

# Online social integration of migrants: Evidence from Twitter

Jisu Kim \*,<sup>†</sup> Soazic Elise Wang Sonne<sup>‡</sup>, Kiran Garimella<sup>§</sup>,  
André Grow<sup>†</sup>, Ingmar Weber \*\* and Emilio Zagheni<sup>†</sup>

<sup>†</sup>Max Planck Institute for Demographic Research, Konrad-Zuse-Straße 1, Rostock, 18057, Germany, <sup>‡</sup>World Bank, 1818 H Street, NW Washington DC, 20433, USA, <sup>§</sup>Rutgers university, 4 Huntington Street, New Brunswick, New Jersey, 08901, USA and \*\*Saarland University, Saarland Informatics Campus, Saarbrücken, 66123, Germany

\*Corresponding author. Email: kim@demogr.mpg.de

## Abstract

As online social activities have become increasingly important for people's lives, understanding how migrants integrate into online spaces is crucial for providing a more complete picture of integration processes. We curate a high-quality data set to quantify patterns of new online social connections among immigrants in the United States. Specifically, we focus on Twitter and leverage the unique features of these data, in combination with a propensity score matching technique, to isolate the effects of migration events on social network formation. The results indicate that migration events led to an expansion of migrants' networks of friends on Twitter in the destination country, relative to those of similar users who did not move. Male migrants between 19 and 29 years old who actively posted more tweets in English after migration also tended to have more local friends after migration compared to other demographic groups, indicating the impact of demographic characteristics and language skills on integration. The percentage of migrants' friends from their country of origin decreased in the first few years after migration and increased again in later years. Finally, unlike for migrants' friends' networks, which were under their control, we did not find any evidence that migration events expanded migrants' networks of followers in the destination country. While following users on Twitter in theory is not a geographically constrained process, our work shows that offline (re)location plays a significant role in the formation of online networks.

**Keywords:** big data, gender, international migration, social integration, Twitter

## 1. Introduction

The integration of migrants is a multidimensional process that involves various aspects of migrants' lives and of host societies, including the economic, educational, cultural, civil, and other characteristics of both the host and the sending countries. While understanding

<https://doi.org/10.1093/migration/mnad017>

© The Author(s) 2023. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

integration processes is challenging given their complexity, it is nonetheless important, as integration enables migrants to transition more smoothly to their new home. In addition, from the point of view of the host country, the successful integration of migrants creates a more harmonious society and reduces inequality, social conflicts, and social polarisation (Gowricharn 2002; Van der Meer and Tolsma, 2014; Abdurahman 2017). Measures of integration often rely on employment, education, and health statistics or related information mostly based on surveys. In some contexts, field experiments have been carried out to study the social integration of migrants (see, for instance, Kearns and Whitley 2015; Yoon, 2022; Wessendorf and Phillimore 2019; Glorius et al. 2020). However, these established data sources have limitations. For example, the sample size of these surveys is often small. Moreover, conducting a survey is costly, and the limited scope of such surveys may reduce the diversity of immigrants' countries of origin that can be studied. In recent years, a growing number of studies on the integration of migrants have used social media data, such as data from Twitter (Sirbu et al. 2021; Kim et al. 2022a, 2022b, 2023), to complement analyses based on survey data. Researchers have shown that because of the detailed information on social networks that can be obtained from social media data—including information on users' friends and followers, and on the opinions and information the users shared—analysing these data can provide new insights into integration processes. Most importantly, as online social activities have become important to people's lives and well-being, understanding the online space has become crucial for researchers seeking to paint a more complete picture of integration processes and to understand the social determinants of inequalities. For many users of an online social platform such as Twitter, who tend to be younger than the general population, the value of their online community has increased, and may even be comparable to that of their offline community (Lehdonvirta and Räsänen 2011).

Among the various aspects of migrant integration, we focus here on social integration. More specifically, we are interested in investigating whether recent immigrants have integrated socially into online spaces through Twitter. We aim to answer the following questions in this study: *After migration, does the number of friends/followers an immigrant has in an online space who are based in the destination country increase? Do immigrants lose their connections with their friends/followers in their origin country? How do the characteristics of migrants whose number of online friends living in the destination country increase differ from those of migrants who do not experience such an increase?*

It is important to study the social integration of migrants, as it is one of the main dimensions of migrants' experiences, and it can facilitate various aspects of their integration process. Migrants' social networks often become their main sources of information, ranging from trivial information like 'where to go for a nice dinner in the area'; to more valuable information, such as 'my company is looking for an employee'. Thus, the barriers to information on the local community migrants often encounter may be removed through the sharing of knowledge by local friends (Rauch 2001; Rauch and Trindade 2002; Comola and Mendola 2015; Chiswick and Wang 2016). At the same time, understanding migrants' social connections to people in their country of origin is also crucial, as it can shed light on immigrants' economic activities, such as the sending of remittances; and on the diffusion of culture and global trade at the macro level (Rauch 2001; Docquier and Rapoport 2012; Bahar et al., 2022). In this study, we focus on immigrants residing in

the USA. Studying the social integration of immigrants in the USA is particularly interesting, as it is known for being ‘the most diverse country in the world’.<sup>1</sup>

Our methodology includes ensuring a careful design of the data collection process, identifying migrants using geo-tagged tweets, and, finally, determining the dates when migrants made social connections. Using Twitter data that we obtained from archive.org, we start by identifying migrants on Twitter along with their likely date of migration, based on our definition of a migrant: namely, ‘a user who tweeted at least one geo-tweet per month in one country (home country) for 12 consecutive months and one geo-tweet per month for 12 consecutive months in another country (destination country)’. This approach allows us to identify recent immigrants. The day a migrant started to tweet in their destination country is identified as the person’s date of migration. Secondly, using the method developed by Meeder et al. (2011), we infer the dates when social links were established to study how the composition of the migrant’s social network changed *before* and *after* migration. In this work, we prioritise having reliable and high-quality information on users over the size of the data. Hence, we match migrants with non-migrants with similar characteristics on Twitter using the 1:1 propensity score matching technique (Stuart et al. 2011). This matching technique allows us to also provide a *placebo* date of migration to non-migrants using the same date of actual migration for the matched migrants. This enables us to compare the friendship patterns of these two groups, and to identify, in a quasi-causal sense, the effects of the migration event on the expansion of migrants’ social links on Twitter. With the enriched data that we curated, we then examine the online social integration patterns of both migrants and non-migrants by looking at how many new US-based and origin-based friends/followers were added to the migrants’ networks after their likely date of migration.

The rest of this article is organised as follows. In the next section, we provide background information, drawing from the sociological literature that deals with social integration, and from the emerging literature that leverages digital trace data. In Section 3, we describe the data and methods we use. In several subsections, we provide details on our data and explain how we identified migrants, validated the data, detected new social connections, and matched migrants and non-migrants. In Section 4, we present our results. We offer our concluding remarks in Section 5.

## 2. Background

The topic of migrant integration has long attracted the interest of social scientists across a number of disciplines, including sociology, economics, anthropology, and psychology, as the process of integration involves a broad spectrum of individual and societal characteristics. One of the concepts related to integration that sociologists have studied extensively is that of ‘social capital’ (Bourdieu 1986; Portes 1998). Bourdieu and Wacquant defined this concept as follows: ‘Social capital is the sum of the resources, actual or virtual, that accrue to an individual or a group by virtue of possessing a durable network of more or less institutionalised relationships of mutual acquaintance and recognition’ (Wacquant and Bourdieu 1992). The core idea is that social connections are valuable assets that enable members of a society to forge mutually beneficial relationships. Social capital facilitates

the exchange of information and knowledge and increases economic, social, and cultural opportunities for individuals by, for example, providing them with job security and social ties. Barriers to obtaining relevant information that immigrants would otherwise face are reduced through having ties in the host society (Schiff 1992; Putnam 1995; Uslaner 2003).

The existing body of literature on the social integration of migrants has examined its mechanisms from different societal points of view while relying mainly on surveys, interviews, or questionnaires. These surveys usually include elements specific to the contexts where they were conducted. For instance, to examine the determinants of social integration in the context of internal migration in China, Chen and Wang (2015) asked respondents 'whether they actively participated in local community activities, whether they had adapted to local social norms and customs and whether they are socially connected with the local community'. Based on the responses to these questions, the authors found that education and receiving higher income were key factors for social integration, as they provided migrants with the time and resources they needed to integrate into local society. They also reported that migrants from geographically close regions were more likely to feel socially integrated. Similarly, Kearns and Whitley (2015) examined different aspects of social integration from the point of view of trust, reliance, safety, and a sense of community in the UK. They then analysed the associations between these social integration factors and the effects of time, place, and migrant type; and of functional factors, such as educational qualifications, language proficiency, and employment status. The respondents were asked, for example, to what extent they felt that their 'neighbourhood is a place where neighbours look out for each other'; and to what extent they agreed with the following statements: 'It is likely that someone would intervene if a group of youths were harassing someone in the local area'; and 'someone who lost a purse or wallet around the area would be likely to have it returned without having anything missing'. The findings of this study are similar to those of Chen and Wang (2015): that is, all functional factors were positively associated with social integration. In particular, they found that higher educational levels were positively associated with greater use of local amenities, and that being employed allowed migrants to access the resources they needed to interact with other members of the community. As the authors were looking at migrants in the UK, they also found that the migrants' language proficiency was positively associated with their level of trust, as it facilitated their communication with their neighbours. Similarly, Yoon (2022) examined the current integration status of North Korean migrants in South Korea. Among other indications of integration, the study investigated the social connections of these migrants, which the author defined using three criteria: number of South Korean friends ('social bridges'), number of North Korean friends who could help in an emergency ('social bonds'), and access to support services from the government ('social links'). The results indicated that North Koreans had a strong sense of belonging and trust. Indeed, the author observed that 'North Korean migrants had a strong belief that they could reach the same status as South Koreans', and were often identified as South Korean citizens.

These are examples of studies that have provided detailed information on migrants' experiences with social integration. However, there are also drawbacks to relying on surveys or interviews only. For instance, most studies of integration have been conducted in specific settings (e.g., in their neighbourhoods (Chen and Wang, 2015), migrants in the

UK (Kearns and Whitley 2015; Wessendorf and Phillimore 2019), and North Koreans in South Korea (Yoon, 2022). This means that the geographical coverage of existing studies is often limited, as they tend to be concentrated on country-specific contexts. Moreover, in a survey, respondents may feel compelled to choose an answer even if their views do not fit any of the choices that they are given, or to agree with given statements (e.g., acquiescence bias). Additionally, information on respondents' social networks can be difficult to obtain. In many cases, only a small part of respondents' social groups are captured, or their networks can differ depending on when the interview was conducted. For example, if a migrant had met a group of friends recently, they may remember them the most clearly, which could influence their survey responses.

In recent years, social media data such as data from Facebook and Twitter have been leveraged to study a number of dimensions of migrant integration (Dubois et al. 2018; Stewart et al. 2019; Chi, 2020; Mazzoli et al. 2020; Bailey et al. 2022; Kim et al. 2022a, 2022b, 2023). The use of social media data can overcome some of the limitations of traditional data sources and can complement them. For instance, Dubois et al. (2018) and Stewart et al. (2019) extracted data from the Facebook advertising platform to study cultural assimilation. In both studies, the authors were able to obtain aggregated information on individual users' interests in specific areas through Facebook data. Dubois et al. (2018) quantified the level of assimilation of migrants in Germany by introducing a score that provides a proxy for the migrants' assimilation to the local population's interests. Using these data, the authors were able to list the detailed interests of Germans in various domains, including music, film, city, companies, and so on. They found that European migrants in Germany had a higher assimilation score than Arabic-speaking migrants and Turkish speakers. Going deeper into the data, they also found that among Arabic-speaking migrants, men, university graduates, and individuals aged 18–24 years had higher assimilation scores. Focusing on Mexican immigrants of both the first and the second generation in the USA, Stewart et al. (2019) looked at changes in their musical tastes as a specific case study of cultural assimilation. Using assimilation score metrics similar to those applied in Dubois et al. (2018), they examined the extent to which the migrants' musical tastes had converged to those of the host population as evidence of their cultural assimilation. They were able to distinguish between the first and the second generation of migrants through the option of targeting specific users on the Facebook advertising platform. In general, they found that Mexican migrants of both the first and the second generation showed a high degree of assimilation to the musical tastes of Anglos and African Americans. They also found that younger migrants and the first generation of Mexican migrants had higher assimilation scores. However, in contrast to previous findings, they found that women assimilated more than men when it came to musical tastes. In addition, although the authors concentrated on the case study of Mexican immigrants, they also showed that their methodology could be used to study the wider immigrant population in the USA. Also using Facebook data, Chi et al. (2020) studied the relationship between international social ties and human mobility through social network analysis. In this study, the authors had access to the social networks of individual users, which were not limited to one geographic setting but could be worldwide. The results showed that long-term international migrants created 83% of the international ties. Most recently, Bailey et al. (2022) conducted a study on the social integration of Syrians in Germany,

which also utilised Facebook data. The study examined three elements associated with social integration, namely the number of connections with native German friends in the area of residence, the utilisation of the German language, and involvement in local social groups. The findings indicated that there are regional disparities in how local institutions influence the social integration of Syrians in their respective communities.

Cultural integration has also been studied using Twitter data by Kim et al. (2022a). They looked at the host country-specific hashtags used by migrants as a proxy for *destination attachment* level and at the origin country-specific hashtags used by migrants as a proxy for *origin attachment* level, while also taking into consideration the preservation of links to the migrants' country of origin. With *destination attachment* and *origin attachment* indexes, they summarised integration patterns into four categories: marginalisation, separation, assimilation, and integration. Their results highlighted several factors that were positively related to different integration patterns. In particular, they observed that language proximity and the distance between the origin and the destination country played important roles in how the immigrants maintained their links to their origin country, and in how they integrated into the destination society. In addition, the authors were able to observe the behaviours of different migrants in various countries, including immigrants in the USA and Great Britain, and Italian emigrants in various destination countries. These kinds of patterns are difficult to observe using traditional data sources. Mazzoli et al. (2020) employed Twitter data to examine the spatial integration of migrants. In their study, the authors showed how Twitter data can be used to observe the places most commonly visited by Venezuelan migrants based on their tweet activity after working hours. Drawing on these observations, they built a segregation index by directly comparing the population distribution of migrants to natives, and the spatial segregation of Venezuelan migrants in Bogotá, Lima, and São Paulo. The results of their segregation index also showed that Venezuelan migrants were spatially segregated from the locals.

In this article, we further advance the literature that uses Twitter data to study the integration of migrants. In contrast to the above-mentioned studies, we focus on a different dimension of migrant integration (i.e., social integration) that has not been previously explored in the literature. Thanks to the availability of Twitter data, we are able to obtain information on the social networks of individual users. These data are different from publicly available Facebook data, which contain no information on the users' social networks. Our analyses build on the work of Chi et al. (2020), as we are also interested in the composition of the social networks of immigrants. Importantly, however, we have collected and harmonised information about the temporal order of the social links based on when they were created. Using these data, we have implemented various methodologies to build a high-quality data set that enables us to quantify how many new social connections migrants added on Twitter before and after migration. In addition, we further advance the literature by inferring whether the migrants' newly added social connections were or were not local friends/followers based on information provided on their profiles. These results allow us to determine to what degree migrants were socially integrated (on the Twitter platform) into the destination country. Moreover, in contrast to previous studies, we are interested in exploring the social integration patterns of migrants in an online setting, which is an important perspective given that the amount of time people spend on online social platforms has increased dramatically in the last decade.

### 3. Data and method

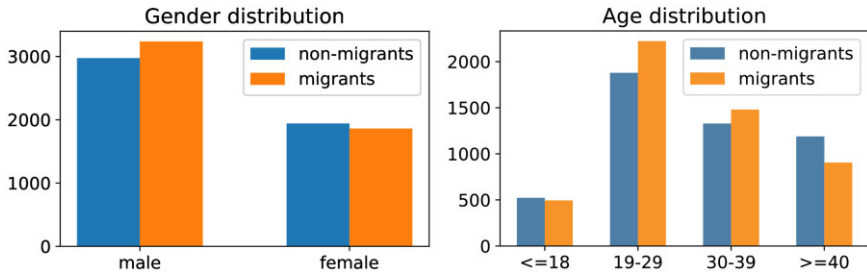
In order to study the online social integration of migrants, we focus on one specific platform: Twitter. We chose Twitter because the data available for this platform enable us to consider changes in the structure of the users' social networks *before* and *after* migration. The methodology includes: (1) the careful design of the data collection approach and of data pre-processing; (2) the high-quality identification of migrants using geo-tagged tweets; and (3) a solid statistical framework for analysing pre- and post-changes with respect to a properly selected control group.

#### 3.1 Data

As a first step, we retrieved Twitter data from a large and publicly available archive of tweets (<https://archive.org/details/twitterstream>) that includes a random sample of 1% of all tweets.<sup>2</sup> This database of Twitter tweets includes, among other information, the text of each tweet, the uniquely identifying user IDs, the geo-location—if enabled by the user—of where the tweet was sent from, the language of the tweet, and the date and time when the tweet was created. As it is most crucial for our analysis to obtain location information, we filtered out tweets from 2017 to 2018 that were not geo-tagged. At this stage, the database included about one million users who had geo-tagged tweets from the time period we are interested in. To obtain the timeline of the geo-tagged tweets of these one million users, we collected their most recent 3,200 tweets, their profile and social network information, as well as both their 'friends' (Twitter users that a specific user follows) and 'followers' (Twitter users that follow a specific user) using the Twitter API.<sup>3</sup> The time series of geo-tagged tweets of all of the users considered were important information for us, as they helped us to determine whether a user was or was not a migrant.

In order to augment the data with additional information, we used the deep learning algorithm<sup>4</sup> developed by Wang et al. (2019) to estimate the demographic characteristics of users (i.e., age and gender where the gender classification provides a binary category; and age in four broad categories:  $\leq 18$ , (18, 30), (30, 40), and (40, 99)). Note that age was grouped into four categories following the age ranges used in censuses and surveys, but also because producing estimates at a higher level of granularity would result in extremely high uncertainty, even with the help of human coders (Wang et al. 2019). This algorithm is used to input the user's profile photo, biography, screen names, and names from the Twitter profile. To be more specific, the authors combined three models (image model, text model, and multimodal model) to process different types of information needed to estimate demographic characteristics. First, the image model was used to process the profile image using DenseNet (Huang et al. 2017). Secondly, the text model was used to process the text inputs from the biography and the username using the Long-Short Term Memory (LSTM) architecture (Hochreiter and Schmidhuber 1997). Finally, the multimodal model was used to combine the first two models 'into a new model with a modality drop-out layer' that allowed them to better regularise the model, including in cases in which not all of the input sources were available. To put it simply, the model picked up hints from the biographical information of users based on words or images that signal gender and age, such as 'mother of two wonderful kids' or '26 y/o dude'. The authors





**Figure 1.** Distribution of demographic characteristics of users by migration status.

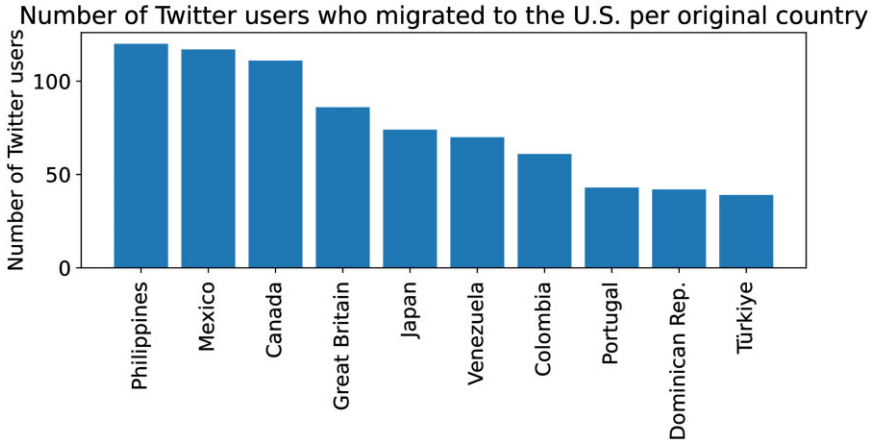
reported that when this algorithm's performance score on Twitter data was compared with other state-of-the-art algorithms, such as the Microsoft classifier<sup>5</sup> and Face++,<sup>6</sup> it outperformed in all three categories. To be more specific, the gender classification obtained a Macro-F1 score of 0.92, 0.9 for organisation status classification, and 0.43 for age classification. Another advantage of this algorithm is that it also takes 32 languages<sup>7</sup> into account, which allows us to obtain the demographic information of users who are from various countries, without requiring any preprocessing of images or languages prior to inputting the data into the algorithm. As Fig. 1 shows, males made up the majority of our data set of Twitter users, regardless of migration status. Looking at the age distribution, we see that a sizable proportion of our population of users were between 19 and 29 years old, regardless of migration status. We also observe that more migrants were concentrated in the central age groups, while more non-migrants were concentrated in the 40+ age group.

As an additional approach, we examined the biographies of users on their profiles to extract any occupation-related information. While the number of job-related words was insufficient to be reported in our study, we did discover that musicians, engineers, producers, photographers, managers, directors, and travellers were some of the most frequent occupation-associated terms that we extracted.

### 3.2 Identifying migrants

Provided with the timeline of geo-tagged tweets, we identify users as migrants if these users 'have tweeted at least one geo-tagged tweet per month in one country (home country) for 12 consecutive months and one geo-tagged tweet per month for 12 consecutive months in another (destination) country'. In other words, we determine whether a user moved to a different country within the time frame of interest based on the tweet timeline of that user. For example, if a user who tweeted at least one geo-tagged tweet per month for 12 consecutive months from Korea was tweeting from the USA in the later period for 12 consecutive months, then we would identify this user as a migrant from Korea to the USA. Our definition of migrants is adapted from the United Nations, which defines a migrant as 'a person who moves to a country other than that of his or her usual place of residence for a period of at least a year'.<sup>8</sup> Following this definition, we look at users' geo-tagged tweets over a period of 12 consecutive months to ensure that we are not observing





**Figure 2.** Top 10 countries of origin of the migrants in the USA in the analysed Twitter data set.

short-term visitors, such as students or tourists. We note also that the definition of migrants employed here focuses on users who moved to the USA during the time period we are interested in. Thus, users who migrated in the earlier period and did not change their location thereafter would not be captured by this definition. During this process, we also identify the likely date of migration, which is the date when the user started tweeting from the destination country.

By applying this approach, we were able to identify about 5,000 migrants coming from 133 different countries who migrated to 154 countries. Of these destination countries, the USA attracted the largest number of migrants (1,300), followed by Great Britain (440) and Brazil (321). Thus, we identified 1,300 migrants to the USA, the majority of whom came from the Philippines, followed by Mexico, Canada, and the UK, as shown in Fig. 2.

In parallel, we randomly sampled about 5,000 non-migrants from the initial data (one million users), excluding the 5,000 identified migrants. These non-migrants were selected to serve as a control group in the later stages of the analysis. Here, we define non-migrants as the users who tweeted in one country only during the same window of observation as the migrants, irrespective of whether they were or were not migrants. However, in this process, we also label the users who tweeted from several countries but did not stay longer than two months as non-migrants, based on the assumption that they were making short trips, and had not migrated. Similarly, long-term migrants who migrated before their move can be inferred from their Twitter timeline were classified as non-migrants. Implicitly, this means that our analytical setup assumes that the characteristics of long-term migrants were similar to those of US residents. Like the migrants, most of the non-migrants in our data came from the USA, Brazil, Indonesia, and Great Britain.

### 3.3 Data validation

To validate the quality of both the inferred user demographics (age and gender), as well as the inferred date of migration, we manually checked 620 individual users in our data by

looking at their public profiles on Twitter. Here, we checked for clues in their name, profile photo, self-declared bio, and tweets related to their demographics. Furthermore, we looked for any signs that users had migrated, such as any mention of a ‘new home’, changes in the geo-tags of their tweets, or photos of what seems to be a new home location. Note that the number of individual users for whom we manually checked the age category, mostly because their profile photo was missing, was lower (555). With this information, we computed an F1 score for each identified category to measure the performance of our estimation and classification. The F1 score measures the performance of our estimation, comparing our classification with true data (i.e., the 620 individual users we manually checked). It provides a score between zero and one, with one signifying perfect precision and recall. The precision score tells us, for instance, how many males among all the males we identified were indeed male; while the recall score tells us, for instance, how many males among all the males we identified were correctly predicted. In the following, we look at the weighted average F1 score, which considers the number of samples from each class.

Starting with the demographic characteristics of the users in the table, we obtained a weighted average F1 score of 0.92 for our gender categories, demonstrating the validity of our gender estimation procedure (Table 1). Similarly, we obtained a weighted average F1 score of 0.85 for our age categories, which generally validates our age estimation procedure (Table 2). The F1 score for age was slightly lower than the F1 score for gender due to the relatively low precision score for the population under the age of 18 years. Lastly, we obtained an accuracy level of 62% for the correct identification of the migration date of the users in our data. Of the remaining 38%, about 52% were users for whom we lacked any clear information about whether they truly were or were not migrants, while the rest were users who either were correctly identified as migrants, but for whom the date of the migration was incorrectly estimated, or were short-term visitors (e.g., on vacation). It should be pointed out that despite the challenges involved in identifying migrants using Twitter data, as highlighted by [Armstrong et al. \(2021\)](#), we were able to accurately identify a considerable proportion of the migrants in our data.

We also note, however, that although individual users are less likely to be misclassified by manual checks than by trained algorithms, it is still possible that we introduced bias into the process. For instance, we may have incorrectly identified the age categories of users who looked young or old to a human observer, or of users who posted outdated

**Table 1.** Average precision, recall, and F1 scores for gender comparing our predicted class label of gender with true data (i.e., manually checked data)

	Precision	Recall	F1 score	Support
Female	0.93	0.85	0.89	229
Male	0.92	0.96	0.94	391
Accuracy	0.92	0.92	0.92	0.92
Macro avg	0.93	0.91	0.92	620
Weighted avg	0.92	0.92	0.92	620

**Table 2.** Average precision, recall, and F1 scores for age comparing our predicted class label of gender with true data (i.e., manually checked data)

	Precision	Recall	F1 score	Support
≤18	0.41	0.81	0.54	16
19–29	0.89	0.83	0.86	251
30–39	0.88	0.81	0.84	182
≥40	0.79	0.88	0.83	106
Accuracy	0.83	0.83	0.83	0.83
Macro avg	0.74	0.83	0.77	555
Weighted avg	0.85	0.83	0.84	555

photos. Moreover, we cannot rule out the possibility that certain self-presentation behaviours (For instance, use of multiple languages, or frequent changes of home locations) were tied to the migration event, imposing certain limitations. Hence, it is possible that in certain circumstances, an automatic evaluation would have performed better than we would otherwise expect.

### 3.4 Detecting new social connections

To study how the composition of a migrant's social network changed before and after migration, we need to understand when the social links were created. To do so, we applied the method developed by Meeder et al. (2011). In their study, they inferred the date of link creation based on two pieces of information: first, the date when the account was created; and, secondly, the list of friends<sup>9</sup>/followers,<sup>10</sup> which the Twitter API<sup>11</sup> provides in the reverse chronological order of link creation between two users,<sup>12</sup> as shown in Fig. 3. Based on this information, the date of link creation for an alter user  $u$  with the ego user  $v$  would be the maximum of the account creation date of all the users who followed the ego user  $v$  before the user  $u$ . This method was shown to have a high accuracy level, with the typical error for inferring the date and the time of link creation being up to a few minutes only. It should, however, be noted that this analysis was conducted using data from celebrities who had large networks with many followers and friends and that larger networks can provide tighter temporal bounds. In contrast, the users in our data did not have as many friends or followers as these celebrities. For this reason, we expect the precision of our estimates to be lower. Nevertheless, as we do not require exact time up to the second of the link creation date, this may not be an issue here.

During this process, we also collected information on the locations of the users' friends and followers at a country level. To obtain their locations, we relied on self-reported information on the users' Twitter profiles instead of on geo-tagged tweets. We did so because we assume that most users state their usual location on their profile, and do not change it when they go on short trips or vacations. Nevertheless, the profile location information is not fully reliable, as it is self-reported. Users are, for instance, free to state their



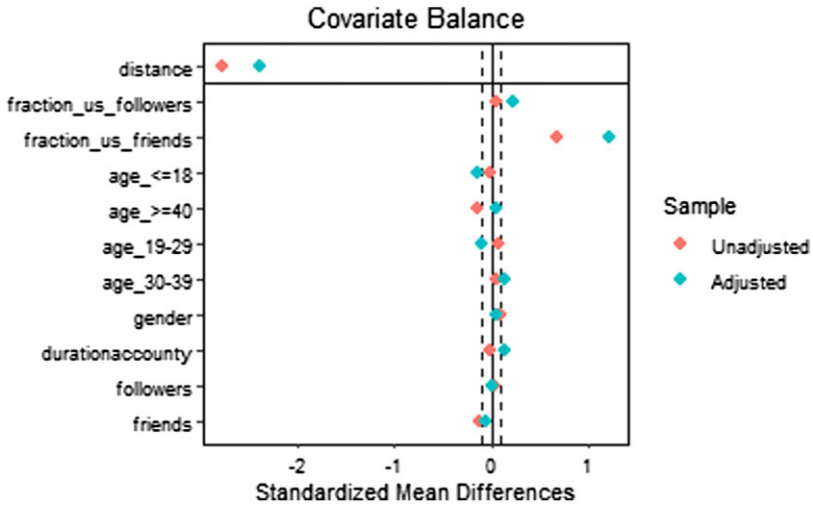
**Figure 3.** Example of friendship network data on Twitter: We have users 1 to 4 (alter users  $u$ ) following Jisu (ego user  $v$ ) and the dates of when these users created their accounts. The link creation column shows an illustrative example of the dates when the user 1 and 2 started following Jisu.

country of residence, city, or even coordinates. To obtain uniform location information at a country level, we used a dictionary of country names taken from <http://www.geonames.org/>. This enabled us to obtain information on how many friends and followers were added before and after migration, and where these friends and followers were located.

### 3.5 Matching migrants and non-migrants

In this project, we prioritise having reliable and high-quality information on users over the size of the data set. Hence, after obtaining complete data for both migrants and non-migrants, we performed 1:1 propensity score matching using the logistic regression model (Stuart et al. 2011) to find the best possible match between migrants and non-migrants who had similar levels of activity on Twitter and similar characteristics. Comparing the friendship patterns among this placebo group with those among migrants over time (in particular, the number of links to other users in the USA) enabled us to isolate the effect of the migration event. The application of this matching technique was particularly important in the context of the USA, as it is the country with the largest Twitter user population. Because of this large user reservoir, there was a certain baseline rate for the creation of US-based friend/follower links, irrespective of migration. Comparing similar migrants and non-migrants helped to account for this baseline rate. As was mentioned previously, we have chosen to focus on migrants in the USA. Our data include about 1,300 migrant individuals who migrated to the USA. Therefore, we performed the matching technique to select non-migrants (i.e., individuals who never changed their usual country of residence) whose characteristics were similar to those of the migrants in the USA. Here, the degree of similarity was measured based on the users' characteristics that we could observe from Twitter (i.e., age, gender, age of Twitter account, number of followers, friends, country of origin) and on the fraction of the users' friends and followers who were based in the USA.

Through this process, we were able to match up 1,163 migrants and 1,163 non-migrants. In Fig. 4, we can see that the covariates were more balanced after matching. The 'distance' in the first row of Fig. 4 is the overall difference in propensity scores. We can also see that the overall distance was reduced after matching, which suggests that we were able to pair migrants and non-migrants with similar characteristics and that the overall



**Figure 4.** Histograms showing the balance for categorical variables before ('unadjusted') and after ('adjusted') matching where *fraction\_us\_friends* and *fraction\_us\_followers* mean fraction of US-based friends and followers, *age\_cate\_0* means age below 18 years, *age\_cate\_1* means age between 19 and 29 years, *age\_cate\_2* means age between 30 and 39 years and *age\_cate\_3* means age above 40 years. The *durationaccounty* signifies account age.

balance improved after matching. It should, however, be noted that the balance improved in most of the covariates except for the fraction of US-based followers and friends and the first two age categories. This may be because there were large differences in these variables between migrants and non-migrants from the beginning.

## 4. Results

We begin this section by comparing the distributions of the fraction of followers and friends for both migrants and non-migrants. Here, we compare the distribution of the fraction of followers (or friends) instead of the absolute number of followers (or friends), because it can provide a clearer picture of how the structure of the network changed after migration. By considering the total number of new followers (or friends) added before and after migration regardless of the user's location, we are able to observe whether an immigrant was indeed adding new social links in the USA, or whether only a small proportion of the user's new social connections were based in the USA.

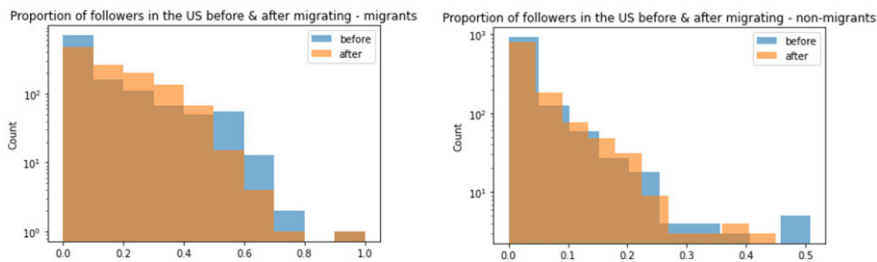
Moreover, we compare the distribution of the followers and the friends of migrants and non-migrants. Using the inferred migration date of the matched migrants, we assign a 'placebo' date of migration to non-migrants to enable us to compare the shift in the distribution of the fraction of friends for both migrants and non-migrants. This allows us to purely observe the effects of the migration event. If we observe an equal shift in the distribution of the fraction of the user's US-based friends before and after the placebo date of

migration, this would indicate that the migration event did not play a role in the expansion of Twitter user’s social network in the USA.

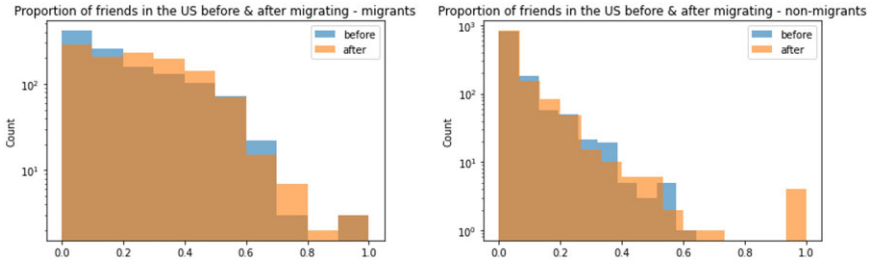
In Fig. 5, we observe a shift towards the right in the fraction of US-based followers for both migrants and non-migrants. For migrants, the average of this fraction increased from 0.14 to 0.165. Using the placebo migration date that we assigned to non-migrants, we also show for non-migrants the distribution of the fraction of followers residing in the USA before and after the placebo migration date. Looking at the right panel of Fig. 5, we can see that this distribution increased from 0.036 to 0.042. The results of the  $t$ -tests for both groups indicate that the means were different before and after migration ( $t$ -statistic value of  $-4.14$  for migrants and  $-2.63$  for non-migrants with a  $p$ -value of  $3.55 \times 10^{-5}$  and  $0.008$ , respectively). Thus, we observe that the percentage change for both groups was approximately the same. Our finding that there was an equal shift in the distribution suggests that the migration event did not necessarily play a role in the expansion of the Twitter users’ follower networks in the USA.

When we compare the distribution of the fraction of friends residing in the USA before and after migration for migrants (left panel) and for non-migrants (right panel) in Fig. 6, we see a different pattern. In the left panel, we observe a shift in the distribution towards the right, indicating that the fraction of friends residing in the USA increased after migration for migrants. To be more specific, the mean value of the fraction of friends who were based in the USA increased, on average, from 0.214 before migration to 0.25 after migration. We further tested whether the difference between the means of the two groups was significant. An independent two-sided  $t$ -test showed that the means were significantly different from each other ( $t$ -statistic value of  $-4.35$  and a  $p$ -value of  $1.42 \times 10^{-5}$ ).

For non-migrants, we look at the fraction of their friends who were based in the USA before and after the *placebo* date of migration. As the right panel of Fig. 6 illustrates, we see no shift in the distribution after migration for non-migrants, meaning that no new US-based friends were added. Before the *placebo* date of migration, about 6.1% of the migrants’ friends were from the USA. This share increased only slightly after the *placebo* date of migration to 6.2%. However, the results of the  $t$ -test showed that the average fractions of US-based friends before and after migration were not statistically different from each other ( $t$ -statistic value of  $-0.2$  with a  $p$ -value of  $0.84$ ). This indicates that there is not enough evidence to conclude that the two distributions had different means. As we



**Figure 5.** Distribution of fraction of followers residing in the USA before and after migration for migrants (left) and for propensity score-matched non-migrants (right).



**Figure 6.** Distribution of fraction of friends residing in the USA before and after migration for migrants (left) and for propensity-score matched non-migrants (right).

observe a clear shift in the distribution for migrants but not for non-migrants, our assumption that the migration event played a significant role in expanding Twitter users' friends networks in the USA is confirmed.

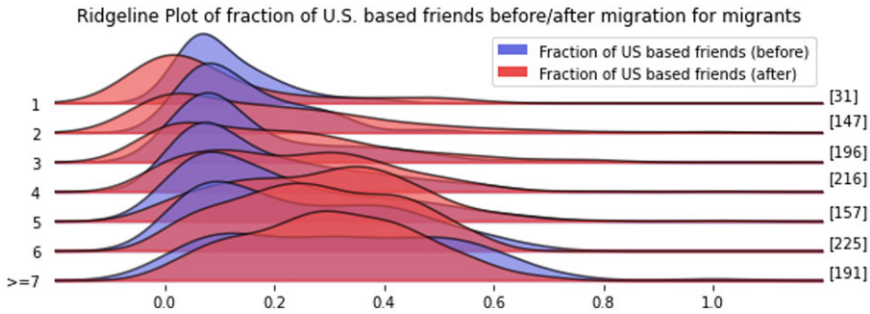
To understand further the magnitude of the migration effect, we used a Difference-in-Differences (DiD) model, which compared changes in the fraction of friends between a treatment group (i.e., migrants) and a control group (i.e., non-migrants). As shown in Table A.1 in the Appendix, our DiD analysis revealed that migration had a positive effect on the fraction of US-based friends among migrants. Specifically, the DiD estimator was 0.01, which indicates that, on average, the fraction of US-based friends increased by 1 percentage point after the migration. To better understand the significance of the effect size of 1%, it is useful to consider the base level, which is approximately 21–25%. Therefore, a 1% absolute increase represents a 4% relative increase. This result suggests that migration may lead to an increase in the number of US-based friends in a social network. However, it is important to note that the DiD approach is subject to certain limitations. In particular, migration may not be an exogenous shock, meaning that it may not be a completely random event that is not influenced by other factors. For instance, individuals who choose to migrate may have different characteristics or motivations than those who do not migrate, and these factors may also affect the composition of their social networks. As a result, it is possible that the effect of migration on the fraction of the US-based friends may be confounded by these other factors. We therefore note that there is a potential of the endogeneity of migration events in this model.

To assess the impact of our matching algorithm on the results of our study, we resampled our two groups of populations. In our previous approach, we matched migrants to non-migrants based on several criteria, including age, gender, age of Twitter account, number of followers, friends, country of origin, and the proportion of users' friends and followers based in the USA. In our new approach, we first matched the populations by country of origin to reduce heterogeneity in users' backgrounds. To achieve this, we used an *exact* matching algorithm. As a result of this process, we were able to identify and select 1,267 migrants and 3,007 non-migrants for our study. Then, we conducted a second matching algorithm and added a new matching element, the fraction of tweets in English, to further refine the matching. This enables us to match users who tweet in the same language as the destination country they migrated to. To achieve this, we



consider only the users whose fraction of tweets in English is equal to or higher than 0.2, which is equal to the median proportion of English tweets in our data. In the second phase of our matching process, we employed two distinct matching algorithms, namely Coarsened Exact Matching (CEM)<sup>13</sup> (Iacus et al. 2012) and Optimal Matching<sup>14</sup> (Hansen et al. 2006), to pair migrants with non-migrants. The CEM is a method that creates groups of observations with similar values on a set of pre-specified covariates. It reduces the dimensionality of the covariates by collapsing them into coarser categories, and then matches observations within these groups. Optimal pair matching is a matching technique that pairs treated units with control units, aiming to optimise an overall criterion rather than selecting the closest match. The criterion used is the sum of the absolute pair distances in the matched sample, where the distance is estimated using logistic regression. Using CEM, we identified 405 migrants and non-migrants who possessed similar characteristics. Meanwhile, the Optimal Matching algorithm enabled us to identify the best possible matches between individuals that minimise the overall distance between matched pairs of individuals based on a range of characteristics. Here, the distance is measured using logistic regression, resulting in a total of 1,628 pairs of migrants and non-migrants. Finally, we repeated the analysis using both matching algorithms and they produced consistent qualitative findings. Although the numerical values may have varied slightly between the two matching algorithms, the similarity in their overall qualitative findings lends support to the validity of the analysis. Therefore, we report the results obtained from the optimal matching algorithm as it has produced a greater number of paired migrants and natives. However, all resulting figures are included in the Appendix (see Figs A.4–A.6 for the results of CEM).

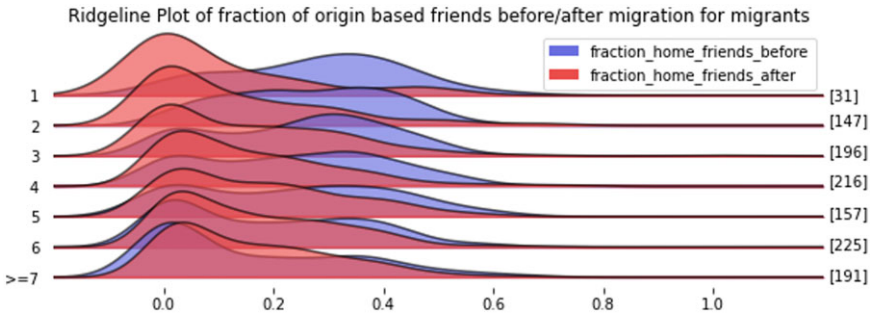
As shown in Fig. A.1, our focus on similarity in the background by country of origin resulted in greater variability in the standardised mean differences. But most importantly, our robustness check demonstrated a similar distribution of US-based followers and friends. Interestingly, we observed a clear shift in the distribution of followers in the USA following migration, as depicted in Fig. A.2. The fraction of US-based followers for migrants increased from an average of 0.34–0.4, while for non-migrants, it rose from 0.046 to 0.062. The *t*-tests for both groups showed that the means were significantly different before and after migration (*t*-statistic value of  $-4.24$  for migrants and  $-3.79$  for non-migrants, with *p*-values of  $1.2 \times 10^{-5}$  and 0.0002, respectively). The data presented in Fig. A.3 provides insights into how migration affects the distribution of friends in the USA. The graph highlights that only migrants experienced a notable shift towards the right, indicating an increase in the proportion of US-based friends. On average, the fraction of US-based friends for migrants increased from 0.25 to 0.28, while for non-migrants, it rose only slightly from 0.07 to 0.073. Furthermore, the *t*-tests conducted on migrant and non-migrant groups indicated that there was a significant difference in means only for migrants, with a *t*-statistic value of  $-3.14$  and a *p*-value of 0.0008. In contrast, there was no significant difference for non-migrants, with a *t*-statistic value of  $-0.64$  and a *p*-value of 0.26. In summary, our robustness tests confirmed that despite using different matching criteria, our results consistently indicated that migration did not influence the growth of Twitter users' follower networks. However, it did have an impact on their friend networks in the USA.



**Figure 7.** Ridgeline plot of fraction of US-based friends after migration for migrants. The y-axis represents number of years spent in the USA. The numbers in squared bracket indicate the number of populations considered for each number of years since migration.

To further improve our understanding of when the migrants began to have more US-based friends on Twitter, we split the distributions by how many years migrants have been residing in the USA. In other words, we examine how the fraction of friends migrants had before and after migration differed for those, for instance, who have been residing in the USA for two years. In Fig. 7, we observe that as the number of years migrants spent in the USA increased, the fraction of friends they had after migration shifted towards the right. This indicates that migrants who had been in the USA for longer periods of time had more US-based friends on Twitter. More specifically, we observe a clear shift for migrants starting in the third year since migration. On the other hand, we also observe that the distributions of the fraction of friends before and after migration were identical in the first year since migration and in the seventh year since migration or later. The results of the  $t$ -tests also indicate that the average fraction of US-based friends before and after migration were identical for these years ( $t$ -statistic value of  $-0.8$  and  $p$ -value of  $0.43$  for the first year and  $t$ -statistic value of  $-1$  and  $p$ -value of  $0.31$  for the seventh year or later). These results suggest that the immigrants needed some time to expand their friendship networks on Twitter after migration. Interestingly, immigrants who had been in the USA for more than seven years had a relatively large fraction of US-based friends before migrating but also added a larger additional fraction of friends after migrating. Note that here we are observing different sets of users as the number of years spent in the USA increased (the number of observed populations for each year is indicated in the squared bracket), meaning that different sets of users are contributing to the changes that we observe each year.

We further test the robustness of the ridgeline plot of the fraction of US-based friends after migration. We do this to check whether the changes in the fraction of US-based friends were caused by changes in the composition of the underlying population. To address this issue, we redo plot 7 with the calendar year of migration, instead of for the years for which we have a comparable number of users. We find that the number of users was similar (about 200 users) for the calendar years 2013–2016. Hence, we report the ridgeline plot for these four years. In Fig. A.7 in the Appendix, we observe somewhat similar patterns across the years. In particular, we see that the users had, on average, a higher fraction



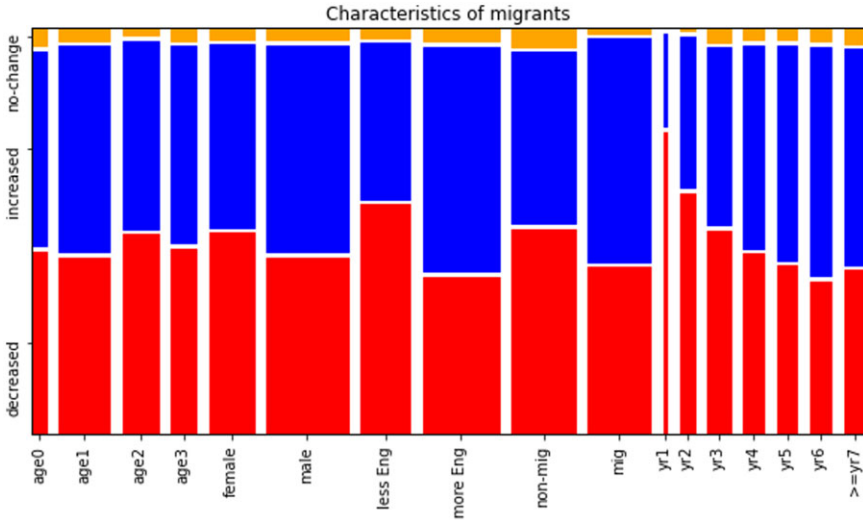
**Figure 8.** Ridgeline plot of fraction of origin-based friends after migration for migrants. The y-axis represents number of years spent in the USA. The numbers in squared bracket indicate the number of populations considered for each number of years since migration.

of US-based friends after migration than before migration. This tells us that the period effect on overall changes in the fraction of US-based friends before and after the migration event was small (see also Fig. A.8 in the Appendix for further robustness test).

We also examine how the migrants' connections to their origin country changed before and after migration. In Fig. 8, we show the same distributions, but for the fraction of origin-based friends of migrants. We see that the pattern differs from that in the previous figure, as the fraction of origin-based friends after migration stayed relatively small throughout the analysed period. This tells us that migrants did not add as many origin-based friends after migration as they did before migration. Note, however, that we observe a different pattern for migrants in the seventh year since migration or later: i.e., they had similar fractions of origin-based friends both before and after migration. The  $t$ -test results also tell us that the means were not significantly different from each other ( $t$ -statistic value of 0.076 and a  $p$ -value of 0.94). Interestingly, as more time since migration passed, the migrants also tended to add more origin-based friends on Twitter.

As we have observed a clear shift in the proportion of the migrants' friends who were based in the USA after migration, we further extend the analysis to study the characteristics of these users based on several features that we can observe from the data. To do so, we first divide the users according to their fraction of US-based friends into three groups: those who had a larger fraction (called 'more friends'), a smaller fraction (called 'less friends'), and no change (called 'no change') in the fraction of US-based friends after migration. We then look at the frequency of the following categorical variables: whether a user was or was not a migrant (named 'mig' and 'non-mig'), age category, gender, whether more English tweets were posted after migration (named 'more Eng' and 'less Eng'), and the number of years since migration. With this contingency table, we perform a correspondence analysis that allows us to study the dependencies between our variables of interest, i.e., the fraction of US-based friends in the three groups, and other categorical variables.

First, the chi-square test tells us that the variables have a statistically significant association (chi-square value of 219.19 and  $p$ -value of 0.0004). To help us understand these associations in more detail, Fig. 9 illustrates the relationships between these categories. Here, we observe that having more friends on Twitter after migration is most closely



**Figure 9.** Correspondence analysis mosaic plot of Twitter characteristics of users: x-axis shows different characteristics of users and y-axis shows whether these characteristics are associated with either decrease, increase or no-changes in the proportion of US-based friends. *age* indicates age groups from 0 to 3, *more/less Eng* indicates whether a user has tweeted more or less in English after migration, *mig* and *non-mig* indicate migrants, and non-migrants, *yr* indicates years since migration from one to seven and above.

associated with tweeting more in English after migration; being male; being between the ages of 19 and 29 years (*age1* from the figure); having a varying number of years since migration, including being in the fourth, fifth, sixth, or seventh year or later since migration; and being a migrant. In contrast, we find that having fewer friends on Twitter after migration is most closely associated with tweeting less in English after migration; being between the ages of 30 and 39 years or being at the age of 40 years or older (*age2* and *age3* from the figure); being female; being in the first, second, or third year since migration; and being a non-migrant. Note, however, that the third year since migration is the period when we begin to observe a shift in the fraction of US-based friends before and after migration as is shown in Fig. 7. Moreover, for migrants who migrated more than seven years ago, we observe no significant difference in the two distributions.

## 5. Discussion

In this study, through a careful design of the data collection process and various analyses, we sought to identify migrant users of Twitter and to understand the evolution of their new social connections on the platform after migration. Thanks to the availability of these data, we were able to observe that the fraction of a migrant's friends who were based in the destination country increased, on average, from 0.21 to 0.25 after migration. Our comparison of the behaviour of matched non-migrants and migrants suggests that this

increase was indeed due to the migration event. To be more specific, we found that the longer a migrant stayed in the USA, the more US-based friends they had on Twitter. Interestingly, however, we observed that a clear shift in the fraction of friends did not start until the third year since migration. Further analysis indicated that this increase in the fraction of friends was associated with several user characteristics. For instance, we observed that being a male migrant between ages 19 and 29 years who tweeted more in English after migration, having migrated four or more years ago, and being a migrant were most closely associated with an increase in the fraction of friends in the destination country. Additionally, when we looked at the fraction of origin-based friends in general, we noticed that migrants did not add as many origin-based friends after migration as they did before migration. Nonetheless, after the migrants had been living in the destination country for a longer period of time, they again tended to add more origin-based friends on Twitter. However, unlike the patterns we observed for friend connections, we found no evidence that the migration event was associated with an expansion of the Twitter users' follower networks in the destination country.

Thanks to the availability of Twitter data, we were able to capture the process of social integration among users of this specific platform. We were able to obtain data that contained detailed information on migrants' (online) social networks, including precise information on how many friends were added to their networks before and after migration. These kinds of data are not available from surveys. Nevertheless, Twitter data has some limitations. First, it is important to mention that as we prioritised having reliable and high-quality data on users, we obtained a relatively small data set that does not allow us to generalise our findings. Additional experiments with Twitter users should be done in order to provide further validation of our results for this platform. Moreover, there are many challenges that arise when seeking to identify migrants from geo-located tweets. Thus, it is possible that we identified members of a highly mobile population, such as short-term visitors (e.g., students) and business and leisure travellers (Armstrong et al. 2021), rather than migrants. On the related note, we relied here on the methodology developed by Wang et al. (2019) to obtain demographic information. As was mentioned previously, although the algorithm had a high accuracy level, it did not allow us to break down age categories into finer classes. Furthermore, the binary gender category we used does not reflect everyone's reality.

Secondly, people's social networks on Twitter differ from their real-world social networks, and from the network typologies of other social media platforms. On Twitter, individual users are free to follow other users as long as the accounts are publicly open. Hence, users can also interact with other users they do not know personally. This is different from other social media platforms such as Facebook, where the interactions are between close and mutual friends. To tackle this issue, we plan to further investigate how much of this social integration on Twitter is reflected in real life. Xie et al. (2012) examined the extent to which users' Twitter networks are reflected in their offline, real-life social networks. We also intend to explore this issue in order to better understand the magnitude of this social network effect in real life. Furthermore, based on the users' self-reported information, we were able to infer the current locations of their friends. However, this did not provide us with any information on whether these friends were or were not migrants themselves. In the future, we plan to also investigate the countries of

origin of immigrants' friends in order to better understand whether they are integrating into the host society, or are remaining within their community of origin. This is also a crucial factor to consider when studying social integration, as the failure of immigrants to integrate into the local population could lead to social division between different communities.

Lastly, it is also important to note that while we use the number of social connections on Twitter as a measure of social integration, we cannot disregard the possibility that these connections may hold different meanings in various contexts. For instance, some individuals may use Twitter primarily for information-seeking purposes (Cha et al. 2012), in which case following someone may indicate a shared interest or opinion rather than a deep social connection. However, we still believe that this measure can provide insights into certain forms of integration, such as cultural integration. This is because an interest in certain types of information may reflect an intention to adopt aspects of the host society's culture, particularly for migrants (Kim et al. 2022a).

Without a doubt, the analysis of new social links that this article has presented does not reflect all aspects of social integration. Nevertheless, we believe that our work shows how offline (re)location plays a significant role in the formation of immigrants' online networks.

Finally, and importantly, we want to emphasise that no personal information has been published at any stage of the research. All of the results are aggregated at the country level to mask any factors that can be traced back to a specific Twitter user.

## ACKNOWLEDGEMENT

This collaboration began during the Migration session of the BIGSSS Summer School in Computational Social Science, held in Sardinia, Italy, in 2019. More information can be found at <http://css.bigsss-bremen.de/migration2019>.

## Funding

I.W. is an Alexander von Humboldt (AvH) Professor for Artificial Intelligence and receives funding from the AvH Foundation.

*Conflict of interest statement.* None declared.

## Notes

1. <https://www.pewresearch.org/fact-tank/2020/08/20/key-findings-about-u-s-immigrants/>
2. <https://developer.twitter.com/en/docs/twitter-api/tweets/volume-streams/api-reference/get-tweets-sample-stream>
3. <https://developer.twitter.com/en/docs/twitter-api/v1/accounts-and-users/follow-search-get-users/overview>

4. <https://github.com/euagendas/m3inference>
5. <https://azure.microsoft.com/en-us/services/cognitive-services/face/#overview>
6. <https://www.faceplusplus.com>
7. The list of supported languages can be found here: <https://github.com/euagendas/m3inference/wiki/Languages>.
8. Recommendations on Statistics of International Migration, Revision 1, Statistical Papers, Series M, No. 58, United Nations, New York, 1998, Glossary.
9. “*friends* refers as the Twitter users that a specific user follows (e.g., following).” as defined by Twitter (<https://developer.twitter.com/en/docs/twitter-api/v1/accounts-and-users/follow-search-get-users/overview>).
10. “*Followers* refers to the Twitter users that follow a specific user.”-Idem.
11. Application Programming Interface
12. <https://developer.twitter.com/en/docs/twitter-api/v1/accounts-and-users/follow-search-get-users/api-reference/get-followers-list>
13. The CEM is a method that creates groups of observations with similar values on a set of pre-specified covariates. It reduces the dimensionality of the covariates by collapsing them into coarser categories, and then matches observations within these groups.
14. The Optimal Matching algorithm finds the best possible matches between observations based on a measure of distance. It compares each observation in the treatment group to every observation in the control group and assigns weights to the pairs based on the distance.

## References

- Armstrong, C. et al. (2021) ‘Challenges When Identifying Migration from Geo-Located Twitter Data’, *EPJ Data Science*, 10/1: 1.
- Abdirahman, O. (2017) ‘Integration and Social Cohesion: Key Elements for Reaping the Benefits of Migration’, *Global Compact Thematic Paper*. <https://migrationnetwork.un.org/resources/integration-and-social-cohesion-key-elements-reaping-benefits-migration>.
- Bailey, M., Drew, J., Koenen, M., Theresa, K., Russel, D. and Stroebel, J. (2022) ‘The Social Integration of International Migrants: Evidence from the Networks of Syrians in Germany’, *CESifo Working Paper No. 9680*.
- Bahar, D., Rapoport, H. and Turati, R. (2022) ‘Birthplace Diversity and Economic Complexity: Cross-Country Evidence’, *Research Policy*, 51/8: 103991.
- Bourdieu, P. (1986) ‘The Forms of Capital’, in Richardson J. (ed.) *Handbook of Theory and Research of the Sociology of Education*, 1st ed, pp. 241–58. Westport CT: Greenwood.
- Cha, M. et al. (2012) ‘The World of Connections and Information Flow in Twitter’, *IEEE Trans Syst Man Cybern Part A Syst Hum*, 42/4: 991–8.
- Chen, Y. and Wang, J. (2015) ‘Social Integration of New-Generation Migrants in Shanghai China’, *Habitat International*, 49: 419–25.



- Chi, G. et al. (2020) 'Who Ties the World Together? Evidence from a Large Online Social Network', in *Complex Networks and Their Applications VIII: Volume 2 Proceedings of the Eighth International Conference on Complex Networks and Their Applications*. Springer International Publishing.
- Chiswick, B. R. and Wang, Z. (2016) 'Social Contacts, Dutch Language Proficiency and Immigrant Economic Performance in the Netherlands: A Longitudinal Study', *IZA Discussion Paper No. 9760*.
- Comola, M. and Mendola, M. (2015) 'Formation of Migrant Networks', *The Scandinavian Journal of Economics. Economics*, 117/2: 592–618.
- Docquier, F. and Rapoport, H. (2012) 'Globalization, Brain Drain, and Development', *Journal of Economic Literature*, 50/3: 681–730.
- Dubois, A., Zagheni, E., Garimella, K. and Weber, I. (2018) 'Studying Migrant Assimilation through Facebook Interests', in *Social Informatics: 10th International Conference, SocInfo 2018, St. Petersburg, Russia, September 25–28, 2018, Proceedings, Part II* 10, pp. 51–60. Springer International Publishing.
- Glorius, B. et al. (2020) 'Is Social Contact with the Resident Population a Prerequisite of Well-Being and Place Attachment? The Case of Refugees in Rural Regions of Germany', *Frontiers in Sociology*, 5: 578495.
- Hochreiter, S. and Schmidhuber, J. (1997) 'Long Short-Term Memory', *Neural Computation*, 9/8: 1735–80.
- Huang, G., Liu, Z., Van Der Maaten, L. and Weinberger, K. Q. (2017) 'Densely Connected Convolutional Networks', in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4700–08.
- Gowricharn, R. (2002) 'Integration and Social Cohesion: The Case of The Netherlands', *Journal of Ethnic and Migration Studies*, 28/2: 259–73. p
- Hansen, B. B. and Klopfer, S. O. (2006) 'Optimal Full Matching and Related Designs via Network Flows', *Journal of Computational and Graphical Statistics*, 15/3: 609–27.
- Iacus, S. M., King, G. and Porro, G. (2012) 'Causal Inference without Balance Checking: Coarsened Exact Matching', *Political Analysis*, 20/1: 1–24.
- Kearns, A. and Whitley, E. (2015) 'Getting There? The Effects of Functional Factors, Time and Place on the Social Integration of Migrants', *Journal of Ethnic and Migration Studies*, 41/13: 2105–29.
- Kim, J., Pratesi, F., Rossetti, G., Sirbu, A. and Giannotti, F. (2023) 'Where Do Migrants and Natives Belong in a Community: A Twitter Case Study and Privacy Risk Analysis', *Social Network Analysis and Mining*, 13/15. <https://doi.org/10.1007/s13278-022-01017-0>.
- Kim, J. et al. (2022a) 'Origin and Destination Attachment: study of Cultural Integration on Twitter', *EPJ Data Science*, 11/1: 55.
- , Sirbu, A., Rossetti, G. and Giannotti, F. (2022b) 'Characterising different communities of Twitter users: migrants and natives', in *Complex Networks & Their Applications X: Volume 1, Proceedings of the Tenth International Conference on Complex Networks and Their Applications COMPLEX NETWORKS 2021* 10, pp. 130–141. Springer International Publishing.
- Lehdonvirta, V. and Räsänen, P. (2011) 'How Do Young People Identify with Online and Offline Peer Groups? A Comparison between UK, Spain and Japan', *Journal of Youth Studies*, 14/1: 91–108.

- Mazzoli, M. et al. (2020) 'Migrant Mobility Flows Characterized with Digital Data', *PLoS One*, 15/3: e0230264.
- Meeder, B., Karrer, B., Sayedi, A., Ravi, R., Borgs, C. and Chayes, J. (2011) 'We Know Who You Followed Last Summer: Inferring Social Link Creation Times in Twitter', in *Proceedings of the 20th International Conference on World Wide Web*, pp. 517–26.
- Portes, A. (1998) 'Social Capital: Its Origins and Applications in Modern Sociology', *Annual Review of Sociology*, 24/1: 1–24.
- Putnam, R. D. (1995) 'Bowling Alone: America's Declining Social Capital', *Journal of Democracy*, 6/1: 65–78.
- Rauch, J. E. (2001) 'Business and Social Networks in International Trade', *Journal of Economic Literature*, 39/4: 1177–203.
- and Trindade, V. (2002) 'Ethnic Chinese Networks in International Trade', *Review of Economics and Statistics*, 84/1: 116–30.
- Schiff, M. (1992) 'Social Capital, Labor Mobility, and Welfare: The Impact of Uniting States', *Rationality and Society*, 4/2: 157–75.
- Sirbu, A. et al. (2021) 'Human Migration: The Big Data Perspective', *International Journal of Data Science and Analytics*, 11/4: 341–60.
- Stewart, I., Flores, R. D., Riffe, T., Weber, I. and Zagheni, E. (2019) 'Rock, Rap, or Reggaeton?: Assessing Mexican Immigrants' Cultural Assimilation using Facebook Data', in *The World Wide Web Conference*, pp. 3258–64.
- Stuart, E. A. et al. (2011) 'MatchIt: Nonparametric Preprocessing for Parametric Causal Inference', *Journal of Statistical Software*, 42/8.
- Uslaner, E. M. (2003) 'Volunteering and Social Capital: How Trust and Religion Shape Civic Participation in the United States', in *Social Capital and Participation in Everyday Life*, pp. 104–17. Routledge.
- Van der Meer, T. and Tolsma, J. (2014) 'Ethnic Diversity and Its Effects on Social Cohesion', *Annual Review of Sociology*, 40/1: 459–78. p
- Wacquant, L. J. and Bourdieu, P. (1992) *An Invitation to Reflexive Sociology*, pp. 1–59. Cambridge: Polity.
- Wang, Z., Hale, S., Adelani, D. I., Grabowicz, P., Hartman, T., Flöck, F. and Jurgens, D. (2019) 'Demographic Inference and Representative Population Estimates from Multilingual Social Media Data', in *The World Wide Web Conference*, pp. 2056–67.
- Wessendorf, S. and Phillimore, J. (2019) 'New Migrants' Social Integration, Embedding and Emplacement in Superdiverse Contexts', *Sociology*, 53/1: 123–38.
- Xie, W., Li, C., Zhu, F., Lim, E. P. and Gong, X. (2012) 'When a Friend in Twitter is a Friend in Life', in *Proceedings of the 4th Annual ACM Web Science Conference*, pp. 344–47.
- Yoon, I. J. (2022) 'Social Integration and Well-Being of North Korean Migrants in South Korea', *Journal of Social Issues*, 78/3: 627–44.

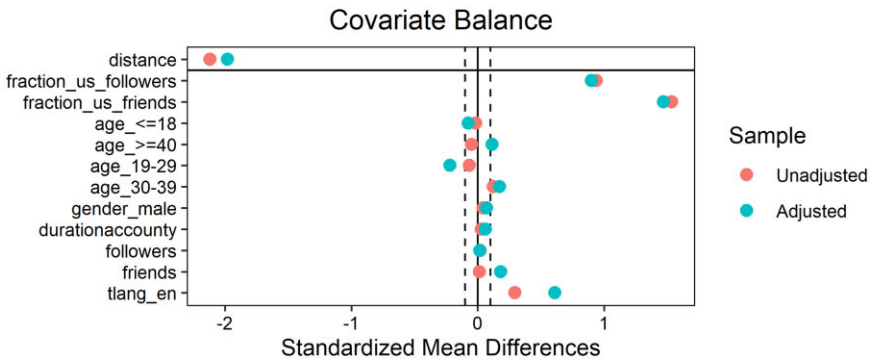
# Appendix

**Table A.1.** Diff-in-Diff model model

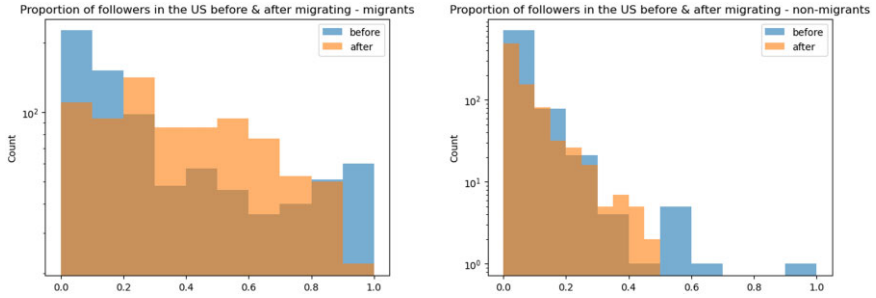
Variable	Model 1
Intercept	0.05*** (0.00)
Migrant or non-migrant (mig_nat)	0.001*** (0.00)
Time	0.002*** (0.00)
Mig_nat * time	0.01*** (0.00)
Adjusted R <sup>2</sup>	0.014

Coefficients, Standard errors in parentheses.

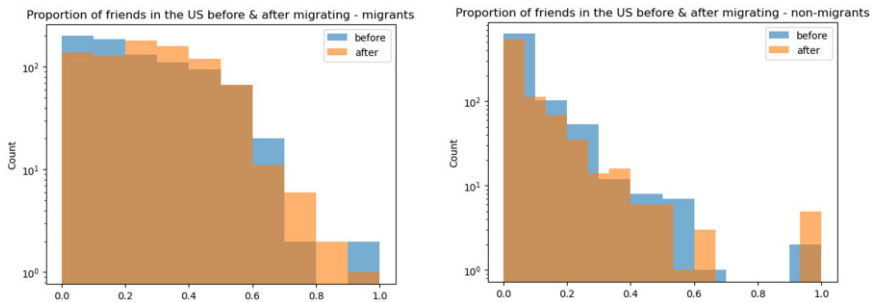
\*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .



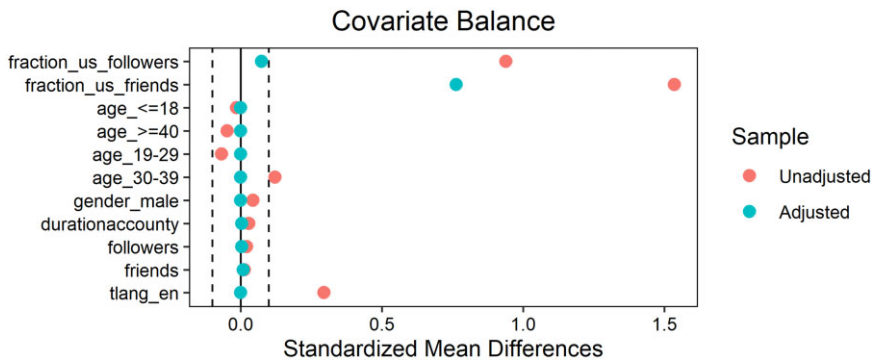
**Figure A.1.** Histograms showing the balance for categorical variables before ('unadjusted') and after ('adjusted') matching as a result of *optimal* matching algorithm, where fraction\_us\_friends and fraction\_us\_followers mean fraction of US-based friends and followers, age\_≤ 18 means age below 18 years, age\_19–29 means age between 19 and 29 years, age\_30–39 means age between 30 and 39 years and age\_≥40 means age above 40 years. The durationaccounty signifies account age. The tiang\_en means proportion of tweets in English.



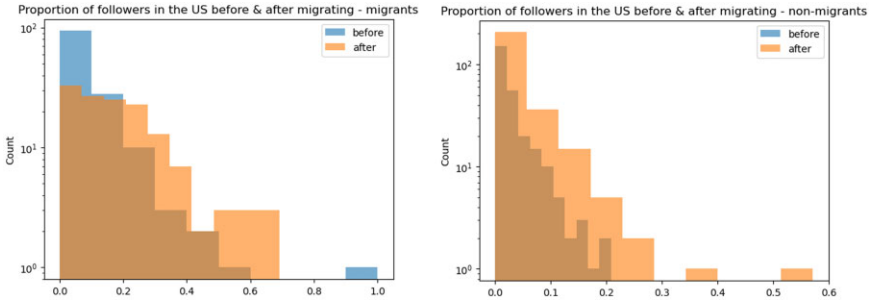
**Figure A.2.** Distribution of fraction of followers residing in the USA before and after migration for migrants (left) and for propensity-score matched non-migrants (right) as a result of *optimal* matching algorithm.



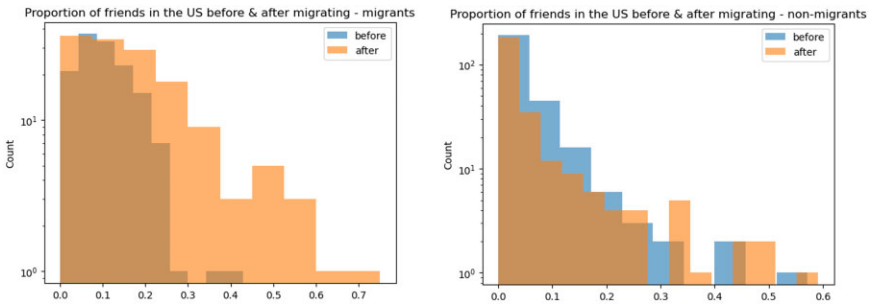
**Figure A.3.** Distribution of fraction of friends residing in the USA before and after migration for migrants (left) and for propensity-score matched non-migrants (right) as a result of *optimal* matching algorithm.



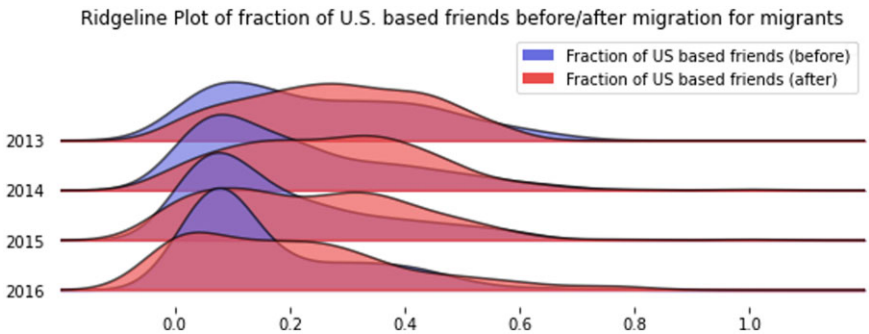
**Figure A.4.** Histograms showing the balance for categorical variables before ('unadjusted') and after ('adjusted') matching as a result of CEM matching algorithm, where `fraction_us_friends` and `fraction_us_followers` mean fraction of US-based friends and followers, `age_≤18` means age below 18 years, `age_19–29` means age between 19 and 29 years, `age_30–39` means age between 30 and 39 and `age_≥40` means age above 40 years. The `durationaccounty` signifies account age. The `tlang_en` means proportion of tweets in English.



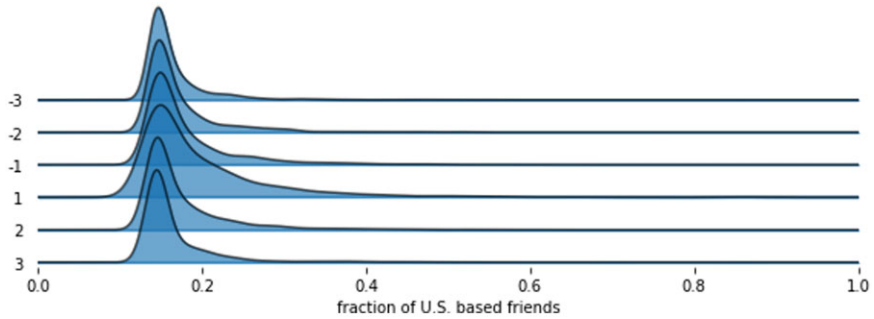
**Figure A.5.** Distribution of fraction of followers residing in the USA before and after migration for migrants (left) and for propensity-score matched non-migrants (right) as a result of CEM matching algorithm. For both migrants and non-migrants, we observe that distributions of the fraction of friends before and after migration are different from each other with  $t$ -statistics value of  $-5.31$  and  $-2.65$  with  $p$ -value of  $1.19e-07$  and  $0.008$  for migrants and non-migrants respectively.



**Figure A.6.** Distribution of fraction of friends residing in the USA before and after migration for migrants (left) and for propensity-score matched non-migrants (right) as a result of CEM matching algorithm. For migrants, we observe that distributions of the fraction of friends before and after migration are different from each other with  $t$ -statistic value of  $-4.53$  for migrants with  $p$ -value of  $0.00$ . The average fraction of US-based friends has increased from  $0.11$  to  $0.17$  for migrants. For non-migrants, we do not observe differences in distributions before and after migration (with  $t$ -statistic value of  $-0.13$  with  $p$ -value of  $0.89$ ).



**Figure A.7.** Ridgeline plot of fraction of US-based friends before and after migration for migrants per calendar year of migration event.



**Figure A.8.** Ridgeline plot of fraction of US-based friends before and after migration for migrants each year. The number of users is the same at each y-axis as we are observing how many US-based friends have an immigrant added each year. For example, 1 on the y-axis means first year of migration, and  $-1$  means a year before migration. We observe that immigrants on Twitter have added most of the new US-based friends in the first year after migration with an average of 0.077. Then we observe a decline in the fraction of US-based friends being added from second year and above (average of 0.045 in the second year and 0.033 in the third year). Furthermore, immigrants tend to also add some fraction of US-based friends prior to migration. More precisely, we observe that immigrants added, on average, 5.2% of new US-based friends a year prior to migration, 4.2% in the two years prior to migration, and 3% in the three years prior to migration.