

Designing a Better Shopbot

Alan L. Montgomery

Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, alan.montgomery@cmu.edu

Kartik Hosanagar

The Wharton School, University of Pennsylvania, Pennsylvania 19104, kartik@wharton.upenn.edu

Ramayya Krishnan, Karen B. Clay

The Heinz School, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213 {rk2x@andrew.cmu.edu, kclay@andrew.cmu.edu}

A primary tool that consumers have for comparative shopping is the shopbot, which is short for shopping robot. These shopbots automatically search a large number of vendors for price and availability. Typically a shopbot searches a predefined set of vendors and reports all results, which can result in time-consuming searches that provide redundant or dominated alternatives. Our research demonstrates analytically how shopbot designs can be improved by developing a utility model of consumer purchasing behavior. This utility model considers the intrinsic value of the product and its attributes, the disutility from waiting, and the cognitive costs associated with evaluating the offers retrieved. We focus on the operational decisions made by the shopbot: which stores to search, how long to wait, and which offers to present to the user. To illustrate our model we calibrate the model to price and response time data collected at online bookstores over a six-month period. Using prior expectations about price and response time, we show how shopbots can substantially increase consumer utility by searching more intelligently and then selectively presenting offers.

Key words: intelligent agents; utility theory; information retrieval; stochastic modeling

History: Accepted by Jagmohan S. Raju; received August 1, 2001. This paper was with the authors 12 months for 3 revisions.

1. Introduction

Comparison shopping engines, better known as shopbots, are a key consumer innovation offered through the Internet. Shopbots automatically and efficiently search a large number of online stores, providing a consumer with nearly all offered prices for a particular product at a given point in time. For example, a consumer who visited the popular shopbot Dealtime.com in October 2001 wanting to purchase Tom Clancy's novel *Bear and the Dragon* would have waited 30 seconds before being presented with a list of 57 offers that ranged in price from \$16.45 to \$40.22. To complete a purchase, the consumer must follow one of the hyperlinks from the Dealtime.com site to the online store. Compared with the traditional sequential search process described by Stigler (1961), shopbots clearly lower the cost of acquiring price information.

Consumers also incur costs from using a shopbot. These costs include the time to use the shopbot, the time to wait for the shopbot to respond, and the additional cognitive effort expended in this more expansive search process. If the shopbot user ends up purchasing from his or her favorite store, then using a shopbot is a less-efficient shopping strategy. So when deciding whether to use a shopbot, consumers must evaluate whether the expected gains from search outweigh the added costs. In other words, is a consumer

better off using a shopbot or simply visiting his or her favorite store?

To address this problem, we develop an analytical model of consumer utility to compare the expected gains of search against the expected costs. Gains include higher utility from identifying lower prices or faster delivery, and costs include the waiting time for the shopbot to respond and the cognitive effort required to compare alternatives. In our model the shopbot can make the following operational decisions that influence these benefits and costs: which stores to query for offers, how long to wait for these stores to respond, and which items to report to a user. This cost-benefit framework is used to evaluate the utility of three alternatives: the traditional shopping process such as visiting a favorite store; the current shopbot design, which visits all stores and presents all results; and an optimized shopbot design.

Using this model and data on prices, wait times, and customer preferences over attributes, we show that directly visiting a favorite retailer may yield higher utility than using a shopbot that searches all stores and presents all offers. Furthermore, we show how a shopbot's design can be optimized to lead to a higher chance of using a shopbot over simply visiting a favorite retailer. This improved design improves the response time of the current generation of shopbots by using past prices to anticipate future prices and by

selectively searching stores likely to yield good offers. Using data collected from 28 stores between August 1999 and January 2000, we find that shopbots can predict prices with a high degree of accuracy. Additionally, our improved design filters offers to reduce the number presented to the consumer and subsequently the consumer's cognitive effort.

Our research contributes to two distinct literatures—the study of intelligent agents and the study of choice when the choice set is unknown. In the last decade, research by computer scientists on shopbots, and more broadly on intelligent agents, has exploded. While many novel technical problems have been addressed, the economic behavior of these intelligent agents is often quite primitive. For example, the Kasbah agent has a utility function composed of a single element—price—and buys if the price is below a certain threshold (Chavez and Maes 1996). Drawing on the economics and marketing literatures, we use a multielement, compensatory utility function, allowing the shopbot to make decisions that better reflect actual consumer behavior. We feel that economic and marketing approaches to consumer behavior have much to contribute to this developing work in helping intelligent agents make better decisions on behalf of consumers.

In contrast to traditional choice models (Ben Akiva and Lerman 1985)—where all offers are known to consumers—in our setting consumers must choose whether to search using a shopbot based only on their prior expectation of the shopbot's choice set. The consumer will know the shopbot's choice set only after the choice to use a shopbot is made. In other words, the consumer must make a decision not only about which item to choose from a set, but also the metachoice of which choice set to choose based only on their prior expectations of the choice sets.

2. Literature Review

Shopbots present a retailing problem that has been studied primarily by computer scientists interested in the design of intelligent agents. However, their interest lies in the technical challenges of agent construction (Jennings et al. 1998). This work can be classified into knowledge representation and agent learning (Poole 1997), agent adapting (Doorenbos et al. 1997, Pham and Karmouch 1998), and multiagent communication (Chavez and Maes 1996). The most closely related research in management science to our shopbot design problem is an empirical study of choice at the shopbot Dealtime.com by Smith and Brynjolfsson (2001). Although direct work on shopbots has not been plentiful in the management science community, there are many streams of research about consumer behavior that can enhance our understanding of shopbot design. We briefly discuss these research streams and their relevance to our shopbot design problem.

Choice with Uncertain Sets. Our analysis shares some commonalities with the optimal cutoff rule defined by Feinberg and Huber (1996), in which a decision maker must decide a screening mechanism for which alternatives to consider with imperfect information. Instead of assuming that consumers will consider all alternatives offered in a choice problem, consumers may limit their attention to a smaller set to reduce effort by explicitly considering the relative gains of adding another item (Hauser and Wernerfelt 1990; Roberts and Lattin 1991, 1997; Siddarth et al. 1995). The shopbot faces a similar problem, except that it must act on behalf of the consumer in forming the consideration set, which in turn will become the consumer's choice set. In this case the shopbot must explicitly value search on behalf of the consumer (cf. Moorthy et al. 1997). The basic advantage of shopbots or, more generally, electronic markets is to lower consumer search costs (Bakos 1997).

Information Overload and Cognitive Costs. Electronic environments tend to make it easy to present a large amount of information to consumers. But from the consumer's vantage point, combing through a long list of items can present a taxing cognitive problem, resulting in information overload. The notion of cognitive overload has a long history in consumer behavior research (Jacoby et al. 1974, Jacoby 1984, Keller and Staelin 1987). Johnson and Payne (1985) show that consumers are willing to trade off cognitive effort in the decision-making process for accuracy, where cognitive effort can be defined in terms of elementary information processes (Bettman et al. 1990). These information-processing arguments can be used to improve the design of electronic interfaces (Hoque and Lohse 1999). Our approach is to directly incorporate measures of the cognitive costs of decision making (Shugan 1980) to offset the benefits of consumption.

Decision-Support Systems. Computer-aided decision-support systems, such as shopbots, might be able to reduce cognitive effort and allow consumers to make better decisions by ordering and filtering alternatives (Sproule and Archer 2000). Widing and Talarzyk (1993) found that ordering alternatives by a linear weighted average of the product's attributes was superior to having users select cutoffs that would eliminate poor choices or randomly order the items. Häubl and Trifts (2000) use recommendation agents (also using weighted averages of the product's attributes) to increase both the accuracy and speed of consumer decisions. Hoch and Schkade (1996) advocate using decision-support systems to exploit the relative strengths of human decision makers with technology that can overcome their weaknesses.

3. Operational Decisions in Shopbot Design

The typical shopping process for a consumer begins with a consumer identifying a product to purchase, such as a book. We use books here because they are commonly purchased online, although our technique can be adapted to other consumer products and services like music, travel, and auctions. Identification of the book may happen after surfing the Web, reading a book review, or browsing an online or physical bookstore.

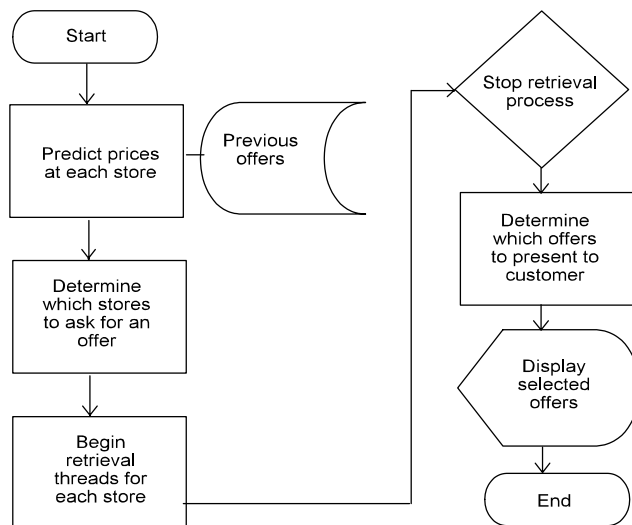
Once the consumer has identified a product to purchase, he or she must decide where to purchase the product either using a shopbot or directly visiting an online store. Suppose the consumer elects to use a shopbot. Currently shopbots query all stores at the time a request is made and report all results to consumers.¹ Query time can be substantial, with the modal time for Pricescan.com and Dealttime.com being about 45 seconds in 2000. The tails of the distribution are fat, especially for Dealttime.com, which times out at three minutes in 10% of the searches. In comparison, individual stores respond more rapidly, with modal response times of two seconds, but also have the potential for slow response.

We conjecture that many consumers may choose to go directly to an online bookstore and avoid a shopbot because they are impatient. Usability research shows that delays of more than 10 seconds result in a loss of user attention (Nielsen 2000, p. 44), and consumers place high emphasis on download times (Udo and Marquis 2001). Instead of querying all stores, shopbots could use prior expectations about prices to identify stores that are most likely to have low prices or high utility. More sophisticated shopbots could reduce both the average and the tails of their distribution by reducing the number of queries or interrupting searches, making shopbots more attractive to consumers.

An additional problem with current shopbot design is the number of alternatives presented. Every additional alternative presented will force the user to expend cognitive effort. Given that consumers are cognitive misers, additional time spent on cognitive activity is more taxing than simply waiting. One suggestion would be to show consumers the cheapest offer. However, it is unlikely that the shopbot could be so confident, because there is a random component to utility. Nor is the best offer synonymous with the cheapest offer, because other attributes such as deliv-

¹ Some shopbots may have direct access to price information from stores because of special marketing agreements. In these circumstances query time may be substantially lessened because it requires a lookup in a local database rather than a query to a store and appropriate parsing of the HTML document that is returned.

Figure 1 Flow Diagram Illustrating the Operational Decision Process for a Shopbot



ery and the store’s identity may affect a consumer’s utility. We propose that shopbots could sort the retrieved offers by utility and offer a consumer fewer but more relevant choices by eliminating unlikely alternatives. Fewer alternatives reduce the cognitive burden to the consumer; however, that also reduces the utility of the offer set.

For our purposes we assume that our customer is interested in searching using a shopbot. First, the shopbot presents the consumer with choices such as location (country), state (for sales tax purposes), and currency. Once the consumer makes his or her choices and initiates the search, the shopbot queries all of the bookstores in its search set and tabulates the information from stores that respond within a specified period. The shopbot then presents the consumer with prices and shipping options. These offers are often ranked by price or the placement fees paid by the store to the shopbot. Shopbots may filter these stores for availability and add information about the online bookstores (e.g., MySimon.com also provides Gomez ratings). Note that if a consumer prefers that the data be ranked by some other criteria, he or she can rerank it based on other data fields such as unit price, store name, delivery type, and so forth.

Figure 1 illustrates the decision framework we assume for the shopbot’s operational decisions. First, the shopbot is given the book title to search. Second, the shopbot makes predictions about the price and response time at each store it will consider querying. These prior predictions can be used to determine which stores to search and how long to wait for a response. We assume that only realized offers (i.e., those offers retrieved from a store) can be presented to consumers. This forces shopbots to actively

search and not rely on prior expectations.² Finally, once the process threads³ that search the stores are started, the shopbot must decide whether to wait for all threads to finish their retrieval or to prematurely stop the retrieval process. Once all offers are collected, the shopbot decides which of these offers to present to the customer. Separating the presentation and query decisions allows the shopbot to respond to poor initial predictions.

4. A Utility Model of Consumer Interaction with a Shopbot

In this section, we formalize the operational decisions made by the shopbot that were introduced in the last section. We begin in §4.1 by presenting a random utility model that allows us to quantify the value of a set of offers and balance this value against the cognitive effort necessary to compare offers within this set. A model of the shopbot decision process is given in §4.2. The optimal shopbot design is considered in §§4.3 and 4.4.

We assume that each customer initiates a search for a single book and then chooses a single item from the set of choices. The shopbot makes the decisions described in §3 about how many stores to query, how long to wait for responses, and how many of these responses to present to the consumer. The shopbot's goal is to maximize the consumer's utility. Throughout this discussion we model the behavior of a single consumer. However, we believe heterogeneity in consumer response is likely to be present. Hence, a shopbot should estimate the parameters at an individual level. Although it is outside of the scope of this research to estimate the parameters of the utility model, we discuss how this could be done in §8.

We assume that there is a universe of S stores that can be queried by the shopbot. The shopbot must choose which stores to query and how many seconds (t^*) to allow the query to run. The decision of which stores to query (\mathbf{q}) is encoded as an $S \times 1$ vector of binary decision variables, where $q_i = 1$ if the store is to be queried and $q_i = 0$ otherwise. The time it takes for the i th store to respond is assumed to be a random variate T_i and the observed value is t_i . The corresponding $S \times 1$ vector of observed response times is \mathbf{t} . The $S \times 1$ vector \mathbf{r} records whether the store responds to the query within t^* seconds, that is $r_i = 1$ if $t_i < t^*$

and $r_i = 0$ otherwise. Notice that \mathbf{r} is a random variate that is not chosen by the shopbot directly, but is a function of realized retrieval times (\mathbf{t}) and the shopbot's choice of \mathbf{q} and t^* . For example, suppose the i th alternative has been queried, $q_i = 1$. If the realized time is less than the time at which all uncompleted threads are to be canceled, $t_i < t^*$, then the offer will be retrieved, $r_i = 1$, otherwise even though the store was queried, its offer will not be retrieved, $r_i = 0$.

Once this query is completed, the shopbot decides which offers to present to the customer. We define the vector \mathbf{p} to represent which offers to present, where $p_i = 1$ if the offer is presented to the consumer and $p_i = 0$ otherwise. The total number of stores queried, stores retrieved, and offers presented equals Q , R , and P , respectively. The sum of the vector elements yields the number of chosen items, $P = \sum p_i$, $Q = \sum q_i$, and $R = \sum r_i$. Notice the following inequality holds: $P \leq R \leq Q \leq S$. We assume that only queried stores can be retrieved, and only those offers that are retrieved can be offered. This relationship can be represented by the inequality $\mathbf{p} \leq \mathbf{r} \leq \mathbf{q}$, which is defined in terms of the element-wise inequalities, $0 \leq p_i \leq r_i \leq q_i \leq 1$. The requirement that only retrieved offers are presented reflects current shopbot design, but could be relaxed in future research to allow predicted offers to be reported.

4.1. A Random Utility Model

Consumer utility from the i th store is a random variate, U_i , and $\mathbf{U} = \{U_1, U_2, \dots, U_S\}$. We assume that utility is specific to an individual and likely to be heterogeneous, but avoid adding a subscript h to denote each consumer for clarity. We argue that the consumer will choose only one item within the set, and therefore the utility from a set of choices is equal to the utility from the best alternative or the maximum of the set. We define the operator $\mathbf{U}(\mathbf{p})$ to denote the set comprised of the elements of the set \mathbf{U} that correspond with those elements in \mathbf{p} that equal 1. Additionally, we define the operator $\langle\langle \mathbf{p} \rangle\rangle$ to denote the set of the indices that correspond to the nonzero elements of \mathbf{p} . For example, if $\mathbf{p} = [1 \ 1 \ 0 \ 0 \ 1]'$ then $\langle\langle \mathbf{p} \rangle\rangle = \{1, 2, 5\}$ and $\mathbf{U}(\mathbf{p}) = \{U_1, U_2, U_5\}$.

The utility of the i th offer (U_i) is modeled as the sum of the utility derived directly from the product (\hat{U}_i) and the disutility associated with the waiting time for the online stores to respond to the shopbot's query (W), the overhead of launching Q threads on the shopbot, and the cognitive effort (C) associated with evaluating the set of alternatives:

$$U_i = \hat{U}_i - \xi W - \omega Q - \lambda C. \quad (1)$$

Because the latter three terms decrease utility, we assume $\xi, \omega, \lambda > 0$. Additionally, these disutility terms are not subscripted by i because they are identical for all items offered. They are a function of the

² This assumption could be relaxed so that expected offers could be shown to the consumer. However, this adds an additional layer of complexity, because the consumer must now evaluate the probability than an offer will be available.

³ A process thread refers to a task within a program that can be run independently from the main body of the program. Most operating systems allow multiple threads to be run at the same time. In our shopbot example, the query for each store would be launched as a separate thread run in parallel.

set of offers queried and presented. We consider the construction of each of these variables below.

Utility. We assume that product utility (\hat{U}_i) is the sum of a component due to a linear function of the attributes of the product (\bar{U}_i) and a stochastic component (ϵ_i):

$$\hat{U}_i = \bar{U}_i + \epsilon_i, \quad \text{where } \bar{U}_i = \sum_{j=1}^A \beta_{ij} a_{ij}, \quad (2)$$

β_{ij} denotes the weight, and a_{ij} denotes the value of the j th attribute of the i th product. For example, in our application to bookstores, the attributes of a book are price, delivery, tax, shipping time, and store brand name. Note that store attributes like store brand name, return policies, and availability can be included as well as product attributes. For example, a dummy variable corresponding with store brand name can capture a consumer's affinity for a popular store like Amazon. (We discuss this further in §5.3 and Table 6.)

This compensatory utility model allows increases in one attribute to offset a decrease associated with another attribute versus a noncompensatory approach that assumes consumers have rigid thresholds or reservation prices. A primary benefit of the compensatory approach is that it can better capture trade-offs that consumers make. For example, the higher price of a book at Amazon may be offset by faster delivery.

The stochastic component is due to unobservable factors or random evaluation error by the consumer and represents the fact that we cannot predict utility with certainty. We assume that these ϵ_i are independently and identically distributed and follow an extreme value distribution with a zero location parameter and a scale parameter of θ . The cumulative distribution of $\hat{U}_i = \bar{U}_i + \epsilon_i$ is

$$\Pr[\hat{U}_i \leq x] = \exp\{-e^{-(x-\bar{u}_i)/\theta}\}. \quad (3)$$

The mean and variance are $E[\hat{U}_i] = \bar{U}_i + \gamma\theta$ and $\text{Var}[\hat{U}_i] = \pi^2\theta^2/6$, where γ is Euler's constant (i.e., $\gamma \approx 0.57722$). The choice of the extreme value distribution is motivated by its extensive use in choice models (McFadden 1980). Usually the scaling parameter θ is set to unity to address the identification of the parameter estimates in the utility function when a constant is included.

Waiting Time. The response time (W) is the time associated with retrieving the set of offers. This is the time for the slowest store to respond, $\max(\mathbf{T}\langle \mathbf{q} \rangle)$, unless this time exceeds the interrupt time set by the shopbot when the retrieval threads are launched, t^* , in which case any remaining threads will be ignored. Hence, response time (W) is defined as

$$W = \min(t^*, \max(\mathbf{T}\langle \mathbf{q} \rangle)). \quad (4)$$

Server Overhead. We assume that the time for the shopbot's server to start and service the threads that

handle the query to an online store is proportional to the number of stores queried (Q) and is measured by ω . The term ωQ measures the disutility that consumers experience as the result of the total delay in response time from launching Q threads. For most situations one would expect that ω would be insignificant, because the time to start a thread is on the order of a few milliseconds; hence $\omega \approx 0$. However, when the system is above its operating capacity during peak periods there could be a measurable delay for the consumer (e.g., perhaps a second or more). An interpretation of ω is delay due to the load on the shopbot's server. The inclusion of $\omega > 0$ prevents the situation in which the optimal shopbot design would launch a query at every store. Under this solution the shopbot's server would always be overutilized and response seriously degraded.

Cognitive Costs. A metric for evaluating the cognitive costs (C) associated with comparing P alternatives each with A attributes was proposed by Shugan (1980):

$$C = (A - 1)(P - 1). \quad (5)$$

The motivation for this formulation is that cognitive costs are proportional to the number of alternative and attribute pairs. For example, a set of three alternatives with four attributes will take six comparisons.

Consumer Heterogeneity. Although we have suppressed the subscript h for individuals in our notation, past marketing research shows that individual heterogeneity is likely to be present and important. Some consumers may value fast delivery, while others think more highly of one store than another. This heterogeneity can be represented by assuming that the β , or more properly the individual level value β_h for individual h , follows a multivariate distribution: $\beta_h \sim N(\boldsymbol{\beta}, \mathbf{V}_\beta)$. There has been a great deal of work in estimating this type of hierarchical Bayesian model in choice occasions (see Rossi et al. 1996). Moreover, an individual consumer's preferences may vary across choice occasions, sometimes needing quick delivery and other times, lower prices. Hence, more flexible assumptions about these distributions can be added.

4.2. The Shopbot's Decision Problem

The shopbot's decision problem is to maximize utility in a two-step process. First, the shopbot needs to make decisions about which stores to query (\mathbf{q}) and how long to wait (t^*) in order to maximize the expected utility of the offers (\mathbf{p}) that it expects to make to a consumer:

$$\max_{\mathbf{q}, t^*} E[\max(\mathbf{U}\langle \mathbf{p} \rangle)]. \quad (6)$$

At this stage the shopbot does not know either prices or utility with certainty. Instead, it must predict utility and its components using past information. In addition, the shopbot does not need to decide which stores it will present because this decision can be made after the offers are retrieved, but it does need to predict which \mathbf{p} is likely to be used.

Second, after t^* seconds have elapsed and the shopbot has actually retrieved a set of offers (\mathbf{r}) it needs to make a decision about which offers to present (\mathbf{p}) to the consumer:

$$\max_{\mathbf{p}} E[\max(\mathbf{U}(\mathbf{p})) | \bar{\mathbf{U}}(\mathbf{r})]. \quad (7)$$

Notice that at this point the shopbot knows the portion of utility due to the product attributes ($\bar{\mathbf{U}}$) but not the portion due to the evaluation error (ϵ_i).

In our notation current shopbot design can be represented as a query to all stores ($\mathbf{q} = \mathbf{v}$), where $\mathbf{v} = [1 \ 1 \ \dots \ 1]$, if the search is interrupted it is not completed within 30 seconds ($t^* = 30$), and all retrieved offers are presented ($\mathbf{p} = \mathbf{r}$). However, there are 2^S possible combinations that the shopbot could consider for query. In addition, for any proposed retrieval set there are many combinations to show consumers. In total, there are

$$\sum_{i=1}^S \sum_{j=1}^i \binom{S}{i} \binom{i}{j}$$

possible sets that the shopbot could offer the consumer. A universe of 10 stores yields 58,025 combinations, while a universe of 30 stores yields more than 205 trillion combinations. In the next two subsections we consider the solution to this problem. Because this is a two-step optimization problem we begin with the final stage decision and then consider the initial decision assuming the final stage decision will be optimized.

4.3. Deciding Which Retrieved Offers to Present to the Consumer

We begin by considering the shopbot's decision at the final stage, where the shopbot must decide which offers to present to the consumer by choosing \mathbf{p} , which implicitly defines P . At this stage the shopbot has already decided which stores to query (this decision will be considered in §4.4) and retrieved a set of R offers (\mathbf{r}). Because the offers have been retrieved, the product attributes (price, delivery cost, etc.) are known at this point; hence the random variate representing utility, \bar{U}_i , will take the observed value \bar{u}_i , where $i \in \langle \mathbf{r} \rangle$. There is still uncertainty associated with the random variates \hat{U}_i due to the consumer's evaluation error ϵ_i . The optimization problem that corresponds with this stage given in (7) can be

written as

$$\begin{aligned} & \max_{\mathbf{p}} E[\max(\mathbf{U}(\mathbf{p})) | (\bar{\mathbf{U}})\langle \mathbf{r} \rangle] \\ & = \max_{\mathbf{p}} E[\max(\hat{U}(\mathbf{p})) - \xi W - \omega Q \\ & \quad - \lambda(A-1)(P-1) | \bar{\mathbf{U}}(\mathbf{r})]. \end{aligned} \quad (8)$$

This simplification uses the fact that the disutility terms for waiting time, server overhead, and cognitive costs are functions of the set of alternatives retrieved, so they can be extracted from the maximization function. The expectation is computed and is conditional on the information set $\{\bar{\mathbf{U}}(\mathbf{r}), \xi, \omega, \lambda, \theta\}$, which is suppressed for notational convenience throughout this subsection.

Clearly, the best alternatives to present are those with the highest expected utilities. To find the set of \mathbf{p} offerings that maximize the expected utility of this set, we order the retrieved offers by their expected utilities. The utilities of the alternatives in the offer set are denoted by $\{\bar{u}_{R:R}, \dots, \bar{u}_{R-P+1:R}\}$, where $\bar{u}_{R:R} \geq \bar{u}_{R-1:R} \geq \dots \geq \bar{u}_{R-P+1:R}$. This reduces our problem to a decision about how many offers to present (P). Given P , the elements of \mathbf{p} are determined by the relation $p_i = 1$ if $r_i = 1$ and $\hat{U}_i \geq \hat{U}_{R-P+1:R}$, otherwise $p_i = 0$.

The properties of the extreme value distribution imply that the maximum variate, $\hat{U} = \max(\hat{U}_{R:R}, \hat{U}_{R-1:R}, \dots, \hat{U}_{R-P+1:R})$, will also follow an extreme value distribution with location of $\theta \ln(\exp\{\bar{u}_{R:R}/\theta\} + \exp\{\bar{u}_{R-1:R}/\theta\} + \dots + \exp\{\bar{u}_{R-P+1:R}/\theta\})$ and scale of θ . (Note that this simplification depends upon the i.i.d. assumption of ϵ_i .) Hence, the expected utility of the offer set is

$$\begin{aligned} & E[\max(\mathbf{U}(\mathbf{p})) | P, \bar{\mathbf{U}}(\mathbf{r})] \\ & = \theta \ln \left(\sum_{i=1}^P \exp\{\bar{u}_{R-i+1:R}/\theta\} \right) + \theta \gamma - \xi W - \omega Q \\ & \quad - \lambda(A-1)(P-1). \end{aligned} \quad (9)$$

The ordering of the items to include is determined by the sorted order of $\bar{u}_{R-i+1:R}$. To determine the number of elements to include in this set, notice that the two terms that involve P , $\theta \ln(\sum \exp\{\bar{u}_{R-i+1:R}/\theta\})$ and $-\lambda(A-1)(P-1)$, are monotonically increasing and monotonically decreasing in P , respectively. Therefore, we can find the optimal value for P by first evaluating at $P=1$ and subsequently incrementing P until the expected utility begins to decline. This yields the following stopping rule to determine the optimal value P^* :

$$\begin{aligned} & E[\max(\mathbf{U}(\mathbf{p})) | P = P^* + 1, \bar{\mathbf{U}}(\mathbf{r})] \\ & < E[\max(\mathbf{U}(\mathbf{p})) | P = P^*, \bar{\mathbf{U}}(\mathbf{r})] \end{aligned}$$

$$\begin{aligned} &\Rightarrow \frac{\exp\{\bar{u}_{R-P^*:R}/\theta\}}{\sum_{i=1}^{P^*} \exp\{\bar{u}_{R-i+1:R}/\theta\}} \\ &< \exp\left\{\frac{\lambda(A-1)}{\theta}\right\} - 1. \end{aligned} \quad (10)$$

In other words, find the largest P such that the relative gain from adding this alternative exceeds the added cognitive costs to the consumer of its evaluation. The vector of offers to present (\mathbf{p}) is implicitly defined by setting those elements that correspond with the indices of $\{\bar{u}_{R:R}, \dots, \bar{u}_{R-P^*:R}\}$ to unity and zero otherwise.

Special Case with Identical Offers. To proceed further, we assume that all offers have the same expected value, $\bar{u}_{R-i+1:R} = \bar{u}$, but the evaluation in the utility term is stochastic. Therefore, the question is not which offers to present (because they are identical), but how many. Under this assumption we can simplify (9) as

$$\begin{aligned} E[\max(\mathbf{U}(\mathbf{p})) | P, \bar{u}_i = \bar{u}] \\ = \bar{u} + \theta(\gamma + \log(P)) - \xi W - \omega Q \\ - \lambda(A-1)(P-1). \end{aligned} \quad (11)$$

If we allow P to take noninteger values, we can differentiate (11) with respect to P and find the optimal value

$$P^* = \frac{\theta}{\lambda(A-1)}. \quad (12)$$

Notice that the optimal set size increases as cognitive costs decrease ($\lambda \rightarrow 0_-$) or the variance (θ) of utility increases. Additionally, we can then show that (11) creates an upper bound of (9) by setting \bar{u} to the maximum value, $\bar{u}_{R:R}$, and that (12) bounds the solution in (10).

Example with Identical Offers. To illustrate these relationships, suppose the average book generates 10 utils with a standard deviation of about two utils, which corresponds with $\bar{u} = 9.1$ and $\theta = 1.6$. Additionally, assume that a book has four attributes (brand, price, shipping cost, and delivery time), $A = 4$, and set ξ and ω to zero without loss of generality in this example. This means $P = 1.6/3\lambda$. Suppose each additional unit of cognitive effort decreases utility by 0.1 utils ($\lambda = 0.1$). The optimal number of offers to present to the customer is 5.3, or about five books. Doubling the effects associated with cognitive effort ($\lambda = 0.2$) will reduce the set to $P^* = 2.7$ or about three books. Again, suppose $\lambda = 0.1$. Then, the utility generated by offering five books (the optimal number) is 11.4. If the shopbot were to present 20 books, then utility would drop by 20%. In summary, ignoring cognitive costs and naively presenting all offers retrieved is not optimal.

4.4. Deciding Which Stores to Query and How Long to Wait for a Response

At this initial stage the shopbot must decide which stores to query, \mathbf{q} , and whether those queries should continue until completion or whether they should be interrupted prematurely at time t^* as stated in (6). We assume that the shopbot will make the optimal decision about which retrieved offers to present as discussed in §4.3. Neither the retrieval times nor the offers are known, as in the previous subsection, so both are assumed to be stochastic variates. The utility from the product attributes (price, delivery time, availability, etc.) is the random variate \bar{U}_i and not the observed value \bar{u}_i , as in the previous subsection, where $i \in \langle\langle \mathbf{q} \rangle\rangle$. We assume that response time and utility are independent, so discount stores will return responses as quickly as expensive stores. Empirically this assumption will be justified by the discussion in §5. Formally we can make use of the result from the previous section and rewrite (6) as follows:

$$E[\max(\mathbf{U}(\mathbf{p}))] = E[E[\max(\mathbf{U}(\mathbf{p})) | \bar{\mathbf{U}}(\mathbf{r})]]. \quad (13)$$

Equation (13) makes use of the following relation: $E[X] = E[E[X | Y]]$.

The solution of the inner expectation was given in (9), and we will assume that \mathbf{p} is chosen optimally by sorting the retrieved values and using the first P^* , where P^* is a function $\mathbf{U}(\mathbf{r})$ as defined in (10). The outer expectation requires integrating over the distribution of product utility and retrieval times and summing over all possible permutations of retrieval sets weighted by their probability:

$$\begin{aligned} &E[E[\max(\mathbf{U}(\mathbf{p})) | \bar{\mathbf{U}}(\mathbf{r})]] \\ &= \sum_{\mathbf{r} \in \mathcal{F}} \Pr[\mathbf{r}] E \left[\theta \ln \left(\sum_{i=1}^{P_r^*} \exp\{\bar{\mathbf{U}}(\mathbf{r})_{R_r-i+1:R_r}/\theta\} \right) \right. \\ &\quad \left. - \lambda(A-1)(P_r^*-1) | \mathbf{r} \right] + \theta\gamma - \xi E[W] - \omega Q, \end{aligned} \quad (14)$$

where \mathcal{F} denotes all possible 2^Q permutations of the query set, $\Pr[\mathbf{r}]$ denotes the probability that query set \mathbf{r} is retrieved, $R_r = \sum r_i = \mathbf{r}'\mathbf{t}$, which is the number of retrievals made from query set \mathbf{r} , P_r^* denotes a stochastic variate that represents the optimal number of offers to present given the retrieved set \mathbf{r} , and $\bar{\mathbf{U}}(\mathbf{r})_{R_r-i+1:R_r}$ denotes the R_r-i+1 ordered statistic from the set $\bar{\mathbf{U}}(\mathbf{r})$ with R_r elements. The expectation is computed conditional on the information set $\{\mathbf{q}, t^*, \xi, \omega, \lambda, \theta, \beta\}$, which is suppressed for notational convenience hereafter.

Additionally, to avoid problems when the rare occurrence of no offers being retrieved, we assume that one inferior offer is always available. In our bookstore example this will be a special order that is twice

the list price and takes six months to receive. As long as any offer is retrieved this offer is dominated and will not be considered or presented.

The probability that any member of \mathcal{F} is retrieved can be computed in the following manner. Consider the probability that all items are retrieved. It equals $\Pr[t_{i[1]} \leq t^*, t_{i[2]} \leq t^*, \dots, t_{i[Q]} \leq t^*]$, where $i[j]$ denotes the j th element of $\langle\langle \mathbf{q} \rangle\rangle$. The probability that all items except the first one is retrieved is $\Pr[t_{i[1]} > t^*, t_{i[2]} \leq t^*, \dots, t_{i[Q]} \leq t^*]$. The remaining elements can be computed in a similar manner. Notice that \mathcal{F} has 2_Q members and its evaluation leads to a computational problem because of the large number of combinations. A set of 10 stores yields 1,024 combinations to evaluate, while 30 stores will lead to more than one billion combinations. Before considering the optimal solution we discuss distributional assumptions for waiting times and utility.

Waiting Times Are Exponentially Distributed.

Equation (14) requires the computation of $E[W]$. To proceed further we assume that retrieval times follow an exponential distribution and are independently and identically distributed. This assumption is supported by our empirical analysis of §5.2. The probability of observing a query to a selected store is denoted by $\tau \equiv \tau(t^*) = \Pr[T_i < t^*] = 1 - \exp\{-t^*/\zeta\}$. The expected time to observe a queried store is $E[T_i] = \zeta$ with variance $\text{Var}[T_i] = \zeta^2$. The expected time to observe the set of queries is $E[\max(\mathbf{T}(\mathbf{q}))] = \zeta \sum_j (1/j)$, where j goes from 1 to Q . For example, if the expected time to observe one store is two seconds, then the expected time to observe a set of ten stores is 5.9 seconds. The density function of the maximum variate is $\Pr_{\max(\mathbf{t}(\mathbf{q}))}[t] = Q\tau(t)^{Q-1} \exp\{-t/\zeta\}/\zeta$. We can derive the expected value of W , which can alternatively be described as the maximum from a distribution of exponential variates censored at t^* :

$$\begin{aligned} E[W] &= E[\min(\max(\mathbf{T}(\mathbf{q})), t^*)] \\ &= \int_0^{t^*} Q\tau(t)^{Q-1} \exp\{-t/\zeta\}/\zeta dt \\ &\quad + t^*(1 - \tau(t^*)^Q). \end{aligned} \tag{15}$$

If $t^* < E[\max(\mathbf{t}(\mathbf{q}))]$ then $E[W] \approx t^*$. The probability of observing an individual member \mathbf{r} from set \mathcal{F} is

$$\Pr[\mathbf{r}] = \tau^R (1 - \tau)^{Q-R} \quad \text{where } R = \sum_i r_i. \tag{16}$$

Utility Is Logistically Distributed. The expectation in (14) implicitly requires the integration of (9) over the distribution of $\bar{\mathbf{U}}(\mathbf{r})$. To proceed further we make additional assumptions about the distribution of product utility. The natural choice is to assume that the attributes, such as price, are normally distributed at each store, and therefore utility itself is normally

distributed. The problem is that the distribution of the order statistics from a normal distribution does not yield closed-form solutions. A reasonable approximation to the normal distribution is the logistic distribution.⁴ If we assume that the attribute component of utility is identically and independently logistically distributed across stores, $\bar{U}_i \sim L(\mu, \sigma^2)$ for $i = 1, 2, \dots, S$. The cumulative distribution function of \bar{U}_i is defined as

$$\Pr[\bar{U}_i \leq x] = \left[1 + \exp\left\{-\frac{x - \mu}{\sigma}\right\} \right]^{-1}. \tag{17}$$

The mean and variance of the logistic distribution are $E[\bar{U}_i] = \mu$ and $\text{Var}[\bar{U}_i] = \sigma^2 \pi^2/3$. The mean of the maximum variate of a logistic distribution is $E[\max(\bar{U}_{R,R})] = \sigma(\psi(R) + \gamma) + \mu$, where $\psi(x) = \Gamma'(x)/\Gamma(x)$. Balakrishnan (1992) presents a full discussion of the properties of the logistic distribution.

We can now bound the component of the expectation in (14) that involves the logarithmic function:

$$\begin{aligned} E\left[\theta \ln\left(\sum_{i=1}^{P^*} \exp\{\bar{U}_{R-i+1,R}/\theta\}\right) \mid \mathbf{r}\right] \\ \geq E[\bar{U}_{R,R} \mid \mathbf{r}] = \sigma(\psi(R) + \gamma) + \mu. \end{aligned} \tag{18}$$

This lower bound can be derived by factoring $\bar{U}_{R,R}$ and showing that the term in the logarithmic function always exceeds unity. Intuitively, we are focusing on the “best” single product; if we could know the consumer’s choice with certainty ($\theta \rightarrow 0$), then this inequality becomes an equality.

Solution. A general analytical solution to the maximization of expected utility given in (12) with respect to Q and t^* is not known. However, we can derive a reasonable approximation for certain cases. First, we consider the case where some queries may be prematurely interrupted when their response time is longer than t^* . Second, we consider the more restrictive case where response time is not interrupted, but all queries that are launched are allowed to execute until completion.

Time is bounded (t^ constraint is active).* If we make four simplifications—(a) any uncompleted query threads are interrupted at time t^* ; (b) P^* is replaced with the approximation defined in (12); and (c) we use the lower bound of expected utility from Equation (18); and (d) we use t^* to approximate W ,

⁴ The distribution of the logistic distribution has longer tails than the normal and is more closely approximated by a t distribution (Mudholkar and George 1978). Also, the variance of the standard logistic is $\pi^2/3$, so the variance parameter of the normal distribution should be scaled by $\pi/3$ before comparing it to the logistic distribution’s scale parameter (for more discussion see David 1981, pp. 77–78).

then the expected utility in Equation (14) yields

$$E[E[\max(\mathbf{U}\langle\mathbf{p}\rangle) | P^*, \mathbf{U}\langle\mathbf{r}\rangle]] \approx \sum_{j=1}^Q \binom{Q}{j} \tau^{Q-j} (1-\tau)^j \{\sigma(\psi(j) + \gamma) + \mu\} + (1-\tau)^Q \mu_0 - (\theta - \lambda(A-1)) + \theta\gamma - \xi t^* - \omega Q, \quad (19)$$

where μ_0 is the utility from the alternative when no offers are made, and we assume that $\mu_0 < \mu$. The assumption that times and utility are i.i.d. is critical in this simplification, because both the expected utility and the probability the set is realized depend only upon the size of the set and not the stores selected. The optimal solution for t^* given Q is

$$t^* = \zeta \ln\left(\left\{\frac{\zeta\xi}{Q(-\mu_0 + \mu + \gamma\sigma + \sigma\psi(Q))}\right\}^{-1/Q}\right). \quad (20)$$

The solution for the optimal value of Q can be found by substituting (20) into (19) and enumerating the values of Q beginning with unity until the expected utility begins to decline.

Time is not bounded ($t^* \rightarrow \infty$). If we make three simplifications—(a) all query threads that are launched are allowed to run to completion then $\tau \rightarrow 1$; (b) the value of P^* is approximated using (12); and (c) we approximate $E[W] = E[\max(T_1, \dots, T_Q)] \approx \zeta \ln(Q)$ —then the expected utility found in (12) yields

$$E[\max(\mathbf{U}\langle\mathbf{p}\rangle)] \approx E\left[\theta \ln\left(\sum_{i=1}^{P^*} \exp\{\bar{U}_{Q-i+1:Q}/\theta\}\right)\right] - \lambda(A-1)(P^* - 1) + \theta\gamma - \xi\zeta \ln(Q) - \omega Q. \quad (21)$$

Replacing the expectation of the logarithmic term in (21) using the lower bound in (18) and using the approximation $\psi(x) \approx \ln(x + 0.5)$ when $x > 2$ yields the following solution:

$$Q^* = \frac{-2\zeta\xi + 2\sigma + \omega + \sqrt{8\zeta\xi\omega + (2\sigma - 2\zeta\xi + \omega)^2}}{4\omega}. \quad (22)$$

If either ξ or ζ are small relative to σ and $\omega \ll \sigma$, then $Q^* \approx \sigma/\omega$. In other words, the number of sites to query is directly proportional to the variance of utility and inversely related to the waiting time associated with the computational overhead of starting additional threads.

5. Empirical Study of Online Book Prices, Store Response Times, and Utility

In this section, we consider calibrating the parameters of the shopbot model proposed in §4. We begin by formulating a predictive model of price in §5.1, present

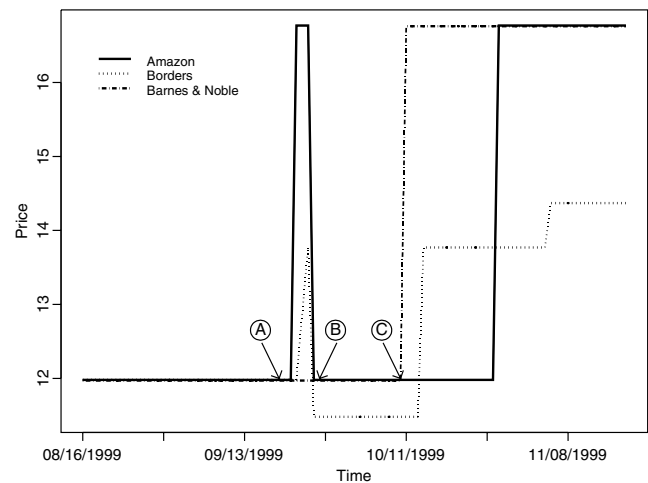
an analysis of store response times in §5.2, and discuss the part-worths of the utility function, the disutility of waiting, and cognitive costs in §5.3. A simulation study of utility using these parameters is presented in §6.

5.1. Predicting Prices

We analyzed the prices at 28 online bookstores for 60 books that were on the August 8, 1999, *New York Times* bestseller list over the course of a six-month period, from August 1999 to January 2000. During this period, even if the book fell off the bestseller list we continued to collect prices throughout the period. These books were part of a broader data set that included computer bestsellers and random books. We elected to focus on *New York Times* bestsellers because they had the highest variability in price. The data on prices were collected using automated agents from two major comparison-shopping engines and a number of individual stores. (For a further discussion of this data set and its construction, see Clay et al. 2001.)

Figure 2 illustrates typical price behavior for a selected book at several online bookstores. A striking feature of the price series is the persistence of prices. In fact, prices may remain at the same level for several weeks. Note that most price increases were precipitated by the book falling off the *New York Times* bestseller list. The book being added to the bestseller list results in a price decline. However, these effects are not always automatic and there can be delays of several days or weeks before any change results. Additionally, some stores, like Buy.com, will respond to changes in prices at another store with high likelihood, while Amazon.com seems to act more

Figure 2 Price Changes at the Three Top Online Booksellers for a Fiction Hardcover Book (August 1999–November 1999)



Note. The three events marked by A, B, and C denote when the bestseller dropped off the bestseller list on September 19, moved back on September 26, and moved off again on October 10, respectively.

like a price leader. In summary, price changes occur fairly infrequently, and only one quarter of all prices changed again within an eight-day period. The average time between price changes in our dataset is about four weeks.

These observations suggest that a shopbot can predict prices with a fair degree of precision. The predictability of prices means that shopbots can leverage information from previous retrievals to improve searches by selectively ignoring high-priced stores (or stores with low expected utility). We propose a formal statistical model to capture these stylized facts:

$$\text{relprice}_{sbt} = \begin{cases} \text{relprice}_{s,b,t-1} & \text{with probability } \rho_{sbt} \\ \phi_{s0} + \phi_{st}\text{relprice}_{s,b,t-1} + \delta'_{sb}\mathbf{x}_{sbt} + \epsilon_{sbt} & \text{otherwise,} \end{cases} \quad (23)$$

where relprice_{sbt} denotes the price at store s for book b on day t relative to its list price ($\text{relprice}_{sbt} = \text{price}_{sbt}/\text{listprice}_{sb}$) and \mathbf{x}_{sbt} denotes a vector of covariates used to forecast price changes. We assume that list prices do not change through time and are known. Relative prices are used to enable direct comparisons across stores and books. Also, note that we use the actual price of the product excluding shipping costs and tax, which were found to be deterministic in our sample.

If prices are changed, the magnitude of the price change is modeled as an autoregressive transfer function in (23). In our application the covariates used for forecasting (\mathbf{x}_{sbt}) are four indicator variables: *uph*, *downh*, *upp*, and *downp* that indicate, respectively, if the book is hardcover and moved onto the bestseller list, if the book is hardcover and moved off of the bestseller list, if the book is paperback and moved onto the bestseller list, and if the book is paperback and moved off of the bestseller list. The MLE estimates for predicting the relative prices are given in Table 1. For the most part when the book moves off of the bestseller list, the expected price increases, and when it moves back onto the bestseller list the expected price declines. The magnitude of the effect is larger for hardcover than paperback books. The overall fit of the models is moderate due to the fact that some price changes occur for inventory management issues, periodic price revisions, or other unobserved factors.

We assume that the number of days between a price change follows a negative binomial model (Hausman et al. 1984)

$$\Pr[n_{sbt}] = \frac{\Gamma(\gamma_{sbt} + n_{sbt})}{\Gamma(\gamma_{sbt})\Gamma(n_{sbt} + 1)} \left(\frac{\eta}{1 + \eta}\right)^{\gamma_{sbt}} (1 + \eta)^{-n_{sbt}}, \quad (24)$$

where n_{sbt} is the number of days between price changes, and while the γ parameter is allowed to

Table 1 Estimates of the Effects on Magnitude of Price Changes

Store	Constant	Lag price	<i>uph</i>	<i>downh</i>	<i>upp</i>	<i>downp</i>
1BookStreet	0.24 (0.05)	0.07 (0.12)	-0.27 (0.13)	0.17 (0.06)	-0.26 (0.07)	-0.02 (0.06)
Amazon.com	0.17 (0.04)	0.69 (0.08)	-0.45 (0.03)	0.23 (0.04)	-0.58 (0.03)	0.43 (0.04)
Barnes&Noble	0.43 (0.04)	-0.30 (0.10)	.	-0.03 (0.05)	-0.23 (0.07)	0.19 (0.05)
Borders	0.24 (0.07)	0.08 (0.10)	.	0.11 (0.11)	-0.30 (0.25)	0.32 (0.10)
Buy.com	0.87 (0.29)	0.01 (0.08)	-0.82 (1.54)	-0.53 (0.58)	-0.74 (1.03)	-0.41 (0.57)

Note. Standard errors of the estimates are given in parentheses below the estimates.

vary as a function of the covariates, $\gamma_{sbt} = \exp\{\mathbf{z}'_{sbt}\boldsymbol{\alpha}_s\}$. This equation implicitly defines the probability of a price change from (23). The expected number of days until a price change is γ_{sbt}/η and the variance is $\gamma_{sbt}(1 + \eta)/\eta^2$. In our application the covariates used for predicting when price changes occur (\mathbf{z}_{sbt}) include the days since the book’s bestseller status has changed and variables that correspond with the number of days since a price change at the major online bookstores. The maximum likelihood estimates of the parameters for the top five online stores are given in Table 2. Note that days since a change in the bestseller status has a significantly positive impact for Amazon.com and Buy.com. This implies that the longer it has been since the bestseller status has changed, the less likely a price change is; in other words, if a price change is to happen, it will occur soon after a price change. For most stores the timing of price changes does not appear to be statistically significant.

Our primary interest in model (23) is its ability to predict future prices given past pricing information. To predict price we can join the two conditional components together. Consider the expectation of the one-step ahead, price forecast:

$$\begin{aligned} E[\text{relprice}_{s,b,t+1} | \text{relprice}_{sbt}] &= \rho_{s,b,t+1}\text{relprice}_{sbt} + (1 - \rho_{s,b,t+1}) \\ &\cdot (\phi_{s0} + \phi_{s1}\text{relprice}_{sbt} + \delta'_{sb}\mathbf{x}_{s,b,t+1}). \end{aligned} \quad (25)$$

This expectation is equal to the probability that a price change has occurred times the conditional expectation plus the probability that a price change has occurred times the probability of no price change. The probability of a price change can be computed from the negative binomial distribution given in (24). Additionally, if \mathbf{x} and \mathbf{z} are not known, then forecasts can be used instead. Subsequent forecasts can be created by recursively applying (25). The eventual forecast model under stationarity, or the forecast when no

Table 2 Maximum Likelihood Estimates for Time Between Price Changes

Store	δ	Constant	Days since change in bestseller status	Number of days since change in price at ...				
				1BookStreet	Amazon.com	Barnes&Noble	Buy.com	Borders
1BookStreet	0.043 (0.007)	0.005 (0.177)	0.009 (0.005)		-0.005 (0.006)	0.003 (0.003)	0.000 (0.003)	0.005 (0.002)
Amazon.com	0.059 (0.009)	0.032 (0.148)	0.006 (0.003)	0.006 (0.004)		-0.001 (0.004)	0.002 (0.003)	0.003 (0.004)
Barnes&Noble	0.053 (0.008)	0.052 (0.153)	-0.003 (0.005)	0.010 (0.002)	-0.005 (0.005)		0.003 (0.002)	0.009 (0.002)
Borders	0.034 (0.005)	0.017 (0.152)	0.002 (0.006)	0.003 (0.005)	-0.002 (0.006)	0.001 (0.005)	-0.001 (0.003)	
Buy.com	0.039 (0.005)	-0.158 (0.128)	0.008 (0.002)	0.003 (0.002)	0.005 (0.002)	-0.006 (0.002)		0.004 (0.002)

Note. Standard errors are given in parentheses below the estimates.

information other than the list price is known, follows the usual autoregressive relationship:

$$E[\text{relprice}_{sbt}] = \frac{\phi_{s0} + \delta'_{sb} E[\mathbf{x}_{sbt}]}{1 - \phi_{s1}} \quad \text{and} \quad (26)$$

$$\text{Var}[\text{relprice}_{sbt}] = \frac{\sigma_s^2}{1 - \phi_{s1}^2}.$$

Table 3 presents the mean and standard deviation of expected prices under the assumption that

Table 3 The Mean and Standard Deviation of Prior Price Expectations for Each Store Without Any Previous Price Information

Store	Mean	Std. Dev.
1BookStreet	0.76	0.13
A1 Books	0.75	0.06
Alldirect.com	0.63	0.05
AlphaCraze.com	0.64	0.09
Amazon.com	0.63	0.13
Baker's Dozen online	0.99	0.06
Barnes&Noble	0.63	0.13
BCY Bookloft	0.72	0.07
Bigwords.com	0.77	0.06
Book Nook Inc.	0.99	0.05
Bookbuyer's Outlet	0.62	0.13
Books.com	0.70	0.09
Booksamillion.com	0.59	0.12
Booksnow.com	0.88	0.06
Borders	0.62	0.13
Buy.com	0.52	0.10
Cherryvalleybooks	0.89	0.02
Classbook.com	0.96	0.06
Codys Books	0.99	0.06
Computerlibrary.com	0.99	0.06
Fatbrain	0.65	0.15
HamiltonBook.com	0.70	0.07
Kingbooks.com	0.73	0.04
page1Book.com	0.99	0.07
Rainy Day Books	0.89	0.05
Rutherfords	0.89	0.05
Varsitybooks.com	0.75	0.05
WordsWorth	0.83	0.10

Note. These estimates are relative to the list price, and the predictions are normally distributed.

no information is known about the store (as defined by Equation (26)). For example, without any specific information about past book prices, other than, say, a list price of \$19.99, our best guess is that Buy.com will be cheapest with an expected price of \$10.39 and a standard deviation of \$2.00. Note that cheaper bookstores tend to have higher variance than more expensive stores.

The accuracy of the price predictions depends on the frequency with which the shopbot chooses to gather prices from the stores. For example, if the shopbot performs a weekly search of prices at Amazon.com on Friday, then prices are known with certainty at that moment, but the shopbot must predict prices between these weekly updates. The quality of the predictions will depend upon the frequency with which the shopbot queries the sites. The correlations between the actual and predicted price (standardized by its list price) if prices are collected every 3, 7, 14, and 30 days, or only once, are 0.987, 0.950, 0.914, 0.819, and 0.297, respectively. Clearly, good price forecasts can be generated without having to query the store. As the frequency of price sampling goes down, the forecasts start to deteriorate. Even with month-old price information, however, there is a correlation of 0.82, which implies a shopbot should be able to make an educated guess about the cheapest stores with even fairly old information. Decreasing the frequency of sampling lessens demands on network traffic (and increases demands on an internal database to look up past prices) but would also lessen the predictive accuracy. This trade-off between accuracy and speed needs to be considered by a shopbot when choosing an optimal frequency to gather prices.

5.2. Analysis of Store Response Times

Stores typically respond quickly to requests from a user or shopbot. The time for Amazon.com's server to respond to 5,700 queries during April to July 2000 between midnight and 6 AM were analyzed. For the most part responses are quick, with almost

80% of requests being retrieved in less than two seconds. However, at certain times network congestion or server overload can cause substantial delays or even no response. In fact, in the remaining 20% of requests Amazon.com could take up to 90 seconds to respond. Additionally, about 4% of requests are not returned within a 180-second period. The response times at BCY Bookloft and Barnes&Noble are similar. We found that response time can be modeled well by assuming that there is probability p that the store responds. Given that the store responds, the time between the user request and the store response is well described by a gamma distribution. Estimates of the probability of no response and the gamma parameters for three major stores are given in Table 4.

If we assume that retrieval times across stores are independent, then we can evaluate the time to retrieve information from not only a single store but a set of stores. The time to retrieve a set of Q stores will be determined by the time for the slowest store to respond. The mean of the maximum response time can be computed directly. To simplify these calculations we ignore the fact that some stores will not respond and assume that response time is identically and independently distributed as a gamma distribution with location and scale parameters of 0.5 and 5, respectively. Therefore, the mean and standard deviation of a single response is 2.5 and 3.5 seconds, respectively, which is similar to our sample in Table 4. The expected time for five stores to be queried is 6.9 seconds, while a set of 10 stores takes 9.5, and 30 stores would take more than 14 seconds. (These numbers are sample estimates using a simulation with 100,000 draws.) Note that as the number of stores increases the mean goes up, proportional to the logarithm. Therefore, a simple retrieval strategy of searching for all stores may not be a good one, because the benefits of retrieving an additional store need to be balanced against the expected benefits of retrieval.

5.3. Calibrating the Utility Model, Disutility of Waiting, and Cognitive Costs

In this subsection, we consider the calibration of the parameters associated with our utility model. Our purpose is to choose reasonable values that will be used for illustrative purposes in a simulation study that will be presented in §6. We do not claim that

Table 5 Parameter Estimates of a Multinomial Logit Choice Model for Consumer Purchases at Dealtime.com

Parameters	Estimates
Price	
Item price	-0.194 (0.001)
Shipping cost	-0.368 (0.002)
Expected days until delivery	-0.019 (0.001)
Effect of branded retailers	
Amazon.com	0.477 (0.020)
Barnes&Noble	0.177 (0.023)
Borders	0.266 (0.020)

Notes. The standard errors of the estimates are given in parentheses.
Source. Smith and Brynjolfsson (2001)

these settings are correct or even representative. Their purpose is to illustrate potential effects of different operational strategies on consumer utility.

The parameters for the utility function measure the implicit trade-offs consumers are willing to make when evaluating a product. These parameters can be estimated directly from previous purchases at the shopbot or through a conjoint task. We use the maximum likelihood estimates reported by Smith and Brynjolfsson (2001) and given in Table 5. They measured the utility associated with the following attributes about each book: total price, expected number of days until delivery, and an indicator for brand if it is sold by one of the three large booksellers (Amazon.com, Borders, and Barnes&Noble). Smith and Brynjolfsson (2001) use actual purchase data from a panel from more than 20,000 unique visitors during late 1999.

One method for interpreting these coefficients is to make relative comparisons. For example, these estimates imply that for every additional day that it takes to have a book delivered, the store needs to decrease its price by \$0.098 ($= -0.019 / -0.194$) to keep utility unchanged. To properly compensate a consumer for an extra two weeks in delivery time, the store would need to decrease the price by about \$1.37. Note that consumers are almost twice as sensitive to a dollar paid for shipping as they are to the item price (this implies it is better to charge higher prices and bundle shipping costs with the book price).

Table 4 Estimates of Gamma Distribution and Probability of Response for Selected Bookstores

Store	Number of observations	Probability of response	Gamma parameters		Moments	
			Location (α)	Scale (σ)	Mean	Std. Dev.
BCY Bookloft	7,803	0.979	0.452	5.61	2.53	3.77
Amazon.com	5,739	0.960	0.775	3.47	2.69	3.05
Barnes&Noble	2,224	0.950	0.443	5.94	2.63	3.95

Also, the value of Amazon.com's brand name can be imputed from the model parameters to be \$2.46 more ($=0.477/-0.194$) than the value of an less-well-known bookstore. All three major booksellers show substantial brand equity.

Our model postulates a certain amount of disutility due to waiting one second (ω), waiting for the shopbot server to launch a thread due to system congestion (ξ), and the cognitive costs associated with making comparisons in a final offer set (λ). Unfortunately, the data from Smith and Brynjolfsson (2001) do not include waiting times, nor do we have access to any data that would allow us to empirically answer this question. However, previous studies of Internet behavior indicate that there is disutility to waiting. Konana et al. (2000) conjecture that there is a direct trade-off between waiting time and costs. Dellaert and Kahn (1999) show experimentally that waiting can negatively affect evaluations of websites. Their results also suggest that waiting is not purely a function of time but can be mediated by other factors. However, for simplicity we have assumed that waiting is a simple trade-off between time and dollars. Johnson et al. (2002) argue that the more a website is used, the faster users can use the website in the future due to improved knowledge about the website's design. They estimate that after five visits to Amazon.com, the user reduces his or her time costs by almost \$1.50 per session (or about \$0.40/minute) versus having to learn a new online bookstore.

Our analytical framework permits arbitrary values for ω and ξ , but to illustrate our technique we choose plausible settings for the simulation in the following section. We assume that the value of time to a consumer is \$0.01/second. This translates into a yearly wage of about \$70,000. Therefore, every additional second of waiting diminishes utility by 0.002 utils ($=\$0.01/\text{sec} \times -0.194 \text{ util}/\$ \approx -0.002 \text{ util}/\text{sec}$), so we let $\xi = 0.002$. Suppose the overhead for launching an additional thread is 10 milliseconds; a corresponding value for ω is 0.00002. If the system is at high utilization, then launching and servicing an additional thread could take a substantial amount of time, and ω could be substantially higher during high utilization. Finally, we assume that a consumer can make one comparison per second, but attaches 10 times the value to a second of cognitive effort versus a second of waiting: $\lambda = 0.02$. For example, suppose we ask a consumer to evaluate a list of three items with four attributes; this would require $(3 - 1) \times (4 - 1) = 6$ comparisons and have an implicit cost of \$0.60 ($=0.02 \text{ util}/\text{comparison} \times 6 \text{ comparisons} \div -0.194 \text{ util}/\$$). In comparison, if the consumer was simply waiting six seconds, the implicit cost of time would be \$0.06.

6. Empirical Illustration of Optimal Shopbot Design

In this section, we consider a simulated example using the data presented in the previous section to show how the design of the shopbot can influence consumer utility. In these simulations we can solve the shopbot's operational problem empirically, which allows us to move beyond some of the assumptions that we were forced to make in §4 to derive analytical results. Specifically, we assume prices are normally distributed and not logistically distributed; we compute the order statistics of utility empirically and do not need to rely upon approximations; the offers are no longer assumed to be identically distributed; and the distribution of time to retrieve an offer is gamma distributed and not exponentially distributed. These simulations allow us to assess the probability that a consumer will prefer shopping at their favorite store rather than using the present or optimal shopbot designs. Furthermore, we perform a couple of simulations to assess the sensitivity of these results to our parameter settings.

The set of 28 online stores is given in Table 3. Each store will likely return several offers with different shipping terms. To construct an actual offer, the shopbot needs to search the online store, find the price, and add the shipping cost. For example, if the list price of a book is \$19.99, and the actual price at 1BookStreet is \$15.19, then the book could be delivered by USPS parcel post with a delivery range of 6–21 days (expected time to deliver is 13.5 days) for a total cost of \$15.19, or via UPS 2nd day with an expected delivery in nine days for a total cost of \$27.14. If just these two offers were presented to the consumer, the utility of these offers would follow an extreme value distribution with locations of -3.204 and -7.516 , respectively, and a common scale parameter of unity. The probability that a consumer would chose the first offer would be 98.7%. Clearly, the first alternative dominates the second and illustrates why not all offers need to be presented to the consumer.

We now consider the consumer's utility under three scenarios using the parameters specified in the previous section. The first scenario is that the shopbot searches all stores and presents all results (this is the current decision rule). The second scenario is that the shopbot knows prices with certainty (we assume that the price of the book is equal to its mean and the list price is \$19.99). The third scenario that we consider is the case where prices are not known with certainty, but instead the shopbot assumes a priori prices are normally distributed with the means and standard deviations given in Table 3. In both of the latter two scenarios we assume that the shopbot will select the optimal set of offerings to present to the consumer. To simplify calculations we

assume that all stores respond to a query (as opposed to a 95% probability that the store will respond). This is not a strong assumption, because all stores have similar probabilities of responding and the probability of no response is independent of the offer returned. Incorporating the probability that the store will not respond reduces the utility of all scenarios. Additionally, we assume that a priori we can determine the optimal order of stores to query by sorting on the expected utility (see the online appendix at mansci.pubs.informs.org/companion.html for a discussion). Finally, in our simulation if one offer from a store is retrieved, then we assume that all offers from that store are retrieved without any extra delay or cost, because delivery costs are deterministic.

The expected utility for the three scenarios is plotted in Figure 3 against the number of stores that are to be queried. The stores are ordered according to their best offer. For example, if only one store could be searched, then the shopbot would only check 1BookStreet. Presumably this would be the consumer's favorite store (denoted by \blacklozenge). Note that the utility for the current shopbot peaks after five stores (denoted by \blacksquare) and then starts to decline quickly due to the high cognitive effort placed on the consumer of comparing so many alternatives. In fact, if consumers were asked to choose between the current shopbot design that queries all stores and presents all offers (denoted by \bullet) and simply visiting their favorite store, we would expect that consumers would choose their favorite store with a probability of 90%. In contrast, the optimal shopbot designs are not penalized for querying a larger number of stores, because

they will only select the best offers to present to the consumer. The optimal designs under the assumption that prices are known or normally distributed (denoted by \blacktriangle or \blacktriangledown , respectively) searched 10 and 16 stores, respectively, and would be preferred by consumers over simply visiting their favorite store with a probability of 76% and 78%, respectively. Even if current shopbots were scaled back so that they searched fewer stores (but still selected the best stores), consumers would prefer the faster search and smaller decisions sets with a probability of 64%.

To help understand what offers would be presented, we list the optimal offer set in Table 6 from the optimal shopbot design when prices are assumed to be known (the solution that corresponds to \blacktriangledown). The 10 stores that were queried would yield 32 separate delivery options, but only 12 of these would be presented to the consumer. We are not simply identifying the cheapest bookstores, but those stores that yield the highest utility. For example, Amazon.com and Borders have higher prices but were included due to the brand equity of their store names; computer-library.com and 1BookStreet are included because of their free shipping policies.

If prices are not known with certainty, then it is better to search at a larger number of stores. Suppose prices are assumed to be normally distributed with the mean and standard deviation given in Table 3 (this solution is denoted by \blacktriangle in Figure 3). Our simulation shows that it is best to search at 16 stores, which would yield 55 possible offers. Because prices are not known with certainty, the offer set cannot be determined until after the prices are realized. For

Figure 3 Expected Utility Based on Number of Stores Queried Using Several Different Shopbot Designs




Table 6 Offers That Will Be Presented to the Consumer for the Scenario in Which Prices Are Known

Offer	Store	Delivery	Service (days)	Price (\$)	Shipping (\$)	Total (\$)
1	1BookStreet	USPS parcel post	6–21	15.19	0.00	15.19
21	Amazon.com	USPS priority mail	5–10	12.59	3.95	16.54
51	Buy.com	Standard shipping	N/A	10.39	3.95	14.34
48	Borders	Standard	5–10	12.39	3.90	16.29
26	Barnes&Noble	Standard ground	4–7	12.59	3.99	16.58
25	Barnes&Noble	USPS	5–9	12.59	3.95	16.54
42	Booksamillion.com	Standard ground	N/A	11.79	3.95	15.74
16	AlphaCraze.com	USPS special rate	5–15	12.79	3.50	16.29
64	Computerlibrary.com	N/A	N/A	19.79	0.00	19.79
69	HamiltonBook.com	USPS	N/A	13.99	3.00	16.99
17	AlphaCraze.com	USPS	N/A	12.79	3.95	16.74
70	Kingbooks.com	USPS book rate	16	14.59	2.50	17.09

Note. Prices assumed to equal their expected values.

example, the previous best offer (#1) would be included with a 95% probability, while the last offer (#70) would be included with a 38% probability. Additionally, there is some chance previously excluded offers, like Amazon.com’s second-day air, would now be included. (We report the full results in the online appendix.) Even though a larger number of offers may be potentially included, on average we would expect to see only nine or 10 offers presented to the consumer. The ability to select a smaller number of offers demonstrates an important reason why the optimal shopbot design performs so much better than the present shopbot design.

The parameter settings play an important role in determining the benefit of the improved shopbot. First, consider the case where ω , ξ , and λ are scaled by a factor of 10 to reflect that time is more valuable. If prices are unknown, the shopbot would search no more than five stores and most likely present only the best offer retrieved. Again, the high cost of time means that the shopbot needs to be much more intelligent in anticipating the tastes of the consumer. In contrast, if time is less valuable, being scaled by a factor of 0.1, then the current shopbot design performs more comparably to the optimal shopbot design. Finally, consider the case where a consumer is indifferent between expending time in a cognitively taxing activity (such as comparing results) and simply waiting ($\xi = \lambda = 0.002$). Under this assumption, the current shopbot design performs more comparably to the optimal shopbot design and the shopbot is more likely to be preferred by the consumer than simply visiting a favorite store. (A computer program is available upon request from the authors for readers interested in further assessing the sensitivity of these results to different parameter values.)

7. Discussion

Our model has provided several insights into how improved design could increase the value of shopbots

and their subsequent use by consumers. Specifically, shopbot design can be improved by selectively presenting and querying stores. Our empirical analysis shows that book prices at online stores can be predicted with a high degree of accuracy without having to query a store but instead relying upon past prices. Embedded within our framework is a compensatory utility model that aids the shopbot in understanding user preferences. This utility model allows the shopbot to predict the expected gains to the consumer from more search and to balance these gains against the cost of searching and presenting too much information.

In practice we find that despite high price variability online (Smith and Brynjolfsson 2001), most Internet shoppers continue to search in the traditional way, for example, visiting a single store (Johnson et al. 2003) or making price comparisons on their own. In Table 7 we report Internet usage information using Media Metrix’s home panel during a five-year period from July 1997 through May 2002. Although shopbot usage has grown, it is still low compared with other retail usage. From July 2001 through May 2002 less than 6% of Internet home users visited a shopbot, while 66% visited an online retailer and 28% visited an online bookstore. On the other hand, shopbots are effective in increasing the number of bookstores that consumers visit. During the same period online book shoppers (i.e., the user visited at least one bookstore during the month) visited 1.4 bookstores, while shopbot users visited 2.0 bookstores. Clearly, shopbots are helpful in leading to increased search.

It is our opinion that shopbots are still underutilized by consumers, although the trends in Table 7 show that shopbot usage has grown substantially. While there are many contributing explanations for shopbot growth, we would point out that during this time shopbots have made significant design changes that are consistent with our recommendations. Specifically, shopbots have improved their response times

Table 7 Various Measures of Home Internet Users' Behaviors Using Media Metrix Panel Data

Time period	Percentage of Internet users during a month who visit a . . .			Number of bookstores visited by visitor to a . . .		Probability of shopbot visit given visit last month	Probability of shopbot visit given visit last month
	Retailer	Bookstore	Shopbot	Bookstore	Shopbot		
July 1997–June 1998	39	9	0.1	1.2	1.5	3	6
July 1998–June 1999	48	15	0.7	1.2	1.9	10	20
July 1999–June 2000	59	23	3.6	1.3	2.0	22	20
July 2000–June 2001	62	26	5.5	1.4	2.1	29	31
July 2001–May 2002	66	28	5.7	1.4	2.0	26	24

Note. The values are computed by averaging monthly values. Note that the final period does not include June 2002, which was unavailable to the authors. Also, we define shopbot usage as a visit to Dealttime.com, Bottomdollar, Pricescan, or MySimon. Retailer refers to online resellers of consumer goods which include bookstores and shopbots, while bookstore refers only to retailers of books.

and refined their information displays. In May 2000 our queries to Dealttime.com resulted in an average response time of 89 seconds, and Pricescan had a 33-second response time. By March 2002 these response times dropped significantly to 18 and 16 seconds, respectively. Moreover, retention rates, given in the last two columns of Table 7 as conditional probabilities of repeat visits, have also increased, suggesting that consumers are also more satisfied. Although the increase in usage and retention is consistent with our prescriptions for improved shopbot design, we do not have information about marketing expenditures such as advertising and cannot control for their effects. We would certainly expect that increased advertising expenditures is another important contributing factor to increased shopbot usage. Finally, this analysis only measures a correlation and cannot prove causation; hence, we can only suggest that these trends are consistent with our prescriptions.

To further explore our proposal, an experimental test was constructed by Basartan (2001). She constructed a simulated shopbot in which response time and the number of offers displayed was varied for a sample of 190 students. First, the students completed a conjoint exercise to estimate their utility model (Green et al. 1981). Second, the students were asked to shop at several simulated shopbots and evaluate whether they would prefer shopping at the presented shopbot or Amazon.com. At one shopbot, students were given the top six alternatives after a 2-second wait. Seventy-seven percent of participants said that they would prefer shopping at this shopbot versus at Amazon.com. At another shopbot students were given the top 19 alternatives after a 30-second wait, and only 61% of participants said that they would prefer shopping at this shopbot versus at Amazon.com. (The difference in response is statistically significant at the 0.001 level.) This decrease in preference for shopbots corroborates our argument that shoppers may prefer visiting their favorite store if a shopbot provides too many alternatives and long waiting times.

8. Conclusions and Future Research

Our approach to solving the shopbot design problem has taken design elements from computer science and statistics and combined them with models of consumer behavior from economics and marketing. This research represents a cross-disciplinary approach that we believe is necessary in the emerging area of e-commerce research. We believe a dominant research theme in this area is to use models of consumer behavior to better improve the design of software and websites. At the same time there are many areas that we have only touched upon that need to be researched further. In this section, we discuss some of the issues related to implementation and consumer behavior.

Implementation

The primary challenge for a shopbot implementing the utility-based design that we have proposed is measuring user preferences, which we have assumed are known. Shopbots have two approaches to learn about consumers. First, the shopbot can learn from previous choices made by a visitor, if there are any. Most directly we would refer to Smith and Brynjolfsson's (2001) empirical work in estimating a multinomial logit model using Dealttime.com data. More advanced techniques for estimating individual-level multinomial probit models have also been demonstrated by Rossi et al. (1996). Perhaps even better would be to allow for heterogeneity within a user through time (Allenby et al. 1998), because users may be delivery sensitive for some books. Second, instead of relying solely upon past data, a shopbot could adaptively ask questions before conducting a search so as to learn about consumer utility using hybrid conjoint analysis (Green et al. 2001). An essential trade-off that must be considered is whether the added predictive ability is worth the effort to train the agent (Alba et al. 1997).

We have also ignored two related and potentially important issues—profit-maximizing behavior of the shopbot and the related strategic issue of cooperation between shopbots and online retailers (Iyer and

Pazgal 2003). Baye and Morgan (2001) consider the price equilibrium introduced by a shopbot or information gatekeeper that charges consumers for access to price information. Greenwald and Kephart (1999) study the effectiveness of various pricing algorithms that can be used by pricebots (adaptive agents that automatically set prices) in marketplaces with significant shopbot presence. It is possible that the decision environment created by shopbots could lead to new competitive retail structures, because shopbots earn revenue not only from choice, but also from consideration (e.g., shopbots can be paid if a visitor clicks on a link to a retailer).

We have not explicitly modeled the shopbot profit function, but instead focused upon one of its input components, consumer utility. It is straightforward to argue using our results that if consumer utility can be increased, *ceteris paribus* then shopbot profits will increase. Shopbot profitability differs from conventional retailer profitability in that shopbots can earn revenue from advertising, referrals, and preferred placement in a list of offers. This means that shopbots need to balance the increased revenue they can earn from giving some high-priced retailers priority in the listings against the potential backlash by consumers. These differences make shopbot profitability an interesting area for further study.

Effects of Consumer Behavior on Shopbot Design

Research has found that the context in which a choice is evaluated can impact its likelihood of being chosen. Degeratu et al. (2000) show that ordering lists by prices can lead consumers to become more price sensitive. Lynch and Ariely (2000) show how quality information can mediate the greater price sensitivity that may result from making price comparisons easier. Simonson (1999) surveys how product assortment can influence buyer preferences and choices. For example, including a higher-quality version of a product can increase the chance that a lower-quality item will be purchased. One possibility is that a consumer's preferences are not fixed but instead constructed during the choice task, because of limited processing capacity (Bettman et al. 1998). These studies point out that the set of alternatives and the context in which they are presented by the shopbot could themselves influence user preferences.

We also know that when consumers are confronted with demanding cognitive tasks they may use heuristics (for further discussion see Payne et al. 1993). For example, to find the best product in a long list of alternatives consumers may use an elimination by aspect strategy (Tversky 1972) to help reduce cognitive effort. Directly modeling these processes could result in better predictions.

Furthermore, an improved understanding of how consumers perceive waiting time is needed. We have

assumed a simple framework in which disutility from waiting is proportional to the time spent waiting. However, filler tasks could be performed that could alter consumers' perceptions of the time spent waiting. These filler tasks could be used to actively collect information related to the query or could be totally unrelated and simply occupy the user while the search is proceeding. All of these comments point to the fact that further study in integrating consumer behavior models to improve shopbot design could be quite fruitful. We believe our research is only a first step in this direction.

An online appendix to this paper is available at mansci.pubs.informs.org/ecompanion.html.

Acknowledgments

This material is based on work supported by the National Science Foundation under Grant No. 0118767. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- Alba, J., J. Lynch, B. Weitz, C. Janiszewski, R. Lutz, A. Sawyer, S. Wood. 1997. Interactive home shopping: Consumer, retailer, and manufacturer incentives to participate in electronic markets. *J. Marketing* 61(3) 38–53.
- Allenby, G. M., N. Arora, J. L. Ginter. 1998. On the heterogeneity of demand. *J. Marketing Res.* 35(3) 384–389.
- Bakos, Y. 1997. Reducing buyer search costs: Implications for electronic marketplaces. *Management Sci.* 43(12) 1676–1692.
- Balakrishnan, N., ed. 1992. *Handbook of the Logistic Distribution*. Dekker, New York.
- Basartan, Y. 2001. Amazon versus the shopbot: An experiment about how to improve the shopbots. Ph.D. summer paper, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA.
- Baye, M., J. Morgan. 2001. Information gatekeepers on the Internet and the competitiveness of homogeneous product markets. *Amer. Econom. Rev.* 91 454–474.
- Ben-Akiva, M., S. R. Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press, Cambridge, MA.
- Bettman, J. R., E. J. Johnson, J. W. Payne. 1990. A componential analysis of cognitive effort in choice. *Organ. Behavior Human Performance* 45(1) 111–139.
- Bettman, J. R., M. F. Luce, J. W. Payne. 1998. Constructive consumer choice processes. *J. Consumer Res.* 25(December) 187–217.
- Chavez, A., P. Maes. 1996. Kasbah: An agent marketplace for buying and selling goods. *Proc. First Internat. Conf. Practical Application Intelligent Agents Multi-Agent Tech.*, London, U.K.
- Clark, C. E. 1961. The greatest of a finite set of random variables. *Oper. Res.* 9(2) 145–162.
- Clay, K. B., R. Krishnan, E. Wolff. 2001. Prices and price dispersion on the Web: Evidence from the online book industry. *J. Indust. Econom.* 49(4) 521–539.
- David, H. A. 1981. *Order Statistics*, 2nd ed. John Wiley and Sons, New York.
- Degeratu, A., A. Rangaswamy, J. Wu. 2000. Consumer choice behavior in online and traditional supermarkets: The effects of brand name, price, and other search attributes. *Internat. J. Res. Marketing* 17(1) 55–78.

- Dellaert, B. G. C., B. E. Kahn. 1999. How tolerable is delay: Consumers' evaluations of Internet web sites after waiting. *J. Interactive Marketing* 13(1) 41–54.
- Doorenbos, B., O. Etzioni, D. Weld. 1997. A scalable comparison-shopping agent for the World Wide Web. *First International Conference on Autonomous Agents*. Marina del Rey, CA.
- Feinberg, F. M., J. Huber. 1996. A theory of cutoff formation under imperfect information. *Management Sci.* 42(1) 65–84.
- Green, P. E., S. M. Goldberg, M. Montemayor. 1981. A hybrid utility estimation model for conjoint analysis. *J. Marketing* 45(Winter) 33–41.
- Green, P. E., A. M. Krieger, Y. Wind. 2001. Thirty years of conjoint analysis: Reflections and prospects. *Interfaces* 31(3) S56–S73.
- Greenwald, A., J. O. Kephart. 1999. Shopbots and pricebots. *Proc. Sixteenth Internat. Joint Conf. Artificial Intelligence*. Stockholm, Sweden.
- Häubl, G., V. Trifts. 2000. Consumer decision making in online shopping environments: The effects of interactive decision aids. *Marketing Sci.* 19(1) 4–21.
- Hauser, J. R., B. Wernerfelt. 1990. An evaluation cost model of consideration sets. *J. Consumer Res.* 16(4) 393–408.
- Hausman, J., B. H. Hall, Z. Griliches. 1984. Econometric models for count data with an application to the patents—R&D relationship. *Econometrica* 52(4) 909–938.
- Hoch, S. J., D. A. Schkade. 1996. A psychological approach to decision support systems. *Management Sci.* 42(1) 51–64.
- Hoque, A. Y., G. L. Lohse. 1999. An information search cost perspective for designing interfaces for electronic commerce. *J. Marketing Res.* 36(3) 387–394.
- Iyer, G., A. Pazgal. 2003. Internet shopping agents: Virtual collocation and competition. *Marketing Sci.* 22(1) 85–106.
- Jacoby, J. 1984. Perspectives on information overload. *J. Consumer Res.* 10(March) 432–435.
- Jacoby, J., D. E. Speller, C. A. Kohn. 1974. Brand choice behavior as a function of information load. *J. Marketing Res.* 11(February) 63–69.
- Jennings, N. R., K. Sycara, M. Wooldridge. 1998. A roadmap of agent research and development. *J. Autonomous Agents Multi-Agent Systems* 1(1) 7–38.
- Johnson, E. J., J. W. Payne. 1985. Effort and accuracy in choice. *Management Sci.* 31 394–414.
- Johnson, E. J., S. Bellman, G. Lohse. 2002. Cognitive lock in and the power law of practice. *J. Marketing* 67(2) 62–75.
- Johnson, E. J., W. Moe, P. Fader, S. Bellman, G. Lohse. 2003. On the depth and dynamics of world wide Web shopping behavior. *Management Sci.* Forthcoming.
- Keller, K. L., R. Staelin. 1987. Effects of quality and quantity of information on decision effectiveness. *J. Consumer Res.* 14(September) 200–213.
- Konana, P., A. Gupta, A. B. Whinston. 2000. Integrating user preferences and real-time workload in information services. *Inform. Systems Res.* 11(2) 177–196.
- Lynch, J. G., D. Ariely. 2000. Wine online: Search costs affect competition on price, quality, and distribution. *Marketing Sci.* 19(1) 83–103.
- McFadden, D. 1980. Econometric models of probabilistic choice among products. *J. Bus.* 53 13–26.
- Moorthy, S., B. T. Ratchford, D. Talukdar. 1997. Consumer information search revisited: Theory and empirical analysis. *J. Consumer Res.* 23(4) 263–277.
- Mudholkar, G. S., E. O. George. 1978. A remark on the shape of the logistic distribution. *Biometrika* 65 667–668.
- Nielsen, J. 2000. *Designing Web Usability*. New Riders Publishing, Indianapolis, IN.
- Payne, J. W., J. R. Bettman, E. J. Johnson. 1993. *The Adaptive Decision Maker*. Cambridge University Press, Cambridge, MA.
- Pham, V. A., A. Karmouch. 1998. Mobile software agents: An overview. *IEEE Comm.* 36(7).
- Poole, D. 1997. The independent choice logic for modelling multiple agents under uncertainty. *Artificial Intelligence* 94(1–2) 7–56.
- Roberts, J. H., J. M. Lattin. 1991. Development and testing of a model of consideration set composition. *J. Marketing Res.* 28(4) 429–440.
- Roberts, J. H., J. M. Lattin. 1997. Consideration: Review of research and prospects for future insights. *J. Marketing Res.* 34(3) 406–410.
- Rossi, P. E., R. E. McCulloch, G. M. Allenby. 1996. The value of purchase history data in target marketing. *Marketing Sci.* 15(4) 321–340.
- Shugan, S. M. 1980. The cost of thinking. *J. Consumer Res.* 7 99–111.
- Siddarth, S., R. E. Bucklin, D. G. Morrison. 1995. Making the cut: Modeling and analyzing choice set restriction in scanner panel data. *J. Marketing Res.* 32(3) 255–266.
- Simonson, I. 1999. The effect of product assortment on buyer preferences. *J. Retailing* 75(3) 347–370.
- Smith, M. D., E. Brynjolfsson. 2001. Consumer decision-making at an Internet shopbot: Brand still matters. *J. Indust. Econom.* 99(4) 541–558.
- Sproule, S., N. Archer. 2000. A buyer behavior framework for the development and design of software agents in e-commerce. *Internet Res. Electronic Networking Appl. Policy* 10(5) 396–405.
- Stigler, G. 1961. The economics of information. *J. Political Econom.* 69(June) 213–225.
- Tversky, A. 1972. Elimination by aspects: A theory of choice. *Psych. Rev.* 84 327–352.
- Udo, G. J., G. P. Marquis. 2001. Factors affecting e-commerce web site effectiveness. *J. Computer Inform. Systems* 42(2) 10–16.
- Widing II, R. E., W. Wayne Talarzyk. 1993. Electronic information systems for consumers: An evaluation of computer-assisted formats in multiple decision environments. *J. Marketing Res.* 30 125–141.