Print Date: 28 Dec 2007

# The Great Equalizer?
# An Empirical Study of Consumer Choice at a Shopbot

Erik Brynjolfsson, Michael D. Smith, and Alan L. Montgomery

*December 2007*

**Working draft.**

Erik Brynjolfsson (erikb@mit.edu) is the George and Sandi Schussel Professor in Management Science and Co-Director Center for eBusiness Center at MIT, Sloan School of Management, Massachusetts Institute of Technology, 50 Memorial Drive, Cambridge, MA 02142, Michael D. Smith (mds@cmu.edu) is an Assistant Professor of Information Technology at the Heinz School and Alan L. Montgomery (alan.montgomery@cmu.edu) is an Associate Professor of Marketing at the Tepper School of Business, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA, 15213. We thank Dealtime.com, now known as Shopping.com, for generously collecting the data necessary to conduct this research and for providing valuable insight into the shopbot market. We also acknowledge financial support from the Marketing Science Institute and (Smith) from the National Science Foundation through grant Award IIS-0448516.

# The Great Equalizer?
# An Empirical Study of Consumer Choice at a Shopbot

**Abstract:**

Shopbots are computer agents that aid consumers by comparing prices across online stores. A consumer visits a shopbot web site, inputs a product to search for, such as a book, and then the shopbot automatically queries available online stores, and tabulates and presents the results to the consumer. Consumers consider a subset of these stores by clicking upon a hyperlink displayed with the offer. This evoked consideration set is important to shopbot owners because shopbot owners are generally compensated directly for consideration as opposed to purchase, in the industry this is referred to as pay-for-click.

In our research we estimate a multivariate probit model to predict a consumer's evoked consideration set. Our utility model includes the product attributes, retailer attributes, and the position of the alternative in the table shown to the consumer. In addition, we include covariates for learning over time, the coefficient of variation of the prices, and the list price of the product to explain consumer-specific variation in the size of the evoked consideration set. Our model is calibrated on clickstream data collected at a major shopbot over a 23-month period. A weakness of the traditional multivariate probit model is that the marginal distribution of the number of items in the evoked consideration set is a function of the latent utility which predicts which offers will be chosen. Unfortunately, the standard model provides a poor approximation to this distribution, so we propose a model for the number of items that will be selected and then conditional upon this value we predict which offers will be chosen. Our findings indicate this model better captures consumer choice.

Additionally we find that positioning and advertising have substantial impact on which items a consumer will choose. The size of the evoked consideration decreases with increasing consumer experience, larger offer sets leads to less search, and more expensive products and offer sets with more variation in prices result in larger evoked consideration sets. These results suggest that shopbot owners can influence the composition of a consumer's evoked consideration set — and therefore the shopbot's revenue — by changing the order of offers, the number of offers shown, and the display of logos next to a particular retailer's name.

**Keywords:** Shopbots, Multivariate Probit Models, Clickstream Data, Computer agents, Internet Marketing, Utility theory

## 1. Introduction

The growth of electronic commerce has lowered consumers' costs associated with searching for price and product information. One type of web site that can dramatically aid consumers in searching across many stores are shopbots. Shopbots are Internet-based services that provide one-click access to price and product information from numerous competing retailers. However, they may also strip away many of the accoutrements of a retailer's brand name by listing only summary information from both well- and lesser-known retailers. Further, every retailer at a shopbot is "one click away", reducing switching costs accordingly. In each instance these factors should serve to increase competition and reduce retailer margins in markets served by shopbots — an effect that should be felt most strongly for homogeneous physical goods (cf., Bakos 1997). These features have led some to call the Internet "The Great Equalizer" since traditional marketing techniques to increase switching costs and take advantage of geographic location are dramatically reduced, especially by shopbots.

An added benefit for managers and researchers of shopbots is the ability to directly observe search using clickstream data. Consumer search costs have long been recognized as a crucial component in understanding consumer behavior (Stigler 1961). However, direct observation of search can be intrusive and costly. In an online context we can observe not only what consumers buy but also the consideration set of products viewed before the purchase. In a physical store this is analogous to an analyst standing behind the consumer and recording their movements in the store and which products they observe. In the past researchers have had to rely upon complex statistical procedures to make inferences about consideration sets (Mehta, Rajiv, and Srinivasan 2003). That is not to say this is a perfect source of information, since consumers may internally evaluate products from memory or the external environment that are not included within clickstream data. Therefore it is more appropriate to say that we observe the evoked consideration set.

Another unusual element of shopbots is their profit structure, which allows them to earn revenue from commissions, referrals, and advertising. Notice that revenue from referrals and advertising imply that the shopbot earns money if they can simply encourage a consumer to consider a store. Traditional retailers require consumers to actually complete a purchase before they earn revenue. Hence, the search process by which consumers select the offers that they will consider but not necessarily purchase become central to a model.

To address this question we study consumer choice at Dealtime, a leading Internet shopbot. Specifically, we observe the complete set of offers presented to each consumer and which offers were reviewed over a two-year period. We develop a multivariate probit model to predict which offers a consumer will review which we refer to as a consumer's evoked consideration set. This contrasts with traditional choice models that focus on the single best choice. Our utility model includes the product attributes, retailer attributes, and the position of the alternative in the table shown to the consumer. In addition, we include covariates for learning over time, the coefficient of variation of the prices, and the list price of the product to explain consumer-specific variation in the size of the evoked consideration set.

An interesting aspect of our data is that the modal value for the number of items in the evoked consideration set is one. However, consumers may select none or all of the offers presented. We find that implied marginal distribution for the number of items that will be selected from the traditional multivariate probit model does not well capture our dataset. To overcome this weakness we propose a new multivariate choice model that first models the number of items that will be chosen. Conditional upon the number of items selected we model which items will be selected. Our model assumes that there is a latent random variable which follows a multivariate normal distribution, and the offers with the highest latent values will be selected. If only one offer is selected then our model is equivalent to the multinomial probit model.

Our approach to analyzing electronic markets differs from recent empirical studies in that it examines the responses of actual consumers to prices set by retailers, not just the retailers' pricing behavior. Research analyzing retailer pricing strategies has been used to characterize the relative efficiency of electronic and physical markets (Bailey 1998; Brynjolfsson and Smith 2000), retailer differentiation strategies (Clay, Krishnan, Wolff, Fernandes 1999), and price discrimination strategies (Clemons, Hann, and Hitt 1998). However, retailer pricing strategies provide only second-order evidence of consumer behavior in electronic markets.

Our findings suggest that a retailer's brand name, on-site advertising, and positioning in the table have a significant impact on the products included in the evoked consideration set. Moreover, the size of the evoked consideration decreases with increasing consumer experience, and more expensive products and consideration sets with more variation in prices result in larger evoked consideration sets. These results suggest that shopbot owners can influence the composition of a consumer's evoked consideration set — and therefore the shopbot's revenue — by changing the order of offers, the number of offers shown, and the display of logos next to a particular retailer's name.

The remainder of this paper is organized in five parts. §2 considers past academic research that could be applied to consumer search at shopbots. §3 addresses the data we collect how it was collected and its strengths and limitations. §4 presents the empirical models we use to analyze our data. §5 presents our results. §6 concludes, discusses implications of our results, and areas for future research.

## 2. Prior Research related to Consumer Search at Shopbots

This paper relates to the academic literature on the impact of shopbots on electronic markets (see Smith (2002) for a review of this literature). Within this literature, our paper is most closely related to Smith and Brynjolfsson (2001) who use the same dataset used in this paper to analyze how shopbot customers respond to brand and partitioned prices at an Internet shopbot. The paper finds that shopbot consumers are willing to pay $1.72 more to purchase from heavily branded retailers than other retailers, and that consumers are more sensitive to changes in shipping price and tax than to changes in item price. Brown and Goolsbee (2000) use survey data and observed prices to show that the introduction of shopbots for consumer term life insurance led to a decrease in prices for these policies, and that initially life insurance led to an initial increase in price dispersion for these policies, which fell as use of life insurance shopbots spread. Ellison and Ellison (2001) use data from an Internet shopbot for computer chips to show that retailers can manipulate shopbot listings to increase consumer search costs. Finally Baye et al (2004) use data from Kelkoo.com on PDAs to show that consumer demand is a function of both the number of competing sellers at the shopbot and the firm's location on the screen relative ranking in the list of prices.

More generally, in the context of consumer search online, this paper relates the literature on consumer search. While consumer search has been studied widely in analytic models, historically consumer search has been difficult to test empirically. Recently, however, Sorensen (2000) finds that frequently purchased prescriptions show lower dispersion and price-cost margins than other prescriptions, which is consistent with the predictions of standard search cost models. Sorensen (2001) uses a structural model to show that search intensity is generally low, but is higher for maintenance medications than other types of medications. In the context of Internet markets, there is also a growing literature using consumer behavior to directly measure consumer search costs in online markets. Johnson et al. (2002) use Media Metrix data to analyze consumer search across different sites. They find that the amount of time consumers spend on web sites declines as consumers gain more experience — and the sites with the steepest declines are those that have the highest consumer loyalty. Bajari and Hortacsu (2003) quantify the implied cost of entering an eBay auction to be $3.20. Hann and Terwiesch (2003) use consumer data from a reverse auction site to show that consumers' perceived cost of rebidding is between 3.54 and 6.08 EUR. Finally, Brynjolfsson et al. (2005) show that consumer benefit from evaluating additional offers at an Internet shopbot are $6.55, while their cost of evaluating these offers is a maximum of $6.45.

There is also a related literature analyzing how information environments impact consumer choice. Within this literature, Lynch and Ariely (2000) use an experiment to show that lower search costs for quality (price) information reduces (increases) consumer price sensitivity. Similarly, Diehl et al (2003) find that when heterogeneous options are ordered according to consumer's quality preferences, consumers tend to select lower priced offers than they would in an environment where offers are presented without regard to fit or price.

## 3. Data

The data used in our study comes from the Dealtime[1]. We focus on consumer search for one class of products, namely books. Although Dealtime offers searches of many categories, such as computers, electronics, office supplies, and toys. An advantage of considering only books is that it is a homogeneous category with known and easily comparable features. Being the largest and leading shopbot Dealtime affords us the ability to have a large panel of consumer choice. Although having data from only one shopbot limits our ability to generalize to other sites. Additionally, one would expect that shopbot users are not representative of Internet shoppers, but are most likely to be price conscious.

Dealtime operates similarly too many other Internet shopbots. A consumer who wants to purchase a book would visit the site, search for the book's title or author, identifying a unique ISBN[2] as the basis for their search. Using these ISBN Dealtime then queries up to 60 distinct book retailers checking to see if they have the book in stock and their price and delivery times. The prices and delivery times are queried in real-time and thus represent the most up-to-date data from the retailer. Because the prices are gathered directly from the retailers, they are the same prices that are charged to consumers who visit the retailer site directly. (Some online retailers may provide price tables directly to a shopbot, which allows the shopbot to retrieve the offer from locally cached databases.)

---

[1] Dealtime was originally introduced in 1996 as acses.com, and later renamed dealpilot.com. In 1999 the site was renamed EvenBetter.com under new management. Most recently, it has been named shopping.com.

[2] International Standard Book Numbers (ISBNs) uniquely identify the individual version of the book (e.g., binding type, printing, and language). Because EvenBetter's search results are based on a single ISBN, all of the products returned in response to a search are physically identical.

The offers are then displayed in a table that is presented to the consumer (see Figure 1). These tables list the total price for the book and the elements of price (item price, shipping cost, and applicable sales taxes) along with the retailer's name and the book's delivery information. If a retailer provides multiple shipping options at multiple prices (e.g., express, priority, and book rate) the table lists separate offers for each shipping option. The tabular format of the offers enables easy comparison (Morwitz, Greenleaf, and Johnson 1998).

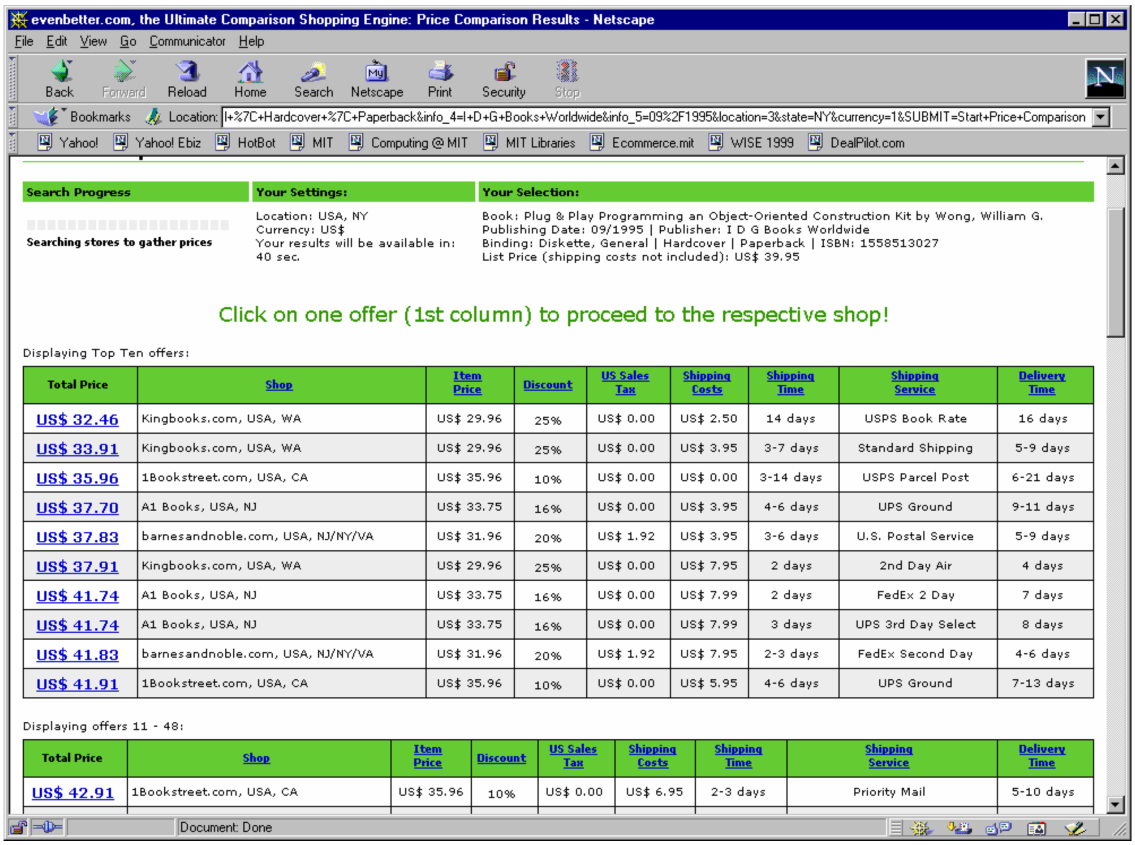
evenbetter.com, the Ultimate Comparison Shopping Engine: Price Comparison Results - Netscape

| Search Progress | Your Settings: | Your Selection: |
|---|---|---|
| Searching stores to gather prices | Location: USA, NY<br>Currency: US$<br>Your results will be available in: 40 sec. | Book: Plug & Play Programming an Object-Oriented Construction Kit by Wong, William G.<br>Publishing Date: 09/1995 \| Publisher: I D G Books Worldwide<br>Binding: Diskette, General \| Hardcover \| Paperback \| ISBN: 1558513027<br>List Price (shipping costs not included): US$ 39.95 |

Click on one offer (1st column) to proceed to the respective shop!

Displaying Top Ten offers:

| Total Price | Shop | Item Price | Discount | US Sales Tax | Shipping Costs | Shipping Time | Shipping Service | Delivery Time |
|---|---|---|---|---|---|---|---|---|
| US$ 32.46 | Kingbooks.com, USA, WA | US$ 29.96 | 25% | US$ 0.00 | US$ 2.50 | 14 days | USPS Book Rate | 16 days |
| US$ 33.91 | Kingbooks.com, USA, WA | US$ 29.96 | 25% | US$ 0.00 | US$ 3.95 | 3-7 days | Standard Shipping | 5-9 days |
| US$ 35.96 | 1Bookstreet.com, USA, CA | US$ 35.96 | 10% | US$ 0.00 | US$ 0.00 | 3-14 days | USPS Parcel Post | 6-21 days |
| US$ 37.70 | A1 Books, USA, NJ | US$ 33.75 | 16% | US$ 0.00 | US$ 3.95 | 4-6 days | UPS Ground | 9-11 days |
| US$ 37.83 | barnesandnoble.com, USA, NJ/NY/VA | US$ 31.96 | 20% | US$ 1.92 | US$ 3.95 | 3-6 days | U.S. Postal Service | 5-9 days |
| US$ 37.91 | Kingbooks.com, USA, WA | US$ 29.96 | 25% | US$ 0.00 | US$ 7.95 | 2 days | 2nd Day Air | 4 days |
| US$ 41.74 | A1 Books, USA, NJ | US$ 33.75 | 16% | US$ 0.00 | US$ 7.99 | 2 days | FedEx 2 Day | 7 days |
| US$ 41.74 | A1 Books, USA, NJ | US$ 33.75 | 16% | US$ 0.00 | US$ 7.99 | 3 days | UPS 3rd Day Select | 8 days |
| US$ 41.83 | barnesandnoble.com, USA, NJ/NY/VA | US$ 31.96 | 20% | US$ 1.92 | US$ 7.95 | 2-3 days | FedEx Second Day | 4-6 days |
| US$ 41.91 | 1Bookstreet.com, USA, CA | US$ 35.96 | 10% | US$ 0.00 | US$ 5.95 | 4-6 days | UPS Ground | 7-13 days |

Displaying offers 11 - 48:

| Total Price | Shop | Item Price | Discount | US Sales Tax | Shipping Costs | Shipping Time | Shipping Service | Delivery Time |
|---|---|---|---|---|---|---|---|---|
| US$ 42.91 | 1Bookstreet.com, USA, CA | US$ 35.96 | 10% | US$ 0.00 | US$ 6.95 | 2-3 days | Priority Mail | 5-10 days |

Document: Done

**Figure 1.** Dealtime web page presented in October 2000 from a typical book search.

This table also illustrates some of the decisions that the shopbot must make that can alter a consumer's perception of the choices. By default Dealtime sorted offers by total price until 2000 and afterwards by the presence of a logo (which denotes sponsorship) with price as a secondary key. However, the consumer can re-sort the table by selecting any of the columns. Each offer has a hyperlink that will take the consumer directly to the product page on the corresponding store. Dealtime can monitor which offers are selected using redirects. Notice that to complete the purchase the consumer makes the purchase using the store's regular ordering system, which is carried out independently from Dealtime's system. We do not have information about what item (if any) was purchased; we only observe information about which offers were considered when the user selects the associated hyperlink.

Dealtime does not charge a consumer directly to access its system, but earns revenue indirectly from each search. First the shopbot can offer advertising space to the queried stores. For example, a store could place a logo or display a graphical hyperlink that would offset it from the other stores. These advertisements would be similar in function to in-store retail displays. Alternatively, instead of embedding the advertisements within the offer list, the shopbot can also sell banner advertisements. Additionally, the shopbot can sell priority placement in the table of offers, for example a retailer could pay a premium to guarantee that they will appear towards the top of this list. Finally, a shopbot can earn a commission from the retailer if a consumer decides to actually purchase, for example the shopbot may be an affiliate for a bookstore which may not pay the shopbot at all for listings.

The shopbot is also able to make design decisions that can influence the choices made by consumers. First, the shopbot chooses what stores to query. Second the shopbot can choose what attributes to display. For example the shopbot can decide whether to provide or omit attributes such as store rating, shipping price, expected delivery time, and availability. Effectively this gives the shopbot the ability to alter the packaging of a product, which is quite different from the role of traditional retailers. Finally, the shopbot can alter the order offers are presented to the consumers. A key difference between shopbots and traditional retail store design is that the shopbot earns revenue by simply encouraging consumers to consider purchasing an item. In contrast, a traditional brick and mortar store only earns revenue if a consumer actually completes a purchase.

### 3.1. Dataset Characteristics

We receive three types of information from Dealtime, which are summarized in Table 1. The first is the complete set of attributes presented for each offer made to the user and a user identifier. Notice from Table 1 that Dealtime separates price into three components: item cost, shipping cost, tax, and total cost (the sum of item cost, shipping cost, and tax). Second Dealtime tracks whether a user has visited the store through the use of a redirection hyperlink. Ordinarily a server will not know the next page that a user views unless the page requested resides on the same server. To be able to monitor which link is selected Dealtime inserts a hyperlink that resides on its server, but will immediately redirect the user's browser to the appropriate link at the online store.

The third type of information collected by Dealtime is a user identifier. This identifier is stored using a cookie. Whenever a user requests a page from Dealtime's server, the user's browser will report the value of any previous written cookies generated by Dealtime (unless it is configured not to relay this information), and if no cookie exists then Dealtime's server will assign a new, unique identifier to the user. This allows Dealtime to identify a user across time and even if the user's IP address changes. Since identification is automatic this overcomes a major limitation of brick and mortar stores that must require their shopper to use a loyalty card to track identity. Also, identity is known even if no purchase is made, unlike a loyalty card.

At the same time there are several limitations when using cookies to identify consumers. First, users may have more than one computer, and thus more than one cookie. Second, some computers (e.g., a computer in a University computer lab available to all students) may be shared by more than one user (while having a single cookie number). Third, consumers may periodically destroy their cookies, making it difficult to track behavior from cookie to cookie. Finally, cookies cannot be used to track behavior at other web sites.

| | |
|---|---|
| *User Data (each row is associated with a user id), $\boldsymbol{z}_i$* | |
| Commercial ISP | Identifies whether the IP is in a major ISP (aol.com, msn.com, etc.) |
| Foreign | If the IP address is registered to US or foreign domain |
| Top Level Domain | Type of user (edu, com, net/org, gov/mil/us) |
| *Book Data (each row is associated with an ISBN), $\boldsymbol{w}_{it}$* | |
| List Price | Publisher's List Price |
| Binding | Book binding (hardcover, paperback, or other—such as an audio cassette) |
| Category | Topic of book (art, business, computer, fiction, juvenile, etc.) |
| *Session Data (each row associated with a user id and session id), $w_{it}$* | |
| Number of offers | Number of offers presented |
| Number of attributes | Number of attributes presented |
| Time Since Last Search | The date of all searches are recorded |
| Time of Day | Time of day for search (day or evening) |
| Weekend | Day of week for search (weekday or weekend) |
| Cumulative visits | Cumulative number of visits |
| Sort Column | What column was used to sort the data (price, logo, and other) |
| New Release | If the publisher's date of book searched is within a year of the session |
| CurrentBestseller | Was the book on the New York Times bestseller list at the time of the search |
| PastBestseller | Was the book ever on the New York Times bestseller list |
| *Offer Data (each row is associated with a user and session id), $\boldsymbol{x}_{ijtk}$* | |
| Total Price | Total price for the offer (item price plus sales tax plus shipping cost) |
| Item Price | The price for the item |
| Shipping Cost | The price for shipping |
| State Sales Tax | Sales tax (if applicable) |
| No Tax | =1 if there is no sales tax on the offer |
| Retailer | Retailer Name (Amazon, Barnes & Noble, Borders, etc.) |
| Shipping Time | Time to ship product from retailer to consumer (Min, Max, Average) |
| Acquisition Time | Time for retailer to acquire product (Min, Max, Average) |
| Delivery Time | Shipping time plus acquisition time (Min, Max, Average) |
| Shipping Method | Priority (1-day or 2-day), Standard (3-7 day), Book Rate (>7 day) |
| Delivery NA | =1 if retailer can't quote acquisition time on book |
| Position | The rank of the offer in the comparison table |
| Logo | Logo displayed with retailer name (none, bold type, or graphic) |
| Click-Through | =1 if the consumer clicked on this offer |

**Table 1.** Description of variables in our dataset.

We augment the Dealtime data with several additional datasets. The first is data about each book purchased, using data from Bowker's *Books in Print*. Specifically, for each book we know the type of book: hardcover or paperback, category of the book, and publisher's list price. Second, we have the books that were listed on the New York Times bestseller list. Finally, we executed reverse domain lookups for each IP address to determine whether the address was associated with a commercial online service or ISP, such as aol.com, msn.com, compuserve.com, etc. This may serve as a proxy for connection speed and location of the user, for example we expect that AOL users are more likely to be at home using a dial-up connection, as opposed to at work using a high-speed connection.

## 4. Modeling the Evoked Consideration Set

Our problem is to predict which offers a consumer will consider from a set of alternatives, we refer to this as the evoked consideration set. If a consumer were always to choose a single item our problem would be the usual choice problem of choosing one alternative from a set of alternatives. This choice problem has been extensively explored using the popular multinomial logit model (Guadagni and Little 1981 and Ben-Akiva and Lerman 1985) or multinomial probit model (Rossi et al. 1996). However, in our problem the user can select multiple offers or even none at all. Again in our problem choice does not equate with purchase, clicking on a shopbot offer may mean the consumer is gathering information before making a purchase decision. Moreover, choice or click-through is interesting not just as an indicator of consumer behavior, but because it directly impacts the shopbot's revenue. Many shopbots are paid a referral fee if a user clicks.

We observe an indicator variable ($y_{itk}$) that is set when consumer $i$ selects an offer $j$ during session $t$, where $i$=1,…,$K$, $t$=1,…,$M_i$, and $j$=1,…,$N_{it}$, (For notational convenience we drop the subscripts on $N$ and $M$, but it should be understood that these dimensions depend upon the session and user, respectively.) We assume that there is a random latent utility ($u$) associated with each offer and when this utility exceeds a certain threshold the offer is selected:

$$y_{itj} = \begin{cases} 1 & \text{if } u_{itj} \geq \lambda_{itj} \\ 0 & \text{otherwise} \end{cases}, \tag{1}$$

where $\lambda_{itj}$ is a threshold that determines if the offer is selected.

The latent utility ($u$) associated with each offer follows a linear model with random coefficients:

$$u_{itj} = \boldsymbol{\beta}_{it}'\mathbf{x}_{itj} + \varepsilon_{itj} \, . \tag{2}$$

The $\mathbf{x}$ vector captures the observable characteristics of each offer described in Table 1, such as price, shipping terms, and positional information. $\varepsilon_{itj}$ represents a random component which we assume follows a spatial model, specifically a first-order autoregressive-moving average or ARMA(1,1) model:

$$\varepsilon_{itj} = \phi_{it}\varepsilon_{it,j-1} + \xi_{itj} - \psi_{it}\xi_{it,j-1}, \quad \xi_{itj} \sim \mathrm{N}\left(0, \frac{1-\phi_{it}^{\,2}}{1+\psi_{it}^{\,2} - 2\phi_{it}\psi_{it}}\right), \;\; \varepsilon_{it0} \sim \mathrm{N}\left(0, \frac{(\psi_{it} - \phi_{it})^2}{1+\psi_{it}^{\,2} - 2\phi_{it}\psi_{it}}\right). \tag{3}$$

There are two special cases of this general formulation that are of interest in our analysis: a simple autoregressive model or AR(1) ( $\psi_{it} = 0$ ) and an independent model ( $\phi_{it} = 0, \psi_{it} = 0,$ ). To insure identification we assume that $\mathrm{E}[\varepsilon_{itj}] = 0$ and $\mathrm{Var}[\varepsilon_{itj}] = 1$. (2) can be expressed in matrix form as follows:

$$\mathbf{u}_{it} = \mathbf{x}_{it}\boldsymbol{\beta}_{it} + \boldsymbol{\varepsilon}_{it}, \quad \boldsymbol{\varepsilon}_{it} \sim \mathrm{N}(\mathbf{0}, \boldsymbol{\Psi}_{it}) \, . \tag{4}$$

where $\mathbf{u}_{it} = \begin{bmatrix} u_{it1} & \cdots & u_{itN} \end{bmatrix}'$, $\mathbf{x}_{it} = \begin{bmatrix} \mathbf{x}_{it1}' & \cdots & \mathbf{x}_{itN}' \end{bmatrix}'$, and the ($k$,$l$)th element of the error covariance matrix is (Box and Jenkins 1970):

$$[\boldsymbol{\Psi}_{it}]_{k,l} = \begin{cases} 1 & \text{if } k = l \\ \dfrac{(\phi_{it} - \psi_{it})(1 - \phi_{it}\psi_{it})}{1 + \psi_{it}^{\,2} - 2\phi_{it}\psi_{it}} \phi_{it}^{\,|k-l|-1} & \text{otherwise} \end{cases}. \tag{5}$$

The motivation for this correlated error structure is that if a consumer unexpectedly clicks upon an offer then its neighboring offers are also more likely to be selected. The order in which the offers are presented yields a natural order in which users may influence a user's consideration, this contrasts with brand choice problems that lack a natural spatial context (since shelf layout is rarely measured, and when it is there is little variation in its design with which to measure its impact). We could consider a more general spatial model, but in our problem this would be computationally difficult due to the large number of choices and the need for the covariance matrix to be defined for varying set sizes as occurs in our data.

The benefit of the ARMA(1,1) model is that it is quite versatile, can approximate a large number of spatial models, and has a known inverse (Tiao and Ali 1971). For further discussion about more general spatial error structures see Yang and Allenby (2002), who consider a spatial autocorrelation model in a binary probit model.

## 4.1. Multivariate Probit

A user may select none, one, or more items. A standard approach to model (1) is to assume that the threshold is constant for a session, which yields a multivariate probit model:

$$\lambda_{itj} = \boldsymbol{\gamma}_i' \mathbf{w}_{it} \tag{6}$$

where $\mathbf{w}_{it}$ is the vector of session covariates given in Table 1. Our covariates ( $\mathbf{w}_{it}$ ) include the variability of prices, the number of days since the item was last searched, time effects to capture potential learning effects, the list price of the book (e.g., more expensive books will result in more search), the coefficient of variation for prices offered (e.g., more dispersion, suitability normalized, will result in more search), and the number of items presented to the consumer. Note that identifiability requires that the covariates used in (6) must be different than those in (2), or equivalently the threshold in (5) is set to zero and the covariates are absorbed into model (2). For a marketing application of a multivariate probit model see Manchanda et al (1999). The probit model occurs when the errors are i.i.d. ( $\phi_{it} = \psi_{it} = 0$ ), which yields independent binary probit models for each offer.

## 4.2. Conditional Order Threshold Probit Models

A problem that with the multivariate probit model is that the marginal distribution of the number of items selected is implicitly defined by the model and may not well capture the observed distribution of choice for our dataset given in Figure 2. To illustrate this problem consider the simple case were $u_{itj}$ are i.i.d. (i.e., $\phi_{it} = \theta_{it} = 0$ and $\mathrm{E}[u_{itj}] = \overline{u}$ ) and $\lambda_{itj} = 0$ , then the probability an offer is selected equals $\zeta = 1 - \Phi(-\bar{u})$. Hence, the distribution of the number of offers selected ($P$) follows a binomial distribution, where the mean and variance are:

$$\mathrm{E}[P] = \zeta N, \mathrm{Var}[P] = \zeta(1-\zeta)N \tag{7}$$

This implies that as $N$ increases the expected number of offers selected will increase and when $N$ is large the marginal distribution is approximately normal.

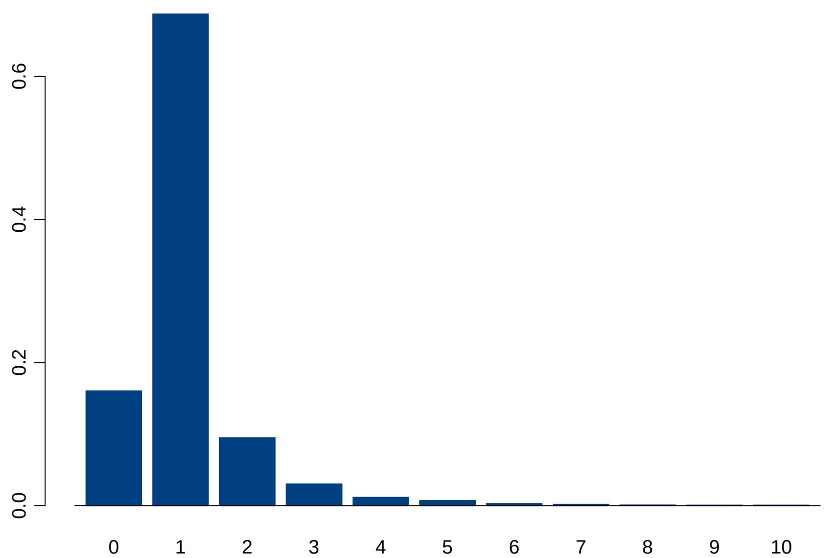


**Figure 2.** The observed frequency of the total number of offers selected during a session.

Nor does a correlated error change this property. Consider a simple bivariate probit ($N$=2) with zero mean ($\bar{u} = 0$), unit variance, and correlation of $\upsilon$, hence the probability an item is selected is $\zeta = .5$ and the expectation and variance of $P$ is:

$$\mathrm{E}[P] = 1, \mathrm{Var}[P] = \frac{1}{2} + \frac{1}{\pi}\arcsin(\upsilon) \tag{8}$$

Notice that the mean is unchanged from (7). Although the variability now depends upon the correlation, for example the variance equals 0, .5, or 1 when $\upsilon = -1,\ 0,\ \text{or}\ 1$, respectively.

The implied marginal distribution of the multivariate probit model conflicts with our exploratory analysis. Therefore we propose a model that we call the conditional order choice model that is more flexible. We assume that the number of items chosen follows a specified marginal distribution, and conditional upon the number of items chosen the user will select the top $P$ ordered latent utilities. For example, if we assume that $P$ follows a Poisson model with a location parameter of $\theta$ then the corresponding mean and variance are:

$$\mathrm{E}[P] = \theta, \mathrm{Var}[P] = \theta \tag{9}$$

Recall that the Poisson distribution is the limiting distribution of the binomial distribution when $\zeta N \to \theta$ as $N \to \infty$. The Poisson distribution is more consistent with our observed marginal distribution in Figure 2. More generally, we can postulate any discrete distribution for $P$.

Conditional upon the number of items chosen we assume that the user chooses the best $P$ alternatives, which is equivalent to setting the threshold for offer $j$ equal to the $P$th order statistic of the set of the remaining alternatives:

$$\lambda_{itj} = \mathbf{u}_{it}^{-j} \langle P_{it} \rangle, \text{ where } \mathbf{u}_{it}^{-j} = \{u_{it1}, \ldots, u_{it,j-1}, u_{it,j+1}, \ldots, u_{itN}\} \tag{10}$$

The operator $\mathbf{z}\langle P \rangle$ denotes the $P$th order statistic, $z_{(P)}$, from the set $\mathbf{z} = \{z_1, \ldots, z_N\}$:

$$\mathbf{z}\langle P \rangle = \begin{cases} \infty & \text{if } P = 0 \\ z_{(P)} & \text{if } 1 \leq P \leq N, \text{ where } z_{(1)} \geq z_{(2)} \geq \cdots \geq z_{(N)} \\ -\infty & \text{if } P > N \end{cases} \tag{11}$$

We also define this operator to handle the extreme cases of either none or all items being selected, since the usual order statistics are not defined for these cases.

When $N$=1 then our model is identical to a multinomial probit model. To illustrate this relationship consider our observational equation when $N$=1:

$$\begin{aligned} u_{itj} \geq \lambda_{itj} &= \max(u_{it1}, \ldots, u_{it,j-1}, u_{it,j+1}, \cdots, u_{itN}) \\ &\Rightarrow u_{itj} \geq u_{it1}, \ldots, u_{itj} \geq u_{it,j-1}, u_{itj} \geq u_{it,j+1}, \ldots, u_{itj} \geq u_{itN} \end{aligned} \tag{12}$$

This expression is identical to the observational equation of the multinomial probit model.

Although we initially suggested a Poisson distribution for $P_{it}$ we can introduce any distribution defined over discrete values. We considered truncated versions of the Poisson or Poisson-log normal distribution (initially introduced by Preston 1948) but found these distributions could not capture the marginal distribution found in Figure 2. We found a more suitable candidate distribution in the discretized log-normal distribution (Li, Liechty, and Montgomery, 2005). We modify this distribution by truncating it because the number of items selected cannot exceed the available number of items. Formally, we define the following distribution for $P_{it}$ :

$$\Pr[P_{it} = p \mid \theta, N_{it}] = \Pr[\ln(p) < \theta_{it} \leq \ln(p+1)] \text{ where } 0 \leq p \leq N_{it} - 1 . \tag{13}$$

The latent random variable ( $\theta_{it}$ ) is assumed to follow a log normal regression:

$$\ln(\theta_{it}) = \gamma_i{}'\mathbf{w}_{it} + \alpha_{it}, \quad \alpha_{it} \sim \mathrm{N}(0, \tau^2) \tag{14}$$

If $N_{it}$ is large (or if there is no truncation on $P_{it}$ ) and we ignore the discretization process then the mean and variance of the expected number of items is:

$$\mathrm{E}[P_{it}] \approx \exp\left\{\gamma_i{}'\mathbf{w}_{it} + \tfrac{1}{2}\tau^2\right\}, \quad \mathrm{Var}[P_{it}] \approx \exp\left\{2\gamma_i{}'\mathbf{w}_{it} + \tfrac{1}{2}\tau^2\right\}\left[\exp\left\{\tau^2\right\} - 1\right] \tag{15}$$

Other models of multivariate choice could be considered. Recently Kim et al. 2002 considered multivariate choice in the context of shopping for yogurt. However, the purchase context of their problem does not translate well to our situation, since we are not modeling the selection of which item to purchase but the precursor of which items to consider. Another alternative model would be to sequentially model choice. For example, we could think about a user making a decision about whether to make a choice, and if they do choose to continue then determining which item to consider, and then repeating this procedure until the user decides to stop. Unfortunately, we do not have the order in which offers were selected, so this sequential method is not possible with our dataset.

## 4.3. Hierarchical Specification

To explain variation in coefficients for each session we introduce the following hyper-distributions that explain variation in coefficients using user and session covariates:

$$\beta_{it} \sim \mathrm{N}(\mathrm{K}_i \mathbf{w}_{it}, \mathrm{H}), \tag{16}$$

$$\gamma_i \sim \mathrm{N}(\boldsymbol{\Gamma} \mathbf{z}_i, \boldsymbol{\Omega}) \tag{17}$$

$$\boldsymbol{\delta}_{it} \sim \mathbf{N}(\bar{\boldsymbol{\delta}}, \mathbf{X}) \cdot \mathrm{I}(-1 \leq \phi_{it} \leq 1, -1 \leq \theta_{it} \leq 1) \ , \ \ \boldsymbol{\delta}_{it} = \begin{bmatrix} \phi_{it} & \theta_{it} \end{bmatrix}' \tag{18}$$

The priors and MCMC algorithm used to estimate the model are discussed in the Appendix.

## 5. Empirical Results

The model was estimated using the dataset described in section 2. Tables 2 and 3 provide descriptive statistics for the covariates used in the model.

| Type | Description | Mean | Std | Min | Max |
|---|---|---|---|---|---|
| | Number of Clicks | 1.12 | 1.10 | 0 | 41 |
| | Number of Offers | 35.35 | 15.94 | 2 | 73 |
| When | Daytime | .53 | .50 | | |
| | Weekday | .71 | .45 | | |
| Search | Cuml. Visit | .40 | .89 | 0 | 13 |
| | Sorted by Price | .56 | .50 | | |
| Book | New Release | .33 | .47 | | |
| | Bestseller | .01 | .09 | | |
| | Log(Publisher's List Price) | 3.38 | .80 | -1.43 | 8.96 |
| | Std Dev Total Price | 9.26 | 12.84 | .71 | 882.27 |
| | Hard Cover | .42 | .49 | | |
| | Soft Cover | .53 | .50 | | |
| Topic | Business | .08 | .28 | | |
| | Computer | .17 | .38 | | |
| | Leisure | .23 | .42 | | |
| | Science | .15 | .36 | | |
| | Education | .13 | .34 | | |
| | Art | .11 | .31 | | |

**Table 2.** Description of session covariates.

| Type | Description | Mean | Std | Min | Max |
|---|---|---|---|---|---|
| Price | Item Price | 36.56 | 62.41 | 0 | 7867.50 |
| | Shipping Cost | 9.40 | 6.59 | 0 | 360.00 |
| | Tax | .27 | 1.67 | 0 | 239.25 |
| | Tax N/A | .89 | .32 | | |
| Delivery | Average Days to Deliver | 7.18 | 10.84 | | |
| Position | Rank | 21.77 | 14.28 | 1 | 73 |
| | Rank^2/100 | 6.78 | 7.48 | .01 | 53.29 |
| | Rank^3/1000 | 24.99 | 37.71 | .001 | 389.02 |
| Logo | Logo Bold | .05 | .22 | | |
| | Logo Graphic | .15 | .36 | | |
| Retailer | Amazon | .07 | .26 | | |
| | B&N | .06 | .24 | | |
| | Borders | .05 | .22 | | |

**Table 3.** Description of offer covariates.

The average of the hyper-distribution for the session and offer level parameters are given in Table 4 and 5. Table 4 provides the effects of changes on the number of items likely to be considered using our Truncated Poisson-log normal model. Since the dependent variable is given in logarithmic terms one can think of the parameters as approximately the percentage change from the mean for a one-unit change in the covariate. For example for each additional cumulative visit for the user one would expect a .25 decrease in the number of items selected. Notice that users are less likely to browse during the daytime but more in the evening. The cumulative number of visits indicates that there is a potential learning effect for shopbots. As users continue to revisit the site they select fewer and fewer offers, the standardized beta indicates this has the largest impact of all variables. Additionally, users tend to search less for bestsellers but more as the price of the book or variability of price increases.

| Type | Description | Mean | Std Err | StdBeta |
|---|---|---|---|---|
| | Intercept | .227 | .038 | |
| | Number of Offers | -.003 | .0002 | -1.9 |
| When | Daytime | -.034 | .010 | -.7 |
| | Weekday | .046 | .010 | .9 |
| Search | Cuml. Visit | -.253 | .009 | -9.1 |
| | Sorted by Price | -.009 | .006 | -.2 |
| Book | New Release | -.015 | .009 | -.3 |
| | Bestseller | -.077 | .035 | -.3 |
| | Log(Publisher's List Price) | .022 | .011 | .7 |
| | Std Dev Total Price | .00012 | .00076 | .1 |
| | Hard Cover | -.057 | .0123 | -1.2 |
| | Soft Cover | -.019 | .012 | -.4 |
| Topic | Business | .002 | .027 | .0 |
| | Computer | .011 | .018 | .2 |
| | Leisure | -.016 | .018 | -.3 |
| | Science | -.034 | .015 | -.5 |
| | Education | -.042 | .021 | -.6 |
| | Art | -.042 | .015 | -.5 |

**Table 4.** Estimates of hyper-distribution for conditional model of number of items chosen.

Table 5 reports the effects of the offer covariates on the latent utility of the item. Notice that the parameters in Table 4 dictate the number of items selected. Notice that as the price or shipping costs increase that the attractiveness of the offer decreases. If tax is not present (either a non-taxed or unreported number) the utility is substantially decreased, indicating the consumers are more sensitive to a dollar of tax than to a dollar for the unit price of the book. As delivery increases the book also becomes less attractive. The most important effects relate to the position of the item. We fit a cubic function which indicates that on average the most attractive position is at the top. This decreases quickly but then begins to increase towards the end of the list. Amazon, B&N, and Borders all show positive branding effects. However, the logos on average seem to decrease the attractiveness of the offers. This effect is potentially compounded by the fact that logo's were not present in the first half of the data, but were present in the second half of the dataset.

| Type | Description | Mean | Std | StdBeta |
|---|---|---|---|---|
| Price | Item Price | -.062 | .002 | -1.2 |
| | Shipping Cost | -.037 | .003 | -.3 |
| | Tax | .553 | .022 | 1.0 |
| | Tax N/A | 6.617 | .237 | 2.5 |
| Delivery | Average Days to Deliver | -.010 | .001 | -.1 |
| Position | Rank | -.649 | .008 | -9.1 |
| | Rank^2/100 | 2.528 | .043 | 19.0 |
| | Rank^3/1000 | -.278 | .005 | -10.6 |
| Logo | Logo Bold | -1.527 | .063 | -.3 |
| | Logo Graphic | -1.392 | .083 | -.6 |
| Retailer | Amazon | .506 | .070 | .2 |
| | B&N | .586 | .036 | .2 |
| | Borders | .077 | .074 | .0 |

**Table 5.** Hyper-parameter estimates for effects on latent choice utilities.

## 6. Conclusions

As Internet shopbot technologies mature, consumer behavior at shopbots will become an increasingly important topic for consumers, retailers, financial markets, and academic researchers. With regard to consumer behavior, our findings demonstrate that, while shopbots substantially weaken the market positions of branded retailers, brand name and retailer loyalty still strongly influence consumer behavior at Internet shopbots. Our findings also suggest that consumers use brand name as a signal of reliability in service quality for non-contractible aspects of the product bundle. These results may derive from service quality differentiation, asymmetric market information regarding quality, or cognitive lock-in among consumers.

With regard to retailers, our results suggest several differential-pricing strategies for shopbot markets. First, it is likely that a consumer's willingness to take the extra time to use a shopbot is a credible signal of price sensitivity. Thus, retailers may use this information as part of a price discrimination strategy — charging lower prices to shopbot consumers than consumers who visit their web site directly. Second, our findings suggest that partitioned pricing strategies that increase demand among web site direct consumers may decrease demand among shopbot consumers. Because of this, retailers should adopt different pricing strategies for shipping cost for shopbot consumers than they would for web site direct consumers. Lastly, the reliability of our models when compared to actual consumer behavior suggests that retailers may be able to use shopbot data to provide personalized prices to consumers.

For financial markets, our findings may help to focus the debate on the size and sustainability of market valuations for Internet retailers. Using Amazon.com as an example, our shopbot data indicate that the retailer maintains a 5.0% margin advantage over unbranded retailers and a 6.8% margin advantage among repeat visitors. Both of these statistics are likely to represent lower bounds on the actual margin advantages among their entire consumer base. A margin advantage of this magnitude, if sustainable and applicable across their entire product line, implies a very large capital value. The relevant questions then become whether companies such as Amazon.com can sustain current positions of competitive advantage, how much it will cost to sustain these positions, and whether they can transfer competitive advantage in one product category to other product categories to expand their revenue base.

Finally, for academic researchers, our results demonstrate the feasibility of using Internet shopping data to better understand consumer behavior in electronic markets. Future research in this regard may be able to extend these results to better understand how web-site direct and shopbot consumers respond to partitioned prices, to evaluate the cognitive processing costs of shopbot consumers, and to empirically analyze the application of personalized pricing strategies to shopbot consumers.

Moreover, our results suggest that the quantity and quality of data available in Internet markets may introduce a revolution the analysis of consumer behavior rivaling that of the scanner data revolution in the 1980s.

# References

Bajari, P. and A. Hortacsu (2003). "The winner's curse, reserve prices, and endogenous entry: empirical insights from eBay auctions." RAND Journal of Economics, 34:329-355.

Bakos, J. Yannis. 1997. Reducing Buyer Search Costs: Implications for Electronic Marketplaces. Management Science 43(12) 1613-1630.

Bakos, J. Yannis. 1998. The Emerging Role of Electronic Marketplaces on the Internet. Communications of the ACM 41(8)

Baye, M., J. Gatti, P. Kattuman, J. Morgan. 2004. Estimating Firm-Level Demand at a Price Comparison Site: Accounting for Shoppers and the Number of Competitors, Working Paper, Indiana University, Bloomington, IN.

Bearse, Peter M., Hamparsum Bozdogan, Alan M. Schlottmann. 1997. Empirical econometric modelling of food consumption using a new informational complexity approach. Journal of Applied Econometrics 12(5) 563-592.

Brown, Jeffrey R., Austan Goolsbee. 2000. Does the Internet make markets more competitive? Evidence from the life insurance industry. Journal of Political Economy 110(3) 481-507.

Brynjolfsson, Erik, Astrid A. Dick, Michael D. Smith. 2004. Search and Product Differentiation at an Internet Shopbot. Working Paper, Carnegie Mellon University, Pittsburgh, PA.

Brynjolfsson, Erik, Michael Smith. 2000. Frictionless Commerce? A Comparison of Internet and Conventional Retailers. Management Science Forthcoming April.

Clay, Karen, Ramayya Krishnan, Eric Wolff, Danny Fernandes. 1999. Retail Strategies on the Web: Price and Non-price Competition in the Online Book Industry. Working Paper, Carnegie-Mellon University, Pittsburgh, PA.

Clemons, Eric K., Il-Horn Hann, Lorin M. Hitt. 1998. The Nature of Competition in Electronic Markets: An Empirical Investigation of Online Travel Agent Offerings. Working Paper, The Wharton School of the University of Pennsylvania, Philadelphia, PA.

Diehl, K., L.J. Kornish, J.G. Lynch, Jr. (2003). "Smart Agents: When Lower Search Costs for Quality Information Increase Price Sensitivity."Journal of Consumer Research 30:56-71.

Ellison, Glen; Sara Fisher Ellison. 2001. Search, Obfuscation, and Price Elasticities on the Internet. Working Paper, MIT, Cambridge, MA, January.

Greenwald, Amy R., Jeffrey O. Kephart. 1999. Shopbots and Pricebots. The Proceedings of International Joint Conference on Artificial Intelligence.

Hann, I-L, C. Terwiesch. 2003. Measuring the Frictional Costs of Online Transactions: The Case of a Name-Your-Own-Price Channel, 49(11) 1563-1579.

Johnson, E.J., S. Bellman and G.L. Lohse (2002). "Cognitive Lock In and the Power Law of Practice." Columbia University, mimeo.

Lynch, John G., Jr., Dan Ariely. 2000. Wine online: Search cost and competition on price, quality, and distribution. Marketing Science 19(1) 83-103.

Milgrom, Paul, John Roberts. 1986. Price and Advertising Signals of Product Quality. The Journal of Political Economy. 94(4) 796-821.

Rossi, Peter E., Robert E. McCulloch, Greg M. Allenby. 1996. The value of purchase history data in target marketing. Marketing Science 15(4) 321-340.

Smith, Michael. 2002. The Impact of Shopbots on Electronic Markets. Journal of the Academy of Marketing Science 30(4) 442-450.

Smith, Michael, Erik Brynjolfsson. 2001. Customer Decision Making at an Internet Shopbot: Brand Still Matters. The Journal of Industrial Economics 49(4) 541-558.

Smith, Michael D.; Bailey, Joseph; Brynjolfsson, Erik. 2000. Understanding Digital Markets: Review and Assessment. E. Brynjolfsson, B. Kahin, eds. Understanding the Digital Economy, MIT Press, Cambridge, MA. Forthcoming.

Sorensen, A.T. (2000). "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs." Journal of Political Economy, 108:833-850.

Sorensen, A.T. (2001). "Price Dispersion and Heterogeneous Consumer Search for Retail Pre- scription Drugs." Working Paper, University of California, San Diego.

# Appendix: Prior and MCMC Algorithm

To complete the Bayesian specification of the model we define the following priors:

$$\bar{\boldsymbol{\delta}} \sim \mathbf{N}(\bar{\bar{\boldsymbol{\delta}}}, \mathbf{V}_{\bar{\boldsymbol{\delta}}}) \tag{1}$$

$$\mathbf{vec}(\boldsymbol{\Gamma}) \sim \mathbf{N}(\bar{\boldsymbol{\gamma}}, \mathbf{V}_{\Gamma}) \tag{2}$$

$$\mathrm{vec}(\mathbf{M}_k) \sim \mathbf{N}(\bar{\boldsymbol{\mu}}_k, \mathbf{V}_{\mathrm{M}_k}) \tag{3}$$

$$\lambda_k^2 \sim \mathrm{IG}(\upsilon_{\lambda_k}, \upsilon_{\lambda_k}) \tag{4}$$

$$\mathbf{X} \sim \mathrm{W}^{-1}(\upsilon_{\mathrm{X}}, \mathbf{V}_{\mathrm{X}}) \tag{5}$$

$$\boldsymbol{\Omega} \sim \mathrm{W}^{-1}(\upsilon_{\Omega}, \mathbf{V}_{\Omega}) \tag{6}$$

$$\Pi_k \sim \mathrm{W}^{-1}(\upsilon_{\Pi_k}, \mathbf{V}_{\Pi_k}) \tag{7}$$

$$\tau^2 \sim \mathrm{IG}(\upsilon_{\tau}, \lambda_{\tau}) \tag{8}$$

To estimate this model we implement an Monte Carlo Markov Chain (MCMC) algorithm which sequentially samples from each of the following distributions. For notational simplicity we assume that each draw is made conditional upon all data and parameters except for the one under consideration. Also, the last draw of each parameter is used.

- Each element of $\mathbf{u}_{it}$ is drawn conditional upon all other values (Rossi, Allenby, and McCulloch 1996):

$$u_{itj} \mid \mathbf{u}_{it}^{-j} \sim \mathrm{N}(\bar{u}_{itj} + \boldsymbol{\psi}_{it}^{j}\boldsymbol{\Psi}_{it}^{-j}(\mathbf{u}_{it}^{-j} - \bar{\mathbf{u}}_{it}^{-j}), [\boldsymbol{\Psi}_{it}]_{jj} - \boldsymbol{\psi}_{it}^{j}\boldsymbol{\Psi}_{it}^{-j}\boldsymbol{\psi}_{it}^{j\,\prime}) \cdot \mathrm{I}(\Xi_{itj}) \,. \tag{9}$$

where $\bar{\mathbf{u}}_{it} = \mathbf{x}_{it}\boldsymbol{\beta}_{it}$, $\boldsymbol{\psi}_{it}^{j}$ is the $j$th row vector of $\boldsymbol{\Psi}_{it}$ with the $j$th element deleted, and $\boldsymbol{\Psi}_{it}^{-j}$ denotes $\boldsymbol{\Psi}_{it}$ with the $j$th row and column deleted. Additionally, the truncation region is:

$$\Xi_{itj} = \begin{cases} u_{itj} \geq \lambda_{itj} & \text{if } y_{itj} = 1 \\ u_{itj} < \lambda_{itj} & \text{if } y_{itj} = 0 \end{cases}, \tag{11}$$

where $\lambda_{itj} = 0$ for the multivariate probit model and $\lambda_{itj} = \mathbf{u}_{it}^{-j}\langle P_{it}\rangle$ for the ordered conditional probit model.

- $\boldsymbol{\beta}_{it}$ is drawn from a generalized regression:

$$\boldsymbol{\beta}_{it} \sim \mathrm{N}\left(\Delta^{-1}\left(\mathbf{x}_{it}'\boldsymbol{\Psi}_{it}^{-1}\mathbf{u}_{it} + \mathbf{H}^{-1}\mathbf{B}_i\mathbf{w}_{it}\right), \Delta^{-1}\right) \text{ , where } \Delta = \mathbf{x}_{it}'\boldsymbol{\Psi}_{it}^{-1}\mathbf{x}_{it} + \mathbf{H}^{-1} \tag{12}$$

We follow Tiao and Ali (1971) who propose an efficient algorithm of computing $\mathbf{x}_{it}'\boldsymbol{\Psi}_{it}^{-1}\mathbf{u}_{it}$ and $\mathbf{x}_{it}'\boldsymbol{\Psi}_{it}^{-1}\mathbf{x}_{it}$.

- Draw $\boldsymbol{\delta}_{it}$ using a slice sampler:

Let $\ell_{it}(\boldsymbol{\delta}_{it})$ define the likelihood function of an ARMA(1,1) model for the latent utility vector:

$$\ell_{it}(\boldsymbol{\delta}_{it}) = (2\pi)^{-\frac{N_{it}}{2}}\left|\boldsymbol{\Psi}_{it}^{-1}\right|\exp\left\{-\tfrac{1}{2}\mathbf{d}_{it}'\boldsymbol{\Psi}_{it}^{-1}\mathbf{d}_{it}\right\}, \text{ where } \mathbf{d}_{it} = \mathbf{u}_{it} - \mathbf{x}_{it}\boldsymbol{\beta} \tag{13}$$

We follow Agarwal and Gelfand (2003) to draw $\boldsymbol{\delta}_{it}$ using a slice sampler. First we draw a new latent variable $v_{it}$:

$$v_{it} \mid \boldsymbol{\delta}_{it} \sim e_{it}\,\mathrm{I}[v_{it} > -\ln(\ell_{it}(\boldsymbol{\delta}_{it}))] \tag{14}$$

where $e_{it}$ is a standard exponential distribution. Second we can draw the $\boldsymbol{\delta}_{it}$ given $v_{it}$:

$$v_{it} \mid \boldsymbol{\delta}_{it} \sim \mathbf{N}(\bar{\boldsymbol{\delta}}, \mathbf{X}) \cdot \mathrm{I}(-1 \le \phi_{it} \le 1, -1 \le \theta_{it} \le 1, \ln(\ell_{it}(\boldsymbol{\delta}_{it})) > -v_{it}] \tag{15}$$

We employ Neal's (2003) rectangular adaptive rejection algorithm since the boundaries of the slice are not easily identified.

- Draw $\theta_{it}$ for Conditional Order Probit Model using a slice sampler:

Damien, Wakefield, and Walker (1999) first proposed estimating a poisson-log normal distribution using a slice sampler. We extend their method to allow for a truncated poisson distribution. To draw $\theta_{it}$ we introduce two latent values $a_{it}$ and $b_{it}$. $a_{it}$ is a truncated exponential distribution:

$$a_{it} \mid \theta_{it} \sim e_{it}\,\mathrm{I}[a_{it} > \theta_{it}] \tag{16}$$

where $e_{it}$ is a standard exponential distribution. $b_{it}$ is a uniform distribution:

$$b_{it} \mid \theta_{it} \sim \mathrm{I}[0 < b_{it} < \frac{1}{\mathrm{Q}(1 + N_{it}, \theta_{it})}] \tag{17}$$

$\ln(\theta_{it})$ can then be drawn as a truncated normal:

$$\ln(\theta_{it}) \mid b_{it} \sim N(\boldsymbol{\gamma}_i'\mathbf{w}_{it} + \tau^2 P_{it}, \tau^2) \cdot \mathrm{I}[a > \ln(\theta_{it}),\ b_{it} < \frac{1}{\mathrm{Q}(1 + N_{it}, \theta_{it})}] \tag{18}$$

Since the last truncation term cannot be inverted directly we follow Neal (2003) to implement a rectangular adaptive rejection algorithm.

- Draw $\boldsymbol{\gamma}_i$ for Multivariate Probit Model

For identification reasons $\boldsymbol{\gamma}_i$ these parameters can instead be treated as a subvector of $\boldsymbol{\beta}_{it}$ and we redefine the threshold to zero.

- Draw $\boldsymbol{\gamma}_i$ for Conditional Order Probit Model

$$\boldsymbol{\gamma}_i \sim \mathrm{N}\left(\Delta^{-1}\left(\tau^{-2}\mathbf{w}_i' \ln(\boldsymbol{\theta}_i) + \boldsymbol{\Omega}^{-1}\Gamma\mathbf{z}_i\right), \Delta^{-1}\right) \text{, where } \Delta = \tau^{-2}\mathbf{w}_i'\mathbf{w}_i + \boldsymbol{\Omega}^{-1} \tag{19}$$

where $\mathbf{w}_i = \begin{bmatrix}\mathbf{w}_{i1} & \cdots & \mathbf{w}_{iM_i}\end{bmatrix}'$ and $\boldsymbol{\theta}_i = \begin{bmatrix}\theta_{i1} & \cdots & \theta_{iM_i}\end{bmatrix}'$.

- Each row of $\mathrm{B}_i$ is drawn from a multivariate regression:

$$\mathrm{row}_k(\mathrm{B}_i) \sim \mathrm{N}\left(\Delta^{-1}\left(\lambda_k^{-2}\mathbf{w}_i'\beta_{ik} + \Pi_k^{-1}\mathbf{z}_i\right), \Delta^{-1}\right) \text{, where } \Delta = \lambda_k^{-2}\mathbf{w}_i'\mathbf{w}_i + \Pi_k^{-1} \tag{20}$$

where $\beta_{ik} = \begin{bmatrix}\beta_{i1k} & \cdots & \beta_{iT_ik}\end{bmatrix}'$.

- $\boldsymbol{\Gamma}$ is drawn from a multivariate regression:

$$\mathrm{vec}(\boldsymbol{\Gamma}) \sim \mathrm{N}\left(\Delta^{-1}\left(\mathbf{z}'(\mathbf{I}_K \otimes \boldsymbol{\Omega}^{-1})\boldsymbol{\gamma} + \mathbf{V}_\Gamma{}^{-1}\overline{\boldsymbol{\gamma}}\right), \Delta^{-1}\right) \text{, where } \Delta = \mathbf{z}'(\mathbf{I}_K \otimes \boldsymbol{\Omega}^{-1})\mathbf{z} + \mathbf{V}_\Gamma{}^{-1} \tag{21}$$

where $\mathbf{z} = \begin{bmatrix}\mathbf{z}_1' & \cdots & \mathbf{z}_K'\end{bmatrix}'$ and $\boldsymbol{\gamma} = \begin{bmatrix}\boldsymbol{\gamma}_1' & \cdots & \boldsymbol{\gamma}_K'\end{bmatrix}'$.

- $\overline{\boldsymbol{\delta}}$ is drawn from a multiple regression:

$$\overline{\boldsymbol{\delta}} \sim \mathrm{N}\left(\Delta^{-1}\left(\sum_{i=1}^{K}\mathbf{X}^{-1}\boldsymbol{\delta}_i + \mathbf{V}_{\bar{\delta}}^{-1}\overline{\overline{\boldsymbol{\delta}}}\right), \Delta^{-1}\right) \text{, where } \Delta = K\mathbf{X}^{-1} + \mathbf{V}_{\bar{\delta}}^{-1} \tag{22}$$

- $\mathrm{M}_k$ is drawn from a multivariate regression:

$$\mathrm{vec}(\mathrm{M}_k) \sim \mathrm{N}\left(\Delta^{-1}\left(\mathrm{B}_k'(\mathbf{I}_K \otimes \Pi_k{}^{-1})\mathbf{z} + \mathbf{V}_{\mathrm{M}_k}{}^{-1}\mu\right), \Delta^{-1}\right) \text{, where } \Delta = \mathrm{B}_k'(\mathbf{I}_K \otimes \Pi_k{}^{-1})\mathrm{B}_k + \mathbf{V}_{\mathrm{M}_k}{}^{-1} \tag{23}$$

where $\mathrm{B}_k = \begin{bmatrix}\mathrm{B}_{1k}' \cdots & \mathrm{B}_{Kk}'\end{bmatrix}'$.

- Draw $\tau$

$$\tau \sim IG\left(\tfrac{1}{2}\left(\sum_{i=1}^{K} M_i + \upsilon_\tau\right), \tfrac{1}{2}\left(\sum_{i=1}^{K}\mathbf{E}_i'\mathbf{E}_i + \lambda_\tau\right)\right), \text{ where } \mathbf{E}_i = \ln(\theta_i) - \mathbf{w}_i\gamma_i \tag{24}$$

- Draw $\lambda_1, \lambda_2, \ldots, \lambda_R$

$$\lambda_k^2 \sim IG\left(\tfrac{1}{2}\left(\sum_{i=1}^{K} M_i + \upsilon_{\lambda_k}\right), \tfrac{1}{2}\left(\sum_{i=1}^{K}\mathbf{E}_i'\mathbf{E}_i + \upsilon_{\lambda_k}\right)\right), \text{ where } \mathbf{E}_{it} = \beta_{it} - \mathrm{B}_{ik}\mathbf{w}_{it} \tag{25}$$

- Draw $\boldsymbol{\Omega}$

$$\boldsymbol{\Omega} \sim \mathrm{W}^{-1}\left(K + \upsilon_{\Omega}, \sum_{i=1}^{K} \mathbf{E}_i'\mathbf{E}_i + \mathbf{V}_{\Omega}\right), \text{ where } \mathbf{E}_i = \boldsymbol{\gamma}_i - \boldsymbol{\Gamma}\mathbf{z}_i \tag{26}$$

- Draw $\mathbf{X}$

$$\mathbf{X} \sim \mathrm{W}^{-1}\left(\sum_{i=1}^{K} M_i + \upsilon_{\mathrm{X}}, \sum_{i=1}^{K}\sum_{t=1}^{M_i} \mathbf{E}_{it}'\mathbf{E}_{it} + \mathbf{V}_{\mathrm{X}}\right), \text{ where } \mathbf{E}_{it} = \boldsymbol{\delta}_{it} - \overline{\boldsymbol{\delta}} \tag{27}$$

- Draw $\Pi_k$

$$\Pi_k \sim \mathrm{W}^{-1}\left(K + \upsilon_{\Pi_k}, \sum_{i=1}^{K} \mathbf{E}_i'\mathbf{E}_i + \mathbf{V}_{\Pi_k}\right), \text{ where } \mathbf{E}_i = \mathrm{vec}(\mathrm{B}_{ik}) - \mathrm{M}_k\mathbf{z}_i \tag{28}$$