

Joint modelling compared with two stage methods for analysing longitudinal data and prospective outcomes: A simulation study of childhood growth and BP

Statistical Methods in Medical Research

2017, Vol. 26(1) 437–452



© The Author(s) 2014

Reprints and permissions:

sagepub.co.uk/journalsPermissions.nav

DOI: 10.1177/0962280214548822

smm.sagepub.com



A Sayers,¹ J Heron,¹ ADAC Smith,^{1,2} C Macdonald-Wallis,^{1,2}
MS Gilthorpe,³ F Steele⁴ and K Tilling^{1,2}

Abstract

There is a growing debate with regards to the appropriate methods of analysis of growth trajectories and their association with prospective dependent outcomes. Using the example of childhood growth and adult BP, we conducted an extensive simulation study to explore four two-stage and two joint modelling methods, and compared their bias and coverage in estimation of the (unconditional) association between birth length and later BP, and the association between growth rate and later BP (conditional on birth length). We show that the two-stage method of using multilevel models to estimate growth parameters and relating these to outcome gives unbiased estimates of the conditional associations between growth and outcome. Using simulations, we demonstrate that the simple methods resulted in bias in the presence of measurement error, as did the two-stage multilevel method when looking at the total (unconditional) association of birth length with outcome. The two joint modelling methods gave unbiased results, but using the re-inflated residuals led to undercoverage of the confidence intervals. We conclude that either joint modelling or the simpler two-stage multilevel approach can be used to estimate conditional associations between growth and later outcomes, but that only joint modelling is unbiased with nominal coverage for unconditional associations.

Keywords

lifecourse epidemiology, joint model, multilevel model, measurement error, growth

¹School of Social and Community Medicine, University of Bristol, Bristol, UK

²MRC Integrative Epidemiology Unit, University of Bristol, Bristol, UK

³Division of Epidemiology and Biostatistics, School of Medicine, University of Leeds, Leeds, UK

⁴Department of Statistics, London School of Economics, London, UK

Corresponding author:

Kate Tilling, School of Social and Community Medicine, University of Bristol, Canynge Hall, 39 Whatley Road, Bristol BS8 2PS, UK.

Email: kate.tilling@bristol.ac.uk

I Introduction

Increasingly in epidemiology and medical research, there is interest in the relationship between both baseline level of, and change in, an exposure and a future outcome. Examples include changes in biomarkers in relation to disease incidence or progression (such as prostate specific antigen (PSA) in relation to progression of prostate cancer);¹ changes in physiological variables in relation to disease outcomes (e.g. changes in renal function and subsequent arterial stiffness);² and associations between physical and cognitive changes in later life.³ One change hypothesis that has been widely explored is the association between birth size, childhood growth and adult outcomes such as blood pressure (BP),^{4–8} and it is this example we focus on here – however, the methods are applicable to associations between linear change in any continuous exposure and a distal outcome.

We consider the example of a study aiming to investigate the association between birth length and linear growth rate during childhood on adult blood pressure (BP). We are interested in the total effect of birth length (B) on BP, and in the effect of linear growth rate (G) on BP, conditioning on birth length (Figure 1). Additional complexities that may occur in real data, such as non-linear growth, and an interaction between birth length and growth rate in their effect on the outcome are not considered. Studies relating birth length and growth to BP typically collect repeated data on height (length measured at birth and on at least one further occasion during childhood), and a later (adult) measure of BP. In our example, we assume that length is measured within 2 weeks of birth, and then height is measured at approximately 2.5, 5, 7.5 and 10 years of age.

A common approach is first to summarise the repeated measures of height, and then relate these summaries to the subsequent outcome. One simple method for carrying out the summary stage is to use the observed birth length and the latest childhood measure of height (or birth length and change between birth length and final height) as exposures;⁹ an alternative is to regress height on age within each individual to estimate their birth length and growth rate.¹⁰ More complex methods include using multilevel or other repeated measures methods to model the trajectories of height, and extracting summaries of these trajectories such as birth length and growth.^{11,12} Irrespective of the method used to summarise changes in exposure, linear regression models are then used to relate the summaries of change in height to the adult systolic BP, meaning that this stage can be carried out using standard statistical software. All the two-stage approaches share the potential problem that uncertainty in the estimates of birth length and growth are not taken into account in the confidence intervals for their associations with BP, meaning that standard errors may be underestimated.




Figure 1. A schematic representation of causal associations between birth length, growth rate and BP.

This problem can be viewed in a measurement error framework, where the regression of blood pressure on birth height and growth is biased by the measurement error/intra-individual variation in height. Measurement error/intra-individual variation in the exposure (height) would tend to attenuate the coefficients between BP and height towards the null in the simple methods. It has been noted in the measurement-error literature that the two-stage multilevel method (also termed the regression calibration, or ‘RC’ method) provides consistent conditional effect estimates when the model relating exposures to outcome is linear,¹³ and that the individual regression method results in biased estimates of the conditional effects. However, more research is needed into the performance of all two-stage methods in estimating unconditional effects, and into their relative bias and coverage under different conditions.

In contrast to these two-stage approaches, the joint modelling approach aims to model birth length, growth and BP simultaneously, often using a bivariate growth model,¹⁴ or in a structural equation modelling framework.¹⁵ Joint modelling approaches are becoming more widely implemented in mainstream statistical software. However, joint models are more complex than the two-stage approaches outlined earlier, and it is not known whether any bias or under-coverage in the two-stage methods is large enough to warrant this extra modelling complexity.

In this simulation study, we compare the bias and coverage under different study conditions of six methods to estimate the association between birth length, linear growth and later BP. The two-stage methods we examine are: (1) a simple approach, which uses the observed birth length, and the difference between birth length and latest height measure (here, at approximately 10 years), divided by the difference between the ages at the two measurements, as an estimate of growth rate; (2) an individual regression (OLS) approach which estimates an individual’s birth length and growth rate from the parameters of the model regressing that individual’s height measures on age; (3) a multilevel growth model (MLM) where estimates of the individual random effects (shrunken residuals (or (4) inflated residuals) are used as estimates of birth length and growth rate. These are compared to two versions of a joint modelling approach: (5) a bivariate MLM, in which child growth and the adult outcome are modelled simultaneously and re-inflated residuals used to estimate the association between growth and outcome and (6) a structural equation model (SEM) which provides an alternative parameterisation of the bivariate model. In this study, we compare the bias and efficiency of using joint modelling and two-stage methods under different assumptions about the between- and within-individual variances of birth length and growth rate, and the relationship between them.

2 Methods

2.1 Simulation study design

We simulated a study where length is measured close to birth and at approximate ages of 2.5, 5, 7.5 and 10 years ($I = 5$ measurement occasions per individual). The exact ages (Age_{ij}) at which height was measured for individual j at occasion i were drawn from a multivariate normal distribution with mean ages of 0, 2.5, 5, 7.5 and 10 years, and no covariance between occasions. The standard deviations of age at measurement at occasion i ($\sigma_{age\ i}$) were 0.5, with the exception of age 0 ($\sigma_{age\ 0}$) which is 0.025 (approximately 95% of individuals had measurements within ± 2 weeks of zero). A single draw from the age distribution was used for all simulations, and age of measurement was checked to ensure that it was monotonically increasing across measurement occasions for each individual. The number of individuals in the study (J) was set to 1000, to illustrate a medium/large cohort study, and all individuals had all five measurements of height.

We simulated height for individual j at occasion i (H_{ij}) from age for that individual at that occasion (Age_{ij}) using a random intercept and random linear slope growth model

$$\begin{aligned} H_{ij} &= \beta_0 + u_{0j} + (\beta_1 + u_{1j})\text{Age}_{ij} + e_{hij} \\ \begin{bmatrix} u_{0j} \\ u_{1j} \end{bmatrix} &\sim N(0, \Omega_u) : \Omega_u = \begin{bmatrix} \sigma_{u0}^2 & \\ & \sigma_{u1}^2 \end{bmatrix} \\ \begin{bmatrix} e_{hij} \end{bmatrix} &\sim iid N(0, \sigma_{eh}^2) \end{aligned} \quad (1)$$

Here, β_0 corresponds to the population average birth length, and β_1 to the population average linear growth rate. Random variables u_{0j} and u_{1j} correspond to the deviations from the population birth length and linear growth rate, respectively, for individual j , and e_{hij} represents the occasion level (level 1) residual for height for individual j at occasion i .

We then simulated a linear association between BP (BP_j) and deviation from the population average birth length (u_{0j}) and deviation from the population average growth rate (u_{1j}) for individual j

$$\begin{aligned} \text{BP}_j &= \gamma_0 + \alpha_3 u_{0j} + \alpha_4 u_{1j} + e_{\text{BP}j} \\ \begin{bmatrix} e_{\text{BP}j} \end{bmatrix} &\sim iid N(0, \sigma_{e\text{BP}}^2) \end{aligned} \quad (2)$$

Here, γ_0 represents the population average systolic BP, α_3 and α_4 represent the effect of a 1 cm increase in birth length and a 1 cm y^{-1} increase in growth rate on systolic BP (mmHg) for individual j , conditional on both being included in the model.

2.2 Two-stage methods

Two-stage methods of analysis attempt to (1) summarise the growth process, i.e. birth length and growth rate, and the covariance between these (where the method considers this explicitly) and (2) use these summaries to investigate the associations of birth length and growth rate with the outcome of interest (BP). We begin by describing the second stage as this is common to all two-stage methods. We then outline the first stage, which differs for each approach.

The second stage starts by assuming that we have estimates of the birth length (\hat{B}_j) and linear growth rate (\hat{G}_j) of each individual j ($j = 1, \dots, J$) from the first stage. These are then related to the outcome of interest using two simple linear regression models. The first regression model assesses the total effect of the estimated birth length (\hat{B}_j) of the j th individual ($j = 1, \dots, J$) on their adult BP (BP_j).

$$\text{BP}_j = \alpha_0 + \alpha_1 \hat{B}_j + \varepsilon_{1j} \quad \text{and} \quad \varepsilon_{1j} \sim N(0, \sigma_{\varepsilon_1}^2) \quad (3)$$

The second regression models the association between the estimated growth rate (\hat{G}_j) and BP (BP_j) conditional on birth length (\hat{B}_j), where α_4 is the parameter of interest. α_3 is not of specific interest and cannot be directly interpreted given the potential of the reversal paradox.¹⁶

$$\text{BP}_j = \alpha_2 + \alpha_3 \hat{B}_j + \alpha_4 \hat{G}_j + \varepsilon_{2j} \quad \text{and} \quad \varepsilon_{2j} \sim N(0, \sigma_{\varepsilon_2}^2) \quad (4)$$

As the second stage of the analysis is estimated using OLS regression, the parameter estimates and their standard errors are obtained using standard methods.

2.3 The first stage

The first stage attempts to summarise the growth trajectory for individual j by estimating the true birth length (B_j) and linear growth rate (G_j) for each individual. The three two-stage methods considered here (the simple approach, the individual regression approach, and a multilevel model (MLM) for growth) differ in how they estimate birth length B_j and growth rate G_j .

2.3.1 The simple approach

The simplest approach summarises the growth trajectory using the first observed measurement as an estimate of birth length, and the difference between the latest height measure (here, height at age of approximately 10 years) and birth length, divided by the elapsed time between measurements, as an estimate of growth rate. This idea is equivalent to that suggested by Lucas et al.,⁸ with the exception that they suggest adjusting for final height, which they recognise as a simple reparameterisation. This method assumes that the first observed height measure (H_{1j}) is the best estimate of birth length. Justifiability of this assumption would depend on the timing of this first measurement. This simple method also assumes that the later measures of height provide no further information about birth length. Similarly, the estimate of growth rate is characterised by the difference between height at the first (H_{1j}) and last measurement occasion (H_{Ij}), and assumes that intermediate measurements are uninformative.

$$\hat{B}_{j\text{SIMPLE}} = H_{1j} \quad (5)$$

$$\hat{G}_{j\text{SIMPLE}} = \frac{(H_{Ij} - H_{1j})}{(\text{Age}_{Ij} - \text{Age}_{1j})} \quad (6)$$

If length/height is measured with error, then H_{1j} is a composite of an individual's 'true' birth length (B_j) and measurement error e_{1j} , i.e. $H_{1j} = B_j + e_{1j}$. Therefore, unless there is no measurement error, the variance of $\hat{B}_{j\text{SIMPLE}}$ will be greater than the true variance of B_j , and the relationship between BP and estimated birth length will be biased towards the null (see Appendix 1). The variance of $\hat{G}_{j\text{SIMPLE}}$ will also be greater than the true variance of G_j . The correlation between $\hat{B}_{j\text{SIMPLE}}$ and $\hat{G}_{j\text{SIMPLE}}$ will be estimated as more negative than the true correlation between birth length and growth rate (i.e. if the true covariance is positive, then this method will underestimate the covariance between birth length and growth rate), due to mathematical coupling.^{17,18} However, the covariance between each of $\hat{B}_{j\text{SIMPLE}}$ and $\hat{G}_{j\text{SIMPLE}}$ and BP will be unbiased. Thus, whether the association between growth rate and BP (conditional on birth length) is overestimated or underestimated will depend on the relative sizes of the variances and covariance of birth length and growth rate.

2.3.2 The individual regression (OLS) approach

An alternative approach is to fit J separate linear regression models (one for each individual) of height for individual j at time i (H_{ij}) on age of that individual at each time-point (Age_{ij}), via ordinary least squares (OLS).

$$\exists j \in (1, J): H_{ij} = \beta_{0j} + \beta_{1j}\text{Age}_{ij} + \varepsilon_j \quad \text{and} \quad \varepsilon_j \sim N(0, \sigma_{\varepsilon_3}^2) \quad (7)$$

The intercept and slope parameters of each individual regression are used to estimate birth length and growth rate for that individual

$$\begin{aligned}\hat{B}_{jOLS} &= \beta_{0j}, \\ \hat{G}_{jOLS} &= \beta_{1j}\end{aligned}\tag{8}$$

\hat{B}_{jOLS} and \hat{G}_{jOLS} are then used as estimates of birth length and growth rate in the second-stage models (3) and (4). Whilst the OLS method of regression will give unbiased estimates of B_j and G_j for individuals with fully observed measurements, the method fails to take into account the dependence between measurements, therefore violating the assumption of independence of ε_{ij} , and underestimating the residual variance. As each individual trajectory is considered independently from the rest of the population, the between individual variance is overestimated, i.e. the variances of \hat{B}_{jOLS} and \hat{G}_{jOLS} will be greater than the true variances of B_j and G_j . The sample covariance between estimated birth length (\hat{B}_{jOLS}) and growth rate (\hat{G}_{jOLS}) will be more negative than the true covariance (i.e. if the true covariance is positive, then this method will underestimate the covariance between birth length and growth rate). However, the covariance of the estimated birth length (\hat{B}_{jOLS}) and growth rate (\hat{G}_{jOLS}) with BP will be unbiased. Thus, as with the simple method, the association between birth length and BP will be biased towards the null, whereas that between growth rate and BP (conditional on birth length) could be biased in either direction.

2.3.3 MLM approach

A further, more parsimonious, method would be to specify a MLM with random intercepts and slopes. This method simultaneously models the growth trajectories of all individuals, assuming that birth length and growth rate are normally distributed around the population mean birth length (β_{0MLM}) and growth rate (β_{1MLM}) with standard deviations σ_{u0} , σ_{u1} and covariance σ_{u01} .

$$\begin{aligned}H_{ij} &= \beta_{0jMLM} + \beta_{1jMLM} \text{Age}_{ij} + \varepsilon_{ij} \\ \beta_{0jMLM} &= \beta_{0MLM} + u_{0j} \\ \beta_{1jMLM} &= \beta_{1MLM} + u_{1j} \\ \begin{bmatrix} u_{0j} \\ u_{1j} \end{bmatrix} &\sim N(0, \Omega_u): \Omega_u = \begin{bmatrix} \sigma_{u0}^2 & \\ \sigma_{u01} & \sigma_{u1}^2 \end{bmatrix} \\ \varepsilon_{ij} &\sim N(0, \sigma_{\varepsilon 4}^2)\end{aligned}\tag{9}$$

The individual birth lengths and growth rates can be estimated by adding the level 2 shrunken residuals \hat{u}_{0j} and \hat{u}_{1j} (random effects estimates) to the estimates of the population mean birth length and growth rate, $\hat{\beta}_{0MLM}$ and $\hat{\beta}_{1MLM}$, to give the estimated birth length (\hat{B}_{jMLM}) and growth rate (\hat{G}_{jMLM})

$$\begin{aligned}\hat{B}_{jMLM} &= \hat{\beta}_{0MLM} + \hat{u}_{0j} \\ \hat{G}_{jMLM} &= \hat{\beta}_{1MLM} + \hat{u}_{1j}.\end{aligned}\tag{10}$$

The individual level residuals \hat{u}_{0j} and \hat{u}_{1j} from the MLM above are the best linear unbiased predictors (BLUPS) for the individual random effects. They are, however, shrunk, in that they are a weighted average of the individual and the population average intercept and slope, with the weighting depending on the timing of measures, and the between- and within-individual variation (see Appendix 1). Thus, the sample distribution of the variances of the residuals will be smaller than the true variances of birth length and growth rate. The sample covariance between estimated birth length ($\hat{B}_{j\text{MLM}}$) and growth rate ($\hat{G}_{j\text{MLM}}$) will tend to be more positive than the true covariance (i.e. if the true covariance is positive, this method will overestimate it). In addition, the covariances between BP and birth length ($\hat{B}_{j\text{MLM}}$) and growth rate ($\hat{G}_{j\text{MLM}}$) will no longer be unbiased. However, we show that when both estimates ($\hat{B}_{j\text{MLM}}$ and $\hat{G}_{j\text{MLM}}$) are included in the linear regression model for the outcome, the estimated associations between outcome and both intercept and slope are unbiased (see Appendix 1).

2.3.4 MLM re-inflation approach

In order to mitigate the potential problem of under-estimation of the variances, shrunken residuals, \hat{u}_{0j} and \hat{u}_{1j} , can be transformed (re-inflated) so that their sample variance and covariances reflect the model based estimates of Ω_u . Re-inflated residuals are then added to the estimates of the population mean birth length and growth rate, $\hat{\beta}_{0\text{MLM}}$ and $\hat{\beta}_{1\text{MLM}}$, to give the estimated birth length ($\hat{B}_{j\text{MLMINFL}}$) and growth rate ($\hat{G}_{j\text{MLMINFL}}$) for each individual j . The process of re-inflation requires multiplying the estimated shrunken residuals by an upper triangular matrix of equal order to the number of random coefficients. A brief description of the process and an example of the code to perform the transformation is given in Appendix 1, and a more detailed description can be found in the original manuscript by Carpenter et al.¹⁹ The variances and covariances of estimated birth length ($\hat{B}_{j\text{MLMINFL}}$) and growth rate ($\hat{G}_{j\text{MLMINFL}}$) will be unbiased compared to the model based variances and covariances.

2.4 Joint modelling methods

2.4.1 Bivariate MLM approach

In the bivariate model, the growth trajectory and adult BP models are estimated simultaneously.^{14,20} This model can be thought of as a combination of a two-level model for height, and a single level model for BP. The specification of the growth model is identical to the MLM presented previously. Additionally, a single level variance component model is estimated for BP and replaces the need to separately estimate equations (3) and (4).

$$\begin{aligned}
 H_{ij} &= \beta_{0\text{BVM}} + \beta_{1\text{BVM}} \text{Age}_{ij} + \varepsilon_{ij} \\
 \beta_{0\text{BVM}} &= \beta_{0\text{BVM}} + u_{0j} \\
 \beta_{1\text{BVM}} &= \beta_{1\text{BVM}} + u_{1j} \\
 \text{BP}_j &= \beta_{2\text{BVM}} \\
 \beta_{2\text{BVM}} &= \beta_{2\text{BVM}} + u_{2j}
 \end{aligned} \tag{11}$$

The single level BP model estimates the population mean BP, and the residuals \hat{u}_{2j} estimate the individual deviation from that mean BP. As both models are estimated simultaneously, the variances and covariances of the residuals \hat{u}_{0j} (birth length), \hat{u}_{1j} (growth rate) and \hat{u}_{2j} (BP) are jointly estimated, and Ω_u is now a 3×3 symmetric covariance matrix.

$$\begin{bmatrix} u_{0j} \\ u_{1j} \\ u_{2j} \end{bmatrix} \sim N(0, \Omega_u): \Omega_u = \begin{bmatrix} \sigma_{u0}^2 & & \\ \sigma_{u01} & \sigma_{u1}^2 & \\ \sigma_{u0BP} & \sigma_{u1BP} & \sigma_{BP}^2 \end{bmatrix} \quad (12)$$

$$\begin{bmatrix} \varepsilon_{ij} \end{bmatrix} \sim N(0, \sigma_{\varepsilon}^2)$$

Using a similar method to that of the univariate MLM, the shrunken residuals \hat{u}_{0j} , \hat{u}_{1j} and \hat{u}_{2j} can be simultaneously re-inflated so that their variances and covariances reflect the true variance–covariance matrix. These inflated residuals can then be added to the population estimates of birth length ($\hat{\beta}_{0BVM}$), growth rate ($\hat{\beta}_{1BVM}$) and BP ($\hat{\beta}_{2BVM}$) to give estimates of birth length (\hat{B}_{jBVM}), growth rate (\hat{G}_{jBVM}) and BP (\hat{BP}_{jBVM}) for each individual.

In a similar way to the two-stage approaches, but here using estimated instead of measured BP, these three estimates are then substituted into equations (3) and (4) and used to estimate the parameters of interest α_1 and α_4 . Whilst deriving the parameters occurs in two stages, the growth and BP model is a single joint model. Thus, the variances and covariances of the estimated birth length, growth rate and BP will be unbiased estimates of the true variances and covariances, and the estimates of the parameters of interest α_1 and α_4 will be unbiased. However, this approach will not take into account the uncertainty in estimating the birth length, growth rate and BP from the bivariate model.

There are alternative methods which could be used to directly estimate α_1 and α_4 from the bivariate MLM. One method would be to manipulate the variances and covariances of Ω_u , using a moment-based approach.²¹ In simple cases, this is fairly straightforward; however, calculating the corresponding standard errors is more complex. Simulation and moment-based methods have been developed,²² but these methods require normality assumptions to be made, and may only be appropriate in large samples.

In order to take into account the uncertainty in the residuals without relying on normality assumptions, a non-parametric bootstrap with replacement, with 1000 replicates, was used. The standard deviation of the 1000 bootstrap replicates was used as an estimate of the standard error. Normal theory confidence intervals were constructed using the observed point estimates and bootstrap standard errors, and percentile based confidence intervals were also calculated. This bootstrap estimation was only carried out for the baseline experimental scenario (see further in the text).

2.4.2 Structural equation model

Using a structural equation modelling (SEM) framework, the bivariate outcome model can be reformulated, and framed in terms of measurement and structural models.

The measurement model (13) describes a model of linear dependence of the height H_{ij} of an individual j at a given Age_{ij} on two latent factors for birth length (B_{jSEM}) and growth rate (G_{jSEM}). The relationship is described by loadings λ_{0i} and λ_{1i} (where $\lambda_{0i} = 1$ and $\lambda_{1i} = Age_{ij}$).

$$H_{ij} = \lambda_{0i} B_{jSEM} + \lambda_{1i} G_{jSEM} + \varepsilon_{1ij} \quad (13)$$

The structural model (14) directly relates the individual's BP to the latent factors representing birth length and growth rate.

$$BP_j = \alpha_2 + \alpha_3 B_{jSEM} + \alpha_4 G_{jSEM} + \varepsilon_{2j} \quad (14)$$

Thus, the structural model has the same form as the model for BP in the two-stage approaches, (4), but with latent variables for birth length and growth replacing their sample estimates. One of the parameters of interest, the association between growth rate and BP conditional on birth length (α_4), is estimated directly by this procedure. The other parameter of interest, the unconditional relationship between BP and birth length (α_1) can be estimated by re-fitting the SEM but changing the structural model such that only birth length and BP are correlated

$$BP_j = \alpha_0 + \alpha_1 B_{jSEM} + \varepsilon_{2j} \quad (15)$$

Alternatively, the unconditional relationship and standard error between BP and birth length (α_1) can be estimated from the original SEM, where r_{u01} is the estimated correlation between the birth length and growth rate

$$\alpha_1 = [\alpha_3 + \alpha_4(r_{u01}(\sigma_{u1}/\sigma_{u0}))] \quad (16)$$

For fixed measurement occasions ($Age_{ij}=Age_i$) and constraints on the loadings ($\lambda_{0i}=1$ and $\lambda_{1i}=Age_i$), this structural equation model is equivalent to the bivariate growth model described earlier. However, the structural equation model estimates the parameters of interest (and, importantly, their standard errors) directly, rather than requiring them to be derived from the model variance/covariance matrix.

2.5 Experimental scenarios

Five different experimental scenarios were chosen to explore the performance of different methods (in terms of bias and coverage) under different assumptions about the magnitude of variation in birth length, growth rate, correlation between birth length and growth rate, measurement error in the growth model and measurement error in the BP model. The number of replications for each of the scenarios outlined below was 1000.

Assuming a non-factorial design, standard deviations were set at $\sigma_{u0}=2.5$, $\sigma_{u1}=0.5$, $\sigma_{eh}=2.0$ and $\sigma_{eBP}=10$, and correlation $\rho_{u01}=0.1$. All other residual correlations were set equal to 0.

- (1) σ_{u0} (standard deviation of birth size, cm): 1.5, 2.0, (2.5), 3.0, 3.5.
- (2) σ_{u1} (standard deviation of growth rate, $\text{cm}\cdot\text{y}^{-1}$): 0.2, 0.3, 0.4, (0.5), 0.6, 0.7, 0.8, 0.9, 1.0.
- (3) σ_{eh} (residual standard deviation in growth model, cm): 0.1, 1, (2), 3, 4, 5.
- (4) ρ_{u01} (correlation of birth size and growth rate): -0.6 , -0.4 , -0.2 , -0.1 , 0, (0.1), 0.2, 0.4, 0.6.
- (5) σ_{eb} (residual standard deviation in BP, mmHg): 8, 9, (10), 11, 12.

Values in parentheses are simulation defaults held constant in the other experimental scenarios, and the baseline experimental scenario held all parameters at these values.

Model parameters were fixed for all simulations at:

- (1) β_0 (Birth Length) = 50
- (2) β_1 (Growth Rate) = 9
- (3) α_2 (Mean BP) = 120
- (4) α_3 (Birth Length BP conditional association) = 0.5
- (5) α_4 (Growth Rate BP conditional association) = 2.0.

2.6 Summary statistics of interest

The statistics of primary interest are the estimated association between birth length and BP ($\hat{\alpha}_1$) and the estimated association between growth rate and BP conditional on birth length ($\hat{\alpha}_4$).

The expected unconditional association between birth length and BP is given by:

$\alpha_1 = [\alpha_3 + \alpha_4(\rho_{u01}(\sigma_{u1}/\sigma_{u0}))]$. For k th parameter ($k=1, 4$), we investigate the relative bias of estimated coefficients $((\hat{\alpha}_k - \alpha_k)/\alpha_k) \cdot 100$. We also estimate coverage, where if $|\hat{\alpha}_k - \alpha_k| \leq 1.96 \cdot SE(\hat{\alpha}_k)$ the estimate was covered, and if $|\hat{\alpha}_k - \alpha_k| > 1.96 \cdot SE(\hat{\alpha}_k)$ it was not.

2.7 Simulation implementation

The data were generated using Stata 12.1. The first stage of the simple and individualised regression approaches were also conducted in Stata 12.1,²³ whereas the multilevel and bivariate MLMs were fitted in MLWIN 2.25²⁴ via the runmlwin²⁵ Stata command. The second stage of all two-stage methods was conducted in Stata 12.1. The structural equation model was fitted in Mplus²⁶ via R using the MplusAutomation package.²⁷ Full details of the syntax used to fit the models are listed in Appendix 1.

3 Results

Results for scenarios 1 (varying σ_{u0}) and 2 (varying σ_{u1}) for the simple and individualised growth trajectories, MLMs and joint models of growth and disease are presented in Figures 2–4,

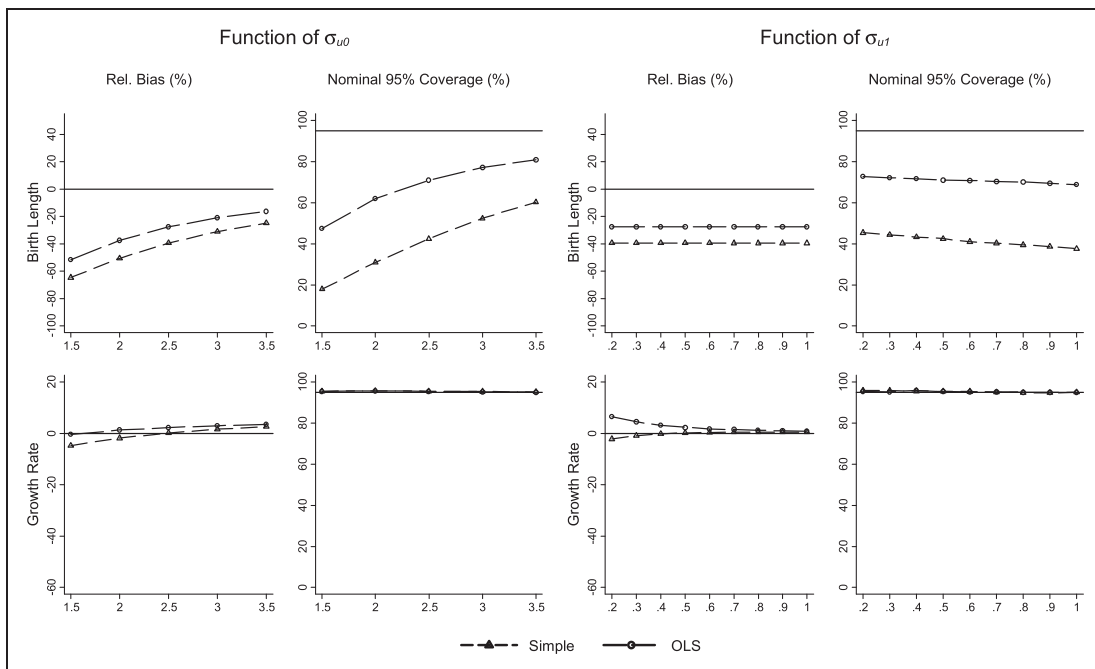


Figure 2. Performance of the simple and OLS methods displayed as relative bias, and nominal 95% coverage plotted as functions of σ_{u0} (standard deviation of birth length) and σ_{u1} (standard deviation of growth rate) for the effect of birth length on BP, and the effect of growth rate on BP conditional on birth length.

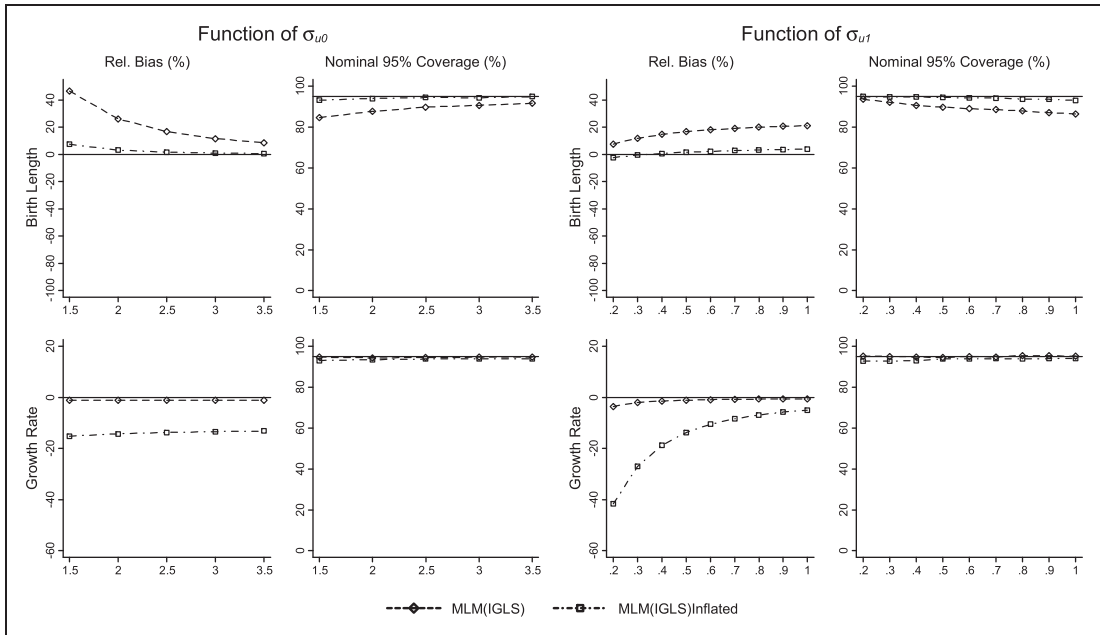


Figure 3. Performance of the multilevel model (MLM(IGLS)) and MLM with re-inflated residuals (MLM(IGLS) Inflated) methods displayed as relative bias, and nominal 95% coverage plotted as functions of σ_{u0} (standard deviation of birth length) and σ_{u1} (standard deviation of growth rate) for the effect of birth length on BP, and the effect of growth rate on BP conditional on birth length.

respectively. Each figure illustrates relative bias and nominal coverage plotted as a function of σ_{u0} and σ_{u1} . The upper panels present bias and coverage for the association of birth length and BP, and the lower panels present bias and coverage for the association between growth rate and BP conditional on birth length.

If a model were to perfectly recover the estimates of interest, we would expect 0% relative bias which did not fluctuate with changes in the experimental values. Similarly, we would also expect nominal coverage to be approximately 95% across the range of experimental values used.

3.1 Two-stage simple and OLS methods

For the association between birth length and BP, both the simple and OLS methods are biased towards the null, and the simple method demonstrates more bias than the OLS method. The simple method has 18% nominal coverage rising to 60% as σ_{u0} increases, and conversely falls from 45% to 37% as σ_{u1} increases. Coverage using the OLS method is higher, rising steadily from 48% to 80% as σ_{u0} increases; conversely, it falls from 72% to 67% as σ_{u1} increases. The rise and fall in coverage is primarily a function of changing bias, as opposed to changes in confidence interval width.

The conditional growth rate/BP association shows less relative bias than the birth length/BP association for both methods. The simple method is modestly biased towards the null, with the bias reducing as σ_{u0} and σ_{u1} increase. The OLS method is biased away from the null with the size of the bias increasing with σ_{u0} and decreasing with σ_{u1} . Coverage is greatly improved in comparison to the birth length/BP association for both the simple and OLS methods. Coverage of the simple and

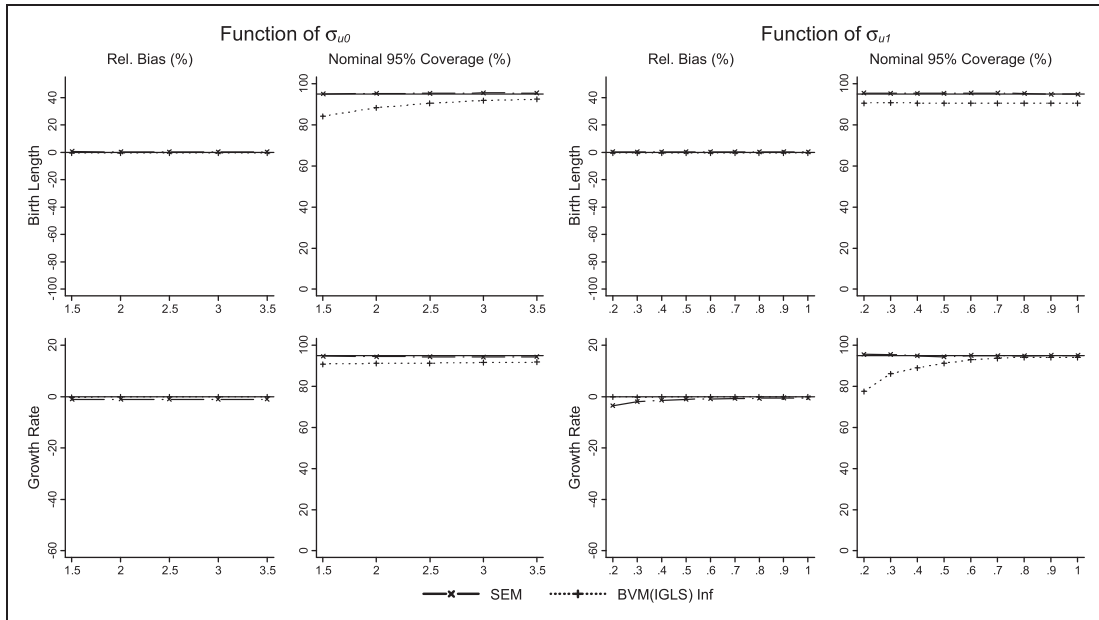


Figure 4. Performance of the structural equation model (SEM) and bivariate growth model with re-inflated residuals (BVM (IGLS) Inf) methods displayed as relative bias, and nominal 95% coverage plotted as functions of σ_{u0} (standard deviation of birth length) and σ_{u1} (standard deviation of growth rate) for the effect of birth length on BP, and the effect of growth rate on BP conditional on birth length.

OLS method is approximately at expected levels in both scenarios, with modest improvements from 93% to 97% with increasing σ_{u0} .

A similar pattern of results (OLS more favourable than simple methods) can be seen for scenarios 3 (varying σ_{eh}), 4 (varying ρ_{u01}) and 5 (varying σ_{eBP}) (see Supplementary material, available at <http://smm.sagepub.com/>). For example, scenario 3 illustrates that simple estimates for birth length are consistently biased towards the null in comparison to OLS estimates. Both methods show little bias when σ_{eh} is small; however, this increases as σ_{eh} becomes larger. Estimates of birth length/BP association are unaffected by changes in ρ_{u01} (scenario 4) and σ_{eBP} (scenario 5), they are both consistently biased towards the null, the simple method more so than the OLS. Estimates of growth rate/BP associations (conditional on birth length) are biased towards the null when birth length and growth rate are negatively correlated, with that bias decreasing to zero and then increasing bias away from the null as birth length and growth rate are increasingly positively correlated. See Supplementary Figure 1 (see Supplementary material, available at <http://smm.sagepub.com/>).

3.2 Two-stage MLMs

Whilst the same MLM is estimated for both two-stage analyses, the difference in results between using shrunken or re-inflated residuals is substantial. For the association between birth length and BP, the use of shrunken residuals results in bias away from the null, with increasing bias as σ_{u0} decreases or as σ_{u1} increases. Whilst the coverage is modestly below expected levels, it increases from

85% to 93% as σ_{u0} increases, and falls from 94% to 87% as σ_{u1} increases. Re-inflating the residuals results in an attenuation of bias, and nominal coverage is at expected levels (95%) across the range of either σ_{u0} or σ_{u1} .

The converse is demonstrated for the effect of growth rate. As expected, shrunken residuals are unbiased for changes in σ_{u0} or σ_{u1} , and coverage is at expected levels (95%) across the ranges investigated. However, the re-inflation process caused an increase in bias towards the null (approximately 15%) across the range of σ_{u0} investigated. The effect is more obvious in response to changes in σ_{u1} , where 40% bias is observed when σ_{u1} is small, which slowly attenuates to 5% as σ_{u1} increases. Coverage is only modestly below expected levels (93%) in both scenarios.

A similar pattern of results (inflation unbiased for birth length, shrunken residuals unbiased for growth rate) is seen for scenarios 3 (σ_{eh}), 4 (ρ_{u01}) and 5 (σ_{eBP}). The exception is for scenario 4, where inflated residuals yield biased results with respect to the birth length BP association when ρ_{u01} is negative, see Supplementary Figure 2.

3.3 Joint models of growth and disease

Estimates from the bivariate MLM are unbiased (less than 1% relative bias) with respect to estimates of the birth length–BP association and the growth rate–BP association, as expected. The SEM method demonstrates unbiased results for the BP birth length association, and a small but persistent bias of 1% towards the null for the growth rate BP association in scenario 1, and a larger bias (3.5%) towards the null when σ_{u1} is small which attenuates to less than 1% as σ_{u1} increases. Nominal coverage of the SEM approach is at expected levels across all experimental scenarios, whereas nominal coverage using the two-stage approach of the bivariate growth model results in coverage slightly below expected levels. In the bivariate MLM, coverage for the association between birth length and BP ranges from 84% to 92% with increasing σ_{u0} , whereas it is approximately constant at 90% across all values of σ_{u1} . The pattern is nearly reversed with regard to the conditional association between growth rate and BP: changes in σ_{u0} have little effect on nominal coverage of the bivariate MLM, which is approximately 91% across the range, whereas coverage is lower (78%) when σ_{u1} is small which then steadily increases to near expected levels (94%) when σ_{u1} becomes larger. Under-coverage is corrected, in the baseline scenario, using a non-parametric bootstrap. This approach resulted in nominal coverage using normal approximation confidence intervals (birth length 95.5%, growth rate 94.6%) or percentile confidence intervals (birth length 95.7%, growth rate 94.6%).

A similar pattern of results is seen for scenarios 3 (σ_{eh}), 4 (ρ_{u01}) and 5 (σ_{BP}). SEM methods show a small bias (1%) towards the null and nominal coverage is at expected levels for all scenarios. The bivariate growth model two-stage method is unbiased, but nominal coverage falls with increasing σ_{eh} , and also falls with increasing ρ_{u01} correlation, see Supplementary Figure 3.

4 Discussion

We have shown algebraically that the two-stage process using a MLM to estimate growth parameters and then relating these to the distal outcome in a second stage will give unbiased estimates of the conditional associations between both growth parameters and outcome. Our simulations confirmed this, and also showed that using the same process to estimate the unconditional association between birth length and outcome leads to bias. We have shown that the two-stage bivariate MLM (with re-inflation) is unbiased in all the scenarios investigated, although with under-coverage of confidence intervals. We have also demonstrated that SEM

produces a small bias towards the null in the estimation of the association between growth rate and BP (due to the different distribution of individual ages at birth compared to the other time-points). The simple and OLS two-stage methods result in biased estimates of the association between BP and birth length, and less biased estimates of the association between BP and growth rate, conditional on birth length.

The simple two-stage method illustrated substantial bias in the presence of measurement error, and underperformed in comparison to the OLS method with regards to estimates of the effect of birth length on BP. The OLS method demonstrated considerable bias with regards to the estimates of birth length on BP and nominal coverage was also poor. The MLM two-stage approach with inflated residuals demonstrated the least bias with regards to the association of birth length with BP, and considerably outperformed the use of shrunken residuals, which was nearly as biased (although in the opposite direction) as the OLS method. Therefore, the two-stage method of choice when investigating the effects of birth length on BP would be the MLM with inflated residuals. However, the process of re-inflating residuals is not unique, and the transformation does not necessarily preserve the relationship between the empirical residual and the outcome of interest. Using the lower triangular matrices of the Cholesky decomposition during re-inflation results in inflated residuals for birth length which are a simple linear transformation of the uninflated residuals. However, using the upper triangular Cholesky decomposition results in inflated residuals for birth length which are a linear combination of the uninflated residuals for growth rate and those for birth length and so does not result in unbiased associations.

In terms of the association of BP with growth rate (conditional on birth length), the most biased two-stage method was the MLM with inflated residuals, which biases results towards the null. The bias occurs because the association between growth rate and BP is distorted during the inflation process, which is a complex transformation that depends on the shrunken residuals of both birth length and growth rate. The simple method also illustrates biases towards the null, but to a lesser extent than those with inflated residuals. The OLS method illustrates biases worse than the simple method under some circumstances, despite the intuitive incorporation of all relevant data. However, the two-stage method using the MLM with shrunken residuals led to unbiased results, since the consistent shrinkage of both birth length and growth rate residuals preserve the association between BP and growth rate conditional on birth length. Nominal coverage is preserved at expected levels by the inflated standard errors due to the reduced residual variance.

The bivariate growth model, which simultaneously generates growth and BP residuals, which are in turn re-inflated, leads to unbiased results in all scenarios. However, this method results in under-coverage of 5% in scenarios 1 and 2. The unbiased result and less than optimal coverage needs to be balanced against the minor bias yielded by SEM method (in this example, due to unbalanced data), and the full propagation of uncertainty and correct 95% nominal coverage, or the computationally intensive nature of the non-parametric bootstrap, which fully incorporates the uncertainty from the growth model.

Approaches for tackling this problem have been suggested, in the context of the measurement error. It has been noted that using parameters from a linear MLM (the 'RC' method) results in unbiased conditional effect estimates when the model relating exposures to outcome is linear.¹³ The same paper noted that the individual regression method resulted in biased estimates of the conditional effects. However, for non-linear models for the outcome (e.g. a logistic model for a binary distal outcome), this method remains biased (although with reduced bias compared to the individual regression method), and alternatives have been proposed.²⁸ Given the difficulties in estimating joint models with a binary outcome, more research is needed into the size of the bias

when using the two-stage multilevel approach in the non-linear case, and the ease of implementation of alternatives.

5 Future work

Whilst this simulation study highlights how variation in specific parts of the data-generating process affect the estimation of the effect of either birth length or growth rate on BP, we have not explored a full factorial experimental design, therefore combinations of unfavourable scenarios may result in unacceptable bias and poor coverage. Additionally, we have only explored these effects when the number of observations is the same in each individual, and not explored the consequences of imbalance and missing data for the simple and OLS methods. We briefly explored the consequences of equalising the size of the effects between growth parameters and BP, and found similar associations, i.e. biases in birth length and BP association were greater than biases in the growth rate and BP associations (results available from authors upon request). Furthermore, we did not vary the direction and magnitude of the association between growth parameters and BP and therefore we are unable to explore the potential reversal paradox described by Tu and colleagues.¹⁶ Similarly, we did not investigate the effect of population size or frequency of measurement in relation to the observed biases, and the stability of the six methods with small numbers of individuals and/or infrequent measurements may be different to that seen here.

We did not examine the effect of violating model assumptions. In particular, whilst growth (and change in other anthropometric variables) is often linear over short periods of time, non-linear models will generally be required for examining change over longer periods. More complex data-generating processes could also be considered, for example by creating an interaction between birth length and growth rate in their association with the outcome. SEM or path analyses could be used to examine the mediation of the association between birth length and BP by growth rate.

6 Conclusions

The joint modelling approach which takes into account and incorporates the variation of growth process into the estimation of effects on subsequent outcomes is clearly the preferred method, giving unbiased estimates of both the conditional and unconditional associations of birth length and growth with BP. Given the requirement for specialist software and the greater technical difficulty in fitting the joint models, this option may not be viable for all researchers without specialist training. An alternative would be to use the two-stage MLM approach to estimate conditional associations, where an experienced analyst can derive the residuals, and less-experienced researchers can use them as exposures in standard regression models. Where measurement error in the repeated outcome is low, this approach may result in little bias even for unconditional associations.

This simulation study could change the interpretation of previously reported null findings, as many of the methods commonly used result in biases towards the null and poor nominal coverage. Thus, reanalysis with more suitable methods may reduce both the type I error rate and the heterogeneity in the current literature.

Acknowledgement

We would like to thank Professor Harvey Goldstein for helpful discussions of this topic.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: AS, ADACS, JH, MSG, FS and KT were supported by the Medical Research Council [grant number G1000726]. FS was also supported by the ESRC [grant number RES-576-25-0032]. CM-W is supported by a Fellowship from the Medical Research Council (MR/J011932). ADACS, CM-W and KT work in a Unit that receives funding from the UK Medical Research Council and the University of Bristol (MC_UU_12013/5).

References

- O'Brien MF, Cronin AM, Fearn PA, et al. Evaluation of prediagnostic prostate-specific antigen dynamics as predictors of death from prostate cancer in patients treated conservatively. *Int J Cancer* 2011; **128**: 2373–2381.
- Elias MF, Davey A, Dore GA, et al. Deterioration in renal function is associated with increased arterial stiffness. *Am J Hypertens* 2014; **27**(2): 207–214.
- Clouston SA, Brewster P, Kuh D, et al. The dynamic relationship between physical function and cognition in longitudinal aging cohorts. *Epidemiol Rev* 2013; **35**(1): 33–50.
- Barker DJ, Osmond C, Forsen TJ, et al. Trajectories of growth among children who have coronary events as adults. *New Engl J Med* 2005; **353**: 1802–1809.
- Forsen T, Osmond C, Eriksson JG, et al. Growth of girls who later develop coronary heart disease. *Heart* 2004; **90**: 20–24.
- Forsen TJ, Eriksson JG, Osmond C, et al. The infant growth of boys who later develop coronary heart disease. *Ann Med* 2004; **36**: 389–392.
- Huxley R, Neil A and Collins R. Unravelling the fetal origins hypothesis: is there really an inverse association between birthweight and subsequent blood pressure? *Lancet* 2002; **360**: 659–665.
- Lucas A, Fewtrell MS and Cole TJ. Fetal origins of adult disease—the hypothesis revisited. *BMJ* 1999; **319**: 245–249.
- Belfort MB, Rifas-Shiman SL, Rich-Edwards J, et al. Size at birth, infant growth, and blood pressure at three years of age. *J Pediatr* 2007; **151**: 670–674.
- Kark M, Tynelius P and Rasmussen F. Associations between birthweight and weight change during infancy and later childhood, and systolic blood pressure at age 15 years: the COMPASS study. *Paediatr Perinatal Epidemiol* 2009; **23**: 245–253.
- Howe LD, Tilling K, Matijasevich A, et al. Linear spline multilevel models for summarising childhood growth trajectories: A guide to their application using examples from five birth cohorts. *Stat Meth Med Res* 2013.
- Tapp RJ, Ness A, Williams C, et al. Differential effects of adiposity and childhood growth trajectories on retinal microvascular architecture. *Microcirculation* 2013; **20**: 609–616.
- Wang CY, Wang N and Wang S. Regression analysis when covariates are regression parameters of a random effects model for observed longitudinal measurements. *Biometrics* 2000; **56**: 487–495.
- Steele F, Vignoles A and Jenkins A. The impact of school resources on pupil attainment: a multilevel simultaneous equation modelling approach. *J Royal Stat Soc A* 2005; **170**: 801–824.
- Llabre MM, Spitzer SB, Saab PG, et al. Piecewise latent growth curve modeling of systolic blood pressure reactivity and recovery from the cold pressor test. *Psychophysiology* 2001; **38**: 951–960.
- Tu YK, West R, Ellison GT, et al. Why evidence for the fetal origins of adult disease might be a statistical artifact: the “reversal paradox” for the relation between birth weight and blood pressure in later life. *Am J Epidemiol* 2005; **161**: 27–32.
- Blance A, Tu YK and Gilthorpe MS. A multilevel modelling solution to mathematical coupling. *Stat Meth Med Res* 2005; **14**: 553–565.
- Rasbash J and Goldstein H. Mathematical coupling: a simpler approach. *Int J Epidemiol* 2004; **33**: 1401–1402; discussion 2–3.
- Carpenter JR, Goldstein H and Rasbash J. A novel bootstrap procedure for assessing the relationship between class size and achievement. *J Roy Stat Soc C* 2003; **52**: 431–443.
- Goldstein H and Kounali D. Multilevel multivariate modelling of childhood growth, numbers of growth measurements and adult characteristics. *J Royal Stat Soc A* 2009; **172**: 599–613.
- Fisher R. *Chapter 5: Tests of significance of means, differences of means, and regression coefficients*. Edinburgh: Oliver and Boyd, 1925.
- Palmer T and Macdonald-Wallis C. REFFADJUST: Stata module to estimate adjusted regression coefficients for the association between two random effects variables, <http://ideas.repec.org/c/boc/bocode/s457403.html> (2012).
- StataCorp. *Stata statistical software: Release 11*. College Station, TX: StataCorp LP, 2009.
- Rasbash J, Charlton C, Browne WJ, et al. *MLwiN version 2.24*. Bristol: Centre for Multilevel Modelling, University of Bristol, Bristol, 2009.
- Leckie G and Charlton C. runmlwin – A program to run the MLwiN multilevel modelling software from within stata. *Journal of Statistical Software* 2013; **52**(11): 1–40.
- Muthén BO and Muthén LK. MPlus V6, www.statmodel.com.
- Hallquist M. *Automating Mplus model estimation and interpretation*. Vienna: R Foundation for Statistical Computing, 2012.
- De la Cruz R, Marshall G and Quintana FA. Logistic regression when covariates are random effects from a non-linear mixed model. *Biometrical J* 2011; **53**: 735–749.