OPEN ACCESS

University of
BRISTOL

Leckie, G., & Goldstein, H. (2015). A multilevel modelling approach to measuring changing patterns of ethnic composition and segregation among London secondary schools, 2001-2010. Journal of the Royal Statistical Society Series A, 178(2), 405-424. DOI: 10.1111/rssa.12066

Peer reviewed version

Link to published version (if available):
10.1111/rssa.12066

Link to publication record in Explore Bristol Research
PDF-document

University of Bristol - Explore Bristol Research
General rights

A Multilevel Modelling Approach to Measuring Changing Patterns of Ethnic Composition and Segregation among London Secondary Schools, 2001-2010

George Leckie and Harvey Goldstein

University of Bristol, UK

Summary.

Multilevel binomial logistic regression has recently been proposed for the special case of statistically modelling changing composition and segregation of two groups of individuals over two occasions among organisational units, enabling inferences to be made about the underlying social processes which generate these patterns. A simulation method can then be used to reexpress the model parameters in the metric of any desired two-group segregation index. In the current paper, we generalise this combined modelling and simulation approach by proposing multilevel random-coefficient multinomial logistic regression for the general case of statistically modelling multiple groups of individuals over multiple occasions and multiple organisational scales. We illustrate this combined approach with an application to modelling changing three-group white-black-Asian ethnic composition and segregation among London secondary schools and local authorities during the first decade of the 21st century.

Keywords: Ethnic composition; Ethnic segregation; Multilevel models; Multinomial logistic regression; Segregation indices; Variance functions

Address for correspondence: George Leckie, Centre for Multilevel Modelling, Graduate School of Education, University of Bristol, 2 Priory Road, Bristol, BS8 1TX, UK.

E-mail: g.leckie@bristol.ac.uk

1. Introduction

Goldstein and Noden (2003) propose multilevel binomial logistic regression for the special case of statistically modelling the composition and segregation of two groups of individuals (e.g., white and black students) over two occasions among organisational units (e.g., schools or neighbourhoods). This approach accounts for the randomness in the observed proportions that lead standard segregation indices to be biased and noisy measures of true segregation, especially when units are small. Their approach also allows statistical tests of differences in composition and segregation across areas, or over time, and can adjust and explain these differences for characteristics of the individuals or organisational units. Leckie et al. (2012) further develop this work by presenting a simulation method which can be used to reexpress the model segregation parameters in the metric of any desired two-group segregation index.

The above work was concerned with two-group segregation, but researchers are often interested in studying segregation between three or more groups of individuals, for example measuring residential segregation (Farrell, 2008; Iceland, 2004) and school segregation (Reardon et al., 2000) between white, black and Asian individuals. Interest lies not only in calculating the overall degree of segregation, but also pairwise segregation between each possible pair of groups. Analysing three or more groups allows for richer descriptions of segregation and may reveal important differences between subgroups. For example, black-white segregation might be found to be increasing over time while Asian-white segregation is simultaneously decreasing. Similarly, the relative importance of organisational scale or spatial aggregation for how segregation operates may also differ by ethnic group. For example, black-white segregation might operate largely at the school-district-level while Asian-white

segregation might operate largely at the school-level. As with two-group segregation, many competing multigroup segregation indices have been proposed with considerable and on-going debates as to their ideal properties (see Reardon and Firebaugh, 2002, for a summary).

The aim of the current paper is to generalise the two-group, two-occasion approach described above to the general case of statistically modelling segregation among multiple groups over multiple occasions and multiple organisational scales. We propose multilevel random-coefficient multinomial logistic regression for this purpose. Specifically, we include random time trends to enable hypothesis testing for whether particular pairs of groups are becoming significantly more segregated over time. We show how including multiple levels of random-effects allows us to simultaneously analyse multigroup segregation at different organisational scales or spatial aggregations (micro-, meso-, and macro-segregation). We discuss how covariates can then be introduced into the model to produce adjusted measures of segregation and to test hypotheses about why levels of segregation change over time. We show how our simulation method can be extended to present estimated levels and trends in multigroup segregation along with their uncertainty in the metric of any desired multigroup segregation index. We illustrate these developments with an application to modelling changing patterns of three-group white-black-Asian ethnic segregation among secondary schools in London across the first decade of the 21st century. While we illustrate our models in terms of ethnicity, the models can be applied to measuring composition and segregation among social and other groupings.

Several other inferential approaches have been proposed whose benefits overlap those presented here. Ransom (2000) derives exact sampling distributions for two indices

3

and develops an asymptotic test for the change in segregation across two occasions. Allen et al. (2009) propose a bootstrap procedure for comparing pairs of segregation index values, while Rathelot (2012) propose a parametric approach. While each of these approaches could, in principle, be extended to multigroup segregation indices, such extensions have not yet been pursued. Nor do these approaches allow for dependence on covariates. In some cases a segregation index is used as a response variable in a linear regression model, with aggregated variables as covariates. However, the regression coefficients fail to account for the varying reliability of each segregation index value. The bootstrap has been proposed to address the latter problem (Willms and Paterson, 1995), but it is not clear how one would extend this linear regression approach to multigroup segregation where there are often multiple interdependent pairwise segregation index values per area. In addition, this approach cannot incorporate covariates measured at the individual or organisational unit level.

Our approach also allows the analyst to directly address issues of public policy interest. Individuals' attitudes about race and ethnicity, multiculturalism and citizenship are often arguably shaped by individuals' school experiences and the degree to which children of different backgrounds interact. For example, there is a continuing debate as to whether social and ethnic segregation has changed as a result of policies designed to encourage parents to exercise choice of school. Likewise, the recent promotion of academies and free schools may also be expected to affect the extent of ethnic and social segregation. The ability to statistically model changes in composition and segregation is an important part of such debates.

The next section introduces the data used in the application. Section 3 describes the multilevel random-coefficient multinomial logistic regression model. Section 4 extends our earlier simulation method to reexpress the model segregation parameters in the metric of any desired segregation index. Section 5 presents the application. Section 6 concludes.

2. Data and Descriptive Analysis of Changing Patterns of Ethnic Composition and Segregation

A number of studies have described the changing patterns of ethnic composition in schools in England (for example, Burgess and Wilson, 2005; Burgess et al., 2005; Johnston et al., 2004, 2005, 2006, 2008). These studies show non-white students are concentrated in cities, particularly London, and that the proportion of non-white students is increasing over time. Johnston et al. (2006) report that nearly half (44%) of England's non-white secondary school students lived in London in 2003. Hamnett (2011) reports the percentage of non-white London secondary school students increased from 40.3% in 1999 to 52.6% in 2009. In terms of changing patterns of ethnic segregation, Johnston et al. (2008) use the index of isolation to report a general picture of considerable, but stable, ethnic segregation in England over the period 1997-2003. However, they highlight some notable exceptions; in particular, they report that black Africans in London became slightly more segregated over this period.

2.1 Annual school census data

The data are drawn from the annual school census (ASC), a census of all schools in the state education system in England (http://www.education.gov.uk/researchandstatistics/national-pupil-database). We

examine the changing ethnic composition of London secondary schools across 10 consecutive entry cohorts of students, 2001-2010, at the point at which they entered their schools (age 11, national curriculum year seven).

We analyse three ethnic groups: 'white', 'black' and 'Asian'. The white group consists of white British, white Irish and other white students. The black group consists of black African and black Caribbean students. The Asian group consists of Indian, Pakistani and Bangladeshi students. We focus on three ethnic groups in order to illustrate the approach as simply as possible. However, the approach readily extends to the case of analysing four or higher numbers of ethnic groups. For example, black African students could be distinguished from black Caribbean students. Similarly, mixed-ethnicity and other students not covered by our three-group classification (12 to 16% of all students in each cohort) could also be included. Researchers will learn more by considering a range of ethnic classifications rather than just one. Plewis (2011) provides a detailed description of the full ethnic classification available in these data.

Over the period to which our data relate, London schools came under the responsibility of 32 local authorities (LAs) (equivalent to school districts in the US). The most central 12 LAs form inner-London, while the 20 surrounding LAs form outer-London. Across the 10 cohorts 395 schools are represented and 81% are present for all 10 years. An advantage of the multilevel approach is its ability to handle this imbalance. In total we observe 3667 school-cohorts with, on average, 161 students per school-cohort.

2.2. Changing macro-level patterns of ethnic composition and segregation

The data show the ethnic composition of London students as a whole changed between 2001 and 2010. The proportion of white students fell from 0.64 to 0.53 while the proportion of black and Asian students increased from 0.20 to 0.24 and from 0.16 to 0.23, respectively. Fig. 1 shows that this pattern differs across inner- (left panel) and outer-London (right panel) suggesting there is ethnic segregation at this most macro of spatial aggregations. Here and throughout the rest of this section, we define 'ethnic segregation' very simply as variation in the ethnic composition of the units under consideration, here inner- and outer-London. Specifically, white students disproportionately attended schools in outer-London while black students disproportionately attended schools in inner-London. Asian students were similarly represented in both inner and outer-London schools. These patterns reflect the differential residential concentrations of these ethnic groups across inner- and outer-London during this period.

2.3. Changing meso-level patterns of ethnic composition and segregation

Fig. 2 plots the proportions of white, black and Asian students between 2001 and 2010 separately for the 32 LAs. The first two rows present the 12 inner-London LAs; rows three to six present the 20 outer-London LAs. The figure shows that the London-wide decline in the proportion of white students, and increased proportions of black and Asian students, applies effectively, albeit to varying degrees, across all LAs. However, more immediately apparent is the substantial heterogeneity in the average proportions of each ethnic group across LAs and the stability in these differences over time. Some LAs tend to have especially high proportions of non-white students in all years: black students are the largest ethnic group in Lambeth and Southwark; while in Tower

7

Hamlets and Newham Asian students are the largest ethnic group. In contrast, Bromley, Havering and Richmond upon Thames have especially low proportions of non-white students in all years. Only in Lewisham (from white to black) and Redbridge (from white to Asian) do we see the dominant ethnic group change between the start and the end of the decade. The substantial variation in the ethnic composition of students across LAs suggests considerable ethnic segregation at this meso-level of analysis, but it is less clear as to the degree to which this might have changed over time.

2.4. Changing micro-level patterns of ethnic composition and segregation

The school proportions of white, black and Asian students between 2001 and 2010 also vary considerably within their LAs (see supplementary materials). While some of this variability will reflect randomness, there appears to be substantial ethnic segregation even at this most micro level of analysis. In these data, where there are so few schools per LA (between 2 and 21), describing the extent to which school-level ethnic segregation might vary from one LA to the next would be of limited value.

3. Multilevel Modelling of Longitudinal Multigroup Segregation

Multigroup segregation data in their simplest form have one record per organisational unit and consist of counts of the number of individuals in each unit who belong to each of several mutually exclusive and exhaustive groups of interest. In our case we have counts of the number of white, black and Asian students in each of 3667 school-cohorts. Data of this form can be viewed as multinomial grouped data where the number of 'trials' is given by the total number of individuals in each unit, while the 'number of successes' in each 'outcome category' is given by the number of individuals in each group. A natural approach to statistically modelling such data is to fit multilevel

8

multinomial logistic regression models (Goldstein, 2011; Raudenbush and Bryk, 2002; Snijders and Bosker, 2012). While we shall present our discussion in terms of these grouped data, we note that expanding the data to one record per individual results in multinomial individual data. An advantage of fitting models to the expanded data is that one can then enter individual-level covariates into the model to examine whether there is segregation in the characteristic being studied over and above that stemming from other characteristics associated with the selection of individuals into units.

3.1. The multilevel random-coefficient multinomial logistic regression model

The simplest model which captures the main features of the data described in Section 2 is a three-level (school-cohorts within schools within LAs) random-coefficient multinomial logistic regression model for the observed numbers of white, black and Asian students in each school-cohort. This can be written as

$$n_{tjk}^{[W]}, n_{tjk}^{[B]}, n_{tjk}^{[A]} \sim \text{Multinomial}\left(\pi_{tjk}^{[W]}, \pi_{tjk}^{[B]}, \pi_{tjk}^{[A]}, n_{tjk}\right),$$

$$\log\left(\frac{\pi_{tjk}^{[B]}}{\pi_{tjk}^{[W]}}\right) = \underbrace{\beta_0^{[B]} + \beta_1^{[B]}t + \beta_2^{[B]}z_k + \beta_3^{[B]}z_k \times t}_{\text{black−white fixed−part}} + \underbrace{v_{0k}^{[B]} + v_{1k}^{[B]}t + u_{0jk}^{[B]} + u_{1jk}^{[B]}t + c_{tjk}^{[B]}}_{\text{black−white random−part}},$$

$$\log\left(\frac{\pi_{tjk}^{[A]}}{\pi_{tjk}^{[W]}}\right) = \underbrace{\beta_0^{[A]} + \beta_1^{[A]}t + \beta_2^{[A]}z_k + \beta_3^{[A]}z_k \times t}_{\text{Asian−white fixed−part}} + \underbrace{v_{0k}^{[A]} + v_{1k}^{[A]}t + u_{0jk}^{[A]} + u_{1jk}^{[A]}t + c_{tjk}^{[A]}}_{\text{Asian−white random−part}},$$

$$(1)$$

$$\begin{pmatrix} v_{0k}^{[B]} \\ v_{1k}^{[B]} \\ v_{0k}^{[A]} \\ v_{1k}^{[A]} \end{pmatrix} \sim \text{N}\left\{ \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{v0}^{2[B]} & & & \\ \sigma_{v01}^{[B]} & \sigma_{v1}^{2[B]} & & \\ \sigma_{v00}^{[BA]} & \sigma_{v10}^{[BA]} & \sigma_{v0}^{2[A]} & \\ \sigma_{v01}^{[BA]} & \sigma_{v11}^{[BA]} & \sigma_{v01}^{[A]} & \sigma_{v1}^{2[A]} \end{pmatrix} \right\},$$

$$
\begin{pmatrix} u_{0jk}^{[B]} \\ u_{1jk}^{[B]} \\ u_{0jk}^{[A]} \\ u_{1jk}^{[A]} \end{pmatrix} \sim N \left\{ \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{u0}^{2[B]} & & & \\ \sigma_{u01}^{[B]} & \sigma_{u1}^{2[B]} & & \\ \sigma_{u00}^{[BA]} & \sigma_{u10}^{[BA]} & \sigma_{u0}^{2[A]} & \\ \sigma_{u01}^{[BA]} & \sigma_{u11}^{[BA]} & \sigma_{u01}^{[A]} & \sigma_{u1}^{2[A]} \end{pmatrix} \right\},
$$

$$
\begin{pmatrix} c_{tjk}^{[B]} \\ c_{tjk}^{[A]} \end{pmatrix} \sim N \left\{ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_c^{2[B]} & \\ \sigma_c^{[BA]} & \sigma_c^{2[A]} \end{pmatrix} \right\},
$$

where $n_{tjk}^{[W]}$, $n_{tjk}^{[B]}$ and $n_{tjk}^{[A]}$ denote the observed counts of white, black and Asian students in cohort t ($t = 2001, \dots, 2010$) in school j ($j = 1, \dots, 395$) in LA k ($k = 1, \dots, 32$), respectively. The total number of students per school-cohort is denoted n_{tjk}. The three corresponding underlying proportions of white, black and Asian students, $\pi_{tjk}^{[W]}$, $\pi_{tjk}^{[B]}$ and $\pi_{tjk}^{[A]}$, are then related to the model parameters via black-white and Asian-white contrast equations.

Fig. 1 showed the proportion of white (black and Asian) students decreased (increased) over the 10 years in both inner- and outer-London, and that white students disproportionately attended schools in outer-London while black students disproportionately attended schools in inner-London. We capture these features of the data through the fixed-part of the black-white and Asian-white contrasts. Each contrast includes a constant, a linear time trend t (coded $0, 1, \dots, 9$ for the ten consecutive cohorts), a binary indicator for inner-London LAs z_k, and their cross-level interaction $z_k \times t$. Thus, in the black-white contrast, the parameters $\beta_0^{[B]}$ and $\beta_1^{[B]}$ give the intercept and slope for the outer-London time trend while the sums $\beta_0^{[B]} + \beta_2^{[B]}$ and $\beta_1^{[B]} + \beta_3^{[B]}$ give the intercept and slope for the inner-London time trend. The Asian-white contrast

is similarly defined. While one could extend the fixed-part of the model to include more flexible functions of time (Fitzmaurice, Laird and Ware, 2011; Hedeker and Gibbons, 2006), we focus here on linear time trends to illustrate our approach as simply as possible.

Fig. 2 showed considerable heterogeneity across LAs in both the initial proportions of each ethnic group and in the rates of change of these proportions over the 10 years. We capture these features of the data through the LA-level random-part of the model. For example, in the black-white contrast, we include LA-specific random-intercept and -slope effects, $v_{0k}^{[B]}$ and $v_{1k}^{[B]}$. Adding these amounts to the relevant inner- or outer-London time trend gives a unique linear time trend for each LA.

Section 2.4 described that even within LAs, there appears to be considerable heterogeneity across schools in both the initial proportions of each ethnic group and in the rates of change of these proportions over the 10 years. We capture these features of the data through the school-level random-part of the model. For example, in the black-white contrast, we include school-specific random-intercept and -slope effects, $u_{0jk}^{[B]}$ and $u_{1jk}^{[B]}$. Adding these amounts to the relevant LA time trend, gives a unique linear time trend for each school.

We include school-cohort-specific random-intercept effects $c_{tjk}^{[B]}$ at the lowest level of analysis in order to capture any remaining overdispersion (extramultinomial variation) in the data (Skrondal and Rabe-Hesketh, 2007). Note that this approach does not allow for underdispersion because variances are nonnegative. However, in most applications,

underdispersion is less prevalent than overdispersion and it is unlikely here where we make only limited adjustments for covariates.

The random-effects are assumed multivariate normally distributed at each level and independent across levels. While inference in multilevel models is typically robust to moderate departures from normality, severe skewness or outliers can pose problems. The plausibility of these assumptions can be readily examined by plotting quantile-quantile plots (normal score plots) or other residual diagnostic plots of the posterior (shrunken) estimates of the random effects.

3.2. Variance functions

The variability in the four LA random-effects is summarised by the 4×4 LA-level covariance matrix. This matrix provides a direct summary of the degree to which students are segregated across LAs at each point in time, having adjusted for segregation associated with the changing overall proportions of each ethnic group in inner- and outer-London. The corresponding school-level covariance matrix provides a direct summary of the degree to which students are segregated across schools at each point in time, over and above the influence of LA-level segregation. To facilitate the interpretation of these covariance matrices, we derive variance functions which map out how the degree of black-white and Asian-white segregation changes over time at each level of the model.

The LA-level black-white and Asian-white variance functions are written as

$$\text{Var}\left(v_{0k}^{[s]} + v_{1k}^{[s]}t \Big| t\right) = \sigma_{v0}^{2[s]} + 2\sigma_{v01}^{[s]}t + \sigma_{v1}^{2[s]}t^2, \quad s = \text{B, A.} \tag{2}$$

The school-level black-white and Asian-white variance functions are written as

$$\text{Var}\left(u_{0jk}^{[s]} + u_{1jk}^{[s]}t \,\middle|\, t\right) = \sigma_{u0}^{2[s]} + 2\sigma_{u01}^{[s]}t + \sigma_{u1}^{2[s]}t^2, \quad s = \text{B}, \text{A}. \tag{3}$$

These functions can be further manipulated to provide additional insights. For example, summing (2) and (3) when $s = \text{B}$ gives a new function measuring how combined LA- and school-level black-white segregation changes over time, while dividing (2) by the sum of (2) and (3) would quantify the changing relative importance of LA- and school-level black-white segregation over time.

We can perform global tests for the presence of LA- or school-level segregation by comparing the fit of this model to a model with no LA- or no school-level random-effects, respectively. However, Fig. 2 and the underlying data suggest that there is substantial ethnic segregation at each level of analysis and so of more interest is to test whether the degree of ethnic segregation at each level has changed significantly over time. We can do this by comparing the fit of the current model to a model where we retain the random-intercept effects at a given level, but remove the random-slope effects. The simpler model then allows for segregation at each level, but constrains the degree of segregation at the specified level to be constant over time. More nuanced tests of segregation can be carried out by testing the joint significance of different combinations of parameters in each of the variance functions.

Equations (2) and (3) are constrained quadratic functions of time. Specifically, they are of the form $a + bt + ct^2$ where $a \geq 0$ (we do not envisage negative intercept variances) and $c \geq 0$ (we do not envisage negative time trend variances). These quadratic

13

functions are therefore convex and always return positive values. The coefficient a gives the degree of segregation in the first cohort ($t = 0$). The interpretation of b, however, is somewhat more subtle due to the constraint on the quadratic term. When $b > 0$, segregation strictly increases with time; when $b < 0$ segregation initially decreases with time, but, since $c \geq 0$, will at some future point in time reach a minimum and then increase from there on. This turning point may lie beyond the end of the observation period in which case segregation would be seen to strictly decrease within the range of the data. Plotting variance functions typically aids their interpretation.

Given the above discussion, modelling black-white and Asian-white segregation as quadratic functions of time may seem unduly restrictive, particularly when there are many time points. However, there is nothing to stop us including higher-order polynomial LA- and school-specific time trends in the model in order to allow for more flexible variance functions. For example, including LA-specific quadratic time trends would result in variance functions which additionally include cubic and quartic components. A comparison of the fit of this model to that of the simpler linear time trend model then provides evidence as to whether this added complexity is required. We can even explore whether different order polynomials are required in each contrast. More generally, including time trends of polynomial degree p implies variance functions of polynomial degree $2p$. We may also include fractional polynomials. A different approach altogether may be required if the degree of segregation is not expected to change in a smooth continuous fashion over time. We can, for example, directly specify variance functions which are step functions of time to cater for such eventualities, but we shall not entertain such possibilities further here.

When describing patterns of ethnic segregation, it is also substantively interesting to examine the degree to which LAs (schools) where blacks are overrepresented relative to whites are also the LAs (schools) where Asians are overrepresented relative to whites. We can examine this by deriving LA- and school-level correlation functions which describe how the correlation in the (log of the) black-white and Asian-white ratios change over time at each level of analysis. First we must calculate the LA-level black-white and Asian-white covariance function, written as

$$\text{Cov}\left(v_{0k}^{[B]} + v_{1k}^{[B]}t, v_{0k}^{[A]} + v_{1k}^{[A]}t \middle| t\right) = \sigma_{v00}^{[BA]} + \left(\sigma_{v01}^{[BA]} + \sigma_{v10}^{[BA]}\right)t + \sigma_{v11}^{[BA]}t^2. \qquad (4)$$

The school-level black-white and Asian-white covariance function takes the same mathematical form and is written as

$$\text{Cov}\left(u_{0jk}^{[B]} + u_{1jk}^{[B]}t, u_{0jk}^{[A]} + u_{1jk}^{[A]}t \middle| t\right) = \sigma_{u00}^{[BA]} + \left(\sigma_{u01}^{[BA]} + \sigma_{u10}^{[BA]}\right)t + \sigma_{u11}^{[BA]}t^2. \qquad (5)$$

The associated correlation functions are derived by dividing each covariance function by the square root of the product of the two corresponding variance functions.

We have discussed black-white and Asian-white segregation; however, Asian-black segregation may also be of interest. The LA-level Asian-black variance function is given by (2) (where $s = B$) plus (2) (where $s = A$) minus two times (4). The school-level Asian-black variance function is similarly defined in terms of (3) and (5).

3.3. Adjusting for covariates

Having modelled the multilevel and longitudinal patterns of LA- and school-level segregation, we might then be interested in explicitly comparing these patterns across a limited number of population subgroups. For example, while we have modelled how the overall proportions of white, black and Asian students differ between inner- and outer-London, it might be interesting to additionally examine whether segregation is higher in inner- or outer-London. For example, black-white school-level segregation may be decreasing over time in inner-London, but increasing over time in outer-London. The simplest approach to capture such patterns is to fit the model separately to each subgroup, but this precludes testing for differences in segregation across subgroups. A more appropriate approach is to model the subgroups jointly by interacting the parameters in the model with a series of binary subgroup indicator variables.

More generally, covariates can be included in the fixed-part of the model to explain the variation in ethnic proportions across schools and across LAs. The resulting variance functions then measure the degree of segregation at each point in time having adjusted for these variables. Additionally including covariates in the random-part of the model allows us to explicitly model how segregation changes as a function of these covariates. For example, school admissions policies are known to differ across LAs with some LAs selecting children into schools based on their academic ability to a greater extent than other LAs. LA-level variation in school admissions policies might then be expected to give rise to LA differences in school-level ethnic segregation if students' test scores are associated with ethnicity.

3.4. Estimation

We fit all models using Markov chain Monte Carlo (MCMC) methods as implemented in the MLwiN multilevel modelling package (Rasbash et al., 2009). We have chosen to call MLwiN from within Stata using the runmlwin command (Leckie and Charlton, 2013). Estimates obtained using the quasi-likelihood methods implemented in MLwiN are used as initial values for all parameters. All models are run for a burn-in of 50,000 iterations followed by a monitoring chain of 500,000 iterations. Minimally informative prior distributions are specified for all parameters. We use hierarchical centring parameterisations (Browne, 2012; Browne et al., 2009) to produce chains that exhibit better mixing. Informal visual assessments of the parameter chains and standard MCMC convergence diagnostics suggest that the sampler was run for sufficiently long. The MCMC approach allows the fit of models to be compared via the deviance information criterion (DIC; Spiegelhalter et al., 2002): models with smaller DIC values are preferred to those with larger values, with differences of five or more considered substantial (Lunn et al., 2012).

4. Simulating multigroup segregation indices based on the fitted multilevel model

In this section, we illustrate how our simulation method can be used to reexpress the model segregation parameters into the metric of a multigroup segregation index. Specifically, we reexpress the overall level of segregation (combined macro-, meso- and micro segregation) in each cohort implied by the model into the metric of the multigroup information theory index (Theil, 1972; Theil and Finezza, 1971). This index is widely applied in studies of residential multigroup ethnic segregation (Farrell, 2008; Fischer, 2003; Fischer et al., 2004; Iceland, 2004) and in studies of multigroup ethnic school segregation (Reardon et al., 2000).

4.1. Calculating the multigroup information theory index on observed data

The multigroup information theory index is a measure of evenness – the extent to which groups are evenly distributed across organisation units (Massey and Denton, 1988). The index can be interpreted as the difference between the ethnic diversity of the school system (the degree to which ethnic groups are equally represented) and the weighted average diversity of individual schools. A value of zero corresponds to no segregation (complete integration) – the ethnic composition of every school is the same – while a value of one corresponds to complete segregation (no integration) – the ethnic composition of each school consists of only one ethnic group. The index, denoted H_t, is calculated as

$$H_t = \frac{E_t - \sum_j \frac{n_{tjk}}{n_t} E_{tjk}}{E_t},$$

where E_t and n_t denote the ethnic diversity and total number of students across London in cohort t, while E_{tjk} and n_{tjk} denote the corresponding quantities in cohort t in school j in LA k. (Note that summing over j implicitly sums over k as we have defined j as a unique identifier rather than a nested identifier.) The London-wide diversity score is calculated as

$$E_t = -\sum_s p_t^{[s]} \ln\left(p_t^{[s]}\right),$$

where $p_t^{[s]}$ is the proportion of London school students in ethnic group s while the analogous diversity score for school j in LA k is calculated as

$$E_{tjk} = -\sum_s p_{tjk}^{[s]} \ln\left(p_{tjk}^{[s]}\right).$$

When $p_{tjk}^{[s]} = 0$, $\ln\left(p_{tjk}^{[s]}\right)$ is set to zero since $\lim_{p_{tjk}^{[s]} \to 0}\left\{-p_{tjk}^{[s]} \ln\left(p_{tjk}^{[s]}\right)\right\} = 0$. In each case, the

diversity score ranges from zero (no diversity; only one ethnic group is represented) to $\ln(S)$ (maximum diversity; each of the S ethnic groups is equally represented). When $E_{tjk} < E_t$, school j is less ethnically diverse than London as a whole and vice versa.

The index provides a one number summary of ethnic segregation between multiple groups. A limitation of this index is that it says nothing about how segregated pairs of groups are (pairwise segregation). For this purpose a two-group version of this index is sometimes employed to compare each ethnic group in turn to every other ethnic group in turn (for example, black-white then black-Asian then Asian-white). Alternatively, the above formulas are sometimes used to compare each ethnic group in turn to the remaining ethnic groups combined (for example, black-non-black, Asian-non-Asian, white-non-white).

4.2. *Simulating model-implied values of the multigroup information theory index*
The simulation method consists of repeating the following steps a large number of times

1. Simulate a new set of random-effects for each LA from the estimated LA-level covariance matrix.
2. Simulate a new set of random-effects for each school from the estimated school-level covariance matrix.

3. Simulate a new set of random-effects for each school-cohort from the estimated school-cohort-level covariance matrix.

4. Use the model contrast equations to compute a new set of proportions $\pi_{tjk}^{[s]}$ based on the observed covariates and the simulated random effects.

5. Compute the overall proportions of each ethnic group

$$\pi_t^{[s]} = \sum_j \frac{n_{tjk}}{n_t} \pi_{tjk}^{[s]}.$$

6. Compute the overall diversity score

$$E_t = -\sum_s \pi_t^{[s]} \ln\left(\pi_t^{[s]}\right).$$

7. For each school-cohort, compute the school-cohort-level diversity score

$$E_{tjk} = -\sum_s \pi_{tjk}^{[s]} \ln\left(\pi_{tjk}^{[s]}\right).$$

8. Compute the overall index score

$$H_t = \frac{E_t - \sum_j \frac{n_{tjk}}{n_t} E_{tjk}}{E_t}.$$

The index point estimate is given by its mean value across the repetitions while its simulation uncertainty is summarised by the 95% interval calculated by taking the 2.5[th]

20

and 97.5th percentiles of the rank ordered list of index values as the lower and upper bounds of the interval. The above simulation method can be readily extended to reexpress differences in model-implied segregation associated with different choices of covariate values. For example, we might test whether the degree of segregation changed significantly between the start and end of the decade.

As they are currently stated, the above steps do not propagate the uncertainty in the estimated model parameters. We fit the models using MCMC methods and can therefore use the parameter chains for this purpose, since they are samples from the joint posterior distribution of the parameters. Specifically, we can repeat the above method a large number of times where on each repetition we use the values of the parameters at the corresponding iteration of the parameter chains. The LA-level variance-covariance parameters will be estimated relatively imprecisely (there are only 32 LAs) and so propagating the parameter uncertainty will prove particularly important at this highest level of analysis.

5 Application of the modelling and simulation approach to the data

We start by fitting four three-level (school-cohorts within schools within LAs) multinomial logistic regression models to establish whether the overall degree of LA- and school-level ethnic segregation among London schools changed significantly over the first decade of the 21st century.

Model 1 is a random-intercept version of the model presented in Section 3. This model captures the main features of the data shown in Fig. 1 by including a time trend, a binary indicator for inner-London LAs, and their cross-level interaction in the fixed-part of the

model. However, only a random-intercept is entered at each level in the random-part of the model and so the model assumes both the degree of LA- and school-level segregation to be constant over time. Model 2 extends Model 1 by allowing the slope of the linear time trend to vary randomly across LAs, as suggested in Fig. 2. By including time in the random-part of the model at the LA-level, but not at the school-level, this model allows the degree of LA-level segregation to change over time, but continues to assume the degree of within-LA school-level segregation is constant. In contrast, Model 3 only allows the slope of the linear time trend to vary randomly across schools, allowing the degree of within-LA school-level segregation to change over time, but restricting the degree of LA-level segregation to be constant. Model 4, described in detail in Section 3, allows the time trend to vary randomly at both levels, allowing both LA-level and school-level segregation to change over time.

Table 1 summarises the fit of these four models. Model 2 is preferred to Model 1 (the DIC improves by 294 points) indicating that the degree of LA-level segregation changed significantly between 2001 and 2010. Model 3 is preferred to both Model 1 (957 points) and Model 2 (663 points) suggesting that there was a change in school-level segregation over and above that implied by changing LA-level segregation. Model 4, which explicitly allows for different changing patterns of segregation at each level, is preferred to both Model 2 (694 points) and Model 3 (31 points).

Table 2 presents the posterior means and standard deviations (analogous to frequentist parameter estimates and standard errors) of the fixed- and random-part parameters for Model 4. The fixed-part of the model confirms the pattern of results suggested in Fig. 1. The similar sized negative black-white ($\beta_0^{[B]} = -1.814$) and Asian-white ($\beta_0^{[A]} =$

−2.097) intercepts confirm that, at the beginning of the decade, whites were the largest ethnic group in outer-London, followed by similar sized black and Asian ethnic groups. Of more interest is how these proportions change over time. In outer-London, the black-white ($\beta_1^{[B]} = 0.080, z = 3.11, p = 0.002$) and Asian-white ($\beta_1^{[A]} = 0.085, z = 3.18, p = 0.001$) ratios are shown to increase significantly over time and so the black and Asian shares of the population increased relative to the white share. In inner-London, the black-white ratio was significantly higher at the beginning of the decade compared to in outer-London ($\beta_2^{[B]} = 1.309, z = 3.23, p = 0.001$) reflecting the higher proportion of black students schooled in inner-London at that time. This initial difference did not change significantly over time ($\beta_3^{[B]} = -0.031, z = -0.63, p = 0.528$). In contrast, the Asian-white ratio did not differ significantly at the beginning of the decade between inner- and outer-London ($\beta_2^{[A]} = 0.318, z = 0.68, p = 0.496$). This initial difference also did not change significantly over time ($\beta_3^{[A]} = -0.001, z = -0.03, p = 0.976$). In sum, the fixed-part of the model shows London became significantly more diverse between 2001 and 2010, but not differentially so between inner- and outer-London. In other words, the degree to which students were ethnically segregated across inner- and outer-London at the beginning of the decade did not change over the course of the decade.

Turning our attention to the random-part of the model, we see the magnitude of the school-cohort variances are substantively small relative to the LA- and school-level variability. Thus, having accounted for the LA and school linear time trends, there is very little remaining overdispersion in the data. The individual LA- and school-level random-part parameters are harder to interpret directly and so we interpret them

below via their implied variance functions. However, what we do see straightaway is that that the LA-level variance-covariance parameters are estimated less precisely than the school-level variance-covariance parameters reflecting the lower number of units at this level (32 LAs schools compared to 395 schools). The LA-level variance function should therefore be interpreted with some caution.

Fig. 3 plots the total, LA- and school-level variance functions (derived from Equations 2 and 3 and the school-cohort variances). We do not plot the school-cohort-level variance function as it is modelled as constant with respect to time and so does not contribute to any changing patterns of segregation. The total variance functions suggest that black-white, Asian-white and especially Asian-black segregation all increased strongly between 2001 and 2010. This overall pattern appears to be driven by increasing segregation at the LA level; school-level segregation remains effectively constant for the three ethnic contrasts. However, these patterns disagree with those resulting from fitting a series of simpler cohort-specific variance-components models to the data. The latter models show LA-level segregation to decrease slightly over the period (see Fig. A2 in the supplementary materials). Centring time around its midpoint as opposed to around 2001 leads to greater agreement. However, essentially, the LA-level random-part of Model 4 is overly complex (10 parameters) given the number of data points at this level (32 LAs) leading the resulting estimated variance-covariance matrix to be untrustworthy. We therefore simplify this part of the model by removing the random-slopes, returning us to the next best fitting model, Model 3, which assumes constant LA-level variance (4 parameters). A limitation of this simplification is that the slight decrease in LA-level segregation suggested by cohort-specific variance-components

24

models will now appear at the school-level. Thus, any changing LA-level segregation will be conflated with changing school-level segregation.

The fixed-part parameters for Model 3 (see Table A1 in the supplementary materials) are effectively the same as those for Model 4 and so we don't interpret them again here. Fig. 4 plots the total, LA-level and school-level segregation functions based on this model. We see straightaway that the total degree of segregation remained broadly stable over the decade. At the LA-level, Asian and white students appear more segregated than do Asian and black students or white and black students, but these differences are not significant. Within their LAs, Asian and white students appear most segregated, followed by Asian and black students, and lastly black and white students. Comparing across levels, black-white segregation is approximately equal within and between LAs, while Asian-black and especially Asian-white segregation appear higher within LAs compared to between LAs. Thus, even within LAs, where schools are located only a short distance apart, there is substantial variation in the proportions of white, black and Asian students among schools. The finding that Asian-white segregation, relative to the other ethnic contrasts, occurs disproportionally within LAs suggests that Asian-white segregation may be driven relatively more by school-level processes than LA-level processes compared to the other ethnic contrasts. Table 3 presents the estimated change in school-level segregation between 2001 and 2010, together with 95% credible intervals. Asian-black and especially Asian-white school-level segregation both decreased over the decade, but neither significantly; black-white segregation remained very constant over this period.

Fig. 5 plots the total, LA- and school-level black-white-Asian-white correlation functions (derived in each case by dividing the corresponding covariance function by the square root of the product of the two associated variance functions). We again see stable patterns over time. The LA-level correlation of 0.58 shows that LAs where blacks were overrepresented relative to whites tended to be the same LAs where Asians were overrepresented relative to whites. We see a very similar pattern within LAs. The strength of the association at this-level drops slightly over time from 0.57 to 0.54, but this decline is not significant.

Fig. 6 presents quantile-quantile plots of the LA- (top panel) and school-level (bottom-panel) studentised residuals (posterior means scaled by their posterior standard deviations) (Langford and Lewis, 2002). The figure shows the LA and school data points mostly lie on or close to the 45 degree line suggesting that the multivariate normality assumptions are broadly acceptable. There is however some suggestion that the school residuals have somewhat heavier tails than would be expected from normal distributions and that some schools might be outlying. An examination of these potential outliers proves substantively revealing (see supplementary materials). For example, the school with the highest Asian-white random-intercept effect turns out to be of Sikh denomination and so admits effectively no white (or black) students in any year. In contrast, the school with the second highest Asian-white random slope effect is a Jewish school that changed their admissions policies towards the end of the decade to start admitting students practising other faiths, especially Asian students. As a third example, the school with the lowest black-white random-slope effect actually saw a substantial decrease in the proportion of black students and a substantial increase in the proportion of white students between 2001 and 2010. Further examination reveals

26

the school was placed in special measures in 1998, but over subsequent years improved greatly and is now the most oversubscribed school in its LA. While these schools are in some senses unusual, separate analysis shows that the model results are not unduly sensitive as to whether we include or exclude them from the analysis.

Finally, Fig. 7 provides an illustration of the simulation method described in Section 4. Specifically, we calculate the model-implied value of the multigroup information theory index for the case of measuring total segregation (combined LA- and school-level segregation) between the three ethnic groups in each cohort. This contrasts our earlier presentation of variance functions which reported pairwise segregation. As might be expected given the pairwise patterns reported earlier, the figure shows that while overall segregation reduced between 2001 and 2010, this reduction was substantively very small and not significant (posterior mean: -0.012; 2.5th percentile: -0.031; 97.5th percentile: 0.008). Note that the model-implied values of the index are slightly lower than the values derived from the observed proportions, reflecting the randomness and uncertainty contained in the latter.

6. Discussion

In this paper we have extended the multilevel modelling approach to measuring segregation to the general case of modelling longitudinal multigroup segregation. We have described how multigroup segregation data are an example of multinomial grouped data and can therefore be modelled using multilevel multinomial logistic regression. We have demonstrated how time trends can be included in these models with random-coefficients to statistically test whether segregation is increasing or decreasing over time and we have shown that by including multiple levels of random-

27

effects we can measure segregation simultaneously at multiple organisational scales. We also extended and illustrated our simulation approach to reformulating model-implied levels of segregation into the metric of any desired multigroup segregation index. Finally, by introducing covariates into our model, we can derive a segregation index conditional on particular covariate values.

A limitation of the current approach is that the school-level covariance matrix is modelled as constant across LAs. However, we may expect school-level ethnic segregation to be more pronounced in some LAs than others. In theory, we can fully relax this assumption by specifying a separate covariance matrix for each LA, but this would be computationally intensive and the resulting estimates would be very unreliable due to the low number of schools per LA. A more appealing approach is to treat the covariance matrices as exchangeable by specifying them as being drawn from a distribution, with unknown hyper-parameters to be estimated. We are actively exploring this new class of model (Leckie et al., 2014) and are implementing this and related extensions in the new Stat-JR software (Charlton et al., 2013) being developed at the Centre for Multilevel Modelling (CMM).

With regard to the application, the results suggest that the London state secondary school education system has become ethnically more diverse during the first decade of the 21st century but that patterns of segregation have remained largely stable. A weakness of our approach shown in the application is that a relatively high number of parameters are required to model changing segregation at any given level. In our case 32 LAs were too few to fully model the changing patterns of LA-segregation. An important strength of our approach is the ability to identify potential outlying units

through graphically examining the model residuals at each level of the analysis. Applying this approach in our application revealed several schools with unusual changing ethnic compositions. It will often be interesting to follow up such schools to better understand the particular school policies and local circumstances which combine to bring about such anomalies. More broadly, we encourage researchers to use our approach to explore hypotheses regarding the drivers of segregation at each level of analysis through introducing appropriate covariates into the model.

While we have focused on modelling ethnic segregation in the education system, the methods we describe can be equally applied to modelling residential ethnic segregation, for example across wards and output areas (Simpson, 2007). It would then be desirable to extend the models to allow for spatial correlation between the area random-effects. Another interesting extension would be to model residential and schooling segregation jointly using cross-classified (crossed random-effects) versions of our model. This would allow an exploration of whether the schooling system exacerbates or mitigates residential segregation and the extent to which this might vary across different local authorities.

Acknowledgements

References

Allen, R., Burgess, S. and Windmeijer, F. (2009). More Reliable Inference for Segregation Indices, Working Paper No 09/216, Bristol: University of Bristol.

Browne, W. J. (2012) *MCMC Estimation in MLwiN, v2.26*. Centre for Multilevel Modelling, University of Bristol.

Browne, W. J., Steele, F., Golalizadeh, M. and Green, M. J. (2009) The use of simple reparameterizations to improve the efficiency of Markov chain Monte Carlo estimation for multilevel models with applications to discrete time survival models. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **172**, 579-598.

Burgess, S. and Wilson, D. (2005) Ethnic segregation in England's schools. *Transactions of the Institute of British Geographers*, **30**, 20-36.

Burgess, S., Wilson, D. and Lupton, R. (2005) Parallel lives? Ethnic segregation in schools and neighbourhoods. *Urban Studies*, **42**, 1027-1056.

Charlton, C.M.J., Michaelides, D.T., Parker, R.M.A., Cameron, B., Szmaragd, C., Yang, H., Zhang, Z., Frazer, A.J., Goldstein, H., Jones, K., Leckie, G., Moreau, L. and Browne, W.J. (2013). *Stat-JR version 1.0*. Centre for Multilevel Modelling, University of Bristol & Electronics and Computer Science, University of Southampton.

Farrell, C. R. (2008). Bifurcation, Fragmentation or Integration? The Racial and Geographical Structure of US Metropolitan Segregation, 1990-2000. *Urban Studies*, **45**, 467-499.

Fischer, C. S., Stockmayer, G., Stiles, J. and Hout, M. (2004) Distinguishing the geographic levels and social dimensions of US metropolitan segregation, 1960–2000, *Demography*, **41**, 37–59.

Fischer, M. J. (2003) The relative importance of income and race in determining residential outcomes in US urban areas, 1970–2000, *Urban Affairs Review*, **38**, 669–696.

Fitzmaurice, GM, Laird NM, Ware JH (2011) *Applied longitudinal analysis*. 2nd edition. Hoboken, New Jersey: Wiley.

Goldstein, H. (2011) *Multilevel statistical models.* 4th edition. Chichester, UK: Wiley.

Goldstein, H. and Noden, P. (2003). Modelling social segregation. *Oxford Review of Education*, **29**, 225-237.

Hamnett, C. (2011) Concentration or Diffusion? The Changing Geography of Ethnic Minority Pupils in English Secondary Schools, 1999–2009. *Urban Studies*, **49**, 1741-1766.

Hedeker, D. and Gibbons, R. D. (2006) *Longitudinal data analysis*. Hoboken, New Jersey: Wiley.

Iceland, J. (2004) Beyond black and white Metropolitan residential segregation in multi-ethnic America, *Social Science Research*, **33**, 248-271.

Johnston, R., Wilson, D. and Burgess, S. (2004) School segregation in multiethnic England. *Ethnicities*, **4**, 237-265.

Johnston, R., Wilson, D. and Burgess, S. (2005) England's multiethnic educational system? A classification of secondary schools. *Environment and Planning A*, **37**, 45-62.

Johnston, R., Burgess, S., Wilson, D. and Harris, R. (2006) School and residential ethnic segregation: An analysis of variations across England's local education authorities. *Regional Studies*, **40**, 973-990.

Johnston, R., Burgess, S., Harris, R. and Wilson, D. (2008) 'Sleep-walking towards segregation'? The changing ethnic composition of English schools, 1997–2003: an entry cohort analysis, *Transactions of the Institute of British Geographers*, **33**, 73-90.

Langford, I. H. and Lewis, T. (2002) Outliers in multilevel data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **161**, 121-160.

Leckie, G., French, R., Charlton, C. and Browne, W. (2014). Modeling Heterogeneous Variance-Covariance Components in Two-Level Multilevel Models. Centre for Multilevel Modelling, University of Bristol.

Leckie, G. and Charlton, C. (2013) runmlwin - A Program to Run the MLwiN Multilevel Modelling Software from within Stata. *Journal of Statistical Software*, **52** (11), 1-40.

Leckie, G., Pillinger, R., Jones, K. and Goldstein, H. (2012). Multilevel Modeling of Social Segregation. *Journal of Educational and Behavioral Statistics*, **37**, 3-30.

Lunn, D., Jackson, C., Best, N., Thomas, A., Spiegelhalter, D. (2012) *The BUGS Book: A Practical Introduction to Bayesian Analysis*. Chapman & Hall/CRC Texts in Statistical Science.

Massey, D. S. and Denton, N. A. (1988). The dimensions of residential segregation. *Social Forces*, **67**, 281-315.

Plewis, I. (2011) Contextual variations in ethnic group differences in educational attainments. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **174**, 419-437.

Ransom, M. R. (2000). Sampling distributions of segregation indexes. *Sociological Methods & Research*, **28**, 454-475.

Rasbash, J., Charlton, C., Browne, W. J., Healy, M. and Cameron, B. (2009) *MLwiN Version 2.1*. Centre for Multilevel Modelling, University of Bristol.

Rathelot, R. (2012). Measuring segregation when units are small: a parametric approach. *Journal of Business & Economic Statistics*, **30**, 546-553.

Raudenbush, S. W. and Bryk, A. S. (2002) *Hierarchical linear models: Applications and data analysis methods*. 2nd ed. Thousand Oaks, CA: Sage.

Reardon, S. F. and Firebaugh, G. (2002). Measures of multigroup segregation. *Sociological methodology*, **32**, 33-67.

Reardon, S. F., Yun, J. T. and Eitle, T. M. (2000) The changing structure of school segregation: measurement and evidence of multiracial metropolitan-area school segregation, 1989–1995, *Demography*, **37**, 351–364.

Simpson, L. (2007). Ghettos of the mind: the empirical behaviour of indices of segregation and diversity. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **170**, 405-424.

Skrondal, A. and Rabe-Hesketh, S. (2007). Redundant overdispersion parameters in multilevel models for categorical responses. *Journal of Educational and Behavioral Statistics*, **32**, 419-430.

Snijders, T. A. B. and Bosker, R. J. (2012) *Multilevel Analysis: an Introduction to Basic and Advanced Multilevel Modeling*. 2nd ed. London: Sage.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P. and van der Linde, A. (2002) Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **64**, 583-639.

Theil, H. (1972) Statistical decomposition analysis: with applications in the social and administrative sciences. Amsterdam: North-Holland Publishing Company.

Theil, H. and Finizza, A. J. (1971) A note on the measurement of racial integration of schools by means of informational concepts. *The Journal of Mathematical Sociology*, **1**, 187-193.

Willms, J. D. and Paterson, L. (1995). A multilevel model for community segregation. *The Journal of Mathematical Sociology*, **20**, 23-40.

Table 1. DIC statistics for Models 1-4

Model	Trend in LA segregation	Trend in school segregation	DIC
1	No	No	845690
2	Yes	No	845396
3	No	Yes	844733
4	Yes	Yes	844702

Table 2. Model 4 posterior means and standard deviations

Parameter	Black-white equation	Asian-white equation
Fixed-part		
Intercept	-1.814 (0.255)	-2.097 (0.305)
Time	0.080 (0.026)	0.085 (0.027)
Inner	1.309 (0.406)	0.318 (0.467)
Inner × Time	-0.031 (0.050)	-0.001 (0.047)
Random-part – LA		
LA intercept variance	1.214 (0.335)	1.526 (0.448)
LA intercept-slope covariance	-0.026 (0.028)	-0.012 (0.031)
LA slope variance	0.017 (0.005)	0.017 (0.005)
LA intercept-intercept covariance	0.835 (0.318)	
LA intercept-slope covariance	-0.022 (0.032)	
LA slope-intercept covariance	-0.014 (0.027)	
LA slope-slope covariance	0.001 (0.003)	
Random-part – School		
School intercept variance	0.956 (0.078)	1.977 (0.160)
School intercept-slope covariance	-0.025 (0.005)	-0.041 (0.008)
School slope variance	0.006 (0.001)	0.006 (0.001)
School intercept-intercept covariance	0.773 (0.091)	
School intercept-slope covariance	-0.012 (0.007)	
School slope-intercept covariance	-0.016 (0.005)	
School slope-slope covariance	0.002 (0.000)	
Random-part – School-cohort		
School-cohort variance	0.025 (0.002)	0.024 (0.002)
School-cohort covariance	0.015 (0.002)	
DIC	844702	

Table 3. Model 3 Estimated change in school-level segregation (variances) between 2001 and 2010. Posterior means and 95% credible intervals.

Level	Segregation	Posterior mean	2.5th percentile	97.5th percentile
	black-white	0.011	-0.199	0.220
School	Asian-white	-0.285	-0.591	0.015
	Asian-black	-0.086	-0.345	0.154
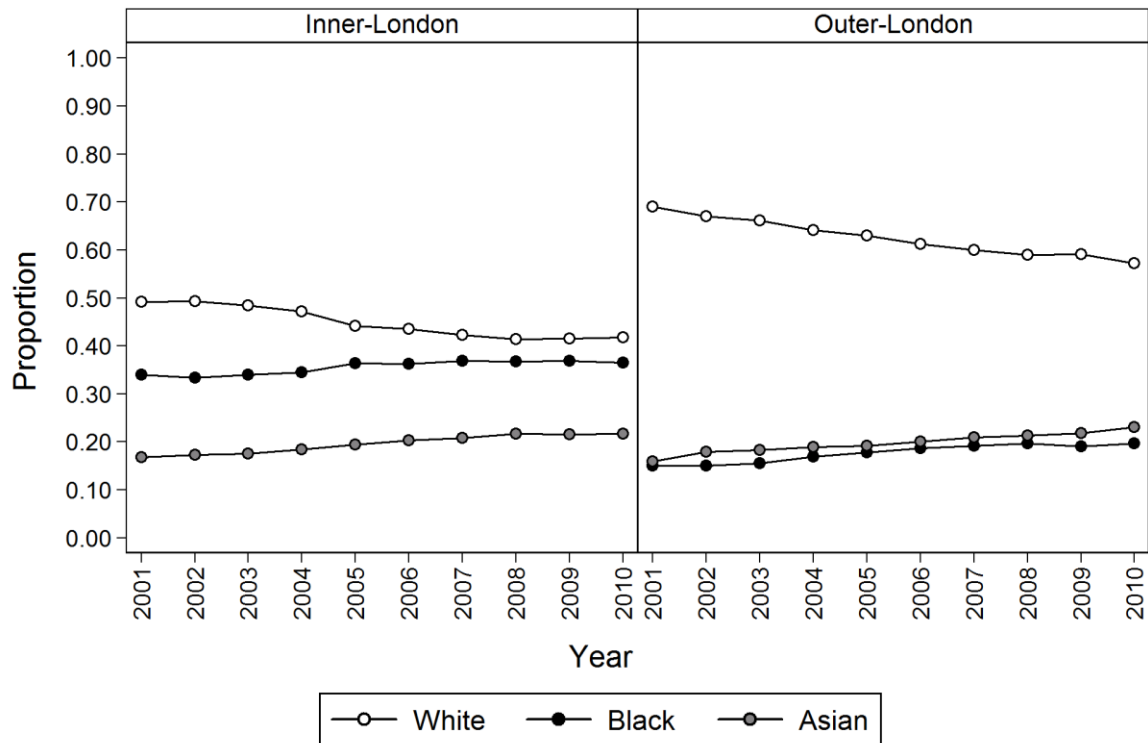
Fig. 1. *Change in observed proportions of white, black and Asian students between 2001 and 2010 plotted separately for inner- (left panel) and outer-London (right panel)*
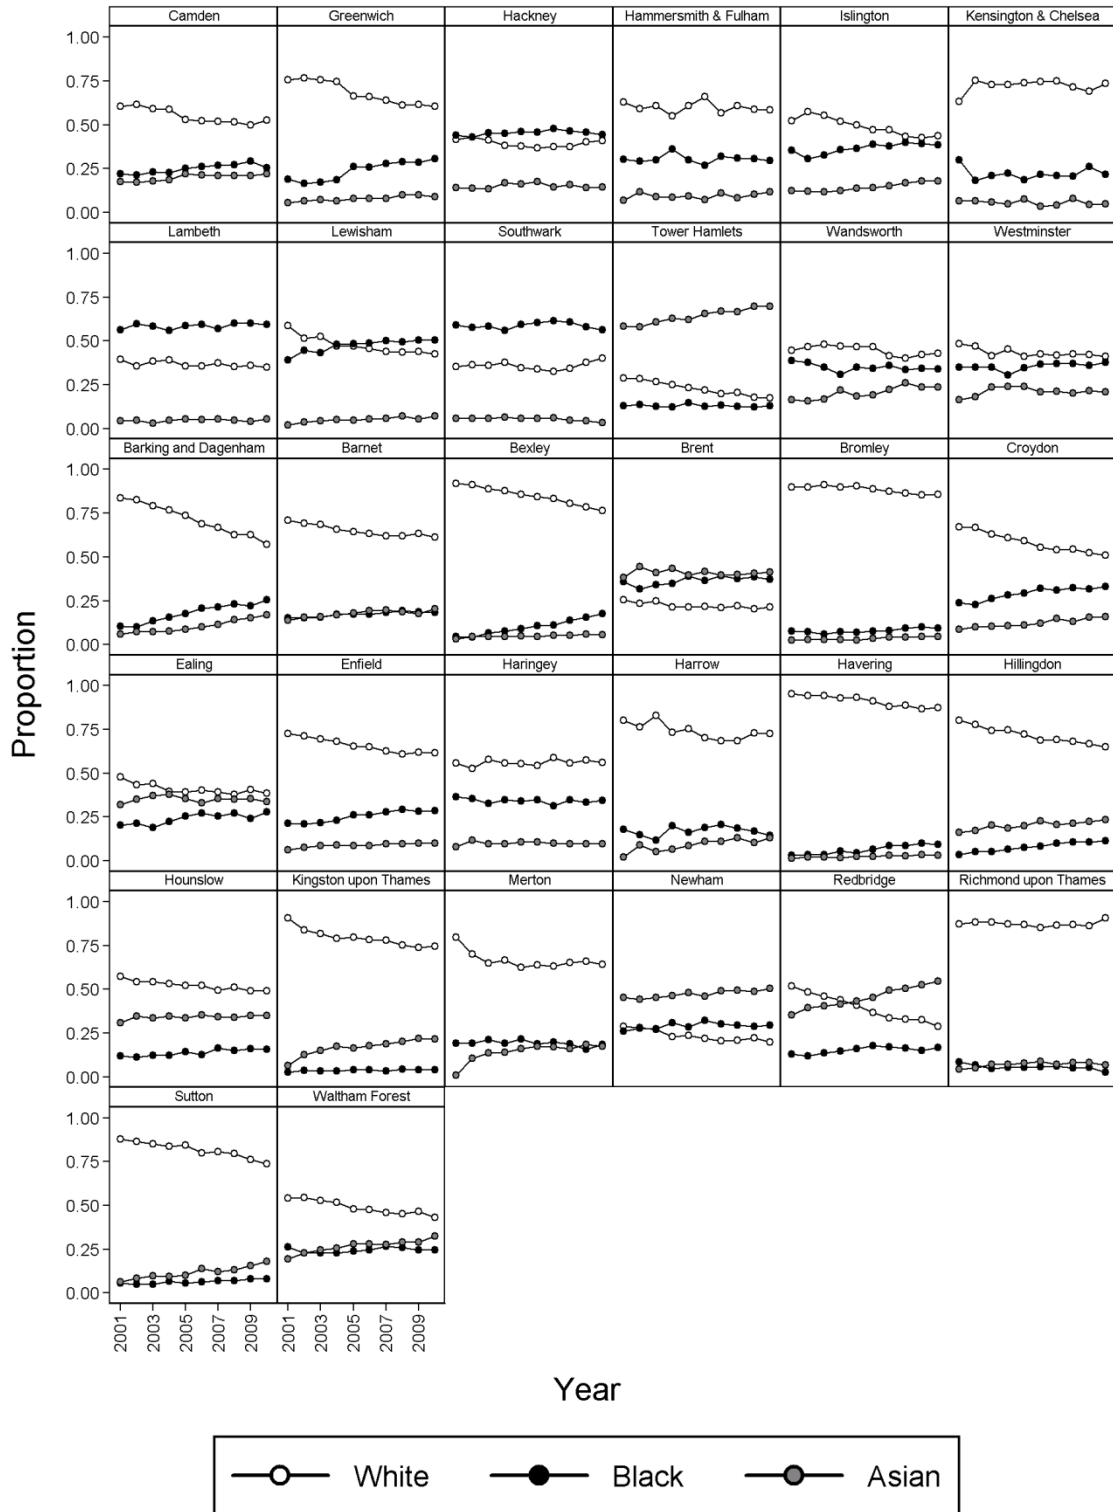
Fig. 2. *Change in observed proportions of white, black and Asian students between 2001 and 2010 plotted separately for each LA. The first two rows present the 12 inner-London LAs; rows three to six present the 20 outer-London LAs.*
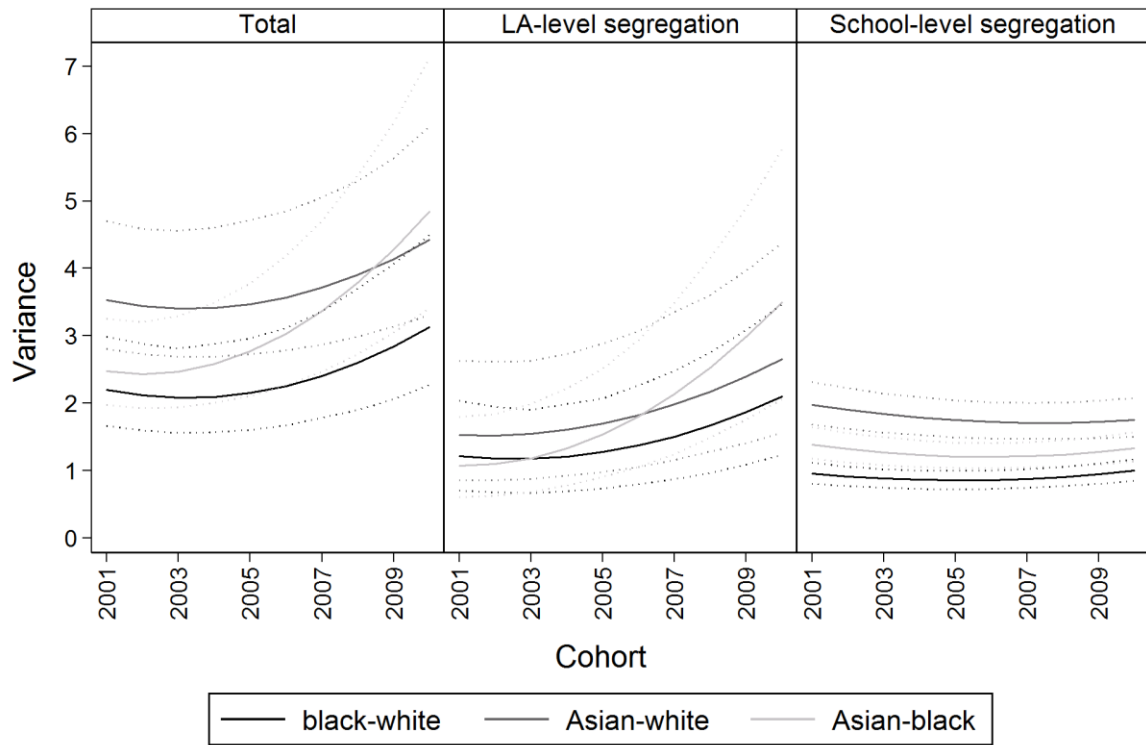
Fig. 3. Model 4 total, LA- and school-level black-white, Asian-white and Asian-black variance functions. Posterior means with 95% credible intervals.
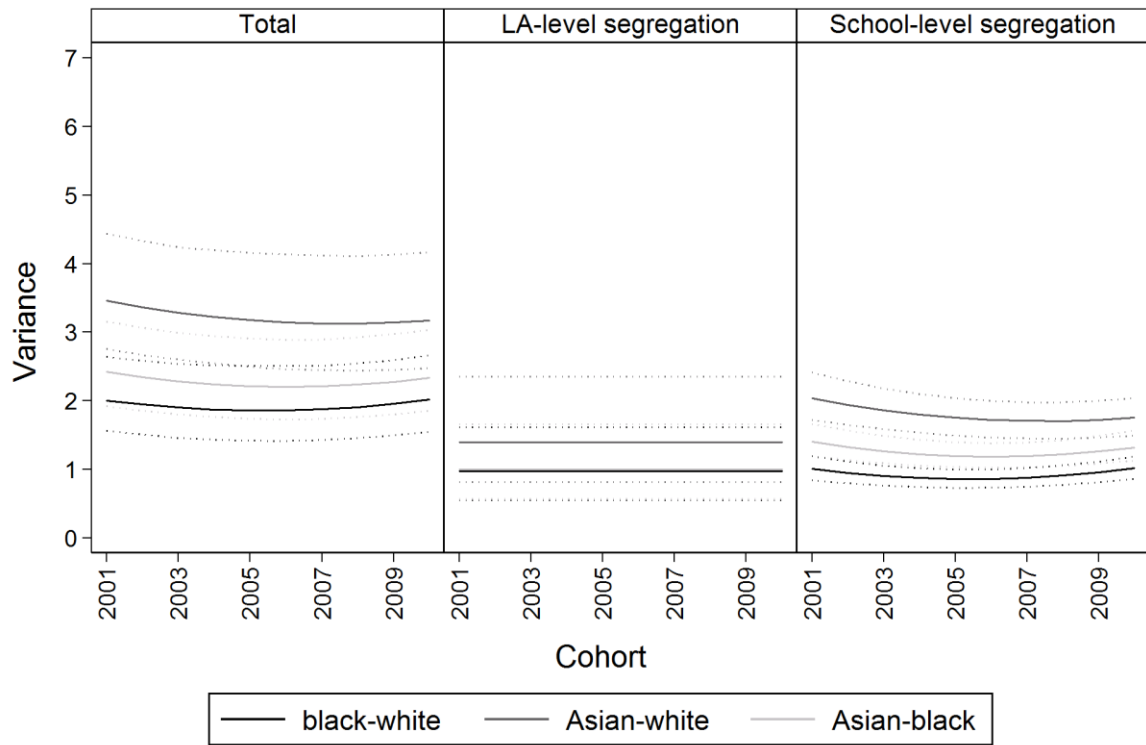
Fig. 4. Model 3 total, LA- and school-level black-white, Asian-white and Asian-black variance functions. Posterior means with 95% credible intervals.
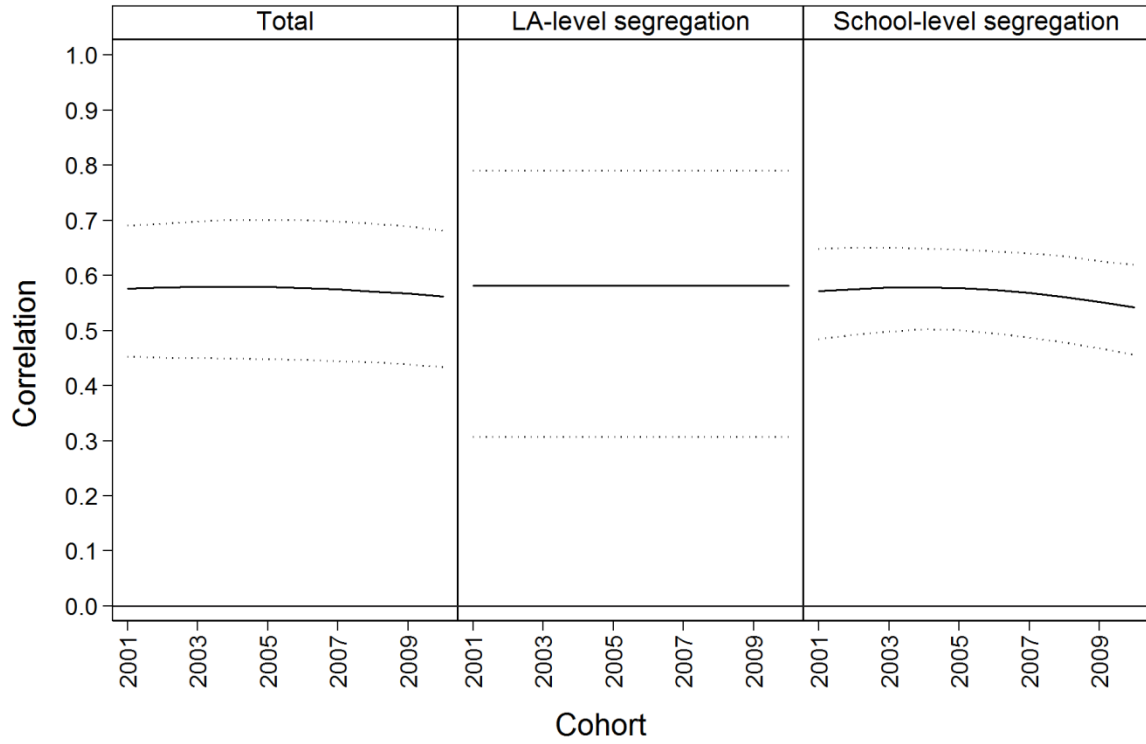
Fig. 5. Model 3 total, LA- and school-level black-white-Asian-white correlation functions. Posterior means with 95% credible intervals.
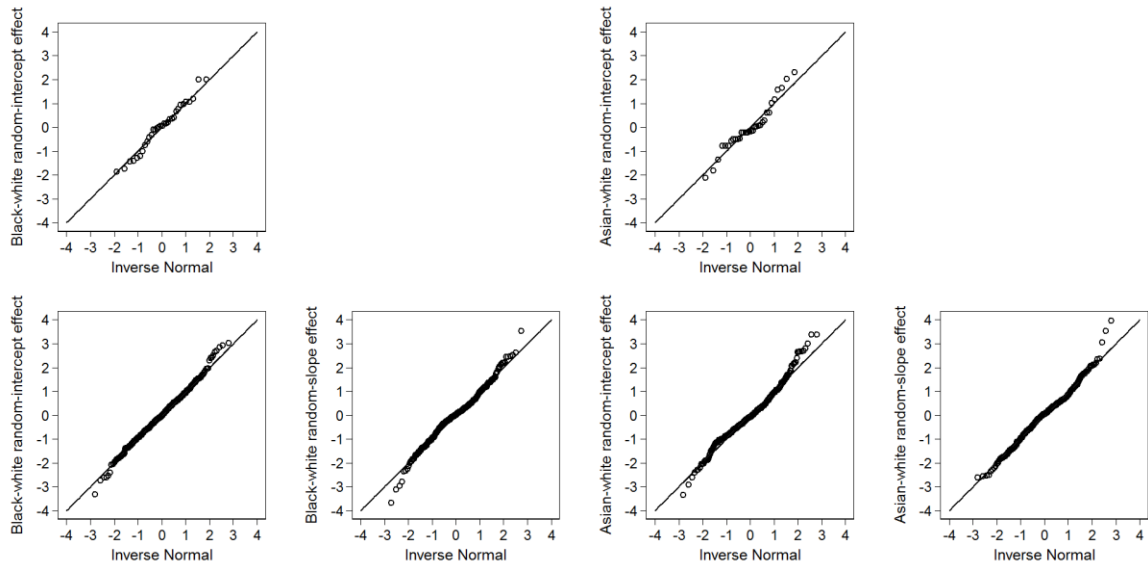
Fig. 6. Model 3 quantile-quantile plots of the LA-level (top panel) and school-level (bottom-panel) residuals
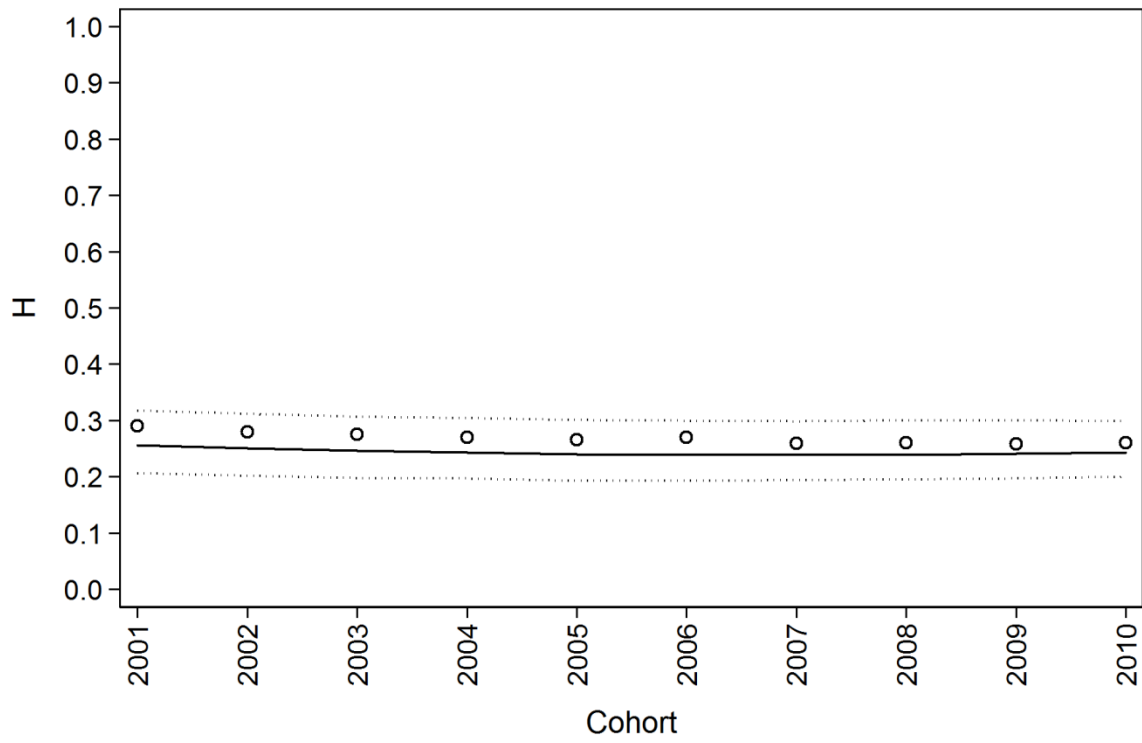
Fig. 7. Model 3 total overall *H*-index function plotted against cohort. Posterior mean plotted with 95% credible intervals. *H* based on observed proportions plotted as points.