X vs. Y: An Analysis of Intergenerational Differences in Transport Mode Use Among Young Adults

2 March 2019

Rico Krueger
Research Centre for Integrated Transport Innovation, School of Civil and Environmental Engineering, UNSW Australia, Sydney NSW 2052, Australia
r.krueger@student.unsw.edu.au
ORCID ID: 0000-0002-5372-741X

Taha H. Rashidi (corresponding author)
Research Centre for Integrated Transport Innovation, School of Civil and Environmental Engineering, UNSW Australia, Sydney NSW 2052, Australia
rashidi@unsw.edu.au

Akshay Vij
Institute for Choice, University of South Australia
140 Arthur Street, North Sydney NSW 2060, Australia
vij.akshay@gmail.com
ORCID ID: 0000-0001-7888-4609

Abstract

Recent research has contrasted the travel patterns of young adults of Generation Y (or, synonymously, the Millennial Generation) with the travel patterns of earlier generations of young adults such as those belonging to Generation X. Young adults of Generation Y are found to drive less and, in some contexts, are found to exhibit more multimodal travel patterns and to use public transit more often. Potential causes for these observed shifts in transport mode use have also been theorised: One view is that period effects in the form of contemporaneous changes in socio-cultural, socio-economic and socio-technical factors are responsible for the observed shifts in transport mode use; another view is that members of Generation Y have inherently different preferences and values due to formative socio-cultural, socio-economic and historical experiences. Motivated by this yet-to-be-resolved dialectic, this paper uses a hierarchical Bayesian multivariate Poisson log-normal model to examine intergenerational differences in transport mode use among young adults. The model is applied to 23 waves of the German Mobility Panel and captures between-cohort and between-period variation of parameters of interest. The trained model informs a counterfactual prediction exercise aiming to decompose intergenerational differences in transport mode use into demography-, cohort-, and period-specific effects. Our findings suggest that all three sets of effects account for intergenerational differences in transport mode use, while the absolute and relative importance of each set of effects vary across transport modes. For the period from 1998 to 2016, two thirds of the decline in car use can be ascribed to period effects; nearly all of the increase in public transit use and 42% of the increase in bicycling can be ascribed to cohort effects.

Keywords: intergenerational differences, Generation X, Generation Y, Millennials, transport mode use, hierarchical Bayesian multivariate Poisson log-normal model, multi-level model

1. Introduction

In recent years, the analysis of intergenerational differences in travel behaviour has attracted considerable attention. In particular, recent research has focussed on contrasting the travel patterns of the current generation of young adults, who are considered to belong to Generation Y (or, synonymously, the Millennial Generation), which comprises the birth cohorts from the mid-1980s to the early 2000s, with the travel patterns of earlier generations of young adults such as those who are considered to belong to Generation X, which comprises the birth cohorts from the mid-1960s to the early 1980s (e.g. Hjorthol, 2016; Klein and Smart, 2017; Kuhnimhof et al., 2011, 2012a,b; McDonald, 2015; Thigpen and Handy, 2018).

The literature widely consents that the day-to-day travel behaviour of young adults of Generation Y differs from that of previous generations of young adults: In comparison with their equally-aged counterparts of prior generations, young adults of Generation Y are observed to drive less (Kuhnimhof et al., 2011, 2012a; McDonald, 2015); and in some contexts, they are found to exhibit more multimodal travel patterns and to use public transport more frequently (Brown et al., 2016; Grimsrud and El-Geneidy, 2014; Kuhnimhof et al., 2011, 2012b). Moreover, members of Generation Y are observed to delay driving licence acquisition and to be less likely to own a car (Hjorthol, 2016; Delbosc, 2017; Delbosc and Currie, 2013; Klein and Smart, 2017; Kuhnimhof et al., 2011).

Two contrasting explanations for these observed intergenerational differences in young adults' travel behaviour have emerged in the literature (also see Delbosc and Ralph, 2017; McDonald, 2015): One view is that period effects in the form of contemporaneous changes in socio-cultural, socio-economic and socio-technical factors account for the observed intergenerational differences in travel behaviour. In particular, the observed decline in car use has been attributed to lower disposable incomes as well as increased vehicle ownership and running costs (Bastian et al., 2016; Klein and Smart, 2017; Ralph, 2015). In addition, it has been hypothesised that young adults of Generation Y substitute the use of information and communication technologies for physical travel (Kroesen and Handy, 2015; Polzin et al., 2014). The second, contrasting view is that members of Generation Y have inherently different preferences, attitudes and values, which manifest in less car-dependent lifestyle choices. For example, young adults of Generation Y are observed to prefer to live in dense, urban, walkable neighbourhoods, where personal mobility need not be predicated on private car use (Myers, 2016). Moreover, young adults of Generation Y are observed to delay entering into life stages that are typically associated with increased car use such as living independently, marriage and parenthood (Garikapati et al., 2016).

Motivated by this yet-to-be-resolved dialectic, the current paper employs a hierarchical Bayesian multivariate Poisson log-normal model to analyse intergenerational differences in transport mode use among young adults of Generations X and Y. The model is applied to 23 waves of the German Mobility Panel and allows for a decomposition of the variation of parameters into cohort- and period-specific terms. The trained model informs a counterfactual prediction exercise aiming to decompose intergenerational differences in transport mode use frequencies into demographic as well as cohort- and period-specific effects.[1]

[1] We note that counterfactual predictions are widely used in applied economics and the quantitative social sciences to decompose group differences into different explanatory factors (see e.g. Fortin et al., 2011). In fact, the studies by McDonald (2015) and Vij et al. (2017) also rely on counterfactual predictions to explain intergenerational differences in transport mode use.

We organise the remainder of this paper as follows: We review pertinent literature and describe the data used for our analysis. We outline the modelling approach and present the estimation results. To evaluate the predictive ability of the trained model, we perform posterior predictive checks. Then we carry out the counterfactual prediction exercise and finally, we conclude.

2. Literature review

Initial research on intergenerational differences in young adults' travel behaviour primarily relied on descriptive analyses of a variety of mobility indicators such as mode-specific travel distances and trip rates, licence-holding and car availability (e.g. Delbosc and Currie, 2013; Kuhnimhof et al., 2011, 2012a,b). Descriptive analyses of mobility indicators are pivotal in identifying trends in travel behaviour, yet they do not afford profound insights into the behavioural processes underlying the observed changes in young adults' travel behaviour. For this reason, a growing body of literature employs disaggregate approaches to explain observed changes in young adults' travel behaviour (Brown et al., 2016; Delbosc and Currie, 2014; Hjorthol, 2016; Klein and Smart, 2017; McDonald, 2015; Thigpen and Handy, 2018; Vij et al., 2017).

Most disaggregate studies employ some form of period-cohort analysis to isolate the effects of periods and cohorts on the examined dependent variables (Brown et al., 2016; Delbosc and Currie, 2014; Hjorthol, 2016; Klein and Smart, 2017; McDonald, 2015; Thigpen and Handy, 2018; Vij et al., 2017). Periods are the points in time, at which measurements are taken; cohorts are groups of individuals who were born in the same exogenously-defined time period. Period effects refer to external changes in socio-cultural, socio-economic and socio-technical factors and are experienced by all individuals, who are observed at the time the period effects in question occur (Yang, 2008). Cohort effects, on the other hand, pertain to intergenerational differences in values, beliefs, attitudes and preferences and are the consequence of unique socio-cultural, socio-economic and historical experiences shared among individuals belonging to the same cohort (Ryder, 1965). Period and cohort effects may be mediated by demographic effects, i.e. changes in lifestyles across cohorts (McDonald, 2015). For example, the decision to take up residence in the urban core rather than in a suburb may be reflective of cohort-specific preferences or period-specific constraints. Thus, demographic effects can be viewed as indirect effects absorbing period and cohort effects that manifest in shifts in observable lifestyle orientations across cohorts.

Given the observed intergenerational differences in young adults' travel behaviour, the question arises as to what extent these changes are associated with either cohort or period effects: Do young adults of Generation Y have inherently different travel preferences? Or, are the observed changes in young adults travel behaviour reflective of broader societal shifts transcending generational cohorts? By and large, evidence on the absolute and relative importance of cohort and period effects in explaining the observed intergenerational differences in young adults' travel behaviour is mixed. The findings reported in the literature differ across behaviours, contexts and methods.

Two studies relying on household travel surveys conducted in Melbourne, Australia and respectively, Norway suggest that declining licensing rates among young adults can be attributed to shifts in underlying lifestyle preferences (Delbosc and Currie, 2014; Hjorthol, 2016). In a retrospective study relying on data collected from a convenience sample, Thigpen and Handy (2018) find that driving licence acquisition remains subject to unobserved intergenerational heterogeneity, when systematic differences in parental, attitudinal, social and regulatory variables are accounted for; the authors are, however, unable to decompose the unobserved heterogeneity into cohort- and period-specific terms. Using data sourced from the United States Panel Study of Income Dynamics, Klein and Smart (2017) examine intergenerational differences in car ownership levels and find that declining car ownership levels among members of Generation Y can be attributed to period effects in the form of economic constraints.

Analysing data sourced from two cross-sections of the United States National Household Travel Survey, Brown et al. (2016) find that higher levels of public transit use among young adults are mostly linked to life-cycle, demographic and locational factors rather than to other period- or cohort-specific effects. Using three cross-sections of the United States National Household Travel Survey, McDonald (2015) estimates a linear regression model to explain young adults' daily travel distances by car. In a counterfactual prediction exercise asking what the level of car use would have been, had demographic, cohort and period effects not been present, the author demonstrates that the decline in demand for car travel can be decomposed into three distinct sources, namely lifestyle-related demographic shifts that cut across all cohorts, other period-specific effects and cohort-specific effects. Lifestyle-related demographic shifts are found to account for 10% to 25% of the decline in car travel; cohort effects are shown to account for 35% to 50% of the decline; and period effects account for the remaining 40% of the decline. Finally, drawing from two cross-sectional household travel surveys of the San Francisco Bay Area, United States, Vij et al. (2017) develop a latent class mode choice model to explain observed changes in transport mode use in terms of underlying preference shifts. By means of counterfactual prediction, the authors show that preference shifts have occurred across all considered birth cohorts and have not been limited to one particular generational cohort. The authors conclude that the observed overall reduction in car dependency in the considered study area can be attributed to broader socio-cultural changes rather than to cohort-specific effects.

As a whole, the literature does not give a clear indication as to what extent either cohort or period effects account for the observed intergenerational differences in young adults' travel behaviour. While one study attributes the observed intergenerational differences to changes in socio-demographic factors rather than to cohort- or period-specific effects (Brown et al., 2016), another study suggests that young adults of Generation Y respond to period-specific effects in the form of economic constraints (Klein and Smart, 2017). Other studies underline the possibility of a cultural shift that may or may not transcend generational cohorts (Delbosc and Currie, 2014; Hjorthol, 2016; Thigpen and Handy, 2018; Vij et al., 2017). Finally, one study takes an intermediate position by demonstrating that the observed intergenerational differences can be linked to both cohort and period effects as well as lifestyle-related demographic shifts (McDonald, 2015). Another observation is that disaggregate studies of intergenerational differences in young adults' travel behaviour have focused on different aspects of travel behaviour. A comparatively large set of studies investigates intergenerational differences in licence-holding and car ownership (Delbosc and Currie, 2014; Hjorthol, 2016; Klein and Smart, 2017; Thigpen and Handy, 2018). Two studies focus on either public transit use or car use (Brown et al., 2016; McDonald, 2015). One study of intergenerational differences in transport mode use is not limited to any one particular mode (Vij et al., 2017).

In that vein, the current paper aims to advance the literature's understanding of the absolute and relative importance of period and cohort effects in explaining intergenerational differences in transport mode use among young adults. Our analysis draws from a rich data source containing information about the day-to-day travel behaviours of young adults of Generations X and Y in Germany during a 23-year period. Amongst other industrialised countries, Germany has been witness to changing travel patterns among young adults (Kuhnimhof et al., 2011, 2012a,b). A flexible hierarchical Bayesian multivariate Poisson log-normal model is used to jointly characterise daily trip rates by car, public transit, bicycle and walking.

3. Data

3.1. Primary data source

The primary data source for our analysis is the German Mobility Panel (MOP), a longitudinal household-based travel survey intended to capture long-term trends in the day-to-day travel behaviour of the German population. The MOP has been conducted annually since 1994 and is designed as a rotating panel survey to compensate for attrition and fatigue. Households regularly stay on the panel for three consecutive years before being replaced by new households. The collected data include information about socio-economic attributes of participating households and their members as well as a trip diary, in which participating household members record details about all trips undertaken during a week-long period in autumn. Zumkeller and Chlond (2009) describe the design of the MOP in more detail.

3.2. Sample composition

In the present paper, we consider 23 MOP waves covering the period from 1994 to 2016. Given our interest in the travel behaviour of young adults, we restrict our analysis to records from 1,846 individuals aged between 20 and 29 years old at the time of data collection. We define six five-year birth cohorts ranging from 1965 to 1996, whereby the sixth cohort includes seven years, as the size of a potential seventh cohort would be too small, when five-year brackets are used to define birth cohorts. We note that defining five-year birth cohorts is a convention in demography (e.g. Yang, 2008). The resulting sample has a cross-classified structure where the data for each observation period consist of varying proportions of records from individuals from different birth cohorts (see Figure 1).

Each individual record is associated with a maximum of three trip diaries, each of which covers a maximum of seven consecutive days. We exclude weekends from the analysis so that a maximum of five diary days remain in each trip diary. For model validation, the sample is partitioned into training and hold-out samples: We construct the training sample by randomly selecting a maximum of three diary days from each trip diary, and in cases where a trip diary includes fewer than three diary days, all reported diary days are assigned to the training sample. Any remaining diary days are assigned to the hold-out sample. The constructed training and hold-out samples include 9,001 and, respectively, 5,721 diary days.

We observe that the sample sizes of individuals for each cohort and period are relatively small (see Figure 1). In addition, it has been noted that self-selection and attrition are disproportionately high among young adults (Eisenmann et al., 2018). For this reason, the considered data may not give a perfect representation of the population of young adults in Germany at all time periods, even if sampling weights were applied. There is, however, no apparent reason as to why any self-selection and attrition effects should be non-orthogonal to either cohort or period effects. Therefore, we argue that any limitations pertaining to the limited representativity of the sample are offset by the richness and consistency of the considered microdata.

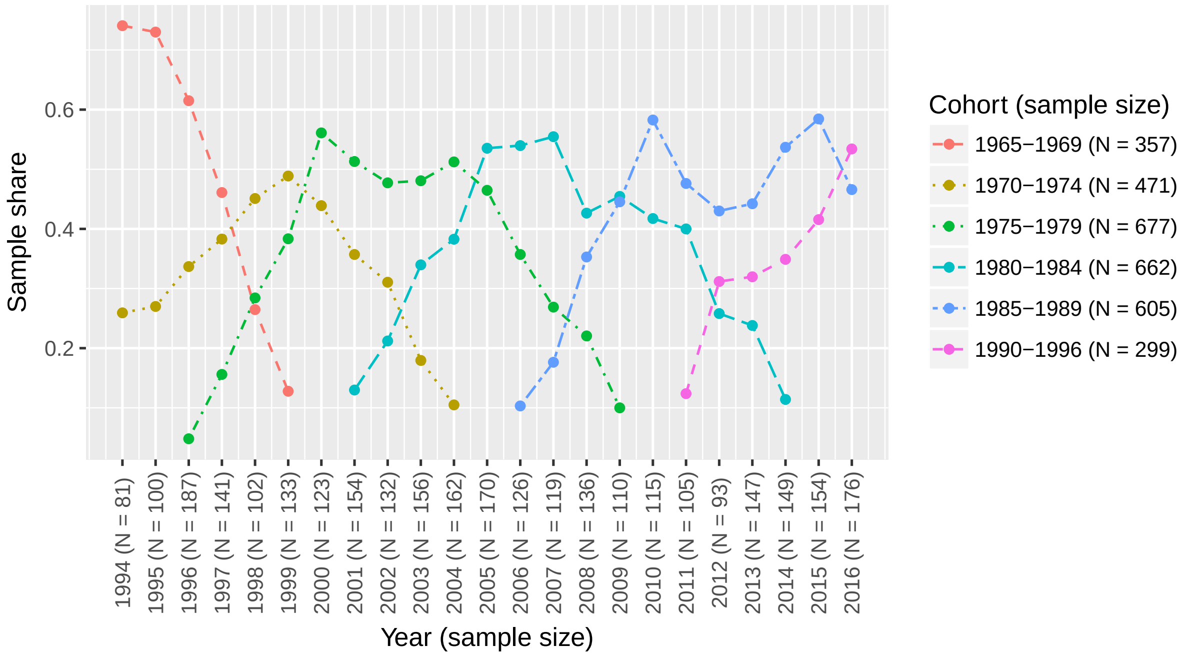

Figure 1: **Sample composition by observation period and birth cohort (zero sample shares are not shown)**

3.3. Dependent variable

The dependent variable considered in our analysis is a four-variate vector of category counts corresponding to daily trip frequencies by the four transport modes car, public transit, bicycling and walking. In the current paper, we examine daily mode use frequencies rather than another variable such as daily travel distances by mode, as we are interested in explaining intergenerational differences in mode-specific demand for mobility rather than in the mode-specific consumption of transportation infrastructure and services.[2] The daily mode use frequencies are derived from the trip diary data of the MOP; business-related trips are excluded from the analysis. We acknowledge that tour-based approaches can provide benefits over trip-based approaches (e.g. Krizek, 2003), but for consistency with the original design of the MOP and with published studies relying on MOP data, we adhere to a trip-based analysis approach.

Figure 2 shows the evolution of the mean number of trips by different transport modes and by cohorts from 1994 to 2016 for the training sample; an additional category represents the mean number of trips by all transport modes combined. In the considered observation period, the overall daily trip rates decreased substantially from an average of 4.4 daily trips in 1998 to an average of 3.3 of daily trips in 2016. Moreover, the mean number of daily trips by car declined sharply from an average of 3.0 trips in 1998 to an average of 1.7 trips in 2016. In turn, the daily trip rates by public transportation increased considerably from an average of 0.1 daily trips to an average of 0.6 trips

[2]Clearly, considering other mobility indicators such as transport mode choices or mode-specific travel distances is equally valid and may likely offer different perspectives on the problem studied.

in 2016. Daily trip rates for bicycling and walking remained relatively stable during the considered observation period. By and large, these observations are consistent with the literature (Kuhnimhof et al., 2011, 2012b) and indicate an overall decrease in trip rates in combination with a modal shift away from automobile-dependent travel to increasing public transport use. Figure 2 also suggests substantial inter-cohort differences in transport mode use. For example, the mean daily trip rates by public transportation are generally greater for individuals born in or after 1980. However, Figure 2 also suggests that the transport mode use of specific cohorts is not decoupled from the overall time trend. For the cohort of young adults born between 1975 and 1979, the mean number of daily trips by car decreased from an average of 3.3 daily trips in 1998 to an average of 2.0 daily trips in 2005; in the same period, the overall number of daily trips by car also declined from an average of 3.0 daily trips to an average of 2.2 daily trips.

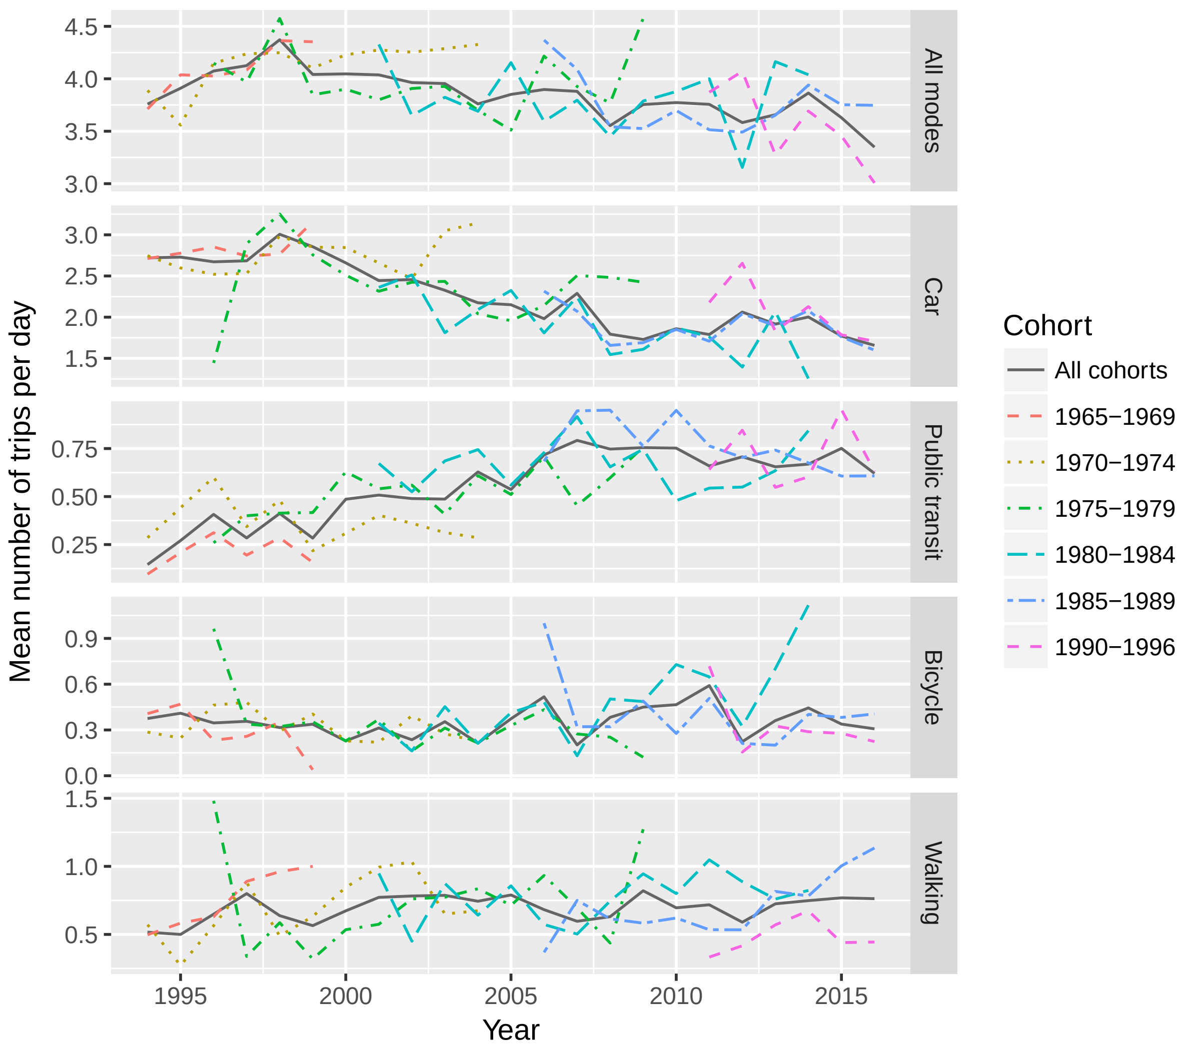

Figure 2: **Mean number of trips by transport mode and cohort from 1994 to 2016 (training sample only)**

3.4. Explanatory variables

For our analysis, we consider a variety of explanatory variables pertaining to the socio-economic attributes and the long-term travel decisions (household car ownership and licence-holding) of the studied individuals and their household. We restrict our analysis to variables that have been consistently surveyed in all 23 considered MOP waves. Notably, income information has not been surveyed until 2002 and is therefore not considered in our analysis.

To control for changes in mobility costs, we supplement the MOP data with information about the price levels of petrol and public transportation services in each period. Information on petrol prices is sourced from International Energy Agency (2016) and adjusted for inflation by dividing the nominal prices for each year by the consumer price index values in the respective years. The consumer price index for passenger transport services is used as a proxy for price levels of public transportation services and is sourced from Federal Statistical Office Germany (2018); the price index for passenger transport services is adjusted for inflation by dividing the nominal index values of each year by the general consumer price index values in the respective years. The general consumer price index assumes 2010 as reference year and is sourced from Federal Statistical Office Germany (2018).

Figure 3 shows the evolution of the means of all considered explanatory variables by cohort from 1994 to 2016. The variable "education level 1" indicates whether an individual has completed secondary school without having obtained a university-entrance qualification; the variable "education level 2" indicates whether an individual has completed secondary education and has also obtained a university-entrance qualification. The variable "currently in education" denotes whether an individual is currently attending an educational institution or is completing vocational training. The variable "household with young children" denotes whether children under the age of ten years live in the individual's household. The variable "household in urban location" captures accessibility information.[3] For the most part, the means of the considered explanatory variables are relatively stable during the observation period. Notable exceptions are the means of the variables "education level 1" and "education level 2", which decrease and, respectively, increase substantially over periods and across cohorts. The trend towards higher educational attainment is mirrored in the evolution of the means of other lifestyle-related explanatory variables: The means of the variables "employed (full-time)" and "currently in education" also decrease and, respectively, increase considerably over periods and across cohorts. Altogether, Figure 3 suggests a profound shift in observable lifestyle orientations across cohorts: On average, individuals belonging to later birth cohorts exhibit comparatively higher rates of participation in education, attain higher levels of education and therefore enter the workforce at a comparatively greater age. These observed changes in demographic factors are consistent with the literature (see for example Brown et al., 2016; Chatterjee et al., 2018; Kuhnimhof et al., 2012b; Garikapati et al., 2016; McDonald, 2015, and the literature referenced therein).

[3]The MOP data contain two kinds of accessibility information: First, the MOP data include information about the settlement structure of the subnational entity, in which the respondent resides. This information is provided for different subnational entities (regions, district, municipalities) in the form of scales with five to seventeen categories. Second, the MOP data contain self-reported information about the ease of accessibility of different forms of public transportation and other places of interest (shopping centres, cinemas, gyms etc.) from the respondent's residence. This second type of accessibility information comes with the known limitations of self-reported data. The variable "household with young children" is derived from a nine-point scale characterising the settlement structure of the municipality, in which the respondent resides, and subsumes categories nos. 1, 2, 3, 5, and 6 of this scale. During the development of the hierarchical Bayesian multivariate Poisson log-normal model (see Section 5.1), we explored incorporating more refined measures of accessibility, using both types of accessibility information, but did not obtain better outcomes in terms of model tractability and fit. In the interest of parsimony, the variable "household in urban location" is therefore the only variable capturing accessibility information in the final model specification presented in Section 5.

Figure 3: **Means of independent variables by cohort from 1994 to 2016**

4. Modelling approach

4.1. Overview and background

Motivated by the cross-classified multi-level structure of the MOP data, we devise a hierarchical Bayesian multivariate Poisson log-normal model for the analysis of daily trip frequencies by transport mode. The multivariate Poisson log-normal model (Chib and Winkelmann, 2001) is a flexible multivariate count data model, which accommodates repeated observations for the same analytical unit, unrestricted correlation across categories of counts and overdispersion in the marginal counts. By contrast, other well-established multivariate count data models are less flexible (see e.g. Zhang et al., 2017): The multinomial model does not account for overdispersion and can only capture negative correlation in the response data; the negative multinomial model can account for overdispersion, but only allows for positive correlation in the response data.

The term hierarchical Bayesian refers to hierarchical models that are estimated using Bayesian methods. Hierarchical models (Gelman and Hill, 2006; Gelman et al., 2013) appreciate the multi-level structure inherent to many empirical datasets. For example, in a typical household travel survey, observations are nested within individuals and individuals are nested within households. In a hierarchical model, parameters are allowed to vary across the units of a level and the parameters are given their own models to infer the distribution of the parameters at each level. A special case of a hierarchical model is the cross-classified multi-level model, which recognises that a unit can belong to multiple clusters at a time. The MOP data are an instance of a dataset with a cross-classified multi-level structure: Observations are nested within individuals and individuals are nested within cohorts; however, observations are also nested within periods. The application of a hierarchical model that reflects the cross-classified multi-level structure of the MOP data induces correlation between observations within nests and allows for the decomposition of the total variance of parameters of interest into level-specific terms.

The complexity of many hierarchical models precludes the application of classical inference methods. The Bayesian approach, however, lends itself well to the estimation of hierarchical models, as the unit-level parameters can be viewed as unknown quantities whose posterior distribution is learnt by combining prior information about the distribution of the latent quantities with the observed data (see e.g. Gelman et al., 2013, for a general treatment).

In the following subsection, we present the basic formulation of the multivariate Poisson log-normal model, before we extend the basic model in Section 4.3 to account for the cross-classified multi-level structure of the MOP data.

4.2. Multivariate Poisson log-normal model

The multivariate Poisson log-normal model (Chib and Winkelmann, 2001) is established as follows: For each individual $n \in \{n, \ldots, N\}$ in the sample, R_n independent measurements indexed by $r \in \{1, \ldots, R_n\}$ are taken. Each measurement is a J-variate vector $\boldsymbol{y}_{n,r} = \left\{y_{n,r,j}\right\}_{j=1}^{J}$ of counts on J categories indexed by $j \in \{1, \ldots, J\}$. Furthermore, $\boldsymbol{\xi}_n = \left\{\xi_{n,j}\right\}_{j=1}^{J}$ is a collection of individual- and category-specific random effects drawn from a multivariate normal distribution with zero mean and unrestricted covariance matrix $\boldsymbol{\Sigma}$, i.e.

$$\boldsymbol{\xi}_n \sim \mathrm{N}(\mathbf{0}, \boldsymbol{\Sigma}), \quad n = 1, \ldots, N. \tag{1}$$

The random effects $\boldsymbol{\xi}_n$ play an integral role in the multivariate Poisson log-normal model: They induce correlation between measurements for the same individual and allow for correlation across categories. Moreover, the random effects account for overdispersion when $\Sigma_{j,j} > 0$, where $\Sigma_{l,k}$ denotes element (l, k) of $\boldsymbol{\Sigma}$.

Conditional on category-specific parameters $\boldsymbol{\beta}_j$ and random effect $\xi_{n,j}$, the probability distribution of count $y_{n,r,j}$ pertaining to category j, individual n and measurement r is Poisson, i.e.

$$P(y_{n,r,j}|\boldsymbol{\beta}_j, \xi_{n,j}, \boldsymbol{X}_{n,r,j}) = \frac{\left(\lambda_{n,r,j}\right)^{y_{n,r,j}} \exp\left(-\lambda_{n,r,j}\right)}{y_{n,r,j}!}, \tag{2}$$

with

$$\ln \lambda_{n,r,j} = \boldsymbol{X}_{n,r,j}\boldsymbol{\beta}_j + \xi_{n,j}, \tag{3}$$

where $\boldsymbol{X}_{n,r,j}$ is a row-vector of covariates including a constant. The conditional probability of the sequence of measurements for individual n is obtained by iterating (2) over categories and measurements:

$$P(\boldsymbol{y}_n|\boldsymbol{\beta}_j, \xi_n, \boldsymbol{X}_n) = \prod_{r=1}^{R_n}\prod_{j=1}^{J} P(y_{n,r,j}|\boldsymbol{\beta}_j, \xi_{n,j}, \boldsymbol{X}_{n,r,j}), \tag{4}$$

where $\boldsymbol{y}_n = \left\{\boldsymbol{y}_{n,r}\right\}_{r=1}^{R_n}$ and $\boldsymbol{X}_n = \left\{\left\{\boldsymbol{X}_{n,r,j}\right\}_{j=1}^{J}\right\}_{r=1}^{R_n}$. The unconditional probability is obtained by integrating out the unobserved random effects $\boldsymbol{\xi}_n$ from (4):

$$P(\boldsymbol{y}_n|\boldsymbol{\beta}_j, \boldsymbol{\Sigma}, \boldsymbol{X}_n) = \int P(\boldsymbol{y}_n|\boldsymbol{\beta}_j, \boldsymbol{\xi}_n, \boldsymbol{X}_n) f(\boldsymbol{\xi}_n|\boldsymbol{\Sigma}) d\boldsymbol{\xi}_n. \tag{5}$$

Iterating (5) over all individuals in the sample yields the likelihood of the multivariate Poisson log-normal model:

$$\mathscr{L}(\boldsymbol{y}|\boldsymbol{\beta}, \boldsymbol{\Sigma}, \boldsymbol{X}) = \prod_{n=1}^{N} P(\boldsymbol{y}_n|\boldsymbol{\beta}_j, \boldsymbol{\Sigma}, \boldsymbol{X}_n). \tag{6}$$

We note that the multivariate Poisson log-normal model receives its name from the fact that (5) can be viewed as a multivariate log-normal mixture of Poissons, because $\left\{\lambda_{n,r,j}\right\}_{j=1}^{J} \sim \mathrm{N}\left(\left\{\boldsymbol{X}_{n,r,j}\boldsymbol{\beta}_j\right\}_{j=1}^{J}, \boldsymbol{\Sigma}\right)$ such that $\left\{\ln \lambda_{n,r,j}\right\}_{j=1}^{J}$ is multivariate log-normally distributed (Aitchison and Ho, 1989; Chib and Winkelmann, 2001).

To assess the instantaneous effect of an explanatory variable on the predicted count, marginal effects can be calculated. In the multivariate Poisson log-normal, marginal effects are given by

$$\int \frac{\partial \lambda_{n,r,j}}{\partial \boldsymbol{X}_{n,r,j}} f(\xi_{n,j}|\Sigma_{j,j}) d\xi_{n,j}. \tag{7}$$

This integration cannot be expressed in closed form and therefore, marginal effects need to be simulated or can be obtained as a by-product of posterior sampling.

Finally, we highlight that the multivariate Poisson log-normal model is an appealing model to characterise transport mode use frequencies, as the conditional distribution of counts in the multivariate Poisson log-normal model can be represented as a two-stage process where an individual first decides on a budget for consumption and subsequently divides the budget between categories (Gopalan et al., 2013). In lieu of

$$y_{n,r,j}|\xi_{n,j} \sim \text{Poisson}\left(\lambda_{n,r,j}\right), \quad j = 1, \ldots, J, \tag{8}$$

which is a re-expression of (2), it is possible to write

$$z_{n,r}|\boldsymbol{\xi}_n \sim \text{Poisson}\left(\sum_{j=1}^{J} \lambda_{n,r,j}\right), \tag{9}$$

$$\boldsymbol{y}_{n,r}|\boldsymbol{\xi}_n \sim \text{Multinomial}\left(z_{n,r}, \left\{\frac{\lambda_{n,r,j}}{\sum_{j'=1}^{J} \lambda_{n,r,j'}}\right\}_{j=1}^{J}\right), \tag{10}$$

where $z_{n,r} \equiv \sum_{j=1}^{J} y_{n,r,j}$. However, to add a caveat, we note that the multivariate Poisson log-normal model gives a probabilistic representation of total and category-specific demand, but it is not a utility-consistent model per se (Bhat et al., 2015).

4.3. Extension

The basic formulation of the multivariate Poisson log-normal model assumes that the regression parameters $\boldsymbol{\beta}_j$ are invariant across analytical and observational units. To accommodate unobserved heterogeneity in the regression parameters $\boldsymbol{\beta}_j$, we devise a hierarchical formulation that reflects the cross-classified structure of the MOP data. To be specific, the generic category-specific parameters $\boldsymbol{\beta}_j$ are replaced by category- and measurement-specific parameters $\boldsymbol{\beta}_{n,r,j}$. We let

$$\boldsymbol{\beta}_{n,r,j} = \boldsymbol{\beta}_{0,j} + \boldsymbol{v}_{c(n),j} + \boldsymbol{w}_{t(n,r),j}, \tag{11}$$

where $\boldsymbol{\beta}_{0,j}$ is a fixed parameter representing the mean of $\boldsymbol{\beta}_{n,r,j}$. $\boldsymbol{v}_{c(n),j}$ and $\boldsymbol{w}_{t(n,r),j}$ are random parameters representing cohort- and, respectively, period-specific perturbations around the mean. Here, we use $c \in \{1, \dots, C\}$ to index cohorts and $t \in \{1, \dots, T\}$ to index periods. $c(n)$ is a mapping indicating individual n's cohort such that $c(n) = c$ if n belongs to cohort c. Similarly, $t(n,r)$ provides a mapping from measurement r for individual n to the period in which the measurement was taken. $\boldsymbol{v}_{c,j}$ and $\boldsymbol{w}_{t,j}$ are both realisations from normal distributions with zero means and covariance matrices whose off-diagonal elements are restricted to zero, i.e.

$$\boldsymbol{v}_{c,j} \sim \mathrm{N}\left(\mathbf{0}, \mathrm{diag}\left(\boldsymbol{\zeta}_j\right)\right), \quad c = 1, \dots, C, \quad j = 1, \dots, J \tag{12}$$

$$\boldsymbol{w}_{t,j} \sim \mathrm{N}\left(\mathbf{0}, \mathrm{diag}\left(\boldsymbol{\eta}_j\right)\right), \quad t = 1, \dots, T, \quad j = 1, \dots, J, \tag{13}$$

with $\boldsymbol{\zeta}_j$ and $\boldsymbol{\eta}_j$ denoting vectors of variances. Hence, the distribution of $\boldsymbol{\beta}_{n,r,j}$ is

$$\boldsymbol{\beta}_{n,r,j} \sim \mathrm{N}\left(\boldsymbol{\beta}_{0,j}, \mathrm{diag}\left(\boldsymbol{\zeta}_j\right) + \mathrm{diag}\left(\boldsymbol{\eta}_j\right)\right), \tag{14}$$

such that the total variation of $\boldsymbol{\beta}_{n,r,j}$ is decomposed into within-cohort variance $\boldsymbol{\zeta}_j$ and within-period variance $\boldsymbol{\eta}_j$.

4.4. Model inference

The Bayesian inference approach requires us to specify prior distributions for all parameters. For fixed parameters, we employ $\mathrm{N}(0,5)$ as a weakly-informative prior. The choice of a prior for scale parameters in a hierarchical model requires careful consideration, especially when the number of groups is small (Gelman et al., 2013). In the present application, we follow Chung et al. (2013) and employ $\mathrm{Gamma}(2, 0.1)$ as a weakly-informative prior for all scale parameters. Since $\mathrm{Gamma}(2, 0.1)$ is not consistent with zero, the choice of this prior reflects the belief that the scale parameters in question are non-zero, which in turn concurs with our motivation to employ a hierarchical model in the current empirical context. Following the examples of Gelman (2006) and Polson and Scott (2012), we also explored the use of half-Cauchy priors. The half-Cauchy distribution has appreciable mass near the origin and thus favours small values, but its heavy tail also allows for large estimates if the likelihood provides support for such estimates. In the present application, we found that the Markov Chains of the posterior draws of some scale parameters exhibited poor mixing, when a half-Cauchy prior was used, whereas all chains converged nicely, when the $\mathrm{Gamma}(2, 0.1)$ prior was employed. For numerical reasons, we do not directly estimate the covariance matrix $\boldsymbol{\Sigma}$, but rather a scale vector $\boldsymbol{\sigma}$ and the Cholesky factor $\boldsymbol{L}$ of the correlation matrix $\boldsymbol{\Omega}$. To this end, we exploit the relationship $\Sigma = \boldsymbol{D}\boldsymbol{\Omega}\boldsymbol{D}$ with $\boldsymbol{D} = \mathrm{diag}(\boldsymbol{\sigma})$ and $\boldsymbol{\Omega} = \boldsymbol{L}\boldsymbol{L}'$. We let $\boldsymbol{\sigma} \sim \mathrm{Cauchy}(0, 5)$ and a suitable prior for $\boldsymbol{L}$ is the LKJ-Cholesky distribution (Lewandowski et al., 2009) with scale four.

We implement the hierarchical Bayesian multivariate Poisson log-normal model in Stan (Carpenter et al., 2016), a probabilistic programming language, which interfaces the No-U-Turn sampler (Hoffman and Gelman, 2014). For the estimation of all models presented in this paper, the sampler is executed with four parallel Markov chains and 4,000 iterations for each chain, whereby the initial 2,000 iterations of each chain are discarded for burn-in. Convergence is assessed by considering the Gelman-Rubin diagnostic values (Gelman and Rubin, 1992) and by visually inspecting the trace plots of the posterior draws of selected parameters.

5. Results

5.1. Final model specification

The final model specification presented in this section is the product of an extensive specification search and represents a constrained implementation of the general modelling approach outlined in the previous section. To be specific, the final model specification does not allow for unobserved heterogeneity in all regression parameters and treats selected regression parameters as fixed. The decision to treat a regression parameters as either fixed or random is motivated by both theoretical and empirical considerations. In the cases of covariates that pertain to price levels, car availability and education, we do not hypothesise that the sensitivities to these variables vary across either cohorts or periods all else unchanged. In the cases of the covariates indicating licence-holding and part-time employment, we found that a lack of variation in sensitivities with respect to these variables caused the sampler to slow down and the respective Markov chains exhibited poor mixing, when the respective regression parameters were treated as random; consequently, the regression parameters in question are treated as fixed parameters in the final model specification.

Furthermore, we note that age is not considered as an explanatory variable in our analysis because of the well-known age-period-cohort identification problem. Any model aiming to jointly identify age, period and cohort effects is not identifiable, because the identity age = period − cohort implies perfect collinearity between age, period and cohort effects. Mechanical solutions to the age-period-cohort identification problem are generally ineffective and therefore, constraints must be imposed on one set of effects. In the present application, we assume that age effects are fixed to zero. We argue that for the stratum of young adults aged between 20 and 29 years old, economic factors and lifestyle factors have more prominent effects on travel behaviour than any effects that are immediately associated with an individual's age. For more details on the issues associated with the joint modelling of age, period and cohort effects, we refer to the literature (Bell and Jones, 2013, 2014).

5.2. Overall model evaluation and model selection

To quantify the benefits of accommodating unobserved heterogeneity between both cohorts and periods, we allow the parameters to vary one level at a time and also estimate a basic null model, which only accounts for unobserved inter-individual differences, but not for unobserved between-cohort and between-period differences. The null model is therefore equivalent to a standard multivariate Poisson log-normal model (see Section 4.2). As shown in Table 1, we contrast the performance of the model variants by considering in-sample and predictive log-likelihood values as well as the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC). Moreover, we apply likelihood-ratio tests to assess whether the goodness of fit of the hierarchical Bayesian multivariate Poisson log-normal model with both cohort and period random effects is statistically better than the goodness of fit of each of the competing models. Note that the competing models are nested within the hierarchical Bayesian multivariate Poisson log-normal model with both cohort and period random effects, as the competing models can be obtained by constraining some of the scale parameters to zero.

As seen in Table 1, the hierarchical Bayesian multivariate Poisson log-normal model with both cohort and period random effects outperforms the competing models in terms of in-sample and predictive log-likelihood values; the likelihood-ratio tests indicate that the goodness of fit of the hierarchical Bayesian multivariate Poisson log-normal model with both cohort and period random effects is statistically better than the goodness of fit of each of the competing models. We also observe that the hierarchical Bayesian multivariate Poisson log-normal model with only period random effects yields better in-sample and predictive log-likelihood values than the hierarchical Bayesian multivariate Poisson log-normal model with only cohort random effects; a possible explanation for these differences is that the model with only period effects allows for greater distributional flexibility than the model with only cohort effects, as there are 23 periods and only six exogenously-defined birth cohorts in the considered data.

The information criteria give divergent indications of which model variant is superior. The Akaike Information Criterion suggests that the model with both cohort and period random effects is the best model. However, the Bayesian Information Criterion, which favours parsimony, indicates that the model with only period effects should be preferred. We acknowledge that the model with only period random effects may provide a more parsimonious representation of the unobserved heterogeneity of the data. Yet, the model with both cohort and period random effects performs better in terms of in-sample fit and out-of-sample predictive ability. For these reasons, our subsequent discussion and analysis focus on the more complex model variant.

	Null	**CRE**	**PRE**	**CRE + PRE**
No. of parameters	50	74	74	98
Log-likelihood				
In-sample	−31,649.7	−31,553.2	−31,171.2	−31,116.9
Hold-out	−21,976.4	−21,933.1	−21,833.1	−21,817.2
AIC	63,399.4	63,254.4	62,490.4	62,429.8
BIC	63,754.7	63,780.2	63,016.2	63,126.1
Likelihood-ratio test w.r.t. CRE + PRE				
χ^2	1065.6	872.6	108.6	
df	48	24	24	
p	< 0.001	< 0.001	< 0.001	

Note: CRE = cohort random effects; PRE = period random effects.

Table 1: **Comparison of model specifications with different levels of parameter variation**

5.3. Parameter estimates and marginal effects

The estimation results for the hierarchical Bayesian multivariate Poisson log-normal model with cohort and period random effects are given in Table 2. The reported estimation results include the posterior means of the fixed regression parameters as well as of the location and scale hyper-parameters of the random regression parameters. Based on the posterior means of the scale hyper-parameters, the total variation of the random regression parameters is decomposed into individual-, cohort-, and period-specific terms. To assess the instantaneous effect of the considered explanatory variables on mode use frequencies, marginal effects can be calculated. As can be seen from (7), marginal effects are computed for each observational unit. The rightmost column of Table 2 gives the mean marginal effects for the entire sample. In addition, Figure 4 shows the mean marginal effects by cohort. In the subsequent paragraphs, we describe noteworthy estimation results for each category of counts.

First, we consider the estimation results for the sub-model characterising daily car use frequencies. We observe that for all random regression parameters, with the exception of the constant, the between-cohort variation is greater than the between-period variation. Next, we consider the estimates of the mean marginal effects. Interestingly, the overall mean marginal effect of the variable "female" is positive and the evolution of the corresponding cohort-specific means is comparatively stable. However, the mean marginal effect of the variable "education level 2" changes quite substantially across cohorts. While the mean marginal effect of this covariate is approximately zero for birth cohort 1965–1969, the mean marginal effect is approximately −0.25 for birth cohort 1990–1996. Interestingly, the mean marginal effects for the variables "driving licence" and "no. of cars in household" decrease substantially across the considered birth cohorts. This implies that on average, given the same mobility resources and all else being equal, individuals belonging to later birth cohorts undertake fewer daily trips by car compared to individuals from earlier birth cohorts.

Second, we consider the estimation results for the sub-model characterising daily trips by public transportation. Again, we observe that for all random regression parameters, with the exception of the constant, the between-cohort variation is greater than the between-period variation. The mean marginal effects of most explanatory variables are relatively stable across the examined birth cohorts. A noteworthy exception is the mean marginal effect of the variable "household in urban location". For birth cohort 1965–1969, the mean marginal effect is approximately 0.25, but individuals born in or after 1980, the mean marginal effect is approximately 0.75. This implies that on average, given the same residential location and all else being equal, individuals born in or after 1980 undertake more daily trips by public transportation than individuals born before 1980.

Third, we consider the estimation results for the sub-model characterising daily trips by bicycle. Again, we observe that for most random regression parameters, the between-cohort variation is greater than the between-period variation. Noteworthy exceptions are the random regression parameters pertaining to the intercept and the explanatory variable "currently in education". For the most part, the mean marginal effects are relatively stable across the considered birth cohorts. A noteworthy exception is the mean marginal effect of the variable "education level 2". For birth cohort 1965–1969, the mean marginal effect of this variable is approximately 0.25, but for individuals born in or after 1980, the mean marginal effect is approximately 0.50. This implies that on average, given the same level of education and all else being equal, individuals born in or after 1980 undertake more daily trips by bicycle than individuals born before 1980.

Fourth, we consider the estimation results for the sub-model characterising daily walking trips. We observe that for all random regression parameters, with the exception of the constant, the between-cohort variation is greater than the between-period variation. All mean marginal effects are relatively stable across the considered birth cohorts.

Finally, Figure 5 visualises the estimated correlation structure across daily mode use frequencies. On the whole, the estimated correlation parameters are consistent with common mode use patterns: For example, counts of daily trips by car are negatively correlated with counts of daily trips by each of the remaining modes; likewise, counts of daily trips by public transit and counts of daily walking trips are weakly and positively correlated. Moreover, we stress that both positive and negative correlation parameters are estimated, which in turn underlines the benefits of the multivariate Poisson log-normal model, which is able to capture unrestricted correlation patterns in the response data.

Covariate	Parameter estimates: Location	Scales: IRE	CRE	PRE	PoTV due to...: BCV	BPV	MME
Car							
Constant	−0.0469	0.6860*	0.1930*	0.0820*	0.0724	0.0131	
Female	0.0694		0.0986*	0.0726*	0.6483	0.3517	0.1519
Employed (full-time)	0.1209		0.1761*	0.0827*	0.8191	0.1809	0.2650
Currently in education	0.0367		0.1079*	0.1035*	0.5212	0.4788	0.0632
Education level 1	0.1277		0.1591*	0.0552*	0.8925	0.1075	0.2894
Education level 2	−0.0338		0.2058*	0.0905*	0.8379	0.1621	−0.0743
Employed (part-time)	0.1177*						0.2631
Household with young children	0.1203*						0.2688
Household in urban location	−0.1963*						−0.4386
Petrol price [2010 EUR/litre]	−0.3808						−0.8509
Driving licence	0.7061*						1.5780
No. of cars in household	0.2198*						0.4911
Public transit							
Constant	−3.7126*	1.7387*	0.6010*	0.1312*	0.1062	0.0051	
Female	0.1390		0.4057*	0.1904*	0.8196	0.1804	0.1093
Employed (full-time)	−0.1855		0.4512*	0.2910*	0.7063	0.2937	−0.0620
Currently in education	0.4968*		0.2928*	0.2318*	0.6148	0.3852	0.2718
Education level 1	−0.6380*		0.4649*	0.1810*	0.8683	0.1317	−0.3353
Education level 2	−0.0689		0.5055*	0.2233*	0.8367	0.1633	−0.0385
Employed (part-time)	0.2590*						0.1448
Household with young children	−0.4670*						−0.2611
Household in urban location	1.1213*						0.6269
PT price index	0.5872						0.3282
Bicycling							
Constant	−5.5075*	2.8790*	1.0766*	0.5647*	0.1187	0.0327	
Female	−0.0911		0.9536*	0.4730*	0.8025	0.1975	−0.0459
Employed (full-time)	0.0861		0.7345*	0.6598*	0.5535	0.4465	0.0275
Currently in education	0.2668		0.5900*	0.8186*	0.3419	0.6581	0.0753
Education level 1	0.7190		1.2543*	0.5079*	0.8591	0.1409	0.2987
Education level 2	1.1992*		0.9139*	0.5131*	0.7603	0.2397	0.4413
Employed (part-time)	0.5839*						0.2047
Household with young children	−0.1447						−0.0507
Household in urban location	0.3428*						0.1202
Walking							
Constant	−0.4855*	1.3117*	0.4020*	0.0744*	0.0856	0.0029	
Female	0.1669		0.4003*	0.1510*	0.8755	0.1245	0.1257
Employed (full-time)	−0.5362*		0.1703*	0.0995*	0.7455	0.2545	−0.3850
Currently in education	−0.4206*		0.2552*	0.2304*	0.5510	0.4490	−0.2850
Education level 1	−1.0123*		0.3554*	0.1280*	0.8853	0.1147	−0.6971
Education level 2	−0.4934*		0.3393*	0.1376*	0.8588	0.1412	−0.3379
Employed (part-time)	−0.2880*						−0.2028
Household with young children	0.5734*						0.4038
Household in urban location	0.2213*						0.1558

Note: The reported parameter estimates are posterior means. * at least 95% of the posterior mass exclude zero. IRE = individual random effect; CRE = cohort random effect; PRE = period random effect; PoTV = proportion of total variation; BCV = between-cohort variation; BPV = between-period variation; MME = mean marginal effect.

Table 2: **Estimation results and mean marginal effects for the hierarchical Bayesian multivariate Poisson log-normal model with cohort and period random effects**

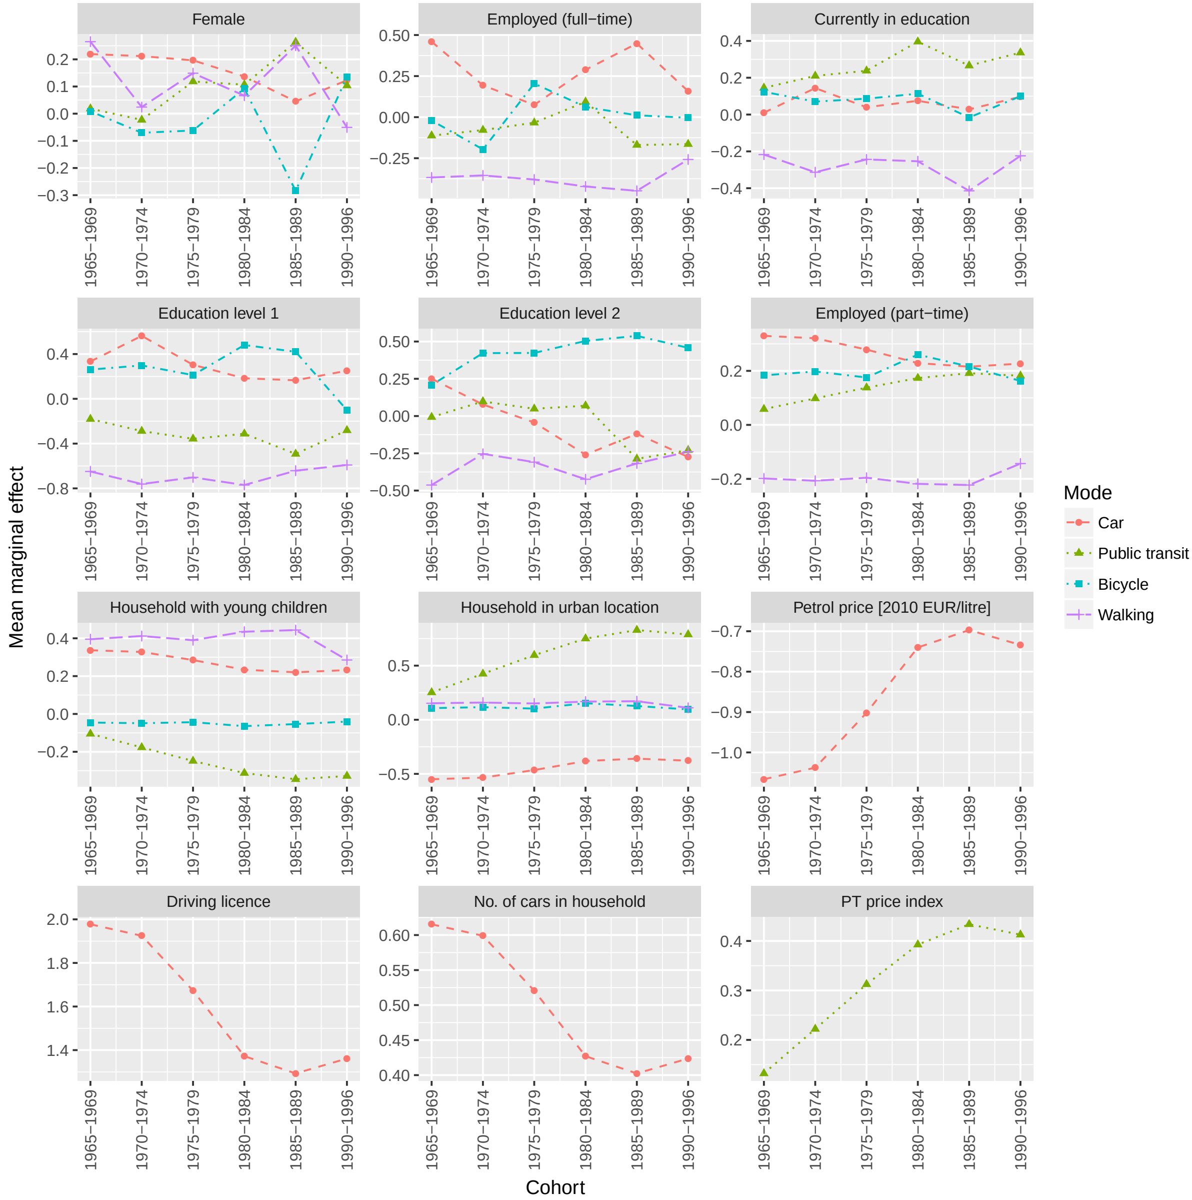

Figure 4: **Mean marginal effects by cohort for the hierarchical Bayesian multivariate Poisson log-normal model with cohort and period random effects**

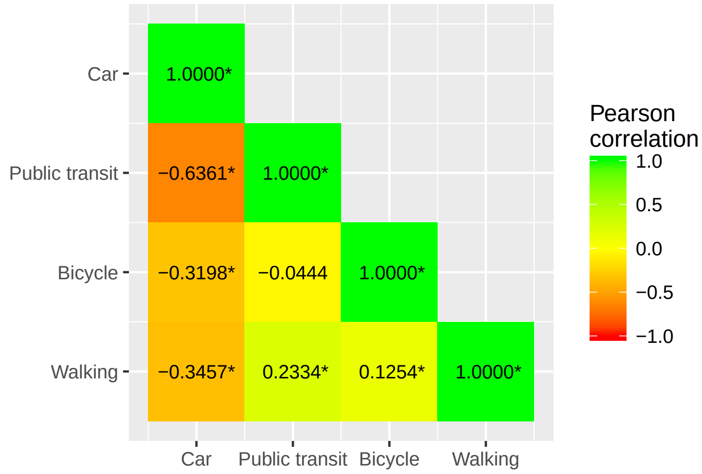

Figure 5: **Estimated correlation between daily mode use frequencies for the hierarchical Bayesian multivariate Poisson log-normal model with cohort and period random effects (the reported values are posterior means; * at least 95% of the posterior mass exclude zero)**

6. Posterior predictive checking

To assess the predictive ability of the trained hierarchical Bayesian multivariate Poisson log-normal model, we perform posterior predictive checks (Gelman et al., 2013, 1996). In a posterior predictive check, model predictions are compared to observed data or to quantities derived from observed data. In the present application, posterior predictive checks are performed on the hold-out sample and cumulative trip counts by mode and period are considered as test quantities. The observed test quantities are given by

$$z_{t,j} \equiv \sum_{n=1}^{N} \sum_{r=1}^{R_n} [t(n,r) = t] \cdot y_{n,r,j}, \quad t = 1, \dots T, j = 1, \dots J, \tag{15}$$

where $[a = b]$ denotes the Kronecker delta, which gives one if $a = b$ is true and zero otherwise. We argue that these test quantities have pivotal policy-relevance, as cumulative trip counts by mode inform modal split calculations.

First, we conduct a graphical posterior predictive check, in which the posterior distribution of the test quantities for unseen data is compared with the actual values of the test quantities for the same unseen data. As seen in Figure 6, the observed values closely overlap with the posterior distributions of the corresponding predicted values. In the case of car and public transit, the observed values are always located within the 90% central credible intervals of the corresponding predicted values. By contrast, in the cases of bicycling and walking, few observed values are located outside of the 90% central credible intervals of the corresponding predicted values.

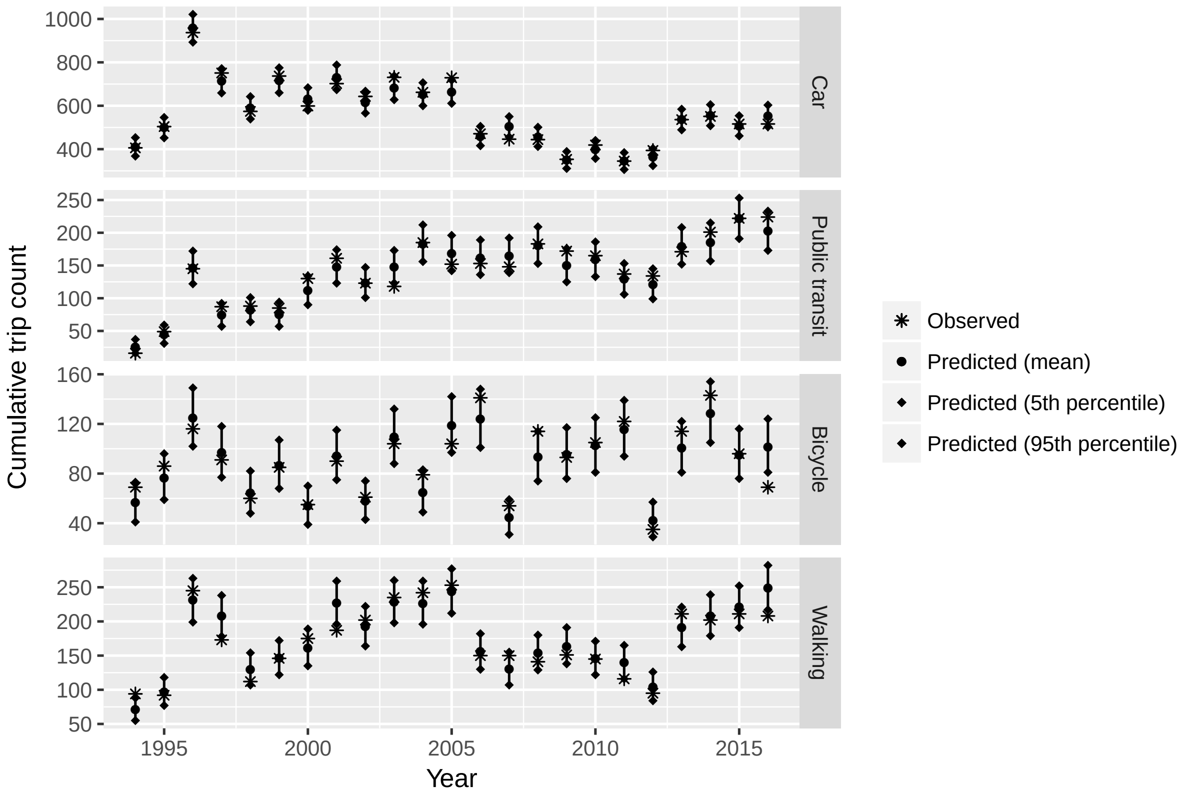

Figure 6: **Posterior predictive check on the cumulative trip counts by mode and period for the hold-out sample**

To quantify the overlap between the predicted and the observed test quantities, we perform a numerical posterior predictive check. As a test statistic, we consider Pearson's (1900) χ^2-discrepancy:

$$\chi^2(\boldsymbol{z}, \boldsymbol{\theta}) = \sum_{t=1}^{T} \sum_{j=1}^{J} \frac{\left(z_{t,j} - \mathbf{E}(z_{t,j}|\boldsymbol{\theta})\right)^2}{\mathbf{E}(z_{t,j}|\boldsymbol{\theta})}, \tag{16}$$

where $\mathbf{E}(z_{t,j}|\boldsymbol{\theta}) = \sum_{n=1}^{N} \sum_{r=1}^{R_n} [t(n,r) = t] \cdot \mathbf{E}(y_{n,r,j}|\boldsymbol{\theta}) = \sum_{n=1}^{N} \sum_{r=1}^{R_n} [t(n,r) = t] \cdot \lambda_{n,r,j}$ and $\boldsymbol{\theta}$ denotes the collection of all model parameters. The test statistic $\chi^2(\boldsymbol{z}, \boldsymbol{\theta})$ represents the normalised sum of squared deviations between the observed and predicted cells of a contingency table giving trip frequencies by mode and period. $\chi^2(\boldsymbol{z}, \boldsymbol{\theta})$ can be assumed to be approximately χ^2_{df}-distributed with $df = (T-1) \cdot (J-1) = (23-1) \cdot (4-1) = 66$ degrees of freedom. As both the test statistic and the tail-area probability of χ^2_{df} depend on the model parameters $\boldsymbol{\theta}$, the posterior predictive p-value (Meng, 1994) is given by:

$$p_b(\boldsymbol{z}) = \int P\left(\chi^2_{df} \geq \chi^2(\boldsymbol{z}, \boldsymbol{\theta})\right) P(\boldsymbol{\theta}|\boldsymbol{y}, \boldsymbol{X}) d\boldsymbol{\theta}. \tag{17}$$

This posterior predictive distribution is easily obtained as a by-product of posterior sampling. In the present application, the posterior mean of $p_b(\boldsymbol{z})$ is equal to 1.000. Hence, we cannot reject the null hypothesis that the observed cumulative mode use frequencies are the same as the predicted ones.

To conclude, graphical and numerical posterior predictive checks show that the trained hierarchical Bayesian multivariate Poisson log-normal model performs very well at predicting the considered test statistic for external data. However, to add caveat, we acknowledge that the hold-out data for the posterior predictive check come from the same set of individuals whose other observations were used to train the model. Ideally, posterior predictive checks would be performed on external data from a new set of individuals. Yet, in the present application, the relatively small sample sizes of individuals for each cohort and period preclude such an approach.

7. Counterfactual analysis

7.1. Overview

Having established that the hierarchical Bayesian multivariate Poisson log-normal model with cohort and period random effects fares well at prediction, we follow the example of McDonald (2015) and perform a counterfactual prediction exercise to decompose intergenerational differences in daily trip frequencies by different modes into three sources of change, namely demography- as well as period- and cohort-specific effects.

In principle, period effects refer to external changes in socio-cultural, socio-economic and socio-technical factors and are experienced by all individuals, who are observed at the time the period effects in question occur (Yang, 2008). By contrast, cohort effects pertain to inter-generational differences in values, beliefs, attitudes and preferences and are the consequences of unique socio-cultural, socio-economic and historical experiences shared among individuals from the same birth cohort (Ryder, 1965). Demographic effects refer to changes in lifestyles across cohorts and mediate cohort and period effects (McDonald, 2015). For example, differences in education levels across cohorts may be reflective of both cohort and period effects: On the one hand, different generational cohorts may attach varying levels of importance to higher education; on the other hand, changes in the labour market as well as in the cost and quality of higher education over time may affect young adults' participation in higher education.

Here, we argue that demographic effects manifest in differences in transport mode use across cohorts due to observable differences in demographic characteristics across cohorts, whereas period effects manifest in differences in transport mode use over time due to changes in macro-economic factors, i.e. the price levels of petrol and public transportation services, as well as in unobserved behavioural factors, i.e. the regression parameters, across periods. Cohort effects manifest in differences in transport mode use across cohorts due to changes in unobserved behavioural factors, i.e. the regression parameters, across cohorts.

7.2. Method

The counterfactual prediction exercise is implemented as follows: To account for changes in demographic factors across Generations X and Y, we randomly select two subsamples of 250 individuals each from the total sample. The first random subsample is composed of young adults born before 1980 and therefore reflects the typical demographic characteristics of Generation X; analogously, the second random subsample is composed of young adults born in or after 1980 and therefore reflects the typical demographic characteristics of Generation Y. For each of these two subsamples, we predict the mean number of daily trips by different transport modes in different time periods and for different birth cohorts. Recall that the regression parameters in the hierarchical Bayesian multivariate Poisson log-normal model are cohort- and period-specific (see Section 4.3). In the counterfactual prediction exercise, we configure the regression parameters such that each individual from each random subsample is effectively pseudo-observed in each time period and in each birth cohort. Technical details concerning the computation of the counterfactual predictions are provided in Appendix A.1.

With the exception of the explanatory variables pertaining to the price levels of petrol and public transportation services, the values of the explanatory variables remain unchanged in each prediction for each subsample. For the price levels of petrol and public transportation services, the values of the period in question are used to systematically account for fluctuations in these variables over time. As a consequence, the reported period effects are inclusive of fluctuations in the price levels of petrol and public transportation services. We highlight that the counterfactual analysis does not explicitly account for changes in other factors that may change over periods and cohorts such as differences in disposable incomes or changes in the transport system, because this information is not available to us. Recall that the model includes constants, which capture the average effect of all excluded factors. Since the constants are allowed to vary across cohorts and periods, the average effect of all excluded factors can be absorbed into both cohort- and period-specific effects. However, excluded factors may also capture changes in demographic characteristics across cohorts such that the measurement of demographic effects is also affected. We re-iterate that demographic effects cannot be entirely separated from period and cohort effects, as differences in demographic characteristics across cohorts can be manifestations of both period-specific constraints and cohort-specific preferences.

The counterfactual predictions inform a Blinder-Oaxaca decomposition for non-linear models (Bauer and Sinning, 2008; Blinder, 1973; Oaxaca, 1973) in order to systematically decompose intergenerational differences in daily trip frequencies by different modes into demography- as well as period- and cohort-specific terms. In principle, a Blinder-Oaxaca decomposition can be carried out for any combination of target and reference cohorts and periods, but for the sake of brevity, we focus on differences between the two larger generational cohorts, when examining cohort differences, and on differences between the years 1998 and 2016, when examining period effects. We note that 1998 is the year in which the greatest mean number of trips by car are observed (see Section 3.3), while 2016 is the most recent survey year available for our analysis.

The Blinder-Oaxaca decomposition is implemented as follows: We treat the prediction results for subsample 1, year 1998, cohort 1965–1979 as a baseline, i.e. we ask what the mean predicted daily trip frequencies would be for a sample of individuals with demographic characteristics reflective of those born between 1965 and 1979, if the sample was observed in 1998 and if the sensitivities to the explanatory variables were the same as for individuals born between 1965 and 1979. The baseline is then compared against three counterfactual scenarios: In the first scenario (DE), we ask what the mean predicted daily trip frequencies would be, if the sample was replaced by another sample of individuals with demographic characteristics reflective of those born between 1980 and 1996, while the sample was still observed in 1998 and while the sensitivities to the explanatory variables were still the same as for individuals born between 1965 and 1979. In the second scenario (DE + PE), we ask what the mean predicted trip frequencies would be, if the sample from the first scenario was observed in 2016, while the sensitivities to the explanatory variables were still the same as for individuals born between 1965 and 1979. Finally, the third scenario (DE + PE + CE) asks what the mean predicted daily trip frequencies would be for a sample of individuals with demographic characteristics reflective of those born in or after 1980, if the sample was observed in 2016 and if the sensitivities to the explanatory variables were the same as for individuals born in or after 1980. Stated succinctly, the first scenario (DE) asks what the mean predicted daily trip frequencies would be, if only demographic effects were at play; the second scenario (DE + PE) additionally accounts for period effects; and the third scenario (DE + PE + CE) factors in all three—demographic, period and cohort—effects. Technical details concerning the implementation of this specific Blinder-Oaxaca decomposition are provided in Appendix A.2.

7.3. Results

Figure 7 visualises the mean predicted number of daily trips by mode, cohort, period and subsample. We report the mean predicted number of daily trips for each of the six exogenously-defined birth cohorts as well as for two additional, larger generational cohorts comprising the years from 1965 to 1979 and from 1980 to 1996 respectively; the two larger generational cohorts therefore broadly correspond to Generations X and Y.[4] Moreover, Table 3 enumerates the results for the Blinder-Oaxaca decomposition. Based on Figure 7 and Table 3, we discuss noteworthy prediction results for each transport mode.

First, we consider the prediction results for the mean daily trip frequencies by car. Figure 7 suggests that the decline in the mean predicted number of daily trips by car from 1998 to 2016 can be largely attributed to period effects and to smaller extents to demographic and cohort effects. The results given in Table 3 confirm this observation. When all three—demographic, period and cohort—effects are factored in, the decrease in the mean predicted number of daily trips by car amounts to 1.50 relative to the base case. Two thirds of the total decrease can be ascribed to period effects, while the remainder of the decrease can be attributed in roughly equal parts to demographic and cohort effects.

Second, we consider the prediction results for the mean daily trip frequencies by public transit. Figure 7 suggests that the increase in the mean number of daily trips by public transit from 1998 to 2016 can be largely ascribed to cohort effects. The results given in Table 3 confirm this observation. Relative to the baseline, the increase in the mean predicted number of daily trips by public transit is 0.14, when all three effects are factored in. Cohort effects explain the greatest share of this increase. The contribution of cohort effects to the total increase is 0.17 daily trips, but the total change is slightly less than 0.17 daily trips due to the presence of negative period effects.

Third, we consider the prediction results for the mean number of daily trips by bicycle. As can be seen in Table 3, the increase in the mean predicted number of daily trips by bicycle is 0.52 relative to the base case, when all three effects are factored in. Cohort effects account for 42% of the total increase, while demographic and period effects each account for 29% of the total increase.

Furthermore, we only find a marginal change in the mean predicted number of daily walking trips from 1998 to 2016 relative to the baseline, when all three effects are factored in, as period effects are neutralised by negative cohort effects.

[4] Note that the cohort effects reported for the two larger generational cohorts, which broadly correspond to Generations X and Y, are unaffected by the decision to consider five-year birth cohorts (see Section 3.2), as it is not critical whether the aggregation occurs over many small or over few large birth cohorts.

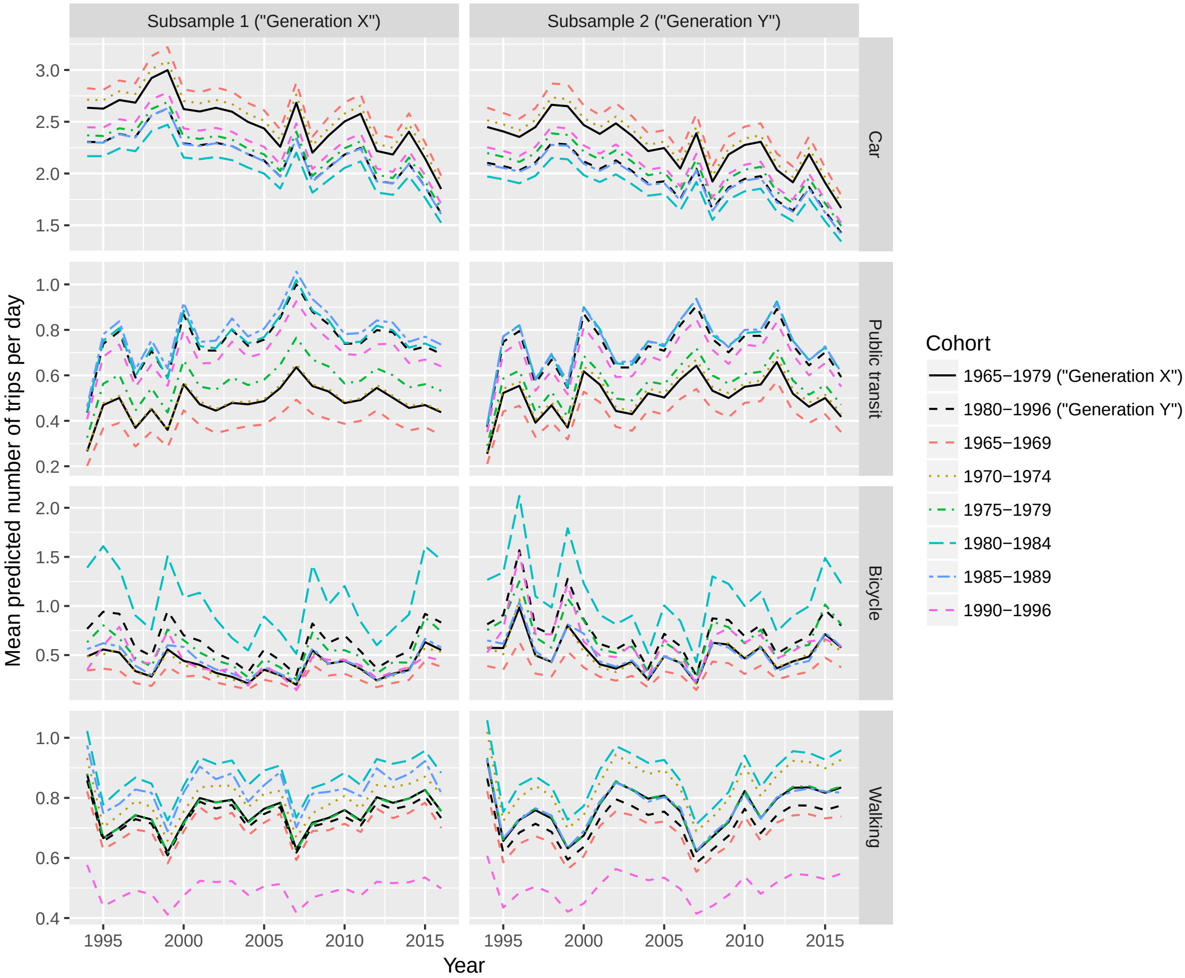

Figure 7: **Mean predicted number of daily trips by mode, cohort, period and subsample**

		Counterfactual scenarios						
	Baseline	**DE**		**DE + PE**		**DE + PE + CE**		
Mode	$\bar{z}_0$	$\bar{z}_1$	$\bar{z}_1 - \bar{z}_0$	$\bar{z}_2$	$\bar{z}_2 - \bar{z}_1$	$\bar{z}_3$	$\bar{z}_3 - \bar{z}_2$	$\bar{z}_3 - \bar{z}_0$
Car	2.92	2.66	−0.26	1.67	−1.00	1.43	−0.24	−1.50
Public transit	0.45	0.47	0.02	0.42	−0.05	0.59	0.17	0.14
Bicycle	0.28	0.43	0.15	0.58	0.15	0.80	0.22	0.52
Walking	0.73	0.73	0.00	0.83	0.10	0.77	−0.06	0.05

Note:
Baseline (subsample 1, year 1998, cohort 1965–1979);
DE = demographic effects (subsample 2, year 1998, cohort 1965–1979);
DE + PE = demographic + period effects combined (subsample 2, year 2016, cohort 1965–1979);
DE + PE + CE = demographic + period + cohort effects (subsample 2, year 2016, cohort 1980–1996);
$\bar{z}_i$ = mean predicted number of daily trips in scenario i

Table 3: **Detailed results for the counterfactual analysis**

8. Conclusions

In this paper, we employ a hierarchical Bayesian multivariate Poisson log-normal model to analyse intergenerational differences in transport mode use among young adults of Generations X and Y. The model is applied to 23 waves of the German Mobility and recognise the cross-classified multi-level structure of the considered data by allowing parameters of interest to vary across periods and cohorts. The trained model informs a counterfactual prediction exercise aiming to decompose changes in daily trip rates by car, public transit, bicycle and walking into demographic as well as into cohort- and period-specific effects.

Our study adds to a body of literature reasoning about the absolute and relative importance of period and cohort effects in explaining intergenerational differences in travel behaviour among young adults of Generation Y and prior generations of young adults. While one view is that period effects in the form of contemporaneous changes in socio-cultural, socio-economic and socio-technical factors are responsible for the observed shifts in transport mode use, another, contrasting view is that the observed changes in travel behaviour can be ascribed to cohort effects, i.e. due to formative socio-cultural, socio-economic and historical experiences, members of Generation Y have inherently different values and preferences, which manifest in travel choices that differ from those of previous generations of young adults.

The results of our counterfactual prediction exercise suggest that three sets of effects, i.e. demographic as well as cohort- and period-specific effects, can explain intergenerational differences in transport mode use among young adults of Generations X and Y, while the absolute and relative importance of each set of effects vary across transport modes. In the case of private car use, we find that the decline in the mean predicted number of daily trips from 1998 to 2016 can be predominantly ascribed to period effects, while the absolute contribution of demographic and cohort-specific effects is also appreciable. In the case of public transit use, the increase in the mean predicted number of daily trips from 1998 to 2016 can be largely ascribed to cohort effects, while the role of other effects is negligible. In the case of bicycling, the increase in the mean predicted number of daily trips from 1998 to 2016 can be attributed to all three sets of effects, while cohort effects account for the greatest proportion of the total change. In the case of walking, the total change in the mean predicted number of daily trips from 1998 to 2016 is marginal and period effects are neutralised by cohort effects.

The findings of this study are subject to the following caveats: The availability and inclusion of control variables may affect the measurement of observed and latent period and cohort effects. This is because the model includes constants, which capture the average effect of all excluded factors. Since the constants are allowed to vary across cohorts and periods, the average effect of excluded factors can be absorbed into both period and cohort effects. However, excluded factors may also capture changes in demographic characteristics across cohorts so that the measurement of demographic effects is also affected. Therefore, studies relying on different data sources and model specifications may reach different conclusions about the absolute and relative importance of demography-, cohort- and period-specific effects in explaining intergenerational differences in transport mode use, and the findings of this current study should not be overgeneralised. Moreover, it is known that the Blinder-Oaxaca decomposition, which is used to separate demography- as well as period- and cohort-specific effects, is sensitive to the selection of the reference groups (Oaxaca and Ransom, 1999). We have also highlighted that is difficult to entirely separate demographic effects from period and cohort effects, as differences in demographic characteristics across cohorts can be reflective of both period-specific constraints and cohort-specific preferences. More research is thus needed to unfold what mechanisms determine changes in demographic characteristics across cohorts.

Moreover, there are other several directions in which future research may build on the work presented in the current paper. First, the literature widely consents that the travel patterns of the current generation of young adults differ from those of previous generations of young adults. Yet, it has been questioned whether members of the current generation of young adults will continue to embrace comparatively less car-dependent lifestyles, as they mature and reach middle adulthood (e.g. Brown et al., 2016; Garikapati et al., 2016); once sufficient amounts of data become available, the hierarchical Bayesian multivariate Poisson log-normal modelling approach employed in the current paper can be applied to examine intergenerational differences in the daily transport mode use frequencies among members of Generation Y and members of other generations in life stages other than young adulthood. Second, another avenue for future research is to give period and cohort effects a hierarchical structure, which explicitly represents relationships between generic and mode-specific effects, to allow for a decomposition of period and cohort effects into generic and mode-specific terms. Third, our analysis does not afford insights into the concrete socio-economic and behavioural factors underlying the measured demography-, cohort- and period-specific effects—the measured effects are generic in that they do not pertain to any specific socio-economic and behavioural factors. Thus, future research may explore ways to triangulate different data sources, which provide information about changes in travel behaviour and other factors of interest and to incorporate this additional information into the disaggregate model framework. Finally, our analysis highlights that young adults' travel behaviours can be subject to substantial intergenerational heterogeneity. Therefore, it is critical to better understand the mobility needs and preferences of current and future generations of young adults to allow for the design of effective interventions aiming to encourage sustainable travel behaviours throughout the life course, and to allow for an equitable provision and distribution of transportation infrastructure investments.

Acknowledgements

We are grateful to three anonymous reviewers for their critical assessment of our work. Moreover, we would like to thank the Clearing House for Transport Data at the German Aerospace Center (DLR) for providing access to the microdata of the German Mobility Panel. RK and THR acknowledge support from the Australian Research Council (DE170101346).

Conflict of interest

On behalf of all authors, the corresponding author states that there is no conflict of interest.

Author contributions

R Krueger: Conception and design, date preparation and analysis, manuscript writing and editing.
TH Rashidi: Conception and design, manuscript editing, supervision.
A Vij: Conception and design, manuscript editing.

References

Aitchison, J. and Ho, C. H. (1989). The multivariate Poisson-log normal distribution. *Biometrika*, 76(4):643–653.

Bastian, A., Börjesson, M., and Eliasson, J. (2016). Explaining "peak car" with economic variables. *Transportation Research Part A: Policy and Practice*, 88:236–250.

Bauer, T. K. and Sinning, M. (2008). An extension of the blinder–oaxaca decomposition to nonlinear models. *AStA Advances in Statistical Analysis*, 92(2):197–206.

Bell, A. and Jones, K. (2013). The impossibility of separating age, period and cohort effects. *Social Science & Medicine*, 93:163–165.

Bell, A. and Jones, K. (2014). Current practice in the modelling of age, period and cohort effects with panel data: a commentary on Tawfik et al. (2012), Clarke et al. (2009), and McCulloch (2012). *Quality & Quantity*, 48(4):2089–2095.

Bhat, C. R., Paleti, R., and Castro, M. (2015). A New Utility-Consistent Econometric Approach to Multivariate Count Data Modeling. *Journal of Applied Econometrics*, 30(5):806–825.

Blinder, A. S. (1973). Wage discrimination: reduced form and structural estimates. *Journal of Human resources*, pages 436–455.

Brown, A., Blumenberg, E., Taylor, B., Ralph, K., and Voulgaris, C. T. (2016). A Taste for Transit? Analyzing Public Transit Use Trends among Youth. *Journal of Public Transportation*, 19(1).

Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., and Riddel, A. (2016). Stan: A probabilistic programming language. *Journal of Statistical Software*.

Chatterjee, K., Goodwin, P., Schwanen, T., Clark, B., Jain, J., Melia, S., Middleton, J., Plyushteva, A., Ricci, M., Santos, G., and Stokes, G. (2018). Young people's travel – What's changed and why? Review and analysis.

Chib, S. and Winkelmann, R. (2001). Markov Chain Monte Carlo Analysis of Correlated Count Data. *Journal of Business & Economic Statistics*, 19(4):428–435.

Chung, Y., Rabe-Hesketh, S., Dorie, V., Gelman, A., and Liu, J. (2013). A Nondegenerate Penalized Likelihood Estimator for Variance Parameters in Multilevel Models. *Psychometrika*, 78(4):685–709.

Delbosc, A. (2017). Delay or forgo? A closer look at youth driver licensing trends in the United States and Australia. *Transportation*, 44(5):919–926.

Delbosc, A. and Currie, G. (2013). Causes of Youth Licensing Decline: A Synthesis of Evidence. *Transport Reviews*, 33(3):271–290.

Delbosc, A. and Currie, G. (2014). Changing demographics and young adult driver license decline in Melbourne, Australia (1994–2009). *Transportation*, 41(3):529–542.

Delbosc, A. and Ralph, K. (2017). A tale of two millennials. *Journal of Transport and Land Use*, 10(1).

Eisenmann, C., Chlond, B., Hilgert, T., Von Behren, S., and Vortisch, P. (2018). Deutsches Mobilitätspanel (MOP) – Wissenschaftliche Begleitung und Auswertungen Bericht 2016/2017: Alltagsmobilität und Fahrleistung. Technical report.

Federal Statistical Office Germany (2018). GENESIS-Online.

Fortin, N., Lemieux, T., and Firpo, S. (2011). Decomposition methods in economics. In *Handbook of Labor Economics*, volume 4, pages 1–102. Elsevier.

Garikapati, V. M., Pendyala, R. M., Morris, E. A., Mokhtarian, P. L., and McDonald, N. (2016). Activity patterns, time use, and travel of millennials: a generation in transition? *Transport Reviews*, 36(5):558–584.

Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Analysis*, 1(3):515–534.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian Data Analysis, Third Edition*. CRC Press.

Gelman, A. and Hill, J. (2006). *Data analysis using regression and multilevel/hierarchical models*. Cambridge university press.

Gelman, A., Meng, X.-L., and Stern, H. (1996). Posterior Predictive Assessment of Model Fitness via Realized Discrepancies. *Statistica Sinica*, 6(4):733–760.

Gelman, A. and Rubin, D. B. (1992). Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 7(4):457–472.

Gopalan, P., Hofman, J. M., and Blei, D. M. (2013). Scalable Recommendation with Poisson Factorization. *arXiv:1311.1704 [cs, stat]*. arXiv: 1311.1704.

Grimsrud, M. and El-Geneidy, A. (2014). Transit to eternal youth: lifecycle and generational trends in Greater Montreal public transport mode share. *Transportation*, 41(1):1–19.

Hjorthol, R. (2016). Decreasing popularity of the car? Changes in driving licence and access to a car among young adults over a 25-year period in Norway. *Journal of Transport Geography*, 51:140–146.

Hoffman, M. and Gelman, A. (2014). The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1).

International Energy Agency (2016). Energy prices in national currency per toe.

Klein, N. J. and Smart, M. J. (2017). Millennials and car ownership: Less money, fewer cars. *Transport Policy*, 53:20–29.

Krizek, K. J. (2003). Neighborhood services, trip purpose, and tour-based travel. *Transportation*, 30(4):387–410.

Kroesen, M. and Handy, S. L. (2015). Is the Rise of the E-Society Responsible for the Decline in Car Use by Young Adults? *Transportation Research Record: Journal of the Transportation Research Board*, 2496:28–35.

Kuhnimhof, T., Armoogum, J., Buehler, R., Dargay, J., Denstadli, J. M., and Yamamoto, T. (2012a). Men Shape a Downward Trend in Car Use among Young Adults—Evidence from Six Industrialized Countries. *Transport Reviews*, 32(6):761–779.

Kuhnimhof, T., Buehler, R., and Dargay, J. (2011). A New Generation. *Transportation Research Record: Journal of the Transportation Research Board*, 2230:58–67.

Kuhnimhof, T., Buehler, R., Wirtz, M., and Kalinowska, D. (2012b). Travel trends among young adults in Germany: increasing multimodality and declining car use for men. *Journal of Transport Geography*, 24:443–450.

Lewandowski, D., Kurowicka, D., and Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of Multivariate Analysis*, 100(9):1989–2001.

McDonald, N. C. (2015). Are Millennials Really the "Go-Nowhere" Generation? *Journal of the American Planning Association*, 81(2):90–103.

Meng, X.-L. (1994). Posterior Predictive \$p\$-Values. *The Annals of Statistics*, 22(3):1142–1160.

Myers, D. (2016). Peak Millennials: Three Reinforcing Cycles That Amplify the Rise and Fall of Urban Concentration by Millennials. *Housing Policy Debate*, 26(6):928–947.

Oaxaca, R. (1973). Male-female wage differentials in urban labor markets. *International economic review*, pages 693–709.

Oaxaca, R. L. and Ransom, M. R. (1999). Identification in detailed wage decompositions. *Review of Economics and Statistics*, 81(1):154–157.

Pearson, K. (1900). X. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 50(302):157–175.

Polson, N. G. and Scott, J. G. (2012). On the Half-Cauchy Prior for a Global Scale Parameter. *Bayesian Analysis*, 7(4):887–902.

Polzin, S. E., Chu, X., and Godfrey, J. (2014). The impact of millennials' travel behavior on future personal vehicle travel. *Energy Strategy Reviews*, 5:59–65.

Ralph, K. M. (2015). *Stalled On The Road To Adulthood? Analyzing the Nature of Recent Travel Changes for Young Adults in America, 1995 to 2009*. PhD thesis, UCLA.

Ryder, N. B. (1965). The Cohort as a Concept in the Study of Social Change. *American Sociological Review*, 30(6):843–861.

Thigpen, C. and Handy, S. (2018). Driver's licensing delay: A retrospective case study of the impact of attitudes, parental and social influences, and intergenerational differences. *Transportation Research Part A: Policy and Practice*, 111:24–40.

Vij, A., Gorripaty, S., and Walker, J. L. (2017). From trend spotting to trend 'splaining: Understanding modal preference shifts in the San Francisco Bay Area. *Transportation Research Part A: Policy and Practice*, 95:238–258.

Yang, Y. (2008). Social inequalities in happiness in the United States, 1972 to 2004: An age-period-cohort analysis. *American Sociological Review*, 73(2):204–226.

Zhang, Y., Zhou, H., Zhou, J., and Sun, W. (2017). Regression Models for Multivariate Count Data. *Journal of Computational and Graphical Statistics*, 26(1):1–13.

Zumkeller, D. and Chlond, B. (2009). Dynamics of Change: Fifteen-Year German Mobility Panel.

A. Counterfactual analysis

A.1. Counterfactual predictions

We explain in more detail how the counterfactual predictions are obtained. In the hierarchical Bayesian multivariate Poisson log-normal model, the mean predicted number of daily trips by mode j for an individual n_S from a subsample S in period t is

$$\bar{y}\left(\boldsymbol{X}_{n_S tj}, P(\boldsymbol{\beta}_{ctj}, \xi_{n_S j}|\boldsymbol{y})\right) = \int\int \exp\left(\boldsymbol{X}_{n_S tj}\boldsymbol{\beta}_{ctj} + \xi_{n_S j}\right) P(\boldsymbol{\beta}_{ctj}, \xi_{n_S j}|\boldsymbol{y}) d\boldsymbol{\beta}_{c(n_S)tj} d\xi_{n_S j}, \tag{18}$$

where $P(\boldsymbol{\beta}_{ctj}, \xi_{n_S j}|\boldsymbol{y})$ denotes the posterior distribution of $\left\{\boldsymbol{\beta}_{ctj}, \xi_{n_S j}\right\}$. The mean predicted number of daily trips by mode j in period t for sample S is then

$$\bar{z}\left(\boldsymbol{X}_{1:N_S,t,j}, P(\boldsymbol{\beta}_{ctj}, \xi_{1:N_S,j}|\boldsymbol{y})\right) = \frac{1}{N_S}\sum_{n_S}^{N_S} \bar{y}\left(\boldsymbol{X}_{n_S tj}, P(\boldsymbol{\beta}_{ctj}, \xi_{n_S j}|\boldsymbol{y})\right) \tag{19}$$

We use X to label the first subsample and Y to label the second subsample. Moreover, we partition the explanatory variables $\boldsymbol{X}_{1:N_S,t,j}$ into two groups $\boldsymbol{X}_{t,j,M}$ and $\boldsymbol{X}_{1:N_S,t,j,D}$ such that

$$\boldsymbol{X}_{1:N_S,t,j} = \left\{\boldsymbol{X}_{tj,M}, \boldsymbol{X}_{1:N_S,j,D}\right\} \tag{20}$$

and

$$\bar{z}\left(\boldsymbol{X}_{1:N_S,t,j}, P(\boldsymbol{\beta}_{ctj}, \xi_{1:N_S,j}|\boldsymbol{y})\right) = \bar{z}\left(\boldsymbol{X}_{tj,M}, \boldsymbol{X}_{1:N_S,j,D}, P(\boldsymbol{\beta}_{ctj}, \xi_{1:N_S,j}|\boldsymbol{y})\right). \tag{21}$$

Here, $\boldsymbol{X}_{tj,M}$ includes macro-economic factors, i.e. price levels of petrol and public transportation services, while $\boldsymbol{X}_{1:N_S,j,D}$ includes demographic factors, i.e. all other non-price explanatory variables (see Section 3.4). Moreover, $\boldsymbol{X}_{1:N_S,t,j,D}$ includes the mode-specific constant.

A.2. Blinder-Oaxaca decomposition

We perform a Blinder-Oaxaca decomposition for non-linear models (Bauer and Sinning, 2008; Blinder, 1973; Oaxaca, 1973) to decompose the total change $\Delta_j^{\mathrm{DE+PE+CE}}$ in daily trip frequencies for mode j

into three terms, namely Δ_j^{DE}, Δ_j^{PE} and Δ_j^{CE}. To be specific, we have

$$\begin{aligned}\Delta_j^{\text{DE + PE + CE}} =& \bar{z}\left(\boldsymbol{X}_{2016,j,M}, \boldsymbol{X}_{1:N_Y,j,D}, P(\boldsymbol{\beta}_{c_Y,2016,j}, \xi_{1:N_Y,j}|\boldsymbol{y})\right) - \\ & \bar{z}\left(\boldsymbol{X}_{1998,j,M}, \boldsymbol{X}_{1:N_X,j,D}, P(\boldsymbol{\beta}_{c_X,1998,j}, \xi_{1:N_X,j}|\boldsymbol{y})\right) \\ =& \Delta_j^{\text{DE}} + \Delta_j^{\text{PE}} + \Delta_j^{\text{CE}},\end{aligned} \tag{22}$$

where

$$\begin{aligned}\Delta_j^{\text{DE}} =& \bar{z}\left(\boldsymbol{X}_{1998,j,M}, \boldsymbol{X}_{1:N_Y,j,D}, P(\boldsymbol{\beta}_{c_X,1998,j}, \xi_{1:N_Y,j}|\boldsymbol{y})\right) - \\ & \bar{z}\left(\boldsymbol{X}_{1998,j,M}, \boldsymbol{X}_{1:N_X,j,D}, P(\boldsymbol{\beta}_{c_X,1998,j}, \xi_{1:N_X,j}|\boldsymbol{y})\right),\end{aligned} \tag{23}$$

$$\begin{aligned}\Delta_j^{\text{PE}} =& \bar{z}\left(\boldsymbol{X}_{2016,j,M}, \boldsymbol{X}_{1:N_Y,j,D}, P(\boldsymbol{\beta}_{c_X,2016,j}, \xi_{1:N_Y,j}|\boldsymbol{y})\right) - \\ & \bar{z}\left(\boldsymbol{X}_{1998,j,M}, \boldsymbol{X}_{1:N_Y,j,D}, P(\boldsymbol{\beta}_{c_X,1998,j}, \xi_{1:N_Y,j}|\boldsymbol{y})\right),\end{aligned} \tag{24}$$

$$\begin{aligned}\Delta_j^{\text{CE}} =& \bar{z}\left(\boldsymbol{X}_{2016,j,M}, \boldsymbol{X}_{1:N_Y,j,D}, P(\boldsymbol{\beta}_{c_Y,2016,j}, \xi_{1:N_Y,j}|\boldsymbol{y})\right) - \\ & \bar{z}\left(\boldsymbol{X}_{2016,j,M}, \boldsymbol{X}_{1:N_Y,j,D}, P(\boldsymbol{\beta}_{c_X,2016,j}, \xi_{1:N_Y,j}|\boldsymbol{y})\right).\end{aligned} \tag{25}$$

Here, c_X is the set of all birth cohorts that make up generational cohort X. We let

$$\bar{z}\left(\boldsymbol{X}_{tj,M}, \boldsymbol{X}_{1:N_X,j,D}, P(\boldsymbol{\beta}_{c_X,t,j}, \xi_{1:N_X,j}|\boldsymbol{y})\right) = \frac{1}{|c_X|}\sum_{c\in c_X} \bar{z}\left(\boldsymbol{X}_{tj,M}, \boldsymbol{X}_{1:N_X,j,D}, P(\boldsymbol{\beta}_{ctj}, \xi_{1:N_X,j}|\boldsymbol{y})\right) \tag{26}$$

and proceed similarly for generational cohort Y.