

MIT Open Access Articles

Recommending Products When Consumers Learn Their Preference Weights

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Dzyabura, Daria and John R. Hauser. "Recommending Products When Consumers Learn Their Preference Weights." *Marketing Science* 38, 3 (May 2019): 365-541 © 2019 INFORMS

As Published: <http://dx.doi.org/10.1287/mksc.2018.1144>

Publisher: Institute for Operations Research and the Management Sciences (INFORMS)

Persistent URL: <https://hdl.handle.net/1721.1/124909>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Recommending Products When Consumers Learn Their Preference Weights

by

Daria Dzyabura

and

John R. Hauser

April 2018

Daria Dzyabura is an Assistant Professor of Marketing, Leonard N. Stern School of Business, New York University, Tisch Hall, 40 West Fourth Street, 805, New York, NY 10012, (212) 992-6843, ddzyabur@stern.nyu.edu.

John R. Hauser is the Kirin Professor of Marketing, MIT Sloan School of Management, Massachusetts Institute of Technology, E62-538, 77 Massachusetts Avenue, Cambridge, MA 02139, (617) 253-2929, hauser@mit.edu.

Recommending Products When Consumers Learn Their Preference Weights

Abstract

Consumers often learn the weights they ascribe to product attributes (“preference weights”) as they search. For example, after test driving cars, a consumer might find he or she undervalued trunk space and overvalued sunroofs. Preference-weight learning makes optimal search complex because, each time a product is searched, updated preference weights affect the expected utility of all products and the value of subsequent optimal search. Product recommendations, that take preference-weight learning into account, help consumers search. We motivate a model in which consumers learn (update) their preference weights. When consumers learn preference weights, it may not be optimal to recommend the product with the highest option value, as in most search models, or the product most likely to be chosen, as in traditional recommendation systems. Recommendations are improved if consumers are encouraged to search products with diverse attribute levels, products that are undervalued, or products where recommendation system priors differ from consumers’ priors. Synthetic data experiments demonstrate that proposed recommendation systems outperform benchmark recommendation systems, especially when consumers are novices and when recommendation systems have good priors. We demonstrate empirically that consumers learn preference weights during search, that recommendation systems can predict changes, and that a proposed recommendation system encourages learning.

Keywords: Recommendation systems, learned preferences, multiattribute utility, consumer search

1. Introduction

Our basic premises are that (1) in some product categories consumers learn their preference weights by searching products and (2) many recommendation systems can anticipate a consumer's (true) preference weights better than novice consumers. (By preference weights, we refer to the weights that a consumer places on an attribute level in an additive multiattribute utility function.) We argue that premise #1 changes a consumer's optimal search path and that premises #1 and #2 change the recommendation that a recommendation system should make.

1.1. Motivation of Premise #1: Preference-Weight Learning

Consider Candace and Dave who were moving to a new city. They wanted a condominium with two bedrooms, two bathrooms, a good school district, hardwood floors, an island kitchen, adequate lighting, proximity to work, and a full-service concierge. They had full access to MLS listings and could search attributes of condominiums easily, but sought a realtor's recommendations. Based on a realtor's recommendation, they visited one condominium that had a playground across the street. Seeing the playground, Candace and Dave realized how convenient this feature would be for them. Although they always valued playgrounds, they had not previously considered proximity to playgrounds to be an important decision criterion for a home. Their preference weight for playgrounds was substantially larger after seeing the condominium near a playground than before. As they decided what to search further, playground proximity was weighted more heavily. Playground proximity influenced their selection of condominiums to search and their choice of which condominium to buy.

Preference-weight learning applies broadly, for example, to a high school student learning about what to value in undergraduate research programs during college search, consumers learning about new features as they search for new automobiles, first-time parents learning how nannies affect their lifestyle, and even many singles dating and searching for partners. See Cook (2012), Finkel, et al. (2012), and Sheehy (2013) for further examples. The common thread in all examples is that novice consumers, the kind most likely to seek recommendations, revised their preference weights after carefully evaluating products, services, or people during costly search. Consumers updated prior beliefs about preference weights after search and used the updated preference weights for subsequent search, and for choice at the end of search. Preference-weight learning appears common even in information-rich environments where consumers have easy access to attribute levels for most products in the market.

Premise #1 is different from the vast literature in marketing on search with Bayesian learning because consumers learn their preference weights during search, rather than simply learning the utility of products, services, or people. This difference is important because learned (or updated) preference

weights affect the relative utility of every product that has (or doesn't have) the corresponding attribute level. In contrast, in most analytical models, when consumers learn product utilities by searching, the learning affects primarily the utility of the product that is searched.

1.2. Motivation for Premise #2: Recommendation System Knowledge

In our vignette, Candice and Dave received a recommendation from a realtor. The realtor was sufficiently experienced to know that a young couple with small children would value playground proximity, and observed that Candice and Dave were not putting sufficient weight on it. The realtor had a better understanding of Candice and Dave's true preferences than did Candice and Dave at the start of the search. The phenomenon is common. For example, Rogers (2013, page F4) suggests: "Often people don't know what they want. [...] You may think you want X, but if you're shown Y, you may love Y better than you ever loved X. [...] Even (or especially) in these days of consumer online access, some of an agent's value lies in her being able to offer a buyer a choice different from his preconception." College counselors help high school students learn what to value in colleges, childcare agencies help new parents learn what to value in nannies, and automotive websites help consumers learn what to value in new automobile purchases. A recommender often has knowledge of the attribute levels of products on the market, but also knowledge of how attribute levels influenced consumers' purchases in the past. This knowledge translates to insight on the tradeoffs consumers make among attribute levels. Good recommenders use this knowledge to guide future consumers' search.

Depending on the availability of information on product attributes, a recommender's value lies in helping the consumer learn product attributes and/or helping the consumer learn preference weights, or both. Attribute-level search is well-studied. We focus on preference-weight learning recognizing that many recommenders do both.

Existing automated recommendation systems value the ability to accurately predict preferred products. Machine learning methods such as collaborative filters, content-based filters, and Bayesian-update systems learn preference weights from past users and apply them to new users (review in §2). When preference weights are relatively homogeneous or when preference weights can be tied to observable characteristics, recommendation systems can efficiently learn consumers' true preference weights. While automated recommendation systems use the knowledge of consumers' preference weights to make recommendations, we know of no automated recommendation system that recommends products that help consumers learn their own preferences.

1.3. Product Categories for which Preference Weight Learning is Likely to be Relevant

We expect preference-weight learning to be relevant in product or service categories in which

the product or service is multiattributed, infrequently purchased, costly to experience without purchase or extensive search, and sufficiently valuable to justify extensive and costly search. The set of categories includes real estate, colleges, automobiles, childcare, mates, vacations, furniture, sailboats, or even high-cost industrial equipment. Henceforth, we use “product” to refer to all relevant categories. In these categories, consumers routinely seek advice from human recommenders and/or automated recommendation systems. Anecdotes abound to suggest that human recommenders take preference-weight learning into account in these categories. We hope to extend that capability to automated recommendation systems.

1.4. Overview of Model and Results

We assume the consumer learns (updates) his or her beliefs about the preference weights for attribute levels after searching a product with the corresponding attribute levels. Search is costly and a purchase, if any, occurs at the end of search. The benefit the consumer derives from purchase, if any, is based on the consumer’s true utility which he or she experiences from purchase and consumption. In searching optimally, the consumer weighs his or her vision of this anticipated utility (possibly discounted) against anticipated search cost. To focus on preference-weight learning, we assume that the information about attribute levels of all products is available at low cost from sources such as multiple-listing services (real estate), Kelly Blue Book (automobiles), or popular press ratings (colleges).

We demonstrate that existing search-theory solutions may not be optimal when preference weights are learned, even when the consumer searches without a recommendation. For example, it is not always best to search products with high option values. We next introduce a recommendation system and argue that the typical criterion used to evaluate recommendation systems, highest predicted utility, does not always identify recommendations that maximize a consumer’s net utility when preference-weights are learned. Instead, a knowledgeable recommendation system should recommend products with diverse attributes (thus providing a theoretical explanation for recent trends) and/or recommend products that are undervalued by the consumer. Such products are effective recommendations even if the products have a low probability of being the chosen product, because the recommendations make the consumer’s subsequent search more efficient. We use numerical examples to illustrate how benevolent recommendation systems can direct the consumer to the optimal-net-utility choice. We also illustrate non-benevolent recommendation systems which lead consumers to profitable products that may not be best for the consumer. Our arguments and examples suggest practical preference-weight-learning recommendation systems. We test the proposed recommendation systems (and a standard benchmark) with synthetic data, and explore when each system is most advantageous. We close the

paper with an empirical test to demonstrate that our premises are reasonable and that a proposed recommendation system helps consumers learn preference weights better than either the standard benchmark recommendation system or allowing consumers to choose without a recommendation.

2. Related Literatures

We build on two related literatures: the recommendation-system literature, mostly from computer science, and the sequential-search literature, mostly from economics and marketing.

2.1 Literature on Recommendation Systems

Traditionally, the primary goal of a (Top- N) recommendation system is to recommend N items that maximize a user's utility (Adomavičius and Tuzhilin 2005). Typically, the recommendation system observes a utility surrogate, a rating or a rank, for some users and some items and attempts to extrapolate the surrogate to all users and items. As a result, most recommendation systems are evaluated on the accuracy of that extrapolation (Herlocker, et al. 2004; Liu, Zhao, Xiang, and Yang 2010; McNee and Konstan 2006; Vargas and Castells 2011; Zhang and Hurley 2008). This focus was most notable in the \$1M Netflix Challenge that began in 2006 and finished in 2009. The Netflix Challenge sought the recommendation system algorithm that best predicted held-out user ratings. Successful recommendation systems focus on similarities among users (collaborative filters), similarities among items (content-based filters), attribute-based utility models or hybrids to recommend products with high expected utility (Adomavičius and Tuzhilin 2005; Jacobs, Donkers, and Fok 2016; Moon and Russel 2008; Urban and Hauser 2004). Although some recommendation systems attempt to match attribute-based utility, attributes are typically defined with taxonomies such as genre (Ansari, Essegaiar, and Kohli 2000). A related literature in marketing uses attribute-level preferences to predict whether consumers will choose a recommended product (Chung and Rao 2012; De Bruyn, et al. 2008; Ghose, Ipeirotis, and Li 2012; Häubl and Trifts 2000; Lu, Xiao, and Ding 2016; Ying, Feinberg, and Wedel 2006). Although both the recommendation-system literature and the marketing literature demonstrate that recommendation systems can learn consumer preference weights well, many authors have criticized the focus on predictive accuracy as, in practice, providing recommendations that are too similar to previously purchased items, for example, recommending the same author after a book is purchased (Fleder and Hosanagar 2009; McNee and Konstan 2006; Zhang and Hurley 2008), or too obvious, for example recommending bread, milk, eggs, and bananas to all grocery store shoppers (Herlocker et al 2004).

In response, researchers have proposed algorithms, and metrics to evaluate those algorithms, that include goals that complement predictive accuracy (Bodapati 2008; Herlocker et al. 2004). New algorithms avoid recommending items that the consumer would have bought without a recommendation.

They augment predictive accuracy with diversity, novelty, and serendipity (Adamopoulos and Tuzhilin 2014; Celma and Herrera 2008; Castells, Vargas, and Wang 2011; Ge, Delgado-Battenfeld, and Jannach 2010; Vargas and Castells 2011; Zhou, et al. 2010; Ziegler, et al. 2005). Diverse items are items that are not similar to one another; novel items are items the consumer would not have chosen without a recommendation; serendipitous items are items that are unexpected, relevant, and useful. To achieve these goals, recommendation systems penalize recommendations that are similar to “accurate” recommendations or recommend products from the “long tail.” Product attributes, when used, are used to define product similarity metrics. Diversity, novelty, and serendipity are based on products, not the levels of attributes of the products. We augment this literature by studying how a recommendation system might incorporate preference-weight learning when making recommendations. Our results challenge the traditional recommendation-system focus on products that are likely to be chosen or likely to have high utility. On the other hand, our analyses provide a theoretical explanation for why diversity, novelty, and serendipity work better in practice than pure predictive accuracy. We suggest modifications to recommendation systems that are based on a reinterpretation of these concepts and highlight when such modifications benefit consumers.

2.2. Literature on Sequential Search

Papers in marketing and economics recognize the importance of the consumer’s search for information, and have studied it empirically and theoretically. For example, Kim, Albuquerque, and Bronnenberg (2010) use Amazon’s view-rank data to infer consumer preferences for observed attributes of camcorders. Consumers know these attribute levels without search, but consumers search to resolve the unobserved utility of the products (error term). Bronnenberg, Kim, and Mela (2016) study observed online search and find that search is over a relatively small region of attribute space that declines with subsequent search. The final choice is rarely the first item searched. Chen and Yao (2016), Hong and Shum (2006), Honka (2014), and Seiler (2013) analyze search-path data to infer price distributions and/or search costs. Although a few authors consider non-sequential search, Bronnenberg, Kim, and Mela (2016) report strong evidence to support sequential search.

Much of this literature is based on theory derived by Weitzman (1979), who studied search over products whose utilities are independently distributed. Weitzman derives the optimal search strategy for this model, which is based on an option value index—the upper tail of the utility distribution. The optimal strategy is to search the products with the highest indices as long as they are above the reservation value. See extensions by Adam (2001) and Bikhchandani and Sharma (1996). Branco, Sun, and Villas-Boas (2012) focus on the optimal search for multiple attributes of a single product. The optimal

strategy in this setting is also index-based: the consumer searches as long as utility is bounded between purchase and not-purchase thresholds. Ke, Shen, and Villas-Boas (2016) extend the model to derive appropriate bounds for two products.

Our analyses are consistent with this literature in the sense that the consumer’s optimal search path is the solution to a dynamic program (a Bellman Equation). However, we modify the recursion to allow consumers to update their preference weights for attributes as they search. Because products share attribute levels, the optimal search strategy is no longer indexable (e.g., Weitzman’s solution). High option value or high variance in product utility matters less; strategies to learn preference weights efficiently matter more.

Finally, our model of preference-weight learning is consistent with examples of preference-weight learning in the marketing science literature. Greenleaf and Lehmann (1995) demonstrate that consumers delay purchases to learn preference weights, and She and MacDonald (2013) show that “trigger features” cause consumers to update preference weights. Hauser, Dong, and Ding (2014) show that, as consumers become more expert, their preference weights stabilize. Predictions, even one-to-three weeks later, improve. Dzyabura, Jagabathula, and Muller (2017), and Dzyabura and Jagabathula (2017) demonstrate that preference weights change when consumers evaluate physical products rather than their online descriptions. Most of these changes persist if consumers go back to the online channel after evaluating physical products, which is consistent with learning preference weights.

3. Model of Consumer Search with Preference-Weight Learning

To model preference-weight learning, we decompose product utilities into components corresponding to product attribute levels and allow consumers to learn their preference weights for attribute levels as they search products. We start by defining the utility of a product and then present the model of search.

3.1. Consumer Utility is Defined on Attributes Levels

Let $j = 1, \dots, J$ index products and $i = 1, \dots, I$ index attributes. For ease of exposition, we begin with binary attributes such that a product either has or does not have an attribute. When it is clear in context, we refer to binary attributes simply as attributes. Later in §4.1.2, when we discuss multilevel attributes, we introduce terminology to distinguish attribute levels from attributes. Let $x_{ij} = 1$ if product j has binary attribute i and let $x_{ij} = 0$ otherwise. Let \vec{x}_j be the binary vector that describes product j .

Let u_j be the utility of product j and let w_i be the relative preference weight that the consumer

places on attribute i such that:

$$(1) \quad u_j = u(\vec{x}_j) = \sum_{i=1}^I w_i x_{ij}.$$

Let $\vec{w} = [w_1, \dots, w_I]$ be the vector of preference weights for all the attributes. The consumer has a prior belief about the values of w_i . The prior probability density for the consumer's prior beliefs is denoted by $f_i^0(w_i)$ for each attribute i . This belief can be updated when the consumer observes a product with attribute i . We assume the prior distributions (and any updated distributions) are independent over i .

To focus on preference-weight learning, we assume consumers know, or can search at negligible cost, whether or not a product has an attribute. That is, we assume they know \vec{x}_j . This simplification is not unrealistic. Zillow, Trulia, or multiple-listing services (MLS) provide attribute levels of new homes; *U.S. News & World Reports* and *Business Week* provide attribute levels for colleges; Autotrader, Edmunds, and Kelly Blue Book provide attribute levels for automobiles; and travel websites, dating websites, and Amazon provide attribute levels for other products. We focus on situations where attribute levels are easy to observe, but more-costly search is necessary for consumers to experience attributes and learn their preference weights. We revisit this assumption and the assumption of discrete attribute levels in §9.2.

3.2. Consumers Learn Preference Weights during Search

Consumers engage in sequential search, searching one product at a time. Searching represents sufficient effort by a consumer to examine and evaluate product j , for example, by test driving a car, visiting a condominium for sale, visiting a college, or interviewing a caregiver. The cost of searching a product is $c > 0$. There is a true value of w_i , which we label w_i^r . (r for revealed.) During search, the consumer updates his or her beliefs about w_i toward w_i^r . For infrequently-purchased products (as in §1.3), the consumer fully learns his or her preference weights when consuming the chosen product, thus the utility the consumer ultimately gets is computed according to w_i^r , that is, $\sum_{i=1}^I w_i^r x_{ij}$.

Let t index sequential searches and let s_t be the t^{th} product search. Whenever $x_{ij}^t = 1$, the consumer receives a signal about the true value of his or her preference weight. This signal takes the form of a probability density function, $g(w_i^r | w_i, s_t = j)$, for the true value, w_i^r . Using Bayes Theorem, the consumer's evaluation during search of product j enables the consumer to update beliefs about w_i . If $x_{ij}^t = 0$, the consumer cannot update beliefs about w_i . Let $F_t(j = s_t) = \{i: x_{ij}^t = 1\}$, then the consumer's belief distribution about preference weight i after the t^{th} search is:

$$(2) \quad f_i^{t+1}(w_i) \equiv f_i^{t+1}(w_i | w^r, s_t = j) = \begin{cases} f_i^t(w_i) & \text{for all } i \notin F_t(j) \\ \frac{g(w_i^r | w_i, s_t = j) f_i^t(w_i)}{\int g(w_i^r | w_i, s_t = j) f_i^t(w_i) dw_i} & \text{for all } i \in F_t(j) \end{cases}.$$

We refer to $f_i^t(w_i)$ as the prior beliefs before the t^{th} search and $f_i^{t+1}(w_i)$ as the posterior beliefs after the t^{th} search. Equation 2 allows the Bayesian updating to remain general. For the numeric examples and synthetic data experiments we make a functional assumption—we use normally-distributed priors leading to normally-distributed posteriors. Appendix 2 summarizes the formulae.

As t increases, we expect the mean of $f_i^t(w_i)$ to approach w_i^r and the variance to approach zero. The rate at which the posterior beliefs converge can vary. For example, one of our colleagues was dead set against swimming pools while searching for vacation homes in a beach community, but when she saw the perfect swimming pool layout for her children she immediately updated her preference weights and bought the house. Another colleague searched three homes in a community before he fully updated his preference weight for a neighborhood-owned swimming pool. We explore differences in the rate at which preferences are updated in our numerical examples and synthetic data experiments.

3.3. Optimal Search

If no recommendations are made, the consumer searches sequentially and optimally. If recommendations are made, the consumer searches all recommended products, updates his or her preferences, and searches optimally thereafter. The consumer is forward looking; therefore, the consumer solves a dynamic programming recursion to select the next product to search, or to stop and purchase. The state is the set of products already searched, S_t , and the beliefs prior to search, $f^t(\vec{w})$. When prior distributions are independent over i , and when an observation of an attribute only updates preference weights for that attribute, posterior beliefs are independent over attributes. Probabilistic independence implies: $f^t(\vec{w}) = \prod_{i=1 \text{ to } I} f_i^t(w_i)$.

If $J(S_t, f^t)$ is the continuation value, the Bellman Equation for search without recommendations recognizes that this value is the maximum over choosing the outside option, denoted by U^* , choosing the maximum utility product without searching based on f^t , or continuing to search. The value of continuing to search is the maximum over all unsearched products taking into account that preferences will be updated through further search (if further search is optimal). Expectations are based on f^t , which is the consumer's belief about his or her preference weights when the search decision is made. The resulting Bellman Equation is:

$$(3) \quad J(S_t, f^t) = \max \left\{ U^*, \max_{j=1 \text{ to } J} E[\sum_{i=1 \text{ to } I} w_i x_{ij} | f^t], \max_{k \notin S_t} \{-c + \beta E[J(S_t \cup \{k\}, f^{t+1} | f^t)]\} \right\},$$

where $\beta \leq 1$ represents the discount factor. A key part of this optimization problem, relative to the search-theory literature, is that the belief distribution, f^t , is part of the state space and is defined over product attributes (rather than products). Naturally, Equation 3 could be extended to include uncertainty over the \vec{x}_j 's for research to marry preference-weight learning to attribute-level search.

3.4. Recommendations

We model product recommendations as follows. A recommendation is a single product, which is recommended to the consumer at a given time. We assume that the consumer follows that recommendation and searches the recommended product. A human recommender might be a real estate agent offering to show the home buyer a particular property, or a childcare agency scheduling interviews for parents with potential caregivers. We are particularly interested in automated recommendation systems. After searching the recommended product, the consumer continues to search optimally given his or her beliefs, according to Equation 3.

For simplicity, we prefer to make the assumption that the consumer searches the recommended product, but the assumption also can be motivated with an assumed utility bonus ($U_{bonus} > 0$) to represent a belief by the consumer that the recommender knows something about the recommended product (not known to the consumer). The consumer trusts that, for some reason to be revealed during search, consumption utility is increased by the utility bonus. In other words, the consumer believes that the expected utility of recommended product ℓ is $E[\sum_{i=1 \text{ to } I} w_i x_{i\ell} | f^t] + U_{bonus}$. If the bonus is sufficient, but not so large that the consumer purchases product ℓ without searching, then it is optimal for the consumer to search the recommended product. At the end of search, knowledge gained by searching overwhelms any ephemeral utility bonus. The consumer's consumption utility is based on \vec{w}^r .

3.5. Preference-Weight Learning versus Attribute-Level Learning

Equation 3 is reminiscent of the Bellman equations used in classical models of optimal sequential search, but with key differences. For example, Equation 3 shares Weitzman's (1979, p. 643) assumption that "the sum of search costs is paid during search, whereas the maximum reward is collected after search has been terminated." However, unlike in Weitzman's model, the updates in Equation 3 do not necessarily reveal the value of the searched product with a single observation. Slower learning is assumed in learning models such as in Chick and Frazier 2012, when each time the consumer searches, the consumer observes a random draw from a distribution, rather than the true parameter value. More crit-

ically, in search models, such as Weitzman, and search with learning models, such as Chick and Frazier 2012, product utilities are independent. Independence is often assumed in search models because independence enables relatively simple search policies called index policies. With an index policy for search, the optimal search strategy continues until the revealed product (or attribute) exceeds an index determined by a simpler Bellman equation. Index strategies break the curse of dimensionality. When independence during search is lost due to preference-weight learning, we know of no optimal or near-optimal index strategies.

Index strategies are also important for a different, but related, set of dynamic programs: multi-armed bandit problems. Multiarmed bandit problems share the property of search with learning problems that multiple alternatives are each described by reward distributions, but multiarmed bandit problems differ because rewards can be obtained in each period rather than only at the end of search. The consumer decides which alternative (“arm”) to try by balancing immediate rewards versus the long-term benefits from learning about the reward distribution in order to choose better in the future. The consumer sequentially tries alternatives, but unlike search problems, the consumer receives a potential payoff every time an arm is pulled, rather than only receiving the reward from the chosen product (or outside option) at the end of search. When the alternatives are independent, Gittins (1979) demonstrated that the optimal policy is to choose the alternative with the largest index. Whittle (1988) extended index policies to restless bandits where the value of the non-chosen alternatives can change independently. Many bandit problems have been shown to be indexable (Gittins, Glazebrook, and Weber 2011), including the partially-observable Markov processes in website and banner morphing (Hauser, et al. 2009). When there is a switching cost between arms, the bandit is said to have memory, and multiple indices might be required (Jun 2004). In some cases, structured interdependence is allowed and greedy policies perform near optimally (Mersereau, Rusmevichientong, and Tsitsiklis 2009). In the marketing literature, multi-armed bandit models have been used to model consumer purchases in CPG categories (Lin, Hauser and Zhang 2014), to optimize advertising creative (Schwartz, Bradlow and Fader 2017), and in revenue management for optimal pricing (Schwartz, Misra, and Abernethy 2017). Research on multi-armed bandits is extensive and ongoing. In time, concepts from this literature might provide new insights to search models with preference-weight learning.

The very nature of preference-weight learning induces interdependence among product utilities—every time a preference weight is updated the utilities of all products with the corresponding attribute level, including those already searched, can change. Index policies are unlikely to be optimal. We know of no simplifying policy that provides an optimal solution for Equation 3, nor for recommendations that

influence a consumer who bases his or her optimal search on Equation 3. Even in simple two-products-two-attribute markets, relationships to determine the optimal policy are complex. The curse of dimensionality applies. (See Appendix 1, which illustrates Equation 3 with and without a recommendation system.)

Interdependence and the curse of dimensionality require that realistic recommendation systems rely on heuristic policies. For example, if all attribute combinations were feasible in a product space with ten attributes at six levels each, a single recommendation would require that we evaluate almost 60 million candidate recommendations. This number would grow to over three quadrillion pairwise recommendations for a Top-2 recommendation system (which recommends two products sequentially). Without any special structure, the memory requirement of the dynamic program in Equation 3 grows exponentially with the size of the problem. It is no surprise that applied recommendation systems rely on heuristics. To be consistent with applications and with the recommendation systems literature, we explore recommendation system modifications that are likely to be feasible in applied situations. These modifications are necessarily heuristic. We develop potential heuristics by examining the structure of consumer search and the implied recommendation policies.

4. Insights About Search and Recommendation

Equation 3 implies that the consumer must make tradeoffs among search costs and the reward to purchasing without search versus the knowledge gained from further search and improved rewards that come from a later purchase. Equation 3 becomes even more challenging when a recommendation system is introduced. The complexities, due to the interdependence imposed by preference-weight learning, requires that we seek qualitative insights toward potential heuristic modifications to recommendation systems. This section provides insights; §7 develops and tests recommendation systems which implement modifications based on the insights.

4.1. Recommendation System Beliefs

A key component of the Candice-and-Dave condominium vignette was that the realtor had beliefs about Candice and Dave’s preference weights and that these beliefs were better (closer to the true preferences) than Candice and Dave’s initial beliefs. Let $f^{rec}(w_i)$ be the recommendation system’s beliefs about the consumers’ preference weights. For some attributes, the recommendation system’s priors may match the consumer’s priors ($f_i^{rec} = f_i^t$), but for other attributes the recommendation system may believe that the consumer beliefs are not accurate ($f_i^{rec} \neq f_i^t$). In this notation, we do not require that the recommendation system know the consumer’s true beliefs, w_i^t , perfectly. The recommendation system’s beliefs may be a probability distribution. However, if the recommendation system is to be val-

uable, $f^{rec}(\vec{w})$ should, in some sense, be closer to \vec{w}^r than is $f^t(\vec{w})$. We quantify “closer” below. (We assume the recommendation system cannot credibly inform the consumer of its beliefs about the consumer’s preference weights; it must do so implicitly by recommending products to the consumer.)

By assumption, if the recommendation system recommends product j at time t , then $s_t = j$. The key difference between the recommendation system and the consumer is that the recommendation system expects the first (or first few) updates to be based on f^{rec} . When the number of potential recommendations is small, as in the example in Appendix 1, the optimal recommendation can be found by exhaustive enumeration. When the number of products is more typical, we must rely on heuristic policies.

4.2. Aspect Diversity

4.2.1. Consumer Search in a Full-Factorial Product Space

Consider first a full-factorial product space in which, for I binary attributes, all 2^I possible combinations of attribute levels are available as products. In Equation 3, the search costs are product-based, not attribute based, thus searching a product j with $x_{ij} = 1$ for I attributes is no more costly than searching products with $x_{ij} = 1$ for only one attribute. On the other hand, if the consumer were to search the product with I attributes present, the consumer would gain more information than he or she would by searching a product with one attribute present. (The consumer can still decide to purchase the one-attribute-present product. Unlike bandit problems, there is no immediate reward during search. Search produces knowledge to improve net consumption utility after a product or the outside option is chosen.) If the optimal solution at $t = 1$ were for the consumer to search a product, then the best product to search is the product j with all I attributes present, $x_{ij} = 1$ for all i . The insight has intuitive appeal. Realtors often recommend that consumers search homes outside of their price range in order to experience and learn about as many attributes as possible, such as proximity to playgrounds, swimming pools, full-service concierge, island kitchens, and media rooms. Automobile dealers often maintain test-drive cars with many “options.” (For interested readers, a formal proof in a two-product-two-attribute product space is available from the authors.)

4.2.2. Full-Factorial Product Spaces are Rare, Especially with Multilevel Attributes

The value of searching a fully-attributed product to learn (update) preference weights raises a conundrum. It seems obvious that the consumer should always begin by searching the product with all I attributes present—no recommendations are needed. But recommendation systems are popular and highly researched. The answer to the conundrum is simple. In most realistic cases, full-factorial product

spaces are not available. Highly featured cars are priced high, and may be outside the consumer’s price range, leading to “sticker shock.” Many possible attributes describe homes—a fully-attributed property is exceedingly rare. The same can be said for caregivers, colleges, jobs, furniture, and dating opportunities.

Recommendations are particularly valuable in product spaces with multilevel attributes in which, by definition, a fully-attributed alternative is not available: a vehicle cannot simultaneously be a sedan, a coupe, a station wagon, a cross-over, a minivan, an SUV, and a truck. Henceforth, to avoid confusion, we follow Tversky (1972) and refer to a binary attribute level as an aspect. A multilevel attribute is then a collection of aspects with the constraint that exactly one of the aspects in an attribute has $x_{ij} = 1$.

4.2.3. Aspect Diversity

Despite the sparsity of full-factorial product spaces, the insights about fully-attributed products are useful. When consumers learn preference weights, it is valuable to search a product that has a diverse set of aspects. This insight is related to recent recommendation-system trends toward diverse, novel, or serendipitous products, and may provide a partial theoretical explanation for those trends. However, aspect diversity is subtly different because it focuses on learning about preferences for aspects rather than for the products themselves.

Recommendation systems provide value by recommending products that improve the consumer’s search, resulting in a higher terminal utility net of search costs, relative to what the consumer would have searched on his or her own. Thus, it is valuable to recommend products with aspects that have a high true preference weight but which the consumer would not otherwise have searched. This value comes into Equation 3 primarily in the continuation value, $J(S_t \cup \{k\}, f^{t+1} | f^{rec})$. If product k has aspects the consumer would not have otherwise searched, the recommendation helps the consumer to make better subsequent search decisions and, ultimately, identify a higher utility product (net of search costs) to purchase and consume. We call a recommendation system modification that implements this modification, “Aspect Diversity.” There are many ways to implement an aspect diversity heuristic. We have found that it is effective to modify the standard recommendation-system criterion (maximize expected utility based on f^{rec}) to include a penalty for recommending aspects the consumer would otherwise have searched. If s_t^* indicates the product the consumer would have searched without a recommendation, then the aspect diversity criterion becomes:

$$(4) \quad \vec{x}_{rec} = \arg \max_{j=1 \text{ to } J} E[\sum_{i=1 \text{ to } I} w_i x_{ij} | f^{rec}] - \lambda \cdot \sum_{i=1 \text{ to } I} x_{is_t^*} \cdot x_{ij}.$$

The parameter, λ , determines how much emphasis the recommendation system places on diverse aspects. When $\lambda = 0$, the recommendation system recommends the maximum-expected-utility product; when $\lambda \rightarrow \infty$, the recommendation focuses primarily on diversity. As is typical in recommender systems applications, λ is a tuning parameter that is best chosen by experience or experimentation.

4.3. Recommendation Systems that Use Information about Expected Preference Weights

Equations 2 and 3 rely on the consumer's priors, f^t , the recommendation system's priors f^{rec} , and the signal distribution, $g(w_i^T | w_i, s_t = j)$. In theory, recommendations might depend upon the full distributions. For example, recommendations may vary depending upon the rate at which the consumer learns preference weights. However, full distributions (f^t , f^{rec} , and g) might be hard for the recommendation system to observe—the recommendation system might be much better at observing or predicting the mean of the distributions. For such cases, we begin with recommendation system modifications that rely only on the expected values of the preference weights and later consider recommendation system modifications based on the full distributions.

4.3.1. Conceptual Example

Consider three condominiums that are similar on all attributes except playground proximity, type of kitchen (traditional, open-concept, island), and type of service (live-in superintendent, full-service concierge). Suppose Candice and Dave's prior beliefs are such that they do not value playground proximity, and prefer an island kitchen in a building with a full-service concierge. Suppose the recommendation system believes that Candice and Dave undervalue playground proximity and correctly value an island kitchen, but overvalue a full-service concierge. Finally, suppose that the condominium stock in Candice and Dave's target market is limited—not all combinations of playground proximity, kitchens, and service are available. Intuitively, the recommendation that would lead to the largest shift in Candice and Dave's beliefs about their preference weights is a condominium, if available, that has playground proximity, but a live-in superintendent. By improving Candice and Dave's understanding of their preference weights, the recommendation may lead them to purchase and consume a condominium with playground proximity. Candice and Dave may learn to forego a full-service concierge and allocate their limited budget to other high-valued aspects. (If nothing else, condo fees are less with a live-in superintendent than a full-service concierge.)

4.3.2. Undervalued Products

The intuitive recommendation in the condominium example is based on both components of the Equation 3, $E[\sum_{i=1}^{to} w_i x_{ij} | f^{rec}]$ and $J(S_t \cup \{k\}, f^{t+1} | f^{rec})$. By recommending a condominium near a playground, the recommendation system improves the expected reward because

$E[w_{pp}x_{pp,j}|f^{rec}] > E[w_{pp}x_{pp,j}|f^t]$ for playground proximity (pp). It does so by shifting the priors towards a high valued attribute level that the consumer currently undervalues. If the mean of f^{rec} is closer to \bar{w}^r than the mean of f^t , then the expected purchase and consumption utility is also improved. The continuation value comes into play because Candice and Dave can search more effectively for the next condominium when further search is optimal. They can do so because their objective function better matches their true utility ($U(S_t \cup \{k\}, f^{t+1}|f^{rec})$ increases). The recommendation is a good recommendation if these gains are greater than the search cost, c .

The suggested modification to a recommendation system is to recommend undervalued products, where “undervalued” implies that the expected utility based on the recommendation system’s beliefs is higher than the expected utility based on the consumer’s priors, that is, $E[\sum_{i=1}^{to} w_i x_{ij}|f^{rec}] > E[\sum_{i=1}^{to} w_i x_{ij}|f^t]$. We implement the heuristic by identifying those products for which the difference between recommendations system’s beliefs about the consumer’s expected utility and the consumer’s beliefs about the expected utility is maximized:

$$(5) \quad \vec{x}_{rec} = \arg \max_{j=1 \text{ to } J} E[\sum_{i=1}^{to} w_i x_{ij}|f^{rec}] - E[\sum_{i=1}^{to} w_i x_{ij}|f^t].$$

We call a recommendation system that implements this modification “Undervalued Products.”

4.4. Differences in the Distributions of f^{rec} and f^t

In some cases, the recommendation system might be able to measure the variance (or the full distributions) of f^{rec} and f^t . If so, we might improve recommendations by allowing the recommendation system to use its knowledge of the preference-weight belief distributions and the signal, not just the means of the distributions. For example, a recommendation system might favor recommending products with aspects for which preference weights are updated more rapidly (less posterior variance) or products with aspects for which the recommendation system has tighter beliefs. We propose two such recommendation-system heuristics.

Option Value Discrepancy. We draw insight from search theory and focus on the option values for the attribute levels of a searched product. However, unlike in optimal search models, we take both f^{rec} and f^t into account. Because option values are easiest to understand for multilevel attributes, we temporarily modify our notation to accommodate multilevel attributes. In particular, let $x_{i\ell j} = 1$ if product j has attribute i at level ℓ . We define $w_{i\ell}$ similarly. Then we implement the heuristic as follows:

$$(6) \quad \vec{x}_{rec} = \arg \max_{j=1 \text{ to } J} \sum_{i=1}^I \sum_{\ell: x_{i\ell j}=1} \left[\int_{w_i^*}^{\infty} w_{i\ell} df^{rec}(w_{i\ell}) - \int_{w_i^*}^{\infty} w_{i\ell} df^t(w_{i\ell}) \right],$$

where w_i^* is the highest expected value over ℓ of $w_{i\ell}$. (To make Equation 6 feasible, we use the expected value for w_i^* rather than integrate over all possible outcomes.) We call a recommendation system that implements this modification, “Option Value Discrepancy.”

Kullback-Liebler. We might also attempt to quantify the difference between the two distributions, f^{rec} and f^t , and use the quantified difference. A formal measure of distance between two distributions is the Kullback-Liebler divergence from f^t to f^{rec} , $D_{KL}(f^t \parallel f^{rec})$. Kullback-Liebler divergence measures the non-symmetric difference between two probability distributions. If a larger divergence between the two distributions leads to the most learning, then, returning to the x_{ij} notation, a heuristic based on this concept is the following where $D_{KL,j}$ is the Kullback-Liebler divergence for product j .

$$(7) \quad \vec{x}_{rec} = \arg \max_{j=1 \text{ to } J} \sum_{i:x_{ij}=1} D_{KL,j}(f_i^t \parallel f_i^{rec}).$$

We call a recommendation system that implements this modification, “Kullback-Liebler.” We provide more specific formulae in §7, when we make distributional assumptions about f^{rec} and f^t .

5. Preference-Weight Learning Criteria

Preference-weight learning introduces new twists to existing literatures in recommendation systems and search theory. In particular, the modifications in §4 are based on criteria that differ from the typical criteria in the recommendation systems literature and in the search theory literature. In this section, we explore those differences.

5.1. Compared to Typical Recommendation System Criteria

In the recommendation systems literature and in the marketing science literature, most recommendation systems are evaluated on their ability to predict the products that consumers will choose. This is a reasonable criterion when the consumer knows his or her preference weights because, in such cases, the goal is to recommend the product that will deliver maximum purchase and consumption utility. This criterion will not maximize purchase and consumption utility when consumers learn preference weights.

Consider the condominium example. If the recommendation system were to recommend the condominium that maximizes expected utility based on f^t , the Candice and Dave would choose a condominium without playground proximity in a full-service building. Candice and Dave would never know what they missed. Our example in §4.3.1 also implies that the best recommendation may not be the product Candice and Dave are most likely to choose. Instead, the best recommendation might be a condominium that the recommendation system believes Candice and Dave will never buy. Searching a con-

dominium with a less-desired kitchen, but with playground proximity and their true preferred level of service, might be the most efficient way for Candice and Dave to learn their preference weights for playground proximity and level of service. After updating their preference weights, the Candice and Dave are more likely to search the product that maximizes purchase and consumption utility net of search costs. The trends of product diversity, product novelty, and product serendipity can be interpreted as improving the continuation value in Equation 3 because consumers learn their preference weights.

5.2. Compared to Typical Search Theory Criteria

Because most search-theory analyses assume independence among products, efficient optimal policies are based on indices (e.g., Weitzman 1979; Branco, Sun, and Villas-Boas 2012). Although the details vary, the indices tend to favor the upper tail of the probability distribution of product utility rather than the expected value of the product utility or the expected value of a level of the attribute. Upper-tail criteria (indices) represent the option value from the searched product (choose it, continue to search, or choose an already searched product). The option value represents the expected gain in utility if the product turns out to have higher utility than the current best option. If the expected utilities of two products are equal, these criteria favor products with high variance in utility distributions (if the distributions are from the same distributional family).

These criteria do not necessarily apply when consumers learn their preference weights, in part, because of the interdependence in Equation 3. Assume that Candice and Dave are now sure that their preference weight for playground proximity, $w_{playground}$, is high, but they don't know how high. They have decided on an open-concept kitchen and a live-in superintendent, but are now considering whether they want a condominium with an eat-in kitchen, a media room, or both. (Attribute level weights are net of added price.) There are three types of condominiums left to search:

- (a) {playground proximity = good, media room = yes, kitchen = not eat-in}
- (b) {playground proximity = good, media room = no, kitchen = eat-in}, and
- (c) {playground proximity = bad, media room = yes, kitchen = eat-in}.

Candice and Dave do not know whether they prefer a media room or an eat-in kitchen, but they know that having both has less preference weight than playground proximity, $w_{playground} > w_{media} + w_{eat-in\ kitchen}$. Candice and Dave may choose to search condominiums in type (c) to resolve their preference weights for media rooms and eat-in kitchens, even though the option values of both the type (a) and type (b) condominiums are higher. After resolving the uncertainty by evaluating type (c) condominiums, they can then choose a condominium in either type (a) or type (b). Their actual decision, and the

decision on whether to continue to search, depends upon the preference-weight distributions, the search costs, and the discount rate. We argue in §5.3 that prior preference-weight distributions exist such that the best strategy is to search type (c) condominiums and that such distributions are reasonable.

The Candice-and-Dave example provides a counterexample to maximizing the option value of the searched products. The optimal product to search, a condominium of type (c), has no option value if $\Pr(w_{playground} > w_{media} + w_{eat-in\ kitchen}) = 1$, because a condominium of type (c) would never have the highest utility. It also provides a counterexample to maximizing the option value of preference-weight distributions. The option value of the weight on playground proximity can be higher than the option value for the weights on the other attributes, but resolving playground-proximity uncertainty does not change Candice and Dave’s decision. The key mathematical insight is that, for preference-weight learning, the continuation value in Equation 3 is not separable in either products (j) or attributes (i) if the product space is not full-factorial (if all 2^I products are not available). Resolving uncertainty by updating f^t to f^{t+1} occurs for all aspects in the searched products which, in turn, affects the option values of all products searched and to be searched.

The Candice-and-Dave vignette illustrates search without a recommendation. It is easy to embellish the vignette for recommendations. If $f^{rec} = f^t$, the best recommendation remains type (c) condominiums. But suppose that f^{rec} is such that the recommendation system believes that Candice and Dave undervalue eat-in kitchens, then the motivation for recommending a search of type (c) condominiums is even stronger.

5.3. Formal Demonstration of Intuition

§5.1 and §5.2 provide intuitive examples to illustrate that neither the consumer’s optimal search nor the best recommendations are based on the traditional criteria. It remains to demonstrate that we can actually choose f^t and/or f^{rec} such that traditional criteria do not apply. Appendix 1 contains two formal proofs for a product space with three aspects. The first formalization addresses optimal search when there is no recommendation system. The second formalization adds a recommendation system. Together the two results formalize the intuition from §5.1 and §5.2. Because we need only establish the existence of consumer search and recommendation systems for which standard criteria do not apply, the formalizations allow the consumer to update beliefs fully with a single observation of an aspect.

Specifically, Appendix 1 establishes that the criteria for the optimal search without a recommendation and for the optimal recommendation under preference-weight learning:

- *Differ from typical criteria in recommendation systems.*

- *the best recommendation may not be the product with the highest expected utility.*
- *the best recommendation may be a product that the consumer is unlikely to choose.*
- *Differ from typical criteria in optimal search*
 - *the best product to search (recommend) may not be the product with the highest option value.*
 - *the best product to search (recommend) may not have the largest variance in utility.*

The formal demonstrations tell us what the recommendation should not do. In §4 we proposed three criteria for new recommendation systems:

- *Modification 1. Recommend products with diverse aspects.*
- *Modification 2. Recommend products with undervalued expectations ($E[f^{rec}]$ vs. $E[f^t]$).*
- *Modification 3. Recommend products to move the consumer's priors closer to the recommendation system's priors.*

The suggested modifications are consistent with, and provide a marketing-science explanation for, recent trends in the recommendation system literature (diversity, novelty, and serendipity). However, preference-weight-learning suggests that diversity should be with respect to aspects, novelty should be based on not-yet-searched aspects, and serendipity should focus on products with aspects that are undervalued.

6. Recommendations When Consumers Learn Preference Weights

Before we test the proposed recommendation system modifications, we gain insight by exploring how recommendations influence the consumer's search path.

6.1. Structure of the Synthetic Data

We examine a product space defined by three six-level attributes. The eighteen ($18 = 3 \times 6$) aspects in our product space are sufficient to illustrate interesting phenomena, but not so complex as to make calculating post-recommendation search infeasible. With three six-level attributes, $6 \times 6 \times 6 = 216$ feasible products exist in the product space. Let L_k denote the set of aspects that correspond to levels of attribute k .

The consumer's prior beliefs, f_i^0 , for all aspects, i , are normally distributed and independent over aspects. We denote the means and standard deviations of the aspect-based normal distributions by \bar{w}_i and σ_i such that $f_i^0 = \mathcal{N}(\bar{w}_i^0, \sigma_i^0)$. The signal obtained about the preference weight by searching a product with the corresponding aspect is also normally distributed, $g = \mathcal{N}(w_i^r, \sigma_i^s)$. These assumptions imply that the posterior distributions, updated from product search, are normally distributed,

$f_i^t = \mathcal{N}(\bar{w}_i^t, \sigma_i^t)$. The formulae for the posterior distributions are standard and given in Appendix 2. The posterior mean is a convex combination of the prior mean and the true mean. For ease of notation, we assume both the consumer and the recommender system know the preference weight for a base level (aspect) of each attribute. We set the preference weight of the base level to 0. (For example, such assumptions are necessary for identification in choice-based conjoint analysis.) The specific values of \bar{w}^0 , $\bar{\sigma}^0$, and \bar{w}^r are given in Appendix 2. Patterns similar to those discussed in this section emerge for a wide range of parameter values and numbers of attributes.

To simplify the search problem, we assume all 216 products are available. We call such product spaces “fully-crossed” to distinguish them from the much larger full-factorial aspect spaces. In a fully-crossed product space, the maximum-utility product is the product with the maximum-preference-weight level for each attribute. In a fully-crossed product space, the search problem can be broken down by attribute. Returning to the multilevel attribute notation introduced in §4.4, for each attribute i , the consumer searches the product that maximizes (over levels) the option value of searching level ℓ in each attribute i , $\int_{w_i^*}^{\infty} w_{i\ell} df^t(w_{i\ell})$, where w_i^* was defined in §4.4. Because the product space is fully-crossed, there exists a product that has attributes at the maximizing levels. (Note that these option values apply attribute-by-attribute and only in a full-crossed product space. This does not contradict the option-value result in §5.3.)

The attribute-level-maximization policy is heuristic rather than optimal because the lower limit of the integral is based on expected values rather than a full (infeasible) solution to the Bellman equation. The attribute-level-maximization heuristic approaches optimality as the signal variances, $(\sigma_i^s)^2$, approach zero. The attribute-level-maximization heuristic outperforms other known heuristics, such as that by Chick and Frazier (2012), which is optimized for search with attribute-level learning among independent products. (Details available from the authors.)

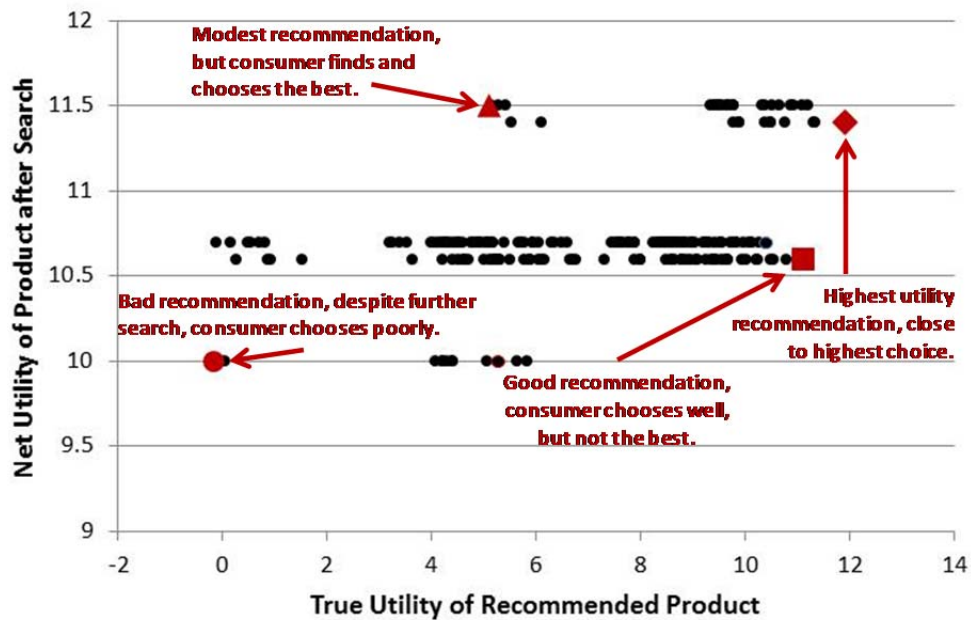
In this section, we evaluate the impact of a single recommendation for each consumer. A single recommendation is sufficient to affect the consumer’s search path and illustrate the phenomena made possible by preference-weight learning. We allow multiple recommendations when we test the proposed recommendation-system modifications in §7.

6.2. Comparison of Recommended and Chosen Products

We evaluate recommendations of each of the 216 possible products. We assume U_{bonus} is such that the consumer searches the recommended product and then continues searching (near) optimally. At the (endogenous) end of search, the consumer either chooses a product or the outside option. We summarize the results in Figures 1 and 2.

The vertical axis of Figure 1 plots the consumer’s net pay off—the utility of the chosen product, minus the incurred search costs. The horizontal axis of Figure 1 represents the true utility of the recommended product (based on \vec{w}^r). The consumer learns (updates) f^1 after he or she searches the recommended product and, perhaps, continues to search. In Figure 1, all recommendations lead the consumer to purchase one of three products as indicated by the horizontal clusters. The net utility of the chosen product differs slightly because search costs differ. (“Three products” is not a general result. Different parameter values give different numbers of post-search products.)

Figure 1. Net Utility of Product Chosen After Search vs. Utility of Recommended Products



We first examine the product recommendation in Figure 1 that is represented by the diamond (◆). In this case, the recommendation system recommended the highest-utility product and the consumer chose that product, but ultimately did so after incurring more search costs than would have been incurred for other recommendations. This suggests that even if a recommender system has perfect knowledge of consumer utility, the highest utility product may not be the best recommendation if the consumer has to learn his or her own preferences. The recommendation indicated by a triangle (▲) also leads the consumer to the highest-utility product, but does so by recommending a much lower-utility product. After receiving this recommendation (▲), the consumer learns efficiently that some attributes are more important and some are less important than previously thought.

The recommendation indicated by a square (■) is interesting. For such recommendations, the consumer is satisfied with the recommended product and has no incentive to search further. Such rec-

ommendations by non-benevolent recommenders (or poorly-designed recommendation systems) exploit the consumer's naïveté and lead the consumer to purchase a product that is not the highest true utility. The consumer would not update his or her priors sufficiently, never learn of better-than-expected-utility products, but would be satisfied with the chosen product at the time of purchase. The consumer might even thank the recommendation system for a recommendation. For example, a non-benevolent realtor might have incentives to recommend his or her own listing to obtain both seller and buyer commissions (Ansari, Essegai, and Kohli 2000). Similarly, and with similar concerns, short-term gains might be tempting if paid advertising supported a recommendation system. The other interesting feature about the recommendation denoted by a square (■) is that other recommendations (▲) exist, with lower initial utility, that lead to higher post-search utility.

Recommendations, such as indicated by a circle (●), lead the consumer to choose low-utility products. After following such recommendations, the consumer, acting optimally based on his or her priors, updates some of his or her preference weights, but never updates preference weights sufficiently to find the highest-utility product. Post search, the consumer believes falsely that he or she has found the product that has the highest utility. Despite the opportunity loss, the consumer might be satisfied. Without a recommendation and without search, the consumer in this example would have chosen the product the consumer believes would maximize utility. The consumer would have expected to receive utility of 11.4 from this purchase, but, upon consumption, would only have received utility of 9.4. Had the consumer searched without a recommendation, and chosen optimally after searching this product, the consumer would have received utility of 10.7, basically in the middle tier of products that could have been recommended.

Figure 2 provides a different perspective based on the same set of candidate recommendations. Figure 2 compares net post-search utility to recommendation system beliefs. In this illustration, the mean of the recommendation system beliefs, f^{rec} , is a convex function of the mean of the consumer's prior beliefs, f^0 , and the consumer's true beliefs, $f(\bar{w}^r)$. As anticipated by §5 and as represented by a hexagon (⬡), a recommended product, thought by the recommendation system to be the highest-utility product, does not lead the consumer to the highest-utility product. After receiving that recommendation (●), the consumer, acting optimally based on priors, would not search sufficiently to find the true highest-net-utility product. A recommendation system would have served the consumer better had it recommended the product indicated by a diamond (◆), or even the product indicated by a triangle (▲). In the latter case (▲), the recommendation itself would not have been a high-utility product, but the recommendation would have caused the consumer to update his or her beliefs and continue searching

until the highest-net-utility product was found.

In Figures 1 and 2, the utilities of the recommended products and the net utilities of the chosen products are correlated ($\rho = 0.43, \rho = 0.28$, respectively), but the relationship is well below 1.0. Preference-weight learning drives the lack of perfect correlation. Detailed examination of the search path reinforces the insights obtained from §4—the best recommendations are those that encourage the consumer to search products that reveal diverse aspects and undervalued aspects. Figure 2 reinforces the insight that good recommendations provide valuable information. Figures 1 and 2 allow the consumer to be surprised either positively (preference weight higher than priors) or negatively (preference weight lower than priors). Both forms of preference-weight learning are valuable to consumers.

Figure 2. Net Utility of Product Chosen after Search vs. Recommendation System’s Beliefs about the Utility of the Recommended Product



7. Preference-Weight-Learning Recommendation Systems Evaluations

To explore whether the modifications suggested in §4 improve recommendation system performance, we expand the analyses of §6 to 5,000 consumers in multiple experimental conditions. Because the value of preference-weight learning depends upon differences in f^{rec} and f^t , our experimental conditions vary with respect to the quality of consumers’ priors (naïveté), the quality of the recommendation system’s priors (recommendation system knowledge), and the rate at which consumers update their preference weights. By design, consumers learn their preference weights as they search products. Our synthetic-data experiments are proof-of-concept experiments: we examine whether the

recommendation system modifications improve recommendation system performance when consumers learn their preferences. We complement these synthetic-data experiments with an empirical demonstration in §8.

We expect preference-weight learning to be particularly relevant for naïve consumers who are new to a product category. Naïveté is more likely for infrequent purchases such as automobiles, housing, college choice, and nannies. Naïveté is also more likely for consumers who feel they need recommendations. Because the preference-weight-learning modifications use more-complete knowledge about f^{rec} and f^t than do typical recommendation system benchmarks, we expect the incremental value of the preference-weight-learning modifications to generally increase with recommendation system knowledge. The exception is low recommendation-system knowledge where we expect that no recommendation system does well. We expect that faster updating should favor the preference-weight-learning modifications.

7.1. Product Space and True Consumer Utilities

We simulate a product space of three six-level attributes using the same structure that we used in §6. We consider recommendation systems that recommend two products sequentially. For each experimental condition and for each of 5,000 consumers in that experimental condition, we draw true aspect preference weights from a mixture of two normal distributions: one with a low mean to represent unimportant aspects and one with a high mean to represent important aspects. The consumer’s prior beliefs are normally distributed and independent over aspects and depend on naïveté as described in §7.2. The variances of the consumers’ priors (v_i in Appendix 2) are drawn i.i.d. from an exponential distribution. The specific values of the parameters of the preference-weight distributions are given in Appendix 2. For readers wishing to explore other parameter values, other recommendation system modifications, or other combinations of naïveté, recommendation system knowledge, or the rate of updating, the software is available from the authors.

7.2. Characteristics Varied in the Synthetic Data Experiments

Consumer Naïveté. For each consumer, we set the prior means equal to the true means for a fraction of the aspects. For the remaining aspects, we redraw the prior means randomly. A consumer is more naïve (less expert) if a larger fraction of the consumer’s prior beliefs are redrawn randomly. We vary this fraction (η in Appendix 2). An expert has naïveté equal to zero and a novice has naïveté equal to 1.

Recommendation System Knowledge. For some aspects we set the recommend system’s priors to the consumer’s priors (probability P^c). For the remaining aspects we set the mean of the recommen-

dation system's priors to the true partworths with probability P^{rec} and we set the recommendation system's priors randomly with probability $1 - P^{rec}$. Larger P^{rec} implies greater recommendation system knowledge. We maintain the P^c parameter to recognize that, even with $P^{rec} = 1$, recommendation systems are unlikely to ever know the true preference weights perfectly. The variances of the recommendation system priors are constant for all consumers. (In theory, we can manipulate recommendation system knowledge by manipulating the means, the variances, or both. Exploratory simulations suggest it is sufficient to manipulate the means.)

There are a variety of well-established methods by which recommendation systems learn consumers' priors and consumers' true preferences. Collaborative filters, content-based filters, and statistical model-based systems, as reviewed in §2, are just a few examples. Methods for such firm-side learning are well-established in theory and in practice, but research on improvements continues. In a related context, Hauser, Liberali, and Urban (2014) illustrate that, in automated systems, firm-side learning about the best “morph” can be decoupled from learning about consumers' cognitive styles. (The analogy is recommendation policy \Leftrightarrow morph and consumer beliefs \Leftrightarrow cognitive style. The mathematical structure, while related, is not identical.) Recommendation system knowledge is likely to be greater if consumers are more homogeneous, if consumer's priors can be measured directly, if true preferences can be estimated as a function of observable characteristics of the consumer, or if the recommender system has access to more data. (Manipulating the heterogeneity of the environment provides an alternative method to manipulate recommendation system knowledge. Results are similar.)

Rate of Updating. The rate at which the consumer converges to the true utility parameters depends on the variance of the signal, $(\sigma_i^s)^2$ in Appendix 2. We vary this parameter in our experiments. In some experimental conditions, consumers update their priors rapidly (low signal variance) and in other experimental conditions consumers update their priors slowly (high signal variance).

7.3. Preference-Weight Learning Recommendation Systems and Benchmark

We test four recommendation system modifications that implement the insights and equations from §4. We compare these modifications to the typical benchmark recommendation system. The basic concepts are reviewed here. Recommendation system priors are $f_i^{rec} = \mathcal{N}(\bar{w}_i^{rec}, \sigma_i^{rec})$.

Maximum Expected Utility Benchmark. The maximum-expected-utility recommendation system recommends products that it expects will give the consumer the highest utility (based on f^{rec}).

Aspect Diversity. The “Aspect Diversity” modification modifies the benchmark recommendation system by subtracting a penalty proportional to the number of aspects in common with the product the consumer would have searched without a recommendation, Equation 4.

Undervalued Products. The “Undervalued Products” modification compares the consumer’s mean prior beliefs (f^t) to the recommendation system’s mean beliefs (f^{rec}) and maximizes the difference in expected utility, Equation 5. Preliminary experiments suggest that we improve recommendations for this modification when we make the recommendation less sensitive to small variations. When the variation is below a threshold, it is likely the consumer has learned his or her preferences—we revert to the benchmark.

Option Value Discrepancy. The “Option Value Discrepancy” modification recommends products for which the attribute-based option values, as calculated by the recommendation system are larger than the attribute-based option values as calculated by the consumer, Equation 6. Preliminary experiments suggest a threshold does not improve this modification.

Kullback-Liebler. The “Kullback-Liebler” modification compares the Kullback-Liebler divergence for every aspect in product j and recommends the product for which the divergence is largest, Equation 7. When the consumer’s prior beliefs and the recommendation system’s beliefs are normally distributed, the Kullback-Liebler Divergence, $D_{KL,j}$, for each product, j , is given by Equation 8:

$$(8) \quad D_{KL,j}(f^t \parallel f^{rec}) = \frac{1}{2} \left\{ \sum_{i: x_{ij}=1} \left[\left(\frac{\sigma_i^t}{\sigma_i^{rec}} \right)^2 + \frac{(\bar{w}_i^{rec} - \bar{w}_i^t)^2}{(\sigma_i^{rec})^2} + 2 \ln \frac{\sigma_i^{rec}}{\sigma_i^t} - 1 \right] \right\}.$$

7.4. Results of the Synthetic Data Experiments

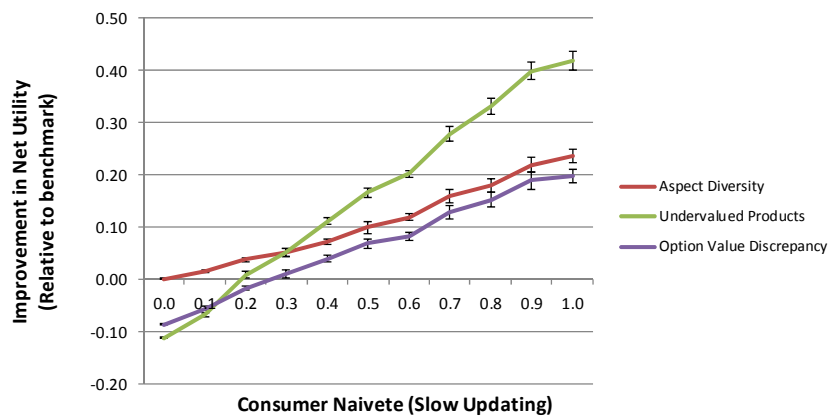
We vary consumer naïveté and recommendation system knowledge in 10 equal steps each from 0.1 to 1.0 for each level of signal variance and for each of five recommendation systems. In each of the experimental conditions, each of 10,000 consumers receives two sequential recommendations. The consumer searches the first recommended product and updates his or her preference weights. The consumer then searches the second recommended product. The consumer either purchases a product or continues to search (near) optimally until the optimal stopping rule is reached. Performance is the difference between net utility achieved by a consumer who searches the product recommended by the recommendation-system modification and the net utility the consumer would have achieved without a recommendation. Net utility is the true utility of the chosen product, if any, minus the incurred search costs.

To visualize the patterns that emerge from the synthetic data experiments, we plot performance of the system while varying one parameter, and holding two other experimental variables constant. Because consumer preference weights are redrawn for each of the 10,000 consumers, we subtract from net utility, the utility of the product the consumer would have chosen without a recommendation. Be-


cause we are interested in performance relative to the benchmark (maximum expected value recommendations), we subtract benchmark performance from the performance of each recommendation system in our plots. (The performance of the benchmark, relative to no recommendation, is plotted in Appendix 2. As expected, the benchmark performance increases for both consumer naïveté and recommendation system knowledge.)

Figure 3 plots relative performance versus consumer naïveté holding recommendation system knowledge constant; Figure 4 plots relative performance versus recommendation system knowledge holding consumer naïveté constant. Each figure contains a plot for slow consumer updating (high signal variance) and for fast consumer updating (low signal variance). We chose intermediate values for the experimental variable that is held constant: naïveté = 0.6 and recommendation system knowledge = 0.6. Figures 3 and 4 illustrate how naïveté and recommendation system knowledge affect relative performance as a function of consumer naïveté, recommendation system knowledge, and the rate of updating.

Figure 3. Improvement in Net Utility as a Function of the Consumer’s Naïveté




(a) Slow Updating of Consumer Priors

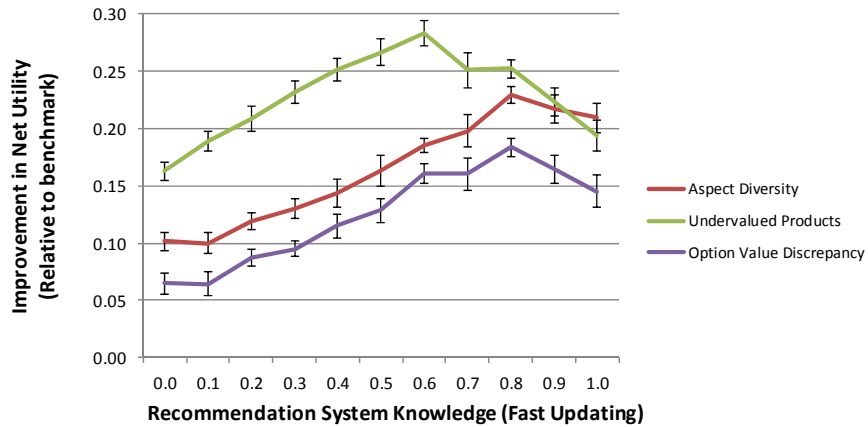


(b) Fast Updating of Consumer Priors

Figure 4. Improvement in Net Utility as a Function of Recommendation System Knowledge



(a) Slow Updating of Consumer Priors



(b) Fast Updating of Consumer Priors

Figures 3 and 4 suggest that “Aspect Diversity”, “Undervalued Products”, and “Option Value Discrepancy” improve recommendations relative to the improvement achieved by the recommendation system benchmark. We did not plot the performance of “Kullback-Liebler” to keep Figures 3 and 4 readable. Although “Kullback-Liebler” takes the full distributions of f^{rec} and f^t into account, rather than just their means, its complexity appears to be handicap. It does not perform as well as the simpler recommendation systems. The performance of the Kullback-Liebler recommendation system is available from the authors.

“Diverse Aspects” does better than the recommendation system benchmark over the entire range of naiveté. “Undervalued Products” and “Option Value Discrepancy” do extremely well for highly naïve consumers, but are slightly counterproductive relative to the benchmark for expert consumers. The latter is not surprising. There is little value to exploration of preference weights when consumers

already know their preference weights. Furthermore, when consumers know their preference weights, they are more likely than the tested recommendation system to choose the best product to search. Overall, the best recommendation system for naïve consumers is “Undervalued Products.” The best recommendation for expert consumers is “Diverse Aspects,” although it does not do much better than the benchmark. The proposed recommendation systems do slightly better when consumers update their preference weights rapidly, however the relative improvement is slight and the pattern is similar to that of the slow updating case.

When we vary recommendation system knowledge holding consumer naïveté constant at an intermediate value (0.6), all of the proposed recommendation systems beat the benchmark. “Undervalued Products” tends to be best except for very high recommendation system knowledge. Once again, the algorithms do slightly better when consumers update their preference weights rapidly.

7.5. Summary of Synthetic Data Experiments

Figures 3 and 4 demonstrate that situations exist in which recommendation system modifications based on preference-weight learning improve the consumer’s net utility more than the typical recommendation-system benchmark. (Improvement is relative to no recommendation.) All three proposed modifications do well over most of the ranges and the performance appears to be robust with respect to the rate of updating. With refinement and tuning, we expect the relative performances of all of the modifications to improve further. For example, we might test a heuristic that combines the best features of the Undervalued-Product and the Aspect-Diversity modifications, or we might tune either or both modifications. We might modify “Kullback-Liebler” to be more robust. For readers wishing to explore refinements, software is available from the authors.

8. Empirical Demonstration

We began §1 with premises that (1) consumers learn preference weights while searching products and (2) recommendation systems can anticipate consumers’ (true) preference weights. Based on these premises we proposed recommendation systems that take preference-weight learning into account. Synthetic data suggest that “Undervalued Products” shows promise relative to a benchmark recommendation system, “Maximum Expected Utility,” and relative to allowing consumers to choose without a recommendation. (For the remainder of this section, for simplicity, we call the three methods to search products Undervalued, Max Expected, and Choice, respectively.) In this section we describe an empirical study in which we demonstrate that the two premises are reasonable and that Undervalued helps consumers learn preference weights better than either Max Expected or Choice.

8.1. Context of the Empirical Demonstration

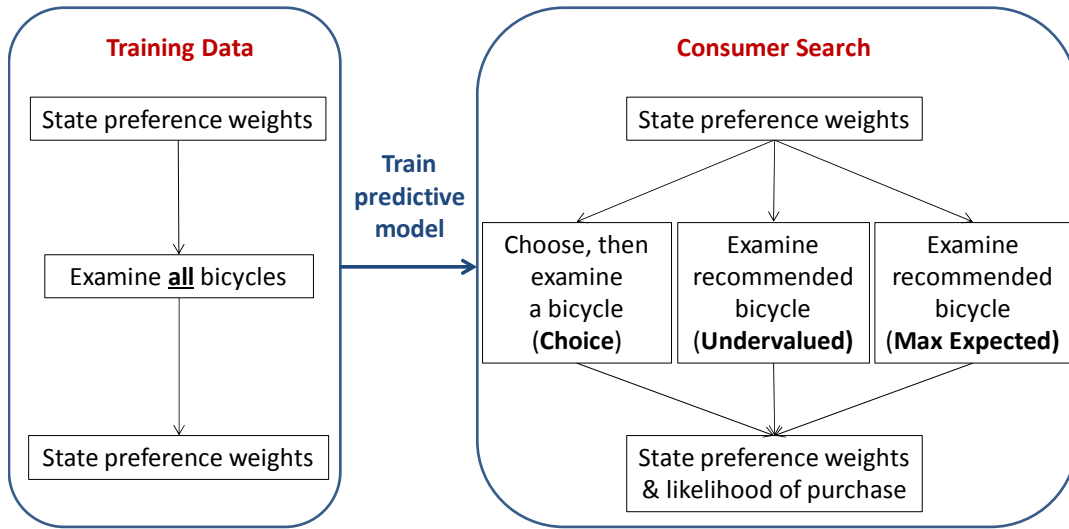
We consider a product space of five bicycles and six attributes. We created rich realistic descriptions of the five bicycles which, when studied by consumers, simulate product search. The bicycles vary on five of the six attributes of gel seat, folding ability, hydraulic brakes, nighttime visibility, and a low bar for step-through access. These attributes were chosen based on discussions with potential bicycle customers and a review of the bicycles available on the market. If a bicycle had the attribute, the rich realistic descriptions highlighted the attribute and its benefits in text and pictures. By design, none of the rich realistic descriptions mentioned the sixth attribute, variety of colors. With only five bicycles, the product space was not full factorial. There was no obviously dominant search strategy.

To encourage respondents to search and evaluate the bicycles seriously, respondents were asked to provide ratings and qualitative comments about any bicycles that they evaluated. Following standard procedures, before any analyses, we eliminated respondents who answered too fast, answered the same for all questions, or provided nonsense qualitative answers (13 respondents were eliminated). The sample was drawn from Amazon Mechanical Turk. Respondents received the standard honorarium for completing the tasks.

8.2. Basic Design

Figure 5 summarizes the study design. In the first stage, prior to providing recommendations or allowing the respondent to choose which bicycle to search, we collected data and trained a model by which the recommendation system can predict consumers' true preference weight distribution, $f^{rec}(\vec{w})$ – left side of Figure 5. Next, respondents were assigned randomly to one of three experimental cells: Choice (92 respondents), Undervalued (110 respondents), or Max Expected (106 respondents)—right side of Figure 5. The Choice cell acts as a control: if assigned to Choice, respondents chose which bicycle to search based on attribute-level summaries. If a respondent was assigned to Undervalued or Max Expected, the recommendation system chose the bicycle for the respondent to search based on the algorithms in §7.3. Respondents, who received a recommendation, then searched the recommended bicycle. In all three cells, respondents first stated their prior preference weights (\vec{w}^0), searched a bicycle (viewed a rich realistic description), rated their likelihood of purchase, then stated their preference weights (\vec{w}^1) again. Preference weights were measured with a ten-point scale where the most important attribute is given a ten. Likelihood of purchase was measured with a three-point scale. The study was pretested with 37 respondents to assure the respondent tasks were clear and relevant.

Figure 5. Basic Design of the Empirical Demonstration



8.3. Training a Predictive Model for $f^{rec}(\vec{w})$

To train $f^{rec}(\vec{w})$ we measure respondents' preference weights after they have searched all available bicycles (training data). We also measure consumer characteristics such as age, gender, where they live (urban, suburban, rural, other), and how they plan to use the bicycle. Using linear regression, we predict $\vec{w}^r = E[\vec{w} | f^{rec}]$ as a function of these variables and the respondents' prior preference weights (\vec{w}^0). The last variable accounts for heterogeneity in preference weights that is not tied to the observed characteristics. The predictive model is trained on 93 respondents. These respondents did not participate in the consumer search on the right side of Figure 5. In practice, a recommendation-system's underlying predictive model would be based on tens of thousands of respondents (or more), hence our empirical demonstration is conservative.

8.4. Results of the Empirical Demonstration

Manipulation checks. We first check whether the recommendation systems recommended bicycles that were different between recommendation systems and different from bicycles that respondents chose to search on their own. The distributions of bicycles differ among the three experimental cells ($\chi^2 = 117.8, d. f. = 8, p < 0.01$), Undervalued differs from Max Expected ($\chi^2 = 24.0, d. f. = 4, p < 0.01$) and from Choice ($\chi^2 = 77.9, d. f. = 4, p < 0.01$). The attribute, variety of colors, serves as a control attribute—we do not anticipate that its preference weight would change after search. It did not: $w_t^1 - w_t^0 = 0.12, 0.12, \text{ and } 0.07$ for Undervalued, Max Expected, and Choice, respectively—no differences were statistically significant from zero, nor statistically different between pairs of experimental cells. Finally, we expect Undervalued to sacrifice the overall expected consumption utility of the recommended bicycle relative to Max Expected (and relative to Choice), but hope the sacrifice is not substan-

tial. Our surrogate for expected consumption utility is the stated likelihood of purchase. The average values were 2.2, 2.3, and 2.3 for Undervalued, Max Expected, and Choice, respectively—no differences were statistically significant between pairs of experimental cells. The data pass the manipulation checks—respondents in each of the three experimental cells chose or were recommended different bicycles, the preference weights for the control attribute did not change, and the surrogate for expected consumption utility behaved as anticipated.

Change in preference weights. We test the first premise by comparing preference weights before and after search using root-mean-square-change (RMSC) and the sum over all attributes. Additionally, we test whether Undervalued is better at helping consumers to learn preference weights. Specifically, we test whether preference weights change more for Undervalued than for the other two experimental cells. Changes are as predicted. RMSC = 0.81, 0.39, and 0.46 and the sum over attributes is 1.69, 0.20, and 0.53, respectively for Undervalued, Max Expected, and Choice. Sums are significantly different from zero for Undervalued and Choice (p 's < 0.05). The change in the Undervalued experimental cell is significantly larger than the change in either Max Expected or Choice for both RMSC and for the sum of changes (p 's < 0.05), but the change in Max Expected is not significantly different than that in Choice ($p = 0.34$). On an attribute-by-attribute basis, the average preference weights (other than for the control attribute) increase significantly or marginally significantly when respondents search bicycles recommended by Undervalued (p 's < 0.05 for gel seat, folding ability, and a low bar, p 's < 0.10 for hydraulic brakes and nighttime visibility). These changes are larger than for those in the other two experimental cells. No attribute-level preference weights change significantly for Max Expected and only hydraulic brakes increase significantly for Choice ($p = 0.03$). Qualitative comments were consistent, e.g., “I changed the folding ability up to ten. I think I really want that in my next bike.”

Thus, the data demonstrate that (1) product search can cause preference weights to change (at least for Undervalued) and (2) searching Undervalued recommendations causes preference weights to change more than searching benchmark recommendations (Max Expected) or searching without a recommendation (Choice).

Moving preference weights toward their true values. Empirically, we cannot observe \vec{w}^r , but we can observe $E[\vec{w}|f^{rec}]$. To evaluate preference weight learning, we compare the observed preference-weight change, $\vec{w}^1 - \vec{w}^0$, to the predicted preference-weight change, $E[\vec{w}|f^{rec}] - \vec{w}^0$. (We compare differences to control for heterogeneity in the \vec{w}^0 .) We expect the two differences to move in the same direction if (1) the search changes preference weights and (2) the predictive model is reasonable. We test the movement with regressions (for each attribute and each experiment cell) in which the ob-

served change is a function of the predicted change. We focus on directional movement recognizing that the predictive models are based on simple regressions using training data for four variables from 93 respondents. Typical commercial recommendation systems might be based on hundreds of variables, 10s of thousands of respondents, and state-of-the-art machine learning methods. Users of commercial recommendation systems would likely search more products than the one recommended product in our tests.

For the Undervalued experimental cell, the regression coefficient is significant for gel seat, folding ability, and hydraulic brakes (p 's < 0.05) and marginally significant for low bar ($p < 0.10$). All coefficients are positive. For the Max Expected experimental cell, all coefficients are positive, but none are significant. For the Choice experimental cell, the gel seat coefficient is positive and significant ($p < 0.05$), and the low bar coefficient is positive and marginally significant ($p < 0.10$). The coefficients is negative, but not significant for visibility. Thus, when respondents search bicycles recommended by Undervalued, their preference weights move in the direction that is predicted for the true preference weights. This movement is greater than the movement when respondents search benchmark recommendations (Max Expected) or when respondents search on their own (Choice).

To summarize, the empirical demonstration is consistent with both premises (consumers can learn preference weights and a recommendation system can predict that learning). The empirical demonstration also suggests that a proposed recommendation system leads to greater preference-weight learning than either a benchmark recommendation system or search without recommendations. We consider this empirical demonstration a proof of concept that we hope will be refined in subsequent tests.

9. Summary and Discussion

9.1. Summary

When consumers update their preference weights as they search, the optimal search strategy becomes more complex because interdependence among products is introduced by preference-weight learning. Whenever a preference weight is updated, the update changes the utilities of all searched products and the expected utilities of all products yet to be searched. Optimal policies in classic search models, such as the policy derived in Weitzman (1979), are no longer optimal. When we introduce a recommendation system, the policies by which the optimal recommendations should be chosen are even more complex. In contrast to traditional recommendation systems literature, but consistent with recent developments with respect to diversity, novelty, and serendipity, preference-weight learning suggests that recommendation systems should not be evaluated solely on accuracy (probability of

choice or expected utility of the recommended product).

Despite the complexity, we gain insight about effective recommendation systems for consumers who learn their preference weights by using intuition and examining the Bellman equation. These insights suggest modifications to recommendation systems—recommend products with diverse aspects, recommend products that the consumer undervalues, and recommend products most likely to update the consumer’s prior beliefs. Recommendation systems based on these modifications perform well in synthetic data experiments, especially for naïve consumers and when recommendation systems can predict consumers’ preference weights. The proposed recommendation systems perform well as long as they are not too complex.

We demonstrate empirically that consumers update their preference weights from searching even a single rich realistic description of a bicycle. Preference-weight learning is greater when consumers search a bicycle recommended by one of the proposed recommendation systems—greater than search based on a benchmark recommendation system or search without a recommendation. Furthermore, the actual change in preference weights moves as predicted, especially when consumers search products recommended by the proposed recommendation system.

9.2. Generality

We expect the insights from the numerical examples, the synthetic data experiments, and the empirical demonstration to scale to larger product spaces. The proposed modifications are stylized proof-of-concept recommendation systems; we expect more-sophisticated modifications to achieve even greater improvements, especially when tuned to specific settings. Our formal model assumes that consumers know attribute levels. Relaxing this abstraction greatly complicates the model, but does not change the insight that preference-weight learning affects search. Preference-weight learning and attribute-level learning are complementary phenomena.

We assumed that attributes are described by discrete levels (aspects). We did not restrict the attributes to be horizontal (type of kitchen) or vertical (live-in superintendent versus full-service concierge), although we did assume that learning about one aspect does not affect the priors with respect to another aspect. Continuous attributes, such as the square footage in a condominium or the proximity to a playground, can be handled in one of two ways—both are common for typical preference-weight measurement methods such as conjoint analysis. Either we discretize the continuous attribute or assume a parametric form (linear, quadratic, logarithmic) for the preference function. Discretized attributes are handled with no modification. For parametric preferences, we need to specify updating rules. For example, the consumer’s preferences might be linearly increasing in square footage. The consumer

updates that linear weight in a matter analogous to the updating rules in Equation 2.

9.3. Future Research

The analyses in this paper demonstrate that preference-weight learning can have a major impact on the study of consumer search and on the design of recommendation systems. Further avenues of research are promising. Researchers might explore interdependence among preference weights for aspects and model how learning about one aspect informs preference weights about another aspect. Interdependence is especially interesting for discretized vertical attributes. Recommendation systems might be developed that identify the consumer's relative naïveté and change algorithms depending upon that naïveté. Ensembles of recommendation systems might do well. We focused on situations in which consumers have ready access to attribute-level information. Combining preference-weight learning and attribute-level search is a complex and interesting challenge.

In §6, non-benevolent recommendation systems influenced outcomes to the benefit of the recommendation system, yet left the consumer satisfied. Such recommendation systems could prove interesting. We assumed that forward-looking consumers solve Equation 3. An alternative interpretation is that heuristic solutions to Equation 3 approximate consumer search. For example, Lin, Zhang, and Hauser (2015) illustrate situations in which consumers use a cognitively simple learning strategy that approximates a complex dynamic program.

We focused on benevolent recommendation systems that maximize consumers' net purchase and consumption utility. One justification for an assumption that the consumer searched recommended products, is the utility bonus. It is possible that the utility bonus is not sufficient—the consumer might decline searching further if the expected utility of the recommended product is too low. This phenomenon might hamper the recommendation system's reputation and provide incentives to the recommendation system to avoid low-expected-utility recommendations even when such recommendations are in the best interests of the consumer. Such situations are worth studying.

Our examples and references from the marketing-science literature motivate preference-weight learning. Research on the underlying mechanism could improve insight. There are rich literatures in both recommendation systems and in marketing about firm-side learning of preference weights. When such algorithms are coupled with the modifications suggested in this paper, we expect that recommendations to consumers will improve.

Finally, our empirical demonstrate was a proof-of-concept. Experiments might explore recommendations in other product categories with more attributes, in situations where commercial prediction algorithms are used, for more extensive search, with endogenous stopping rules.

References

- Adam K (2001) Learning while searching for the best alternative. *Journal of Economic Theory*, 101(1):252-280.
- Adamopoulos P, Tuzhilin A (2014) On unexpectedness in recommender systems: or how to better expect the unexpected. *ACM Transactions on Intelligent Systems and Technology* 5(4):54:1-32.
- Adomavičius G, A Tuzhilin (2005) Towards the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering* 17(6):734-749.
- Ansari A, S Essegaiier, R Kohli (2000) Internet recommendation systems, *Journal of Marketing Research*, 37(August):363-376.
- Bikhchandani S, Shama S (1996) Optimal search with learning. *Journal of Economic Dynamics and Control* 20(1):333-359.
- Branco F, Sun M, Villas-Boas, JM (2012) Optimal search for product information. *Management Science* 58(11):2037-2056.
- Bodapati A (2008) Recommendation systems with purchase data. *J. of Marketing Research*, 16:77–93.
- Bronnenberg BJ, Kim J, Mela CF (2015) Zooming in on choice: How do consumers search for cameras online? Forthcoming *Marketing Science*.
- Castells P, Vargas S, Wang J (2011) Novelty and diversity metrics for recommender systems: choice, discovery and relevance. *International Workshop on Diversity in Document Retrieval (DDR 2011)* at the 33rd European Conference on Information Retrieval (ECIR 2011). Dublin, Ireland, April.
- Celma Ò, Herrera P (2008) A new approach to evaluating novel recommendations. *RecSys '08*, Proceedings of the 2008 ACM Conference on Recommender Systems. Lausanne, Switzerland.
- Chick SE, Frazier P (2012) Sequential sampling with economics of selection procedures. *Management Science* 58(3):550-569.
- Chen Y, Yao S (2016) Sequential search with refinement: model and application with clickstream data, *Management Science* forthcoming, pubsonline.informs.org/doi/10.1287/mnsc.2016.2557.
- Chung J, Rao VR (2012) A general consumer preference model for experience products: Application to Internet recommendation services. *Journal of Marketing Research* 49(3):289-305.
- Cook J (2014) MiniDates schedules real-life (legitimately) blind dates for you. *TechCrunch*, May 30, 2012.
- De Bruyn A, Liechty JC, Huizingh EKRE, Lilien GL (2008) Offering online recommendations with minimum customer input through conjoint-based decision aids. *Marketing Science* 27:3, 443-460.
- Dzyabura D, Jagabathula S (2017) Offline assortment optimization in the presence of an online chan-

- nel. *Management Science*. <http://pubsonline.informs.org/doi/pdf/10.1287/mnsc.2016.2708>
- Dzyabura D, Jagabathula S, Muller E (2018) Using online preference measurement to infer offline purchase behavior. NYU Stern working paper. Forthcoming, *Marketing Science*.
- Erdem T, Keane MP (1996) Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science* 15(1):1–20.
- Finkel EJ, PW Eastwick, BR Karney, HT Reis, S Sprecher (2012) Online dating: a critical analysis from the perspective of psychological science. *Psychological Science in the Public Interest*, 13:3-66.
- Fleder D, Hosanagar K (2009) Blockbuster culture's next rise or fall: the impact of recommender systems on sales diversity. *Management Science* 55(5):697-712.
- Ge M, Delgado-Battenfeld C, Jannach D (2010) Beyond accuracy: Evaluating recommender systems by coverage and serendipity. *RecSys '10*. Proceedings of the 2008 ACM Conference on Recommender Systems, Barcelona, Spain. 257–260.
- Ghose A, Ipeirotis PG, Li, B (2012) Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content. *Marketing Science* 31(3):493-520.
- Gittins, JC (1979) Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society, Series B (Methodological)* 41(2):148-177, plus commentary.
- Gittins, JC, Glazebrook K, Weber R (2011) *Multi-armed bandit allocation indices*. John Wiley & Sons, London.
- Greenleaf EA, DR Lehmann (1995) Reasons for substantial delay in consumer decision making. *Journal of Consumer Research* 22(2):186-199.
- Häubl G, Trifts V (2000) Consumer decision making in online shopping environments: The effects of interactive decision aids. *Marketing Science* 19(1):4-21.
- Hauser JR, S Dong, M Ding (2014) Self-reflection and articulated consumer preferences. *Journal of Product Innovation Management* 31(1):17-32.
- Hauser JR, GL Urban, G Liberali, M Braun (2009) Website morphing. *Marketing Science* 28(2):202-224.
- Hauser JR, G Liberali, GL Urban (2014) Website morphing 2.0: Switching costs, partial exposure, random exit, and when to morph. *Management Science* 60(6):1594–1616.
- Herlocker J, Konstan J, Terveen L, Riedl J (2004) Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems* 22(1):5–53.
- Hong H, Shum M (2006) Using price distributions to estimate search costs. *RAND Journal of Economics* 37(2):257–275.
- Honka E (2014) Quantifying search and switching costs in the US auto insurance industry. *The RAND*

- Journal of Economics* 45(4):847-884.
- Jacobs BJD, Donkers B, Fok D (2016) Model-based purchase predictions for large assortments. *Marketing Science* 35(3):389-404.
- Jun T (2004) A survey on the bandit problem with switching costs. *De Economist* 152(4):513-541.
- Ke TT, Shen Z-JM, Villas-Boas, JM (2016) Search for information on multiple products. Forthcoming *Management Science*.
- Kim JB, Albuquerque P, Bronnenberg BJ (2010) Online demand under limited consumer search. *Marketing Science* 29(6):1001-1023.
- Lin S, J Zhang, JR Hauser (2014) Learning from experience, simply. *Marketing Science* 34(1):1-19.
- Liu NN, Zhao M, Xiang E, Yang Q (2010) Online evolutionary collaborative filtering. *RecSys2010*, Barcelona Spain, 95-102.
- Lu S, Xiao L, Ding M (2016) A video-based automated recommender (VAR) system for garments. *Marketing Science* 35(3):484-510.
- McNee S, Riedl J, Konstan J (2006) Accurate is not always good: how accuracy metrics have hurt recommender systems. *CHI EA 2006*, Extended Abstracts on ACM Human Factors in Computing Systems, Quebec, Canada, 1097-1101.
- Mersereau AJ, Rusmevichientong P, Tsitsiklis JN (2009) A structured multiarmed bandit problem and the greedy policy. *IEEE Transactions on Automatic Control* 54(12):2787-2802.
- Moon S, Russell GL (2008) Predicting product purchase from inferred customer similarity: An autologistic model approach. *Management Science* 54(1):71-82.
- Rogers A (2013) After you read the listings, your agent reads you. *New York Times* March 26, 2013: F4.
- Schlaifer R (1959) *Probability and statistics for business decisions*. (New York, NY: McGraw-Hill).
- Schwartz EM, ET Bradlow, PS Fader (2017) Customer acquisition via display advertising using multi-armed bandit experiments. *Marketing Science* April 20, 2017.
- Schwartz EM, K Misra, J Abernethy (2017) dynamic online pricing with incomplete information using multi-armed bandit experiments. University of Michigan working paper.
- Seiler S (2013) The impact of search costs on consumer behavior: A dynamic approach. *Quantitative Marketing and Economics* 11(2):155-203.
- She J, EF MacDonald (2013) Trigger features on prototypes increase preference for sustainability. *Proceedings of the 25th ASME International Conference on Design Theory and Methodology*, Portland, OR. August 04, 2013.
- Sheehy K (2013) Study: High school grads choosing wrong college majors. *U.S. News & World Report* No-

- vember 11. At <http://www.usnews.com/education/blogs/high-school-notes/2013/11/11/study-high-school-grads-choosing-wrong-college-majors>.
- Tversky A (1972) Elimination by aspects: a theory of choice. *Psychological Review* 79 (4):281-99.
- Urban GL, Hauser JR (2004) 'Listening-in' to find and explore new combinations of customer needs. *Journal of Marketing* 68:72-87.
- Vargas S, P Castells (2011) Rank and relevance in novelty and diversity metrics for recommender systems. *RecSys '11*. Proceedings of the Fifth ACM Conference on Recommender systems Chicago, IL.
- Weitzman ML (1979) Optimal search for the best alternative. *Econometrica* 47(3):641-654.
- Whittle P (1988) Restless bandits: activity allocation in a changing world. *Journal of Applied Probability* 25:287-298.
- Ying Y, Feinberg F, Wedel M (2006) Leveraging missing ratings to improve online recommendation systems. *Journal of Marketing Research* 43(3):355-365.
- Zhang M, Hurley N (2008) Avoiding monotony: improving the diversity of recommendation lists. *RecSys '08*. Proceedings of the 2008 ACM Conference on Recommender Systems, Lausanne, Switzerland, 123-130.
- Zhou T, Z Kuscsik, J-G Liu, M Medo, JR Wakeling, Y-C Zhang (2010) Solving the apparent diversity-accuracy dilemma of recommender systems. *PNAS* 107(10):451-4515.
- Ziegler C-N, McNee, SM, Konstan JA, Lausen G (2005) Improving recommendation lists through topic diversification. *WWW '05*. ACM Proceedings of the 14th International Conference on World Wide Web, Chiba, Japan, 22-32.

Appendix 1: Formal Demonstrations

Consumer Search Criteria. *When consumers learn their preference weights, there exist product spaces in which the consumer, searching optimally without a recommendation may use search criteria that differ from the typical criteria in either search theory or recommendation systems. In particular, unlike typical criteria in search theory, the consumer may choose to search a product that does not have the highest option value (highest variance). Unlike typical recommendation systems criteria, the consumer may choose to search a product that does not have the highest expected utility. The product may even have a low or zero probability of being chosen.*

Proof. The formal analysis is based on binary attributes. However, to show the result we need at least three binary attributes and not a full-factorial product space. Thus, we consider a product space of $\vec{x}_1 = (0, 0, 0)$, $\vec{x}_2 = (1, 0, 0)$, $\vec{x}_3 = (0, 1, 0)$, $\vec{x}_4 = (0, 0, 1)$, $\vec{x}_5 = (1, 1, 0)$, $\vec{x}_6 = (1, 0, 1)$, and $\vec{x}_7 = (0, 1, 1)$. In a proof available from the authors, and as illustrated in §4, we demonstrate that the consumer will prefer searching those products that reveal two attributes, \vec{x}_5 , \vec{x}_6 , or \vec{x}_7 , rather than those products that reveal only one attribute.

Because this result is an existence proof, we need only show an example. (We actually show a class of examples.) For our example, we assume zero signal variance such that the posterior distributions are $\delta(w_i^r)$ when the consumer searches a product with $x_{ij} = 1$. We consider distributions for the consumer's prior beliefs in which the consumer prior beliefs assure that $w_1^r, w_2^r, w_3^r \geq 0$ and $\min\{w_3^r\} \geq \max\{w_1^r, w_2^r\}$. Our result is not limited to such distributions, but such distributions suffice. There are many distributions that satisfy these properties. For example, the conditions hold for uniformly distributed beliefs, $w_i \sim \mathcal{U}[a_i, b_i]$, with parameters $a_1, a_2 \geq 0$ and $a_3 \geq \max[b_1, b_2]$. The assumption of non-negative true importances simplifies the tree of conditions in the dynamic program and assures that the consumer weakly prefers \vec{x}_5 to \vec{x}_1, \vec{x}_2 , and \vec{x}_3 , weakly prefers \vec{x}_6 to \vec{x}_1, \vec{x}_2 , and \vec{x}_4 , and weakly prefers \vec{x}_7 to \vec{x}_1, \vec{x}_3 , and \vec{x}_4 . We need only consider a product space of \vec{x}_5, \vec{x}_6 , and \vec{x}_7 . These are really the most interesting products for our purposes. The outside option is $U^* = u(\vec{x}_1) = 0$.

We first consider searching on \vec{x}_5 . With $\min\{w_3^r\} \geq \max\{w_1^r, w_2^r\}$, the consumer would never choose \vec{x}_5 , but may consider searching \vec{x}_5 to learn w_1^r and w_2^r efficiently. We assume $\beta \leq 1$ is sufficiently large to justify search. After searching \vec{x}_5 , the consumer will either choose the outside option, one of \vec{x}_5, \vec{x}_6 , and \vec{x}_7 , or search \vec{x}_6 or \vec{x}_7 , and, perhaps, if the consumer searches, the consumer will continue thereafter. Using Equation 3 in the text we obtain:

$$(A7) \quad J(\{5\}, f^0) = \max\left\{0, \max\{w_1^r + w_2^r, w_1^r + \bar{w}_3, w_2^r + \bar{w}_3\}, \max_{k=6,7}\{-c + \beta E_{w_3}[J(\{5, k\}, f^1 | f^0)]\}\right\}$$

With $\min\{w_3^r\} \geq \max\{w_1^r, w_2^r\}$, there is no value to searching to reveal w_3^r , because knowing w_3^r does not change the consumer's decision. Using $w_1^r, w_2^r, w_3^r \geq 0$, we eliminate the outside option as a choice. Using $\min\{w_3^r\} \geq \max\{w_1^r, w_2^r\}$, we eliminate \vec{x}_5 as a choice. Hence, we obtain the cost of searching \vec{x}_5 as:

$$(A8) \quad \begin{aligned} -c + \beta E_{w_1, w_2}[J(\{5\}, f^0)] &= -c + \beta E_{w_1, w_2}[\max\{\max\{w_1^r + \bar{w}_3, w_2^r + \bar{w}_3\}, -c + 0\}] \\ &= -c + \beta E_{w_1, w_2}[\max\{w_1, w_2\}] + \bar{w}_3 \end{aligned}$$

The last step uses the consumer's prior beliefs to compute expected values for the w_i^r 's that are revealed by search.

We now consider searching on \vec{x}_6 . Using similar reasoning to Equation A7 we obtain:

$$(A9) \quad J(\{6\}, f^0) = \max\left\{0, \max\{w_1^r + \bar{w}_2, w_1^r + w_3^r, \bar{w}_2 + w_3^r\}, \max_{k=5,7}\{-c + \beta E[J(\{6, k\}, f^1 | f^0)]\}\right\}$$

We first examine the value of further search. Searching either \vec{x}_5 or \vec{x}_7 reveals w_2^r , so the consumer is indifferent between searching \vec{x}_5 or \vec{x}_7 . We replace the value of further search by the value of searching one of the two products. As in the case of searching \vec{x}_5 , further search beyond \vec{x}_6 and \vec{x}_5 or \vec{x}_6 and \vec{x}_7 has no value. Thus, we have, if the consumer were to choose to search:

$$(A10) \quad \begin{aligned} -c + \beta E_{w_2}[J(\{5,6\}, f^1)] &= -c + \beta E_{w_2}[J(\{6,7\}, f^1)] = -c + \beta E_{w_2}[\max\{w_1^r + w_3^r, w_2 + w_3^r\}] \\ &= -c + \beta E_{w_2}[\max\{w_1^r, w_2\}] + w_3^r \end{aligned}$$

If, after searching only \vec{x}_6 , the consumer were to choose without search, then w_1^r and w_3^r are revealed, but the consumer expects \bar{w}_2 if \vec{x}_5 or \vec{x}_7 are chosen. Recall that the outside option and \vec{x}_5 are dominated if consumer beliefs follow the example class of distributions. Thus, the value of choosing without search, that is, the second internal max in Equation A9, is given by the following.

$$(A11) \quad \max\{w_1^r + \bar{w}_2, w_1^r + w_3^r, \bar{w}_2 + w_3^r\} = \max\{w_1^r, \bar{w}_2\} + w_3^r$$

Putting it all together and factoring out w_3^r , we obtain:

$$(A12) \quad J(\{6\}, f^0) = \max\{\max\{w_1^r, \bar{w}_2\}, -c + \beta E_{w_2}[\max\{w_1^r, w_2\}]\} + w_3^r$$

And the value of searching \vec{x}_6 is:

$$\begin{aligned}
\text{(A13)} \quad -c + \beta E_{w_1, w_3} [J(\{6\}, f^0)] &= -c + \beta E_{w_1} [\max\{\max\{w_1^r, \bar{w}_2\}, -c + \beta E_{w_2} [\max\{w_1^r, w_2\}]\}] + \bar{w}_3 \\
&= -c + \beta E_{w_1} [\max\{w_1^r, \bar{w}_2, -c + \beta E_{w_2} [\max\{w_1^r, w_2\}]\}] + \bar{w}_3
\end{aligned}$$

Suppose $w_1^r \geq \bar{w}_2$, then, for all w_1^r , $\max\{w_1^r, \bar{w}_2, -c + \beta E_{w_2} [\max\{w_1^r, w_2\}]\} \leq \beta E_{w_2} [\max\{w_1^r, w_2\}]$ for sufficiently large $\beta \leq 1$. This is true because, for the last internal max, $\max\{w_1^r, w_2\} \geq w_1^r$. (It is certainly true for $\beta = 1$.) Hence $\beta E_{w_1} [\max\{w_1^r, \bar{w}_2, -c + \beta E_{w_2} [\max\{w_1^r, w_2\}]\} | w_1^r \geq \bar{w}_2] \leq \beta E_{w_1, w_2} [\max\{w_1, w_2\} | w_1^r \geq \bar{w}_2]$. Suppose $w_1^r < \bar{w}_2$, then, for all w_1^r , $\max\{w_1^r, \bar{w}_2, -c + \beta E_{w_2} [\max\{w_1^r, w_2\}]\} \leq \beta E_{w_2} [\max\{w_1^r, w_2\}]$. This is true because $\bar{w}_2 \leq \beta E_{w_2} [\max\{w_1^r, w_2\}]$ for sufficiently large $\beta \leq 1$. Thus, $\beta E_{w_1} [\max\{w_1^r, \bar{w}_2, -c + \beta E_{w_2} [\max\{w_1^r, w_2\}]\} | w_1^r < \bar{w}_2] \leq \beta E_{w_1, w_2} [\max\{w_1, w_2\} | w_1^r < \bar{w}_2]$ for sufficiently large $\beta \leq 1$. Putting these expressions together, establishes that $E_{w_1} [\max\{w_1^r, \bar{w}_2, -c + \beta E_{w_2} [\max\{w_1^r, w_2\}]\}] \leq \beta E_{w_1, w_2} [\max\{w_1, w_2\}]$. Comparing Equations A13 to A8, we have the result that $-c + \beta E_{w_1, w_3} [J(\{6\}, f^0)] \leq -c + \beta E_{w_1, w_2} [J(\{5\}, f^0)]$.

We have demonstrated that the consumer prefers to search \vec{x}_5 rather than \vec{x}_6 . The consumer's preference for searching \vec{x}_5 rather than \vec{x}_7 follows by symmetry. We have also demonstrated that, after the first product is searched, the consumer does not search further. Finally, we can easily show that the consumer will choose to search whenever $-c + \beta E_{w_1, w_2} [\max\{w_1, w_2\}] \geq \max\{\bar{w}_1, \bar{w}_2\}$. This must hold for some c and for sufficiently large $\beta \leq 1$. By design, all possible realized values of w_3 are greater than any possible realized value of either w_1 or w_2 , hence both the expected utility and the option value of $u(\vec{x}_6)$ and $u(\vec{x}_7)$ dominate the expected utility and option value of $u(\vec{x}_5)$. By option value we mean that $E[u(\vec{x}_5) | u(\vec{x}_5) \geq U^*] < E[u(\vec{x}_6) | u(\vec{x}_6) \geq U^*] = E[u(\vec{x}_7) | u(\vec{x}_7) \geq U^*]$. We have nowhere restricted the variances of prior beliefs. The result holds even if the variance of w_3 is greater than the variance of w_1 and the variance of w_2 . Finally, all possible realized values of w_3 are greater than any possible realized value of either w_1 or w_2 , the consumer will never choose \vec{x}_5 after searching \vec{x}_5 . This completes the proof. ■

Recommendation System Criteria. For recommendation systems that take preference-weight learning into account, the optimal product to recommend may not satisfy the typical criterion that it be the highest expected utility (most likely to be chosen). The product might even have a low or zero probability of being chosen. Furthermore, the product may not have the largest option value as might be expected from optimal search theory.

Proof. This result is also based on binary attributes. We consider a three-product product space: $\vec{x}_1 = (0, 0)$, $\vec{x}_2 = (1, 0)$, $\vec{x}_3 = (0, 1)$. Because this result is an existence proof, we need only show an

example. We begin by deriving general conditions and then show that there exists a class of examples that satisfy the general conditions.

We first examine the consumer's optimal search path in the absence of a recommendation. The consumer makes decisions based on the consumer's prior beliefs and revealed values. We are particularly interested in the case where one of the attributes is undervalued by the consumer, but not by the recommendation system: $\beta \int_0^\infty w_2 df_2^0(w_2) < c$, but $\beta \int_0^\infty w_2 df_2^{rec}(w_2) > c$. The other attribute is not undervalued: $f_1^{rec}(w_1) = f_1^0(w_1)$ and $\beta \int_0^\infty w_1 df_1^0(w_1) = \beta \int_0^\infty w_1 df_1^{rec}(w_1) > c$. We are interested in the case where the consumer searches (at least) \vec{x}_2 , so we further assume that $\beta \int_0^\infty w_1 df_1^0(w_1) - c > \bar{w}_1$. For simplicity we set the utility of the outside option to the utility of choosing $\vec{x}_1 = (0, 0)$ such that $U^* = u(\vec{x}_1) = 0$.

We evaluate whether the consumer searches \vec{x}_2 . The expected value of searching \vec{x}_2 is:

$$(A14) \quad E[J(\{\vec{x}_2\}, f^1)] = \max \left\{ 0, \bar{w}_1, -c + \beta \int_{\max\{0, w_1^r\}}^\infty w_2 df_2^0(w_2) \right\}$$

Recognizing that $\beta \int_{\max\{0, w_1^r\}}^\infty w_2 df_2^0(w_2) \leq \beta \int_0^\infty w_2 df_2^0(w_2) < c$ for $\beta \leq 1$, the consumer will not search \vec{x}_3 after \vec{x}_2 , hence $E[J(\{\vec{x}_2\}, f^1)] = \max\{0, \bar{w}_1\}$. Furthermore, because $\beta \int_0^\infty w_1 df_1^0(w_1) - c > \bar{w}_1$, $\beta E[J(\{\vec{x}_2\}, f^1)] - c > \bar{w}_1$, hence the consumer prefers to search \vec{x}_2 rather than simply choose \vec{x}_2 .

The expected value of searching \vec{x}_3 is:

$$(A15) \quad E[J(\{\vec{x}_3\}, f^1)] = \max \left\{ 0, \bar{w}_2, -c + \beta \int_{\max\{0, \bar{w}_2\}}^\infty w_1 df_1^0(w_1) \right\} \leq \max\{0, \bar{w}_2\} = \int_0^\infty w_2 df_2^0(w_2) < c$$

which implies the consumer will not choose to search \vec{x}_3 (without a recommendations). Putting these together, if no recommendations are make, the consumer would search \vec{x}_2 and then either purchase \vec{x}_2 if $w_1^r \geq 0$ or accept the outside option, $U^* = 0$, if $w_1^r < 0$. The expected payoff based on the consumer's beliefs is given in Equation A16. Because $f_1^{rec}(w_1) = f_1^0(w_1)$, this is also the recommendation system's beliefs about what the consumer will achieve without any recommendation.

$$(A16) \quad E[J(\emptyset, f^0)|f^{rec}] = -c + \beta \int_0^\infty w_1 df_1^0(w_1) = -c + \beta \int_0^\infty w_1 df_1^{rec}(w_1).$$

We now analyze the cases where the recommendation system recommends a product, \vec{x}_2 or \vec{x}_3 , and the consumer follows that recommendation. We assume the consumer acts optimally after the recommendation. (We will later consider what would happen if the recommendation system could recommend both products.) We seek to establish a case where the recommendation system recommends \vec{x}_3 even though $\int_0^\infty w_2 df_2^{rec}(w_2) < \int_0^\infty w_1 df_1^{rec}(w_1)$.

We first consider the expected payoff if the recommendation system recommends \vec{x}_2 only. If \vec{x}_2 is the recommendation, then the consumer does not deviate from the optimal path that the consumer would have chosen without a recommendation. Define $J(\{\vec{x}_2\}_{rec}, f^1)$ as the continuation value given that the consumer searched \vec{x}_2 due to a recommendation. Then:

$$(A17) \quad E[J(\{\vec{x}_2\}_{rec}, f^1) | f^{rec}] = -c + \beta \int_0^\infty w_1 df_1^{rec}(w_1).$$

We now consider the expected payoff if the recommendation system recommends \vec{x}_3 . After searching the recommended product, the consumer may either stop and purchase \vec{x}_3 , take the outside option ($U^* = 0$), or search the remaining product, \vec{x}_2 . (Our condition that $\beta \int_0^\infty w_1 df_1^0(w_1) - c > \bar{w}_1$ assures that the consumer will not purchase \vec{x}_2 without searching. It's easy to show that Equation A18 does not change that.) From the consumer's perspective:

$$(A18) \quad E[J(\{\vec{x}_3\}_{rec}, f^1 | f^0)] = \max \left\{ 0, \bar{w}_2, -c + \beta \int_{\max\{0, \bar{w}_2\}}^\infty w_1 df_1^0(w_1) \right\}$$

The consumer will search \vec{x}_2 if $\beta \int_{\max\{0, \bar{w}_2\}}^\infty w_1 df_1^0 > c + \bar{w}_2$. We are interested in the recommendation system's expectation of this payoff, or $E[J(\{\vec{x}_3\}_{rec}, f^1) | f^{rec}]$. This value depends on whether or not the consumer would choose to search \vec{x}_2 after searching \vec{x}_3 . Define \tilde{w}_2 to be the maximum observed value of w_2^r for which it would be optimal for the consumer to search \vec{x}_2 after \vec{x}_3 . This value is defined implicitly by $\beta \int_{\tilde{w}_2}^\infty w_1 df_1^0 = c + \tilde{w}_2$. If $w_2^r > \tilde{w}_2$, the consumer purchases \vec{x}_3 and does not search \vec{x}_2 , because $w_2^r > -c + \beta \int_{\tilde{w}_2}^\infty w_1 df_1^0(w_1)$.

To compute the expected value we consider three regions for the outcome of the \vec{x}_3 search. Each outcome corresponds to a different action by the consumer after searching \vec{x}_3 . These regions are: $w_2^r \leq 0$, $0 < w_2^r \leq \tilde{w}_2$, and $w_2^r > \tilde{w}_2$. Note that \tilde{w}_2 is defined by f^0 , but we will compute expectations based on the recommendation system's beliefs, f^{rec} .

Case 1: $w_2^r \leq 0$. In this region, the consumer expects to search \vec{x}_2 after \vec{x}_3 because $\beta \int_0^\infty w_1 df_1^0(w_1) > c$. After searching \vec{x}_2 , there are no products left to search; the consumer purchases \vec{x}_2 if $w_1^r > 0$ and takes the outside good if $w_1^r \leq 0$. In this region, the recommendation believes:

$$(A19) \quad -c + \beta E[J(\{\vec{x}_3\}_{rec}, f^1) | w_2^r \leq 0, f^{rec}] = -(1 + \beta)c + \beta \int_0^\infty w_1 df_1^{rec}(w_1)$$

Case 2: $0 < w_2^r \leq \tilde{w}_2$. In this region, by the definition of \tilde{w}_2 , the consumer expects to search \vec{x}_2 after \vec{x}_3 , after which there are no products left to search. The net expected payoff to the consumer is $-(1 + \beta)c + \beta \max\{0, w_1^r, w_2^r\}$ according to the recommendation system's beliefs.

Case 3: $\tilde{w}_2 < w_2^r$. In this region, by the definition of \tilde{w}_2 , the consumer chooses \vec{x}_3 and does not search \vec{x}_2 . But because $w_2^r > -c + \beta \int_{\tilde{w}_2}^{\infty} w_1 df_1^0(w_1) \geq -c + \beta \int_{w_2^r}^{\infty} w_1 df_1^0(w_1)$ and $w_2^r > 0$, we know:

$$(A20) \quad E[J(\{\vec{x}_3\}_{rec}, f^1 | f^{rec})] = \max \left\{ 0, w_2^r, -c + \beta \int_{\max\{0, w_2^r\}}^{\infty} w_1 df_1^0(w_1) \right\} \\ \geq -c + \beta \int_{\max\{0, w_2^r\}}^{\infty} w_1 df_1^0(w_1)$$

Thus, the net payoff in Case 3 is at least as large as that which the consumer would obtain by searching \vec{x}_2 after \vec{x}_3 , thus the next payoff is at least as large as $-(1 + \beta)c + \beta \max\{0, w_1^r, w_2^r\}$. By combining Cases 2 and 3, which occur according to the recommendation system's beliefs with probability, $\Pr(w_2 > 0) = \int_0^{\infty} w_2 df_2^{rec}(w_2)$, we obtain a lower bound on the recommendation system's beliefs for Cases 1, 2, and 3 as $-(1 + \beta)c + \beta \max\{0, w_1^r, w_2^r\}$:

$$(A21) \quad -c + \beta E[J(\{\vec{x}_3\}_{rec}, f^1) | f^{rec}] \\ \geq -(1 + \beta)c + \beta \int_{w_2=-\infty}^0 \int_{w_1=0}^{\infty} w_1 df_1^{rec}(w_1) df_2^{rec}(w_2) \\ + \beta \int_{w_2=0}^{\infty} \int_{w_1=-\infty}^0 w_2 df_1^{rec}(w_1) df_2^{rec}(w_2) + \beta \int_{w_2=0}^{\infty} \int_{w_1=w_2}^{\infty} w_1 df_1^{rec}(w_1) df_2^{rec}(w_2) \\ + \beta \int_{w_2=0}^{\infty} \int_{w_1=0}^{w_2} w_2 df_1^{rec}(w_1) df_2^{rec}(w_2)$$

Note that we would also obtain the right-hand side of Equation A21 if the recommendation system were to recommend both \vec{x}_2 and \vec{x}_3 . Thus, we have shown that recommending \vec{x}_3 alone weakly dominates recommending both products for sufficiently large $\beta \leq 1$. We now rearrange the limits of integration to obtain.

$$\begin{aligned}
& -c + \beta E[J(\{\vec{x}_3\}_{rec}, f^1) | f^{rec}] \\
& \geq -(1 + \beta)c + \beta \int_{w_2=-\infty}^{\infty} \int_{w_1=0}^{\infty} w_1 df_1^{rec}(w_1) df_2^{rec}(w_2) \\
& + \beta \int_{w_2=0}^{\infty} \int_{w_1=-\infty}^0 w_2 df_1^{rec}(w_1) df_2^{rec}(w_2) \\
(A22) \quad & + \beta \int_{w_2=0}^{\infty} \int_{w_1=0}^{w_2} (w_2 - w_1) df_1^{rec}(w_1) df_2^{rec}(w_2) \\
& \geq \beta E[J(\emptyset, f^0) | f^{rec}] - c + \beta \int_{w_2=0}^{\infty} \int_{w_1=-\infty}^0 w_2 df_1^{rec}(w_1) df_2^{rec}(w_2) \\
& + \beta \int_{w_2=0}^{\infty} \int_{w_1=0}^{w_2} (w_2 - w_1) df_1^{rec}(w_1) df_2^{rec}(w_2)
\end{aligned}$$

Equation A22 is a general condition for when the recommendation system will recommend \vec{x}_3 to the consumer. All that remains to complete the proof is to establish at least one example where the last two integrals in Equation A2 exceed c . We can do this for many distributions. We do it for at least one.

We consider uniform distributions, all of which have the zero mean: $f_1^0(w_1) = f_1^{rec}(w_1) = \mathcal{U}[-b_1, b_1]$, $f_2^{rec}(w_2) = \mathcal{U}[-\alpha b_1, \alpha b_1]$, and $f_2^0(w_2) = \mathcal{U}[-b_2, b_2]$, where $\alpha < 1$ assures the option value of \vec{x}_3 is less than the option value of \vec{x}_2 and the variance of $u(\vec{x}_3)$ is less than the variance of $u(\vec{x}_2)$. If $\beta = 1$, we assure $\beta \int_0^{\infty} w_2 df_2^0(w_2) < c$ with $b_2 < 4c/\beta$. We assure $\beta \int_0^{\infty} w_1 df_1^0(w_1) = \beta \int_0^{\infty} w_1 df_1^{rec}(w_1) > c$ and $\beta \int_0^{\infty} w_2 df_2^{rec}(w_2) > c$ with $\alpha b_1 > 4c/\beta$. The value of the next-to-last integral is $\beta \alpha b_1/8$ and the value of the last integral is $\beta \alpha^2 b_1/24$. Thus, for any positive $\alpha < 1$, the result holds as long as $\alpha^2 b_1/24 + \alpha b_1/8 > c/\beta$. Because $\alpha < 1$, we have that $\alpha^2 < \alpha$. Thus, we require only that $\alpha^2 b_1 > 6c/\beta$. This condition, and the condition $\alpha b_1 > 4c/\beta$, are satisfied for many values of α and b_1 . For example, if $\alpha = 1/2$, then it is sufficient that $b_1 > 24c/\beta$.

$Prob\{choose x_3\} = \frac{1}{4} + \frac{\alpha}{8} < \frac{1}{2} - \frac{\alpha}{8} = Prob\{choose x_2\}$ for all $\alpha < 1$. Finally, replacing the support of

w_2 with $f^{rec}(w_2) = \mathcal{U}[-\alpha b_1 - \epsilon, \alpha b_1 - \epsilon]$, condition A22 becomes $-c/\beta + \frac{(\alpha b_1 - \epsilon)^2}{8\alpha b_1} + \frac{(\alpha b_1 - \epsilon)^3}{24\alpha b_1^2} > 0$.

We choose $\epsilon > 0$ so that $E^{rec}[u(\vec{x}_3)] < E^{rec}[u(\vec{x}_2)]$. With $\alpha = 1/2$, $b_1 = 3$, and $c/\beta = 0.1$, setting $\epsilon = 0.05$ suffices. This completes the proof. ■

Appendix 2: Details of Synthetic Data Experiments

Formulae for Bayesian updating. We assume the priors are normally distributed, $f_i^0 = \mathcal{N}(\bar{w}_i^0, \sigma_i^0)$, and the signal from searching product j is $g(w_i^r | w_i, s_t = j) = \mathcal{N}(w_i^r, \sigma_i^s)$ when $x_{ij} = 1$, where σ_i^s is the standard deviation of the signal. Then, $f_i^1(w_i) \equiv f_i^1(w_i | w^r, s_t = j) = \mathcal{N}(\bar{w}_i^1, \sigma_i^1)$ with $\bar{w}_i^1 = [(\sigma_i^s)^2 \bar{w}_i^0 + (\sigma_i^0)^2 w_i^r] / [(\sigma_i^s)^2 + (\sigma_i^0)^2]$ and $\sigma_i^1 = \sqrt{(\sigma_i^s)^2 (\sigma_i^0)^2 / [(\sigma_i^s)^2 + (\sigma_i^0)^2]}$. Subsequent updates follow similar formulae. All posteriors remain normally distributed. See Schlaifer (1959, p. 441).

Figures 1 and 2. Figures 1 and 2 are based on a product space with three six-level attributes. All $6 \times 6 \times 6 = 216$ products are available. The consumer's prior beliefs are normally distributed, independently over aspects, with $f_i^0 = \mathcal{N}(\bar{w}_i^0, \sigma_i^0) \forall i$. The specific values are given in Table A1.

Table A1. Generating Parameters for the Synthetic Data in Figures 1 and 2

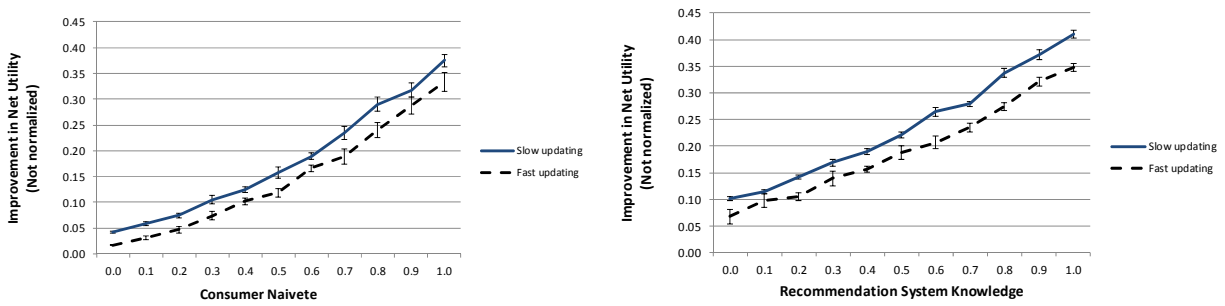
Attribute	Level	Prior mean, \bar{w}_i^0	Prior variance, $(\sigma_i^0)^2$	True weight, w_i^r
1	1	3.362	0.216	2.405
1	2	1.011	0.002	2.293
1	3	1.632	0.046	-1.957
1	4	1.743	0.143	2.412
1	5	3.783	0.204	3.276
1	6	0.437	0.097	3.833
2	1	3.288	0.042	2.833
2	2	3.937	0.297	2.632
2	3	0.049	0.485	3.654
2	4	2.688	0.267	2.681
2	5	3.346	0.082	2.946
2	6	2.746	0.054	3.058
3	1	0.204	0.232	-0.836
3	2	1.441	0.314	2.529
3	3	5.140	0.311	3.610
3	4	-0.709	0.616	-0.176
3	5	0.294	0.275	3.314
3	6	0.106	0.002	4.419

One of the 216 products is recommended to the consumer, after which the consumer searches optimally. The vertical axis in both figures is the utility of the product, net of search costs, chosen by the consumer after optimal search. The horizontal axis in Figure 1 is the true utility of the recommended product. The horizontal axis in Figure 2 is the recommendation system's beliefs about the utility of the recommended product. Each point in Figure 2 represents one of the 216 feasible recommended products.

Figure 3. Consumers in Recommendation-System Synthetic-Data Experiments. Consumer beliefs are normally distributed and independent over aspects. For each consumer we draw the true aspect importances, \vec{w}_r , from a mixture of two normal distributions, one with mean 0.0 and variance 0.5; the other with mean 4.0 and variance 0.2. The mixing parameter, 0.7, favors the second normal distribution to illustrate a case where most of the aspects are important. We draw the variances, v_i , of the consumer’s prior beliefs i.i.d. from an exponential distribution, $\lambda e^{-\lambda v}$, with parameter, $\lambda = 10$. This gives $E[v] = 0.1$ and $E[\sigma_i^0] = 0.32$. In Figure 3, we vary naïveté with a parameter, η . We set the mean of the consumer’s priors to a “naïve” value with probability, η , and equal to w_i^r with probability, $1 - \eta$. The naïve value is redrawn randomly from the same generating distribution—a mixture of two normal distributions. For Figure 4, we hold consumer naïveté constant at $\eta = 0.6$.

Figure 4. Recommendation System Knowledge. We vary recommendation system knowledge with two parameters, P^c and P^{rec} . For some aspects we set the recommend system’s priors to the consumer’s priors (probability P^c). For the remaining aspects we set the mean of the recommendation system’s priors to the true partworths with probability P^{rec} and randomly with probability $1 - P^{rec}$. For the synthetic data experiments we set $P^c = 0.3$. For Figure 3, we hold recommendation system knowledge constant at $P^{rec} = 0.6$.

Figure A1. Benchmark Performance as a Function of the Experimental Variables.



Benchmark Performance. In Figure A1, as expected, the value of a the benchmark recommendation system (maximize expected utility) increases when consumers are more naïve and when recommendation system knowledge increases. The vertical axes in Figures 3 and 4 are relative to this benchmark. The improvement in performance due to the benchmark is larger when consumers update their preferences more slowly. The explanation is simple; the vertical axis in Figure A1 is the net utility achieved using the benchmark recommendation system, minus the utility achieved without a recommendation. The rate of updating has a larger impact on the net utility achieved without a recommendation than it does on the net utility achieved by following the benchmark recommendations.