

Demand sensing in e-business

K RAVIKUMAR*, ATUL SAROOP, H K NARAHARI and
PANKAJ DAYAMA

General Motors India Science Laboratory, International Technology Park,
Bangalore 560 066, India
e-mail: ravikumar.karumanchi@gm.com

Abstract. In this paper, we identify various models from the optimization and econometrics literature that can potentially help sense customer demand in the e-business era. While modelling reality is a difficult task, many of these models come close to modelling the customer's decision-making process. We provide a brief overview of these techniques, interspersing the discussion occasionally with a tutorial introduction of the underlying concepts.

Keywords. Demand sensing; discrete choice models; reinforcement learning; latent demand modelling; econometrics; fuzzy sets.

1. Introduction

Many a time the performance of internet giants like Amazon.com makes one wonder, "Can these giants ever be defeated?", "Would there ever be a day when you would no longer be ordering books on Amazon?". Who would be the Hercules for defeating the Titans this time around? We predict that *demand sensing* will be the competitive advantage that companies of tomorrow will compete on. The one who understands her customers best in terms of product design and customization, pricing and delivery-time promises would be the market-leader in times to come. Companies are already competing on product design and pricing factors, but competing on delivery-time promises and being able to customize without compromising on inventory costs would lead to high reliance on the black magic art of *Demand Sensing*. In this paper, we try to explain the science behind the art of demand sensing.

Demand sensing, as the name indicates, refers to sensing customer purchase behaviour or, more generally, customers' choice behaviour. The scope of demand sensing can range from estimation of the price a potential customer would be willing to pay for an existing or new product and identification of his economic segment, to understanding his *latent consideration set*, the set of new products or the set of new features in products that the customer will be interested in. However, as opposed to demand forecasting, which concerns itself primarily with estimating demand for future periods of time, demand sensing pertains to assessing the current state of various factors related to customers' choice behaviour by capturing different demand signals. Specifically, when these signals are captured during the pre-buy phase of the

*For correspondence

customer and correspond to future demand, demand sensing does help in estimating demand for future periods as well.

Accurate assessment of demand and market shares is critical for many businesses and public organizations. Examples of its application include: predicting demand for a new product under alternative pricing strategies, designing a business plan for a new technology, and analysing competitive scenarios for introducing a new service. Further, advance demand sensing helps enterprises execute various supply-chain related activities in a coordinated and a cost-effective manner. With the advent of internet-enabled business practices, and with emerging paradigms like Customer Relationship Management (CRM), it has become possible for the modern enterprises to track purchase behaviour of its existing and potential customers with better precision than that an offline mode can offer.

Online market environments are a market researcher's paradise. Not only can the researcher get detailed access to data from individual customers, she can also configure the market environment to measure the signals from the customers. Given enough processing power, in future, market researchers may not even find the need to segment customers for differentiation. Such environments allow for individual customers being treated separately. Even with the current state-of-the-art technology, the market researcher can capture the customer's search process through various web-pages, which is a rich source of data for studying various stages of decision-making followed by the customer. However, with impending improvements in technology supporting online markets, we foresee that in future, marketers and market researchers would be able to observe many more signals from prospective customers, allowing them to reap much higher benefits from demand sensing.

In this paper, we provide a tutorial introduction to different modelling techniques that exist in the optimization and econometrics literature that can potentially help analyse purchase or choice behaviour of a customer.

Starting with an exposition, in §2, on typical consumer demand behaviour and associated demand management issues, we subsequently describe different models for a customer's decision-making process. While modelling reality is extremely difficult, many of these models come close to reality under different conditions. In §3, we describe single-stage and *static* decision models that model decision-making when customers attempt to compare all available alternatives and try to come up with the alternative that maximizes their utility. However, when the number of available alternatives is large, comprehending and comparing all of them at once becomes a difficult proposition. Under these conditions two-stage decision-making models come closer to reality. Two-stage decision-making models assume that customers first shortlist alternatives into a smaller consideration set, and then choose from the elements of their consideration set to form their choice set. Two-stage decision-making models are discussed in § 3.4.

Sections 4 to 6 bring in temporal dimension to choice behaviour and discuss appropriate *dynamic* models. In §4, we describe fuzzy decision-making methods that model a customer's decision-making, keeping in mind that the definition of consideration and choice sets is fuzzy in nature, that the customer's memory is imperfect and that product categories are graded in structure. We discuss a two-stage fuzzy decision-making model with consideration and choice stages in the decision-making process. However, we note that as the fuzzy spread of the consideration set increases, the customer becomes indifferent between including or excluding a brand from the consideration set; and the model then approximates a single-stage choice model.

Sections 3 and 4 assume that the alternatives that a customer might consider for purchase are explicitly known to the market researcher. This assumption might not typically be accurate

in many situations. To overcome this shortcoming, many researchers have tried to design survey instruments for finding out the complete consideration set of a customer. However, the validity and reliability of these approaches is also not beyond doubt. Hence, we describe, in §5, methods for *latent demand modelling* that incorporate consideration sets that are latent and probabilistic from the researcher's perspective.

One way to determine a customer's willingness to pay for a product is to experiment with offering the product at various prices with appropriate product differentiation. Such experiments are greatly facilitated by the internet. Reinforcement learning (RL), a stochastic approximation-based simulation technique, when coupled with nonlinear pricing, provides an appropriate framework for learning market demand curve through online experimentation. Section 6 details the RL procedure in a retailer market setting.

In §7, we identify important areas for future research.

2. Demand behaviour

2.1 *The willingness-to-pay curve*

The willingness-to-pay curve is the most direct way of describing whether customers will purchase the product and how much of the product will they purchase. Each potential customer is represented by a horizontal segment of the curve, with the price-level corresponding to the segment representing the customer's *indifference* point. The customer will be willing to make the purchase if the payment required to buy the product is below that amount. The horizontal dimension of the curve represents the quantity of the product that the customers will be willing to buy at that per-unit amount.

The willingness-to-pay curve is likely to resemble a series of steps and plateaus as depicted in figure 1. The step-like shape of willingness-to-pay curves is especially conspicuous when the products are more innovative, more distinctive, or constitute a larger or less frequent purchase. It is due to the fact that different groups of customers usually adopt the product in different contexts, to satisfy different needs, and in comparison with different alternatives. As the price of a product becomes lower, it will typically be used for completely different

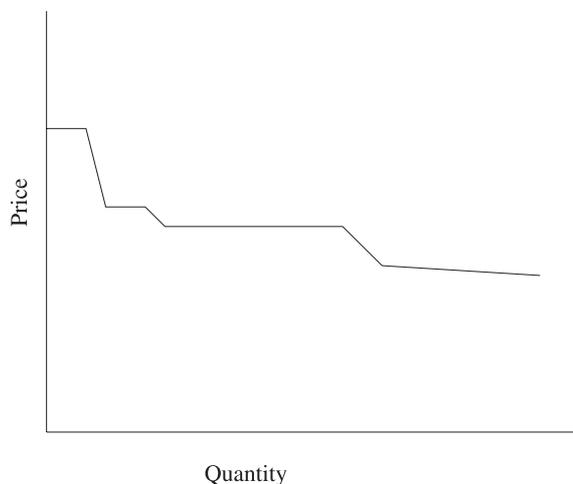


Figure 1. Willingness-to-pay curve.

purposes than it was used when its price was higher. Each new use will tend to create a new step or plateau in the willingness-to-pay curve.

The curve will also reveal situations where a relatively large change in a set price will have little effect on the quantity sold. This is the case whenever there is a large difference between the levels of successive steps or plateaus. On the other hand, when the current price is near one of the plateaus in the willingness-to-pay curve, a relatively small change in a set price will have a huge effect on the quantity sold. Often the plateau will be defined by a competing product. The step-like shape of the willingness-to-pay curves becomes less conspicuous when the products are less innovative, less distinctive, or constitute a smaller or more frequent purchase. This is so because, in these cases, there will be more gradations of alternatives and more gradations of uses. This is often the situation when a product has been around for a while, has gone through a large drop in price, is available in numerous variations, and is used constantly in a wide range of contexts. Mass-produced customer goods, whether they are desktop computers or toothbrushes, tend to fall in this category. In these situations *small* variations in the contours of a willingness-to-pay curve matter most. Volumes in these cases are typically high, margins are typically low, and fine-tuning the pricing decisions can have a huge impact.

There are two ways to determine the levels and contours of a willingness-to-pay curve. The first way is to look inside the customer's operations to see what the alternatives are and how they would affect the customer's balance sheet. For business customers, the result will generally be a set of hard numbers. It is possible to arrive at those numbers with considerable accuracy surveying various such business customers based on the utility value of the product to those customers. For private customers, it is not a hard set of numbers, because it will depend more on subjective judgments. But it can still be estimated by looking at alternative expenditures to which a purchase of the product under consideration will be compared. Such type of estimation procedures are carried out through *discrete choice models* that attempt to construct customer utility functions. These models are described in little more detail in §3.

Another way to determine the levels and contours of a willingness-to-pay curve is to offer the product at various prices. Every time a product is offered to a different group of customers at a given price, the acceptance rate indicates what portion of the customers has a willingness-to-pay higher than the designated price. Bargaining situations, while more time-consuming, can be used to locate the willingness-to-pay of a single customer with more precision. Sometimes this type of experimental test of the willingness-to-pay is necessary even when it is possible to take an inside look at the customer's operations.

Both these ways of sensing demand are greatly facilitated by the internet. The internet makes it much easier to collect information on the operations of potential customers and to determine what alternatives are available to them. The internet also makes it much easier to conduct pricing experiments, offering special prices to select populations or engaging in the bargaining process with a single customer.

Every observation about customers needs to be translatable into an observation pertaining to the level and shape of the willingness-to-pay curve. Early adopters are usually customers with high willingness-to-pay. If there is a gap or chasm between the early adopters and the next group of customers, it indicates a big difference between the first step/plateau and the next step/plateau of the willingness-to-pay curve.

In order to estimate the willingness-to-pay curve of different customer segments, companies follow various mechanisms such as price discount options, product versioning and bundling of different products. While exercising these options, companies should understand cost implications of the incremental demand. In other words, companies must seek to balance the

trade-offs between costs of different product/service options and their value to customers. We refer to the decisions resulting from these explicit trade-offs as Differentiated Service Policies. To see how these policies are implemented in practice, we cite some illustrative examples below.

Amazon.com offers different types of shipment promises, such as “usually ships in 24 hours”, “ships in three to five days”, or “special order”. The underlying reason can be attributed to the economics of inventory. Although Amazon offers approximately 4.5 million titles, it cannot afford to keep all those books in the inventory. It holds the most popular titles in its own distribution centres and typically can ship those books in 24 hours. A second tier of books is stocked by book wholesalers. Some wholesalers can fill in orders in 24 to 48 hours, enabling Amazon to meet its promise of “ships in two to three days”. Other wholesalers may take a few more days, so Amazon promises “three to five” days. At some point in time in the life-cycle of the book, wholesalers stop carrying inventory and only publishers can fulfill an order – usually from inventory, but sometimes through a new print run. The longer leadtime from the publisher forces Amazon to extend its promise to “two to three weeks”. So simple leadtime-based differentiation policy allows Amazon to address the inherent conflict between marketing and operations.

Airlines offer price-based differentiated service policy. Airlines offer widely varying prices for the same seat depending on when the customer books his reservation. The lowest prices are intended to capture incremental sales to travellers who might not take up the trip otherwise. However, such low-priced tickets should not cannibalize the sale of high-priced tickets to business travellers who may not have a choice. To ensure that the low-cost seats go to incremental travellers, airlines have traditionally imposed Saturday-night-stay requirements that the typical business traveller rarely meets. Over time, the airlines have become quite sophisticated in pricing the seats, dynamically adjusting the quantity of openings allocated to each price option based on real-time demand patterns from the reservation system. In the automotive industry too, there have been efforts (Rusmevichientong *et al* 2004) that try to optimally price vehicles through use of price ladders.

Differentiated service policies can also include factors such as different fill rates, delivery methods, quantity and price. A variation on “Pareto’s Law” or the “80/20” rule offers the key to one of the most common but least publicized techniques for differentiated service policies. Employing this concept, inventory managers apply an “ABC” segmentation of the items on hand. “A items” encompass the 5 to 10 percent of items accounting for the majority of the sales, “B items” capture the next 10 to 15 percent of the items, and “C items” cover the 80 percent of the items that typically generate only 10 to 20 percent of the sales. The classification allows the inventory manager to set “safety-stock levels” based on different expected fill rates that balance inventory carrying cost against the potential lost margin from missed shipments. In many industries, distributors typically set safety stock levels to fill orders immediately 99 percent of the time for A items, 95 percent of the time for B items, and 90 percent of the time for C items. Such a policy provides another example of differentiation in service policies that offers a cost-effective compromise between marketing and operations.

Menu pricing takes the same concept to a more sophisticated level by identifying customers by *how* they buy as opposed to differentiating on “*what*” they buy. Increasingly common among manufacturers of customer goods, menu pricing addresses the cost/value trade-offs in the retail industry value chain. Manufacturers offer better service or significant price discounts, sometimes even both, to encourage their customers to change their buying behaviour and improve the economies of the overall value chain. With menu pricing, customers choose their

preferred buying practice. We will discuss a form of menu pricing, that is, the *nonlinear pricing*, in §6 and detail how such pricing mechanism can be used to *learn* customer demand.

2.2 Cost of complexity and value of variety

The traditional corporate strategists' view that focus is the only viable approach to fending off competitors has natural limitations. At some point in a firm's evolution, its focused market, whether premised on geography, product, service, or segment, would inevitably become saturated. With globalization, the challenge for companies is not achieving a single point of focus but harmonizing among multiple points of focus. As pointed out in Oliver *et al* (2004), no company is immune from the new customer mantra: "I want what I want". In every industry, customers are demanding ever-higher levels of customization with products and services tailored to their needs. In an economy characterized by greater and greater information transparency and laden with information technology and operational advances that make customization possible, they stand an excellent chance of getting it as well. It was argued in (Oliver *et al* 2004) that smart customizers need to:

- Understand the sources of value that customization provides to their customers
- Evolve toward *virtuous variety*, the ever-changing point at which customization adds value to both company and customers.
- Tailor their business streams, aligning them to customer needs in order to provide value at the *lowest cost*

The present day's businesses are stymied by the challenges of optimizing complexity, managing the trade-off between customers' demand for variety in products and services, and the ballooning costs of meeting those needs.

Companies need to understand their "participation choice", the decisions about the markets they serve and the ways they serve them, with a clear sight of the true value of a customer segment and the cost impact of the programs they develop. When companies provide customization for too many customers without addressing the critical question "*Is the additional complexity worth the costs incurred?*", standardization opportunities or economies-of-scale are undermined. Unwarranted customization takes resources away from the highest-priority accounts and limits a company's ability to invest in more of the "right" opportunities. Typically, for most of the companies, 80 percent of the business they do is basic and stable, and transferable across product or service segments; only 20 percent of it requires some form of tailoring to customer demand. Tailored business streams match differentiated flows to differentiated markets as depicted in figure 2.

Consider the case of customized prices that is prevalent in electronic retail market. Taking better account of customers' price sensitivities can allow firms to set more advantageous prices and adjust them more dynamically, sometimes even on customer-by-customer basis. To derive maximum benefit from the pricing policies, there are ultimately two things that need to be taken into account. One is the willingness-to-pay curve of the customers and the other is the opportunity cost curve of the firm's supply-chain. It is imperative to understand (i) how these two curves are shaped, (ii) how they line up against each other, and (iii) how they are likely to change over time.

Opportunity cost is an "indifference point" or "break-even point" for the supplier. As opposed to the willingness-to-pay curve described earlier, the opportunity cost curve does not represent a whole population. It represents only the supply costs associated with one firm. In many situations, this curve is more likely to resemble a series of deepening troughs with sharp

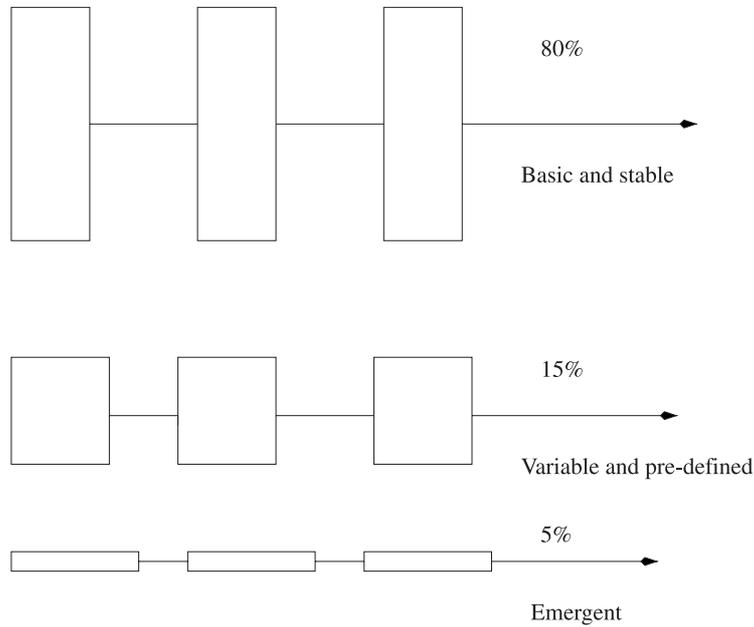


Figure 2. Tailored business streams.

peaks in between (see figure 3). It reflects the fact that the expenditure necessary to produce increasing quantities of a product do not increase in a steady manner. After relatively high investments in new production facilities and worker training have been made, the cost-per-unit will tend to fall with increasing economies-of-scale and with effects of organizational learning. But as soon as the utilization of resources touches a certain level, often about 80% capacity, cost-per-unit will once again increase as depicted in the increasing utilization curve in figure 3. The main reason for such behaviour is that with high resource utilization, the total supply-chain becomes less capable coping-up with variations in the levels of production. Even with extensive sharing of information, small fluctuations at the demand side of the chain lead to large fluctuations elsewhere down the supply-chain. Queue lengths grow, lead-times

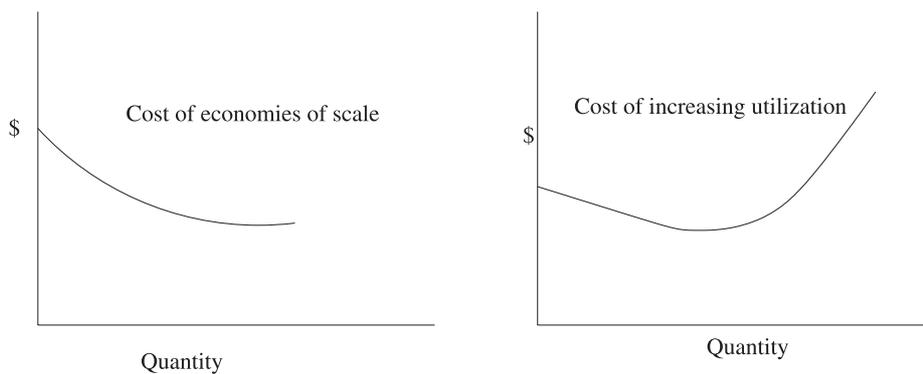


Figure 3. Opportunity cost curves.

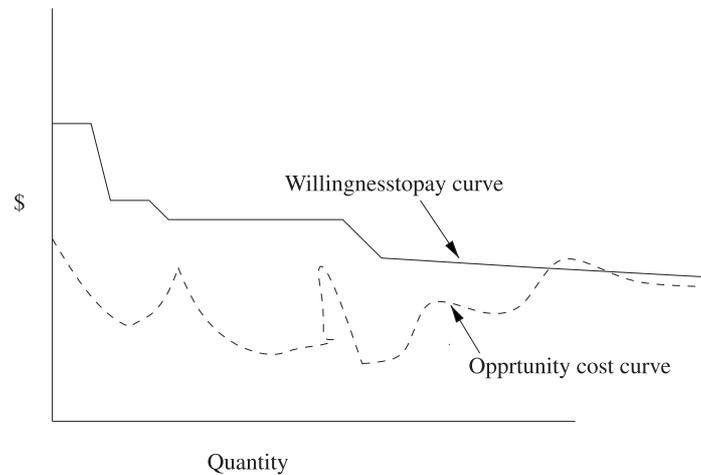


Figure 4. The pricing structure.

increase, and a phenomenon known as the bull-whip effect (Lee *et al* 1997) sets in. Additional expenditure is necessary to make up the shortages, to discount the surpluses, and to manage the growing leadtimes. Eventually, as the system gets overloaded, new facilities would need to be built, leading to a sharp spike in the costs. Hence, the opportunity cost curves resembles the shape depicted in figure 3.

Similar curves can be drawn for economies-of-scope, with a firm offering a variety of products. Product transformation curves, which represent various combinations of products that can be produced with a given set of resources, need to be studied to plan for new product introduction or for change in output levels of individual products. Pindyck & Rubinfeld (2000) give an interesting exposition of costs involved in economies-of-scope.

For successful implementation of customization, it is imperative to understand both the willingness-to-pay curve and the opportunity cost curve. If, via customization, a company is able to charge each customer the maximum that she is willing to pay, the profit would be the customer's willingness-to-pay minus the company's opportunity cost for supplying that customer. The area between the two curves would measure the value obtained from following such a differentiated pricing structure (see figure 4). Interestingly, in such a situation, the customer would not be deriving any value. However, in most practical cases, the value created is divided between the supplier and the customer. The price can be anywhere between the customer's willingness-to-pay and the supplier's opportunity cost without anyone losing on the deal. The portion of the value captured by each individual party decides the value of opportunity created by customization. It is in the manufacturer's interest to allow the customer to capture enough portion of the value so as to attract more of them, thereby making it harder for a competitor to lure them away.

3. Discrete choice models for demand sensing

In this section, we discuss econometrics-based approaches to demand modelling. One such modelling approach that has gained ground amongst the marketing research community is the discrete choice modelling, an approach that effectively fuses statistics and mathematical methods to develop predictive models for customer choice behaviour with considerable

accuracy. The cost-less data acquisition capabilities of today's web-technologies have helped researchers to use these models to their fullest potential. Discrete choice models are powerful but complex. The art of finding the appropriate model for a particular application requires an analyst to have close familiarity with the reality and a strong understanding of the methodology and theoretical background of the model.

Given the fact that the methodology has reached a stage of maturity, giving a comprehensive survey of chronological developments is beyond the scope of this paper. Instead, we provide a foundational introduction to the methodology, (giving references to some important contributions in the passing), and introduce concepts needed for the development of appropriate models needed to explain choice behaviour in the e-business era. For a detailed and excellent exposition on these models, we direct readers to (Train 2003). The treatise that appears below derives its motivation from this work.

Discrete choice models describe decision makers' choices over possible alternatives. The alternatives constitute the *choice set*. The alternatives of the choice set must be *mutually exclusive* and *collectively exhaustive* and the cardinality of the set must be *finite*. The first and second properties are not as restrictive as they might seem at the outset. By defining the alternatives appropriately, it is always possible to make the choice set satisfy these requirements. For instance, suppose two alternatives *A* and *B* are not mutually exclusive. Then the alternatives can be redefined as *A only* and *B only*, and *both A and B*, which are necessarily mutually exclusive. Similarly, if all the alternatives are not exhaustive, one can expand the choice set to include the extra alternative, namely, *none of the other alternatives*, and thus make it exhaustive.

Finiteness of the choice set is the defining characteristic of discrete choice models and distinguishes their realm of applications from that for regression models. Discrete choice models are usually derived under an assumption of utility-maximizing behaviour of the decision maker. Thurstone (1927) originally developed in terms of psychological stimuli, leading to a binary probit model of whether respondents can differentiate between the levels of stimulus. Marschak (1960) interprets stimuli as utility and provide a derivation from utility maximization. These models are referred to as Random Utility Models (RUMs) and are derived as described below.

3.1 Random utility models

Assume that a customer, labelled *n*, faces *J* alternatives and each alternative offers certain level of real-valued utility to him. Let the utility for *n* from the alternative *j*, $j = 1, 2, \dots, J$, be U_{nj} . The customer chooses the alternative that provides the greatest utility. In other words, the behaviour model of the customer is to select the alternative j^* such that

$$U_{nj^*} > U_{nk} \forall k \neq j^*.$$

However, marketing researchers will not be able to observe $\{U_{nj}, n > 0, j = 1, 2, \dots, J\}$ but instead can observe only some attributes x_{nj} corresponding to the alternative *j* and attributes, s_n corresponding to the customer *n* that will help estimate U_{nj} . Let $V_{nj} = h(x_{nj}, s_n)$ be an estimate for U_{nj} for some function *h*. The functional form of *h* in general can be prescribed at an individual customer level or at a customer segment level. Since there are aspects of utility that cannot be observed, $V_{nj} \neq U_{nj}$ in general. In other words, V_{nj} can be written as,

$$U_{nj} = V_{nj} + \epsilon_{nj},$$

for some random error ϵ_{nj} with a specified distribution.

Let $\epsilon_n = [\epsilon_{n1}, \epsilon_{n2}, \dots, \epsilon_{nj}]$ be the random vector of errors and the joint density of ϵ_n be denoted by $f(\cdot)$. Now, the probability P'_{ni} that customer n chooses alternative i is given by:

$$\begin{aligned} P'_{ni} &= P(U_{ni} > U_{nj}, \forall j \neq i) \\ &= P(V_{ni} + \epsilon_{ni} > V_{nj} + \epsilon_{nj}, \forall j \neq i) \\ &= P(\epsilon_{nj} - \epsilon_{ni} < V_{nj} - V_{ni}, \forall j \neq i) \end{aligned} \quad (1)$$

With an abuse of notation, we use ϵ_n to denote realized values of the random variable ϵ_n above. Now, we can rewrite (1) as:

$$P'_{ni} = \int_{\epsilon_n} I_{\{\epsilon_{nj} - \epsilon_{ni} < V_{nj} - V_{ni}, \forall j \neq i\}} f(\epsilon_n) d\epsilon_n, \quad (2)$$

(2) is a multi-dimensional integral over the density of the *unobserved* portion of U . Different choice models differ in their specifications of this density and only a few of them yield closed form expressions for the above integral.

In *logit* model, the most popular discrete choice model, it is assumed that ϵ_{ni} are *i.i.d* extreme value for all i . The critical part of the assumption is that unobserved factors are uncorrelated across alternatives. Though this assumption is restrictive, it provides a convenient way of evaluating the above integral. Also, it assumes that choices across periods are independent. Correlation can be factored in the generalized extreme value (GEV) model. The GEV model partitions alternatives into groups, where members of a group are correlated while the alternatives across groups are uncorrelated.

The *probit* models assume that $\epsilon_n \sim N(0, \Sigma)$ where Σ is the covariance matrix. That is, unobserved factors are assumed jointly normal (even over time). Σ can accommodate any pattern of correlation and heteroskedasticity. The flexibility of the probit model in handling correlations over alternatives and time is its main advantage. Its only functional limitation stems from its dependence on normal assumption; which will not be an appropriate assumption, for example, in cases when customers' utilities are strictly positive.

A fully general discrete choice model is the *mixed logit* model that allows unobserved factors to follow any distribution. The defining characteristic of mixed logit model is that the unobserved factors can be decomposed into two parts, one that accounts for correlation and heteroskedasticity and the other that is *i.i.d* extreme. The mixed logit model can approximate any discrete choice model.

3.1a Logit model: The relation of the logit formula to the distribution of unobserved was developed by Marley (as cited in (Train 2003)), who showed that the extreme value distribution leads to the logit formula. We present here the analysis underlying the derivation of the logit formula. To this end, we need the following definition:

DEFINITION 1

For a given set of *i.i.d* random variables $X_i, i = 1, 2, \dots, N$, an *extreme value distribution* is the distribution of the extreme order statistic $X^{(1)}$ (the maximum of X_i 's) or $X^{(N)}$ (the minimum of X_i 's).

Remark 1. If the parent distribution is well-behaved, (*i.e.*, $F(\cdot)$ is continuous and has an inverse), only a few models are needed to describe the limiting distribution for the above order statistics, depending on whether one is interested in maximum or minimum, and whether the

initial distribution has a bounded tail or not. In the following, we refer to the appropriate limiting distributions as the extreme value distributions.

Remark 2. The largest member of a sample of size N , for large N , has the Gumbel distribution, if the parent distribution has finite moments and an unbounded tail that decreases at least as fast as the exponential function (as does the normal distribution for example). Its *pdf* is given by

$$f(x|\theta_1, \theta_2) = \frac{1}{\theta_2} e^{-z-e^{-z}} \quad (3)$$

where $z = \frac{(x-\theta_1)}{\theta_2}$, and θ_1, θ_2 are *location* and *scale*¹ parameters, respectively, and $\theta_1 > 0$.

Now, let the customer, labelled n , face J alternatives. The utility that the customer obtains from alternative j is decomposed into (i) a part labelled V_{nj} that is observable up to some parameters, and (ii) an unknown part ϵ_{nj} that is treated as *random*. So the actual utility to the customer is: $U_{nj} = V_{nj} + \epsilon_{nj} \forall j$. Assume that each ϵ_{nj} is i.i.d extreme value. The density of each unobserved component of utility is

$$f(\epsilon_{nj}) = e^{-\epsilon_{nj}} e^{-e^{-\epsilon_{nj}}}, \quad (4)$$

and the distribution function is

$$F(\epsilon_{nj}) = e^{-e^{-\epsilon_{nj}}}. \quad (5)$$

The variance of this distribution is $\pi^2/6$. The mean of the extreme value distribution is non-zero. However, it is important to note that, since only difference in utilities matter for choice probability, the difference between two *i.i.d* random errors has a mean of zero. Also, interestingly, this difference follows a *logistic* distribution, that is, if $\epsilon_{nji}^* = \epsilon_{nj} - \epsilon_{ni}$, then,

$$F(\epsilon_{nji}^*) = e^{\epsilon_{nji}^*} / (1 + e^{\epsilon_{nji}^*}). \quad (6)$$

Assuming extreme value distribution for the errors is the same as assuming that the errors are independently normal, the extreme value distribution has a slightly fatter tail than normal, but the difference is indistinguishable empirically.

3.1b Logit choice probabilities: The probability that customer n chooses alternative i is

$$\begin{aligned} P'_{ni} &= P(V_{ni} + \epsilon_{ni} > V_{nj} + \epsilon_{nj} \forall j \neq i) \\ &= P(\epsilon_{nj} < \epsilon_{ni} + V_{ni} - V_{nj} \forall j \neq i) \end{aligned} \quad (7)$$

Because of *i.i.d.* assumption on ϵ 's, it follows that

$$P'_{ni} = \int \left(\prod_{j \neq i} e^{-e^{-(\epsilon_{ni} + V_{ni} - V_{nj})}} \right) e^{-\epsilon_{ni} - e^{\epsilon_{ni}}} d\epsilon_{ni}. \quad (8)$$

¹Note that a probability density is a **location, scale density** if it can take the form $P(X \leq x) = F(x|\theta_1, \theta_2) = \Phi[(x - \theta_1)/\theta_2]$ where Φ is a proper density and does not depend on any unknown parameters.

With some algebraic manipulations, (8) reduces to

$$P'_{ni} = e^{V_{ni}} / \sum_j e^{V_{nj}}. \quad (9)$$

Generally, $V_{ni} = h(\mathbf{x}_{ni})$ for a h on observed parameter vector \mathbf{x}_{ni} . h can be obtained either from a linear regression or from a *neural network* or from any such function approximation.

Under fairly general conditions, any function can be approximated by a function that is linear in parameters. Let $V_{ni} = \alpha^T \mathbf{x}_{ni}$. With such specification, the choice probability of alternative i is

$$P'_{ni} = e^{\beta^T \mathbf{x}_{ni}} / \sum_j e^{\beta^T \mathbf{x}_{nj}}. \quad (10)$$

The above linearization yields a log-likelihood function with the above choice probabilities that is globally concave in parameters β . Hence it is convenient to work with the linear approximation particularly in the context of numerical maximization procedures.

Figure 5 plots the relation between logit probability and representative utility. Its shape is *sigmoid* and has implications for the impact of changes in explanatory variables. If the representative utility of an alternative is very low compared to other alternatives, a small increase in the utility of the alternative has little effect on the probability of its being chosen. Similarly, if one alternative is far superior to the others on observed attributes, a further increase in its representative utility has little effect on the choice probability. The point at which the increase in V_{ni} has the greatest effect is when the probability is close to 0.5. The sigmoid shape is shared by most discrete choice models and has important implications in decision-making. For example, it illustrates the fact that improving level of comfort/appeal through modifying feature in an item when the demand for that feature is so poor that few customers opt for it would be less effective in terms of improved sales than making the same level of improvement in comfort through modifying a feature that has good demand.

To motivate practical use of the logit model, consider a decision maker's problem of estimating a household's choice between a gas and an electric heating system. Suppose that utility of the household from each choice depends only on the purchase price, the operating cost, and

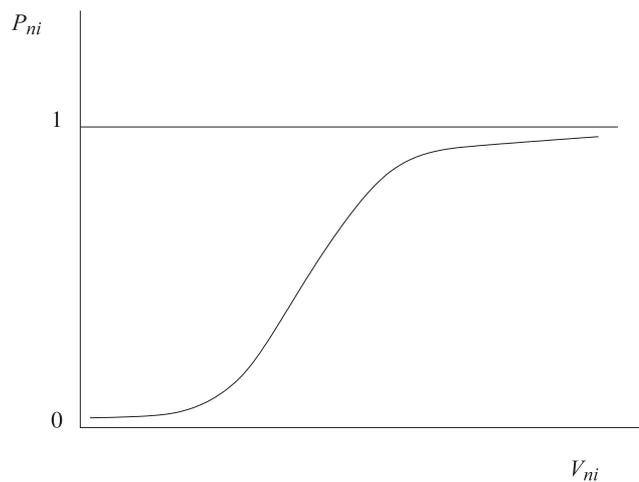


Figure 5. Choice probability of alternative i .

the household's view of the convenience and quality of heating with each type of the system and the relative aesthetics of the systems within the house. The first two factors are observable by the decision maker but not the other two. If one models the observed part of the utility by a linear function, then the utility to the household of each type of the system can be written as

$$U_{\text{gas}} = \alpha_1 P_g + \alpha_2 C_g + \epsilon_{\text{gas}}, \quad (11)$$

$$U_{\text{electric}} = \alpha_1 P_e + \alpha_2 C_e + \epsilon_{\text{electric}}, \quad (12)$$

where α_1 and α_2 are scalar parameters for purchase price P and operating cost C respectively. Note that $\alpha_1 < 0$ and $\alpha_2 < 0$. The unobserved portion of the utility varies over households depending on how each household views quality, aesthetics, and convenience. If these are *i.i.d* extreme value, then the household's probability of choosing a gas system is

$$P'_{\text{gas}} = e^{\alpha_1 P_g + \alpha_2 C_g} / (e^{\alpha_1 P_g + \alpha_2 C_g} + e^{\alpha_1 P_e + \alpha_2 C_e}). \quad (13)$$

As in economics, the ratio α_2/α_1 represents the household's willingness-to-pay for operating cost reductions and it is the increase in the purchase price that keeps the household's utility constant given a reduction in operating costs. In other words, $\partial P/\partial C = -\alpha_2/\alpha_1$.

3.1c Scope of logit models: The wide use of logit models is mainly due to their simplicity and power to model different choice behaviours. Most important among them are its ability to represent taste variations in observed characteristics, proportional substitution across alternatives and repeated choices over time.

Taste variations – The value that customers place on each alternative varies across customer segments, or more generally across customers. For example, size of a car is probably more important to households with many members than to smaller households. Low-income households are more concerned about the price of the car than any other attribute. Further, there can be differences in individual preferences though all the individuals belong to same economic segment. Logit models can capture tastes that vary systematically with respect to *observed* characteristics.

To be more precise, consider customers' choices among different brands and models of cars. Suppose that the only observable characteristics are price and roominess. The value that a customer places on these alternatives differs among customers. A linear model can be developed to capture individual customer preferences associating a customer specific parameter in the utility model:

$$U_{ni} = \alpha_n R_i + \beta_n P_i + \epsilon_{ni}, \quad (14)$$

where α_n and β_n are parameters specific to customer n . However, estimating these individual parameters is not possible in general. If these parameters are *uniformly* related to some observable characteristics in a linear or nonlinear fashion, the same model can still be applied to estimate these parameters. For instance, assume that preference to roominess is a function of members, M_n in the household of the customer n , then

$$\alpha_n = \gamma M_n, \quad (15)$$

and suppose that the willingness-to-pay is *inversely* proportional to his income I_n as given below:

$$\beta_n = \rho / I_n. \quad (16)$$

Then customer n 's utility can be rewritten in the form of a standard logit model as,

$$U_{nj} = \gamma(M_n R_i) + \rho(P_i/I_n) + \epsilon_{ni}. \quad (17)$$

If the tastes vary purely randomly in a non-uniform manner or vary in unobservable characteristics, then the logit model is not an appropriate choice to address such issues. For instance, in (15), if there is an additional random term, say to account for size of individual, then this additional term adds to the unobserved variables and the new error terms need not be from an *i.i.d* sequence.

Substitution effect – When the attributes of one alternative improve, the probability of its being chosen rises; a fraction of the people who would have chosen other alternatives will choose the improved alternative instead. In many decision situations it is important to analyse the substitution effects. For instance, when a cell-phone manufacturer plans to introduce a new product, he would be more interested in knowing what percentage of customers will be drawn away from other products rather than from competitor's phones.

Logit models rely on the property of *Independence from Irrelevant Alternatives* (IIA), to estimate levels of substitution or more generally, cross elasticities of alternatives. Note that ratio of choice probabilities of two alternatives under the logit model,

$$P'_{ni}/P'_{nj} = e^{V_{ni}}/e^{V_{nj}}, \quad (18)$$

are independent of other alternatives.

More flexible models are needed if IIA assumption is violated, which is the case when introduction of a new alternative disturbs choice probabilities of all other alternatives. Consider, for example, the famous red-bus-blue-bus problem (Chipman 1960). A traveller has a choice of commuting by a car or a blue bus. Assume that the observed utility for both the alternatives are the same. Then $P'_{\text{bluebus}} = P'_{\text{car}} = 1/2$ and the ratio of choice probabilities is 1. Now assume that a red bus is introduced as another alternative for commuting and the traveller considers the red bus to be exactly like the blue bus. Then the choice probabilities for the two variations of the buses are the same and, hence, the ratio of probabilities, $P'_{\text{bluebus}}/P'_{\text{redbus}}$ should equal 1. However, note that according to logit model, $P'_{\text{car}}/P'_{\text{bluebus}} = 1$ which with $P'_{\text{bluebus}}/P'_{\text{redbus}} = 1$ yields $P'_{\text{car}} = P'_{\text{bluebus}} = P'_{\text{redbus}} = 1/3$. However, we should expect the choice probability of car to remain the same if the new bus introduced is exactly the same as the old bus. That is, we would expect, $P'_{\text{car}} = 1/2$ and $P'_{\text{bluebus}} = P'_{\text{redbus}} = 1/4$. Hence, the ratio $P'_{\text{car}}/P'_{\text{bluebus}} (= 2)$ actually changes with the introduction of the new bus, rather than remaining constant as required by the logit model. A similar result holds even when the new bus introduced has the feature, say, express service. Introduction of such alternative reduces the choice probability of regular bus by a greater proportion than it reduces the choice probability of the car, and hence, the ratio of probability of car and the old bus do not remain the same. Thus, more flexible models are needed to assist in understanding such substitution effects.

More generally, logit model is an appropriate choice only when a change in attributes of one alternative changes the probabilities of other alternatives by the same percentage. This is referred to as *proportional substitution*. For a detailed study on the substitution effects, the reader is urged to refer to (McFadden 1978).

Repeated choices – In many settings, market researchers can observe numerous choices made by each customer. Data on the current and past vehicle purchases of sampled customers might

be obtained by market researcher who is interested in dynamics of car choice. In market surveys, respondents are often asked a series of hypothetical choice questions, called *stated preference* experiments. A series of questions is asked over different variations of attributes corresponding to a product. Data that represents *repeated* choices like these is called panel data.

If unobserved factors that affect the choice are independent over the repeated choices, then logit model can be used to analyse panel data in a similar way as cross-sectional data. An additional situation-dependent dimension can be added to the cross-section level model.

The utility that customer n obtains from alternative i in period or situation t is

$$U_{nit} = V_{nit} + \epsilon_{nit}. \quad (19)$$

If ϵ_{nit} is distributed extreme value, independent over all n , i , and t , the choice probabilities take the form

$$P'_{nit} = e^{V_{nit}} / \sum_j e^{V_{njt}} \quad (20)$$

and V_{nit} can be approximated as in the earlier case by a linear function or neural network or by any other function approximation model.

Dynamic aspects of behavioural response can be captured by specifying observed utility in each period or situation to depend on observed variables from other periods. For example, a lagged price response is represented using price in period $t - 1$ as explanatory variable in the utility model for period t . Also, for example, to capture *inertia* or habit formation in customer's choice behaviour, such as customer's tendency to stay with the alternative that he has previously chosen unless another alternative offers sufficiently high utility to warrant a switch, can be modelled as,

$$V_{nit} = \alpha y_{ni(t-1)} + \beta x_{nit}, \quad (21)$$

where y_{nit} is the indicator variable assuming value 1 if the alternative i is chosen in period t and value 0, otherwise. If $\alpha > 0$, the utility of i is higher in period t if it is chosen in period $t - 1$. If $\alpha < 0$, then it pays not to choose the alternative previously selected. y_{nit} is assumed to be uncorrelated with the current error ϵ_{nit} . The situation is analogous to linear regression, where a lagged dependent variable can be added without inducing bias as long as the errors are independent over time.

If observed factors are correlated over time, one would naturally expect unobserved factors or errors also to be correlated. In such situations, one is advised to use other models such as probit or mixed logit, described in the sequel.

3.2 Simulation-based discrete choice models

As pointed out earlier, models which are more flexible than logit are needed to address issues related to correlations in unobserved factors. Also, it is often true that different *customer segments* exhibit different choice behaviours. So it is important to include segment-specific information in the model for estimation of choice probabilities. These features can be incorporated in *mixed logit* model to be detailed below.

3.2a Mixed logit model: Mixed logit model is a very general and flexible model that can approximate any random utility model (McFadden & Train 2000). It generalizes the standard logit by allowing for random taste variations, substitution patterns that are more general than

the ones that satisfy IIA and also, allows for correlations in errors. As one would expect, any such generalized model cannot yield closed form expressions for choice probabilities and often one has to rely upon simulation-based techniques. Interestingly, in the case of mixed logit, simulation of choice probabilities turns out to be computationally simple.

The first application of mixed logit model was apparently the automobile demand models jointly created by (Boyd & Mellman 1980; Cardell & Dunbar 1980). With the advent of Internet-enabled commerce and emerging web-based technologies, it has become easy for market researchers to gather customer specific data. Further, improvements in the speed of computers and in simulation methodologies have given enough scope and power for mixed logit models to be utilized to the fullest extent.

Mixed logit models are defined on the basis of the *functional form* of their choice probabilities. Any behavioural specification that takes the following specific functional form is referred to as mixed logit model.

A *mixed logit* model is any model whose choice probabilities can be expressed in the following form,

$$P'_{ni} = \int_{\beta} L_{ni}(\beta) f(\beta) d\beta, \quad (22)$$

where L_{ni} is the logit probability, with respect to parameter value β , that customer n chooses alternative i , as given in (8), and $f(\beta)$ is the density function of the parameter vector β . (Recall the abuse of notation for random variables and their realized values.)

If the observed utility, V_{ni} is linear in values of attributes, then with $V_{ni} = \beta^T x_{ni}$, the choice probabilities under mixed logit become

$$P'_{ni} = \int_{\beta} \left(e^{\beta^T x_{ni}} / \sum_j e^{\beta^T x_{nj}} \right) f(\beta) d\beta. \quad (23)$$

Hence, mixed logit is the weighted average of logit probabilities evaluated at different β 's with $f(\beta)$ as the mixing function. $f(\beta)$ can be a discrete distribution. Suppose that f has support on $\{p_1, \dots, p_M\}$, then the choice probability for alternative i is

$$P'_{ni} = \sum_{k=1}^M p_k \left(e^{\beta^T x_{ni}} / \sum_j e^{\beta^T x_{nj}} \right). \quad (24)$$

Equation (24) above is the *latent class model*, the popular model in marketing and psychology. This model is useful if customer population can be grouped into M segments and all customers of a given segment exhibit the same choice behaviour.

In most cases, $f(\beta)$ is assumed to be continuous. For example, if the density of β is normal with mean μ and covariance Ω , then, choice probabilities are,

$$P'_{ni} = \int_{\beta} \left(e^{\beta^T x_{ni}} / \sum_j e^{\beta^T x_{nj}} \right) \psi(\beta | \mu, \Omega) d\beta, \quad (25)$$

where $\psi(\cdot)$ is the normal density function with parameters μ and Ω .

Note that there are two sets of parameters in the mixed logit model; β , that enters the logit formula and the other set that describes the density $f(\beta)$. Let the parameters that define the density of β be denoted by θ or more appropriately, the density of β be denoted by $f(\beta|\theta)$. With this specification the choice probabilities $P_{ni} = \int L_{ni}(\beta) f(\beta|\theta) d\beta$ are functions of

θ . This is a convenient representation if one wants to obtain information about β 's for each sampled customer. θ helps represent variations in customers' tastes.

Many choice models invoke mixed logit function whenever the utilities of customers are represented using random coefficients. That is, assume that the utility customer n derives from alternative i is of the following form:

$$U_{ni} = \beta_n x_{ni} + \epsilon_{ni}, \quad (26)$$

where β and ϵ are random, the latter being *i.i.d* extreme value. The coefficients β_n vary over customer population according to the density $f(\beta)$ and this density is a function of parameters θ . This specification is the same as for standard logit model except that β varies over customers rather than being fixed. If β_n 's are observed, then the choice probability would be standard logit. That is, *conditional* on β_n the probability is

$$L_{ni}(\beta_n) = e^{\beta_n^T x_{ni}} / \sum_j e^{\beta_n^T x_{nj}}, \quad (27)$$

which when unconditioned will yield the mixed logit formula.

The most common specification for density of β is normal or log-normal. Log normal specification is useful if the coefficient has same sign for all the customers, for example, in case of coefficient to price that is known to be negative for everyone.

Remark 3. Mixed logit does not follow IIA assumption for substitution of alternatives, and hence can model different substitution patterns. It is easy to see that percentage change in one alternative given a change in the k th attribute of other alternative is given by

$$\epsilon_{ni x_{nj}^k} = -(1/P_{ni}) \int \beta^{(k)} L_{ni}(\beta) L_{nj}(\beta) f(\beta) d\beta, \quad (28)$$

where $\beta^{(k)}$ is the k th component in the β vector. This *elasticity* is different for each alternative.

3.2b Simulation-based estimation: Mixed logit is well-suited to simulation methods for estimation. Let customer utility be,

$$U_{ni} = \beta_n V_{ni} + \epsilon_{ni}, \quad (29)$$

where β_n are distributed with density $f(\beta|\theta)$, where θ is the vector of parameters that define f . Assume that the functional form of f is known and it remains to estimate the parameter θ . Recall that the choice probabilities are

$$P_{ni} = \int_{\beta} L_{ni}(\beta) f(\beta|\theta) d\beta, \quad (30)$$

where

$$L_{ni}(\beta) = e^{\beta^T x_{ni}} / \sum_j e^{\beta^T x_{nj}} \quad (31)$$

These probabilities can be obtained through simulation for any given value of θ as follows.

- (1) Draw a value of β from $f(\beta|\theta)$. Let $\beta^{(k)}$ denote the k th draw.
- (2) Calculate the logit formula $L_{ni}(\beta^{(k)})$ for the k th draw,

- (3) Repeat the above steps and average the results. The average yields the simulated probability

$$\check{P}'_{ni} = (1/M) \sum_{k=1}^N L_{ni}(\beta^{(k)}), \quad (32)$$

where M is the number of draws. \check{P}'_{ni} is an unbiased estimator of P'_{ni} by construction.

Remark 4. The following hold true:

- \check{P}'_{ni} sums to one over all alternatives.
- The variance of \check{P}'_{ni} decreases as N increases.
- \check{P}'_{ni} is strictly positive and hence, $\ln \check{P}'_{ni}$ is well-defined, so is useful in estimating the log likelihood function.
- \check{P}'_{ni} is smooth (twice differentiable) in θ and in x and hence, facilitates search for maximum likelihood function and the calculation of elasticities.

Above simulated probabilities can be inserted into the log-likelihood function to give a simulated log-likelihood:

$$\check{L}L = \sum_{n=1}^n \sum_{j=1}^J d_{nj} \ln \check{P}'_{nj}, \quad (33)$$

where $d_{nj} = 1$ if n chooses j and zero otherwise. The maximum likelihood estimator is the value of θ that maximizes $\check{L}L$. This maximal value is the estimate of the parameter θ that makes the observed data most likely. For advances in simulation techniques for different choice models, the reader is directed to (McFadden & Train 2000). As indicated in this section, probit models assume that unobserved error vector has multi-variate normal distribution. Similarly, GEV assume generalized extreme value distribution for errors. In these cases too, simulation-based methods are needed for estimation of choice probabilities. For efficient estimation techniques for these models, refer Train (2003).

3.3 Discrete choice models – data related issues

Primarily, two different data sets are used in marketing surveys; *revealed-preference data* that reflects actual choices made by customers in the real world and *stated-preference data*, the data collected in experimental/survey situations where customers respond over hypothetical choice situations. While the revealed preference data are limited to the choice situations and attributes of alternatives that currently exist or have existed historically, the stated-preference data offers an advantage that experiments can be designed to contain rich variations in attributes to *sense* a customer's behaviour. Also, as pointed out by Train (2003), even in situations that currently exist, there may be insufficient variation in relevant factors to allow for estimation with revealed-preferences.

Web technologies have enabled cost-effective implementation of various forms of experiments over customers to gather data pertaining to stated preferences. Customer's choice behaviour is generally elicited from questionnaires of different formats. For example, a customer may be asked to *rank* alternatives presented to him. Models that can handle such questionnaire will be able to track choice behaviour with a better precision than the models that

work on single alternative selected on utility-maximization principle. In the cases with ranked data, choice probabilities cannot be directly derived from the above models. With a modelling artifice, the logit and mixed-logit models can still be adapted to deal with such situations. For details, refer McFadden & Train (1996).

3.4 Two-stage decision-making models for buying

In §3, we have dealt with single-stage discrete choice models, where the customer makes a single-step buying decision by considering all possible alternatives in her choice set. After constructing the mutually exclusive and collectively exhaustive choice set, she then proceeds to compare them. Such models explain decision-making by customers to a large extent when the number of alternatives is small. As the number of alternatives grows large, it becomes difficult for the customers to comprehend the complete set of possible alternatives in a single shot. To account for such situations, Manski (1977) proposed the use of a two-stage decision-making model, wherein the customer first builds a consideration set out the available alternatives in the first stage (called as the *Consideration Stage*) and compares the selected alternatives in the second stage (called as the *Choice Stage*) for making the final choice. See Shocker *et al* (1991) and Roberts & Lattin (1991) for a review of these models.

Shocker *et al* (1991) cite some empirical research that supports the existence of consideration sets. Hulland (1992) conclude through experimentation that given enough prior choices, many customers consider only previously selected alternatives in their consideration set. They note that in many cases customers would be unwilling to re-evaluate previously rejected alternatives.

Models of consideration sets – Customer consideration sets can be modelled in different ways. Roberts & Lattin (1991) construct customer consideration sets by incrementally building them up, considering addition of one alternative a time. However, they also note that the marginal utility of adding an extra alternative to the set approaches zero as the set grows in size. To ensure that the consideration set hence built is finite, alternatives are added only till their addition makes substantial difference to the utility of the overall set. This means that the consideration set hence obtained becomes dependent on the order in which individual alternatives were considered for addition.

Manski (1977) and Andrews & Srinivasan (1991) construct a model with crisp consideration sets, where these sets are not visible to the market researcher. The probability of a set M being the consideration set for the customer is specified for all possible consideration sets. These probabilities are calculated, in turn from the outcome of the choices made by the customer. The probability of choosing an alternative j , hence, is

$$p_j = \sum_{M \in U} p(j|M) \cdot p(M),$$

where $p(M)$ denotes the probability of M being the consideration set and $p(j|M)$ is the probability of choosing alternative j given that M is the consideration set. However, even with moderate numbers of alternatives available, the model becomes computationally difficult to solve. To overcome this problem, Siddarth *et al* (1995) construct a model where customers choose a consideration set M with a probability q and choose any other consideration set with probability $(1 - q)$. Let $p_i(U)$ denote the probability of choosing alternative i when the consideration set can be any member of the universal set U . Then, we can state that

$$p_i(U) = q \cdot p_i(\{M\}) + (1 - q) \cdot p_i(U - \{M\}).$$

Ben-Akiva & Boccara (1995) use a constraint-based approach for generating choice sets for customers. They state that only those alternatives that meet certain criteria enter the choice set of a customer. Their model allows for criteria that are specific to the alternative under consideration. They estimate the probability of an alternative satisfying its criteria through the probability of its availability to customers in the population being surveyed. On the other hand, Bronnenberg & Vanhonacker (1996) use a threshold-based approach where a customer includes an alternative in her choice set only if its features (measured in terms of brand salience) exceed a customer-specific threshold.

However, two-stage decision-making models are also not beyond doubt. Chiang *et al* (1999) state that as the market researcher rarely gets visibility into the customer's decision-making process, it is probably better to build a model where probability of a set M being the customer's consideration set, $p(M)$ has mass points over the universal set U , and is modelled as a Dirichlet distribution with weak prior information about the individual $p(M)$'s. They build an integrated model that accommodates consideration set and parameter heterogeneities across the population of customers. Their results show that if consideration set heterogeneity is ignored in modelling, the effect of marketing mix on customer response and market share is understated. With inclusion of consideration set heterogeneity, the effect of brand specific intercepts and the importance of past purchase effects are lowered. So, omitting consideration set heterogeneity understates the impact of marketing variables and overstates the impact of preferences. However, results show that a pure consideration set model is fairly robust to inclusion or exclusion of parameter heterogeneity.

With regards to impact of promotions, the proposed model concurs with the results of (Chintagunta *et al* 1991) in the sense that including effects of heterogeneity (consideration set heterogeneity in this model and parameter heterogeneity in Chintagunta *et al*'s model) lead to promotions showing bigger impact on customer choices and market share. Chiang *et al* build a multinomial Logit model through which choice of brand is conditioned on consideration sets and parameters. So, at an occasion t , the choice of household i , denoted by y_{it} would have the following probabilities associated:

$$\text{Prob}(y_{it} = j | \{\beta, b_i, C_i\}) = \begin{cases} \frac{\exp[x_j(\beta + b_i)]}{\sum_{k \in C_i} \exp[x_k(\beta + b_i)]}, & \text{if } j \in C_i, \\ 0, & \text{otherwise,} \end{cases} \quad (34)$$

where x_j denotes the attributes of alternative j , while β and b_i represent the alternative choice parameters and random effects representing parameter heterogeneity respectively.

The consideration set models described in this section are crisp set models that are at best probabilistic in nature. Consideration sets are modelled as crisp sets and probabilities of a particular set being the customer's consideration set and of the customer choosing an alternative given her consideration set are estimated. In §4, we describe models developed by Fortheringham (1988), Bronnenberg & Vanhonacker (1996) and Wu & Rangaswamy (2003) among others that consider fuzzy consideration sets. Notably amongst these, the model proposed by Wu & Rangaswamy (2003), though based on two-stage decision-making, is also able to approximate single-stage decision-making models described in this section, when the fuzziness present during decision-making is high.

4. Fuzzy sets approach

While making the buying decision, the customer is faced with the choice amongst brands that fall into a graded structure of product categories. Compounded with the fact that customers

have an imperfect memory of their past preferences and purchases, it is reasonable to view product categories, consideration (and choice) sets as fuzzy sets and the decision-making problem as a fuzzy set decision-making problem. Using fuzzy sets, a modeler can account for effects of repeated customer exposures to various brands and can create a model that generalizes both single-stage and two-stage decision-making models.

In literature, buying decision has been modelled using fuzzy sets by researchers including Fortheringham (1988), Bronnenberg & Vanhonacker (1996) and Wu & Rangaswamy (2003). In their work, Fortheringham (1988) describe a model with fuzzy consideration sets, where as a part of two-stage decision-making process, customers select a store from a fuzzy cluster of stores. Having selected a store, they choose a product from that specific store. Wu & Rangaswamy (2003) also adopt a two-stage decision-making model with consideration stage and choice stage to construct their model. They posit that in the consideration stage, customers form their consideration utilities for alternatives. Those alternatives that have utilities larger than a threshold are then included in the choice stage. Both utilities and thresholds are modelled as fuzzy real numbers. During the choice stage, the customers are assumed to maximize their fuzziness-adjusted utility (similar to the concept of risk-adjusted utility concept of Roberts & Urban (1988)). However, as the modeler is unable to specify the customer's utility function completely, a random utility framework similar to that described in §3.1 is used.

While making the purchase decision, the customer is swayed by many factors. On typical online markets, customers search for products using specific attributes, maintain a personal list of brands that they liked on earlier purchase instances and are usually loyal to brands that they have been purchasing. Loyalty towards specific brands are in turn swayed by the word of mouth, advertising and other image factors revolving around the brands. Further, loyalty towards specific brands and contents of personal lists evolve over time as the customer receives more and more information about the brands.

For measuring loyalty of a customer towards a brand Guadagni & Little (1983) follow a method where the loyalty measure $L_{j(t)}^i$ of a customer i for a brand j at a purchase instant t is expressed as an update of her loyalty at the earlier purchase instance of $(t - 1)$ by the signals received during current purchase instance. Since the modeler does not have access to the exact signals being received by customers during a real-life purchase, and comes to know only about whether the customer bought brand j at purchase instant t or not, the loyalty measure can be modelled as,

$$L_{j(t)}^i = \lambda L_{j(t-1)}^i + (1 - \lambda) Y_{j(t)}^i, \quad (35)$$

where λ denotes a parameter representing strength of prior beliefs and

$$Y_{j(t)}^i = \begin{cases} 1, & \text{if customer } i \text{ purchases brand } j \text{ at purchase instance } t, \\ 0, & \text{otherwise,} \end{cases} \quad (36)$$

denotes whether the customer bought brand j at purchase instance t or not.

On the other hand, while carrying out in-depth studies with captive subjects, researchers do have access to the kind of signals a potential customer receives prior to a purchase instance and they can control them as well (as in the case of Roberts & Urban (1988)). In such cases, the customer society can be modelled as a graph, with individual customers as nodes and interactions amongst customers being denoted by arcs. Let us assume that a potential customer comes in contact with n other existing customers (owners) of brand j before making the final purchase decision at each purchase instance. The effect of recommendations of n owners on the loyalty measure $L_{j(t)}^i$ of customer i can be modelled as follows. Let us assume that the

loyalty measure $L_{j(t)}^i \sim N(\mu_{L(t)}, \sigma_{L(t)}^2)$, that is, it is distributed normally with mean $\mu_{L(t)}$ and variance $\sigma_{L(t)}^2$. Let us say that the incoming signals from each of the existing owners of brand j are distributed normally with mean \bar{x}_j and variance $\sigma_{x_j}^2$. Then, the mean and variance of the loyalty measure at purchase instance t can be given as

$$\mu_{L(t)} = [(\lambda/n)\mu_{L(t-1)} + \bar{x}_j]/[(\lambda/n) + 1], \tag{37}$$

$$\sigma_{L(t)}^2 = \left(\frac{\lambda}{\lambda + n}\right)^2 \sigma_{L(t-1)}^2 + \left(\frac{n}{\lambda + n}\right)^2 \sigma_{x_j}^2. \tag{38}$$

On many online markets, customers maintain a personal list of brands that they like (as in the case of Wu & Rangaswamy (2003) using purchase data from PeaPod). The choice of brands from personal lists can be represented as a binary variable $Z_{i(t)}^j$ where

$$Z_{i(t)}^j = \begin{cases} 1, & \text{if customer } i \text{ makes a purchase from personal list at instance } t, \\ 0, & \text{otherwise,} \end{cases} \tag{39}$$

and its effect on fuzziness reduction of customer’s utilities from various brands, along with the effects of customer searches can be modelled as,

$$\Delta_{i(t)} = \exp\left(\sum_{p \in P} \omega_p S_{p(t)} + \omega_i Z_{i(t)} + \omega_L L_{i(t)}\right), \tag{40}$$

where $\Delta_{i(t)}$ is the change in the fuzzy spread of customer utility measures, $S_{p(t)}$ denotes the variable denoting that the customer searched on attribute p , ω_p measures the effects of processing on p on fuzziness reduction, $Z_{i(t)}$ is a binary or dummy variable denoting whether choice is from the personal list or not, ω_i measures effect of personal list on fuzziness reduction, $L_{i(t)}$ is a loyalty measure while ω_L measures the carry-over effect of past experience.

Like discussed in §3.1, to account for the modeler’s inability to specify the customer’s utility function completely, let us use a random utility framework, so that

$$U_{ij}(\tilde{\chi}_j) = h(\tilde{\chi}_j, m_{ij}) + \epsilon_j, \tag{41}$$

where U_{ij} is the utility of brand j for customer i , $\tilde{\chi}_j$ are the attributes of brand j , m_{ij} measures the degree of consideration of customer i for brand j and ϵ_j is the error term. Similar to discussion in §3.1, h is a function calculating an estimate of the customer’s utility. If we assume a Logit choice model of §3.1a, as the customer would choose the alternative that maximizes his utility U_{ij} , choice probabilities of individual brands can be given as

$$p_{ij} = \exp(h(\tilde{\chi}_j, m_{ij})) / \sum_{k \in C_i} \exp(h(\tilde{\chi}_k, m_{ik})). \tag{42}$$

Note that if the membership values m_{ij} ’s were either 0 or 1, we would land up with a situation very similar to that discussed by Roberts & Urban (1988). So, a measure of fuzziness-adjusted utility (similar to the concept of risk-adjusted utility introduced by Roberts & Urban (1988) can now be used. For this, we assume that $h(\cdot, \cdot)$ is exponential in nature and both individual attributes and consideration set membership functions have forms similar to the normal distribution. Also, we assume that customers try to maximize their fuzziness-adjusted utility

during the choice stage. For estimating the parameters of the model, the likelihood expression can be written as,

$$\text{Likelihood} = \prod_t \prod_i p_{ij(t)}^{y_{j(t)}^i}, \quad (43)$$

whose parameter values can be estimated by using the EM algorithm. As discussed by (Dempster 1977; McLachlan & Krishnan 1997; Lange 1999), EM algorithm is a method of estimating parameters in a model when there is absent data and/or parameters from the observed data set. The absence of data/parameters could be either due to missing values or the data/parameter not being observable as such. The EM algorithm can be described as the following series of steps.

- (1) Start with an initial estimate of parameters ($\theta^{(0)}$).
- (2) Fill in the missing data using likelihood calculations according to parameters $\theta^{(0)}$.
- (3) Calculate the likelihood expression and maximize it with respect to parameters θ . Assign the corresponding parameter values to $\theta^{(1)}$.
- (4) Repeat the procedure with $\theta^{(1)}$ as the initial estimate of parameters.

EM algorithm leads to an increase in the complete data likelihood. The complete data, however are artificially constructed through likelihood calculations on the incomplete data. However, Dempster *et al* (1977) show that each iteration of EM algorithm either maintains or improves the incomplete data likelihood as well.

Having discussed fuzzy models for decision-making in this section, some notes are in order. First, the advantage of using a fuzzy sets approach originates from being able to model the uncertainties and fuzziness inherent in the decision-making process. Also, the model described here was developed by extending a two-stage decision making model of §3.4. However, if the fuzzy spread of the consideration sets becomes higher, the customer becomes indifferent between including or excluding a brand from the consideration set. Under such a scenario, the two stage fuzzy choice model becomes equivalent to a single stage choice model.

5. Latent demand modelling

During the buying process, a customer makes decisions based on many parameters that are not available to the researchers. For example, in using a centralized server for carrying out simulation jobs, users may not submit requests during specific time periods due to server's slow response. They would usually adjust their utilization to accommodate server load. However, if this usage data was used to model demand, it would be erroneous considering that the demand of not all users came to the forth. So, if the capacity of the server in this case is increased, the users who had earlier been using the server according to the load, would start using it according to their convenience. Such a hidden demand is referred to as "latent demand" in literature. Econometricians use latent variable models to model the effect of variables that can not be measured directly and hence are latent.

Further, during our discussions of §§ 3.4 and 4, we found that it is never easy for the modeler to figure out the specific consideration (or choice) set of a customer just by looking at the choices she makes. Although, discrete choice modelers have tried to overcome this shortcoming by designing a variety of survey instruments, but their accuracy has never been beyond doubt. In traditional markets and especially in online markets, there are situations

where customers (or potential customers) reveal information about their decision process that can be advantageously used for modelling latent demand. For example, in online auctions the losing bids provide information about the willingness-to-pay of losing bidders. Pre-purchase visits of customers provide a rich source of data for modelling demand in durables market.

In this section, we discuss the modelling of decision-making situations where the key deciding factors and their importance is not available to the modeler. We also discuss the econometric methodologies used for modelling of such variables. As far as the latent consideration sets are concerned, several studies including (Manski 1977; Ben-Akiva & Boccara 1995; Siddarth *et al* 1995; Bronnenberg & Vanhonacker 1996; Haab & McConnell 1996) have tried to model latent and probabilistic consideration (or choice) sets. We discussed some of these studies in §3.4.

In particular, Bronnenberg & Vanhonacker (1996) use scanner purchase data to develop a two-stage choice model, where the customer forms a choice set from her consideration set in the first stage and then chooses an alternative from the choice set in the second stage. The formation of the choice set in the first stage is modelled through incorporation of thresholds on alternative salience. An alternative is included in the choice set only if its salience exceeds the customer-specific threshold. Salience of an alternative i for customer n in period t can be expressed as,

$$S_{nit} = s_{nit} + \xi_{nit}, \quad (44)$$

and the customer-specific threshold Θ_{nit} can be expressed as,

$$\Theta_{nit} = \theta_{nit} + \phi_{nit}, \quad (45)$$

whereas the utility that the customer derives out of the alternative is represented as,

$$U_{nit} = V_{nit} + \epsilon_{nit}, \quad (46)$$

where s_{nit} , θ_{nit} and V_{nit} are the deterministic components of alternative salience, customer-specific thresholds and brand utilities respectively. ξ_{nit} , ϕ_{nit} and ϵ_{nit} represent their random components that are independent and identically distributed (i.i.d) draws from a type I extreme value distribution.

Under the distributional assumptions, the probability that customer n includes alternative i in her choice set in period t would be given as,

$$\pi_{nit} = 1/(1 + e^{(\theta_{nit}-s_{nit})}), \quad (47)$$

and the probability that an alternative i is finally chosen by customer n in period t is given as,

$$P_{nit} = \pi_{nit} \cdot e^{V_{nit}} / \sum_{\forall j} \pi_{njt} \cdot e^{V_{njt}}. \quad (48)$$

Note that alternative salience is different than the utility the customer derives out of an alternative, which is calculated only for those alternatives that show up in the choice set. Bronnenberg & Vanhonacker (1996) express the deterministic component of alternative salience, s_{nit} , in terms of variables associated with promotion display of brand, the shelf space allocated to it, recency of its choice by the customer, price range membership of the alternative, inherent preference associated with the alternative for the customer, alternative's unpromoted price and the price discount on offer due to the discount scheme.

For finding the effect of recency of choice by the customer on her alternative selection into the choice set, Bronnenberg and Vanhonacker use a latent class model of Kamakura & Russell (1989) to partition customers into two segments: loyal customers and in-store sensitive customers. They find that loyal customers, defined as ones with positive effects of recent purchase behaviours on the inclusion of an alternative into their choice set, show an exponential decay in choice set sizes across their population. On the other hand, in-store sensitive customers, who have a negative correlation with their recently purchased alternatives, show a unimodal distribution of their choice set sizes.

Construction of choice models through latent variables and a double-hurdle choice formation process lead to models that better fit data. von Haefen (2003) note that models based on Kuhn-Tucker demand framework (Wales & Woodland 1983) possess the power to address the customer's extensive margin choice (whether to consume the brand) and her intensive margin choice (how much to consume) with a single comparison of the brand's virtual and market prices. But while modelling customer decision-making with consideration sets, this power turns out to be restrictive. To overcome this shortcoming, a general form of utility can be used to break up the Kuhn-Tucker demand framework, so that an individual's extensive margin choice decision and intensive margin choice decision are separated out.

Under their framework, if $f_i(\cdot)$ denotes the consideration of the individual customer for consumption of the i^{th} good, and if the i^{th} good is not consumed, either $f_i(\cdot) \leq 0$ or that $f_i(\cdot) > 0$ and the virtual gain from consumption of i^{th} good is less than its market price. The virtual gain from consumption of good i is calculated by comparing utility gains from its consumption and consumption of an essential Hicksian good (denoted by z). Virtual gain from consumption of good i can hence be written as $(p_z \times (\partial U / \partial x_i)) / (\partial U / \partial z)$. For his model, von Haefen finds that a price or quality change that results in the expansion of consideration set does not in itself increase utility. Only if the increase is concomitant with a change in consumption would the individual experience a welfare gain.

5.1 Latent demand information in online auctions

Online auctions are a source of excellent information about demand, however, most of it lies hidden in the values of losing bids. Note that losing bids are never utilized as such in the auctioning procedure. But, if used properly, data from online auctions can form an important asset for the market researcher. Bids from a single-item auction can be used for estimating the willingness-to-pay functions of the individual prospective customers. Bids of a multi-item auction contain even more information, and can be used to find out the complementarities existing between different items.

Saroop & Bagchi (2000, 2002) and Bagchi & Saroop (2003) use bids from an unfinished single-item single-unit auction to predict the final price of an ongoing eBay auction. Saroop & Bagchi (2000) use a heuristic for estimating the final price of an ongoing auction. The salient point of this heuristic is that it tries to avoid the effect of bids placed by non-serious bidders on the estimate of the final price. Bidders who would end up as being "High Valuers" are first identified according to the times at which they place their bids. Since these bidders would not want the going price in the auction to rise very high, they would desist from bidding. On the other hand, they would also not want too many bidders to be left in the fray in the closing stages. Keeping the two conflicting objectives in mind, the authors contend that the high valuers would try to keep the bidding at an "appropriate level" that balances off these two considerations. As can be expected, the appropriate level curve for an auction would depend on the item that has been put on auction. An estimate of the appropriate level applicable to a particular auction can be found by finding it out for similar auctions that took place in the past.

Saroop & Bagchi (2002) use an array of neural networks for carrying out the prediction of the final price in an ongoing auction. The learning architecture used there tries to accommodate the fact that the proceeds of auctions would behave according to the class of the auction. The architecture dynamically learns characteristics of a number of classes from past auctions, and then utilizes the learned patterns to predict the final price in the ongoing auction.

Bapna *et al* (2002) use the bids from an ongoing auction to carry out a predictive calibration of the minimum bid increment that would maximize the auctioneer's revenue. They build a predictive model that estimates the final price as a function of the bids of "evaluator" bidders. To recognize evaluator bidders in an ongoing auction, they have used a logit regression function of the bid value and the number of times a bidder has placed bids. The predictions of final price are then used to dynamically decide upon the minimum bid increment. The bid increment is adjusted so as to minimize the difference between the valuations of the winning bidders and the price(s) they pay.

Huberman *et al* (2000) show a method of finding the complementarities and supplementarities amongst goods put up on auction using information hidden in unsuccessful bids. The interactions amongst good valuations is modelled using a parameter α in the expression for value of a set S of goods. Value of a set S of goods is expressed as

$$\text{Value}(S) = \left(\sum_{i \in S} v_i \right)^{\alpha_S} \quad (49)$$

where v_i denotes the value of an individual good i . α_S is evaluated using the following heuristic.

- (1) If bids on the bundle S occur more than once, then use a weighted mean of individually calculated α values.
- (2) If S itself and none of its proper subsets has been bid upon, then assume $\alpha_S = 1$, with no complementarities or supplementarities.
- (3) If proper subsets of S have been bid upon, then use a weighted mean of the α values calculated from each of the bids. The weights are assigned in such a manner that larger subsets have a higher weight.

Note that in this model, $\alpha < 1$ would denote that bidders value the bundle S less than they value the individual items separately. $\alpha = 1$ would denote that there is no particular value of bundling, except for the value of the individual items making up the bundle. $\alpha > 1$ would denote that the auctioneer can derive positive value from the bidders by bundling the items together into a bundle S . Such an analysis can help auctioneers to decide upon the optimal bundle compositions that should be put up on auction to maximize revenue from the sale of the items involved.

6. Demand sensing through learning models

In this section, we describe how nonlinear pricing can be used to learn market demand curve and how self-selecting mechanisms can help in customer segmentation. (Raju *et al* 2004) devise various learning models to understand the demand behaviour and simultaneously, to derive an appropriate response mechanism that will minimize system-wide costs. Ravikumar *et al* (2004) develop a learning model to model demand in a competitive service market

environment. To motivate the underlying learning procedures, we consider the case of a retailer of customer goods who experiments with price-quantity pairs under capacity and service-level constraints as described by Raju *et al* (2004c). The capacity constraints come from the shelf-space in the store and from the replenishment leadtimes of the suppliers of the commodities. Note that derivation of willingness-to-pay curves for customers and even demand curves in general, intrinsically assume an infinite supply of goods. However, in reality, suppliers or retailers often face constraints on supply and work within them. Hence, models that accommodate such constraints are often closer to reality.

6.1 *Nonlinear pricing by a retailer with capacity constraints*

In a competitive retail market, there are typically several retailers selling an identical product. Any attempt by any of the retailers to sell its product at more than the market price leads customers to desert the high-priced retailer in favor of its competitors. On the other hand, in a monopolized market there is only one retailer selling a given product and when a monopolist raises its price, it loses some, but not all, of its customers. In reality, retail markets are somewhere in between these two extremes. Every retailer, in spite of being part of a competitive market, still enjoys limited monopoly power, originating, for example, from locational advantage of his store or from its customer service strategy.

If a retail store has some degree of monopoly power, it has more options available to it than a retail store that operates in a perfectly competitive market. Such a retailer can further *differentiate* its products from competitors to enhance its market power. In this paper we consider a specific form of differentiation, namely the *nonlinear pricing*, where price per unit of sale is not constant but depends on how much a customer buys. Volume discounts for large purchases is an example of such a pricing scheme. This form of differentiation appeals for two reasons. One, due to retailers in general having limited flexibility with regard to price changes because of the lean margins left by manufacturers; quantity discounts offer a simple and feasible implementation of differentiation that can improve the retailer's profit. Secondly, volume discounts enable a retailer to construct price-quantity packages that give the customers an incentive to *self select*. Otherwise, the retailer has to know the willingness-to-pay and demand curves of the customers to set the right price for the right customer. And even if the retailer knew something about the statistical distribution of willingness-to-pay, it might be hard for him to prevent the *gaming*: a high-willingness-to-pay customer pretending to be a low-willingness-to-pay customer. The retailer may have no effective way to separate them apart. Nonlinear pricing will help get around this problem by offering two different price-quantity package, one targeted towards the high-demand customer and one toward the low-demand customer.

Now assume that the retailer uses nonlinear pricing-based differentiation mechanism. Because of its intrinsic ability to offer a customer with self-selection, it becomes possible for the retailer to *learn* customer segmentation from their preferences over price-quantity packages. To make the developments simple, let us confine to the case with two types of price-quantity packages offered by the retailer; price for unit quantity and *buy two and get one free*. Such offers can be realistically implemented as follows: Customers who access the web-page of the retailer would initially see unit price offer against the item requested for and also will see an alert depicting "discounts available at higher volumes". Customers who select the former option or package can be safely assumed to be not so price-sensitive and hence would be willing to pay high price for the item if any additional service offer also comes as part of the package. Assume that the retailer offers a lead time commitment to such customers in the event of no stock being available at the time of request. Customers who choose this option

will *learn* over time such additional service benefits from the retailer and adhere to the same retailer for future purchases. We call such customers “captives” to the retailer. Assume further that each captive places an order (or purchases, if stock is readily available) only if he derives strictly positive utility from the price quote and lead-time quote of the retailer, else leaves the market. On the other hand, the customer who wishes to choose the second offer will have low willingness-to-pay per item (provided this price is affordable to him). These customers would be willing to bear with the inconveniences imposed by the retailer. Dynamic pricing policies of airlines reflect the same phenomenon. Consider the *waiting-based* inconvenience by the retailer as detailed below.

Captives get preference over the second class of customers with regard to supply of items in the absence of stock. The retailer can implement this priority in the following way. The customer who clicks on the alert will be next led to a transaction page that exhibits the volume offer in the event of availability of stock after serving the pending “captive” orders. Either when no stock is available or when replenishment quantity to arrive is just sufficient to meet the pending orders, then the customer will be requested to place his orders in his shopping cart. The customer will revisit his shopping cart after a random time interval to check for the availability status and will cancel his order if he again observes stock-out, else will purchase the quantity at the offer made provided the price quote is rewarding enough for him and balks from the system otherwise. The retailer serves these customers in the order according to the time-stamps recorded at the shopping-carts. We call such customers as *shoppers* in the sequel. This type of behaviour models, for example, the customers who purchase objects for a future resale.

The retailer follows the standard (q, r) policy for replenishment with values for q and r so chosen as to ensure all captives are given identical expected lead-time quotes and dynamically chooses his unit price in the above duopoly setting so as to maximize his profits in the presence of uncertainty with regard to arrival pattern of customers, their purchase behaviour and under uncertain replenishment leadtimes. Further, the retailer is assumed to incur unit purchase price and holding cost per unit per unit time for keeping items in his inventory and cost per unit time per back-logged request. Assume zero ordering costs.

Since each customer, captive or otherwise, is served according to the time-stamps recorded, backlogged requests can be modelled using *virtual* queues. Backlogged captive orders form a priority queue at each retailer whereas the shoppers requests in their respective shopping carts form an *orbit* queue with single retrial. It is important to note that only the retailer but not the customers will be able to observe these queues.

Model – In order to illustrate use of reinforcement learning, we need the following mathematical formalism for the above model.

- Customers arrive at the market according to a Poisson process with rate λ . f fraction of these customers are captives at the retailer and the remaining fraction constitute the shoppers
- The retailer posts per unit price p as part of his menu to the arriving customers
- The retailer turns away any arriving request if the total demand backlogged exceeds N at that point in time
- The retailer has finite inventory capacity I_{\max} and follows a fixed reorder policy for replenishing; when the inventory position, current inventory level plus the quantity ordered, falls below a level r , would order for replenishment of size $(I_{\max} - r)$. This is the classical (q, r) policy of inventory planning

- The replenishment leadtimes at the retailer are *i.i.d* exponential with mean $1/\mu$
- The captive measures his utility of a price quote p and lead-time quote w (equal to the expected replenishment leadtime, $1/\mu$) by,

$$U_c(p, w) = [(1 - \beta)(p_c - p) + \beta(w_c - w)]\Theta(p_c - p)\Theta(w_c - w), \quad (50)$$

where $\Theta(x) = 1$ if $x \geq 0$ and is zero otherwise, and $0 \leq \beta \leq 1$. $p_c \sim U(0, p_{\max}]$ and $w_c \sim U(0, w_{\max}]$ with $U(\cdot)$ denoting the uniform distribution over the specified interval for given p_{\max} and w_{\max} .

- The shopper measures his utility of the unit price option p on the menu by,

$$U_s(p) = (p_s - p)\Theta(p_s - p), \quad (51)$$

where $p_s \sim U(0, p_{\max})$ for a given p_{\max} and p_s is the shopper's maximum willingness-to-pay per item.

- Every shopper upon placing an order in his shopping cart at the retailer will revisit the retailer again after an interval of time distributed exponentially with rate μ_s .
- The retailer sets his unit price p from a finite set A .
- The retailer incurs a holding cost rate of H_I per unit item per unit time and a cost of H_q per each back-logged request. The purchasing price per item is P_c . For simplicity, assume zero reorder costs.

The retailer dynamically resets the prices at random intervals so as to maximize his expected profits. Under the above Markovian distributional assumptions, the dynamic pricing problem can be reduced to a semi-Markov decision process. However, in reality, retailers do not in general have knowledge about the distributions underlying the model and also about customers' behaviour. In such cases, these retailers learn about their most beneficial prices over time using reinforcement learning (RL), the right paradigm for learning in Markov decision processes. In the following sections, we describe the appropriate RL-based learning schemes.

The backlogged requests of the captives and the shoppers are assumed to form separate *virtual* queues at the retailer and are referred to as queue 1 and queue 2 respectively in the following. Queue 2 is an orbit queue, where each shopper who has registered his orders in the shopping cart would make a retrial after an interval of time that is assumed to be exponentially distributed with mean $1/\mu_s$. Since the two retailers act independently under our no-information assumption, we isolate, for notational convenience, one retailer and develop a model for price dynamics at that retailer.

Let $\mathbf{X}(t) := (X_1(t), X_2(t), I(t))$ be the state of the system at the retailer with $X_i(\cdot)$'s representing the number of backlogged requests in queue i and $I(\cdot)$ representing the inventory level at the retailer at time t . The retailer posts unit price quote and the volume discount alert on his web-page and resets the prices only at transition epochs, that is, whenever a purchase happens (and hence the inventory drops) or when a request is backlogged in either of the queues. Recall that the captive (shopper) will purchase or make a backlog request only when U_b in (50) (U_s in (51)) is positive. It is easy to see that price dynamics can be modelled as a continuous time Markov decision process model. Below we give the state dynamics.

At time 0, the process $\mathbf{X}(t)$ is observed and classified into one of the states in the possible set of states (denoted by S). After identification of the state, the retailer chooses a pricing action from A . If the process is in state i and the retailer chooses price $p \in A$, then

- (i) the process transitions into state $j \in S$ with probability $P_{ij}(p)$,

(ii) and further, conditional on the event that the next state is j , the time until next transition is a random variable with probability distribution $F_{ij}(\cdot|p)$.

After the transition occurs, pricing action is again chosen by the retailer and (i) and (ii) are repeated. Further, in state i , for the action chosen p , the resulting reward, $S_p(\cdot)$, the inventory cost, $H(i)$ and the backorder cost $C(i, j)$ costs are as follows: Let $i = [x_1, x_2, i_1]$ and $j = [x'_1, x'_2, i'_1]$.

$$S_p(i, p, j) = \begin{cases} p, & \text{if } x'_1 = x_1 + 1, \\ p, & \text{if } i'_1 = i_1 - 1, \\ 2p, & \text{if } i'_1 = i_1 - 3, \\ 0, & \text{otherwise,} \end{cases}$$

$$C(i, j) = [i_1 - i'_1]^+ P_c,$$

$$H(i) = x_1 H_q + i_1 H_l.$$

The following remark is in order.

Remark 5. The complementarity condition given below holds.

$$I(t)X_1(t) = 0 \forall t \text{ for } l = 1, 2.$$

Let p below represent the retailer's price in the observed states and p' be the price posted by the second retailer. Then the following transitions occur:

- $[0, x_2, i_1] \rightarrow [0, x_2, i_1 - 1]$ with rate $f_1 \lambda P(U_c(p, \frac{1}{\mu}) > 0) \forall x_2, i_1$
- $[0, x_2, i_1] \rightarrow [0, x_2 - 1, i_1 - 3]$ with rate $(1 - f_1 - f_2) \mu_s P(U_b(p) > 0) \forall x_2, i_1$
- $[0, x_2, i_1] \rightarrow [0, x_2, i_1 - 3]$ with rate $(1 - f_1 - f_2) \lambda P(U_b(p) > U_b(p') \geq 0) \forall i_1$.
- $[x_1, x_2, 0] \rightarrow [x_1 + 1, x_2, 0]$ with rate $f_1 \lambda P(U_c(p, \frac{1}{\mu}) > 0) \forall x_2$
- $[x_1, x_2, 0] \rightarrow [x_1, x_2 + 1, 0]$ with rate $(1 - f_1 - f_2) \lambda P(U_b(p) > U_b(p') \geq 0) \forall x_1$
- $[x_1, x_2, 0] \rightarrow [(x_1 - r)^+, x_2, (r - x_1)^+]$ with rate $\mu \forall x_2$

Discounted optimality – Let $\pi : S \rightarrow A$ denote a stationary deterministic pricing policy, followed by the retailer, that selects an action only based on the state information. Let $t_0 = 0$ and let $\{t_n\}_{n \geq 1}$ be the sequence of successive transition epochs under policy π and $X(t_n -)$ denote the state of the system just before t_n . For any $0 < \alpha < 1$, let the long-term discounted expected profit for the policy π be

$$V_\pi(i) = E_\pi \left[\sum_{n=1}^{\infty} e^{-\alpha t_{n-1}} (S_p(X(t_n -), \pi(X(t_n -), X(t_n))) - C(X(t_n -), X(t_n)) - \int_{t_{n-1}}^{t_n} H(X(t_n -)) e^{-\alpha t} dt | X_1 = i) \right]. \quad (52)$$

α above discounts the rewards to the initial time epoch.

Let $V_\alpha^*(i) = \max_{\pi} V_\pi(i)$. The retailer's problem is to find a $\pi^* : S \rightarrow A$ such that $V_{\pi^*}(i) = V_\alpha^*(i)$. The domain of optimization above can be in general the set of all non-anticipative policies (randomized, non-randomized or history dependent). But from the Markovian assumptions it follows that there always exists a stationary deterministic optimal policy and hence,

we have restricted the domain to stationary policies only (and they are finite in number for our finite state model).

From the Bellman's optimality condition, it follows that

$$V_\alpha^*(i) = \max_p \left\{ \bar{R}_\alpha(i, p) + \sum_{j \in S} P_{ij}(p) \int_0^\infty e^{-\alpha t} V_\alpha^*(j) dF_{ij}(t|p) \right\}, \quad (53)$$

where

$$\bar{R}_\alpha(i, p) = \sum_{j \in S} P_{ij}(p) \left[R(i, p, j) + \int_0^\infty \int_0^t e^{-\alpha s} H(i) ds dF_{ij}(t|p) \right]. \quad (54)$$

Equation (53) above can be solved algorithmically using any fixed point iteration scheme, such as the *value iteration* for the typical Markov decision processes. Such schemes are assured of convergence because of the contraction property coming from the discount factor, α . Also, one can construct the optimal policy π^* by assigning $\pi^*(i)$ to equal the maximizer on the RHS of (53). However, as we assumed earlier, the retailers do not have any information about underlying distributions involved at various stages and hence the conditional averaging that appears in (53) cannot be performed to derive the optimal value, and hence the optimal pricing policy. This motivates the retailers to use online learning to converge to optimal policy π^* eventually. One can think of devising a *learning* scheme based on any fixed point iteration methods on V^* . Even if we assume that this can be done, we would still need to know about P_{ij} 's above to construct π^* . To obviate this difficulty, the *Q-learning* has been proposed by Watkins & Dayan (1992) for Markov decision processes. Below we proceed to modify it to suit to our continuous time case. To motivate the algorithm, consider the following *Q-value* associated with an action p in state i :

$$Q(i, p) = \bar{R}_\alpha(i, p) + \sum_{j \in S} P_{ij}(p) \int_0^\infty e^{-\alpha t} V_\alpha^*(j) dF_{ij}(t|a). \quad (55)$$

In other words, $Q(i, p)$ represents the long-term expected profit starting from state i when the first action to be followed in state i is p , which is possibly different from the optimal action. It is easy to see that $V_\alpha^* = \max_p Q(i, p)$, and the maximizer is the optimal action to perform in state i . Thus, if one gets somehow the Q -values, it easy to construct an optimal policy. In online learning, these Q -values are obtained from learning through actual execution of various actions in state i and measuring their relative merits. For details, please refer to (Watkins & Dayan 1992).

Below we give the actual learning update rule involved in Q-learning for our continuous time MDP.

Let $t_0 = 0$ and start with an initial arbitrary guess, $Q_0(\cdot)$ of $Q(i, p)$ above for all i and p .

- *Step 1:* At any n -th transition epoch at time t_n , observe the state i and select the price action $p_0 \in \operatorname{argmax}_p Q(i, p)$ with probability $1 - \epsilon$ and any other price in A with probability ϵ for some $\epsilon > 0$.
- *Step 2:* If $X(t_n) = i$ and the price action chosen is p , then update its Q -value as follows:

$$Q_{n+1}(i, p) = Q_n(i, p) + \gamma_n [S_p(i, p, \cdot) - H(i)((1 - e^{-\alpha T_{ij}})/\alpha) + e^{-\alpha T_{ij}} \max_b Q_n(j, b) - Q_n(i, p)], \quad (56)$$

where j above is the state resulting from the action p in i and $S_p(\cdot)$ is the reward collected from such action. T_{ij} is the average of sampled transition times between states i and j . γ_n above is called *learning parameter* and should be such that $\sum_n \gamma_n = \infty$ and $\sum_n \gamma_n^2 < \infty$.

Repeat steps (1)–(2) infinitely. Convergence is slow as is typical of any RL-algorithm. The speed of convergence can be drastically improved using function approximations to Q -values based on some observed features (Bertsekas 1997).

Remark 6. It is important to note that in the above model, customers' choice behaviour, capacity constraints, and the replenishment policies decide transition dynamics. Observations over these transitions, but not the knowledge of functional form of customers' utilities, are required for the implementation of RL. Hence, reinforcement learning offers a valuable experimental framework for learning the demand functions.

7. Future research

We have identified the following areas of further research.

- Existing approaches to figuring out demand function include econometric estimation models and procedures. These are mainly developed for macro-economic situations. Micro-economic studies suitable for adaptation to a single firm need, on the other hand, an extensive sets of data, that include competitor's prices, sales volume, and advertising expenditure in order to achieve a reasonable statistical accuracy. The availability of data may be a non-issue for a modern enterprise operating on computerized data collection systems. However, it is the requirement on data from competitors that is making the existing approaches difficult to apply in practice. More often than not, competitive information is available through third-party market research firms. This data are usually available at aggregate levels, aggregated by time, such as monthly and quarterly reports, aggregated by product models, aggregated by geography etc, making them inappropriate for the level of data granularity that the existing models require. These data-related issues open up a few interesting research directions.
 - Development of robust models that are capable of *learning* demand behaviour from aggregated data sets.
 - Development of Maximum Likelihood Estimators of choice probabilities that conform to different forms of data aggregation
- As cited by Roberts & Urban (1988), Bell & Raiffa (1979) show that if the customer obeys the axioms of von Neumann–Morgenstern utilities for lotteries and if a utility function exists, then the customer must display constant risk aversion, which means that her utility function should either be linear or negative exponential in nature. Similar results for situations when customers are averse to fuzziness in decision-making need to be studied in detail.
- There have been studies like (Mitra 1995) that study the consistency of consideration sets over purchase instances. Longitudinal studies are carried out over human subjects. However, Online markets provide a rich set of data for customers making decisions over a length of time. An in-depth analysis needs to be carried out to model the process of consideration set modification over time.

- The signals that can be captured in an online market are much more numerous than in traditional markets. However, we are still concentrating on the use of signals that are captured by the state of the technology as of today. While there exist concerns with respect to use of Internet technology to sense demand in view of customer privacy, adapting the technology according to the requirements of demand sensing methods is an issue that needs to be resolved at both the levels of technology research and privacy laws to enhance the accuracy of demand sensing methods.
- Many marketing managers in their attempt to predict demand behaviour, collect data from customers through formal or informal surveys. Out of privacy concerns, the customers' stated preference data might not reflect their true purchase behaviours. The stated preferences need to be evaluated incorporating different degrees of fuzziness over the consideration sets. In order to address security concerns of customers, we envision a future where technologies emerge to support *privacy preserving transformations* of data. While mining such transformed data remains an interesting data mining challenge, sensing demand from such data will be a statistical challenge.

References

- Andrews R L, Srinivasan T C 1991 Studying consideration effects in empirical choice models using scanner panel data. *J. Marketing Res.* 32: 30–41
- Bagchi A, Saroop A 2003 *Internet auctions: Some issues and problems*, (Springer-Verlag) Lecture Notes in Computer Science (eds) G Goos, J Hartmanis, J van Leeuwen, vol. LNCS-2918, pp 112–119
- Bapna R, Goes P, Gupta A, Karuga G 2002 Predictive calibration of online multi-unit ascending auctions. *Proc. WITS-2002, Twelfth Annual Workshop on Information Technology and Systems*, Barcelona, Spain
- Bell D E, Raiffa H 1979 Managerial value and intrinsic risk aversion. Working Paper, Technical Report 79–65, Harvard Business School, Boston, MA
- Ben-Akiva M, Boccara B 1995 Discrete choice models and latent choice sets. *Int. J. Res. Marketing* 12: 9–24
- Bertsekas D P 1995 *Nonlinear programming* (Belmont, MA: Athena Scientific)
- Boyd J, Mellman J 1980 The effect of fuel economy standards on the US automotive market: A hedonic demand analysis. *Transportation Res.* A14: 367–378
- Bronnenberg, Bart J, Vanhonacker W 1996 Limited choice sets, local price response and implied measures of price competition. *J. Marketing Res.* 33: 163–173
- Cardell S, Dunbar F 1980 Measuring the societal impacts of automobile downsizing. *Transportation Res.* A14: 423–434
- Chiang J, Chib S, Narasimhan C 1999 Markov chain Monte Carlo and models of consideration set and parameter heterogeneity. *J. Econometrics* 89: 223–248
- Chintagunta P K, Jain D C, Vilcassim N J 1991 Investigating heterogeneity in brand preferences in logit models for panel data. *J. Marketing Res.* 28: 417–429
- Chipman J 1960 The foundations of utility. *Econometrica* 28: 193–224
- Dempster A, Laird N, Rubin D 1977 Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc.* 39: 1–38
- Fortheringham A S 1988 Consumer store choice and choice set definition. *Marketing Sci.* 7: 299–310
- Guadagni P M, Little J D C 1983 A logit model of brand choice calibrated on scanner data. *Marketing Sci.* 2: 203–238
- Haab T, McConnell K 1996 Count data models and the problem of zeros in recreation demand analysis. *Am. J. Agri. Econ.* 78: 89–102
- Herman B A, Hogg T, Swami A 2000 Using unsuccessful auction bids to identify latent demand. Downloaded from <http://www.hpl.hp.com/research/idl/papers/auctions/auctions.pdf>

- Hulland J S 1992 An empirical investigation of consideration set formation. *Adv. Consumer Res.* 19: 253–254
- Kamakura W A, Russell G 1989 A probabilistic choice model for market segmentation and elasticity structure. *J. Marketing Res.* 26: 379–390
- Lange K 1999 *Numerical analysis for statisticians* (New York: Springer)
- Lee H, Padmanabhan P, Whang S 1997 The paralyzing curse of the bullwhip effect in a supply chain. *Sloan Manage. Rev.* : 93–102
- Manski C 1977 The structure of random utility models. *Theor. Decision* 8: 229–254
- Marschak J 1960 *Stanford symposium on mathematical methods in the social sciences* (ed.) K Arrow, (Stanford, CA: University Press) pp 312–329
- McFadden, Daniel 1978 Estimation techniques for the elasticity of substitution and other production parameters. *Production economics: A dual approach to theory and applications* (eds) M Fuss, D McFadden, (Amsterdam: Elsevier Science, North-Holland) vol. 2 pp. 73–124
- McFadden D, Train K 1996 Consumers' evaluation of new products: Learning from self and others. *J. Polit. Econ.* 104: 683–703
- McFadden D, Train K 2000 Mixed MNL models for discrete response. *J. Appl. Econometrics* 15: 447–470
- McLachlan G, Krishnan T 1997 *The EM algorithm and extensions* (New York: Wiley)
- Mitra, Anusree 1995 Advertising and the stability of consideration sets over multiple purchase occasions. *Int. J. Res. Marketing* 12: 81–94
- Oliver K, Moeller L, Lakenan B 2004 Smart customization: Profitable growth through tailored business streams. Downloaded from http://www.strategy-business.com/media/file/sb34_04104.pdf
- Pindyck R S, Rubinfeld D 2000 *Microeconomics* (Englebert cliffs, NJ: Prentice Hall)
- Raju C V L, Narahari Y, Ravikumar K 2004a Learning dynamic prices in electronic retail markets with customer segmentation. *Ann. Oper. Res.* 41 (to appear)
- Raju C V L, Narahari Y, Ravikumar K 2004b Learning nonlinear dynamic prices in electronic markets with price sensitive customers, stochastic demands, and inventory replenishments. *Electron. Commerce Res.* (communicated) (<http://www.orie.cornell.edu/~paatrus/psfiles/applied-pricing.pdf>)
- Raju C V L, Narahari Y, Ravikumar K 2004c Learning nonlinear dynamic prices in multi-seller electronic markets with price sensitive customers, stochastic demands, and inventory replenishments. *IEEE Trans. Syst. Man Cybernetics* (communicated)
- Ravikumar K, Sundar D K, Batra G, Saluja R 2004 Learning in dynamic pricing games of electronic service markets. *Mach. Learning* (communicated) (<http://doi.ieeecomputersociety.org/10.1109/COEC.2003.1210269>)
- Roberts J H, Lattin J 1991 Development and testing of a model of consideration set composition. *J. Marketing Res.* 28: 429–440
- Roberts J H, Urban G 1988 Modelling multiattribute utility, risk and belief dynamics for new consumer durable brand choice. *Manage. Sci.* 34: 167–185
- Rusmevichientong P, Salisbury J, Truss L T, Roy B, Glynn P 2004 Opportunities and challenges in using online preference data for vehicle pricing: A case study at general motors. *J. Revenue Pricing Manage.* (communicated)
- Saroop A, Bagchi A 2000 Decision support system for timed bids in internet auctions. *Proc. WITS-2000, Tenth Annu. Workshop on Information Technology and Systems*, Brisbane, Australia, pp 229–234
- Saroop A, Bagchi A 2002 Artificial neural networks for predicting final prices in eBay auctions. *Proc. WITS-2002, Twelfth Annu. Workshop on Information Technology and Systems*, Barcelona, Spain, pp 19–24
- Shocker A, Ben-Akiva M, Boccara B, Nedungadi P 1991 Consideration set influences on consumer decision-making and choice: Issues, models and suggestions. *Marketing Lett.* 2: 181–197
- Siddarth S, Bucklin R, Morrison D 1995 Making the cut: Modelling and analysing choice set restriction is scanner panel data. *J. Marketing Res.* 32: 255–266
- Thurstone L 1927 A law of comparative judgement, *Psychol. Rev.* 34: 273–286

- Train K E 2003 *Discrete choice methods with simulation*. (Cambridge: University Press)
- von Haefen R H 2003 Latent consideration sets and continuous demand system models. Downloaded from <http://cals.arizona.edu/rogervh/consider073103.pdf>
- Wales T, Woodland A 1983 Estimation of consumer demand systems with binding non-negativity constraints. *J. Econometrics* 21: 263–285
- Watkins C, Dayan P 1992 *Q-learning*, *machine learning* 8: 279–292
- Wu J, Rangaswamy A 2003 A fuzzy set model of search and consideration with an application to an online market. *Marketing Sci.* 22: 411–434