mcmsupply: An R package for estimating modern contraceptive method supplies

Hannah Comiskey and Niamh Cahill

Hamilton Institute and Department of Mathematics & Statistics,
Maynooth University, Maynooth, Ireland

March 15, 2024

Chapter 1

Mathematical description of process to estimate and project contraceptive method supply shares at national and subnational administration levels over time using Bayesian hierarchical penalised splines

1.1 Introduction

This document is a summary of the extensions made to the model described in Comiskey et al., (2023) [4]. Described below are the statistical models used to estimate national and subnational contraceptive method supply shares over time using Bayesian hierarchical penalised spline models with multi-country and single-country datasets. These models are utilised in the mcmsupply R package [5].

1.2 Terminology

Contraceptive method supply shares: The proportion of modern contraceptives supplied by the public, private commercial medical and private other sectors over time.

Multi-country national model: This model estimates the contraceptive method supply shares at the national administration level over time for many countries simultaneously.

Single-country national model: This model estimates the contraceptive method supply shares at the national administration level over time for a single country.

Multi-country subnational model: This model estimates the contraceptive method supply shares at the subnational administration level over time for many countries simultaneously.

Single-country subnational model: This model estimates the contraceptive method supply shares at the subnational administration level over time for a single country.

1.3 Overall set-up

The outcome of interest is the components of a compositional vector

$$\boldsymbol{\phi_{q,t,m}} = (\phi_{q,t,m,s=1}, \phi_{q,t,m,s=2}, \phi_{q,t,m,s=3})$$

where, $\phi_{q,t,m,s}$ is the proportion supplied by the public sector (s =1), the private commercial medical sector (s=2) and the other private sector (s=3) of modern contraceptive method m, at time t, in population q (national or subnational).

We begin by defining a regression model for $\phi_{q,t,m,1}$. The logit-transformed proportion, logit$(\phi_{q,t,m,1})$, is modelled through a latent variable $\psi_{q,t,m,1}$, with a penalized basis-spline (P-spline) regression model:

$$\text{logit}\left(\phi_{q,t,m,1}\right) = \psi_{q,t,m,1} = \sum_{k=1}^{K} \beta_{q,m,1,k} B_{q,k}(t), \tag{1.1}$$

where,
$\psi_{q,t,m,1}$ is the latent variable capturing the logit proportions of the public sector (s=1) supply share of method m, at time t, in population q. $B_{q,k}(t)$ refers to the k^{th} basis function evaluated in population q, at time t . $\beta_{q,m,1,k}$ is the k^{th} spline coefficient for the public sector supply (s=1) of method m in population q.

Similarly, we model the latent variable, $\psi_{q,t,m,2}$, to capture the logit-transformed ratio of the private commercial medical supply share to the total private sector share. The model is specified as follows:

$$\text{logit}\left(\frac{\phi_{q,t,m,2}}{1-\phi_{q,t,m,1}}\right) = \psi_{q,t,m,2} = \sum_{k=1}^{K} \beta_{q,m,2,k} B_{q,k}(t), \tag{1.2}$$

where, $\beta_{q,m,2,k}$ is the k^{th} spline coefficient for the ratio of private commercial medical sector (s=2) to total private sector for method m in population q.

The basis functions $B_k(t)$ are constructed using cubic splines. The basis are fitted over the years 1990 to 2025. We align the knot placement of the basis splines with the most recent survey year in each country. As the most recent survey year varies by country (in the case of national-level data) or province (in the case of subnational data), the basis splines $B_{q,k}(t)$ also vary by location.

To estimate the spline coefficients, $\beta_{q,m,s,k}$, we use a random walk model of order 1 on spline coefficients such that the first-order differences, $\boldsymbol{\delta_{q,m,s}}$, are penalized. This model choice is motivated by prior work that used constant projections past the most recent data point [10]. The $\boldsymbol{\delta_{q,m,s}}$ vector is of length h where h=K-1, and K is the total number of knots used in the set of basis functions. It is defined as,

$$\boldsymbol{\delta_{q,m,s}} = (\beta_{q,m,s,2} - \beta_{q,m,s,1}, \beta_{q,m,s,3} - \beta_{m,s,2}, ..., \beta_{m,s,K} - \beta_{q,m,s,K-1}). \tag{1.3}$$

We assume that in population q, for method m and sector s, the value of spline coefficient at knot index k*, aligning with the year t*, the most recent survey available, is $\alpha_{q,m,s}$. By doing this, we are assuming that the $\alpha_{q,m,s}$ parameter will act as the spline coefficient for the reference spline at k*. We are then able to calculate the remaining spline coefficients from the reference index (k*) using the estimated $\boldsymbol{\delta_{q,m,s}}$.

$$\beta_{q,m,s,k} = \begin{cases} \alpha_{q,m,s} & k = k^*, \\ \beta_{q,m,s,k+1} - \delta_{q,m,s,k} & k < k^*, \\ \beta_{q,m,s,k-1} + \delta_{q,m,s,k-1} & k > k^*. \end{cases} \tag{1.4}$$

Where,
$\alpha_{q,m,s}$ is the most recently observed supply share on the logit scale for sector s , method m, in population q. This proxies as an intercept in the model.
k is the knot index along the set of basis splines $B_{q,k}(t)$
k^* is the index of the knot that corresponds with t^*, the year index where the most recent survey occurred in population q.
$\delta_{q,m,s,k-1}$ is the first order difference between spline coefficients $\beta_{q,m,s,K}$ and $\beta_{q,m,s,K-1}$

We assume a smooth transition between spline coefficients. Thus, we centre our rates of change, $\delta_{q,1:M,s,k}$, on 0, with a variance-covariance matrix, Σ_{δ_s}, that captures the correlations that exist between the rates of change in supply shares for each pair of methods.

$$\delta_{q,1:M,s,h} \mid \Sigma_{\delta_s} \sim MVN(\mathbf{0}, \Sigma_{\delta_s}), \tag{1.5}$$

From the latent variable vector, $\boldsymbol{\psi_{q,t,m}}$, it is possible to infer the compositional vector $\boldsymbol{\phi_{q,t,m}}$, such that,

$$\begin{aligned} \phi_{q,t,m,1} &= \text{logit}^{-1}\left(\psi_{q,t,m,1}\right), \\ \phi_{q,t,m,2} &= \left(1 - \phi_{q,t,m,1}\right)\text{logit}^{-1}\left(\psi_{q,t,m,2}\right), \\ \phi_{q,t,m,3} &= 1 - \left(\phi_{q,t,m,1} + \phi_{q,t,m,2}\right). \end{aligned} \tag{1.6}$$

At the national level, the likelihood of the logit-transformed observed data, logit(Y_i), the observed logit-transformed proportions of modern contraceptive method supplied by the public and commercial medical sectors (s=1 and s=2) for method m, at time t in country c, are modelled using Multivariate Normal distributions such that,

$$\text{logit}(\boldsymbol{Y_i}) \mid logit(\boldsymbol{\phi}_{c[i],t[i],m[i],s=1:2}) \sim MVN(logit(\boldsymbol{\phi}_{c[i],t[i],m[i],s=1:2}), \Sigma_{Y_i}), \tag{1.7}$$

Where, $logit(\boldsymbol{\phi}_{c[i],t[i],m[i],s=1:2})$ is the vector of logit-transformed public (s=1) and private commercial medical (s=2) supply proportions for the country, time-point, method associated with observation i.
The variance-covariance matrix, Σ_{Y_i}, utilizes the standard errors (SE) and covariances calculated using the DHS survey microdata associated with the logit-transformed observations, $logit(Y_i)$. Details of the delta-method transformation can be found in the appendix (Chapter **??**, section 4).

At the subnational level, the likelihood of the logit-transformed observed data, logit(y_i), the observed logit-transformed proportion of modern contraceptive method supplied by the public and commercial medical sectors (s=1 or s=2) for method m, at time t are modelled using Normal distributions, such that,

$$\text{logit}(y_i) \mid logit(\phi_{q[i],t[i],m[i],s[i]}) \sim N(logit(\phi_{q[i],t[i],m[i],s[i]}), SE_i^2). \tag{1.8}$$

Where,
$logit(\phi_{q[i],t[i],m[i],s[i]})$ is the logit-transformed supply proportion for the population q, time t, method m, and sector s associated with observation i.
The variance, SE_i^2, utilizes the standard error (SE) calculated using the DHS survey microdata associated with logit-transformed observation y_i. The variance is transformed onto the logit scale using the delta-method [3].

1.4 Estimating parameters for national and subnational models using multi-country and single-country datasets

Summaries of the national and subnational-level models can be found in Figure 1.1 and Figure 1.2. A table of parameters and their interpretations can be found in Table 1.1.

1.4.1 Modelling $\alpha_{q,m,s}$ hierarchically with a multi-country dataset

In this approach, we take advantage of the geographic nature of the dataset. We pool data to estimate precise intercepts at higher geographic levels that then go on to inform more granular level intercepts, until we reach our geographic level of interest (national or subnational), where less data is present.

National-level model

At the national-level, the hierarchical distributions to capture the most recently observed DHS level in country c, for method m, supplied by sector s, are given by:

$$\begin{aligned}
\alpha_{c,m,s}^{country} \mid \theta_{r[c],m,s}^{subcon.}, \sigma_{\alpha,s}^2 &\sim N\left(\theta_{r[c],m,s}^{subcon.}, \sigma_{\alpha,s}^2\right), \\
\theta_{r,m,s}^{subcon.} \mid \theta_{w,m,s}^{world}, \sigma_{\theta,s}^2 &\sim N\left(\theta_{w,m,s}^{world}, \sigma_{\theta,s}^2\right), \\
\theta_{w,m,s}^{world} &\sim N\left(0, 10^2\right), \\
\sigma_{\alpha,s} &\sim Cauchy\left(0,1\right)_+, \\
\sigma_{\theta,s} &\sim Cauchy\left(0,1\right)_+.
\end{aligned} \tag{1.9}$$

Where, the geographic hierarchy begins at the world level $\theta_{w,m,s}^{world}$, which informs the subcontinental intercepts, $\theta_{r,m,s}^{subcon.}$, which in turn inform individual country intercepts, $\alpha_{c,m,s}^{country}$. Vaguely informative Cauchy priors are given to the standard deviation terms of the country- and subcontinental-terms [6]. The standard deviation terms capture the cross-country ($\sigma_{\alpha,s}$) and cross-subcontinent ($\sigma_{\theta,s}$) variation within the data.

Subnational-level model

At the subnational-level, we include an additional layer of geographic intercepts to capture the most recently observed DHS level in subnational province p, for method m,supplied by sector s. While, we use the same notation to explain the hierarchical set up of this approach, the estimates of the country-level and above parameters will be different from the national-level model to the subnational-level model. In the subnational instance, the hierarchical distributions are given by:

$$\begin{aligned}
\alpha^{prov.}_{p,m,s} \mid \alpha^{country}_{c[p],m,s}, \sigma^2_{\alpha_p,s} &\sim N\left(\alpha^{country}_{c[p],m,s}, \sigma^2_{\alpha_p,s}\right),\\
\alpha^{country}_{c,m,s} \mid \theta^{subcon.}_{r[c],m,s}, \sigma^2_{\alpha_c,s} &\sim N\left(\theta^{subcon.}_{r[c],m,s}, \sigma^2_{\alpha_c,s}\right),\\
\theta^{subcon.}_{r,m,s} \mid \theta^{world}_{w,m,s}, \sigma^2_{\theta,s} &\sim N\left(\theta^{world}_{w,m,s}, \sigma^2_{\theta,s}\right),\\
\theta^{world}_{w,m,s} &\sim N\left(0, 10^2\right),\\
\sigma_{\alpha_p,s} &\sim Cauchy\left(0,1\right)_+,\\
\sigma_{\alpha_c,s} &\sim Cauchy\left(0,1\right)_+,\\
\sigma_{\theta_s} &\sim Cauchy\left(0,1\right)_+.
\end{aligned} \tag{1.10}$$

In this instance, we mirror the geographic hierarchy of the national model, and add an additional layer to reflect the province-level intercepts, $\alpha^{prov.}_{p,m,s}$, of the subnational level model, and cross-provincial variation ($\sigma_{\alpha_p,s}$).

1.4.2 Modelling $\alpha_{q,m,s}$ using informative priors with a single-country dataset

In this approach, priors for higher-population level intercept parameters are informed from the multi-country national- or subnational-level models (i.e., the models that used multi-country datasets).

National-level model

$\alpha^{country}_{c,m,s}$ is the national-level intercept for country c, method m and sector s, informed by the posterior median estimates of the subcontinental level model interceptand the associated variance parameter estimated from the multi-country national model, such that

$$\alpha^{country}_{c,m,s} \mid \hat{\theta}^{subcon.}_{r[c],m,s}, \hat{\sigma}^2_{\alpha_c,s} \sim N\left(\hat{\theta}^{subcon.}_{r[c],m,s}, \hat{\sigma}_{\alpha_c,s}\right), \tag{1.11}$$

where, $\hat{\theta}^{subcon.}_{r[c],m,s}$ is the posterior median UNSD subcontinental population intercept for region r, method m, sector s, associated with country c estimated from the national-level multi-country model and $\hat{\sigma}^2_{\alpha_c,s}$ is the posterior median of the sector specific cross-country variation associated with the $\alpha_{c,m,s}$ intercept estimated from the national-level multi-country model.

Subnational-level model

$\alpha^{prov.}_{p,m,s}$ is the subnational-level intercept for subnational province p, method m, sector s, informed by the posterior median estimates of the country-level model interceptand the associated variance parameter estimated from the multi-country subnational model, such that

$$\alpha^{prov.}_{p,m,s} \mid \hat{\alpha}^{country}_{c,m,s}, \hat{\sigma}^2_{\alpha_p,s} \sim N\left(\hat{\alpha}^{country}_{c,m,s}, \hat{\sigma}^2_{\alpha_p,s}\right), \tag{1.12}$$

where, $\hat{\alpha}^{country}_{c[p],m,s}$ is the posterior median national-level population intercept for country c, method m, sector s, associated with the subnational province p, estimated from the subnational-level multi-country model and $\hat{\sigma}^2_{\alpha_p,s}$ is the posterior median of the cross-province variation associated with the $\alpha_{p,m,s}$ intercept estimated from the subnational-level multi-country model.

1.4.3 Modelling Σ_{δ_s} using cross-method correlations with a multi-country dataset

In this approach, we decompose the Σ_{δ_s} into its variance and correlation matrices and estimate the components separately. This is a two-model run approach which involves estimating the correlations using a model run with correlations set to 0.

For both the national and subnational models, a multivariate normal prior centred on 0 was assigned to the vector of length M of first-order differences of the spline coefficients, $\delta_{q,1:M,s,h}$, for population q (national or subnational), using all methods supplied by sector s at first-order difference h,

$$\delta_{q,1:M,s,h} \mid \Sigma_{\delta_s} \sim MVN(\mathbf{0}, \Sigma_{\delta_s}), \tag{1.13}$$

where,

$$\Sigma_{\delta_s} = \begin{bmatrix} \sigma^2_{\delta_{1,s}} & \hat{\rho}_{1,2,S}\sigma_{\delta_{1,S}}\sigma_{\delta_{2,s}} & \cdots & \cdots & \hat{\rho}_{1,M,S}\sigma_{\delta_{1,s}}\sigma_{\delta_{M,S}} \\ \hat{\rho}_{2,1,s}\sigma_{\delta_{2,s}}\sigma_{\delta_{1,s}} & \sigma^2_{\delta_{2,s}} & \cdots & \cdots & \hat{\rho}_{2,M,S}\sigma_{\delta_{2,s}}\sigma_{\delta_{M,s}} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ \hat{\rho}_{M,1,s}\sigma_{\delta_{M,s}}\sigma_{\delta_{M,s}} & \cdots & \cdots & \cdots & \sigma^2_{\delta_{M,s}} \end{bmatrix}. \tag{1.14}$$

The correlation terms of the covariance matrix, $\rho_{i,j,s}$, were estimated using a maximum a posteriori estimator for the correlation matrix as described in Azose and Raftery, 2018 [1]. This approach involves fitting a model where the covariance terms in $\sigma_{\delta_{j,s}}$ are set equal to zero. The national and subnational models deviate in terms of the geographic level that they estimate these correlations at. For the national model, we estimate the correlations based on country-level rates of change, $\delta_{c,1:M,s,h}$, terms. This captures the correlations across all countries. While in the subnational model, we estimate the correlations using provincial-level rates of change, $\delta_{p,1:M,s,h}$. In this instance, the correlations captured are across the subnational provinces of all countries.

National-level model

In the national model, the deviation terms of the Σ_{δ_s} matrix are given vague uniform priors,

$$\sigma_{\delta_{m,s}} \sim Uniform\,(0, 10)\,. \tag{1.15}$$

The 0-covariance model estimates are used to estimate the correlation between methods across time and all countries at the national-level. Specifically, for sector s, the correlation between method i and method j is calculated as follows,

$$\hat{\rho}_{i,j,s} = \frac{\sum_{c=1}^{C}\sum_{h=1}^{K-1}\tilde{\delta}_{c,m[i],s,h}\tilde{\delta}_{c,m[j],s,h}}{\sqrt{\sum_{c=1}^{C}\sum_{h=1}^{K-1}\tilde{\delta}^2_{c,m[i],s,h}}\sqrt{\sum_{c=1}^{C}\sum_{h=1}^{K-1}\tilde{\delta}^2_{c,m[j],s,h}}}, \tag{1.16}$$

Where, $\tilde{\delta}_{c,m[j],s,h}$ are the estimated first order differences of the spline coefficients for country c, method m, sector s, at the h-th difference between spline coefficients. They are given by the posterior medians of $\delta_{c,m,s,h}$ from the zero-covariance run, after subsetting the period considered to periods with data within a country. C represents the total number of countries involved in the study. K is the number of knots in the basis functions. h represents the number of differences (h=K-1) between the spline coefficients.

Subnational model

In the subnational model, the deviation terms of the Σ_{δ_s} matrix are given vague Cauchy priors. This prior is suggested as a weakly informative prior in the paper titled Prior distributions for variance parameters in hierarchical models by Gelman, Bayesian Analysis (2006) [6].

$$\sigma_{\delta_{m,s}} \sim Cauchy\,(0, 1)_+\,. \tag{1.17}$$

The 0-covariance model estimates are used to estimate the strength of the correlations between methods across time and all provinces in all countries at the subnational-level. Specifically, for sector s, the correlation between method i and method j is calculated as follows,

$$\hat{\rho}_{i,j,s} = \frac{\sum_{p=1}^{P}\sum_{h=1}^{K-1}\tilde{\delta}_{p,m[i],s,h}\tilde{\delta}_{p,m[j],s,h}}{\sqrt{\sum_{p=1}^{P}\sum_{h=1}^{K-1}\tilde{\delta}^2_{p,m[i],s,h}}\sqrt{\sum_{p=1}^{P}\sum_{h=1}^{K-1}\tilde{\delta}^2_{p,m[j],s,h}}}, \tag{1.18}$$

Where, $\tilde{\delta}_{p,m[j],s,h}$ are the estimated first order differences of the spline coefficients for province p, method m, sector s, at the h-th difference between spline coefficients. They are given by the posterior medians of $\delta_{p,m,s,h}$ from the zero-covariance run, after subsetting the period considered to periods with data within each province. P represents the total number of subnational provinces across all countries involved in the study. K is the number of knots in the basis functions. h represents the number of differences (h=K-1) between the spline coefficients.

1.4.4 Modelling Σ_{δ_s} using informative priors with a single-country dataset

For estimation of the method supply shares using a single-country national or subnational dataset, we set Σ_{δ_s} as the median estimate of the MxM variance-covariance matrix from the corresponding (national or subnational) multi-country model, $\Sigma_{\delta_s}^{\textbf{global}}$ and we estimate the first-order difference spline coefficients using a Multivariate Normal prior centred on zero such that,

$$\delta_{1:M,s,h} \mid \hat{\Sigma}_{\delta_s}^{\textbf{global}} \sim MVN(\mathbf{0}, \hat{\Sigma}_{\delta_s}^{\textbf{global}}). \tag{1.19}$$

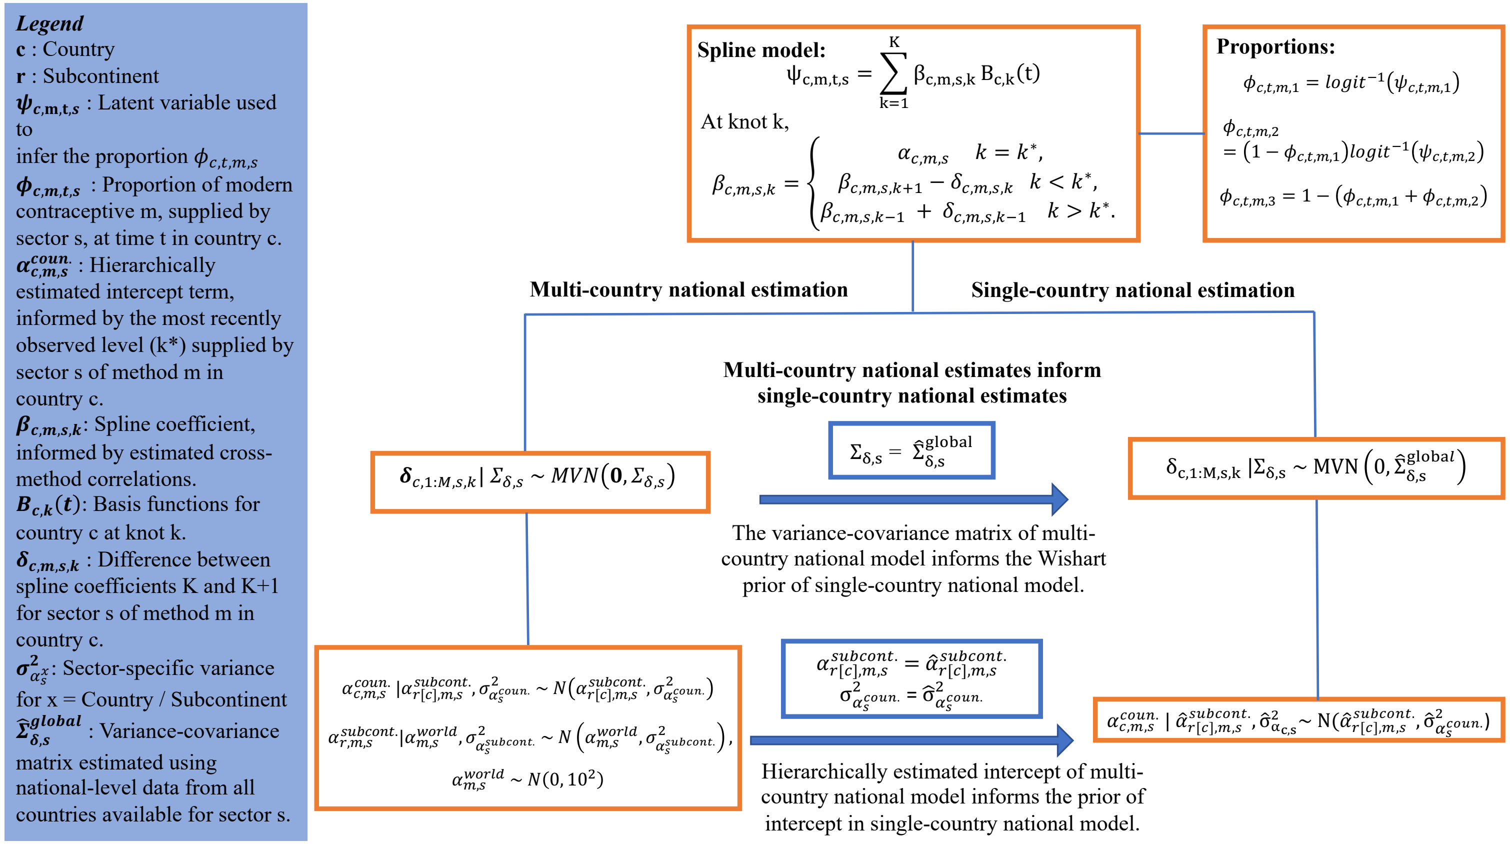

Figure 1.1: Schematic linking the multi-country and single-country national-level modelling approaches.

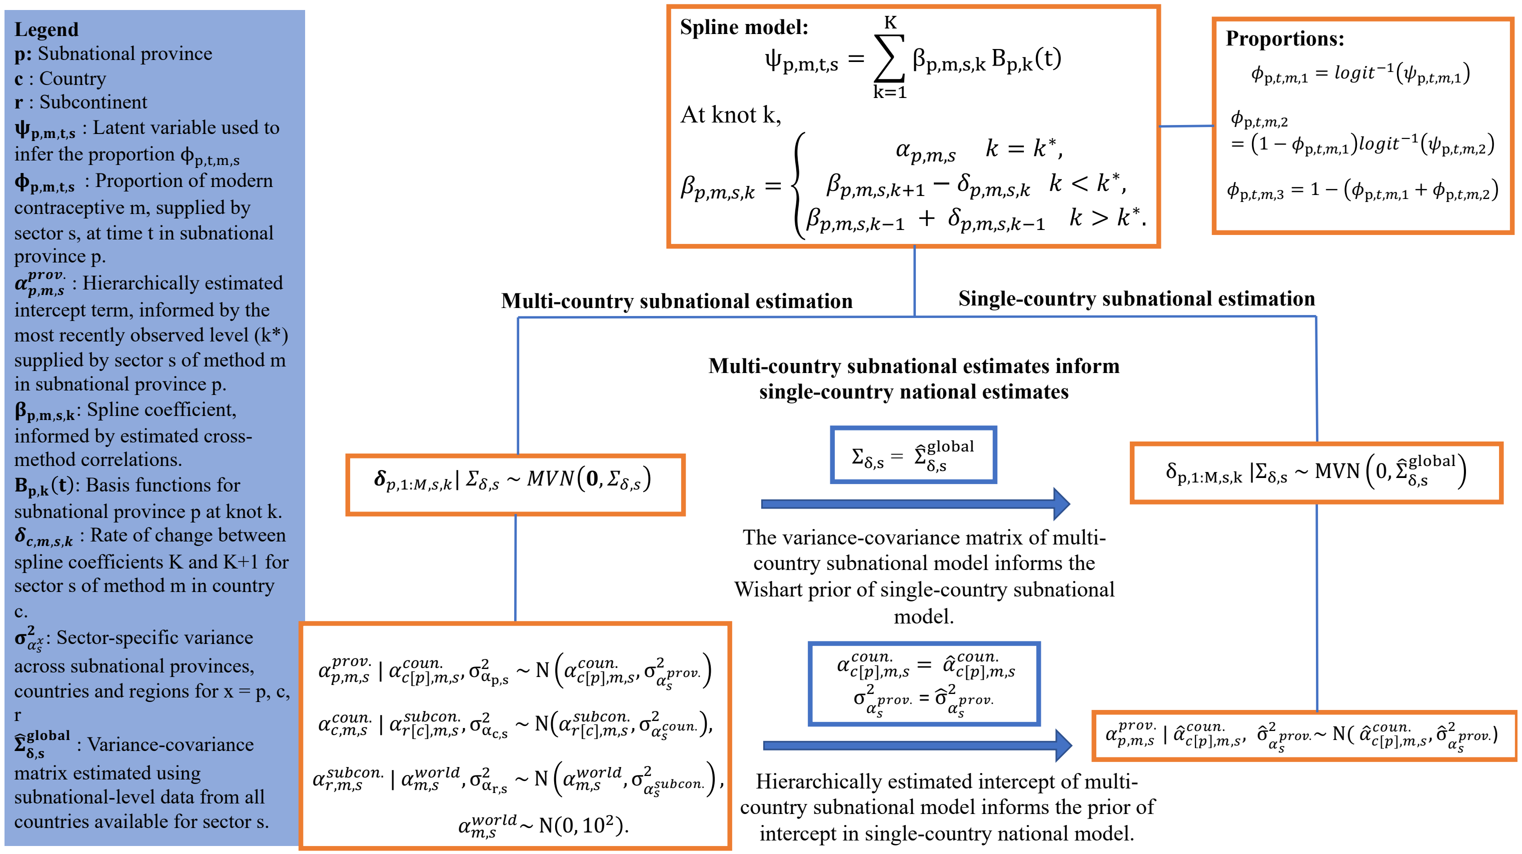

Figure 1.2: Schematic linking the multi-country and single-country subnational-level modelling approaches.

Model Parameter	Interpretation
$\phi_{q,t,m,s}$	The proportion supplied by sector s, of modern contraceptive method m, at time t, in population q (national or subnational).
$\psi_{q,t,m,s}$	The latent variable used to model $\phi_{q,t,m,s}$ on the logit scale
$\beta_{q,m,s,k}$	The k^{th} spline coefficient for sector s, method m in population q
$B_{q,k}(t)$	The set of basis functions for population q, evaluated at knot k for time t.
$\alpha^{pop.}_{q,m,s}$	The most recently observed supply share on the logit scale for sector s , method m, in population q. This proxies as an intercept in the model.
$\delta_{q,m,s,k-1}$	The first order difference between spline coefficients $\beta_{q,m,s,K}$ and $\beta_{q,m,s,K-1}$
Σ_{δ_s}	Variance-covariance matrix used in the MVN prior of $\delta_{q,1:M,s,h}$
$\sigma_{X,s}$	Standard deviation terms relating to the intercept parameter X for sector s. X may be at the provincial-, country-, or subcontinental-level.
$\rho_{i,j,s}$	Correlation between the rates of change in supply shares for method[i] and method[j] in sector s.

Table 1.1: A table of parameters names and their interpretations across the national and subnational models. The indexing refers to sector s, of modern contraceptive method m, at time t, in population q (national or subnational).

Chapter 2

Validation of the multi-country and single-country subnational models and the single-country national model

2.1 Data

2.1.1 National data source

In this study we consider countries involved in the FP2030 initiative. A database of the public and private sector breakdown of modern contraceptive supply with their associated standard errors at the national administration level was created using data from the DHS [11]. Table 2.1 lists the thirty countries used in the national level method supply share database. The total number of surveys carried out and the year of the most recent survey is listed for each country. A full description of the national level data can be found in Comiskey et al., 2023 [4].

UNSD intermediate world regions	Country	Total Number of Surveys	Recent Survey Year
Southern Asia	Afghanistan	1	2015
Western Africa	Benin	5	2017
Western Africa	Burkina Faso	4	2010
Middle Africa	Cameroon	5	2018
Middle Africa	Congo	1	2005
Middle Africa	Democratic Republic of Congo	2	2013
Western Africa	Cote d'Ivoire	3	2011
Eastern Africa	Ethiopia	5	2019
Western Africa	Ghana	5	2014
Western Africa	Guinea	4	2018
Southern Asia	India	4	2005
Eastern Africa	Kenya	5	2014
Western Africa	Liberia	4	2019
Eastern Africa	Madagascar	4	2008
Eastern Africa	Malawi	5	2015
Western Africa	Mali	5	2018
Eastern Africa	Mozambique	3	2011
South-Eastern Asia	Myanmar	1	2015
Southern Asia	Nepal	5	2016
Western Africa	Niger	4	2012
Western Africa	Nigeria	5	2018
Southern Asia	Pakistan	4	2017
South-Eastern Asia	Philippines	6	2017
Eastern Africa	Rwanda	6	2019
Western Africa	Senegal	10	2019
Western Africa	Sierra Leone	3	2019
Eastern Africa	Tanzania	6	2015
Western Africa	Togo	2	2013
Eastern Africa	Uganda	5	2016
Eastern Africa	Zimbabwe	5	2015

Table 2.1: Summary of DHS microdata used for the validation of the national level estimation. This table provides the United Nation Statistics Division (UNSD) intermediate world region names, country names, the number of DHS surveys per country available and the year of the most recent DHS survey available. Just over 46% of countries have data available after 2015.

2.1.2 Subnational data source

In this study we consider countries involved in the FP2030 initiative. A database of administration-1 level Demographic and Health Survey (DHS) data observations for the supply of modern contraceptive methods by the public and private sectors and their associated standard errors was created using the Integrated Public Use Microdata Series (IPUMS) project, IPUMS-DHS [7]. Like the national-level study [4], the modern methods of contraception considered in this study are female sterilisation, oral contraceptive pills (OC pills), implants (including Implanon, Jadelle and Sino-implant), intra-uterine devices (IUD, including Copper- T 380-A IUD and LNG-IUS), and injectables (including Depo Provera (DMPA), Noristerat (NET-En), Lunelle, Sayana Press and other injectables). The variables contained within the IPUMS-DHS database are consistent over time and space. IPUMS-DHS uses integrated geography variables for a country across sample years to address issues with subnational boundaries changing over time and enable comparisons over time. Table 2.2 lists the 23 countries captured in this subnational database. The total number of administration level 1 (admin-1) subnational regions, the number of DHS surveys each country has in the database, and the year of the most recent survey in the database is listed for each country. Just under half of the countries included have survey data available after 2015, highlighting the need for annual up-to-date estimates of the contraceptive supply shares.

Country	Number of admin-1 level subnational provinces	Number of IPUMS-DHS surveys	Recent survey year
Benin	6	4	2017
Burkina Faso	13	4	2010
Cameroon	3	3	2004
Congo Democratic Republic	5	2	2013
Cote d'Ivoire	15	3	2011
Ethiopia	10	4	2016
Ghana	8	5	2014
Guinea	3	4	2018
India	27	4	2015
Kenya	8	5	2014
Liberia	5	2	2013
Madagascar	6	4	2008
Malawi	3	5	2016
Mali	4	5	2018
Mozambique	11	3	2011
Nepal	5	5	2016
Niger	6	4	2012
Pakistan	6	4	2017
Rwanda	7	6	2014
Senegal	4	9	2017
Tanzania	6	6	2015
Uganda	4	4	2016
Zimbabwe	10	5	2015

Table 2.2: Summary information regarding the countries considered for subnational modelling. The name, number of subnational administration level 1 (admin-1) regions, the total number of DHS surveys present in the data, and the year of the most recent DHS survey in the data for each country are listed.

2.1.3 Calculating the variance-covariance matrices of the DHS national level supply share observations using DHS microdata

Using the 'svymean' function from the 'survey' package in R [9], the proportions of each method supplied the public and privates sectors for the specific surveys in the countries listed above were calculated along with their associated variance-covariance matrices. The 'survey' package uses the Taylor series linearisation method to approximate the standard error of the calculated proportions [2] [8]

2.1.4 Standard error calculation

Using the 'survey' package in R, the national-level method supply shares and their associated variance-covariance matrices were calculated using the 'svymean' function. The subnational-level method supply share databases use the 'svyciprop' function to calculate the standard errors associated with each observation [9]. Using DHS design factors, we impute calculated standard errors when the calculated standard error is 0. A full description of how the standard errors are calculated and the imputation technique used to estimate standard errors can be found in the supplementary material of Comiskey et al., (2023).

Measure	**Range (% over all methods)**	**Median SE size (% over all methods)**	**Largest mean SE (method, %)**		**Smallest mean SE (method, %)**	
Result	0.015 , 18.19	2.23	IUD	4.10	Injectables	2.03

Table 2.3: Summary table for the calculated standard errors of the data observations.

From Table 2.3, the calculated national-level standard errors range from 0.015 to 18.19 percentage points. The median standard error size across all method is 2.23 percentage points. On average, they tend to be largest for IUDs where the mean standard error size is approximately 4 percentage points and smallest in injectables where the mean standard error size is approximately 2 percentage points.

Measure	**Range (% over all methods)**	**Median SE size (% over all methods)**	**Largest mean SE (method, %)**		**Smallest mean SE (method, %)**	
Result	0.0 , 22.0	3.84	OC pills	5.3	Implants	3.6

Table 2.4: Summary table for the calculated standard errors of the subnational IPUMS-DHS data observations.

From Table 2.4, the calculated subnational-level standard errors range from 0 to 22 percentage points. The median standard error size across all method is 3.8 percentage points approximately. On average, they tend to be largest for OC pills where the mean standard error size is approximately 5 percentage points and smallest in implants where the mean standard error size is approximately 4 percentage points. The calculated standard errors of the subnational IPUMS-DHS data are larger than those calculated using the DHS national-level data (Table 2.3) . At the national level, the median standard error (across all methods) is 2.23%. This is almost half the size of the subnational median. Similarly, at the national-level IUDs have the largest mean standard error (4.10%). This is almost 1% smaller than the 5.3% observed for OC pills at the subnational level.

2.2 Out-of-sample validation results

2.2.1 Errors and coverage

We calculate sector specific error terms, $e_{j,s}$, to describe the difference between the observed data point j in sector s, $y_{j,s}$, and the median estimate from the posterior predictive distribution, $y_{j,s}$ such that,

$$e_{j,s} = y_{j,s} - \hat{y}_{j,s}. \tag{2.1}$$

We evaluated the results of the validation using different measures of accuracy and prediction interval calibration. To evaluate the accuracy of our model, we considered the root mean square error (RMSE) for each sector's set of estimates.

Let,

$$\text{RMSE}_s = \sqrt{\frac{\sum_{j=1}^{N_s} e_{j,s}^2}{N_s}}, \tag{2.2}$$

where, N_s is the number of observations in the sector s. $e_{j,s}$ is the error for observation j in sector s which is described above. The RMSE can be interpreted as the average error observed across all countries, time points and methods in the test set.

We also evaluated the mean error (eq. 3.6) and the median absolute errors (eq. 3.7). The mean error is the average difference between the observed proportion and true proportion estimated by the model and is an effective measurement of bias within the model. When the mean error is positive, this indicates systematic under-prediction by the model and conversely, a negative mean error indicates that the model is over-estimating the observed data. Median absolute error is the 50^{th} percentile of absolute differences between the observed proportion and true proportion estimated by the model. Median absolute error captures the overall variation within the model estimates.

$$\text{Mean error}_s = \frac{\sum_{j=1}^{N_s} e_{j,s}}{N_s}, \tag{2.3}$$

$$\text{Median absolute error}_s = \text{Median}(|\boldsymbol{e_s}|), \tag{2.4}$$

where, $\boldsymbol{e_s}$ is is a vector of length N_s, containing the complete set of errors estimates for all observations belonging to sector s.

Coverage assumes that if our model is correctly calibrated, then for each sector the model should be able to capture the test set of out-of-sample observations with 95% accuracy, where the remaining 5% of incorrectly estimated observations are approximately evenly distributed above and below the estimated 95% prediction interval. To examine the bias of our models estimates, we examined the location of the incorrectly estimated test set observations. We consider the proportion of test observations located above and below the estimated prediction intervals. By examining the breakdown of locations, we are evaluating the tendency of the model to under- or over-estimate the test set. If a higher proportion of observations are located below the prediction intervals, this indicates that the model is tending to over-estimate the test set. Similarly, if a higher proportion of the incorrectly estimated observations are located above the prediction intervals, the model is tending to under-estimate the test set.

2.2.2 Multi-country national model with cross-method correlations

The validation results (out-of-sample validation results and a comparison of the model-based estimates to the direct estimates) with a discussion of these results for the multi-country national model with cross-method correlations can be found in the supplementary materials associated with Comiskey et al., (2023).

2.2.3 Multi-country subnational model with cross-method correlations

Sector	95% coverage (%)	Root mean square error (RMSE) (%)	Proportion of incorrectly estimated observations located above and below the prediction interval (PI) boundary (%)		95% PI width (%)	Mean error (%)	Median absolute error (%)
Commercial medical	95.3	15.5	Above	2.84	68.4	-2.32	7.25
			Below	1.90			
Other	98.1	6.13	Above	1.18	33.3	-0.17	1.0
			Below	0.71			
Public	97.2	15.4	Above	0.71	72.2	2.49	7.08
			Below	2.13			

Table 2.5: Out-of-sample validation results for the test set using multi-country subnational model with cross-method correlations. Coverage is the proportion of the test set observations that are captured within the 95% prediction interval (PI) produced by the model. RMSE is root mean square error. The 95% PI width reflects the median PI width for each observation estimated by the model.

The multi-country subnational model has been evaluated using various out-of-sample model validation measures to gauge its effectiveness at estimating the method supply shares at a subnational level, while also considering the prediction intervals it uses to produce these estimates. It is performing reasonably well considering the complex nature of the data. It has an overall coverage of approximately 97%. The results for the out-of-sample validation are found in Table 2.5. The target coverage is 95%. The model is reasonably well calibrated to the data with the public sector having 97% coverage and the commercial medical sector having 95% coverage. The private other sector has 98% coverage of the test set. The public and private other sectors have coverage of that test set that is slightly higher than expected. The private commercial medical is showing optimal coverage at 95%. The root mean square error (RMSE) for the private commercial medical sector and the public sector are both at approximately 15 percentage points. The private other sector has an RMSE of approximately 6 percentage points. We also considered where the incorrectly estimated test set observations lie with respect to the prediction interval bounds to assess the bias of the model. In theory, if the model is unbiased and well calibrated then we would expect an equal proportion of incorrectly estimated observations above and below the prediction interval boundaries. Both the commercial medical and other sector has a higher proportion of observations above the prediction interval boundary. This would imply that the model tends to under-estimate the observations in these sectors. In the public sector, there is a higher proportion of incorrectly estimated test set observations below the prediction interval. This implies that for this sector, the model tends to over-estimate the public sector. The median width of the 95% prediction intervals is largest in the public sector at 72 percentage points. The private other sector has the smallest median 95% prediction interval width at 33 percentage points. The mean error for the private commercial medical is approximately -2 percentage points. The mean error for the public sector is the absolute largest of all three sectors at approximately 2.5 percentage points. The private other sector has a mean error of less than 1 percentage point. The median absolute error of the private commercial medical sector is the largest at approximately 7 percentage points while the median absolute error of the private other sector was the smallest at approximately 1 percentage point.

2.2.4 Comparison of the model-based estimates to the direct estimates

In Figure 2.1, we consider the observed standard errors calculated using the DHS microdata versus the corresponding standard deviations of the model-based estimates for the proportions. From this figure, we can see that in the commercial medical and public sectors, the estimated standard deviation terms are smaller than the observed standard error terms calculated using the DHS microdata. For both sectors, we can see that most of the observations have standard errors up to 15 percentage points approximately. These same observations when estimated within the model have corresponding standard deviations of approximately up to 10 percentage points. The use of this model results in a considerable reduction in the uncertainty of these observations. The observed SEs of approximately 3 percentage points correspond to the observations where the standard errors were imputed. The outlier observation at 25 percentage points in both the commercial medical and public sector corresponds to IUDs in Maputo City, Mozambique. IUDs in Maputo only has one observation in 1997, whereas the other methods have more recent survey observations to inform model estimates. The lack of data for IUDs in this instance causes the model estimates to have larger uncertainty.

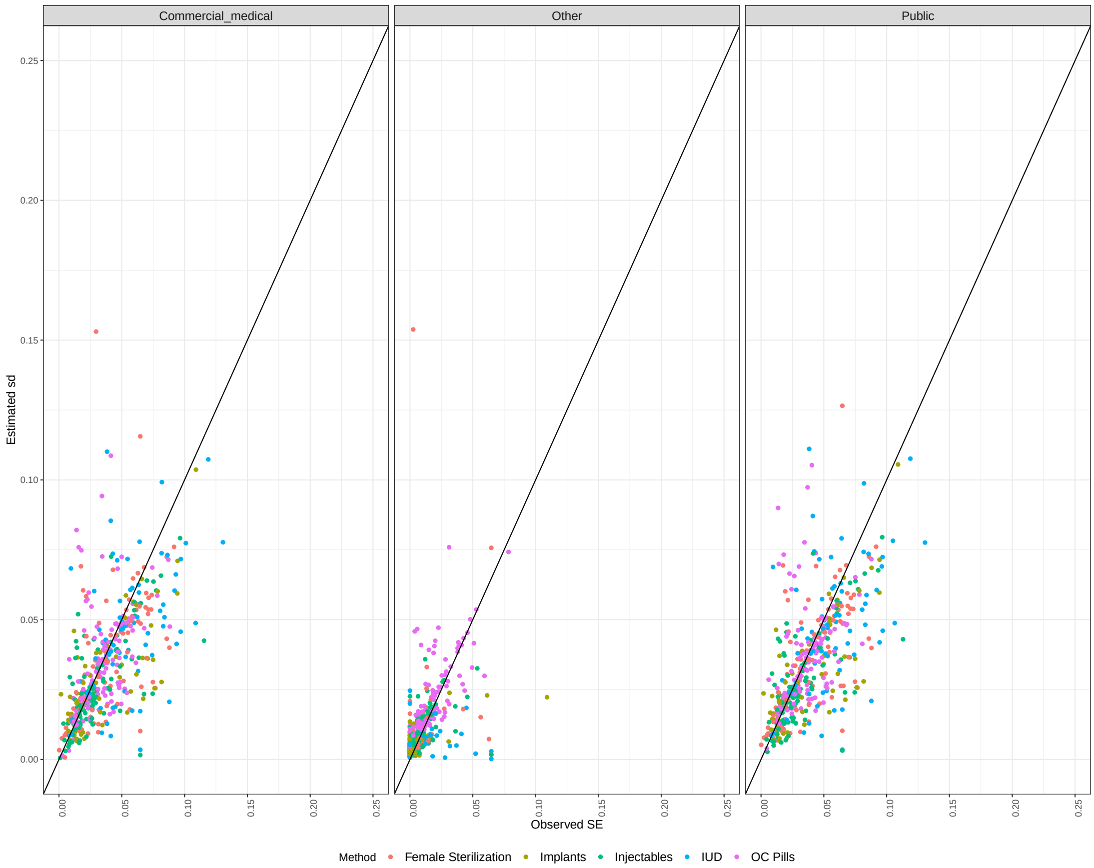

Figure 2.1: A scatter plot across the three sectors comparing the standard error of the direct estimates to the model-based estimates at the subnational level.

In Figure 2.2, we consider the sample size with respect to the ratio of observed to estimated proportions, both of which are on the log scale for clarity. From this figure we see that as the sample size increases, the ratio of observed to estimated data point tends towards 1 after approximately log(5), which corresponds approximately to a sample size of 148. Therefore, the model's ability to capture the observed data point increases as the sample size associated with each observation increases. This aligns with the same property that is seen in many small area estimation models.

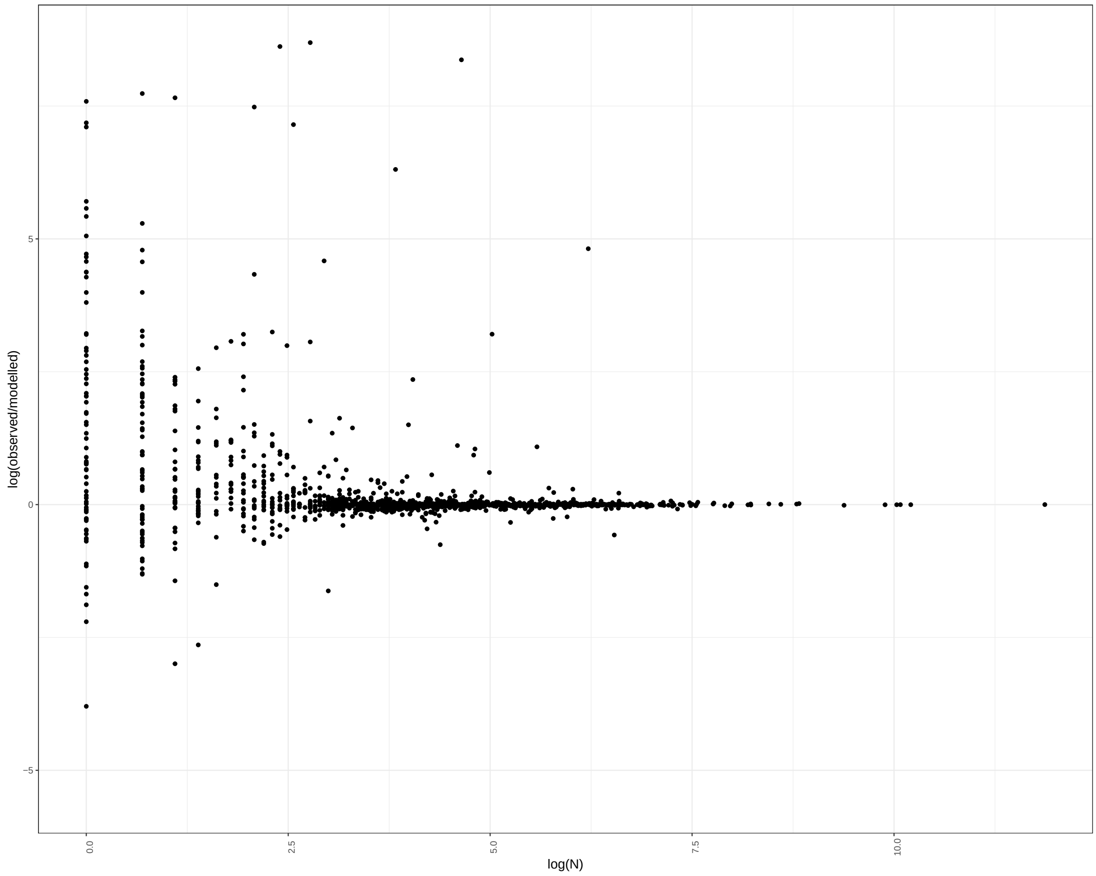

Figure 2.2: A scatterplot of the estimates with log of the sample size on the x-axis and the log ratio of direct to model-based estimates on the y-axis.

2.2.5 Model comparison

Multi-country subnational model with 0-covariance

For the subnational multi-country model, we used a 0-covariance model to validate the use of the cross-method correlations within the estimation of $\delta_{q,1:M,s,h}$, the first-order differences between spline coefficients. This approach is similar to that described in Comiskey et al., (2023). In this instance, the off-diagonal elements of Σ_{δ_s} are set to 0. The variance-covariance matrix Σ_{δ_s} informs the multi-variate normal prior of $\delta_{q,1:M,s,h}$.

As before, we describe the first-order differences between spline coefficients, $\delta_{q,1:M,s,h}$, using a Multi-variate Normal prior centred on 0 with variance-covariance matrix Σ_{δ_s}.

$$\delta_{q,1:M,s,h} \mid \Sigma_{\delta_s} \sim MVN(\mathbf{0}, \Sigma_{\delta_s}), \tag{2.5}$$

such that, Σ_{δ_s} is a diagonal matrix with 0 on the off-diagonal elements;

$$\Sigma_{\delta_s} = \begin{bmatrix} \sigma^2_{\delta_{1,s}} & 0 & \dots & \dots & 0 \\ 0 & \sigma^2_{\delta_{2,s}} & \dots & \dots & 0 \\ 0 & \dots & \dots & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & \dots & \dots & \dots & \sigma^2_{\delta_{M,s}} \end{bmatrix}. \tag{2.6}$$

Having 0-covariance between the variance terms of Σ_{δ_s} implies that the rates of change in method supplies act independently of one another. We compare the validation results for this simpler 0-covariance model with the model using cross-method correlations to investigate the impact of including cross-method correlations in the model estimation process.

Sector	95% coverage (%)	Root mean square error (RMSE) (%)	Proportion of incorrectly estimated observations located above and below the prediction interval (PI) boundary (%)		95% PI width (%)	Mean error (%)	Median absolute error (%)
Commercial medical	96.4	17.4	Above	2.37	68.9	-3.02	7.35
			Below	1.18			
Other	98.1	6.77	Above	1.42	33.4	0.768	0.976
			Below	0.474			
Public	97.9	14.8	Above	0.474	72.8	2.30	7.21
			Below	1.66			

Table 2.6: Out-of-sample validation results for the test set using multi-country subnational 0-covariance model. Coverage is the proportion of the test set observations that are captured within the 95% prediction interval (PI) produced by the model. The MAE is the median absolute error. RMSE is root mean square error

The coverage of the 0-covariance model (Table 2.6) is higher than that of the cross-method correlation model (Table 2.5). The 0-covariance model has 98% coverage in both the public and private other sectors. The commercial medical sector has 96% coverage (Table 2.6). The coverage of the model with cross-method correlations is Commercial medical = 95%, Other = 98%, Public = 97% (Table 2.5).

To evaluate the bias and variance produced by the 0-covariance and cross-method correlation models, we consider the mean errors, median absolute errors (MAE) and root mean square errors (RMSE). Across all three sectors, the RMSE of the 0-covariance model is larger than the cross-method correlation model. The private commercial medical sector has the largest

RMSE with an average error of approximately 17 percentage points(Table 2.6). The RMSE of the private commercial medical sector in the cross-method correlation model is approximately 2 percentage points smaller at 15 percentage points (Table 2.5). In both models, the private other sector has the smallest RMSE. In the 0-covariance model, it is approximately 7 percentage points whereas in the cross-method correlation model it is approximately 6 percentage points. Overall, the cross-method correlation model performs better in this model validation measure than the 0-covariance model. In both models the mean error on the private commercial medical sector is negative (-3.02 percentage points in the 0-covariance model and -2.32 percentage points in the cross-method correlation model) and the mean error of the private other and public sectors are positive. This implies that both models over-predict the test set of the commercial medical sector and under-predict the remaining two sectors. The median absolute errors (MAE) of both models are very similar. Both the 0-covariance and cross-method correlation models see the largest MAE in the private commercial medical sector (7.35 percentage points in the 0-covariance model and 7.25 percentage points in the cross-method correlation model). In both models, the private other sector has an MAE of less than 1 percentage point (0.1 percentage points in the 0-covariance model and 0.1 percentage points in the cross-method correlation model).

When considering the median prediction interval widths, we see that the cross-method correlation model has slightly smaller sized prediction interval widths as compared to the 0-covariance model for the commercial medical and public sectors (68 percentage points and 72 percentage points Table 2.5; 69 percentage points and 73 percentage points Table 2.6). The median prediction interval width of the private other sector is the same in both models at 33 percentage points.

Lastly, when considering the location of the incorrectly estimated test set observations, the cross-method correlation model and 0-covariance models both tend to over-estimate the public sector (as there is higher proportion of incorrectly estimated observations below the prediction interval) and under-estimate the private commercial medical and private other sectors (as there is higher proportion of incorrectly estimated observations above the prediction interval) (Table 2.5) (Table 2.6).

Overall, the multi-country subnational model with cross-method correlations is the most suitable model to describe this complex data. It captures the complex shape and relationships without over-fitting it or missing the shape. It incorporates information regarding the correlations between the rates of change across the contraceptive methods. The coverage, RMSE and median 95% prediction intervals widths produced by the cross-method correlation model are similar but slightly better than those of the 0-covariance model. The strength of the full model is seen in the absence of data for a particular contraceptive method, where model estimates can still be informed by the behaviour of related methods to produce realistic estimates.

2.3 Single-country national and subnational model validation

We validate the single country models indirectly using the multi-country model estimates. The idea here is that, by comparing the median estimates of the single-country model to the validated multi-country model estimates when can get a gauge of the reliability of our single-country model estimates. If the single-country model estimates align with the multi-country model estimates, then they too are validated.

2.4 Single-country national model validation

In Figure 2.3, it is clear that the single-country national model median estimates align approximately with the multi-country national median estimates. Therefore, we can conclude that the single-country national model estimates are as valid and reliable as those estimated by the multi-country national model.

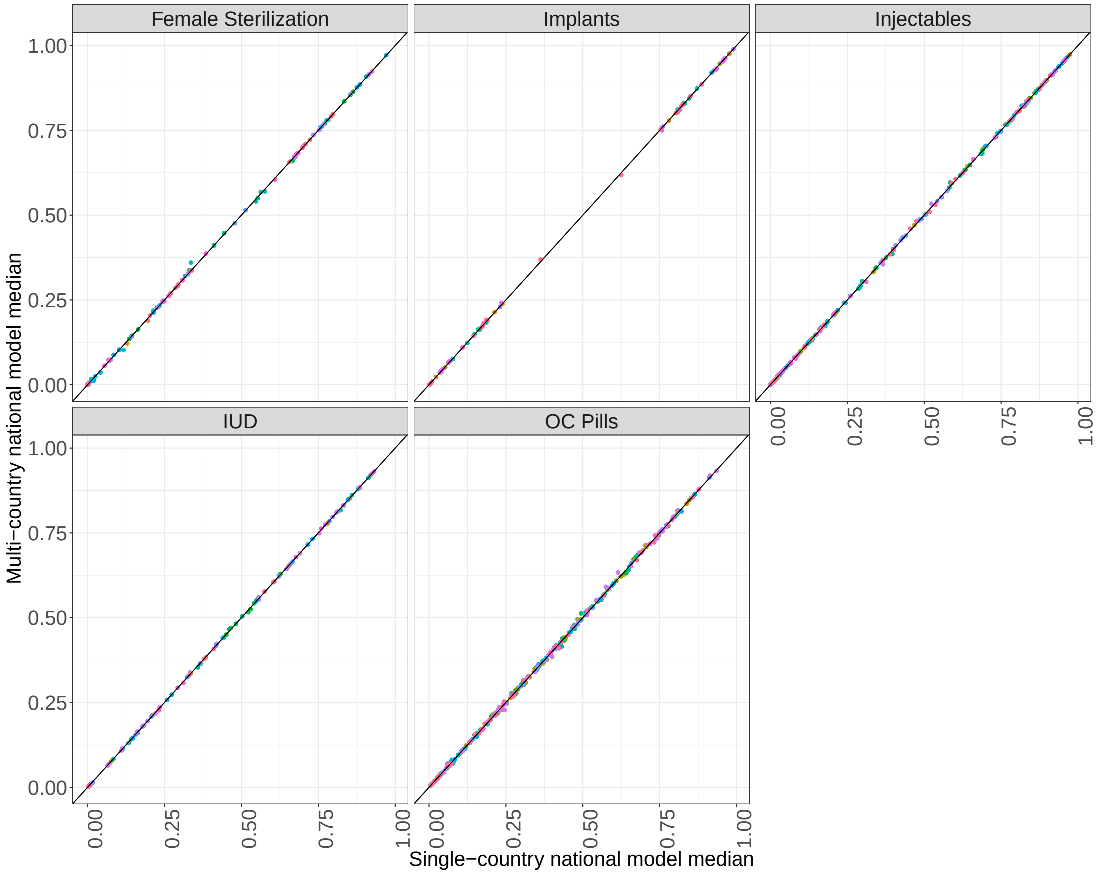

Figure 2.3: A scatterplot comparing the median estimates of each country, method, sector, and time point estimated by the single-country model (x-axis) and the multi-country model (y-axis). Each panel represents a different method and each colour represents a country. The diagonal line capture the 1:1 agreement between the two modelling approaches.

2.5 Single-country subnational model validation

In Figure 2.4, we can see that when comparing the single-country estimates to the multi-country estimates, the majority of observations fall inside the ±5% boundary. This means that the estimates and projections produced by both models have a difference of up to ±5%. There are few observations in IUDs that are outliers to this. These belong to Mozambique, where there is only one survey in 1997 taken in the City of Maputo. Therefore, we can conclude that the single-country subnational model median estimates are as valid and reliable as those estimated by the multi-country subnational model.

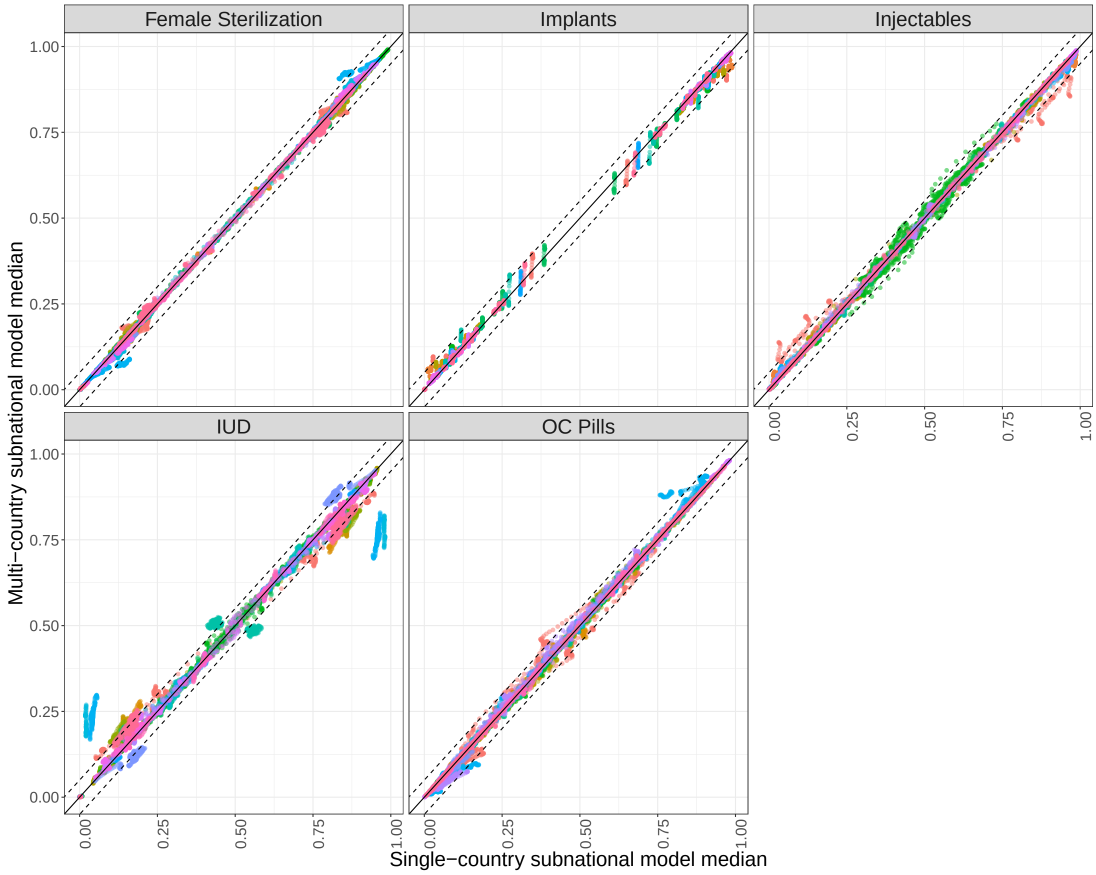

Figure 2.4: A scatterplot comparing the median estimates of each country, method, sector, and time point estimated by the single-country subnational model (x-axis) and the multi-country subnational model (y-axis). Each panel represents a different method and each colour represents a country. The outer dashed lines represent the +5% and -5% from complete agreement between the two models. The diagonal solid line capture the 1:1 agreement between the two modelling approaches.

Bibliography

[1] J. J. Azose and A. E. Raftery. Estimating large correlation matrices for international migration. *Annals of Applied Statistics*, 12(2):940–970, 2018.

[2] D. A. Binder. On the Variances of Asymptotically Normal Estimators from Complex Surveys. *International Statistical Review / Revue Internationale de Statistique*, 51(3):279, 1983.

[3] G. Casella and R. L. Berger. *Statistical Inference*. Duxbury, second edition, 2002.

[4] H. Comiskey, L. Alkema, and N. Cahill. Estimating the proportion of modern contraceptives supplied by the public and private sectors using a Bayesian hierarchical penalized spline model. *Arxiv*, 2023.

[5] H. Comiskey, L. Alkema, and N. Cahill. Package 'mcmsupply', 2023.

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

[7] E. Heger Boyle, M. King, and M. Sobek. IPUMS-Demographic and Health Surveys: Version 9 [dataset], 2022.

[8] P. Lavrakas. Taylor Series Linearization (TSL). *Encyclopedia of Survey Research Methods*, 2013.

[9] T. Lumley. Package 'survey' Title Analysis of Complex Survey Samples. Technical report, CRAN, 2021.

[10] Track20. Monitoring Progress in Family Planning Estimated Modern Use (EMU): A New Service Statistics-Based Family Planning Indicator. Technical report, Avenir Health, Glastonbury, Conneticut, 2017.

[11] U.S. Agency for International Development (USAID). The Demographic and Health Surveys (DHS) Program. `https://dhsprogram.com/Methodology/Survey-Process.cfm`, 2022.