Demand Estimation Using Binary Regression Quantiles[1]

George Deltas
University of Illinois
1206 South Sixth Street
Champaign IL 61820
Phone: (217) 333-4586
e-mail: deltas@uiuc.edu
http://www.staff.uiuc.edu/ ˜deltas

Gregory Kordas
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA
Phone: (215) 898-6777
e-mail: kordas@econ.upenn.edu

April 2004

Abstract

This paper develops a new semi-parametric procedure to estimate demand for a product using binary choice data. This procedure nests as special cases homoskedastic and heteroskedastic parametric models and also linear index semi-parametric models. Its main advantage over these models is that it allows not only an arbitrary shape for the demand curve, but also allows consumer and product characteristics to cause demand shifts that vary in size (non-parametrically) along the demand curve. The technique is computationally intensive but does not require large datasets. We demonstrate this new estimation approach using survey data on the willingness to pay for MP3 players, and show that the added flexibility of the approach yields differences from standard parametric methods that are both quantitatively important, but also economically interpretable and relevant from a decision-maker's point of view.

JEL Classification Codes: C14, C25, D12.

Keywords: Discrete Choice, Maximum Score, Quantile Regression.

Very preliminary version. All comments welcome.

---

[1]Very Preliminary and Incomplete. We would like to thank Roger Koekner and seminar participants at the University of Illinois and at pERE for comments and suggestions. Financial Assistance by the University of Illinois College of Commerce is gratefully acknowledged. We alone are responsible for all errors.

1. Introduction

Discrete choice models are a staple of modern demand estimation techniques. Discrete choice models can be broadly classified as binary or multinomial, and as parametric or semi-parametric. Binary models examine a consumer's choice to purchase or not to purchase a particular product, while multinomial models examine a consumer's choice among a larger set of possible products. Parametric models specify a parametric distribution (up to a number of finite parameters) for the consumer's willingness to pay for each product. Semi-parametric models do not impose a particular distribution, but allow the form of the distribution to be estimated by the data.

This paper introduces a new semi-parametric technique to estimate binary choice models and derive the associated demand function. This technique offers certain advantages over existing parametric and semi-parametric techniques that are used to estimate binary choice models. We describe these advantages next. The simplest parametric techniques, e.g., logit regression, impose a particular functional form on the demand function. In particular, the shape of the demand curve is that of the inverse of the cumulative distribution function of the consumer willingness to pay for the product. In the case of standard parametric families, such as the logit and probit models, this results in demand curves that asymptote to infinity, possibly with relatively fat upper tails, have a relatively flat portion, and then plummet for low enough prices. Moreover, any shift in the demand curve will shift the entire demand up and down by the same amount (or by the same proportion, if the log of price, instead of price, is used as a dependent variable). That is, the standard parametric models do not allow for demand shifts that differentially affect the high-willingness-to-pay and low-willingness-to-pay consumers. Heteroskedastic and random coefficients variants of these standard models allow more flexibility in the shape of the demand curve and on the nature of demand shifts, but this flexibility is still rather limited.

Semi-parametric demand estimation eliminates the need to parametrically specify the shape of the demand function. The shape of the demand function is now obtained from the observed decisions of consumers to purchase the product at various prices. Typically, however, demand *shifts* still affect all portions of the demand curve equally (or proportionately, if log price is used as a regressor). Fully non-parametric estimation of the demand curve, which would allow for arbitrary demand shape and arbitrary demand shifts has excessive data requirements and does not appear to be of practical use. The technique we introduce in this paper combines key aspects of the flexibility of fully non-parametric estimation without requiring large datasets. In particular, we impose only a weak assumption on the distribution of willingness to pay: that the quantiles of this distribution are linear to any variables that affect the willingness to pay for the product. This allows flexibility both in the shape of the demand function, and also in the nature of demand shifts. In particular, there is full flexibility in the shape of the demand function. Thus, our approach "nests" semi-parametric demand estimation. Further, there is substantially flexibility in the nature of demand shifts. For example, it is possible that a variable that shifts up the demand for high-willingness-to-pay individuals, can shift it down for average-willingness-to-pay inviduals, and shift it up again for low-willingness-to-pay indviduals. Many possible geometries of demand shifts are possible (too many to reasonably enumerate).

In this paper we outline the binary regression quantile approach of estimating demand curves, and implement it using a contingent valuation survey dataset of willingness to pay for MP3 players. We condition the demand on MP3 players on two consumer characteristics: on whether the consumer already owns a portable CD player and on how many hours a day he listens to music. We compare the estimates of the (conditional on characteristics) demand function obtained via our technique with that obtained using logit regression. To help describe the differences in the estimates, it would be useful to define the concept of propensity to purchase a product. Consumers of given characteristics have a high propensity to purchase a product if their willingness to pay for the product is high, relative to other consumers with the same characteristics (and conversely,consumers with a low propensity to purchase a product are those with low willingness to pay for the product, relative to other consumers with the same characteristics. Note that in a logit model, a consumer characteristic affects willingness to pay for a product equally for all consumers regardless of their propensity to purchase that product (while in quantile-based estimates of the demand function, this is not necessarily the case).

We find that the differences between the two sets of estimates are economically meaningful. First, the shape of the demand curve differs substantially from that imposed by the logit specification. For most types of consumers, the logit speficication over-estimates willingness to pay for consumers with moderately high propensity to purchase MP3 players, and underestimates it for consumers with low propensity to purchase. The logit model is (on average) reasonably accurate in estimating willingness to pay for consumers with high or moderate propensity to purchase MP3 players. Notice that the pattern of diviation between the two estimates of the demand curve is such that a heteroskedastic logit model would still fail to match the shape of the demand curve estimated using our semi-parametric method.

Second, we find that demand shifts are equal across the demand curve, but differ in ways that are easily interpretable. For example, the demand of consumers who own a portable CD player differs from the demand of people who do not, primarily for consumers with relatively low propensity to purchase an MP3 player. Amongst consumers with a high propensity to purchase an MP3 player, ownership of a portable CD player is not a signal of increased willingness to pay. To put it an other way, the highest willingness to pay consumers amongst those who own and those who do not own a portable CD player are very similar in their preferences for MP3 players. Number of hours spent listening music affects the willingness to pay onfly for those with above average propensity to purchase an MP3 player. The effect for those of below average propensity to purchase an MP3 player is effectively zero. This finding has a straightforward economic interpretation. There are some consumers who because of lifestyle choices value having an MP3 player. These consumers are more likely to have high willingness to pay, holding the amount of time spent listening to music constant, than other consumers who perhaps only listen to music in their rooms. For this first set of consumers, the higher the amount of time spent listening to music, the greater the willingness to pay for an MP3 player (because the more valueable it is to them). For the second set of consumers, willingness to pay for an MP3 player is not likely to be affected much by how long they listen to music: if you listen to music only in your apartment, whether you spend one hour or three hours doing so, it is not likely to affect the value of an MP3 player to you. Observe that the relative effect of these two characteristics on willingness to pay as a function of propensity to purchase is opposite.

Third, economic decisions based on the logit-derived demand curves would differ substantially from economic decisions based on the quantile-derived demand curves. For example, the profit maximizing monopoly price varies substantially across the two estimates of the market demand. For the quantile-based estimates of the demand curve, price responds strongly to increases in marginal cost for low values of marginal cost, but responds less at higher levels of marginal costs. The response of the optimal price to marginal cost is much smoother for the logit-based estimates of the demand curve. Not only are optimal prices different across the two specifications, but the implications from maximal profit levels economically important, with the effect on maximal profits (of using the correct estimate of the demand curve) often of the order of 20 percent or higher. Similarly, consumer surplus for different price levels varies between the quantile and logit based demand curves.

Literature review to be completed.

The rest of this paper is structured as follows. Section 2 outlines the modeling and estimation framework, section 3 describes the empirical application, while section 4 concludes.

2. Modeling and Econometric Framework

2.1. Preliminaries

Consider an individual, i, whose willingess to pay for a product is described by the random utility model

$$V_i = \beta_i X_i + \gamma_i Z + \epsilon_i \tag{1}$$

where X_i is a vector of individual characteristics, Z is a vector of product characteristics, β_i and γ_i are parameter vectors, and ϵ_i is a distrubance term which accounts for unobserved product and consumer characteristics. The individual is assumed to make a 'now-or-never' purchase decision. He will purchase the product if the current price faced by this consumer, p_i, is less than his willingness to pay for it, V_i. Unlike models in which a consumer chooses one of a large number of alternatives, this paper does not consider his decision of purchasing any other (imperfect) substitute products. This limitation is imposed by the state of the statistical technology. However, one can *indirectly* (and imperfectly) account for the existence of other options by including the number of other competing products and measures of their prices and characteristics in the vector of explanatory variables, Z. Note that the parameter vector γ can only be identified if the product characteristics vary over time. It follows that in cross-section data $\gamma_i Z$ will be subsumed in the mean of the disturbance term, ϵ_i.

In its simplest form the model in (1) is estimated under the assumption that β_i and γ_i take the same values for every individual and that the distribution of ϵ is also the same for every individual. Under these simplifying assumptions a change in individual or product characteristics affects the mean the consumer willingess to pay but leaves all other central moments unaffected. Equivalently, a change in X_i or Z shifts all quantiles of the distribution of willingess to pay equally. This is an outcome of the simplifying assumptions which imply homogeneity of consumer response to product and individual characteristics. Such homogeneity of consumer response is problematic since the i.i.d. error term ϵ_i does "all the work" in explaining differences in consumer purchase decisions (conditional on characteristics). As a consequence, estimates of consumer demand and welfare gains from the introduction of a new product often appear to be implausibly large. Such a problem could be mitigated by introducing consumer response heterogeneity.

Traditionally, heterogeneity of consumer response to individual and product characteristics has been accomodated by postulating that β_i and γ_i have a non-degenerate distribution. This specification implies a change in individual or product characteristics affects not only the mean but also higher moments of the willingness to pay. Once the distribution of β_i and

γ_i is estimated the distribution (and quantiles) of V_i can be obtained (typically) via simulation. This paper adopts a different approach. The parameter vectors β_i and γ_i are assumed to take the same values for all individuals. In contrast, the distribution of the distrubance term ϵ is assumed to differ across individuals. In particular, it is assumed that

$$\epsilon_i = \alpha + \sigma(X_i, Z)u_i$$

where α is the median of ϵ_i, $\sigma(\cdot)$ is a function of individual and product characteristics, and u_i is a disturbance term with median equal to zero and distribution function $F(\cdot)$. Denote the τ quantile u by $Q_\tau(u)$, i.e., $Q_\tau(y) = F^{-1}(\tau)$. Then, using similar notation to denote the quantiles of other random variables, the τ quantile of the willingess to pay for the product is given by

$$Q_\tau(V_i) = \alpha + \beta X_i + \gamma Z + \sigma(X_i, Z)Q_\tau(u)$$

The impact of a change in the j^{th} individual characteristic, $x_{j,i}$, on the τ quantile of willingness to pay is given by

$$\frac{\partial Q_\tau(V_i)}{\partial x_{j,i}} = \beta + \frac{\partial \sigma(X_i, Z)}{\partial x_{j,i}}\, Q_\tau(u)$$

This formulation implies a differential response, across different quantiles, of the willingness to pay to changes in consumer characteristics. In principle, such a model can be estimated using standard methods by imposing parametric assumptions on $F(\cdot)$ and $\sigma(\cdot)$. The parameter estimates can then be used to obtain the quantiles (and distribution) of willingness to pay for a consumer with a given set of characteristics.

In this paper, the quantiles of the willingness to pay for the product are estimated directly without any parametric assumptions on $F(\cdot)$ and $\sigma(\cdot)$ and without estimating the parameter vectors β and γ. Rather, the quantiles of the willingness to pay are postulated to depend linearly on product and individual characteristics (and/or functions of such characteristics). In particular, we estimate the model

$$Q_\tau(V_i) = \theta_\tau W_i$$

where θ_τ is a $(1 \times K)$ vector of parameters, W is a $(K \times 1)$ vector of consumer and product characteristics. The parameter vector θ_τ is estimated for M different quantiles using observations from N different individuals. These parameter estimates are then used to estimate the quantiles of V_i. From these, we finally obtain estimates of the consumer demand and the consumer suprlus associated with this particular product.

2.2. Binary Regression Quantiles/Maximum Score

Let the indicator variable y_i take the value of 1 if consumer i purchases the product and the value 0 if he does not, that is, $y_i = 1 \Leftrightarrow V_i - p \geq 0$ and $y_i = 0$ otherwise. Following Manski (1975, 1985) and Kordas (2000) we can estimate the parameter vector θ_τ by maximizing the generalized score function

$$S_\tau(\theta) = \frac{1}{N}\sum_{i=1}^{N}[sgn(\theta_\tau W_i - p_i) - (1 - 2\tau)]sgn(\theta_\tau W_i - p_i) \tag{2}$$

with respect to θ. Note that the above expression indicates that implied parameter on the product price is equal to 1. This represents no restriction on the model as the binary nature of the data requires a normalization. In parametric models such a normalization is typically obtained by fixing the scale of the disturbance term. In Maximum Score estimation no parametric assumptions are made with regards to the disturbance term. Therefore, to identify the parameter vector one must normalize the value of one of the parameters (i.e., fix the value of one the parameters to equal 1) or a function of the values of the entire parameter vector (i.e., fix the norm of the parameter vector to equal 1). In this paper, the implicit normalization is to fix the parameter of p_i to -1. This, in turn, implies that the units of the parameter vector θ are such that the latent variable, V_i, represents willingness to pay, i.e., money metric utility.[2]

[2] This, of course, is not the case for standard parametric models of discrete choice. For example, in

Equation (2) can be equivalently written as

$$S_\tau(\theta) = \frac{1}{N}\sum_{i=1}^{N}[\, y_i - (1-\tau)]\ \mathbf{1}_{(\theta W_i - p_i \geq 0)} \tag{3}$$

where $\mathbf{1}_{(condition)}$ is an indicator variable that takes the value of 1 if the *condition* is true and the value of 0 if the *condition* is false. The score is a step function and its maximization requires use of gradient-free methods, such as simulated annealing. Standard errors can be obtained via bootstrapping.[3] Alternatively, one can estimate the parameter vector θ_τ by smoothing the score function

$$\hat{S}_\tau(\theta, \sigma_N) = \frac{1}{N}\sum_{i=1}^{N}[\, y_i - (1-\tau)]\ \ K\left(\frac{\theta W_i - p_i}{\sigma_N}\right) \tag{4}$$

where $\lim_{N\to\infty} \sigma_N = 0$ and $K(\cdot)$ is a continuous function satisfying the standard kernel assumptions:

- $|K(x) < B|$ for all x in the support of $K(\cdot)$ and for some finite B.
- $\lim_{x\to-\infty} K(x) = 0$.
- $\lim_{x\to\infty} K(x) = 1$.

Smoothing the score function allows the estimation of θ using standard maximization techniques, such as Newton-Raphson and BHHH. In practice, however, simulated annealing is recommended because as the score function is not guaranteed to be concave. Use of the smoothed maximum score also allows us to analytically compute the covariance matrix of

the probit model one needs to divide by the coefficient of price to obtain impact of the other regressors on consumer willingness to pay for the product.

[3] We note that, since the criterion function is a step-function, bootstrapping is not guaranteed to give correct standard errors, even though the approach retains intuitive appeal. Moreover, for large enough datasets, in which the flat regions are of relatively small area each, bootstrap standard errors are going to be approximately correct.

the estimates, $\hat{\theta}_\tau$, of the parameter vector. The asymptotic distribution of $\hat{\theta}_\tau$ is derived in Kordas (2000).[4]

The estimates, $\hat{\theta}_\tau$, of the parameter vectors can be used to estimate the consumer demand for a product, conditional on a set of consumer and product characteristics, and the consumer surplus that the consumers derive from the consumption of that product. We turn to this next.

2.3. Consumer Demand and Welfare Estimation

A consumer will purchase the product if $V_i - p_i \geq 0$. The unobserved product and consumer heterogeneity subsumed in the disturbance term ϵ of equation (1) implies that the consumer's decision is non-deterministic, even conditional on the vector of observed product and consumer characteristics. In particular, the probability that a consumer with characteristics X_i will purchase a product with characteristics Z at a price p_i is given by

$$\begin{aligned} Pr[V_i - p_i \geq 0] &= Pr[\beta X_i + \gamma Z - p_i + \epsilon_i \geq 0] \\ &= Pr[\epsilon_i \geq -\beta X_i + \gamma_i Z + p_i] \\ &= \int_{-\beta X_i - \gamma Z + p_i}^{\infty} f(\epsilon)\, d\epsilon \\ &= \int_{-\infty}^{\infty} \mathbf{1}_{(\epsilon \geq -\beta X_i - \gamma Z + p_i)} f(\epsilon)\, d\epsilon \end{aligned}$$

Using the distribution quantile as the variable over which the integration takes place we obtain

$$\begin{aligned} Pr[V_i - p_i \geq 0] &= \int_0^1 \mathbf{1}_{(\beta X_i + \gamma Z - p_i + \epsilon \geq 0)}\, dF(\epsilon) \\ &= \int_0^1 \mathbf{1}_{(Q_\tau(V_i) - p_i)}\, d\tau \end{aligned}$$

The above integral is approximated by the sum

$$Pr[V_i - p_i \geq 0] \approx \frac{1}{M} \sum_\tau \mathbf{1}_{(Q_\tau(V_i) - p_i)}$$

[4] See Horowitz (1992) for further discussion on smoothed maximum score estimation and a derivation of the asymptotic distribution of the parameter vector for the median smoothed maximum score.

where M is the number of (equally spaced) estimated quantiles of V_i.[5] Therefore, the conditional probability that a consumer with any given set of characteristics will purchase the product is computed directly using the estimated quantiles obtained from equation (3) or equation (4). The market demand can then be obtained by aggregation over the entire set of consumers.

Similarly, since the latent variable V_i expresses money metric utility, the expected consumer surplus of an individual with characteristics X_i who is faced with the choice of purchasing a product with characteristics Z at a price p_i is given by

$$\begin{aligned} CS(X_i, Z, p_i) &= \int_{-\beta X_i - \gamma Z + p_i}^{\infty} (V_i - p_i)\ f(\epsilon)\ d\epsilon \\ &= \int_{-\infty}^{\infty} \mathbf{1}_{(\epsilon \geq -\beta X_i - \gamma Z + p_i)}\ (V_i - p_i)\ f(\epsilon)\ d\epsilon \end{aligned}$$

Using, as above, the distribution quantile as the variable over which integration takes place we obtain

$$\begin{aligned} CS(X_i, Z, p_i) &= \int_0^1 \mathbf{1}_{(\beta X_i + \gamma Z - p_i + \epsilon \geq 0)}\ (V_i - p_i)\ dF(\epsilon) \\ &= \int_0^1 \mathbf{1}_{(Q_\tau(V_i) - p_i \geq 0)} (V_i - p_i)\ d\tau \end{aligned}$$

This integral is approximated by the sum

$$CS(X_i, Z, p_i) \approx \frac{1}{M} \sum_\tau \mathbf{1}_{(Q_\tau(V_i) - p_i \geq 0)} (Q_\tau(V_i) - p_i)$$

where M is the number of equally spaced estimated quantiles of V_i. The consumer surplus for the entire market is obtained by aggregation over the entire set of consumers in the market.

3. Empirical Illustration

The dataset was obtained by a consumer survey that was distributed in-class to University of Illinois undergraduates in the Fall of 2000. The students were first shown a standard MP3

[5] Extension to non-equally spaced quantiles is straightforward.

player and were given information about the player's capabilities (most students appeared to have already been well informed about MP3 players). Then, the students were asked to fill the consumer survey, which inquired about (i) a few key characteristics of the students (gender and class year), (ii) a few questions that are potentially relevant to their willingness to pay for an MP3 player (whether or not they have a portable CD-player, how many hours per day do they listen to music, whether they have heard of MP3 players before, and whether or not they own a personal computer), and (iii) whether they would be willing to purhase an MP3 player at a particular price had they *not* owned an MP3 player and had they *not* been able to purchase an MP3 from any other source.[6] The offered price differed from student to student, and ranged from 50 dollars to 450 dollars.[7]

Summary statistics about the collected variables are presented in Table 1. About nine out of ten students own a personal computer and had already been informed about MP3, which reduces the value of these variables as potential regressors. Indeed, neither variable appears to affect willingness to pay for an MP3 player in an all-inclusive logit model (though knowledge of MP3 players, a proxy for intrinsic interest in these devices, comes close to be statistically significant). As mentioned above, current ownership of an MP3 player was only employed to ensure that respondents could abstract from current ownership in deciding whether or not they would be willing to purchase an MP3 at a given price. Indeed, MP3 ownership (which is low given that that the data was collected in the year 2000) is never statistically significant in a logit model. Notice that three times as many people would be willing to purchase an MP3 player than currently own one. The reason is that we have a large spread in the prices at which we offer an MP3 player to these respondents, and that

[6]To ensure that ownership of MP3 players did not contaminate students' answers, we also included as a cross-check a question about whether they currently owned an MP3 player (about 1 in 20 respondents did). We then investigated whether this variable was indeed not statistically significant, as we would have hoped; it's p-value was approximately 0.200. It seems, indeed, that students were able to abstract from the whether or not they currently own an MP3 player.

[7]Offered prices were equally likely to be 50, 100, 150, 200, 250, 300, 350, 400, and 450 dollars. These price points generate a reasonable spread given the prices of MP3 players in the first half of the year 2000.

the low range of these prices is substantially lower than the price of MP3 players at the time of the survey. Indeed, the responses of the respodents when we limit ourselves to prices that are similar to the (then) current prices of MP3 players appear to be very close to the actual ownership rate of MP3 players. Thus, students seem to have taken the survey seriously and responded truthfully.

We first run logit regressions to explain the decision of a respondent to purchase or not to purchase an MP3 player at the offered price. In the first regression, we included all explanatory variables. Only ownership of a portable CD player and number of hours spent listening to music were statistically significant. As explained above, current ownership of an MP3 player was only meant as a "survey quality assurance" variable, and was then removed from the regression. Ownership of a personal computer and being informed about MP3 players do not have much discriminating power across respondents, and were also removed from the regression. Class year was not expected to have any effect on willingness to pay for MP3 players, and did not, so it was also removed from the set of regressors. Of the remaining three variables, hours spent listening to music and ownership of a portable CD player continue to be statistically significant, and gender continued to be statitistically insignificant (price, of course, is highly significant in all specifications). Though this is would constitute the best formulated set of regressors (including all those that are on a priori grounds reasonable candidates for inclusion), we report the results of the specification that excludes gender. We use this more parsimonious model because binary quantile regressions are computationally very intensive, especially when computing standard errors via bootstrapping. Even with this most parsimonious set, it takes one day to run one replication in a Sun station.[8]

The logit results of our model are listed in the first column of Table 2. The second column of Table 2 shows the results of the same regression, re-scaling the coefficients so that the coefficient of price is equal to minus one. This permits an easier comparison between the logit results and the results of the quantile regressions, and expresses the effects of the various

[8] At this point, we are still in the middle of bootstrap replications.

variables in terms of dollars. An extra hour per day of listening to music increases mean willingness to pay for an MP3 player by 20.5 dollars. Owners of portable CD players are willing to spend an additional 91.5 dollars for MP3 players than non-owners of portable CD players. The logit results mask considerable heterogeneity in the impact of these variables as a function of propensity to purchase an MP3 player. Recall that propensity to purchase an MP3 player is the intensity of a consumer's preference for an MP3 player *conditional* on that consumer's characteristics. One can think of propensity a consumer to purchase as the ranking of that consumer's willingness to pay among consumers of the same characteristics. The effects of variables on the consumer willingness to pay for consumers with different propensity to purchase is given by the quantile regression results, a sample of which is reported in the remaining collumns of Table 2.[9] Number of hours listening to music has a substantial impact on willingness to pay for consumers with high propensity to prepensity to purchase and MP3 player, and a much lower impact for lower quantiles. This makes sense; if a consumer is not of a type that values MP3 players (perhaps he does not like headsets, or does not listen to music on the road, or does not like converting songs), then how much time he spends listening to music is not likely to be important is his purchase decision. On the other hand, the extent of listening to music is going to be much more important for a consumer type whose life-style preferences are compatible with the use of an MP3 player. Similarly, the effect of owning a portable CD player vary substantially across quantiles.

[9] We do not report results for low quantiles because there is some dependence of parameter estimates values on starting values for these quantiles. All of these parameter values correspond to the same value of the criterion function. Parameter estimates are non-unique because the objective function is a step function (and thus the optimum corresponds to a plateau). For low quantiles, the size of this plateau tends to be large because most consumers of low propensity to purchase choose not to purchase, and thus, there are locally a large set of parameter values that correspond to the same set of "predictions." The non-uniqueness of the parameter vector does not materially affect the estimated demand curves (as reported in the figures below), especially for prices above 50 dollars, because the demand curve estimates is for the most part obtained from the parameter estimates for high quantiles. The results in Table 2 and the figures below are those obtained by using the logit parameter estimates as starting values.

Of greater interest, however, are the estimates of the demand for consumers of different characteristics. Figure 1 plots the demand curves implied by the logit specification for three different hypothetical types of consumers. The first is the set of consumers with no portable CD player who listens to music for one hour each day, the second is the set of consumers who owns a portable CD player and who spends one hour each day listening to music, and the third is the set of consumers who own a portable CD player and spend 3 hours a day listening to musics. The three demand curves have the shape of the inverse logistic cummulative distribution (truncated at zero), and differ from each other by a parellel shift. Figures 2 through 5 compare these logit-based estimates with those obtained from binary quantile regressions. For the first group of consumers, the logit-based estimate of the demand curve underestimates the rapid decline of the demand curve after an small initial "plateau" at high prices, and then misses the relatively broad demand for MP3 players at low prices. For the second group of consumers, the differences between the logit-based and quantile-based demand estimates are more pronounced; the logit-based demand for MP3 players is twice that predicted by the quantile regressions for prices around 150 dollars, and two-thirds of that predicted by the quantile regressions for prices below 50 dollars. The differences between the two sets of estimates are less pronounced (in proportional terms) for the third set of consumers.

The flexibility in the predicted patterns of demand shifts that quantile-estimation allows is demonstrated in Figure 5, which shows the quantile-based estimates of the demand curve for the three types of consumers. The demand shift from owning a CD player is qualitatively very different from the demand shift from spending more time listening to music. The former is more pronounced at the bottome end of the demand curve, the latter more pronounced at the higher end of the demand curve, though neither factor is important amongst the consumers with the highest propensity to purchase MP3 players. This pattern has intuitive appeal. Those who have a high value for MP3 players, relatively to people of the same characteristics, have a similar willingness to pay for MP3 regardless of whether they own a CD-player or of how much time they spend listening to music. For the with somewhat lower propensity for MP3 players, duration of listening to music is very important, but ownership of a portable CD player is not an important indicator of willingness to pay. For the bottom half of the market, intensity of music listening is not an important indicator of willingness to pay, because many of these consumers have preferences or lifestyles that are less compatible to MP3 listening. Ownership of a portable CD player is a more important indicator for this group because it better picks the sub-set of consumers with lifestyles somewhat more conducive to use of MP3 players, by picking the portion of consumers who like to listen to music on-the-go.

The differences in the demand estimates are also economically meaningful. MP3 player production is not a monopoly, but we perform a benchmark exercise that is meant to capture the importance of using our approach to estimate demand for a firm that is planning the introduction of a product. In Table 3 we show the optimal monopoly price (at various levels of marginal cost and for homogeneous markets that consist of the three consumer types) computed by a decision maker who uses the logit model and by a decision maker that uses the quantile-based estimates of the demand curve. A decision maker who uses the quantile-based estimates of the demand would charge a higher price for low values of the marginal cost, would increase it rapidly if marginal cost were somewhat higher, but would keep the same for further increases in marginal cost. This closely matches what one would expect by looking at the shape of the demand curves in Figure 5. A decision maker who uses the logit model, on the other hand, would charge a price that varies smoothly with marginal cost. Since our approach is more flexible in that it nests logit as a special case, it is more reasonable to consider the quantile-based estiamates as being closer to the "true" demand. Under this pressumption, we calculated the percentage loss from incorrectly using the logit-based estimates of the demand. Table 4 reports the profits the monopolist would earn if he charged the price price implied by the logit demand, when the true demand is given by Figure 5, as a percent of the profit he would obtain if he used the quantile-based estimates of the demand. For medium values of marginal cost, profits would be halfed by using the incorrect estimates of the demand curve!

The second panel of Table 3 reports the profits as forecasted by a manufacturer who is contemplating to introduce the product (here we assume that the manufacturer who uses logit demand to compute the optimal price, also uses the logit demand to calcuate his profits). The difference in profits is substantial. Assuming that the introduction of the product involves a fixed cost, a monopolist who uses the quantile-based demand, and a monopolist who uses the logit based demand would often reach different conclusions about the desirability of introducing the product. The last panel of Table 3 shows the calculation of the consumer surplus at different prices for each of the three benchmark sets of consumers, a calculation that may be important for policy analysis. Again, the differences between the two sets of estimates is substantial.

4. Concluding Remarks

to be completed

References

Horowitz, J [1992] "A Smoothed Maximum Score Estimator for the Binary Response Model," *Econometrica*, vol. 60, pages 505-531.

Manski, C. [1975] "Maximum Score Estimation of the Stochastic Utility Model of Choice," *Journal of Econometrics*, vol. 3, pages 205-228.

Manski, C. [1985] "Semiparametric Analysis of Discrete Response: Asymptotic Properties of the Maximum Score Estimator," *Journal of Econometrics* vol. 32, pages 65-108.

Kordas, G. [2000] *Binary Regression Quantiles*, Ph.D. Dissertation, University of Illinois, Urbana-Champaign.

Table 1. Descriptive Statistics

Variable	Mean	Standard Dev.	Minimum	Maximum
Hours Listening to Music	2.45	1.72	0	6
Owns Portable CD	0.74	0.44	0	1
Owns Personal Computer	0.89	0.31	0	1
Female	0.44	0.50	0	1
Class Year	3.16	0.76	1	5
Owns MP3	0.06	0.23	0	1
Is Informed about MP3	0.91	0.28	0	1
Offered Price	238.72	127.51	50	450
Willing to Purchase	0.18	0.38	0	1

Notes: Class Year takes the value of 1 for freshman, 2 for sophomore, 3 for junior, 4 for senior, 5 for graduate student. The sample size for all variables is 257.

Table 2. Estimation Results

Variable	Logit regression		Quantile regression			
	standard	rescaled	90th percentile	80th percentile	70th percentile	60th percentile
Intercept	-0.825	-67.458	-19.046	71.583	34.575	-75.018
	0.576					
Hours	0.253	20.664	54.884	16.261	13.168	15.001
	0.107					
Owns Portable CD	1.119	91.560	111.559	23.582	64.798	168.881
	0.504					
Price	-0.012	-1.000	-1.000	-1.000	-1.000	-1.000
	0.002					
Log-likelihood	-89.055					
Sum of weighted absolute deviations			36.500	45.200	60.300	75.400

Note: Standard errors are listed below the parameter estimates in italics. The number of observations is equal to 257. Percentiles are measured from low to high willingness to pay (conditional on consumer characteristics).

Table 3. Comparison of Results: Optimal Pricing and Equilibrium Profits.

	Optimal Monopoly Price					
	Quantile-based Estimates			Logit-based Estimates		
Marginal Cost	CD=0, Hrs=1	CD=1, Hrs=1	CD=1, Hrs=3	CD=0, Hrs=1	CD=1, Hrs=1	CD=1, Hrs=3
0	80	105	135	96	116	130
50	270	105	135	140	153	163
100	270	300	345	185	194	211
150	270	300	345	232	239	244
200	270	300	345	280	285	289
250	270	300	345	329	334	335
	Equilibrium Monopoly Profits					
	Quantile-based Estimates			Logit-based Estimates		
Marginal Cost	CD=0, Hrs=1	CD=1, Hrs=1	CD=1, Hrs=3	CD=0, Hrs=1	CD=1, Hrs=1	CD=1, Hrs=3
0	15	43	58	14	34	48
50	11	23	37	8	22	32
100	9	18	22	5	13	20
150	6	14	18	3	8	12
200	4	9	13	1	4	7
250	1	5	9	1	2	4
	Consumer Surplus					
	Quantile-based Estimates			Logit-based Estimates		
Price	CD=0, Hrs=1	CD=1, Hrs=1	CD=1, Hrs=3	CD=0, Hrs=1	CD=1, Hrs=1	CD=1, Hrs=3
50	19	46	71	22	54	77
100	10	24	44	13	34	50
150	6	14	25	7	20	31
200	4	10	16	4	11	18
250	1	5	10	2	6	10
300	0	1	5	1	4	6

Notes: Logit-based equilibrium profits are obtained assuming logit specification is correct.

Table 4. Profit Loss From Incorrectly Imposing Logit Specification (in percent).

Marginal Cost	CD=0, Hrs=1	CD=1, Hrs=1	CD=1, Hrs=3
0	44.2%	45.8%	96.3%
50	40.9%	41.1%	52.6%
100	50.0%	47.0%	75.5%
150	68.3%	59.3%	58.9%
200	0.0%	85.0%	61.4%
250	0.0%	0.0%	89.5%

Notes: Profits are calculating using the quantile-based estimate of the demand function at prices that are optimal using the logit demand specification.

Figure 1. Logit-Based Demand Estimates

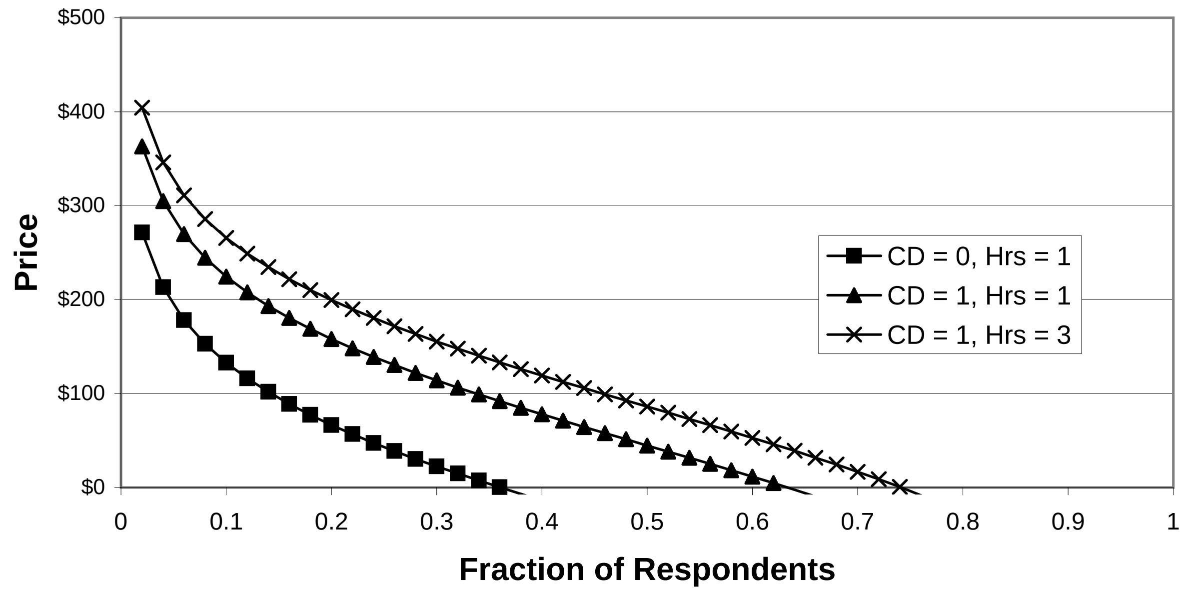

Figure 2. Demand Estimates; CD=0, Hrs=1

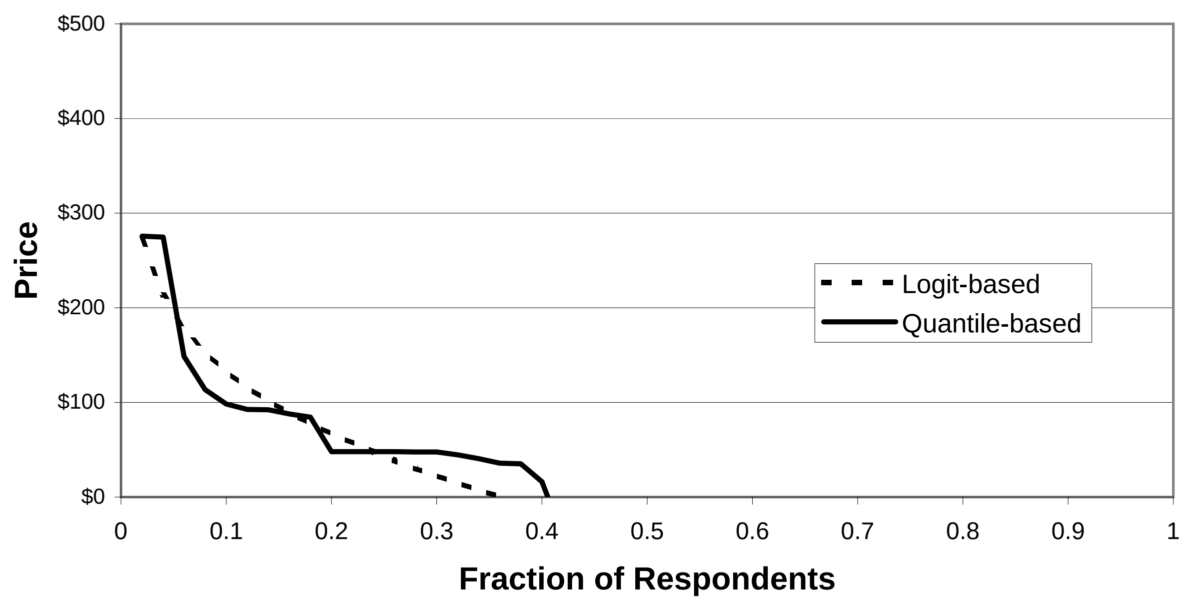

Figure 3. Demand Estimates; CD=1, Hrs=1

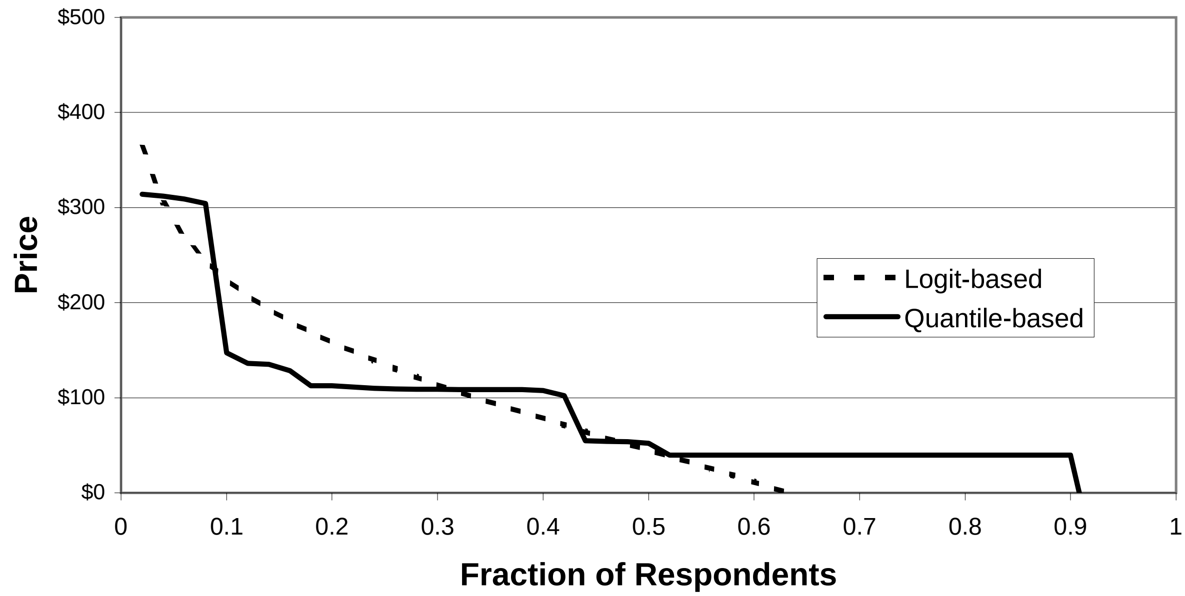

Figure 4. Demand Estimates; CD=1, Hrs=3

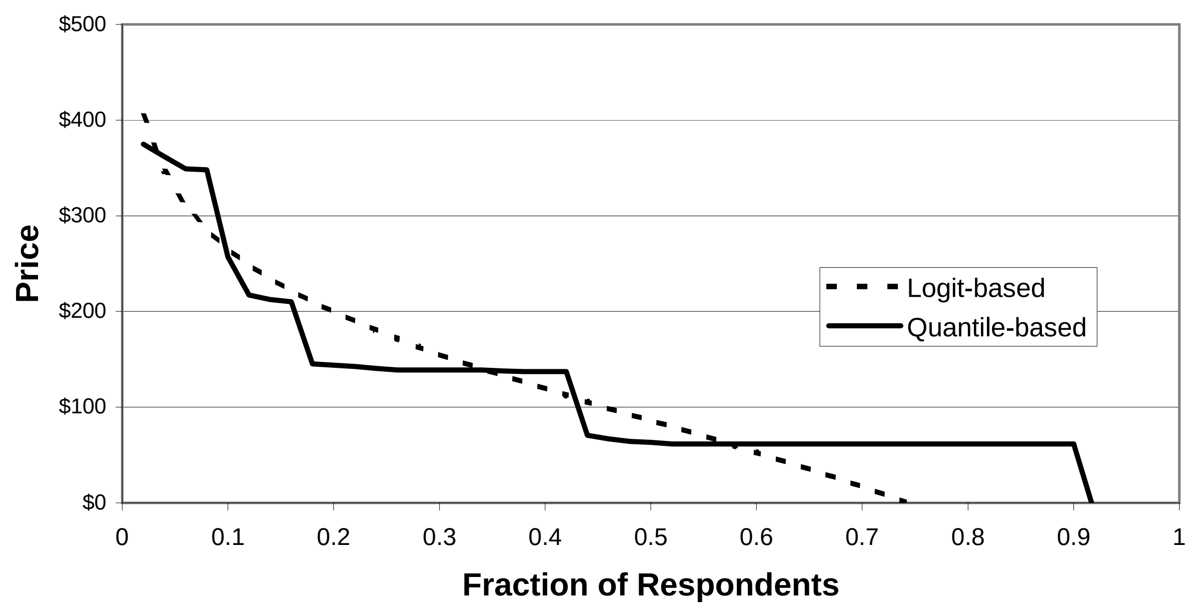

Figure 5. Quantile-Based Demand Estimates