

Learning Individual Reproductive Behavior from Aggregate Fertility Rates via Neural Posterior Estimation

Daniel Ciganda^{1,2}, Ignacio Campón^{1,2}, Iñaki Permanyer^{3,4}, and Jakob H Macke^{5,6}

¹Max Planck Institute for Demographic Research

²Statistics Institute, University of the Republic

³Center For Demographic Studies, Autonomous University of Barcelona

⁴ICREA, Catalan Institution for Research and Advanced Studies

⁵Machine Learning in Science, University of Tübingen & Tübingen AI Center

⁶Department Empirical Inference, Max Planck Institute for Intelligent Systems

Abstract

Age-specific fertility rates (ASFRs) provide the most extensive record of reproductive change, but their aggregate nature obscures the individual-level behavioral mechanisms that drive fertility trends. To bridge this micro-macro divide, we introduce a likelihood-free Bayesian framework that couples a demographically interpretable, individual-level simulation model of the reproductive process with Sequential Neural Posterior Estimation (SNPE). We show that this framework successfully recovers core behavioral parameters governing contemporary fertility, including preferences for family size, reproductive timing, and contraceptive failure, using only ASFRs. The framework’s effectiveness is validated on cohorts from four countries with diverse fertility regimes. Most compellingly, the model, estimated solely on aggregate data, successfully predicts out-of-sample distributions of individual-level outcomes, including age at first sex, desired family size, and birth intervals. Because our framework yields complete synthetic life histories, it significantly reduces the data requirements for building microsimulation models and enables behaviorally explicit demographic forecasts.

Keywords— simulation-based inference, individual-level modelling, age-specific fertility rates, reproductive behavior, microsimulation

1 Introduction

Globally declining fertility rates are driving a profound demographic transformation, presenting first-order societal challenges and underscoring the need for robust models to understand and anticipate future fertility trends. Age-specific fertility rates (ASFRs), defined as the number of births to women of a given age per person-year of exposure, remain the workhorse input for these tasks because the underlying birth counts are collected continuously in every country’s vital-statistics system. ASFRs series therefore provide unparalleled temporal and geographic coverage.

Yet ASFRs are only marginal aggregates. They reveal *when* births occur but not *why*. A long tradition of parametric models has been developed to describe ASFRs curves, smooth observed schedules, extrapolate fitted parameters, or convert projected total fertility rates into age patterns for cohort-component population forecasts (Coale and Trussell, 1974; Brass, 1974; Chandola et al., 1999; Schmertmann, 2003; Mazzuco and Scarpa, 2015). While existing approaches can excel at capturing aggregate ASFRs shapes, their parameters are typically macro-level abstractions. As such, they offer limited insight into the individual behavioural mechanisms driving fertility decisions (Hoem and Hoem, 1989). Consequently, forecasts based on these models often cannot explicitly account for how evolving individual choices regarding the timing of parenthood, contraceptive practice, and birth-spacing and stopping decisions shape future fertility trends.

This disconnect is striking in light of the explosion of micro-level fertility research based on rich individual data sources. Studies using birth histories, time-use diaries, and linked administrative records have uncovered nuanced relationships between education, labor market trajectories, union stability, and childbearing. Yet virtually none of that behavioural insight is incorporated into fertility projection models. The challenge of understanding contemporary fertility lies in bridging this micro-macro divide.

This paper proposes a step toward this goal. We introduce a demographically interpretable individual-level model that captures key features of modern reproductive behavior, including stopping, spacing, and contraceptive failure. Estimating such a richly parameterised model from coarse data is a significant challenge, as many combinations of micro-parameters can generate similar ASFRs curves. We overcome this challenge by pairing our individual-level, mechanistic model with Sequential Neural Posterior Estimation (SNPE), a powerful simulation-based inference method. This approach allows us to demonstrate that core behavioral parameters governing the overall level and timing of fertility can be recovered from ASFRs alone. We then fit the model to U.S. National Survey of Family Growth cohorts and to Demographic and Health Survey cohorts from Colombia, the Dominican Republic, and Peru, representing four settings that span markedly different fertility regimes.

The framework’s effectiveness is confirmed through rigorous validation. First, posterior predictive checks show that the observed aggregate fertility rates are highly plausible under the model, falling well within the 95% posterior predictive intervals. Most significantly, we conduct out-of-sample validation, demonstrating that the model can successfully predict distributions of individual-level characteristics, such as the empirical distributions of age at first sexual intercourse, desired family size, and birth intervals, none of which inform the estimation step.

By establishing a demographically interpretable and statistically robust link between individual re-

productive behaviors and population-level fertility outcomes, our framework offers a promising path towards a more nuanced understanding of fertility dynamics. This approach not only enables more accurate, behaviorally grounded population forecasts but also lowers the barrier to building detailed microsimulation models by significantly reducing their data requirements.

1.1 Previous Work

Efforts to model the reproductive process mechanistically have a long and distinguished tradition in demography. Pioneered by early work on fecundability (Gini, 1924) and significantly advanced by work on 'natural fertility' (Henry, 1953), this foundation led to a first wave of mathematical and early simulation models of the reproductive process (Brass, 1958; Singh, 1963; Ridley and Sheps, 1966; Potter, 1972; Bongaarts, 1977, among others). These models sought to capture the biological and behavioral drivers of family building with increasing realism.

However, this vibrant tradition of individual-level modeling did not continue to evolve within demography at the same pace or in the same manner as in other scientific fields. While empirical fertility research using micro-data flourished (Xie, 2000), the further development of these process-based simulators saw less sustained momentum. More critically, this stream of work remained largely disconnected from the concurrent revolution in statistical methodology for simulation-based inference. Techniques like Approximate Bayesian Computation (ABC) or other simulation-based inference (SBI) approaches that became influential in fields like ecology and genetics (Beaumont, 2010; Hartig et al., 2011), were not systematically integrated in demography. This methodological lag created a persistent gap in formally linking mechanistic understanding to empirical observations in fertility research.

Our own previous work showed that such a bridge is feasible in natural-fertility contexts (historical or religious communities without deliberate birth control) where a handful of biological parameters govern the reproductive process and can be inferred from ASFRs alone (Ciganda and Todd, 2024). The present paper tackles the far more ambitious and policy-relevant task of modeling contemporary fertility, where stopping, spacing, and imperfect contraception introduce complex decision-making and heterogeneity in reproductive dynamics.

1.2 Organization of the Article

The remainder of this paper is structured as follows: Section 2 describes our individual-level computational model of the reproductive process. Section 3 details the data sources and the construction of the main demographic quantities used in our analysis. Section 4 outlines the SNPE methodology. The Results section presents our findings on parameter identifiability via cross-validation, model fit through posterior predictive checks, and the out-of-sample validation against micro-level data. Finally, we discuss the implications of our findings, limitations, and avenues for future research in the Discussion.

2 Model

The transition from a natural to a modern fertility regime in human populations is marked by the emergence of *stopping* and *spacing* behaviors (deciding whether and when to have children), representing a shift from a reproductive process largely determined by physiological factors to one shaped by individual preferences. As a result, fertility patterns become more diverse, with family formation increasingly tailored to individual or couple-level choices rather than biological constraints alone.

Building on these insights, our model represents the dynamics of reproductive behavior in contemporary populations by placing fertility intentions and desires at its core. Specifically, it conceptualizes childbearing decisions as the interplay between physiological capacity, access to contraception, and individuals' subjective preferences, such as desired family size and timing.

We employ an individual-level, discrete-time simulation model. A cohort of N women is simulated from the age of potential sexual activity until the end of their reproductive capacity, with their state updated in monthly time steps.

The model begins by assigning each woman key characteristics that will shape her reproductive life. These include two critical ages: the *age at sexual initiation*, when she first becomes exposed to the risk of childbearing, and the *age at intentional reproduction*, representing the point at which she actively starts trying to conceive. Because these two ages often differ, there is typically a period during which a woman is sexually active but not yet intending to have children, an interval associated with the risk of unplanned pregnancies.

Additionally, each woman is assigned a *desired family size*, reflecting the number of children she ultimately hopes to have, and a *desired number of months between births*, which influences her birth spacing behavior.

In each monthly step, a woman's probability of conception is evaluated. This probability depends on her current age, her current parity relative to her desired family size, whether she has reached her age at intentional reproduction, and, if she has had prior births, whether her desired birth spacing interval has elapsed.

Specifically, when a woman is past her age at intentional reproduction, has not yet reached her desired family size, and is past her desired birth spacing interval (if applicable), we assume she is actively trying to conceive and uses no contraception. Conversely, if she has not yet reached her age at intentional reproduction, or has already achieved her desired family size, or is within her desired spacing interval, she is assumed to be using contraception. This results in a reduced, but non-zero, probability of conception, accounting for imperfect contraceptive effectiveness and the possibility of unplanned pregnancies throughout the reproductive lifespan.

Whenever a conception occurs, the woman is considered non-susceptible for nine months of pregnancy, followed by an additional three-month period of postpartum amenorrhea, during which her probability of another conception is zero. After giving birth, a further delay is introduced to represent the woman's desired number of months between births, mirroring real-life behaviors around birth spacing.

2.1 Implementation

The dynamics described above are implemented by assigning the following characteristics to each woman i and defining the processes that govern her reproductive trajectory:

1. **Age at Sexual Initiation (X_i):** Each woman's age at sexual initiation, X_i , is drawn from a lognormal distribution. This distribution is parameterized by a mean age μ_s and a standard deviation σ_s (both in years), which are converted to the distribution's underlying parameters (the mean of logarithms, $\mu_{\ln,s}$, and the standard deviation of logarithms, $\sigma_{\ln,s}$) using the standard formulas: $\mu_{\ln,s} = \ln(\mu_s^2 / \sqrt{\mu_s^2 + \sigma_s^2})$ and $\sigma_{\ln,s} = \sqrt{\ln(1 + \sigma_s^2 / \mu_s^2)}$. The value for each woman is drawn in years and then converted to months.
2. **Age at Intentional Reproduction (R_i):** R_i (in months) is also drawn from a lognormal distribution. The mean of this distribution is set to $\mu_r = \mu_s + \delta_r$, where δ_r (in years) is an estimated model parameter. The parameters μ_r and σ_r (in years) are used to derive the underlying lognormal parameters $\mu_{\ln,r}$ and $\sigma_{\ln,r}$ using the same conversion formulas.
3. **Desired Family Size (D_i):** D_i is drawn from a Weibull distribution, parameterized by its mean μ_d and standard deviation σ_d , which are used to approximate the Weibull shape parameter $\alpha_d = (\sigma_d / \mu_d)^{-1.086}$ and scale parameter $\lambda_d = \mu_d / \Gamma(1 + 1/\alpha_d)$. Samples are rounded to the nearest integer.
4. **Desired Birth Spacing (B_i):** B_i (in months) is drawn from a lognormal distribution with a mean desired spacing μ_b and a standard deviation σ_b (both in months). These parameters are used to derive the underlying mean and standard deviation for the lognormal distribution.

The age-specific monthly probability of conception (fecundability), $\phi(x)$ at age x (in years), is a crucial component of the model. It is well established that fecundability typically increases from puberty, reaches a peak during the early adult years, and progressively declines thereafter towards menopause (Bendel and Hua, 1978; Larsen and Yan, 2000; Weinstein et al., 1990). To capture this characteristic age pattern, we model $\phi(x)$ over a reproductive window from $x_{min} = 10$ to $x_{max} = 50$ years. Age x is first rescaled to $x_s = (x - x_{min}) / (x_{max} - x_{min})$, so $x_s \in [0, 1]$. Instead of using a standard polynomial basis, we employ Bernstein basis polynomials, which provide flexibility and ensure the curve is well-behaved at the boundaries of the reproductive window. Specifically, $\phi(x)$ is modeled as a linear combination of two Bernstein basis polynomials of degree 3:

$$\phi(x) = \beta_1 [3x_s(1 - x_s)^2] + \beta_2 [3x_s^2(1 - x_s)]$$

The terms $3x_s(1 - x_s)^2$ and $3x_s^2(1 - x_s)$ correspond to the two intermediate basis functions ($B_{1,3}(x_s)$ and $B_{2,3}(x_s)$) of a Bernstein basis of degree 3. The exclusion of the first and last basis functions from the model inherently forces the curve $\phi(x)$ to be zero at the start ($x_s = 0$) and end ($x_s = 1$) of the reproductive window. The parameters β_1 and β_2 are estimated and control the level and shape of the fecundability curve, allowing for a variety of empirically plausible age-fecundability profiles.

The actual probability of conception in a given month depends not only on this baseline age-specific fecundability but also on a woman's reproductive intentions and contraceptive use. As outlined earlier,

if a woman i with current parity k_i is *actively trying to conceive* (i.e., her current age is $\geq R_i$, her parity $k_i < D_i$, and sufficient time B_i has elapsed since her last birth), her probability of conception is $\phi(x)$.

Conversely, if a woman i with current parity k_i is sexually active but *not* actively trying to conceive, she is assumed to use contraception. Her monthly probability of conception is then $\kappa_{eff}\phi(x)$. The parameter $\kappa \in (0, 1)$ is an estimated baseline probability of contraceptive failure. To model an *increased contraceptive vigilance once desired parity (D_i) is met*, the effective failure rate κ_{eff} is defined as κ if her current parity $k_i < D_i$, and is reduced to κ^2 if $k_i \geq D_i$. This formulation reflects intensified efforts to avoid further births, substantially reducing the risk of an unplanned pregnancy (assuming $\kappa < 1$) after a woman reaches her desired family size.

Table 1 summarizes the estimated model parameters.

Table 1: Model parameters and their descriptions.

Parameter	Description
μ_s	Mean age at sexual initiation (in years), for a lognormal distribution.
σ_s	Standard deviation of the age at sexual initiation.
δ_r	Mean gap (in years) between sexual initiation and intentional reproduction.
σ_r	Standard deviation of the age at start of intentional reproduction.
μ_d	Mean desired family size, for a Weibull distribution.
σ_d	Standard deviation of desired family size.
μ_b	Mean desired birth spacing (in months), for a lognormal distribution.
σ_b	Standard deviation of the desired birth spacing.
κ	Probability of contraceptive failure per month.
β_1	Defines the peak of the age-specific fecundability curve.
β_2	Defines the decline of the age-specific fecundability curve.

3 Data

Our analysis draws on two large, nationally representative sources. For countries outside the United States, we use the Demographic and Health Surveys (DHS), a leading initiative of the United States Agency for International Development (USAID) that has conducted over 300 surveys in more than 90 low-and middle-income countries. Each DHS round collects complete retrospective birth histories, allowing the reconstruction of every woman’s fertility trajectory. Data for the United States are drawn from several early cycles of the National Survey of Family Growth (NSFG). These surveys, conducted by the National Center for Health Statistics (NCHS), are designed to provide detailed national data on family life, marriage, divorce, contraception, and importantly for this study, complete pregnancy and birth histories. We utilize these data as compiled and harmonized by the Integrated Fertility Survey Series (IFSS), applying IFSS sampling weights to generate population-level estimates. For both DHS and NSFG data, to achieve adequate sample sizes while ensuring cohorts of women were exposed

to broadly similar social conditions, we pool available survey rounds and organize respondents into 10-year birth cohorts. The final cohorts analyzed correspond to women born between 1938–1948 for the United States, and between 1966–1975 for Colombia, the Dominican Republic, and Peru.

A key consideration in selecting empirical settings was to align them with the behavioral logic embedded in our model. Thus, we narrowed the DHS universe to countries where adolescent childbearing is primarily unintended. This was operationalized by calculating, for each potential cohort, the proportion of births to women under 18 that were declared as unplanned, retaining only those populations where this share exceeded one-half. This selection criterion identifies populations where, for a significant fraction of teenagers, sexual activity and intentional reproduction are decoupled—a precondition for our model’s assumption that early fertility is largely viewed as undesirable. It also serves to exclude contexts where prevailing norms, such as marriage-centered childbearing from a young age, would imply a data-generating process markedly different from the one our model specifies.

Because DHS sample sizes and the number of survey rounds differ widely across countries, we implemented a further filter based on cohort size. We selected DHS countries where the chosen cohort included a minimum of 4,000 women, a threshold established to guarantee reasonably precise estimates of both age-specific total and unplanned fertility rates across the entire reproductive span. This resulted in the inclusion of three DHS countries: Colombia, Peru, and the Dominican Republic. Together with the United States, these four nations constitute our study sample, offering a valuable cross-section of different fertility patterns and socio-economic contexts that allows for a robust examination of our model’s performance across diverse demographic and social settings.

We compute the conventional age-specific fertility rate by dividing the number of births at each age by the number of person-years of exposure contributed at that age. To test how information on birth intentionality improves our estimates, we also construct an age-specific unplanned fertility rate (ASUFRs). This rate is defined as the number of *unplanned* births at a given age divided by person-years of exposure. The classification of a birth as unplanned is based on a harmonized strategy across the DHS and NSFG that combines direct survey questions about birth timing with a parity check (i.e., whether the birth exceeded the mother’s ideal family size). This hybrid approach is designed to mitigate recall bias and ex-post rationalization, which are known to affect estimates based solely on intention questions (Casterline and El-Zeini, 2007).

4 Simulation-Based Inference

Inferring parameters for complex process-based models is often challenging due to intractable likelihood functions. Neural Posterior Estimation (NPE) offers a powerful simulation-based inference approach that circumvents these difficulties by directly learning an approximation to the posterior distribution $p(\theta | x)$ from simulated data pairs (θ, x) . In practice, NPE proceeds by drawing parameters θ from a prior distribution $p(\theta)$, using the simulator to generate the corresponding data $x \sim p(x | \theta)$, and then training a flexible neural density estimator, $q_\phi(\theta | x)$ (with network parameters ϕ), on these simulated pairs to approximate the true posterior.

While basic NPE is effective, our study employs a powerful sequential variant known as Automatic Posterior Transformation (APT), sometimes referred to as SNPE-C (Greenberg et al., 2019). Like

other sequential neural posterior estimation methods, APT enhances inferential efficiency by drawing parameter samples from an iteratively updated proposal distribution, $\tilde{p}(\theta)$, rather than solely from the fixed prior. The key challenge in this approach is to correct for the use of a proposal distribution, as naively training on the resulting samples would lead to an incorrect posterior estimate. APT solves this by analytically transforming the posterior estimate itself. For each simulated data point drawn from a proposal $\tilde{p}_r(\theta)$, it re-weights the target density using the known ratio $\tilde{p}_r(\theta)/p(\theta)$. The network is then trained by minimizing the loss with respect to this transformed "proposal posterior" (see line 10 of Algorithm 1). A key advantage of this formulation is that because the loss is valid for any proposal, the network can be trained on the cumulative set of all simulations generated across all previous rounds. This procedure allows the network to learn the true posterior without relying on the potentially unstable importance weights used by other methods.

A key strength of the APT framework is its flexibility, as it supports a wide range of density estimators, including mixture-density networks and various types of normalizing flows. In this study, we employ a Neural Spline Flow (NSF), an expressive and powerful density estimator based on monotonic rational-quadratic splines (Durkan et al., 2019).

This adaptive refinement progressively concentrates simulation effort on regions of high posterior probability, leading to a more accurate posterior approximation with fewer overall simulations. SNPE and related neural posterior estimation techniques have found successful application in diverse scientific domains, including neuroscience (Gonçalves et al., 2020; Groschner et al., 2022; Deistler et al., 2022), astrophysics (Dax et al., 2025), cognitive science (von Krause et al., 2022), and exoplanet searches (Vasist et al., 2023), among others.

Algorithm 1 summarizes the complete APT procedure employed in this study, which was implemented using the `sbi` Python package (Boelts et al., 2024).

Algorithm 1: Automatic Posterior Transformation (APT)

Input: Simulator with implicit likelihood $p(x|\theta)$, observed data x_o , prior $p(\theta)$, conditional density family $q_{F(x,\phi)}(\theta)$, number of rounds R , simulations per round N .

Output: The final posterior approximation $q_{F(x_o,\phi)}(\theta)$.

Set initial proposal distribution $\tilde{p}_1(\theta) \leftarrow p(\theta)$;

for $r = 1$ **to** R **do**

for $j = 1$ **to** N **do**

 Sample $\theta_{r,j} \sim \tilde{p}_r(\theta)$;

 Simulate $x_{r,j} \sim p(x|\theta_{r,j})$;

end

 Update network weights ϕ by minimizing the loss over all collected data $\{(\theta_{i,j}, x_{i,j})\}$ from rounds $i = 1 \dots r$:

$\phi \leftarrow \arg \min_{\phi} \sum_{i=1}^r \sum_{j=1}^N -\log \tilde{q}_{x_{i,j},\phi}(\theta_{i,j})$;

 where $\tilde{q}_{x,\phi}(\theta) := q_{F(x,\phi)}(\theta) \frac{\tilde{p}(\theta)}{p(\theta)} \frac{1}{Z(x,\phi)}$ with $Z(x,\phi) = \int_{\theta} q_{F(x,\phi)}(\theta) \frac{\tilde{p}(\theta)}{p(\theta)} d\theta$;

 Set the next proposal to the current posterior estimate for the observed data x_o :

$\tilde{p}_{r+1}(\theta) \leftarrow q_{F(x_o,\phi)}(\theta)$;

end

The code and data used in this study, along with detailed instructions for replication, can be found on GitHub at https://anonymous.4open.science/r/asfr_snpe-BEFF (anonymized repository for peer review).

4.1 Inference Scenarios

Simulation-Based Inference typically relies on summarizing potentially high-dimensional data into lower-dimensional, sufficiently informative statistics. This can be done either through explicit, pre-defined summaries or via learned embeddings. The ability to work with a low-dimensional yet descriptive representation of the data is often essential for making parameter inference tractable, especially for complex simulators.

When working with aggregate data, as in our study, the dimensionality reduction is an inherent feature of the observation process. For instance, vital statistics or administrative registers implicitly aggregate thousands of individual fertility trajectories into relatively low-dimensional summaries like age-specific fertility rates. We thus inherit a pre-summarized dataset, limiting our ability to engineer or learn optimal summary statistics directly from the underlying micro-level process.

A critical question then arises: do these aggregate data provide sufficient information to tightly constrain the posterior distributions of the parameters in our individual-level model? The core challenge is that different combinations of micro-parameters (e.g., those governing desired family size, timing of intentional reproduction, or contraceptive failure) could potentially yield very similar aggregate fertility schedules, resulting in high posterior uncertainty.

While ASFRs represent the baseline data available in most settings, demographic analyses can sometimes draw upon richer sources of information. We therefore test how the parameter estimation process can be strengthened by incorporating such additional information. This can be done in two primary ways: either by augmenting the data with more detailed summary statistics or by encoding external knowledge into the model through more informative priors. Given that our model explicitly simulates the planned and unplanned nature of births, information on birth intentionality is a particularly relevant source of additional data. However, since this type of data is costly to collect and often unavailable, it is valuable to assess whether similar gains in estimation accuracy can be achieved by instead using informative priors on key behavioral parameters.

We therefore evaluate parameter recovery under three increasingly informative inference scenarios:

Scenario 1: ASFRs with Weak Priors: This baseline scenario uses only simulated age-specific fertility rates as summary statistics. All model parameters are assigned weakly informative priors.

Scenario 2: ASFRs with Informative Priors. This scenario tests the impact of incorporating plausible external information. The summary statistics are still limited to ASFRs, but we use narrower, more informative priors for the mean desired family size (μ_d), the mean gap to intentional reproduction (δ_r), and the mean desired birth spacing (μ_b). All other parameters retain their weakly informative priors.

Scenario 3: ASFRs and ASUFRs with Weak Priors. This scenario tests the impact of adding more detailed data. We revert to the weakly informative priors from Scenario 1, but we augment the summary statistics to include both simulated ASFRs and age-specific unplanned fertility rates.

The informative priors used in Scenario 2 simulate a context where a researcher might leverage external information to improve estimates. For the mean desired family size (μ_d), we construct the prior directly from the empirical survey distribution for the U.S. case, mimicking a situation where such summary data is readily available. For the remaining two parameters, we simulate the process of knowledge transfer from a data-rich to a data-poor setting by using the posteriors from our most data-intensive setup (Scenario 3) as the informative priors for Scenario 2. This mimics a situation where results from a detailed prior study on a similar cohort/country are used to inform a new analysis where only ASFRs are available.

4.2 Prior Specification

For our baseline scenario (Scenario 1), we selected weakly informative priors for all eleven model parameters to allow the data to drive the inference. The specific parameterization of these priors was chosen to cover a broad yet plausible range of values based on existing demographic literature and the inherent constraints of each parameter.

For strictly positive quantities like means and standard deviations (e.g., μ_s, σ_s), we used broad Gamma distributions. For example, the prior for the mean age at sexual initiation (μ_s) contains 95% of its probability mass concentrated between 14.4 and 28.9 years. This range comfortably contains existing estimates for the mean age at onset of sexual activity across diverse Western countries and cohorts, which typically falls between 16 and 19 years (Wellings et al., 2006; Kerry L. et al., 2017). Likewise, the prior for mean desired family size (μ_d) is centered at 4.5 children, with a 95% probability interval of [1.9, 8.1], again comfortably encompassing the available estimates for the cohorts studied (Westoff and Ryder, 2015; Bongaarts and Lightbourne, 1992; Sobotka and Beaujouan, 2014). For parameters where reliable external estimates are lacking, such as the gap to intentional reproduction (δ_r) and mean desired birth spacing (μ_b), we employed Uniform priors to define very wide yet plausible ranges (0 to 8 years for δ_r and 10 to 100 months for μ_b). The contraceptive failure probability (κ) was assigned a Beta distribution. Its prior reflects a wide range of plausible contraceptive efficacy levels for our study cohorts, centered around a 20% failure rate but with significant probability assigned to both lower and higher values. Finally, the priors for the Bernstein coefficients (β_1 and β_2) were selected to ensure the resulting age-fecundability curves would generously contain available estimates of conception probabilities at early, peak, and later reproductive ages (Schwartz and Mayaux, 1982; Dunson et al., 2002; Wesselink et al., 2017; ACO, 2014).

5 Results

This section presents the results from the validation experiments designed to test our framework under the three information scenarios defined above. We first assess the model’s ability to recover known parameters from simulated data using a cross-validation procedure (Section 5.1). We then apply the framework to empirical data from our four study countries, evaluating both the in-sample fit through posterior predictive checks and, most critically, the out-of-sample predictive performance against micro-level data not used for estimation.

5.1 Cross-Validation

Our validation procedure involves repeatedly generating datasets with known parameters and then attempting to recover these parameters using our inference procedure. Specifically, we repeat the following steps for I iterations:

1. **Generate Ground-Truth Data:** In each iteration i (for $i = 1, \dots, I$), we draw a distinct "true" parameter vector, denoted $\theta^{*(i)}$, from the prior distribution $p(\theta)$ defined for the given scenario. Using our individual-level simulation model, we then generate a corresponding vector of aggregate summary statistics, denoted $x_0^{(i)}$, which serves as the pseudo-observed data for this iteration. The pair $(\theta^{*(i)}, x_0^{(i)})$ constitutes a ground-truth dataset where the data-generating parameters are known.
2. **Perform Full Posterior Inference:** For each pseudo-observed dataset $x_0^{(i)}$, we independently run our complete Sequential Neural Posterior Estimation procedure. This involves training a neural density estimator over R rounds with N_{sim} new simulations drawn from the evolving proposal distribution in each round, to obtain an approximation of the posterior distribution $q_\phi(\theta | x_0^{(i)})$.
3. **Evaluate Parameter Recovery:** From each resulting posterior approximation $q_\phi(\theta | x_0^{(i)})$, we use the posterior mean as the point estimate for the parameters, denoted $\hat{\theta}^{(i)}$. We then assess the accuracy of recovery by comparing this estimate $\hat{\theta}^{(i)}$ with the known true generating parameters $\theta^{*(i)}$.

By performing these I independent inference runs, we obtain a distribution of estimation errors (e.g., $\hat{\theta}^{(i)} - \theta^{*(i)}$) for each model parameter under each of the three informational scenarios. We then summarize these error distributions using the normalized root-mean-squared error (RMSE). Lower RMSE values indicate better parameter recovery and allow for a direct comparison of how effectively each informational setup constrains the model parameters.

Figure 1 presents the results of our cross-validation analysis, performed over 25 folds for the seven core model parameters (Cross-validation results for the four heterogeneity parameters are presented in Figure 5 in the Appendix). The first three columns display scatterplots comparing the known "true" parameter values (horizontal axis) against the posterior mean estimates recovered by our model (vertical axis) for each of the three scenarios. Perfect parameter recovery would place all points directly on the dashed 45° reference line. The final column summarizes the performance for each parameter across scenarios by plotting the normalized root-mean-squared error (RMSE).

The results demonstrate that parameter recovery is already reasonably robust in the baseline scenario. In the scatterplots for this scenario, the estimated parameter values consistently cluster around the true values on the 45-degree line, with some parameters exhibiting tighter clustering than others. This indicates that the model can successfully disambiguate the majority of its behavioral components from the aggregate fertility schedule alone. As expected, parameter recovery improves systematically as more information is incorporated, with the RMSE declining for nearly every parameter in Scenarios 2 and 3.

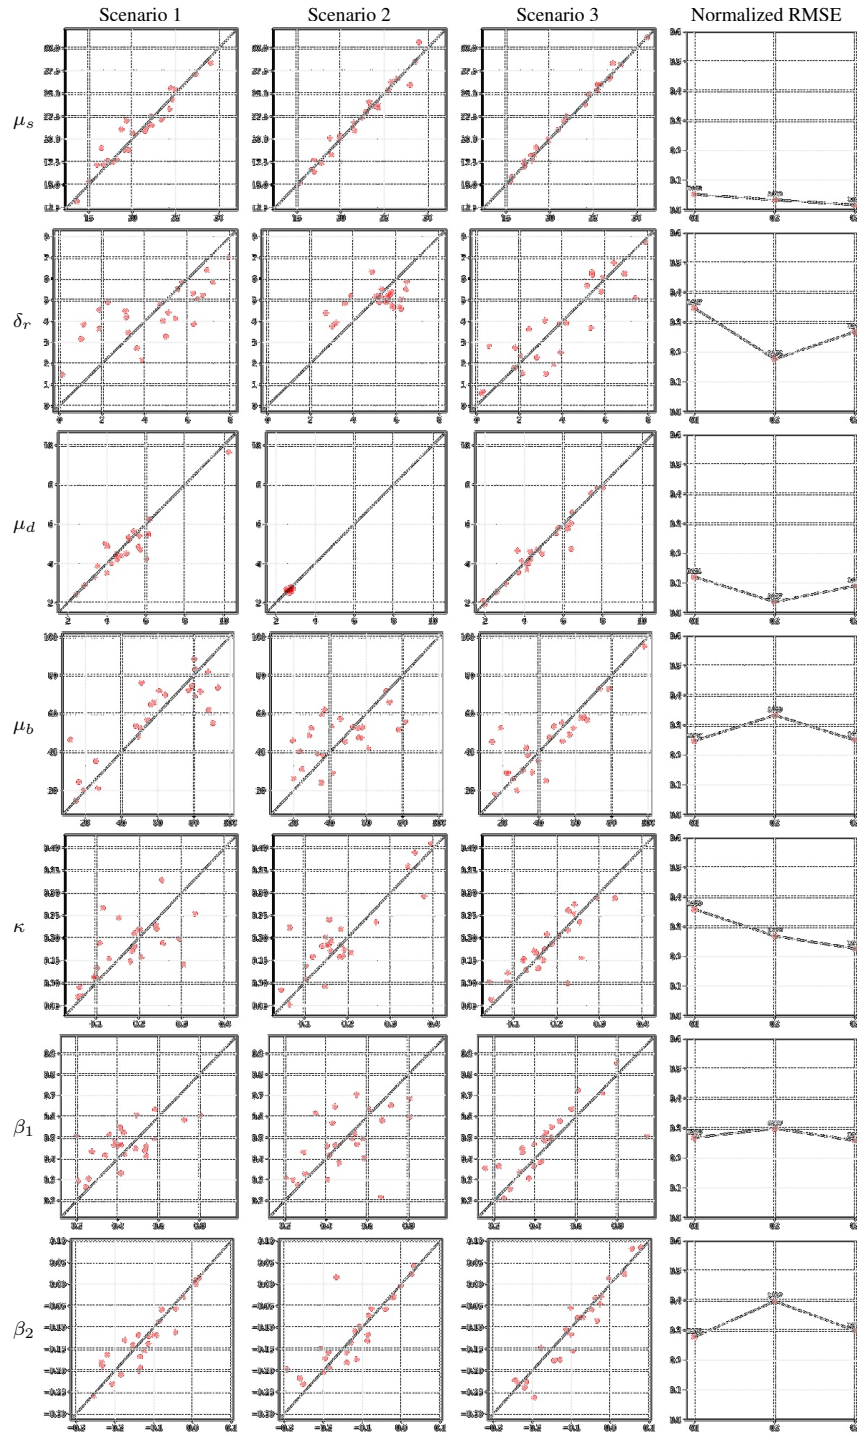


Figure 1: Cross-validation of seven core model parameters over 25 folds. Columns 1–3 display scatterplots of true parameter values (horizontal axis) versus posterior mean estimates (vertical axis) for Scenarios 1, 2, and 3, with the dashed 45° reference line indicating perfect recovery. Column 4 reports the scenario-specific root-mean-squared error. Cross-validation for the four heterogeneity parameters is shown in the Appendix.

5.2 Model Validation

The cross-validation experiments (Section 5.1) confirmed the capacity of our SBI framework to recover key micro-level parameters from aggregate ASFRs. To further scrutinize the model’s performance, we conduct two additional forms of validation. First, posterior predictive checks, to assess the model’s ability to replicate the aggregate data used for inference. The second validation stage examines the model’s capacity to predict out-of-sample individual-level characteristics of the observed cohorts, such as distributions of desired family size, age at first sexual intercourse, and birth interval durations, which were not used during the parameter estimation process.

We present the main validation results using the posteriors obtained under Scenario 1 (ASFRs with weak priors). We focus on this baseline scenario as it provides the most stringent test of our model’s capacity to infer individual-level mechanisms solely from aggregate fertility rates. Results for the scenarios incorporating additional information (Scenarios 2 and 3) are presented in the Appendix.

5.2.1 Posterior Predictive Checks

To perform posterior predictive checks, we draw 5,000 samples from the joint posterior distribution for each of the four countries (United States, Colombia, Dominican Republic, and Peru). For each parameter sample, we simulate a full cohort of individual reproductive histories using our model and then recompute the ASFRs.

Figure 2 presents observed ASFRs against their corresponding posterior predictive distributions. Observed ASFRs are represented by purple open circles. Red open circles represent the mean of the 5,000 simulated ASFRs schedules, while the shaded red areas depict the 95% posterior predictive intervals.

Across all four diverse demographic contexts, the model demonstrates a strong capacity to replicate the observed data. The posterior predictive means closely track the observed ASFRs, capturing key features such as the location and height of peak fertility, and the age pattern of decline. Notably, across all countries, almost all observed points fall well within the 95% posterior predictive interval, indicating that the model, conditioned on the inferred parameters, generates aggregate outcomes highly consistent with the observed data.

These results confirm that our framework, even with minimal data inputs, can consistently and accurately reproduce the population-level fertility schedules it was intended to explain.

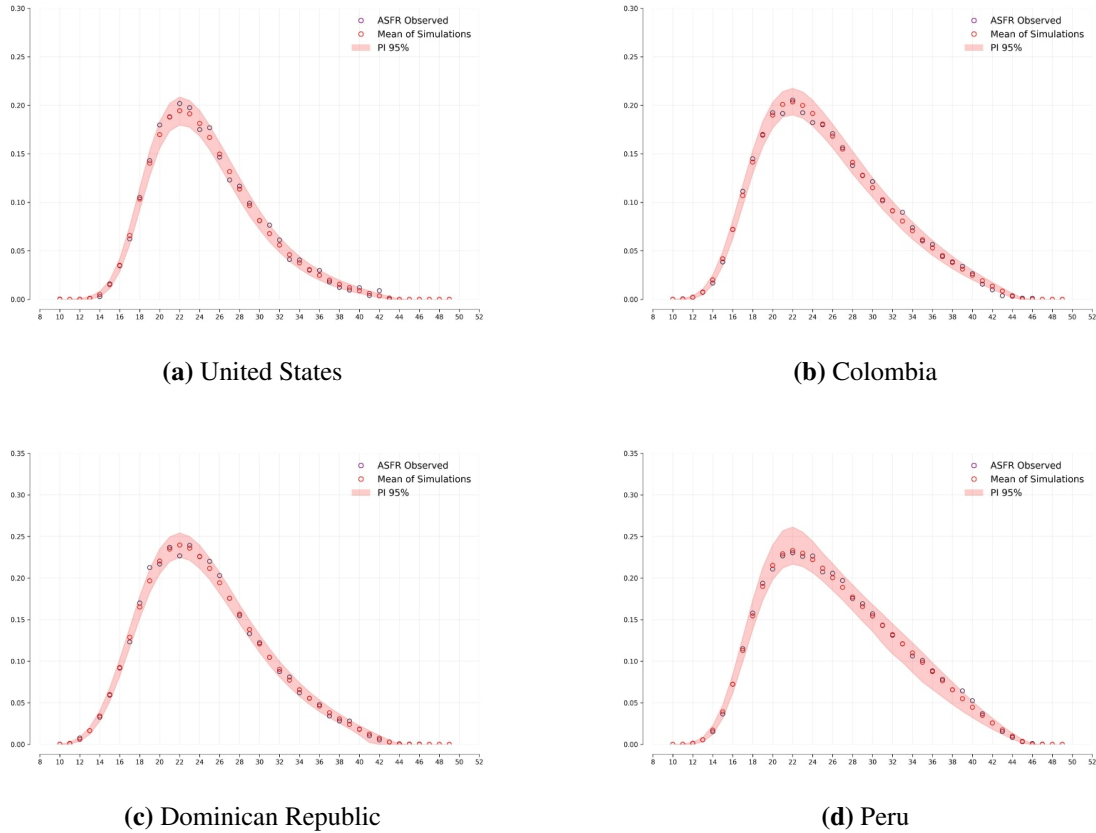


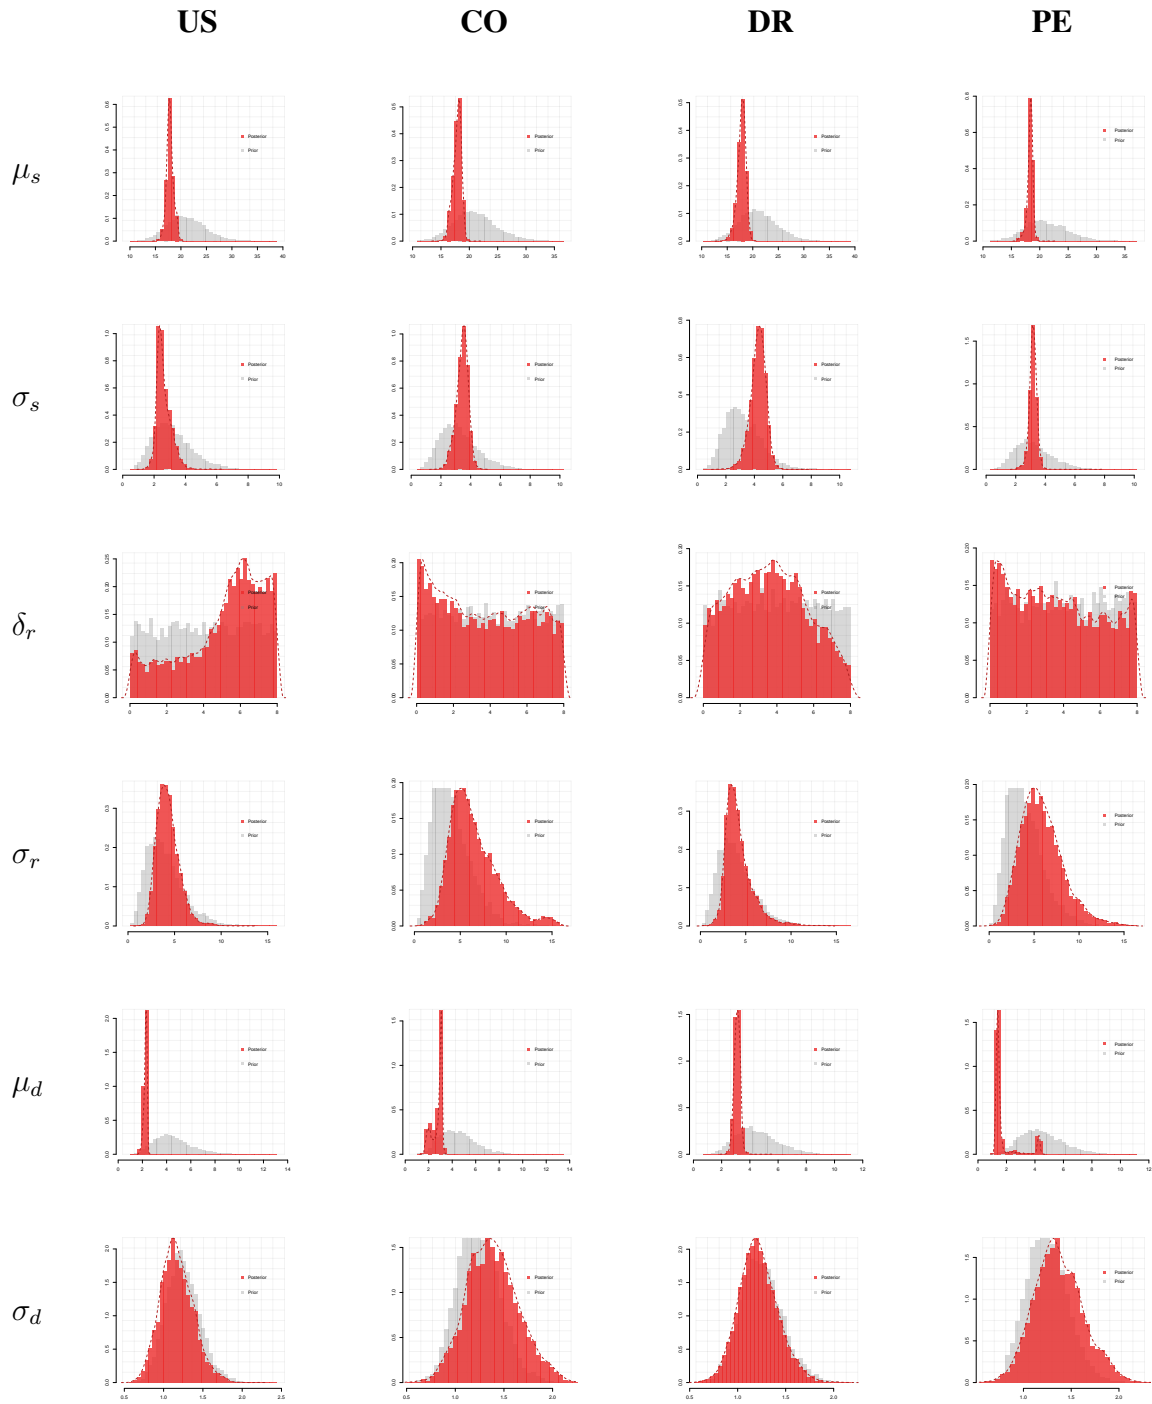
Figure 2: Posterior Predictive Checks for Age-Specific Fertility Rates. Comparison of observed rates (purple open circles) with posterior predictions for four countries: (a) United States, (b) Colombia, (c) Dominican Republic, and (d) Peru. Red open circles indicate the mean of 5,000 posterior predictive simulations. Red shaded areas represent the 95% posterior predictive intervals. Parameters were estimated under Scenario 1 (ASFRs with weakly informative priors).

5.3 Estimated Marginal Posteriors

Figure 3 displays the marginal posterior distributions for all eleven model parameters, estimated under Scenario 1 for each of the four countries. The results clearly show that the aggregate ASFRs data provide substantial information for constraining most parameters. In nearly all cases, the posterior distributions (in red) are considerably sharper and more concentrated than their corresponding wide priors (in grey), indicating a significant reduction in uncertainty after conditioning on the observed data.

The degree to which parameters are constrained varies logically across the different components of the model. Parameters related to the timing of sexual initiation (μ_s, σ_s), the age-pattern of fecundability (β_1, β_2), desired family size (μ_d, σ_d), and contraceptive failure (κ) are all well-constrained, exhibiting sharp, unimodal posterior distributions. The posteriors also reveal plausible cross-national differences. For example, the mean desired family size (μ_d) for the U.S. cohort is lower than in the other three

countries, as expected. Similarly, the contraceptive failure probability (κ) is estimated to be lowest in the U.S., a finding that aligns with the socio-demographic context of these cohorts. In contrast, parameters governing the timing between life-course events, like the gap to intentional reproduction (δ_r) and the spacing between births (μ_b, σ_b), show higher posterior uncertainty. Their posteriors are wider and less distinguished from their priors, suggesting that while ASFRs alone are sufficient to identify the core components of the model, pinpointing the precise timing of *intended* fertility is more challenging without additional data on birth intentionality. However, it is important to note that this remaining uncertainty in specific parameters does not significantly degrade the model's overall predictive performance, as evidenced by the successful posterior predictive checks (Figure 2). Finally, the estimation for these timing-related parameters improves substantially when direct information on birth intentionality is included, as shown in the results for Scenario 3 (see Figure 7 in the Appendix).



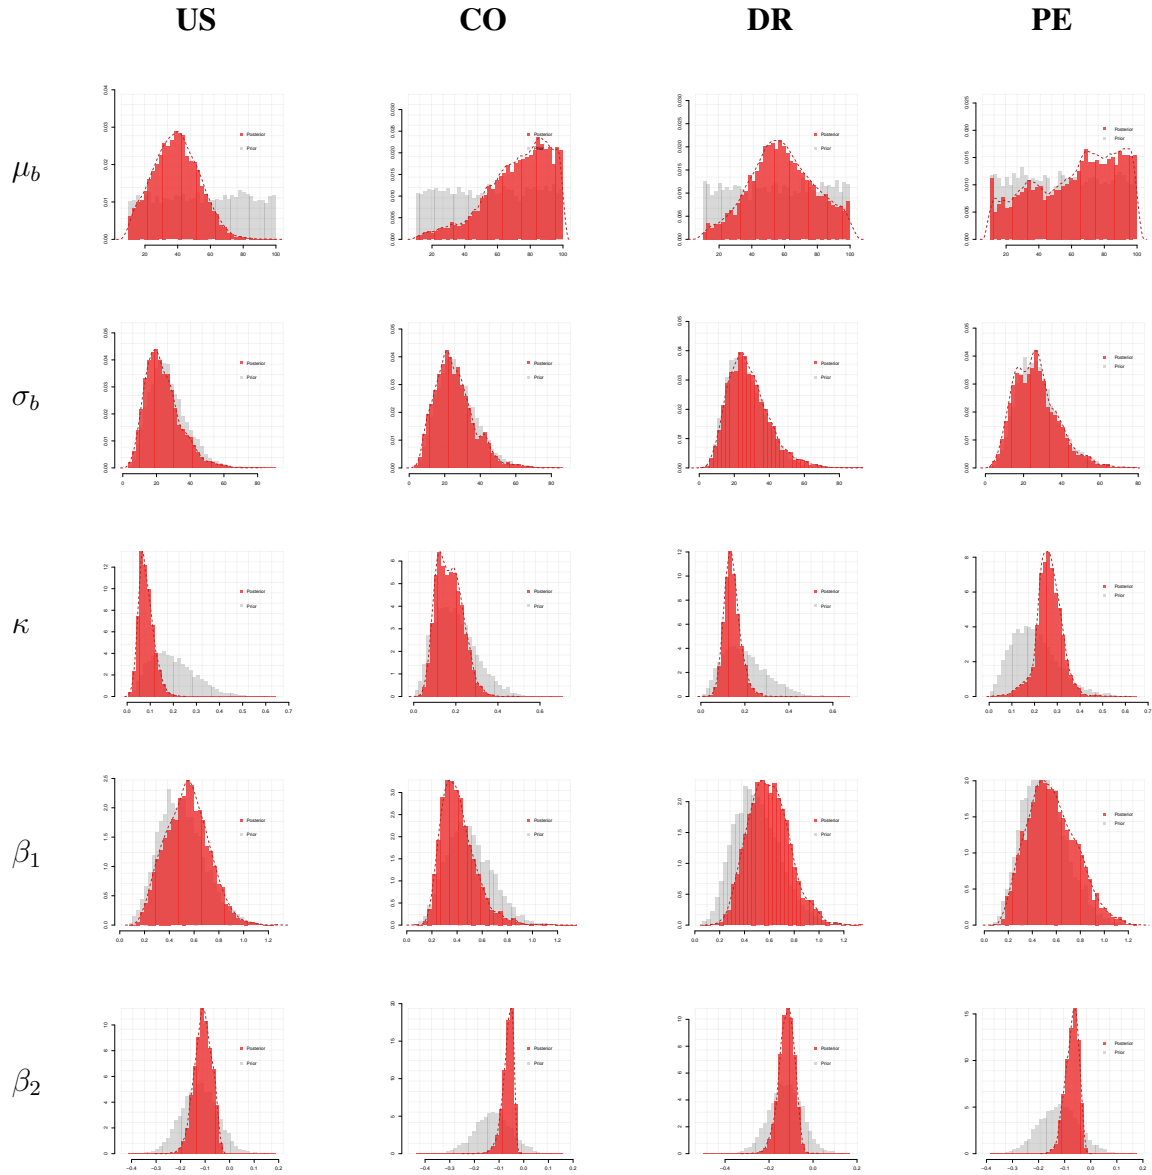


Figure 3: Marginal posterior distributions of the eleven key model parameters for the United States, Colombia, Dominican Republic, and Peru, inferred from empirical data under Scenario 1.

5.3.1 Out-of-Sample Validation

To assess the model’s ability to generate realistic individual-level reproductive behaviors not directly used in parameter estimation, we use parameter values representative of the posterior distribution (specifically, the posterior mean obtained under Scenario 1, as described in Section 5.2) to drive new simulations. These simulations yield full synthetic reproductive histories. We then compare the distributions of key micro-level outcomes from these synthetic life courses, namely, age at first sexual activity, desired family size, and inter-birth intervals, with their empirical counterparts derived from

the survey data (Figure 4). This side-by-side comparison provides a direct appraisal of the model's realism in capturing individual-level reproductive patterns.

Overall fit. Across the four countries and three micro-level indicators, the distributions generated from the model's synthetic life courses generally provide a strong approximation of the empirical distributions observed in the survey data. This indicates that the parameters inferred solely from aggregate ASFRs enable the reconstruction of distinct micro-behavioral dimensions without any specific tuning towards these individual-level targets. Quantitatively, with the notable exception of desired family size in Peru, the Jensen-Shannon divergence (base 2) between simulated and observed distributions never exceeds 0.068 bits.

Age at first sex. The simulated distributions reproduce the overall shape and cross-country ordering of sexual debut: the mean age is highest in the United States and Peru, slightly lower in Colombia, and lowest in the Dominican Republic. The model naturally smooths over sharp heaping in the survey data (e.g., at age 18), and for the United States, the predicted curve is displaced slightly to the left of the observed one. Even with these minor mismatches, the alignment is striking, given the relative simplicity of the model and of the data used for its estimation.

Desired family size. Overall, the simulated preferences closely match the shape of the observed distributions, reproducing the main peak around two to three children and the gradual taper into larger family ideals. A clear deviation appears at exactly two children, where the empirical data show a sharp bump (a clear outcome of strong social norms) that the model's simple two-parameter form inevitably smooths over. A more significant discrepancy is observed for Peru, where the model's predicted distribution is shifted to the left of the empirical one (JS divergence: 0.144 bits). This specific mismatch is consistent with the higher posterior uncertainty observed for the Peruvian cohort and may reflect underlying issues like lower data quality or greater-than-modeled heterogeneity in reproductive behavior in this cohort. It is important to note, however, that the alignment of these distributions improves substantially when additional information is incorporated (see Figure 9 in the Appendix).

Birth intervals. The synthetic data recovers the principal spacing pattern: a broad peak around two to three years followed by a gradual decline towards longer gaps. The model undershoots the cluster of very short gaps and shows a slightly heavier tail beyond five years. Apart from these edges of the distribution, the match is close across the main mass of birth intervals.

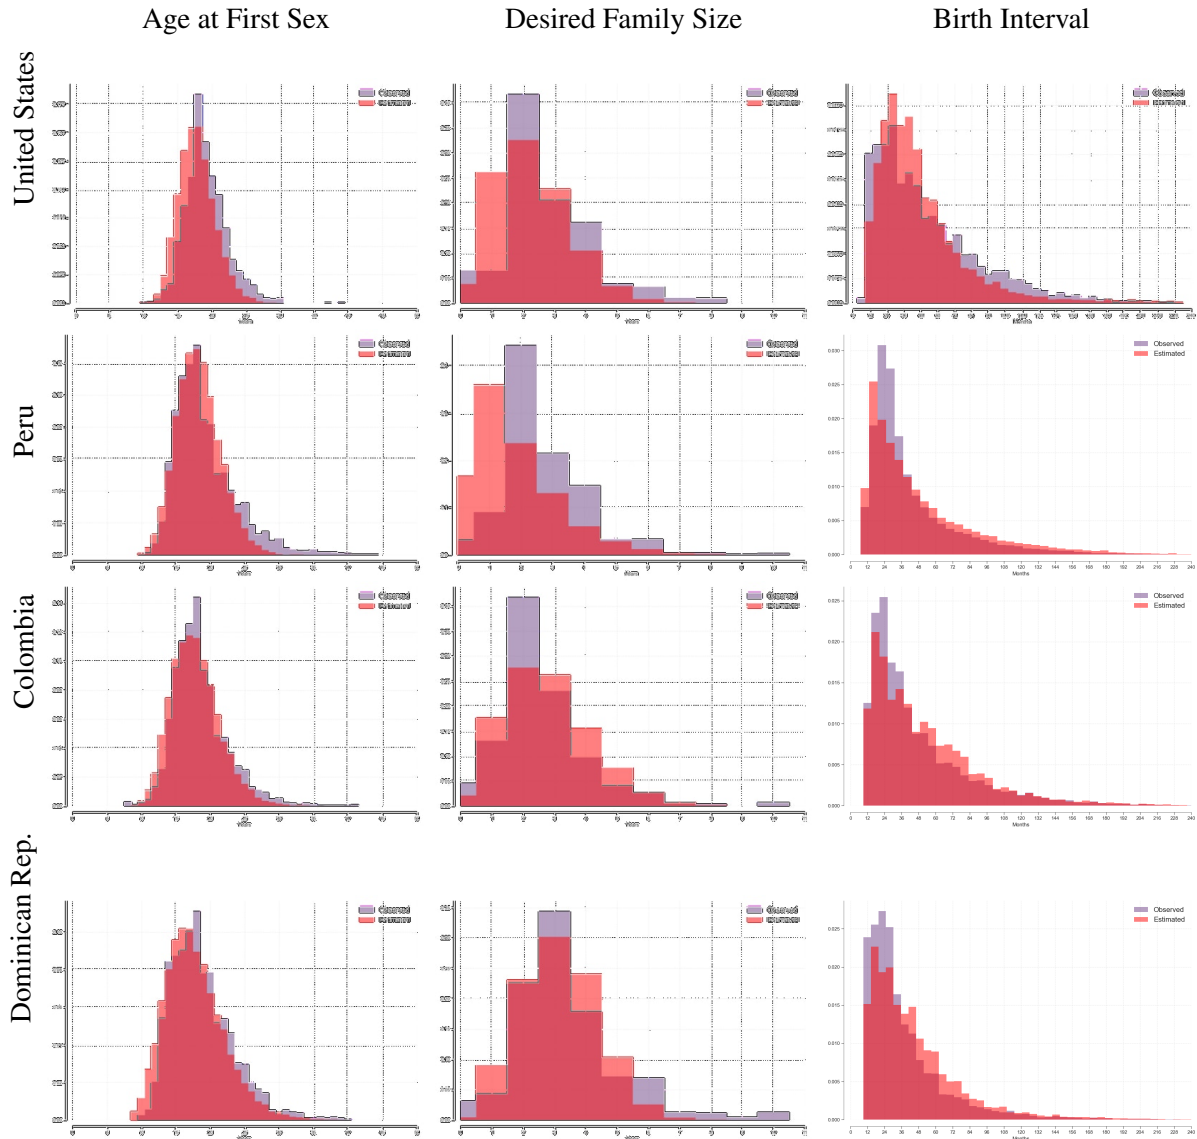


Figure 4: Validation on micro-level outcomes. Purple bars denote the observed survey distributions, while red bars show the model-implied counterparts generated from posterior draws. Columns display key behavioral dimensions: age at first sex, desired family size, and birth interval, whereas rows correspond to each study cohort.

6 Discussion

This study has introduced a likelihood-free Bayesian framework capable of inferring complex individual-level reproductive behaviors solely from aggregate population data. By integrating a demographically interpretable individual-level simulation model with Sequential Neural Posterior Estimation, we demonstrated that the core parameters governing the reproductive process can be recovered from age-specific fertility rates alone. While incorporating additional information (either through the model

priors or more detailed data) serves to further reduce parameter uncertainty, our findings show that the information contained in ASFRs alone provide a robust baseline estimation. The model's ability to accurately reproduce observed fertility schedules, as confirmed by our posterior predictive checks, offers strong support for this conclusion. However, the most compelling validation of our approach lies in the model's ability, using these inferred parameters, to accurately predict out-of-sample micro-level distributions for behaviors such as age at first sexual intercourse, desired family size, and birth intervals, none of which informed the original estimation. This establishes a new, statistically rigorous bridge between macro-level fertility observations and their underlying micro-behavioral drivers.

Methodologically, this work serves as an empirical demonstration of SNPE's utility for parameter inference in interpretable, yet complex, social science simulation models where direct likelihood calculation is prohibitive. The ability to recover these parameters from a single, widely available data series like ASFRs offers transferable insights for simulation-based inference across other disciplines facing similar inverse problems. Substantively, our framework provides a novel toolkit for demographers, enabling the estimation of behaviorally meaningful parameters (e.g., mean desired family size, age at intentional reproduction, contraceptive failure rates) even in data-scarce contexts, thereby offering a mechanistic complement to traditional micro-data analyses.

This capability to infer interpretable behavioral parameters significantly enhances population forecasting, moving beyond traditional methods that often rely on extrapolating aggregate trends or employ macro-level parameters with limited behavioral grounding. Our approach enables the generation of 'behaviorally explicit population forecasts' where future scenarios are constructed from hypothesized changes in these underlying, micro-level behavioral parameters. Crucially, by first distilling complex fertility dynamics into these more causally proximate behavioral components, our framework also facilitates a clearer understanding of how external covariates, such as education or economic conditions, shape reproductive outcomes. Machine learning models, for instance, may more effectively identify systematic relationships when predicting these interpretable behavioral parameters rather than attempting to directly predict highly aggregated and confounded ASFRs. This improved ability to model behavior mechanisms provides more robust and theoretically sound inputs for scenario-based population projections. Furthermore, by reducing the data requirements for microsimulation, our framework makes the construction of 'demographic digital twins' a more accessible and feasible tool for policy exploration.

Despite these advances, the present framework has limitations that define important avenues for future development. First, the individual-level model, while behaviorally detailed, necessarily simplifies the complex dynamics of human reproduction. Second, the data generating process encoded by our model is suited to contemporary populations where reproduction is substantially decoupled from sexuality and where such articulable birth control decision-making is prevalent. This may limit its direct applicability to contexts with more natural fertility regimes or where reproductive agency and intentions are structured differently. Third, the framework as presented assumes relatively homogeneous cohorts. Applying the model to highly heterogeneous populations where reproductive behavior differs systematically by socioeconomic characteristics would likely require disaggregated input data. A final consideration is that while ASFRs provide a robust foundation for estimation, applications that demand higher precision or focus on specific mechanisms may require supplementary data that are not always easily available.

These limitations point toward several exciting directions for future work. First, strategies can be developed to enhance the precision and utility of the framework, particularly for analyses relying solely on ASFRs. One important avenue is the incorporation of stronger, empirically-grounded priors on the key behavioral parameters, allowing researchers to leverage existing knowledge from demographically similar populations. Another promising direction is the development of methods to generate plausible supplementary data; for instance, one could use the learned associations between covariates (e.g., education, age) and birth intentionality from available survey to produce plausible estimates in data-scarce contexts. A second major direction involves adapting the framework's core structure to broaden its applicability and complexity. A key extension would be to explicitly model population heterogeneity. In the many contexts where data such as education-specific ASFRs are available, the framework could be adapted to infer distinct behavioral parameters for different subgroups. For populations closer to natural fertility, the model could be simplified to focus on a smaller set of parameters, while for contexts where Western notions of "unplanned" births are less salient, research could explore identifying and incorporating alternative summary statistics that better reflect local modes of fertility regulation.

Appendix

Figure 5 presents the cross-validation results for the four heterogeneity parameters, revealing varied patterns of recovery. The standard deviation of age at sexual initiation (σ_s) is recovered quite well across all scenarios, showing a low initial error that improves further with richer data. The recovery for the standard deviation of age at intentional reproduction (σ_r) clearly demonstrates the value of additional information; while the baseline model has some difficulty with higher parameter values in Scenario 1, the estimation improves significantly in Scenarios 2 and 3. In contrast, parameter recovery is less precise for the dispersion of desired family size (σ_d) and birth spacing (σ_b), where the estimates show considerable scatter even when additional information is incorporated.

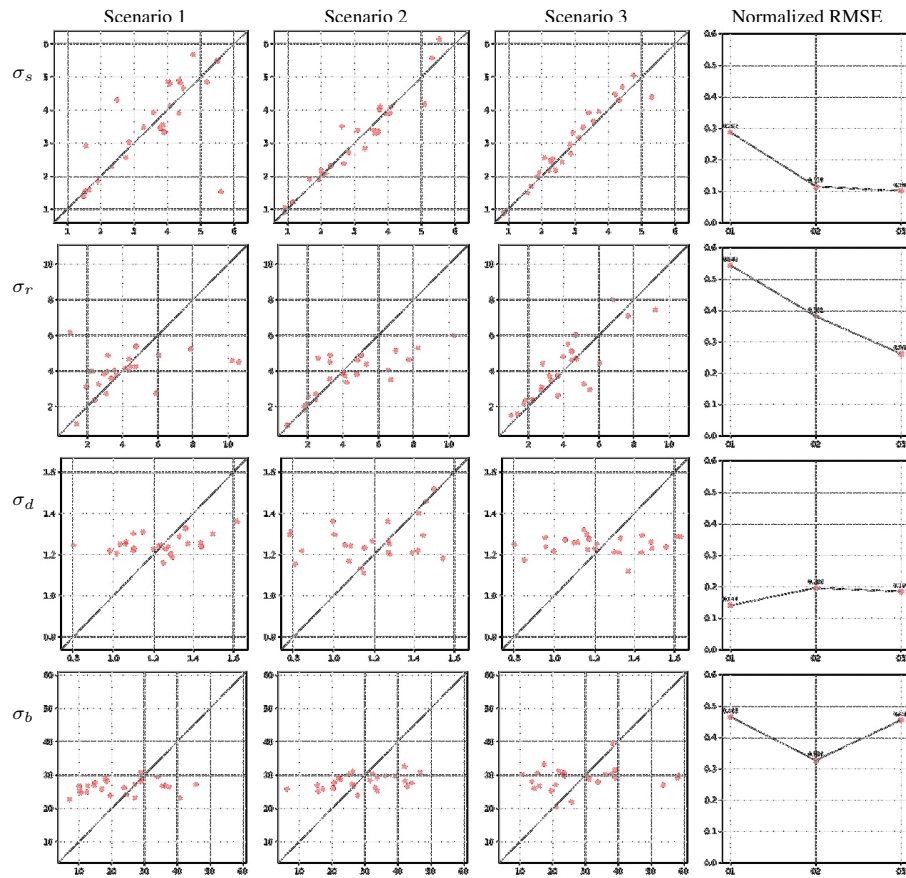


Figure 5: Cross-validation of the four heterogeneity parameters over 25 folds. These parameters govern the standard deviation of age at sexual initiation (σ_s), age at intentional reproduction (σ_r), desired family size (σ_d), and birth spacing (σ_b).

Figure 6 presents the full posterior predictive checks for all scenarios, now including the fit to the Age-Specific Unplanned Fertility Rate (ASUFR) alongside the ASFRs. While the model provides a strong fit to the observed ASFRs across all scenarios, the varying fit to the ASUFRs allows for a clearer visualization of the performance gains from incorporating different sources of information. Notably, even in the baseline scenario which uses no data on birth intentionality, the posterior predictive mean

tracks the ASUFRs, although the associated predictive interval is considerably wide. This predictive uncertainty is progressively reduced across the subsequent scenarios. The use of informative priors in Scenario 2 reduces the distance between the observed data and the predictive posterior means and leads to a sharper and more precise predictive distribution. The fit becomes most precise in Scenario 3, where the inclusion of ASUFR data allows the model to replicate the observed pattern of unplanned fertility with high accuracy.

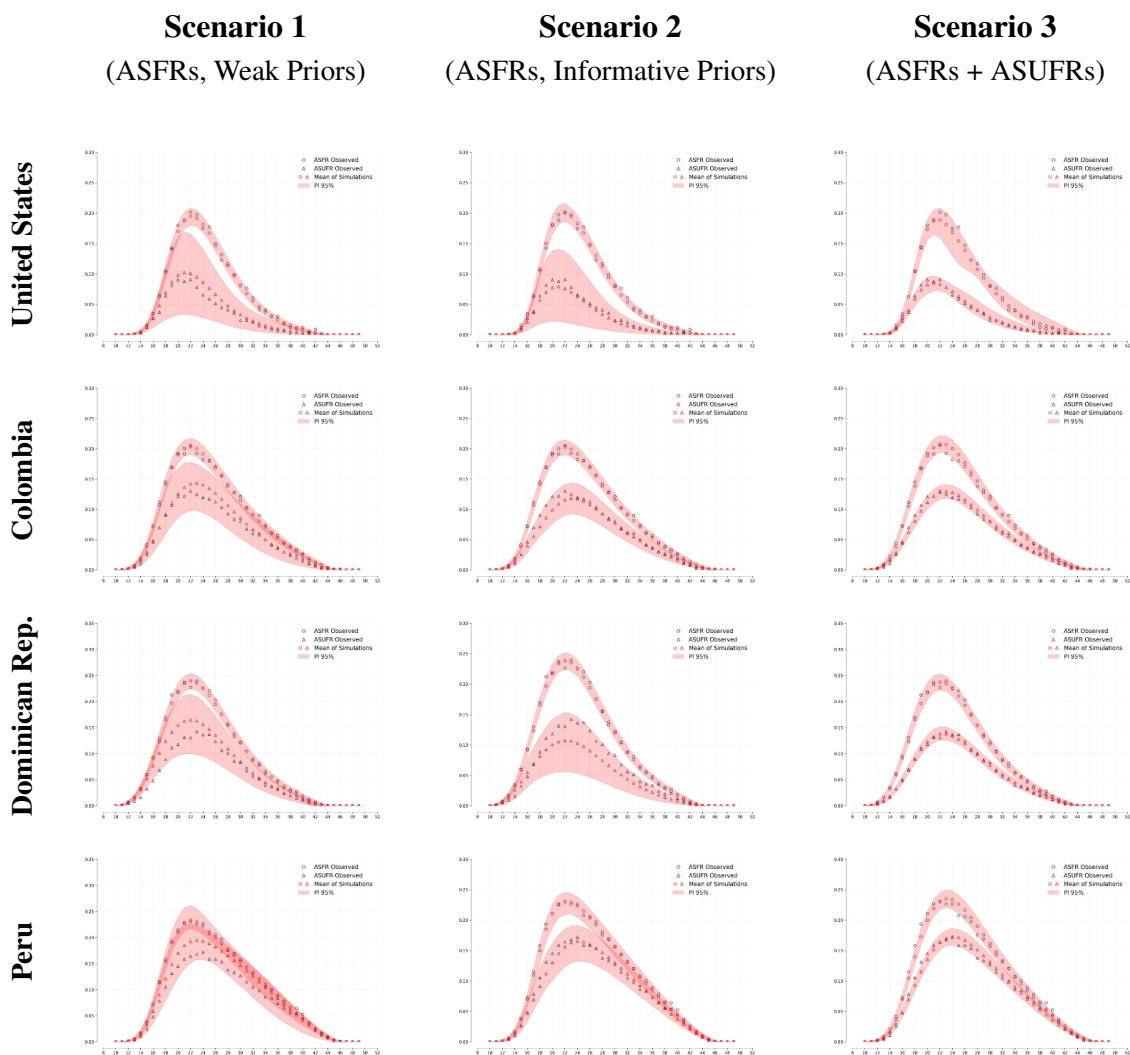


Figure 6: Full posterior predictive checks for both Age-Specific Fertility Rates (ASFRs) and Unplanned Fertility Rates (ASUFRs) across all three inference scenarios. Each panel compares the observed rates with the 95% posterior predictive interval generated under a given scenario (columns) for each country (rows).

Figure 7 compares the marginal posterior distributions across the three inference scenarios for each of the eleven model parameters and the four study countries (United States, Colombia, Dominican

Republic, and Peru). In each panel, the posteriors for Scenario 1 (light purple), Scenario 2 (orange), and Scenario 3 (red) are overlaid, illustrating the progressive reduction in parameter uncertainty as additional information is incorporated. Overall, the posteriors become narrower and more peaked from Scenario 1 to Scenario 3, reflecting improved precision with richer inputs, though the extent of this sharpening varies by parameter and country.

Notable patterns emerge across scenarios. For parameters related to the timing of sexual initiation (e.g., μ_s and σ_s) and fecundability (β_1 and β_2), the posteriors are already relatively well-constrained in Scenario 1 but show further refinement in Scenarios 2 and 3, with minimal overlap between the distributions. In contrast, timing-related parameters like the gap to intentional reproduction (δ_r) and desired birth spacing (μ_b and σ_b) exhibit higher initial uncertainty in Scenario 1, often with wider and flatter posteriors; this uncertainty decreases markedly in Scenarios 2 and 3, where the added information (priors or data) leads to substantial narrowing. The contraceptive failure rate (κ) and desired family size parameters (μ_d and σ_d) demonstrate consistent narrowing, particularly benefiting from the added information in Scenario 3, which reduces posterior variance by emphasizing unplanned fertility dynamics.

Differences in estimation behavior are evident across countries/datasets. For instance, the United States shows the most consistent posteriors across scenarios, with distributions narrowing progressively but remaining in similar locations; this stability is noteworthy given the higher data quality in the U.S. surveys, which may contribute to robust identifiability even under weaker priors. In contrast, Colombia and Peru exhibit more pronounced shifts and sharpening from Scenario 1 to the others, with posteriors in Scenario 1 often broader and sometimes multimodal, suggesting greater sensitivity to additional information in these contexts. The Dominican Republic falls in between, with moderate improvements in precision but less marked relocation of posterior mass.

Figure 7: Comparison of marginal posterior distributions across the three inference scenarios. Each panel shows the posterior for a given parameter (rows) and country (columns). The distributions for Scenario 1 (light purple), Scenario 2 (orange), and Scenario 3 (red) are overlaid to show how parameter uncertainty is reduced as more information is incorporated.

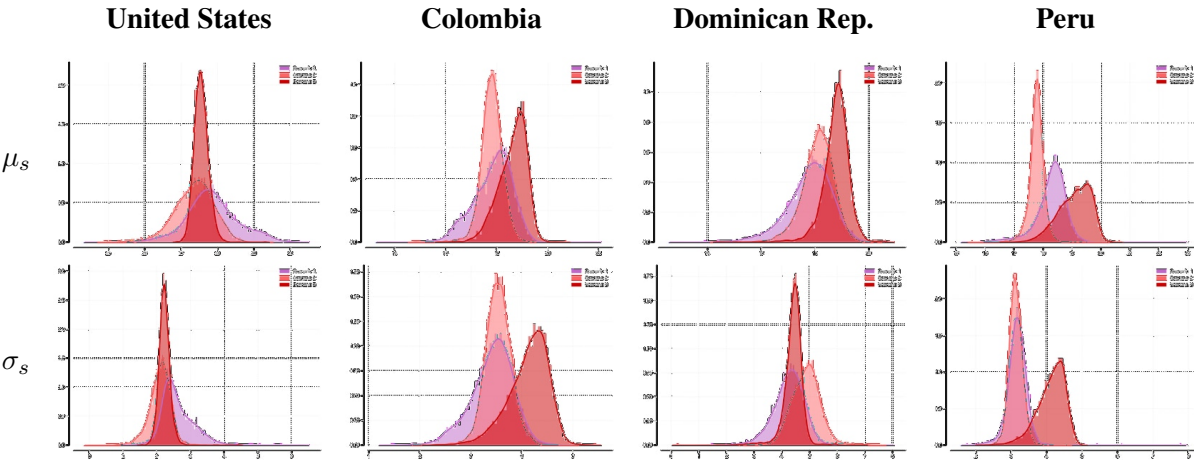


Figure 7 – continued from previous page

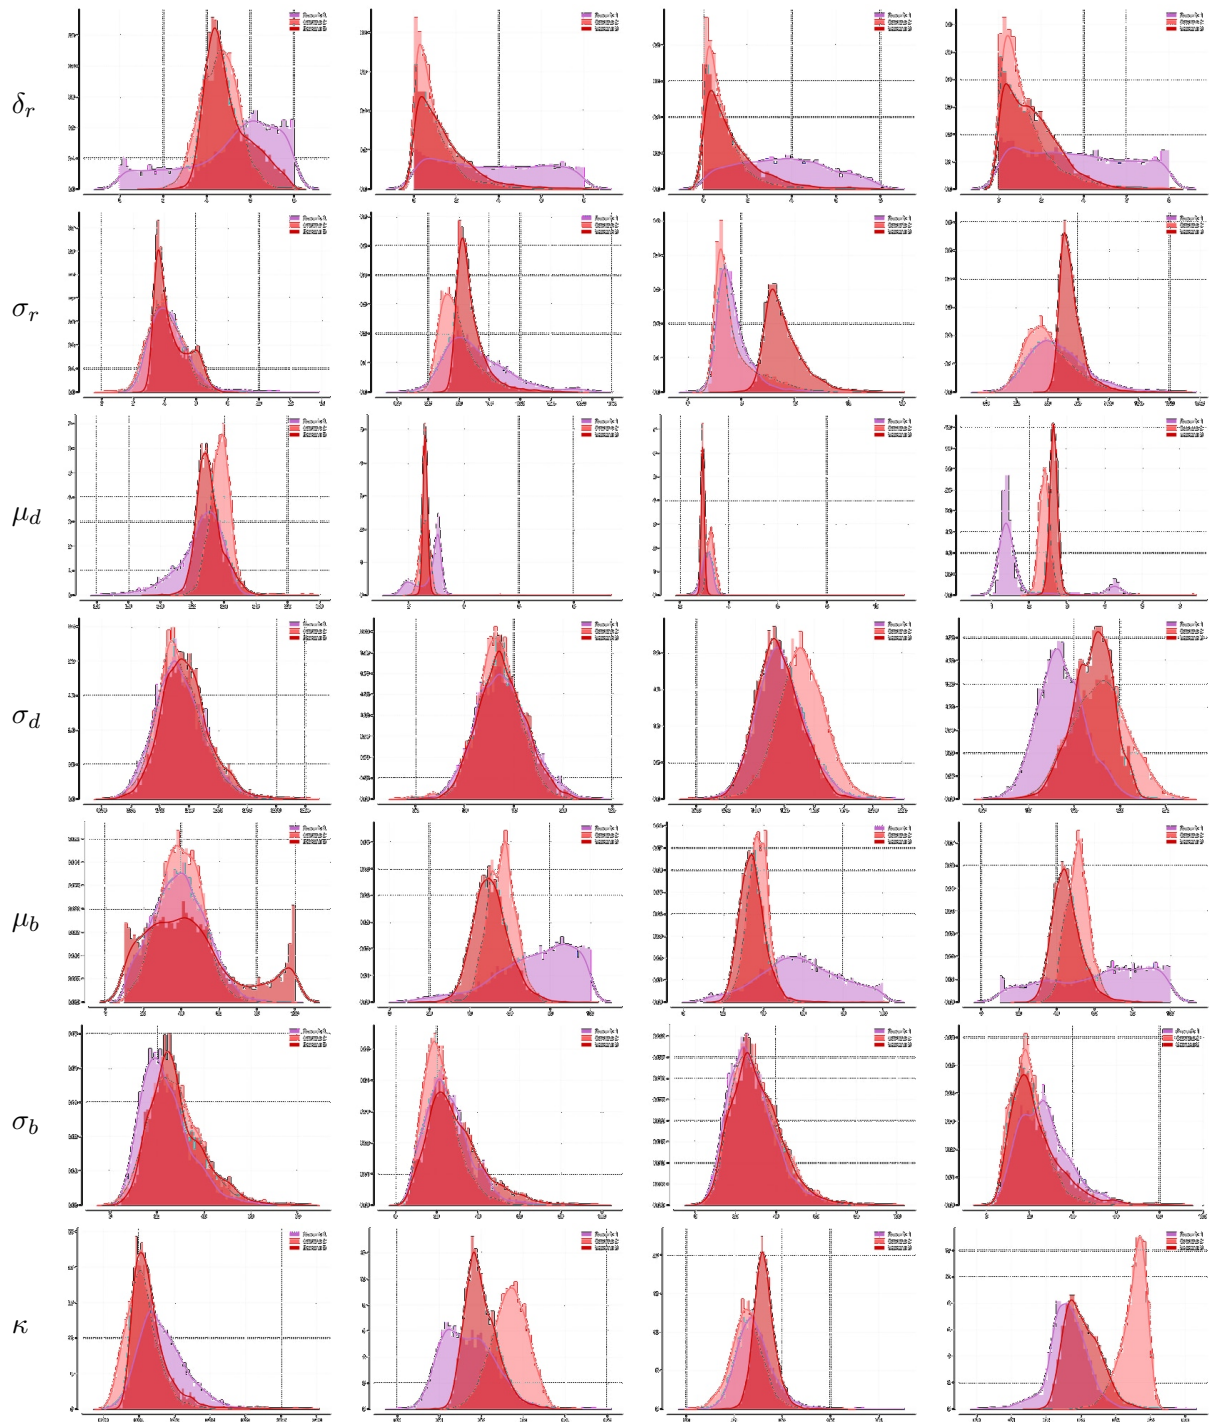


Figure 7 – continued from previous page

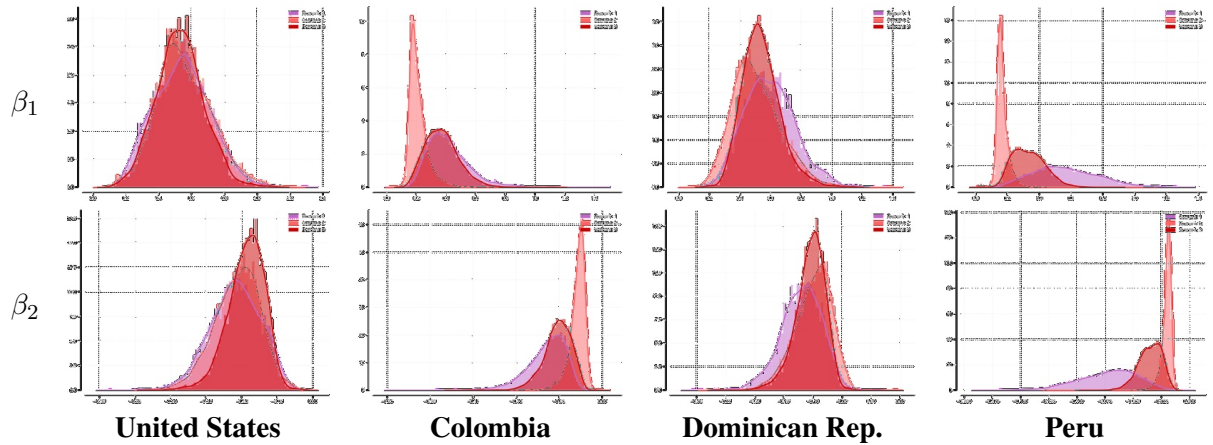


Figure 8 presents the out-of-sample validation for the distribution of Age at First Sex across all four countries and three inference scenarios. Overall, the model demonstrates a strong ability to predict this key micro-level distribution, even under the baseline Scenario 1 which uses only aggregate ASFRs for estimation. In most cases, the model-implied posterior predictive distribution (red) closely approximates the shape and location of the observed survey data (purple).

Key aspects include a general smoothing of observed heaping (e.g., at age 18) by the model across scenarios, with initial mismatches like leftward shifts in the United States and overestimation of early debuts in Colombia and Peru in Scenario 1. For the Latin American countries, Scenario 3 shows clear enhancements, with red lines adhering more closely to the bars, better resolving tails and peaks. In contrast, the United States exhibits consistent predictions across all scenarios, with the red line remaining slightly left-shifted and showing no substantial improvement.

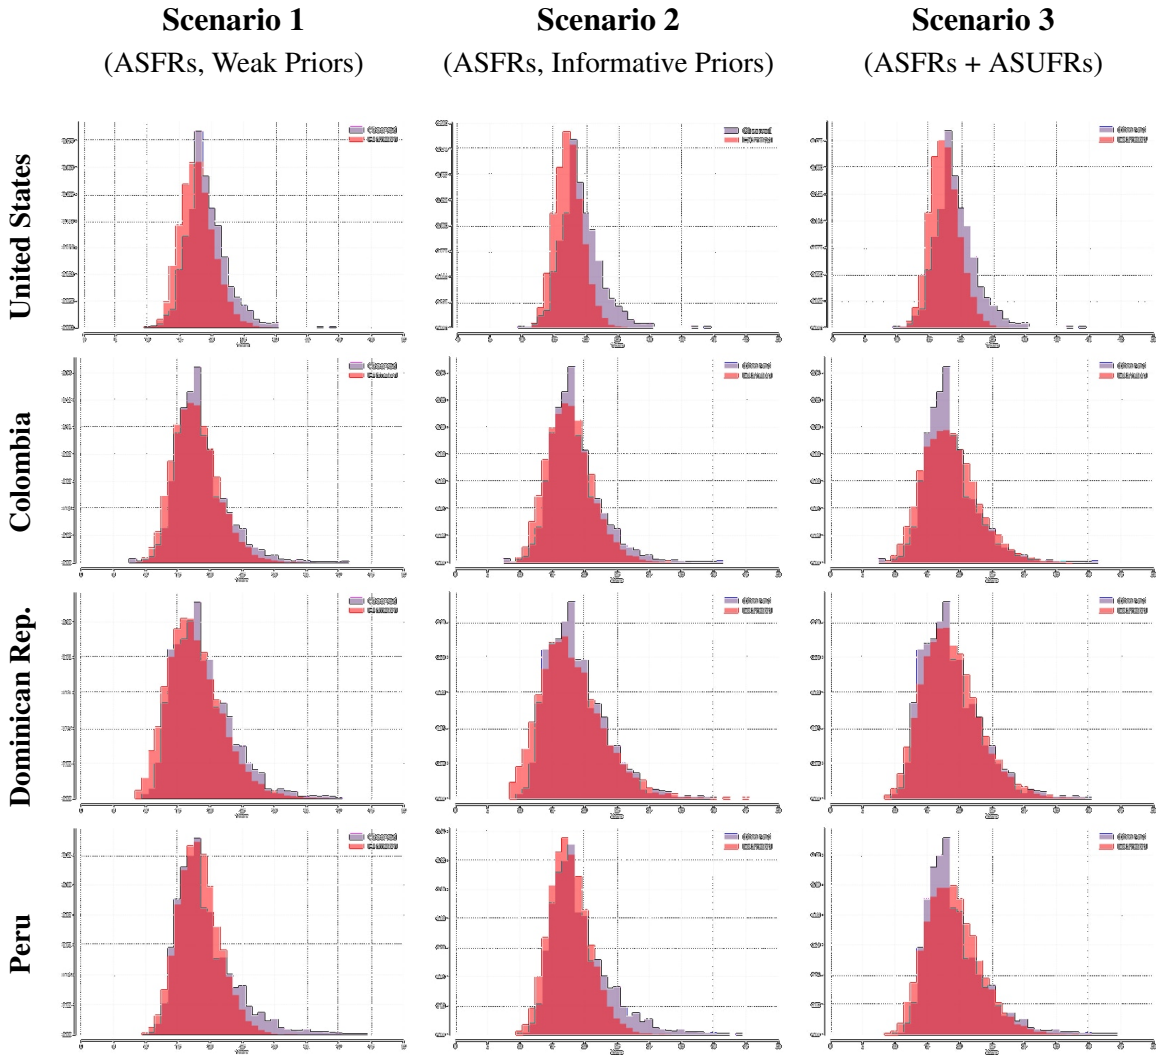


Figure 8: Out-of-sample validation for the distribution of Age at First Sex. Each panel compares the observed survey distribution (purple bars) with the estimated posterior predictive distribution (red line) for each country (rows) and inference scenario (columns).

Figure 9 shows the out-of-sample validation for Desired Family Size. Overall, the framework successfully recovers the general shape and rightward skew of the distribution in most cases. As noted in the main text, a consistent deviation is the model’s treatment of the two-child norm; using a smooth parametric distribution, the model captures the central tendency of desired family size but smooths over the sharp empirical peak at exactly two children that results from a strong social norm.

The results for Peru, however, most clearly illustrate the value of incorporating additional information. In Scenario 1, the model struggles to correctly identify the mode of the distribution. The fit improves progressively from Scenario 2 to Scenario 3, where the predicted distribution aligns much more closely with the observed data. This highlights that, for some populations, incorporating external knowledge can be critical to achieve an accurate estimation of specific characteristic of the data.

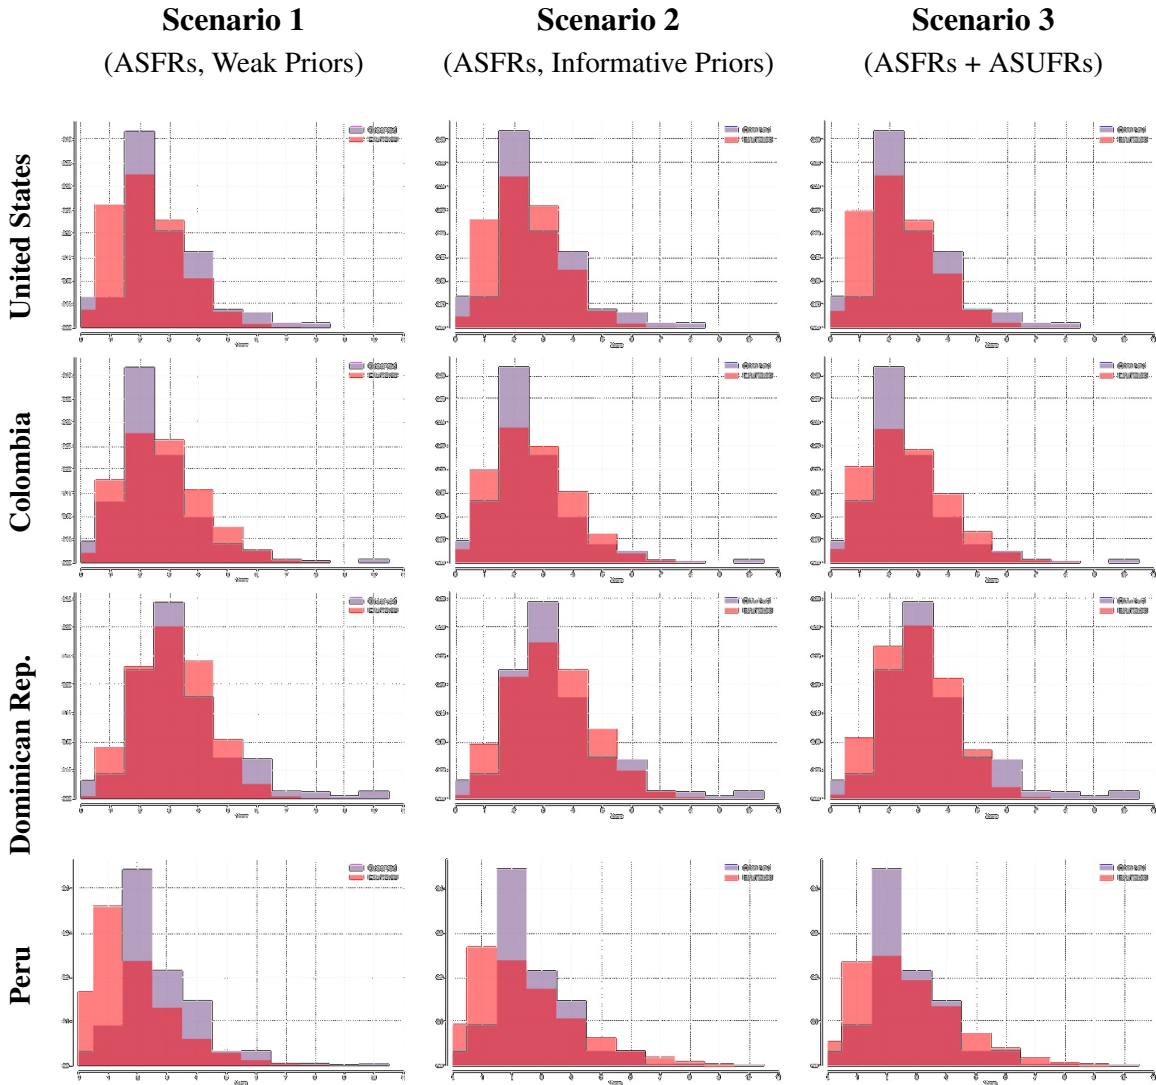


Figure 9: Out-of-sample validation for the distribution of Desired Family Size. Each panel compares the observed survey distribution (purple bars) with the model-implied posterior predictive distribution (red line) for each country (rows) and inference scenario (columns).

Figure 10 presents the out-of-sample validation for the distribution of Birth Intervals. The model is broadly successful in recovering the principal spacing pattern across all countries, correctly identifying a peak in intervals around two to three years, followed by a long, declining tail.

The most notable discrepancy occurs at this peak. The model's predicted distribution (red line) often appears smoother and wider than the observed survey data (purple bars), which can exhibit a sharper peak at the most common spacing intervals. The value of incorporating additional information is most evident for the Dominican Republic. In Scenario 1, the location of the predicted peak is noticeably offset from the observed data, but it becomes progressively more accurate in Scenarios 2 and 3, eventually aligning almost perfectly with the empirical mode. For the other countries, the improvements

across scenarios are more subtle, with the baseline model already providing a reasonable approximation of the distribution's central tendency, given model and data constraints.

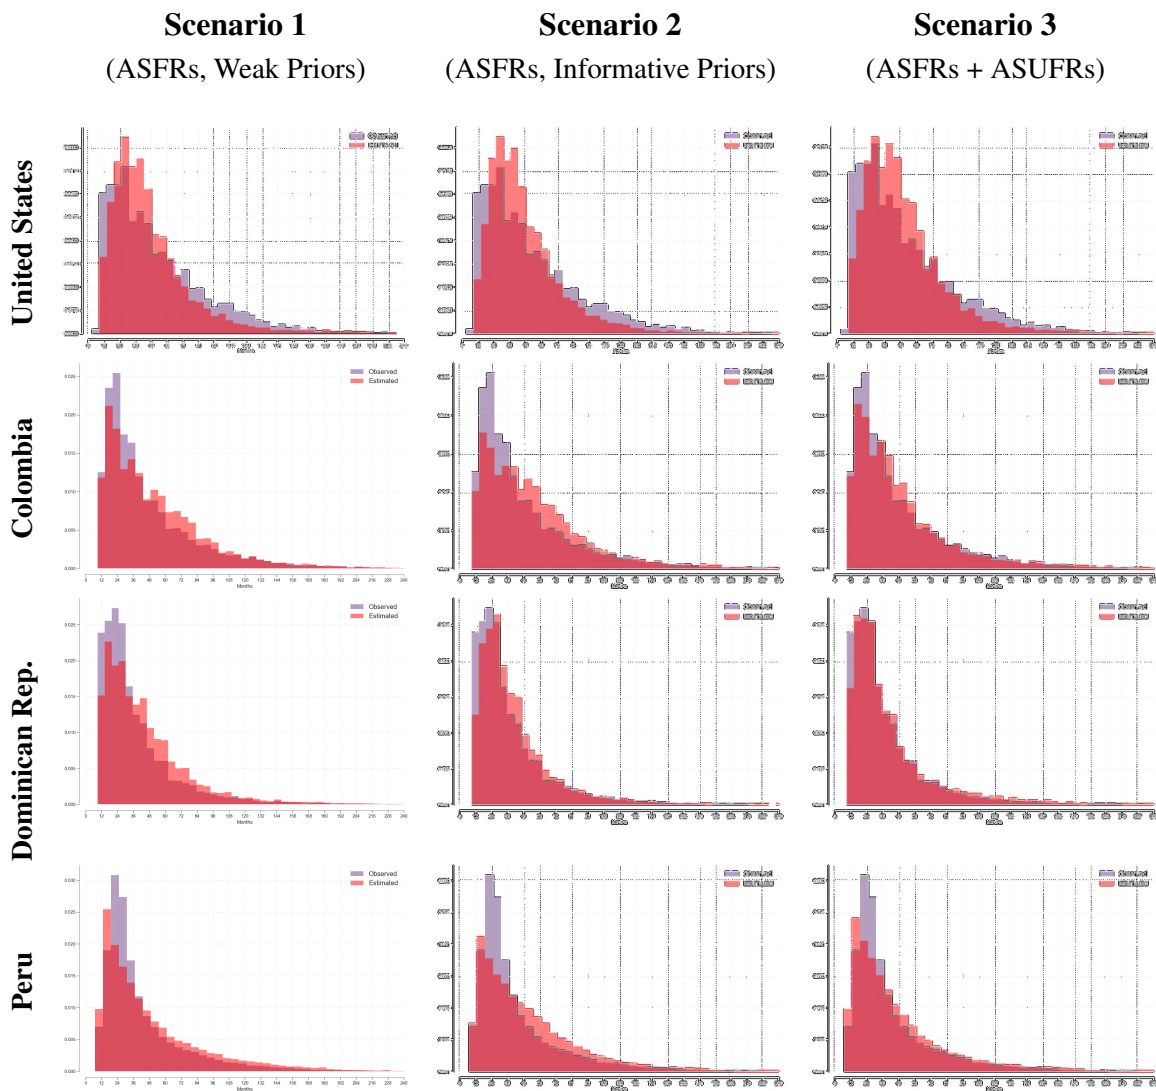


Figure 10: Out-of-sample validation for the distribution of Birth Intervals. Each panel compares the observed survey distribution (purple bars) with the model-implied posterior predictive distribution (red line) for each country (rows) and inference scenario (columns).

References

(2014). Female age-related fertility decline. *Committee Opinion No. 589, American College of Obstetricians and Gynecologists* 123(3), 719–721. Published jointly in *Obstetrics & Gynecology* and *Fertility & Sterility*.

- Beaumont, M. A. (2010). Approximate bayesian computation in evolution and ecology. *Annual review of ecology, evolution, and systematics* 41, 379–406.
- Bendel, J.-P. and C.-i. Hua (1978). An estimate of the natural fecundability ratio curve. *Social biology* 25(3), 210–227.
- Boelts, J., M. Deistler, M. Gloeckler, Á. Tejero-Cantero, J.-M. Lueckmann, G. Moss, P. Steinbach, T. Moreau, F. Muratore, J. Linhart, et al. (2024). sbi reloaded: a toolkit for simulation-based inference workflows. *arXiv preprint arXiv:2411.17337*.
- Bongaarts, J. (1977). A dynamic model of the reproductive process. *Population Studies* 31(1), 59–73.
- Bongaarts, J. and R. Lightbourne (1992). Fertility preferences in latin america: trends and differentials in seven countries. *Notas De Población* 20(55), 79–102.
- Brass, W. (1958). The distribution of births in human populations in rural taiwan. *Population Studies* 12(1), 51–72.
- Brass, W. (1974). Perspectives in population prediction: Illustrated by the statistics of england and wales. *Journal of the Royal Statistical Society Series A: Statistics in Society* 137(4), 532–570.
- Casterline, J. B. and L. O. El-Zeini (2007). The estimation of unwanted fertility. *Demography* 44(4), 729–745.
- Chandola, T., D. A. Coleman, and R. W. Hiorns (1999). Recent european fertility patterns: Fitting curves to 'distorted' distributions. *Population Studies*, 317–329.
- Ciganda, D. and N. Todd (2024). Modelling the age pattern of fertility: an individual-level approach. *Royal Society Open Science* 11(11), 240366.
- Coale, A. J. and T. J. Trussell (1974). Model fertility schedules: variations in the age structure of childbearing in human populations. *Population index*, 185–258.
- Dax, M., S. R. Green, J. Gair, N. Gupte, M. Pürner, V. Raymond, J. Wildberger, J. H. Macke, A. Buonanno, and B. Schölkopf (2025). Real-time inference for binary neutron star mergers using machine learning. *Nature* 639(8053), 49–53.
- Deistler, M., J. H. Macke, and P. J. Gonçalves (2022). Energy-efficient network activity from disparate circuit parameters. *Proceedings of the National Academy of Sciences* 119(44), e2207632119.
- Dunson, D. B., B. Colombo, and D. D. Baird (2002). Changes with age in the level and duration of fertility in the menstrual cycle. *Human reproduction* 17(5), 1399–1403.
- Durkan, C., A. Bekasov, I. Murray, and G. Papamakarios (2019). Neural spline flows. *Advances in neural information processing systems* 32.
- Gini, C. (1924). Premières recherches sur la fécondabilité de la femme. In North-Holland: (Ed.), *Proceedings of the International Mathematical Congress.*, Volume Vol. 2., Toronto.

- Gonçalves, P. J., J.-M. Lueckmann, M. Deistler, M. Nonnenmacher, K. Öcal, G. Bassetto, C. Chintaluri, W. F. Podlaski, S. A. Haddad, T. P. Vogels, et al. (2020). Training deep neural density estimators to identify mechanistic models of neural dynamics. *elife* 9, e56261.
- Greenberg, D., M. Nonnenmacher, and J. Macke (2019). Automatic posterior transformation for likelihood-free inference. In *International Conference on Machine Learning*, pp. 2404–2414. PMLR.
- Groschner, L. N., J. G. Malis, B. Zuidinga, and A. Borst (2022). A biophysical account of multiplication by a single neuron. *Nature* 603(7899), 119–123.
- Hartig, F., J. M. Calabrese, B. Reineking, T. Wiegand, and A. Huth (2011). Statistical inference for stochastic simulation models—theory and application. *Ecology letters* 14(8), 816–827.
- Henry, L. (1953). Fondements théoriques des mesures de la fécondité naturelle. *Revue de l'Institut International de Statistique / Review of the International Statistical Institute* 21(3), 135–151.
- Hoem, B. and J. M. Hoem (1989). The impact of women's employment on second and third births in modern sweden. *Population Studies* 43(1), 47–67.
- Kerry L. D., Lindsay Mallick, and Courtney Allen (2017). Sexual and reproductive health in early and later adolescence: DHS data on youth age 10–19. Technical report, USAID, The DHS Program, Rockville, Maryland, USA.
- Larsen, U. and S. Yan (2000). The age pattern of fecundability: an analysis of french canadian and hutterite birth histories. *Social biology* 47(1-2), 34–50.
- Mazzuco, S. and B. Scarpa (2015). Fitting age-specific fertility rates by a flexible generalized skew normal probability density function. *Journal of the Royal Statistical Society Series A: Statistics in Society* 178(1), 187–203.
- Potter, R. G. (1972). Additional births averted when abortion is added to contraception. *Studies in family planning* 3(4), 53–59.
- Ridley, J. C. and M. C. Sheps (1966). An analytic simulation model of human reproduction with demographic and biological components. *Population Studies* 19(3), 297–310.
- Schmertmann, C. P. (2003). A system of model fertility schedules with graphically intuitive parameters. *Demographic research* 9, 81–110.
- Schwartz, D. and M. J. Mayaux (1982). Female fecundity as a function of age: results of artificial insemination in 2193 nulliparous women with azoospermic husbands. *New England Journal of Medicine* 306(7), 404–406.
- Singh, S. (1963). Probability models for the variation in the number of births per couple. *Journal of the American Statistical Association* 58(303), 721–727.

- Sobotka, T. and É. Beaujouan (2014). Two Is Best? The Persistence of a Two-Child Family Ideal in Europe. *Population and Development Review* 40(3), 391–419.
- Vasist, M., F. Rozet, O. Absil, P. Mollière, E. Nasedkin, and G. Louppe (2023). Neural posterior estimation for exoplanetary atmospheric retrieval. *Astronomy & Astrophysics* 672, A147.
- von Krause, M., S. T. Radev, and A. Voss (2022). Mental speed is high until age 60 as revealed by analysis of over a million participants. *Nature human behaviour* 6(5), 700–708.
- Weinstein, M., J. W. Wood, M. A. Stoto, and D. D. Greenfield (1990). Components of age-specific fecundability. *Population Studies* 44(3), 447–467.
- Wellings, K., M. Collumbien, E. Slaymaker, S. Singh, Z. Hodges, D. Patel, and N. Bajos (2006). Sexual behaviour in context: a global perspective. *The Lancet* 368(9548), 1706–1728.
- Wesselink, A. K., K. J. Rothman, E. E. Hatch, E. M. Mikkelsen, H. T. Sørensen, and L. A. Wise (2017). Age and fecundability in a north american preconception cohort study. *American journal of obstetrics and gynecology* 217(6), 667–e1.
- Westoff, C. F. and N. B. Ryder (2015). *The contraceptive revolution*. Princeton University Press.
- Xie, Y. (2000). Demography: Past, present, and future. *Journal of the American Statistical Association* 95(450), 670–673.