


ISSN 1440-771X
ISBN 0 7326 1071 0

MONASH UNIVERSITY



AUSTRALIA

An EM Algorithm for Modelling Variably-Aggregated Demand

Simone Grose and Keith McLaren

Working Paper 2/2000

March 2000

**DEPARTMENT OF ECONOMETRICS
AND BUSINESS STATISTICS**

AN EM ALGORITHM FOR MODELLING VARIABLY- AGGREGATED DEMAND

SIMONE GROSE*

AND

KEITH MCLAREN

DEPARTMENT OF ECONOMETRICS AND BUSINESS STATISTICS
MONASH UNIVERSITY, AUSTRALIA.

MARCH 2000

* Corresponding author: Dr. S. D. Grose, Department of Econometrics and Business Statistics, Monash University, Clayton, Australia 3800. Email: Simone.Grose@BusEco.monash.edu.au, Tel: +61-3-9905-2964, Fax: +61-3-9905-5474.

ABSTRACT

The response of consumer demand to prices, income, and other characteristics is important for a range of policy issues. Naturally, the level of detail for which consumer behaviour can be estimated depends on the level of disaggregation of the available data. However, it is often the case that the available data is differently aggregated in different time periods, with the information available in later time periods usually being more detailed. The applied researcher is thus faced with choosing between detail, in which case the more highly aggregated data is ignored; or duration, in which case the data must be aggregated up to the “lowest common denominator”. Furthermore, since parametric demand systems invariably involve a large number of parameters, with the number increasing at least linearly with the number of expenditure categories, it may well be that only the second option is feasible. That is, there is simply not enough data available at the finer aggregation level for the chosen model to be estimated.

This paper develops an EM algorithm for the estimation of a consumer demand system involving variably aggregated data. The methodology is based on the observation that more highly aggregated data does in fact contain information on the finer subcategories. It is therefore possible, under certain simplifying assumptions, to derive the distribution of the unobserved fine-level expenditures conditional on the observed but more highly aggregated data. The expectation of the log-likelihood is then taken with respect to this conditional distribution. Under the assumption of multivariate normality both these steps can be performed analytically, resulting in an EM criterion that can be maximised iteratively at comparatively little cost. The technique is applied to an ABS dataset containing historical information relating to private final consumption expenditures on up to 18 commodities.

KEYWORDS: EM Algorithm, Singular demand systems, Linear expenditure system, Missing data.

JEL classification: C32, C51, D12, E21

1. Introduction.

Grose and McLaren (1999) considered the problem of estimating a consumer demand system in a situation in which data is not available on all commodity categories in all time periods. Such a situation can arise quite routinely, since it is common practice to collect, or at least publish, consumption expenditure and price data for more highly disaggregated commodity sets over time, with the result that data is available at differing levels of disaggregation in different time periods. For example, expenditure data may initially be collected for categories “Food”, “Durables” and “Other”; where “Other” is later split into “Other goods” and “Other services”. It is evident, however, that not only does the expenditure on “Other” equal the sum of expenditures on “Other goods” and “Other services”; but that the *expected* expenditure on “Other” also equals the sum of *expected* component expenditures. Furthermore, if an additive stochastic component is assumed, then this also obeys a similar “summing” rule. In other words, the specification of an economic and statistical model for the most disaggregated data naturally implies a corresponding model applying to the data at any level of aggregation.

A known form of aggregation thus implies an “aggregated” model, in which the *observed* expenditure shares are expressed in terms of the model for the complete system. The conventional assumption of additive, normally distributed disturbances then allows the straightforward derivation of the likelihood function implied by differing degrees of commodity disaggregation¹. Nonetheless, straightforward derivation of the likelihood function does not necessarily equate with straightforward estimation. It turns out that it is no longer possible to “concentrate” the likelihood, leaving all the model parameters, including the $\frac{1}{2}n(n+1)$ parameters of the covariance matrix, to be estimated numerically. The result is an optimization problem found by Grose and McLaren to be infeasible for any realistic sample size.

¹ The “aggregated” likelihood is reviewed in more detail in Section 2.

The strategy considered in Grose and McLaren (1999) (hereafter GM) was to reduce the dimension of the optimization problem by resorting to a reduced-order parameterization of the covariance matrix. This paper considers, instead, a “data augmentation” approach, in which the EM algorithm is used to iteratively maximize the aggregated likelihood, without the need to restrictively reparameterize the model covariance structure.

2. The “aggregated” likelihood.

Consider the standard time series specification of an N -commodity demand system, in which the t^{th} observation on the N -vector of expenditure shares is modelled as a function of the N -vector of prices $\tilde{\mathbf{p}}'_t$, income (assumed equal to total expenditure) in the t^{th} period m_t , a k -vector of parameters θ , and an additive, serially independent, zero mean disturbance with constant $N \times N$ positive-semidefinite variance-covariance matrix $\tilde{\Sigma}$ of rank $n = N - 1$. That is,

$$\tilde{\mathbf{w}}'_t = \tilde{\mathcal{M}}'(\tilde{\mathbf{p}}'_t, m_t, \theta) + \tilde{\mathbf{u}}'_t; \quad \tilde{\mathbf{u}}_t \sim (0, \tilde{\Sigma}), \quad (2.1)$$

where $\tilde{\mathcal{M}}(\tilde{\mathbf{p}}'_t, m_t, \theta)$ is the N -vector of *expected* expenditure shares, conditional on prices and income, for given θ^2 . The “full rank” system, involving n of the N categories, is then written as

$$\mathbf{w}'_t = \mathcal{M}'(\tilde{\mathbf{p}}'_t, m_t, \theta) + \mathbf{u}'_t; \quad \mathbf{u}_t \sim (0, \Sigma); \quad (2.2)$$

where we assume, without loss of generality, that $\mathbf{w}_t = \mathbf{J} \tilde{\mathbf{w}}_t$, where $\mathbf{J} = [\mathbf{I}_n \mid \mathbf{0}]$, \mathbf{u}_t and \mathcal{M} are defined analogously.

The further assumption that \mathbf{u}_t is serially independent and distributed n -variate normal³ then implies the conventional log-likelihood⁴

² Quantities pertaining to the complete $T \times N$ system will be indicated by a “ $\tilde{\cdot}$ ” over the symbol for the corresponding “full rank” quantity.

³ The multivariate normal (MVN) assumption is quite standard, even after the transformation to expenditure share form. Other possibilities that take the bounds on the expenditure share disturbance into account are naturally rather less tractable. See, for instance, Fry, Fry and McLaren (1996).

⁴ Constants will generally be ignored when writing down likelihood functions.

$$\ell(\theta, \Sigma) = -\frac{T}{2} \ln |\Sigma| - \frac{1}{2} \text{tr}(\Sigma^{-1} \mathbf{U}' \mathbf{U}), \quad (2.3)$$

where

$$\mathbf{U} = \mathbf{W} - \mathcal{M}(\tilde{\mathbf{P}}, \mathbf{m}, \theta) \quad (2.4)$$

$\begin{matrix} T \times n & T \times n & \mathcal{M} & \begin{matrix} T \times N & T \times 1 \end{matrix} \end{matrix}$

is the $T \times n$ matrix of disturbances, \mathbf{W} is the $T \times n$ matrix of observed expenditure shares, and $\mathcal{M}(\cdot)$ is the $T \times n$ matrix of expected expenditure shares, conditional on the $T \times N$ matrix of prices $\tilde{\mathbf{P}}$, the $T \times 1$ vector of total expenditures \mathbf{m} , and the vector of “mean” parameters, θ .

Suppose now that the data is divided according to the differing degrees of expenditure category aggregation in each of $S > 1$ subperiods $\mathcal{T}_1, \dots, \mathcal{T}_S$. With a slight change in notation, let the N -vector of complete, but possibly partially unobserved, expenditure shares⁵ be now denoted by $\tilde{\mathbf{x}}$, let $\tilde{\mathbf{w}}$ denote the vector of *observed*, but more highly aggregated, expenditure shares, and note that $\tilde{\mathbf{w}}_t$ is naturally a linear combination of the elements of $\tilde{\mathbf{x}}_t$. It therefore follows (see GM for details) that, after selecting a category common to all subperiods for omission, we have, for the r^{th} subperiod \mathcal{T}_r , $r = 1, \dots, S$, the full-rank “aggregated” model

$$\mathbf{w}'_t = \mathbf{x}'_t \mathbf{A}'_r = \mathcal{M}(\tilde{\mathbf{p}}'_t, m_t, \theta) \mathbf{A}'_r + \mathbf{u}'_t \mathbf{A}'_r; \quad \mathbf{A}_r \mathbf{u}_t \sim N(0, \mathbf{A}_r \Sigma \mathbf{A}'_r),$$

where $\tilde{\mathcal{M}}(\cdot)$ denotes the expectation of the *complete* N -commodity expenditure set $\tilde{\mathbf{x}}$, $\mathbf{x} = \mathbf{J} \tilde{\mathbf{x}}$ is the underlying n -vector of expenditure shares after exclusion of the common commodity, $\mathbf{w} = \mathbf{A}_r \mathbf{x}$ is the n_r -vector of expenditure shares actually observed in the r^{th} subperiod (except for the excluded common commodity), and \mathbf{A}_r is the $n_r \times n$ full row rank “aggregation matrix” taking \mathbf{x} into \mathbf{w} . Assuming T_r such observations then yields the r^{th} subperiod log-likelihood:

$$\ell_r = -\frac{T_r}{2} \ln |\Sigma_r| - \frac{1}{2} \text{tr}(\Sigma_r^{-1} \mathbf{U}'_r \mathbf{U}_r), \quad (2.5)$$

where

$$\Sigma_r = \mathbf{A}_r \Sigma \mathbf{A}'_r, \quad n_r \times n_r$$

⁵ It will be convenient, in the following, to treat expenditures and expenditure shares interchangeably, assuming throughout that likelihood (2.3) is appropriate in either case. The actual choice of endogenous

$$\mathbf{U}_r = \mathbf{W}_r - \mathcal{M}(\tilde{\mathbf{P}}_r, \mathbf{m}_r, \theta) \mathbf{A}'_r, \quad (2.6)$$

$\begin{matrix} T_r \times n_r & T_r \times n_r & T_r \times N & T_r \times 1 & n \times n_r \end{matrix}$

\mathbf{W}_r and \mathbf{U}_r are the $T_r \times n_r$ matrices of observed expenditure shares and their disturbances for the r^{th} subperiod, and $\mathcal{M}(\tilde{\mathbf{P}}_r, \mathbf{m}_r, \theta) \equiv \mathcal{M}_r(\theta)$ is the $T_r \times n$ matrix of expected expenditure shares, conditional on the $T_r \times N$ matrix of prices⁶ $\tilde{\mathbf{P}}_r$, the $T_r \times 1$ vector of total expenditures \mathbf{m}_r , and the k -vector of mean parameters θ . The “aggregated” log-likelihood follows:

$$\ell(\theta, \Sigma) = \sum_{r=1}^S \ell_r = -\frac{1}{2} \sum_{r=1}^S T_r \ln |\mathbf{A}_r \Sigma \mathbf{A}'_r| - \frac{1}{2} \sum_{r=1}^S \text{tr}(\mathbf{U}'_r \mathbf{U}_r (\mathbf{A}_r \Sigma \mathbf{A}'_r)^{-1}). \quad (2.7)$$

Now, in the case of likelihood (2.3) there is a well-known closed-form expression for the MLE of Σ ; namely $\hat{\Sigma} = T^{-1} \mathbf{U}' \mathbf{U}$. The FOC⁷ for Σ in likelihood (2.7), on the other hand, has no such closed form solution; and hence there is no “aggregated” analogue of the usual profile likelihood that might be feasibly maximized with respect to θ . Grose and McLaren dealt with this problem by reparameterizing Σ as per De Boer and Harkema (1986), resulting in a feasible estimation procedure at the price of a somewhat restrictive covariance structure. We now consider a method of maximizing (2.7) as it stands.

3. “Modelling” the missing data: the EM algorithm.

We are aware that, if expenditure data were observed for all commodity categories in all time periods, then the covariance matrix can be concentrated out of the log-likelihood. This results in an estimation problem that, while still infeasible for “small” sample sizes (such as, it must be noted, the 18 category, 27 annual observations example of Grose and McLaren (1999)), *is* practical in larger samples, such as those that become available if we are willing to use quarterly data. In other words, if we had a complete expenditure set, spanning an adequate timeframe, our model could be

variable makes no difference to the present argument, and thus may as well be reserved for the model-specification stage.

⁶ Note that it is assumed that, although expenditure data is not available for all N commodities in all time periods, price data is.

estimated without imposing assumptions on the cross-commodity covariance matrix. The main problem with the estimation of such a model when some of the expenditure data are missing is due to our inability to “remove” the covariance matrix from the consequent, more complex, likelihood.

This suggests an alternative strategy based on the idea of *data augmentation* and the expectation-maximization (EM) algorithm. Briefly, this involves using the chosen model to simulate, or *impute* values to, the missing expenditure shares, conditional on given values of the parameters θ and Σ , and the data actually observed. This “completes” the expenditure dataset, enabling maximization of the likelihood, conditional on the imputed data. The log-likelihood function of the parameters given the *observed* data is then just the expectation of the “augmented” log-likelihood, taken with respect to the conditional density of the missing data; and can therefore be estimated by averaging over multiple such imputed datasets.

Let us denote the complete (but only partially observed) expenditure set by \mathbf{X} . \mathbf{W} will denote the observed expenditure data, and \mathbf{Z} the unobserved remainder. That is,

$$\mathbf{X} = (\mathbf{W}, \mathbf{Z}) = (\text{observed}, \text{missing}).$$

Also let $\omega = (\theta, \Sigma)$ denote the complete set of unknown parameters.

Obviously, the “complete” log-likelihood $\ell(\omega | \mathbf{X})$ is as per (2.3) for a model of the form (2.2) with MVN errors; that is,

$$\ell(\omega | \mathbf{X}) = \text{const} - \frac{T}{2} \ln |\Sigma| - \frac{1}{2} \text{tr}(\mathbf{U}(\theta) \Sigma^{-1} \mathbf{U}'(\theta)); \quad (3.1)$$

where $\mathbf{U}(\theta) = \mathbf{X} - \mathcal{M}(\theta)$, and expected expenditure $\mathcal{M}(\theta)$ is assumed known⁸ for given θ . We note that, in our missing data context, $\ell(\omega | \mathbf{X})$ involves not only the

⁷ See Grose and McLaren (1999), §3.1.

⁸ Strictly speaking, $\mathcal{M}(\theta)$ is the expected expenditure *conditional* on the observed values of any explanatory variables (such as price, and total expenditure). This conditioning is thus implicit in (3.1); allowing us to evade, in the usual manner, the regression aspect of the model. It is not made explicit for reasons of notational convenience and clarity. We note, however, that the requirement that the expectation of \mathbf{X} be known except for θ *does* imply that the complete $T \times N$ matrix of prices and the T -vector of total expenditures must be available.

unknown parameters θ and Σ , but also the unknown data \mathbf{Z} ; and thus could not in general be maximized consistently with respect to \mathbf{Z} even if this were thought desirable.

Accordingly, let $Q(\omega, \omega^*)$ denote the *expected* log-likelihood of the parameters ω ; with the expectation taken with respect to the missing data \mathbf{Z} , conditional on the observed data \mathbf{W} , and some given ω^* . That is,

$$Q(\omega, \omega^*) = E_{\mathbf{Z}|\mathbf{W}, \omega^*} \ell(\omega|\mathbf{X}) = \int_{\mathcal{Z}} \ell(\omega|\mathbf{X}) \mathcal{P}(\mathbf{Z}|\mathbf{W}, \omega^*) d\mathbf{Z}; \quad (3.2)$$

where “ $\mathcal{P}(\cdot)$ ” denotes a probability density function, and $\mathcal{P}(\mathbf{Z}|\mathbf{W}, \omega^*)$ is the density of \mathbf{Z} conditional on \mathbf{W} and ω^* .

Now, assuming some pre-specified ω^* , $Q(\omega, \omega^*)$ can (in principle) be maximized with respect to ω , yielding a “first approximation” $\hat{\omega}$; which can then be used to refine the conditional density of \mathbf{Z} . A new expectation $Q(\omega, \hat{\omega})$ and a new maximizing $\hat{\omega}$ follow. In summary, beginning with some plausible starting value $\hat{\omega}^{(0)}$, the iterative EM scheme is based on the construction, and subsequent maximization, of $Q(\omega, \hat{\omega}^{(i)})$ with respect to ω , yielding the next iterate $\hat{\omega}^{(i+1)}$, the corresponding maximum $Q(\hat{\omega}^{(i+1)}, \hat{\omega}^{(i)})$, and a “new” criterion $Q(\omega, \hat{\omega}^{(i+1)})$.

It can be shown⁹ that $Q(\hat{\omega}^{(i+1)}, \hat{\omega}^{(i)}) \geq Q(\hat{\omega}^{(i)}, \hat{\omega}^{(i)})$ implies $\ell(\hat{\omega}^{(i+1)}|\mathbf{W}) \geq \ell(\hat{\omega}^{(i)}|\mathbf{W})$; that is, the scheme results in a non-decreasing sequence of log-likelihood values; and that, subject to certain regularity conditions, this sequence converges to a maximum of the likelihood function. Hence the EM algorithm (eventually) maximizes (or at least finds a maximum of) the observed log-likelihood.

The crucial point here is that it be possible to maximize $Q(\omega, \omega^*)$, where it is not possible (or is very difficult) to maximize $\ell(\omega|\mathbf{W})$ directly. The latter, as we noted in §2, is the situation with regard to our “aggregated” log-likelihood if we do not wish to reparameterize Σ , because we cannot concentrate Σ out of the likelihood. Equally, however, maximization of the expected log-likelihood will itself be feasible only if it turns out to be possible to concentrate Σ out of $Q(\omega, \omega^*)$. On this point observe, for

⁹ See Dempster, Laird and Rubin (1977) for the proofs of these results in a general setting. A somewhat more accessible summary can be found in McLachlan and Krishnan (1997).

later reference, that for a criterion of the form (3.1) the $n \times n$ matrix of scores with respect to the elements of Σ is

$$\frac{\partial \ell}{\partial \Sigma} = \Sigma^{-1} \{ \mathbf{U}'\mathbf{U} - T\Sigma \} \Sigma^{-1} - \frac{1}{2} \mathbf{D}_{\Sigma^{-1} \{ \mathbf{U}'\mathbf{U} - T\Sigma \} \Sigma^{-1}}; \quad (3.3)$$

where \mathbf{D}_Q denotes the diagonal matrix with nonzero elements set equal to the diagonal elements of Q . The corresponding first order condition (FOC) for Σ is thus

$$\mathbf{U}'\mathbf{U} - T\Sigma = 0. \quad (3.4)$$

3.1 The conditional density of the missing data

Suppose, for any period t , the complete expenditure model is written

$$\mathbf{x}_t | \omega \sim N(\boldsymbol{\mu}_t(\theta), \Sigma); \quad (3.5)$$

where, in the r^{th} subperiod, only $\mathbf{w}_t = \mathbf{A}_r \mathbf{x}_t$, $n_r \leq n$, is observed.

For convenience dropping the time subscript, partition the n -vector \mathbf{x} as

$$\mathbf{x} = \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{pmatrix} \begin{matrix} m_r \times 1 \\ n_r \times 1 \end{matrix} \sim N \left(\begin{pmatrix} \boldsymbol{\mu}_1(\theta) \\ \boldsymbol{\mu}_2(\theta) \end{pmatrix}, \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix} \right),$$

where the commodities are ordered such that the $m_r = (n - n_r)$ -vector \mathbf{x}_1 equals the unknown, or “missing”, expenditure set \mathbf{z} . The remainder \mathbf{x}_2 , though not directly observed, could be determined from knowledge of \mathbf{x}_1 and \mathbf{w} . The r^{th} subperiod transformation from \mathbf{x} to (\mathbf{z}, \mathbf{w}) ,

$$\mathbf{x} = \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{pmatrix} \begin{matrix} m_r \times 1 \\ n_r \times 1 \end{matrix} \rightarrow \begin{pmatrix} \mathbf{z} \\ \mathbf{w} \end{pmatrix} \begin{matrix} m_r \times 1 \\ n_r \times 1 \end{matrix}$$

is therefore defined by $\mathbf{z} = \mathbf{x}_1 = \mathbf{J}_r \mathbf{x}$ and $\mathbf{w} = \mathbf{A}_r \mathbf{x}$; where \mathbf{J}_r denotes the $m_r \times n$ matrix $\begin{bmatrix} \mathbf{I}_{m_r} & \mathbf{0} \end{bmatrix}$. We thus have

$$\begin{pmatrix} \mathbf{z} \\ \mathbf{w} \end{pmatrix} = \mathbf{B}_r \mathbf{x}, \text{ where } \mathbf{B}_r = \begin{bmatrix} \mathbf{J}_r \\ \mathbf{A}_r \end{bmatrix},$$

implying

$$\begin{pmatrix} \mathbf{z} \\ \mathbf{w} \end{pmatrix} \sim \mathcal{N}(\mathbf{B}_r \boldsymbol{\mu}, \mathbf{B}_r \boldsymbol{\Sigma} \mathbf{B}_r') \equiv \mathcal{N}\left(\begin{pmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \end{pmatrix}, \begin{bmatrix} \boldsymbol{\Omega}_{11} & \boldsymbol{\Omega}_{12} \\ \boldsymbol{\Omega}_{21} & \boldsymbol{\Omega}_{22} \end{bmatrix}\right),$$

with $\mathbf{v}_1 = \mathbf{J}_r \boldsymbol{\mu} \equiv \boldsymbol{\mu}_1$, $\mathbf{v}_2 = \mathbf{A}_r \boldsymbol{\mu}$, $\boldsymbol{\Omega}_{11} = \mathbf{J}_r \boldsymbol{\Sigma} \mathbf{J}_r' \equiv \boldsymbol{\Sigma}_{11}$, $\boldsymbol{\Omega}_{12} = \mathbf{J}_r \boldsymbol{\Sigma} \mathbf{A}_r'$, and $\boldsymbol{\Omega}_{22} = \mathbf{A}_r \boldsymbol{\Sigma} \mathbf{A}_r' \equiv \boldsymbol{\Sigma}_r$. The standard result regarding the conditional distribution of a subvector of normal variates then instantly yields

$$\mathbf{z} | \mathbf{w}, \boldsymbol{\omega} \sim \mathcal{N}(\mathbf{v}_1 + \boldsymbol{\Omega}_{12} \boldsymbol{\Omega}_{22}^{-1} (\mathbf{w} - \mathbf{v}_2), \boldsymbol{\Omega}_{11} - \boldsymbol{\Omega}_{12} \boldsymbol{\Omega}_{22}^{-1} \boldsymbol{\Omega}_{21}),$$

where $(\boldsymbol{\Omega}_{11} - \boldsymbol{\Omega}_{12} \boldsymbol{\Omega}_{22}^{-1} \boldsymbol{\Omega}_{21})^{-1} = (\boldsymbol{\Omega}^{-1})_{11}$ is the top-left $m_r \times m_r$ partition of $\boldsymbol{\Omega}^{-1} = \mathbf{B}_r^{-1} \boldsymbol{\Sigma}^{-1} \mathbf{B}_r^{-1}$. Hence, for $t \in \mathcal{T}_r$,

$$\mathbf{z}_t | \mathbf{w}_t, \boldsymbol{\omega} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{z}_t | \mathbf{w}_t}, \boldsymbol{\Sigma}_{\mathbf{z} | \mathbf{w}}), \quad (3.6)$$

where $\boldsymbol{\mu}_{\mathbf{z}_t | \mathbf{w}_t} = \mathbf{J}_r \boldsymbol{\mu}_t + \mathbf{J}_r \boldsymbol{\Sigma} \mathbf{A}_r' (\mathbf{A}_r \boldsymbol{\Sigma} \mathbf{A}_r')^{-1} (\mathbf{w}_t - \mathbf{A}_r \boldsymbol{\mu}_t)$ is the conditional mean of \mathbf{z}_t , the $m_r \times m_r$ matrix $\boldsymbol{\Sigma}_{\mathbf{z} | \mathbf{w}} = \mathbf{J}_r \boldsymbol{\Sigma} \mathbf{J}_r' + \mathbf{J}_r \boldsymbol{\Sigma} \mathbf{A}_r' (\mathbf{A}_r \boldsymbol{\Sigma} \mathbf{A}_r')^{-1} \mathbf{A}_r \boldsymbol{\Sigma} \mathbf{J}_r'$ is its conditional variance, and, for the model of §2 $\boldsymbol{\mu}_t \equiv \mathcal{M}(\tilde{\boldsymbol{\rho}}_t', m_t, \boldsymbol{\theta})$. The density of the $T_r \times m_r$ matrix of missing data \mathbf{Z}_r , conditional on the $T_r \times n_r$ matrix of observed data \mathbf{W}_r and $\boldsymbol{\omega} = (\boldsymbol{\theta}, \boldsymbol{\Sigma})$, then follows

$$\begin{aligned} \mathcal{P}(\mathbf{Z}_r | \mathbf{W}_r, \boldsymbol{\omega}) &= \prod_{t \in \mathcal{T}_r} \mathcal{P}(\mathbf{z}_t | \mathbf{w}_t, \boldsymbol{\omega}) \\ &= (2\pi)^{-m_r T_r / 2} \left| \boldsymbol{\Sigma}_{\mathbf{z} | \mathbf{w}}^{-1} \right|^{T_r / 2} \exp \left\{ -\frac{1}{2} \sum_{t \in \mathcal{T}_r} (\mathbf{z}_t - \boldsymbol{\mu}_{\mathbf{z}_t | \mathbf{w}_t})' \boldsymbol{\Sigma}_{\mathbf{z} | \mathbf{w}}^{-1} (\mathbf{z}_t - \boldsymbol{\mu}_{\mathbf{z}_t | \mathbf{w}_t}) \right\} \\ &= (2\pi)^{-m_r T_r / 2} \left| \boldsymbol{\Sigma}_{\mathbf{z} | \mathbf{w}}^{-1} \right|^{T_r / 2} \exp \left\{ -\frac{1}{2} \text{tr} \left\{ (\mathbf{Z}_r - \mathcal{M}_{\mathbf{Z}_r | \mathbf{W}_r}) \boldsymbol{\Sigma}_{\mathbf{z} | \mathbf{w}}^{-1} (\mathbf{Z}_r - \mathcal{M}_{\mathbf{Z}_r | \mathbf{W}_r})' \right\} \right\} \end{aligned}$$

where $\mathcal{M}_{\mathbf{Z}_r | \mathbf{W}_r}$ is the $T_r \times m_r$ matrix with t^{th} row equal to $\boldsymbol{\mu}_{\mathbf{z}_t | \mathbf{w}_t}'$.

3.2 MCEM

Simulation of the missing data in each subperiod is thus quite straightforward for given \mathbf{W}_r , $r = 1, \dots, S$, and $\omega^* = (\theta^*, \Sigma^*)$. Having obtained an imputed \mathbf{Z}_r (denoted \mathbf{Z}_r^* to emphasize its dependence on ω^*), the augmented $T_r \times n$ expenditure set for $t \in \mathcal{T}_r$, \mathbf{X}_r^* , then follows via the reverse transform

$$\mathbf{x}_t = \mathbf{B}_r^{-1} \begin{pmatrix} \mathbf{z}_t \\ \mathbf{w}_t \end{pmatrix} \Rightarrow \mathbf{X}_r^* = \begin{bmatrix} \mathbf{Z}_r^* & \mathbf{W}_r \end{bmatrix} \mathbf{B}_r^{-1'}, \quad r = 1, \dots, S.$$

$T_r \times n$ $T_r \times m_r$ $T_r \times n_r$

Finally, combining \mathbf{X}_1^* , ..., \mathbf{X}_S^* yields the corresponding ‘‘realisation’’, \mathbf{X}^* , of the complete $T \times n$ expenditure set.

The expected log-likelihood $Q(\omega, \omega^*)$ can now be estimated by the *average* log-likelihood, taken over M such independently realised augmented datasets $\mathbf{X}_1^*, \dots, \mathbf{X}_M^*$,

$$\bar{Q}(\omega, \omega^*) = \frac{1}{M} \sum_{j=1}^M \ell(\omega | \mathbf{X}_j^*);$$

leading to the Monte-Carlo EM (MCEM) algorithm as follows.

For $i = 0, 1, 2, \dots$, some starter $\hat{\omega}^{(0)} = (\hat{\theta}^{(0)}, \hat{\Sigma}^{(0)})$, and prespecified tolerance ε :

- (1) Draw $\mathbf{Z}_{r1}, \mathbf{Z}_{r2}, \dots, \mathbf{Z}_{rM} \stackrel{iid}{\sim} \mathcal{P}(\mathbf{Z}_r | \mathbf{W}_r, \hat{\omega}^{(i)})$ for each subperiod $r = 1, 2, \dots, S$.
- (2) Combine $\mathbf{Z}_{r1}, \mathbf{Z}_{r2}, \dots, \mathbf{Z}_{rM}$ with \mathbf{W}_r to generate $\mathbf{X}_{r1}, \mathbf{X}_{r2}, \dots, \mathbf{X}_{rM}$, and hence M realizations of the augmented expenditure data, $\mathbf{X}_1^*, \mathbf{X}_2^*, \dots, \mathbf{X}_M^*$.
- (3) Maximize the average log-likelihood $\bar{Q}(\omega, \hat{\omega}^{(i)})$ with respect to ω , so obtaining $\hat{\omega}^{(i+1)}$.
- (4) Repeat (1) – (3) until $|\bar{Q}(\hat{\omega}^{(i+1)}, \hat{\omega}^{(i)}) - \bar{Q}(\hat{\omega}^{(i)}, \hat{\omega}^{(i)})| < \varepsilon$.

Notice, however, that

$$\frac{\partial \bar{Q}(\omega, \omega^*)}{\partial \Sigma} = \frac{1}{M} \sum_{j=1}^M \frac{\partial \ell(\omega | \mathbf{X}_j^*)}{\partial \Sigma};$$

where $\ell(\omega|\mathbf{X}_j^*)$ is as per (3.1). The required derivative is thus of the form (3.3), with \mathbf{U} replaced by $\mathbf{U}_j = \mathbf{X}_j^* - \mathcal{M}(\theta)$, implying

$$\begin{aligned}\frac{\partial \bar{Q}(\omega, \omega^*)}{\partial \Sigma} &= \frac{1}{M} \sum_{j=1}^M \left(\Sigma^{-1} \{ \mathbf{U}_j' \mathbf{U}_j - T \Sigma \} \Sigma^{-1} \right) - \frac{1}{2M} \sum_{j=1}^M \mathbf{D}_{\Sigma^{-1} \{ \mathbf{U}_j \mathbf{U}_j - T \Sigma \} \Sigma^{-1}} \\ &= \Sigma^{-1} \{ \bar{\mathbf{S}}^* - T \Sigma \} \Sigma^{-1} - \frac{1}{2} \mathbf{D}_{\Sigma^{-1} \{ \bar{\mathbf{S}}^* - T \Sigma \} \Sigma^{-1}}\end{aligned}$$

where

$$\bar{\mathbf{S}}^* = \frac{1}{M} \sum_{j=1}^M \mathbf{U}_j' \mathbf{U}_j .$$

Thus $\frac{\partial \bar{Q}(\omega, \omega^*)}{\partial \Sigma} = 0$ yields $\hat{\Sigma} = \bar{\mathbf{S}}^* / T$; implying the ‘‘concentrated’’ criterion

$$\bar{Q}_c(\theta, \omega^*) = \frac{1}{M} \sum_{j=1}^M \ell_c(\theta | \mathbf{X}_j^*) = \text{const} - \frac{T}{2} \ln |\bar{\mathbf{S}}^*(\theta)|. \quad (3.7)$$

Consequently, Step (3), which we recall in our context is likely to be infeasible, would now be replaced by maximization of $\bar{Q}_c(\theta, \hat{\omega}^{(i)})$ (or, equivalently, minimisation of $\ln |\bar{\mathbf{S}}^*(\theta)|$) with respect to θ , yielding $\hat{\theta}^{(i+1)}$, $\bar{\mathbf{S}}^{(i+1)}$, and hence $\hat{\Sigma}^{(i+1)}$.

4. EM: performing the E-step analytically.

MCEM is thus feasible, because it allows us to concentrate Σ out of the criterion, and relatively simple to implement. However, it is, as it happens, extremely time-consuming, due to the very large¹⁰ number of imputation draws needed in the MC expectation stage to achieve convergence with an acceptable degree of precision. We would clearly prefer an analytic expression for the expected log-likelihood.

We accordingly require a closed form solution for (3.2), where $\ell(\omega | \mathbf{X})$ is given by (3.1), and

$$\mathcal{P}(\mathbf{Z} | \mathbf{W}, \omega) = \prod_{r=1}^S \mathcal{P}(\mathbf{Z}_r | \mathbf{W}_r, \omega).$$

4.1 “Vectorizing” the conditional density

To rewrite the conditional density in a form more suited to our purpose, let us define the “*vecr*” operator¹¹, and the accompanying notation:

$$\underline{\underline{\mathbf{u}}}_{nT \times 1} \triangleq \text{vecr}(\mathbf{U}) = \text{vec}(\mathbf{U}');$$

so that $\text{tr}(\mathbf{U}\Sigma^{-1}\mathbf{U}') \equiv \underline{\underline{\mathbf{u}}}' [\mathbf{I}_T \otimes \Sigma^{-1}] \underline{\underline{\mathbf{u}}}$. Similarly defining $\underline{\underline{\mathbf{x}}} = \text{vecr}(\mathbf{X})$ and

$\underline{\underline{\mu}}_{\mathbf{x}}(\theta) = \text{vecr}(\mathcal{M}(\theta))$; so that $\underline{\underline{\mathbf{u}}} = \underline{\underline{\mathbf{x}}} - \underline{\underline{\mu}}_{\mathbf{x}}(\theta)$, then implies the “complete” model

$$P(\underline{\underline{\mathbf{x}}} | \omega) = (2\pi)^{-nT/2} |\Sigma|^{-T/2} \exp\left\{-\frac{1}{2}(\underline{\underline{\mathbf{x}}} - \underline{\underline{\mu}}_{\mathbf{x}}(\theta))'(\mathbf{I}_T \otimes \Sigma^{-1})(\underline{\underline{\mathbf{x}}} - \underline{\underline{\mu}}_{\mathbf{x}}(\theta))\right\}. \quad (4.1)$$

We further define

$$\underline{\underline{\mathbf{w}}} = \begin{pmatrix} \underline{\underline{\mathbf{w}}}_1 \\ \vdots \\ \underline{\underline{\mathbf{w}}}_S \end{pmatrix} = \begin{pmatrix} \text{vecr}(\mathbf{W}_1) \\ \vdots \\ \text{vecr}(\mathbf{W}_S) \end{pmatrix}_{\substack{n_1 T_1 \times 1 \\ \vdots \\ n_S T_S \times 1}} \quad \text{and} \quad \underline{\underline{\mathbf{z}}} = \begin{pmatrix} \underline{\underline{\mathbf{z}}}_1 \\ \vdots \\ \underline{\underline{\mathbf{z}}}_S \end{pmatrix} = \begin{pmatrix} \text{vecr}(\mathbf{Z}_1) \\ \vdots \\ \text{vecr}(\mathbf{Z}_S) \end{pmatrix}_{\substack{m_1 T_1 \times 1 \\ \vdots \\ m_S T_S \times 1}};$$

so that the division of \mathbf{X} into “observed” and “missing” (\mathbf{W}, \mathbf{Z}) is equivalent to the transformation

$$\underline{\underline{\mathbf{x}}}_{nT \times 1} \rightarrow \begin{pmatrix} \underline{\underline{\mathbf{z}}} \\ \underline{\underline{\mathbf{w}}} \end{pmatrix}_{\substack{\sum_{r=1}^S m_r T_r \times 1 \\ \sum_{r=1}^S n_r T_r \times 1}}. \quad (4.2)$$

Transformation (4.2) is necessarily accomplished in two stages. The first of these is “aggregation”; that is,

¹⁰ In the model under consideration at least 10^4 imputation draws in each iteration were needed to achieve an adequately accurate MC expectation.

¹¹ The “observation-by-observation” operator “*vecr*” is here preferred to the more commonly employed “*vec*” as it allows us to extrapolate easily from any given subperiod r to the complete sample.

$$\underline{\mathbf{x}}^\dagger = \tilde{\mathbf{B}} \underline{\mathbf{x}},$$

where

$$\tilde{\mathbf{B}}_{nT \times nT} = \begin{bmatrix} \mathbf{I}_{T_1} \otimes \mathbf{B}_1 & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & \mathbf{I}_{T_S} \otimes \mathbf{B}_S \end{bmatrix}.$$

$\tilde{\mathbf{B}} \underline{\mathbf{x}}$ yields a nT -vector $\underline{\mathbf{x}}^\dagger$ comprised explicitly of the “observed” and “missing” data; but still co-mingled together, as we see below. The second stage, from $\underline{\mathbf{x}}^\dagger$ to $\underline{\mathbf{x}}^\ddagger$, therefore requires a permutation operation, in which the missing data is moved to the “top” of the vector. This is accomplished by the operation

$$\underline{\mathbf{x}}^\ddagger = \tilde{\mathbf{P}} \underline{\mathbf{x}}^\dagger = \begin{pmatrix} \underline{\mathbf{z}} \\ \underline{\mathbf{w}} \end{pmatrix};$$

where

$$\tilde{\mathbf{P}}_{nT \times nT} = \begin{pmatrix} \tilde{\mathbf{P}}_z \\ \tilde{\mathbf{P}}_w \end{pmatrix} \begin{matrix} \sum_{r=1}^S m_r T_r \times nT \\ \sum_{r=1}^S n_r T_r \times nT \end{matrix},$$

$$\tilde{\mathbf{P}}_z = \begin{bmatrix} \mathbf{I}_{T_1} \otimes \mathbf{J}_1 & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & \mathbf{I}_{T_S} \otimes \mathbf{J}_S \end{bmatrix},$$

and

$$\tilde{\mathbf{P}}_w = \begin{bmatrix} \mathbf{I}_{T_1} \otimes \bar{\mathbf{J}}_1 & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & \mathbf{I}_{T_S} \otimes \bar{\mathbf{J}}_S \end{bmatrix},$$

in which \mathbf{J}_r and $\bar{\mathbf{J}}_r$ denote the $m_r \times n$ and $n_r \times n$ “selection matrices” $\mathbf{J}_r = \begin{bmatrix} \mathbf{I}_{m_r} & \mathbf{0} \end{bmatrix}$ and $\bar{\mathbf{J}}_r = \begin{bmatrix} \mathbf{0} & \mathbf{I}_{n_r} \end{bmatrix}$, respectively, $r = 1, \dots, S$.

It can be confirmed, by direct computation, that $\underline{\mathbf{x}}^\ddagger = \tilde{\mathbf{P}} \underline{\mathbf{x}}^\dagger = \tilde{\mathbf{P}} \tilde{\mathbf{B}} \underline{\mathbf{x}}$ does indeed equal $(\underline{\mathbf{z}}' \ \underline{\mathbf{w}}')'$. That is, with $\underline{\mathbf{x}} = \text{vectr}(\underline{\mathbf{X}})_{T \times n}$, $\underline{\mathbf{x}}^\dagger = \tilde{\mathbf{B}} \underline{\mathbf{x}}$, and $\underline{\mathbf{x}}^\ddagger = \tilde{\mathbf{P}} \underline{\mathbf{x}}^\dagger$, we have

$$\begin{aligned}
\tilde{\mathbf{x}} = & \begin{pmatrix} \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_{T_1} \\ \vdots \\ \mathbf{x}_{T_1+T_2} \\ \vdots \\ \mathbf{x}_{1+\sum_r^{s-1} T_r} \\ \vdots \\ \mathbf{x}_T \end{pmatrix}_{n \times 1}, \quad \tilde{\mathbf{x}}^\dagger = \begin{pmatrix} \mathbf{z}_1 \\ \mathbf{w}_1 \\ \vdots \\ \mathbf{z}_{T_1} \\ \mathbf{w}_{T_1} \\ \vdots \\ \mathbf{z}_{T_1+T_2} \\ \mathbf{w}_{T_1+T_2} \\ \vdots \\ \mathbf{z}_{1+\sum_r^{s-1} T_r} \\ \mathbf{w}_{1+\sum_r^{s-1} T_r} \\ \vdots \\ \mathbf{z}_T \\ \mathbf{w}_T \end{pmatrix}_{m_1 \times 1}, \quad \text{and } \tilde{\mathbf{x}}^\dagger = \begin{pmatrix} \mathbf{z}_1 \\ \vdots \\ \mathbf{z}_{T_1} \\ \mathbf{z}_{T_1+T_2} \\ \vdots \\ \mathbf{z}_{1+\sum_r^{s-1} T_r} \\ \mathbf{z}_T \\ \mathbf{w}_1 \\ \vdots \\ \mathbf{w}_{T_1} \\ \mathbf{w}_{T_1+T_2} \\ \vdots \\ \mathbf{w}_{1+\sum_r^{s-1} T_r} \\ \vdots \\ \mathbf{w}_T \end{pmatrix}_{m_1 \times 1} = \begin{pmatrix} \tilde{\mathbf{z}}_1 \\ \tilde{\mathbf{z}}_2 \\ \vdots \\ \tilde{\mathbf{z}}_S \\ \tilde{\mathbf{w}}_1 \\ \tilde{\mathbf{w}}_2 \\ \vdots \\ \tilde{\mathbf{w}}_S \end{pmatrix}.
\end{aligned}$$

Now, $\tilde{\mathbf{P}}$ is clearly $nT \times nT$ orthogonal, implying $|\tilde{\mathbf{P}}| = 1$. Furthermore, it is generally possible to order our commodities so that $\mathbf{A}_r = [\mathbf{L}_r \quad \mathbf{I}_{n_r}] \quad \forall r = 1, \dots, S$ with \mathbf{L}_r an $n_r \times m_r$ matrix of ones and zeros. It therefore follows that

$$\mathbf{B}_r = \begin{bmatrix} \mathbf{I}_{m_r} & \mathbf{0} \\ \mathbf{L}_r & \mathbf{I}_{n_r} \end{bmatrix}, \quad \mathbf{B}_r^{-1} = \begin{bmatrix} \mathbf{I}_{m_r} & \mathbf{0} \\ -\mathbf{L}_r & \mathbf{I}_{n_r} \end{bmatrix}, \quad \text{and } |\mathbf{B}_r| = 1,$$

implying $|\tilde{\mathbf{B}}| = 1$ also. The density of $\tilde{\mathbf{x}}^\dagger \triangleq \begin{pmatrix} \mathbf{z} \\ \tilde{\mathbf{w}} \end{pmatrix} = \tilde{\mathbf{P}}\tilde{\mathbf{B}}\tilde{\mathbf{x}}$ is therefore simply

$$\mathcal{P}(\tilde{\mathbf{x}}^\dagger | \omega) = (2\pi)^{-nT/2} |\Sigma|^{-7/2} \exp\left\{-\frac{1}{2}(\tilde{\mathbf{x}}^\dagger - \underline{\mu}_{\tilde{\mathbf{x}}^\dagger}(\theta))' \tilde{\Omega}^{-1}(\tilde{\mathbf{x}}^\dagger - \underline{\mu}_{\tilde{\mathbf{x}}^\dagger}(\theta))\right\}, \quad (4.3)$$

with mean
$$\tilde{\boldsymbol{\mu}}_x^\dagger = \tilde{\mathbf{P}}\tilde{\mathbf{B}}\tilde{\boldsymbol{\mu}}_x \equiv \begin{pmatrix} \tilde{\boldsymbol{\mu}}_z \\ \tilde{\boldsymbol{\mu}}_w \end{pmatrix}$$

and variance
$$\tilde{\boldsymbol{\Omega}} = \tilde{\mathbf{P}}\tilde{\mathbf{B}}[\mathbf{I}_T \otimes \boldsymbol{\Sigma}]\tilde{\mathbf{B}}'\tilde{\mathbf{P}}' = \begin{bmatrix} \tilde{\boldsymbol{\Omega}}_{zz} & \tilde{\boldsymbol{\Omega}}_{zw} \\ \tilde{\boldsymbol{\Omega}}_{wz} & \tilde{\boldsymbol{\Omega}}_{ww} \end{bmatrix}.$$

The conditional pdf of the missing data, $\mathcal{P}(\mathbf{Z}|\mathbf{W}, \omega)$, is now just

$$\mathcal{P}(\tilde{\mathbf{z}}|\tilde{\mathbf{w}}, \omega) = (2\pi)^{-\frac{1}{2}\sum_r^{s_r} m_r T_r} \left| (\tilde{\boldsymbol{\Omega}}^{-1})_{zz} \right|^{1/2} \exp\left\{-\frac{1}{2}(\tilde{\mathbf{z}} - \tilde{\boldsymbol{\mu}}_{z|w})' (\tilde{\boldsymbol{\Omega}}^{-1})_{zz} (\tilde{\mathbf{z}} - \tilde{\boldsymbol{\mu}}_{z|w})\right\}, \quad (4.4)$$

where
$$\tilde{\boldsymbol{\mu}}_{z|w} = \tilde{\boldsymbol{\mu}}_{z|w}(\omega) = \tilde{\boldsymbol{\mu}}_z(\theta) + \tilde{\boldsymbol{\Omega}}_{zw}\tilde{\boldsymbol{\Omega}}_{ww}^{-1}(\tilde{\mathbf{w}} - \tilde{\boldsymbol{\mu}}_w(\theta)) \quad (4.5)$$

is the conditional mean of \mathbf{z} given \mathbf{w} and $\omega = (\theta, \boldsymbol{\Sigma})$, and

$$(\tilde{\boldsymbol{\Omega}}^{-1})_{zz} = \left(\tilde{\boldsymbol{\Omega}}_{zz} - \tilde{\boldsymbol{\Omega}}_{zw}\tilde{\boldsymbol{\Omega}}_{ww}^{-1}\tilde{\boldsymbol{\Omega}}_{wz} \right)^{-1}$$

is its conditional variance.

On the computational side, since $\tilde{\mathbf{P}}^{-1} = \tilde{\mathbf{P}}'$ it follows that

$$\tilde{\boldsymbol{\Omega}}^{-1} = \tilde{\mathbf{P}}\tilde{\mathbf{B}}'^{-1}[\mathbf{I}_T \otimes \boldsymbol{\Sigma}^{-1}]\tilde{\mathbf{B}}^{-1}\tilde{\mathbf{P}}' = \begin{bmatrix} \tilde{\mathbf{P}}_z\tilde{\mathbf{B}}'^{-1}[\mathbf{I}_T \otimes \boldsymbol{\Sigma}^{-1}]\tilde{\mathbf{B}}^{-1}\tilde{\mathbf{P}}'_z & \tilde{\mathbf{P}}_z\tilde{\mathbf{B}}'^{-1}[\mathbf{I}_T \otimes \boldsymbol{\Sigma}^{-1}]\tilde{\mathbf{B}}^{-1}\tilde{\mathbf{P}}'_w \\ \tilde{\mathbf{P}}_w\tilde{\mathbf{B}}'^{-1}[\mathbf{I}_T \otimes \boldsymbol{\Sigma}^{-1}]\tilde{\mathbf{B}}^{-1}\tilde{\mathbf{P}}'_z & \tilde{\mathbf{P}}_w\tilde{\mathbf{B}}'^{-1}[\mathbf{I}_T \otimes \boldsymbol{\Sigma}^{-1}]\tilde{\mathbf{B}}^{-1}\tilde{\mathbf{P}}'_w \end{bmatrix};$$

implying
$$(\tilde{\boldsymbol{\Omega}}^{-1})_{zz} = \tilde{\mathbf{P}}_z\tilde{\mathbf{B}}'^{-1}[\mathbf{I}_T \otimes \boldsymbol{\Sigma}^{-1}]\tilde{\mathbf{B}}^{-1}\tilde{\mathbf{P}}'_z.$$

The latter simplifies, on partitioning \mathbf{B}_r^{-1} as $[\bar{\mathbf{L}}_r' \quad \bar{\mathbf{J}}_r']$ with $\bar{\mathbf{L}}_r = [\mathbf{I}_{m_r} \quad -\mathbf{L}_r']$, as the block-diagonal matrix

$$(\tilde{\boldsymbol{\Omega}}^{-1})_{zz} = \begin{bmatrix} \mathbf{I}_{T_1} \otimes \bar{\mathbf{L}}_1 \boldsymbol{\Sigma}^{-1} \bar{\mathbf{L}}_1' & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{T_s} \otimes \bar{\mathbf{L}}_s \boldsymbol{\Sigma}^{-1} \bar{\mathbf{L}}_s' \end{bmatrix}. \quad (4.6)$$

Similarly, substituting $\underline{\mu}_z = \tilde{\mathbf{P}}_z \tilde{\mathbf{B}} \underline{\mu}_x \equiv \tilde{\mathbf{P}}_z \underline{\mu}_x$ and $\underline{\mu}_w = \tilde{\mathbf{P}}_w \tilde{\mathbf{B}} \underline{\mu}_x$ into (4.5) and exploiting the usual partition of \mathbf{B}_r (see §3.1) yields

$$\underline{\mu}_{z|w} = \tilde{\mathbf{D}}(\Sigma) \underline{\mu}_x(\theta) + \tilde{\mathbf{G}}(\Sigma) \underline{w};$$

where

$$\tilde{\mathbf{G}}(\Sigma) = \begin{bmatrix} \mathbf{I}_{T_1} \otimes \mathbf{J}_1 \Sigma \mathbf{A}'_1 (\mathbf{A}_1 \Sigma \mathbf{A}'_1)^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{T_s} \otimes \mathbf{J}_s \Sigma \mathbf{A}'_s (\mathbf{A}_s \Sigma \mathbf{A}'_s)^{-1} \end{bmatrix},$$

and

$$\tilde{\mathbf{D}}(\Sigma) = \begin{bmatrix} \mathbf{I}_{T_1} \otimes (\mathbf{J}_1 - \mathbf{J}_1 \Sigma \mathbf{A}'_1 (\mathbf{A}_1 \Sigma \mathbf{A}'_1)^{-1} \mathbf{A}_1) & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{T_s} \otimes (\mathbf{J}_s - \mathbf{J}_s \Sigma \mathbf{A}'_s (\mathbf{A}_s \Sigma \mathbf{A}'_s)^{-1} \mathbf{A}_s) \end{bmatrix}.$$

4.2 E-step: the expected log-likelihood

We are now in a position to establish the following closed-form solution for $Q(\omega, \omega^*)$.

PROPOSITION 1. Let $\underline{\mathbf{x}}_a^*$ denote the “mean-augmented” expenditure set

$$\underline{\mathbf{x}}_a^* \equiv \underline{\mathbf{x}}_a(\omega^*) = \begin{pmatrix} \underline{\mu}_{z|w}^* \\ \underline{w} \end{pmatrix},$$

where $\underline{\mu}_{z|w}^* = \underline{\mu}_{z|w}(\omega^*) = \tilde{\mathbf{D}}(\Sigma^*) \underline{\mu}_x(\theta^*) + \tilde{\mathbf{G}}(\Sigma^*) \underline{w}.$

Also let $\underline{\mathbf{u}}_a^*$ denote the vector of “augmented residuals”

$$\underline{\mathbf{u}}_a^* = (\tilde{\mathbf{P}}\tilde{\mathbf{B}})^{-1}(\underline{\mathbf{x}}_a^* - \underline{\mu}_x^*) = (\tilde{\mathbf{P}}\tilde{\mathbf{B}})^{-1}\underline{\mathbf{x}}_a^* - \underline{\mu}_x(\theta),$$

and let $\mathbf{U}^*(\theta)$ denote the $T \times n$ matrix constructed such that $\text{vecr}(\mathbf{U}^*) = \underline{\mathbf{u}}_a^*.$

Finally, let $\mathbf{U}^*{}'\mathbf{U}^* = \mathbf{S}^*(\theta)$ and $\bar{\mathbf{L}}_r'(\bar{\mathbf{L}}_r \Sigma^{*-1} \bar{\mathbf{L}}_r')^{-1} \bar{\mathbf{L}}_r = \mathbf{V}_r^*$. Then the EM criterion for the problem at hand is

$$Q(\omega, \omega^*) = -\frac{T}{2} \ln |\Sigma| - \frac{1}{2} \text{tr} \{ \Sigma^{-1} \tilde{\mathbf{S}}^*(\theta) \}; \quad (4.7)$$

where

$$\tilde{\mathbf{S}}^*(\theta) = \mathbf{S}^*(\theta) + \sum_{r=1}^S T_r \mathbf{V}_r^*.$$

Proof of the proposition requires the following preliminary Lemma.

LEMMA 1. Let \mathbf{x} denote an n -vector of independent variables partitioned as $\begin{pmatrix} \mathbf{z} \\ \mathbf{w} \end{pmatrix}$, where \mathbf{z} is $m \times 1$, $m < n$. Let the n -vector \mathbf{a} and the $n \times n$ symmetric matrix \mathbf{A} be partitioned conformably; ie; $\mathbf{a} = \begin{pmatrix} \mathbf{a}_z \\ \mathbf{a}_w \end{pmatrix}$, and $\mathbf{A} = \begin{bmatrix} \mathbf{A}_{zz} & \mathbf{A}_{zw} \\ \mathbf{A}_{wz} & \mathbf{A}_{ww} \end{bmatrix}$. Then, for \mathbf{b} an m -vector of constants, and matrix \mathbf{B} $m \times m$ symmetric positive definite:

$$\begin{aligned} \int \cdots \int (\mathbf{x} - \mathbf{a})' \mathbf{A} (\mathbf{x} - \mathbf{a}) e^{-(\mathbf{z} - \mathbf{b})' \mathbf{B} (\mathbf{z} - \mathbf{b}) / 2} dz_1 \cdots dz_m \\ = (2\pi)^{m/2} |\mathbf{B}|^{-1/2} \left\{ \text{tr}(\mathbf{A}_{zz} \mathbf{B}^{-1}) + (\mathbf{b}^* - \mathbf{a})' \mathbf{A} (\mathbf{b}^* - \mathbf{a}) \right\}; \end{aligned}$$

where $\mathbf{b}^* = \begin{pmatrix} \mathbf{b} \\ \mathbf{w} \end{pmatrix}$.

PROOF OF LEMMA 1: The result proceeds directly from Graybill (1983, p. 342, Thm. 10.5.1), on rewriting the quadratic forms involved as

$$\begin{aligned} (\mathbf{x} - \mathbf{a})' \mathbf{A} (\mathbf{x} - \mathbf{a}) &= (\mathbf{z} - \mathbf{a}_z)' \mathbf{A}_{zz} (\mathbf{z} - \mathbf{a}_z) + 2(\mathbf{z} - \mathbf{a}_z)' \mathbf{A}_{zw} (\mathbf{w} - \mathbf{a}_w) \\ &\quad + (\mathbf{w} - \mathbf{a}_w)' \mathbf{A}_{ww} (\mathbf{w} - \mathbf{a}_w) \\ &= \mathbf{z}' \mathbf{A}_{zz} \mathbf{z} + 2\mathbf{z}' \{ \mathbf{A}_{zw} (\mathbf{w} - \mathbf{a}_w) - \mathbf{A}_{zz} \mathbf{a}_z \} \\ &\quad + \{ \mathbf{a}_z' \mathbf{A}_{zz} \mathbf{a}_z - 2\mathbf{a}_z' \mathbf{A}_{zw} (\mathbf{w} - \mathbf{a}_w) + (\mathbf{w} - \mathbf{a}_w)' \mathbf{A}_{ww} (\mathbf{w} - \mathbf{a}_w) \} \end{aligned}$$

and

$$(\mathbf{z} - \mathbf{b})' \mathbf{B} (\mathbf{z} - \mathbf{b}) / 2 = \frac{1}{2} \mathbf{z}' \mathbf{B} \mathbf{z} - \mathbf{z}' \mathbf{B} \mathbf{b} + \frac{1}{2} \mathbf{b}' \mathbf{B} \mathbf{b}.$$

PROOF OF PROPOSITION 1.

In the notation of §4.1 the EM criterion $Q(\omega, \omega^*)$ is now

$$Q(\omega, \omega^*) \triangleq \int_{\mathbf{z}} \ell(\omega | \mathbf{x}^\dagger) \mathcal{P}(\mathbf{z} | \mathbf{w}, \omega^*) d\mathbf{z},$$

with $\ell(\omega | \mathbf{x}^\dagger) = \ln \mathcal{P}(\mathbf{x}^\dagger | \omega)$. This becomes, on substituting (4.4) and (4.3),

$$Q(\omega, \omega^*) = c - \frac{T}{2} \ln |\Sigma| - \frac{1}{2} (2\pi)^{-\frac{1}{2} \sum_r m_r T_r} \left| (\tilde{\Omega}^{-1})_{zz}^* \right|^{\frac{1}{2}} \mathcal{J}(\omega, \omega^*),$$

with $c = -nT \ln(2\pi)/2$, $(\tilde{\Omega}^{-1})_{zz}^* = (\tilde{\Omega}^{-1})_{zz}|_{\omega^*}$, and

$$\mathcal{J}(\omega, \omega^*) = \int_{\mathbf{z}} (\mathbf{x}^\dagger - \mu_{\mathbf{x}}^\dagger)' \tilde{\Omega}^{-1} (\mathbf{x}^\dagger - \mu_{\mathbf{x}}^\dagger) \exp \left\{ -\frac{1}{2} (\mathbf{z} - \mu_{\mathbf{z}|\mathbf{w}}^*)' (\tilde{\Omega}^{-1})_{zz}^* (\mathbf{z} - \mu_{\mathbf{z}|\mathbf{w}}^*) \right\} d\mathbf{z}.$$

Applying Lemma 1 to $\mathcal{J}(\omega, \omega^*)$ then yields

$$(2\pi)^{-\frac{1}{2} \sum_r m_r T_r} \left| (\tilde{\Omega}^{-1})_{zz}^* \right|^{\frac{1}{2}} \mathcal{J}(\omega, \omega^*) = \text{tr} \left\{ (\tilde{\Omega}^{-1})_{zz} (\tilde{\Omega}^{-1})_{zz}^{*-1} \right\} + (\mathbf{x}_a^* - \mu_{\mathbf{x}}^\dagger)' \tilde{\Omega}^{-1} (\mathbf{x}_a^* - \mu_{\mathbf{x}}^\dagger),$$

where the quadratic form simplifies as

$$\begin{aligned} (\mathbf{x}_a^* - \mu_{\mathbf{x}}^\dagger)' \tilde{\Omega}^{-1} (\mathbf{x}_a^* - \mu_{\mathbf{x}}^\dagger) &= (\mathbf{x}_a^* - \mu_{\mathbf{x}}^\dagger)' (\tilde{\mathbf{P}}\tilde{\mathbf{B}})^{-1} (\mathbf{I}_T \otimes \Sigma^{-1}) (\tilde{\mathbf{P}}\tilde{\mathbf{B}})^{-1} (\mathbf{x}_a^* - \mu_{\mathbf{x}}^\dagger) \\ &= \mathbf{u}_a'^* (\mathbf{I}_T \otimes \Sigma^{-1}) \mathbf{u}_a^* \\ &\equiv \text{tr}(\Sigma^{-1} \mathbf{U}^*{}' \mathbf{U}^*). \end{aligned}$$

Furthermore, $(\tilde{\Omega}^{-1})_{zz}$ is block-diagonal, with typical block $\mathbf{I}_T \otimes \bar{\mathbf{L}}_r \Sigma^{-1} \bar{\mathbf{L}}_r'$, so that

$$\text{tr} \left\{ (\tilde{\Omega}^{-1})_{zz} (\tilde{\Omega}^{-1})_{zz}^{*-1} \right\} = \sum_{r=1}^S T_r \text{tr} \left(\bar{\mathbf{L}}_r \Sigma^{-1} \bar{\mathbf{L}}_r' (\bar{\mathbf{L}}_r \Sigma^{-1} \bar{\mathbf{L}}_r')^{-1} \right).$$

Thus we finally have (ignoring the constant)

$$Q(\omega, \omega^*) = -\frac{T}{2} \ln |\Sigma| - \frac{1}{2} \text{tr}(\Sigma^{-1} \mathbf{U}^{*'} \mathbf{U}^*) - \frac{1}{2} \sum_{r=1}^S T_r \text{tr}(\Sigma^{-1} \bar{\mathbf{L}}_r' (\bar{\mathbf{L}}_r \Sigma^{*-1} \bar{\mathbf{L}}_r')^{-1} \bar{\mathbf{L}}_r).$$

Collecting the trace terms then establishes the Proposition.

4.3 M-step: the concentrated expected log-likelihood

It is evident, from (4.7), that $Q(\omega, \omega^*)$ takes the familiar form of a Gaussian “log-likelihood”. The FOC for Σ will, accordingly, be of the form (3.4), implying

$$\hat{\Sigma} = T^{-1} \tilde{\mathbf{S}}^*(\theta).$$

Hence we can, as before, construct a *concentrated* expected log-likelihood

$$Q_c(\theta, \omega^*) = -\frac{T}{2} \ln |\tilde{\mathbf{S}}^*(\theta)|, \quad (4.8)$$

requiring maximization with respect to θ only. We now have the EM algorithm for the maximization of the observed log-likelihood in a variably-aggregated demand system as simply:

For $i = 0, 1, 2, \dots$, some starting value $\hat{\omega}^{(0)} = (\hat{\theta}^{(0)}, \hat{\Sigma}^{(0)})$, and prespecified tolerance ε :

- (1) Compute $\mu_{z|w}(\hat{\omega}^{(i)})$, and hence $\mathbf{x}_a^{(i)} = \mathbf{x}_a(\hat{\omega}^{(i)})$.
- (2) Construct $\mathbf{u}_a^{(i)}(\theta) = \mathbf{u}_a(\theta, \hat{\omega}^{(i)})$, $\tilde{\mathbf{S}}^{(i)}(\theta) = \tilde{\mathbf{S}}(\theta, \hat{\omega}^{(i)})$ and hence $Q_c(\theta, \hat{\omega}^{(i)})$.
- (3) Maximize $Q_c(\theta, \hat{\omega}^{(i)})$ with respect to θ , so obtaining $\hat{\theta}^{(i+1)}$, $\tilde{\mathbf{S}}(\hat{\theta}^{(i+1)}, \hat{\omega}^{(i)})$, and hence $\hat{\Sigma}^{(i+1)} = T^{-1} \tilde{\mathbf{S}}(\hat{\theta}^{(i+1)}, \hat{\omega}^{(i)})$.
- (4) Repeat (1) – (3) until $|Q_c(\hat{\omega}^{(i+1)}, \hat{\omega}^{(i)}) - Q_c(\hat{\omega}^{(i)}, \hat{\omega}^{(i)})| < \varepsilon$.

4.4 Practical issues

4.4.1 No missing data in the final subperiod

A few practical issues now arise. To begin with, notice that, in contrast to the treatment of GM, there is no explicit need, in our EM setup, for a complete set of expenditure data to exist in one of the subperiods. The only obvious limit on the number of expenditure categories that could be “modelled” in this way is the availability of the corresponding price data. In practice, as one might expect, there is a limit on the amount of data which can be said to be “missing”, in this case imposed by the effect on the rate of convergence of increasing the proportion of “missing” to “observed”.

Typically, however, the observed expenditures will be held to be “complete” for one of the subperiods, generally the last; in which case $n_s = n$, $m_s = 0$, $\mathbf{B}_s = \mathbf{I}_n$, the lower-right submatrix of $\tilde{\mathbf{P}}_w$ reduces to the $nT_s \times nT_s$ identity matrix, while the corresponding rows of $\tilde{\mathbf{P}}_z$, and subsequent matrices such as $\tilde{\mathbf{D}}$ and $\tilde{\mathbf{G}}$, disappear entirely. Thus, for the 3-subperiod example of §5 following we will have

$$\tilde{\mathbf{P}}_z = \begin{bmatrix} \mathbf{I}_{T_1} \otimes \mathbf{J}_1 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{T_2} \otimes \mathbf{J}_2 & \mathbf{0} \end{bmatrix};$$

and

$$\tilde{\mathbf{P}}_w = \begin{bmatrix} \mathbf{I}_{T_1} \otimes \bar{\mathbf{J}}_1 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{T_2} \otimes \bar{\mathbf{J}}_2 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I}_{nT_3} \end{bmatrix}.$$

4.4.2 Standard errors

We now turn to the matter of obtaining appropriate standard errors at the end of the EM estimation. It can be shown (see Tanner (1993), p.47) that the “observed” Hessian,

$$H(\omega | \mathbf{W}) = \frac{\partial^2 \ell(\omega | \mathbf{W})}{\partial \omega \partial \omega'},$$

is the sum of the expected augmented Hessian and the variance of the augmented score; ie,

$$\begin{aligned}
H(\omega|\mathbf{W}) &= E_{\mathbf{Z}|\mathbf{W},\omega}\{H(\omega|\mathbf{X})\} + V_{\mathbf{Z}|\mathbf{W},\omega}\{q(\omega|\mathbf{X})\} \\
&\equiv \left. \frac{\partial^2 Q(\omega, \omega^*)}{\partial \omega \partial \omega'} \right|_{\omega=\omega^*} + E_{\mathbf{Z}|\mathbf{W},\omega}\{q(\omega|\mathbf{X})q(\omega|\mathbf{X})'\};
\end{aligned}$$

so that the Hessian matrix naturally returned at the end of the EM estimation need only be “adjusted” by the expectation, taken with respect to $\mathcal{P}(\mathbf{Z}|\mathbf{W}, \hat{\omega})$, of the outer product of the score, also evaluated at $\hat{\omega}$. The latter, in our case, is most simply “estimated” by the MC method discussed in §3.2. That is, we generate, with $\omega = \hat{\omega}$, M augmented expenditure sets \mathbf{X}_j , compute the corresponding $q(\omega|\mathbf{X}_j)q(\omega|\mathbf{X}_j)'$, and obtain the average of the latter. $q(\omega|\mathbf{X})$ would in this case be given by Grose and McLaren (1999) equations (2.4) and (A.1).

Alternatively, since we in fact already have a functional form for $\ell(\omega|\mathbf{W})$ (equation (2.7)), the observed Hessian $H(\omega|\mathbf{W})$ could (in theory) be obtained directly, either by differentiation (analytic or numeric) of the analytic observed score $q(\omega|\mathbf{W})$ ¹², or by simply computing the numeric second derivatives of $\ell(\omega|\mathbf{W})$.

4.4.3 Choice of demand system for estimation

A final practical issue involves the feasibility of estimating any given demand system with the available data, even after augmentation. In short, for the algorithm to be feasible it must be possible to maximize the full N -commodity likelihood using the augmented dataset. In our case, with up to 18 commodities, this means that we must restrict our attention to either the extremely parsimonious LES, or to models, such as linearized AIDS, that can be estimated by iterative GLS (see Grose and McLaren (2000)). It must be noted, however, that the latter is already a fairly computer intensive procedure, and its insertion into an EM “loop” would be time consuming in the extreme.

¹² The derivatives of the “aggregated” log-likelihood with respect to Σ and θ are given in GM.

5. EM estimation of the LES.

The linear expenditure system (LES) was estimated for Grose and McLaren (1999)'s quarterly data with 3 subperiods of differing aggregation. The data consists of 95 quarterly observations on up to 18 expenditure categories, covering the period 1974(3rd quarter) – 1998(1st quarter). The 1st subperiod (1974(3) – 1985(3)) comprises 45 quarterly observations on 12 expenditure categories; the 2nd (1985(4) – 1989(3)) involves 16 quarterly observations on 16 categories; and the 3rd (1989(4) – 1998(1)) involves 34 quarterly observations on 18 categories. The 12, 16, and 18 commodity sets are as listed in Grose and McLaren (1999, §4.3). All expenditures are in A\$ per capita; prices are measured by the IPD for each expenditure category, and equal unity in 1989/90. The “omitted” category in all subperiods was, once again, Food.

The LES with “free” covariance matrix was estimated by EM as described above, with the expectation stage performed analytically as per §4, and the maximization stage performed by numerically maximizing (using Newton-Raphson) the expected log-likelihood with respect to $\theta = (\tilde{\gamma}, \beta)$ after concentrating Σ out as per §4.3.

An attempt to check for a global maximum¹³ was made by starting the algorithm at different starting values. For the first “pass”, the EM algorithm was started at $\tilde{\gamma} = \mathbf{0}$, $\beta = T_3^{-1} \mathbf{W}_3' \mathbf{1}_{T_3} = \bar{\mathbf{W}}_3$ (the mean expenditure shares in the final subperiod), and $\Sigma = T_3^{-1} \text{diag}((\mathbf{W}_3 - \bar{\mathbf{W}}_3)'(\mathbf{W}_3 - \bar{\mathbf{W}}_3))$; that is, the initial Σ is a diagonal matrix, with elements set equal to the variance of the corresponding expenditure share in the final subperiod. Results are given in Table 2, with standard errors estimated as per §4.4.2.

For the second pass we started the EM algorithm at the MLE's obtained by maximizing the observed (“aggregated”) log-likelihood assuming De Boer and Harkema's covariance matrix as per Grose and McLaren (1999). While the algorithm now took slightly longer to converge, and resulted in a slightly smaller maximized log-observed likelihood value, the parameter estimates obtained were very similar. By way of a final comparison, the algorithm was repeated, this time starting at the MLE's of θ obtained using the “complete” $T \times N$ expenditure set recently made available by

¹³ While it can be shown that EM will converge to a maximum of the expected log-likelihood, this is not necessarily a global maximum.

the ABS. The parameter estimates obtained were once again very similar to those obtained from the mean expenditure share starters.

Conclusion

The problem considered here is the ML estimation of a consumer demand system in the situation where not all expenditures are observed for all commodity categories in all time periods. The major problem with the estimation of such a system is that, while the likelihood function can be written down simply enough, it cannot be maximized directly because of the $\frac{1}{2}n(n+1)$ covariance parameters that must now also be included in the objective function. It may be worth noting that the complete log-likelihood cannot be satisfactorily maximized even with 95 observations available on all 18 expenditure categories, *unless* Σ is concentrated out.

The strategy considered in this paper takes a “data augmentation” approach to our unobserved expenditure data. This leads to an EM method, in which we maximize the *expectation* of the complete log-likelihood, with the expectation taken with respect to the conditional density of the missing data. This has the advantage that it effectively reformulates the optimization problem such that Σ *can* be concentrated out; and thus avoids the need to restrictively reparameterize the covariance matrix. We find that convergence, while slow, seems quite sure; in contrast to earlier attempts to maximize the observed log-likelihood directly. Furthermore, starting the algorithm at quite different parameter values did not result in substantially different estimates, giving us some confidence in concluding that the mode obtained is global, and hence the MLE.

The prime disadvantage of the EM approach is the slow rate of convergence, which would become even slower if the proportion of missing data were increased. This means that, while it may be feasible to estimate a model such as AIDS, the time taken for each M-step if we do may make the procedure impractical for routine application. Finally, EM naturally requires sufficient data to permit conventional maximization at each iteration; and so cannot be applied to very small (eg. annual) datasets any more than can ordinary ML. In other words, it is worth checking whether the chosen model can actually be estimated at the desired level of commodity disaggregation, before

considering the additional complication of having some expenditure categories partially unobserved.

On the other hand, if *additional* data is available at a coarser level of commodity aggregation then the EM algorithm presented here may well represent a convenient means of estimating the model for the full commodity set that avoids the need to resort to a more restrictive respecification.

References

De Boer, P.M.C. and R. Harkema (1986), “Maximum likelihood estimation of sum-constrained linear models with insufficient observations”. *Economics Letters*, 20, 325–329.

Dempster, A.P., N.M. Laird and D.B. Rubin (1977), “Maximum likelihood from incomplete data via the EM algorithm”. *Journal of the Royal Statistical Society B*, 39, 1–38.

Fry, J.M., T.R.L. Fry and K.R. McLaren (1996), “The Stochastic Specification of Demand Share Equations: Restricting Budget Shares to the Unit Simplex”. *Journal of Econometrics*, 73(2), 377–385.

Graybill, F.A. (1983), *Matrices with Applications in Statistics*, 2nd ed. Wadsworth, Belmont California.

Grose, S.D. and K.R. McLaren (1999), “Estimating consumer demand with missing expenditure data”. *Econometric Society Australasian Meeting 1999, Proceedings*, UTS.

Grose, S.D. and K.R. McLaren (2000), “SUR estimation of AIDS”. *Working Paper, Econometrics and Business Statistics, Monash University*, Melbourne.

McLachlan, G.J. and T. Krishnan (1997), *The EM Algorithm and Extensions*. Wiley, New York.

Tanner, M.A. (1993), *Tools for Statistical Analysis*. Springer-Verlag, New York.

Table 1. Abbreviations for expenditure categories

<i>Food</i>	FOD
<i>Cigarettes and Tobacco</i>	CGT
<i>Alcohol and spirits</i>	ALC
<i>Clothing and footwear</i>	CFF
<i>Dwelling rent</i>	RNT
<i>Purchase of motor vehicles</i>	MVP
<i>Household appliances</i>	HAP
<i>Other household durables</i>	HDU
<i>Postal and telecommunications</i>	TEL
<i>Gas, electricity and fuel</i>	GEF
<i>Fares</i>	FRS
<i>Operation of motor vehicles</i>	MVO
<i>Health</i>	MED
<i>Entertainment and recreation</i>	REC
<i>Financial services</i>	FIN
<i>Other goods</i>	OGD
<i>Other services</i>	OSV
<i>Net expenditure overseas</i>	NEO

Table 2. EM estimation of the LES. Unrestricted covariance matrix; 3 subperiods; 18 expenditure categories; quarterly data

Started from $\gamma = 0$, $\beta = \text{mean expenditure share}$.

Expenditure category	β		γ		σ^\ddagger
	estimate	standard error	estimate [†]	standard error	estimate
FOD	0.1191	0.0049	0.2636	0.0108	0.00472
CGT	0.0263	0.0013	0.0135	0.0018	0.00308
ALC	0.0580	0.0036	0.0455	0.0060	0.00736
CFF	0.0845	0.0052	0.0455	0.0073	0.01184
RNT	0.0408	0.0065	0.4765	0.0100	0.02344
MVP	0.0708	0.0085	-0.0039	0.0163	0.00609
HAP	0.0624	0.0030	-0.0089	0.0033	0.00835
HDU	0.0922	0.0046	-0.0320	0.0079	0.01020
TEL	0.0245	0.0011	0.0047	0.0013	0.00228
GEF	0.0075	0.0037	0.0534	0.0065	0.00369
FRS	0.0336	0.0022	0.0257	0.0041	0.00164
MVO	0.0282	0.0021	0.1541	0.0037	0.00147
MED	0.0220	0.0053	0.1545	0.0082	0.01428
REC	0.0178	0.0043	0.1117	0.0072	0.00975
FIN	-0.0070	0.0036	0.1009	0.0047	0.01297
OGD	0.1710	0.0045	-0.0616	0.0051	0.01453
OSV	0.1879	0.0083	-0.0134	0.0121	0.01842
NEO	-0.0396	0.0032	-0.0014	0.0008	0.01665

Initial observed log-likelihood	2487.26	Final observed log-likelihood	6438.54
Initial expected log-likelihood	7168.01	Final expected log-likelihood	8199.75
Number of EM iterations	769	Final $\ \theta^{(i+1)} - \theta^{(i)}\ $	0.000105
Time to convergence	12.9 minutes	Number of observations	95

[†] For estimation purposes the matrix of price ratios (that is, the ratio of price (an index, =1 in 1989/90) to total expenditure per capita (in Australian \$)) has been scaled up by 10^3 . Estimates of γ are thus in units of thousands of 1989/90\$.

[‡] σ is the square root of the corresponding diagonal element of the cross-commodity covariance matrix.

Table 3. EM estimation of the LES. Unrestricted covariance matrix; 3 subperiods; 18 expenditure categories; quarterly data. Summary statistics.

Started from the MLE's obtained by maximizing the observed ("aggregated") log-likelihood assuming De Boer and Harkema's covariance matrix with quarterly data.

Initial observed log-likelihood	5588.39	Final observed log-likelihood	6434.90
Initial expected log-likelihood	7414.93	Final expected log-likelihood	8189.93
Number of EM iterations	776	Final $\ \theta^{(i+1)} - \theta^{(i)}\ $	0.0001284
Time to convergence	12.75 minutes	Number of observations	95

Table 4. EM estimation of the LES. Unrestricted covariance matrix; 3 subperiods; 18 expenditure categories; quarterly data. Summary statistics.

Started from the MLE's obtained by assuming complete expenditure information available in all subperiods.

Initial observed log-likelihood	6306.13	Final observed log-likelihood	6438.12
Initial expected log-likelihood	7921.85	Final expected log-likelihood	8198.16
Number of EM iterations	520	Final $\ \theta^{(i+1)} - \theta^{(i)}\ $	9.558×10^{-5}
Time to convergence	8.4 minutes	Number of observations	95