

Mater Certa Est, Pater Numquam: What Can Facebook Advertising Data Tell Us about Male Fertility Rates?

Francesco Rampazzo,¹ Emilio Zagheni,² Ingmar Weber,³ Maria Rita Testa,⁴ Francesco Billari⁵

¹University of Southampton & Max Planck Institute for Demographic Research,

²University of Washington & Max Planck Institute for Demographic Research,

³Qatar Computing Research Institute, ⁴Wittgenstein Centre, ⁵Università Bocconi

f.rampazzo@soton.ac.uk, emilioz@uw.edu, iweber@hbku.edu.qa, maria.rita.testa@oeaw.ac.at, francesco.billari@unibocconi.it

Abstract

In many developing countries, timely and accurate information about birth rates and other demographic indicators is still lacking, especially for male fertility rates. Using anonymous and aggregate data from Facebook's Advertising Platform, we produce global estimates of the Mean Age at Childbearing (MAC), a key indicator of fertility postponement. Our analysis indicates that fertility measures based on Facebook data are highly correlated with conventional indicators based on traditional data, for those countries for which we have statistics. For instance, the correlation of the MAC computed using Facebook and United Nations data is 0.47 ($p = 4.02e - 08$) and 0.79 ($p = 2.2e - 15$) for female and male respectively. Out of sample validation for a simple regression model indicates that the mean absolute percentage error is 2.3%. We use the linear model and Facebook data to produce estimates of the male MAC for countries for which we do not have data.

Introduction

To count and project populations for the world, or a given country, we need timely and accurate information about birth and death rates. This information exists for developed countries, but it is missing for many developing countries that do not have adequate vital registration systems (Bongaarts and Blanc 2015). As a result, estimates are often obtained using indirect methods applied to survey data, especially for developing countries and male fertility (United Nations and Social Affairs 2015; 2017). The lack of up-to-date information has severe repercussions in the implementation and monitoring of policies.

This paper aims to understand whether data from Facebook's Advertising platform can be used to obtain estimates of the mean age at childbearing (MAC). The MAC is an important descriptive measure in demographic research, which can help with understanding patterns in fertility behaviour, such as postponement of childbearing. For this, we investigate the extent to which the MAC produced using the Facebook Advertising Platform data are congruent with figures from the United Nations. The significance of the study is to determine whether Facebook is a viable data source for studying two struggling areas of demographic research:

male fertility which has been thus far neglected, and the fertility of developing countries which is hampered by the lack of accurate data (Schoumaker 2017).

Fertility Data

The three main sources of fertility data are (i) vital registration systems, (ii) censuses, and (iii) surveys. These three sources can vary immensely in quality between and within countries and, as Moultrie et al. (2013) suggest, "*population statistics, like other demographic statistics, whether they are obtained by enumeration, registration, or other means, are subject to error*".

Vital registration systems have high coverage in developed countries, but not in developing countries. This type of data is generally the key source for estimating births, but in Sub-Saharan Africa they lack representativeness and accuracy (AbouZahr et al. 2015). With improvements in vital registration systems, the percentage of infants whose births had been registered increased from 58% to 65% (Mikkelsen et al. 2015), but UNICEF (2015) is still reporting that worldwide one in three infant children (ages 0 to 5) go un-registered.

In countries where vital registrations data is lacking, censuses and surveys are used for estimating births. However, these sources are not timely and accurate: there may be years in between censuses or surveys and lifetime fertility events may be underreported or completely omitted. Indeed, data from these sources in Sub-Saharan Africa are far from accurate (AbouZahr et al. 2015). Outdated surveys and censuses provide fragmented and contradictory information. (Mikkelsen et al. 2015; AbouZahr et al. 2015).

The diffusion of Internet and social media appears to be faster than the improvements made in vital registration. Although there is need to re-purpose statistical methods for producing statistical inference from Internet data (Zagheni and Weber 2015; Billari and Zagheni 2017), Facebook can otherwise be thought of as a regularly updated, although non-representative, digital census that could fill data gaps.

In addition to the limitations already discussed, current fertility data sources usually miss one side of the total picture: the male. Only a few developed countries produce statistics for male fertility through civil registration and vital statistics systems (CRVS) (Schoumaker 2017). Male fertility has recently started to be analysed for Germany and Greece (Dudel and Klüsener 2016; Tragaki and Bagavos 2014), but

also for developing countries with Demographic and Health Surveys (Schoumaker 2017). Facebook can provide information in regards to this yet under-analyzed aspect.

Related Work

Demography, being a data driven discipline among the social sciences, is one of the fields of research which can benefit from the abundance of Internet data. The focus of demography is on fertility, mortality and migration. Migration studies, in comparison to studies of fertility and mortality, are benefiting the most from social media data. Indeed, migration has been studied through Yahoo, Twitter, LinkedIn and Facebook (Zagheni and Weber 2012; State et al. 2015; 2014; Zagheni, Weber, and Gummedi 2017). The use of social media data in migration studies can shed new light on a branch of research in which the data available contains many deficiencies. Fertility research through Internet data has manifested itself into studies of fertility by age and location in the US (Ojala et al. 2017), in seasonality of mating-related Web searchers and consequential fertility (Markey and Markey 2013). (Hitsch, Hortaçsu, and Ariely 2010), (Bellou 2015), and (Billari 2016) focus on the impact of the diffusion of internet on the postponement in timing of marriages and births. The area of fertility desires and intentions have also been explored through Twitter (Adair et al. 2014). It is interesting that the area of mortality research is also employing social media data. In fact, the Web is rich with websites of family trees and these have been used to study life expectancy (Fire and Elovici 2015). The study of causes of death and the health of populations has also been affected by the Internet (Gittelman et al. 2015).

Facebook Advertising Data

Facebook’s advertising platform permits advertisers to selectively show their advertisement to Facebook users matching criteria specified by the advertiser. The criteria are divided into four sections of variables: (i) demographics, (ii) location, (iii) interests, and (iv) behaviour. The demographic section contains information based on traits such as age, gender, relationship status, education, workplace, job titles and more. Under this section, we can find information concerning parenthood. Parents on Facebook are further sub-divided by the age of their children. Facebook allows advertisers to target individuals by their location from a country level, to a neighbourhood level. This information can be disaggregated further between individuals that are travelling and whom are living in the specific area. Through the section on interests, advertisers can target individuals based on their interests, which are inferred based on pages liked and other signals. The behaviours section provides information concerning purchase behaviours, device usage and other activities. It is worth noting that Facebook gathers these estimations based on information from sites other than just facebook.com, as long as those sites have a Facebook Like or share functionality, which means a connection to facebook.com.

Dataset

Our research focuses on individuals with a Facebook profile, aged between 15 and 49 years old. The dataset was collected on January 2, 2018. Facebook’s advertising platform does not provide data for Cuba, Iran, North Korea, Syria, and Sudan. We are interested in women and men of a reproductive age, between 15-49 years old, who had a child in the last 12 months. Facebook already prepares an aggregated estimate for this category. The age variable is divided into 5 years gaps (15-19, 20-24, 25-29, 30-34, 35-39, 40-44, 45-49). We downloaded information about the total population of women and men in each age group as a proxy for exposure population. We used the Facebook Application Programming Interface (API) to download the data (Araújo et al. 2017).

In addition to the above online data, we obtained United Nations estimates for fertility as “ground truth” data with which to compare the Facebook data.

Measures

Fertility analysis measures are generally computed considering one sex only, the female population. The reason for this standard is that only women in their reproductive ages can give birth. Total Fertility Rate (TFR) and Mean Age at Childbearing (MAC) are the typically computed measures. In this paper, we are not interested in creating a model for studying the distortion linked to the digital divide, Facebook penetration (Fatehkia, Kashyap, and Weber) or language of the TFR. Our focus is on the MAC, which is computed separately for both sexes. The Spearman correlation coefficient and the Mean Absolute Percentage Error are the measures used for comparing the Facebook estimates to the ground truth data. To guard against overfitting, i.e. overly optimistic estimates, we used a leave-one-out-cross-validation (LOOCV): one observation is removed from the sample, while the remaining observations are used for training a model and predict the value for the removed observation. This approach is repeated as many times as the number of observations for calculating the average correlation and MAPE by continents.

Mean Age at Childbearing:

The Mean Age at Childbearing (MAC) is computed as the sum of the Age Specific Fertility Rate (ASFR) weighted by the mid-point of each age group, divided by the sum of the ASFR. MAC can be computed as follows:

$${}_nF_x[0, T] = \frac{\text{Births between 0 and T to women aged } x \text{ to } x+n}{\text{Women risk having a child in the period 0 to T}}$$

$$MAC = \frac{\sum_{x=15}^{49} (x + 0.5) \times {}_nF_x}{\sum_{x=15}^{49} ({}_nF_x)}$$

Where $(x + 0.5)$ is the mid-point for each age interval and $(x + 0.5) \times {}_nF_x$ is the age-specific fertility rate for women or men whose age corresponds to the age group of which $(x + 0.5)$ is the mid-point.

Results

Mean Age at Childbearing

The calculations of the Mean Age at Childbearing (MAC) for females and males are presented in the next two sections. We have not included in the analysis those countries in which Facebook Advertising Data reports user count estimates equal to 20, which is the default lower bound response, as 0 is never reported.

Female The correlation of the Female MAC with the corresponding UN estimates is 0.47 ($p = 4.02e - 08$). The calculation is made on 138 countries out of 194 available on the UN data¹. Dividing the result by continents, we can see that the correlation is negative for Africa and South America, while for the other continents it is positive. The highest correlation is obtained in Europe.

Table 1: Spearman Correlation and MAPE for Female MAC by continent and by LOOCV.

Continent	Female MAC		LOOCV		N
	Correlation	MAPE	Correlation	MAPE	
Africa	-0.27*	6.64%	0.11	3.06%	28
Asia	0.52***	6.80%	0.46***	3.92%	33
Europe	0.69***	6.54%	0.66***	2.98%	41
North America	0.62***	6.52%	0.55***	3.15%	16
Oceania	0.29	4.98%	-0.47	3.80%	8
South America	-0.26	6.88%	-0.62**	2.77%	12

*p<0.1; **p<0.05; ***p<0.01

Male The correlation for Male MAC is 0.79 ($p = 2.2e - 15$). The calculation has been made on 82 countries out of 164 available countries in the UN data with the latest available estimates for the period 2006-2015. Only three African countries are included in this calculations. We calculated the correlation for Female MAC approximately for the same sample (71 countries) and the correlation is equal to 0.75 ($p = 6.21e - 14$).

Table 2: Spearman Correlation and MAPE for Male MAC by continent and by LOOCV.

Continent	Male MAC		LOOCV		N
	Correlation	MAPE	Correlation	MAPE	
Africa	1.00	6.44%	0.5	6.74%	3
Asia	0.75***	6.88%	0.46***	3.92%	17
Europe	0.71***	5.35%	0.70***	2.41%	40
North America	0.87***	2.01%	0.81***	5.89%	13
Oceania	0.50	4.82%	-0.5	3.25%	3
South America	0.08	6.51%	-0.48	3.63%	6

*p<0.1; **p<0.05; ***p<0.01

¹Countries for which Facebook reported the default result, indistinguishable from no users, were not included in the analysis.

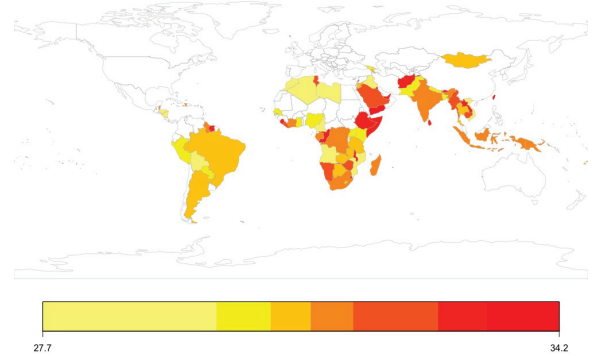


Figure 1: Prediction of Male MAC 2017 for countries without UN ground truth data.

Modeling male fertility

We fit a simple linear regression to model male fertility:

$$MAC_{UN} = \beta_0 + \beta_1 MAC_{FB} + \epsilon$$

where MAC_{UN} is the MAC calculated with the United Nation data, and MAC_{FB} is the MAC estimated through Facebook's Advertising data. In Table 3, we report the result of the linear regression model. The R^2 (0.676) indicates a good fit of the model. On average, Facebook data is underestimating male MAC.

Table 3: Linear Regression for Male MAC.

	Estimate
Intercept	7.451*** (1.936)
Facebook MAC	0.811*** (0.063)
N	82
R^2	0.676
Adjusted R^2	0.671
Residual Std. Error	0.949 (df=79)
F Static	164.4*** (df=1;79)

*p<0.1; **p<0.05; ***p<0.01

To validate our results, we performed an out-of-sample exercise. The simple linear regression has been run ten times, each time including 72 randomly selected observations as training data and 10 as a test data set. The average value of the MAPE for the predictions on the test set is equal to 2.3%, indicating that the model has high predictive capacity. For reference, the (standard deviation)/(average) is 11.72%.

Then, we predicted the Male MAC through a linear regression for those countries (79) without ground truth data. The prediction results are shown in Figure 1. As the map shows, our method helps to fill "data gaps" in many developing countries where these kinds of estimates are currently unavailable, potentially having big implications for policy making.

Conclusion and Discussion

This paper provides the basis for running more detailed (male) fertility analysis through Facebook Advertising Data as it shows the feasibility to estimate Mean Age at Childbearing (MAC). Our work highlights the limitations as well as the advantages of this data source. One advantage of Facebook data is that MAC estimates can be produced instantaneously, in particular for under-studied dimensions such as male fertility. There are further possibilities with this data combining the analysis of fertility with other targeting variables provided by Facebook, such as education, relationship status or interests in certain topics, such as religious content. We believe that this is a promising and new direction for future work on more multi-faceted fertility research at a global scale. Another advantage is that, due to Facebook's global reach, we can study fertility in developing countries. This has shown promising results for certain countries and is a promising starting point to fill data gaps for countries where ground truth data are missing or not up-to-date. Moreover, the Internet penetration rate is increasing in these countries and globally, which will lead to more Internet users and therefore, more data obtainable through Facebook or other online advertising platforms.

References

- AbouZahr, C.; de Savigny, D.; Mikkelsen, L.; Setel, P. W.; Lozano, R.; Nichols, E.; Notzon, F.; and Lopez, A. D. 2015. Civil registration and vital statistics: progress in the data revolution for counting and accountability. *The Lancet* 386(10001):1373–1385.
- Adair, L. E.; Brase, G. L.; Akao, K.; and Jantsch, M. 2014. #babyfever: Social and media influences on fertility desires. *Personality and Individual Differences* 71:135–139.
- Araújo, M.; Mejova, Y.; Weber, I.; and Benevenuto, F. 2017. Using facebook ads audiences for global lifestyle disease surveillance: Promises and limitations. In *WebSci*, 253–257.
- Bellou, A. 2015. The impact of Internet diffusion on marriage rates: evidence from the broadband market. *Journal of Population Economics* 28(2):265–297.
- Billari, F. C., and Zagheni, E. 2017. Big Data and Population Processes: A Revolution? *SocArXiv*.
- Billari, F. C. 2016. Internet and the Timing of Births. In *Giornate di Studio sulla Popolazione 2017*.
- Bongaarts, J., and Blanc, A. K. 2015. Estimating the current mean age of mothers at the birth of their first child from household surveys. *Population Health Metrics* 13:25.
- Dudel, C., and Klüsener, S. 2016. Estimating male fertility in eastern and western Germany since 1991: A new lowest low? *Demographic Research* 35(53):1549–1560.
- Fatehkia, M.; Kashyap, R.; and Weber, I. Using facebook ad data to track the global digital gender gap. *World Development* 107:189–209.
- Fire, M., and Elovici, Y. 2015. Data Mining of Online Genealogy Datasets for Revealing Lifespan Patterns in Human Population. *ACM TIST* 6(2):28:1–28:22.
- Gittelman, S.; Lange, V.; Gotway Crawford, C. A.; Okoro, C. A.; Lieb, E.; Dhingra, S. S.; and Trimarchi, E. 2015. A New Source of Data for Public Health Surveillance: Facebook Likes. *Journal of Medical Internet Research* 17(4).
- Hitsch, G. J.; Hortaçsu, A.; and Ariely, D. 2010. Matching and Sorting in Online Dating. *The American Economic Review* 100(1):130–163.
- Markey, P. M., and Markey, C. N. 2013. Seasonal variation in internet keyword searches: a proxy assessment of sex mating behaviors. *Archives of Sexual Behavior* 42(4):515–521.
- Mikkelsen, L.; Phillips, D. E.; AbouZahr, C.; Setel, P. W.; de Savigny, D.; Lozano, R.; and Lopez, A. D. 2015. A global assessment of civil registration and vital statistics systems: monitoring data quality and progress. *The Lancet* 386(10001):1395–1406.
- Moultrie, T. A.; Dorrington, R. E.; Hill, A. G.; Hill, K.; Timaeus, I. M.; and Zaba, B. 2013. Tools for Demographic Estimation.
- Ojala, J.; Zagheni, E.; Billari, F. C.; and Weber, I. 2017. Fertility and its meaning: Evidence from search behavior. In *ICWSM*, 640–643.
- Schoumaker, B. 2017. Measuring male fertility rates in developing countries with Demographic and Health Surveys: An assessment of three methods. *Demographic Research* 36(28):803–850.
- State, B.; Rodriguez, M.; Helbing, D.; and Zagheni, E. 2014. Migration of Professionals to the U.S. In *SocInfo*, 531–543.
- State, B.; Park, P.; Weber, I.; and Macy, M. 2015. The mesh of civilizations in the global network of digital communication. *PLOS ONE* 10(5):1–9.
- Tragaki, A., and Bagavos, C. 2014. Male fertility in Greece: Trends and differentials by education level and employment status. *Demographic Research; Rostock* 31:137–159.
- UNICEF. 2015. The State of the World's Children 2014 In Numbers: Every Child Counts.
- United Nations, D. o. E., and Social Affairs, P. D. 2015. World population prospects: The 2015 revision, methodology of the united nations population estimates and projections. Working Paper No. ESA/P/WP.242.
- United Nations, D. o. E., and Social Affairs, P. D. 2017. World population prospects: The 2017 revision. II:Demographic Profile (ST/ESA/SER.A/400)(2).
- Zagheni, E., and Weber, I. 2012. You Are Where You e-Mail: Using e-Mail Data to Estimate International Migration Rates. In *WebSci*, 348–351.
- Zagheni, E., and Weber, I. 2015. Demographic research with non-representative internet data. *International Journal of Manpower* 36(1):13–25.
- Zagheni, E.; Weber, I.; and Gummadi, K. 2017. Leveraging Facebook's Advertising Platform to Monitor Stocks of Migrants. *Population and Development Review* 43(4):721–734.