Model-based estimates in demography and global health: quantifying the contribution of population-period-specific information

Leontine Alkema, Guandong Yang and Krista Gile *
University of Massachusetts Amherst

Extended abstract prepared for PAA 2019

Abstract

Sophisticated statistical models are used to produce estimates for demographic and health indicators even when data are limited, very uncertain or lacking. We aim to provide a standardized approach to answer the question: To what extent is a model-based estimate of an indicator of interest informed by data for the relevant population-period as opposed to information supplied by other periods and populations and model assumptions? We propose a data weight measure to calculate the weight associated with population-period data y relative to the model-based prior estimate obtained by fitting the model to all data excluding y. In addition, we propose a data-model accordance measure which quantifies how extreme the population-period data are relative to the prior model-based prediction. We illustrate the insights obtained from the combination of both measures in the estimation of modern contraceptive use.

*This project is funded by the Bill and Melinda Gates Foundation

1 Introduction

Sophisticated statistical models are used to produce estimates for demographic and global health indicators to provide levels and trends even when data are limited, very uncertain or lacking. Common model features include covariates, parametrization of transitions, hierarchical models, temporal and/or spatial smoothing.

We aim to answer the following question: To what extent is a model-based estimate of an indicator of interest informed by data for the relevant population-period as opposed to information supplied by other periods and populations and model assumptions? Answering this question is relevant to avoid misuse and misinterpretation; a country case study would not be informative unless data are informative. Moreover, answering this question allows for highlighting where more information is needed for data-driven monitoring. Our goal is to provide a standardized approach to answering this question, so that the approach can be used for a range of models and be included in global health reporting guidelines.[1]

We propose a new measure, referred to as the data weight, that quantifies the weight associated with country-period data y relative to the model-based prior estimate obtained by fitting the model to all data excluding y. In addition, we propose a data-model accordance measure which quantifies how extreme the population-period data are relative to the prior model-based prediction. By combining both measures, we are able to identify model-based estimates where data and model are in disagreement as well as settings where more data are needed. We illustrate the insights obtained from the combination of both measures in the estimation of modern contraceptive use.

2 Methods

2.1 Notation

Let μ denote the indicator of interest and y data for the relevant population-period. Information on μ is summarized in terms of probability distributions. This arises naturally when modeling is carried out using a Bayesian approach.

The information available on μ prior to observing y, based on the model and data from other population-periods only, is summarized in distribution $p(\mu|z)$, with z being data from other periods and populations. For ease of notation, we leave out the conditioning on z and refer to $p(\mu)$ as the model-based prior distribution, summarizing information "coming from the model", which implicitly includes z. We aim to compare $p(\mu)$ to the posterior distribution $p(\mu|y)$, which summarizes the information about μ after updating the prior with information from the population-period of interest. The posterior is proportional to prior times likelihood, $p(\mu|y) \propto p(\mu)p(y|\mu)$. The normalized likelihood is defined as $q(\mu) = p(y|\mu)/\int_\mu p(y|\mu)d\mu$.

2.2 Computation and implication for proposed measures

We assume that for a given model and parameter of interest μ, we can obtain a finite number of samples from model-based prior $p(\mu)$ and posterior $p(\mu|y)$ to calculate data weight and data-model accordance measures. For example, samples from model-based prior $p(\mu)$ can be obtained based on a model fit after leaving out y, and samples from $p(\mu|y)$ can be obtained by fitting the model to the full data set.

For more complex models, the (normalized) likelihood function is generally not available in closed form nor easy to sample from, hence we assume $p(y|\mu)$ to be unknown. We approximate the likelihood by the ratio of the approximated posterior and prior distribution, $\hat{p}(y|\mu) \propto \hat{p}(\mu|y)/\hat{p}(\mu)$, where $\hat{p}(\mu|y)$ and $\hat{p}(\mu)$ refer to the posterior and prior density approximated using samples. The approximation $\hat{p}(y|\mu)$ may be inaccurate or infinite for μ with posterior support, $p(\mu|y) > 0$, but with the sample-based estimate for the prior density $\hat{p}(\mu) \approx 0$. While it may be possible in theory to obtain a more accurate approximation based on obtaining a larger sample from the prior density with posterior support, we assume that this is not realistic in practice, as it would require refitting of the model. Given this setting, we focus on measures that can be calculated based on a limited sample from the model-based prior and posterior distributions only. This restriction prevented us from using several more conventional measures such as divergence or relative entropy methods. These measures proved too sensitive to the tail behavior of the estimated densities.

Where possible, the calculation of the proposed measures takes into account that $\hat{p}(y|\mu)$ will be inaccurate or infinite for μ with $p(\mu|y) > 0$ and $\hat{p}(\mu) \approx 0$. In extreme cases, calculation of the measures is not possible and upper/lower bounds of the measures are presented instead.

2.3 Data weight measure

We aim to define a data weight $0 < w < 1$ that quantifies the weight given to information from the data as opposed to prior when calculating posterior mean $\hat{\mu}$. We propose the following measure:

$$w = \frac{1/\hat{\sigma}_L^2}{1/\hat{\sigma}_L^2 + 1/\hat{\sigma}_0^2},$$

where $\hat{\sigma}_L$ refers to the uncertainty associated with the data (the standard deviation of $\mu \sim q(\mu)$) and $\hat{\sigma}_0$ refers to the uncertainty associated with the model-based prior (the standard deviation of $\mu \sim p(\mu)$). This measure is motivated by the setting where the model-based prior and data are normally distributed. Specifically, when the model-based prior on μ is normally distributed:

$$\mu \sim N(\mu_0, \sigma_0^2),$$

with mean μ_0 and variance σ_0^2, and the data are normally distributed

$$y|\mu, \sigma_L \sim N(\mu, \sigma_L^2),$$

with mean μ and variance σ_L^2, then the posterior mean $\hat{\mu}$ is given by the weighted average of the data y and prior mean μ_0:

$$\hat{\mu} = E(\mu|y, \sigma_L^2) = w \cdot y + (1 - w) \cdot \mu_0,$$

with w as defined above. With this definition, the following holds true for the model-based point estimate of μ:

- $w = 0 \Rightarrow$ not informed by data y
- $w = 1 \Rightarrow$ not informed by model-based prior $p(\mu)$
- $w \uparrow \Rightarrow$ more weight given to data relative to model-based prior.

Figure 1 illustrates examples of model-based prior and posterior distributions with – from left to right– increasing data weights.

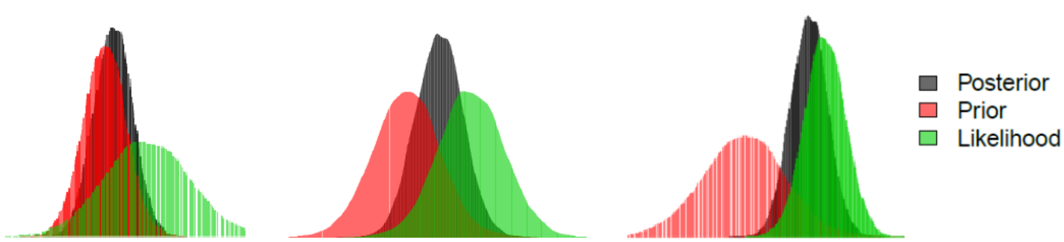

Figure 1: Illustration of model-based prior and posterior distributions with, from left to right, increasing data weights.

2.4 Data-model accordance measure

Let $\tilde{\mu}$ denote the point estimate for μ as suggested by data y. We set a equal to twice the one-sided prior probability tail area associated with $\tilde{\mu}$ such that

- $a = 0 \Rightarrow \tilde{\mu}$ is very extreme/not supported under the prior,

- $a = 1 \Rightarrow \tilde{\mu}$ is aligned with prior,
- $a \uparrow \quad \Rightarrow \tilde{\mu}$ is increasingly likely under the prior.

Figure 2 illustrates examples of model-based prior and posterior distributions with – from left to right– increasing data-model accordance.

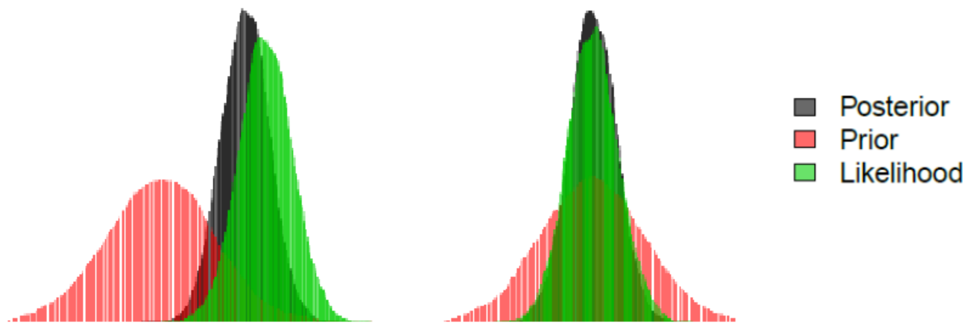

Figure 2: Illustration of model-based prior and posterior distributions with, from left to right, increasing data-model accordance.

2.5 Combining data weight and data-model accordance measures

When combining the data weight measure with the data-model accordance measure, model-based estimates can be divided into 4 categories (see Figure 3 (right)):

- Model-based estimates that are data driven (top row, large data weight) are broken down into
  - "interesting case study" findings (yellow) , when model-based prior and data suggest different point estimates (low data-model accordance);
  - "data driving the estimate; estimate is expected under the model" findings (green) when model and data suggest similar point estimates (high data-model accordance).
- Model-based estimates that are not data driven (bottom row, low data weight) are broken down into
  - "more data are needed" findings (red) when model and data suggest different point estimates (low data-model accordance);
  - "data are limited but there is no strong evidence that data and model are in conflict" findings (white) when model and data suggest similar point estimates (high data-model accordance).

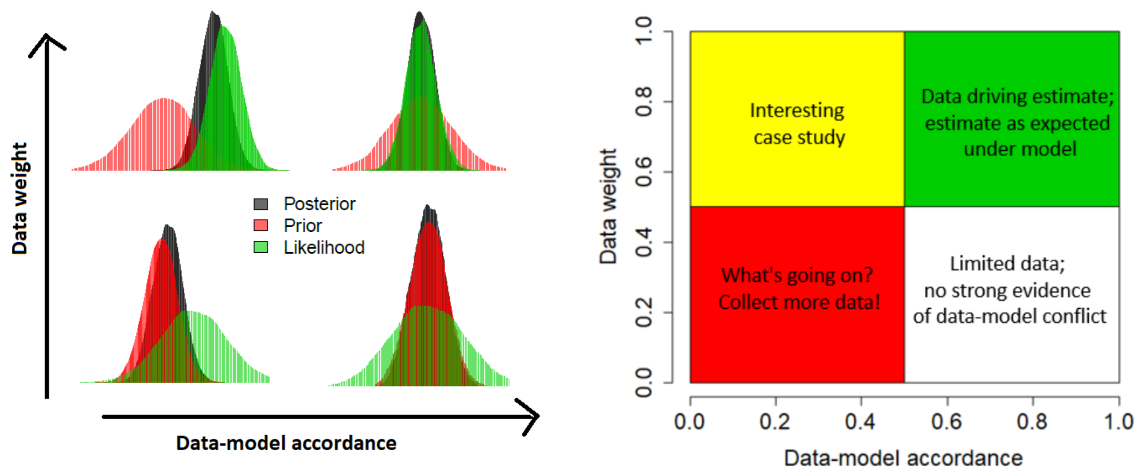

Figure 3: Left: Examples of model-based prior distributions, posteriors, and associated likelihood functions for different data weight and data-model accordance combinations. Right: Definition of categories (quadrants) based on data weight and data-model accordance.

3 Case study: modern contraceptive use estimates for 2017

We use the estimation of modern contraceptive use for 2017, as produced by the Family Planning Estimation Model (FPEM), as a case study to illustrate the application of the two measures.[2;3] In summary, in FPEM total contraceptive use is modeled using a logistic growth curve and ARIMA(1,1,0) time series process. Country-specific parameters of the logistic growth curves are estimated using hierarchical models, where countries are organized into subregions, regions and the world. A similar approach is used for estimating the ratio of modern to total use. The likelihood function accounts for sampling and non-sampling errors (by source type), as well as biases associated with a subset of observations. Estimates are illustrated in Figure 4 for Mauritania.

We assess the following question: What is the contribution of country-specific data after 2012 to the model-based estimates of contraceptive use μ for 2017? The following steps are taken to answer this question for each country:

(1) Construct $p(\mu)$ by fitting the model to a data set that excludes data past 2012 for the respective country.

(2) Compare $p(\mu)$ to posterior $p(\mu|y)$ and calculate w and a.

Preliminary results for all FP2020 countries are given in Figure 5. We selected one country from each quadrant and illustrate the estimates from those countries in

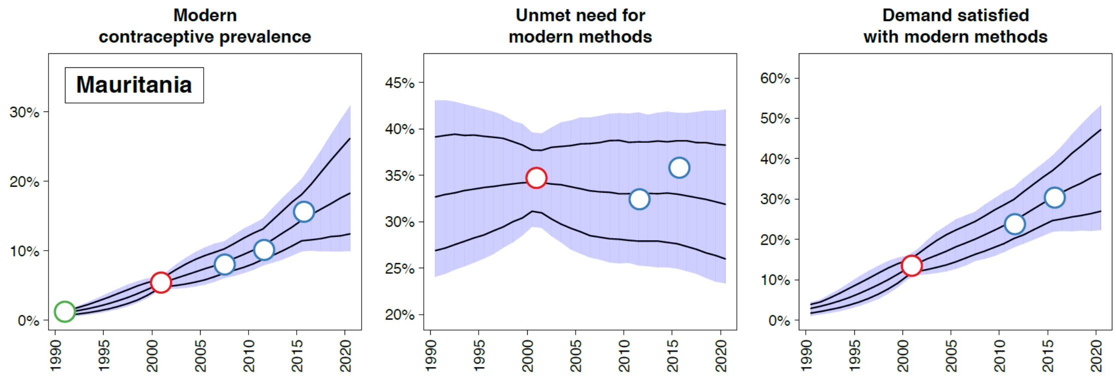

Figure 4: FPEM estimates for Mauritania

Figure 6. Mauritania shows up as an example country with limited data; the data after 2012 are subject to large non-sampling errors. Moreover, the estimates are suggested by the model divert from those suggested by the data. The analysis highlights that additional information is needed for Mauritania for improved monitoring. On the contrary, Kenya is an example of a country where estimates are data-driven. The model and data suggest different trends (the data suggest a greater increase than indicated by the model alone), hence Kenya is flagged as a country that can be used for a case study.

References

[1] Stevens GA, Alkema L, Black RE, Boerma JT, Collins GS, Ezzati M, et al. Guidelines for accurate and transparent health estimates reporting: the GATHER statement. PLoS medicine. 2016;13(6):e1002056.

[2] Cahill N, Sonneveldt E, Stover J, Weinberger M, Williamson J, Wei C, et al. Modern contraceptive use, unmet need, and demand satisfied among women of reproductive age who are married or in a union in the focus countries of the Family Planning 2020 initiative: a systematic analysis using the Family Planning Estimation Tool. The Lancet. 2018;391(10123):870–882.

[3] Alkema L, Kantorova V, Menozzi C, Biddlecom A. National, regional, and global rates and trends in contraceptive prevalence and unmet need for family planning between 1990 and 2015: a systematic and comprehensive analysis. The Lancet. 2013 2017/09/05;381(9878):1642–1652. Available from: http://dx.doi.org/10.1016/S0140-6736(12)62204-1.

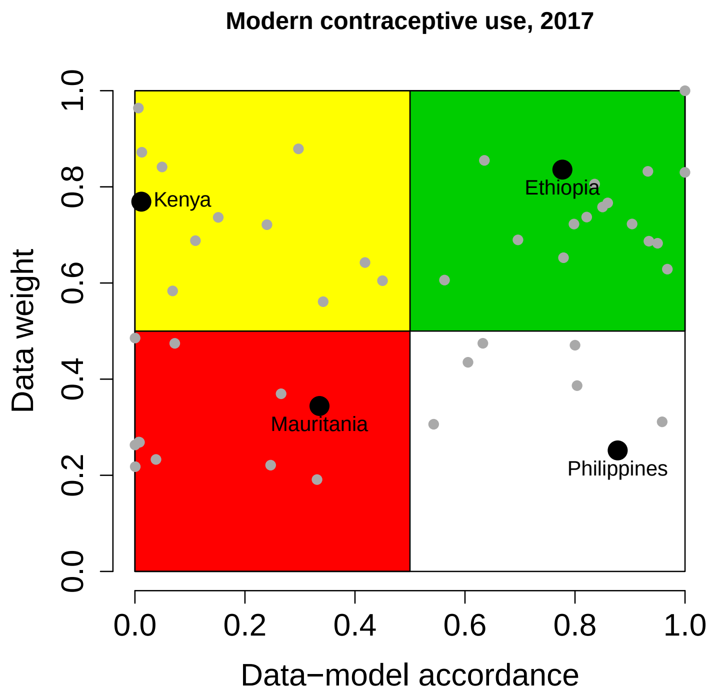

Figure 5: Preliminary results for estimating modern contraceptive use for 2017 for all FP2020 countries, considering the contribution of data after 2012.

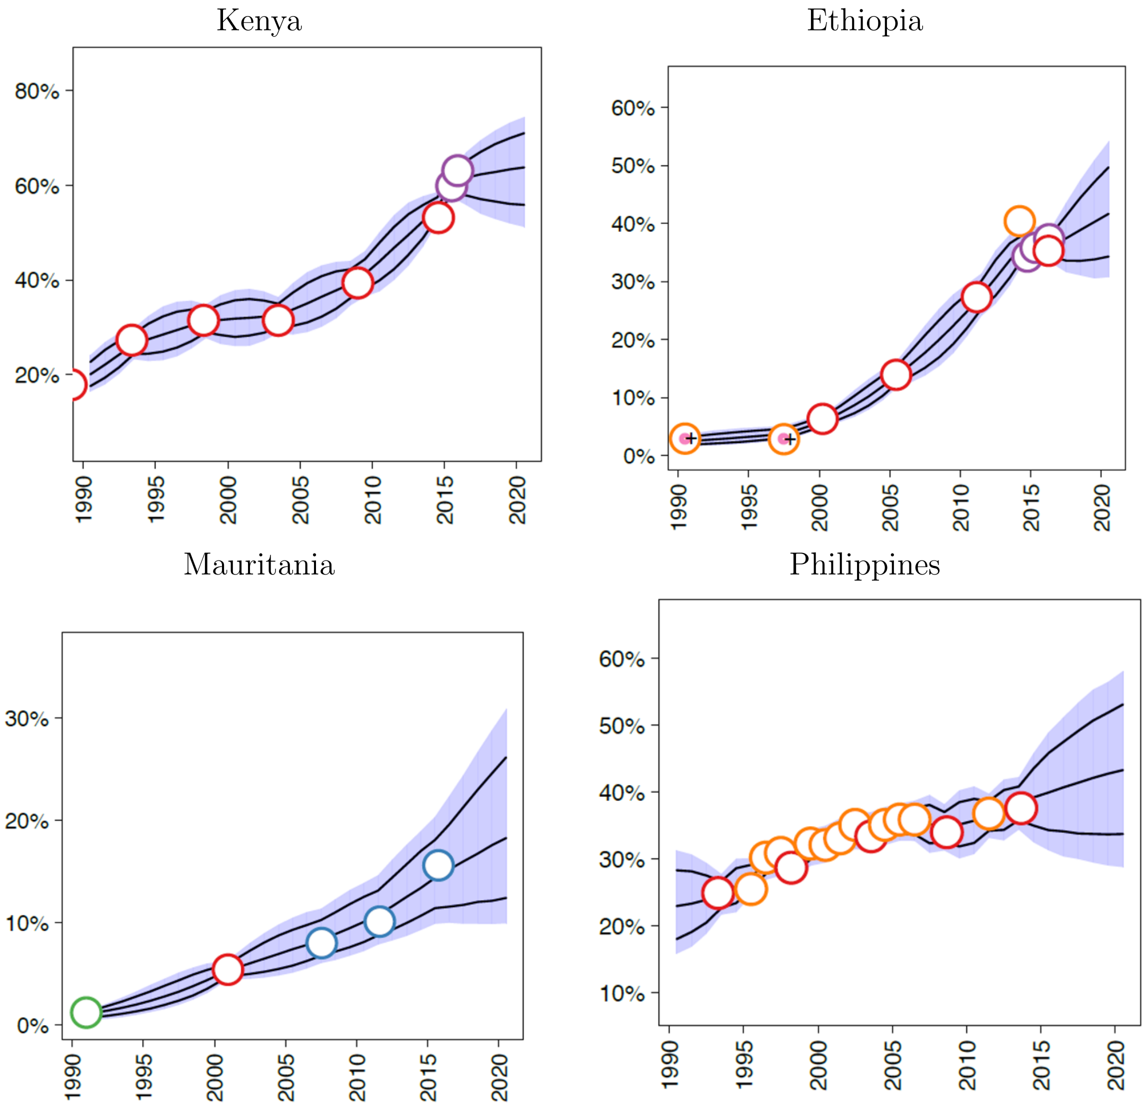

Figure 6: Modern contraceptive use estimates for Kenya, Ethiopia, Mauritania and the Philippines.