Split-Second Decisions During Online Information Search: Static vs. Dynamic Decision Thresholds for Making the First Click

May 2017

PLEASE DO NOT CITE WITHOUT PERMISSION FROM AUTHORS

Joy Lu

Doctoral Candidate in Marketing

University of Pennsylvania

727.9 Jon M. Huntsman Hall

3730 Walnut Street, Philadelphia, PA 19104

(215) 898-9894

tonglu@wharton.upenn.edu

J. Wesley Hutchinson

Professor of Marketing

University of Pennsylvania

746 Jon M. Huntsman Hall

3730 Walnut Street, Philadelphia, PA 19104

(215) 898-6450

jwhutch@wharton.upenn.edu

Acknowledgement: The authors thank Eric Bradlow, Raghurem Iyengar, Robert Meyer, Barbara Mellers, Peter Fader, Aviv Nevo, and Elisabeth Honka for providing helpful comments and suggestions. Support from the Wharton Behavioral Lab and the Jay H. Baker Retailing Center PhD Grants is greatly appreciated.

Split-Second Decision Making During Online Information Search: Static vs. Dynamic Decision Thresholds for Making the First Click

Abstract

Just like the physical entrance of a brick-and-mortar store, the landing page of an online retailer serves to draw people deeper into the website. As with most sequential behaviors, landing page decisions are likely to have important downstream consequences. Our objective is to explore the internal decision processes and external characteristics of the shopping environment that determine how extensively shoppers search on the landing page. Using eye-tracking data collected during an incentive-compatible online shopping experiment, we build a sequential sampling model that captures how product information is acquired through "split-second" eye fixation decisions. The model assumes that the shopper clicks on a fixated link when the perceived value of that link crosses a decision threshold; otherwise, the shopper continues to search the landing page. First, we investigate whether this decision threshold varies over time, either in a "next-step" way that takes into account the potential gains of continuing search with one more fixation, or in a "whole-search" way that heuristically accounts for how the value of search changes over time. We find that whole-search dynamics better fit the data compared to next-step dynamics and the static threshold. Second, through counterfactual simulations, we demonstrate that under the assumptions of our model, the length of search changes significantly if products are ordered by individual shopper preferences on the page (i.e., worst-first or best-first), and then verify these predictions empirically in a follow-up incentive-compatible experiment. Our work contributes to the emerging stream of research that treats eye fixations as endogenous decisions and provides implications for how retailers should position products within the layout of an online store.

Keywords: *eye-tracking, search models, sequential sampling, information processing*

1. Introduction

Marketers have long held that increasing foot traffic within physical stores and the click through rate within virtual environments can lead to more purchases. Brick-and-mortar retailers create attention-grabbing window displays to draw people farther into the store and consider more products (Tice 2012). Just like the physical entrance of a brick-and-mortar store, the landing page of an online retailer serves to draw people deeper into the website. Online retailers will often fill the landing page with attractive product images that link to more detailed product information, or with category links that help guide shoppers to more organized groups of products. Thus, the landing page is especially important because the first click reflects both the shopper's goals and the retailer's desire to influence downstream shopper decisions.

Despite these intuitions and the rapid growth of online retail, to our knowledge, no work has been done on how people make decisions within the landing page of an online store. We address this gap in the literature by modeling the internal decision processes and external characteristics of the shopping environment that determine how extensively shoppers search on this first page before clicking on either a product or category link that takes them deeper into the virtual store. How long shoppers spend in certain areas of a store may have downstream consequences regarding their consideration set and final purchases. For example, Hui et al. (2013) recorded the shopping pathways of shoppers within a grocery store using radio frequency identification (RFID) tracking and found that increasing travel distance within a store leads to more unplanned purchases. More generally, Bronnenberg, Kim, and Mela (2016) find strong state dependence in the search path through the attribute space of digital cameras in online search such that early search is highly predictive of the characteristics of the camera eventually purchased, and that searched attribute levels converge towards these chracteristics.

To study the search processes of shoppers on the landing page of an online store, we conducted an incentive-compatible eye-tracking experiment with female undergraduates shopping for 5 minutes at the website of the clothing retailer American Apparel. Shoppers began at the landing page of the website and found a grid of products from three categories (10 tops, 10 sweaters, and 10 dresses), plus links to every product category available in the store (see Figure 1). Using eye-tracking, we were able to observe where on the page participants were looking at each moment in time. Our units of analysis are the locations of gaze during eye fixations, defined as when the eyes are still for at least 50 ms – long enough for individuals to process visual information about the fixated area.

[Insert Figure 1 about here]

In order to identify the information being gathered by the shopper at each eye fixation, we divided up the landing page into 31 areas of interest (or AOIs). There was one AOI for each of the 30 products on the page, and one AOI containing the names of the product categories. Each AOI contained one or more links that would bring the shopper to a new webpage within the store's website. Shoppers could either make an eye fixation on the information contained within an AOI or click on a link within an AOI. For the remainder of the paper, we refer to the AOIs containing product information and links as Product AOIs, and the AOI containing the category names and links as the Category AOI. The landing page was identical for all shoppers in our study; however, shoppers were randomly assigned to either hedonic or utilitarian shopping goals. After clicking on the link, they were free to explore any pages on the store's website and add items to their virtual shopping cart.[1]

In our data, we find that shoppers who spent more time searching on the landing page subsequently explored less of the rest of the store's website (about one less unique webpage per 20 additional eye fixations[2]) and put less in their shopping carts (about \$0.58 less per additional eye fixation[3]). Our findings suggest that the initial decision about how to leave

[1]Fixations that fell outside of the 31 AOIs (i.e., on empty white space, the store logo, etc.) were removed. Note that in this research we only model the landing page fixations and first click. However, we use final purchases to demonstrate the downstream consequences of landing page search. The entire shopping trip will be examined in future research.

[2]$t = -13.10, p < 0.01, R^2 = .11$

[3]$t = -1.95, p = .06, R^2 = .04$

the landing page is an important determinant of subsequent search and that, consistent with Hui et al. (2013), perhaps it is better to send shoppers deeper into the store than encourage them to linger on the landing page.

1.1. Research Contributions

Our main contribution is to the emerging stream of research that endogenizes the visual search process (Chandon et al. 2009; Stüttgen, Boatwright, and Monroe 2012; Shi, Wedel, and Pieters 2013; Towal, Mormann, and Koch 2013; Yang, Toubia, and de Jong 2015). We model eye fixations on the landing page as a series of "split-second" decisions (50-500 ms) that depend on the shopper's valuations of the AOIs and the physical effort of making eye movements (i.e., shoppers need to use more effort to fixate on farther AOIs). We use a sequential sampling framework to model how shoppers collect information through eye fixations. Sequential sampling models have been shown to be psychologically and neurologically valid accounts of how consumers make fast decisions (Busemeyer and Townsend 1993; Ratcliff and Smith 2004; Otter et al. 2008; Gluth, Rieskamp, and Büchel 2013). In our model, as shoppers fixate on AOIs, the AOIs accumulate perceived value or "attractiveness" (which can be positive or negative), and when the attractivenesss of one of the AOIs crosses a decision threshold, the shopper clicks on a link.

Since we conducted our study within a controlled experiment, we were able to measure each shopper's "true" enduring preferences for all products they saw on the landing page in a survey conducted 60-90 days after the shopping experience (i.e., liking rating on a 7-point Likert scale). These individual level preference measures were used to specify the information learned through eye fixations. Each fixation moved the perceived value of an AOI away from a neutral starting point toward the true value.

We directly estimate the value of the decision threshold, and test variations that represent contrasting dynamics identified in prior empirical and theoretical research on "boundedly rational," forward looking decision making. First, we allow the threshold to vary in the short run as a function of the expected gains from continued search with one more fixation, which we refer to as "next-step" dynamics. Prior work has demonstrated that it may be more accurate to consider consumers as boundedly rational, one-step forward-looking rather than infinitely forward-looking (Gabaix 2006; Yang, Toubia, and de Jong 2015). Second, we allow the threshold to vary over time in the long run, which we refer to as "whole-search" dynamics. As time progresses, individuals may employ a valid heuristic by allowing their decision thresholds to decrease because more time spent searching implies lower value of available options. For a variety of decisions, whole-search dynamics have been characterized as explore/exploit decision-making (March 1991; Huang and Hutchinson 2013), shifting from search mode to choice mode (Dellaert and Haübl 2012), and impatience or focus on accumulated search costs (Adam 2001; Brown, Flinn, and Schotter 2011; Häubl, Dellaert, and Donkers 2010; Koulayev 2014). In summary, we allow the decision threshold to vary according to two boundedly rational mechanisms: the next-step dynamics of the decision threshold allow shoppers to think carefully in the short term, while the whole-search dynamics allow shoppers to think heuristically across the full time horizon of search on the landing page.

We assess how well models with different decision threshold specifications are able to capture shopper behavior on the landing page by simulating fixation pathways using the individual-level posteriors of the parameters obtained through hierarchical Bayes estimation. We are specifically interested in accurately predicting the length and breadth of search, as measured by the number of fixations and the number of unique AOIs seen, as well as whether shoppers ultimately clicked on a link in a Product or Category AOI. Note that we do not assume that shoppers are searching optimally; however, through counterfactual simulation with different decision threshold values, we can test ex-post how closely shoppers behave compared to optimal search.

We also conducted counterfactual simulations to predict shopper behaviors on the landing page if products were positioned in a different order based on their individual-level preferences. Our simulations predicted that if products were positioned in a "worst-first" order with the least-liked products displayed first (from left to right, top to bottom), shoppers would search longer and farther down the webpage but would be more likely to click on the Category AOI rather than on one of the Product AOIs. On the other hand, a "best-first" order was predicted to result in less search but a higher likelihood of clicking on one of the Product AOIs. To empirically test these predictions, we conducted a second incentive-compatible experiment with participants shopping at the landing page of a different retailer, Forever 21, where the order of products was customized to each shopper based on her individual-level preferences measured in a pre-screening survey. As predicted, we found that shoppers presented with a worst-first positioning of products made more fixations and looked at products farther down the landing page compared to shoppers presented with a best-first positioning.

2. Model Overview

We are interested in modeling the "scan path" of eye fixations within the landing page of an online store as a sequence of split-second choices leading up to the decision to click on a link. Path data (e.g., grocery shopping, eye tracking, web browsing, etc.) can be used by marketers to understand the dynamic decision processes and goal orientations of consumers as they interact with their virtual or physical shopping environment (Hui, Fader, and Bradlow 2009). Our model captures how shoppers fixate on different AOIs on the landing page and how they decide to stop searching by clicking on a link within one of the 31 Product or Category AOIs.

2.1. Process Models of Visual Attention and Information Search

We use eye fixations as a measure of visual search on the landing page. The stream of eye movements is divided up into fixations, defined as when the eyes are still for at least 50 ms, and saccades, defined as when the eyes are rapidly moving. Only fixated visual information can be processed at a detailed level. For example, a single fixation can encode 5-7 letters when reading text and one product image when identifying image features (Rayner 1998). More generally, the human eye is in constant motion, constructing the visual world from the information obtained by each fixation, and directing the eyes to the parts of the world that are most likely to provide goal relevant information (Corbetta, Patel, and Shulman 2008; Hutchinson, Lu, and Weingarten 2016).

Fixations indicate the information accumulated about AOIs, and shoppers can re-fixate or revisit AOIs to gather more information. After each fixation on an AOI, we assume that shoppers update their beliefs about the AOI. These beliefs can be about objective information like the price of a product, or subjective valuations like the shopper's preference for the product. These beliefs form the shopper's perception of the attractiveness of each AOI and evolve over time through a sequential sampling process. The shopper clicks on an AOI once the attractiveness of one of the AOIs crosses a decision threshold.

A large body of literature within psychology and neuroscience supports a sequential sampling process for how people make fast decisions (split-second or within a few seconds). Random walk models like Decision Field Theory (Busemeyer and Townsend 1993), diffusion models (Ratcliff and Tuerlinckx 2002; Satomura, Wedel, and Pieters 2014), and Poisson race models (Van Zandt, Colonius, and Proctor 2000; Huang and Hutchinson 2008) all fall under the umbrella of sequential sampling models (Townsend and Ashby1983; Ratcliff and Smith 2004; Otter et al. 2008). Sequential sampling models assume that preference evolves as a stochastic process over time until a decision threshold is reached. Treating the decision process as a stochastic accumulation of information can explain the differences in reaction times for simple versus complex choices, speed-accuracy tradeoffs, memory retrieval, and serial position effects. Branco, Sun, and Villas-Boas (2012) describes a similar model for product search (albeit across a much longer time horizon) in which product knowledge is accumulated over time until it reaches a threshold for purchase or no purchase. Krajbich, Armel, and Rangel (2010) developed the attentional drift diffusion model that assumes the latent evolving preferences described by sequential sampling models can be measured via eye movements; the direction of "drift" or accumulation of evidence is indicated by which option is being looked at.

In our model, the sequential sampling of information via eye movements is able to capture the differences in how long shoppers spend searching on the landing page before they click on a link within an AOI once the decision threshold is reached. Note that traditional sequential sampling models treat the accumulation of information as an exogenous random walk. For example, Huang and Hutchinson (2008) model the retrieval of memories of ad information as a Poisson arrival process, while Krajbich, Armel, and Rangel (2010) treat eye fixations on left vs. right options within binary choices as exogenously determined. We endogenize the accumulation of information by modeling eye fixations as split-second decisions that depend on the value and physical position of the AOI options.

2.2. Eye Movements as Split-Second Decisions

Research in natural scene perception has long held that eye movements are not random, but actually planned sequences of voluntary motor decisions (Zingale and Kowler 1986) and that people are capable of employing top-down gaze control processes to adjust their fixation patterns (Corbetta and Shulman 2002; Jovancevic-Misic and Hayhoe 2009). In addition, when making decisions, people typically demonstrate a "gaze cascade" by re-fixating with increasing frequency on more attractive options, resulting in a positive feedback loop of visual attention and preference (Shimojo et al. 2003); this suggests that visual attention is positively correlated with an individual's moment-to-moment subjective valuations.

Within the domain of consumer choice, researchers have recently attempted to endogenize the search process using more sophisticated models that treat fixations as decisions that depend on state variables. Within a store environment, factors such as shelf position, facings, and price displays influence where consumers focus visual attention, which Chandon et al. (2007) modeled as a decision pathway of "consider" and "look" choices leading up to a purchase. Similarly, the drift diffusion model has been extended to incorporate the value and visual salience of options to predict visual pathways (Towal, Mormann, and Koch 2013). A shopper's sequence of eye movements has also been modeled as following specific patterns within latent states (Liechty, Pieters, and Wedel 2003; Stüttgen, Boatwright, and Monroe 2012; Shi, Wedel, and Pieters 2013; Wedel, Pieters, and Liechty 2008), and as a one-step forward-looking utility maximization process (Yang, Toubia, and de Jong 2015). Here, we assume that shoppers make the decision to fixate on AOIs based on the attractiveness of the AOIs and the costs of making the physical eye movements towards the AOIs.

2.3. Relationship to Optimal Search

In our model, each decision "event" for the shopper occurs in two stages. In the first stage, she decides probabilistically whether the current AOI she is fixating on has crossed the decision threshold. If yes, then she clicks on the currently fixated product; if not, she moves on to the second stage, in which she chooses the next AOI to fixate on, depending on the attractiveness of the AOI and the physical effort of making the eye movement.

The decision threshold model that we propose shares similarities and differences with problems of optimal sequential search and discrete choice in which individuals maximize the discounted present value of rewards. In sequential search with recall, the optimal solution is characterized by the reservation values of each option, defined as the value of the "sure thing" that would make the cost of searching an option equivalent to the expected benefit (Weitzman 1979; Kim, Albuquerque, and Bronnenberg 2010; Honka and Chintagunta 2016). Shoppers search options in descending order of reservation values and learn the true value or attractiveness of each option with a single search, and stop when there are no more options with reservation values that exceed the best option on-hand. Kim, Albuquerque, and Bronnenberg (2010) in particular assume a full-information model in which the consumers have full knowledge about the products and their attribute values, which allows them to derive an expression for the reservation utility of each option.

Similarly, in our model, we assume that shoppers examine the options on the landing page based on their attractiveness. However, there are two important distinctions. First, our model assumes that shoppers can revisit (i.e., re-fixate on) options multiple times to gather more information and learn the true attractiveness values of options, rather than having full information before search or resolving uncertainty with a single search. Second, shoppers have a single decision threshold for all options that determines when to end search on the landing page, rather than separate reservation utilities for each option that are computed before search.

To allow for multiple visits of the same option, Lippman and McCardle (1991) proposed a model of uncertain search that incorporates uncertainty resolution into the standard sequential search problem. The optimal solution involves an upper and lower threshold that determines whether enough information about an option has been sampled to accept or reject the option. Koulayev (2014) allows reservation utilities to decrease across search periods via increasing search costs, so individuals can choose a previously-seen option. However, these models are not able to capture the flexibility observed in our data (and in eye-tracking data in general), in which shoppers can also make multiple non-consecutive re-visits to different options to gather information.

Finally, sequences of discrete choices (i.e., products, brands, investments) can be characterized by a multi-armed bandit, where options are represented by "arms" that yield an uncertain payoff when "pulled," and arms can be re-visited at any time. Just as in sequential search problems, the optimal solution of a multi-armed bandit problem involves an index assigned to each arm and individuals choosing the arm with the largest index (Gittins 1979; Gittins and Jones 1979). Lin, Zhang, and Hauser (2015) demonstrate that index strategies are close to optimal (relative to dynamic programming solutions) within a restless bandit problem of discrete brand choice. In multi-armed bandits, individuals receive payoffs with each choice, but in our model of landing page search (and in search problems in general) the only payoff comes from the single accepted choice when search ends (i.e., with a click on a link within an AOI).

2.4. Boundedly Rational Search

In our decision threshold model, shoppers are weighing the value of an option against the expected gains from continued search, with this tradeoff captured by the parameterization of the decision threshold. Shoppers decide which AOIs to fixate on by weighing AOI attractiveness against the effort of making the eye movements, and click on an AOI once the attractiveness of the AOI crosses the decision threshold. Additionally, our model assumes that shoppers accumulate evidence about the attractiveness of an AOI over multiple fixations, so they often revisit AOIs during search. Our model endogenizes the fixation and click choices of shoppers, but we do not assume that each shopper's decision threshold is optimal. This allows us to determine through counterfactual simulation if they could have obtained better search outcomes with different threshold values.

We specify the decision threshold stopping rule to take into account various behavioral influences on consumer decisions in product search, often referred to as boundedly rational search processes. A static threshold remains the same over time and resembles a satisficing model (Simon 1955, Reutskaja et al. 2011; Stüttgen, Boatwright, and Monroe 2012).

To incorporate next-step dynamics that capture short-run time horizons, we allow the decision threshold to vary base on the expected gains if the next fixation was the last opportunity to search (i.e., "one-step" forward-looking). Gabaix et al. (2006) demonstrated that this type of limited time horizon model predicted a sequence of choices that more closely resembled how lab participants behaved during a product search task, compared to the sequence predicted by the Weitzman index strategy. Yang, Toubia, and de Jong (2015) applied this approach to visual search within a conjoint experiment.

Consumers are limited in their time, cognitive capacity, and others resources that can be spent on product search (Shugan 1980; Payne, Bettman, and Johnson 1993), so they may employ choice heuristics (Hey 1981; Hey 1982; Gilbride and Allenby 2004) or adjust their search strategies to make it easier to navigate information across the whole search process

(Bettman and Kakkar 1977) and make “fast and frugal” decisions that are not far from optimal in many situations (Gigerenzer and Goldstein 1996). Researchers have observed that lab participants demonstrate declining reservation utilities over time during stylized search tasks, even when the underlying sampling distributions were stationary, possibly because they focus on the elapsed time or accumulated search costs (Häubl, Dellaert, and Donkers 2010; Brown, Flinn, and Schotter 2011). Declining reservation utilities may also be explained by rising search costs (Koulayev 2014), or consumers shifting from from explore mode to exploit mode (March 1991; Huang and Hutchinson 2013), or from search mode to choice mode (Dellaert and Häubl 2012). To capture whole-search dynamics, in particular the decline due to heuristic inferences about the decreasing gains of continuing search or consideration of increasing search costs, we allow the decision threshold to change over time across the whole search process on the landing page.

3. Model Specification

Let $i = 1, ..., I$ denote shoppers and let $m = 1, ..., M_i$ denote shopper i's sequence of events. At each event m, shopper i chooses an alternative $y[m] \in Y = \{\text{click}\} \cup \{\text{fixate on AOI } j\}$, where $j \in \{1, 2, ..., 31\}$ represents the possible AOI choices. Each event m occurs in two stages. In the first stage, with probability $P_{iC}[m]$, the shopper decides to click on the AOI that was fixated on during the previous event m. If she doesn't click, with probability $1 - P_{iC}[m]$, then search continues and she moves to the second stage. In the second stage, the shopper selects the next AOI upon which to fixate, with probability $P_{ijF}[m]$ for each AOI j.

3.1. Stage 1: Click vs. Continue Search on the Landing Page

Clicking on an AOI ends search on the landing page and brings the shopper to another page in the store. The shopper's subjective estimate of the value of clicking on AOI j at event m is $W_{ij}[m] = V_{ij}[m] + \epsilon_{iW}$. The shopper's decision threshold is $W_{iR}[m] = R_i[m] + \epsilon_{iW}$

and represents the value of continuing search on the landing page.

The probability $P_{iC}[m]$ that the shopper will click on an AOI is modeled as a binary logit that weighs the value of clicking against the decision threshold, assuming the error term ϵ_{iW} follows the type-1 extreme value distribution. Note that we use an indicator function to restrict clicks to the AOI that was fixated on during the previous event $m-1$. Intuitively, it makes sense that shoppers would only select the AOI they had just fixated on instead of choosing without looking. Also, in our data we did in fact find that all shoppers selected a currently fixated AOI.

$$P_{iC}[m] = \sum_{j=1}^{31} \frac{e^{V_{ij}[m]}}{e^{V_{ij}[m]} + e^{R_i[m]}} \underbrace{\mathbb{1}(y[m\text{-}1] = j)}_{\text{Previously fixated AOI}} \tag{1}$$

3.2. Stage 2: Fixation Decision

If the shopper chooses to continue search instead of clicking on the AOI fixated at $m-1$, she must then choose which AOI to fixate on next. The value of fixating on AOI j is given by $U_{ij}[m] = V_{ij}[m] + D_{ij}[m] + \epsilon_{iU}$. The term $V_{ij}[m]$ represents the "attractiveness" of AOI j and includes attributes like price and shoppers' self-reported liking ratings of products. The term $D_{ij}[m]$ accounts for the physical and mental effort of eye fixations, which depend on the relative positions of the AOIs. Assuming the error term ϵ_{iU} follows the type-1 extreme value distribution, the probability of fixating on AOI j is modeled as a multinomial logit multiplied by the probability of not clicking and continuing search:

$$P_{ijF}[m] = \frac{e^{V_{ij}[m]+D_{ij}[m]}}{\sum_k e^{V_{ik}[m]+D_{ik}[m]}} \underbrace{(1 - P_{iC}[m])}_{\text{Prob. no click}} \tag{2}$$

Note that since shoppers are assumed to only select previously fixated AOIs, $D_{ij}[m]$ does not appear in the value of *clicking on* an AOI since we assume that it takes no effort to click on a fixated AOI.

3.3. AOI Attractiveness and Search Costs

The value of fixating on an AOI, $U_{ij}[m]$, consists of a linear combination of S attributes that contribute to the AOI's attractiveness and T costs that account for the physical effort of fixating. Let X_i be the $J \times S \times M$ matrix of attribute values and let β_i be the vector of coefficients for the attributes (recall that J is the number of AOIs and M is the number of events). Let Z_i be the $J \times T \times M$ matrix of costs and let γ_i be its vector of coefficients. The attractiveness $V_{ij}[m]$ and search cost $D_{ij}[m]$ associated with fixating on AOI j are given by:

$$V_{ij}[m] = \sum_{s=1}^{S} \beta_{is} x_{ijs}[m] \tag{3}$$

$$D_{ij}[m] = \sum_{t=1}^{T} \gamma_{it} z_{ijs}[m] \tag{4}$$

3.4. Updating of Perceived Product Attribute Values

Possible attributes include shoppers' subjective valuations of the products as well as objective product features like price. We assume that the perceived attribute values start at the same neutral point for all AOIs and are updated to approach their true values as information is accumulated over time. True values might be larger or smaller than the neutral starting value. If the shopper fixates on AOI j at event m, then the perceived attribute value for AOI j is updated according to the following updating rule, where x^* is the true value:

$$x_{ijs}[m+1] = (1 - \alpha_s) x_{ijs}[m] + \alpha_s x^*_{ijs} \tag{5}$$

We allow each attribute s to have its own updating rate α_s since prior research using mouse cursor tracking to measure the rate of information processing of alternatives during choice has demonstrated that people may process attributes at different speeds (Sullivan et al. 2015).

3.5. Decision Threshold

For each shopper we can estimate a decision threshold that defines the stopping rule for information search. The stopping rule determines when shoppers exit the landing page by clicking on an AOI (See Equation 1). Our specification of the decision threshold is flexible enough to facilitate comparison across several types of behavioral patterns proposed in prior research.

$$\begin{aligned} R_i[m] &= \rho_{i0} \\ &+ \rho_{i1}\Bigg(\sum_{j \in \text{Seen}} P_{ijF}[m](V_{ij}[m+1] + D_{ij}[m]) + \sum_{k \notin \text{Seen}} P_{ikF}[m](V_{ik}[m] + D_{ik}[m])\Bigg) \\ &+ \rho_{i2}\log(m) \end{aligned} \tag{6}$$

The components of the reservation utility represent different psychological mechanisms. The first term represents a constant decision threshold. The second term captures next-step dynamics by allowing the decision threshold to vary based on shoppers' one-step forward-looking expectations. Note that for AOIs that shoppers have already "Seen" (i.e., at least one fixation; $j \in$ Seen), the anticipated values of the AOIs at the next event $m+1$ can be calculated by extrapolating the learning process described in Equation 5. However, we can't assume that shoppers are able to perform this extrapolation for AOIs that haven't been seen ($k \notin$ Seen), so the anticipated values of these AOIs are not adjusted. The third term captures whole-search dynamics by allowing the decision threshold to change exogenously over time, as measured by the number of fixation events that have occurred.

3.6. Likelihood Function

In summary, after each eye fixation, the consumer's choice is modeled as a two-stage process: first, she decides whether or not to click on the link within the currently fixated

AOI, which brings her to another webpage; then if search continues, she selects the next AOI upon which to fixate. The probability of clicking on an AOI depends on how its attractiveness value weighs against a decision threshold, which determines the search stopping rule. The probability of fixating on an AOI depends on both its attractiveness and the physical search costs of making the eye movement from the previously fixated AOI.

This two-stage process results in a series of eye fixations ending in a single click for each shopper. The events for all shoppers are captured by the following likelihood function, where $\theta = \{\beta_i, \gamma_i, \rho_i\}$:

$$\mathcal{L}(\theta) = \prod_{i=1}^{I} \left(\prod_{m=1}^{M_i - 1} P_{ijF}[m] \right) P_{iC}[M_i] \tag{7}$$

4. Experimental Design

We conducted an eye-tracking experiment in which female undergraduates at a large university were paid \$10 to browse the online store of American Apparel, plus a 1 in 20 chance of winning a randomly chosen item from their cart up to a maximum price of \$75 to incentivize shoppers to select items they would actually purchase. When shopping for clothing, consumers rely heavily on visual features, so eye-tracking was well-suited for capturing the accumulation of product information. Participants in our study were informed that their eye movements and mouse activity would be unobtrusively recorded. Participants were calibrated to the eye-trackers using the standard procedures recommended for SensoMotoric Instruments eye-tracking systems.

To test our model, we wanted to use an experimental manipulation in which shoppers would have different search goals in order to induce predictable variation in the value of different products and the value of continuing search. To achieve this, shoppers were randomly assigned either a utilitarian or hedonic shopping goal before they entered the store. The instructions were based on prior literature that suggests consumers may be motivated by utilitarian, practical considerations or hedonic, pleasure-seeking ones (Holbrook and Hirschman 1982; Dhar and Wertenbroch 2000):

Utilitarian goal instructions: "You have decided to go shopping for a new top. Think of this as a shopping trip for practical, every day, casual clothing, something you might wear to work or class. Your goal is to find a top."

Hedonic goal instructions: "You have decided to go shopping for a new top. Think of this as a shopping trip for frivolous, indulgent, fun clothing, something you might wear to a party. Your goal is to find a top, and also to enjoy the shopping experience."

Based on findings from neuroscience demonstrating that different neural pathways are responsible for goal-directed vs. stimulus-driven attention (Corbetta and Shulman 2002), we expected utilitarian shoppers to be efficient and focus attention on goal-relevant products (i.e., tops). In contrast, we expected hedonic shoppers to be more easily distracted by attractive products that were not goal-relevant (i.e., dresses), which would lengthen the search process. Note that we estimate the same model for shoppers in both conditions, but we estimate separate parameters on the dynamic components of the decision threshold (ρ_1 and ρ_2) for each goal condition to see whether the decision threshold varies differently shoppers with different goals.

Participants were brought to the landing page of American Apparel and had 5 minutes to shop. The virtual landing page (shown in Figure 1) was a modified version of the actual store website's "New Arrivals" page and remained unchanged across the data collection period. Row 1 was visible upon arrival to the website, while the remaining rows could be seen by scrolling down with the computer mouse.

Shoppers were free to visit any page in the store beyond the landing page and could select items by adding them to their virtual cart. Eye-tracking data and mouse clicks were collected throughout the entire shopping trip. In a brief survey about their general shopping habits, very few participants indicated that they had shopped at American Apparel in the past month, so prior knowledge of the product selection was low.

We restricted data collection to a 3-week period so that the product selection and prices would not change substantially. Since we collected extensive in-depth behavioral and physiological measurements for each shopping trip, data collection took on average 30 minutes for each shopper. This allowed us to collect eye-tracking data from 84 online participants during the 3-week time period. Out of the 84 participants, 8 were excluded from the analysis because of eye-tracker calibration issues, which left 76 participants total for analysis.

Approximately 2-3 months later, participants were contacted by email and invited to complete a follow-up survey for an additional $10 and a chance to win a $100 American Apparel gift card. They rated the 30 product pictures from the landing page on overall liking and other preference measures. Participants also rated how much they liked the products they had selected in the shopping task, along with a random sample of 30-35 products that other participants had selected, so we could assess how well they had been able to find products that they liked in the store. The 60+ products were presented in a random order with no identifying information about their source.This survey was conducted after a lengthy delay to reduce the effects of product exposure and to elicit measures from participants that presumably reflected their stable preferences. Measuring individual preferences for products either before or after the main experimental task is common within research employing psychophysical or neurophysiological methods in order to understand the split-second decision processes of each individual (Krajbich, Armel, and Rangel 2010; Reutskaja et al. 2011).

5. Descriptive Analysis

Table 1 contains the descriptive statistics for several features of the shopping trip, comparing shoppers in the utilitarian and hedonic goal conditions. The statistics from the landing page fixation pathways include the total number of fixations, the number of fixations on Product AOIs, the total number of AOIs seen, and the lowest AOI seen (calculated by numbering the Category AOI 0 and the Product AOIs from 1-30, left to right, top to bottom), and the percentage of shoppers who clicked on the Category AOI. On average, shoppers fixated on between a quarter to a third of the Product AOIs and got about a third to halfway down the page. Most shoppers ended up clicking on the Category AOI rather than on one of the Product AOIs. On average, shoppers ended up finding 2 products, valued at around $90 total.

[Insert Table 1 about here]

Consistent with our hypothesis that utilitarian shoppers would be more goal-driven (i.e., focus on the first two rows, which contained tops), while hedonic shoppers would be more distracted by attractive products farther down the page (i.e., sweater and dresses), we found that hedonic shoppers searched longer and more extensively on the landing page. Hedonic shoppers made more total fixations (t(51.4) = 1.94, p = 0.06), more fixations on Product AOIs (t(50.8) = 2.13, p = 0.04), and looked at more unique AOIs (t(59.2) = 2.21, p = 0.03). Hedonic shoppers also made more fixations farther down the page. This is illustrated by Figure 2 panel A, which plots the number of fixations for shoppers by whether they were given a utilitarian or hedonic goal. In both goal conditions, a large proportion of fixations were concentrated in the first row, presumably because shoppers had to use the computer mouse to scroll down to see the remaining five rows. Within the bottom three rows, hedonic shoppers made about 5 more fixations on average than utilitarian shoppers ($t(49.7) = 2.15, p = 0.04$).

[Insert Figure 2 about here]

These descriptive results will be used to assess how well the model captures important aspects of shopping behavior. Specifically, we use the parameter estimates of our model to simulate search pathways on the landing page for each shopper. We then compare the simulated and observed shopping patterns as a way of quantifying in-sample model fit.

6. Estimation Method and Results

The main goal of our estimation procedure was to compare the performance of our model using a decision threshold that has three components: a threshold intercept ρ_0, an endogenously dynamic component ρ_1 that represents next-step forward-looking, and an exogenously dynamic component ρ_2 that represents whole-search heuristics. More specifically, four versions of the model were estimated: a static version and three dynamic versions of the model (next-step only, whole-search only, and next-step and whole-search combined). For each version of the model, we obtained individual-level parameter estimates following a hierarchical Bayes procedure, except for the dynamic components ρ_1 and ρ_2, which were estimated at the goal condition level. We then used the parameter estimates to simulate each shopper's fixation pathway. We compared the relative performance of the models by calculating the deviance information criterion (DIC) for each set of parameter estimates, and also by looking at how closely simulated descriptive statistics matched observed descriptive statistics.

6.1. AOI Attributes and Search Costs

The AOI attributes we included in the estimation of attractiveness, $V_{ij}[m]$, were the liking ratings from the follow-up survey, the product prices, and binary variables for being below vs. above the \$75 budget and for Category vs. Product AOIs. The shopper's perceived value of each attribute s was initialized at a neutral point and updated by its own learning rate α_s (see Equation 5). From a grid search on a population-level version of the full model with log-likelihoods obtained through maximum likelihood estimation, we selected neutral starting values and updating rates for the AOI attributes that were used for all shoppers. We used a grid search because it is empirically difficult to disentangle low/high learning rates from weak/strong preferences.

The 7-point liking rating (1="Don't Like At All", 7 = "Like Very Much") was centered at 0 such that the liking for each product started at the neutral point -1 and gradually approached its "true" value (ranging from -3 to +3) at a rate of $\alpha_s = 0.1$ as shoppers made more fixations to gather information about the product. On average, among all the products that shoppers rated in the follow-up survey, only 8% were liked more than the products in their own carts. This suggests that shoppers were pretty good at finding products they liked and lends validity to the liking rating as a measure of stable product preferences.

For products that were within the $75 budget, prices were mean-centered, with $25 as the neutral starting value and updated at a rate of $\alpha_s = 0.9$. Although higher prices are typically viewed as undesirable, in our experiment shoppers were incentivized to select products that were close to $75, without going over, so we expect price to contribute positively to Product AOI attractiveness. We also included a variable that indicated whether products were out-of-budget, which began at 0 for all products and approached 1 for products that cost over $75 at a rate of $\alpha_s = 0.5$.

Since the Category AOI did not possess any inherent measurable attributes, we included a variable that started at 0 and approached 1 for fixations on the Category AOI at a rate of $\alpha_s = 0.5$. This Category variable is scaled by the coefficient estimate and is the only variable that contributes to the perceived attractiveness of the Category AOI; all other attributes are fixed at 0 for the Category AOI. The Category variable is fixed at 0 for the Product AOIs.

Effort and distance were binary variables that captured the cognitive and physical search costs, $D_{ij}[m]$ incurred when visually navigating the store environment. We assumed that AOIs that were *not* immediately visible upon entering the landing page required high cognitive effort to fixate on, while AOIs in the first row of Product AOIs and the Category AOI required low cognitive effort. Thus the "effort" variable separated immediately available AOIs from those requiring the shopper to scroll with the computer mouse.

Distance referred to whether or not the AOI was positioned adjacent to the previously fixated AOI, and indicated whether AOIs required "near" or "far" eye movements. Given these definitions, effort was absolute and remained the same across fixation decision, while distance was relative to the previous fixation and varied over time. We used contrast coding (-1 and +1) to refer to the two levels of the search cost variables.

6.2. Parameter Estimation

In order to capture the heterogeneity across individual shoppers' sequences of fixations on the landing page, we estimated our model using a standard hierarchical Bayes method (Gelman et al. 2014). For the first-stage priors, we assumed the set of parameters $\{\beta_i, \gamma_i, \rho_i\}$ followed a multivariate normal distribution with mean μ and precision Ω. For the second-stage priors, we assumed that μ followed a conjugate multivariate normal distribution with mean μ_0 (a vector of zeros) and precision Ω_0 (an identity matrix), and that Ω followed a conjugate Wishart distribution with degrees of freedom r (the number of parameters in $\{\beta_i, \gamma_i, \rho_i\}$ plus 3) and inverse scale matrix R (where R^{-1} was an identity matrix).

We made one modification to the general model to address the specific characteristics of our data. Since we observed multiple "continue search" (or fixation) decisions for each shopper but only a single click decision on the landing page, if we allowed both the constant (ρ_0) and dynamic (ρ_1 and ρ_2) components of the decision threshold to be fully heterogeneous, there would be no single maximizing solution for the parameters. So the estimation of ρ_1 and ρ_2 is not heterogeneous, but instead separately estimated at the population level for the shopping goal conditions (utilitarian vs. hedonic).

We simulated the parameters using a Markov chain Monte Carlo (MCMC) sampler in the programming language R. For each model, we ran three MCMC chains from different starting values for 10,000 iterations each. After checking for convergence, we used the first 5,000 iterations as burn-in and thinned the chains to reduce auto-correlation, leaving us with 1,500 posterior samples for each parameter. For the model estimation, we had to exclude 8 participants from the analysis for not completing the follow-up survey regarding product preferences, leaving 69 shoppers total, with the number of observations per shopper ranging from 1 to 93 with a median of 19.

Table 2 shows the parameter estimates and compares the model fit across all model variations, including the posterior means, as well as the 95% Credible Intervals and variances for the heterogeneous parameters. The decision threshold was either static with only the constant term (ρ_0) or included the next-step component (ρ_1) and/or the whole-search component (ρ_2). Across all four model variations, the sign of the parameter estimates generally match a prior predictions and intuitions. The AOI attributes of liking, price, and category have positive coefficients, while the out-of-budget attribute has a negative coefficient. Effort and distance, which represent search costs, have negative coefficients.

[Insert Table 2 about here]

The coefficient ρ_1 on the next-step forward-looking term is positive for shoppers with a utilitarian goal and negative for shoppers with a hedonic goal, but not significantly different from 0 based on the 95% CI.The coefficient ρ_2 on the whole-search time-varying term in the decision threshold is negative, which suggests that shoppers got more impatient to choose an AOI as search went on. So it appears that shoppers were not taking into account the expected value of continuing search with one more fixation on the landing page when making their click decisions, but rather taking into account the passage of time in a heuristic way.

To compare the fit across the model variations, we calculated the DIC using the equation $-4\text{E}_\theta \log(p(y|\theta)) + 2\log(p(y|\hat{\theta}))$, where y represents the observed data and $\hat{\theta}$ represents the estimated parameter values (Gelman et al. 2014). We also used the individual posterior estimates of the parameters to simulate the pathway of eye fixations ending in a click on the landing page (1500 simulated pathways per shopper). This allowed us to assess how well the models were able to capture the descriptive statistics described earlier in Section 5, including the total number of fixations, the number of unique AOIs seen, the lowest AOI seen, and whether shoppers clicked on a Product AOI or the Category AOI (see Table 2).

For example, using the simulations from the full model, we successfully replicated the fixation patterns going down the rows of products, including the drop-off in fixations between the first row and subsequent rows of Product AOIs, as well as the difference in number of fixations between utilitarian and hedonic shoppers (see Figure 2, compare panels A and B). The full model was also able to accurately simulate whether shoppers ended up clicking on a Category or Product AOI. For each shopper, we calculated the percentage of simulated shopping pathways that resulted in clicking a link in the Category AOI. In the full model, the percentage of shoppers whose simulated shopping pathways resulted in the "correct" AOI click (either the Category AOI or a Product AOI) in over half of the 1500 simulations was 81%.

6.3. Detailed Comparison of Decision Thresholds

Adding the next-step and whole-search dynamic components to the decision threshold on their own improved in-sample fit in terms of DIC (see Table 2). Specifically, adding the whole-search dynamics also greatly improved the root mean squared error (RMSE) of the simulated total fixations and Product AOI fixations, while there was a much smaller improvement after adding the next-steps component. This is further illustrated in Figure 3, which plots the simulated versus observed number of fixations for each shopper across the four model variations. Recall that the number of fixations is the measure of search length and therefore the key variable in comparing the different decision threshold specifications. In Figure 3, we see that the static model systematically overestimated the number of fixations, and adding the whole-search component corrected most of this upward bias. However, adding both next-step and whole-search components does not further improve the model fit.

[Insert Figure 3 about here]

Since the full model with both next-step and whole-search terms does not have the lowest DIC and the next-step coefficients (ρ_1) were close to 0 for online shoppers, we can conclude that the improvement in model fit in the full model with all dynamic components (compared to the static model) is mainly driven by the coefficients on the whole-search term. It appears that shoppers were relying mostly on the passage of time, rather than performing a mental calculation of the expected value of continuing search.

To illustrate how the decision threshold changes over time, Figure 4 plots the value of the decision threshold across all four model variations for each shopper at each fixation, with shoppers separated by shopping goal condition (utilitarian vs. hedonic). To compute the decision threshold, we used the mean of each shopper's individual posterior estimates for the heterogeneous parameters and the mean of the population-level posterior estimates for the non-heterogeneous parameters.

[Insert Figure 4 about here]

In the static model (see Figure 4 Panel A), the decision threshold is constant over time with just the heterogenous intercept term (ρ_i) for all shoppers, with little separation between utilitarian and hedonic shoppers. Adding next-step dynamics (ρ_1) to the decision threshold (see Panel B) results in small perturbations from fixation to fixation, since we assume that shoppers are computing the expected value or attractiveness of the selected AOI if search were to continue with one more fixation. Adding whole-search dynamics (ρ_2) results in the decision threshold starting much higher than in the static and next-step dynamic models (see Panel C), and then exhibiting a smooth logarithmic decrease across the whole search process on the landing page. This is consistent with prior research that suggests consumers engage in an initial exploration period before honing in or exploiting a particular option (Huang and Hutchinson 2013), shifting between search and choice modes (Dellaert and Häubl 2012), as well as declining reservation utilities due to impatience or rising search costs (Häubl, Dellaert, and Donkers 2010; Brown, Flinn, and Schotter 2011; Koulayev 2014). In the full model with both next-step and whole-search dynamics (see Panel D), we see both short run perturbations in the decision threshold and long-run decreases overall, as well as greater separation between utilitarian and hedonic shoppers. The decision threshold for utilitarian shoppers appears to decrease more rapidly than for hedonic shoppers, resulting in fewer fixations and early click decisions. Note that both next-step and whole-search dynamics do a better job at separating hedonic and utilitarian shoppers than does the static model, even though the latter has heterogeneous intercepts.

7. Counterfactual Analyses

We conducted two counterfactual analyses. First, since we did not assume that shoppers were searching optimally, we were interested in assessing “how well” they searched, in terms of clicking on high value AOIs. To do this, we varied the parameters of the decision thresholds to assess how well they were searching by testing whether they could have clicked on better AOI choices if they had used different decision thresholds.

Second, we were interested in how search patterns would change if the products on the landing page had been positioned in a different way. So we conducted counterfactual simulations of shopper search pathways with the Product AOIs positioned with “best” vs. “worst” products first, based on each shopper’s individual self-reported product preferences. We then verified the predictions of this counterfactual simulation using a controlled follow-up lab study.

7.1. How Well Were People Searching?

To assess how well shoppers were able to find attractive AOIs to click on, we performed a counterfactual analysis by simulating shoppers’ sequences of fixations using parameters that deviated from those estimated by the full model with both next-step and whole-search dynamics. We chose to use the full model (rather than the winning whole-search model) because we were interested in looking at the impact of varying all three components of the decision threshold and the difference in fit was marginal.

We independently varied the values of the constant term ρ_0, the coefficient ρ_1 on the next-step dynamic term, and the coefficient ρ_2 on the whole-search dynamic term. Recall that ρ_0 was fully heterogeneous across individuals, while ρ_1 and ρ_2 were fixed effects that only differed across the utilitarian vs. hedonic shopping goals.

Figure 5 plots the results of our simulation. For each set of alternative parameter values, we simulated shopper fixation pathways ending with a click on an AOI, and compared the attractiveness $V_{ij}[m]$ of the chosen AOI j at the time it was clicked on (i.e., “perceived” attractiveness, filled circles) to its attractiveness if the shopper had learned complete information about the AOI (i.e., “true” attractiveness, open circles). Recall that attractiveness includes the liking rating, price, whether the product was out-of-budget, and the scaled Category AOI variable. Figure 4 plots how attractiveness changes with the alternative decision thresholds (panels A-C) and number of additional fixations (panels D-F). We varied each of the coefficients independently, with the dotted line indicating the original estimated decision thresholds.

[Insert Figure 5 about here]

In panels A-C, we see that at 0 deviance, the attractiveness of selected AOIs are near their max and are on the flat regions of the curves for ρ_1 and ρ_2, so the returns on increasing the weight on the next-step or whole-search dynamic terms are not huge. Very negative deviations resulted in too little search, chosen AOIs that were low in attractiveness, and perceived values that are higher than true values. This is consistent with the interpretation that, under the assumptions of our model, not enough information had been acquired about the AOIs. Positive deviations increase the length of search, which shrinks the difference between the perceived and true values.[4]

In panels D-F, we see that the true attractiveness of chosen AOIs greatly increases within the first 30 fixations, but there are diminishing returns for continuing search with the next 30 fixations. Shoppers were clearly not "over-searching" but even searching twice as long would have only raised the attractiveness of the chosen AOI from 2 to 2.5.

7.2. How Would Search Be Affected by Product Layouts?

Since we had individual product ratings from each shopper, we were able to examine the implications of our decision threshold model regarding how shopper search patterns would change with different physical product positionings. We conducted a second counterfactual analysis in which we simulated participants' patterns of fixations if the 30 products on the page were positioned "worst-first" or "best-first" according to individual shoppers'

[4] Positive deviations over +1 resulted in decision thresholds increasing too quickly over time and simulations of "never-ending" search, so we artificially cut off search after 500 fixations, resulting in a discontinuity at +1.

self-reported liking ratings of the landing page products.

Table 3 Panel A compares the results from simulations using the original product positioning to those under best-first and worst-first product positionings. Under the worst-first layouts, shoppers (in simulation) make more fixations, view more unique AOIs farther down the page, and are more likely to choose a Category AOI link. Under the best-first positioning, shoppers make fewer fixations and are more likely to click on one of the Product AOI links rather than the Category AOI links.

[Insert Table 3 about here]

To empirically verify the predictions made by our counterfactual simulations, we conducted an follow-up eye-tracking study in which we experimentally manipulated how products were sorted on the landing page. We required that all participants complete a pre-screening survey at least 5 days prior to the eye-tracking portion of the study. In this survey, participants were presented with the images of 60 products that were available for online purchase from the clothing retailer Forever 21. Participants first rated their liking of the product on a 7-point Likert scale. Participants then ranked the products within each point rating to break ties, which gave us a full ranking of all 60 products from each individual.

In the eye-tracking experiment, 132 female undergraduate participants were brought to a modified version of the landing page of the Forever 21 and had 5 minutes to shop, again with the possibility that they might win one of the items in their shopping cart up to $50. The landing page displayed a set of links to specific product categories, as well as 40 of the 60 products that participants had ranked in the pre-screening survey (10 rows of 4 products each).

Participants were randomly assigned to one of three conditions: random product positioning, best-first positioning, and worst-first positioning. In the random positioning conditions, all participants saw the same random configuration of 40 products on the landing page. In the best-first and worst-first conditions, we used each participants' responses from the pre-screening conditions to create a customized landing page with the products ordered according to each participants' own rankings.

Table 3 Panel B shows the results of the experiment. Compared to the random product positioning, shoppers in the best-first condition made fewer fixations, while shoppers in the worst-first positioning made more fixations farther down the page, which is consistent with the predictions of the counterfactual simulation. Note that we did not see the same pattern of Category vs. Product clicks among the Forever21 shoppers. In the Forever21 sample, shoppers were much more likely to click on one of the Product AOIs.

8. Conclusion and Future Directions

We contribute to the growing literature that treats the eye fixations of shoppers as endogenous decisions by building and empirically testing a search model for how online shoppers visually acquire product information on a store's landing page leading up to the decision to click on a link that brings them to a new part of the website. Since we directly estimated the components of the decision threshold that determines when shoppers ended search on the landing page, we were able to say something about the psychological mechanisms involved when shoppers engage in search. Specifically, we found that shoppers' decision thresholds did not significantly vary in the short-run with the expected attractiveness of selected AOIs if search were to continue for one more fixation, but did decrease in the long-run, which can be interpreted as a whole-search forward-looking heuristic. The results of the model estimation suggest that adding the whole-search dynamics to the decision threshold improved the model fit compared to a static decision threshold.

From the fixation pathways generated by simulations using the estimated parameters, we demonstrated that our model can track a number of search characteristics, including the total number of fixations (see Figure 4) and other features like the number of unique AOIs seen and the lowest AOI seen (see Table 2). Our model was also able to capture the differences in the number of fixations on specific rows of products between shoppers with different goals (see Figure 3). We confirmed through counterfactual analyses that shoppers were reasonably good at choosing attractive AOIs under the assumptions of our model.

For external validity, we conducted an analogous in-store eye-tracking experiment at a brick-and-mortar American Apparel location with 23 participants using mobile eye-tracking glasses that recorded video of each shopper's field of view and the position of their eyes at each moment. The analogous "landing page" of the brick-and-mortar store was defined by tracing a perimeter around the product displays that were visible from the end of the entry walkway of the store and included 12 Product AOIs that could be physically touched while standing in this landing page zone.

The in-store Category AOI encompassed all products displays that were positioned beyond the landing page zone but could be seen by looking past the Product AOIs. Search within the landing page ended when in-store shoppers physically touched a product display (analogous to online shoppers clicking on a link within a Product AOI) or walked beyond the landing page zone and further into the store (analogous to online shoppers clicking on a link within the Category AOI).

Using the same model and estimation procedure as for the online shoppers, we obtained similar parameter estimates and simulated fixation pathways for the in-store shoppers. Just as with the online shoppers, we found that the whole-search dynamics (but not the next-step dynamics) improved the model fit over the static decision threshold model. Figure 6 plots the RMSE of total fixations across the 4 model variations, comparing online and in-store samples. We see that for both online and in-store shoppers, the whole-search dynamics contributed to the best model fit (i.e., the lowest RMSE), and for in-store shoppers the next-step dynamics actually decreased the model fit relative to the static model.

[Insert Figure 6 about here]

One challenge in future research will be in formulating models for the complex decision processes within an entire shopping trip, or possibly across multiple shopping trips within a customer's "lifetime". On the landing page, the decision processes of shoppers can be reduced to a sequence of "stay" versus "leave" choices. But as shoppers visit different virtual or physical areas of a store, the set of possible actions grows very large and the stopping rules become more complicated. For example, for any given area of the store, we must take into account the qualitative differences between ending search to visit a new location versus to revisit a prior location, which adds further complexity to incorporating dynamics like time considerations or forward-looking computations (both next-step and whole-search) into models of shoppers' split-second decisions.

Trying to model the behavior of shoppers as they diverge from a common landing page and continue on their own search paths presents an opportunity to test different theories about how shoppers process and retain product information, as well as how they manage the physical and cognitive effort involved in search. For example, shoppers may begin to employ more forward-looking strategies in later stages of the shopping trip as they familiarize themselves with the shopping environment, as suggested by Gopalakrishnan, Iyengar, and Meyer (2014). However, note that the type of micro-level modeling that we employed in this paper quickly becomes intractable for "real" stores because of the number of products and store locations. Thus, the current data could be modeled at a more macro-level or we can conduct future experiments within more limited shopping environments.

9. References

Adam, K. (2001), "Learning While Searching for the Best Alternative," Journal of Economic Theory, 101 (1), 252-281.

Bettman, James R., and Pradeep Kakkar (1977), "Effects of Information Presentation Format on Consumer Information Acquisition Strategies," Journal of Consumer Research, 3 (4), 233-240.

Branco, Fernando, Monic Sun, and J. Miguel Villas-Boas (2012), "Optimal Search for Product Information," Management Science, 58 (11), 2037-2056.

Brown, Meta, Christopher J. Flinn, and Andrew Schotter (2011), "Real-Time Search in the Laboratory and the Market," The American Economic Review, 101 (2), 948-974.

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

Busemeyer, Jerome R., and James T. Townsend (1993), "Decision Field Theory: A Dynamic-Cognitive Approach to Decision Making In An Uncertain Environment," Psychological Review, 100 (3), 432-459.

Chandon, Pierre, J. Wesley Hutchinson, Eric T. Bradlow, and Scott H. Young (2007). *Visual Marketing: From Attention to Action.* Lawrence Erlbaum Assoc., Mahwah, NJ.

Chandon, Pierre, J. Wesley Hutchinson, Eric T. Bradlow, and Scott H. Young (2009), "Does In-Store Marketing Work? Effects of the Number and Position of Shelf Facings on Brand Attention and Evaluation at the Point of Purchase," Journal of Marketing, 73 (6), 1-17.

Corbetta, Maurizio, and Gordon L. Shulman, "Control of Goal-Directed and Stimulus-Driven Attention in the Brain," Nature Reviews Neuroscience, 3 (3), 201-215.

Corbetta, Maurizio, Gaurav Patel, and Gordon L. Shulman, "The Reorienting System of the Human Brain: From Environment to Theory of Mind," Neuron 58 (3), 306-324.

Dellaert, Benedict G. C., and Gerald Häubl (2012), "Searching in Choice Mode: Consumer Decision Processes in Product Search with Recommendations," Journal of Marketing Research, 49 (2), 277-288.

Dhar, Ravi, and Klaus Wertenbroch (2000), "Consumer Choice Between Hedonic and Utilitarian Goods." Journal of Marketing Research, 37 (1), 60-71.

Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg (2006), "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model," American Economic Review, 96 (4), 1043-1068.

Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin (2014), Bayesian Data Analysis, 3rd ed. Boca Raton, FL: CRC Press.

Gigerenzer, Gerd, and Daniel G. Goldstein (1996), "Reasoning the Fast and Frugal Way: Models of Bounded Rationality," Psychological Review, 103 (4), 650-669.

Gilbride, Timothy J. and Greg M. Allenby (2004), "A Choice Model with Conjunctive, Disjunctive, and Compensatory Screening Rules," Marketing Science, 23 (3), 391-406.

Gittins, John C. (1979), "Bandit Processes and Dynamic Allocation Indices," Journal of the Royal Statistical Society. Series B (Methodological), 41 (2), 148-177.

Gittins, John C., and David M. Jones, "A Dynamic Allocation Index for the Discounted Multiarmed Bandit Problem," Biometrika, 66 (3), 561-565.

Gluth, Sebastian, Jörg Rieskamp, and Christian Büchel, "Deciding Not to Decide: Computational and Neural Evidence for Hidden Behavior in Sequential Choice," PLoS Computational Biology, 13 (4), e1005476.

Gopalakrishnan, Arun, Raghuram Iyengar, and Robert J. Meyer (2014), "Consumer Dynamic Usage Allocation and Learning Under Multipart Tariffs," Marketing Science, 34 (1), 116-133.

Häubl, Gerland, Benedict G.C. Dellaert, and Bas Donkers (2010), "Tunnel Vision: Local Behavioral Influences on Consumer Decisions in Product Search," Marketing Science 29 (3), 438-455.

Hey, John D (1981), "Are Optimal Search Rules Reasonable? And Vice Versa? (And Does It Matter Anyway?)," Journal of Economic Behavior & Organization 2 (1), 47-70.

Hey, John D (1982). "Search For Rules For Search," Journal of Economic Behavior & Organization, 3 (1), 65-81.

Holbrook, Morris B., and Elizabeth C. Hirschman (1982), "The Experiential Aspects of Consumption: Consumer Fantasies, Feelings, and Fun," Journal of Consumer Research, 9 (2), 132-140.

Honka, Elisabeth, and Pradeep Chintagunta (2016), "Simultaneous or Sequential? Search Strategies in the US Auto Insurance Industry," Marketing Science.

Huang, Yanliu, and J. Wesley Hutchinson (2008), "Counting Every Thought: Implicit Measures of Cognitive Responses to Advertising," Journal of Consumer Research, 35 (1), 98-118.

Huang, Yanliu, and J. Wesley Hutchinson (2013), "The Roles of Planning, Learning, and Mental Models in Repeated Dynamic Decision Making," Organizational Behavior and Human Decision Processes, 122 (2), 163-176.

Hui, Sam, Eric T. Bradlow, and Peter S. Fader (2009), "Testing Behavioral Hypotheses Using an Integrated Model of Grocery Store Shopping Path and Purchase Behavior," Journal of Consumer Research, 36 (3), 478-493.

Hui, Sam, J. Jeffrey Inman, Yanliu Huang, and Jacob Suher (2013), "Estimating the Effect of Travel Distance on Unplanned Spending: Applications to Mobile Promotion Strategies," Journal of Marketing, 77 (2), 1-16.

Hutchinson, J. Wesley, Joy Lu, and Evan Weingarten (2016), "Visual Attention in Consumer Settings," In International Handbook of Consumer Psychology, edited by Cathrine Janssen-Boyd, Magadlena Zawisza.

Hutchinson, J. Wesley, and Robert J. Meyer (1994), "Dynamic Decision Making: Optimal Policies and Actual Behavior in Sequential Choice Problems," Marketing Letters, 5 (4), 369-382.

Jovancevic-Misic, Jelena, and Mary Hayhoe (2009), "Adaptive Gaze Control in Natural Environments," The Journal of Neuroscience, 29 (19): 6234-6238.

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

Koulayev, Sergei (2014), "Search for Differentiated Products: Identification and Estimation," RAND Journal of Economics, 45 (3): 553-575.

Krajbich, Ian, Carrie Armel, and Antonio Rangel (2010), "Visual Fixations and the Computation and Comparison of Value in Simple Choice," Nature Neuroscience, 13 (10), 1292-1298.

Larson, Jeffrey S., Eric T. Bradlow, and Peter S. Fader (2005), "An Exploratory Look at Supermarket Shopping Paths," International Journal of Research in Marketing, 22 (4), 395-414.

Liechty, John, Rik Pieters, and Michel Wedel (2003), "Global and Local Covert Visual Attention: Evidence from a Bayesian Hidden Markov Model," Psychometrika, 68 (4), 519-541.

Lin, Song, Juanjuan Zhang, and John R. Hauser (2015), "Learning from Experience, Simply," Marketing Science, 34 (1), 1-19.

Lippman, Steven A. and Kevin F. McCardle (1991), "Uncertain Search: A Model of Search among Technologies of Uncertain Values," Management Science, 37 (11), 1474-1490.

March, James G., (1991), "Exploration and Exploitation in Organizational Learning," Organization Science, 2 (1), 71-87.

Otter, Thomas, Joe Johnson, Jrg Rieskamp, Greg M. Allenby, Jeff D. Brazell, Adele Diederich, J. Wesley Hutchinson, Steven MacEachern, Shiling Ruan, and Jim Townsend (2008), "Sequential Sampling Models of Choice: Some Recent Advances," Marketing Letters, 19 (3-4), 255-267.

Payne, John W., James R. Bettman, and Eric J. Johnson (1993), The Adaptive Decision Maker. Cambridge, UK, Cambridge University Press.

Ratcliff, Roger, and Philip L. Smith, "A Comparison of Sequential Sampling Models for Two-Choice Reaching Time," Psychological Review, 111 (2), 333-367.

Ratcliff, Roger, and Francis Tuerlinckx (2002), "Estimating Parameters of the Diffusion Model: Approaches to Dealing with Contaminant Reaction Times and Parameter Variability," Psychonomic Bulletin & Review, 9 (3), 438-481.

Rayner, Keith, "Eye Movements in Reading and Information Processing: 20 Years of Research," Psychological Bulletin, 124 (3), 372-422.

Reutskaja, Elena, Rosemarie Nagel, Colin F. Camerer, and Antonio Rangel (2011), "Search Dynamics in Consumer Choice Under Time Pressure: An Eye-tracking Study," The American Economic Review, 101 (2), 900-926.

Satomura, Takuya, Michel Wedel, and Rik Pieters (2014), "Copy Alert: A Method and Metric to Detect Visual Copycat Brands," Journal of Marketing Research, 51 (1), 1-13.

Shi, Savannah Wei, Michel Wedel and Rik Pieters (2013), "Information Acquisition During Online Decision Making: A Model-Based Exploration Using Eye-Tracking Data," Management Science, 59 (5), 1009-1026.

Shimojo, Shinsuke, Claudiu Simion, Eiko Shimojo, and Christian Scheier (2003), "Gaze Bias Both Reflects and Influences Preference," Nature Neuroscience, 6 (12), 1317-1322.

Shugan, Steven M. (1980), "The Cost of Thinking," Journal of Consumer Research, 7 (2), 99-111. Simon, Herbert A. (1955), "A Behavioral Model of Rational Choice," Quarterly Journal of Economics, 69 (1), 99-118.

Simon, Herbert A. (1955), "A Behavioral Model of Rational Choice," The Quarterly Journal of Economics, 69 (1), 99-118.

Stttgen, Peter, Peter Boatwright, and Robert T. Monroe (2012), "A Satisficing Choice Model," Marketing Science, 31 (6), 878-899.

Sullivan, Nicolette, Cendri Hutcherson, Alison Harris, and Antonio Rangel (2015), "Dietary Self-Control is Related to the Speed with which Attributes of Healthfulness and Tastiness are processed," Psychological Science, 26 (2), 122-134.

Time, Carol (2012), "7 Layout Secrets of the Big Retail Chains," Entrepreneur, available at https://www.entrepreneur.com/article/223808.

Towal, R. Blythe, Milica Mormann, and Christof Koch (2013), "Simultaneous Modeling of Visual Saliency and Value Computation Improves Predictions of Economic Choice," Proceedings of the National Academy of Sciences, 110 (40), E3858-E3867.

Townsend, James T., and F. Gregory Ashby (1983). *The Stochastic Modeling of Elementary Psychological Processes.* Elementary Cambridge, England: Cambridge University Press.

Van Zandt, Trisha, Hans Colonius, and Robert W. Proctor (2000), "A Comparison of Two Response Time Models Applied to Perceptual Matching," Psychonomic Bulletin & Review, 7 (2), 208-256.

Wedel, Michel, Rik Pieters, and John Liechty (2008), "Attention Switching During Scene Perception: How Goals Influence the Time Course of Eye Movements Across Advertisements," Journal of Experimental Psychology: Applied, 14 (2), 129-138.

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

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

Zingale, Carolina M., and Eileen Kowler (1987), "Planning Sequences of Saccades," Vision Research, 27 (8), 1327-1341.

10. Tables

Table 1: Descriptive statistics of shopping trip

Shopping Stage	Activity	Utilitarian (N=41)	Hedonic (N=35)
Landing Page	Total fixations	24.6 (SD=21.3)	38.6 (SD=38.2)
	Product AOI fixations	18.0 (SD=20.8)	33.4 (SD=38.01)
	Total AOIs seen	8.1 (SD=7.6)	12.9 (SD=10.9)
	Lowest AOI Seen	10.5 (SD=10.3)	15.0 (SD=13.0)
Exit Landing Page	% clicked on the Category AOI	82.9%	71.4%
Final Shopping Cart	# of products	2.2 (SD=1.2)	2.0 (SD=1.2)
	total value of products	$89.7 (SD=$52.6)	$97.3 (SD=$71.3)

Table 2: Parameter Estimates of Reservation Utility Model

		Static Model			**Next-Step Model**			**Whole-Search Model**			**Full Model**		
	Parameter	μ	**95% CI**	σ^2	μ	**95% CI**	σ^2	μ	**95% CI**	σ^2	μ	**95% CI**	σ^2
AOI Attractiveness	β_1 Liking	0.71	[0.42,1.01]	0.58	0.72	[0.44,1.00]	0.58	0.70	[0.43,0.99]	0.58	0.70	[0.42,1.00]	0.58
	β_2 Price	0.13	[-0.02,0.28]	0.27	0.13	[-0.02,0.29]	0.27	0.13	[-0.02,0.28]	0.27	0.13	[-0.02,0.29]	0.27
	β_3 Out-of-Budget	-0.89	[1.42,-0.37]	1.24	-0.90	[-1.42,-1.05]	1.26	-0.93	[-1.53,0.42]	1.32	-0.87	[-1.44,-0.29]	1.23
	β_4 Category	2.47	[1.74,3.17]	3.86	2.53	[1.79,3.54]	4.26	2.60	[1.89,8.25]	4.55	2.75	[1.98,3.70]	5.42
Physical Movements	γ_1 Effort	-1.24	[-1.43,-1.07]	0.38	-1.24	[-1.52,-0.84]	0.38	-1.19	[-1.43,-0.42]	0.38	-1.23	[-1.41,-1.05]	0.38
	γ_2 Distance	-1.16	[-1.56,-0.85]	1.07	-1.15	[-1.53,-0.84]	1.03	-1.19	[-1.56,-0.88]	1.10	-1.16	[-1.53,-0.84]	1.08
Decision Threshold	ρ_0 Constant	4.41	[3.94,4.93]	0.82	4.32	[3.68,4.96]	0.72	5.70	[4.52,7.13]	0.86	5.61	[4.31,7.18]	0.82
	ρ_1 Time-Varying (Util)	–	–	–	–	–	–	-0.46	[-0.94,-0.05]	–	-0.55	[-1.01,-0.15]	–
	ρ_1 Time-Varying (Hed)	–	–	–	–	–	–	-0.40	[-0.85,-0.004]	–	-0.25	[-0.70, 0.14]	–
	ρ_2 Forward-Looking (Util)	–	–	–	0.09	[-0.15,0.32]	–	–	–	–	0.15	[-0.09,0.38]	–
	ρ_2 Forward-Looking (Hed)	–	–	–	-0.02	[-0.41,0.47]	–	–	–	–	-0.23	[-0.76,0.32]	–
Model Fit Statistics	DIC	8863.331			8805.15			8784.607			8869.29		
	Total fixations (RMSE)	16.72			16.67			12.39			13.75		
	Product AOI fixations (RMSE)	14.95			14.37			12.00			13.13		
	Total AOIs seen (RMSE)	4.29			4.41			4.17			4.21		
	Lowest AOI seen (RMSE)	6.51			6.66			6.50			6.55		
	Click (% Correct)	79.71%			79.71%			79.71%			81.16 %		

Table 3: Search pathways statistics with different product positioning

A. Simulated Results

Positioning	**Total fixations***	**Product AOI fixations**	**Total AOIs seen*****	**Lowest AOI seen****	**% Clicked on Category AOI**
Worst-First	36.80 (SD=30.67)	26.95 (SD=26.93)	9.10 (SD=6.09)	15.78 (SD=8.51)	73.91%
Original	33.33 (SD=13.74)	25.39 (SD=23.04)	8.65 (SD=5.82)	15.38 (SD=6.55)	69.57%
Best-First	31.10 (SD=20.83)	24.27 (SD=20.85)	8.23 (SD=4.49)	14.92 (SD=8.17)	66.67%

B. Experimental Results

Positioning	**Total fixations***	**Product AOI fixations***	**Total AOIs seen****	**Lowest AOI seen***	**% Clicked on Category AOI**
Worst-First	109.4 (SD=59.8)	104.2 (SD=60.2)	27.7 (SD=11.6)	29.7 (SD=11.6)	6.7%
Original	94.5 (SD=82.8)	88.6 (SD=80.5)	21.7 (SD=13.6)	25.2 (SD=15.4)	7.3%
Best-First	78.6 (SD=57.0)	74.7 (SD=56.1)	19.5 (SD=13.9)	22.8 (SD=15.6)	10.1%

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

11. Figures

Figure 1: Experimental Design and Layout of Online Shopping Environment.

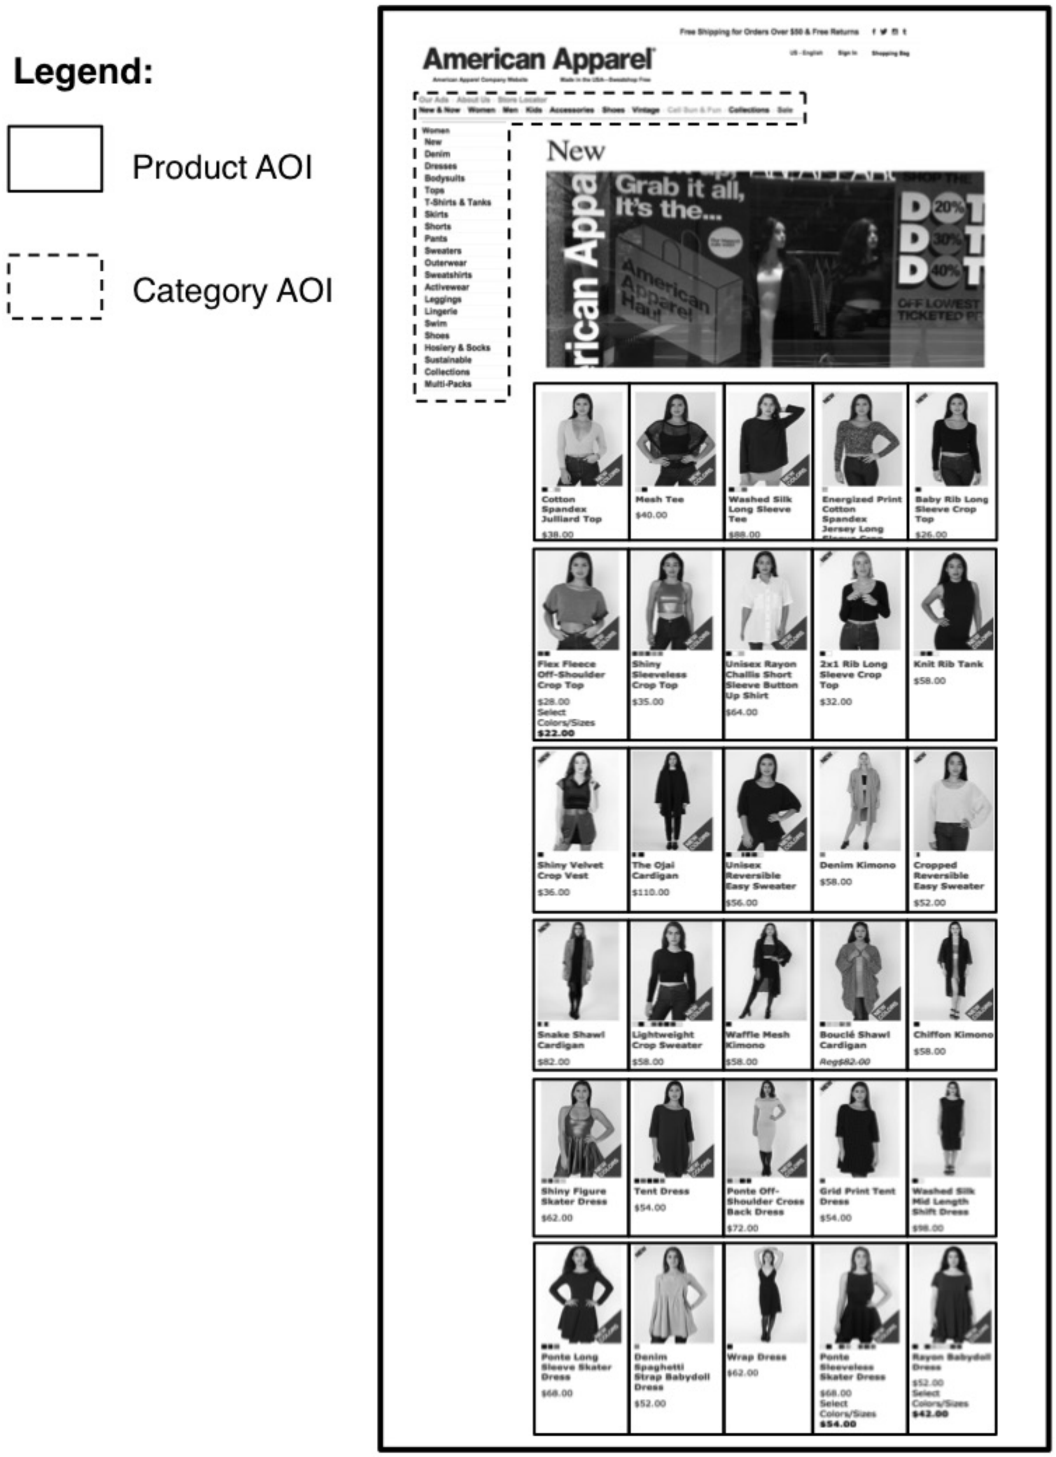

Figure 2: Average number of fixations across rows for observed (A) and simulated data (B). White bars indicate utilitarian shoppers, grey bars indicate hedonic shoppers.

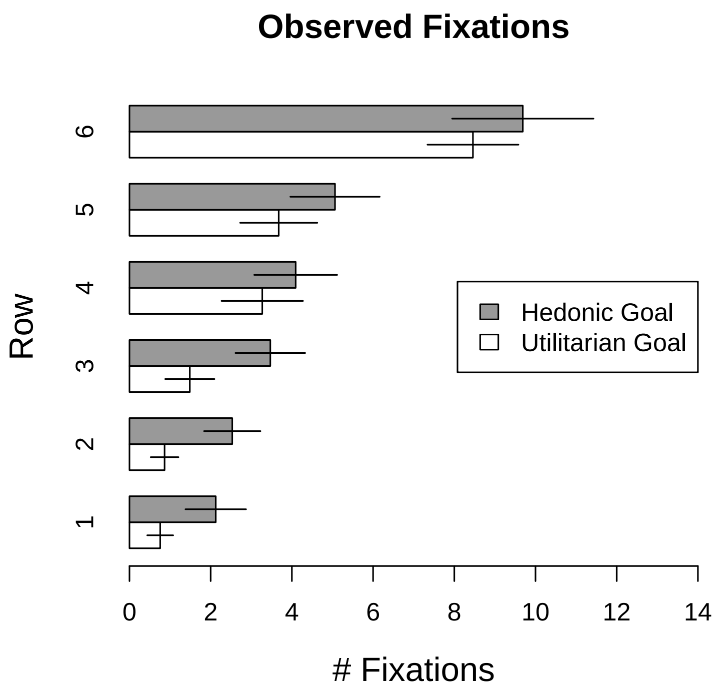

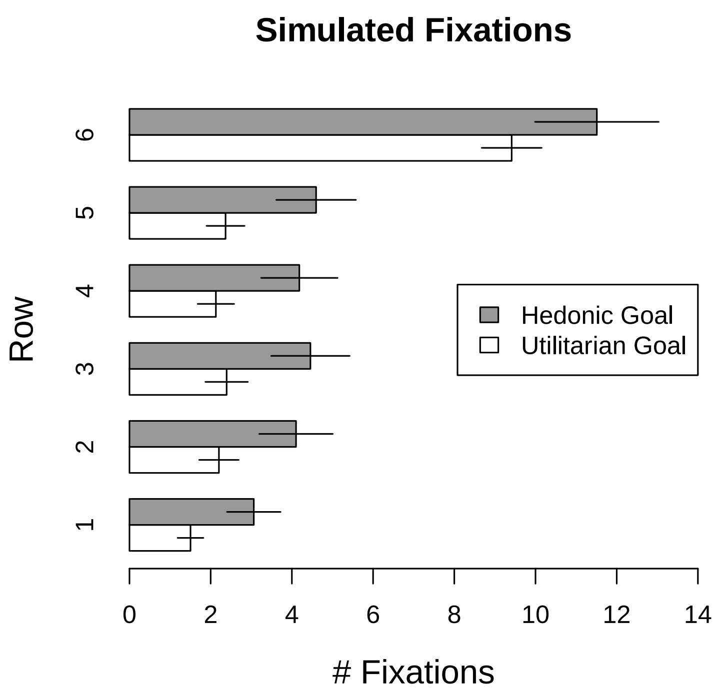

Figure 3: Observed versus simulated number of fixations across model variations

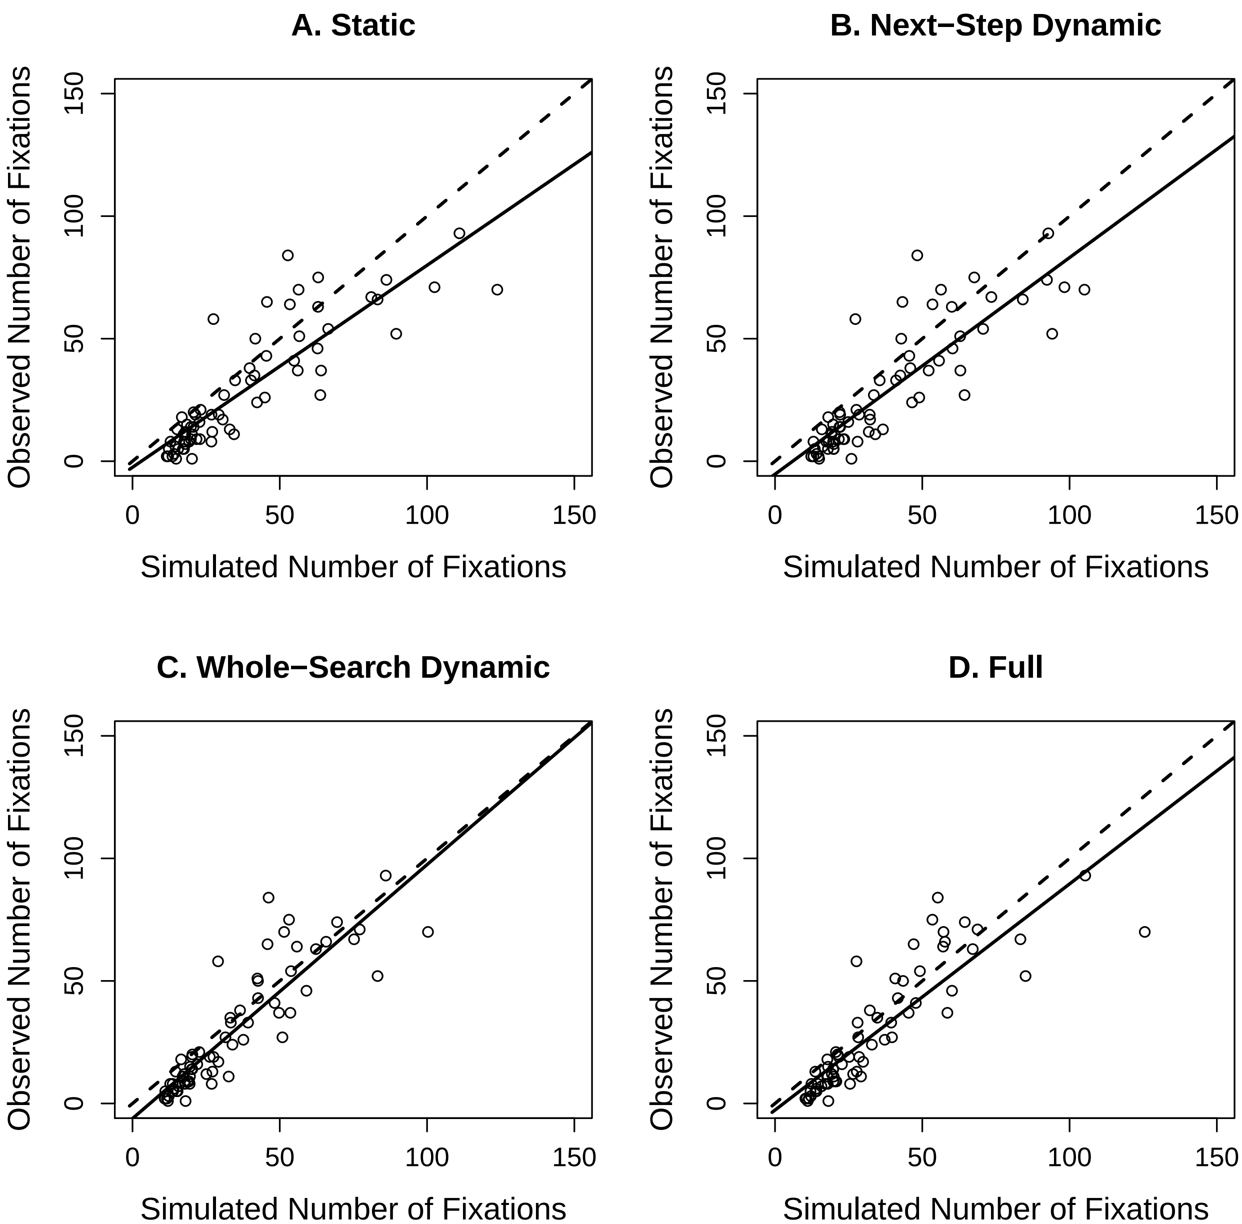

Figure 4: Decision thresholds

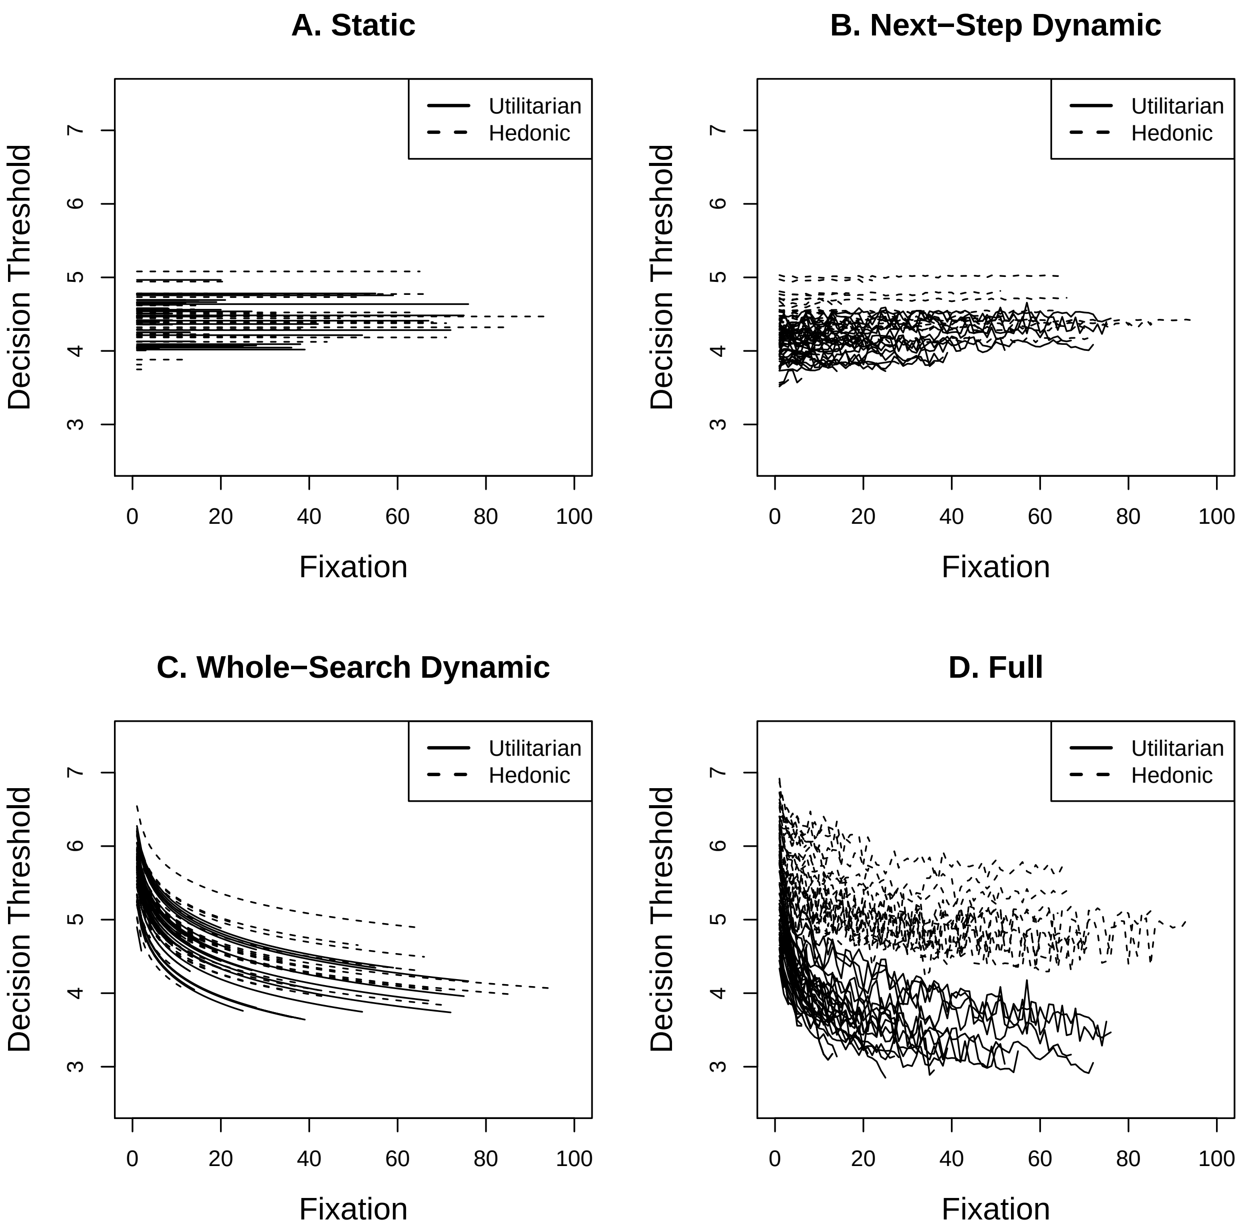

Figure 5: Comparison of Attractiveness of Selected AOIs under Alternative Decision Thresholds.

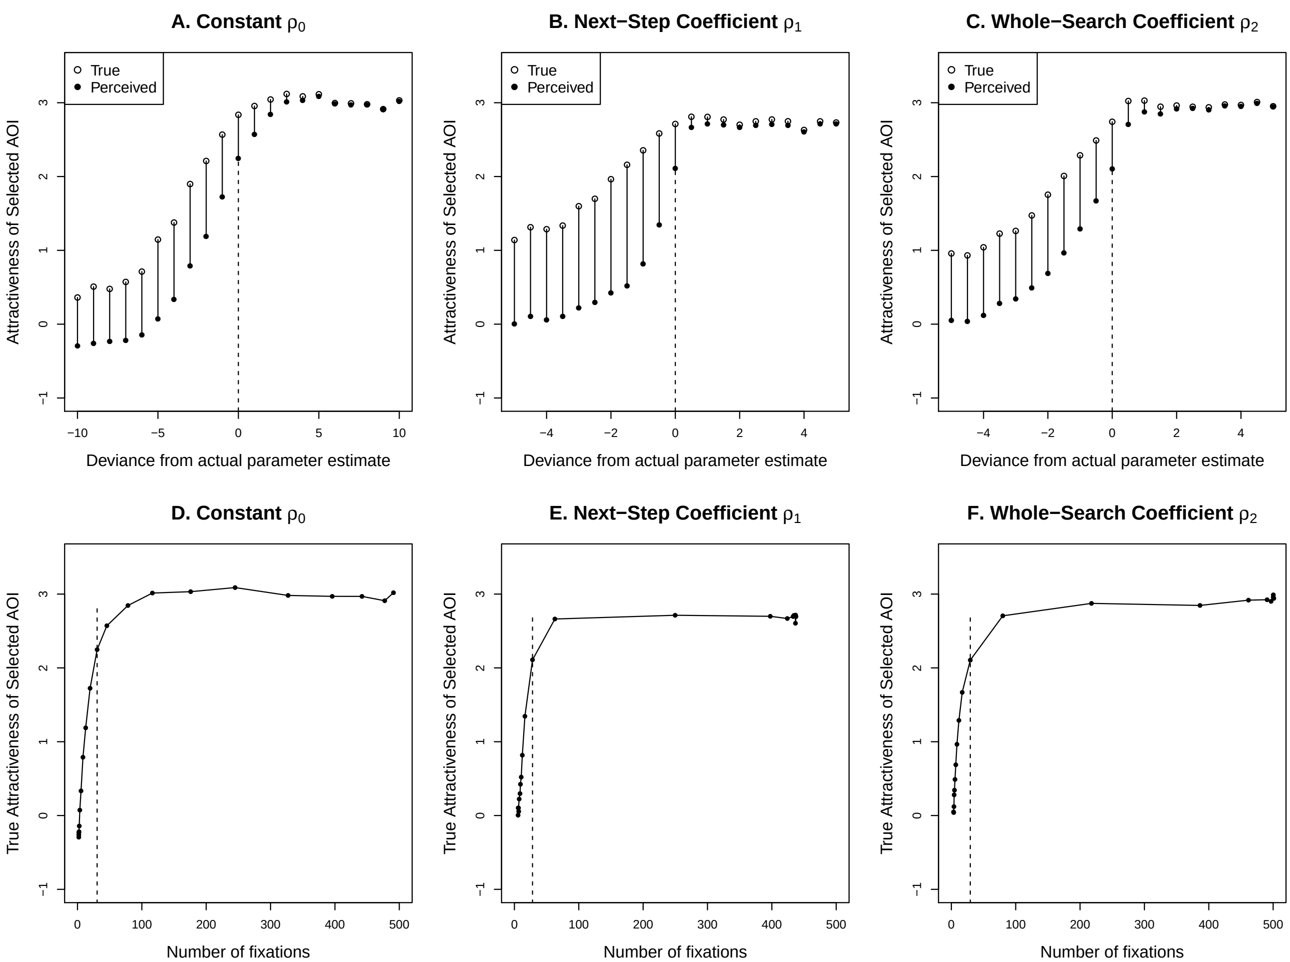

Figure 6: Comparison of RMSE of Total Fixations Across Online and In-Store Samples.