



Combining Social Media and Survey Data to Nowcast Migrant Stocks in the United States

Monica Alexander¹ · Kivan Polimis² · Emilio Zagheni³

Received: 11 April 2020 / Accepted: 20 July 2020 / Published online: 18 August 2020
© The Author(s) 2020

Abstract

Measuring and forecasting migration patterns has important implications for understanding broader population trends, for designing policy effectively and for allocating resources. However, data on migration and mobility are often lacking, and those that do exist are not available in a timely manner. Social media data offer new opportunities to provide more up-to-date demographic estimates and to complement more traditional data sources. Facebook’s Advertising Platform, for example, is a potentially rich data source of demographic information that is regularly updated. However, Facebook’s users are not representative of the underlying population. This paper proposes a statistical framework to combine social media data with traditional survey data to produce timely ‘nowcasts’ of migrant stocks by state in the United States. The model incorporates bias adjustment of Facebook data, and a pooled principal component time series approach, to account for correlations across age, time and space. We use the model to estimate and project migrants from Mexico, India and Germany, three migrant groups with varying levels and trends of migration in the US. By comparing short-term projections with data from the American Community Survey, we show that the model predictions outperform alternatives that rely solely on either social media or survey data.

Keywords Migration · Social media data · Bayesian methods · United States

✉ Monica Alexander
monica.alexander@utoronto.ca

¹ Departments of Statistical Sciences and Sociology, University of Toronto, 100 St. George Street, Toronto, ON M5S 3G3, Canada

² Center for Studies in Demography and Ecology, University of Washington, Seattle, USA

³ Max Planck Institute for Demographic Research, Rostock, Germany

Introduction

Accurate, reliable, and timely estimates of migration indicators, such as flows and stocks, are crucial for understanding population dynamics and demographic change, for designing effective economic, social and health policies, and for supporting migrants and their families. However, data on migration from traditional sources, such as censuses, surveys or administrative registers, are often insufficient. Even when these sources exist, the data available may lack the granularity of information required to understand migration trends, or are not released in a manner that is timely enough to monitor changes in trends.

As migration flows can change substantially over a short period of time—such as in response to a natural disaster or war and conflict—relying on out-dated data is often not sufficient.

Timely and reliable information about migration stocks is important not only to understand migration patterns. It is key also to monitor fertility, population health and mortality. Even when accurate data on births and deaths exist, often demographers are faced with large uncertainty in population counts, which, in various disaggregations, form the denominators for standard demographic rates. Most of the uncertainty is driven by lack of appropriate information on how migration stocks change over time and space.

As a consequence of data availability issues, we need to consider how non-traditional data can be leveraged to complement existing sources in order to improve estimates and predictions of migration indicators over time. Previous work has explored the use of data such as call detail records (Blumenstock 2012; Pestre et al. 2020), air traffic data (Gabrielli et al. 2019), tax file records (Engels and Healy 1981) and other sources like billing addresses or school enrollment (Foulkes and Newbold 2008) to estimate migration. Additionally, an increasingly large body of work has investigated the use of social media data, from websites such as Twitter (Zagheni et al. 2014), Facebook (Zagheni et al. 2017) and LinkedIn (State et al. 2014). Provided that the data can be obtained in a reliable, timely, and ethical way, information about the users of social media websites is potentially an incredibly rich demographic data source.

Data on these populations are essentially collected in real time, and while individual-level information is usually restricted, many of the social media websites provide a certain amount of aggregate-level information through their advertising platforms (Cesare et al. 2018). In particular, Facebook's Advertising platform allows information to be extracted on the relative size of groups by key demographics such as age, sex, location of residence and country of origin, thereby potentially acting as a measure of the relative size of migrant groups in a particular country.

While these data have clear potential for use in demographic research, with respect to timeliness and the size of the sample being considered, there are some notable issues that need to be overcome. In particular, for any given population subgroup of interest, the corresponding users of Facebook or any other social media platform are unlikely to be a representative sample. An additional

challenge is to use these data in a way that meaningfully combines new ‘signal’ or information about migration trends with existing knowledge on probable migration trends from historical data sources.

This paper proposes a statistical framework to combine social media data from Facebook, with traditional survey data from the American Community Survey (ACS), in order to produce timely ‘nowcasts’ of migrant stocks by state in the United States. The framework consists of a Bayesian hierarchical model which incorporates bias adjustment of the Facebook data, a demographic time series approach to account for historical past trends, and a geographic pooling component which allows information about the age structure of migrants to be shared across space. The model also accounts for the different types of uncertainty that are likely to be present in Facebook and traditional survey data. The resulting model produces estimates and short-term projections of migrant stocks by US state of destination and country of origin, and is shown to outperform valid modeling alternatives.

The remainder of the paper is structured as follows. First, we briefly discuss previous demographic research which incorporates social media. Then we outline the data sources used, and in particular how the Facebook data were collected. The “[Model](#)” section discusses the model set-up, assumptions, and computation. We then present results for Mexican, Indian and German migrants by US state, and validate model performance against reasonable alternatives. Finally, the strengths and limitations of the model are discussed, together with avenues for future research.

Background

The lack of good-quality data on migration is a global problem, with data sparsity issues prevalent in both developed and developing countries (Landau and Achiume 2017). This has prompted scholars to investigate the use of other types of data to monitor migration trends. In particular, with the rise of social media use around the world, new data that have potential for demographic research have emerged.

Scholars began using social media and web data to estimate and track demographic indicators over time in the early 2010s. The earliest papers illustrated how geo-located data from email services and web-based applications such as Twitter, Google Latitude, Foursquare or Yahoo! can be used (Ferrari et al. 2011; Noulas et al. 2011; Zagheni and Weber 2012). Initial research focused on evaluating spatial mobility of populations at a city or regional level. For example, Ferrari et al. (2011) used Twitter data to study patterns of urban movement in New York. In the first effort to tackle global trends, Zagheni and Weber (2012) linked the geographic locations of IP addresses of Yahoo! emails to the user’s self-reported demographic data to estimate age- and sex-specific migration flows in a large number of countries around the world.

Recent efforts have focused on using data from social media and networking websites such as Twitter, LinkedIn and, more recently, Facebook and Instagram. These websites provide public access to an Application Programming Interface (API), which makes it possible to send requests and receive responses from these websites for data like tweet hashtag counts, the number of jobs in a certain industry,

or number of cell phone users in a particular area. Researchers have utilized these APIs to extract publicly available demographic and location data for use in social research, in particular to study outcomes such as migration (Yildiz et al. 2017; Zagheni et al. 2017), fertility (Rampazzo et al. 2018), gender equality (Fatehka et al. 2018; Garcia et al. 2018), and health (Araujo et al. 2017). For instance, Garcia et al. (2018) used Facebook data to create an index of the internet gender divide in 217 countries, showing that this indicator encapsulated gender equality indices in education, health and economic opportunity. Yildiz et al. (2017) used a combination of geo-located tweets and image recognition software to obtain estimates of internal migration in England.

In work relevant to this paper, Zagheni et al. (2017) presented a proof of concept for estimating migration stocks in the United State by age, sex and state, using Facebook's Advertising Platform. More recently, Alexander et al. (2019) used the same type of data to track changes in migrants over time, in the context of estimating out-migration from Puerto Rico following Hurricane Maria in September 2017.

The main gap in the literature is related to the lack of a suitable statistical model for combining 'traditional' data sources on migrants—in the form of censuses, nationally representative surveys, or other vital statistics—with migration information from social media data. The goal of this paper is thus to develop a probabilistic framework that allows representative and historical time series to be combined, in a sound statistical framework, with non-representative—but timely—sources from social media.

Data

Facebook Advertising Data

Facebook for Business has developed a targeted advertising platform, called Ads Manager that provides a graphical user interface to allow advertisers to micro-target specific audiences. Demographic characteristics that can be targeted include information directly reported by Facebook users, such as age or sex, and information indirectly inferred from using Facebook's platform or affiliated websites, such as location and behavioral interests. Before launching an advertisement, an advertiser can select a variety of characteristics (e.g., Australians living in California, who are female, and aged 30–35) and get an estimate of the 'potential reach' (monthly active users) to this subgroup. These estimates can be obtained, in a programmatic way, for a variety of different migrant groups (i.e., by age and sex) to whom Facebook refers as expatriates ('expats').

We use the estimates of potential reach by expat group, age and sex to track sizes of migration stocks over time. These estimates can be obtained before the launch of an advertisement, and as such are obtained free of charge. We use the Ads Manager back-end application, Facebook's Marketing API, to extract estimates of potential reach over time programmatically with the Python module pySocialWatcher (Araujo et al. 2017). With pySocialWatcher, we collected data across 11 age groups (10 UN age groups from 15–19 to 60–65; an 11th group for the entire available Facebook

population of 13–65 was also used) and three gender groups (female, male, and total population). Data was collected using Amazon Web Services (AWS) EC2 Instance servers.

As part of a broader project on using social media in demographic research, we started data collection in January 2017 and collect a new wave of data every 2–3 months. For each wave of data collection we obtained state-level estimates of all Facebook users (by age, sex, and gender) as well as state-level estimates of 90 expat groups.¹

American Community Survey

The American Community Survey (ACS) is an annual survey of the U.S. Census Bureau, designed to supplement the decennial census. Based on the long-form version of the census, the ACS collects information on topics including population, housing, employment and education from a nationally representative sample. Data on migrant stocks can be readily obtained from the ACS. In particular, in every year of the ACS, the survey has contained a question asking the birthplace of the person; if it is inside the United States, the state is recorded, and if it is outside the United States, the country is recorded. This birthplace variable is recorded as a three digit code to indicate the US state or country of birth. In addition to the birthplace variable, the ACS has information on current state of residence. Thus, we can tabulate the number of migrants from a particular country living in a particular state by looking at the combination of these two variables. From a modeling perspective, we are interested in the proportion of migrants from a particular origin of the total population by 5-year age groups (15–19, 20–24, ..., 50–54) in each state.

We calculated the migrant stock proportions using the 1-year ACS for each year between 2001 and 2017 using micro-data available through the Integrated Public Use Microdata (IPUMS) US project (Ruggles et al. 2000). Standard errors around the calculated proportions based on sampling variation were calculated based on ACS accuracy guidelines (US Census Bureau 2020) and using the delta method.

Model

We are considering two data sources of migration trends in the US: data from Facebook's Advertising Platform, and the ACS. The overall goal of the modeling strategy is to combine information from both these sources to produce estimates of current and future migrant stocks. To do this, the model should have three main characteristics. Firstly, we want to adjust for biases in Facebook data to effectively use up-to-date information on migration patterns from this source. Secondly, we want to be able to incorporate longer time series of information from the ACS. Finally, the data should be combined in a probabilistic way, in order to objectively weigh information

¹ The waves used in this paper are: Wave 1: January 2017; Wave 2: April 2017; Wave 3: June 2017; Wave 4: October 2017; Wave 5: January 2018; Wave 6: March 2018.

from both sources. We propose a Bayesian hierarchical model which achieves these goals. In this section we describe the model in detail.

For a particular migrant group, define $\rho_{x,t,s}$ to be the proportion of migrants of the total populations in age group x at time t and in state s . This quantity $\rho_{x,t,s}$ is the main parameter of interest to be estimated. We have observations of this proportion, which will be denoted $p_{x,t,s}$. The observed proportions are either from Facebook ($p_{x,t,s}^{FB}$) or from the ACS ($p_{x,t,s}^{ACS}$). The $p_{x,t,s}$ s are observed, and it is assumed that these are somehow related to the latent proportions, $\rho_{x,t,s}$, with some associated error.

Facebook Bias Adjustment

The first goal is to adjust the Facebook data to account for the non-representativeness of the Facebook user population. Previous research from Zagheni et al. (2017) showed that, while the bias in the Facebook migrant data is substantial, it is also relatively systematic by age and migrant group and can be modeled.

Following their approach, we introduce a regression model which relates the proportions of migrants in Facebook, $p_{x,t,s}^{FB}$, to the proportions in the ACS in a similar time period, $p_{x,t,s}^{ACS}$, plus a series of age and state variables. In particular, for a particular migrant group, express $p_{x,t,s}^{ACS}$ on a log scale, as

$$\log p_{x,t,s}^{ACS} = \alpha_0 + \alpha_1 \log p_{x,t,s}^{FB} + \beta \mathbf{X} + \varepsilon_{FB} \quad (1)$$

where \mathbf{X} is a covariate matrix containing an indicator variable for each age group (15–19, 20–24, ..., 50–54) and each of the 50 states plus Washington D.C. This means that we estimate a fixed effect for each age group and state. In addition, we assume that the error is i.i.d. and that

$$\varepsilon_{FB} \sim N(0, \sigma_{FB}^2). \quad (2)$$

Estimates of the coefficients α_0 , α_1 and the vector of β 's are obtained using the first wave of the Facebook data and the 2016 ACS data. Once obtained, these coefficient estimates are then used to adjust subsequent waves of Facebook data, i.e. we calculate

$$\log p_{x,t,s}^* = \hat{\alpha}_0 + \hat{\alpha}_1 \log p_{x,t,s}^{FB} + \hat{\beta} \mathbf{X} \quad (3)$$

where $\log p_{x,t,s}^*$ is a 'bias-adjusted' version of the Facebook data. This is taken to be our 'best guess' of what the migrant stocks in group x, t, s are, based on the Facebook data alone. Note that an estimate of σ_{FB}^2 is also obtained, that is, the variance of the error terms, which becomes important in the final model (see the "[Bringing It All Together](#)" section).

Time Series Modeling of ACS Using Principal Components

In addition to using data from Facebook, we also want to incorporate the relatively long historical time series of information on migrant stocks obtained from the ACS.

A reasonable short-term forecast based on ACS should model historical trends and project them forward.

There are many different time series models that could be used in this context. Perhaps the simplest approach would be to project forward a moving average of the time series for each age group and state combination. Alternatively, we could use a classical Box–Jenkins approach and model the time series of migrant stocks in each age group and state separately using an appropriately specified ARIMA model. However, these methods would not place any constraints on the age structure of migration. Given this demographic context, we expect that the age distribution of migration displays strong patterns and changes in a relatively regular way over time. This is because of regularities in the age at migration as well as historical trends which include different waves of migrants, who also age over time. As such, we chose to incorporate this prior knowledge into our model through a principal components approach.

Principal component-based models have a long history in demographic modeling, with the most well-known example being the Lee–Carter mortality model (Lee and Carter 1992). The idea is that a set of age-specific demographic rates observed over time can be expressed as a combination of a series of characteristic patterns (or principal components). The Lee–Carter approach uses the mean age-specific mortality schedule and first principal component, which is interpreted as age-specific contributions to mortality change over time. This model can easily be extended to include higher-order principal components, which various researchers have done.

Apart from the Lee–Carter model and variants (e.g. Li et al. (2004), Lee (2000), Renshaw and Haberman (2006)), principal component models were recently used to estimate and forecast mortality (e.g. Alexander et al. (2017)), fertility (e.g. Schmertmann et al. (2014)) and overall population (Wiśniowski et al. (2015)). Here, we extend this idea to parsimoniously estimate and project migration stocks by age and state.

Model Overview

Age-specific migration schedules are decomposed into independent age and time components. The time component is then projected forward as a time series, taking autocorrelated error into account. We propose a log-linear model for $p_{x,t,s}$:

$$\log p_{x,t,s}^{ACS} = \beta_{t,s,1}Z_{x,1} + \beta_{t,s,2}Z_{x,2} + \varepsilon_{x,t,s} \quad (4)$$

where $Z_{x,1}$ and $Z_{x,2}$ are the first and second ‘principal components’, $\beta_{t,s,1}$ and $\beta_{t,s,2}$ are state and time-specific coefficients, to be estimated, and $\varepsilon_{x,t,s}$ is an error term. The principal components are obtained via Singular Value Decomposition (SVD), as outlined in the next section. To obtain estimates of $\beta_{ts,1}$ and $\beta_{ts,2}$, we impose some smoothing over time and pooling of information across space, as outlined in the "[Time Series Modeling of ACS Using Principal Components](#)" section. Finally, as discussed in the "[Time Series Modeling of ACS Using Principal Components](#)" section we place a time series model on the error term, $\varepsilon_{x,t,s}$, to account for autocorrelation.

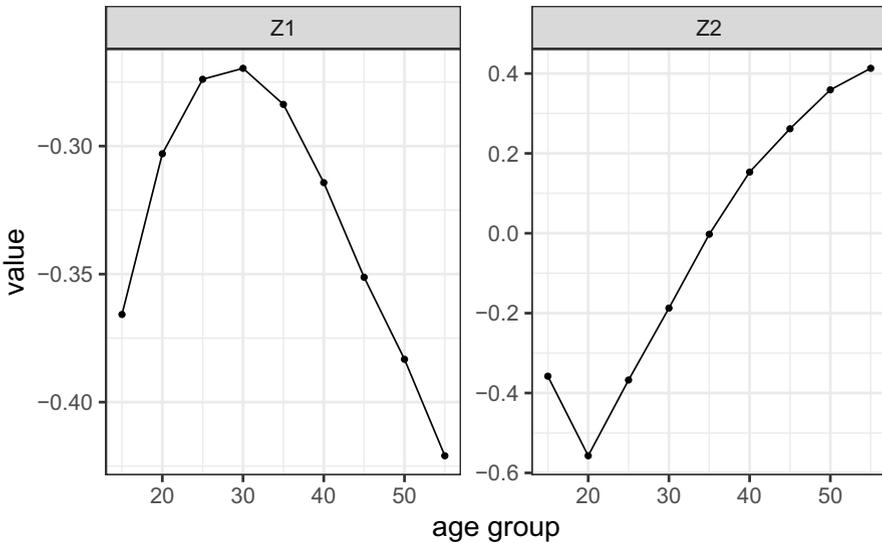


Fig. 1 Principal Components for Mexico

Obtaining the Principal Components

The principal component terms $Z_{x,1}$ and $Z_{x,2}$ aim to capture the main sources of systematic variation in migration patterns across age. They are obtained by first creating a matrix of (logged) historical age-specific migration schedules based on ACS data from 2001 to 2016. Singular Value Decomposition (SVD) is then performed on this matrix to obtain principal components of the age-specific migration. In particular, let \mathbf{X} be a $N \times G$ matrix of log-migration stock rates, where N is the number of state-years and G is the number of age groups. In this case, we had $N = 51$ states + DC $\times 16$ years = 816 observations of $G = 9$ age groups (15–19, 20–24, ..., 50–54). The SVD of \mathbf{X} is

$$\mathbf{X} = \mathbf{U}\mathbf{D}\mathbf{V}', \quad (5)$$

where \mathbf{U} is a $N \times N$ matrix, \mathbf{D} is a $N \times G$ matrix and \mathbf{V} is a $G \times G$ matrix. The first two columns of \mathbf{V} (the first two right-singular values of \mathbf{X}) are $Z_{x,1}$ and $Z_{x,2}$.

For example, Fig. 1 shows the resulting $Z_{x,1}$ and $Z_{x,2}$ for the Mexican migrant group in the US. These were obtained via the following steps:

1. Calculate $p_{x,t,s}^{ACS}$, i.e. the proportion of migrants in age group x , year t and state s for each age group, year and state in ACS 2001–2016.
2. Create \mathbf{X} where each element is $\log p_{x,t,s}^{ACS}$, every row is a state-year and every column is an age group.

3. Perform SVD on \mathbf{X} and extract the first two columns² of \mathbf{V} .

The principal components shown in Fig. 1 can be interpreted as a baseline migration age schedule ($Z_{\cdot,1}$) and age-specific contributions to change over time ($Z_{\cdot,2}$). In the model, the coefficient on $Z_{\cdot,1}$ ($\beta_{t,s,1}$) moves the overall level of Mexican migrants up or down, depending on the year and state. The coefficient on $Z_{\cdot,2}$ allows the age distribution to shift to older or younger ages. For $Z_{\cdot,2}$, the sign changes from negative to positive at age 35. This means that the larger and more positive the value of $\beta_{t,s,2}$, the older the migrant age distribution.³

Sharing Information Across Time and Space

The model specified in Eq. 4 requires the estimation of two coefficients, $\beta_{ts,1}$ and $\beta_{ts,2}$ for each time t and state s . One option would be to estimate each of these coefficients separately for every year and state. However, we would like to incorporate the knowledge that trends in migration over time are likely to exhibit relatively regular patterns. In addition, for the coefficient on the second principal component—which allows for the age distribution of migrants to shift to the left or right, we would like to share information about the patterns in migration across geographic space.

The coefficient on the first principal component, $\beta_{t,s,1}$, is modeled as a random walk, i.e.

$$\beta_{t,s,1} \sim N(\beta_{t-1,s,1}, \sigma_{\beta_1}^2) \quad (6)$$

This allows for information about the level of migration within each state to be smoothed over time. The random walk structure allows for the estimate in the current time period, $\beta_{t,s,1}$, to be partially informed by the previous period.

For the coefficient on the second principal component, we place the following hierarchical structure on the β 's:

$$\beta_{t-1,s,2} \sim N(\Phi_t, \sigma_{\beta}^2) \quad (7)$$

$$\Phi_t \sim N(\Phi_{t-1}, \sigma_{\Phi}^2) \quad (8)$$

The Φ_t term represents essentially a national mean; as such the $\beta_{t,s,2}$'s are a draw from a national distribution with some mean and variance. In this way, information about how the age distribution is ageing over time is shared across states. The more information about migration there is available for a particular state (i.e., the larger the migrant population), the less the estimate of $\beta_{t,s,2}$ is influenced by the overall mean. Conversely, states with smaller migrant populations where the trends over time are less clear from the data are partially informed by patterns in larger states.

² We used the 'svd' function in R.

³ Note that the interpretation of $Z_{\cdot,1}$ and $Z_{\cdot,2}$ is similar to the interpretation of the a_x and b_x terms in the usual Lee–Carter model.

Note that the geographical hierarchical structure is not present on the first coefficient, as this represents an overall level of migration. Pooling information across space about the level of migration would artificially increase migrant proportions in smaller states.

Auto-Correlated Error

The final piece of the time series model is the error term $\varepsilon_{x,t,s}$. This term is included in the model to allow for extra variation in migration age schedules that is not otherwise picked up by the principal components. We expect the extra variation to be autocorrelated, and as such we model the error term as an AR(1) process:

$$\varepsilon_{x,t,s} \sim N(\rho_{x,s}\varepsilon_{x,t-1,s}, \sigma_\varepsilon^2) \quad (9)$$

where $\rho_{x,s} \in [0, 1]$.

Projection

The model described above is fit to ACS data from 2001 to 2016. However, estimates in more recent years can easily be obtained by projecting the time series aspects of this model forward. In particular, for time $t + 1$:

- Obtain an estimate for $\beta_{t+1,s,1}$ from $\beta_{t+1,s,1} \sim N(\beta_{t,s,1}, \sigma_\beta^2)$.
- Obtain an estimate for $\beta_{t+1,s,2}$ from $\beta_{t+1,s,2} \sim N(\Phi_{t+1}, \sigma_\Phi^2)$ and $\Phi_{t+1} \sim N(\Phi_t, \sigma_\Phi^2)$.
- Obtain an estimate for $\varepsilon_{x,t+1,s}$ from $\varepsilon_{x,t+1,s} \sim N(\rho_{x,s}\varepsilon_{x,t,s}, \sigma_\varepsilon^2)$.
- Calculate $\log p_{x,t+1,s}^{ACS}$ based on Eq. 4.

Bringing It All Together

The "Facebook Bias Adjustment" and "Time Series Modeling of ACS Using Principal Components" sections described two ways to obtain current 'nowcasts' of migrant stocks. One option would be to take the most recent data obtained from Facebook, adjust using the bias-adjustment model, and take the resulting estimate as our nowcast. Another option would be to project forward the ACS model to the time period of interest. Ideally, we would like to incorporate both sources into our final estimate. Perhaps an option to do this would be to just take an average of the two resulting estimates. However, we would like to weigh the estimates from both sources more objectively, taking different sorts of uncertainty into consideration.

Our solution is to combine both models into one framework, and use the results from both methods as data points for our 'best estimate' nowcast. This is illustrated in Fig. 2. Facebook inputs are calibrated with the ACS via the adjustment model. ACS data are used to obtain principal components based on past migration data. The modeling structure allows for information exchange over time and across geographic space. The key piece of the combined model, which has yet to be explained, is the data model (or likelihood), which allows data from the different sources to have different associated error.

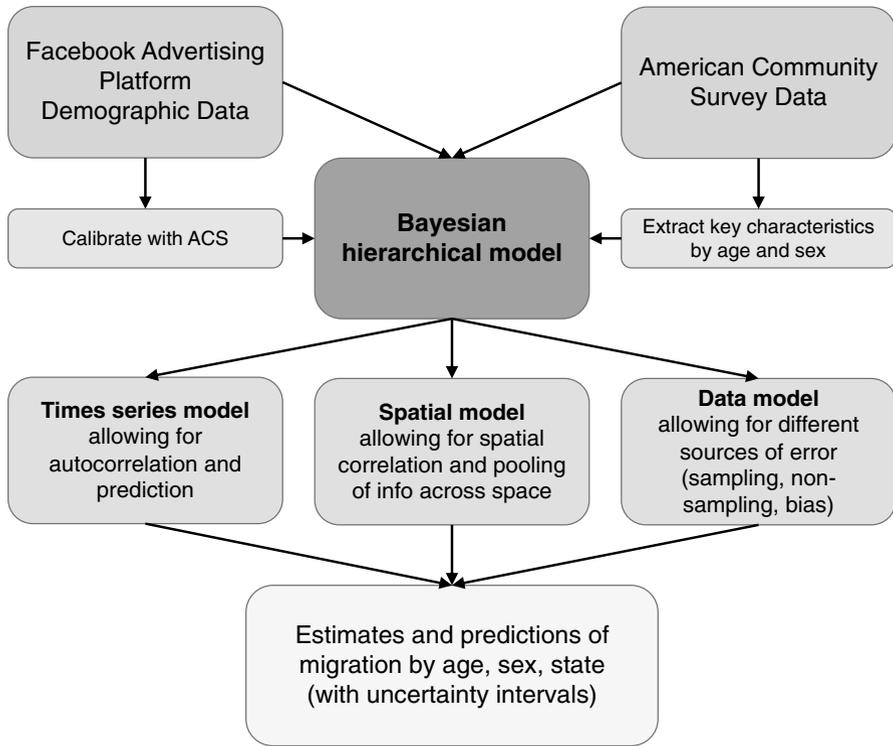


Fig. 2 Modeling framework

Data Model

As outlined above, we observe migrant proportions $p_{x,t,s}$ from either Facebook or the ACS. The data model assumes

$$\log p_{x,t,s} \sim N(\log \rho_{x,t,s}, \sigma_p^2)$$

i.e. the log of the observed proportion is assumed to have mean $\log \rho_{x,t,s}$ and variance σ_p^2 , where σ_p^2 depends on the data source:

$$\sigma_p^2 = \begin{cases} \sigma_s^2, & \text{if ACS} \\ \sigma_s^2 + \sigma_{FB}^2 + \sigma_{ns}^2, & \text{if Facebook} \end{cases}$$

Here, σ_s^2 refers to sampling error, and is assumed to be present in both ACS and Facebook data. For the ACS data, sampling errors are calculated based on guidelines from the US Census Bureau (2020). For Facebook data, the sampling error is calculated assuming the binomial approximation to the Normal distribution and calculating

$$\sigma_s^2 = \frac{P_{x,t,s} \cdot (1 - P_{x,t,s})}{N_{x,t,s}^{FB}}$$

where $N_{x,t,s}^{FB}$ is the total size of the Facebook population in subgroup x, t, s . The sampling errors were then transformed to be on the log scale using the delta method.

For the Facebook data there are two additional error terms. σ_{FB}^2 refers to the error associated with our bias-adjustment model (Eq. 1) and is estimated within this model. This captures the fact that our adjustment model is imperfect and that extra variation remains. Additionally, we allow for a non-sampling error with σ_{ns}^2 , which aims at capturing additional uncertainty like variation in the way potential reach is estimated across waves.

For a given population size, the sampling error is going to be of similar size for ACS and Facebook data. As such, the error term associated with the Facebook data, which is the sum of three terms, will always be bigger than for ACS. In practice, this means that estimates from the model will follow (i.e. give more weight to) the ACS data.

Summary of Full Model

The full model is summarized below. Equation 10 is the data model. Equations 11–15 relate to the ACS time series model. Equation 18 relates to the Facebook regression model. Equations 16 and 17 allow the observation of the proportion of interest to come from a different source (Facebook or ACS), which has a different associated variance. Note that $\rho_{x,t,s}$ is estimated on a yearly basis, but it is assumed that j waves of Facebook data are collected within any 1 year.

$$\log p_{x,t,s} \sim N(\log \rho_{x,t,s}, \sigma^2) \quad (10)$$

$$\log \rho_{x,t,s} = \beta_{t,s,1} Z_{x,1} + \beta_{t,s,2} Z_{x,2} + \varepsilon_{x,t,s} \quad (11)$$

$$\beta_{t,s,1} \sim N(\beta_{t-1,s,1}, \sigma_{\beta_1}^2) \quad (12)$$

$$\beta_{t,s,2} \sim N(\Phi_{t,2}, \sigma_{\beta}^2) \quad (13)$$

$$\Phi_t \sim N(\Phi_{t-1}, \sigma_{\Phi}^2) \quad (14)$$

$$\varepsilon_{x,t,s} \sim N(\rho_{x,s} \varepsilon_{x,t-1,s}, \sigma_{\varepsilon}^2) \quad (15)$$

$$p_{x,t,s} = \begin{cases} p_{x,t,s}^{ACS}, & \text{if } 2001 \leq t \leq 2016 \\ p_{x,t,sj}^*, & \text{if } t \geq 2017 \end{cases} \quad (16)$$

$$\sigma^2 = \begin{cases} \sigma_s^2, & \text{if ACS} \\ \sigma_s^2 + \sigma_{FB}^2 + \sigma_{ns}^2, & \text{if Facebook} \end{cases} \quad (17)$$

$$p_{x,t,sj}^* \sim N(\alpha_0 + \alpha_1 \cdot p_{x,t,sj}^{\text{Facebook}} + X\Gamma, \sigma_{FB}^2) \quad (18)$$

Priors

Weakly informative priors were placed on the coefficients in the Facebook bias-adjustment model, as well as the principal component coefficients in the initial periods:

$$\begin{aligned} \alpha_0 &\sim N(0, 100) \\ \alpha_1 &\sim N(0, 100) \\ \Gamma_0 &\sim N(0, 100) \\ \beta_{1,s,1} &\sim N(0, 100) \\ \Phi_1 &\sim N(0, 100) \end{aligned}$$

In addition, we put weakly informative half-Normal priors on the two standard deviation terms to be estimated:

$$\begin{aligned} \sigma_{FB} &\sim N_+(0, 1) \\ \sigma_{ns} &\sim N_+(0, 1). \end{aligned}$$

Computation

The model was fitted in a Bayesian framework using the statistical software R. Samples were taken from the posterior distributions of the parameters via a Markov Chain Monte Carlo (MCMC) algorithm. This was performed using JAGS software (Plummer et al. 2003). Standard diagnostic checks using trace plots and the Gelman and Rubin diagnostic were used to check convergence (Gelman et al. 2013).

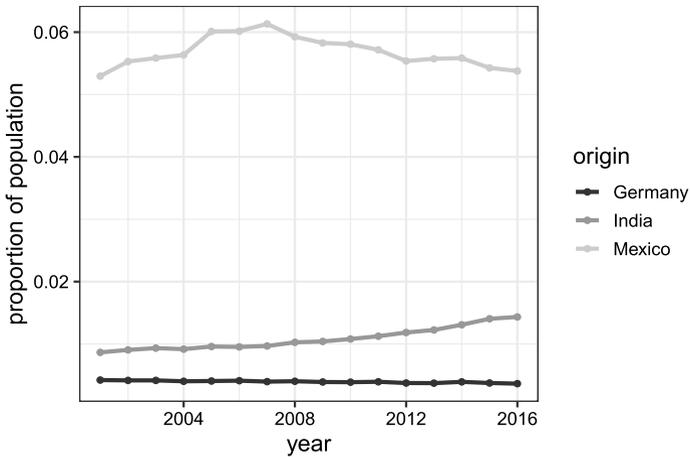
Best estimates of all parameters of interest were taken to be the median of the relevant posterior samples. The 95% Bayesian credible intervals were calculated by finding the 2.5% and 97.5% quantiles of the posterior samples.

All code and data are available on GitHub: <https://github.com/MJAlexander/fb-migration-bayes>.

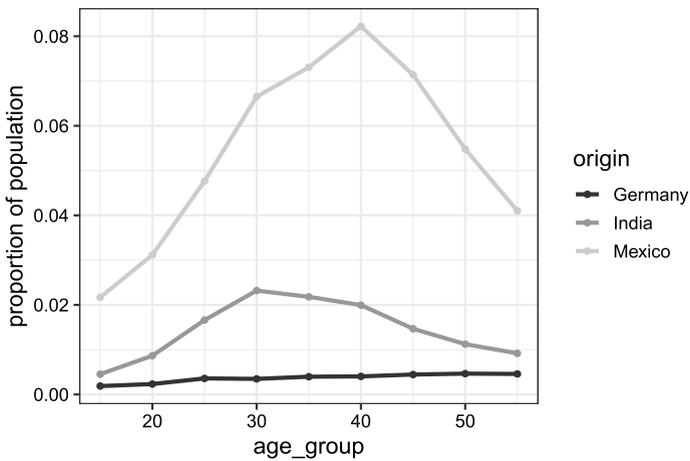
Results

We illustrate the model on male migrants from three different countries: Mexico, India, and Germany. We chose these three migrant groups as they represent a range of different levels and trends over time, as illustrated by the trends in the ACS data shown in Fig. 3a and b.

Firstly, Mexican migrants make up a relatively large share of the overall population, but the proportion has generally declined since around 2007. The age distribution at the national level peaks in the 40–44 year old age group. Secondly, Indian migrants make up a moderate proportion of the total population, but this share is



(a) By proportion of the total population



(b) By proportion of each age group, 2016

Fig. 3 German, Indian and Mexican migrants

increasing over time. The age distribution peaks at younger ages (30–34), compared to Mexicans. Finally, German migrants make up a low and declining share of the population. In contrast to the other migrant groups, the age distribution of German migrants at the national level is relatively flat, increasing slightly across age.

Bias Adjustment of Facebook Data

We firstly illustrate the results of the bias-adjustment step of the Facebook data. Figure 4 shows, for each US state and five-year age group where data are available, the proportion of migrants in each age group for the ACS data in 2016 (black dots), the un-adjusted Facebook data (gray dots), and the estimated bias-adjusted

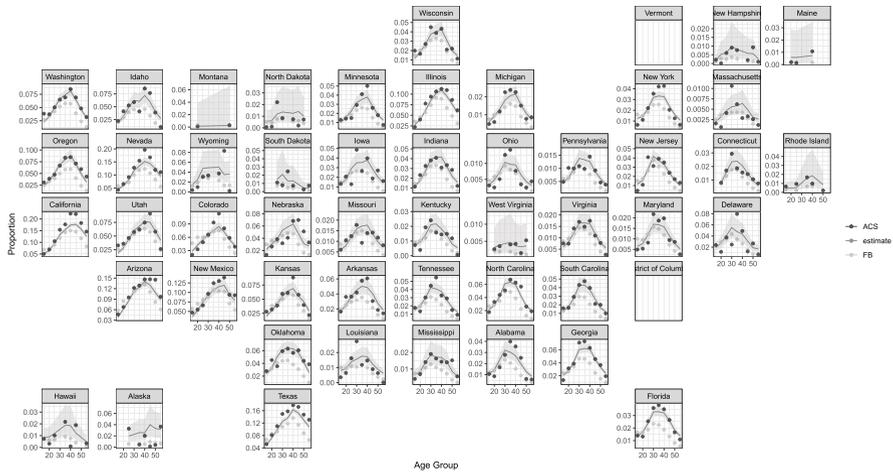


Fig. 4 Bias adjustment of Facebook data for Mexican migrants

Facebook data (line and associated shaded area) for Mexican migrants. Similar plots for migrants from India and Germany are shown in "Appendix A". The interpretation is that if the bias-adjustment step is working reasonably well, the estimate line would be close to the black dots. In general, this appears to be the case. In the case of Mexico, bias-adjusted estimates have a root mean squared error (RMSE) of 0.01, when compared to the ACS data; this is reduced by half from the RMSE of the raw Facebook data.

For all three migrant groups, the raw Facebook data are generally lower than the ACS data, but the bias-adjustment model adjusts these values upwards. In general, across the three migrant groups and across states, the shape of the age distributions in the Facebook and ACS data are similar, with more substantial under-representation in Facebook in the older age groups. These systematic differences mean that the model works well to adjust the raw Facebook data based on age and state effects.

Nowcasts by Age Group and State

Now we move on to short-term projections by age and state. Figure 5 shows the estimated age distribution in 2008 (dark gray) and projected distribution in 2018 (light gray) for Mexican migrants. Similar plots for India and Germany can be found in "Appendix B".

For Mexico (Fig. 5), the relatively high proportions in the border states and on the West coast are apparent, with the highest proportions in California, Texas, Nevada and Arizona. Additionally, the age distribution of Mexican migrants is generally aging over time (shifting to the right), which is consistent with relatively constant stocks.

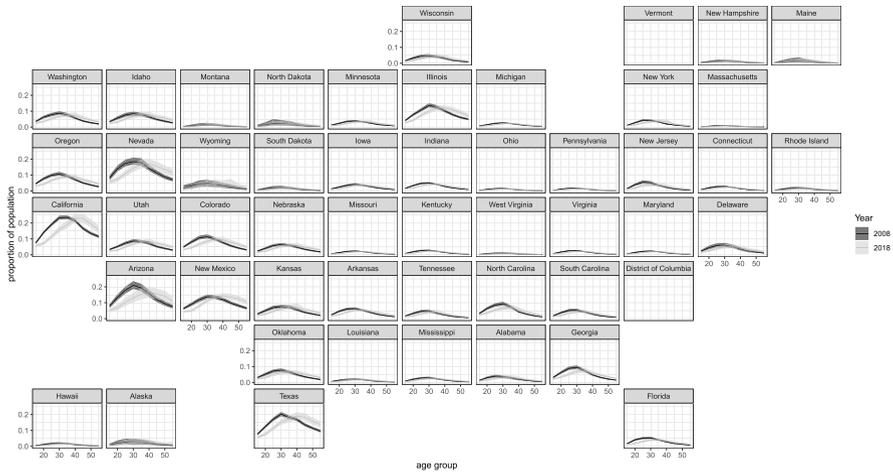


Fig. 5 Estimated and projected age distributions of Mexican migrants by state, 2008 and 2018

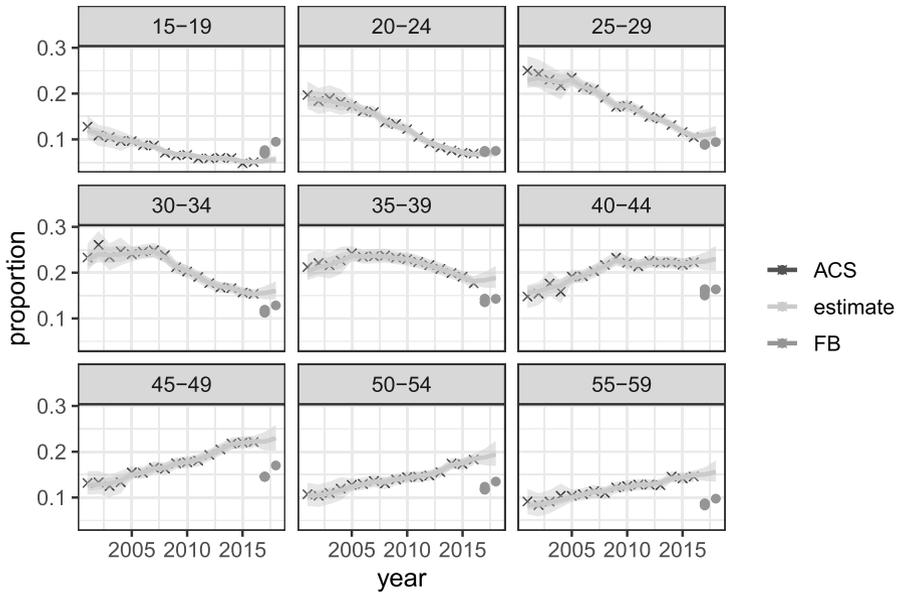
Projected Time Series

Figures 6, 7 and 8 zoom in on two states for each migrant group and show how the Facebook data are used to project forward the time series to the most recent 2 years (2017 and 2018). The full estimated and projected time series from 2001 to 2018 is shown. In the figures, each facet is a 5-year age group. The black crosses area represent the ACS data; these data are broadly available from 2001 to 2016, although some observations are missing (if sample sizes in the ACS were too small to capture information about migrants in that particular state and age group). The medium gray dots represent the (adjusted) Facebook observations, which are available in years 2017 and 2018. The light gray line and associated shaded area is the model estimate and 95% uncertainty intervals.

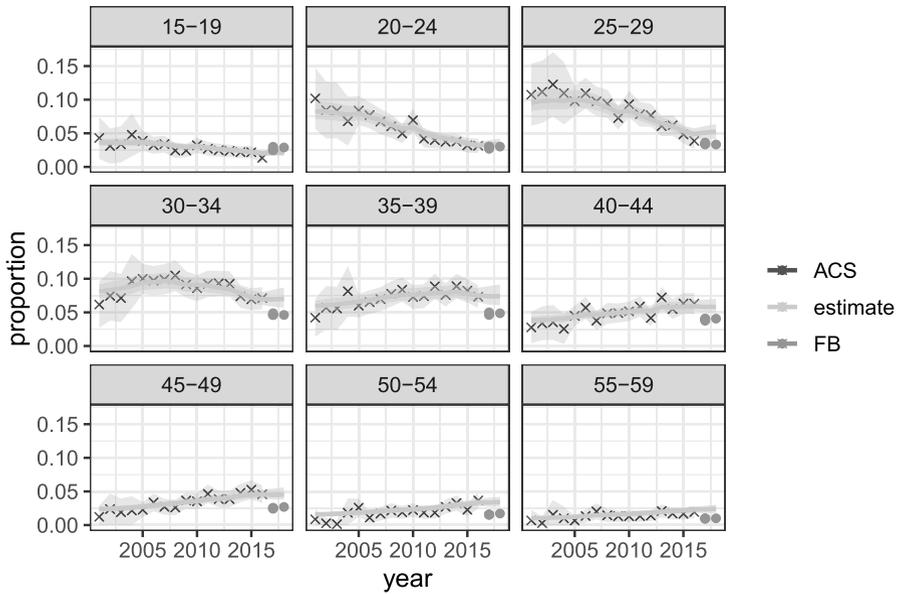
Mexican males in California (Fig. 6a) represent by far the highest proportions of any of the migrant origin/ state combinations considered. The proportion is as high as 0.25 in some age groups, for example 25–29 year olds in 2001 and 40–44 year olds in 2018. As a consequence, the sampling error around the ACS data for this migrant group is relatively small and the model estimates closely follow these data. For the most recent two years, where only Facebook data are available, note that the model estimates do not follow the data as closely and the uncertainty around the model estimates increases. This reflects the fact that there are more sources of error associated with the Facebook data.

In Georgia (Fig. 6b) the levels of Mexican migrants are around half as high as in California. Due to smaller sample sizes in Georgia, the standard errors around the ACS data are much larger, and as such the model estimates do not follow the data as closely. However, the trends for Mexican migrants in California, and Georgia are broadly the same: decreases in the younger age groups, and increases in the older age groups, representing an aging stock of migrants.

For Indian males in California and Georgia (Fig. 7), the proportions are much lower than for the Mexican migrant population, peaking at around 3–4% of the

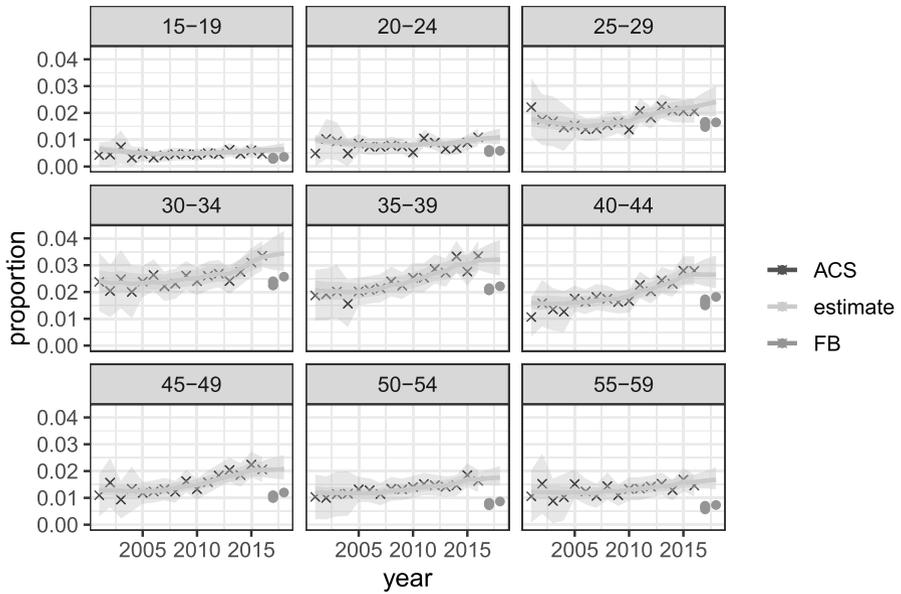


(a) California

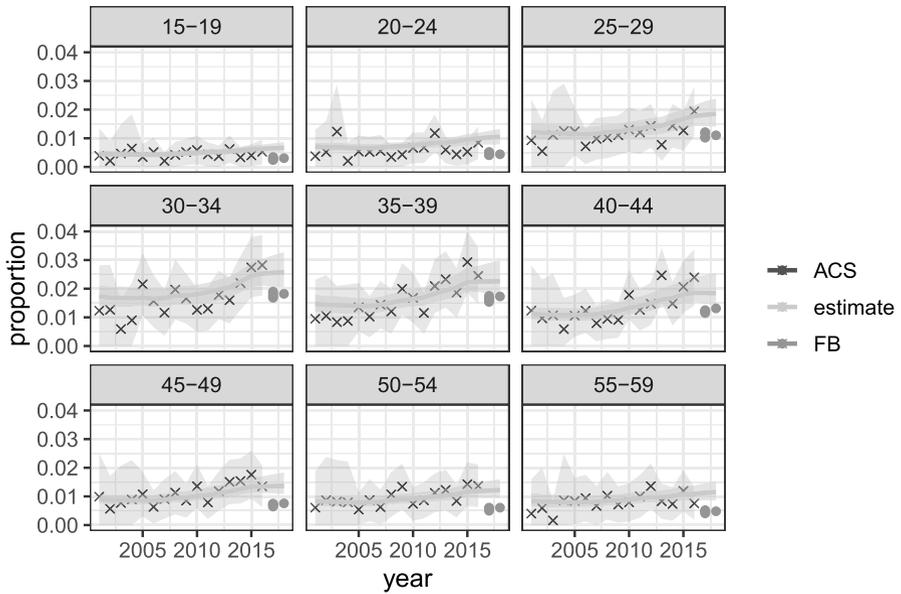


(b) Georgia

Fig. 6 Mexican male migrants by age group, California and Georgia, 2001–2018



(a) California



(b) Georgia

Fig. 7 Indian male migrants by age group, California and Georgia, 2001–2018

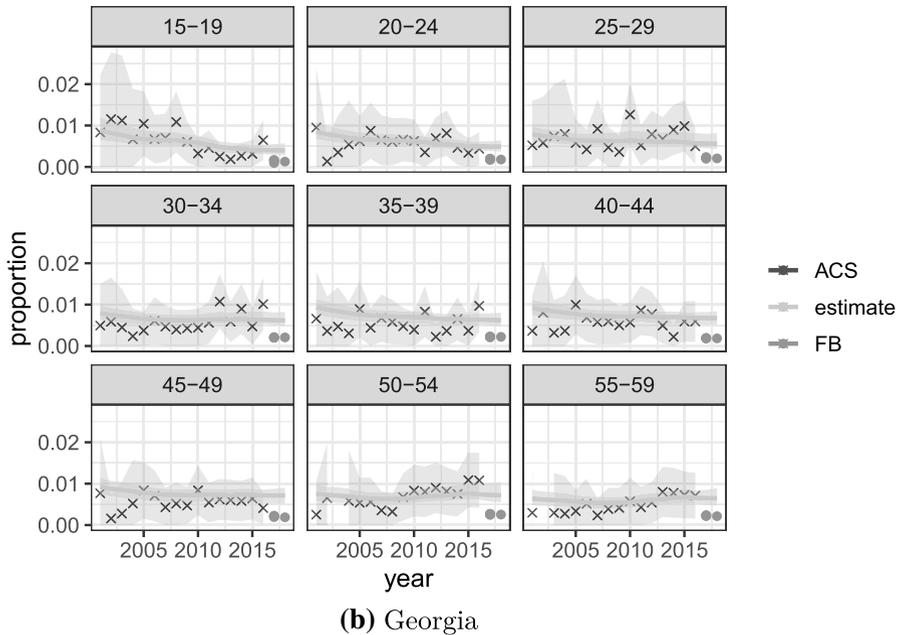
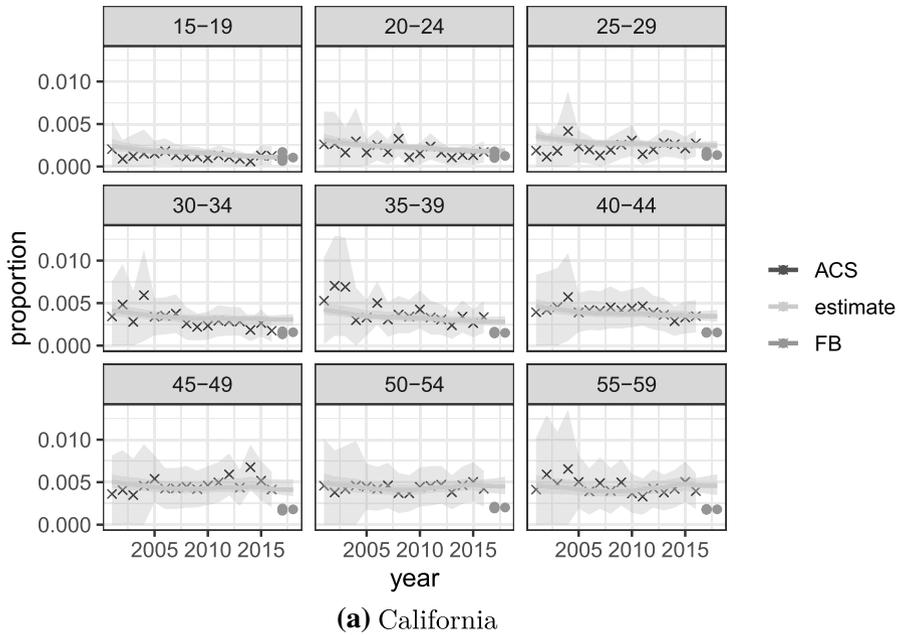


Fig. 8 German male migrants by age group, California and Georgia, 2001–2018

Table 1 Overall RMSE by model and migrant origin

Model	Mexico	India	Germany
Moving average	0.01190	0.00485	0.00325
Facebook	0.01690	0.00800	0.00376
ACS	0.01020	0.00447	0.00263
Combined	0.00988	0.00356	0.00261

population in the 30–39 year old age groups. The proportions are increasing over time, however, particularly in the 25–44 age bracket. Finally, for German male migrants in California and Georgia (Fig. 8), we see low and constant migrant proportions. The uncertainty around the ACS data is already relatively high, and so there is not so much of an increase in uncertainty in the final two years.

Validation

We evaluated the performance of the Bayesian model compared to other reasonable forecasting alternatives. To do this, we ran the model on data from 2001 to 2016, and forecast migration stocks in 2017. We then compared these forecasts to the actual ACS data in 2017. We compared the accuracy of the Bayesian model forecast to forecasts produced by three other models:

1. *Three-year moving average* of the ACS data. This is one of the simplest options available and does not require the Facebook data or any statistical modeling.
2. *Facebook data only* Estimates are based just on the available Facebook data in 2017, after it has been adjusted for biases.
3. *ACS time series model* Here, we ran the Bayesian hierarchical time series model described in the "Model" section above, but just using data from the ACS (no Facebook).

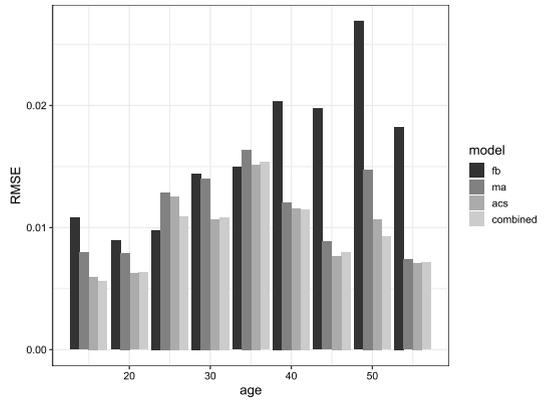
In order to assess model performance, we compare the root mean squared error (RMSE):

$$\text{RMSE} = \sqrt{\frac{\sum_n (\hat{p}_{g,2017} - p_{g,2017}^{\text{ACS}})^2}{N}} \quad (19)$$

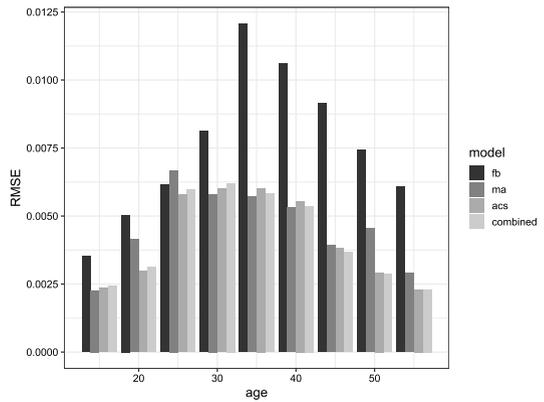
where $\hat{p}_{g,2017}$ is the estimated proportion of migrants from a particular group g , $p_{g,2017}^{\text{ACS}}$ is the equivalent proportion from the ACS and N is the size of the group. Here, the g can refer to any combination of age group, state and migrant origin.

Table 1 shows the overall RMSE for the four models for Mexican, Indian and German migrants. The main result is that in each of the three migrant groups, the Bayesian model presented (which combines the ACS and bias-adjusted Facebook data and thus is referred to as the 'combined model'), produces the lowest RMSE and thus the most accurate forecasts. The overall results also illustrate that the Bayesian hierarchical time series model produces substantially more accurate forecasts compared to a simple

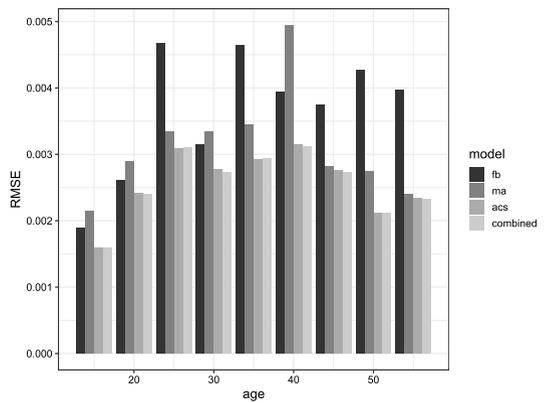
Fig. 9 RMSE by age group and model



(a) Mexico



(b) India



(c) Germany

moving average or the bias-adjusted Facebook data alone, producing RMSEs that are up to an order of magnitude smaller. This gain in accuracy is much larger than the gain moving from ACS-only to the combined model, although there is still a gain in each case.

Figure 9 illustrates the RMSE by age group and model type for each of the three migrant groups. Similar plots by state can be found in "Appendix C". For Mexico, there is generally an incremental decline in the RMSE moving from the Facebook-only model, to moving average, to ACS, to the combined model. For all but the Facebook-only model, the RMSE is highest in the 30–34 year old age group, which is also where the proportion of migrants is highest (see Fig. 6). For India (Fig. 9b), the RMSE is particularly high from the Facebook-only model. For Germany, the gain in accuracy in moving to the hierarchical time series set-up is most noticeable (that is, the improvement over the 'ACS' model compared to the Facebook-only or moving average). This is most likely related to the fact that the proportions of German migrants are in general a lot lower than for Mexico or India, and so there are noticeable gains in pooling information across state, age and time. It should be noted that the combined model does not always have a lower RMSE for every age and state combination, but as shown in Table 1, the combined model performs better overall for every migrant group.

To summarize, the validation exercise comparing the 1-year-out predictions from a range of models to the migrant proportions reported in the ACS illustrates both (i) the strength of the proposed Bayesian hierarchical time series model as a general framework, and (ii) the additional information obtained from including up-to-date Facebook data compared to just historical ACS data alone.

Discussion

As the size and frequency of migration movements continue to increase worldwide, new sources of data are being considered in order to better understand both historical and future migration trends. There is a growing body of work considering the feasibility of using social media data to achieve such goals, from platforms such as Facebook, Twitter and LinkedIn. While the granularity of social media use varies widely (from individual geo-tagged tweets, to aggregated advertising demographic data, as is in this paper), the common challenges of using such data remain: firstly, to adequately adjust for known biases in the social media data, primarily as a consequence of the non-representativeness of the population of social media users; and secondly, to meaningfully combine information from social media data with information from more traditional data sources, such as surveys or censuses.

In this paper we presented a statistical framework to achieve these goals in the context of producing short-term projections of migrant stocks in the United States. The model includes a bias-adjustment process of the Facebook data, and a 'principal components time series' model, which allows for the projection of trends in stocks into the future, considering both Facebook and ACS data. The model allows for different types of uncertainty around the different data sources, and shares information on migration trends over time and pools across geographic space. Illustrative results were presented for three separate migrant groups:

Mexicans, Indians, and Germans. The results of the validation exercise, comparing projections with 2017 ACS data, suggest that the proposed model improves prediction of short-term trends when compared to viable alternatives.

The validation exercise illustrated the substantial gain in accuracy achieved when moving to the Bayesian hierarchical time series model, regardless of whether or not the Facebook data were included. While the benefits of including the Facebook data in this particular case were relatively marginal, more generally the Facebook data has the advantage of being up-to-date and essentially available in real time. Thus, in a situation of a ‘shock’, such as a natural disaster or other event, the collection of Facebook data allows for the production of a timely estimate of the effects of that shock on migration. The combination of these data with past trends allows for the identification of surprising increases or decreases, that are out of the expected bounds based on historical patterns.

There are several limitations of the proposed model, which naturally lead into avenues for future work. Firstly, the bias-adjustment model assumes that the systematic bias in the Facebook data (by age and state) is constant over time. In reality, it is reasonable to believe that the biases in the Facebook data are changing over time, as the composition of the underlying Facebook population changes. The relationship between the age/location composition of the Facebook and the actual population (as measured by the ACS) could be investigated in future work. Secondly, the bias-adjustment model also assumes that the non-sampling error is constant over Facebook’s ‘waves’ of data collection—that is, sources of error that include changes in how the reach population is calculated, or other computational reasons, are assumed to be constant. In practice, and in other work using these data (Alexander et al. 2019), we have observed that this is probably not the case, and needs to be further investigated to better understand non-migration-related fluctuations over time. While we only consider two data sources in this paper, the general statistical framework could easily be extended to include information from other sources. Future work could also include taking advantage of the rich demographic and socioeconomic data available through the Facebook Advertising Platform, including information on education and occupation.

While this work focused on a model for the estimation of migrant stocks, the philosophy of combining social media data with more traditional data sources in one statistical framework—allowing for different sources of uncertainty—can be readily extended to model other demographic indicators. Indeed, the underlying time series model is itself an extension of principal component techniques that were previously used in demography to study mortality and fertility. Our framework shows the strength of combining more traditional demographic modeling techniques, survey data, and novel social media data to gain insights into underlying population processes.

In conclusion, we developed a statistical method for combining social media data with survey data to produce nowcasts of migrant stocks in the United States. We illustrated the method on three migrant groups in the United States with differing patterns. The statistical framework presented could be extended to include other data sources or nowcast other migration indicators in different contexts.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long

as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Bias Adjustment Plots for Migrants from India and Germany

See Figs. 10 and 11.

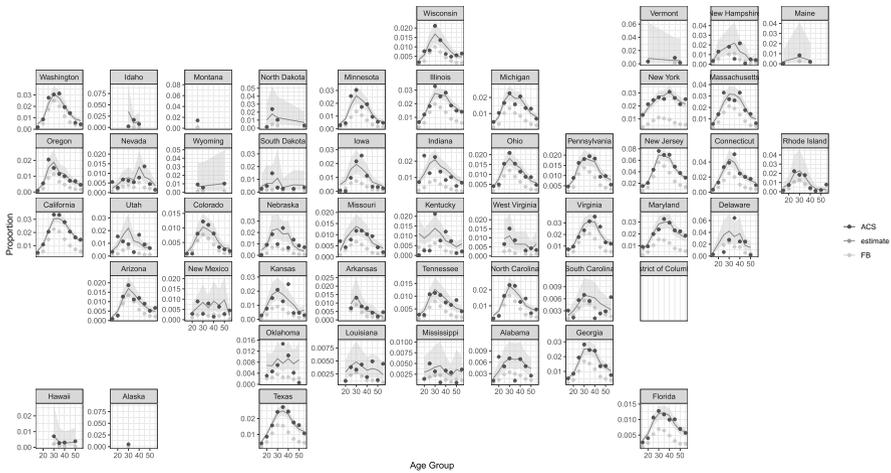


Fig. 10 Bias adjustment of Facebook data for Indian migrants

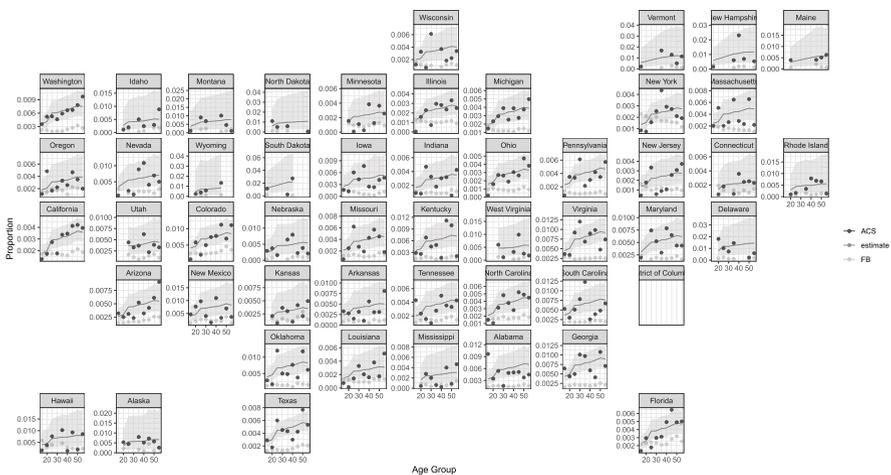


Fig. 11 Bias adjustment of Facebook data for German migrants

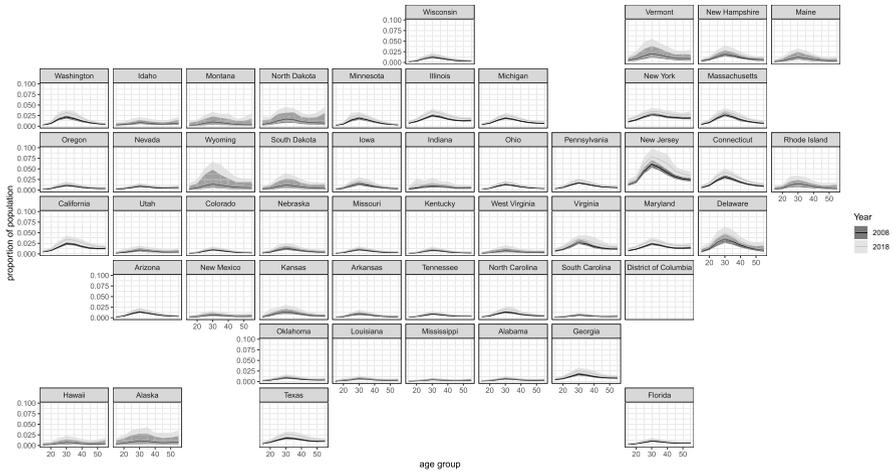


Fig. 12 Estimated and projected age distributions of Indian migrants by state, 2008 and 2018

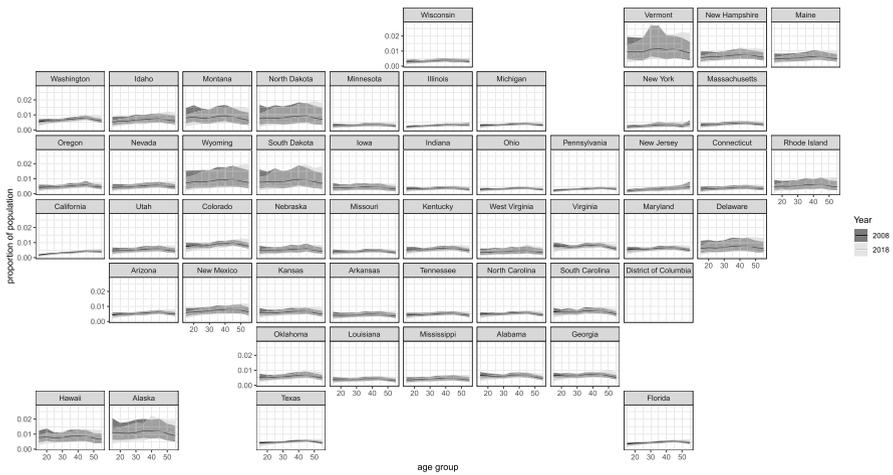


Fig. 13 Estimated and projected age distributions of German migrants by state, 2008 and 2018

Age Distributions in 2008 and 2018 for Indian and German Migrants by State

For India (Fig. 12) we see that the proportions across age groups have generally increased over the decade, with relatively high proportions concentrated in the northeast region of the country. Unlike Mexico, the age distribution is fairly constant, with the highest proportions generally being in the 30–34 age group. The German male migrant populations (Fig. 13) by state show relatively flat, low and unchanging levels over the decade 2008–2018.

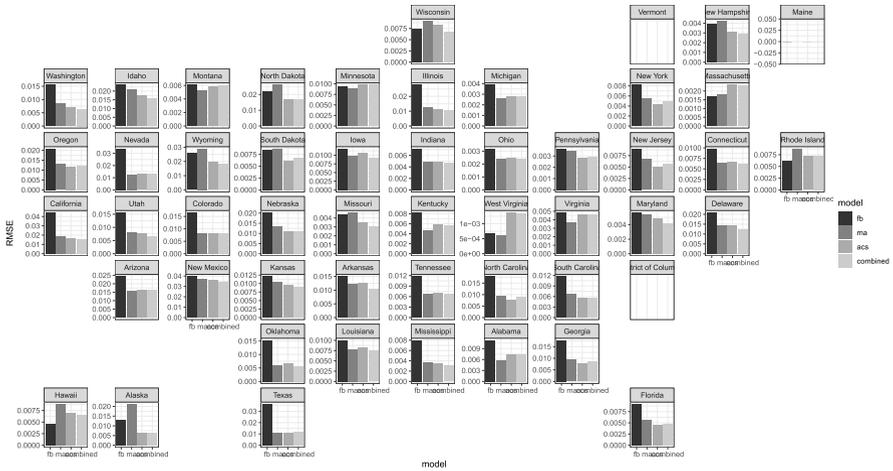


Fig. 14 RSME by state and model type for Mexican migrants in 2018

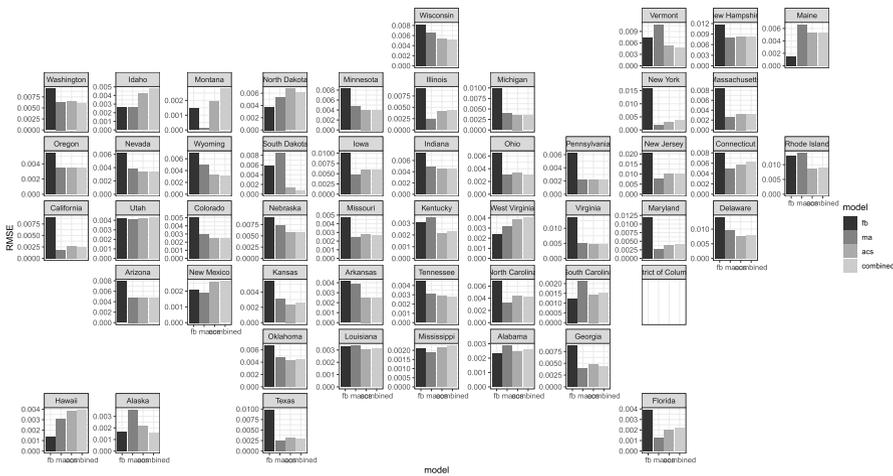


Fig. 15 RSME by state and model type for Indian migrants in 2018

Validation Results by State

The figures below show the RMSE for each model, by state of destination in the US, for Mexican, Indian and German migrants. The figures show that, in general, the combined model outperforms the other three alternative models. For Mexico (Fig. 14), this is particularly the case for the states with relatively high proportions of Mexican migrants. Results are more variable for Indian migrants (Fig. 15), with the combined model performing relatively poorly in the Southern states but well in the mid-West states. For German migrants (Fig. 16), results are again variable, but relatively good in the Southern states.

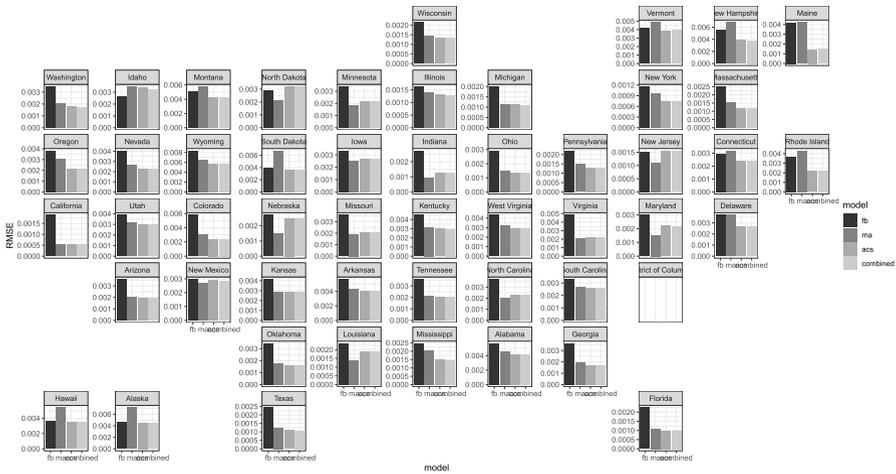


Fig. 16 RMSE by state and model type for German migrants in 2018

References

Alexander, M., Zagheni, E., & Barbieri, M. (2017). A flexible bayesian model for estimating subnational mortality. *Demography*, 54(6), 2025–2041.

Alexander, M., Zagheni, E., & Polimis, K. (2019). The impact of hurricane maria on out-migration from puerto rico: Evidence from facebook data. *Population and Development Review*, 45(3), 617–630.

Araujo, M., Mejova, Y., Weber, I., & Benevenuto, F. (2017b). Using facebook ads audiences for global lifestyle disease surveillance: Promises and limitations. WebSci '17, New York, NY, USA. ACM.

Blumenstock, J. E. (2012). Inferring patterns of internal migration from mobile phone call records: Evidence from rwanda. *Information Technology for Development*, 18(2), 107–125.

Cesare, N., Lee, H., McCormick, T., Spiro, E., & Zagheni, E. (2018). Promises and pitfalls of using digital traces for demographic research. *Demography*, 55(5), 1979–1999.

Engels, R. A., & Healy, M. K. (1981). Measuring interstate migration flows: An origin-destination network based on internal revenue service records. *Environment and Planning A*, 13(11), 1345–1360.

Fatehkia, M., Kashyap, R., & Weber, I. (2018). Using facebook ad data to track the global digital gender gap. *World Development*, 107, 189–209.

Ferrari, L., Rosi, A., Mamei, M., & Zambonelli, F. (2011). Extracting urban patterns from location-based social networks. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-Based Social Networks* (pp. 9–16).

Foulkes, M., & Newbold, K. B. (2008). Using alternative data sources to study rural migration: examples from illinois. *Population, Space and Place*, 14(3), 177–188.

Gabrielli, L., Deutschmann, E., Natale, F., Recchi, E., & Vespe, M. (2019). Dissecting global air traffic data to discern different types and trends of transnational human mobility. *EPJ Data Science*, 8(1), 1–24.

Garcia, D., Kassa, Y. M., Cuevas, A., Cebrian, M., Moro, E., Rahwan, I., et al. (2018). Analyzing gender inequality through large-scale facebook advertising data. *Proceedings of the National Academy of Sciences*, 115(27), 6958–6963.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*. New York: Chapman and Hall/CRC.

Landau, L. B., & Achiume, E. T. (2017). International migration report 2015: Highlights. *Development and Change*, 48(5), 1182–1195.

Lee, R. (2000). The Lee–Carter method for forecasting mortality, with various extensions and applications. *North American Actuarial Journal*, 4(1), 80–91.

- Lee, R. D., & Carter, L. R. (1992). Modeling and forecasting us mortality. *Journal of the American Statistical Association*, 87(419), 659–671.
- Li, N., Lee, R., & Tuljapurkar, S. (2004). Using the Lee–Carter method to forecast mortality for populations with limited data. *International Statistical Review*, 72(1), 19–36.
- Noulas, A., Scellato, S., Mascolo, C., & Pontil, M. (2011). An empirical study of geographic user activity patterns in foursquare. In *Fifth international AAAI conference on weblogs and social media*.
- Pestre, G., Letouzé, E., & Zagheni, E. (2020). The abcde of big data: Assessing biases in call-detail records for development estimates. *The World Bank Economic Review*, 34, S89–S97.
- Plummer, M. et al. (2003). Jags: A program for analysis of bayesian graphical models using gibbs sampling. In *Proceedings of the 3rd international workshop on distributed statistical computing*, Volume 124, pp. 1–10. Vienna, Austria.
- Rampazzo, F., Zagheni, E., Weber, I., Testa, M. R., & Billari, F. (2018). Mater certa est, pater numquam: What can facebook advertising data tell us about male fertility rates? In *Twelfth International AAAI Conference on Web and Social Media*.
- Renshaw, A. E., & Haberman, S. (2006). A cohort-based extension to the Lee-Carter model for mortality reduction factors. *Insurance: Mathematics and Economics*, 38(3), 556–570.
- Ruggles, S., Fitch, C. A., Hall, P. K., & Sobek, M. (2000). *Ipums-usa: Integrated public use microdata series for the United States. Handbook of international historical microdata for population research* (pp. 259–284). Minneapolis: Minnesota Population Center.
- Schmertmann, C., Zagheni, E., Goldstein, J. R., & Myrskylä, M. (2014). Bayesian forecasting of cohort fertility. *Journal of the American Statistical Association*, 109(506), 500–513.
- State, B., Rodriguez, M., Helbing, D., & Zagheni, E. (2014). Migration of professionals to the us. In *International Conference on Social Informatics*, pp. 531–543. Springer.
- US Census Bureau (2020). PUMS accuracy. https://www2.census.gov/programs-surveys/acs/tech_docs/pums/accuracy/.
- Wiśniowski, A., Smith, P. W., Bijak, J., Raymer, J., & Forster, J. J. (2015). Bayesian population forecasting: Extending the lee-carter method. *Demography*, 52(3), 1035–1059.
- Yildiz, D., Munson, J., Vitali, A., Tinati, R., & Holland, J. A. (2017). Using twitter data for demographic research. *Demographic Research*, 37, 1477–1514.
- Zagheni, E., Garimella, V. R. K., Weber, I., & State, B. (2014). Inferring international and internal migration patterns from twitter data. In *Proceedings of the 23rd International Conference on World Wide Web* (pp. 439–444).
- Zagheni, E., & Weber, I. (2012). You are where you e-mail: using e-mail data to estimate international migration rates. In *Proceedings of the 4th annual ACM web science conference* (pp. 348–351).
- Zagheni, E., Weber, I., & Gummedi, K. (2017). Leveraging facebook’s advertising platform to monitor stocks of migrants. *Population and Development Review*, 43(4), 721–734.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.