

Random effects dynamic panel models for unequally spaced multivariate categorical repeated measures: an application to child–parent exchanges of support

Fiona Steele¹ | Emily Grundy²

¹Department of Statistics, London School of Economics & Political Science, London, UK

²Institute for Economic and Social Research, University of Essex, Colchester, UK

Correspondence

Fiona Steele, Department of Statistics, London School of Economics & Political Science, UK.

Email: f.a.steele@lse.ac.uk

Funding information

UK Economic and Social Research Council, Grant/Award Number: ES/P000118/1

Abstract

Exchanges of practical or financial help between people living in different households are a major component of inter-generational exchanges within families and an increasingly important source of support for individuals in need. Using longitudinal data, bivariate dynamic panel models can be applied to study the effects of changes in individual circumstances on help given to and received from non-coresident parents and the reciprocity of exchanges. However, the use of a rotating module for collection of data on exchanges leads to data where the response measurements are unequally spaced and taken less frequently than for the time-varying covariates. Existing approaches to this problem focus on fixed effects linear models for univariate continuous responses. We propose a random effects estimator for a family of dynamic panel models that can handle continuous, binary or ordinal multivariate responses. The performance of the estimator is assessed in a simulation study. A bivariate probit dynamic panel model is then applied to estimate the effects of partnership and employment transitions in the previous year and the presence and age of children in the current year on an individual's propensity to give or receive help. Annual data on respondents' partnership, employment status and dependent children, and data on exchanges of help collected at 2- and 5-year intervals are used in this study.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2020 The Authors. Journal of the Royal Statistical Society: Series C (Applied Statistics) published by John Wiley & Sons Ltd on behalf of Royal Statistical Society

KEYWORDS

autoregressive models, intergenerational exchanges, lagged response models, longitudinal data, unequal spacing

1 | INTRODUCTION

Dynamic models, also known as autoregressive or lagged response models, have been widely used for the analysis of longitudinal data in a broad range of application areas such as unemployment (Arulampalam et al., 2000), political preferences (Stegmueller, 2013), and health and well-being (Pudney, 2008; Steele et al., 2013). A common motivation for using a dynamic model is to investigate the relative contributions of persistence and unobserved heterogeneity as explanations for serial correlation in a response y , where persistence (or inertia) is a causal effect of an individual's past values of y on their current value and unobserved heterogeneity refers to individual differences in y due to unmeasured time-invariant characteristics. Dynamic models are also used to estimate the effects of time-varying covariates on change in y , and generalisations to multivariate responses allow the study of interdependencies among repeated measures on multiple correlated responses.

Standard discrete-time dynamic panel models assume that measurements of the response and time-varying covariates are taken at the same equally spaced occasions. As the correlation between an individual's responses at two points in time t_1 and t_2 will typically decay with $|t_2 - t_1|$, the coefficients of the lagged response and time-varying covariates capturing changes between t_1 and t_2 cannot be assumed invariant to spacing. The naïve treatment of unequally spaced observations as if they are equally spaced will lead to biased coefficients of the lagged response and serially correlated predictors (Millimet & McDonough, 2017; Sasaki & Xin, 2017). The usual approach to this problem is simply to exclude observations to achieve an equally spaced panel, which may lead to a large reduction in the number of measurement occasions available for analysis. Unfortunately, unequal spacing is a common feature of longitudinal studies, which may arise by design or because of wave non-response (see Millimet and McDonough (2017) for a number of examples from developed countries). Moreover, a common design feature of household panel studies is the use of rotating modules to reduce survey costs and respondent burden, which leads to some variables being measured less frequently, and often at unequal intervals, than variables collected in the core questionnaire at each (usually annual) wave.

This paper is concerned with dynamic models for the analysis of responses that are measured less frequently than time-varying covariates, and is motivated by a study of bidirectional exchanges of support between adult children and their parent(s). The data are from the British Household Panel Survey (BHPS) and its successor the UK Household Longitudinal Study (UKHLS). Information about the help that respondents give to and receive from their non-coresident parents was collected in a rotating module which was administered at five waves between 2001 and 2015, with gaps of 2 and 5 years. Our focus is the effects of changes in adult children's circumstances on their capacity to give help to parents and on their need for help from parents, and the reciprocity of child-parent exchanges. We consider the effects of partnership and employment transitions in the previous year and the presence and age of children in the current year on an individual's propensity to give or receive any help using annual data on respondents' partnership and employment status and coresident children.

Although unequally spaced panel data are commonplace, there has been relatively little research on the topic and previous work has considered only linear models for univariate continuous responses with fixed effects to account for unobserved heterogeneity. Apart from the different collection schedules for the responses and covariates, the BHPS/UKHLS data on intergenerational exchanges have

several other features that require extensions to existing approaches: exchanges of help are measured by binary responses, help given and received must be treated as a bivariate response and modelled jointly because of the interest in reciprocity, and there is missing data on some of the time-varying covariates. We therefore propose a random effects estimator for a general class of linear and non-linear dynamic panel models for potentially unequally spaced multivariate and categorical responses. The estimator can be implemented using Markov chain Monte Carlo (MCMC) methods in Bayesian modelling software.

While the proposed approach is motivated by an application to household panel data where covariate information is collected annually, it can be applied to birth cohort studies where waves of data collection are at unequal intervals but time-varying covariates for finer, regularly spaced, time intervals can be derived from retrospectively collected data on partnership, childbearing and employment histories. The same methods are applicable, under a missing at random assumption, to unequal spacing due to wave or item non-response or changes in a respondent's membership of the target population for an analysis; in our application, for example we accommodate gaps of 4 years due to wave non-response. Other applications include medical settings where patients do not attend clinics at the scheduled (equally spaced) times (Rosner & Muñoz, 1988), and pseudo-panel data constructed from repeated cross-sectional surveys conducted at unequal intervals (McKenzie, 2001) where covariate data are available from external sources.

The remainder of the paper is organised as follows. In Section 2, we give a brief overview of previous research on dynamic models and describe how our approach builds on this literature. The proposed random effects estimator is then described in Section 3, starting with a univariate continuous response before generalisations to binary or ordinal responses and multivariate responses. In Section 4, the finite-sample performance of the proposed estimator is assessed in a small-scale simulation study which considers the impact of imbalance in the number of repeated measurements per person and the degree of autocorrelation in the response and in a time-varying covariate. The application to intergenerational exchanges is described in Section 5, and we conclude with remarks on directions for future research in Section 6.

2 | OVERVIEW OF PREVIOUS RESEARCH ON UNEQUAL SPACING

Let $\mathbf{y}_{it} = (y_{1it}, y_{2it}, \dots, y_{Rit})^T$ be a R -variate response at time t ($t = 1, 2, \dots, T$) for individual i ($i = 1, \dots, n$), where the elements of \mathbf{y}_{it} may be observed or latent continuous variables. Suppose that the underlying data generating process (DGP), and model of interest, for \mathbf{y}_{it} takes the form of a first-order dynamic panel model

$$\mathbf{y}_{it} = \mathbf{\Gamma} \mathbf{y}_{t-1,i} + \mathbf{B} \mathbf{x}_{it} + \mathbf{u}_i + \mathbf{e}_{it}, \quad t = 2, 3, \dots, T \quad (1)$$

where $\mathbf{\Gamma}$ is an $R \times R$ coefficient matrix, \mathbf{x}_{it} is a p -vector of covariates with $R \times p$ coefficient matrix \mathbf{B} , \mathbf{u}_i is an R -vector of individual effects capturing unmeasured time-invariant influences (unobserved heterogeneity) and $\mathbf{e}_{it} \sim \mathbf{N}(\mathbf{0}, \mathbf{\Omega}_e)$ is an R -vector of time-varying residuals. The diagonal entries of $\mathbf{\Gamma}$ are the autoregressive effects for each response y_{rit} , representing persistence or state dependence, and the off-diagonal entries are the cross-lagged effects of all other responses on each response. (For simplicity, we assume \mathbf{y}_{it} and \mathbf{x}_{it} are mean-centred and omit the intercepts from \mathbf{B} .) There has been much debate about whether to treat \mathbf{u}_i as fixed or random in panel models and more generally hierarchical models for clustered data (e.g. Clarke et al., 2015). In the case of continuous univariate y_{it} , fixed effects models are

often favoured because of the additional distributional assumptions required in random effects models. In the random effects approach, it is common to assume that \mathbf{u}_i are normally distributed. Although alternatives such as discrete distributions (latent class models) and mixtures of normals are available, evidence from simulation studies for random effects logistic models estimated via maximum likelihood indicates that statistical inference is robust to departures from the normality assumption, especially within-cluster (or time-varying in a longitudinal setting) covariate effects (McCulloch & Neuhaus, 2011). More serious concerns with random effects models are the treatment of the initial conditions \mathbf{y}_{1i} and the assumption that unmeasured time-invariant influences on \mathbf{y}_{it} (represented by \mathbf{u}_i) are uncorrelated with the covariates. A common way to handle the initial conditions problem is to specify a model for \mathbf{y}_{1i} (see Section 3.4), while a straightforward way to relax the assumption that the time-varying covariates are independent of the random effects is to use a correlated random effects model in which the individual means of the time-varying covariates are included as additional predictors (e.g. Mundlak, 1978). In the linear case, the estimates of \mathbf{B} obtained from fitting a random effects model that includes the individual means of \mathbf{x}_{it} are equal to the fixed effects estimates of \mathbf{B} . The Mundlak approach generalises to unbalanced panels (Wooldridge, 2019) and translates to probit models for binary responses (Wooldridge, 2010, pp. 615–617). Advantages of random effects models include the possibility to include time-invariant predictors and their greater flexibility, for example in handling multivariate categorical responses.

Following Millimet and McDonough (2017), we refer to the discrete unit of time for the DGP, that is the length of the interval $(t - 1, t)$, as the unit period. In a discrete-time model for equally spaced panel data, the unknown unit period is usually taken as the time between measurements, referred to as the observation interval. In general, an irregular spacing arises whenever the observation interval differs from the unit period. We now review special cases of Equation (1), which have been considered in previous work.

In the classic dynamic panel model, $R = 1$, y_{it} is observed and the observation interval coincides with the unit period. Several authors have considered departures from this setting where, for univariate observed y_{it} , the observation schedule leads to a form of data coarsening with both y_{it} and \mathbf{x}_{it} observed together at $M < T$ potentially unequally spaced time points. Such observation patterns may arise by design, where waves of a study are planned at unequal intervals, or because of non-response, in which case the observation intervals may vary across individuals. Rosner and Muñoz (1988) consider a simplified form of the classic model without u_i for $M < T$, while Millimet and McDonough (2017) and Sasaki and Xin (2017) specify u_i as fixed effects. Pacini and Windmeijer (2015) consider a setting closest to ours where the data on y_{it} may be missing at random while \mathbf{x}_{it} is observed at every wave. (Further details of the methods of estimation used in previous research are given in Section 3.1.) While we restrict our attention to models with lagged responses as predictors, there has been related work on static panel models without a lagged response but with e_{it} assumed to follow an AR(1) process (Baltagi & Wu, 1999; Jones & Boadi-Boateng, 1991), and in the time-series literature (see Millimet & McDonough, 2017 for a review).

In a separate strand of research, generalisations of the classic dynamic model have been developed to handle categorical and multivariate responses, but only for equally spaced y_{it} . Examples include the semiparametric fixed effects approach of Honoré and Kyriazidou (2000) and random effects probit models (e.g. Wooldridge, 2010), which can be formulated as models for the dependency of an underlying latent continuous response y_{it}^* on the observed lagged response $y_{t-1,i}$. An alternative approach, proposed for a univariate ordinal response, is the random effects latent autoregression model in which the lagged predictor $y_{t-1,i}$ is replaced by $y_{t-1,i}^*$ (Pudney, 2008). Bivariate dynamic models, with the lag of one response included as a predictor of the current value of the other and residual correlation between responses, are commonly used in the social sciences and particularly in psychology where they

are referred to as cross-lagged panel models (e.g. Selig & Little, 2012). Generalisations to bivariate binary data are also available (e.g. Alessie et al., 2004).

This paper brings together these two streams of research to develop a flexible random effects estimator for dynamic panel models for any combination of the following types of data and measurement schedule: (i) unequally spaced measurements of \mathbf{y}_{it} , which may be measured more frequently than the time-varying covariates \mathbf{x}_{it} , (ii) continuous, binary or ordinal responses, (iii) multivariate responses, and (iv) missing data on \mathbf{x}_{it} . Our approach additionally incorporates models for the initial conditions \mathbf{y}_{1i} .

3 | A GENERAL RANDOM EFFECTS MODEL FOR UNEQUALLY SPACED RESPONSES

We begin with the special case of a linear model for a univariate continuous response to fix ideas. We briefly describe how the research reviewed in Section 2 accommodates unequal spacing for a model with individual fixed effects before setting out our proposed random effects approach. This is followed by generalisations for binary or ordinal and multivariate responses within this random effects framework.

3.1 | Linear model for univariate continuous responses

Consider a special case of (1) for a continuous univariate response y_{it} :

$$y_{it} = \gamma y_{t-1,i} + \boldsymbol{\beta} \mathbf{x}_{it} + u_i + e_{it}, \quad t = 2, 3, \dots, T \quad (2)$$

where γ is the autoregression parameter, with $|\gamma| < 1$ for a stationary process, $\boldsymbol{\beta}$ is a row vector of coefficients, u_i is an individual effect and $e_{it} \sim \text{i.i.d. } \mathbf{N}(0, \sigma_e^2)$. It is straightforward to include time-invariant covariates, as described later.

Suppose that the observation schedule for y is such that two consecutive measures are spaced s unit periods apart. The dependency of y_{it} on $y_{t-s,i}$ can be obtained by repeated substitution in Equation (2) to give

$$y_{it} = \gamma^s y_{t-s,i} + \boldsymbol{\beta} \sum_{k=0}^{s-1} \gamma^k \mathbf{x}_{t-k,i} + \left(\frac{1-\gamma^s}{1-\gamma} \right) u_i + \epsilon_{it}, \quad t = s+1, \dots, T \quad (3)$$

using $\sum_{k=0}^{s-1} \gamma^k = \left(\frac{1-\gamma^s}{1-\gamma} \right)$, where $\epsilon_{it} = \sum_{k=0}^{s-1} \gamma^k e_{t-k,i}$ and $\text{var}(\epsilon_{it}) = \left(\frac{1-\gamma^{2s}}{1-\gamma^2} \right) \sigma_e^2$.

Equation (3) has several important features: (i) the coefficient of $y_{t-s,i}$ depends on the gap s between observations, which, for a stationary process, implies that the autocorrelation in y decays with s , (ii) the coefficients of the time-varying covariates are non-linear functions of the parameters of interest γ and $\boldsymbol{\beta}$, (iii) the covariates include not only \mathbf{x}_{it} for the waves at which y_{it} is observed, but also the intermediate values $\mathbf{x}_{t-1,i}, \dots, \mathbf{x}_{t-s+1,i}$, (iv) the form of ϵ_{it} implies that the residual variance depends on s , and (v) the contribution of u_i depends on s . Point (iii) is especially important because excluding intermediate \mathbf{x} values, so that they are absorbed in ϵ_{it} , leads to omitted variable bias if \mathbf{x}_{it} are serially correlated. Most previous research considers settings where \mathbf{x}_{it} is available only when y_{it} is observed, leading to missing data on $\mathbf{x}_{t-1,i}, \dots, \mathbf{x}_{t-s+1,i}$ which is typically handled using some form of imputation.

When $u_i = 0$ in Equations (2) and (3) can be estimated directly via weighted non-linear least squares, with weights $(1 - \gamma^2)/(1 - \gamma^{2s})$ to account for the heteroskedacity in ϵ_{it} (Rosner & Muñoz, 1988). For models with u_i specified as fixed effects, the standard approach to estimation of Equation (2) is to take first differences $y_{it} - y_{t-1,i}$, in order to eliminate u_i , and to use a generalised method of moments (GMM) estimator (Ahn & Schmidt, 1995; Arellano & Bond, 1991). A complication when observations are unequally spaced is that formulating (3) as a model for $y_{it} - y_{t-s,i}$ does not eliminate u_i ; instead u_i has a factor structure with a loading that depends on s . Millimet and McDonough (2017) therefore proposed two alternative approaches to GMM that are suitable for unequally spaced intervals: quasi-differenced GMM, a type of differencing that removes u_i , and a form of non-linear least squares with an instrumental variable for $y_{t-s,i}$ and the individual fixed effects replaced by a linear combination of the individual means of \mathbf{x}_{it} . Both estimators are consistent if it is assumed that the covariates are exogenous and serially uncorrelated. The assumption of serial independence in \mathbf{x}_{it} may be relaxed by imputation of \mathbf{x} for the missing periods or linear interpolation as in Rosner and Muñoz (1988). Sasaki and Xin (2017) further investigate fixed effects dynamic models for unequally spaced panels, and provide details of the general spacing patterns required for model identification, which include those found in many UK and US longitudinal surveys such as the British birth cohort studies and the BHPS/UKHLS data on family exchanges analysed here. They then propose a GMM estimator for these patterns which, in contrast to quasi-differenced GMM, does not assume that time-varying covariates are serially uncorrelated or require imputation of covariates for the missing periods, provided that the covariance between pairs of y and \mathbf{x} values over time is assumed constant for a given width of observation interval.

Turning to settings where y_{it} is observed at $M < T$ waves while \mathbf{x}_{it} is measured at all T waves, Pacini and Windmeijer (2015) note that traditional GMM discards any individual with y_{it} observed for fewer than three consecutive periods, and propose a two-step GMM estimator that requires only that individuals provide at least three, possibly non-consecutive, observations. McKenzie (2001) considers a dynamic model for pseudo-panel data where cohorts, or some other fixed grouping of individuals, are followed over time. Assuming an individual-level DGP of form similar to (2), it is shown that the corresponding pseudo-panel model can be estimated using non-linear least squares provided that time-varying covariates, defined at the cohort level, can be obtained from external sources for the time periods between surveys.

In this paper, we adopt a random effects approach because of its flexibility for handling categorical and multivariate responses, as required for our application, and ease of implementation in Bayesian modelling software using MCMC. We first propose a random effects estimator for situations where univariate continuous y_{it} is subject to data coarsening. We also consider estimation of the effects of transitions in a binary variable x_{it} in the preceding unit period. Let $m = 1, \dots, M$ index the occasions at which y is observed, and denote by t_m the timing of measurement m and $\Delta t_m = t_m - t_{m-1}$ the gap between consecutive observations, measured in the unit periods of the DGP of Equation (2). Using this notation, Equation (3) can be re-expressed as a model for the observed responses $y_{t_m i}$

$$y_{t_m i} = \gamma^{\Delta t_m} y_{t_{m-1}, i} + \beta \sum_{k=0}^{\Delta t_m - 1} \gamma^k \mathbf{x}_{t_m - k, i} + \left(\frac{1 - \gamma^{\Delta t_m}}{1 - \gamma} \right) u_i + \epsilon_{t_m i}, \quad m = 2, \dots, M \quad (4)$$

where $\epsilon_{t_m i} = \sum_{k=0}^{\Delta t_m - 1} \gamma^k \epsilon_{t_m - k, i}$ with $\text{var}(\epsilon_{t_m i}) = \left(\frac{1 - \gamma^{2\Delta t_m}}{1 - \gamma^2} \right) \sigma_e^2$, and we assume $u_i \sim \text{i.i.d. } N(0, \sigma_u^2)$.

Thus for $\Delta t_m = 2$, for example, $\mathbf{x}_{t_m, i}$ and its first-order lag $\mathbf{x}_{t_m - 1, i}$ are included as covariates with coefficients β and $\beta \gamma$, respectively. From (4), it is clear that Δt_m must be an integer, which has implications for the unit period assumed in the DGP. For example, in our application, observations are spaced 2 and 5 years apart, but setting the unit period equal to the minimum observation interval would give

$\Delta t_m = 1$ and 2.5. We therefore assume the DGP operates for annual y_{it} , which also corresponds to the frequency of measurement of x_{it} . The model of Equation (4) assumes a balanced design with fixed measurement occasions, but M and t_m (and therefore Δt_m) may vary across individuals.

In our application, the covariates of primary interest are indicators of transitions (or events) occurring in the year before t . In the simplest case, a transition is a move between two possible states. For a binary state indicator x_{it} , there are four possible transitions ($x_{t-1,i}$, x_{it}), including remaining in the same state: (0, 0), (0, 1), (1, 0) and (1, 1). Taking (0, 0) as the reference, and omitting i subscripts, the effects of transitions can be modelled in Equation (2) as

$$\beta_1 x_{t-1} + \beta_2 (1 - x_{t-1}) x_t + \beta_3 x_{t-1} (1 - x_t) \quad (5)$$

where β_1 contrasts individuals remaining in state 1 (1, 1) with those remaining in state 0 (0, 0), β_2 contrasts (0, 1) with (0, 0) and is therefore the effect of a transition from state 0 to state 1, and β_3 is the effect of a transition from 1 to 0. By expressing the transition indicators in terms of x_{t-1} and x_t , model (4) for the observed y can be generalised to allow for transition effects by including the three functions of x in Equation (5) and their k th-order lags ($k = 1, \dots, \Delta t_m - 1$) with coefficients constrained to $\beta_j \gamma^k$ ($j = 1, 2, 3$). For a time-invariant covariate x , βx in Equation (2) becomes $\beta x (1 - \gamma^{\Delta t_m}) / (1 - \gamma)$ in (4). In the case of a deterministically varying covariate such as age, we use $x_{it} = x_{it-k} + k$ for $k = 1, 2, \dots$ to derive its contribution to Equation (4) (see Rosner and Muñoz (1988) for details).

3.2 | Latent autoregression model for binary and ordinal responses

The dynamic model for equally spaced data of Equation (2) can be generalised to categorical responses. For an ordinal response with C categories, for example a logit or probit model can be specified for $\Pr(y_{it} \leq c)$, $c = 1, \dots, C - 1$. An ordered logit or probit model may also be expressed in terms of an underlying continuous response y_{it}^* , which is related to the observed ordinal y_{it} by

$$y_{it} = c \iff y_{it}^* \in [\tau_{c-1}, \tau_c), \quad c = 1, \dots, C$$

where τ_c are threshold parameters with $\tau_0 = -\infty$ and $\tau_C = \infty$. Binary responses ($C = 2$) are a special case with a single threshold, or intercept, τ_1 .

The standard dynamic model for categorical data includes dummy variables for categories of the lagged observed response as explanatory variables, leading to a state dependence model (Heckman, 1981). Alternatively, the lagged latent response may be included to obtain a latent autoregression model (Pudney, 2008) so that Equation (2) becomes:

$$y_{it}^* = \gamma y_{t-1,i}^* + \beta x_{it} + u_i + e_{it}^*, \quad t = 2, 3, \dots, T \quad (6)$$

where e_{it}^* is assumed to follow a standard normal distribution for a probit model and a standard logistic distribution for a logit model.

The latent autoregression model has been used in studies of subjective well-being (Pudney, 2008), and political preferences (Stegmueller, 2013). In these applications, y_{it}^* is viewed as the 'true' level of well-being or preference which is measured on a discrete scale, and it is argued that the dynamics in y_{it}^* are of greater substantive interest than the effect of past observed survey responses $y_{t-1,i}$ on y_{it}^* . This is typically the case for ordinal variables but, as noted by Pudney, the state dependence model may be more natural in situations where the response is inherently discrete, for example a binary indicator of

employment status. However, for unequally spaced panel data, the advantage of assuming a DGP of the form (6) is that it allows the application of repeated substitution to obtain a model for the dependency of y_{it}^* on $y_{t-s,i}^*$, that has the same form as (3) with $y_{t-s,i}$ replaced by $y_{t-s,i}^*$. It follows that the model for an ordinal (or binary) response at measurement occasion m can be expressed as Equation (4) with $(y_{t_m i}, y_{t_{m-1}, i}, e_{t_m i})$ replaced by $(y_{t_m i}^*, y_{t_{m-1}, i}^*, e_{t_m i}^*)$ where $e_{t_m i}^* = \sum_{k=0}^{\Delta t_m - 1} \gamma^k e_{t_m - k, i}^*$.

3.3 | Multivariate responses

The above models for univariate continuous and ordinal data can be generalised to multivariate responses. We consider the special case of a bivariate continuous response where each response may influence the other over time. Denote by y_{rti} the response on variable r ($r = 1, 2$) at time t for individual i . For a pair of continuous responses, the bivariate generalisation of the DGP Equation (2) is

$$y_{rti} = \gamma_r y_{r,t-1,i} + \delta_r y_{r',t-1,i} + \beta_r \mathbf{x}_{ti} + u_{ri} + e_{rti}, \quad r = 1, 2; \quad r' \neq r; \quad t = 2, 3, \dots, T \quad (7)$$

where $\mathbf{u}_i = (u_{1i}, u_{2i})^T \sim N_2(\mathbf{0}, \Omega_u)$ and $\mathbf{e}_i = (e_{1i}, e_{2i})^T \sim N_2(\mathbf{0}, \Omega_e)$, and γ_r and β_r are the response-specific autoregression parameter and coefficient vector of \mathbf{x} . The model allows for association between the two processes in three ways: (i) a cross-lag effect δ_r , (ii) correlation between time-invariant unmeasured influences on the two responses, $\text{cor}(u_{1i}, u_{2i})$, and (iii) residual correlation between the two responses at the same time, $\text{cor}(e_{1i}, e_{2i})$. For an ordinal (or binary) response, an ordered probit model can be used where (7) is specified in terms of underlying latent variables y_{rti}^* and their lags $y_{r,t-1,i}^*$.

For consecutive responses spaced s unit periods apart, the dependency of y_{rti} on the lag $y_{r,t-s,i}$ and cross-lag $y_{r',t-s,i}$ can be derived by repeated substitution as in Equation (3). However, the inclusion of the cross-lag requires a double substitution which leads to more complex expressions for the coefficients, random effects and residual covariance structure than in the univariate case. Details are given in the supplementary material.

3.4 | Modelling the initial condition

The dynamic model of Equation (4) can be extended to allow for the possibility that y at the first measurement occasion is endogenous. This source of selection bias is widely referred to as the initial conditions problem, and arises when the start of the observation period does not coincide with the start of the process of interest. In our application, for example, the majority of survey respondents would have left the parental home, thereby joining the target population for studying exchanges between non-co-resident relatives before their entry to the panel.

In a random effects model, the inclusion of lagged responses in Equations (2) and (4) may invalidate the standard assumption that the individual effects u_i are uncorrelated with the explanatory variables when the initial condition $y_{t_1 i}$ is assumed exogenous. One way to allow for an endogenous initial condition is to specify a model for the response at $m = 1$ which shares random effects, and must therefore be estimated jointly, with model (4) for $m > 1$ (Kazemi & Crouchley, 2006):

$$y_{t_1 i} = \alpha \mathbf{x}_{t_1 i} + \lambda_1 u_i + \epsilon_{t_1 i} \quad (8)$$

where $\mathbf{x}_{t_1 i}$ is a vector of covariates, which may be time-invariant or specific to the first measurement occasion, α is a vector of parameters, λ_1 is a loading for the individual random effect and $\epsilon_{t_1 i} \sim N(0, \sigma_{\epsilon_1}^2)$.

The inclusion of u_i in both (8) and (4) allows for the presence of unobserved time-invariant factors affecting responses at all occasions, while λ_1 allows their contribution to differ for the first and subsequent observations of y .

3.5 | Estimation

For a univariate continuous response, the contribution to the likelihood for individual i is

$$\mathcal{L}_i = \int_{-\infty}^{\infty} f_1(y_{t_1,i} | \mathbf{x}_{t_1,i}, u_i) \prod_{m=2}^M f(y_{t_m,i} | y_{t_{m-1},i}, \mathbf{x}_{t_m,i}, \mathbf{x}_{t_{m-1},i}, \dots, \mathbf{x}_{t_{m-1}+1,i}, u_i) dG(u_i) \quad (9)$$

where $f_1(\cdot)$ and $f(\cdot)$ are the pdfs of normal distributions with means given by the linear predictors of (8) and (4), respectively, and variances $\text{var}(\epsilon_{t_m,i})$ ($m = 1, \dots, M$), and $G(\cdot)$ is the cdf of a normal distribution with mean zero and variance σ_u^2 . The sample likelihood $\mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \sigma_u^2, \sigma_e^2, \sigma_{\epsilon_1}^2)$ is obtained by taking the product of \mathcal{L}_i over the n individuals in the sample. (The likelihood for the generalisation to multivariate responses is given in supplementary materials.) The model cannot be fitted using standard software for random effects models because of the non-linear constraints on the coefficients, random effects loading and residual variance. However, maximisation of the log-likelihood function can be carried out using standard optimisers available in statistical software. Alternatively, Bayesian estimation methods may be used.

For a univariate binary or ordinal response, the generalisation of the latent autoregression model for unequally spaced observations described in Section 3.2, together with a probit formulation of the model for the initial condition of Equation (8), can be framed as a type of multivariate probit model. The form of the likelihood is given in the supplementary materials. The latent autoregression model for equally spaced measurements on a univariate ordinal response can be estimated using maximum simulated likelihood (Pudney, 2008) or Bayesian methods (Stegmueller, 2013). We use Bayesian estimation via MCMC, extending Stegmueller's approach to handle unequally spaced multivariate repeated measures.

All model estimation in the simulation study and in the real-data application in Sections 4 and 5 was carried out using MCMC (specifically Gibbs sampling) as implemented in JAGS (Plummer, 2003) via the `runjags` R package (Denwood, 2016) (see Supplementary Materials for annotated code). Assuming a bivariate response, the following prior distributions were specified:

1. $U(-1, 1)$ priors for the autoregression parameters $\boldsymbol{\gamma}$ and $N(0, 1000)$ priors for all other regression coefficients and random effect loadings.
2. For univariate data, we assign Gamma priors for the random effect precisions and, for continuous responses, the residual variance: $\sigma_u^{-2}, \sigma_e^{-2} \sim \text{Gamma}(0.001, 0.001)$.
3. For bivariate data, we specify Wishart priors for the random effects precision matrix: $\boldsymbol{\Omega}_u^{-1} \sim \text{Wishart}(\nu, \mathbf{S}^{-1})$ with degrees of freedom $\nu = R = 2$ and scale matrix $\mathbf{S} = \mathbf{I}_2$. These common choices for ν and \mathbf{S} imply a relatively uninformative prior, but may be informative when the variances are small leading to overestimation of the variances which may affect estimates of the other parameters (Schuurman et al., 2016). We therefore consider the sensitivity of estimates to setting \mathbf{S} with diagonal elements equal to variance estimates obtained from fitting univariate models to each response and assuming Gamma priors for the random effects precisions.
4. For continuous bivariate responses, the same Wishart priors were assumed for $\boldsymbol{\Omega}_e^{-1}$. For binary or ordinal bivariate responses, where the residual variances are constrained to 1 for identification, we specify a hyperbolic tangent function for the residual correlation

$$\text{cor}(e_{1it}^*, e_{2it}^*) = \frac{\exp(\theta) - 1}{\exp(\theta) + 1}$$

and assign the prior $\theta \sim N(0, 1000)$.

If the model of interest includes x_t as a predictor, we require data for $(x_{t_m}, x_{t_m-1}, \dots, x_{t_m-s+1})$ when the spacing between two consecutive measurements on y is $\Delta t_m = s$ unit periods. In the application that follows, we additionally require x_{t_m-s} because of interest in transitions in x between $t-1$ and t . In practice, it is likely that there is missing data on at least one of the lags and we treat these missing values as parameters to be estimated in the Bayesian framework. We estimate an autoregressive logistic model for missing binary x_t

$$\begin{aligned} \text{logit}[\text{Pr}(x_{t_m-s,i} = 1)] &= \phi_2 \\ \text{logit}[\text{Pr}(x_{it} = 1 | x_{t-1}, i)] &= \phi_0 + \phi_1 x_{t-1,i} (t = t_m - s + 1, \dots, t_m) \end{aligned}$$

and specify $N(0, 1000)$ priors for ϕ_k ($k = 1, 2, 3$).

4 | SIMULATION STUDY

4.1 | Study design

A simulation study was carried out to assess the finite-sample performance of the methods proposed in Section 3 to handle unequally spaced panel data for continuous and binary responses. Univariate continuous y_{it} were generated from a dynamic model of form (2), while univariate binary y_{it} were generated from the latent autoregression model of Equation (6). We also considered bivariate responses y_{rit} ($r = 1, 2$) with lag and cross-lag effects, generated from Equation (7) and its latent generalisation for binary responses. In each case, the parameters of primary interest are the coefficients of a set of indicators for binary state transitions between $t-1$ and t , parameterised as in Equation (5). All DGPs incorporate a model for the initial condition.

Responses were initially generated for $T = 15$ unit periods, the length of the observation window (in years) in the application of Section 5. Unequal spacing was generated to follow the measurement schedule in the BHPS/UKHLS family network module by selecting observations $y_{t_m,i}$ where $\{t_m\} = (1, 6, 11, 13, 15)$, corresponding to the spacing $\{\Delta t_m\} = (5, 5, 2, 2)$. For bivariate responses, we also considered a measurement schedule that included intervals equal to the unit period by setting $T = 7$ and selecting observations at $\{t_m\} = (1, 3, 5, 6, 7)$ to produce spacing $\{\Delta t_m\} = (2, 2, 1, 1)$.

Our objective is to estimate the parameters of the DGP using the 5 observed unequally spaced y_{t_m} and 15 (or 7) annual measurements of x_t . Combinations of the following simulation conditions were considered, each with $n = 1000$ individuals:

1. Balanced ($M = 5$) and unbalanced ($M_i \leq 5$ with $\bar{M} = 3.6$) panels generated from a dropout model that assumes an increasing probability of dropout with occasion m .
2. Strong ($\gamma = 0.8$) and moderate ($\gamma = 0.4$) autoregression effect, as in Millimet and McDonough (2017).
3. Uncorrelated and serially correlated x_{it} . While Millimet and McDonough (2017) assumed a continuous time-varying covariate, we focus on changes in a binary x_{it} between $t-1$ and t , as in our application. High, moderate and low transition probabilities were considered to assess the impact of the amount of within-person variation in x_{it} . High transition probabilities (corresponding to uncorrelated x_{it}) were generated by taking independent $\text{Bin}(1, 0.5)$ draws of x_{it} for $t = 1, \dots, T$,

leading to transition probabilities of 0.25 in both directions. The moderate case assumed transition probabilities of 0.1, while the low case assumed a probability of 0.1 in one direction and 0.02 in the other, based on annual partnership transition probabilities in the BHPS/UKHLS data. For both the moderate and low scenarios, x_{it} was generated from a logistic regression model with predictor $x_{t-1,i}$ to impose serial correlation.

Three additional conditions were considered for univariate continuous and binary y_{it} : a reduction in the number of individuals to $n = 500$; a reduction in the number of waves in the DGP and observations to $T = 10$ and $M_i \leq 4$, respectively, with spacing $\{\Delta t_m\} = (5, 2, 2)$; and a skewed normal distribution (with a skewness coefficient of 0.75) for the individual random effects in the DGP. In each of these cases, we assumed an unbalanced design with $\gamma = 0.8$ and low transitions probabilities for x_{it} .

Apart from γ , all other parameters of the dynamic model for $t > 1$ were fixed across simulation conditions at the values given in Table 1 and Tables S1–S5. The parameter values of the DGPs for the initial condition are given in the supplementary material. For univariate y_{it} , 500 data sets were generated for each scenario. The number of replications for bivariate responses was reduced to 100 due to the increased estimation times. Joint models of the form (4) and (8), and its bivariate extension (see Supplementary Material), were fitted to the simulated data for $y_{t,m,i}$ ($m > 1$) and $y_{t,i}$, respectively, with

TABLE 1 Simulation results for 500 replicates of unbalanced designs with univariate binary y observed at intervals of (5, 5, 2, 2) for $n = 1000$ individuals. Estimates from dynamic model for $m > 1$, varying the autoregression parameter γ and the annual transition rate for binary x_t

Parameter ^a	True	$\gamma = 0.8$				$\gamma = 0.4$			
		Mean ^b	Mean SE ^c	SD ^d	95% cov ^e	Mean	Mean SE	SD	95% cov
$\Pr(x_t = s x_{t-1} = s') = 0.25$ for $s, s' = 0, 1; s \neq s'$									
β_0	0	-0.005	0.082	0.084	0.930	-0.004	0.068	0.071	0.920
γ	0.8/0.4	0.799	0.033	0.037	0.910	0.385	0.064	0.065	0.940
β_1	-0.2	-0.200	0.106	0.109	0.948	-0.203	0.089	0.093	0.944
β_2	-0.6	-0.613	0.161	0.161	0.948	-0.599	0.103	0.108	0.950
β_3	0.6	0.612	0.161	0.165	0.952	0.606	0.104	0.106	0.948
σ_u^2	1	1.054	0.176	0.182	0.944	1.066	0.167	0.170	0.936
$\Pr(x_t = 1 x_{t-1} = 0) = 0.11$ and $\Pr(x_t = 0 x_{t-1} = 1) = 0.02$									
β_0	0	-0.001	0.078	0.078	0.952	0.001	0.077	0.076	0.950
γ	0.8/0.4	0.797	0.034	0.037	0.910	0.377	0.065	0.068	0.914
β_1	-0.2	-0.205	0.082	0.081	0.950	-0.208	0.081	0.079	0.954
β_2	-0.6	-0.628	0.329	0.335	0.956	-0.613	0.241	0.228	0.970
β_3	0.6	0.599	0.378	0.382	0.952	0.588	0.267	0.275	0.940
σ_u^2	1	1.062	0.179	0.191	0.930	1.080	0.170	0.174	0.930

^aParameters are intercept (β_0), effect of transitions in binary x_t , as in Eq. (5), 11 versus 00 (β_1), 01 versus 00 (β_2), 10 versus 11 (β_3) and random effect variance (σ_u^2).

^bMean of the parameter estimates across replications where the estimates for a given replicate are computed as the posterior means of the chain values.

^cMean of the posterior standard deviations.

^dEmpirical SE, calculated as the standard deviation of the posterior means.

^eProportion of replications where the 95% credible interval includes the true parameter value.

the transition indicators and their lags replacing $x_{t_{m-1}+1,i}, \dots, x_{t_m,i}$ in Equation (4). All models were fitted using MCMC as implemented in JAGS, with the priors given in Section 3.5.

4.2 | Results

Selected results for univariate continuous y_{t_m} with observation intervals of (5, 5, 2, 2) are given in Table S1. We focus on unbalanced designs because the estimators were unbiased with good confidence interval coverage for the balanced case. The proposed random effects estimator for a univariate continuous response performs well in terms of bias and accuracy for different values of the autoregression parameter γ and of the within-person correlation in the binary predictor x_t (as measured by the transition probabilities between $t - 1$ and t). For all scenarios, the mean and empirical standard errors are similar and the confidence interval coverage probabilities are in general close to the nominal 95% level. Performance was also good for the three additional conditions described above ($n = 500$, $M_i \leq 4$ and skewed random effects distribution; see Table S2).

In the case of univariate binary y_{t_m} , accuracy and confidence interval coverage are acceptable for all conditions, but there are slightly larger biases in some parameters (Table 1 and Table S2). We find that γ is underestimated when $\gamma = 0.8$, with a relative bias of 3.8% or 5.8% according to the extent of within-person variation in x_t . When the transition probabilities for x_t are very low, there is some bias in the parameters associated with x_t regardless of the value of γ , with the largest relative bias estimated as 4.7% (for β_2 when $\gamma = 0.8$). We also note that in the binary case the standard errors of the parameters for the x transitions depend on the within-person variance in both x and y , where the latter is varied in the simulations by changing γ as σ_u^2 is fixed throughout; the effects of transitions in x_t are most precisely estimated when γ is small and the transition probabilities for x_t are large. Finally, in all scenarios, there is an upward bias of 5–8% in σ_u^2 , which was also apparent for balanced designs (results not shown). We would expect that increasing the number of observations of y per person, and the number of individuals, would lead to bias reduction. However, we do not consider this here as the fixed effects parameters are of primary substantive interest in our application.

Results for bivariate continuous and binary responses for an unbalanced design with moderate transition probabilities for x_t are given in the supplementary materials. We began with a full bivariate model with correlated residuals, random effects and cross-lag effects with observation gaps of (5, 5, 2, 2). For continuous $y_{r_{t_m}}$, there are small biases for some parameters, but the empirical standard errors were considerably larger than the mean standard errors, suggesting that the estimator is unstable (Table S3). However, the estimators for a simplified model without the cross-lag terms performs well (see Table S4 for binary $y_{r_{t_m}}$). Further investigation of the full bivariate model revealed that the instability in estimation was due to the observation intervals all being larger than the unit period in the DGP. Reducing the observation gaps to (2, 2, 1, 1) led to low bias and close agreement between the mean and empirical standard errors for both continuous and binary responses (Tables S5 and S6). This finding led us to exclude cross-lagged effects from the bivariate models fitted to the BHPS/UKHLS data.

5 | APPLICATION

5.1 | Background and research questions

In all societies, intergenerational transfers have major implications for individual, family and societal well-being (Mason & Lee, 2018). Private transfers between adult children and their parents are an

important element of intergenerational linkages and a means of providing support to those in need and spreading risks across the life course (Kunemud et al., 2005). In contemporary low-mortality societies, generational overlap is much higher than in the past and the co-survival of parents and their adult children often extends for longer than the child-rearing phase of life (Murphy et al., 2006). Associated increases in the numbers and proportions of older people imply an increase in the volume of help needed by those with age-related functional limitations. Needs for assistance may also be increasing in other age groups as a result of delayed transitions to adulthood, precarious employment, and increasingly diverse and complex family life courses (Henretta et al., 2018; Lesthaeghe, 2014). For all these reasons, extending our understanding of factors associated with exchanges of support between adult children and their parents is important.

The three main ‘currencies’ of intergenerational exchange between older parents and adult children are coresidence, provision of practical help and financial transfers. In the United States and some European countries, including the United Kingdom, the long-term trend towards declining intergenerational coresidence has recently been interrupted by a slight increase in coresidence rates between young adults and their parents, a trend attributed to the greater challenges young people now face in establishing and maintaining separate households (Stone et al., 2011). Exchanges of instrumental support (practical or financial help) between people living in different households nevertheless remain a more important component of intergenerational exchanges within families. Although practical help and financial gifts are somewhat different, they are to some extent substitutable, for example better-off adult children who are more engaged in labour market work provide more financial assistance, but less time help, to older parents (Bonsang, 2007).

Previous research indicates that family transfers operate within an exchange framework in which reciprocity, either contemporaneous or over the life course, is an important motivating factor (e.g. Silverstein & Bengtson, 1997) albeit one moderated by family culture (Grundy & Henretta, 2006). Grundy (2005), for example in a British study which used the same measure of help as employed here, found a strong reciprocal element to help given to and received from children by parents in late mid-life. Our analysis extends earlier work on intergenerational exchanges in two major ways. First, we use longitudinal data to examine the dynamics of exchanges of instrumental support between adult children and their parents, and in particular, the relationship between exchanges and time-varying child characteristics. Second, unlike most previous studies which have analysed support given separately from support received, we employ a bivariate model to estimate the degree of reciprocation between parents and children. We investigate the following research questions:

1. To what extent is the giving and receipt of help persistent over time? We expect that exchanges will have a strong persistence component until the occurrence of a ‘shock’ (e.g. changes in the child’s situation).
2. Are adult children with greater needs more likely to receive help from their parents? We consider indicators of need that relate to partnership and employment transitions and presence of young children.
3. Do changes in adult children’s circumstances affect their capacity to provide help to parents? For example, respondents with young children may have limited time and money to offer parental support.
4. Are exchanges reciprocated contemporaneously and over time? The bivariate model allows for correlations between giving and receiving help, which are expected to be positive.

5.2 | Data and measures

The BHPS began in 1991 with a sample of around 5500 households and 10,300 adults. These original sample members (OSMs) have been followed up and interviewed annually, and at wave 18 were asked to join the larger UKHLS. Information on exchanges of help with relatives living outside a respondent's household was collected as part of the rotating 'family network' module which was administered in 2001, 2006 (BHPS waves 11 and 16) and biennially thereafter in 2011–2013, 2013–2015 and 2015–2017 (UKHLS waves 3, 5 and 7). We analyse data for all five waves using a harmonised file with variables coded consistently across the two studies (ISER, 2017).

Respondents with a non-coresident parent were asked whether they had given the following types of help to their parent(s) 'regularly or frequently': lifts in a car, help with shopping, providing or cooking meals, help with personal needs, washing or cleaning, personal affairs, DIY or gardening, and financial help. The same questions were asked about receipt of support from parents, but with 'personal needs' replaced by help with childcare. We define two binary responses indicating whether any support was given to parents ($y_{1t_m,i}$) or received from parents ($y_{2t_m,i}$) by child (respondent) i at year t_m , the timing of measurement occasion $m = 1, 2, \dots, M_i$ where $M_i \leq 5$.

The analysis sample was first restricted to 18,999 person-wave observations contributed by BHPS OSMs when they were aged 18 or over and had at least one non-coresident parent living in the United Kingdom but no coresident parent. There were 1069 observations of $y_{rt_m,i}$ when respondents were living with one parent, mainly from younger respondents who had not yet left the parental home; they were excluded because the nature of their exchanges with the non-coresident parent are likely to differ from those of respondents who do not live with either parent. Individuals may move in and out of the target population over time; for example, a respondent becomes eligible for inclusion after leaving the parental home and becomes ineligible following the death of both parents or a parent moving into their household. These respondents will have $M_i < 5$, as will individuals who do not respond to the survey request at every wave. Our approach accommodates increased gaps between observations for either reason under a missing at random assumption. In the case of wave non-response (or attrition), an alternative approach which allows for the possibility that data are not missing at random would be to estimate a non-response model jointly with the model of interest (e.g. Washbrook et al., 2014). Respondents with measurements of $y_{rt_m,i}$ at only one wave ($n = 2587$) were excluded because they contribute only to the estimation of the parameters of the model for the initial condition. The sample was further restricted to person-waves m such that $\Delta t_m \leq 5$ years, leading to a further drop of 270 observations.

We consider the following covariates: respondent's gender, travel time to the nearest (UK-based) parent, and time-varying measures of a respondent's age (which is highly correlated with parental age), their current partnership and employment status, the presence of children and the age of the youngest child, and partnership and employment transitions in the last year. Travel time was measured on a 5-point ordinal scale, dichotomised at a threshold of 1 h. This variable was collected at the same waves as the parental exchange data but remained constant across the observation period for 89% of respondents. We therefore use a time-invariant variable based on the first observation. All time-varying covariates are based on annual data and were computed for each year between a person's first and last measurements of help given and received. At each wave, the partnership status and age of each individual in a responding household, and their relationship to the household reference person, was collected at the interviewer's first contact with an adult member of the household. This information was used to derive time-varying indicators of partnership status, annual partnership transitions and the presence of biological or adopted children of different ages. Indicators of employment status and transitions were derived from data collected directly from respondents through the individual

questionnaire. In preliminary analysis allowing the effects of employment transitions to differ for men and women, economic activity was grouped into four categories: employed full-time (60.5% of person-waves), employed part-time (17.0%), unemployed (3.3%) and out of the labour market (19.2%). Based on this analysis, a simple binary classification (employed vs. non-employed) was found to be adequate for both men and women.

Complete data on partnership status and the number and age of children were available at all occasions m at which $y_{rt_{m,i}}$ was measured, and data on economic activity were missing for only six records. However, there were some missing data on the values of time-varying variables between the measurements of y , which must be included in the estimation model for the observed data $y_{rt_{m,i}}$ ($m > 1$), as described in Section 3.1. Information on children was missing for at least one of the required lags $x_{t_{m-1}}, \dots, x_{t_{m-1}+1}$ for 467 person-wave observations. This reduced to 189 observations after using a deductive imputation rule where for missing data at year k , and non-missing data at years $k - 1$ and $k + 1$, we set $x_k = x_{k-1}$ when $x_{k-1} = x_{k+1}$. Observations for which the child lags could not be imputed were excluded, but missing data on lags of partnership and employment status for 412 person-waves were treated as additional unknown quantities in the Bayesian model for $y_{rt_{m,i}}$ (see Section 3.5). A small amount of missingness on travel time to the nearest parent (57 respondents) was handled in a similar way.

The final analysis sample contains 4839 individuals who contribute 15,904 repeated measurements of the bivariate response for exchanges of support, a mean of 3.3 observations per respondent. Descriptive statistics for all variables are given in Table 2.

5.3 | Results

In preliminary analysis, we allowed for gender-specific effects of employment transitions on exchanges of support in both directions, distinguishing between full-time and part-time employment for women, and for gender-specific reciprocity (the random effect and residual correlations). As there was little evidence of any gender differences in these parameters, or that the effects of employment transitions differed for part-time and full-time work, the model was simplified by including only a main effect of gender and using a binary indicator for employment status. We also considered whether the effects of partnership transitions depend on the presence and age of children; for example, any increase in parental support following dissolution might be stronger when there are young children involved. However, we found no evidence of any such interaction effects. To allow for the possibility that the indicators for employment and partnership transitions and the presence and age of the youngest children are correlated with the individual random effects, a correlated random effects model was fitted in which the individual means of these time-varying covariates were included as additional predictors (see Section 2). There is little difference in the substantive conclusions between the two models, but the coefficients of the time-varying covariates were estimated less precisely with wider credible intervals when the individual means were included. We therefore conclude that these covariates may be considered independent of the random effects and focus on the model without individual means. (The results for the correlated random effects model are given in Table S8.) Finally, cross-lagged effects were excluded, following the finding from the simulation study in Section 4.2 that their inclusion leads to instability for data structures such as ours where all observed intervals Δt_m are larger than the unit period in the DGP.

The results for the final specification are given in Table 3 and Table S7. These are based on five parallel chains of 20,000 MCMC iterations, each using a different set of starting values and with a burn-in sample of 10,000. Convergence was assessed using a range of graphical diagnostics and the

TABLE 2 Descriptive statistics for response variables and covariates

Variable	Number	Per cent
Measured at t_m ($n = 15,904$)		
Any help given to parent(s) (y_{1t_m})	7575	47.6
Any help received from parent(s) (y_{2t_m})	6774	42.6
Partnership status at t_m		
Partnered	12,175	76.6
Single	3729	23.4
Employment status at t_m		
Employed	12,356	77.7
Non-employed	3548	22.3
Children at t_m		
No children	6217	39.1
Child(ren), youngest < 2 years	1056	6.6
Child(ren), youngest 2–4 years	1527	9.6
Child(ren), youngest 5–10 years	2591	16.3
Child(ren), youngest 11–16 years	2221	14.0
Child(ren), youngest > 16 years	2292	14.1
Respondent age at t_m	Mean = 42.3	SD = 11.0
Individual-level ($n = 4839$)		
Respondent is female	2793	57.7
Travel time to nearest parent more than 1 h	1112	23.2
Transition probabilities (annual)	Probability	
Partnership		
Pr(partner at t single at $t - 1$)	0.111	
Pr(single at t partner at $t - 1$)	0.023	
Employment		
Pr(employed at t non-employed at $t - 1$)	0.190	
Pr(non-employed at t employed at $t - 1$)	0.053	

potential scale reduction factor (PSRF) (Gelman et al., 2004). Visual inspection of trace plots of each parameter for the multiple chains suggested adequate mixing. Final PSRF estimates were close to 1 for all parameters. Furthermore, increasing the chain length led to little change in the running means of the posterior estimates. The results presented are based on a Wishart prior for the random effects precision matrix with \mathbf{S} equal to the identity matrix (see Section 3.5). An alternative prior with the diagonal elements of \mathbf{S} equal to variance estimates obtained from univariate models with Gamma priors led to slightly larger variance estimates and smaller autocorrelation effects, but little change in the other parameters.

The results from the model for the initial condition (Table S7) show the effects of covariates on a respondent's propensity to give and receive help at t_1 , the first year at which they were observed to have a non-coresident parent. Of primary interest, however, are the dynamics of exchanges and the effects of time-varying covariates on exchanges at t conditional on exchanges in the previous year. The autoregressive parameters (γ_r) are the polychoric correlations between y_{rit}^* and $y_{rit-1,i}^*$, conditional on

TABLE 3 Results from bivariate latent dynamic model for any help given to and received from parents at time $t > 1$. Parameters for which the 95% credible interval does not contain zero are in bold

	To parents ($r = 1$)		From parents ($r = 2$)	
	Mean	95% CI	Mean	95% CI
Constant	0.031	(-0.033, 0.096)	-0.029	(-0.096, 0.043)
Lagged response $y_{r,t-1}^*$	0.676	(0.612, 0.735)	0.548	(0.476, 0.618)
Age at t (years)	0.016	(0.012, 0.019)	-0.024	(-0.029, -0.019)
Age squared	-0.0001	(-0.0003, 0.00003)	-0.0003	(-0.0006, -0.0001)
Female	0.113	(0.068, 0.159)	0.139	(0.086, 0.192)
Closest parent >1 h away	-0.475	(-0.573, -0.381)	-0.434	(-0.525, -0.351)
Partnership transitions in $(t-1, t)$				
Partnered at $t - 1^a$	-0.054	(-0.108, -0.0003)	-0.310	(-0.381, -0.241)
Partnership formation	-0.055	(-0.286, 0.180)	-0.108	(-0.339, 0.129)
Partnership dissolution	-0.106	(-0.371, 0.154)	0.463	(0.204, 0.728)
Employment transitions in $(t - 1, t)$				
Not employed at t^b	-0.024	(-0.085, 0.035)	-0.059	(-0.126, 0.010)
Move out of employment	0.099	(-0.081, 0.277)	0.122	(-0.064, 0.305)
Move into employment	0.002	(-0.189, 0.187)	-0.153	(-0.347, 0.040)
Presence of children and age of youngest (years) at t				
Any child ^c	-0.106	(-0.222, 0.013)	0.628	(0.494, 0.755)
Children, youngest age 2-4 ^d	0.131	(-0.023, 0.286)	-0.120	(-0.278, 0.031)
Children, youngest age 5-10 ^d	0.057	(-0.060, 0.172)	-0.165	(-0.289, -0.044)
Children, youngest age 11-16 ^d	0.079	(-0.045, 0.204)	-0.581	(-0.710, -0.448)
Children, youngest age > 16 ^d	0.060	(-0.066, 0.187)	-0.658	(-0.799, -0.514)
Random effect variance (σ_{ur}^2)	0.203	(0.121, 0.296)	0.236	(0.150, 0.336)
Random effect correlation (ρ_{u12})	0.345	(0.286, 0.403)		
Residual correlation (ρ_{e12})	0.467	(0.425, 0.509)		

^aContrast of partnered at $t - 1$ and t versus single at both years.

^bContrast of not employed at $t - 1$ and t versus employed at both years.

^cContrast of any children and youngest <2 years versus no children.

^dEffects of age of youngest child among respondents with children (reference is youngest child <2 years).

covariates and the individual random effect. In answer to research question (i) in Section 5.1, there are moderate to strong lag effects for exchanges in both directions, suggesting a high degree of persistence in behaviour over the observation period (consistent with Grundy and Henretta (2006)) after accounting for changes in the child's circumstances. In contrast, there is a relatively small amount of unobserved heterogeneity; the proportion of the total unexplained variance in y_{rt}^* that can be attributed to unobserved time-invariant characteristics is estimated as 17% and 19% for 'to' and 'from' exchanges, respectively (using $\sigma_{ru}^2 / (1 + \sigma_{ru}^2)$).

The coefficients of the covariates are their effects on the continuous latent variable y_{rti}^* underlying the observed binary response y_{rti} ($r = 1, 2$), conditional on the lag $y_{rt-1,i}^*$. After scaling by the factor $1 / \sqrt{1 + \sigma_{ru}^2}$, the coefficients represent effects in standard deviation units on the latent

propensities to give or receive help. For instance, the female–male difference in the propensity to help parents is estimated as $0.113 / \sqrt{1.203} = 0.103$ sd units. Starting with the time-invariant covariates, we find that women are more likely than men to both give and receive parental support, while a travel time of more than 1 h between an adult child and their parents is associated with a reduction in the propensity of exchanges in either direction. However, the effect of distance should be interpreted with caution due to the possibility that children or parents move closer to each other when either becomes in greater need of help. As expected, older respondents (who tend to have older parents) are more likely to give and less likely to receive support than their younger counterparts, with the quadratic effect implying an acceleration in the negative effect of age on giving help at older ages.

We next consider the effects of changes in a child's circumstances on receiving help from their parents (research question (ii)). There is evidence that the experience of partnership breakdown in the last year is associated with an increase in the propensity to receive parental support. There is no change in the propensity to receive support following the start of a new (coresidential) partnership, but we find that respondents who were partnered at both the current and previous year are less likely to receive help than those who were single in both years. There is little evidence of parents responding to changes in their child's employment status, although the effects of employment transitions are in the expected directions with a move out of employment associated with an increase in the propensity to receive help and a move into employment associated with a decrease. Respondents are more likely to receive help from their parents when they have young children, but the effect of the presence of children decreases with the age of the youngest child to the extent that respondents whose youngest is 11 or older are no more likely to receive parental support than respondents without children. While there is evidence that parents respond to the changing needs of their children when giving support, we find that changes in the child's circumstances do not influence the provision of support to their parents (research question (iii)).

Finally, we turn to question (iv) on reciprocity of exchanges, represented in the model by the random effect and residual correlations between giving and receiving parental support. Both correlations are positive which suggests that children with a higher propensity to help their parents tend to have a higher propensity to receive help in return. The random effect correlation captures unmeasured time-invariant characteristics of the respondent, their parent(s) and the dyad (for example, the quality of the child–parent relationship and family norms of behaviour). Reciprocity due to unmeasured time-varying factors is captured by the residual correlation which measures the association between giving and receiving help in a given year. The relatively large residual correlation may be due to the omission of time-varying covariates relating to the parents' circumstances.

6 | DISCUSSION

Unequal spacing between measurements is commonplace in longitudinal studies, but discrete-time models which include lagged responses as predictors traditionally assume that responses and covariates are available at the same equidistant time points. The general random effects estimator proposed here can handle unequally spaced continuous and categorical multivariate responses and is straightforward to estimate using MCMC methods. A bivariate model is applied in an analysis of the exchanges of support between adult children and their non-coresident parents, with a focus on the effects of changes in children's circumstances and the extent of reciprocity of exchanges. The proposed method makes efficient use of all available data on the responses and covariates of interest:

binary indicators of giving and receiving parental support collected at intervals of 5 and 2 years, and annual panel data on time-varying socio-demographic characteristics. We find that parents are more likely to support their children following a recent partnership dissolution and when there are young grandchildren, but changes in children's circumstances do not appear to influence the provision of support to their parents. We might expect that children are more likely to react to changes in their parent's needs, especially as they grow older, but we are unable to pursue this question in our analysis of exchanges from a child's perspective because limited information was collected on non-coresident relatives. Future research could use the same approach to study the provision and receipt of support from a parental perspective using data collected on exchanges between respondents (parents) and non-coresident adult children.

In our application, time-varying covariates x_t are available at more regular intervals than y_t , but designs where (y_t, x_t) are observed at the same times can be accommodated using data augmentation in MCMC (as we have done to handle partial missingness on x_t). Viewing unequally spaced measures of y_t as a missing data problem, alternative approaches include data augmentation or multiple imputation. However, their implementation would be considerably more computationally intensive than the proposed method, especially for long observation periods with large gaps between the observed y_m . In our application, for example we would have to impute 10 values of y_t for individuals observed at 5 years over the 15-year period. Moreover, it would not be straightforward to specify an imputation model that is consistent with the assumed autoregressive DGP.

Based on the findings of the simulation study, caution should be exercised when including cross-lagged effects in bivariate models for unequally spaced responses. Our results suggest that their inclusion leads to instability and, in the binary case, bias when the observed intervals between responses are all larger than the unit period in the assumed DGP, as in our application. For this reason, we were unable to pursue investigation of reciprocal effects between giving and receiving help over time in our analysis, and we can conclude only that there is a positive association between giving and receiving based on the random effect and residual correlations. Even if cross-lagged effects could be estimated, simple first-order effects are unlikely to adequately capture the complexity of reciprocity in exchanges over time. For example, reciprocity may occur over a longer period than 1 year and the length of the lag between giving and receiving may depend on the age of the respondent and their parent and the individual circumstances of each. Another issue with cross-lagged effects is that their magnitude and direction will depend on the assumed time interval in discrete-time models, which is usually determined by the gap between measurements rather than substantive considerations. An alternative to the discrete-time approach, which has been proposed for longitudinal models with cross-lag effects among multivariate responses, is a continuous-time model estimated via stochastic differential equations (Voelke et al., 2012). An attractive feature of continuous-time models is that they place no requirements on the spacing of measurements, but they are more difficult to implement and interpret than discrete-time models which have a strong tradition in the health and social sciences.

ACKNOWLEDGEMENTS

This research was supported by a UK Economic and Social Research Council (ESRC) grant 'Methods for the Analysis of Longitudinal Dyadic Data with an Application to Inter-generational Exchanges of Family Support' (ref. ES/P000118/1).

REFERENCES

- Ahn, S. & Schmidt, P. (1995) Efficient estimation of models for dynamic panel data. *Journal of Econometrics*, 68, 5–27.
- Alessie, R., Hochguertel, S. & van Soest, A. (2004) Ownership of stocks and mutual funds: A panel data analysis. *The Review of Economics and Statistics*, 86, 783–796.

- Arellano, M. & Bond, S. (1991) Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies*, 58, 277–297.
- Arulampalam, W., Booth, A.L. & Taylor, M.P. (2000) Unemployment persistence. *Oxford Economic Papers*, 52, 24–50.
- Baltagi, B.H. & Wu, P.X. (1999) Unequally spaced panel data regressions with AR(1) disturbances. *Econometric Theory*, 15, 814–823.
- Bonsang, E. (2007) How do middle-aged children allocate time and money transfers to their older parents in Europe? *Empirica*, 34, 171–188.
- Clarke, P., Crawford, C., Steele, F. & Vignoles, A. (2015) Revisiting fixed- and random-effects models: Some considerations for policy-relevant education research. *Education Economics*, 23, 259–277.
- Denwood, M.J. (2016) runjags: An R package providing interface utilities, model templates, parallel computing methods and additional distributions for MCMC models in JAGS. *Journal of Statistical Software*, 71, 1. Available from: <http://runjags.sourceforge.net>.
- Gelman, A., Carlin, J.B., Stern, H.S. & Rubin, D.B. (2004) *Bayesian data analysis*, 2nd edn. Boca Raton, FL: Chapman and Hall/CRC.
- Grundy, E. (2005) Reciprocity in relationships: Socio-economic and health influences on intergenerational exchanges between third age parents and their adult children in Great Britain. *British Journal of Sociology*, 56, 233–255.
- Grundy, E. & Henretta, J.C. (2006) Between elderly parents and adult children: A new look at the intergenerational care provided by the ‘sandwich generation’. *Ageing and Society*, 26, 707–722.
- Heckman, J.J. (1981) Statistical models for discrete panel data. In: Manski, C. & McFadden, D. (Eds.), *Structural analysis of discrete data with econometric applications*. Cambridge: MIT Press, pp. 227–269.
- Henretta, J.C., Van Voorhis, M.F. & Soldo, B.J. (2018) Cohort differences in parental financial help to adult children. *Demography*, 55, 1567–1582.
- Honoré, B. & Kyriazidou, E. (2000) Discrete choice models with lagged dependent variables. *Econometrica*, 68, 839–874.
- ISER. (2017) Understanding Society: Waves 1-7, 2009–2016 and Harmonised BHPS: Waves 1-18, 1991–2009. [data collection]. University of Essex, Institute for Social and Economic Research. UK Data Service. SN: 6614, 9th edn. Available from: <http://doi.org/10.5255/UKDA-SN-6614-10>.
- Jones, R.H. & Boadi-Boateng, F. (1991) Unequally spaced longitudinal data with AR(1) serial correlation. *Biometrics*, 47, 161–175.
- Kazemi, I. & Crouchley, R. (2006) Modelling the initial conditions in dynamic regression models of panel data with random effects. In: Baltagi, B. (Ed.), *Panel data econometrics: Theoretical contributions and empirical applications* (pp. 91–117). Amsterdam: Elsevier.
- Kunemud, H., Motel-Lingebiel, A. & Kohli, M. (2005) Do intergenerational transfers from elderly parents increase social inequality among their middle-aged children? Evidence from the German Ageing Survey. *Journal of Gerontology: Social Sciences*, 60B, S30–S36.
- Lesthaeghe, R. (2014) The second demographic transition: A concise overview of its development. *Proceedings of the National Academy of Sciences of the United States of America*, 111, 18112–18115.
- Mason, A. & Lee, R. (2018) Intergenerational transfers and the older population. In: Hayward, M. & Majmundar, M. (Eds.), *Future directions for the demography of aging: Proceedings of a workshop* (pp. 187–214). Washington, DC: The National Academies Press.
- McCulloch, C.E. & Neuhaus, J.M. (2011) Misspecifying the shape of a random effects distribution: Why getting it wrong may not matter. *Statistical Science*, 26, 388–402.
- McKenzie, D.J. (2001) Estimation of AR(1) models with unequally spaced pseudopanel. *Econometrics Journal*, 4, 89–108.
- Millimet, D.L. & McDonough, I.K. (2017) Dynamic panel data models with irregular spacing: With an application to early childhood development. *Journal of Applied Econometrics*, 32, 725–743.
- Mundlak, Y. (1978) On the pooling of time series and cross section data. *Econometrica*, 46, 69–85.
- Murphy, M., Martikainen, P. & Pennec, S. (2006) Demographic change and the supply of potential family supporters in Britain, Finland and France in the period 1911–2050. *European Journal of Population*, 22, 219–240.
- Pacini, D. & Windmeijer, F. (2015) Moment conditions for AR(1) panel data models with missing outcomes. *Bristol Economics Discussion Papers 15/660*, School of Economics, University of Bristol, UK. Available from: <https://ideas.repec.org/p/bri/uobdis/15-660.html>.

- Plummer, M. (2003) JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In: Hornik, K., Leisch, F. & Zeileis, A. (Eds.) 3rd International Workshop on Distributed Statistical Computing, pp. 20–22.
- Pudney, S. (2008) The dynamics of perception: Modelling subjective wellbeing in a short panel. *Journal of the Royal Statistical Society A*, 171, 21–40.
- Rosner, B. & Muñoz, A. (1988) Autoregressive modelling for the analysis of longitudinal data with unequally spaced examinations. *Statistics in Medicine*, 7, 59–71.
- Sasaki, Y. & Xin, Y. (2017) Unequal spacing in dynamic panel data: Identification and estimation. *Journal of Econometrics*, 196, 320–330.
- Schuurman, N.K., Grasman, R.P.P.P. & Hamaker, E.L. (2016) A comparison of inverse-Wishart prior specifications for covariance matrices in multilevel autoregressive models. *Multivariate Behavioral Research*, 51, 185–206.
- Selig, J.P. & Little, T. (2012) Autoregressive and cross-lagged panel analysis for longitudinal data. In: Laursen, B. & Little, T.D. (Eds.), *Handbook of developmental research methods*. New York: Guilford Press, pp. 265–278.
- Silverstein, M. & Bengtson, V.L. (1997) Intergenerational solidarity and the structure of adult child–parent relationships in American families. *American Journal of Sociology*, 103, 429–460.
- Steele, F., French, R. & Bartley, M. (2013) Adjusting for selection bias in longitudinal analyses using simultaneous equations modelling: The relationship between employment transitions and mental health. *Epidemiology*, 24, 703–711.
- Stegmueller, D. (2013) Modeling dynamic preferences: A Bayesian robust dynamic latent ordered probit model. *Political Analysis*, 21, 314–333.
- Stone, J., Berrington, A. & Falkingham, J. (2011) The changing determinants of UK young adults' living arrangements. *Demographic Research*, 25, 629–666.
- Voelke, M.C., Oud, J.H.L., Davidov, E. & Schmidt, P. (2012) An SEM approach to continuous time modeling of panel data: Relating authoritarianism and anomia. *Psychological Methods*, 17, 176–192.
- Washbrook, E., Clarke, P.S. & Steele, F. (2014) Investigating non-ignorable dropout in panel studies of residential mobility. *Journal of the Royal Statistical Society C (Applied Statistics)*, 63, 239–266.
- Wooldridge, J. (2010) *Econometric analysis of cross section and panel data*, 2nd edn. Cambridge, MA: The MIT Press.
- Wooldridge, J. (2019) Correlated random effects models with unbalanced panels. *Journal of Econometrics*, 211, 137–150.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section.

How to cite this article: Steele F, Grundy E. Random effects dynamic panel models for unequally spaced multivariate categorical repeated measures: an application to child–parent exchanges of support. *J R Stat Soc Series C*. 2020;00:1–21. <https://doi.org/10.1111/rssc.12446>