

Designing a Better Shopbot


(with Karen Clay, Ramayya Krishnan,
and Kartik Hosanagar)

Alan Montgomery
Associate Professor
Carnegie Mellon University
Graduate School of Industrial Administration

e-mail: alan.montgomery@cmu.edu
web: <http://www.andrew.cmu.edu/user/alm3>

Cornell Marketing Workshop, 26 January 2000

What is a shopbot?



RETAILER	PRICE	DESCRIPTION	UNIT	LIST
buy.com	\$13.49	PKA	EA	\$18.44
totalinformation.com	\$19.17	7-28 Logo K&G Ground Shroud	EA	\$27.22
buy.com	\$13.49	PKA	EA	\$25.44
totalinformation.com	\$18.17	7-28 Logo Guarded Shroud	EA	\$22.27
totalinformation.com	\$19.17	4-5 Logo K&G Ground Shroud	EA	\$22.53
totalinformation.com	\$18.17	7-28 Logo Guarded Shroud	EA	\$22.51
totalinformation.com	\$19.17	7-25 Logo K&G Ground	EA	\$22.12

Shopping Robot's automatically search
a large number of stores for a specific
product

Example:

John Grisham's *The Brethren*, list price
\$27.95, prices range between \$13.49
(buy.com) and \$50.75
(totalinformation.com)

Design involves:

- Computer science (agents)
- Economics (value of price search)
- Marketing (consumer behavior)
- Statistical Modeling (uncertainty)



Outline

- The State of Shopbots
- Improving Shopbot Design
- Modeling Consumer Utility
- Application to Online Bookstores
- Simulation Results
- Conclusions

3



The State of Shopbots

Online Retail Shopping

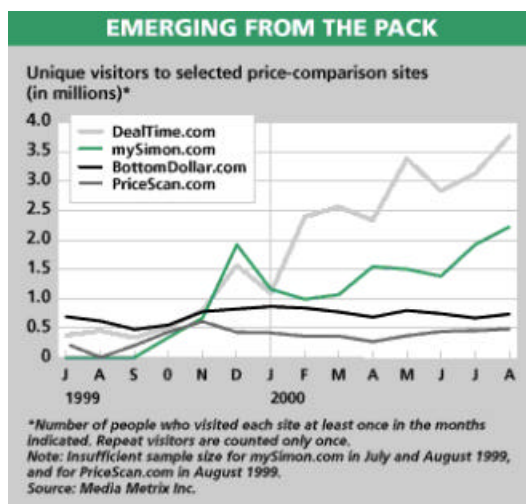
	Millions of Unique Monthly Visitors	% of Web Users
All Digital Media	80,097	100.0
Retail	53,485	66.8
1 amazon.com	14,464	18.1
2 americangreetings.com	7,719	9.6
3 webstakes.com	5,314	6.6
4 bamesandnoble.com	5,281	6.6
5 mypoints.com	5,269	6.6
6 bizrate.com	5,050	6.3
7 directhit.com	3,952	4.9
8 cdnow.com	3,857	4.8
9 ticketmaster*	3,602	4.5
10 apple.com	3,421	4.3
13 dealtime.com	3,131	3.9
32 mysimon.com	1,915	2.4
bottomdollar.com	670	0.8

Source:
July 2000



5

Shopbot Trends



6

Problems with Shopbots

- Less than 10% of shoppers use shopbots, this is up from last year.
- Why don't more people use shopbots?
 - Lack of awareness
 - Lack of benefit (not enough price variation)
 - Lack of information about book (no reviews), you must already know the book
 - Slow response time (the modal time for pricescan and dealpilot is 45 seconds, amazon is <2 seconds)
 - Poor interface, displays too much information

7

The returns of search

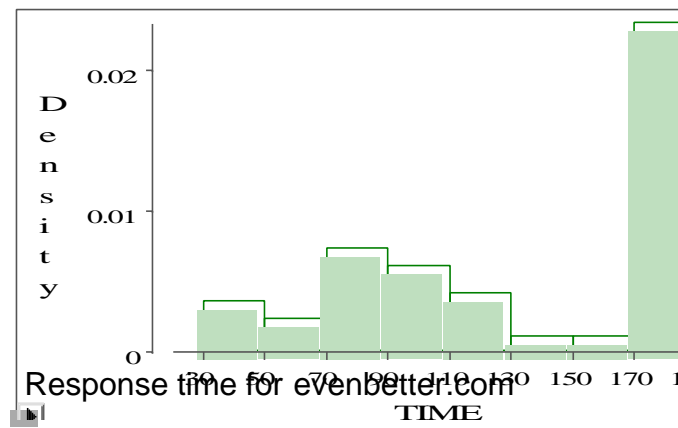
Price Dispersion

Kernel Density Plots on De-meaned Price Data — Books

Source:
Brynjolfsson
and Smith
(2000)

8

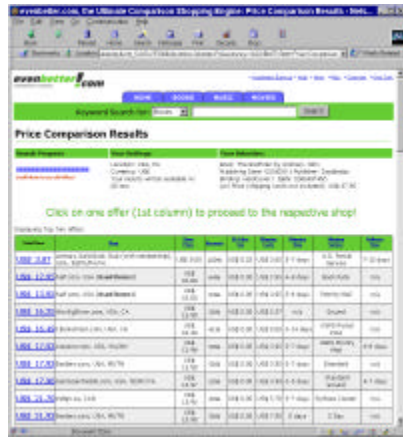
Distribution of Response Times



9

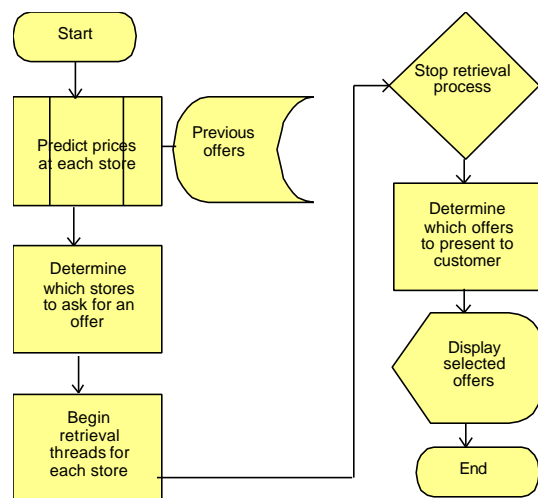
Improving Shopbot Design

Current ShopBot Design



- Enter a book and evenbetter.com will search 30+ bookstores and return all searches ordered by price
- Yahoo currently lists over 100 bookstores
- Makes search quick and simple
- Search is valuable
 - Average range is \$12
 - Amazon lowest only 5% of time (and dropping quickly)

Shopbot Operational Flow



Operational Decisions

- Which stores to search?
Shopbots can form prior expectations about prices that can help eliminate searching at high price stores
- How long to wait?
About 5% of store requests time out, also it may be better to interrupt searches at a certain point
- Which offers to present?
It is very cognitively taxing for consumers to have to search through scores of offers. Consumer research tells us they will use less efficient comparison rules.

13

Model limitations

- Treat this as a batch job
Problem becomes a sequential decision process
- Consumer already knows what they want
Imagine designing a shopbot that could find things the consumer did not identify (make tradeoffs in broader product classes)
- We know consumer preferences
Allow for random component, but assume that part-worths of utility function are known
- Do not explicitly consider shopbot costs
We consider shopbot costs only to the extent they impact waiting time (and therefore utility)
- Do not explicitly consider shopbot profits
We presume the shopbot wants to maximize consumer utility maximize purchase probabilities, however the shopbot may want to lead consumers to purchase specific alternatives

14



Modeling Consumer Utility



Modeling Consumer Interaction with a Shopbot

- Use a compensatory utility model to determine consumer's tradeoff between price, delivery, tax, and waiting time
- Consider the cognitive costs that a consumer incurs in making comparisons
- Use past information from previous web retrievals to intelligently retrieve prices

Compare this approach with IR models that assume noncompensatory rules

Utility Model

- Usual additive utility model for the i th product given P alternatives with A attributes in the set:

$$u_i = \overbrace{\sum_{j=1}^N b_{ij} a_{ij}}^{u_i^*} + e_i - \mathbf{x} \cdot T - \mathbf{w} \cdot Q - \mathbf{l} \cdot C$$

↑
↑
↑
↑

attributes: price,
delivery time, etc.
waiting time
= $\max(t_1, \dots, t_M)$
system
overhead to
process
requests
cognitive costs
(A-1)(P-1)

17

Utility of the Choice Set

- Utility of basket with P choices from M alternatives:

$$\begin{aligned}
 U &= \max(u_{(1)}, \dots, u_{(P)}) \\
 &= \max(u_{(1)}^* - \mathbf{x} \cdot T - \mathbf{w} \cdot Q - \mathbf{l} \cdot C, \dots, u_{(P)}^* - \mathbf{x} \cdot T - \mathbf{w} \cdot Q - \mathbf{l} \cdot C) \\
 &= \max(u_{(1)}^*, \dots, u_{(P)}^*) - \mathbf{x} \cdot T - \mathbf{w} \cdot Q - \mathbf{l} \cdot (A-1)(P-1)
 \end{aligned}$$

↑
present the best offers,
ordered observations

18

Formal Problem

- Sequential Optimization – solved backwards

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

Variables:

- q offers to query
- r offers retrieved
- p offers presented
- t^* time to interrupt query

19

Example

- Shopbot can search the following stores:
 $Price_1 \sim N(10,1)$, $Price_2 \sim N(11,1)$, $Price_3 \sim N(12,1)$
- Query stores 1 & 2
 $\mathbf{q}=[1 \ 1 \ 0]$, $t=[8 \ 12 \ 10]$
- Interrupt query at 10 seconds
 $t^*=10$, $\mathbf{r}=[1 \ 0 \ 0]$
- Present offer from store 1 to customer
 $\mathbf{p}=[1 \ 0 \ 0]$, $U=6$

Objective: Maximize utility of the set offered to consumer

Current Solution: $\mathbf{q}=[1 \ 1 \ \dots \ 1]$, $t^=30$, $\mathbf{p}=[1 \ 1 \ \dots \ 1]$*

20

Proposed Solution

Sequential Optimization Problem:

3. Which offers should be presented, given the retrieval set
2. When should they retrievals be interrupted, given the queries were made
1. Which stores should be queried

21

3. Which offers to present?

Assume that utility errors follows an extreme value distribution with parameters $(0, ?)$. Usual multinomial logit model. Implies that the maximum also has an extreme value distribution:

$$\max(u_{(1)}^*, \dots, u_{(P)}^*) = \mathbf{q} \ln \left(\sum_{i=1}^P \exp \left\{ \frac{\bar{u}_{R-i+1:R}}{\mathbf{q}} \right\} \right) + \mathbf{q} \mathbf{g}, \quad \bar{u}_i = \sum_{j=1}^A \mathbf{b}_{ij} a_{ij}$$

22

Which offers to present?

Can now determine expected utility (conditioned on retrieval information set)

$$E[U] = q \ln \left(\sum_{i=1}^P \exp \left[\frac{\bar{u}_{R-i+1:R}}{q} \right] \right) + qg - x \cdot T - w \cdot Q - I \cdot (A-1)(P-1)$$

Solution:

Start with $P^* = R$

Stop if $E[U|P-1] < E[U|P]$

Let $P = P-1$

23

Special Case

Suppose prices are all identical, how many offers to present?

$$P^* = \frac{q}{I(A-1)}$$

↗ variability (gain to added item)
↘ cognitive cost

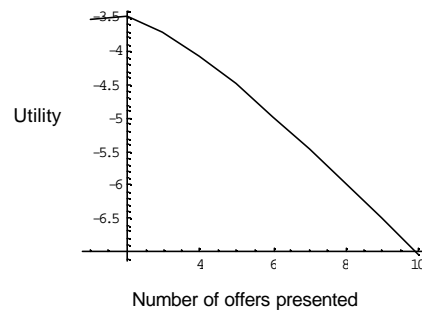
Example

- Average book generates 10 utils with s.d. of 2 utils, i.e., ($U=9.1, \sigma=1.6$), Book has 4 attributes ($A=4$)
- $\sigma=.1$ $P^*=5.3$, $\sigma=.2$ $P^*=2.7$
- Offer set of 20 books: 8.7 utils
- Offer set of 5 books: 11 utils

24

Implications

Do not show all the results! Utility can decline quickly as a result of added items.



25

2. How long to wait?

Just because a store is queried doesn't imply that it will respond, there is a probability h_i of no response, and let t_i represent time to retrieve offer. If $t_i < t^*$ then observation is censored.

Probability of no response:

$$t_i = h_i + (1 - h_i) \Pr[t_i > t^*]$$

Assume the probability of response independent across stores and also retrieved offer.

26

How long to wait?

Must now evaluate utility over all possible sets of retrievals based on the queries made.


$$\sum_{i \in \Omega} \left(\prod t_i^{t_i} (1 - t_i)^{1 - t_i} \right) \cdot E[U_i]$$

Ω is set of all possible combinations, dimension is 2^Q .
Combinatorial explosion, $Q=10$ yields 1,024 combinations, $Q=30$ yields one billion.

27

Implications

Suppose times are gamma distribution (mean=1.5, std=2.7), $E[\max \text{ of } 10 \text{ variates}] = 7$, chance of retrieval=95%.



Act aggressively and truncate response

28

1. Which stores to query?

At this stage the shopbot does not know the price it will retrieve, however it can guess (in fact pretty well).

Order stores based on prior expectations, and the problem now becomes how many stores to query?

29

Example

Suppose there are three stores that may be queried and prices are normally distributed:

$$Utility_1 \sim N(1,1), \quad Utility_2 \sim N(0, s^2), \quad Utility_3 \sim N(0, s^2)$$

If you could only select two stores which ones will yield the expected maximum utility?

Choose {1,2} if $s < 3.67$, otherwise {2,3}

$$E[\max(Utility_1, Utility_2)] = \mu_X \Phi(\Delta/v) + \mu_Y \Phi(-\Delta/v) + v\phi(\Delta/v)$$

$$\text{Where } \Delta = \mu_X - \mu_Y \text{ and } v^2 = \sigma_X^2 + \sigma_Y^2 - 2\rho\sigma_X\sigma_Y$$

30

Which stores to query?

Unfortunately, while normality may be a good distributional assumption in practice, theoretical properties of its order statistics are not tractable.

A reasonable approximation is to assume that utility (prices) are logistically distributed. Furthermore to yield an analytical result we assume that the stores are i.i.d.. After some work (and approximation) we get the following stopping rule:

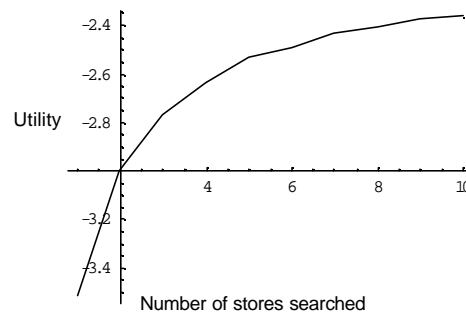
$$Q \geq \frac{s}{w + I(A-1)(P_{R+1}^* - P_R^*)}$$

Increase set size as the gains to search are higher (s) and reduce them as the disutility of waiting time increases.

31

Implications

Suppose there are 10 stores (utility ranges from -4 to -5, std=1)



Larger query size is better, but at a declining rate

32



Application to Online Bookstores



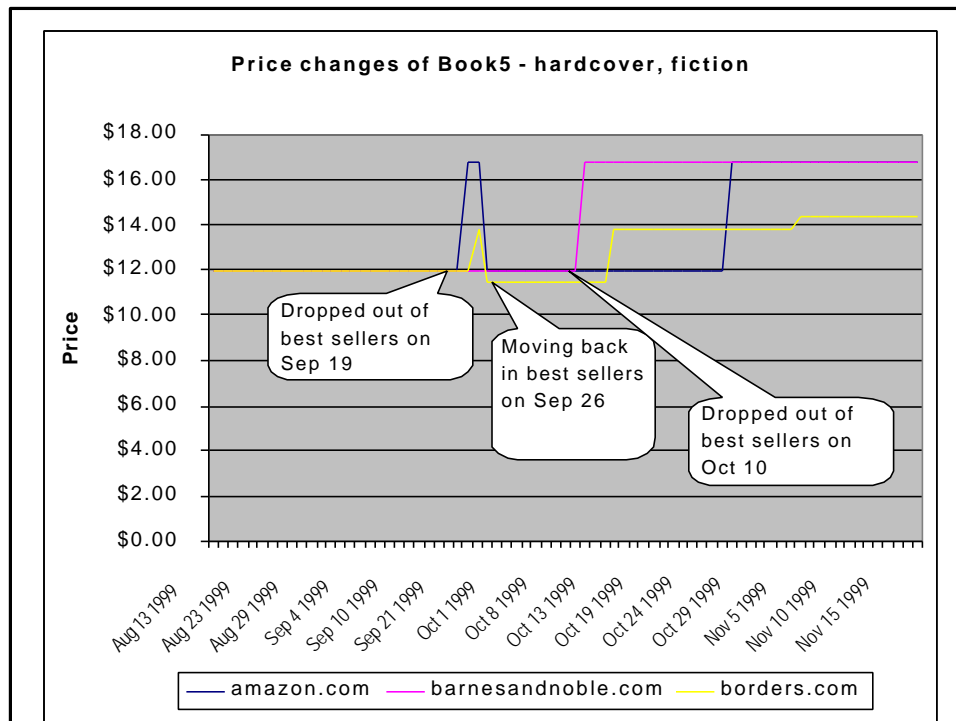
Empirical Components

- Predicting Price Changes
- Response Times for Searching
- Consumer Utility Part-worths

Data


- Automated agents collected data from 2 shopbots and several individual stores
- August 99 - January 00
- 600 books
 - NY Times Bestsellers
 - Randomly selected ISBNs
 - Computer Books

35



Time between price change

Number of days between price changes follows an exponentially declining distribution



Predicting Prices

- Predict days between price changes using a Negative Binomial Model with parameters (γ, δ) , where:

$$\ln(\gamma) = \alpha_0 + \alpha_1 \text{days_since_bestseller_change} \\ + \alpha_2 \text{1BookStreet} + \alpha_3 \text{amazon} + \alpha_4 \text{bn} \\ + \alpha_5 \text{buy.com} + \alpha_6 \text{borders}$$

- Given that prices have changed predict the magnitude of price change using an autoregressive model.

$$\text{RelPrice}(t) = \beta_0 + \beta_1 \text{RelPrice}(t-1) + \beta_2 \text{uphard} \\ + \beta_3 \text{uppaper} + \beta_4 \text{downhard} + \beta_5 \text{downpaper}$$

38

Summary of Results

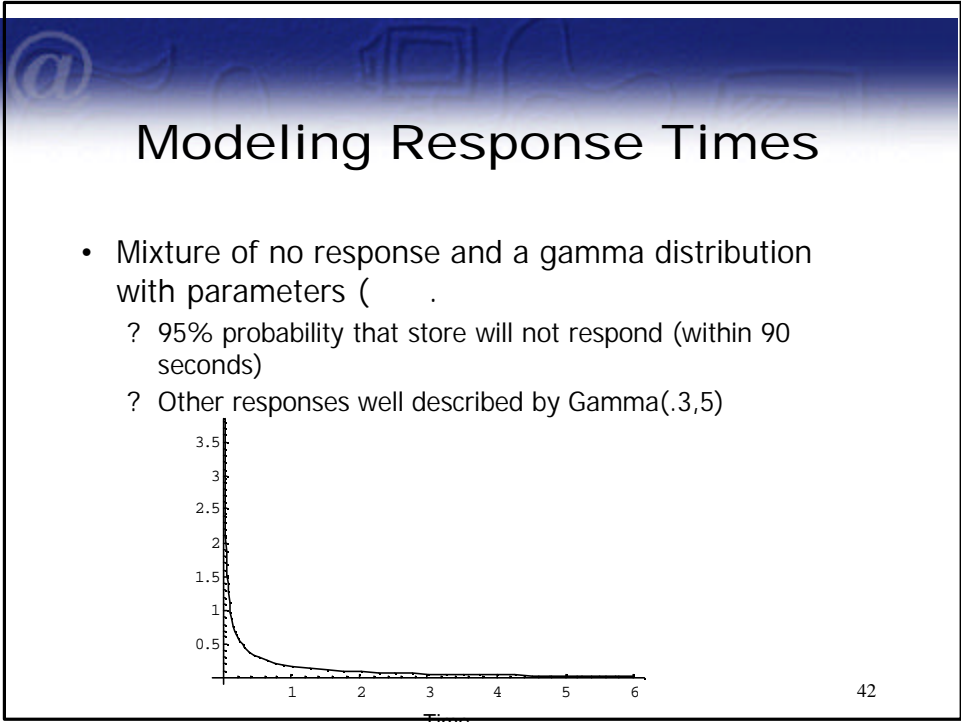
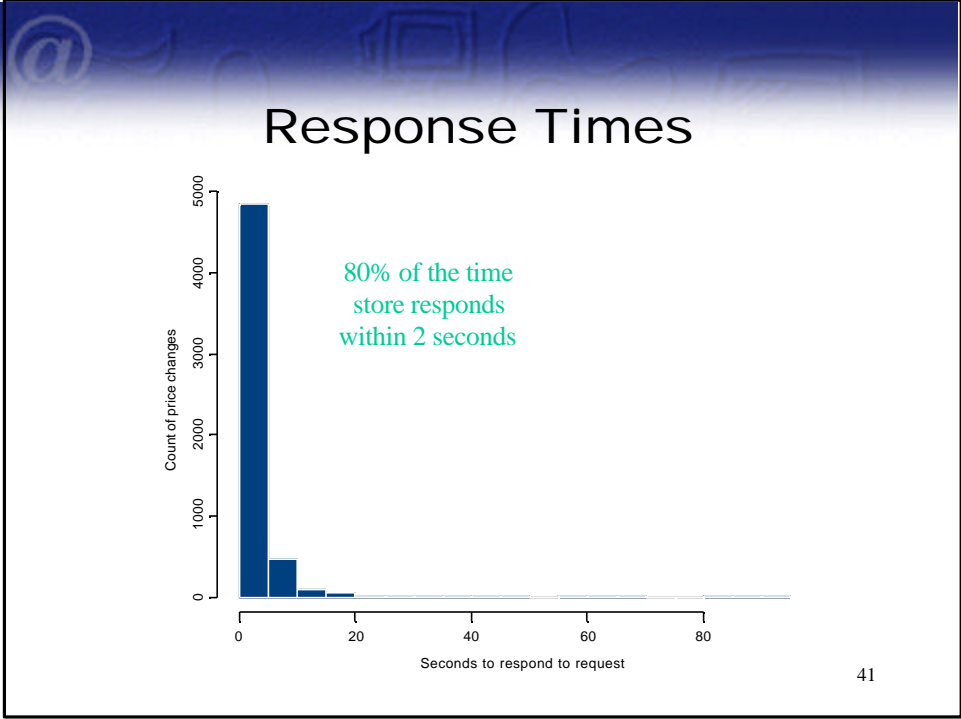
- Amazon and Barnesandnoble responding quickly to change in bestseller status
- Amazon shows some price leadership, but for the most part weak relationships between price changes at stores
- When books move onto the bestseller list prices drop (more for hardcovers)
- When books move off the bestseller list prices rise

39

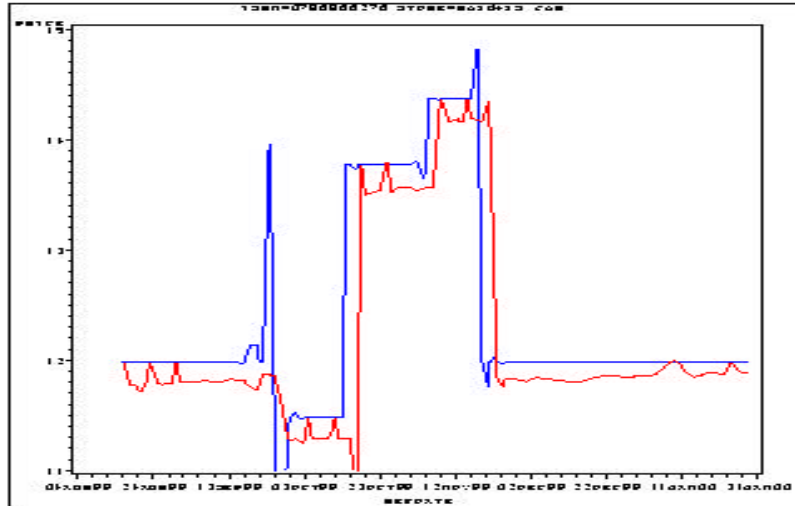
Correlation between Actual and Predicted Prices

Price Collection Frequency	Correlation
Only once (initial time)	.30
Once every 30 days	.82
Once every 14 days	.91
Once every 7 days	.95
Once every 3 days	.99
Once every day	1.00

40



Predicted Price Plot



Shopbot Choice Model

<u>Parameter</u>	<u>Estimate</u>	
<i>Total Price</i>		
Item Price	-.19	(\$1.00)
Shipping Price	-.37	(\$1.95)
U.S. Tax	-.43	(\$2.26)
<i>Delivery Average</i>	-.02	(\$.10)
<i>Delivery "n/a"</i>	-.37	(\$1.94)
<i>"Big 3"</i>		
Amazon	.48	(\$2.52)
BarnesandNoble	.17	(\$.89)
Borders	.27	(\$1.42)

44

Waiting/Cognitive Costs

- Disutility of waiting one second = \$.01 (?=.01)
 - If you make \$70,000/year, your time is worth \$.01/second
- Overhead for launching a thread to search an online store = 10 milliseconds (? =.0001)
- Cost to make a comparison = \$.05 (?=.05)
 - To compare 5 items with 4 attributes implies a total cognitive cost of \$.60

45

Simulation Results

Which stores to query?

<u>Store</u>	<u>Mean</u>	<u>Std.Dev.</u>	<u>Store</u>	<u>Mean</u>	<u>Std.Dev.</u>
buy.com	.52	.10	varsitybooks.com	.75	.05
booksamillion.com	.59	.12	1bookstreet	.76	.13
Bookbuyer's Outlet	.62	.13	bigwords.com	.77	.06
Borders.com	.62	.13	WordsWorth	.83	.10
alldirect.com	.63	.05	booksnow.com	.88	.06
Amazon	.63	.13	Cherryvalleybooks	.89	.02
barnesandnoble.com	.63	.13	Rainy Day Books	.89	.05
AlphaCraze.com	.64	.09	Rutherfords	.89	.05
Fatbrain	.65	.15	Classbook.com	.96	.06
Books.com	.70	.09	Baker's Dozen online	.99	.06
HamiltonBook.com	.70	.07	Book Nook Inc.	.99	.05
BCY Book Loft	.72	.07	Codys Books	.99	.06
kingbooks.com	.73	.04	computerlibrary.com	.99	.06
A1 Books	.75	.06	page1book.com	.99	.07

47

Delivery Options for BarnesandNoble.com

<u>Service</u>	<u>Delivery</u>	<u>Cost</u>
U.S. Postal Service	5-9 days	\$3.95
Standard Ground	4-7 days	\$3.99
FedEx Second Day	3-4 days	\$7.95
UPS 2nd Day Air	3-4 days	\$7.99
FedEx Overnight	2-3 days	\$10.95

48

How long to wait?



49

Best set to offer

(if prices known)

<u>Store</u>	<u>Service</u>	<u>Delivery</u>	<u>Price</u>	<u>Cost</u>	<u>Total</u>
1BookStreet.com	USPS Parcel Post	6-21 days	\$15.19	\$0	\$15.19
Amazon.com	USPS Priority Mail	5-10 days	\$12.59	\$3.95	\$16.54
Buy.com	Standard Shipping	n/a	\$10.39	\$3.95	\$14.34
Borders.com	Standard	5-10 days	\$12.39	\$3.90	\$16.29

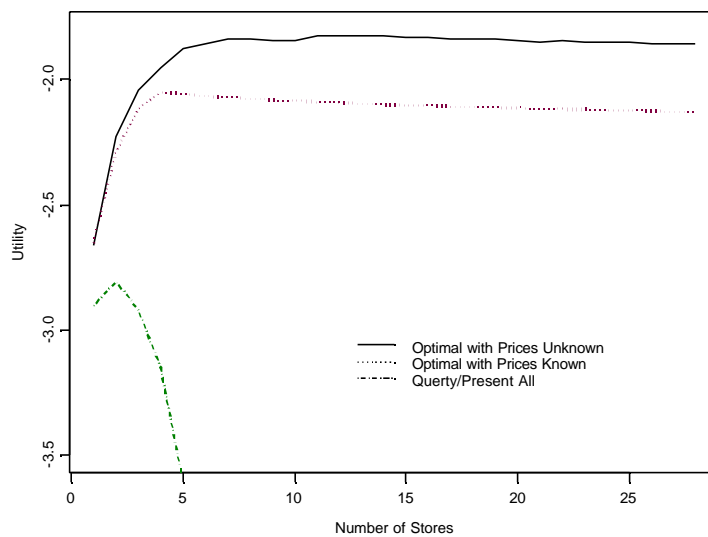
50

Best set to offer (if prices are forecasted)

	Probability	Store	Service	Delivery	ExpPrice	Cost	ExpTotal
Still include these	82%	1BookStreet.com	USPS Parcel Post	6-21 days	\$ 15.19	\$ -	\$ 15.19
	56%	Amazon.com	USPS Priority Mail	5-10 days	\$ 12.59	\$ 3.95	\$ 16.54
	51%	Buy.com	Standard Shipping	n/a	\$ 10.39	\$ 3.95	\$ 14.34
	42%	Borders.com	Standard	5-10 days	\$ 12.39	\$ 3.90	\$ 16.29
Chance of including these	34%	barnesandnoble.com	Standard Ground	4-7 days	\$ 12.59	\$ 3.99	\$ 16.58
	28%	barnesandnoble.com	U.S. Postal Service	5-9 days	\$ 12.59	\$ 3.95	\$ 16.54
	26%	booksamillion.com	Standard Ground	N/A	\$ 11.79	\$ 3.95	\$ 15.74
	21%	Fatbrain.com	UPS Ground	4-8 days	\$ 12.99	\$ 3.95	\$ 16.94
	17%	AlphaCraze.com	USPS Special Rate	5-15 days	\$ 12.79	\$ 3.50	\$ 16.29
	13%	Bookbuyer's Outlet	Standard	n/a	\$ 12.39	\$ 4.50	\$ 16.89

51

Effects on Utility



52

Findings


- Present small number of the best alternatives (~4 items)
- Only need to search small number of stores if prices known with certainty
- If prices are unknown, it's best to continue searching at up to 14 stores. Since not all results presented (and little computational cost) go ahead and search more stores.
- Using traditional shopbot best to only search a couple of stores
- Probability consumer will prefer optimal shopbot over current shopbot >99.9% or \$2.00, or worth \$.20 compared with best result when all results presented from 2 stores

What happens if customers less sensitive to cognitive costs?

53

Utility

(with Low Cognitive Costs)



54



Conclusions



Summary

- Intelligent design of shopbots can dramatically increase the utility that consumers garner from their use
- Instead of passively searching, can incorporate information about utility and price expectations to speed up search and satisfaction
- Incorporates cognitive effort, compensatory utility functions, and information retrieval

An improved shopbot?



- Ask users for filtering questions about preferences or use information from previous history
- Appropriately balance the cost of asking for the information with its benefits
- Allow further search
- Better understand how consumers perceive waiting time based on expectations, provide 'filler' tasks

57

Future Directions

- Learning from past purchases and designing active shopbots (versus the passive design presented)
- Need better understanding of how to quantify cognitive costs and effects of waiting
- Train shopbots to be proactive in seeking out good deals. If bestseller status changed today & shopbot knows a store responds to status change in 2 days, it can make recommendations ("wait 2 days and price at amazon likely to be less by \$10")
- Identify baskets of products or more complex products like travel (airline tickets, car rental, hotel, etc.)
- Applications to information goods (e.g., news stories, recommender systems, search engines)

58

Reflections on E-Commerce Research

- Most of the existing marketing literature focuses on describing how consumers behave
- What is really needed is prescribing how agents should behave so they can work with and/or replace consumers

Can yield some new insights into old problems (e.g., how do consumers shop?)

59

Store Choice Analogy

Which stores to search? → Which store to shop at?

How long to search? → How far to travel?

Which products to present? → How does assortment affect choice?

60

