Journal of Econometrics 21 (1983) 263–285. North-Holland Publishing Company

ESTIMATION OF CONSUMER DEMAND SYSTEMS WITH BINDING NON-NEGATIVITY CONSTRAINTS*

T.J. WALES and A.D. WOODLAND

University of British Columbia, Vancouver V6T 1Y2, Canada

Received November 1980, final version received February 1982

1. Introduction

Although the economic and statistical theory underlying the estimation of systems of demand functions is well developed, very little attention has been paid to the problems which arise when the sample contains a significant proportion of observations in which expenditure on one or more goods is zero. For such a sample the econometric model should allow for zero expenditures to occur with positive probability. However, the econometric model used in most studies assumes that expenditures (or shares) follow a joint normal distribution and this does not allow for a positive probability of zero expenditures.[1] Standard estimation methods for this model, such as Zellner's two-stage estimator for seemingly unrelated regressions and the maximum likelihood estimator, do not take special account of zero expenditures, and consequently yield inconsistent estimates of the parameters. Indeed even if every observation containing zero expenditures on one or more goods was excluded for purposes of estimation, these standard estimators would be biased and inconsistent.[2] Moreover, excluding these observations might significantly reduce the sample size. Regardless of whether or not the complete sample is used, the bias and inconsistency occur because the random disturbances have expectations which are not zero and which depend upon the exogenous variables.

*This paper is a synthesis and extension of two previous papers (Department of Economics discussion papers 78–32 and 79–32). We are indebted to Dorwin Cho, Mukesh Eswaran and David Ryan for exceptional research assistance and to the Canada Council and the British Columbia Department of Labour for financial support. We are grateful to the Bureau of Agricultural Economics, Canberra for providing the data.

[1]See, for example, the recent survey of the systems approach to consumer demand by Barten (1977).

[2]Of course one could obtain the maximum likelihood estimator conditional on positive expenditures on all goods. This estimator is consistent and asymptotically efficient. We do not consider it here since we wish to use an estimator based on all observations.

In the paper we formulate two alternative econometric models of consumer demand which take into account the possibility that expenditures on one or more goods are zero for a significant proportion of the sample. The first model is based upon the Kuhn–Tucker (1951) conditions for the maximization of a utility function subject to the budget constraint, and assumes that preferences are random over the population. The second model assumes that preferences are non-random and is an extension of a limited dependent variable model of the type developed by Tobin (1958) for the case of a single equation, and by Amemiya (1974) for a set of equations. It represents a non-trivial extension of this approach in that it takes account of the restrictions imposed by the budget constraint. Maximization of the likelihood functions for these models yields parameter estimates which are consistent, and asymptotically efficient and normally distributed.

Our Kuhn–Tucker model is developed in the following section, which concludes with a discussion of a special case in which the utility function is quadratic. Our Amemiya–Tobin type model is discussed in section 3, which concludes with a discussion of the Stone–Geary utility function as a special case. In section 4 we report on the application of our models to meat consumption in Australia. Section 5 contains some concluding comments.

2. The Kuhn–Tucker approach

The traditional theory of consumer demand assumes that the individual maximizes a utility function $G(x)$ over the set of non-negative quantities $x=(x_1,\dots,x_M)$ which satisfy the budget constraint $v^T x \leqq 1$, where $v=(v_1,\dots,v_M)^T>0$ is the vector of normalized prices (that is, $v_i=p_i/m$ where p_i is the price of good i and m is income). Formally, this problem is

$$H(v) \equiv \max_x \{G(x): v^T x \leqq 1,\ x \geqq 0\}. \tag{1}$$

For present purposes it is assumed that $G(x)$ is a continuously differentiable, quasi-concave, increasing function.[3] Following the results of Arrow and Enthoven (1961), the necessary and sufficient Kuhn–Tucker conditions for a solution are

$$\begin{aligned} &G_i(x) - \lambda v_i \leqq 0 \leqq x_i, \qquad i=1,\dots,M, \\ &v^T x - 1 \leqq 0 \leqq \lambda, \end{aligned} \tag{2}$$

where λ is the Lagrange multiplier associated with the budget constraint. Whenever conditions such as (2) appear (with a double inequality) they are

[3]By increasing we mean that $x^1 > x^0$ implies that $G(x^1) > G(x^0)$.

taken to imply tha the two terms which have the inequality constraints, when multiplied together yield zero. This convention is used throughout the paper.

Since $G(x)$ is assumed to be increasing the consumer will spend all his income. Consequently, λ will be positive (almost everywhere) and at least one good will be consumed which, without loss of generality, is assumed to be the first good. Thus (2) implies that $G_1(x)=\lambda v_1$ or $\lambda=G_1(x)/v_1$. Using these results the necessary and sufficient conditions for utility maximization may be rewritten as

$$v_1 G_i(x) - v_i G_1(x) \leqq 0 \leqq x_i, \qquad i=2,\dots,M, \tag{3}$$

$$v^T x = 1.$$

Conditions (3) may be interpreted as follows. If $x_i>0$ then $v_1 G_i(x) - v_i G_1(x) = 0$, that is $G_i(x)/G_1(x)=v_i/v_1$, meaning that the marginal rate of substitution between goods i and 1 along the indifference curve at the solution is equal to the price ratio (slope of budget constraint). On the other hand, the marginal rate of substitution between goods i and 1 may (and generally will) be less than the price ratio if, and only if, good i is not consumed.

To allow for individual differences in tastes, we assume that preferences are randomly distributed over the population. A particularly convenient assumption is to let marginal utilities consist of a deterministic ($\bar{G}_i(x)$) and a random (u_i) component as

$$G_i(x, u_i) = \bar{G}_i(x) + u_i, \qquad i=1,\dots,M, \tag{4}$$

which are consistent with the random utility function $G(x,u) \equiv \bar{G}(x) + u^T x$ where $u=(u_1,\dots,u_M)$. Replacing $G_i(x)$ in (3) by $G_i(x,u_i)$ defined in (4) we obtain

$$(v_1 u_i - v_i u_1) + [v_1 \bar{G}_i(x) - v_i \bar{G}_1(x)] \leqq 0 \leqq x_i, \qquad i=2,\dots,M, \tag{5}$$

$$v^T x = 1.$$

From the individual's point of view the utility function $G(x,u)$ is non-stochastic since his u vector is known to him, thus he obtains his optimal consumption vector by solving (5) for x. However, from the researcher's point of view the vector u, and hence x, for each individual are random drawings from a population. By specifying the distribution for the u vectors we can calculate the distribution of the consumption vector x from (5).

We assume that u has a joint normal distribution with zero means and constant covariance matrix Σ. Since the left-hand sides of (5) are linear in u it is convenient to define $y_i \equiv v_1 u_i - v_i u_1$, $i=2,\dots,M$, which also follow a joint normal distribution with zero means and (non-constant) covariance matrix Ω. Recognizing that the budget constraint may be used to eliminate one of the elements of x, say x_1, and defining $\bar{y}_i(\hat{x}) \equiv v_i \bar{G}_1(\hat{x}) - v_1 \bar{G}_i(\hat{x})$ where $\hat{x} = (x_2, \ldots, x_M)$, (5) may be rewritten as[4]

$$y_i - \bar{y}_i(\hat{x}) \leqq 0 \leqq x_i, \qquad i = 2, \ldots, M. \tag{6}$$

For the case where all M goods are consumed, (6) implies that $y_i = \bar{y}_i(\hat{x})$, $i = 2, \ldots, M$, and hence the density function for $\hat{x}$ is obtained as

$$f(\hat{x}) = n(\hat{y}, \Omega) \operatorname{abs}[J(\hat{x})], \tag{7}$$

where $\hat{y} = (y_2, \ldots, y_M)$, $n(\hat{y}, \Omega)$ is the normal density function for $\hat{y}$ which has mean zero and covariance matrix Ω, and $J(\hat{x})$ is the Jacobian of the transformation from $\hat{y}$ to $\hat{x}$. For the case where only one good is consumed, which as will be recalled is the first good, all $M-1$ conditions in (6) are inequalities. Thus the probability of the event $\hat{x} = 0$ is

$$f(0) = \int_{-\infty}^{\bar{y}_M} \cdots \int_{-\infty}^{\bar{y}_2} n(\hat{y}, \Omega) \, dy_2 \ldots dy_M. \tag{8}$$

In general, if the number of goods consumed is K (taken to be the first K goods), the density is

$$f(x_2, \ldots, x_K, 0, \ldots, 0) = \int_{-\infty}^{\bar{y}_M} \cdots \int_{-\infty}^{\bar{y}_{K+1}} n(\bar{y}_2, \ldots, \bar{y}_K, y_{K+1}, \ldots, y_M, \Omega)$$

$$\times \operatorname{abs}[J_K(\hat{x})] \, dy_{K+1}, \ldots, dy_M, \tag{9}$$

where $J_K(\hat{x})$ is the Jacobian of the transformation from $(y_2, \ldots, y_K)$ to $(x_2, \ldots, x_K)$ when $(x_{K+1}, \ldots, x_M) = 0$. Of course, if $K = M$ then (9) reduces to (7) and if $K = 1$ it reduces to (8). It should be noted that eq. (9) is valid only if the transformation from $(y_2, \ldots, y_K)$ to $(x_2, \ldots, x_K)$ when $(x_{K+1}, \ldots, x_M) = 0$ is one-to-one, that is $J_K(\hat{x})$ does not change sign over the region of integration. If the transformation is not one-to-one, then the integration must be performed separately over the sub-regions over which the transformation is one-to-one.

There are $M!/K!(M-K)!$ possible consumption patterns consisting of K positive quantities and $M-K$ zero quantities. Thus there are $\sum_{K=1}^{M} M!/K!(M-K)!$ possible consumption patterns and corresponding density functions which make up the complete density function for $\hat{x}$. For example, when

[4] The choice of the variable to eliminate is arbitrary, the first being chosen for notational simplicity only.

$K=M$ there is just one consumption pattern with density given by (7). When $K=1$ there are M possible consumption patterns each having density (actually probability) of the form given by (8). Since eqs. (3)–(9) are based upon the assumption that $x_1>0$, the densities for the cases where $x_1=0$ are obtained by renumbering the goods so that some other good, whose consumption is positive, becomes good 1.

Given a random sample of N observations on $\hat{x}$, each observation corresponds to just one of the $\sum_{K=1}^{M} M!/K!(M-K)!$ components of the complete density function. The sample likelihood function may then be expressed as

$$L(\hat{x}_1,\ldots,\hat{x}_N)=\prod_{i=1}^{N} f(\hat{x}_i), \tag{10}$$

where $\hat{x}_i$ is the ith observation on $\hat{x}$. Given a functional form for the utility function $G(x,u)$, the parameters of this utility function and the covariance matrix Σ can be estimated by maximizing the likelihood function (10). These maximum likelihood estimates will be consistent, and asymptotically efficient and normally distributed. While the likelihood function involves the evaluation of multiple integrals under the multivariate normal density function, these calculations are feasible for small M using the approximation formulae of Dutt (1976).

A functional form that is particularly convenient for the empirical implementation of the model is the quadratic utility function that has been discussed by Allen and Bowley (1935), Houthakker (1961), Wegge (1968), and others. In the following sections we consider this functional form and, for the case where $M=3$, present empirical results. The quadratic utility function is general in the sense that it is a 'flexible functional form'. Like other flexible functional forms it is not globally quasiconcave and non-decreasing, but this is not a serious disadvantage since it can satisfy these conditions over a subset of the commodity space, which is all that is required in empirical applications. As will be discussed below the Engel curves are piece-wise linear, but we feel this restriction is outweighed by the simplicity in empirical implementation. An alternative more commonly used (but also more restrictive) functional form that could be readily implemented in a similar fashion is the Stone–Geary utility function.

Special case: The quadratic utility function

The quadratic utility function with random components added to the coefficients of the linear terms is

$$G(x,u)=a_{00}+\sum_{i=1}^{M}(a_{i0}+u_i)x_i+\tfrac{1}{2}\sum_{i=1}^{M}\sum_{j=1}^{M}a_{ij}x_ix_j. \tag{11}$$

For this functional form the Kuhn–Tucker conditions corresponding to (6) are

$$y_i - \left[b_{i0} + \sum_{j=2}^{M} b_{ij} x_j \right] \leqq 0 \leqq x_i, \qquad i=2,\dots,M, \tag{12}$$

where

$$b_{i0} \equiv (v_1 a_{i0} - v_i a_{10}) + (a_{i1} - a_{11} v_i/v_1), \qquad i,j=2,\dots,M, \tag{13}$$

$$b_{ij} \equiv (v_1 a_{ij} - v_i a_{1j}) - (a_{i1} - a_{11} v_i/v_1) v_j,$$

and $y_i \equiv v_1 u_i - v_i u_1$ as before.

It is evident from (11) that our model can be interpreted as a type of random coefficients model in which the coefficients $a_{i0}+u_i$ of the x_i consist of a deterministic part a_{i0} and a random component u_i.[5] Since u is assumed to have an M-variate normal distribution with constant covariance matrix Σ, the coefficients of x_i in the utility function have means a_{i0} and constant covariance matrix Σ. For the purpose of identification, a normalization of the parameters of the utility function is required. A convenient normalization is to require that $\sum_{i=1}^{M} (a_{i0}+u_i)=1$ which, since $\mathrm{E}u=0$, implies $\sum_{i=1}^{M} a_{i0}=1$ and $\sum_{i=1}^{M} u_i=0$. Hence the parametric restrictions implied by our normalization are

$$\sum_{i=1}^{M} a_{i0}=1 \quad \text{and} \quad 1^T\Sigma=0, \tag{14}$$

where $1^T=(1,\dots,1)$. Finally, we note that the y_i's have a multivariate normal distribution with means zero and covariances given by

$$\omega_{ij} \equiv \mathrm{E}(y_i y_j) = v_1^2\sigma_{ij} + v_i v_j \sigma_{11} - v_1 v_j \sigma_{1i} - v_1 v_i \sigma_{1j}, \qquad i,j=2,\dots,M. \tag{15}$$

The empirical results presented below are for the case where $M=3$, which we now consider in more detail. There are seven possible consumption patterns, one with all x_i's positive, three with one x_i positive, and three with two x_i's positive. The density for an observation with all x_i's positive is given by (7) where $y_i = b_{i0} + b_{i2}x_2 + b_{i3}x_3 \equiv \bar{y}_i(\hat{x})$, $i=2,3$, and the Jacobian is $J = b_{22}b_{33} - b_{23}b_{32}$. This Jacobian is independent of $\hat{x}$ and hence the transformation from $\hat{y}$ to $\hat{x}$ is one-to-one. When $x_2=x_3=0$, $\bar{y}_i(\hat{x})=b_{i0}$, and

[5] An alternative interpretation is that the individual's utility function is given by (11) where $u=0$, that is $G(x,0)$, and the y_i's introduced in (12) represent random errors of maximization. The latter create a distinction between the observed consumption vector and the consumption vector which maximizes $G(x,0)$. For the data sets used below, it is not possible to distinguish between these interpretations since they are observationally equivalent.

hence density (8) reduces to

$$f(0,0)=\int_{-\infty}^{b_{30}}\int_{-\infty}^{b_{20}} n(y_2,y_3,\Omega)\,\mathrm{d}y_2\,\mathrm{d}y_3. \tag{16}$$

When $x_1=x_2=0$ or $x_1=x_3=0$, similar expressions are obtained by renumbering goods. When two goods are consumed, the density is obtained by integrating under a univariate normal density function. For example, when $x_1, x_2>0$ and $x_3=0$, the density corresponding to (9) with $K=2$ and $M=3$ is

$$f(x_2,0)=\int_{-\infty}^{\bar{y}_3} n(\bar{y}_2,y_3,\Omega)\,\mathrm{abs}\,[J(\hat{x})]\,\mathrm{d}y_3$$

$$=n(\bar{y}_2,\omega_{22})F([\bar{y}_3-\bar{y}_2\omega_{23}/\omega_{22}]/|\Omega|/\omega_{22})\,\mathrm{abs}\,[b_{22}], \tag{17}$$

where F denotes the normal distribution function, $\bar{y}_i=b_{i0}+b_{i2}x_2$ and the last expression is obtained by expressing the joint normal density as the product of the marginal density for y_2 and the conditional density for y_3 given y_2, and carrying out the integration. When $x_1, x_3>0$ or $x_2, x_3>0$ the densities may be obtained by renumbering goods.

3. The Amemiya–Tobin approach

Following the customary approach to the estimation of systems of consumer demand equations, we assume that the consumer attempts to maximize a utility function subject to the budget constraint. The utility maximizing shares are denoted $s_i(v)$, $i=1,\dots,M$, where v is the vector of normalized prices. These deterministic shares are assumed to lie between zero and unity. However, due to errors of maximization by the consumer, errors of measurement of the observed shares, and other random disturbances which influence the consumer's decisions, the observed shares will not coincide with the deterministic shares. The usual approach taken to incorporate these stochastic elements is to add normal disturbances to the deterministic shares implying that the 'shares', so constructed, are normally distributed about the deterministic shares. However, there is nothing in this formulation to ensure that these 'shares' lie between zero and unity. To avoid this difficulty, we assume that the observed shares follow a truncated multivariate normal distribution. By following this procedure we not only respect the constraint that shares lie between zero and unity, but also provide for a pile-up of density at the boundary.

For the purpose of constructing the truncated normal density function, a vector $y^*=(y_1^*,\dots,y_M^*)$ of latent random variables is introduced. Since the shares sum identically to unity their joint density function is degenerate and so we can deal with the joint density function over any $M-1$ shares. Accordingly it is assumed that $(y_1^*,\dots,y_{M-1}^*)$ is distributed normally with mean vector $\mu=[s_1(v),\dots,s_{M-1}(v)]$ and positive definite covariance matrix Ω. Then $y_M^*\equiv 1-\sum_{i=1}^{M-1} y_i^*$ is normally distributed. The vector $y^*=(y_1^*,\dots,y_M^*)$ will not necessarily satisfy the requirement of shares that they be between zero and unity, though they clearly sum to one. Following the idea behind Tobin's limited dependent variable model, we map the density lying outside the unit simplex,

$$S\equiv\left\{(y_1^*,\dots,y_{M-1}^*): y_1^*,\dots,y_{M-1}^*, 1-\sum_{i=1}^{M-1} y_i^*\geqq 0\right\},$$

onto its boundary by defining the vector of shares $y=(y_1,\dots,y_M)$ in terms of y^* as

$$\begin{aligned} y_i&=0, && \text{if} \quad y_i^*\leqq 0,\\ y_i&=y_i^*\Big/\sum_{j\in J} y_j^*, && \text{if} \quad y_i^*>0, \end{aligned} \tag{18}$$

where

$$J\equiv\{j: y_j^*>0\}\cap\{1,2,\dots,M\}.$$

Let there be K shares that are positive with the remaining $M-K$ shares equal to zero and let the categories be ordered so that the positive shares occur first. Then the density for $y=(y_1,\dots,y_K,0,\dots,0)$ is

$$\begin{aligned} f(y_1,\dots,y_K,0,\dots,0)&=\int_{y_1}^{\infty}\int_{\alpha_{K+1}}^{0}\dots\int_{\alpha_{M-1}}^{0} n(\sigma_1(y_1^*),\ \sigma_2(y_1^*),\dots,\sigma_K(y_1^*),\\ &\qquad y_{K+1}^*,\dots,y_{M-1}^*)J(y)\,\mathrm{d}y_{M-1}^*\dots\mathrm{d}y_{K+1}^*\,\mathrm{d}y_1^*,\\ &\qquad\qquad 1\leqq K<M,\\ &=n(y_1^*,y_2^*,\dots,y_{M-1}^*), \qquad K=M, \end{aligned} \tag{19}$$

where

$$\begin{aligned} \alpha_{K+1}&\equiv 1-[\sigma_1(y_1^*)+\sigma_2(y_1^*)+\dots+\sigma_K(y_1^*)],\\ \alpha_j&\equiv\alpha_{K+1}-\sum_{k=K+1}^{j-1} y_k^*, \quad j=K+2,\dots,M-1, \\ \alpha_j(y_1^*)&\equiv y_1^*(y_j/y_1), \qquad j=1,\dots,K, \end{aligned} \tag{20}$$

and where $n(y_1^*, \ldots, y_{M-1}^*)$ is the joint normal density function for the first $M-1$ categories and $J(y)$ is a measure of distance.[6] Since there is always at least one positive share we can, without loss of generality, assume that the first is positive in which case the functions $\sigma_j(y_1^*)$ are well defined. Also if $K<M$ there is at least one share which is zero so we can, again without loss of generality, assume that the last share is zero.

It should be noted that in the case where there are K positive shares and $M-K$ zero shares there are $M!/K!(M-K)!$ combinations to consider. Since any $M-1$ of the y_i^*'s have a joint normal distribution with parameters defined in terms of μ and Ω, we can always arrange the shares so that the positive ones occupy the first K positions. Thus (19) applies to each of these combinations where the normal density n is appropriately defined for each combination.

The likelihood function for a random sample of N observations is

$$L=\prod_{i=1}^{N} f(y_{i1}, \ldots, y_{iM}), \tag{21}$$

where $f(y_{i1}, \ldots, y_{iM})$ is defined by (19). Maximization of (21) yields maximum likelihood estimates of the utility function (or, equivalently the share functions) and covariance parameters.

Several points should be noted about this formulation of the density function for shares. First, although there may be ways other than (18) to allocate the density for y^* to the feasible region S, the one we have chosen is both simple and has the property that the resulting density function is independent of which set of $M-1$ y_i^*'s is used in its derivation. Second, because we take account of the budget constraint (the shares sum to one) the model is not just a simple extension of Tobin's limited dependent variable model for $M>3$.[7] Finally, if we had chosen to estimate a system of expenditure equations rather than share equations, it is evident that the same

[6]Since the joint normal density function for $y_1^*, \ldots, y_M^*$ is degenerate, the calculations involved in (19) may be undertaken in any $(M-1)$-dimensional subspace. If the calculations are undertaken in $y_1^*, \ldots, y_{M-1}^*$ space, where $y_M=0$, the function $J(y)$ is a measure of distance along the line defined parametrically by $\sigma_j(y_1^*)$, $j=1, \ldots, K$, and is

$$J(y) \equiv \left[\sum_{j=1}^{K} (y_j/y_1)^2\right]^{\frac{1}{2}}.$$

To ensure that the last (Mth) observed share is zero, a change of space is required in order to apply (19). If the density function $f(y)$ is to be independent of the space in which the calculation is done (and hence integrate to unity), the function $J(y)$ must be adjusted to take account of the change in space. This is because the distances between corresponding points on corresponding lines in different spaces will be different. However, it should be noted that these distance measures $J(y)$ depend only on the observed y vectors and not upon the parameters of the model, hence they can be ignored in the likelihood function.

[7]However, when $M=2$ the model is a simple extension of Tobin's limited dependent variable model to the case of an upper and lower bound.

type of density function would have to be employed, the only difference being that the upper bound on the sum of expenditures would vary from individual to individual.

This model, like the Kuhn–Tucker model, generates a density for consumption (or expenditure shares) which respects the non-negativity and budget constraints and which allows for a pile-up of density whenever the consumption of one or more goods is zero. There is a difference in the resulting densities for the two models, however, because of a difference in the way the stochastic element is incorporated. In the Kuhn–Tucker model, the consumption vector for an individual is obtained by constrained maximization of a utility function, and may involve zero consumption of one or more goods. Randomness is incorporated by supposing that the parameters of the utility function are randomly distributed over the population. In the Amemiya–Tobin model, individuals have the same utility function. An individual's observed consumption vector is the sum of the utility maximizing consumption vector plus a vector of random disturbances which have a truncated distribution. This truncation allows the observed consumption vector to involve zero expenditures on one or more goods.

Special case: The Stone–Geary utility function

The Stone–Geary utility function is given by

$$G(x)=\prod_{i=1}^{M}(x_i-b_i)^{a_i}, \qquad x_i>b_i, \qquad \sum_{i=1}^{M}a_i=1. \tag{22}$$

For this functional form the demand equations are given by the familiar Linear Expenditure System (LES), which in share form is

$$s_i\equiv v_ix_i=v_ib_i+a_i\left(1-\sum_{k=1}^{M}v_kb_k\right), \qquad i=1,\dots,M. \tag{23}$$

As discussed above these deterministic shares are assumed to lie between zero and one. Addition of a multivariate normal disturbance to (23) yields a vector of latent variables y^*,

$$y_i^*=v_ib_i+a_i\left(1-\sum_{k=1}^{M}v_kb_k\right)+\varepsilon_i, \qquad i=1,\dots,M. \tag{24}$$

These y^*'s are used to construct the truncated normal density, that is, the observed shares (y_i) are defined in terms of the y_i^* according to eq. (18).

Before proceeding with the estimation of these models we consider briefly the possibility of an alternative approach involving the indirect utility function. In many recent empirical studies the approach has been to specify an indirect utility function

$$\tilde{H}(v) \equiv \max_{x} (G(x): v^T x \leqq 1), \tag{25}$$

in which the non-negativity constraints in (1) are ignored. Using Roy's identity the demand functions may be obtained from (25) as

$$x = \tilde{H}_v(v)/v^T \tilde{H}_v(v). \tag{26}$$

Of course $\tilde{H}(v)$ corresponds exactly with $H(v)$ when all the x's are positive, in which case this procedure is appropriate. However, if some of the x's given by (26) are negative then $\tilde{H}(v)$ is inappropriate. At first glance, it might appear that this difficulty could be resolved by setting the negative x's to zero. If this were the case, the model would become

$$\begin{aligned} x_i &= \tilde{H}_i(v)/v^T \tilde{H}_v(v), \quad \text{if} \quad \text{RHS} > 0, \\ &= 0, \qquad\qquad\qquad\quad \text{if} \quad \text{RHS} \leqq 0, \end{aligned} \tag{27}$$

where it is assumed that the demand functions are random, due for example to random preferences. This is a multivariate Tobit model of the type dealt with by Amemiya (1974).

The difficulty with setting the negative x's to zero is that the budget constraint would no longer be satisfied and the demand functions for the remaining x's would not be given by (26), but would have to be obtained from a *new* indirect utility function conditional upon some of the x's being zero. Thus we conclude that the standard indirect utility approach is inappropriate for dealing with non-negativity constraints.

4. Application to meat consumption

In this section we discuss an application of the econometric model outlined above using a survey of household meat consumption carried out by the Bureau of Agricultural Economics in Australia (1970). The survey contains information on the purchases of various types of meat by each household in the sample, together with information on the characteristics of the household members. In the survey, which was undertaken in Melbourne in 1967, the expenditure information is based upon the recall method and covers a one-week period. Because of the short survey period there are categories of meat expenditures involving no purchases, even when the data are aggregated into three broad groups — beef, lamb and other meats.[8] Indeed, meat expenditure is positive for all three meat groups in only 66 percent of households in the survey. The econometric model developed in sections 2 and 3 allows us to estimate an economic model of meat consumption using all observations in the samples.

Each household is assumed to allocate a predetermined level of total meat expenditure m among the three categories — beef, lamb and other meats — by maximizing a utility function with quantities consumed as arguments, subject to the budget constraint that the sum of expenditures is m.[9] Since all households within the survey live in the same city, we assume that they face the same prices. Furthermore, we choose units of measurement such that prices of the three goods are unity, in which case x_i is expenditure on good i and $v_1=v_2=v_3$ for each household in the sample. A consequence of the absence of price variability in the sample is that not all of the parameters of the utility functions can be identified. For the quadratic utility function a convenient normalization which achieves identification of the remaining parameters is

$$\begin{aligned} &a_{ij}=0, \qquad j\neq i, \qquad i,j=1,2,3, \\ &a_{33}=-1. \end{aligned} \tag{28}$$

This leaves seven parameters to be estimated, namely σ_1, σ_2, ρ_{12}, a_{10}, a_{20}, a_{11} and a_{22}. Of course, since there is no price variability in the sample we can only identify the response of quantities to changes in total expenditure on meats, m. The resulting Engel curves, relating the x_i's to m, are linear in m for each combination of goods purchased, and are given by

$$x_i=s_i m-\left[a_{i0}-\sum_{j\in I} s_j a_{j0}\right]\Big/ a_{ii}, \qquad i\in I, \tag{29}$$

[8]Ideally we would like information on the consumption of meat in this period, and there will be some error involved in using reported purchases if these two do not coincide. Unfortunately there is no information on the quantity of meat that individuals are adding to or withdrawing from their freezers. Thus to the extent that these data are used at all in the estimation of preferences the procedures outlined in the text would appear to be the most appropriate.

In addition to this problem there is the possibility that a week is too short a period to permit the revelation of preferences. For the Amemiya–Tobin approach this does not appear serious since any randomness induced by the short period is included in the disturbance term. For the Kuhn–Tucker approach the consequences may be more serious since we are assuming non-stochastic utility-maximizing behaviour.

[9]To allow for the possibility that m is endogenous, the model would have to be extended to include another category of expenditure consisting of 'non-meats', together with the assumption that total expenditure on meats and non-meats is exogenous. However, we have not pursued this approach since the information required was incomplete.

where $s_i = 1/a_{ii}/\sum_{j\in I} 1/a_{jj}$ is the marginal budget share for good i, and I is the set of subscripts for those goods which are consumed. Such curves are consistent with any set of preferences which have a Gorman polar form. For example, the Stone–Geary utility function yields Engel curves which are linear in m. Thus the quadratic utility function provides one way of obtaining the estimating equations, and our results may be reinterpreted within the context of any utility function which has a Gorman polar form.[10] Of course normalizations other than (28) are possible, and while the resulting parameter estimates will be different, the estimates of the Engel curves will not be affected.

For the LES a convenient normalization which achieves identification is to set $\sum_{i=1}^{3} b_i = 0$, this leaves seven parameters to be estimated, namely σ_1, σ_2, ρ_{12}, a_1, a_2, b_1 and b_2. The Engel curves are again linear in m and are given by

$$x_i = b_i - \left(a_i \Big/ \sum_{j\in I} a_j\right) \sum_{j\in I} b_j + \left(a_i \Big/ \sum_{j\in I} a_j\right) m, \qquad i \in I, \tag{30}$$

where as before I is the set of subscripts for those goods which are consumed.

Non-random differences in preferences between households are taken into account by assuming that the a_{i0} parameters of the quadratic utility function and the b_i parameters of the LES are linear functions of a vector of household characteristics. Thus, for example, for the quadratic,

$$a_{i0} \equiv c_{i0} + \sum_{j=1}^{5} c_{ij} z_j,$$

and for LES,

$$b_i \equiv d_{i0} + \sum_{j=1}^{5} d_{ij} z_j,$$

where z_j is the jth characteristic $(j = 1, \ldots, J)$ and the c's are parameters to be estimated. Because of the restriction $\sum_{i=1}^{3} a_{i0} = 1$ for the quadratic, the c parameters have the constraints

$$\sum_{i=1}^{3} c_{i0} = 1 \quad \text{and} \quad \sum_{i=1}^{3} c_{ij} = 0 \quad \text{for all } j,$$

[10]Thus since we have no price variation we could have applied the Kuhn–Tucker procedure directly to the Stone–Geary utility function and would have obtained the same estimates of the Engel curves, but defined in terms of the parameters appearing in (30) below. The reason for not doing so is to illustrate the procedure to follow, when price variation does exist, in estimating a more general utility function than the Stone–Geary, namely the quadratic.

while for the LES we have

$$\sum_1^3 d_{i0}=0 \quad \text{and} \quad \sum_1^3 d_{ij}=0 \quad \text{for all } j.$$

Given the above expressions for a_{i0} and b_i in terms of the z's, it is evident that the Engel curves described by (29) and (30) have exactly the same structure, apart from the parameterization of course. That is, the marginal budget shares are constant (for any given combination of goods purchased) and the intercepts are linear functions of the household characteristics (z's). Thus despite the different parameterizations the non-stochastic structure of the models based on the quadratic and Stone–Geary utility functions is the same.

The following five characteristics are hypothesized to affect household preferences for meat:

z_1 = number of individuals in household 13 years or older;
z_2 = number of individuals in household younger than 13 years;
$z_3=1$ if the household head is a Roman Catholic,
$=0$ otherwise;
$z_4=1$ if the household head was *not* born in Australia or New Zealand,
$=0$ otherwise;
$z_5=1/(R+1)$ if $z_4=1$, where R is the number of years of residence,
$=0$ if $z_4=0$.

Variables z_1 and z_2 are intended to reflect the influence of family size and age composition upon the pattern of household meat consumption. Variables z_4 and z_5 are introduced to take account of possible differences in tastes between Australian born families and others. Variable z_4 allows for a permanent difference in tastes while z_5 allows for the possibility that the taste patterns of non-Australians may alter over time.

After eliminating some households due to lack of complete information the sample consisted of 790 households. Further details regarding the data may be found in appendix A.

We have used the likelihood ratio statistic to test the significance of the demographic variables. Separate tests of the null hypothesis that a group of variables do not affect the shares of meat expenditure were carried out for the following groups: (a) family size and age composition (z_1, z_2), (b) religion (z_3), and (c) country of birth and length of residence (z_4, z_5). The null hypothesis of no effect was rejected at the 1% level of significance for each group of variables, for both estimation procedures. Thus all the demographic variables are significant determinants of the pattern of meat consumption, and are therefore retained in the ensuing analysis.

The effects of the demographic variables upon the consumption of beef, lamb and other meats for each consumption pattern in which more than one meat type is consumed are presented in table 1 for the Kuhn–Tucker method and table 2 for the Amemiya–Tobin approach.[11] It should be recalled that both of these methods yield piece-wise linear Engel curves and thus the results may be readily compared.

Table 1
Marginal budget shares and demographic effects — Kuhn–Tucker approach.[a]

	$\partial x_i/\partial m$	$\partial x_i/\partial z_1$	$\partial x_i/\partial z_2$	$\partial x_i/\partial z_3$	$\partial x_i/\partial z_4$	$\partial x_i/\partial z_5$
x_1, x_2, x_3 positive						
x_1	0.409	0.036	0.052	0.351	0.375	−1.770
	(0.018)	(0.053)	(0.044)	(0.123)	(0.139)	(0.797)
x_2	0.188	0.076	0.096	−0.254	−0.525	−0.465
	(0.010)	(0.038)	(0.033)	(0.092)	(0.105)	(0.619)
x_3	0.403	−0.112	−0.148	−0.097	0.150	2.235
	(0.013)	(0.044)	(0.037)	(0.106)	(0.119)	(0.674)
x_1, x_2 positive						
x_1	0.686	−0.040	−0.049	0.285	0.478	−0.238
x_2	0.314	0.040	0.049	−0.285	−0.478	0.238
	(0.019)	(0.039)	(0.033)	(0.090)	(0.103)	(0.605)
x_1, x_3 positive						
x_1	0.504	0.074	0.100	0.223	0.111	−2.005
x_3	0.496	−0.074	−0.100	−0.223	−0.111	2.005
	(0.018)	(0.044)	(0.037)	(0.105)	(0.119)	(0.669)
x_2, x_3 positive						
x_2	0.3175	0.087	0.112	−0.143	−0.406	−1.027
x_3	0.6825	−0.087	−0.112	0.143	0.406	1.027
	(0.013)	(0.031)	(0.027)	(0.078)	(0.089)	(0.518)

[a]The x_i's and m are measured in dollars per week. Standard errors are contained in parentheses. Since marginal budget shares sum to unity and demographic effects sum to zero, when only two meat types are consumed the standard errors on these shares and effects are the same for both meat types.

Although the deomgraphic variables are significant, as discussed above, they do not appear to have a very large effect upon the consumption of the three meat types. Consider first table 1. When all three meat types are consumed an increase in the number of household members thirteen years or older, z_1, increases beef consumption by 3.6 cents, lamb consumption by 7.6 cents and consequently reduces the consumption of other meats by 11.2

[11]The parameter estimates on which these are based appear in the appendix. Since these values depend on the normalizations chosen to identify the parameters they are not of much interest in themselves. The estimates in tables 1 and 2 are of course independent of these normalizations.

Table 2
Marginal budget shares and demographic effects — Amemiya–Tobin approach.[a]

	$\partial x_i/\partial m$	$\partial x_i/\partial z_1$	$\partial x_i/\partial z_2$	$\partial x_i/\partial z_3$	$\partial x_i/\partial z_4$	$\partial x_i/\partial z_5$
x_1, x_2, x_3 *positive*						
x_1	0.371	0.114	0.038	−0.007	0.024	−0.747
	(0.018)	(0.031)	(0.031)	(0.066)	(0.080)	(0.652)
x_2	0.244	0.012	0.032	−0.123	−0.169	0.298
	(0.018)	(0.934)	(0.034)	(0.066)	(0.079)	(0.659)
x_3	0.385	−0.126	−0.070	0.130	0.145	0.449
	(0.016)	(0.032)	(0.030)	(0.058)	(0.069)	(0.566)
x_1, x_2 *positive*						
x_1	0.603	0.038	−0.004	0.072	0.112	−0.476
x_2	0.397	−0.038	0.004	−0.072	−0.112	0.476
	(0.026)	(0.031)	(0.030)	(0.060)	(0.072)	(0.595)
x_1, x_3 *positive*						
x_1	0.491	0.120	0.054	−0.068	−0.059	−0.601
x_3	0.509	−0.120	−0.054	0.068	0.059	0.601
	(0.019)	(0.026)	(0.026)	(0.053)	(0.063)	(0.515)
x_2, x_3 *positive*						
x_2	0.388	0.056	0.047	−0.126	−0.160	0.009
x_3	0.612	−0.056	−0.047	0.126	0.160	−0.009
	(0.023)	(0.028)	(0.028)	(0.054)	(0.064)	(0.538)

[a]See table 1.

cents. A somewhat surprising result is that Catholics ($z_3=1$) on average spend 35.1 cents more on the consumption of beef, 25.4 cents less on lamb, and 9.7 cents less on other meats (which includes fish) than do others ($z_3=0$).[12]

Variable z_4 allows households with Australian born heads to have different expenditure patterns than other households, while variable z_5 permits an adjustment over time of expenditure patterns of households with non-Australian born heads.[13] If $\partial x_i/\partial z_5=0$ then differences in the expenditure patterns of these two groups persist over time, whereas if $\partial x_i/\partial z_4=0$ the initial differences decrease (asymptotically) as years of residence increase. To interpret the results, we consider first the case of a household whose head is non-Australian and who resided in Australia less than one year. With years of residence equal to zero, $z_5=1$ and so $\partial x_i/\partial z_4+\partial x_i/\partial z_5$ is the difference in the expenditure for category i between Australian and non-Australian headed households. The results indicate that, when all three meat types are

[12]Approximately 20% of the expenditure on the other meats category consists of expenditure on seafood.

[13]Recall that $z_5=0$ for a household with an Australian born head and $z_5=1/(1+R)$ for other households where R is years of residence in Australia of the household head.

consumed, non-Australians have an estimated expenditure for beef which is \$1.39 lower than Australians, \$0.99 lower for lamb, and \$2.39 higher for other meats. As years of residence increases, z_5 declines towards zero so that $\partial x_i/\partial z_4$ is the limiting difference between the expenditures for households with Australian born heads and others for category i. The $\partial x_i/\partial z_4$ estimates for beef, lamb and other meats are 0.375, −0.525 and 0.150, respectively. This means that immigrants reduce their expenditure on other meats and increase their expenditure on beef and lamb as years of residence increase, and eventually spend more on beef but still less on lamb than do households with Australian born heads.

Table 2 may be interpreted in a similar manner, but rather than discuss it in detail we instead make a few general comparisons with table 1. Consider first the signs of the significant coefficients in table 1 (using an asymptotic t-value of 2 as a rough measure). Of the eighteen significant coefficients sixteen have the same sign as the corresponding coefficient in table 2, indicating that the two procedures give the same results in terms of the direction of effects. Second it is interesting to note that the absolute values of all but one of these sixteen coefficients are lower in table 2 than in table 1 although it is not clear why this occurs.

Tables 1 and 2 also present estimates of the marginal budget shares for each of four consumption patterns. The estimated budget shares are of the same order of magnitude in each table but differences between them are large relative to their estimated standard errors.[14] For the 'average household', that is with the z_i's set equal to the sample means and with utility function parameters equal to the estimates provided in appendix B, we have calculated the Engel curves relating expenditure on each meat type to total meat expenditure m.[15] These Engel curves are presented in figs. 1 and 2 (corresponding to tables 1 and 2).

Consider first fig. 1. For m less than \$0.44, lamb (good 2) is the only meat type consumed and its marginal budget share is clearly unity. For m greater than \$0.44 but less than \$0.98, both beef and lamb are consumed with marginal budget shares estimated to be 0.686 and 0.314 respectively. If m exceeds \$0.98 then all three meat types are consumed with marginal budget shares of 0.409, 0.188 and 0.403. While fig. 1 refers to the 'average household', it is evident that individual households with different values for the z_i variables will have different Engel curves, involving a different sequence of consumption patterns as m increases, and with the switches from one pattern to another occurring at different levels of m.

Fig. 2 differs from fig. 1 in that beef is now consumed by the average

[14]Of course this is not a rigorous test of significance.

[15]For a detailed account of the derivation of Engel curves, and of price effects upon demands, for a quadratic utility function using a quadratic programming algorithm, see Wegge (1968). Since there is no price variation in our data, our calculations are much simpler than Wegge's.

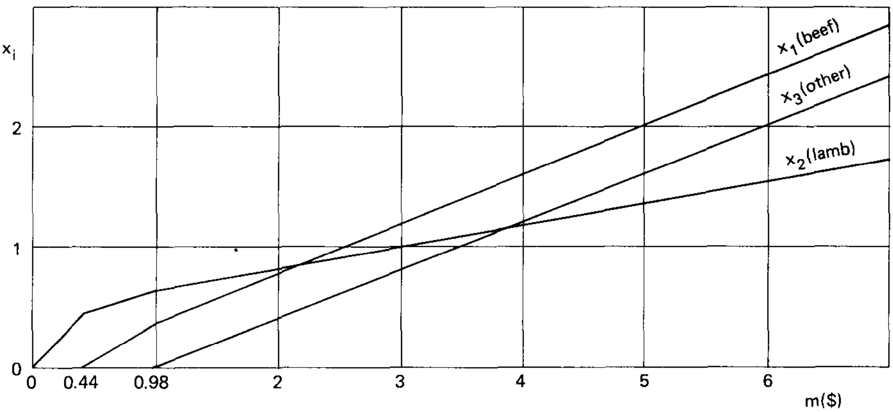

Fig. 1. Kuhn–Tucker approach.

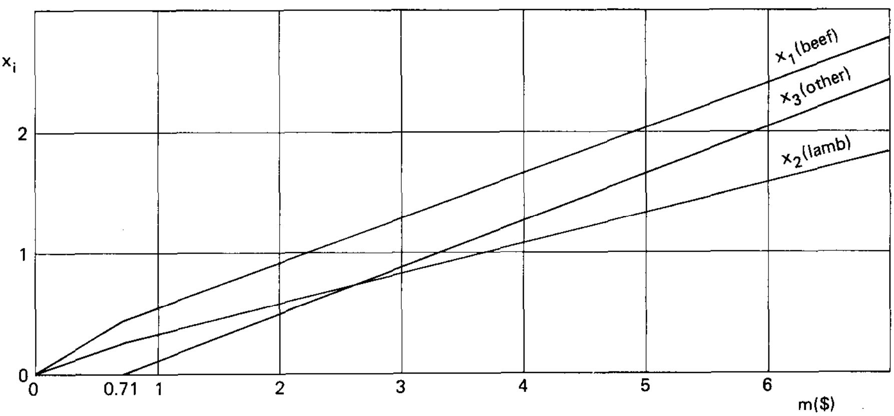

Fig. 2. Amemiya–Tobin approach.

household at very low levels of m, with the marginal budget shares for beef and lamb being 0.686 and 0.314, respectively. For m above $0.71 all three goods are consumed with marginal budget shares for beef, lamb and 'other' being 0.371, 0.244 and 0.385.

The marginal budget shares are very close to those given in fig. 1. Since the mean sample value for m is $5.64, the Engel curves obtained from the two different methods are similar over a wide range of sample observations.

5. Conclusion

In this paper we have formulated two econometric models of consumer demand which explicitly allow expenditure on one or more goods to be zero for a significant proportion of a random sample. The model which we refer to as the Amemiya–Tobin type corresponds more closely to the traditional approach to econometric estimation of systems of demand functions. An individual's observed consumption vector is assumed to be the sum of the utility maximizing consumption vector, obtained without regard to non-negativity constraints, and a vector of random disturbances which have a truncated distribution. This truncation allows the observed vector to involve zero expenditures on one or more goods.

Our other model is based on the full set of Kuhn–Tucker conditions for the maximization of the utility function subject to the budget and non-negativity constraints. Consequently zero consumption on one or more goods is possible. By assuming that preferences are random over the population the density function for the consumption vector of an individual drawn from the population can be derived using the Kuhn–Tucker conditions directly.

Both of our models have been estimated by the method of maximum likelihood for a sample of data on Australian meat consumption. As discussed in the text the results are not very sensitive to the method used. Since both methods involve the same degree of complexity in terms of estimation technique (for example evaluation of the bivariate normal distribution function) this provides no basis for making a choice between them. On theoretical grounds, however, the Amemiya–Tobin approach suffers from the fact that it requires an arbitrary assumption about the allocation of density to the feasible region [eq. (18)]. As discussed in the text our method is simple and results in a density that is independent of which goods are used in its derivation; nevertheless, this arbitrariness still remains. The Kuhn–Tucker approach, on the other hand, does not involve any arbitrary assumptions of this kind, and thus on these grounds may be considered preferable.

The models developed in this paper should prove useful in the analysis of consumer demand when the sample contains a non-trivial proportion of observations for which the consumption of one or more goods is zero. For such samples the traditional econometric approach to the estimation of consumer demand systems is inappropriate. In addition to the survey data analyzed in this paper, there are an increasing number of household surveys of expenditures to which application of our model may be appropriate.

In addition the models are applicable to the analysis of family labour supply. Under the assumption that the household maximizes a utility function with consumption of goods, hours of work of the husband, and hours of work of the wife as arguments then it may be the case that either the husband or wife or both may choose not to work. Indeed, there exist a number of data sets, such as the National Longitudinal Survey and the University of Michigan Survey Research Center's Panel Study of Income Dynamics, in which there is a significant proportion of wives (and a small proportion of husbands) who do not work. Our models may be easily reformulated to apply to these samples.

Finally, although our empirical example involves only three goods the theoretical analysis is presented for an arbitrary number of goods. Nevertheless computational time will increase rapidly as the number of goods increases, hence applications of these methods will likely have to be restricted to systems with a small number of goods.

Appendix A

Summary statistics for Melbourne sample

	s_1	s_2	s_3	m	z_1	z_2	z_3	z_4	R
Mean	0.435	0.437	0.364	5.64	2.65	0.88	0.26	0.33	18.8
Standard deviation	0.217	0.215	0.197	3.18	1.12	1.22	0.44	0.47	15.8

Number of households in each expenditure pattern

	Positive expenditures on goods							
	1, 2, 3	2, 3	1, 3	1, 2	1	2	3	Total
Number	520	53	115	63	18	14	7	790

Notes

(1) For any household s_i is defined as expenditure on the ith meat category divided by m. The mean and standard deviation for any s_i value is calculated after excluding observations for which expenditure on that s_i value is zero. Hence the sum of the average meat shares over the three categories is not equal to one.

(2) The indices 1, 2, 3 refer to three meat categories, beef, lamb and 'other', respectively. The beef category 1 comprises beef, veal, and beef sausage; in terms of the Bureau of Agricultural Economics (B.A.E.) codes it is defined as 101–115, 201–209 and 603. The lamb category 2 comprises lamb and mutton; in terms of the B.A.E. codes it is defined as 301–313 and 401–413. The 'other' category 3 comprises pork, ham, bacon, poultry, game, seafood, offal, small goods (e.g. frankfurters),

sausage (except beef sausage); in terms of the B.A.E. codes it is defined as 501–509, 801–806, 1101–1107, 1109, 1110, 1201–1204, 1208, 1211–1215, 701–709, 901–903, 601, 602 and 604. The items in these three meat categories include fresh, frozen and ready-to-eat (cooked) meat products only, all canned products are excluded (also, pet food is excluded as a meat product).

Appendix B

Estimates of meat preference parameters

	Kuhn–Tucker approach	
Parameter	Maximum likelihood estimate	Asymptotic standard error
σ_1	0.163	0.009
σ_2	0.186	0.010
ρ_{12}	−0.736	0.022
c_{10}	0.288	0.017
c_{20}	0.405	0.020
a_{11}	−0.985	0.072
a_{22}	−2.149	0.126
c_{11}	0.0007	0.006
c_{21}	0.013	0.007
c_{12}	0.0015	0.005
c_{22}	0.017	0.006
c_{13}	0.045	0.015
c_{23}	−0.045	0.016
c_{14}	0.057	0.017
c_{24}	−0.093	0.019
c_{15}	−0.157	0.095
c_{25}	−0.083	0.109
σ_3	0.129	0.005
ρ_{13}	−0.208	0.066
ρ_{23}	−0.509	0.048

Notes

(1) For the purposes of estimation the x_i's and total expenditure m are measured in units of $10 per week.

(2) Standard errors are based on numerical derivatives.

	Amemiya–Tobin approach	
Parameter	Maximum likelihood estimate	Asymptotic standard error
σ_1	0.291	0.008
σ_2	0.289	0.008
ρ_{12}	−0.618	0.024
d_{10}	−0.014	0.004
d_{20}	0.012	0.005
a_1	0.371	0.018
a_2	0.244	0.017
d_{11}	0.011	0.003
d_{21}	0.0012	0.003
d_{12}	0.0038	0.003
d_{22}	0.0032	0.003
d_{13}	−0.0007	0.007
d_{23}	−0.012	0.007
d_{14}	0.0024	0.008
d_{24}	−0.017	0.008
d_{15}	−0.075	0.065
d_{25}	0.030	0.066
σ_3	0.253	0.007
ρ_{13}	−0.444	0.031
ρ_{23}	−0.430	0.032

See Notes (1) and (2) above.

References

Allen, R.G.D. and A.L. Bowley, 1935, Family expenditure (Staples, London).

Amemiya, T., 1974, Multivariate regression and simultaneous equation models when the dependent variables are truncated normal, Econometrica 42, 999–1012.

Arrow, K.J. and A.C. Enthoven, 1961, Quasi-concave programming, Econometrica 29, 779–800.

Barten, A.P., 1977, The systems of consumer demand functions approach: A review, Econometrica 45, 23–52.

Bureau of Agricultural Economics, 1970, Household meat consumption in Melbourne, Beef research report no. 8 (B.A.E., Canberra).

Cragg, J.G., 1971, Some statistical models for limited dependent variables with application to the demand for durable goods, Econometrica 39, 829–844.

Dutt, J.E., 1976, Numerical aspects of multivariate normal probabilities in econometric models, Annals of Economic and Social Measurement 5, 547–561.

Heckman, J.J., 1974, Shadow prices, market wages, and labor supply, Econometrica 42, 679–694.

Houthakker, H.S., 1961, The present state of consumption theory, Econometrica 29, 704–740.

Kuhn, H.W. and A.W. Tucker, 1951, Nonlinear programming, in: J. Neyman, ed., Proceedings of the second Berkeley symposium on mathematical statistics and probability (University of California Press, Berkeley, CA) 481–492.
Tobin, J., 1958, Estimation of relationships for limited dependent variables, Econometrica 26, 24–36.
Wales, T.J. and A.D. Woodland, 1978, A limited dependent variables model for a system of share equations, with an application to meat consumption, Discussion paper no. 78-32 (Department of Economics, University of British Columbia, Vancouver).
Wegge, L., 1968, The demand curves from a quadratic utility indicator, Review of Economic Studies 35, 209–224.