Consumer Optimization Problem Solving

By

Alfred Lorn Norman

This is an extended abstract for the book that is forthcoming at World Scientific Publishers

Chapter 1: Introduction

The processor in a cell phone can execute millions of instructions per second. In contrast, the System 2 (conscious) processor of the human brain is well known to be very slow. For example, a human can take several seconds to make one binary comparison between two pens. This raises a fundamental question: How does a human solve the difficult consumer optimization problem?

We study the simplifications consumers employ to make the consumer problem tractable. Consumers search for goods and services item-by-item, which greatly reduces the number of alternatives to consider. In addition, consumers have operators that can process a set in a single operation. Finally, consumers budget by incremental adjustment.

In considering consumer performance we focus on how close to optimal is consumer performance and not whether consumers optimize as a yes/no question. Given the ordinal nature of utility theory this creates a basic measurement problem. We review the literature on consumer performance.

This is an opportune time to study consumer procedures because the Internet provides a media to make substantial improvements in consumer performance. We design a decision aid and demonstrate its superior performance in an experiment.

We now discuss the organization of the book in greater detail.

Part I: Computational Complexity and the Consumer

Chapter 2: Computational complexity

In this book, formal models of consumer procedures will be specified as algorithms that will be defined in terms of psychological decision rules, information operations, as well as arithmetic operations that will be treated as black boxes with execution costs. Algorithms can be analyzed by computational complexity, which is a measure of the difficulty in solving a problem asymptotically. A basic overview of these techniques used in the book is presented. We also present our analysis of the computational complexity of a discrete consumer utility optimization model and a review of the subsequent literature.

How difficult a problem a processor can solve depends on the speed at which the processor executes an elementary operation. Problems, which a processor can solve in a ``reasonable'' amount of time, are called processor tractable. Given the speed at which digital computers execute operations, algorithms for which the cost function increases as a polynomial of the growth parameter are considered tractable in that it could be executed in a time considered economic. An example would be conventional matrix multiplication of two n x n matrices for which the computational complexity is n^3.

Now, let us consider a human as a computational agent executing System 2 (conscious) operations. Compared with a digital computer an unaided human is generally very slow in executing an operation. For example, we determined that subjects took about 3.2 seconds to execute a binary comparison involving two pens. A linear human algorithm requiring 3.2 seconds to compare two alternatives would take 889 hours to process a million alternatives and as we show, the number of bundles in the marketplace is much larger than a million. This is not tractable for unaided humans. Algorithms that are tractable for an unaided human are sublinear and this defines Simon's concept of bounded rationality in terms of computational complexity.

The most important simplification of the consumer optimization problem that humans make to obtain a tractable problem is to search for their goods and services item-by-item. In order to intuitively understand why this is the case, let us consider two models of a grocery store. The first model is the traditional grocery store where the consumer traverses the aisles filling the shopping cart item by item. An item-by-item grocery store organizes close substitutes such as cereal and bread in the same aisle.

A second model is a bundles grocery store that enables the consumer to compare bundles of groceries. In such a store the consumer walks down a single aisle comparing shopping carts, each loaded with a bundle basket of groceries. In addition to the current practice of listing the price of each item, assume the store lists the price of the entire bundle on the shopping cart. Also, assume the store organizes the giant line of shopping carts containing all the possible bundles in ascending order by bundle price.

To obtain some perspective of the number of alternatives that must be considered, let us assume that the two grocery stores each carry 30 categories of goods and 10 alternatives in each, for example, 10 types of cereal and 10 types of bread. These numbers are very conservative for a modern grocery store. Also, we assume that the consumer wants one item from each category. In the bundle grocery store the number of shopping carts that the consumer must consider is 10 x 10 x 10 x … x 10 = 1030. If it only took 10 seconds to make a binary comparison between two bundles of 30 items, it would take 1.59×10^{23} years to find the preferred bundle using the linear procedure of comparing the first with the second and the preferred with the third and so on. If each bundle were placed in a 3-foot-long shopping cart then the consumer would have to travel 5.68×10^{27} miles just to view all the bundles.

To drastically reduce their display space, sellers organize their goods for item-by -item acquisition, and consumers acquire goods item by item to drastically reduce the number of alternatives they need to consider. In the case above acquiring groceries item-by-item reduces the number of alternatives to consider from 10^{30} to 10 + 10 + 10 + … + 10 = 300 that is the number of alternatives increase additively not geometrically.

Part II: Component studies

Before examining how consumers perform item searches, we present preliminary studies of some components of the consumer problem in the next three chapters.

Chapter 3: Ordering

We use experiments to show that humans can order a small number of items using a linear procedure Binary comparison operators form the basis of consumer set theory. If humans could only perform binary comparisons, the most efficient procedure a human might employ to make a complete preference ordering of n items would be an $n \log_2 n$ algorithm. But, if humans are capable of assigning each item a utility value, they are capable of implementing a more efficient linear algorithm. We consider six incentive systems for ordering three different sets of objects: pens, notebooks, and Hot Wheels. We establish that subjects are using a linear procedure by performing regression analysis, observing the hand motions of the subjects, and talking with the subjects about which algorithm they used. What makes regression analysis possible is that a binary comparison takes seconds not nanoseconds. The fact that humans can order a small number of alternatives such as stores, products, or attributes using a linear procedure is useful to develop efficient search procedures over time.

Chapter 4: Psychological Decision Rules

A consumer entering a new bookstore can face more than 250,000 alternatives. We then consider the efficiency of using various known psychological decision rules in search procedures. In this analysis we assume that the cost of creating such decision rules is fixed and that a consumer uses the same decision rule plus an information operator from start to finish. We show that procedures based on these known rules in search are linear procedures and that such rules are not tractable for a human given the number of alternatives the consumer faces in the marketplace.

We introduce a new psychological decision rule based on a set information operator, and show that it leads to a sublinear procedure. Next, we perform an experiment to show that humans have the ability to employ such a rule effectively, and finally we show that markets in both physical space and cyberspace are organized to facilitate the use of such set rules by consumers. In cyberspace, decision rules can be encoded as decision aids, which reduce the human effort to a click. Set decision rules enable humans to process a large number of alternatives in an item search using a sublinear procedure.

Chapter 5: Repeated Price Search

In this chapter we consider repeated price search. Consumers check few sites in online purchases. We perform repeated price search experiments that are dynamic Bayesian optimization problems for which the subjects could not calculate the optimal strategy. Instead, they used heuristics whose performance is better than random and less than optimal. To investigate online price search performance we surveyed students on their online textbook purchases. Students achieve good performance because they start with a good strategy and because of the online marketplace and meta-search sites. An important factor is that algorithms at sites searched perform calculations that reduce the computational complexity of the search. Finally, new decision aids on the Internet and on smart phones that effectively reduce the effort of a price search to a click are changing price search strategies.

Part III: Item Search and Budgeting

In the next three chapters we present our theory of how consumers solve the economic consumer problem. We relax the assumption that tastes are given and assume that preferences are learned. When pressed, very few economists would assert that consumers acquire a complete preference ordering neatly encoded in their genome at conception. When cardinal utility was formulated in the 19th century, consumers had few choices and the rate of technological advance was slow. Under these conditions it was a reasonable assumption to assume that a consumer had learned his preferences by the time he became an adult consumer. Today consumers face a very large number of alternatives and frequently encounter products that have new features or are completely new. Under such conditions they learn their preferences as part of the item search.

Chapter 6: Item Search: Forecasting

Consumers, in considering their alternatives, need to forecast the future performance of an alternative in its intended use. For frequently purchased goods such as food, consumers can forecast the future performance and learn their preferences by a sequential trial and error (STE) strategy that is a strategy in which consumers sequentially try different alternatives until preferences are formed. We investigate how students use this strategy to choose which restaurants to frequent for lunch.

The longer the time period between purchases and the greater the rate of technological change, the less useful a STE strategy becomes for forecasting and the less the value of prior consumption. With the rapid rate of technological change and the proliferation of goods and services in the marketplace, consumers must develop strategies for forecasting the performance of alternatives with which they may have had little experience.

As consumers have limited ability to test most products, consumers must gather data from others including media ads, data supplied by the sellers and manufacturers, product reviews by current users and experts. We review the development of data sources first in newspapers, magazines and now the Internet. How much data should a consumer acquire to forecast? It is not ``perfect information" because this concept is not easily definable in that the potential limit of data a consumer could acquire and process is determined by the Heisenberg uncertainty principle. Given that acquiring and processing data has real costs, it is readily apparent that consumers acquire and process a truly minuscule fraction of the data that could be provided to them. The information value of data is a function of its processing cost, reliability and capacity to discriminate among alternatives.

In forecasting that requires data acquisition and processing, consumers use a wide variety of procedures that we demonstrate using digital cameras as an example.

Chapter 7: Item Search: Choice

A typical consumer item search has the following steps:

1. Goal and start set
2. Several set decision steps
3. Evaluation of specific alternatives in the final set
4. Selection of preferred item or items.

With repeated searching, the number of alternatives in the start set decreases and in some case it can decrease to one, which means the item search has been reduced to a replacement operation. The amount of System 2 effort involved in making the various steps decreases over time and steps are combined.

Over time a consumer develops a variety of types of item searches for different types of goods that vary greatly depending on the frequency of purchase, the cost, the type of goods, and the rate of technological change. There is also tremendous variation in item search for the same type of product among consumers.

Consumers achieve search efficiency through repeated searches by acquiring a great volume of consumer knowledge about market organization and goods and services. The acquisition of this knowledge starts in early childhood and while prosaic, is extensive. In repeated searches there is no point in remembering preferences over specific alternatives that will be replaced by new products by the time of the next search for that good. An important factor in consumer search efficiency is that consumers learn decision rules and preferences over sets and attributes that will be useful in future shopping. We demonstrate this factor with a study of students' search for jeans.

Chapter 8: Budgeting

An optimal solution of budgeting that considers every possible bundle is intractable. Even allocating money into alternative accounts is an intractable quadratic process. For example, there are 5050 possible allocations in the case of allocating

$100 into three accounts in $1 increments. How do consumers obtain tractable heuristics to solve the budgeting problem?

In order to understand budgeting we conducted several surveys of students who lived in apartments. At our university most students live in a dorm the first year and move into apartments the second year, a move that requires them to allocate money into several accounts. What we observe is that students use an incremental adjustment process. They monitor the flow of funds and make adjustments. What makes this type of budgeting difficult is not the amount of calculation, but rather controlling the flow.

Some students are feast or famine budgeters in that when they receive funds they initially feast and then starve towards the end of their budget period. Other students have good control over the flow of funds and some are soft budgeters in that they can ask their parents for more money if needed. Another area where student budgeters differ greatly is in the amount of budget planning they perform. Only a few use spreadsheets to budget or budgeting software such as Intuit's Quicken. The latest trend is budget applications on cell phones. Budgeting is not a computational intensive process, but rather an intuitive adjustment process.

Part IV: How Close to Optimal and Improving Performance

In the next two chapters we ask how close to optimal is consumer performance and what potential innovations could be created to improve consumer performance.

Chapter 9: How Close to Optimal

Given the difficulty of the consumer optimization problem relative to what is tractable for an unaided human processor, the question to ask is how close to optimal is human consumer performance. This question presents a very difficult measurement problem given the ordinal nature of utility theory and the limitations of revealed preference.

We define the attributes of the reference model we will use to measure performance. We consider a one period model with savings as the variable that connects the periods. We consider both anticipated or decision utility and realized or experienced utility. We review the various measures that have been created to identify mistakes and evaluate performance relative to both anticipated and realized utility.

We then consider two experiments that are good approximations to our reference model. The results demonstrate behavior consistent with the theory. We also review numerous articles concerning the identification of mistakes and performance evaluation with respect to the various simplifications made to solve the consumer problem. We first consider psychological limitations to achieving optimal performance. Then we consider the performance implications of shifting from bundles to item-by-item searches. Finally, we consider decisions involving a subset of consumer purchases in the marketplace.

We assess the current status of research on consumer performance. It is currently impossible to quantitatively determine how close to optimal is the typical consumer. Qualitatively we assert two hypotheses. One is that individual differences are an important factor in the magnitude of differences from optimal. Second is that the magnitude of the deviation depends on the magnitude of change taking place in the life of the consumer.

Chapter 10: Improving Consumer Performance

Given the increasing value of consumers' time and the increasing difficulty of solving the consumer problem, consumers have need of improved data, strategies, and decision aids in order to improve their performance. In this Chapter we consider the role of government, business, and research to achieve this goal.

The development of the Internet provides an opportunity to develop decision aids to improve consumer performance. We develop a decision aid to select digital cameras and perform an experiment to demonstrate that the advice this aid provides is superior to the advice of sales people.