Repeated Price Searches

A. Norman, M. Drake, C. Govil, K. Kuang, S. Keyburn, S. Ler
K. Powers J. Sanghai, J. Schieck, J. Sussman, and K. Yan

December 18, 2006

Abstract

Most students at graduation have bought textbooks for at least 8 semesters or 12 quarters. We surveyed seniors to determine the characteristics of those who buy at the campus bookstore versus those who buy books online. To see how online students improve their performance, we created a simple Bayesian multiperiod optimization problem with an unknown parameter (cheaper store). Subjects did not compute the optimal strategy, but did demonstrate some intuitive understanding of the problem. In order to judge the optimality of actual student strategies, we recorded online prices for 24 days. Performance varied among students. Some were near optimal.

Send correspondence to:

Alfred Lorn Norman
Department of Economics
The University of Texas at Austin
Austin, TX USA 78712
E-Mail: norman@eco.utexas.edu
Phone: 512-346-4253

1 Introduction

Consumers search repeatedly for similar products, such as food, clothing, and textbooks. We examine student textbook search procedures of The University of Texas at Austin (UT) upperclassmen who seek cost savings by buying their books from many sellers rather just buying their textbooks from the campus bookstore, the UT Co-op.

In Section 2, we provide an overview of the textbook market at UT. We use four surveys to characterize those students who go directly to the UT Co-op and those who shop around. The former are willing to pay higher prices in return for the convenience of the Co-op; others search many stores and Web sites to obtain cheaper prices, but with a risk of a late delivery and less quality than stated. The price searchers report saving about 40% relative to the Co-op prices over time. Our question is to what extent they are capable of creating an optimal strategy to achieve this goal.

In Section 3, we formulate a simple Bayesian optimization problem to determine whether subjects are capable of creating an optimal strategy for searching. This problem is a variation of estimation and control problems studied by a group of economists starting in the 1970s. For example, see Chow (1981), Norman (1976), Norman and Jung (1977), and a survey of this literature by Kendrick (2005). We begin the section by showing why previously studied estimation and control problems are too complicated even if the time horizon is only two periods. We then create a computationally tractable model using a binomial distribution and a small number of alternatives, resulting in a problem whose computational complexity is quadratic in the time horizon.

In Section 4, we discuss our experimental design. Subjects solve four search problems before answering survey questions on their own search procedures. All problems require the subject to maximize her expected savings in selecting a textbook from one of two stores for either two semesters (problem 1), or four semesters (problems 24). To determine an optimal strategy, the subject must calculate how many semesters she will check prices before she stops bothering to shop around. The number of arithmetic operations is less than 10 for the two-period problem and less than 50 for the four-period problems. The survey questions focused on determining whether subjects had an intuitive understanding of the problems.

In Section 5, we discuss the results of our first experiment. Only six subjects determined the correct strategy for the first problem. We conclude that most didn't understand the problem–or, if they did, they couldn't perform the calculations correctly. But subjects did demonstrate an intuitive understanding of the second, third, and fourth problems.

In Section 6, we present the results of our fourth questionnaire in which we asked graduating seniors to recommend to freshman the sites they should check for low textbook prices. We also checked the prices for economic textbooks for all undergraduate classes for a period of 21 days for a new, good used, and usable used textbook. In the case of a new and good used textbook the student would obtain the low cost textbook over 90% of the time if they just checked two sources, Half.com and Amazon.com. For usable used textbooks a freshman should check more sites. How close to optimal are the seniors recommendations? The trade-off between price and risk complicates this analysis. Nevertheless, a freshmen, without solving a dynamic optimization problem, can develop a near optimal search strategy by talking to a small number of seniors who buy their textbooks online from many sources.

In Section 7, we conclude. This is the fourth in a series of papers to define a procedural consumer. The others are Norman et al (2001), Norman et al (2003), Norman et al (2004), and Norman et al (2006)

2 Textbook Market and Buyers

In this section, we characterize the UT textbook market and student participation therein.

Students buy textbooks each semester. The professor usually defines exactly which materials are needed for the class, but the student still faces a wide number of textbook choices. Widely used textbooks are frequently available not only in U.S. editions, but also in less expensive softcover foreign editions printed in color on quality paper and in much less expensive softcover foreign editions printed in black and white on newspaper quality paper. In many cases it is illegal to sell these foreign editions in the U.S. Sellers also offer used copies of each edition in varying states of disrepair.

If a student buys from the UT Co-op, she can buy a new U.S. edition at the publisher's list price or a used U.S. edition at 75% of the list price, regardless of the condition of the used book. She can return the book for up to 12 class days, which is important because many students add or drop classes. If she keeps the book, she can sell it back to the Co-op at the end of the semester for half of its new price (regardless of whether the copy she purchased was new or used) if a professor has requested the book for the next semester. At the end of the academic year, she will also receive a 10% rebate towards Co-op purchases, but the most important characteristic of this seller is convenience.

Most of the online sites that students use to purchase textbooks are what we call marketplace websites. In contrast to a traditional, direct sales site like Amazon.com, marketplace sites list third-party retailers, who describe their offering, set a price, and are rated by their previous customers. Amazon Marketplace, Half.com, BookByte.com, and AbeBooks.com are examples of this genre. The most important characteristic of marketplace sites is their low cost, but there is also a risk that a faceless seller will not send a book in good time or will fail to accurately describe his product. Recognizing this problem–even among sellers rated by previous customers–most marketplace websites refund purchases that fail to arrive, although refund policies vary from site to site.

Another variety of product search site are meta-search sites, such as PriceGrabber.com, CampusBooks.com, Bigwords.com, and Froogle.com. These search a variety of sales sites to provide a list of sellers ordered by price. By using such a site, a consumer can search a wide variety of marketplaces for his desired product, covering a broader selection of sellers with less effort. These sites are not so widely used, in part because they may be less well known, and in part because they suffer from spurious results–as with sites that say they charge one dollar for a book, provided the consumer jumps through a series of hoops in the form of arcane point-trading programs or the student goes to a site to find the low cost book indicated by the meta-search site only to find that the book is no longer available.

To determine how students bought their textbooks, we surveyed UT upperclassmen every semester between Fall '04 and Spring '06. For the first three surveys, we offered each student five dollars for completing the survey. For the fourth, students who stated that they only shop at the Co-op were offered a one-page survey to complete for three dollars, whereas students who shop around were offered a four-page survey they could complete for six. The Fall '05 survey included a brief essay detailing the factors why some students choose to buy textbooks only at the UT Co-op while others search many sellers.

From these questionnaires we have created the following table:

Table 1: Percent of the total cost of schoolbooks paid by the following sources:

Source	Co-op: See a below ($n = 42$)	Many Sellers: See a below ($n = 71$)
Money you earned	11.07% (4)	25.00% (9)
Parents and relatives	69.52% (25)	53.94% (21)
Scholarship	8.93% (2)	7.54% (3)
Loan	8.10% (0)	11.41% (3)
Trust	2.38% (1)	2.11% (1)

where for a the % is the mean from that source and the number in the () is the number who obtain 100% of their funding from that source. From questionnaire 4 we present data on parental incomes:

Table 2: Parental Income Brackets

Brackets	UT Co-op	Many Sellers
0-$70K	11	15
$70,001-100K	2	9
$100K+	26	14

The differences between groups with different funding sources are smaller than might be expected. Students who earn their own money might be thought to be more money-conscious, but 40% do not shop for the best price; students who get money from their parents might be thought to be less concerned about textbook price, but 40% shop at multiple websites. Also, 11 out of the 16 students with the lowest household income do not comparison shop.

The reasons for shopping only at the UT Co-op are convenience, risk avoidance, and lack of incentives to shop elsewhere. In the first questionnaire, 78% of students who shopped only at the Co-op said that convenience was the most important factor in their purchasing decision. Seventeen out of the twenty-five students who shop at the Co-op and whose parents cover their textbook costs said that their parents give them no incentive to shop elsewhere. As not all of these respondents said that they were worried about delivery risk, the only concern for several of these students was convenience. Moreover, between questionnaires one and four, we found that 69% of Co-op shoppers said that they wanted to avoid the risk of receiving their books late or not receiving their books at all. Adding these two groups together, we have explained 36 of the 42 Co-op shoppers. Further adding the students who wanted to visually examine textbooks before purchasing, we explain 41 out of the 42 Co-op shoppers' choices through a combination of risk avoidance, lack of incentives to shop elsewhere, and convenience. Interestingly, nine of the eleven students who were in the lowest income bracket and purchased their books from the Co-op indicated that they purchased their books from the Co-op to avoid late delivery.

Now consider the many sellers group. Ninety-five percent of the comparison-shopping students in the first questionnaire indicated that price was their most important concern in textbook shopping. When asked what incentives they have to buy books from many sellers, the three most common responses were: (1) I can spend money I save on other things ($n = 29/38$); (2) I respect my parents and do not want to waste their money ($n = 25/38$): and (3) my parents nag me to buy cheap($n = 5/38$).

3 Estimation and Control Problems

In this section, we outline a multiperiod textbook-buying problem. A small group of economists have studied similar estimation and control problems since the 1970s. These linear optimization problems are often intractable, as shown in Norman (1994). We present an example in Appendix A to demonstrate how such a problem can appear simple, yet be noncomputable.

Almost all economists ignore the procedures to solve these optimization problems, using Milton Friedman's parable of the billiard player(Friedman and Savage (1948)) to denigrate them. Friedman is right to indicate that we have a natural talent for judging distance, velocity, and acceleration; all mammals involved in hunter-prey interactions must have these talents. But that doesn't mean billiards is an innate skill. While an expert billiards player does not integrate the equations of motion, he spends hours and hours in practice adjusting the weights of his human neural network to become an expert. Similarly, people are naturally capable of processing choice problems, but they aren't born knowing the optimal route to find the cheapest cereal that gives them the breakfast experience they desire. Thus, our current assumptions about procedures need revision. How do people actually find things? Are these procedures optimal? Man's processing capabilities will be revealed by the ongoing research into economics and the brain, but there are clear limits. Humans don't multiply nontrivial $n \times n$ matrices in their heads–a process whose computational complexity is at least quadratic. Indeed, if man had polynomial processing capability, there would be no need for computers. Our research–Norman et at (2004)–suggests that humans, when they must make explicit calculations, are sublinear processors.

We designed a multiperiod price search experiment with four trials where the amount of computation did not exceed 50 arithmetic operations for any of the trials. In these trials, the subject needs to buy a new textbook each semester for either two- or four-semester periods. She can choose from two textbook stores, each costing \$5 in time and transportation costs to visit. For the sake of simplicity, it's assumed that she will not incur further transportation costs after visiting each store once. The subject is told that one of the stores is cheaper by \$X 80% of the time and the other store is cheaper by \$X 20% of the time, but she isn't told which store is which. The values of \$X for the four trials in order processed were \$30, \$10, \$20, and \$15.

In each semester, the subject must choose to go to either store A, store B, or both stores A and B. If she goes to store A, she incurs a cost of \$5 and buys the textbook without finding out its price at store B. The converse is true if she goes to store B. If she checks both stores, she will incur a search cost of $10 and buy from whichever store is cheaper. As the subject doesn't know which store will usually be cheaper, she should use Bayesian reasoning. If the subject checks both stores twice and twice finds A to be the cheaper store, the probability that A will be cheaper 80% of the time is 0.94. Given her initial lack of knowledge, the subject should always check both stores in the first semester. After that, her optimal strategy depends on the price difference and the length of the time horizon.

As proved in Appendix B, this problem's computational complexity is at least quadratic in the number of semesters to consider. We only considered one source of information (personal experience) to keep the computations as simple as possible. If subjects can correctly solve a one-source problem, we must consider two-source problems such as students also asking advice from other students who shop online. If they can't, there's no such need.

4 Experimental Design

Our experiment had each subject solve four problems and answer questions on a fifth page. This was informed by a preliminary experiment that involved eight semesters and required more calculations to compute the optimal strategy than most subjects were able to perform.

Subjects were asked to determine the strategy to maximize expected savings in two- and four-semester problems. In each problem, the student faced two stores. Every semester, one of the two stores would be cheaper than the other by \$X. The subject was told that one store had a 0.8 probability of being cheaper in any given semester, but wasn't told which store. The subject was also told that his costs would increase by \$5 for every store at which he checked prices in a semester: he therefore had an incentive to determine which store would usually be cheaper. The experiment was prefaced with an instructions page that carefully spelled out the details of the experiment and clarified them with an example; this page is in Appendix C.

The characteristics of the four problems to be solved by the subjects are shown in the table below:

Table 3
Characteristics of Four Problems

Problem	Stores	Cheaper Store	Number of Semesters	Price Difference
1	A and B	A	2	\$30
2	C and D	D	4	\$10
3	E and F	F	4	\$20
4	G and H	G	4	\$15

The cost of checking prices at a store is always \$5. Expected savings as a function of strategy is shown in the table below, where "Chk n" is the number of semesters that the subject checked prices at both stores and "Calc" is the number of arithmetic operations to compute the indicated results for the row:

Table 4
Expected Savings

Problem	No. Semesters	Price Diff	Chk 1	Chk 2	Chk 3	Chk 4	Calc
1	2	\$30	\$35.4	\$40	N/A	N/A	9
2	4	\$10	\$5.40	\$4.18	\$2.78	\$0.00	43
3	4	\$20	\$35.8	\$39.95	\$40.57	\$40.00	49
4	4	\$15	\$20.60	\$21.14	\$21.06	\$20.00	49

As indicated in Table 4, the optimal strategies for each of the four problems are:

1. Check both stores twice.

2. Check both stores once. For the remaining three semesters, go only to the store that had the lowest price in the first semester.

3. Check both stores three times, then go to the store which was more frequently indicated to have lower prices.

4. Check both stores twice. If the same store is cheaper both times, go to that store in the remaining periods. If not, check both stores in the third period, and return to the two-time-lower-priced store in the fourth period.

The subjects for this experiment were 25 students from an author's freshman economics class. We offered a flat fee of \$10 for participating and as much as \$25 more for answering questions correctly. Each question's reward was indicated on its page. In the preliminary experiment, we asked subjects whether the probability that the best strategy identifies the cheapest store was 0.7, 0.8, or 0.9. Only 12.8% chose the correct answer (0.9). We thus expected no more than twenty percent of our subjects to select the optimal strategy in any one period, but we didn't know whether subjects' intuition would lead them in the optimal direction. For example, would subjects check both stores more often for the third problem than the fourth problem? To gain further insight into subjects' intuition, we asked the following qualitative questions.

Answer questions 1, 2, and 3 below from your experience with the previous 4 problems. There is a correct answer for each and each correct answer will raise your score by one point.

1. The greater the price difference, the more times you should check prices at both stores: __ True, or __ False
2. If, on the 2nd, 3rd, and 4th problems on the previous pages, we had used 8 semesters instead of 4, it generally, but not always, would have been wise to check prices at both stores more often than you did for the equivalent 4 semester problem: __True, or __ False
3. If you check prices at both stores twice and you get the same store as having the lower prices twice, you should be more confident that this store is actually the lower priced store 80% of the time than if you only checked prices at both stores once. __ True or __ False

5 Results

Our results indicate that, although subjects don't make optimal decisions, their intuition concerning the qualitative differences in the optimal strategies is generally correct.

Let us start by considering the results for the first problem, which had two semesters and a price difference of \$30. Only six out of the twenty-five subjects chose this problem's optimal strategy: checking prices at both stores in both semesters. Computing this strategy required few calculations. If the subject checked prices at both stores in both semesters, his savings are $2 \times (\$30 - 2 \times \$5) = \$40$. If the subject only checks both stores once and then returns to the store that had the lower price in the first semester, his expected savings are $(\$30 - 2 \times \$5) + (0.8 \times 0.8 + 0.2 \times 0.2) \times \$30 - \$5 = \$35.40$. Given how many subjects failed to optimize their savings, we conclude that subjects have little quantitative understanding of the problem.

But most subjects correctly answered our first three qualitative questions, as shown in the table below:

Table 5
Responses to questions 1-3

	Question 1	Question 2	Question 3
Correct Ans	True	True	True
No. True	21	15	20
No. False	4	10	5
Significant	Yes	No	Yes

The last row in the table indicates whether the responses were significantly different from random at an α of 0.001. We therefore conclude that the subjects have some qualitative understanding of textbook search: they realize that they should check both stores more often as periods get longer and as price differences increase. This is reflected in subjects' performance in problems 2-4 (each four semesters long, and with price differences of, respectively, \$10, \$20, and \$15):

Table 6
Problems 2-4

	Problem 2	Problem 3	Problem 4
Price Diff	\$10	\$20	\$15
Optimal Checks	1	3	2.32
No. Correct	10	3	2
Avg Checks	0.88	1.52	1.48
No. Skips	7	4	5

Most subjects failed to check both stores the optimal number of times. Some even failed to check both stores in the first semester, as indicated by the "No. Skips" row. But the real issue isn't whether or not they pick the optimal strategy; we want to know if their intuition leads them toward the best strategy.

We thus examine whether the differences in "Avg Checks," the average number of times subjects checked both stores, are statistically significant. Problems 2-4 can be analyzed using Anova with one factor, the difference in prices. The results are shown below:

Table 7
Anova results for problems 2-4

Source	DF	Anova SS	Mean Square	F Value	Pr > F
Avg Checks	2	6.42666667	3.21333333	4.69	0.0138

The Duncan test at $\alpha = 0.05$ shows that the average number of two-store checks for Problem 2 is different from that of Problem 3 and Problem 4. Subjects may have little quantitative understanding of these problems, but they do have qualitative insight into how many times to check prices and for this problem the average number of price checks of both stores is near optimal. El-Gamal and Grether (1995) found in a static experiment that the rule subjects used the most was the Bayes rule. This experiment suggests that subjects have an intuitive understanding of the implications of Bayes rule in a dynamic context without making (or correctly making) the calculations.

6 How Close to Optimal

We are interested in whether consumers converge to an optimal purchase strategy. As we and numerous previous experiments have shown–for example, see Edwards (1968)–people need more evidence than do perfect statisticians to determine the parameters of a model. Again, we would expect an intuitive optimization procedure to converge more slowly to an optimal terminal condition than the optimal strategy. But we assert that student online textbook buying strategies approach optimality during their college years.

We performed four surveys of students who bought their textbooks from many sellers. Of the 92 students in these surveys, the first year they bought textbooks from many sellers was:

Table 8
First year to buy textbooks from many sellers

	First Year	Second Year	Third Year	Fourth Year
Number	28	36	22	6
Cum %	30.4	69.6	93.5	100

Our fourth questionnaire's subjects bought books online for an average of 5.2 semesters and summer sessions. If students' search procedures converge to near optimal, it happens quickly! In Questionnaires 2, 3, and 4, we investigated how students learned which sellers sell textbooks and among these who is likely to have the cheapest prices. We asked the 92 students of Questionnaires 2, 3, and 4, "How did you find out about sellers other than the UT Co-op (i.e., radio advertisements, search engines, friends, etc.)?" Their responses are shown in the following table:

Table 9
Data sources for students using many sellers

Source	Number	%
Friends/Relatives	75	82
Search Engines	43	47
Advertisements	24	26
Professor	6	7

Students learn about sellers over time, so one aspect of their search strategy is the number of sellers' prices they check. The following table shows the number of sellers other than the UT Co-op at which the fourth questionnaire's subjects checked prices.

Table 10
Number of sellers at which subjects checked prices

Number Sites	1	2	3	4	5	6	7
Number Students	1	7	8	12	6	2	1

By using a tabbed browser and textbook ISBN numbers listed on the UT Co-op's website, a student can quickly check prices at a large number of book websites. It would thus appear that subjects' search procedures were suboptimal. But consider the prices we found for textbooks between 30 December 2005 and 19 January 2006 for all 23 of the undergraduate economics courses on the sites listed in Section 2 (except for TexBooks, which was not activated until 9 January, after which we started checking it, too). We considered three quality levels of textbooks: (1) new U.S. edition; (2) good quality U.S. edition with no missing pages, highlighting, or writing; and (3) lowest priced used book, ignoring whether or not it was a U.S. edition. In the table below, we show the frequency that sellers had the lowest price in each of the three categories.

Table 11
Cheapest sites in price survey

Site	New	%	Good Used	%	Usable	%
AbeBooks.com	0	0	8	2	13	3.2
Amazon.com	180	43	175	42	87	21
BookByte.com	0	0	0	0	32	8
eBay.com	18	4	4	1	28	7
Half.com	212	51	199	48	193	48
Texbooks.com	0	0	6	1	35	9
UT Co-op Online	8	2	24	6	17	6

For each of the three categories, either Half.com or Amazon.com was the cheapest seller 94, 90, and 69 percent of the time, respectively. Hence a student interested in either a new or good quality used U.S. text would have a near-optimal strategy for buying economics books if she just checked two sources other than the UT Co-op.

In survey four, we gave the participants a list of sites and asked them to recommend sites at which they would recommend freshmen check textbook prices. We also asked them at which sites they had checked prices, from which they had bought, and of which they were previously unaware.

Table 12
Recommendations from seniors (n=38)

Source	Recommended this site	Checked this site	Bought from site	Unaware of site
SE1: PriceGrabber.com	2	13	0	21
SE2: Froogle.com	5	13	1	20
SE3: BigWords.com	1	5	2	27
SE4: CampusBooks.com	11	20	6	13
OL1: AbeBooks.com	12	16	11	20
OL2: Alibris.com	1	7	2	28
OL3: Amazon.com	31	35	28	0
OL4: BookByte.com	2	7	4	24
OL5: eBay.com	24	34	20	0
OL6: Half.com	31	34	31	2
OL7: Texbooks.com	7	14	7	19
OL8: B & N Online	4	30	5	1
OL9: UT Co-op Online	2	31	10	1
PS1: Half-Price Books	14	26	11	1
PS2: UT Co-op Bookstore	6	17	17	0

We can evaluate their strategy and recommendations by considering our data on which sites had the cheapest price. Of the 38 participants, 18 said that they bought the cheapest usable book they could find, while 20 said that they preferred good quality used books, but would buy new if no such used book were available. Two ways we can evaluate the strategies are: (1) to look at the probability that the strategy finds the cheapest price; and (2) to ask how close the strategy comes to the cheapest price. From the perspective of (1), the strategies are frequently not very good. Eleven of the 20 who wanted to buy good used books didn't recommend both Half.com and Amazon.com, and would therefore find the cheapest book less than 60% of the time. But the real issue is how close a subject comes to the optimal strategy: if a student misses the cheapest book half the time but only pays a cent extra, his strategy is fine. We can estimate how good a strategy is by comparing the student's performance with checking all sites for the lowest prices, just Half.com, just Amazon.com, or just Amazon.com and Half.com. Consider the following table, where proportions given are the proportion of textbook prices to prices at the Co-op and integers in parentheses are how many students used that strategy as compared to how many recommended it:

Table 13
Estimate of how close advised (used) strategy is to optimal

Approximate Strategy	Goal Good Used	Number Advised(Used)	Goal Cheap Usable	Number Advised(Used)
All Relevant Sites	.69	1(3)	0.59	0(1)
Half.com and Amazon.com	.71	10(13)	63	13(15)
Half.com	.78	6(3)	.66	2(1)
Amazon.com	.81	3(1)	.74	1(0)
Neither Half.com nor Amazon.com	1.0	0(0)	1.0	2(1)
Weighted Average of Above	0.75(0.72)	N/A	0.64(0.62)	N/A

The subjects' advised (used) strategy costs no more than 8% (5%) above the estimated optimal when searching for good used books, and costs no more than 9% (5%) more when searching for the cheapest books. Note that participants often checked more sites than they'd recommend others check: 35 out of the 38 went to at least one website that they didn't purchase from and wouldn't recommend to others. As a typical student checks far fewer websites over much less time than we did, this behavior isn't surprising.

If getting advice from friends was an independent drawing, we could ask how many students would a freshman have to ask to get good advice. The table below shows the probability of obtaining good advice from upperclassmen who had bought textbooks online 4 or more times.

Table 14
Probability of obtaining recommendation for both Amazon.com and Half.com

Group	1	2	3	4
Half & Amazon	0.63	0.86	0.95	0.98

This indicates that a freshman only has to talk to a few upperclassman to hear good advice.

Two other factors in evaluating a strategy are timing and risk. To evaluate timing, we normalized our collected data by the relevant new or used UT Co-op price, and performed the following regression:

$$price = a + bt + \epsilon$$

where *price* is the normalized price and t the number of days after 30 December 2005. The results are displayed in the following table:

Table 15

Time trend in the textbook data

Seller	New	Good Used	Low Cost
a	0.53	0.55	0.52
b	0.0086	0.01	0.002
t-value for b	4.88	4.77	1.05

The regression indicates that the prices for all three categories increase slowly over time (although the increase for the low cost case is insignificant). The students surveyed who wanted high-quality texts purchased their books an average of 1.6 days before class started. They would have performed better if they'd started earlier.

Risk factors make it hard to say exactly how optimal a strategy is. A book's quality may be misrepresented in an online listing, or a semi-anonymous seller may not deliver the book as promised. Twenty-five of the 38 students questioned reported that at least one book they'd ordered had been delivered late. In these cases, the average time of arrival was more than four days late. One student never received his order. Another didn't recommend Half.com because he received two textbooks not of the promised quality. Twenty-two participants try to control risk by using the rating system; the average user rating they demanded was 91% positive. Of the data we collected, the average rating of sellers was 95%. Once the risk factor is introduced, it is difficult to precisely specify an optimal strategy. For example, if a student wanted a guaranteed refund should her book not be delivered, then an optimal strategy would be to only check Half.com and books sold directly by the various sites, but not books sold in associated marketplaces. At sites other than Half.com, refunds are largely handled by buyers and sellers, not the website.

From the perspective of pure cost minimization, there are individual differences among the subjects, with some subjects having very good strategies and advice, and others not so good. A student doesn't have to check prices at a lot of sellers to have a good strategy. For widely used introductory textbooks, Amazon.com Marketplace sometimes lists over 400 sellers; Half.com is a similarly large marketplace. As listings are sorted by price, there is enough competition for students to expect a reasonably good price.

7 Conclusion

Decision psychologists since Tversky and Kahneman, for example, see Kahneman, Slovic, and Tversky (1982), have asserted and shown in numerous experiments that human heuristics do not exactly optimize. The most passionate current advocate of the limits of bounded rational humans is Gigerenzer (2002). In this problem limited bounded rational humans are capable of achieving near optimal performance because bounded rational humans solve the problem as a distributed processing problem with an informal network and a time delay in that the market knowledge of online shopping students is transferred to new students. For cases where human procedures approach the results of an optimization model, the optimization model is a useful model for studying outcomes. Procedural economics needs to establish under what conditions and how close to optimal are human procedures. The fact that human distributed processing can provide greater computational resources than individual processing has received scant attention in the literature.

References

1. Birkhoff, Garrett, and Shane MacLane. 1953. *A Survey of Modern Algebra*, New York: The MacMillan Company:

2. Chow, Gregory. 1981. Econometric Analysis by Control Methods, Hoboken: John Wily and Sons:

3. Edwards, Ward. 1968. "Conservation in human information processing". In *Formal Representation of Human Judgment*, ed. B. Kleinmuntz New York: Wiley

4. El-Gamal, Mahnoud and David Grether. 1995. "Are People Bayesian? Uncovering Behavioral Strategies", *Journal of the American Statistical Association,* 90(432): 1137-1145

5. Friedman, Milton and L. J. Savage. 1948. "The Utility Analysis of Choices Involving Risk", *Journal of Political Economy*, LVI(4): 270-304.

6. Gigerenzer, Gerd. 2002. "The Adaptive Toolbox" in *Bounded Rationality* ed. Gerd Gigerenzer, and Reinhard Selten, 37-50, Cambridge, MA: MIT Press

7. Kahneman Daniel, Paul Slovic, and Amos Tversky. 1982. *Judgment under Uncertainty: Heuristics and Biases*, Cambridge, U.K: Cambridge University Press

8. Kendrick, David. 2005. "Stochastic Control for Economic Models: Past, Present and the Paths Ahead", *Journal of Economic Dynamics and Control*, Vol. 29: 3-30,

9. Norman A, A. Ahmed, J. Chou, A. Dalal, K. Fortson, M. Jindal, C. Kurz, H. Lee, K. Payne, R. Rando, K. Sheppard, E. Sublett, and I. White. 2004. "On

the Computational Complexity of Consumer Decision Rules", *Computational Economics*, 23(2), 173-192

10. Norman, A., A. Ahmed, J. Chou, K. Fortson, C. Kurz, H. Lee, L. Linden, K. Meythaler, R. Rando, K. Sheppard, N. Tantzen, I. White and M. Ziegler. 2003. "An Ordering Experiment", *Journal of Economic Behavior and Organization*, Vol 50 249-262

11. Norman, Alfred. and David Shimer. 1994., "Risk, Uncertainty and Complexity", *Journal of Economic Dynamics and Control*, 18, 231-249.

12. Norman, Alfred. 1994. "On the Complexity of Linear Quadratic Control", *European Journal of Operations Research*, 73, 319-330.

13. Norman, Alfred., and Woo Jung. 1977. "Linear quadratic theory for models with long lags," *Econometrica* 45, 905-17.

14. Norman, Alfred. 1976. "First order dual control," *Annals of Economic and Social Measurement* 5/3: 311-21.

Appendix A

We consider the problem presented in Norman (1993), which employs the information-based complexity model. A monopolist has a linear production process, faces a linear inverse demand function, and has a profit function for $t = 1,2,\ldots,T$:

$$p_t = a - dq_t,$$

$$\pi_t = p_t q_t - c(q_t), \tag{1}$$

$$q_t = \beta x_t + \zeta_t,$$

where a and d are known, q_t is the tth observation of net output, x_t is the tth level of the production process, β is the unknown scalar parameter, and ζ_t is the tth unobserved disturbance term. The ζ_t are iid normal with mean zero and known variance one. Since the complexity results are invariant to defining the cost function as a zero, linear, or quadratic function, the cost function is defined as $c(q_t) = 0$ to simplify the notation.

Given a normal prior on β at time $t = 1$, the prior information on β at time t is a normal distribution $N(m_t, h_t)$, where m_t is the mean updated by $h_t = h_{t-1} + x_{t-1}^2$and h_t is the precision updated by $m_t = (m_{t-1}h_{t-1} + q_{t-1}x_{t-1})/h_t$. For this paper let us consider two cases:

1. The agent knows β precisely. He or she has either been given precise knowledge of β or has observed (1) a countable number of times so that his or her prior on β has asymptotically converged to $N(\beta, \infty)$.

2. The agent's prior information on β is represented by $N(m_1, h_1)$, where h_1 has a very small positive value.

The monopolist is interested in maximizing his expected discounted profit over a finite time horizon:

$$J_T = \sup_{x^T} E[\sum_{t=1}^{T} \tau^{t-1} p_t(x_t) q_t(x_t) \mid q^{t-1}, x^{t-1}], \tag{2}$$

where τ is the discount factor, q^{t-1} is $(q_1, q_2, ..., q_{t-1})$ and x^{t-1} is $(x_1, x_2, ..., x_{t-1})$. q^{t-1} and x^{t-1} represent the fact that the decision maker anticipates complete information that is observed exactly and without delay.

First consider the optimization problem where β is a known parameter. The optimal x_t can be exactly determined as a function of the parameters of $f \in F$ without recourse to the information operator as

$$x_t^* = \frac{a}{2d\beta}. \tag{3}$$

The (ϵ=0)-computational complexity of this problem is T^0, polynomial zero, because the control that can be computed in 3 operations needs to be computed only once for the entire time horizon.

Now let us illustrate the computational difficulty with case 2, the simplest nontrivial example having a time horizon of only two periods, $T = 2$. The value function in the first period is

$$J_1(q_1) = \frac{a^2((m_1h_1 + q_1x_1)/h_1)^2}{4d([(m_1h_1 + q_1x_1)/(h_1 + x_1^2)]^2 + (h_1 + x_1^2)^{-1})} - d. \tag{4}$$

While the expectation of $J_1(q_1)$ has the form

$$E\left[\frac{Q_1(q_1)}{Q_2(q_1)}\right] - d, \tag{5}$$

where $Q_1(q_1)$ and $Q_2(q_1)$ are quadratic forms in the normal variable q_1. This expectation cannot be carried out explicitly to give an analytic closed expression. This implies the 0-complexity of this problem with an unknown parameter is transfinite.

Appendix B

Figure 1

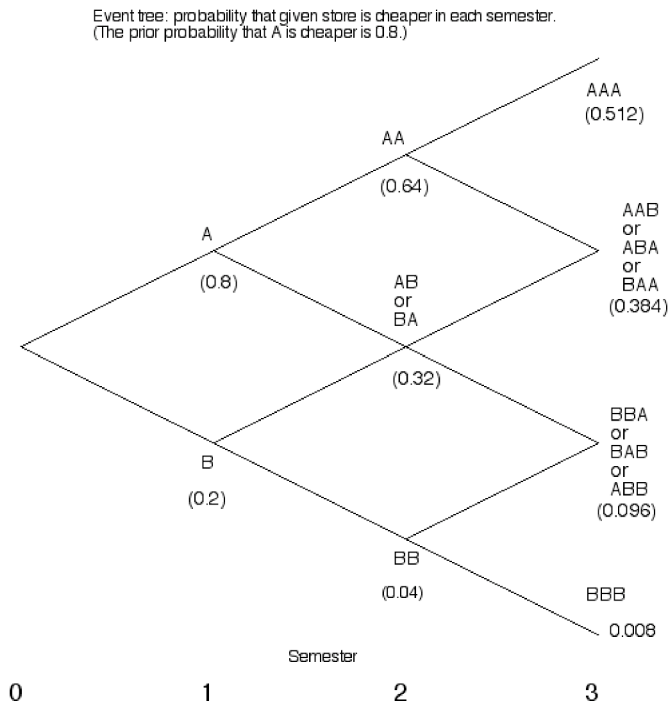

Theorem: If the price difference Z is less than \$25 in the problem shown in Figure 6, the computational complexity of solving this problem is at least a quadratic process at the number of periods n increases.

Proof. From Figure 1 the number of nodes increases by 1 each period so that total number to consider is $1 + 2 + 3 + \ldots$ which sums to a quadratic number of nodes to consider. Now as n increases we need to show that the number of nodes that must be considered to compute the optimal strategy increases as a quadratic function of n.

It is obvious that if one checks both prices m times the consumer should check both prices in the first m periods. If the second law of induction (page 13 of Birkhoff and MacLane (1958)) we assume the proposition $P(j)$ holds for all $j < m$ implies $P(m)$ is true, then $P(n)$ is true for all n. The proposition $P(j)$ is if we check prices j times and find the same store has lower price j times, then there is some value of n that is optimal to check prices in the $j+1$ period. Now let us assume that we have checked prices m times. The possible combinations of stores shown are m Ds and 0 Cs, (m-1) Ds and 1 C, (m-2) Ds and 2 Cs, ..., 1 D and (m-1) Cs, and 0 Ds and m Cs. Let us consider the case of m -1 Ds and we draw another D. The posterior probability that D is the cheaper store increases from $\frac{0.8^{m-1}}{(0.8^{m-1})+0.2^{m-1})}$ to $\frac{0.8^{m}}{(0.8^{m}+0.2^{m})}$ and if we choose buy at D for the rest of the time horizon the expected savings increases from $S_{m-1} = 0.8*\frac{(0.8^{m-1}}{(0.8^{m-1})+0.2^{m-1}))}$ *10 -5 to$S_m = 0.8*\frac{(0.8^{m}}{(0.8^{m}+0.2^{m}))}$*10-5. In the first case the expected savings in the last $n-m+1$ periods are $(n-m+1)*S_{m-1}$ and in the second case the expected savings in these periods are $0+(n-m)*S_m$. As $S_m > S_{m-1}$ there is an n such that the second sum is greater that the first and it is therefore optimal to check prices in the m^{th} period. Now let consider the node with u Ds and v Cs where $u+v=m$ and $u > v$. As the v Cs cancel out v Ds we are left with an equivalent problem of $u-v$ Ds and v * 2 (10 - 2*5) = 0 contributions to expected savings for the v pairs of CDs. By our hypothesis it is optimal to check prices again for this node. The same is true for the other combinations of Ds and Cs. Thus as the time horizon increases the number of nodes that must be considered in computing the optimal strategy increases quadratically.

Appendix C

Instructions

Objective: On each of next four pages you will be asked to buy one new textbook each semester for either 2 or 4 consecutive semesters. You must devise a strategy to find the textbook at the lowest average price - search cost. That is if 1000 students used your strategy the average savings would be greater than any other strategy. On the last page you will be asked several questions.

Reward: You will be paid a flat fee of $10 for participating determining the best strategy is worth $2, $7, $6, and $7 on the next four textbook buying problems. You can earn $1 for answering each of the first 3 questions correctly on the last page.

Problem Setup:

Factors common to the four problems:

Textbook: Buy one textbook each semester for the stated number of semesters.

Stores: There are two stores X and W with one more expensive than the other. The stores in the 4 problems are A & B, C & D, E & F, and G & H

Search Cost: The stores are on opposite sides of town, and it costs $5 to drive to each store from your home. For simplicity, assume cost of driving to both stores, checking prices and buying cheaper priced textbook is $5 + $5 = $10.

Unknown Prices: If you check the prices of 100 textbooks at store X and store W you would find that the one store would have prices that are $Z cheaper than the other 80% of the time for all 2 or 4 semesters in the problem. The other store will have prices that are $Z cheaper prices 20% of the time for all 2 or 4 semesters in the problem. You do not know whether X or W will be the lower priced store 80% of the time for all semesters in the problem unless you drive to find out.

Factors that vary in the four problems:

Semesters: The first problem is 2 semesters and the other three are 4 semesters

Price difference between two stores: The price difference is $30, $10, $20, and $15 for the four problems in the order given.

Strategy:

Consider the following 3 semester problem with a price difference of $18 and the other factors set to the values common to the four problems. You would see the a table similar to the one below except there would be buttons to click:

Period 1	Period 2	Period 3
Buy store X	Buy store X	Buy store X
Buy store W	Buy store W	Buy store W
Check prices at X and W, buy cheaper textbook	Check prices at X and W, buy cheaper textbook	Check prices at X and W, buy cheaper textbook

If you click "buy store X" it costs you $5 to buy the textbook at store X without finding out whether store W was cheaper. If you click "buy store W" it costs you $5

to buy the textbook at store W without finding out whether store X was cheaper. If you click "Check prices at X and W, buy cheaper textbook" it costs you $10 to check prices at store X and store W and buy the cheaper priced textbook. In devising your strategy to find the lowest prices the question to ask is how many times do I check prices at both stores: 0, 1, 2, or 3. It is more costly to check prices at both stores so what do you gain? When you consider the four problems does the number of semesters or the price difference make a difference in how many times you should check prices at both stores?
Details for the four problems:

Think first and then click once for each period as only the first click on a button: "Buy store X", "Buy store W", or "Check prices at X and W, buy cheapest" will be counted.

If you click on "Check prices at X and W, buy cheaper" WAIT FOR POPUP WINDOW THAT TELLS YOU WHICH STORE IS THE CHEAPER.