Nonparametric estimation of a heterogeneous demand function under the Slutsky inequality restriction

Richard Blundell
Joel Horowitz
Matthias Parey

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP54/13

Nonparametric Estimation of a Heterogeneous Demand Function under the Slutsky Inequality Restriction

Richard Blundell*, Joel Horowitz† and Matthias Parey‡§

October 2013

Abstract

Economic theory rarely provides a parametric specification for a model, but it often provides shape restrictions. We consider nonparametric estimation of the heterogeneous demand for gasoline in the U.S. subject to the Slutsky inequality restriction of consumer choice theory. We derive conditions under which the demand function can be estimated consistently by nonparametric quantile regression subject to the Slutsky restriction. The estimated function reveals systematic variation in price responsiveness across the income distribution. A new method for estimating quantile instrumental variables models is also developed to allow for the endogeneity of prices. In our application, shape-constrained quantile IV estimates show similar patterns of demand as shape-constrained estimates under exogeneity. The results illustrate the improvements in the finite-sample performance of a nonparametric estimator that can be achieved by imposing shape restrictions based on economic theory.

*Ricardo Professor of Political Economy, Department of Economics, University College London (UCL), and Institute for Fiscal Studies (IFS). 7 Ridgmount Street, London WC1E 7AE, United Kingdom, email: r.blundell@ucl.ac.uk, tel: 44 (0)20 7679 5863, fax: 44 (0)20 7916 2775 (corresponding author).

†Charles E. and Emma H. Morrison Professor of Market Economics, Department of Economics, Northwestern University, and Cemmap. 2001 Sheridan Road, Evanston, Illinois 60208, USA, email: joel-horowitz@northwestern.edu.

‡Lecturer (Assistant Professor), Department of Economics, University of Essex, and Institute for Fiscal Studies (IFS). Wivenhoe Park, Colchester CO4 3SQ, United Kingdom, email: mparey@essex.ac.uk.

§We thank seminar participants at Cologne, Essex, Manchester, Southampton, the 2010 Conference on Shape Restrictions at Northwestern, and the 2013 Statistical Meeting (DAGStat) in Freiburg for helpful comments.

1 Introduction

Nonparametric methods are frequently used in demand analysis and have become an indispensable tool for understanding consumer behavior (Lewbel, 1997; Blundell, 2005). In addition to increased availability of detailed micro data and advances in computational methods, this reflects concerns that simple parametric approaches may lead to problems of misspecification. These in turn may lead to biases in estimated coefficients and potentially misleading conclusions about the substantive questions of interest. Nonparametric methods, by contrast, are consistent.

Nonetheless, the use of nonparametric methods can be demanding in terms of data requirements. Even with large data sets it may be difficult in practice to estimate precisely the object of interest, especially in the case of multiple regressors (see e.g. Silverman (1986)). Endogeneity concerns may necessitate the use of control function or instrumental variables approaches, potentially adding further complexity. While these problems are well understood in principle, they pose extraordinary challenges for the applied researcher trying to estimate nonparametric models given the typical sample sizes, and raise the question how the researcher can add structure to the estimates without incurring the risk of misspecification inherent in simple parametric models.

In this paper we argue that constraints implied by economic theory provide powerful shape restrictions which allow us to impose meaningful structure on otherwise nonparametric estimates. These restrictions enable the estimates to be informed by an underlying model of behavior rather than by potentially arbitrary functional form assumptions. However, these restrictions apply to individual demand behavior. In this paper we utilize a monotonicity assumption on unobserved heterogeneity together with quantile estimation to recover individual demands.

We present methods for incorporating the Slutsky constraint in the context of demand estimation, and we illustrate these methods using the application of gasoline demand in the U. S. Given the changes in the price of gasoline that have been observed in recent years, and the role of taxation in the gasoline market, understanding the elasticity of demand is of key policy interest. We pay particular attention to the question of how demand behavior varies across the income distribution, and ask whether the welfare implications of price changes are uniform across the income distribution. This is an example where very simple parametric models impose strong restrictions on the behavioral responses allowed for, which may in turn affect resulting policy conclusions.

Imposing Slutsky negativity has several benefits. It ensures that our estimates are consistent with utility maximization, and allows us to recover welfare measures. Violations of the Slutsky constraint might translate into misleading (and potentially even wrongly signed) welfare measures such as deadweight loss. The integrability conditions allow one to verify that a specific parametric specification satisfies the requirements to recover preferences; Stern (1986) presents such an analysis for a range of commonly used specifications in the context of labor supply. Our method can be thought of as a nonparametric analogue to this approach, where admissible candidate estimates are restricted to those satisfying these restrictions of consumer theory. Where prices take only a few discrete values an equivalent approach is to impose the Afriat revealed preference inequalities, see Blundell, Kristensen, and Matzkin (2011).

In terms of statistical precision we expect the additional structure provided by this restriction to improve the finite-sample performance of our estimator, analogous to the way sign restrictions in parametric models reduce the Mean Squared Error (MSE). Nonparametric estimation often requires the choice of bandwidth parameters, such as Kernel bandwidths or number of knots for a spline. These parameters are optimally chosen in a way which balances bias and variance of the estimates. The use of shape restrictions, reducing the variance of the estimates, modifies this trade-off, and therefore allows potentially for smaller optimal bandwidth choices. Shape restrictions can therefore be thought of as a substitute for bandwidth smoothing, helping to recover the features of interest of the underlying relationship.

In previous work, we have investigated gasoline demand, focussing on the conditional mean (Blundell, Horowitz, and Parey, 2012). If unobserved heterogeneity enters in a non-separable manner, the conditional mean represents an average across unobservables, which may be difficult to interpret. As Brown and Walker (1989) and Lewbel (2001) have shown, demand functions generated from random utility functions do not typically yield demand functions where the unobserved tastes are additive. The identification and estimation of consumer demand models that are consistent with unobserved taste variation require analyzing demand models with nonadditive random terms. Under suitable restrictions, quantile estimation allows us to recover demand at a specific point in the distribution of unobservables. Matzkin (2003, 2008) derives general identification results for models that are non-separable in unobserved heterogeneity. This motivates our interest in a quantile estimator. Quantile regression also allows us to study differential effects of price changes and welfare costs across the distribution of unobservables.[1] The quantile estimation allows us to compare heavy users with moderate or light users.

The paper makes a number of contributions. We present a quantile estimator which incorporates shape restrictions. We develop a new estimator for the case of quantile estimation under endogeneity. We apply these methods in the context of gasoline demand. The nonparametric estimate of the demand function is noisy due to random sampling errors. The estimated function is non-monotonic, and there are instances where the estimate, taken at face value, is inconsistent with economic theory. When we impose the Slutsky restriction of consumer theory on the demand function, this approach yields

[1] In the context of alcohol demand, for example, Manning, Blumberg, and Moulton (1995) show that price responsiveness differs at different quantiles.

well-behaved estimates of the demand function and welfare costs across the income and taste distribution. Comparing across income groups and quantiles, our work allows us to document differences in demand behavior across both observables and unobservables.

The paper proceeds as follows. The next section develops our non-separable model of demand behavior and the restrictions required for a structural interpretation. Section 3 presents our estimation method, where we describe the nonparametric estimation method for both the unconstrained estimates and those obtained under the Slutsky constraint. We also present our procedure for quantile estimation under endogeneity. In Section 4 we discuss the data we use in our investigation and present our empirical findings. We compare the quantile demand estimates to those from a conditional mean regression. The endogeneity of prices is considered in Section 5 where we present the results of an exogeneity test and our quantile instrumental variables procedure. Section 6 concludes.

2 Unobserved Heterogeneity and Structural Demand Functions

The model of interest in this paper is

$$W = g(P, Y, U), \tag{1}$$

where W is demand (measured as budget share), P is price, Y is income, and U represents unobserved heterogeneity. We impose two types of restrictions on this demand function: The first set of restrictions addresses the way unobserved heterogeneity enters demand, and its relationship to price and income. We assume that demand g is monotone in the unobserved heterogeneity U. To ensure identification, we for now assume that U is statistically independent of (P, Y). Given these assumptions, we can also assume without loss of generality that $U \sim \mathrm{U}[0;1]$. This allows recovery of the demand function for specific types of households from the observed conditional quantiles of demand: the α quantile of W, conditional on (P, Y), is

$$\mathcal{Q}_\alpha(W|P, Y) = g(P, Y, \alpha) \equiv G_\alpha(P, Y). \tag{2}$$

Thus, the underlying demand function, evaluated at a specific value of the unobservable, can be recovered via quantile estimation.

In contrast, the conditional mean is

$$\begin{aligned} E\left(W|P=p, Y=y\right) &= \int g(p, y, u)\, f_U(u)\, du \\ &\equiv m(p, y), \end{aligned}$$

where $f_U(u)$ is the probability density function of U. Given that we are interested in imposing shape restrictions based on consumer theory, estimating the demand function at a specific value of $U = \alpha$ using quantile methods is attractive because economic theory informs us about $g(\cdot)$ rather than $m(\cdot)$. It is possible therefore that $m(\cdot)$ does not satisfy the restrictions even though each individual consumer does (see also Lewbel (2001)).

We will later relax the assumption of independence between U and the price P, test for endogeneity following the cost-shifter approach in Blundell, Horowitz, and Parey (2012) and present instrumental variables estimates. Imbens and Newey (2009) define the quantile structural function (QSF) as the α-quantile of demand $g(p, y, U)$, for fixed p and y; under endogeneity of prices, the QSF will be different from the α-quantile of $g(P, Y, U)$, conditional on $P = p$ and $Y = y$.

Hausman and Newey (2013) consider the case of multi-dimensional unobserved heterogeneity; they show that in this case neither the demand function nor the dimension of heterogeneity is identified, and use bounds on the income effect to derive bounds for average surplus. In the context of scalar heterogeneity, Hoderlein and Vanhems (2011) consider identification of welfare effects, and allow for endogenous regressors in a control function approach. Hoderlein and Stoye (2013) investigate how violations of the Weak Axiom of Revealed Preference (WARP) can be detected in a heterogeneous population based on repeated cross-sectional data, using copula methods to bound the fraction of the population violating WARP.

We impose the Slutsky constraint by restricting the price and income responses of the demand function g. Preference maximization implies that the Slutsky substitution matrix is symmetric negative semidefinite (Mas-Colell, Whinston, and Green, 1995). Ensuring that our estimates satisfy this restriction is however not only desirable because of the increase in precision from additional structure, it is also a necessary restriction in order to be able to perform welfare analysis. Welfare analysis requires knowledge of the underlying preferences. The question under which conditions we can recover the utility function from the observed Marshallian demand function, referred to as the integrability problem, has therefore been of long-standing interest in the analysis of consumer behavior (Hurwicz and Uzawa (1971)). A demand function which satisfies adding up, homogeneity of degree zero, and a symmetric negative semidefinite Slutsky matrix allows recovery of preferences (Deaton and Muellbauer (1980)). As Deaton and Muellbauer (1980) emphasize, these characteristics also represent the *only* structure that is implied by utility maximization. Slutsky negative semidefiniteness is therefore critical for policy analysis of changes in the prices consumers face. In the context of the two good model considered here, these integrability conditions are represented through the negative compensated price elasticity of gasoline demand.

In previous work household demographics or other household characteristics have been found to be relevant determinants of transport demand. One possibility of accounting for these characteristics would be to incorporate them in a semiparametric specification. However, in order to maintain the fully nonparametric nature of the model, we instead condition on a set of key demographics in our analysis.[2] Thus we address the dimension-reduction problem by conditioning on a particular set of covariates. This exploits the fact that the relevant household characteristics are all discrete in our application. We then estimate our nonparametric specification on this sample which is quite homogeneous in terms of household demographics.

3 Nonparametric Estimation

3.1 Unconstrained Nonparametric Estimation

From equation (2), we can write

$$W = G_\alpha(P, Y) + V_\alpha; \ \ P(V_\alpha \leq 0 \mid P, Y) = \alpha, \tag{3}$$

where V_α is a random variable whose α quantile conditional on (P, Y) is zero. We estimate G_α using a truncated B-spline approximation with truncation points M_1 and M_2 chosen by cross-validation. Thus

$$G_\alpha(P, Y) = \sum_{m_1=1}^{M_1} \sum_{m_2=1}^{M_2} c_{m_1,m_2;\,\alpha}\, B^p_{m_1}(P) B^y_{m_2}(Y),$$

where B^p and B^y (with indices m_1 and m_2) are spline functions following Powell (1981) and $c_{m_1,m_2;\,\alpha}$ is the finite-dimensional matrix of coefficients.

We denote the data by $\{W_i, P_i, Y_i : i = 1, ..., n\}$. The estimator is defined in the following optimization problem:

$$\min_{\{c_{m_1,m_2;\,\alpha}\}} \sum_{i=1}^{n} \rho_\alpha \left(W_i \ - \ G_\alpha(P_i, Y_i)\right), \tag{4}$$

where $\rho_\alpha(V) = (\alpha - \{V < 0\})\, V$ is the check function.

3.2 Estimation Subject to the Slutsky Inequality

The Slutsky condition is imposed on the nonparametric estimate of the conditional quantile function. Writing this condition in terms of shares, and taking price and income

[2] These characteristics include household composition and life-cycle stage of the household, race of the survey respondent, and as well as the urban-rural location of the household. We describe these selection criteria in detail in Section 4.1 below.

to be measured in logs, gives

$$\frac{\partial \hat{G}^C_\alpha(P,Y)}{\partial p} + \hat{G}^C_\alpha(P,Y)\frac{\partial \hat{G}^C_\alpha(P,Y)}{\partial y} \leq \hat{G}^C_\alpha(P,Y)\ \left(1 - \hat{G}^C_\alpha(P,Y)\right), \tag{5}$$

where the superscript C indicates that the estimator is constrained by the Slutsky condition. The constrained estimator is obtained by solving the problem (4), subject to (5), for all (P,Y). This problem has uncountably many constraints. We replace the continuum of constraints by a discrete set, thereby solving:

$$\min_{\{c_{m_1,m_2;\alpha}\}} \sum_{i=1}^{n} \rho_\alpha \left(W_i\ -\ \hat{G}^C_\alpha(P_i,Y_i)\right)$$

subject to

$$\frac{\partial \hat{G}^C_\alpha(p_j,y_j)}{\partial p} + \hat{G}^C_\alpha(p_j,y_j)\frac{\partial \hat{G}^C_\alpha(p_j,y_j)}{\partial y} \leq \hat{G}^C_\alpha(p_j,y_j)\ \left(1 - \hat{G}^C_\alpha(p_j,y_j)\right), \quad j = 1,...,J,$$

where $\{p_j, y_j : j = 1,...,J\}$ is a grid of points. To implement this, we use a standard optimization routine from the NAG library (E04UC). In the objective function we use a check function which is locally smoothed in a small neighborhood around 0 (Chen (2007)). In Appendix Figure A.1, we show that the resulting demand figures are not sensitive to a range of alternative values of the corresponding smoothing parameter. For imposing the constraints, we choose a fine grid of points along the price dimension, at each of the 15 income category midpoints.

We use the bootstrap for inference under the assumption that the Slutsky constraint does not bind in the population.

3.3 Welfare Measures

The estimates of the Slutsky constrained demand function can then be used to recover measures of deadweight loss (DWL). For this purpose, we consider a hypothetical tax change which moves the price from p^0 to p^1. Let $e(p)$ denote the expenditure function at price p and some reference utility level. The DWL of this price change is given by

$$L(p^0,p^1) = e(p^1) - e(p^0) - (p^1 - p^0)\ G_\alpha\left[p^1, e(p^1)\right].$$

$L(p^0,p^1)$ is computed by replacing e and g with consistent estimates. The estimator of e, $\hat{e}$, is obtained by numerical solution of the differential equation

$$\frac{d\hat{e}(t)}{dt} = \hat{G}_\alpha\left[p(t),\hat{e}(t)\right]\frac{dp(t)}{dt},$$

where $[p(t),\hat{e}(t)]$ $(0 \leq t \leq 1)$ is a price-(estimated) expenditure path.

3.4 Quantile Instrumental Variable Estimation

To recognize potential endogeneity of prices, we introduce a cost-shifter instrument Z for prices. In the application this is a distance measure to gulf supply refinery to reflect transport costs. Consider again equation (3) from above, where now we impose the quantile restriction conditional on the distance instrument (and household income):

$$W = G_\alpha(P, Y) + V_\alpha; \;\; P(V_\alpha \leq 0 \mid Z, Y) = \alpha.$$

The identifying relation can be written as

$$P(W - G_\alpha(P, Y) \leq 0 \mid Z, Y) = \alpha.$$

Let $f_{Z,Y}$ be the probability density function of (Z, Y). Then we have

$$\int_{Z \leq z, Y \leq y} P(W - G_\alpha(P, Y) \leq 0 | Z, Y) \; f_{Z,Y}(Z, Y) \; dZ \; dY = \alpha P(Z \leq z, Y \leq y)$$

for all (z, y). An empirical analog is

$$n^{-1} \sum_{i=1}^{n} 1 \left[W_i - G_\alpha(P_i, Y_i) \leq 0 \right] 1 \left[Z_i \leq z, Y_i \leq y \right] = \frac{\alpha}{n} \sum_{i=1}^{n} 1 \left[Z_i \leq z, Y_i \leq y \right].$$

Define

$$Q_n(G_\alpha, z, y) = n^{-1} \sum_{i=1}^{n} \left\{ 1 \left[W_i - G_\alpha(P_i, Y_i) \leq 0 \right] - \alpha \right\} 1 \left[Z_i \leq z, Y_i \leq y \right]. \qquad (6)$$

Estimate G_α by solving

$$\min_{G_\alpha \epsilon \mathcal{H}_n} \int Q_n(G_\alpha, z, y)^2 \; dz \; dy,$$

where $\mathcal{H}_n$ is the finite-dimensional space consisting of truncated series approximations and includes the shape restriction when we impose it.

4 Estimation Results

4.1 Data

The data are from the 2001 National Household Travel Survey (NHTS). The NHTS surveys the civilian non-institutionalized population in the United States. This is a household-level survey conducted by telephone and complemented by travel diaries and odometer readings.[3] We select the sample to minimize heterogeneity as follows: we re-

[3]See ORNL (2004) for further detail on the survey.

strict the analysis to households with a white respondent, two or more adults, at least one child under age 16, and at least one driver. We drop households in the most rural areas, given the relevance of farming activities in these areas.[4] We also restrict attention to those localities where the state of residence is known, and omit households in Hawaii due to its different geographic situation compared to continental U.S. states. Households where key variables are not reported are excluded and we restrict attention to gasoline-based vehicles (rather than diesel, natural gas, or electricity), requiring gasoline demand of at least one gallon; we also drop one observation where the reported gasoline share is larger than 1. We take vehicle ownership as given and do not investigate how changes in gasoline prices affect vehicle purchases or ownership. The results by Bento, Goulder, Jacobsen, and von Haefen (2009) indicate that price changes operate mainly through vehicle miles traveled rather than through fleet composition: they find that more than 95% of the reduction in gasoline consumption in response to an increase in gasoline tax is due to a reduction in vehicle miles traveled.

The resulting sample contains 3,640 observations. The key variables of interest are gasoline demand, price of gasoline, and household income. Corresponding sample descriptives are reported in Table 1; further detail on these variables can be found in Blundell, Horowitz, and Parey (2012).[5]

[TABLE 1 ABOUT HERE]

The nonparametric estimates are shown below for the three income groups whose midpoints in 2001 dollars are $42,500, $57,500 and $72,500. These income levels are chosen to compare the behavior of lower, middle and upper income households.[6]

We use cubic B-splines for our nonparametric analysis.[7] For each quantile of interest, the number of knots is obtained by cross-validation, separately for each quantile.[8] The resulting number of (interior) knots is shown in Panel (1) of Table 2. In particular, at the median, the procedure indicates 4 interior knots in the price dimension and 3 knots in the income dimension. Across the quartiles, we obtain the same number of knots in the

[4]These are households in rural localities according to the Claritas urbanicity index, indicating a locality in the lowest quintile in terms of population density (ORNL (2004, Appendix Q)).

[5]In the nonparametric analysis below, we impose two additional restrictions to avoid low-density areas in the data. For this purpose, we restrict attention to households with (2001) household income of at least $15,000, facing a price of at least $1.20.

[6]These three income points occupy the 19.1-22.8th, 34.2-42.3th, and 51.7-55.9th percentiles of the income distribution in our data (see Table 1).

[7]In the income dimension, we place the knots at equally-spaced percentiles of a normal distribution, where we have estimated the corresponding mean and variance in our data. In the (log) price dimension we space the knots linearly.

[8]Following equation (1) we use the budget share as dependent variable in the cross-validation. Given that our analysis focuses on the demand behavior for the three income levels of interest, we evaluate the cross-validation function only for observations which are not too far from these income points, and use 0.5 (in the log income dimension) as cutoff.

income dimension, while in the price dimension the cross-validation procedure indicates a more restrictive B-spline for the first quartile ($\alpha = 0.25$).

In the subsequent analysis we follow these knot choices for both the unconstrained and the constrained quantile estimates under exogeneity. We have also investigated whether this cross-validation outcome is sensitive to outliers in the share variable. For this purpose, we have repeated the cross-validation procedure, leaving out the 10 highest and the 10 lowest gasoline budget share observations. The results are reported in Panel (2) of Table 2, suggesting that overall the number of knots is not very sensitive to this exercise.

[TABLE 2 ABOUT HERE]

4.2 Quantile estimates under exogeneity of prices

Parametric benchmark specifications using linear quantile estimates can be found in Table 3, where we regress log quantity on log price and log income:

$$\log Q = \beta_0 + \beta_1 \log P + \beta_2 \log Y + U; \quad \mathcal{Q}_\alpha(U|P,Y) = 0.$$

For comparison we also report estimates obtained using an OLS estimator (see column (4)). These indicate a price elasticity of -0.83 and an income elasticity of 0.34. These are similar to those reported by others (see Hausman and Newey (1995); Schmalensee and Stoker (1999); West (2004); Yatchew and No (2001)).

The quantile regression estimates are reported in columns (1)-(3), revealing plausible and interesting patterns in the elasticities across quantiles. At lower quantiles, the estimated price elasticity is much higher (in absolute values) than at higher quantiles.[9] Similarly, the estimated income elasticity declines strongly as we move from the first quartile to the median, and from the median to the third quartile. Thus, low-intensity users appear to be substantially more sensitive in their demand responses to price and income variation than high-intensity users.

A natural question is whether this benchmark specification is appropriately specified. To investigate this, we perform the specification test for the linear quantile regression model developed in Horowitz and Spokoiny (2002). The results are reported in Table 4. We clearly reject our baseline specification at a 5% level. This holds whether we measure our dependent variable as log quantity or as gasoline budget share.

[TABLES 3–4 ABOUT HERE]

We have also augmented the specification reported in Table 3 with squares and cubes of price and income and found these to be significant. This suggests that the parametric

[9] A similar pattern is reported in Frondel, Ritter, and Vance (2012) using travel diary data for Germany.

benchmark model may be misspecified. We therefore now proceed to the nonparametric analysis.

Figure 1 shows the nonparametric estimates. Each panel corresponds to a particular point in the income distribution. The line shown with open markers represents the unconstrained estimates, together with the corresponding bootstrapped confidence intervals (solid lines). As can be seen in panel (b) for the middle income level, for example, the unconstrained estimates show overall a downward-sloping trend, but there are several instances where the estimated demand is upward sloping. A similar pattern is also found in Hausman and Newey (1995). Although here we plot the Marshallian demand estimate, these instances of upward sloping demand also point to violations of the Slutsky negativity when we compensate the household for the increase in prices. The line shown as filled markers represents the estimate constrained by the Slutsky shape restriction. By design, the constrained estimates are consistent with economic theory.

Interestingly, the constrained and the unconstrained estimates are both well contained in a 90% confidence band around the unconstrained ones; this pattern is consistent with the random sampling error interpretation. At the same time, the constrained estimates show that imposing the shape constraint can also be thought of as providing additional smoothing. Focussing on the constrained estimates, we compare the price sensitivity across the three income groups. The middle income group appears to be more price sensitive than either the upper or the lower income group; this is a pattern also found in Blundell, Horowitz, and Parey (2012).

[FIGURE 1 ABOUT HERE]

4.3 Comparison across quantiles and the conditional mean estimates

Figure 2 compares the quantile estimates across the three quartiles, holding income constant at the middle income group. In the unconstrained estimates, the differences in flexibility (corresponding to the cross-validated number of knots in the price dimension) are clearly visible. The constrained estimates, however, are quite similar in shape, suggesting that they may approximately be parallel shifts of each other. This would be consistent with a location-scale model together with conditional homoskedasticity (Koenker (2005)). Under this model, conditional mean estimates would show the same shape as seen in the conditional quartile results, and we turn to this comparison now.

[FIGURE 2 ABOUT HERE]

As noted in the introduction, we have previously investigated gasoline demand, focussing on the conditional mean (Blundell, Horowitz, and Parey (2012)). That analysis used a Kernel regression method, in which the shape restriction is imposed by reweighting the data in an approach building on Hall and Huang (2001). As in the quantile demand results here we found strong evidence of differential price responsiveness across the income distribution, suggesting a stronger price responsiveness in the middle income group. Figure 3 shows the conditional mean regression estimates, where we use the same B-spline basis functions as in the quantile results presented above (Figure 1). The shape of these two sets of estimates is remarkably similar, especially for the constrained estimates; in terms of levels, the mean estimates are somewhat higher than the median estimates (by around 0.1 on the log scale).

[FIGURE 3 ABOUT HERE]

4.4 Welfare measurement

The Slutsky constrained demand function estimates can in turn be used for welfare analysis of changes in prices. For this purpose we consider a change in price from the 5th to the 95th percentile in our sample for the nonparametric analysis, and we report Deadweight Loss measures corresponding to this price change. Table 5 shows the DWL estimates for the three quartiles. In the constrained estimates, we find that the middle income group has the highest DWL at all quartiles. This is consistent with the graphical evidence presented in Figure 1 above. The table also shows the DWL estimates implied by the parametric estimates corresponding to a linear specification. The uniform patterns in the corresponding DWL figures (within each quantile) reflect the strong assumptions underlying these functional forms, which have direct consequences for the way DWL measures vary across these subgroups in the population.

There are two instances (both for the lower-income group) where the unconstrained DWL shows the wrong sign. This underscores that DWL analysis is only meaningful if the underlying estimates satisfy the required properties of consumer demand behavior.

One feature of the estimates in Table 5 is the variation in DWL seen across different quantiles. More generally, we can ask how DWL is distributed over the entire population of types. Such an analysis is presented in Figure 4. In this figure we show for each income group the density of DWL across the range of quantiles (from $\alpha = 0.05$ to $\alpha = 0.95$), comparing unconstrained and constrained estimates.

[TABLE 5 AND FIGURE 4 ABOUT HERE]

5 Price Endogeneity

So far we have maintained the assumption of exogeneity on prices. There are many reasons why prices vary at the local market level. These include cost differences on the supply side, short-run supply shocks, local competition, as well as taxes and government regulation (EIA (2010)). However, one may be concerned that prices may also reflect preferences of the consumers in the locality, so that prices faced by consumers may potentially be correlated with unobserved determinants of gasoline demand.

To address this concern, we follow Blundell, Horowitz, and Parey (2012) and use a cost-shifter approach to identify the demand function. An important determinant of prices is the cost of transporting the fuel from the supply source. The U.S. Gulf Coast Region accounts for the majority of total U.S. refinery net production of finished motor gasoline, and for almost two thirds of U.S. crude oil imports. It is also the starting point for most major gasoline pipelines. We therefore expect that transportation cost increases with distance to the Gulf of Mexico, and implement this with the distance between one of the major oil platforms in the Gulf of Mexico and the state capital (see Blundell, Horowitz, and Parey (2012) for further details and references). Figure 5 shows the systematic and positive relationship between log price and distance (in 1,000 km) at state level.

[FIGURE 5 ABOUT HERE]

In the following, we first present evidence from a nonparametric exogeneity test. We then estimate a nonparametric quantile IV specification, incorporating the shape restriction.

5.1 Exogeneity Test

Building on the work for the conditional mean case in Blundell and Horowitz (2007), Fu (2010) develops a nonparametric exogeneity test in a quantile setting. As Blundell and Horowitz (2007), this approach does not require an instrumental variables estimate, and instead tests the exogeneity hypothesis directly. By avoiding the ill-posed inverse problem, it is likely to have substantially better power properties than alternative tests.

To simplify the computation we focus on the univariate version of the test here. For this purpose, we split the overall sample according to household income, and then run the test for each household income group separately.[10] We select income groups to broadly correspond to our three reference income levels in the quantile estimation; we select a low income group of households (household income between \$35,000 and \$50,000), a middle income group of households (household income between \$50,000 and \$65,000), and an upper income group of households (household income between \$65,000 and \$80,000). Given that we perform the test three times (for these three income groups) we can adjust the size for a joint 0.05-level test. Given the independence of the three income samples,

[10] The test makes use of the vector of residuals from the quantile model under the null hypothesis. Even though we implement the test separately for three income groups, we use the residuals from the bivariate model using all observations, so that these residuals correspond to the main (unconstrained) specification of interest (see e.g. Figure 1).

the adjusted p-value for a joint 0.05-level test of exogeneity, at each of the three income groups, is $1 - (0.95)^{(1/3)} = 0.01695$.

Table 6 shows the test results, where column (1) presents our baseline estimates, and columns (2) and (3) show a sensitivity with respect to the bandwidth parameter choice required for the Kernel density estimation. For the median case, the p-values are above 0.1 throughout and thus there is no evidence of a violation of exogeneity at the median. The evidence for the first quartile is similar. The only instance of a borderline p-value is for the lower income group for the upper quartile, with a baseline p-value of 0.02, which is still above the adjusted cutoff value for a test 0.05-level test. Overall, we interpret this evidence as suggesting that we do not find strong evidence of endogeneity in this application. This finding is also consistent with our earlier analysis focusing on the conditional mean (see Blundell, Horowitz, and Parey (2012)). In order to allow a comparison, we nonetheless present quantile IV estimates in the following.

[TABLE 6 ABOUT HERE]

5.2 Quantile Instrumental Variable Estimates

Figure 6 presents our quantile IV estimates under the shape restriction. These estimates are shown as filled markers, and compared with our earlier shape-constrained estimates assuming exogeneity of prices (see Figure 1), shown as open markers.[11] Overall, the shape of the IV estimates is quite similar to those obtained under the assumption of exogeneity. This is consistent with the evidence from the exogeneity test presented above. As before the comparison across income groups suggests that the middle income group is more elastic than the two other income groups, in particular over the lower part of the price range.

[FIGURE 6 ABOUT HERE]

6 Conclusions

The starting point of this analysis are the following two observations: First, when there is heterogeneity in terms of usage intensity, the patterns of demand may potentially be quite different at different points in the distribution of the unobservable heterogeneity. Under suitable exogeneity assumptions and a monotonicity restriction, quantile methods

[11] To simplify the computation of the IV estimates we set the number of interior knots for the cubic splines to 2 in both the income and the price dimension here, and impose the Slutsky constraint at five points in the income dimension ($37,500, $42,500, $57,500, $72,500, and $77,500). We use the NAG routine E04US together with a multi-start procedure to solve the global minimization problem. The resulting demand function estimates do not appear sensitive to specific starting values. In the implementation of the objective function (see equation (6)), we smooth the indicator function corresponding to the term $1\left[W_i - G_\alpha(P_i, Y_i) \leq 0\right]$ in the neighborhood of 0 using a Gaussian Kernel.

allow us to recover the demand function at different points in the distribution of unobservables. This allows us to estimate demand functions for specific types of individuals, rather than averaging across different types of consumers.

Second, we want to be able to allow a flexible effect of price and income on household demand, and in particular allow price responses to differ by income level. Nonparametric estimates eliminate the risk of specification error but can be poorly behaved due to random sampling errors. Fully nonparametric demand estimates can be non-monotonic and may violate consumer theory. In contrast, a researcher choosing a tightly specified model is able to precisely estimate the parameter vector; however simple parametric models of demand functions can be misspecified and, consequently, yield misleading estimates of price sensitivity and DWL. We argue that in the context of demand estimation, this apparent trade-off can be overcome by constraining nonparametric estimates to satisfy the Slutsky condition of economic theory. We have illustrated this approach by estimating a gasoline demand function. The constrained estimates are well-behaved and reveal features not found with typical parametric model specifications. We present estimates across income groups and at different points in the distribution of the unobservables.

These estimates are obtained initially under the assumption of exogenous prices, and the reader may therefore be concerned about potential endogeneity of prices. We investigate this in two ways. First, we implement an exogeneity test to provide direct evidence on this. As instrument, we use a cost shifter variable measuring transportation cost. The results suggest that endogeneity is unlikely to be of first order relevance. Nonetheless, we investigate the shape of the demand function without imposing exogeneity of prices. For this purpose, we develop a novel estimation approach to nonparametric quantile estimation with endogeneity. We estimate IV quantile models under shape restrictions. The results are broadly similar to the estimates under exogeneity.

The analysis showcases the value of imposing shape restrictions in nonparametric quantile regressions. These restrictions provide a way of imposing structure and thus informing the estimates without the need for arbitrary functional form assumptions which have no basis in economic theory.

References

Bento, A. M., L. H. Goulder, M. R. Jacobsen, and R. H. von Haefen (2009): "Distributional and Efficiency Impacts of Increased US Gasoline Taxes," *American Economic Review*, 99(3), pp. 667–699.

Blundell, R. (2005): "How Revealing is Revealed Preference?," *Journal of the European Economic Association*, 3(2-3), 211–235.

Blundell, R., and J. L. Horowitz (2007): "A Non-Parametric Test of Exogeneity," *Review of Economic Studies*, 74(4), 1035–1058.

Blundell, R., J. L. Horowitz, and M. Parey (2012): "Measuring the price responsiveness of gasoline demand: Economic shape restrictions and nonparametric demand estimation," *Quantitative Economics*, 3(1), 29–51.

Blundell, R., D. Kristensen, and R. Matzkin (2011): "Bounding Quantile Demand Functions Using Revealed Preference Inequalities," *CeMMAP working paper*, CWP21/11.

Brown, B. W., and M. B. Walker (1989): "The Random Utility Hypothesis and Inference in Demand Systems," *Econometrica*, 57(4), 815–829.

Chen, C. (2007): "A Finite Smoothing Algorithm for Quantile Regression," *Journal of Computational and Graphical Statistics*, 16(1), pp. 136–164.

Deaton, A., and J. Muellbauer (1980): *Economics and consumer behavior*. Cambridge University Press, reprinted 1999.

EIA (2010): "Regional gasoline price differences," Energy Information Administration, U.S. Department of Energy. Available at http://www.eia.gov/energyexplained/index.cfm?page=gasoline_regional.

Frondel, M., N. Ritter, and C. Vance (2012): "Heterogeneity in the rebound effect: Further evidence for Germany," *Energy Economics*, 34(2), 461–467.

Fu, J.-Y. M. (2010): "Testing Exogeneity in Nonparametric Instrumental Variables Models Identified by Conditional Quantile Restrictions," Northwestern University, November 2010.

Hall, P., and L.-S. Huang (2001): "Nonparametric Kernel Regression Subject to Monotonicity Constraints," *The Annals of Statistics*, 29(3), 624–647.

Hausman, J. A., and W. K. Newey (1995): "Nonparametric Estimation of Exact Consumers Surplus and Deadweight Loss," *Econometrica*, 63(6), 1445–1476.

——— (2013): "Individual Heterogeneity and Average Welfare," *cemmap working paper*, CWP34/13.

Hoderlein, S., and J. Stoye (2013): "Revealed Preferences in a Heterogeneous Population," Forthcoming in the Review of Economics and Statistics.

Hoderlein, S., and A. Vanhems (2011): "Welfare Analysis using Nonseparable Models," *cemmap working paper*, CWP01/11.

Horowitz, J. L., and V. G. Spokoiny (2002): "An Adaptive, Rate-Optimal Test of Linearity for Median Regression Models," *Journal of the American Statistical Association*, 97(459), 822–835.

Hurwicz, L., and H. Uzawa (1971): "On the Problem of Integrability of Demand Functions," in *Preferences, Utility, and Demand. A Minnesota Symposium*, ed. by J. Chipman, L. Hurwicz, M. K. Richter, and H. F. Sonnenschein. Harcourt Brace Jovanovich.

Imbens, G. W., and W. K. Newey (2009): "Identification and Estimation of Triangular Simultaneous Equations Models Without Additivity," *Econometrica*, 77(5), 1481–1512.

Jones, M. C. (1993): "Simple boundary correction for density estimation," *Statistics and Computing*, 3, 135–146.

Koenker, R. (2005): *Quantile Regression*, Econometric Society Monographs. Cambridge University Press.

Lewbel, A. (1997): "Consumer Demand Systems and Household Equivalence Scales," in *Handbook of Applied Econometrics*, ed. by M. H. Pesaran, and P. Schmidt, vol. II: Microeconomics. Blackwell Publishers.

——— (2001): "Demand Systems with and without Errors," *The American Economic Review*, 91(3), pp. 611–618.

Manning, W. G., L. Blumberg, and L. H. Moulton (1995): "The demand for alcohol: The differential response to price," *Journal of Health Economics*, 14(2), 123–148.

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995): *Microeconomic Theory.* Oxford University Press.

Matzkin, R. L. (2003): "Nonparametric Estimation of Nonadditive Random Functions," *Econometrica*, 71(5), 1339–1375.

——— (2008): "Identification in Nonparametric Simultaneous Equations Models," *Econometrica*, 76(5), 945–978.

ORNL (2004): "2001 National Household Travel Survey. User Guide," Oak Ridge National Laboratory. Available at `http://nhts.ornl.gov/2001/`.

Powell, M. J. D. (1981): *Approximation Theory and Methods.* Cambridge University Press.

Schmalensee, R., and T. M. Stoker (1999): “Household Gasoline Demand in the United States,” *Econometrica*, 67(3), 645–662.

Silverman, B. W. (1986): *Density Estimation for Statistics and Data Analysis*. Chapman and Hall.

Stern, N. (1986): “On the specification of labour supply functions,” in *Unemployment, search and labour supply*, ed. by R. Blundell, and I. Walker, chap. 9, pp. 143–189. Cambridge University Press.

West, S. E. (2004): “Distributional effects of alternative vehicle pollution control policies,” *Journal of Public Economics*, 88, 735–757.

Yatchew, A., and J. A. No (2001): “Household Gasoline Demand in Canada,” *Econometrica*, 69(6), 1697–1709.

Table 1: Sample descriptives

	Mean	*St. dev.*
Log gasoline demand	7.127	0.646
Log price	0.286	0.057
Log income	11.054	0.580
Observations	3640	

Note: See text for details.

Table 2: Cross-validation results

quantile (α)	number interior knots price	income
	(1) Base case	
0.25	1	3
0.50	4	3
0.75	3	3
	(2) Leaving out largest 10 and lowest 10 share observations	
0.25	1	3
0.50	4	4
0.75	1	3

Note: Table shows cross-validation results by quantile.

Table 3: Log-log model estimates

	$\alpha = 0.25$ (1)	$\alpha = 0.50$ (2)	$\alpha = 0.75$ (3)	OLS (4)
$\log(p)$	-1.00	-0.72	-0.60	-0.83
	[0.23]	[0.19]	[0.22]	[0.18]
$\log(y)$	0.41	0.33	0.23	0.34
	[0.02]	[0.02]	[0.02]	[0.02]
Constant	2.58	3.74	5.15	3.62
	[0.27]	[0.21]	[0.26]	[0.20]
N	3640	3640	3640	3640

Note: Dependent variable is log gasoline demand. See text for details.

Table 4: Specification test

Dependent var.	test statistic	critical value		p-value	reject?
		0.05 level	0.01 level		
gasoline share	2.52	1.88	2.69	0.0120	yes
log quantity	2.71	1.82	2.43	0.0020	yes

Note: Test implements Horowitz and Spokoiny (2002) for the median case. The first row reports the test results for gasoline demand measured as budget share, the second row for log quantity. Under the null hypothesis, the model is linear in log price and log income. See text for details.

Figure 1: Quantile regression estimates: constrained versus unconstrained estimates

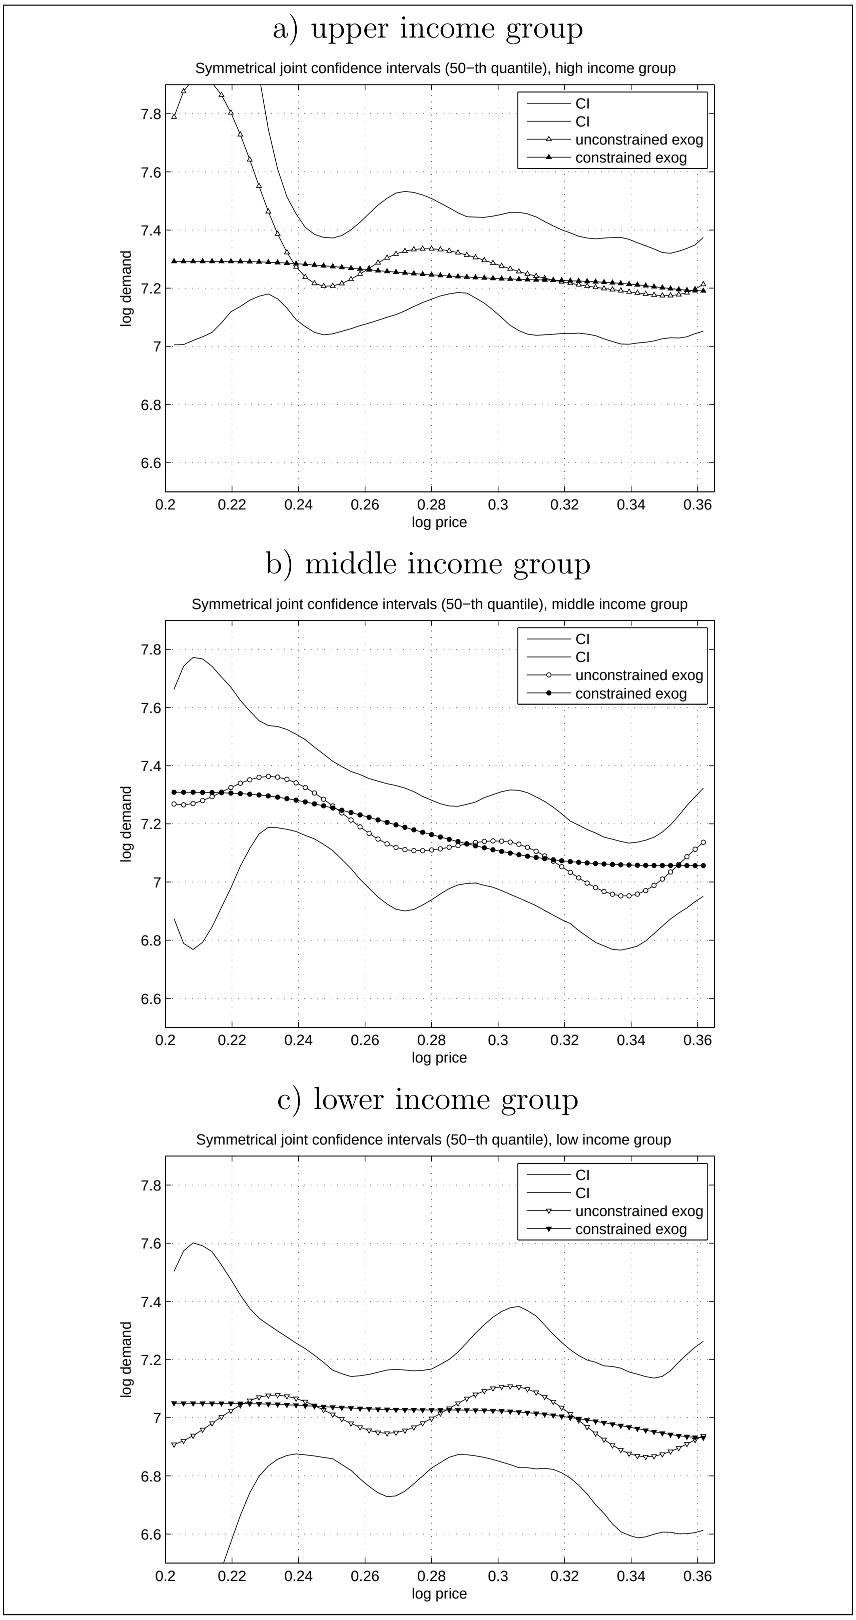

Note: Figure shows unconstrained nonparametric quantile demand estimates (open markers) and constrained nonparametric demand estimates (filled markers) at different points in the income distribution for the median ($\alpha = 0.5$), together with simultaneous confidence intervals. Income groups correspond to \$72,500, \$57,500, and \$42,500. Confidence intervals shown refer to bootstrapped symmetrical, simultaneous confidence intervals with a confidence level of 90%, based on 4,999 replications. See text for details.

Figure 2: Quantile regression estimates: constrained versus unconstrained estimates (middle income group)

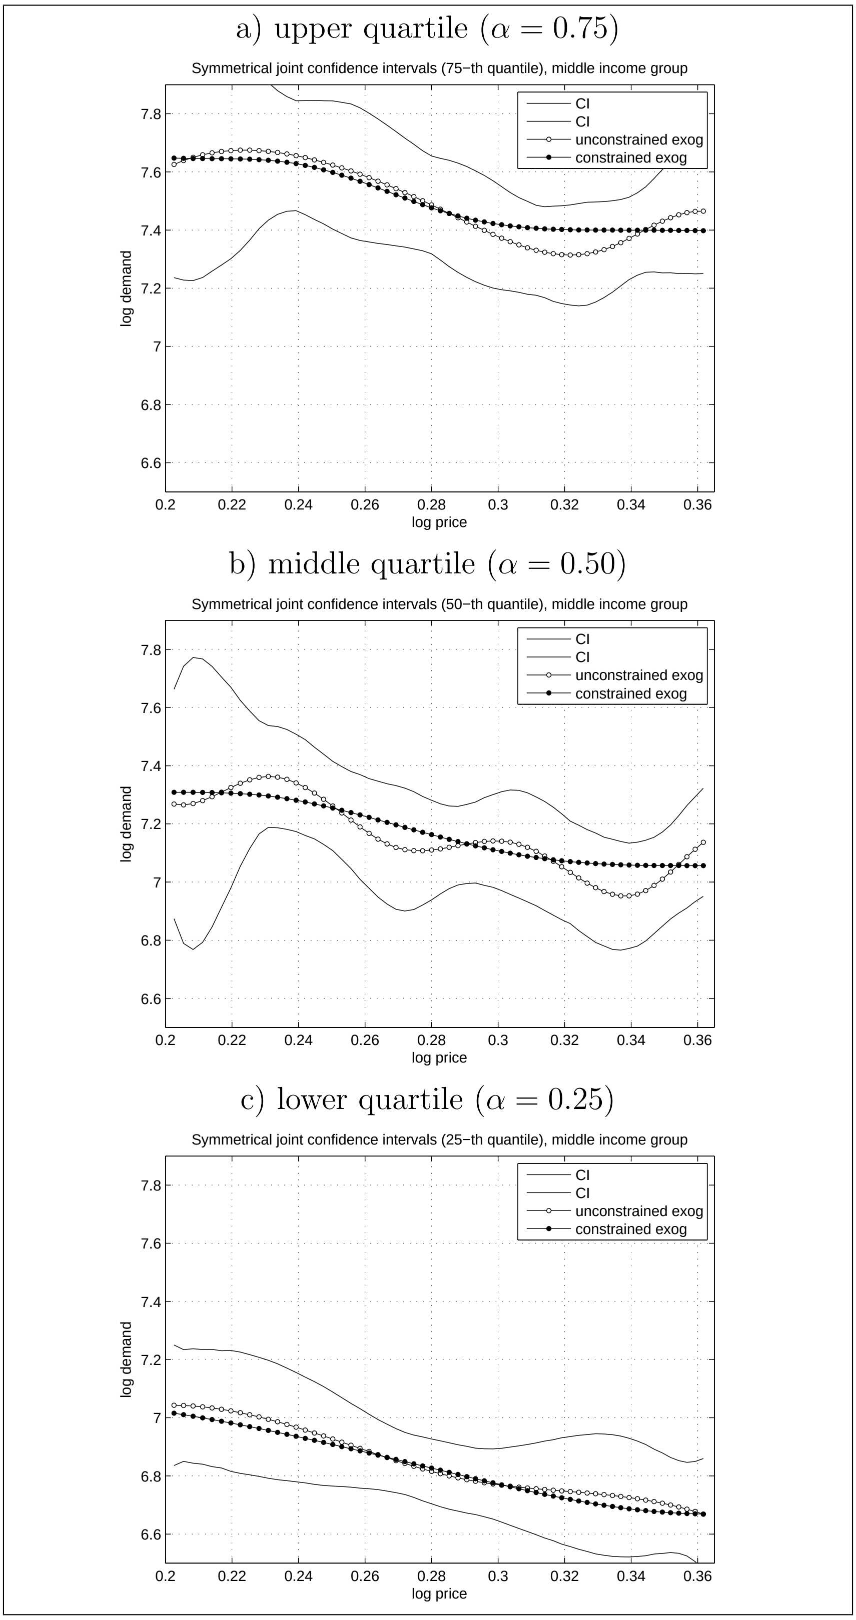

Note: Figure shows unconstrained nonparametric quantile demand estimates (filled markers) and constrained nonparametric demand estimates (filled markers) at the quartiles for the middle income group ($57,500), together with simultaneous confidence intervals. Confidence intervals shown refer to bootstrapped symmetrical, simultaneous confidence intervals with a confidence level of 90%, based on 4,999 replications. See text for details.

Figure 3: Mean regression estimates: constrained versus unconstrained estimates

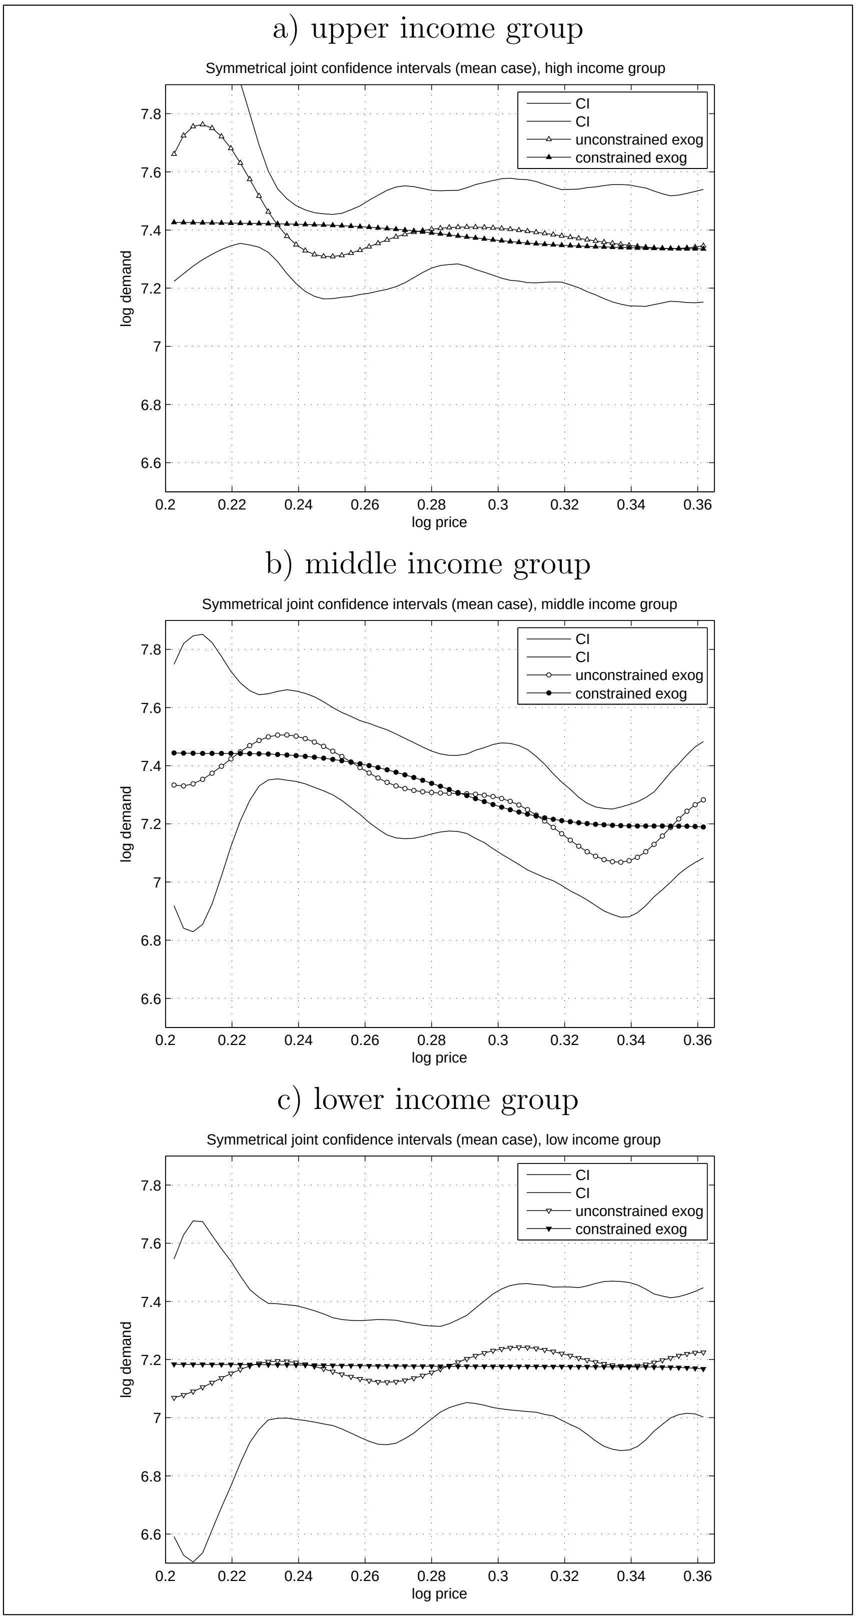

Note: Figure shows unconstrained nonparametric mean regression demand estimates (filled markers) and constrained nonparametric demand estimates (filled markers) at different points in the income distribution, together with simultaneous confidence intervals. Income groups correspond to \$72,500, \$57,500, and \$42,500. Confidence intervals shown refer to bootstrapped symmetrical, simultaneous confidence intervals with a confidence level of 90%, based on 4,999 replications. See text for details.

Table 5: DWL estimates

	unconstrained			constrained			linear quantile estimates		
income	DWL	DWL/tax	DWL/inc*	DWL	DWL/tax	DWL/inc*	DWL	DWL/tax	DWL/inc*
lower quartile ($\alpha = 0.25$)									
72500	11.76	5.72%	1.62	12.74	6.21%	1.76	13.89	7.12%	1.92
57500	33.24	20.01%	5.78	29.18	17.54%	5.08	12.88	7.24%	2.24
42500	-15.40	-8.91%	-3.62	0.85	0.54%	0.20	11.30	7.35%	2.66
median ($\alpha = 0.50$)									
72500	49.64	17.30%	6.85	16.32	5.81%	2.25	20.33	7.26%	2.80
57500	5.86	2.20%	1.02	30.20	12.30%	5.25	19.06	7.36%	3.32
42500	12.81	5.87%	3.01	18.57	8.56%	4.37	16.90	7.45%	3.98
upper quartile ($\alpha = 0.75$)									
72500	23.07	5.71%	3.18	20.64	5.07%	2.85	19.29	4.76%	2.66
57500	15.98	4.35%	2.78	39.40	11.42%	6.85	19.77	5.22%	3.44
42500	-43.60	-11.25%	-10.26	1.17	0.35%	0.28	18.86	5.63%	4.44

Note: Table shows DWL estimates, corresponding to a change in prices from the 5th to the 95th percentile, that is from $1.225 to $1.436. For comparability all three sets of estimates are based on the sample for the nonparametric analysis, and use budget share as dependent variable. * DWL per income figures are rescaled by factor 10^4 for better readibility.

Figure 4: Distribution of DWL, constrained versus unconstrained

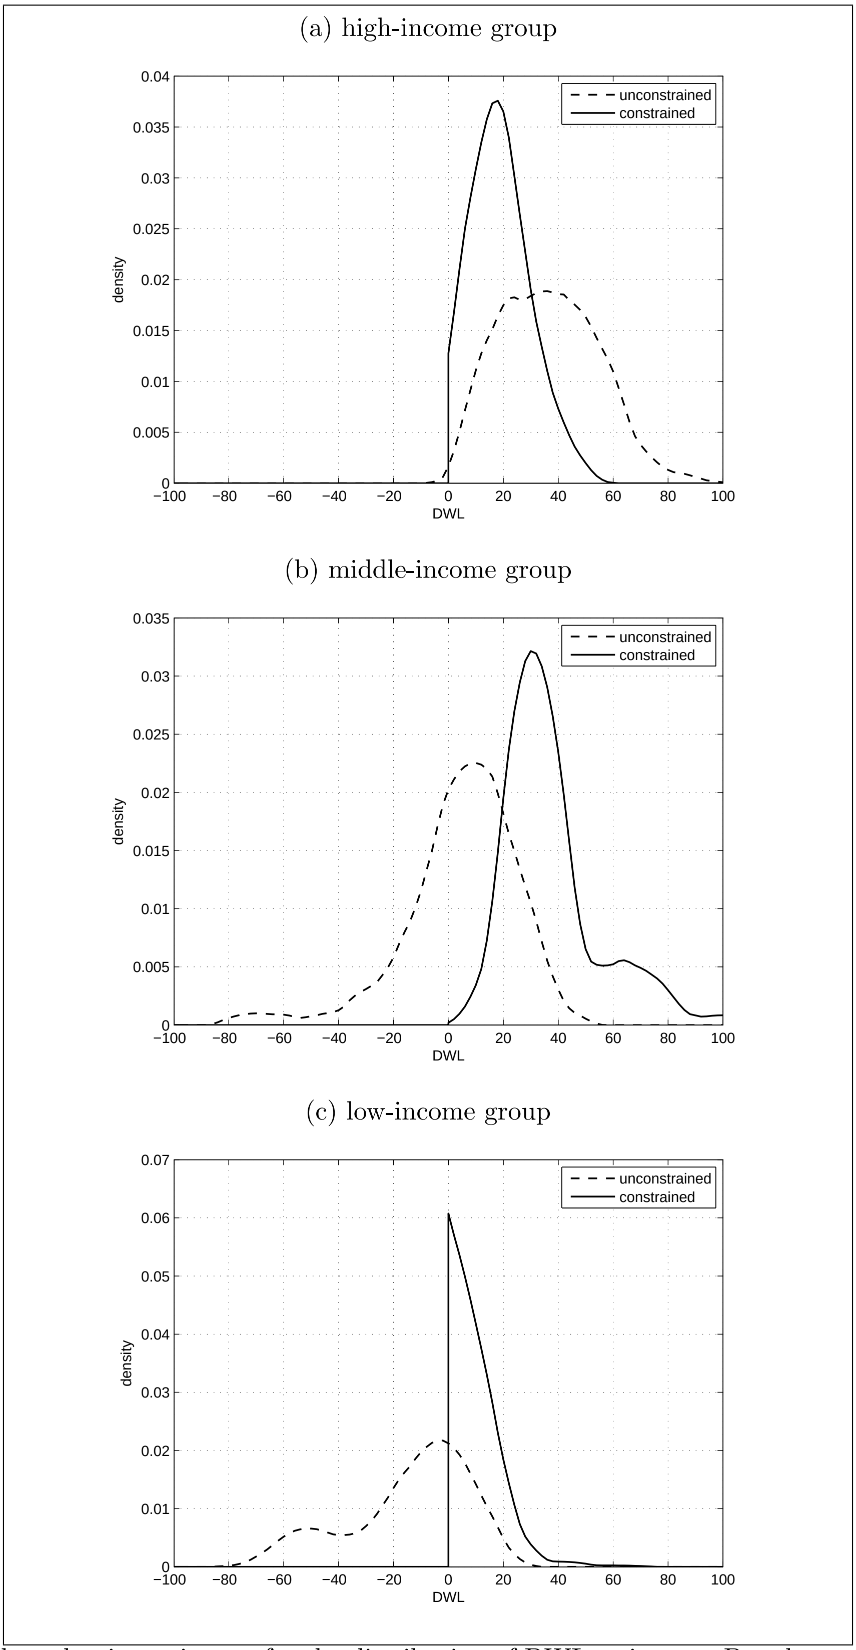

Note: Graphs show density estimates for the distribution of DWL estimates. Based on estimates for the 5th to the 95th percentile ($\alpha = 0.05$ to 0.95 in steps of 0.005). Density estimates computed using an Epanechnikov Kernel. Since DWL is nonnegative in the constrained case, density is renormalized in the boundary area (Jones (1993)). Estimates computed using the same knot choice throughout as crossvalidated for the median.

Figure 5: The Instrument Variable for Price: Distance to the Gulf of Mexico

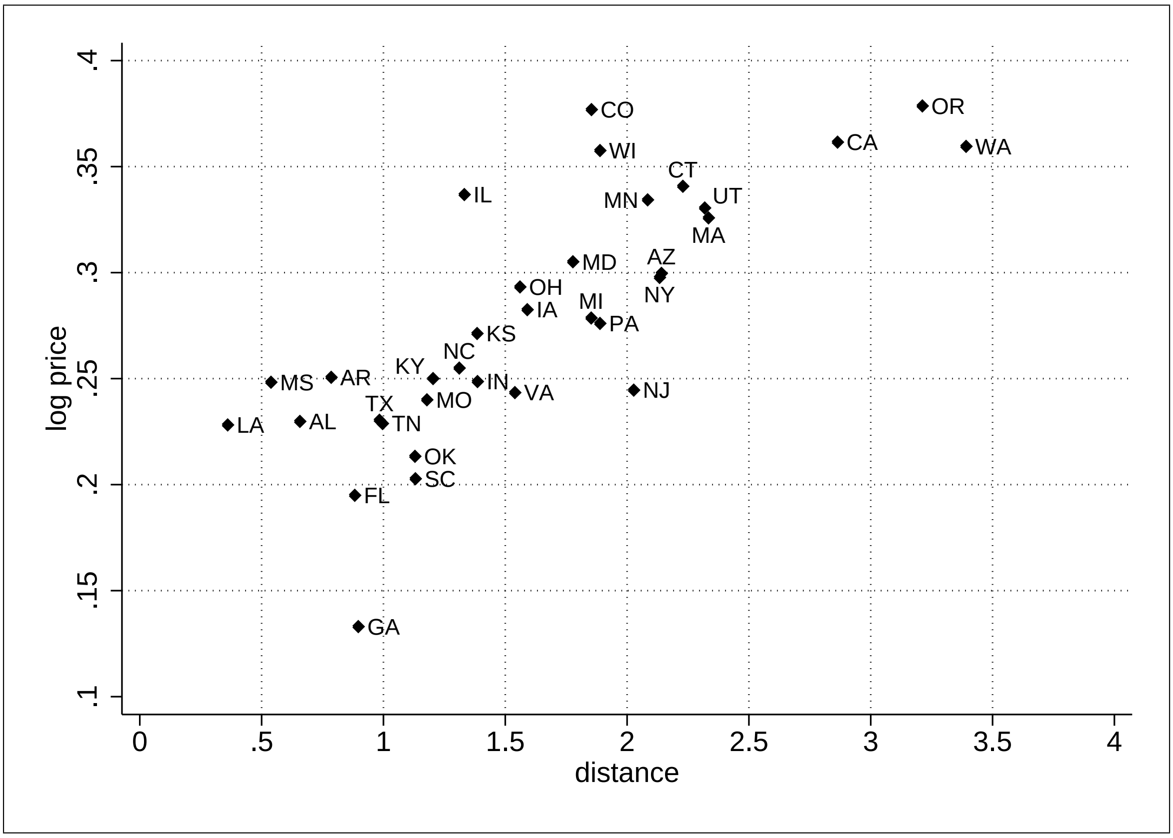

Source: BHP (2012, Figure 5).

Table 6: Exogeneity test (p-values)

	income range	base case	bandwidth sensitivity factor 0.8	bandwidth sensitivity factor 1.25
		(1)	(2)	(3)
	low	0.382	0.350	0.450
first quartile ($\alpha = 0.25$)	middle	0.212	0.196	0.195
	high	0.361	0.290	0.470
	low	0.331	0.224	0.453
median ($\alpha = 0.50$)	middle	0.140	0.171	0.126
	high	0.685	0.643	0.734
	low	0.020	0.024	0.015
third quartile ($\alpha = 0.75$)	middle	0.710	0.825	0.585
	high	0.818	0.848	0.817

Note: Table shows p-values for the exogeneity test from Fu (2010). Endogenous variable is price, instrumented with distance. We run separate tests for three income groups; for this test, these groups are defined as follows: 'low': income between \$35,000 and \$50,000, 'middle': \$50,000 – \$65,000, 'high': \$65,000 – \$80,000. The specification we test is the unconstrained nonparametric quantile estimate as shown e.g. in Figure 1 for the median. In implementing this test, required bandwidth choices for the Kernel density estimates use Silverman's rule of thumb. Columns (2) and (3) vary all bandwidth inputs by the indicated factor.

Figure 6: Quantile regression estimates under the shape restriction: IV estimates versus estimates assuming exogeneity

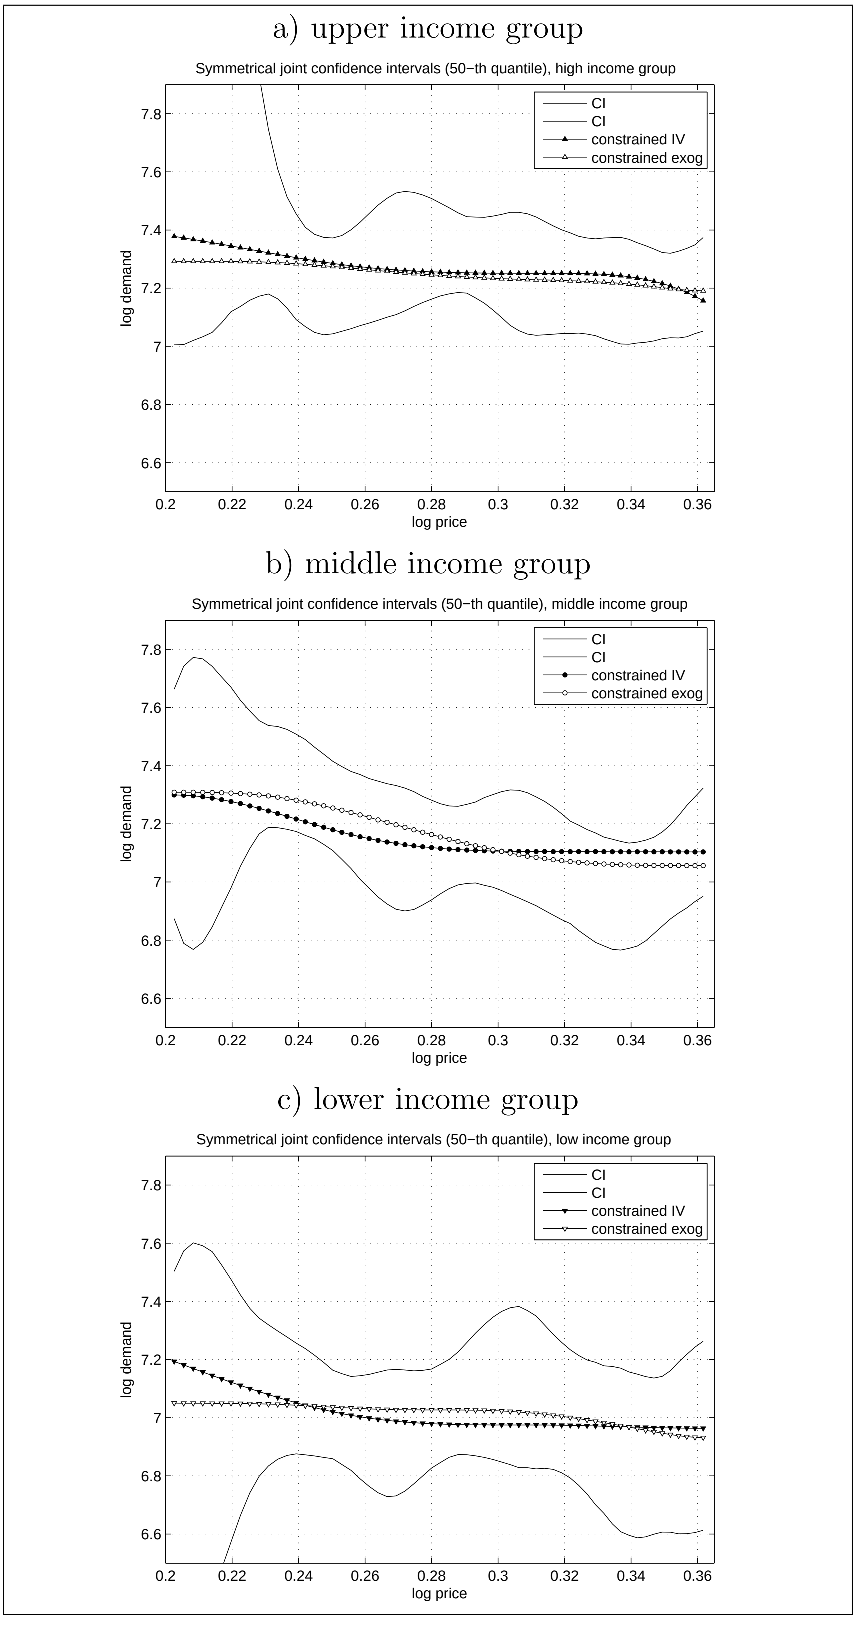

Note: Figure shows constrained nonparametric IV quantile demand estimates (filled markers) and constrained quantile demand estimates under exogeneity (open markers) at different points in the income distribution for the median ($\alpha = 0.5$), together with simultaneous confidence intervals. Income groups correspond to \$72,500, \$57,500, and \$42,500. Confidence intervals shown correspond to the unconstrained quantile estimates under exogeneity as in Figure 1. See text for details.

A Appendix

Figure A.1: Sensitivity of constrained quantile regression estimates to smoothing parameter

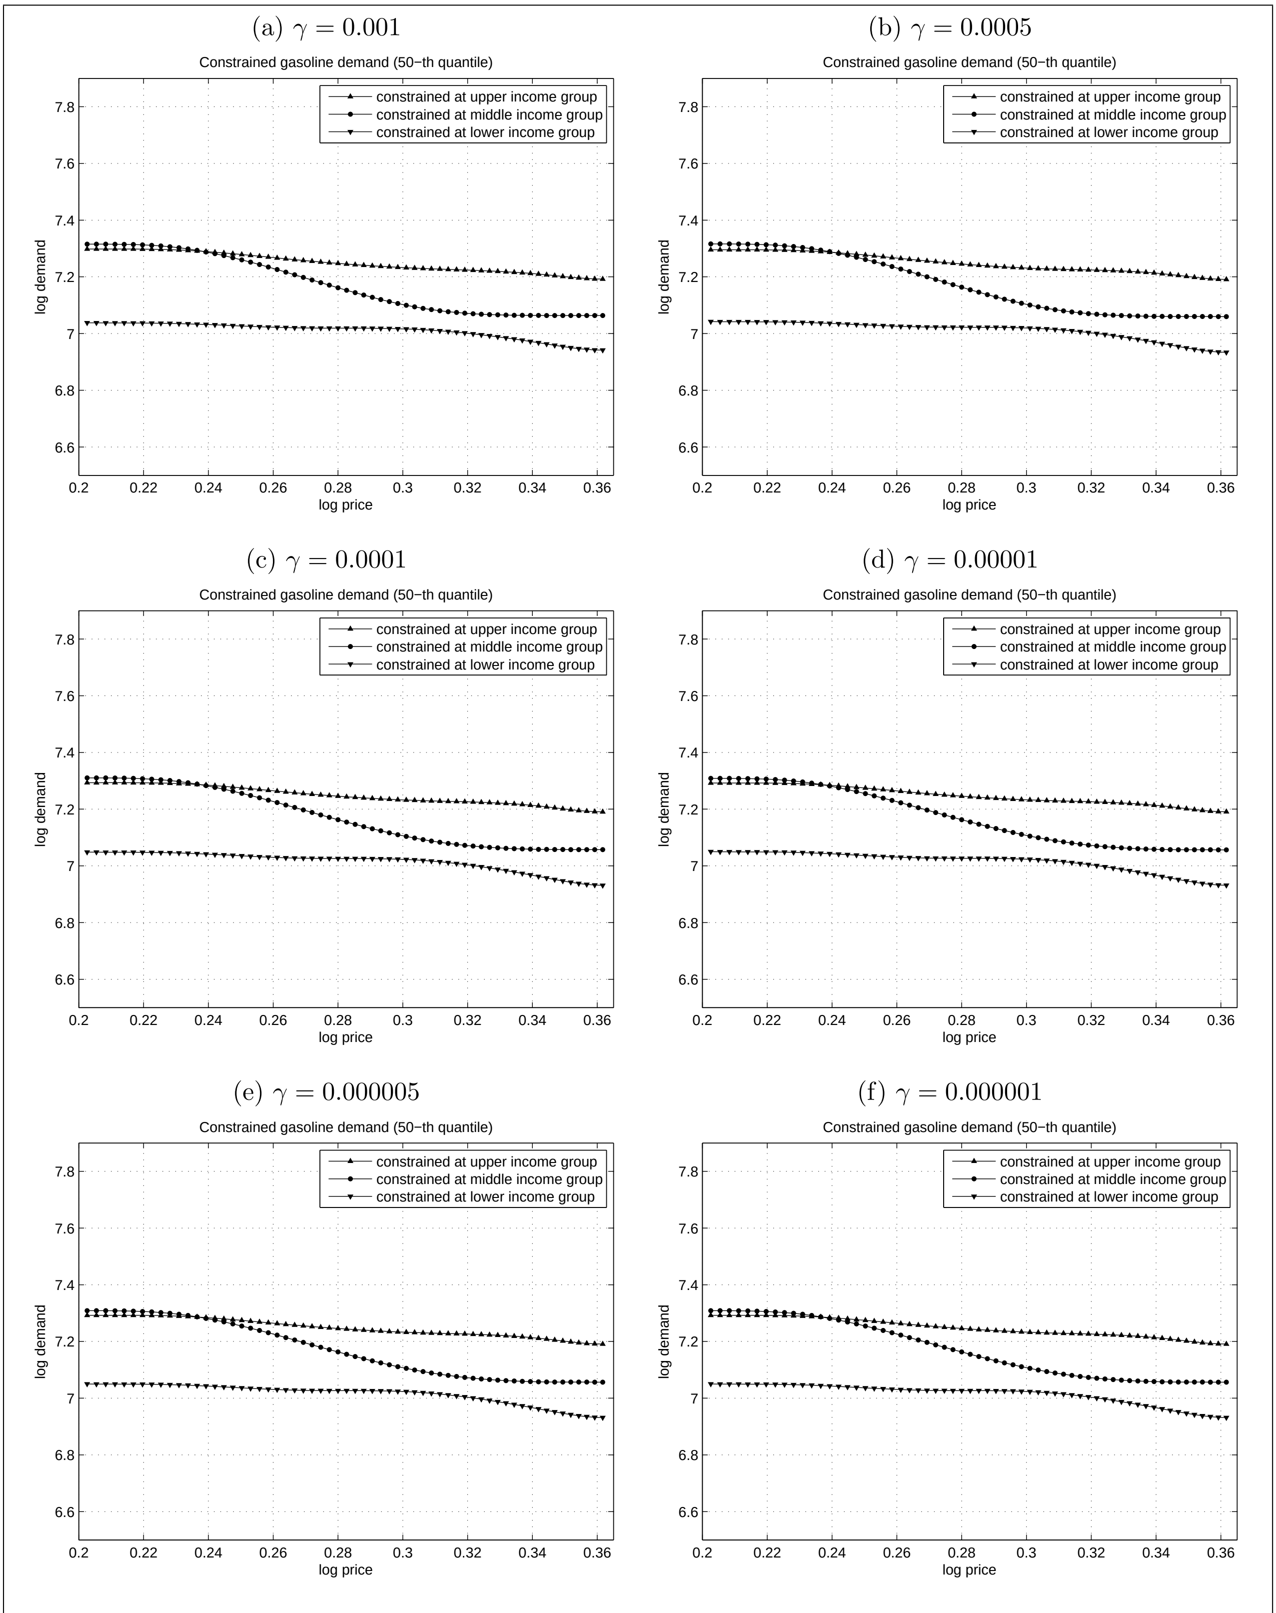

Note: In the computation of the constrained quantile estimates (see Section 3.2), the check function is smoothed in a small neighborhood around 0, using a quadratic approximation over the range $[-(1-\alpha)\gamma;\ \alpha\gamma]$ (see Chen (2007)), where γ is a bandwidth parameter. This figure shows the constrained quantile regression estimates for the median ($\alpha = 0.5$), resulting from alternative choices of γ. The figures presented in the main text correspond to panel (f) of Figure A.1.