



THE UNIVERSITY *of York*

Discussion Papers in Economics

No. 2001/20

Using Simulation-Based Inference with Panel Data in Health Economics

by

Paul Contoyannis, Andrew M Jones and Roberto Leon-Gonzalez

Department of Economics and Related Studies
University of York
Heslington
York, YO10 5DD

USING SIMULATION-BASED INFERENCE WITH PANEL DATA IN HEALTH ECONOMICS

Paul Contoyannis, Andrew M. Jones and Roberto Leon-Gonzalez

*Department of Economics and Related Studies,
University of York,
York, YO10 5DD,
United Kingdom
Fax: +44-1904-433759*

*E-mail: pc107@york.ac.uk, amj1@york.ac.uk, rlg103@york.ac.uk
<http://www.york.ac.uk/res/herc/yshe>*

ABSTRACT

Panel datasets provide a rich source of information for health economists, offering the scope to control for individual heterogeneity and to model the dynamics of individual behaviour. However the qualitative or categorical measures of outcome often used in health economics create special problems for estimating econometric models. Allowing a flexible specification of individual heterogeneity leads to models involving higher order integrals that cannot be handled by conventional numerical methods. The dramatic growth in computing power over recent years has been accompanied by the development of simulation estimators that solve this problem. This review uses binary choice models to show what can be done with conventional methods and how the range of models can be expanded by using simulation methods. Practical applications of the methods are illustrated using data on health from the British Household Panel Survey (BHPS).

Keywords: econometrics, panel data, simulation methods, determinants of health

Acknowledgements: This paper is based on the preconference session ‘Using panel data econometrics in health economics’ presented at the 3rd iHEA Conference, University of York, 21st-25th July 2001. Data from the *British Household Panel Survey (BHPS)* were supplied by the ESRC Data Archive. Neither the original collectors of the data nor the Archive bear any responsibility for the analysis or interpretations presented here. Funding was provided by the ESRC award no: R000238169 – ‘Simulation-based econometric approaches to investigating the interaction of lifestyle and health’. The authors would like to thank Nigel Rice for his advice and comments.

Introduction

Panel datasets, such as the European Community Household Panel (ECHP) and the U.S. PSID and Monitoring the Future panels, provide a rich source of information for health economists. Panel data offer the scope to control for individual heterogeneity and to model the dynamics of individual behaviour. However the measures of outcome used in health economics are often qualitative or categorical. These create special problems for estimating econometric models. Allowing a flexible specification of individual heterogeneity leads to models involving higher order integrals that cannot be handled by conventional numerical methods. The dramatic growth in computing power over recent years has been accompanied by the development of simulation estimators that solve this problem. This review uses binary choice models to show what can be done with conventional methods and how the range of models can be expanded by using simulation methods. Practical applications of the methods are illustrated using data on health from the British Household Panel Survey (the BHPS).

Section 1 gives an overview of binary choice models for panel data and introduces our empirical application to BHPS data for a binary measure of health. It discusses the interpretation of individual effects in panel data models and shows how these can be modelled using the random effects probit model, the conditional logit model and by parameterising the individual effect. Extensions of the random effects model, to allow for serial correlation, can be dealt with by simulation-based inference. Section 2 introduces classical simulation methods. These are designed to approximate higher order integrals and they include the GHK simulator for the truncated multivariate normal distribution. We focus on a particular method of estimation, Maximum Simulated Likelihood (MSL). We present some empirical results and use these to discuss issues that arise in practical applications of MSL. The section concludes with a brief overview of other methods of estimation (MSM, MSS). Section 3 moves to Bayesian MCMC methods. It begins with an introduction to the Bayesian approach to inference before introducing the concept of Markov Chain Monte Carlo (MCMC). Implementation of MCMC involves the use of

Gibbs sampling and Metropolis-Hastings algorithms, along with the use of data augmentation to deal with latent variables. This section concludes with an overview of convergence analysis and methods for model selection and testing.

1. Binary choice models for panel data

1.1 A brief introduction to our model

To illustrate the methods reviewed in this paper we use a panel data model for a binary measure of health applied to data drawn from the British Household Panel Survey (BHPS). The BHPS is a longitudinal survey of private households in Great Britain, with the same respondents questioned each year. The survey contains data on socio-demographic, income and health variables. It is an annual survey of each adult household member (aged 16 and over). The survey was designed to be a nationally representative sample of over 5,000 households, giving approximately 10,000 individual interviewees. The first wave was carried out between 1st September 1990 and 30th April 1991.

Our model applies to a binary dependent variable (“does health limit your daily activities?”). There are repeated measurements for each wave ($t=1, \dots, T$) for a sample of n individuals ($i=1, \dots, n$), and the binary dependent variable y_{it} can be modelled in terms of a continuous latent variable y^*_{it} ,

$$(1) \quad y_{it} = 1(y^*_{it} > 0) = 1(X'_{it}\mathbf{b} + u_{it} > 0)$$

where $1(\cdot)$ is a binary indicator function. X includes variables to capture “permanent” and “transitory” income, measured by the mean of household income across all waves of the panel and deviations around that mean respectively, along with marital status, education and household composition. In our empirical application we restrict the analysis to a sub-

sample of 2,715 men (full details of the sample and variables are given in Contoyannis, Jones and Rice [1]).

The error term u_{it} could be allowed to be freely correlated over time or the correlation structure could be restricted. A common specification is the error components model which splits the error into a time-invariant individual random effect (\mathbf{a}_i) and a time-varying idiosyncratic random error (\mathbf{e}_{it}),

$$(2) \quad y_{it} = 1(y_{it}^* > 0) = 1(X'_{it}\mathbf{b} + \mathbf{a}_i + \mathbf{e}_{it} > 0)$$

The error term could be autocorrelated, for example following an AR(1) process, $\mathbf{e}_{it} = \rho\mathbf{e}_{it-1} + \mathbf{h}_{it}$, or it could be independent over t (giving the random effects model). The simplest possible specification is to assume that the u_{it} are independent over t .

1.2 Individual effects in panel data

To understand the role of individual effects in panel data models, consider the standard linear panel data regression model, in which there are repeated measurements ($t=1, \dots, T$) for a sample of n individuals ($i=1, \dots, n$),

$$(3) \quad y_{it} = X'_{it}\mathbf{b} + u_{it} = X'_{it}\mathbf{b} + \mathbf{a}_i + \mathbf{e}_{it}$$

The presence of \mathbf{a}_i implies clustering within individuals so that a random effects specification can improve the efficiency of the estimates of \mathbf{b} . This stems from the structure imposed on the variance-covariance matrix of the error term,

$$(4) \quad \begin{aligned} \text{Var}[u_{it}] &= E[u_{it}u_{it}] = \mathbf{s}_a^2 + \mathbf{s}_e^2, \quad t=s \\ E[u_{it}u_{is}] &= \mathbf{s}_a^2, \quad t \neq s \end{aligned}$$

These efficiency gains can be exploited by using (4) to construct a generalised least squares (GLS) estimator.

Consistency of the GLS estimator rests on the assumption that the error term is independent of the regressors. Failure to account for correlation between the unobservable individual effects (α) and the regressors (X) will lead to inconsistent estimates of the β s. The least squares dummy variable approach (LSDV) gets around this by conditioning on the individual effects, including a dummy variable for each individual, but this may be prohibitive if there are a large number of cross section observations. Alternatively, the individual effects can be swept from the equation by transforming variables into deviations from their within-group means. Applying least squares to the transformed equation gives the covariance or within-groups estimator of β (CV). Similarly, the model could be estimated in first differences to eliminate the time-invariant individual effects. Identification of β rests on there being sufficient variation within groups. In practice, fixed effects may only work well when there are many observations and much variation within groups.

Now consider a nonlinear model, for example the binary choice model based on the latent variable specification in Equation (2). Assume that the distribution of ε_{it} is symmetric with distribution function $F(\cdot)$. Then,

$$(5) \quad P(y_{it} = 1) = P(e_{it} > -X'_{it}\mathbf{b} - \mathbf{a}_i) = F(X'_{it}\mathbf{b} + \mathbf{a}_i)$$

This illustrates the so-called problem of incidental parameters. As $n \rightarrow \infty$ the number of parameters to be estimated (β, α_i) also grows. In linear models the estimators $\hat{\mathbf{b}}$ and $\hat{\mathbf{a}}$ are asymptotically independent, which means that taking mean deviations or differencing the data allows the derivation of estimators for β whose limits do not depend on $\hat{\mathbf{a}}$. In general, this is not possible in nonlinear models and the inconsistency of estimates of α carries over into the estimates of β . Setting the incidental parameter problem aside, the fixed effect probit model can be estimated by including a dummy variable for each

individual. Heckman [2] presents Monte Carlo evidence that suggests that the small sample bias in the estimates of β is relatively small for values of T of 8 and over. More recently, Greene [3] has championed the use of this 'brute force' approach to fixed effects estimation of nonlinear models.

1.3 Random effects probit model

Assuming that α and ε are normally distributed and independent of X gives the random effects probit model (REP). In this case α can be integrated out to give the sample log-likelihood function,

$$(6) \quad \ln L = \sum_{i=1}^n \left\{ \ln \int_{-\infty}^{+\infty} \prod_{t=1}^T (F[d_{it}(X'_{it}\mathbf{b} + \alpha)]) f(\alpha) d\alpha \right\}$$

where $d_{it} = 2y_{it} - 1$. This expression contains a univariate integral which can be approximated by Gauss-Hermite quadrature. Assuming $\mathbf{a} \sim N(0, \sigma_a^2)$, the contribution of each individual to the sample likelihood function is,

$$(7) \quad L_i = \int_{-\infty}^{+\infty} (1/\sqrt{2\pi}\sigma_a) \exp(-\mathbf{a}^2/2\sigma_a^2) \{g(\mathbf{a})\} d\mathbf{a},$$

where $g(\mathbf{a}) = \prod_{t=1}^T F[d_{it}(X'_{it}\mathbf{b} + \mathbf{a})]$. Use the change of variables, $\mathbf{a} = (\sqrt{2}\sigma_a)z$, to give,

$$(8) \quad L_i = (1/\sqrt{\pi}) \int_{-\infty}^{+\infty} \exp(-z^2) \{g((\sqrt{2}\sigma_a)z)\} dz$$

This expression is suitable for Gauss-hermite quadrature and can be approximated as a weighted sum,

$$(9) \quad L_i \gg (1/\ddot{\theta}p) \sum_{j=1}^m w_j g((\ddot{\theta}2s^2)a_j)$$

where the weights (w_j) and abscissae (a_j) are tabulated in standard mathematical references.

Table 1 shows a summary of results for the pooled and random effects probit models applied to our binary measure of health problems in the BHPS (for brevity we only report the coefficients on the income variables). The pooled probit model treats the data as a single cross section and ignores the fact that there are repeated observations for each individual. The pooled probit estimates provide a useful benchmark for the random effects model. It has been shown that the pooled probit (pseudo-) ML estimator gives consistent estimates of the β s, irrespective of whether the assumed error structure is correct. Of course the pooled probit model does not provide an estimate of s_a^2 and, therefore, information about the structure of the error term and the relative importance of the individual effect.

Table 1: Estimates for the pooled and random effects probits

	Pooled probit	Random effects probit (24 point quadrature)
Ln('permanent income')	-0.573	-0.490
Ln('transitory income')	-0.115	-0.049
σ_α^2	-	0.784
LnL	-6263.5	-4291.2

The income effects are negative, suggesting that those with higher household income are less likely to report limiting health problems. The estimates of the coefficient on the log of permanent income are quite similar for the pooled and random effects probits, although the size of the permanent effect relative to the transitory effect is smaller in the pooled probit. The estimate of the variance of the individual effect in the random effects specification, σ_α^2 , is 0.784. Since the overall error variance has to be set equal to one in

order to identify the probit model, this can also be interpreted as the proportion of the overall error variance that is explained by the time invariant individual effect.

1.4 Extensions and alternatives for the random effects probit

The random effects probit model has two important limitations: it relies on the assumptions that the error components have a normal distribution and that errors are not correlated with the regressors. Normality can be relaxed by using a finite mixture model. The possibility of correlated effects can be dealt with by using conditional (fixed effects) approaches or by parameterising the effect.

The finite mixture model

Deb [4] presents a random effects probit model in which the distribution of the individual effect is approximated by a discrete density. This is an example of a finite mixture model (see [5]). In this case the sample log-likelihood is approximated by,

$$(10) \quad \ln L = \sum_{i=1}^n \ln \left(\sum_{j=1}^C p_j \left\{ \prod_{t=1}^T F[d_{it}(X'_{it}\mathbf{b} + \mathbf{a}_i)] \right\} \right), \quad 0 \leq p_j \leq 1, \quad \sum_{j=1}^C p_j = 1$$

Monte Carlo experiments are used to assess the small sample properties of the estimator. These show that only 3-4 points of support are required for the discrete density to mimic normal and chi-square densities sufficiently well so as to provide approximately unbiased estimates of the structural parameters and the variance of the individual effect.

The conditional logit estimator.

The conditional logit estimator uses the fact that $\sum_t y_{it}$ is a sufficient statistic for \mathbf{a}_i (see e.g., [6]). This means that conditioning on $\sum_t y_{it}$ allows a consistent estimator for β to be derived. Using the logistic function,

$$(11) \quad P(y_{it} = 1) = F(X'_{it}\mathbf{b} + \mathbf{a}_i) = \exp(X'_{it}\mathbf{b} + \mathbf{a}_i) / (1 + \exp(X'_{it}\mathbf{b} + \mathbf{a}_i))$$

it is possible to show that,

$$(12) \quad P[(0,1) | (0,1) \text{ or } (1,0)] = \exp((X_{i2} - X_{i1})' \mathbf{b}) / (1 + \exp((X_{i2} - X_{i1})' \mathbf{b}))$$

This implies that a standard logit model can be applied to differenced data and the individual effect is swept out.

Parameterising the individual effect

Another approach to dealing with individual effects that are correlated with the regressors is to specify $E(\mathbf{a}/X)$ directly. For example, in dealing with a random effects probit model Chamberlain [6,7] suggests using,

$$(13) \quad \mathbf{a}_i = X_i' \mathbf{a} + u_i, \quad u_i \sim iid N(0, \mathbf{S}^2)$$

where $X_i = (X_{i1}, \dots, X_{iT})$, the values of the regressors for every wave of the panel, and $\mathbf{a} = (\mathbf{a}_1, \dots, \mathbf{a}_T)$. Then, by substituting, the distribution of y_{it} conditional on X but marginal to α_i has the probit form,

$$(14) \quad P(y_{it} = 1) = F[(1 + \mathbf{S}^2)^{-1/2} (X_{it}' \mathbf{b} + X_i' \mathbf{a})]$$

The model could be estimated as a random effects probit to retrieve the parameters of interest (β, σ) . Recently Wooldridge [8] has shown that this approach can also be applied in a random effects probit model with state dependence. In this case the initial values of the dependent variable are also included in Equation (13) in order to deal with the problem that the initial conditions are correlated with the individual effect (the so-called 'initial conditions' problem, see Heckman [2]).

1.5 Simulation-based inference

The random effects probit model only involves a univariate integral. More complex models, for example where the error term ε_{it} is assumed to follow an AR(1) process lead to sample log-likelihood functions that involve higher order integrals. Monte Carlo simulation techniques can be used to deal with the computational intractability of nonlinear models, such as the panel probit model and the multinomial probit. Popular methods of simulation-based inference include classical Maximum Simulated Likelihood (MSL) estimation, and Bayesian Markov Chain Monte Carlo (MCMC) estimation. These are discussed in more detail in sections 2 and 3, but the basic principles of simulation-based estimation can be illustrated in the context of the random effects probit model.

The principle behind Maximum Simulated Likelihood (MSL) estimation is to replace population expectations with a sample analogue. As a simple illustration, consider the example of the random effects probit model. An individual's contribution to the sample likelihood function can be written in the form,

$$(15) \quad L_i = \int_{-\infty}^{+\infty} \{g(\mathbf{a})\} f(\mathbf{a}) d\mathbf{a} = E_a[g(\mathbf{a})]$$

Then the individual contribution to the corresponding simulated likelihood function is,

$$(16) \quad L_i = (1/R) \sum_{j=1}^R g(\mathbf{a}_j)$$

The α_j 's are draws from a standard normal and the simulated likelihood is the average of $g(\alpha_j)$ over R draws. The MSL estimator finds the parameter values that maximize the simulated likelihood function. The properties of this estimator are discussed more thoroughly in the next section.

2. Classical simulation methods

2.1 Simulation-based estimation

Recall that the general version of our model is,

$$(17) \quad y_{it} = 1(y_{it}^* > 0) = 1(X'_{it}\mathbf{b} + u_{it} > 0)$$

This implies that the probability of observing the sequence $y_{i1} \dots y_{iT}$ for a particular individual is,

$$(18) \quad Prob(y_{i1}, \dots, y_{iT}) = \int_{a_{i1}}^{b_{i1}} \dots \int_{a_{iT}}^{b_{iT}} f(u_{i1}, \dots, u_{iT}) du_{iT}, \dots, du_{i1}$$

with $a_{it} = -X'_{it}\mathbf{b}$, $b_{it} = \infty$ if $y_{it}=1$ and $a_{it}=-\infty$, $b_{it} = X'_{it}\mathbf{b}$ if $y_{it}=0$. The sample likelihood L is the product of these integrals, L_i , over all n individuals. In certain cases, such as the random effects probit model, L_i can be evaluated by quadrature. In general, the T -dimensional integral L_i cannot be written in terms of univariate integrals that are easy to evaluate. Gaussian quadrature works well with low dimensions but computational problems arise with higher dimensions. Multivariate quadrature uses the Cartesian product of univariate evaluation points and the number of evaluation points increases exponentially. Instead we can use Monte Carlo (MC) simulation to approximate integrals that are numerically intractable. This includes numerous models derived from the multivariate normal distribution (the panel probit, multinomial and multivariate probit, panel ordered probit and interval regression, panel Tobit, etc.). MC approaches use pseudo-random selection of evaluation points and computational cost rises less rapidly than with quadrature.

The principle behind simulation-based estimation is to replace a population value by a sample analogue. This means that we can use laws of large numbers (LLNs) and central

limit theorems (CLTs) to derive the statistical properties of the estimators. The basic problem is to evaluate an integral of the form,

$$(19) \quad \int_a^b \{h(u)\}f(u) du = E_u[h(u)]$$

This can be approximated using draws from $f(u)$, u_r , $r=1, \dots, R$,

$$(20) \quad (1/R) \sum_{r=1}^R h(u_r)$$

This is the direct MC estimate of $E_u[h(u)]$. Direct MC estimators are usually unbiased and consistent in R (due to the LLN) and asymptotically normal (due to the CLT).

The Crude Frequency Simulator

Lerman and Manski [9] proposed a MC algorithm for the evaluation of multivariate normal (MVN) probabilities such as those in the panel probit model. This is rarely used in practical applications but it illustrates a simple way of simulating MVN probabilities directly. The CFS for the probability, P_i , of a sequence of binary outcomes in the panel probit model works as follows:

The CFS algorithm

1. Generate a T vector of pseudo-random independent standard normal variates.
2. Convert this into a $N(0, \hat{\mathbf{a}})$ vector, where $\hat{\mathbf{a}}$ is the covariance matrix of $f(u_{i1}, \dots, u_{iT})$.
3. Determine whether this vector matches the conditions for the observed sequence of outcomes y_{i1}, \dots, y_{iT} .
4. Repeat these steps a large number, R , times.
5. Evaluate the relative frequency of draws that are consistent with the observed outcomes. This gives an approximate value for P_i .

The CFS is computationally simple and cheap. But it has problems. It can easily return zero for P_i . This leads to computational problems when taking logs or ratios. It is discontinuous in the parameters creating a problem for derivative-based optimisation routines. The CFS has higher variance than other unbiased and consistent simulators for MVN probabilities.

An alternative is the GHK (Geweke-Hajivassiliou-Keane) simulator (see Part II of the structured bibliography in the Appendix). The GHK is a smooth recursive conditioning simulator (SRC). The GHK algorithm draws recursively from truncated univariate normals. This relies on the decomposition,

$$(21) \quad f(u_1, \dots, u_T) = f(u_1)f(u_2|u_1)\dots\dots f(u_{T-1}|u_{T-2}, \dots, u_1)f(u_T|u_{T-1}, \dots, u_1)$$

along with the fact that the conditional (in our case truncated) normal density can be written as a univariate normal. The GHK simulator produces probability estimates that are bounded between 0 and 1. The estimates are continuous and differentiable with respect to (β, Σ) , because each contribution is continuous and differentiable. It has a smaller variance than the CFS, because each element is bounded between 0 and 1. The GHK appears to be the most accurate simulator available for a given computation time. Box 1 provides Hajivassiliou's [10] Gauss code for the GHK algorithm.

**BOX 1: GAUSS code for obtaining probability estimates using the GHK algorithm
(source: Hajivassiliou [10])**

```

proc ghk(m,mu,w,wi,c,a,b,r,u);
local j,ii,ta,tb,tt,wgt,v,p;
j= 1;
ii= 1;
ta= cdfn((a[1,1]-mu[1,1])/(c[1,1]+ 1.e-100))*ones(1,r);
tb= cdfn((b[1,1]-mu[1,1])/(c[1,1]+ 1.e-100))*ones(1,r);
tt= cdfni(u[1,.].*ta+ (1-u[1,.]).*tb);
wgt= tb-ta;
do while j< m;
j=j+ 1;
ta= cdfn(((a[j,1]-mu[j,1])*ones(1,r)-c[j,ii]*tt)/(c[j,j]+ 1.e-100));
tb= cdfn(((b[j,1]-mu[j,1])*ones(1,r)-c[j,ii]*tt)/(c[j,j]+ 1.e-100));
tt= tt | cdfni(u[j,.].*ta+ (1-u[j,.]).*tb);
ii= ii | j;
wgt= wgt.*(tb-ta);
endo;
p= sumc(wgt')/r;
retp(p);
endp;

```

2.2 Maximum Simulated Likelihood (MSL)

This is a simple extension of classical maximum likelihood estimation (MLE) and is useful in many cases when the log-likelihood function involves high dimensional integrals. This includes the panel probit with RE+ AR(1). The idea is to replace the likelihood function L_i with a sample average over random draws,

$$(22) \quad l_i = (1/R) \sum_{r=1}^R l(u_{ir})$$

where $l(u_{ir})$ is an unbiased simulator of L_i . The MSL estimates are the parameter values that maximize,

$$(23) \quad \ln l = \sum_{i=1}^n \ln l_i$$

In practice, antithetics can be used to reduce the variance of the simulator. These are based on,

$$(24) \quad l_i = (1/2R) \sum_{r=1}^{2R} l(u_{ir})$$

where $u_{ir} = u_{i1}, \dots, u_{i2R}$ and $u_j = -u_{jR}$ for $j=R+1, \dots, 2R$. Antithetics reduce the variance by using symmetric draws. If the probability simulator is linear in the draws, this approach reduces the variance to zero.

Having an unbiased simulator l_i of L_i (from CFS or GHK) does not imply an unbiased simulator of $\ln L_i$ or the overall sample log-likelihood function (as $E[\ln l_i] \neq \ln(E[l_i])$). Of course MLE is, in general, biased due to nonlinearity. But, unlike MLE, the MSL estimator is not consistent solely in n . This is because the simulator is biased downwards for all

individuals and the bias depends on β . Consistency and asymptotic unbiasedness can be obtained by reducing the error in the simulated sample log-likelihood to zero as $R \rightarrow \infty$ at a sufficient rate with n . Hajivassiliou and Ruud [11] show that a sufficient rate is $R/\sqrt{n} \rightarrow \infty$ as $n \rightarrow \infty$. Hajivassiliou and Ruud also show that this is sufficient for the usual MLE estimate of the covariance matrix to be used without any correction.

2.3 Application of MSL

Table 2 compares the estimates of the income effects and the variance of the individual effect for the random effects probit model computed using both quadrature (with 24 points) and MSL (with 150 replications). Simulation is not required for this model but it provides a useful test of the reliability of the simulation approach. It is clear from the table that the estimates are very similar.

Table 2: Comparison of quadrature and MSL for the random effects probit

	MLE (24 point)	MSL (R= 150)
Ln('permanent income')	-0.490	-0.510
Ln('transitory income')	-0.049	-0.052
σ_α^2	0.784	0.788
LnL	-4291.2	-4290.5

Simulation becomes necessary to move beyond the simple random effects (RE) specification. Table 3 presents estimates of the income effects for models ranging from independent probit equations, through the RE and RE+AR(1), to an unrestricted covariance matrix. The results show that the income effects are largely unaffected by moving to more flexible specifications of the covariance matrix.

Table 3: Estimated income effects under alternative covariance structures

	INP	RE	RE+ AR(1)	Unrestricted
Ln('permanent income')	-0.573 (0.061)	-0.510 (0.049)	-0.509 (0.049)	-0.511 (0.049)
Ln('transitory income')	-0.115 (0.030)	-0.052 (0.023)	-0.057 (0.025)	-0.055 (0.024)

To assess the overall statistical performance of the models Table 4 shows various model selection criteria. The most general model, with an unrestricted covariance matrix, has the largest log-likelihood function. But the model has many more parameters than the RE+ AR(1) and RE specifications. Information criteria can be used to penalise the measure of goodness of fit for the loss of degrees of freedom. The unrestricted model is still preferred according to the Akaike information criterion (AIC) but, when the number of parameters is penalised more heavily with the Bayesian information criterion (BIC) and Consistent AIC (CAIC), the RE+ AR(1) specification is favoured.

Table 4: Model selection criteria using MSL

	lnL	AIC	BIC	CAIC
Unrestricted	-4188.07	8421.12	8731.92	8820.00
RE+ AR(1)	-4214.83	8455.66	8635.23	8626.82
RE	-4290.55	8606.10	8778.76	8851.31

Table 5 assesses the impact of simulation bias on the results by comparing estimates based on different numbers of replications. These show that the estimates are robust for values of R in the range 75-150.

Table 5: Estimates for different R

	‘permanent’	‘transitory’	s_a^2	ρ
R= 75	-0.515	-0.057	0.698	0.521
R= 100	-0.515	-0.059	0.695	0.528
R= 150	-0.510	-0.057	0.691	0.541

Simulation bias can be assessed more formally by using the test statistic proposed by Hajivassiliou [10]. Table 6 shows statistics for both the full sample and for separate sub-samples selected according to individuals’ highest academic qualifications. Statistics are presented with p-values in parentheses. As expected, using only one replication is insufficient to obtain an asymptotically unbiased estimator. However there is a positive relationship between the test statistics and sample size. This suggests that the statistics reflect the relative magnitude of the bias: for smaller values of n, R can be reduced while maintaining the variance due to simulation. As the number of replications is increased the average value of the test statistics gradually reduces.

Table 6: Test statistics for simulation bias

	R= 1	R= 40	R= 75	R= 150
FULL SAMPLE	701.1 (0.00)	30.4 (0.25)	17.7 (0.89)	37.5 (0.07)
DEGREE	72.1 (0.00)	16.1 (0.85)	34.9 (0.05)	32.8 (0.08)
A LEVEL	171.0 (0.00)	24.8 (0.36)	34.0 (0.06)	17.9 (0.77)
O LEVEL	273.7 (0.00)	22.7 (0.48)	30.0 (0.15)	32.8 (0.08)
NO QUALIFICATIONS	257.1 (0.00)	58.1 (0.00)	38.3 (0.02)	22.5 (0.49)

Some guidelines

Theoretical considerations and experience of applying these methods suggests that the following guidelines should be taken into account when putting MSL into practice:-

1. Use fixed draws: The random draws should only be drawn once and not varied as the optimisation algorithm searches. If the random variates do change, the values of (β, Σ) which maximize the simulated log-likelihood will change and the optimisation routine may never converge. This is a general requirement for any simulation-based estimation that uses an iterative optimisation routine. Furthermore, the asymptotic theory for these estimators is based on a given set of draws.
2. Use a smooth and bounded simulator: It is important to use a simulator which is smooth in β (such as the GHK algorithm) so that derivative based optimisation routines may be used. It is also important to use simulators which are bounded by 0 and 1 (e.g. GHK), so the simulated sample log-likelihood can always be evaluated. When using numerical derivatives it is advisable to use alternative step sizes until the estimates of the derivatives are stable.
3. Scale the data and use good initial values: The tails of the multivariate normal density die out very rapidly. This causes potential underflows during computation (the values are smaller than can be expressed and manipulated during computation). This problem is severe when taking logs. Hajivassiliou [10] suggests standardising regressors (to zero mean and unit variance) before estimation. More generally, reasonable starting values will help to prevent this problem.
4. Use antithetics: Use of antithetic variates can reduce simulation variance and bias substantially.
5. Validate the algorithm: Use a simple model to validate a new program by comparing MSL and MLE.
6. Check for bias: Bias is model specific and, while there is guidance in the literature – for example many studies suggest that values of R less than 50 are sufficient - sensitivity analysis is important. As noted above, a test for asymptotic bias is available (Hajivassiliou, [10]). Bias corrections are available but are computationally difficult and may not perform well (e.g., Lee, [12]).

2.4 Other Classical estimators

MSL is not the only classical simulation estimator available. The Method of Simulated Scores (MSS) is a general approach based on the simulated score function. The estimator is implicitly defined as the value of β which satisfies,

$$(25) \quad 1/n \sum_{i=1}^n l_b / l_i = 0$$

where $l_b = [1/R_1 \sum_{r=1}^{R_1} l_{br}]$ is an unbiased simulator for the vector of derivatives L_b and l_i is an unbiased simulator for L_i based on R_2 simulations. In general, the MSS estimator is consistent and asymptotically normal as $n \rightarrow \infty$ and $R_2/\sqrt{n} \rightarrow \infty$. The value of R_1 affects the efficiency of the estimator.

The method of moments (MoM) estimator for the probit model solves the orthogonality condition,

$$(26) \quad 1/n \sum_{i=1}^n Q_i [y_i - F(X'_i \mathbf{b})] = 0$$

for a fixed and exogenous vector of instruments Q (e.g. X). MoM is equivalent to MLE with a suitable choice of Q . If the moments cannot be evaluated analytically (as in the panel probit model) then the residual can be simulated and, for efficiency, so can the optimal instruments. This gives the Maximum Simulated Moments (MSM) estimator (McFadden [13]). Further details of the links between MSL, MSM and MSS can be found in the references given in the Appendix.

Box 2 suggests some convenient sources of software for classical simulation estimation.

BOX 2: Software for classical simulation methods

<http://econ.lse.ac.uk/~vassilis/pub/simulation>

NMRSIM_G.LIB

Library of GAUSS routines for multivariate normal rectangle probabilities and their derivatives. Companion to Hajivassiliou, McFadden and Ruud [14]

SSMLMNP by A. Borsch-Supan & V. Hajivasilliou

Estimation code for simulated maximum likelihood of multinomial probit in Fortran 77

Limdep v.7.0 for Windows (<http://www.limdep.com>)

Simulator for multivariate normal CDF using GHK algorithm. By default, R=100 but can be changed. Up to m=20 variate integral

3. Bayesian MCMC methods

3.1 The Bayesian approach

In Bayesian analysis a prior distribution $p(\mathbf{q})$ is updated with the information contained in the sample (for the RE+AR(1) panel probit model $\theta = \{\beta, \rho, \sigma^2_\alpha\}$). Given a specified likelihood, $p(y/\mathbf{q})$, the posterior density is given by Bayes' theorem,

$$(27) \quad p(\mathbf{q}/y) = p(\mathbf{q})p(y/\mathbf{q}) / p(y)$$

where,

$$(28) \quad p(y) = \int p(\mathbf{q})p(y/\mathbf{q}) d\mathbf{q}$$

$p(y)$ is known as the predictive likelihood and it is used for model comparison. It determines the probability that the specified model is correct. The posterior density $p(\mathbf{q}/y)$ reflects updated beliefs about the parameters. Given the posterior distribution, a 95% credible interval can be constructed that contains the parameter with probability equal to 95%. Point estimates of the parameters of interest, q_i , are provided by the posterior mean,

$$(29) \quad E(q_i/y) = \int q_i p(q_i/y) dq_i$$

3.2 Markov Chain Monte Carlo (MCMC) Methods

MCMC methods are used when it is not possible to obtain the characteristics of the posterior distribution analytically. The methods provide a sample from the posterior distribution. Posterior moments and credible intervals are obtained from this sample. The

method works because, under appropriate conditions, a Markov chain, in which draws are conditional on the previous iteration, should converge to a stationary distribution that is independent of the initial values. After discarding the initial iterations, the remaining values can be regarded as a sample from the posterior distribution.

Gibbs Sampling

A draw from a distribution $p(\mathbf{q}_1, \mathbf{q}_2)$ can be obtained in two steps. First, draw \mathbf{q}_1 from its marginal distribution $p(\mathbf{q}_1)$. Second, draw \mathbf{q}_2 from its conditional distribution given \mathbf{q}_1 , $p(\mathbf{q}_2 / \mathbf{q}_1)$. In many situations it is possible to sample from the conditional distribution $p(\mathbf{q}_2 / \mathbf{q}_1)$ but it is not obvious how to sample from the marginal $p(\mathbf{q}_1)$. The Gibbs sampling algorithm obtains a sample by sampling iteratively from the full conditional distributions. Even though the Gibbs sampling algorithm never draws from the marginal, after a sufficiently large number of iterations the draws can be regarded as a sample from the joint distribution.

To implement Gibbs sampling the vector of parameters \mathbf{q} is subdivided into s groups, $\mathbf{q} = (\mathbf{q}_1, \dots, \mathbf{q}_s)$. The initial values $\mathbf{q}^1 = (\mathbf{q}_1^1, \dots, \mathbf{q}_s^1)$ are fixed arbitrarily. For example, let $\mathbf{q} = (\mathbf{q}_1, \mathbf{q}_2)$, then $\mathbf{q}^k = (\mathbf{q}_1^k, \mathbf{q}_2^k)$ is drawn as follows:

1. \mathbf{q}_1^k is drawn from the distribution $p(\mathbf{q}_1 / y, \mathbf{q}_2^{k-1})$
2. \mathbf{q}_2^k is drawn from the distribution $p(\mathbf{q}_2 / y, \mathbf{q}_1^k)$

The process is repeated until a sufficiently large sample is obtained.

In practice the conditional distributions must be easy to sample from. Also, the number of sub-groups in which \mathbf{q} is subdivided should be kept as small as possible to speed up convergence. Good starting values can substantially save in computing time in complex models.

Data Augmentation

Following Tanner and Wong [15], latent or missing data can be regarded as parameters belonging to \mathbf{q} . Once the latent data is imputed, the model becomes linear and the conditional distributions are easier to sample from. For example, in a simple probit model,

$$(30) \quad y^*_i = X'_i \mathbf{b} + \mathbf{e}_i, \mathbf{e} \sim N(0,1), \mathbf{q} = (\mathbf{b}, y^*)$$

The conditional distribution of y^*_i given \mathbf{b} is $N(X'_i \mathbf{b}, 1)$, truncated to positive values if $y=1$ and truncated to negative values if $y=0$. If the prior for \mathbf{b} is non-informative, the conditional of \mathbf{b} given y^* is a $N(b, V)$ where b and V are the OLS estimators when y^* is observed and the variance of \mathbf{e} is known.

Metropolis-Hastings (M-H) Algorithms

When the conditional distributions cannot be sampled directly, Gibbs sampling can be combined with a Metropolis step. Assume $p(\mathbf{q}_1 / y, \mathbf{q}_2)$ is not easy to sample from. Let $q(\mathbf{q}_1 / \mathbf{q}_1^{k-1}, \mathbf{q}_2^{k-1})$ be a density for \mathbf{q}_1 given \mathbf{q}_2^{k-1} that can be sampled easily. A M-H algorithm generates $\mathbf{q}^k = (\mathbf{q}_1^k, \mathbf{q}_2^k)$ as follows:

1. Draw a candidate value \mathbf{q}_1^* for \mathbf{q}_1^k from the distribution $q(\mathbf{q}_1^* / \mathbf{q}_1^{k-1}, \mathbf{q}_2^{k-1})$.
2. Set $\mathbf{q}_1^k = \mathbf{q}_1^*$ with some probability γ otherwise keep $\mathbf{q}_1^k = \mathbf{q}_1^{k-1}$.
3. Generate \mathbf{q}_2^k from the distribution $p(\mathbf{q}_2 / y, \mathbf{q}_1^k)$.

The probability γ depends on the values of the ratio $\pi(.) / q(.)$ evaluated at the new proposed value and at the previous value in the chain. New candidates can be generated in a simple way from a normal distribution centred at the previous value in the chain and with arbitrary variance. However, if the dimension of \mathbf{q}_1 is greater than 2, this may not work well. Alternatively, new candidates can be generated from a distribution that approximates the conditional density $p(\mathbf{q}_1 / y, \mathbf{q}_2^{k-1})$. Bad approximations may result in low probabilities of acceptance and slow convergence.

3.3 MCMC estimation of the panel probit

To see how MCMC works in practice consider the panel probit with autocorrelated error terms,

$$(31) \quad y_{it} = 1(y_{it}^* > 0) = 1(X'_{it}\mathbf{b} + \mathbf{a}_i + \mathbf{e}_{it} > 0), \quad \mathbf{e}_{it} = \mathbf{r}\mathbf{e}_{it-1} + \mathbf{h}_{it}, \quad \text{var}(\mathbf{a}_i) = \sigma_{\alpha}^2$$

The priors for β , σ_{α}^2 and ρ are normal, inverted gamma and uniform (-1,1) distributions respectively. The parameters and latent data are divided into 5 groups: y_{it}^* , \mathbf{b} , \mathbf{a}_i , σ_{α}^2 , \mathbf{r} . The parameters in each group are generated conditionally on the parameters in the rest of the groups in the following way:

1. Fix the initial values.
2. Generate y_{it}^* from a truncated normal distribution with mean $X'_{it}\mathbf{b} + \mathbf{a}_i$ and variance $\text{var}(\mathbf{e}_{it})$, according to the value of y_{it} .
3. Generate \mathbf{b} from a normal distribution. The mean and variance of this distribution are the Bayesian point estimates in a model in which y_{it}^* and \mathbf{a}_i are observed and $(\sigma_{\alpha}^2, \mathbf{r})$ are known.
4. Generate \mathbf{a}_i from a normal distribution. The mean and variance are the point estimates in a model in which y_{it}^* is observed and $(\mathbf{b}, \sigma_{\alpha}^2, \mathbf{r})$ are known.
5. Generate σ_{α}^2 from an inverted gamma distribution. The parameters of this inverted gamma are the same as in a model in which y_{it}^* and \mathbf{a}_i are observed and (\mathbf{b}, \mathbf{r}) are known.
6. \mathbf{r} can be generated with a M-H step. New candidates can be generated from a normal distribution.

Box 3 shows how this algorithm can be implemented in Gauss.

BOX 3: Gauss code for MCMC estimation of panel probit model

```
iter= 1;
Beta[1,.]= zeros(1,K);
Rho[1]= 0;
Var_ind[1]= 1;
Do while iter< = NIT;
Latent= augment(Beta[iter,.], Effects, Var_ind[iter], Rho[iter]);
Beta[iter+ 1,.]= sampleB(Latent, Effects, Var_ind[iter], Rho[iter]);
Effects= sampleI(Latent, Beta[iter+ 1,.], Var_ind[iter], Rho[iter]);
Var_ind= sampleV(Latent, Effects, Beta[iter+ 1,.], Rho[iter]);
Rho[iter+ 1]= sampleR(Latent, Effects, Beta[iter+ 1], Var_ind[iter+ 1],
Rho[iter]);
iter= iter+ 1;
endo;
```

3.4 Output and convergence analysis

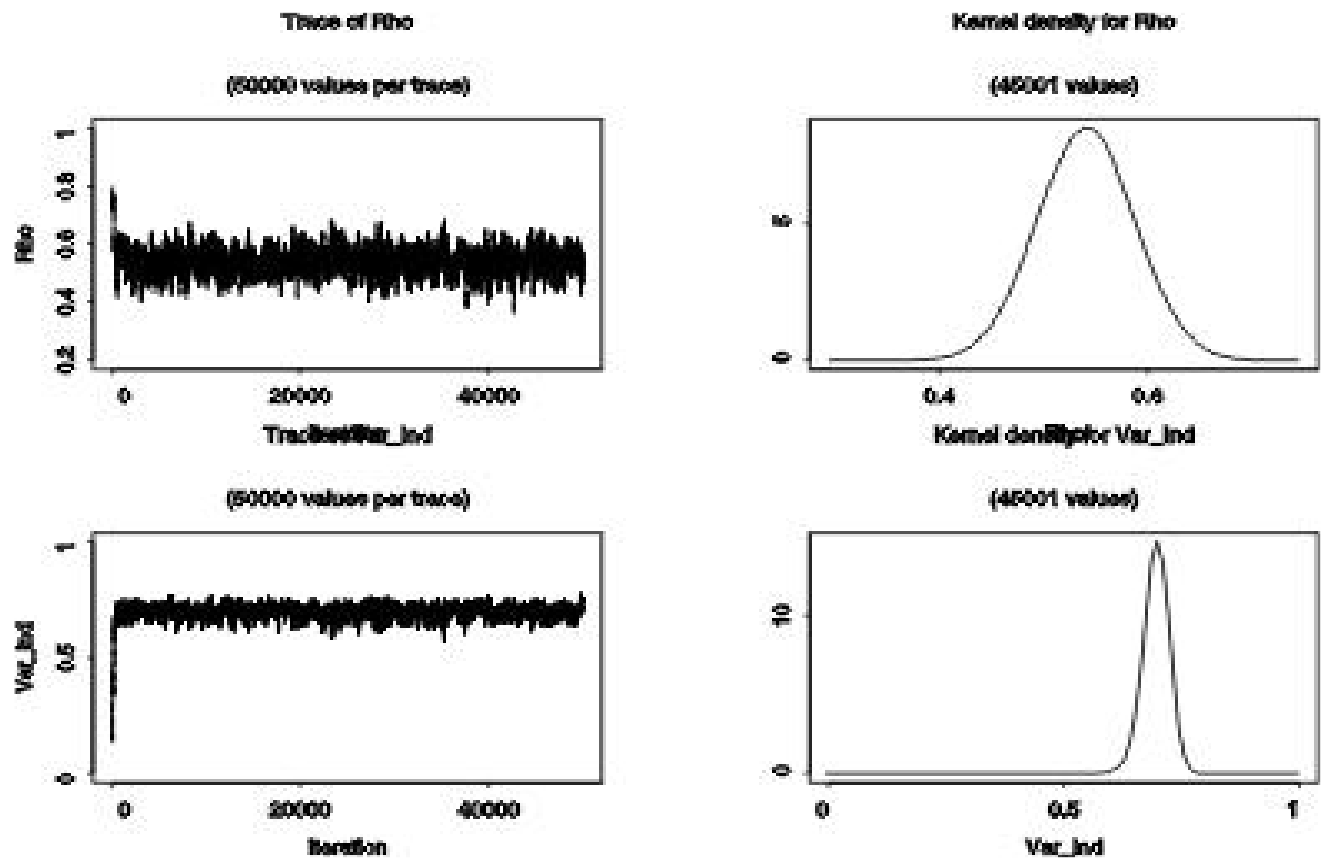
Sample means, standard deviations and correlations estimated from the MCMC draws are all strongly consistent estimates of the corresponding characteristics of the posterior distribution. A sample for any transformation of the parameters $g(\mathbf{q})$ can be obtained by equivalently transforming the sampled values for \mathbf{q} . This allows the calculation of credible intervals and posterior means for $g(\mathbf{q})$. In the panel probit model, $g(\mathbf{q})$ could represent the marginal effect of income on the probability of being ill. Since values in the chain are not independent, sample standard deviations are usually biased. Geweke [16] proposes an alternative method that is implemented in standard packages.

Table 7 shows that the MCMC algorithm produces very similar point estimates and standard errors to those estimated by MSL, for the full range of specifications. Figure 1 shows the results of the MCMC estimation for the autocorrelation parameter \mathbf{r} and the variance of the individual effect σ_α^2 in the RE+AR(1) model. The left-hand panels show the output of the Markov chain for successive iterations, while the right-hand panels show kernel density estimates of the posterior densities for the two parameters.

Table 7: Comparison of MSL and MCMC estimates for the panel probit models

	RE		RE + AR(1)		Un-	restricted
	MSL	MCMC	MSL	MCMC	MSL	MCMC
Ln('permanent income')	-0.510 (0.049)	-0.499 (0.050)	-0.510 (0.049)	-0.504 (0.049)	-0.511 (0.049)	-0.481 (0.051)
Ln('transitory income')	-0.052 (0.023)	-0.051 (0.023)	-0.057 (0.025)	-0.057 (0.024)	-0.054 (0.024)	-0.051 (0.022)
σ_α^2	0.788 (0.012)	0.791 (0.012)	0.691 (0.025)	0.696 (0.025)		
ρ			0.541 (0.040)	0.541 (0.042)		

Figure 1: MCMC parameter estimates and posterior densities.



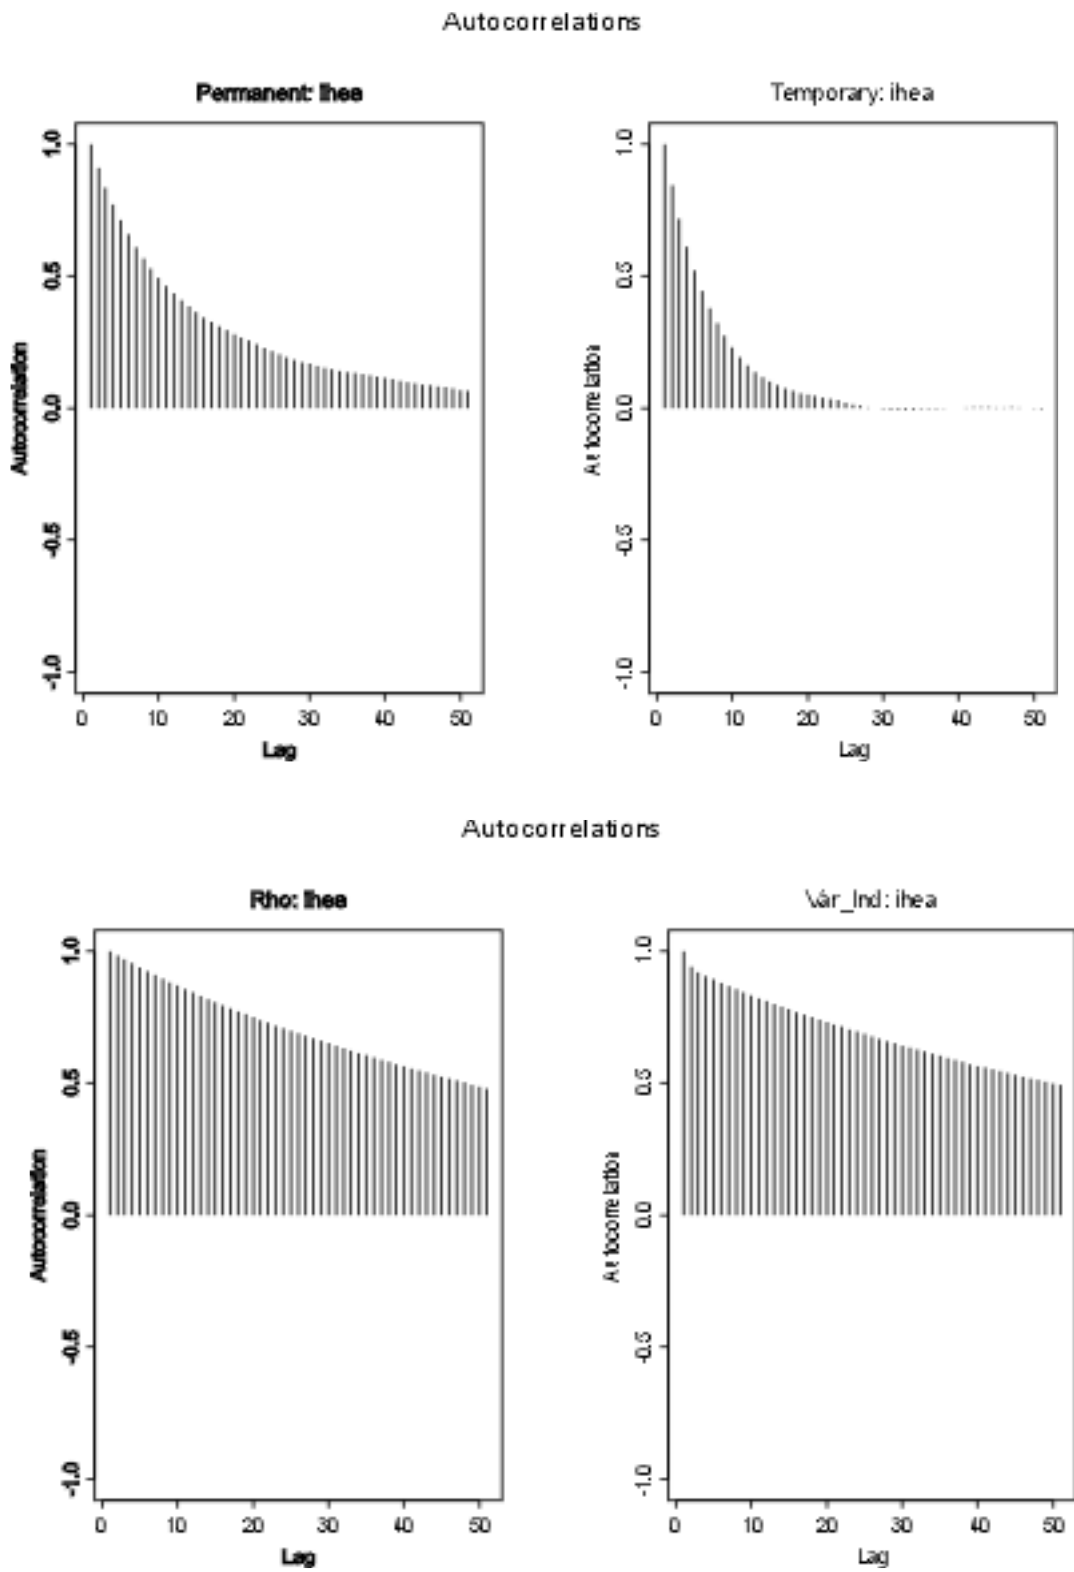
Independently of the initial values, after a sufficient number of iterations, the algorithm should converge to the posterior distribution. An informal way of checking for convergence is to subdivide the chain in several parts and compare the average and standard deviations for each part. Alternatively, averages and standard deviations of chains started at different initial values may be compared. High serial correlation in the chain values indicates that a longer chain will be necessary to obtain precise estimates. Cowles and Carlin [17] review some formal tests of convergence. The freely available software CODA includes the following tests: the Geweke test for convergence; the Heidelberger and Welch test for convergence which also tests whether the posterior mean has been estimated with a given degree of accuracy; and the Raftery and Lewis test gives the length of the chain necessary to calculate a credible interval with a specified degree of accuracy.

Results of applying the CODA software to our MCMC algorithm for the RE+ AR(1) model are shown in Table 8. These show that a longer chain is required to get convergence in the estimates for σ_{α}^2 and ρ than for the income effects. This finding is reflected in the plots of the autocorrelation functions, shown in Figure 2, these ‘die-out’ for the two income parameters but persist for σ_{α}^2 and ρ .

Table 8: Results of convergence analysis

	Geweke	Convergence (H-W)	Posterior mean (H-W)	Length of chain
Ln(‘permanent income’)	0.108	passed	passed	11,000
Ln(‘transitory income’)	0.255	passed	passed	8,000
σ_{α}^2	-0.782	passed	passed	35,000
ρ	0.420	passed	passed	21,000

Figure 2: Autocorrelations for key parameters



3.5 Testing and model selection

Testing in the Bayesian approach consists of comparing the probabilities of different hypotheses. The probability of each model is determined by the predictive likelihood, $\mathbf{p}(y)$, which is the normalising constant in the denominator of Bayes' theorem,

$$(32) \quad \mathbf{p}(y|M_i) = \int \mathbf{p}(\mathbf{q})\mathbf{p}(y|\mathbf{q}, M_i) d\mathbf{q}$$

Given m possible models $\{M_i\}$, and prior probabilities for each model, $\mathbf{p}(M_i)$, the posterior probability for model M_i is,

$$(33) \quad \mathbf{p}(M_i|y) = \mathbf{p}(M_i)\mathbf{p}(y|M_i) / \sum_{j=1}^m \mathbf{p}(M_j)\mathbf{p}(y|M_j)$$

Although the posterior probability depends on the number of models m , which is determined a priori, the ratio of the probabilities of two different models does not depend on m . In the case of equal prior probabilities for each model this ratio is known as the Bayes factor,

$$(34) \quad B_{ij} = \mathbf{p}(y|M_i)/\mathbf{p}(y|M_j)$$

If a model is to be selected, it should be the model with the largest value for the predictive likelihood.

To illustrate the Bayes factors for our models let:

M₁: Random effects and independent time variant errors

M₂: Random effects and AR(1) errors

M₃: Unrestricted variance-covariance matrix

The log of the predictive likelihood ($\ln \mathbf{p}(y|M_j)$) for each of the models is:

M₁: -4293.18 (1.02)

M₂: -3986.85 (0.88)

M₃: -4367.28 (1.19)

Giving Bayes factors:

$\Pr(M_2)/\Pr(M_1) = 1.09E+133$

$\Pr(M_1)/\Pr(M_3) = 1.52E+32$

$\Pr(M_2)/\Pr(M_3) = 1.66E+165$

These results support the RE+AR(1) specification against the simpler RE model and against the more general unrestricted model. This parallels the finding based on the BIC and CAIC criteria in the classical analysis.

The posterior probabilities for each model lead to a procedure to deal with uncertainty about the appropriate model to use. The posterior density for \mathbf{q} takes into account the different possible specifications,

$$(35) \quad p(\mathbf{q} | y) = \sum_{j=1}^m p(\mathbf{q} | y, M_j) p(M_j | y)$$

The posterior mean for \mathbf{q} is a weighted average of the posterior means in each model,

$$(36) \quad E(\mathbf{q} | y) = \sum_{j=1}^m E(\mathbf{q} | y, M_j) p(M_j | y)$$

Chib and Jeliazkov [18] present a method to calculate the predictive likelihood, $p(y)$. From Bayes' theorem,

$$(37) \quad p(\mathbf{q} | y) = p(\mathbf{q}) p(y | \mathbf{q}) / p(y)$$

It follows that any particular value, \mathbf{q}^* , of the parameters satisfies the identity,

$$(38) \quad \ln p(y) = \ln p(q^*) + \ln p(y/q^*) - \ln p(q^*/y)$$

Chib and Jeliazkov [18] propose a method to estimate the posterior ordinate $\ln p(q^*/y)$. The method requires running the algorithm for additional iterations. In order to assess the accuracy of the calculation, they also provide the standard deviation of the estimated value for $\ln p(q^*/y)$. If evaluation of the likelihood, $p(q^*/y)$, involves multiple integrals it can be computed using the methods described in Section 2.

Testing hypotheses about q

When the hypothesis of interest is of the type $q_1 = k$, it is possible to use Verdinelli and Wasserman's [19] method. Unlike the Chib and Jeliazkov method, in many situations this method does not require any additional computations. Their procedure gets more complicated in terms of computing time when the normalising constant of $p(q_1/y, q_2)$ is not known, or when q_1 and q_2 are not independent a priori.

Box 4 lists sources of software for doing Bayesian analysis using MCMC. It also provides a reference for the Gauss code used to estimate the panel probit models.

Overview

This review illustrates the scope for using simulation methods to allow for flexible specifications of heterogeneity in nonlinear models for panel data. It uses binary choice models to show what can be done with conventional methods and how the range of models can be expanded by using classical and Bayesian simulation methods. Practical applications of the methods are illustrated using data on self-reported health from the British Household Panel Survey (the BHPS). Our aim is to provide a brief introduction to simulation methods and to show their relevance for applied analysis in health economics. To provide some guidance for readers who would like to pursue the topic in more detail the Appendix includes a structured bibliography.

BOX 4: Software for Bayesian MCMC analysis

BACC freely available at www.econ.umn.edu/~bacc

Includes, among other models, univariate latent models with flexible distributions for the error term (mixture of normals and student-t distributions). Allows the user to implement more complex procedures building on simpler models. This requires knowledge of C language.

BUGS and CODA freely available at www.mrc-bsu.cam.ac.uk/bugs/

BUGS allows users to easily specify their own MCMC sampling algorithm

CODA provides checks of convergence for the MCMC output

MLWIN is commercial software which includes Bayesian estimation of multi-level models

The Gauss code in Box 3 can be downloaded from <http://www.york.ac.uk/res/herc/yshe>

REFERENCES

1. Contoyannis P, Jones AM, Rice N. The dynamics of health in British households: simulation-based inference in panel probit models. <http://www.york.ac.uk/res/herc/yshe> [26 November 2001].
2. Heckman J. Statistical models for discrete panel data. In *Structural Analysis of discrete data with econometric applications*, Manski CF, McFadden D (eds). MIT Press: Cambridge Mass., 1981; 114-178.
3. Greene W. Fixed and random effects in nonlinear models. <http://www.stern.nyu.edu/~wgreene> [26 November 2001]
4. Deb P. A discrete random effects probit model with application to the demand for preventive care. *Health Economics* 2001; **10**: 371-383.
5. Heckman JJ, Singer B. A method of minimizing the distributional impact in econometric models for duration data. *Econometrica* 1984; **52**: 271-230.
6. Chamberlain G. Analysis of covariance with qualitative data. *Review of Economic Studies* 1980; **47**: 225-238.
7. Chamberlain G. Panel data. In *Handbook of Econometrics Volume 1*, Griliches Z, Intrilligator M. (eds). Elsevier: Amsterdam, 1984; 1247-1318.
8. Wooldridge JM. The initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity. Mimeo, Michigan State University, 2000.
9. Lerman S, Manski C. On the use of simulated frequencies to approximate choice probabilities. In *Structural Analysis of discrete data with econometric applications*, Manski CF, McFadden D. (eds). MIT Press: Cambridge Mass., 1981; 305-319.
10. Hajivassiliou V. Some practical issues in maximum simulated likelihood. In *Simulation-Based inference in econometrics: Methods and Applications*. Mariano R, Schuermann T, Weeks MJ. (eds). Cambridge University Press: Cambridge, 2000; 71-99.
11. Hajivassiliou V, Ruud P. Classical Estimation methods using simulation. In *Handbook of Econometrics Volume 4*. Engle R, McFadden D. (eds). Elsevier: Amsterdam, 1994; 2383-2441.
12. Lee L-F. Simulated maximum likelihood estimation of dynamic discrete choice statistical models. *Journal of Econometrics* 1997; **82**:1-35.

13. McFadden D. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica* 1989; **57**:995-1026.
14. Hajivassiliou V, McFadden D, Ruud P. Simulation of multivariate normal rectangle probabilities and their derivatives: Theoretical and computational results. *Journal of Econometrics* 1996, **72**:85-134.
15. Tanner M, Wong W. The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association* 1987; **82**: 528-550.
16. Geweke J. Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In *Bayesian Statistics 4*. Bernardo JM, Berger J, Dawid AP, Smith AFM. (eds). Oxford University Press: Oxford, 1992; 169-193.
17. Cowles M, Carlin B. Markov Chain Monte Carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association* 1996; **91**:883-904.
18. Chib S, Jeliazkov I. *Marginal Likelihood From the Metropolis-Hastings Output*. *Journal of the American Statistical Association* 2001; **96**:270-281.
19. Verdinelli I, Wasserman L. Computing Bayes Factors Using a Generalization of the Savage-Dickey Density Ratio. *Journal of the American Statistical Association* 1995; **90**: 614-618.

APPENDIX: STRUCTURED BIBLIOGRAPHY

PART I: Methods for nonlinear panel data models

Arellano M, Honoré B. Panel data models: some recent developments. In *Handbook of Econometrics Volume 5*, Heckman JJ, Leamer E. (eds). Elsevier: Amsterdam, 2001; 3229-3296.

Baltagi B. *Econometric Analysis of Panel Data*. John Wiley: Chichester, 2001.

Butler J, Moffitt R. A computationally efficient quadrature procedure for the one-factor multinomial probit model. *Econometrica* 1982; **50**:761-764.

Chamberlain G. Panel Data. In *Handbook of Econometrics Volume 1*. Griliches Z, Intriligator MD. (eds). Elsevier: Amsterdam, 1984; 1247-1318

Guilkey D, Murphy J. Estimation and testing in the random effects probit model. *Journal of Econometrics* 1993; **59**:301-317.

Heckman J. Statistical models for discrete panel data. In *Structural Analysis of discrete data with econometric applications*, edited by C. F. Manski CF, McFadden D. (eds). MIT Press: Cambridge Mass., 1981: 114-178.

PART II: Classical simulation methods

i) Numerical integration and simulation as alternatives for solving integrals.

Geweke J. Monte Carlo simulation and numerical integration. In *Handbook of Computational Economics*, Amman HM, Kendrick DA, Rust J. (eds). Elsevier: Amsterdam, 1996: 731-800.

Pudney S. *Modelling individual choice: The econometrics of corners, kinks and holes*. Appendix 3- Computation of probability integrals. Blackwell: Oxford, 1989.

Rubinstein R. *Simulation and the Monte Carlo method*. John Wiley: Chichester, 1981.

ii) The GHK and other probability simulators

Borsch-Supan A, Hajivassiliou V. Smooth unbiased multivariate probability simulators for maximum likelihood estimation of limited dependent variable models. *Journal of Econometrics* 1993; **58**:347-368

Geweke J. Efficient simulation from the multivariate normal and student- t distributions subject to linear constraints and the evaluation of constraint probabilities. *Prepared for 'Computing Science and Statistics: 23rd symposium'* Seattle, April 22-24 1991.

Hajivassiliou V. Simulating normal rectangle probabilities and their derivatives: Effects of vectorization. *The International Journal of Supercomputer Applications* 1993; **7**:231-53.

Hajivassiliou V, McFadden D, Ruud P. Simulation of multivariate normal rectangle probabilities and their derivatives: Theoretical and computational results. *Journal of Econometrics* 1996; **72**:85-134.

Stern S. A method for smoothing simulated moments of discrete probabilities in multinomial probit models. *Econometrica* 1992; **60**: 943-952.

Vijverberg V. Monte Carlo evaluation of multivariate normal probabilities. *Journal of Econometrics* 1997; **76**:281-307.

Vijverberg V. Rectangular and wedge-shaped multivariate normal probabilities. *Economics Letters* 2000; **68**:13-20.

(iii) Maximum simulated likelihood

Hajivassiliou V. Some practical issues in maximum simulated likelihood. In *Simulation-based inference in econometrics: methods and applications*. Mariano R, Schuermann T, Weeks MJ. (eds). Cambridge University Press: Cambridge, 2000; 71-99.

Hyslop D. State dependence, serial correlation and heterogeneity in intertemporal labor force participation of married women. *Econometrica* 1999; **67**:1255-1294.

Lee L-F. Simulated maximum likelihood estimation of dynamic discrete choice statistical models. *Journal of Econometrics* 1997; **82**:1-35.

Lerman S, Manski C. On the use of simulated frequencies to approximate choice probabilities. In *Structural Analysis of discrete data with econometric applications*. C. F. Manski CF, McFadden D. (eds). MIT Press: Cambridge Mass., 1981; 305-319.

Mealli F, Pudney S. Occupational pensions and job mobility in Britain: Estimation of a random-effects competing risks model. *Journal of Applied Econometrics* 1996; **11**:293-320.

(iv) MSS and MSM

Hajivassiliou V, McFadden D. The method of simulated scores for the estimation of LDV models. *Econometrica* 1998; **66**:863-896.

Keane M. A computationally practical simulation estimator for panel data. *Econometrica* 1994; **62**: 95-116.

McFadden D. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica* 1989; **57**:995-1026

(v) Surveys

Contoyannis P. An introduction to simulation based estimation and inference. *University of York*: <http://www.york.ac.uk/res/herc/docs/simulation.pdf> [26 November 2001].

Geweke J, Keane M. Computationally intensive methods for integration in econometrics. In *Handbook of Econometrics Volume 5*, Heckman JJ, Leamer E. (eds). Elsevier: Amsterdam, 2001; 3463-3568.

Gourieroux C, Monfort A. Simulation-based inference: A survey with special reference to panel data models. *Journal of Econometrics* 1993; **59**:5-33.

Gourieroux C, Monfort A. *Simulation-based econometric methods*. Oxford University Press: Oxford, 1996.

Hajivassiliou V. Simulation estimation methods for limited dependent variable models In *Handbook of Statistics Volume 11*. Maddala GS, Rao CR, Vinod HD. (eds). Elsevier: Amsterdam, 1993; 519-543.

Hajivassiliou V, Ruud P. Classical estimation methods using simulation. In *Handbook of Econometrics Volume 4*. Engle R, McFadden D. (eds). Elsevier: Amsterdam, 1994; 2383-2441.

Keane M. Simulation estimation for panel data models with limited dependent variables. In *Handbook of Statistics Volume 11*. Maddala GS, Rao CR, Vinod HD. (eds). Elsevier: Amsterdam, 1993; 545-571.

McFadden D, Ruud P. Estimation by simulation. *The Review of Economics and Statistics* 1994; **76**:591-608.

Stern S. Simulation-based estimation. *Journal of Economic Literature* 1997, **35**:2006-2039.

PART III: Bayesian simulation methods

(i) Introduction to the Bayesian approach.

Lee PM. *Bayesian statistics: an introduction*. Arnold: London, 1997.

(ii) Introduction to MCMC methods.

Casella G, George E. Explaining the Gibbs sampler. *The American Statistician* 1992; **46**:167-174.

Chib S, Greenberg E. Understanding the Metropolis-Hastings algorithm. *The American Statistician* 1995; **49**:327-335.

Gamerman D. *Markov Chain Monte Carlo*. Chapman & Hall: London, 1997.

Gilks WR, Richardson S, Spiegelhalter, D.J. (eds) *Markov Chain Monte Carlo in Practice*. Chapman & Hall: London, 1996.

Smith A, Roberts G. Bayesian computation via the Gibbs sampler and related Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society B* 1993; **55**:3-23.

Tanner TA, Wong WH. The Calculation of Posterior Distributions by Data Augmentation. *Journal of the American Statistical Association* 1987; **82**: 528-550.

(iii) Bayesian MCMC methods in econometrics.

Bauwens L, Lubrano M, Richard J. *Bayesian inference in dynamic econometric models*. Oxford University Press: Oxford, 1999.

Chib S. Markov chain Monte Carlo methods: computation and inference. In *Handbook of Econometrics Volume 5*, Heckman JJ, Leamer E. (eds). Elsevier: Amsterdam, 2001; 3569-3649.

Geweke J. Using simulation methods for Bayesian econometric models: inference, development, and communication. *Econometric Reviews* 1999; **18**:1-73.

(iv) Some Bayesian MCMC algorithms for limited dependent variable models.

Albert J, Chib S. Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association* 1993; **88**: 669-679.

Chib S. Bayesian analysis of correlated binary data. In *Generalized linear models: a Bayesian perspective*. Dey D, Ghosh S, Mallick B. (eds). Marcel-Dekker: New York, 2000; 113-131.

Chib S, Greenberg E. Analysis of multivariate probit models. *Biometrika* 1998; **85**:347-361.

Cowles K, Carlin BP, Connett JE. Bayesian Tobit modelling of longitudinal ordinal clinical trial compliance data with nonignorable missingness. *Journal of the American Statistical Association* 1996; **91**: 86-98.

Geweke J, Keane, M, and Runkle, D. (1997) Statistical inference in the multinomial multiperiod probit model. *Journal of Econometrics*, 80:125-165

McCulloch RE, Polson NG, Rossi PE. A Bayesian analysis of the multinomial probit model with fully identified parameters. *Journal of Econometrics* 2000; **99**: 173-193.

(v) Analysis of convergence.

Cowles M, Carlin B. Markov Chain Monte Carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association* 1996; **91**:883-904.

Geweke J. Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In *Bayesian Statistics 4*. Bernardo JM, Berger J, Dawid AP, Smith AFM. (eds). Oxford University Press: Oxford, 1992; 169-193.

(vi) Model uncertainty.

Chib S, Jeliazkov I. *Marginal Likelihood From the Metropolis-Hastings Output*. *Journal of the American Statistical Association* 2001; **96**:270-281.

Fernandez C, Steel M. Benchmark priors for Bayesian model averaging. *Journal of Econometrics* 2001; **100**: 381-427.

Verdinelli I, Wasserman L. Computing Bayes factors using a generalization of the Savage-Dickey density ratio. *Journal of the American Statistical Association* 1995; **90**: 614-618.

PART IV: Selected applications in health economics

Bolduc D, Lacroix G, Muller C. The choice of medical providers in rural Benin: a comparison of discrete choice models. *Journal of Health Economics* 1996; **15**:477-498

Borsch-Supan A, Hajivassiliou VA, Kotlikoff LJ, Morris N. Health, children, and elderly living arrangements: a multiperiod multinomial probit model with unobserved heterogeneity and autocorrelated errors. In *Topics in the Economics of Aging*. Wise DA. (ed). University of Chicago Press: Chicago, 1992.

Chib S, Hamilton B. Bayesian analysis of cross-section and clustered data treatment models. *Journal of Econometrics* 2000; **97**:25-50.

Contoyannis P, Jones AM, Rice N. The dynamics of health in British households: simulation-based inference in panel probit models. <http://www.york.ac.uk/res/herc/yshe> [26 November 2001].

Hamilton B. The impact of HMOs on Medicare costs: Bayesian MCMC estimation of a robust panel data Tobit model with survival, *Health Economics* 1999; **8**:403-414.

Hamilton B. Estimating treatment effects in randomized clinical trials with non-compliance: The impact of maternal smoking on birth-weight. *Health Economics* 2001; **10**: 399-410.

Munkin M, Trivedi, P. Simulated maximum likelihood estimation of multivariate mixed-Poisson models, with application. *Econometrics Journal* 1999; **2**: 29-40.

Pudney S, Shields M. Gender and Racial Discrimination in pay and promotion for NHS nurses. *Oxford Bulletin of Economics and Statistics* 2000; **62** Special Issue: 801-835.

Pudney S, Shields M. Gender, race, pay and promotion in the British nursing profession: estimation of a generalized ordered probit model. *Journal of Applied Econometrics* 2000; **15**:367-399.