547

Notes: Table presents regressions of search step size on the product residual index $\hat{\theta}_{j(i,t-1)}$. Sample is 10,000 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 $\lambda = 0$. Step sizes are measured using the absolute difference in log product attributes between the tth and the $t-1$th search. $\hat{\theta}_{j(i,t-1)}$ is constructed as described in Section 4 of the text. Values of $\hat{\theta}_{j(i,t-1)}$ are standardized so that estimated coefficients are the effect of one standard deviation. Other covariates are described in Section 4. All regressions include 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.

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 Table 2 are not statistically different from zero for each of the product attributes. Without spatial learning, there is no mechanism through which product effects ξ_j can affect beliefs about other products.

Figure 5 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 top panel of Figure 5 reports this relationship for log price and log zoom 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[18]. These patterns are generated by the dynamics of spatial learning in the model, and not an artifact of the data.

[18]The same exercise for display and pixels is recorded in Appendix Table A.5. The pattern is flatter in both the data and the simulations for these attributes.

The bottom two panels record 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, although the pattern is less statistically significant and the magnitude of the drift is smaller. The ability of the model to rationalize these patterns and the step size patterns in Table 7 makes the presence spatial learning a plausible explanation for several aspects of non-stationary 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 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 draw 20,000 values of $m(X)$, and simulate search paths under the baseline assumption of $\delta = 1$. Let the simulated search length of simulated consumer i be l_i. We then simulate search paths for these same values of $m(X)$ with the multiplier, δ, set to values between 0 and 2, fixing the length of each consumer's search path at l_i. 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

Figure 5: Simulations: Convergence in Attribute Levels

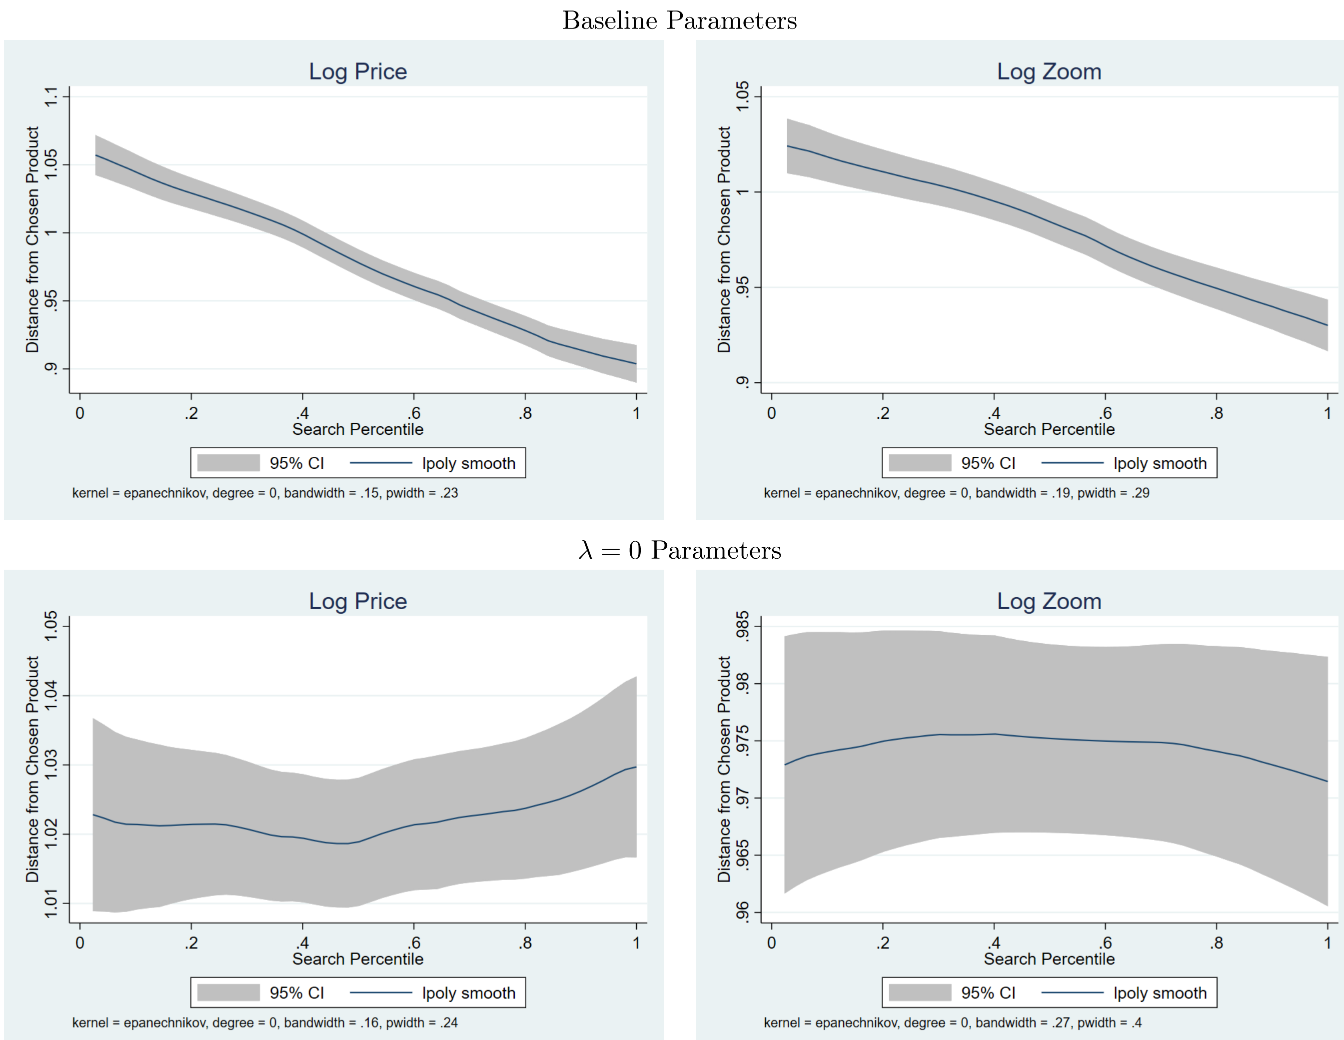

Notes: Figures are constructed using 10,000 search paths simulated at the estimated parameters. The top two panels use the baseline estimates, and the bottom two panels use the estimates under the restriction the $\lambda = 0$. The y-axis records, for the relevant product attribute, the absolute difference in standard deviations of log price and log zoom 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 6: The Value of Learning

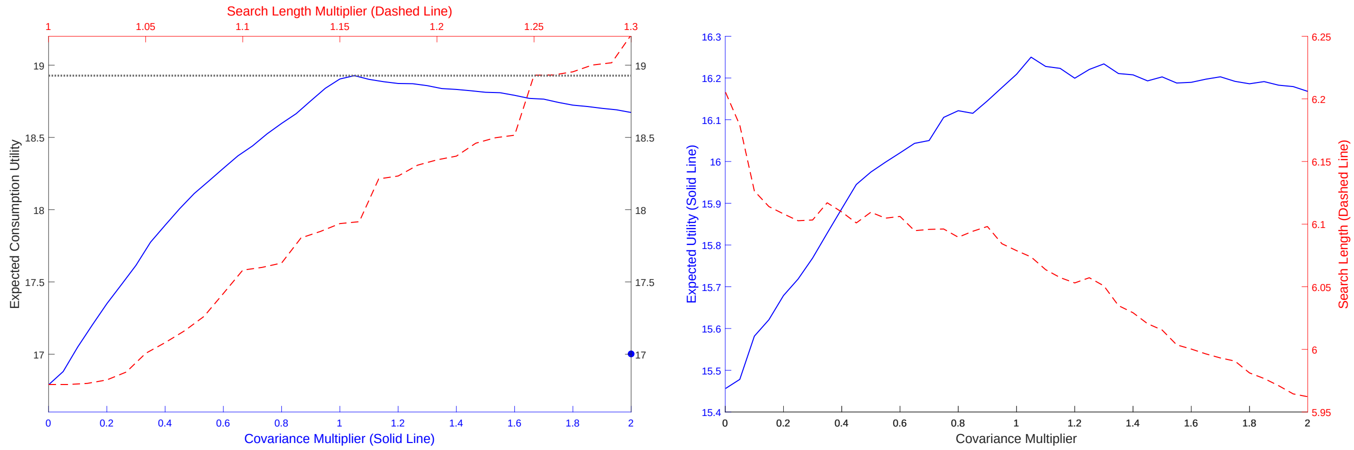

Fixed Search Length | Optimal Search Length

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. Each point on the blue line is a separate average over 20,000 simulations. Search length is held fixed for each simulated consumer at l_i, 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, rounded to the nearest integer, and the covariance multiplier is set to $\delta = 0$. In the right panel the solid blue line records, for values δ along the lower x-axis, the average total utility (consumption utility less search costs) of 20,000 simulations. Search length is not fixed. The dashed red line records the average search length for the same simulations.

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 6 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, she does not update her 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 12% lower than at $\delta = 1$. That is, if consumers do not update their beliefs as they search, the best match from the resulting search paths, fixing search length, is 12% worse than the best product obtained when consumers update beliefs correctly.

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 the same as if $\delta = 0$. This is illustrated by the blue dot at the far right of Figure 6, which simulates a counterfactual in which $\kappa(X, X') = \lambda^2$, the limit of the function given by equation X as $\rho \to \infty$.

To benchmark the value of learning, we ask how much longer a consumer who does not update her 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 6. Expected utility increases with search length and reaches the level of utility obtained by a consumer with correct beliefs (indicated by the horizontal line) at around $\gamma = 1.25$. This means that a consumer that does not learn as she searches has to sample about 25% more products than a consumer who learns 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 6 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 - and the red line records average search length. As in the fixed length simulations, utility is maximized when $\delta = 1$ due to over-and under-extrapolation when $\delta \neq 0$. However, note that total utility declines less steeply as $\delta \to 0$. Utility is about 5.6% lower at $\delta = 0$ than at $\delta = 1$. The loss in utility from under-extrapolation is partially offset by more search, illustrated by the dashed red line. When consumers do not extrapolate across products, the variance of beliefs about unsearched products utilities is not reduced through search (per equation

4), and consumers persistently overestimate the potential gains from continuing to search.

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 and purchase decisions, and that incorrect beliefs can lead to welfare losses. The effect of over- and under-extrapolation on consumption utility highlights one of the innovative features of our model of search - the introduction of path dependence. What a consumer learns from the first product she searches determines what she searches next, and ultimately what she purchases. Affecting consumer beliefs and search paths through information provision is therefore a potentially important channel through which online retail platforms can influence purchase decisions. Examples of information provision that may alter consumers' beliefs and search paths include product recommendations, comparisons, search results rankings, and sponsored search results. For example, search engines and online retail platforms such as Amazon, Google, and eBay frequently place sponsored products or advertisements at the top of search results pages. Indeed, it is well documented (for example, see Ursu (2018)) that highly ranked or salient of products on online platforms are more lilkely to be searched first.

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. This stands in contrast to models of sequential search without spatial learning, in which manipulating the first object viewed has no effect on the sequence of objects searched thereafter, only on the point in the sequence at which the consumer stops searching. This also implies that in a model without learning, a consumer who purchased product A would purchase either A or B in a counterfactual world where he is forced to view B first, whereas in a model with learning, forcing a consumer to view product B first could alter their search path such that they end up purchasing some third product C. For instance, if a consumer learns that they would obtain an unexpectedly high payoff from B, they might search through other products that are similar to B and yield similarly high payoffs, and end up purchasing such a product, C, that they would not have searched at all had they been free to start their search anywhere.

In this section we show that information provision through product recommendations can be used by the search platform to direct search and affect consumer welfare. First, we show that platforms can manipulate consumers beliefs to direct them away from certain regions of the product space. This "search diversion" can be exploited by platforms that want to direct consumers towards high margin products.[19] Next, we characterize the properties of consumer welfare optimizing product recommendations, and show that a platform that wants to optimize consumer experience should show a diverse, representative range of products to the consumer to help speed up learning.

7.1 Search Diversion

To illustrate the effect of information provision on search paths and consumer welfare, we use the estimated model to simulate search paths under different information provision scenarios. We draw 5000 values of $m(X)$ and simulate search paths. For each search path, we add a "focal product", F, at a random location X_F drawn from the set of existing product locations, and "show" the consumer the utility they would obtain from this product before they begin their search. Because they have viewed the focal product before they begin their search, consumers' beliefs at the beginning of their search process differ from the uninformed beliefs described by the Gaussian process prior. 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.

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. Although consumers have rational beliefs about the normal distribution from which ξ_F is drawn - consumers make correct inferences given this distribution - they do not observe ξ_F separately from total utility. Particularly large

[19] 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.

Figure 7: Search Diversion

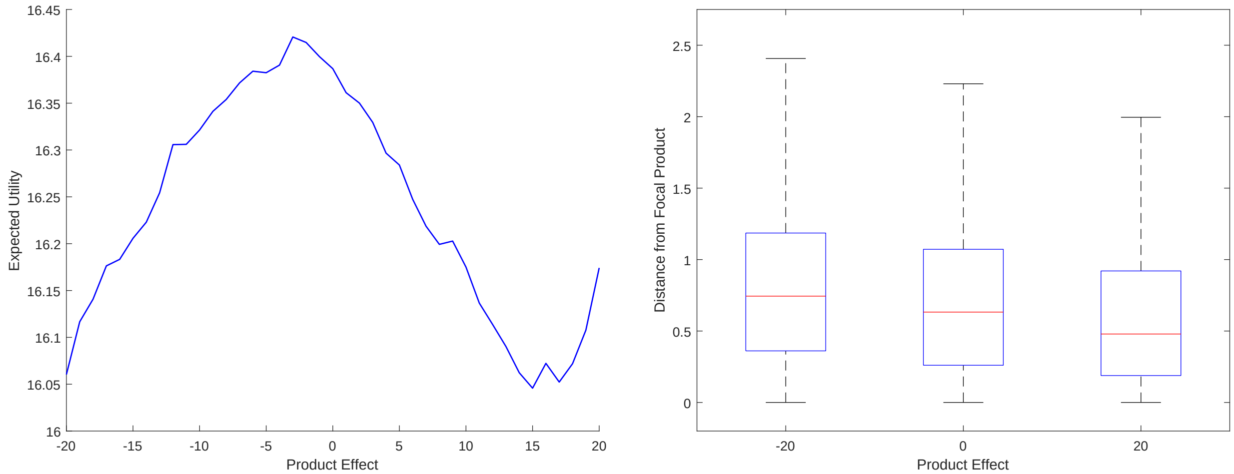

Notes: The left panel records the average utility (consumption utility minus search cost) over 20,000 simulated paths. Consumers update their beliefs before searching based on viewing a product F with X_F drawn at random from the set of existing products. The x-axis records the value of ξ_F, and each point along the x-axis is a separate set of 20,000 simulations.The right panel records box plots of the distribution on $|logprice_{\hat{j}(i)} - logprice_{iF}|$ where $\hat{j}(i)$ is the product purchased by consumer i for simulations using three different values of ξ_F. The Box records the 25th, 50th, and 75th percentiles of the distribution and the whiskers record the upper and lower adjacent values.

(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. For consumers with $m(X_F) > 0$, but $m(X_F) + \xi_F < 0$, this inference will lead them to *incorrectly* revise down their beliefs about the expected utility of nearby products. 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 7 illustrates the effects of information provision on consumption utility. The blue solid line plots the mean utility (consumption utility minus search costs) among the 5000 paths for different values of the focal product effect. Expected utility is maximized locally when ξ_F is close to 0. In this case, the observed utility of focal product is equal to the consumer's Gaussian process draw plus idiosyncratic error: $X_F\beta + m(X_F) + \epsilon_{iF}$. In other words, the product F is *representative* of the part of utility that is correlated across observable dimensions when $\xi_F = 0$.

When ξ_F is increased or decreased from 0, average consumption utility falls. Reducing ξ_F by one standard deviation (i.e. $\sigma_\xi \cong 20$) lowers expected consumption utility by about 15%, as consumers are diverted away from products near F. As ξ_F is increased above 0, expected utility also falls as consumers whose best match is farther away search more products near F. However, past some threshold, expected utility in 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 point at which this takes place depends on the relative importance of these two components of unobserved utility - if λ is large relative to σ_ξ, then cross-product inference about $m(X)$ is more important and the vertical component is less important.

The diversion of search paths that generates these effects on consumption utility is illustrated by the right panel of Figure 7. Recall $\hat{j}(i)$ is the product purchased by simulated consumer i. Each box plot illustrates the distribution of $\left|logprice_{\hat{j}(i)} - logprice_F\right|$ among the 5,000 simulated consumers. The x-axis records the value of ξ_F used in the simulation. A comparison of the three box plots reveals how information provision shifts the distribution of demand across products. When $\xi_F < 0$, demand falls for products that are close to F and rises for products that are further away, and vice versa for $\xi_F > 0$. In a model without spatial learning, any substitution in this counterfactual would be towards the focal product - the demand for all other products would remain the same or fall. The significant search diversion recorded in Figure 7 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 Consumer-Optimal Recommendations

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. However, a platform's objective function when designing its information provision strategy is likely more complex than simply maximizing short run revenue. A forward looking platform may have an interest in maximizing consumer utility to encourage consumers to return to the platform in future. 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 "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". In this subsection we ask what the model tells us about consumer-optimal information provision: what are the characteristics of consumer optimal product recommendations when consumers are spatial learners?

To answer this question, we simulate search paths with different sets of "recommended products". For each simulation, we show the consumer the utility they would obtain from two products, then let them search as normal. As with the simulations described above, the consumer has to pay the search cost if they with to purchase one of the two "recommended" products, so the intervention is purely informational. Platforms are typically designed in such a way that a limited number of products can be highlighted to the consumer before they begin to examine alternatives - for example the products at the top of the results page or highlighted in a separate recommendation panel. This exercise models this type of information provision technique and allows us to characterize the pairs of recommended products that optimize consumer welfare.

We select 1,000 pairs of products at random from the set of available products in the data. We label the two recommended products j_1 and j_2. Each product is characterized by its observable characteristics, X_j and its unobservable product effect, ξ_j. For each pair of recommended products, we draw 1,000 values of $m(X)$, and simulate 1,000 search paths. We simulate both fixed-length paths (fixing the length of each consumer's search path at it's value from a no-recommendation simulation), and optimal (up to the one-period look ahead assumption) search paths in which we give the consumer the option to stop searching at any point. We then use the simulated search paths to run regressions described by specification 15 and 16.

$$Y_r = \alpha + Distance\left(j_1(r), j_2(r)\right)\beta_1 + InvDensity\left(j_1(r), j_2(r)\right)\beta_2 + f(\xi_{j_1(r)}) + f(\xi_{j_2(r)}) + \epsilon_r \quad (15)$$

$$Y_r = \alpha + KL\left(j_1(r), j_2(r)\right)\beta_1 + f(\xi_{j_1(r)}) + f(\xi_{j_2(r)}) + \epsilon_r \tag{16}$$

Where r indexes a simulation, and $j_1(r)$ and $j_2(r)$ are the recommended products for that simulation. $f(\xi)$ is a flexible function of the product effect. In practice we allow it to be piecewise quadratic with different coefficients for $\xi > 0$ and $\xi < 0$.

$Distance\left(j_1(r), j_2(r)\right)$ is the average distance between j_1 and j_2 over the four (standardized) observable product attributes. $InvDensity\left(j_1(r), j_2(r)\right)$ is the average distance between products j_1 and j_2 and all other products. That is, it is a measure of "inverse density" - when $InvDensity\left(j_1(r), j_2(r)\right)$ is smaller, the recommended products are located in a denser region of the product space. $KL\left(j_1(r), j_2(r)\right)$ is the expected Kullback-Leibler divergence between the prior and posterior beliefs about all other products after observing products j_1 and j_2. This is a measure of the informativeness of the recommended products about the unobserved products.[20]

We run these regression for three dependent variables, Y_r - the average consumption utility, the average search length, and the average total utility including search costs. Results are reported in Table 8.

The first column records the effects of recommended product characteristics on consumption utility for fixed search length simulations. Expected consumption utility increases with the distance between j_1 and j_2 and decreases with the average distance between j_1 and j_2 and other products. Both coefficients are statistically significant. These results indicate that, holding fixed product effects ξ_j, consumers are able to obtain better product matches in a fixed number of searches when the recommended products are located in dense regions of the search space, and the two recommended products are not close together. The intuition behind these findings is clear. First, the information value of viewing a recommended product is higher when it is more informative about other products. In the context of spatial learning, consumers learn more when the recommended product is close to other products along the observable attribute dimensions. Second, when the consumer views two products, the information value of j_2 is diminished if it located close to j_1,

[20] The KL divergence is a measure of the difference between two distributions. It can be interpreted as the information gain when moving from one distribution to another (see Kullback and Leibler (1951) and Kullback (1997)). See Appendix D for details.

Table 8: Effects of Recommended Product Characteristics

Search Length:	Fixed		Not Fixed		
Dependent Variable:	Consumption Utility		Cons. Utility	Search Length	Total Utility
$Distance$	0.286***		0.159***	-0.068***	
	(0.054)		(0.049)	(0.020)	
$InvDensity$	-1.091***		-0.561***	0.084*	
	(0.136)		(0.124)	(0.050)	
KL		7.530***			5.869***
		(1.484)			(1.411)
$\lvert\xi_{j_1}\rvert\,(\xi_{j_1} > 0)$	-0.071***	-0.063***	-0.069***	0.016***	-0.117***
	(0.005)	(0.006)	(0.005)	(0.002)	(0.005)
$\xi_{j_1}^2\,(\xi_{j_1} > 0)$	0.006***	0.006***	0.005***	-0.001***	0.006***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\lvert\xi_{j_1}\rvert\,(\xi_{j_1} < 0)$	-0.052***	-0.047***	-0.063***	-0.004	-0.047***
	(0.009)	(0.010)	(0.009)	(0.003)	(0.009)
$\xi_{j_1}^2\,(\xi_{j_1} < 0)$	0.001***	0.001***	0.001***	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
N	1000	1000	1000	1000	1000

Notes: An observation in the regressions is a simulation. Each simulation generates 1,000 paths at the estimated parameter values after consumers observe the utility from two products, j_1 and j_2, drawn randomly as described in the text. for each column, 1,000 pairs of locations, (j_1, j_2) are drawn, and the statistics used in the regression are computed from the resulting search paths. Each regression controls for a peicewise quadratic function in ξ_{j_1} and ξ_{j_2}. Coefficients on ξ_{j_2} are not reported. The first two columns use fixed search lengh simulations. The third through fifth columns use simulations for which the search path length is not fixed. *** indicates significance at the 99% level. ** indicates significance at the 95% level. * indicates significance at the 90% level.

since the posterior variance of beliefs conditional on viewing j_1 is lower in the region of j_1.

The second column shows that consumption utility is significantly positively correlated with the KL divergence in beliefs induced by j_i and j_2. That is, average consumption utility is higher when the recommended products have a larger effect on the posterior beliefs about other products. This finding is consistent with the results in the first column - as described by Appendix Table A.6, KL divergence is positively correlated with *Distance* and negatively correlated with *Density*. Together, these findings indicate that the average consumption utility achieved in a fixed number of searches is maximized when the platform recommends an *diverse* and *informative* set of products.

The bottom four rows for the first two columns of Table 8 record the effect of ξ_{j_1} on expected consumption utility. These results are consistent with the pattern described in Figure 7. As discussed above, product recommendations that are *representative* in the sense that $\xi_j \simeq 0$ maximize expected utility locally.

Economically significant changes in welfare can be generated simply by relocating j_1 and j_2 in product space. Fixing $\xi_{j_1} = \xi_{j_2} = 0$, predicted utility varies substantially over the set of recommended product locations used in the simulations: the best recommendation in terms of consumer utility generates a predicted utility that is 9% higher than the worst recommendation. Of course, as $\xi_j \rightarrow \infty$, the direct effect of recommending a product with a large vertical utility component dominates the learning effect of product location.

The third through fifth columns of Table 8 repeat this exercise without the fixed search length restriction. As in the fixed search length simulations, locating j_1 and j_2 in denser regions of the product space and farther away from each other improve consumer welfare. These gains are achieved both through improved consumption utility and lower search costs. When j_1 and j_2 are more informative - located further apart and in dense regions of the search space - consumers obtain better matches to products and shorten their search paths because they learn more quickly about their preferences over a wider range of the search space.

8 Conclusion

In this paper, we develop a model of search with spatial learning and investigate its implications for platform power in online retail. In this model, 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.

In standard models of sequential search, the consumer's choice of *whether* to continue searching depends on the observed utilities of those objects already searched. With spatial learning, the decision of *which* product to search also depends on past observations. For example, a negative observed utility for a particular product is likely to direct subsequent search away from similar products. 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 document stylized facts using data on consumer search paths in online search for digital cameras that provide support for the intuitive notion that consumers explore the product space to learn about their preferences, and that learning about the payoff from one object provides information on the payoffs from similar objects. Consumers initially take large steps over a wide range of the product attribute space before focusing on products close to the ultimately purchased product in later search. Consumers also take larger steps in attribute space away from products that are rarely purchased.

We argue that these descriptive patterns identify the learning parameters of the model, and we estimate the parameters of the model using the search sequence data. We show using simulations that the estimated model matches these patterns closely and that a model without learning does not. This provides striking evidence that spatial learning drives consumer search patterns. We quantify the value of learning by simulating search paths under different assumptions about consumer beliefs, and show that non-learning consumers would have to search for 25% longer than learning consumers in order to obtain the same utility.

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. This means that, unlike in a model without spatial learning, providing additional information to consumers can reduce consumer welfare on average. In particular, recommendations can reduce welfare even if the platform recommends a better than average product to the consumer because of the effect of viewing non-representative products on consumer beliefs. An interesting direction for future research would be to use data with observed variation in product recommendations or rankings to provide empirical confirmation of these path dependence effects.

A final set of simulations show that when platforms recommend sets (in our case pairs) of products, consumer welfare is maximized when the recommended products are located in densely populated regions of the search space and are far away from each other. Pairs of recommended products that have these characteristics are informative in terms of Kullback-Leibler divergence. This finding complements work by De Los Santos and Koulayev (2017), who show how promoting products that are expected to yield high utility to specific consumers can reduce consumer search costs. In practice, a platform interested in maximizing consumer welfare might be interested in promoting informative products in order to aid consumer learning *and* consumer-specific high value products, depending on the extent to which the platform has ex-ante information about consumer-specific preferences. Exploring the trade-off between targeting consumer preferences and informing consumer beliefs through recommendations is a promising avenue for future research.

Together, these findings point to the importance of cross-product learning in determining consumer search paths. This mechanism should be taken into account debates about the regulation of search platforms and recommendation systems. Our empirical findings suggest that omitting cross-product learning from such an analysis would lead to misleading predictions about changes in demand and welfare under different information environments. Future research might apply this modeling framework to analyzing the behavior of platforms and evaluating the effects of different recommendation algorithms. This line of research has the potential to contribute to the design of regulations that promote “trust in recommendations,” highlighted by Tirole (2019) as one of leading challenges of the digital economy.

References

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

Anghel, Anca-Patricia, “Demand Estimation with Learning and Search Costs,” Working Paper 2020.

Boik, Andre, Shane Greenstein, and Jeffrey Prince, “The Empirical Economics of Online Attention,” Working Paper 22427, National Bureau of Economic Research jul 2016.

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.

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.

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 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.

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

Evans, David and Richard Schmalensee, *Matchmakers : the new economics of multisided platforms*, Boston Massachusetts: Harvard Business Review Press, 2016.

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.

Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg, "Costly information acquisition: Experimental analysis of a boundedly rational model," *American Economic Review*, sep 2006, *96* (4), 1043–1068.

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

Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin, *Bayesian Data Analysis*, New York: Chapman and Hall/CRC, nov 2013.

Hagiu, Andrei and Bruno Jullien, "Why do intermediaries divert search?," 2011, *42* (2), 337–362.

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.

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

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.

Kullback, Solomon, *Information Theory and Statistics*, New York: Dover Publications, 1997.

— **and Richard A. Leibler**, "On Information and Sufficiency," *Annals of Mathematical Statistics*, 1951, *22* (1), 79–86.

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.

Pissarides, Christopher A., "Job Search and Participation," *Economica*, feb 1976, *43* (169), 33.

Powell, Warren B. and Ilya O. Ryzhov, *Optimal Learning* Wiley Series in Probability and Statistics, Hoboken, NJ, USA: John Wiley & Sons, Inc., mar 2012.

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.

Simon, Herbert A, *Designing Organizations for an Information-Rich World*

Tirole, Jean, *Economics for the Common Good*, Princeton: Princeton University Press, 2019.

Ursu, Raluca and Qianyun Zhang, "Search Gaps," *SSRN Electronic Journal*, feb 2020.

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.

Varian, Hal R, "A Model of Sales," *The American Economic Review*, 1980, *70* (4), 651–659.

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

Yang, Liu (Cathy), Olivier Toubia, and Martijn G. De Jong, "A Bounded Rationality Model of Information Search and Choice in Preference Measurement," *Journal of Marketing Research*, apr 2015, *52* (2), 166–183.

Appendix

A. 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 so far, 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

Figure A.1: Belief Updating

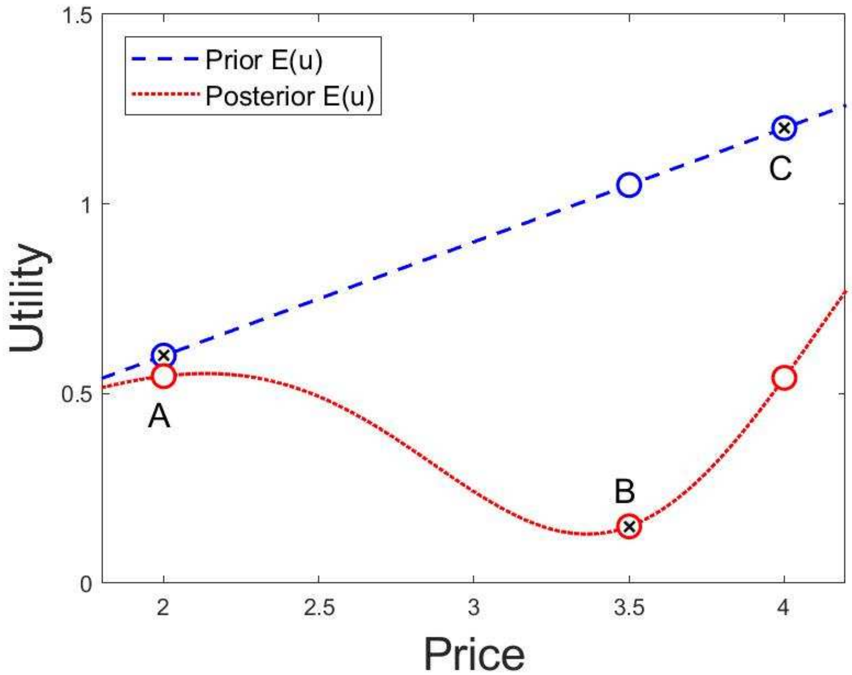

Notes: Black crosses indicate the location of the three products, A, B,and C, in price-utility space. The blue dashed line is the consumer's prior expected utility of hypothetical products at different price levels. This is given by $E(u_j) = (\mu - 1)p_j$. The red dashed line indicates the consumer's 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.

product can affect consumers' beliefs about other products. *Product rankings can therefore be used 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. MCMC Estimation Algorithm

We estimate the model by drawing from the posterior distributions given by equations 17, 18, and 19.

$$P(\tilde{\psi}|\sigma_\xi, \boldsymbol{\xi}) \propto L(\tilde{\psi}, \sigma_\xi, \boldsymbol{\xi}) \tag{17}$$

$$P(\sigma_\xi|\tilde{\psi}, \boldsymbol{\xi}) \propto L(\tilde{\psi}, \sigma_\xi, \boldsymbol{\xi}) k\left(\sigma_\xi; 1+J, \frac{1+\sum \xi_j^2}{1+J}\right) \tag{18}$$

$$P(\xi_j|\tilde{\psi}, \sigma_\xi, \{\xi_k\}_{k\neq j}) \propto L(\tilde{\psi}, \sigma_\xi, \boldsymbol{\xi}) \phi\left(\frac{\xi_j}{\sigma_\xi}\right) \tag{19}$$

To draw from these posteriors we use a Metropolis-Hasings in Gibbs sampler (Gelman, Carlin, Stern, Dunson, Vehtari and Rubin 2013, Chapter 11.3). The algorithm proceeds fom a set or starting values $(\tilde{\psi}, \sigma_\xi, \boldsymbol{\xi})$. The tth iteration of the algorithm is as follows.

1. **Draw Parameters:** Sample proposal $\tilde{\psi}' \sim N(\tilde{\psi}_t, \Omega_\psi)$. Draw $\alpha \sim U[0, 1]$. If $\frac{P(\tilde{\psi}'|\sigma_{\xi t}, \boldsymbol{\xi}_t)}{P(\tilde{\psi}_t|\sigma_{\xi t}, \boldsymbol{\xi}_t)} > \alpha$, set $\tilde{\psi}_{t+1} = \tilde{\psi}'_t$. Otherwise set $\tilde{\psi}_{t+1} = \tilde{\psi}_t$.

2. **Draw Product Effect Variance:** Sample proposal $\sigma'_\xi \sim N(\sigma_{\xi t}, \omega_\sigma)$. Draw $\alpha \sim U[0, 1]$. If $\frac{P(\sigma'_\xi|\tilde{\psi}_{t+1}, \boldsymbol{\xi}_t)}{P(\sigma_{\xi t}|\tilde{\psi}_{t+1}, \boldsymbol{\xi}_t)} > \alpha$, set $\sigma_{\xi t+1} = \sigma'_\xi$. Otherwise set $\sigma_{\xi t+1} = \sigma_{\xi t}$.

3. **Draw Product Effects**: Randomly select 10% of the entries of $\boldsymbol{\xi}$ to be updated. For each of these entries j, sample proposal $\xi'_j \sim N(\xi_j, \omega_\xi)$. For the other entries, $\xi'_j = \xi_j$.

Starting with $j = 1$, iterate through entries of $\boldsymbol{\xi}$. For each entry, draw $\alpha \sim U[0,1]$. If $\frac{P(\xi'_j|\tilde{\psi}_{t+1},\sigma_{\xi t+1},\{\xi_{kt}\}_{k<j},\{\xi_{kt+1}\}_{k>j})}{P(\xi_{jt}|\tilde{\psi}_{t+1},\sigma_{\xi t+1},\{\xi_{kt}\}_{k<j},\{\xi_{kt+1}\}_{k>j})} > \alpha$, set $\xi_{jt+1} = \xi'_j$. Otherwise set $\xi_{jt+1} = \xi_{jt}$.

5000 draws from the chain are dropped for burn in, and the reported estimates are the mean and standard deviations of the 15000 draws from three parallel chains. We visually inspect chains for convergence and adjust proposal variances Ω_ψ, ω_σ, and ω_{xi} manually before running the chain to optimized the speed of convergence.

D. Identification Details

Fix the number of products, J, and let the number of consumers grow large, $N \to \infty$. It is clear that conditional search probabilities $P(j(i,t) = j|j(i,1),...,j(i,t-1))$ and purchase probabilities $P(\hat{j}(i) = j|j(i,1),...,j(i,t-1))$are identified for all products j conditional on all possible sequences of products searched $(j(i,1),...,j(i,t-1))$. In particular, the probability that product j is searched first, given by equation 20, is identified.

$$
\begin{aligned}
P(j(i,1),=j) &= \frac{exp(z_j)}{\sum_{k\in J} exp(z_k)} \\
z_j &= \Phi\left(\frac{\alpha + X_j\beta}{\sigma}\right)(\alpha + X_j\beta) + \phi\left(\frac{\alpha + X_j\beta}{\sigma}\right)\sigma
\end{aligned}
\tag{20}
$$

Where z_j is product j's prior search index, which is a non-linear function of three parameters, α, β, and σ, where σ is the total variance in utility (i.e. $\sigma = \sqrt{\lambda^2 + \sigma_\xi^2 + \sigma_\epsilon^2}$). Suppose X is one dimensional and there are three products with $X^A > X^B > X^C$. $P(j_{i1} = A)$, $P(j_{i1} = B)$, and $P(j_{i1} = C)$ define a system of three non-linear equations with three unknown parameters. If these equations have a unique solution, then (α, β, σ) are identified from the first search probabilities. Intuitively, β is identified by the correlation between product attributes and search probability. The identification of α and σ from the first search is less clear, and comes from the non-linear functional form imposed by Gaussian beliefs.[21] Of course, there is additional variation in the data

[21] In discrete choice demand estimation, we usually think that the scale parameter on utility is not identified. Here, σ is different from the variance of the logit shocks (which is normalized as usual) and is identified because it enters non-linearly in z_j. In particular, as $\sigma \to 0$, $\frac{\partial z_j}{\partial X_j} \to 0$ for $\alpha + X_j\beta < 0$ and $\frac{\partial z_j}{\partial X_j} \to \beta$ for $\alpha + X_j\beta > 0$.

identifying each of these parameters. For example, as $\alpha \to -\infty$, the probability of taking the outside option increases, $P(\hat{j}(i) = 0) \to 1$.

Suppose we can identify (α, β, σ) from the first search. Fix a product A and consider the probabilities of stopping after the first search and purchasing product A. It is clear that $P(\hat{j}(i) = A|j(i,1) = A)$ is increasing in ξ_A. Fixing the other parameters, there is a value of ξ_j which rationalizes the purchase probability for each product. As discussed in Section X, products that are frequently purchased relative to others with similar X_j must have higher ξ_j. [22]

Next, consider the second search probabilities $P(j(i,2) = B|j(i,1) = A)$, $P(j(i,2) = C|j(i,1) = A)$, $P(j(i,2) = A|j(i,1) = B)$ etc. For J products there are $J(J-1)$ such probabilities, which depend on $3 + J$ free parameters, $\{\lambda, \rho, \sigma_\epsilon, \{\xi_j\}_{j=1}^J\}$. Recall Proposition 1: if A is the first product searched, then $\frac{\partial z_B}{\partial \xi_A} > \frac{\partial z_C}{\partial \xi_A} > 0$ if $X^A > X^B > X^C$, and thus decreasing ξ_A lowers the probability of searching B relative to the probability of searching C. The extent to which these effects decline with distance depend on the parameters that control covariance, λ and ρ, (see equations X and Y). Notice that there is a close analogue between this source of identification and the step size effects described in Table X: consumers "jump" away from low-ξ_j products, and the size of these jumps depends on the spatial correlation.

The estimation procedure uses the full search path of each consumer, and therefore contains additional identifying variation beyond the choice probabilities discussed here. Furthermore, notice that this argument does not use the assumption that $\xi_{re} \sim N(0, \sigma_{re})$, which is imposed in estimation and implies that the values of ξ_j contain information about σ_{xi}.

As $\sigma \to \infty$, $\frac{\partial z_j}{\partial X_j} \to 0.5\beta$ for all values of X_j. The size of the disconsintinuous change in in $\frac{\partial z_j}{\partial X_j}$ at $\alpha + X_j\beta = 0$ identifies σ and the position of this change is slope identifies α.

[22] Purchase probabilities $P(\hat{j}(i) = A|j(i,1) = A)$ and $P(\hat{j}(i) = 0|j(i,1) = A)$ are also informative about σ_ϵ. Higher values of σ_ϵ add noise to revealed utilities and reduce the correlation between product attributes, X_j and purchase probabilities. Similarly, ξ_j has a larger effect of purchase probability when σ_ϵ is small.

Figure A.2: Policy Approximation

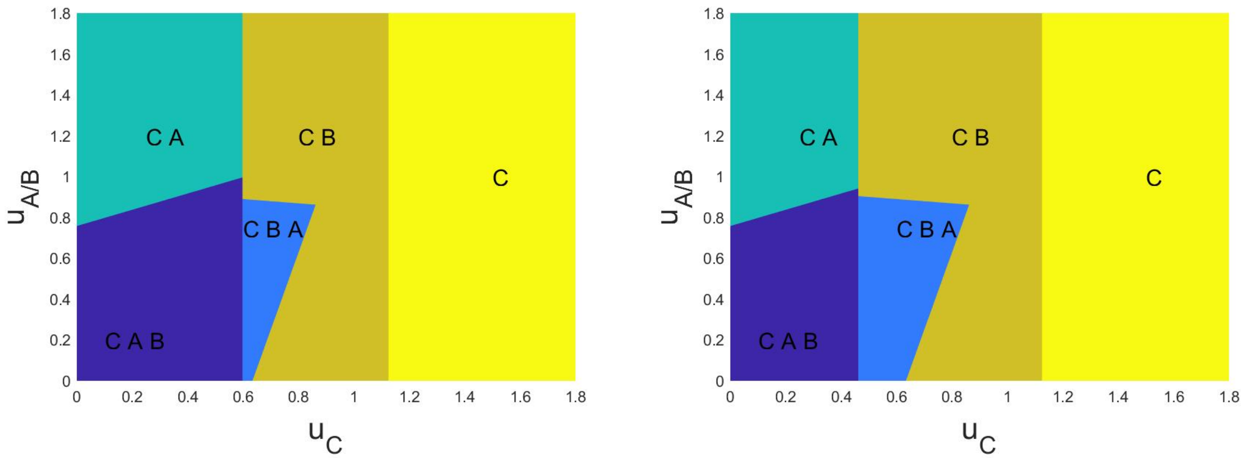

Notes: The left panel shows the optimal policy and the right panel the optimal one-period look ahead policy, in the example from earlier, with $p_A = 2$, $p_B = 3$, and $p_C = 4$. $\mu = 1.5$ $c = 0.45$, and $\Sigma_{ii} = 0.55$, $\Sigma_{AB} = \Sigma_{BC} = 0.15$, and $\Sigma_{AB} = 0.005$.

E. Definition of KL Divergence

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

$$KL(j_1, j_2) = \frac{1}{2} \int \left(\left(\Sigma_0^{-1} \Sigma_1 \right) + (\boldsymbol{\mu_0} - \boldsymbol{\mu}_1)' \, \Sigma_0^{-1} \left(\boldsymbol{\mu_0} - \boldsymbol{\mu}_1 \right) + \ln \left(\frac{\det \Sigma_0}{\det \Sigma_1} \right) \right) dF_0(u_{j_1}, u_{j_2}) \tag{21}$$

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.

F. Additional Tables and Figures

Figure A.3: Covergence to Chosen Attribute Level

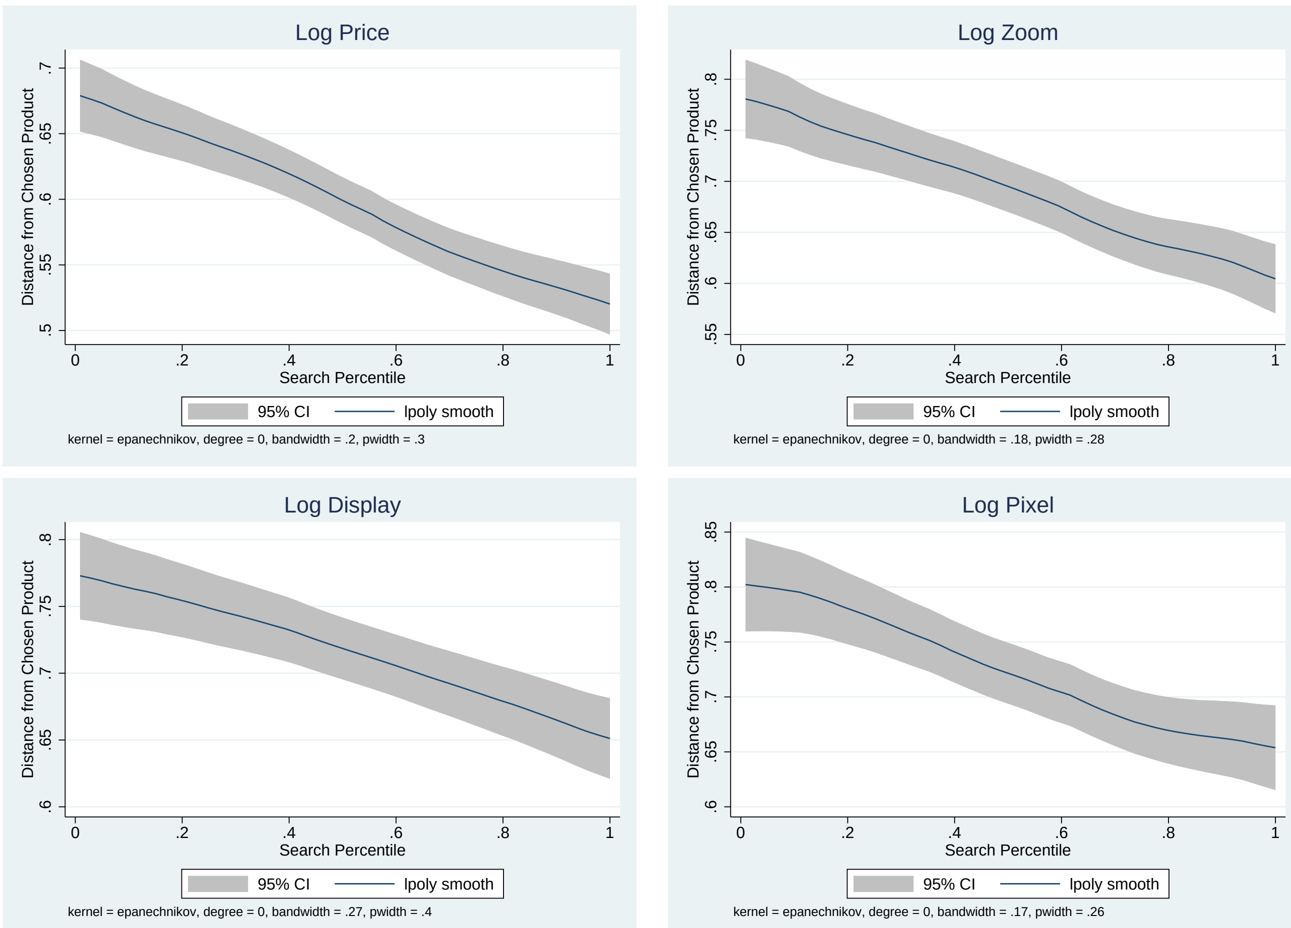

Notes: The y-axis for each panel records, for the relevant product attribute, the absolute difference in standard deviations of the attribute between the searched product and the product ultimately purchased. The x-axis reports the search percentile, as defined in the text. The product ultimately purchased 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.

Figure A.4: Covergence to Chosen Attribute Level: No Revisits

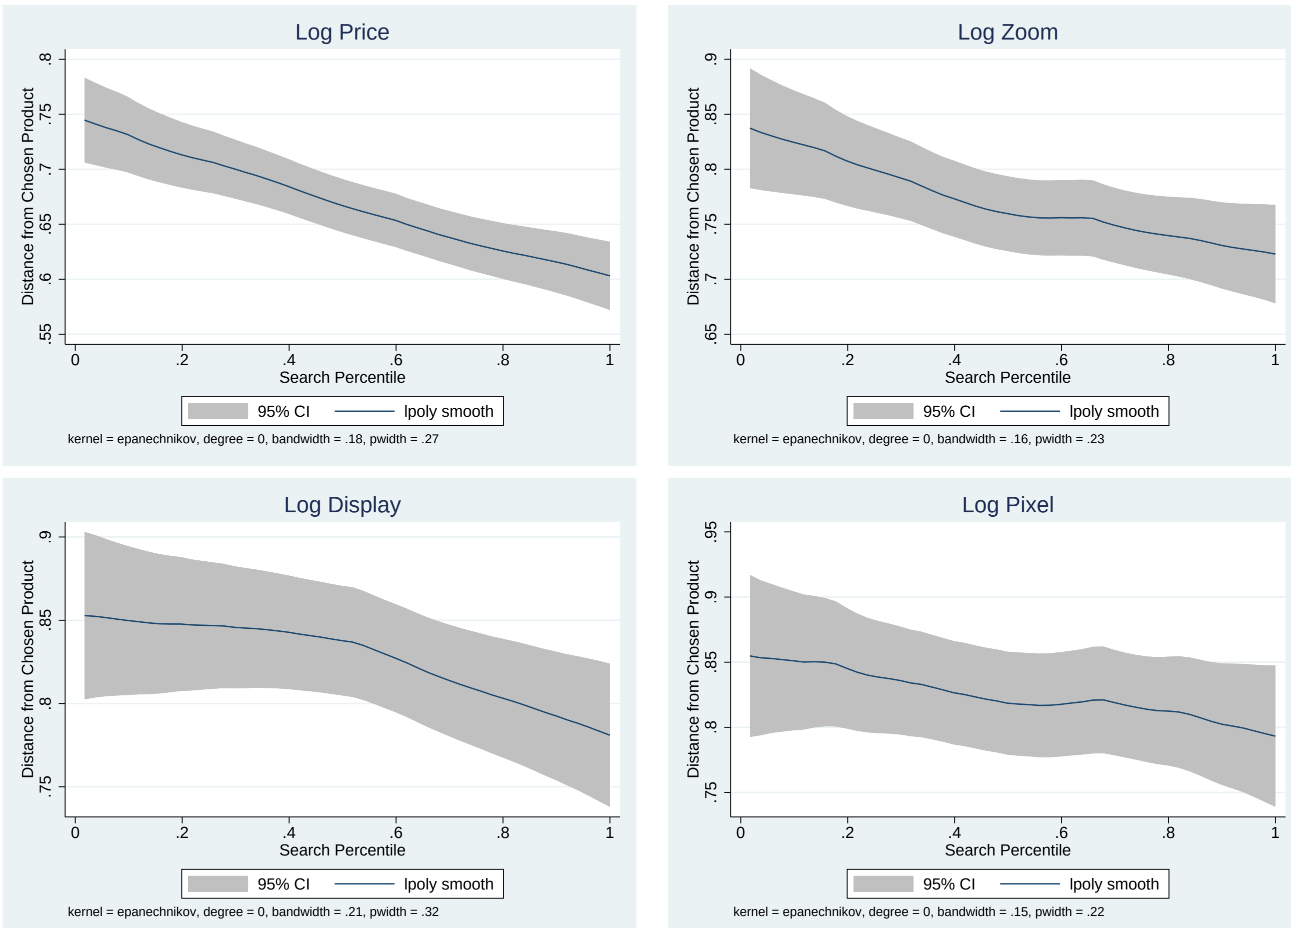

Notes: The y-axis for each panel records, for the relevant product attribute, the absolute difference in standard deviations of the attribute between the searched product and the product ultimately purchased. The x-axis reports the search percentile, as defined in the text. The product ultimately purchased 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.5: Covergence to Chosen Attribute Level: Simulations

Baseline Parameter Estimates

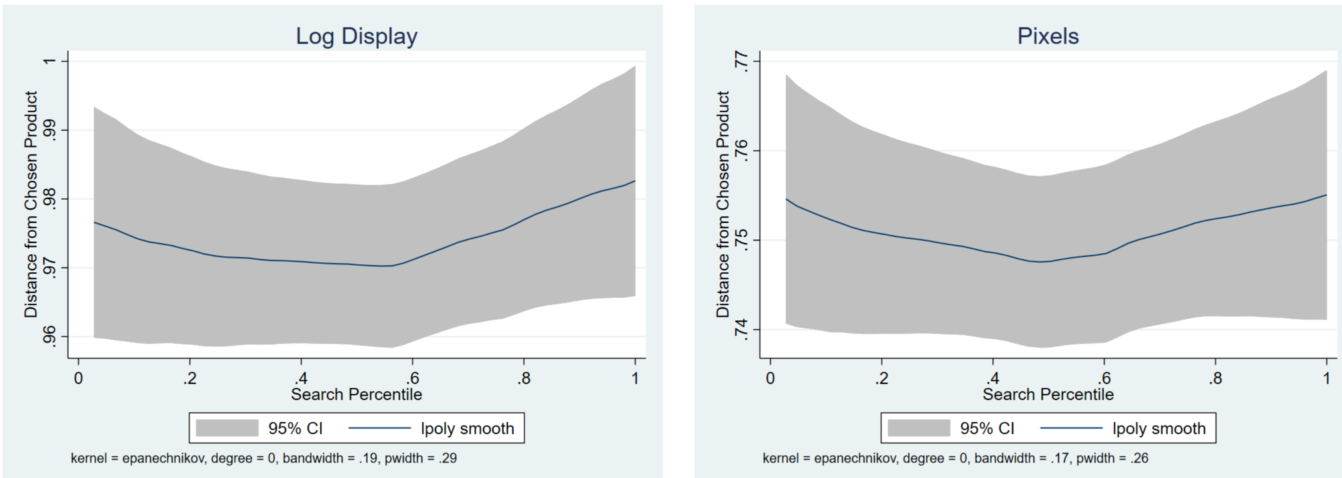

$\lambda = 0$ Parameter Estimates

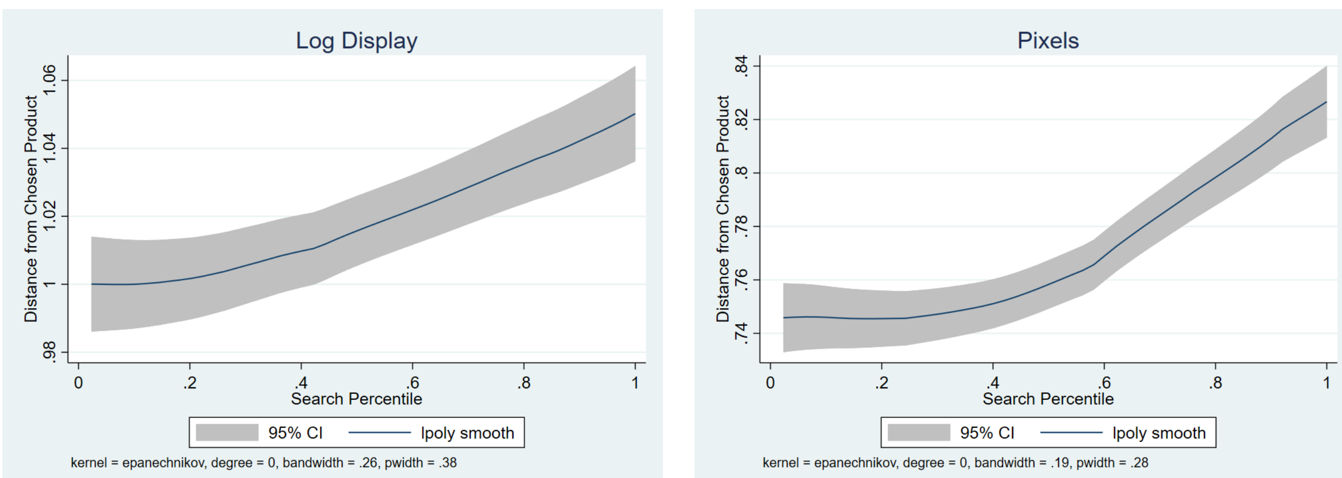

Notes: Figures are constructed using 15,000 search paths simulated at the estimated parameters. The top two panels use the baseline estimates, and the bottom two panels use the estimates under the restriction the $\lambda = 0$. The y-axis records, for the relevant product attribute, the absolute difference in standard deviations of log price and log zoom 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.2: Effect of Product Residuals on Step Size: No Revisits

	$\|price_{it} - price_{it-1}\|$	$\|pixel_{it} - pixel_{it-1}\|$	$\|zoom_{it} - zoom_{it-1}\|$	$\|display_{it} - display_{it-1}\|$
$\hat{\xi}_{j(i,t-1)}$	-.035	-.213***	-.077**	-.154***
	(.024)	(.037)	(.034)	(.039)
$SearchPercentile_{it}$	-.044	-.069	.044	-.024
	(.039)	(.059)	(.055)	(.061)
$Purchased_{it}$	-.097**	.001	-.032	.049
	(.046)	(.069)	(.064)	(.071)
$ProductDensity_{it}$	.172***	2.946***	.365***	22.838***
	(.012)	(.199)	(.028)	(1.214)
N	3385	3385	3385	3385
Consumer FE	Yes	Yes	Yes	Yes
Mean of Dep. Var.	.670	.897	.784	.934

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. 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.

Table A.3: Effect of Rarely Purchased Product on Step Size

	$\|price_{it} - price_{it-1}\|$	$\|pixel_{it} - pixel_{it-1}\|$	$\|zoom_{it} - zoom_{it-1}\|$	$\|display_{it} - display_{it-1}\|$
$BadProduct_{it-1}$	.078**	.000	.183***	.199***
	(.040)	(.061)	(.055)	(.061)
$SearchPercentile_{it}$	-.139***	-.120***	-.089**	-.094**
	(.026)	(.040)	(.036)	(.040)
N	6526	6526	6526	6526
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 10 consumers in the data and purchased with probability less than 5% 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.4: Placebo Tests: Leads and Lags of Product Residuals

	$\hat{\xi}_{j(i,t-3)}$	$\hat{\xi}_{j(i,t-2)}$	$\hat{\xi}_{j(i,t-1)}$	$\hat{\xi}_{j(i,t)}$	$\hat{\xi}_{j(i,t+1)}$
$\lvert price_{it} - price_{it-1}\rvert$	.015	-.002	-.064***	.003	.044**
	(.022)	(.020)	(.019)	(.019)	(.020)
$\lvert pixel_{it} - pixel_{it-1}\rvert$	.001	.025	-.274***	-.011	.009
	(.033)	(.030)	(.029)	(.029)	(.030)
$\lvert zoom_{it} - zoom_{it-1}\rvert$	.032	-.012	-.076***	-.004	-.002
	(.030)	(.028)	(.026)	(.026)	(.027)
$\lvert display_{it} - display_{it-1}\rvert$	.018	.035	-.280***	-.043	.022
	(.033)	(.031)	(.030)	(.029)	(.030)

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.5: Estimated Parameters with $\lambda = 0$

	Estimate	SE		Estimate	SE
β_1 (log price)	-0.667	0.071	α	-20.248	0.535
β_2 (log zoom)	0.008	0.070	c	6.798	0.098
β_3 (log pixels)	3.312	0.118	σ_ξ	20.226	0.506
β_4 (log display)	0.231	0.106	σ_ϵ	3.029	0.119

Notes: Table reports estimated parameters under the restriction $\lambda = 0$ and standard errors. Estimation uses the MCMC procedure described in Section 5. 5000 draws from the chain are dropped for burn in, and the reported estimates are the mean and standard deviations of the 15000 draws from three parallel chains. For more details on the estmiation procedure, see Appendix B.

Table A.6: KL Divergence and Recommended Product Location

Dependent Variable:	KL
$Distance$	0.017***
	(0.001)
$InvDensity$	-0.013***
	(0.002)
N	1000

Notes: Regression of $KL(j_1, j_2)$ on $Distance(j_1, j_2)$ and $InvDensity(j_1, j_2)$, as defined in the text. An observation in the regressions is a simulation. Each simulation generates 1,000 paths at the estimated parameter values after consumers observe the utility from two products, j_1 and j_2, drawn randomly as described in the text. for each column, 1,000 pairs of locations, (j_1, j_2) are drawn, and the statistics used in the regression are computed from the resulting search paths.