Decision Heuristics and Price Search

A. Norman, K. Brehm, M. Drake, A. Dyer, J. Frisby,
C. Govil, C. Hinchey, J. Ke, S. Kejriwal, Kuang,
S. Keyburn, S. Ler K. Powers J. Sanghai, J. Schieck,
J. Sussman, R. Wang and K. Yan

May 29, 2008

Abstract

We surveyed seniors to characterize campus bookstore buyers versus online buyers. We performed a Bayesian multiperiod textbook search problem. While most subjects did not compute the optimal strategy, they did demonstrate some intuitive understanding. We performed a simplier experiment in order to determine subjects search heuristics. Subjects employed heuristics that lead to good performance without demonstrating an understanding of the optimal rule. The performance of students with a course in statistics was statistically superior to those without statistics in the one period, but not two period problems. In developing their buying strategies, students talk to those with prior textbook buyhing experience. In order to judge the optimality of actual student strategies, we recorded online prices for 24 days. Students can achieve good performance by only checking two sites.

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

Microeconomic theory is the study of outcomes from optimization models. Why should economists study procedures? Kahneman, Tversky and numerous other psychologists have devoted their careers demonstrating that economic agent performance is less than optimal. Humans use simplifying heuristics that frequently result in good, but not optimal behavior. Humans make mistakes. For a recent book describing this literature, the reader might view Gigerrenzer and Selten (2002) A study of consumer procedures is intrinsically interesting, and can lead to software decision aids to bring consumer behavior closer to optimal behavior, for example see Norman et al (2008).

One reason consumer use simplifying heuristics to solve consumer problems is that the solution of the optimization problem is frequently not tractable that is unaided humans do not possess enough computational resources to compute it. What is tractable for a computer executing over a trillion operations a second vastly exceeds what is tractable for a human taking several seconds to execute a single binary comparison. Norman et al (2003) demonstrated that humans can order up to 25 alternatives using an algorithm whose executing time is a linear function of the number of alternatives considered. But consumers can face several hundred thousand alternatives in the marketplace that makes the execution of an algorithm whose execution time is linear in the number of alternatives considered intractable, see Norman et al (2004). In such cases, consumers use sublinear heuristics. Norman et al (2001) showed that solving a two stage budgeting problem is computational intractable. Consumers simplify the consumer optimization problem by selecting their goods item-by-item; also see Norman et al (2007). A second reason consumers use simplifying heuristics is that they do not possess the knowledge to implement the operations to obtain the optimal solution.

In this paper we study the interaction of these two reasons why consumers use simplifying heuristics. We study repeated price search; students buying textbooks are our empirical example. In Section 2, we provide an overview of the textbook market at UT. We use four surveys to characterize the UT Co-op buyers and online multisite buyers. The former are willing to pay higher prices in return for the convenience of the Co-op; others search many stores and Web sites saving about 40% relative to the Co-op prices over time, but with a risk of a late or no delivery and overstated quality. Our study investigates the extent to which experimental subjects and online multisite buyers are capable of creating an optimal textbook search strategy?

Repeated price search is a dynamic Bayesian optimization problem, also known as estimation and control problem, a topic economics have studied since the 1970s. For example, see Chow (1981), Norman and Jung (1977), or a survey of this literature by Kendrick (2005). Earlier in this study we created a two store multiperiod Bayesian optimization experiment. Each period, the subject each period had to decide either to check prices at both stores and buy from the cheapest or just go to one store and buy from that store. In a multiperiod problem, if the subject checks prices at both stores, he gathers knowledge to identify the cheaper store. Few subjects demonstrated optimal performance especially in the four period problem, although most demonstrated an intuitive understand of some aspects of the problem. This work is shown in Section 3.

Because it appeared that our subjects did not understand the Bayesian optimization problem we decided to explore their knowledge more deeply by creating a much simpler experiment. This second experiment had five problems that were designed to test subjects intuitive decision heuristics and see how well they understood the problems. Subjects demonstrated that they could achieve good performance even though they were unable to formulate a decision rule that would lead to optimal performance, especially in the case of the two period estimation and control problem.. The performance of students with a course in statistics was statistically superior to those without statistics for the one period, but not two period problems. These results are presented in Section 4.

In Section 5, we present the results of our fourth questionnaire in which we asked graduating seniors to recommend to freshman the websites they should check for low textbook prices. We also checked the prices for economic textbooks for all undergraduate classes online at multiple websites for a period of 21 days looking for a new, good used, and acceptable used textbook for two risk factors. Because many of these sites are marketplaces listing numerous sellers, a students only has to check a few sites to achieve good performance for different risk levels. Section 6 contains the conclusion.

2 Textbook Market and Buyers

Students at UT 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. There are legal issues in selling 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 UT Co-op specified 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. Students buying textbooks at the UT Co-op pay 8.25% sales tax, but at the end of the academic year, they will also receive a 10% rebate towards Co-op purchases. The most important characteristic of this seller is convenience and zero risk of nondelivery.

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. These third party sellers can be students, bookstores with their own websites for sales, or even marketplace websites. 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. These sites search mid to large sized sellers. 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. For example, 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. They also do not search continuously so that a student can go 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.

The reason some students search multisites online is to save about 40% from the UT Co-op prices, but with a risk that the book will be delivered late or not at all or not be in the listed condition. Students buy from the UT Co-op for convenience and risk avoidance. If their parents allow them to buy from the UT Co-op with their credit card they lack incentives to shop around.

A UT student completing a four year program in 8 semesters could buy textbooks online from mulitisites 8 times. How does a student through repeated price searches learn to improve his performance? Also, given the existence of marketplace websites, does a student have to check a large number of alternatives to achieve good performance?

3 First Experiment

We design a Bayesian multiperiod optimization experiment to text whether students are capable of devising a strategy to optimal learn which store (or site) has the lowest relative prices. As these optimization problems are often intractable, Norman (1994), we obtained a computational tractable experiment by using discrete distributions and only three alternatives. The subject needs to buy a textbook each semester and has three choices each period:

Table 1: Choices

Choice	Per 1	Per 2	Per 3	Per 4	⋯
Buy from A	A	A	A	A	
Buy from B	B	B	B	B	⋯
Check prices both, buy from cheaper	A,B	A,B	A,B	A,B	

The subjects solves four multiperiod search problems. The factors that are the same in all four problems are:

1. Probability: 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. 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.

2. Travel Costs: It costs \$5 in time and transportation costs to visit each store. For the sake of simplicity, it's assumed that she will not incur further transportation costs after visiting each store once.

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

Table 2: Characteristics of Four Problems

Problem	Stores	Cheaper Store	Number of Semesters	\$X
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

where \$X is the amount one of the two stores would be cheaper. Our experiment had each subject solve four problems and answer questions on a fifth page. The experiment is at: http://www.eco.utexas.edu/Homepages/Faculty/Norman/00Ashley/

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 3: 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 3, 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.

Let us consider the computational complexity of this problem in terms of arithmetic operations and a growth parameter the time horizon, T. One approach would be to consider the T^3 sequences of choices. As is shown in Appendix B, this problem's computational complexity is at least quadratic in the number of semesters to consider so an algorithm more efficient than T^2 does not exist. We limited the time horizon to four periods to keep the number of arithmetic operations to less than 50.

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. In this experiment we were interested in whether subjects could devise an optimal strategy and paid them if they could. In the next experiment we will explore suboptimal performance and provide a reward based on performance.

The results for experiment 1 are shown below:

Table 4: Results for Problems 1-4

	Problem 1	Problem 2	Problem 3	Problem 4
Price Diff	\$30	\$10	\$20	\$15
Optimal Checks	2	1	3	2.32
No. Correct	6	10	3	2
Avg Checks	1.4	0.88	1.52	1.48
No. Skips	1	7	4	5

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 $(\$30 - 2 \times \$5) + (0.8 \times 0.8 + 0.2 \times 0.2) \times \$30 - \$5 = \$35.40$. On average the subjects checked both prices fewer times than optimal in each problem. Given how many subjects failed to optimize their savings, we conclude that subjects have little quantitative understanding of the problem. But, because the differences between the optimal and the next best supoptimal was small there is a possibility that subjects poor performance is due to arithmetic errors.

We also explored how much qualitative knowledge the subjectgs have of this type of problem. 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

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

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):

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 6: 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.

4 Second Experiment

Because subjects were not able to devise an optimal strategy for a potentially tractable estimation and control problem, we decided to probe more deeply into the actual heuristics they employ in this type of problem. For this purpose we devised a five part experiment. The first page of the experiment gave the subjects instructions. These instructions and the experiment itself can be viewed at: http://www.eco.utexas.edu/Homepages/Faculty/Norman/00Carolyn/. On pages 2-5 the subject solved a one period search problem. The instructions on these pages consisted of three paragraphs: information, travel costs, and anticipated price differences.

To conserve space the pages are combined in the description below.

Information:

(Pages 2 and 3:) In this case one of the stores (either A or B) is cheaper than the other most of the time. But, as you have had no previous experience with these two stores, you have NO information as to which store is cheaper.

(Page 4:) In talking with seniors about the price at A and B, you can assume that if you checked the price at both A and B 100 times you would find that 60 times A would be cheaper and 40 times B would be cheaper.

(Page 5:) In talking with seniors about the price at A and B, you can assume that if you checked the price at both A and B 100 times you would find that 80 times A would be cheaper and 20 times B would be cheaper.

Travel Costs:

(Pages 2, 4 and 5:) If you check the price at only one store and buy from that store, it costs you \$3.50. If you check the price at both stores and buy from the cheapest, your expected cost is \$7.

(Page 3:) If you check the price at only one store and buy from that store, it costs you \$6. If you check the price at both stores and buy from the cheapest, your expected cost is \$12.

Anticipated Price Differences:

(Pages 2 -5:) Now suppose you anticipate that one store will be cheaper than the other by \$X where the values of X are given below. For each price difference indicate whether you would check the price at both stores.

Table 7: Anticipated Price Differences

Anticipated Price Difference	Check price at both store A and store B
\$5	Yes ◯ No ◯
\$10	Yes ◯ No ◯
\$15	Yes ◯ No ◯
\$20	Yes ◯ No ◯
\$25	Yes ◯ No ◯

Note: On pages 2 and 3 the subject has no prior knowledge which store is cheaper and the difference between these pages is the travel cost. On pages 4 and 5 the subject has complete prior knowledge which store is cheaper and the difference between these pages is the probability of store A is cheaper.

The subject is presented with a two period estimation and control problem on page 6.

Information: In this case you will make decisions as to checking the price at two stores, A and B for two semesters. Initially, you have NO information as to which store is cheaper. If you did price-comparison shopping at the two stores 100 times, you would find one of the two stores was cheaper 90 times out of 100. For your second semester decision you have one observation which store is cheaper if you check prices at both stores the first semester.

Travel Costs: If you check the price at only one store and buy from that store, it costs you $7. If you check the price at both stores and buy from the cheapest, your expected cost is $14.
Anticipated Price Differences: Now suppose you anticipate that one store will be cheaper than the other by $X where the values of X are given below.

Two Semester Problem: This is a 2-semester problem. Suppose in the first semester you decide to check prices at both stores in semester one and buy from the cheaper store. You would have one observation as to whether store A or store B is cheaper. Note: In semester 2 use the appropriate row based on what action you took in semester 1. In the tables below Just One means either one and Both means check price at both and buy from cheaper store.

Table 8

Anticipated Price Difference = $5

Semester 1	Just one __ Use line A below Both __ Use line B below
A: Semester 2: If in Semester 1, you checked just one __ then use this row ——————————- B: Semester 2: If in Semester 1, you checked both __ then use rows below ——	Just one __ Both __ __ ————————- Cheaper store semester 1 __ __ More expensive store semester 1 __ Both __

There was a second table the same as above except the Anticipated Price Difference = $15 and a third table with the Anticipated Price Difference = $60.

On pages 2, 4 and 6 there was a textarea in which the subject was asked to “Type a paragraph describing how you solved this problem:” Subjects were given a 4 function calculator and a piece of scrath paper to make calculations if they so desired.

We performed this experiment on two groups of predominatedly economics and business administration majors. One group of 19 had never taken a course in statistics and the other of 19 had taken a basic course in statistics such as economic statistics, business statistics, or AP statistics in high school. The incentives for the experiment were: “Incentives: Assume you are being paid to advise 100 freshmen about buying a certain textbook. If your recommendations are better in the sense of lower costs on average, you earn more money. You will receive a flat fee of $7 for coming to the experiment. There are 26 questions. You will receive (your score)/(perfect score) times $8. If you get over 90% you will receive an additional $4. If you get all questions correct you will earn an addtional $4 so that the maximum possible earnings is $23.”The subjects took abour one half hour to complete the experiment.

Now let us consider the results. There are three questions to resolve concerning performance. Were the performance of both group significantly better than random selection and significantly less that optimal? Also, did a basic course in statistics made a difference?. We measured the performance of each subject as a percent of the optimal score for the listed anticipated prices. In determining the optimal score for each anticipated price difference we determined which of the alternatives had the greatest expected savings from the higher price including the travel costs and added the scores for each page. As we expect the performance of subjects with statistics to be greater than those without the appropriate test is a single tail t test with unequal variances: $H_0 : Ran = \mu_{no} = \mu_{stat} = Opt$
$H_a : Ran < \mu_{no} < \mu_{stat} < Opt$
The results are shown in the table below:

Table 9: Does basic statistics course matter?

Periods	Problem	Pages	Sub	Ran	μ_{no}	μ_{stat}	P(no-stat)
1	No knowledge	2 and 3	19	68.2	94.20	97.94	0.5
1	Complete knowledge	4 and 5	19	87.2	95.37	98.42	0.01
2	Partial knowledge	Page 6	19	63.8	90.67	91.26	0.44

where Sub is the number of subjects, Ran is the performance is they randomly chose alternatives, and μ_{no} and μ_{stat} are the means of students with no statistics and a basic course respectively. An optimal score would be 100 in all calses. With an α of 0.01 both groups performance were statistically better than random and statistically worse than optimal. In all cases the mean performance of subjects with a basic statistics course was greater than those without. As is shown in the table above this difference is significant with both one period problems, but not the two period problem. If we consider pages 2,3,4, and 5 separately, the performance of subjects with a basic statistics course was greater than those without, but the difference was significant at $\alpha = 0.05$ for only pages 2 and 4. With a much larger sample it is possible the difference would be significant for all cases. The subjects with a course in statistics had much less variation in their performance in pages 3 and 5 than those without statistics.

Now let us examine the reasons for less than optimal performance. Because the subjects all had a calculator and scratch paper, the errors they made were conceptual not arithmetic. The explanations on pages 2,4 and 6 provide some insights into how the subjects solved the problem. The following table summarizes the number of subjects that solved the problem correctly and provided a logical explanation of how they solved the problem;

Table 10: Number subjects with the correct response

No. Periods	Problem	Pages	Sub	No Ans	No Epl	St Ans	St Epl
1	No Info	2,3	19	5	5	8	7
1	Complete Info	4,5	19	3	0	7	2
2	Partial Info	6	19	4	0	3	0

where No stands for no statistics course group and St stands for the statistics course group. Ans means that they solved all problems on the respective pages correctly and Epl means they gave the correct answer in the textarea on the respective page. As can be seen only 5 members of the no statistics group could articulate the correct solution to the page 2 problem and none could articulate the correct solution to page 4 or 6. The performance of the statistics group was only slightly better. No one presented a correct response for how to solve the page 6 problem.

The experiment was designed to also test the subjects intuitive understanding of the problems. For pages 2-5 there is a cutoff anticipated price difference such that for all lower anticipated price differences the subject should buy from one store and for that and higher anticipated prices difference the subject should check prices at both stores and buy from the cheaper. The formula for computing this price difference for pages 2 and 3 is: check prices at both stores and buy from the cheaper if the anticipated prices difference is greater than the travel costs to both stores. The formula for pages 4 and 5 is: check prices at both stores and buy from the cheaper if the (1-$P_{cheaper}$) $\times$ Anticipated Price Difference is less than the travel cost to the second store. These optimal shift points are shown in the table below:

Table 11: Optimal Shift Points

Page	2	3	4	5
Shift from one to both at	10	15	10	20

In order to evaluate the results of the experiment, we must categorize the data. We anticipate that subjects will use the following heuristic for the first four sheets: as the anticipated price difference increases, the subject will shift from buying at one to checking both and buying cheaper. Once they shift they should stay shifted at all higher prices. Call the point at which a subject begins to check at both stores the shift point. In order to perform statistical tests on the shift point, we must determine where the shift point is. Here we consider thirty-six subjects with four sheets each, for a total of 76 trials. 56 of these trials strictly follow the heuristic, that is, the player begins checking one and at some point switches to checking both. In a further 16 trials, the subject plays one strategy throughout, either checking one or both. These trials follow the pattern of the heuristic, and so can be coded with a hypothetical shift point above the actual range of price differences given if they bought at one throughout and a hypothetical value below the actual range if they checked both throughout. However, there is no way for us to know if the player who played one strategy for all price differences is aware of the concept of a shift point. We may be claiming that there is a shift point when none really exists. However, all subjects played at least one sheet strictly following the heuristic, so it seem reasonable to interpret this to mean we do not need to worry about players not understanding the concept of a shift point and thus not having one. Finally, there are four sheets where the player used a strategy that did not follow the heuristic. For example, some players began checking both, and then as the price difference rose the player switched to checking one and then switched back to checking both. We are using matched pair tests, so we cannot use the pairs of data in which one sheet from the pair follows a non-heuristic strategy as there is simply no way to assign a shift point value to them that is no totally arbitrary. There are three such pairs, so we must discard them. Inspection of the data reveals no strong outliers, and n$\geq$17 for each population so by the central limit theorem the data should be normal enough to justify using a T-test. Use $\alpha = 0.05$ for all tests. In each pair, the sample variances of the populations differ by at least a quarter, so use unequal variance tests throughout.

Let us perform our first test on the subjects understanding of travel costs.

The first and second pages of the experiment form a pair, in which the only difference is that the second page has a higher travel costs. The optimum shift point increases as the travel cost increases. To test if the subjects understood this, we developed the following null and alternative hypotheses: $H_0 : \mu_2 - \mu_3 = 0$, $H_a : \mu_2 - \mu_3 < 0$, where μ_n is the mean shift point on the nth sheet for the population under consideration. We ran T-tests testing this hypothesis with the data from the first and second sheets for both the stats and no stats subjects. The results are described in the following chart:

Table 12: Tests on intuition–Pages 2 and 3

Test	Group	No. Obs	Pages	μ_2	μ_3	t Var	P-value
1	stat	19	(2,3)	6.84	14.21	unequal	0.0000
6	no stat	18	(2,3)	11.11	15.56	unequal	0.0066

In both cases we reject the null hypothesis with $\alpha = 0.05$ in favor of the alternative hypothesis. The low p-values suggest we can be very confident in our result. Thus, our experiment suggests that the subjects understood that as travel costs increased, they should wait for a higher price difference to switch strategies. This result holds for subjects both with and without statistics.

Let us perform our second test on the subjects' understanding of the probability of one store being cheaper. The third and fourth pages of the experiment form a pair in which the only difference is that on the fourth page, the first store is cheaper more often (80% on page 4 vs. 60% on page 3). The optimum shift point increases as the probability increases. To test if the subjects understood this, we developed the following null and alternative hypotheses: $H_0 : \mu_3 - \mu_4$

$= 0$, $H_a : \mu_3 - \mu_4 < 0$. Again we ran T-tests on the stats and no stats subjects, with the following results:

Table 13: Tests on intuition–Pages 4 and 5

Test	Group	No. Obs	Pages	μ_3	μ_4	t Var	P-value
1	stat	19	(4,5)	11.58	17.37	unequal	0.0023
2	no stat	17	(4,5)	15.59	21.47	unequal	0.0094

In all four cases we reject the null hypothesis with $\alpha = 0.05$ in favor of the alternative hypothesis. The low p-values suggest we can be very confident in our result. Thus, the experiment suggests that the subjects understood that as the likelihood of one store being cheaper increased, they should wait for a higher price difference to switch strategies. This result holds for subjects both with and without statistics.

Thus the data suggest that subjects somehow understand that they need to switch at a higher shift point both when travel costs increases and when it becomes more likely that the first population will be cheaper than the second one. Considering how poorly the students articulated their strategies in the free response text boxes included in the experiment, this is a surprising result.

Now let us consider Page 6, the two period problem. The results are shown in the following table:

Table 14: Two Period Problem–Page 6

P Dif	Correct	% Correct No	% Correct Stat
\$5	1-1	84	100
\$15	B-C	74	53
\$60	B-Bt	32	26

where P Dif is the Anticipated Price Difference and 1-1, B-C, and B-B means just one-just one, both-cheapest, and both-both respectively. As can be seen by the table the performance of both groups is poor and a first course in statistics may well reduce performance in this type of problem. For the anticipated price difference of $60 11 students with no statistics and 13 students with statistics chose B-C. This would be the correct answer if one observation in semester 1 indentified the cheaper store with certainty.

5 Buying Textbooks Online

Given that students performance was less than optimal performancein our experiments and that subjects appeared to have very little understanding of estimation and control problems, we assumed that student textbook buying strategies are far from optimal. As we shall see in this section there is a factor not considered in the experiments, the transfer of knowledge about textbook buying from students with experience to those without experience.

In order to evaluating student textbook buying performance, 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 15
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. 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 16
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. This raises the question how many sites does a subject have to check to achieve good performance.

We recorded the lowest market prices for 23 of the undergraduate economics course textbooks, we found online for 19 days between between 28 December 2007 and 19 January 2008. We considered three publishers: US books, International Color, International Black and White, and two levels of risk: cheapest price with no concern for the reliability rating of the seller and a 95+ rating with at least 30 transactions. For those sites that used a different rating system we used as close an approximation as possible. For the US books we also recorded 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) acceptable used book. For the international editions we just recorded prices for new textbooks.

In order to determine what sites to check we started with the meta search sites. Of these we found CampusBooks.com and directtextbook.com the most useful. From these sites we determined the sites most useful to check on a daily basis. Each day we checked a1.com, abebooks.com, alibris.com, amazon.com, bn.com, biblio.com, eBay.com, express.eBay.com, half.com, textbooks.com, textbooksnow.com, textbooksRus.com, textbooksX.com, and Valorebooks.com. Because smaller sellers who have their own websites also list at the large marketplaces such as amazon.com and half.com, we consider the search comprehensive especially in the case of a 95 rating with 30 transactions. In the table below, we show the frequency that sellers had the lowest price in each of the three categories and two risk levels for US published textbooks. Sellers who had the lowest price less than 5% overall were combined into the Other category.

Table 17
Cheapest sites in price survey: % of 437 data points

Site	New	New R95	Good	Good R95	Fair	Fair R95
Half.com	28	48	30	49	28	49
Amazon.com	35	31	34	30	21	23
AbeBooks.com	5	5	9	7	12	12
Textbooksnow.com	4	5	2	0	9	8
Valore.com	6	0	6	7	6	1
Other	22	11	19	7	24	7

Examining the table a strategy just to check prices at both Half.com and Amazon.com marketplaces results in a lowest price 63%, 79%, 64%,79%, 49%, and 72% of the time for the six categories. We also checked prices for new International Black and White at two risk levels and new International Color at two risk levels. The low cost sites are shown below:

Table 18

Cheapest sites for new international editions (Percent)

Site	NIB	NIB R95	NIC	NIC R95
Abe.com	53	84	28	33
eBay .com	16	14	30	42
TextbooksRUS.com	9	1	21	7
a1.com	22	0	7	0
Valore.com	0	0	4	13
Other	0	0	9	4

Again, a strategy just to check prices at Abe.com and eBay.com results in the lowest price 69%, 98%, 58% and 75% of the time for the four categories.

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 19
Data sources for students using many sellers

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

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 20

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 as is shown in the following table where performance is measured relative to the cheapest price set to 1.

Table 21
Performance of Amazon.com and Half.com Strategies

Strategy	New	New 95	Good	Good 95	Fair	Fair 95
Both	1.05	1.03	1.06	1.02	1.08	1.02
Amazon.com	1.12	1.14	1.13	1.12	1.14	1.12
Half.com	1.12	1.06	1.12	1.06	1.13	1.05

The table shows that just checking both Amazon.com and Half.com would result in a strategy no worse than 8 % higher prices than our lowest. In the cases where the student uses a 95 rating to reduce risk the increase is no greater than 3%. How do these prices compare with the listed UT Co-op prices. In an earlier version of the paper we showed there is a slight upward trend in price data; therefore, we show this comparison for three difference days in the table below:

Table 22
Cheapest Prices relative to UT Co-op (Percent)

Day	New	New 95	Good	Good 95	Fair	Fair 95
28 Dec	58	64	66	69	65	68
6 Jan	61	66	68	77	66	72
19 Jan	60	67	73	83	72	78

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 23
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.

Now we can state that even though subjects performance on our experiments was not optimal and they knew very little about estimation and control problems, they are abe to achieve good performance by talking to seniors who are experience with the online textbook market. Secondly, because the large marketplace sellers such as Amazon.com can list over 400 sellers in their marketplace a students does not have to check a large number of sites in order to obtain a good price. In addition, the large marketplaces have well established rating systems. It would appear that some websites use noncomparable rating systems to mask rather than illuminate the reliability of their sellers. The interested reader might compare the rating system of a1.com with amazon.com.

6 Conclusion

Transfer of knowledge from the experienced to less experienced is an important factor in how humans achieve performance in problems where they do not have good mathematical understanding of the underlying problems. Given the growth of meta search sites on the internet, such as pricegrabber.com, a consumer frequently does not have to search through a large number of sites to find a good price.

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

Appendix A

Figure 1

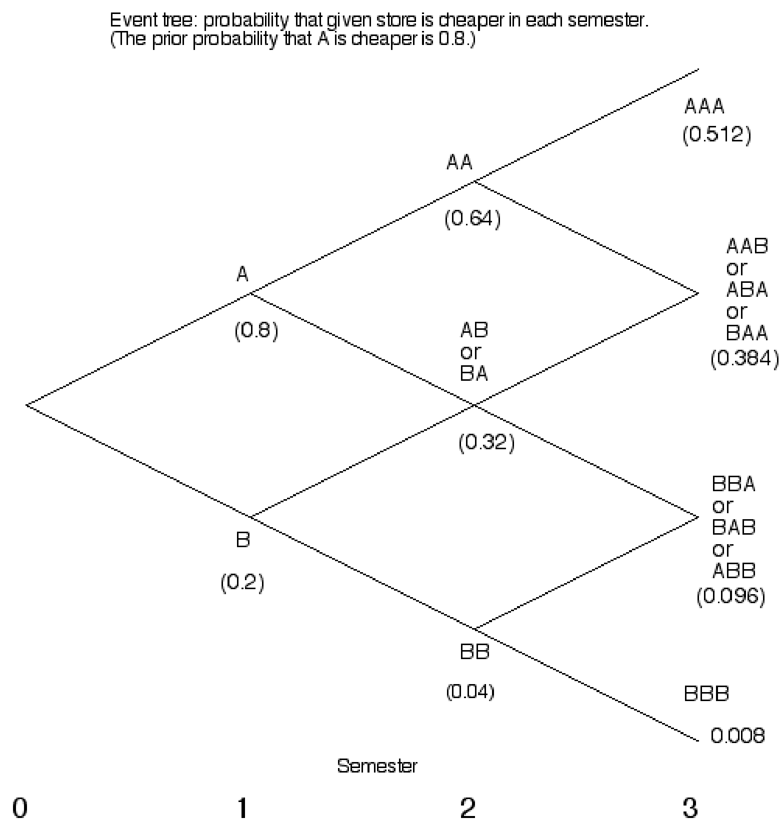

Theorem: If the price difference Z is less than \$25 in the problem shown in Figure 6, the computational complexity of solving this problem is at least a quadratic process at the number of periods n increases.

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.