MAX PLANCK INSTITUTE
FOR DEMOGRAPHIC RESEARCH

Konrad-Zuse-Strasse 1 · D-18057 Rostock · Germany · Tel +49 (0) 3 81 20 81 - 0 · Fax +49 (0) 3 81 20 81 - 202 · www.demogr.mpg.de

MPIDR Working Paper WP 2023-010 | February 2023
https://doi.org/10.4054/MPIDR-WP-2023-010

Future Fertility Scenarios in Finland: A Computational Forecasting Approach

Daniel Ciganda | ciganda@demogr.mpg.de
Julia Hellstrand | hellstrand@demogr.mpg.de
Mikko Myrskylä | office-myrskyla@demogr.mpg.de

Future Fertility Scenarios in Finland: A Computational Forecasting Approach

Daniel Ciganda ✉[1,2], Julia Hellstrand[1,3], and Mikko Myrskylä *[1,3,4]

[1]Max Planck Institute for Demographic Research
[2]Statistics Institute, University of the Republic, Montevideo, Uruguay
[3]Center for Social Data Science and Population Research Unit, University of Helsinki, Helsinki, Finland
[4]Max Planck – University of Helsinki Center for Social Inequalities in Population Health, Rostock, Germany and Helsinki, Finland

Abstract

Until the first decade of the 21st century, the Nordic countries maintained relatively high fertility rates, cementing the idea that a generous welfare system was one of the keys to preventing fertility from falling to unsettlingly low levels. This idea has been called into question by the recent decline in period fertility in the region. Although this trend has been observed across countries, Finland has clearly led the way, with its TFR reaching an all-time low of 1.35 in 2019. In this article, we use a novel computational modeling approach to look for insights into the factors that are driving this process, and to asses the likelihood of various future fertility scenarios in Finland. Rather than extrapolating from macro-level trends, our forecast is based on the simulation of individual trajectories. The main advantage of this approach is that it enables us to generate future fertility scenarios based on various hypotheses about the evolution of the main drivers of fertility change and to better understand the mechanisms influencing fertility outcomes. We find that the social processes that explained most of the variation in fertility rates in the past, like the expansion of higher education or the transition of women into the labor market, will play a diminishing role in the future trajectory of fertility rates, while individual preferences will become the dominant driver. Our forecast does not differ much from the predictions obtained with extrapolation methods in the short run, but it does offer a more optimistic outlook in the long run.

Keywords— fertility, Finland, computational modeling, microsimulation, demography, reproductive process.

*Mikko Myrskylä was supported by the ERC Synergy grant BIOSFER (#101071773) and the Strategic Research Council (SRC), FLUX consortium, decision numbers: 345130 and 345131.

1 Introduction

By the end of the first decade of the 21st century, the trajectories of fertility rates in many of Europe's high-income nations were showing signs of strength. Demographic forecasts obtained at the time predicted an increase in fertility levels in various regions, including in the Nordic countries and in Finland in particular (Myrskylä et al., 2013; Schmertmann et al., 2014). About a decade later, the scenario has changed drastically. Recent forecasts indicate that cohort fertility in Finland could fall from 1.9 to under 1.6 children per woman in the next 10 years, with Norway and Iceland experiencing similar declines (Hellstrand et al., 2020, 2021).

These recent developments have called into question the notion that a strong welfare state with progressive family policies and high levels of gender equality were preconditions for higher fertility levels in post-transitional countries. The decline in fertility rates in Finland and in the Nordic region as a whole remains puzzling, as it cannot be linked to a significant deterioration in living standards or a meaningful change in public spending on children and families.

Although the global recession of 2008 appears to have initially triggered the downward trend in the TFR, fertility rates in Finland continued to fall even after the economy began to recover around 2016 (Hiilamo, 2020).

In this context, a number of studies have focused their attention on different "subjective uncertainty" measures. However, Hellstrand et al. (2022) have shown that fertility is declining much faster among individuals in fields of activity characterized by higher uncertainty, which suggest that objective economic uncertainty still plays a central role in the family formation decisions of young people in Finland.

Although there is still uncertainty about its drivers, the decline has been extensively characterized from a demographic perspective. Recent studies have shown that the largest contribution to the drop in the TFR has come from a reduction in first births among women aged 25–29 (Hellstrand et al., 2020). Far from being a new phenomenon, the decline in birth rates among the younger generations is part of a long-term trend of fertility postponement that began at least four decades ago. However, Hellstrand et al. (2020) have shown that the downward trend in birth rates now extends to women in their thirties, a group whose fertility rates had been increasing since the mid-1970s.

Indeed, one of the most salient features of family formation dynamics in Finland after the global recession has been the acceleration of the postponement of the transition to parenthood. The mean age at birth has increased by 1.3 years since 2008, and set a new record in 2021, even though the TFR has stabilized in the last couple of years.

What remains to be determined is whether this new trend represents a "second wave" of postponement, or whether it is an expression of declining fertility preferences. By "second wave", we mean a process in which younger generations postpone childbearing even further, from the thirties to the early forties, following a previous shift in which the majority of the change occurred from the twenties to the thirties. If this is the case, the recent decline in the TFR may be largely a period phenomenon; i.e., a momentary decline that will fade once the delayed births are actually realized. The second scenario suggests a more long-lasting trend that is driven by a shift in preferences toward smaller families, including a rise in voluntary childlessness.

These two scenarios have very different impacts on the ultimate number of children that the cohorts who have not yet reached the end of their reproductive years will have in the future. This means that when attempting to predict future fertility levels, the accuracy of the results will be greatly influenced by the assumptions made about how fertility preferences will change in the future and the possibility of birth rates recovering at older ages. However, because they rely solely on information on the past trajectory of a given fertility indicator, traditional forecasting methods cannot account for any of these factors.

In this paper, we use a novel approach to model and forecast fertility outcomes that will improve our understanding of not only *where* fertility levels in Finland are likely to be in the near future, but also *why*. In this framework, originally introduced by Ciganda and Todd (2021), reproductive trajectories are modeled at the individual level, as an outcome of individual characteristics like educational attainment and labor force participation. As a result, the evolution of fertility trends, from both a period and a cohort perspective, can be directly linked to changes in the distribution of these individual characteristics across cohorts. In other words, fertility trends can be modeled and forecasted as the outcome of other social processes, like the expansion of higher education or the transition of women into the labor market.

Because of the abundance of information contained at the microlevel, this framework is well suited to generating scenario-based forecasts. It is possible, for example, to forecast the evolution of fertility levels based on a given assumption about the evolution of higher education or based on assumptions about other key aspects of the reproductive process, such as family size preferences or the mean age at birth. Modeling at the individual level also allows for the construction of any fertility indicator, bridging the gap between cohort-based and period-based forecasts.

2 Approach

Building on the idea of fecundability, defined by Gini (1924) as the probability of conception in the absence of contraceptive practices, Henry (1953) developed a model of reproductive trajectories in a population that makes no attempt to limit births. In this type of setting, which Henry defined as "natural fertility," the reproductive process can be roughly simulated by defining four fundamental components: fecundability (the probability of experiencing conception); the duration of the non-susceptibility period; as well as the start and the end of the reproductive process, which are typically defined as marriage and the onset of permanent sterility.

Using this model, it is possible to simulate the time from marriage until a first conception for a couple i. Once this first conception is observed, the process is "paused" until the end of the non-susceptibility period, which includes the length of the pregnancy plus the length of postpartum infertility experienced by the woman in our couple i. After this period has elapsed, the couple becomes exposed to the risk of a conception again. This phase is followed by another period of non-susceptibility, after which the partners go back to being exposed to the risk of a new conception. This process continues uninterrupted until the couple becomes permanently incapable of conceiving.

This straightforward model can be used to produce results that resemble the reproductive trajectory of a non-contracepting couple from marriage until age 50. More specifically, the model will simulate each birth on the trajectory and the corresponding age of the mother at each birth, which is enough data to calculate all commonly used fertility indicators.

Modeling this process in a regulated fertility context is slightly more challenging. One important difference between a natural fertility and a regulated fertility regime is that in the latter there is a dual risk: the risk of a conception when the couple intends to have a child, i.e., fecundability, and the residual risk when the couple is actively trying to prevent a pregnancy, defined by Bongaarts and Potter (1983) as *residual fecundability*.

However, there is another, more significant distinction. When effective contraceptive methods are available, a couple's final family size is primarily determined by the partner's own preferences and decisions rather than their biological capacity to conceive or their age at union formation.

A variety of personal traits influence these preferences and decisions, including the couple's socioeconomic status, living arrangements, psychological characteristics, beliefs, among many others.

Following these ideas, Ciganda and Todd (2021) proposed a framework to model reproductive trajectories in a regulated fertility context as an outcome of individual characteristics. The pseudo-code presented below describes the basic operation of the model contained in that framework.

The reproductive process is represented using a discrete event simulation, in which time does not advance at fixed time intervals, but rather with the realization of events (Zeigler et al., 2000). The entire process is described by four events: starting a cohabiting union; evaluating whether to have a child; having a child; and dying.

The distribution of education levels and labor force participation in the simulated population come from empirical data (see Section 2.1). Based on these input data, the algorithm produces the birth dates for the initial cohort: i.e, the group of women who were born the year the simulation begins.

The central part of the algorithm consists of a while loop that runs until the end of the observation window. In the first operation inside the while loop, the algorithm selects the next event to be simulated from a list of events ordered by their waiting times (the remaining time until their occurrence). Once the next event and its corresponding waiting time are selected, the algorithm updates the simulation clock and the rest of the variables that involve durations, including the age of all women in the population and the remaining times for the events waiting to be simulated.

The births of the initial cohort during the first calendar year of the simulation are the first set of events that the algorithm simulates. When a new birth is simulated, the algorithm first determines whether it is a female birth and then adds the new baby to the population if it is. This newly simulated woman is now assigned a number of characteristics and variables, such as the educational level she will attain, her future labor force status, and the waiting times for other events she will experience in the future.

After increasing the simulated population by one member, the indicators for the mother are updated, including her age at the time of the simulated birth, her current number of children, and the difference between her current number of children and her desired number of children.

Because the formation of a union signals the start of the reproductive process, the first step following union formation is to assign a desired family size to the corresponding woman/couple. The second operation defines the newly partnered woman's intention to have a child. The intention is determined by her educational attainment, employment status, and the time elapsed since the previous birth, if there was one. The intention will determine whether or not the newly formed couple decides to have a child within the next 12 months. If they do, the time it takes to conceive will be determined by the woman's fecundability f, which is a function of her age. If the woman and her partner decide to wait another year, they will still be subject to a risk defined by residual fecundability rf, which is also affected by calendar time, as is the efficacy and availability of contraceptive methods.

A couple must decide again (every year) whether to try to have a child after successfully avoiding death or an unintended pregnancy in the previous year.

The Comfert algorithm

read input data
educational attainment (cohort)
labor force participation by edu. (cohort)

initialization
generate wts to birth in the first year

While *time* < *end time* **do**

choose next event $\rightarrow$ *n_event*
update clock
update ages & waiting times

if *n_event* = *birth* **then**

  if *girl* **then**

    add to population
    assign:
    educational attainment
    years in education
    labor force participation
    wt to union
    wt to death
    wt to evaluate == Inf

  end
  update indicators of mother:
  age at birth
  nr. children
  Desired fam. size - nr. children
  wt to evaluate

end

if *n_event* = *death* **then**

  remove from the population

end

(*continues at top of right column*)

if *n_event* = *union formation* **then**

  assign desired family size D_i
  compute intention $I_{i,t}$

  if $I_{i,t} > x \sim U(0,1)$ **then**
    wt conception $\rightarrow f_x$
  else
    wt conception $\rightarrow rf_{x,t}$
  end
  if *wt conception* > *1 year* **then**
    wt evaluate == 1 year
  end

end

if *n.event* = *evaluate* **then**

  update intention $I_{i,t}$

  if $I_{i,t} > x \sim U(0,1)$ **then**
    wt conception $\rightarrow f_{i,x}$
  else
    wt conception $\rightarrow rf_{i,x,t}$
  end
  if *wt conception* > *1 year* **then**
    wt evaluate == 1 year
  end

end

if *year change* **then**
  compute indicators.
  remove those with age > *max_age*
end

end While

save output
age specific fert. rates; unplanned births;
desired fam. size.

end

wt = waiting time; n_event = next event; max_age = age at end of reproductive period (50); end time = last year for which model is run; f = fecundability; rf = residual fecundability

The preceding description focused on the structure of the computational model; i.e., the set of steps that allows us to generate synthetic reproductive trajectories for a cohort of women. While this is essential to understand the operation of the model, it has to be complemented by the description of the mechanisms and assumptions that translate individual characteristics into reproductive outcomes.

While a detailed description can be found in Appendix A, below we provide a brief summary of the most relevant mechanisms.

First, family size preferences are modeled as a function of a woman's labor force participation status in an attempt to capture the dynamics that lead dual-earner couples to prefer smaller families than couples in which women do not actively participate in the labor force. Second, the risk of an unplanned pregnancy is modeled as a function of women's educational attainment with the aim of capturing the relationship between excess fertility and lower educational attainment, specially among older cohorts. Finally, given that most people wait until they have completed their education to start a family, the transition to parenthood is also modeled as a function of educational attainment.

2.1 Long and Short-Term Drivers of Fertility Dynamics

The approach described in the previous section was originally conceived to model post-baby boom fertility trends as the outcome of three long-term social processes that have radically transformed family dynamics since the second half of the 20$^{\text{th}}$ century: the contraceptive transition, the expansion of higher education, and the transition of women into the labor market - which could also be described as the transition from a male-breadwinner model to a dual-earner model -.

Figure 1 depicts the evolution of the these three processes in Finland. While changes in educational attainment and labor force participation are depicted across birth cohorts, the diffusion of modern contraceptive methods is measured across calendar time. Regardless of the perspective used, the magnitude and the speed of these changes are remarkable.

Although the available data only go back only to the early 1930s cohorts, our estimates suggest that the proportions of women actively participating in the labor market more than doubled in the first four decades of the 20$^{\text{th}}$ century, at least among women with less than tertiary education (Figure 1b).

The prevalence of modern contraceptive methods has also increased at an impressive pace (Figure 1c). Although there is no information before 1971, the available data allow us to reconstruct the timing and the speed of the diffusion of these methods, which went from being practically unknown at the end of the 1950s to being used by roughly 80% of the population by the end of the 1980s.

The process of educational expansion has moved at a similar speed (Figure 1a). The percentage of women with tertiary education increased from less than 10% among those born before 1930 to close to 60% among those born in the late 1970s (Figure 1a). Interestingly, however, this process seems to have peaked, or even reversed, in the last decade. For the first time in about 75 years, younger generations are entering the labor force with levels of education that are the same or lower than those of their predecessors.

Figure 1: Proportion of Women by Education (a), Proportion of Women Actively Participating in the Labor Force by Education (b), Prevalence of Modern Contraceptive Methods (c) | Finland. The information on contraception comes from multiple sex and fertility surveys in Finland, as processed by the United Nations Population Division (United Nations and Social Affairs, 2022). The estimates relate to married or cohabiting women who were asked about the contraceptive method they used the last time they had intercourse. Estimates for educational attainment and labor force participation by education were obtained using Finnish population register data compiled by Statistics Finland. The data include women born in Finland and living in Finland for whom we could access complete education histories. Personal identification numbers were used to link the register data to different register sources, such as information on births, educational attainment, and employment status. This information is available starting in 1987/1989. Educational attainment estimates were derived using the highest educational attainment at age 30+ according to the ISCED 2011 classification. Primary education refers to ISCED 0–2, secondary education to ISCED 3–4, and tertiary education to ISCED 5–8. Estimates of labor force participation by education level were derived based on activity status at age 30–55 at the end of each calendar year.

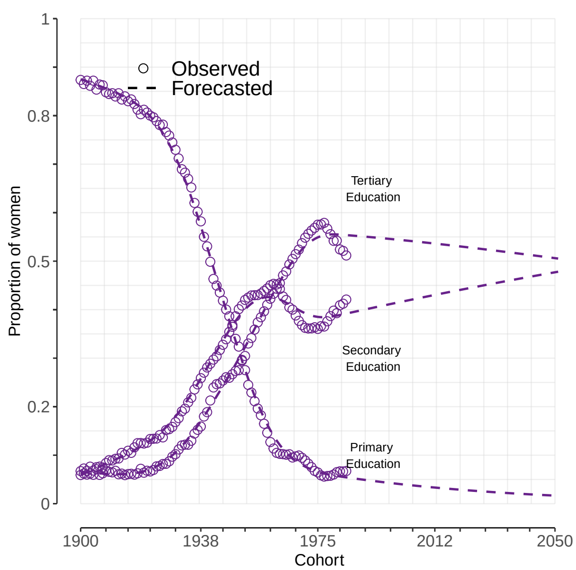

(a) Educational Attainment

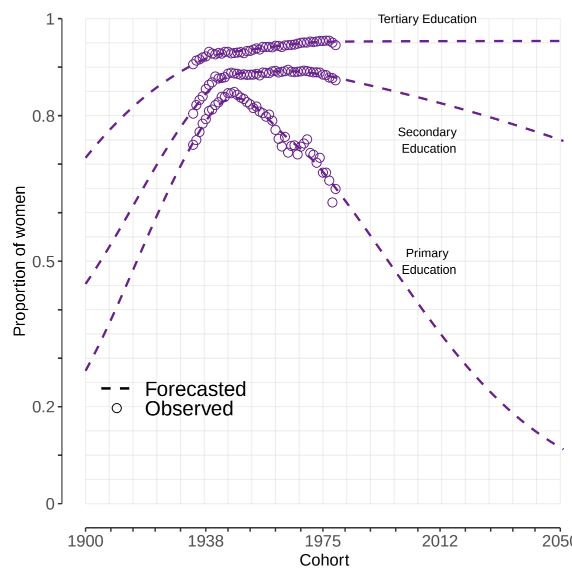

(b) Labor force Participation by Educational Attainment

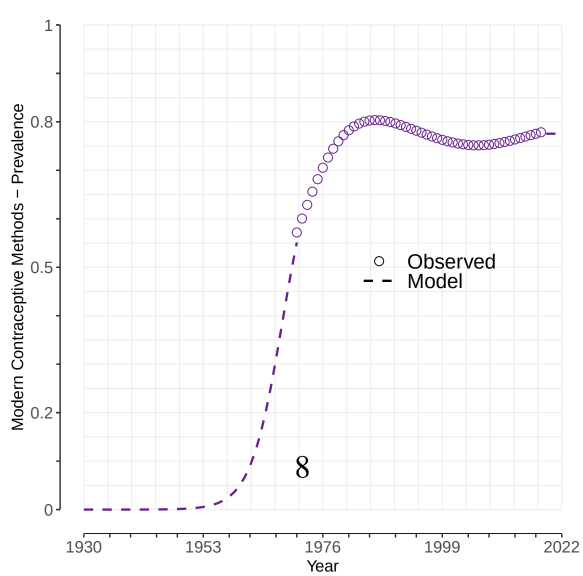

(c) Prevalence of Modern Contraceptive Methods

Taken together, the data in Figure 1 suggest that the three major transitions that have shaped fertility rates since the second half of the 20th century have come to an end. If this is the case, and these processes have indeed reached a long-term equilibrium, short-term shocks such as the shock associated with the 2008 global recession may become the primary source of variation in fertility rates in the future.

We model these dynamics by incorporating lagged period effects on the formation of family size preferences and the timing of the transition to parenthood for the 2010-2018 period (see Appendix A for more details). The resulting model incorporates both cohort dynamics associated with long-term societal transformations as well as shorter-term effects associated with the economic cycle. In other words, the reproductive trajectories of women in our simulated population will be shaped not only by their socioeconomic characteristics but also by the timing of key events in their life course.

Naturally, these simulated trajectories must be simplified versions of their empirical counterparts. While in reality people differ from one another in an almost infinite number of ways, in our model people are heterogeneous with respect to a limited set of characteristics. More importantly, while real individuals dissolve their unions, experience accidents, migrate, and have their plans delayed or altered, our simulated individuals have fixed and distinct preferences and progress toward their objectives in a rather straightforward manner. Given that simplification is an unavoidable part of modeling, the important question is whether the simplified set of mechanisms and characteristics described above can capture the key dynamics of fertility change in Finland. This question is answered in Section 3. Before that we introduce the approach we use to estimate the model.

2.2 Estimation

The parameters of the model are estimated using an Approximate Bayesian Computation (ABC) approach. This approach enables statistical inference on large-scale computational models (see: Beaumont, 2010, 2019). The aim of ABC algorithms is to identify the distribution of parameter values that best explains the observed data. This is typically accomplished by simulating realizations of a model's outcome at various parameter values and retaining those values that result in outcomes that are "close" to the observed data; i.e., the combination of parameter values that reduce the prediction error.

Since we are working with a computationally expensive model, we follow the approach developed by (Gutmann and Corander, 2016), which combines ABC with Bayesian Optimization with the objective of reducing computation times. Their approach could be defined as follows:

Let $\mathcal{X} \subset \mathbb{R}^d$ be the parameter space, where d is the number of parameters. At a given combination of parameter values $\theta \in \mathcal{X}$, we compute Δ_θ the Mean Squared Error (MSE) at this location as a weighted mean of the MSEs for the following vectors:

- Age-specific fertility rates between 1960 and 2021. Obtained from the Human Fertility Database (2011) for the 1960-2021 period and from Statistics Finland for 2021.
- Completed fertility by education for the cohorts born 1924 to 1970, obtained from Finland's national population register.

- The average number of children desired by women of reproductive age from 1997 to 2018. Available from multiple Family Barometer Surveys.

The objective of the algorithm is to find the region of the parameter space with the minimum distance Δ_θ. This is achieved by means of a sequential procedure that starts with the computation of $\Delta_{\theta_1}, ..., \Delta_{\theta_k}$ at a sample of k locations $\theta_1, ..., \theta_k$ from $\mathcal{X}$. This initial sample of locations and distances is then used to build an emulator; i.e., a model of the relationship between the values in $\mathcal{X}$ and the distance Δ_θ. We model this relationship non-parametrically, fitting a Gaussian Process regression with the help of the `mlegp` R package (Dancik and Dorman, 2008).

The emulator allows us to obtain predictions of the distance Δ_θ at any combination of parameters values without the need to run the original model at each of these locations.

In the next step, a new set of predictions is obtained from the emulator. Out of this set, the algorithm chooses the locations where either the error Δ_θ is smallest or the uncertainty is highest. The intention here is to focus on the regions of the parameter space where the minimum distances are more likely to be found, without leaving any region under-explored. The weight given to more uncertain versus more "promising" predictions is defined in the acquisition function $A(\cdot)$. In the final step, the algorithm computes the distance Δ_θ (runs the original model) at the locations selected with the acquisition function. The resulting information is then used in the next iteration to continue training the emulator.

The following pseudo-code describes the basic operation of the algorithm:

> Obtain an initial sample of $\mathcal{X}$ of size n_0, $\theta_1, ..., \theta_{n_0}$
> Compute $\Delta_{\theta_1}, ..., \Delta_{\theta_{n_0}}$
> Set $n = n_0$
>
> **While** $n \leq N$ **do**
>
> Map the relationship $\theta \to \Delta_\theta$ with a Gaussian Process emulator $G(\cdot)$
> Obtain a new sample of $\mathcal{X}$ of size n^*
> Obtain predictions for $G(\theta_k) \qquad k = 1, ..., n^*$
> Compute acquisition function $A(\cdot)$
> Obtain the new locations to be explored θ_j
> Compute Δ_{θ_j}
> Increment n
>
> **end While**
>
> Return:
> The value of θ that minimizes Δ_θ; or
> A fraction p of θ values with lowest Δ_θ

All the code and data needed to reproduce the results presented in the next section are available at `https://github.com/dciganda/comfert_nordic`.

3 Results

3.1 Model Fit

Figure 2 depicts the model's fit in relation to period age-specific fertility rates for selected years. The model is able to closely capture the evolution of the fertility age schedule over time, which is marked by a general decline in the average number of births across ages and the postponement of the mean age at birth, as evidenced by a shift in the mode of the distribution from the early twenties in 1960 to the early thirties at the end of the observation window.

Figure 2: Observed vs. Simulated Period Age-Specific Fertility Rates with 95% Credible Intervals, Selected Years, Finland.

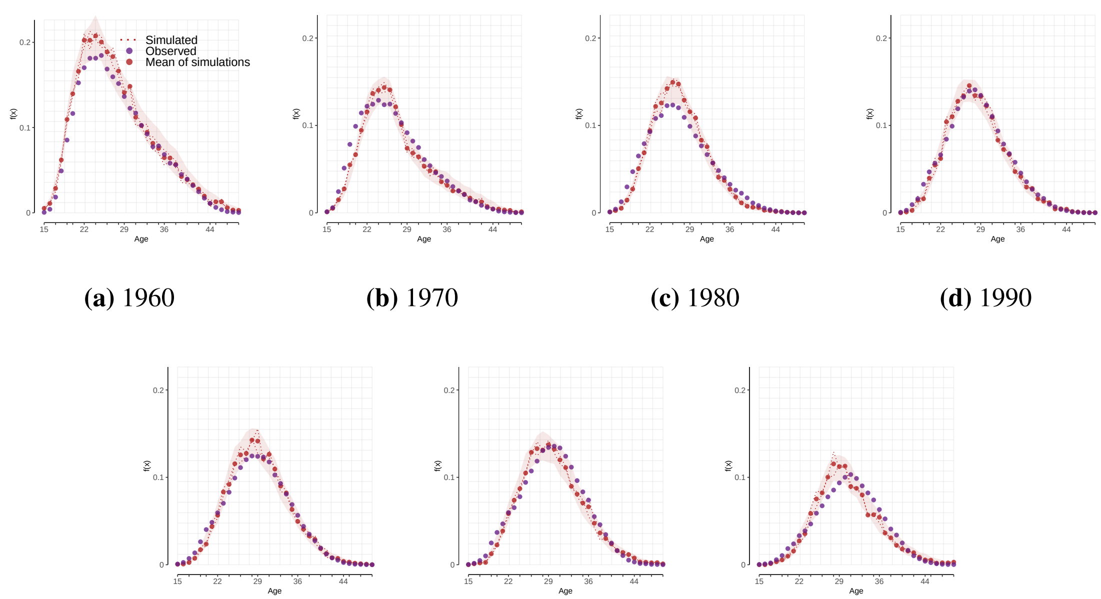

(a) 1960 **(b)** 1970 **(c)** 1980 **(d)** 1990

(e) 2000 **(f)** 2010 **(g)** 2020

Figure 3 presents the model fit with respect to the total fertility rate (Fig. 3a), the completed fertility of the cohorts born from 1925 to 1970 (fig. 3b), the desired number of children (Fig 3c) and the mean age at birth (Fig 3d).

Both sets of results (in Figures 2 and Figure 3) indicate that the underlying reproductive trajectories generated by the model reproduce the essential features of the observed reproductive trajectories of Finnish women, as well as their general evolution across age and time.

Figure 3: Observed and Simulated Total Fertility Rate (a) Completed Cohort Fertility (b) Family Size Preferences (c) and Mean Age at Birth (d) with 95% Credible Intervals | Finland

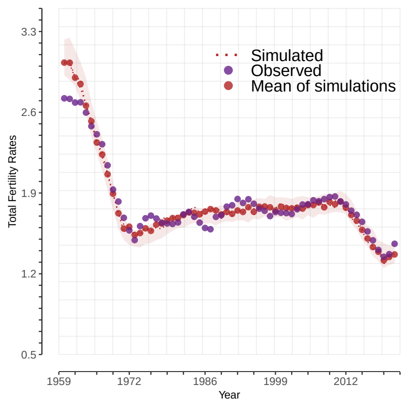

(a) Total Fertility Rate

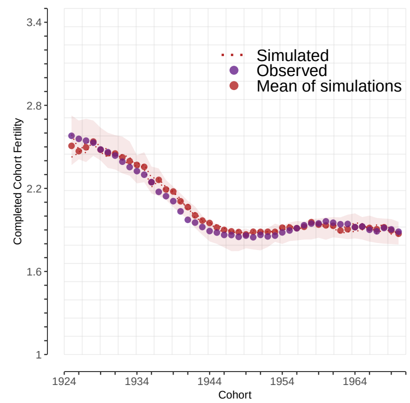

(b) Completed Cohort Fertility

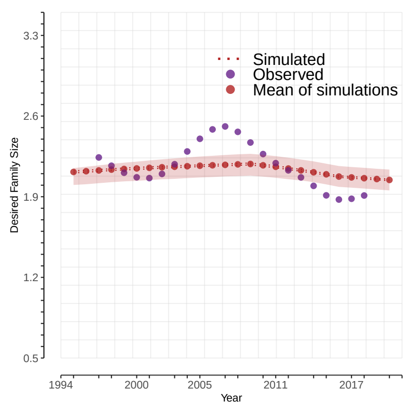

(c) Desired Family Size

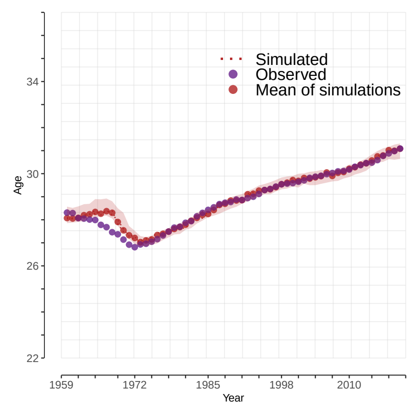

(d) Mean Age at Birth

Unlike the rest of the information displayed in Figure 3, the information on family size preferences (Figure 3c) comes from surveys with a limited sample size, which explains the large year-to-year variations. As a result, the precise value of the estimates is less informative than the overall pattern of the data. This pattern shows a slight increase in the desired number of children until 2008 followed by a steep decline until around 2017 when the pattern stabilizes, albeit below its pre-recession level.

As shown in Figure 3d, the model also closely tracks the evolution of the mean age at birth. The sustained increase in the time spent in school accounts for most of the variation in the mean age at birth until 2008. Following the recession, the model picks up strong period effects leading to the acceleration of the postponement process. Interestingly, this slight change of pace, together with a slight decrease in the desired number of children, result in a very significant change in the total fertility rate (Figure 3a).

3.2 Forecast

Figure 4 presents the future evolution of the TFR, Cohort Completed Fertility (CCF), mean age at birth and the desired family size under the most likely scenario according to our model. These forecasts are based on the assumption that educational attainment trends will follow the trajectory shown in Figure 1.

Although the observed trend in the mean age at birth does not indicate an impending end to the postponement process, the deceleration, or even reversal, of the educational expansion process suggests this is a very likely scenario. As Figure 4d shows, the model predicts this outcome. If this prediction is accurate, the TFR might recover its pre-recession levels in the next decade or so (Figure 4a). But rather than leading to an actual increase in the level of fertility, this recovery will be an artifact that can be linked to the well-known limitations of period fertility rates rates.

As the evolution of cohort rates indicates, in the most likely scenario the average number of children at the end of the reproductive period will continue to decrease slowly for about two decades, dropping slightly below 1.6 children per woman for the cohorts born at the end of the 1980s. Thereafter, it will recover slightly, and stabilize at between 1.6 and 1.7 children per woman.

Both the TFR and the CCF trajectories are heavily dependent on the expected trajectory of family size preferences. As shown in Figure 4c, the average desired number of children is expected to decline gradually for about a decade before stabilizing at a level slightly above 1.9. While we assume that the effect of the economic recession on fertility preferences eventually wanes, the resulting downward trend has its own momentum as existing preferences act as a reference (a baseline) for the new generation reaching childbearing ages. This generates a path-dependent process that extends the decline of family size preferences beyond the immediate aftermath of the recession (see Appendix A for more details on how fertility preferences are modeled).

Figure 4: Forecasted Total Fertility Rate (a) Completed Cohort Fertility (b) Ideal Family Size (c) and Mean Age at Birth (d) | Finland

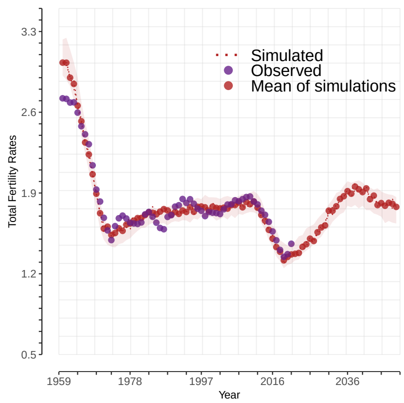

(a) Total Fertility Rate

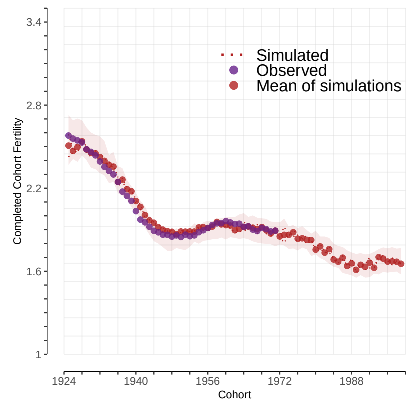

(b) Completed Cohort Fertility

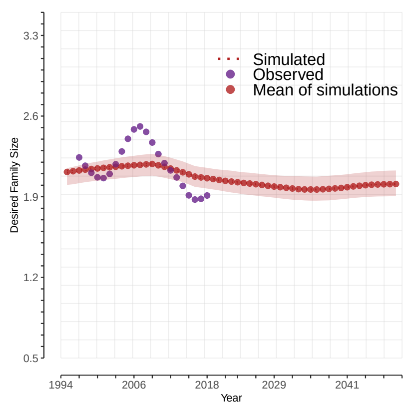

(c) Family Size Preferences

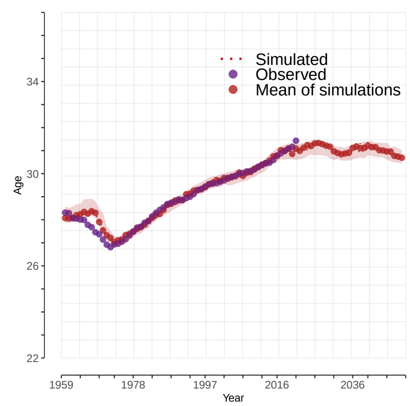

(d) Mean Age at Birth

Finally, in Figure 5, we compare our forecast for the mean cohort number of children with the forecasts obtained using three of the best-performing extrapolation methods (Bohk-Ewald et al., 2018). Although these three methods approach the problem from different angles, they all rely exclusively on information about the past trajectory of fertility rates to produce their forecasts.

While the expected trajectory of cohort fertility rates is similar in all four models over the short term, it begins to diverge significantly for the cohorts born after 1990. The expected stabilization of family size preferences and the deceleration of the postponement process result in higher expected fertility rates for younger cohorts in our model.

Figure 5: Cohort Completed Fertility. Comparison of Approaches | Finland The five-year extrapolation method by Myrskylä et al. (2013) takes the previous five-year trend, extrapolates it into the future and freezes the rates. The Bayesian method developed by Schmertmann et al. (2014) produces a probabilistic forecast by extrapolating trends in fertility rates over time and age, using information on the previous trajectories of age-specific fertility rates in a large number of countries as prior data. Finally, the freeze rate method extrapolates the latest set of observed age-specific fertility rates into the future.

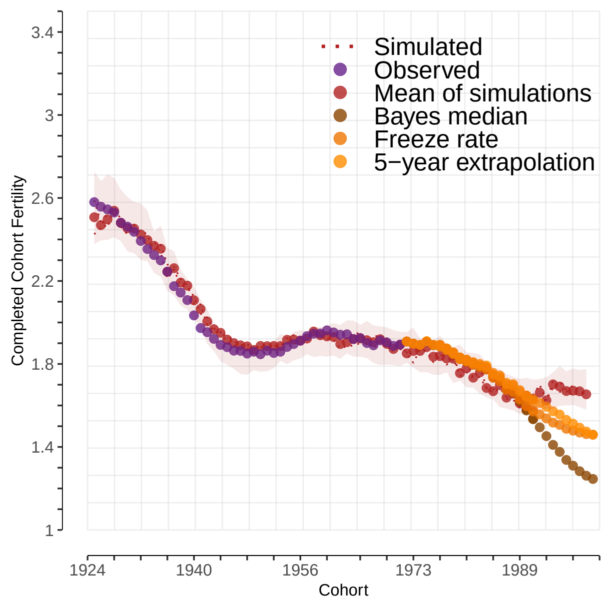

3.3 Scenarios

As was mentioned earlier, our approach allows us to generate a wide range of hypothetical future scenarios in the form of "what if" statements by modifying the assumed future evolution of the main drivers of fertility change in Finland. For example, what if education resumes its upward trend? What if it experiences a reversal? But we could also analyze how fertility levels would evolve under alternative trajectories of other dimensions of the reproductive process. For example, what if preferences go back to replacement level? What if they continue to decline?

To answer these questions, we developed two sets of scenarios: one corresponding to two alternative trajectories of educational attainment and another corresponding to two alternative trajectories of family size preferences. Figure 6 depicts the two sets of scenarios.

Figure 6a shows the two alternative trajectories of the proportion of women by education level and birth cohort. In our first scenario, the recent decline is only temporary and the progress toward higher levels of education continues after a brief interruption. In our second scenario the process of educational expansion indeed reverses, with future generations attaining, on average, lower and lower levels of education than their predecessors.

Figure 6b, displays the alternative family size preferences scenarios. In the first scenario, preferences recover and stabilize at replacement level (an average of 2.1 desired children per woman), while in the second scenario, preferences continue to decline below the levels expected in the most likely scenario represented in our initial forecast. The question of how economic uncertainty affects on family formation behavior has long attracted attention in the fertility literature (Vignoli et al., 2020; Matysiak et al., 2021). The exacerbation of these uncertainties, combined with potential feedback mechanisms leading to a low fertility trap (Lutz et al., 2006), can increase the likelihood of our low preferences scenario becoming reality. Conversely, the implementation of policies aimed at reducing these uncertainties in early adulthood and facilitating work-family balance may increase the chances that our high preferences scenario will prove accurate.

Figure 6: Education Scenario (a) Family Size Preferences Scenario (b) | Finland

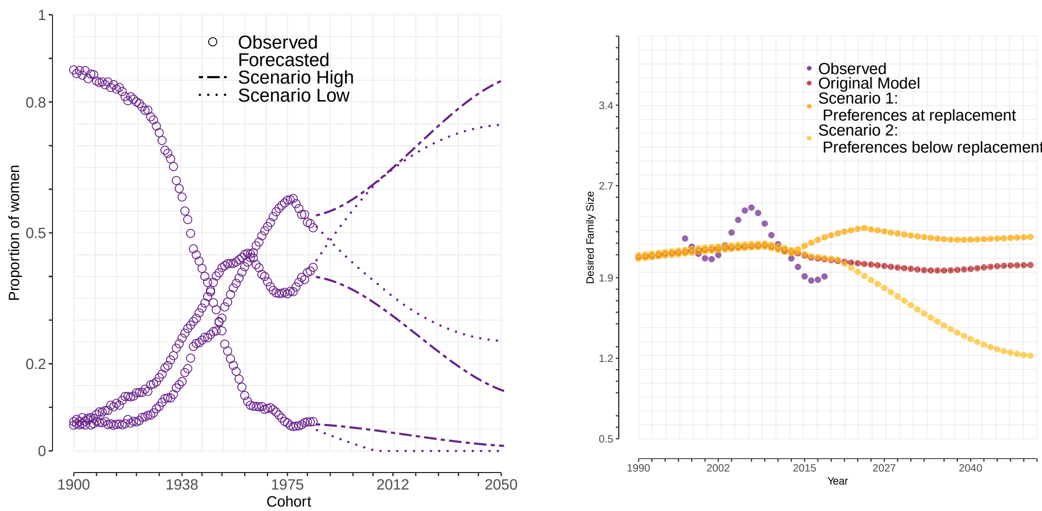

(a) Educational Attainment

(b) Family Size Preferences

Figure 7 displays the resulting cohort mean number of children under the two sets of scenarios. The evolution of the level of fertility under the alternative education scenarios is rather similar. This is not surprising considering that some of the mechanisms by which education influences family formation are no longer as relevant as they previously were given the stage of the fertility transition that Finland is currently in. For example, a higher education level is associated with a higher likelihood of labor force participation and, as a result, a smaller desired family size. However, as shown in Figure 1, In Finland, the transition of women into the labor market is almost complete for women of all education levels. A similar patter can be observed for the use of contraception, which is assumed to differ less between women with different levels of education than it might have varied earlier in the contraceptive transition.

The differences in the two scenarios are more visible when adopting a period perspective. Figure 8 shows the evolution of the TFR under the two alternative scenarios. As expected the TFR is consistently lower in the scenario in which the educational expansion resumes its upward trajectory. This is explained by the fact that the mean age at birth continues to rise in this scenario.

Figure 7: Forecasted Completed Cohort Fertility. Education Scenario (a) Family Size Preferences Scenario (b) | Finland

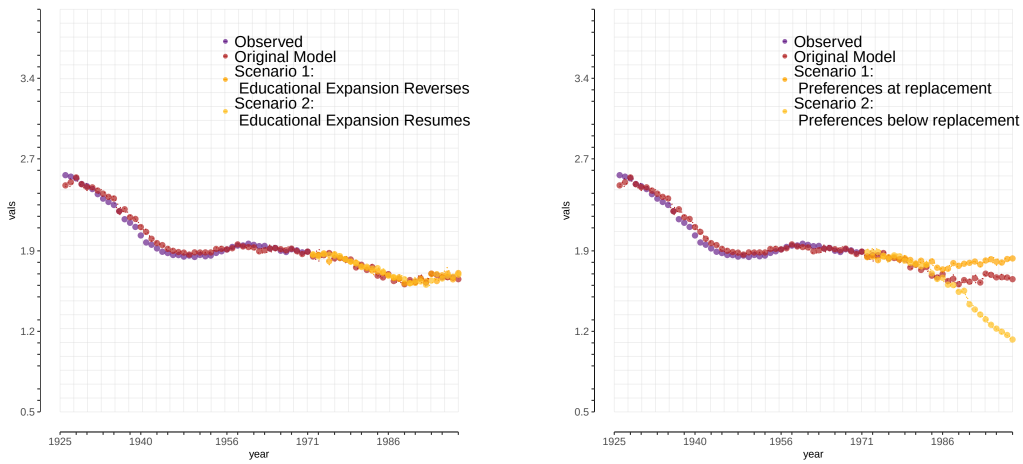

(a) Cohort Fertility - Education Scenarios

(b) Cohort Fertility - Preferences Scenarios

Figure 8: Forecasted Total Fertility Rate. Education Scenario (a) Family Size Preferences Scenario (b) | Finland

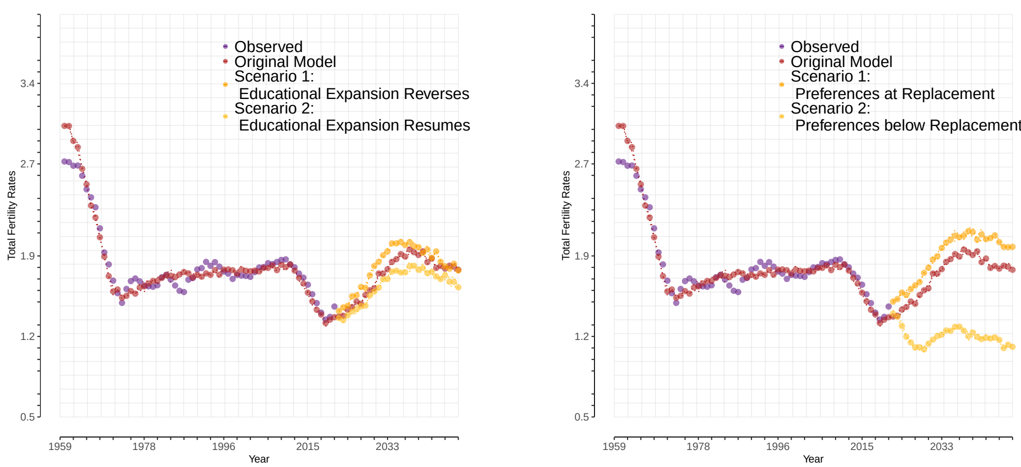

(a) Total Fertility Rate - Education Scenarios

(b) Total Fertility Rate - Preferences Scenarios

The trajectories of both period and cohort rates (Figures 8b and 7b) under the alternative preferences scenarios display more notable differences. In the scenario in which preferences return to replacement level, the cohort mean number of children is expected to recover as well, stabilizing at around the 1.8 level; while in the scenario in which preferences continuously decline, fertility levels are expected to follow a similar path until they reach a very low level of 1.1 children per woman on average.

The sensitivity of fertility indicators to family size preferences is not unexpected. As explained earlier, the main determinant of fertility levels in societies at very advanced stages of the fertility transition are individual preferences.

4 Discussion

The main goal of our article was to illustrate how an individual-level, computational approach to demographic forecasting can broaden our understanding of potential fertility scenarios. Finland provided us with an interesting and challenging case study, as it has recently experienced an unexpected, pronounced downward trend in the TFR, which might hold clues for the future trajectory of fertility rates in other highly developed countries.

Our approach allowed us to place recent developments within the longer-term fertility transition that began in the latter half of the 20th century.

Our analysis suggests that the recent drop in fertility levels is more likely a permanent trend driven by a reduction in the number of children people actually want to have, rather than a temporary event associated to the postponement of births. Moreover, we showed how modest, but simultaneous changes in key aspects of the family formation process can lead to momentous changes in period fertility indicators like the total fertility rate.

A second main insight from our analysis is that the recent reduction in family size preferences does not appear to be driven by the same set of factors than those that explained the majority of the variation in birth rates over the preceding decades.

The early transition of Finnish women into the labor force has drastically reduced one of the main sources of variability in family formation behavior among recent cohorts.

The expansion of higher education, the second major driver of fertility change since the second half of the 20th century, appears to have stalled as well. This is an extremely interesting and surprising development. While the transition of women into the labor force and the adoption of modern contraceptive methods have reached saturation points, the expansion of higher education appears to have peaked sooner than expected, with the proportion of women with tertiary education remaining below the 60% threshold.

The completion of the three transitions that drove fertility change since the second half of the 20th century has important implications for our understanding of family formation dynamics. Researchers have recently highlighted the limitations of established theories to explain recent fertility dynamics in Finland (Hiilamo, 2020). This is not surprising considering that the underlying drivers of fertility change are themselves changing.

Despite the limitations in the quality of the available information, recent data suggest that the post-recession decline in family size preferences may have reached a plateau. As a result, our most likely scenario predicts that the average number of children born to each woman will stabilize in the medium to long term, without recovering its pre-recession level.

When interpreting these results it is important to keep in mind that, as was mentioned earlier, the model we use in this article attempts to capture the stylized facts of fertility change over a relatively long time horizon, therefore it has to rely on a number of simplifying assumptions. Furthermore, while our model is able to reconstruct observed patterns in key dimensions of family formation dynamics, this does not imply that it is the only model capable of explaining the data. Indeed, some of the questions we leave open can be tackled with more narrowly defined models, or with models that incorporate more complex mechanisms or other sources of heterogeneity. In other words, beyond the value and limitations of the specific model presented here, what we would really like to highlight is the advantages of an individual-level computational approach in a forecasting setting.

Compared to aggregate approaches, individual-level modeling allows for a much richer analysis, as it enables us to consider fertility dynamics as a system of interrelationships between key social processes and the various aspects of the reproductive process. As well as helping us asses the likelihood of future scenarios, this modeling approach can facilitate the communication of the dynamics behind each scenario to non-academic audiences.

In this specific case, the inclusion of complementary information like the deceleration of the process of educational expansion or the stabilization of family size preferences allowed us to obtain a slightly less pessimistic forecast than the majority of the forecasts obtained by methods that rely exclusively on information on the past trajectory of a given fertility indicator.

Although forecasts based on expert opinions tend to implicitly incorporate social dynamics, having an explicit model of these relationships also seems like a clear step forward, as it allows us to compare, combine and assess models based on the accuracy of their predictions.

In sum, thanks to advances in computing technology and infrastructure, as well as advances in computational statistics, demographic forecasting can now leverage the scenario-generating capabilities of simulation modeling within a solid, data-driven statistical framework.

A Appendix

A.1 Desired Family Size

The desired family size D_i of woman i in the simulation is formed immediately after a union is established, and is derived from a Gamma distribution. Based on the assumption that couples in which both partners work will tend to prefer a smaller family, we make the desired family size dependent on the labor force participation (active/inactive) of the woman. Our model for D_i also takes into account the fact that individuals live in a society with a given set of norms, and that these social norms influence their preferences as much as their individual characteristics.

In practice, the desired family size D_i of couple i in the model is obtained by rounding to the nearest integer a number drawn from a truncated Gamma distribution with expected value:

$$E_G = \overline{D}_{t[i]} \cdot \big(1 + (\text{-}1)^{w_i} \cdot (1 - p_{t[i]}) \cdot \delta\big) \cdot \delta' \tag{1}$$

where $t[i]$ is the time at union formation for woman i, and $\overline{D}_{t[i]}$ is the average desired family size for all women of reproductive age at time $t[i]$.

$\overline{D}_{t[i]}$ represents the existing social norm regarding the ideal family size at a given time. The amount by which E_G departs from this norm is given by $(\text{-}1)^{w_i} \cdot (1 - p_{t[i]}) \cdot \delta$, where w_i is a dummy variable that indicates the labor force status of woman i (0 = inactive, 1 = active), $p_{t[i]}$ is the proportion of women in the population who share her labor force participation status, and δ is therefore the maximum proportional departure from the norm. The implication of the $(1 - p_{t[i]})$ term is that the larger the fraction of the population who share her labor force participation status is, the closer to the existing social norm the woman will be. For instance, in the extreme case of a homogeneous population of inactive women, the deviation from the norm is zero if newly-cohabiting woman i is inactive and is equal to δ if she is active.

The first two terms of equation 1 model the decline and stabilization of family size preferences driven by cohort changes in labor force participation. The last term, δ', captures period effects. We distinguish two periods in which we allow the value of δ' to be different than 1: 1975-1995 and 2010-2018. Table 1, at the end of this section, presents the estimated value of δ' along with the rest of the parameters of the model.

A.2 Intentions

While the desired family size D_i determines the attempted final parity, we define an intention $I_{i,t}$, which plays a role in determining the timing and the likelihood of each specific birth. More specifically, the intention $I_{i,t}$ is the probability that a couple will decide to try to have a child in the next year. The strength of the intention depends not only on individual, fixed characteristics, but also on the time elapsed since the previous birth.

For a couple who have achieved their desired family size, $I_{i,t} = 0$. Otherwise, the intention is given by:

$$I_{i,t} = (\beta_i - \omega_i \cdot w_i)(1 - e^{-\lambda \cdot d_{i,t}}) \tag{2}$$

where β_i is a baseline that represents the probability of deciding to have a child in a given year for a non-working woman who has not had a child recently and has not yet achieved her desired family size.

For working women, the intention also depends on education through ω_i, which we define as:

$$\omega_i = \frac{\eta}{1 + \exp(\epsilon \cdot (y_i - \tau))} \tag{3}$$

This formulation implies that a penalty on the intention $I_{i,t}$ exists for women who work, but that this penalty is reduced for those who are more educated. Larger values of τ imply that more years of education are needed to reduce the penalty. A very large value of τ implies there is no positive effect of education on the intention to have a child.

Allowing education to have a neutral to positive effect on the intention to have a child, we account for the mechanisms through which having higher educational attainment eases the decision to have a child for a woman who works and has not yet achieved her desired family size. These include the increased ability to outsource childcare, to reduce economic uncertainty, to negotiate the division of housework, and mobilize personal and familial resources to strike a better balance between work and family.

Finally, the last term in equation 2 models how the intention is affected by $d_{i,t}$, the time elapsed since the last birth. In this case, we assume that there is a strong penalty immediately after childbirth that decays over time.

A.3 Conception Risks

When a couple decides to have a child, the partners' waiting time to conception wt_c is determined by their risk of conception. This implies, as in the real world, that a couple's intention is independent of their actual risk, or ability, to conceive. Following earlier models, we represent the risk of conception through the notion of *fecundability*. In the absence of contraception, fecundability is highest among young couples, and decreases with age as the frequency of intercourse and the biological capacity to conceive decline.

As we model in continuous time, we define fecundability as the instantaneous risk of conception. We assume it constant, and therefore draw the waiting time to conception from an exponential distribution. Specifically, instantaneous fecundability is defined as:

$$f_{i,x} = \frac{\phi}{1 + \exp(\kappa \cdot (x_i - \gamma))} \tag{4}$$

We fix at 0.22 mo^{-1} the value of ϕ, the maximum instantaneous fecundability, so as to obtain a 0.93 probability of conceiving within a year ($= 1 - e^{-12\cdot\phi}$). This corresponds to a monthly fecundability of 0.20 ($= 1 - e^{-\phi}$), a value consistent with what has been reported in the literature for non contracepting, young couples (Bongaarts and Potter, 1983; Leridon, 2004). Equation 4 specifies that instantaneous fecundability decreases with rate κ as the age x_i of woman i starts to approach age γ. The values for κ and γ were calibrated to fit the pattern of the evolution of fecundability with age proposed by Coale and Trussell (1974).

If the partners decide not to have a child in the next 12 months, they will be at risk of an unplanned conception during this period due to *residual fecundability*, defined as:

$$rf_{i,x,t} = f_{i,x} \cdot c_{i,t} \cdot a_i \tag{5}$$

where $c_{i,t} \in [0, 1]$ represents the fecundability-reducing effect (effectiveness) of the contraceptive methods available to woman i at time t. $c_{i,t} = 0$ represents perfect effectiveness while $c_{i,t} = 1$ means complete ineffectiveness. $c_{i,t}$ is defined by:

$$c_{i,t} = \frac{(\rho/y_i)^\alpha}{1 + e^{r\cdot(t-(\psi-y_i))}} + \upsilon \tag{6}$$

For a given educational attainment level (fixed y_i), the evolution of $c_{i,t}$ through calendar time is simply a decreasing logistic function with maximum value $(\rho/y_i)^\alpha$, inflection point $\psi - y_i$ and steepness r to which is added a constant υ, the effectiveness of the best contraceptive methods. This definition reflects the notion that while the effectiveness of contraceptive methods improves over time for all women, better educated women more readily adopt efficient contraceptive methods.

Finally, the last term in equation 5, a_i is equal to one if couple i have not yet achieved their desired family size; and fixed value a if the couple have already achieved their desired family size. The assumption here is that the partners will intensify their efforts to prevent additional births after they have reached their goal.

If the waiting time until conception (planned or unplanned) ${}_CT$ is shorter than 12 months, a waiting time until birth ${}_BT$ is created by adding 270 days of gestation to ${}_CT$. Following the birth, there will be a period of six months that corresponds to the period of postpartum subfecundity, abstinence, and increased contraceptive use with the aim of avoiding dangerously close births. After this period of time, the couple will go back to being exposed to the risk of a new conception.

If no conception occurs within the next 12 months, our couple will update their intentions again at the end of that period. Thus, the partners will either decide to try to have a child in the next year, or they will be subject to the risk of an unplanned birth.

A.4 Age at Union Formation

To keep things simple, the preferences, intentions, and risks that control the process described so far are assumed to operate only after the formation of a cohabitation union. The age at which this event occurs is therefore of considerable importance. We follow previous studies that have successfully approached the empirical distribution in the age at marriage using the log-normal distribution (Mode, 1985).

$$\ln(M_i) \sim \mathcal{N}(\mu_i, \sigma^2) \tag{7}$$

with $\mu_i = y_i + (\xi \cdot \xi^{'})$, where y_i is the numbers of years of education of woman i, and ξ is the average waiting time (in years) to the formation of a cohabiting union after the end of the schooling period. $\xi^{'}$ is equal to 1 except for the birth cohorts born 1990 to 2000, those reaching adulthood at the beginning of the 2008 global recession.

A.5 Death

The waiting time until death is sampled using the inverse distribution function method (for a description of the method see: Willekens, 2009), where the distribution of waiting times until death is reconstructed using age-specific cohort mortality rates available at (HMD, 2015).

A.6 Estimated Parameter values

Table 1 shows the combination of parameters values that generates simulated output that that is closer to our data in terms of mean squared error.

Name	Domain	Eq.	Description	Value	Source
δ	desired fam. size	1	effect on D_i of working	0.59	estimated
δ'	desired fam. size	1	period effects on D_i (2010 - 2018)	0.81	estimated
β	intention	2	baseline intention	0.8	fixed
λ	intention	2	rate of reduction penalty after pregnancy	$2.5\text{e}{-8}$ s^{-1}	fixed
η	intention	3	penalty on intention of working	0.54	estimated
τ	intention	3	years of edu. after which η is reduced	13.6 yr	estimated
ϵ	intention	3	speed at which η is reduced	0.5 yr^{-1}	fixed
ϕ	fecundability	4	maximum fecundability	0.22 mo^{-1}	fixed
γ	fecundability	4	inflection point of decline (age)	38.5 yr	fixed
κ	fecundability	4	speed of decline with age	0.6 yr^{-1}	fixed
A	contraception	5	additional effect due to achieved D	0.07	fixed
ρ	contraception	6	minimum effect on conception risk	0.11	estimated
υ	contraception	6	maximum effect on conception risk	0.66	estimated
ψ	contraception	6	inflection year in the diffusion process	1973.3	estimated
r	contraception	6	speed of the diffusion	1.2 yr^{-1}	estimated
α	contraception	6	differential access/use by education	0.15	estimated
ξ	age union form.	7	time to union after schooling	4.8 yr	estimated
ξ'	age union form.	7	period effects on age at union	1.5	estimated

Table 1 | Model parameters and best-fitting values to data from Finland. Available estimates used in the case of ϕ, γ and κ. The rest of the fixed parameters were obtained by calibration. The baseline intention β was set a value lower than 1 to consider other effects not explicitly modeled that might impose a penalty on the intention.

References

Beaumont, M. A. (2010). Approximate bayesian computation in evolution and ecology. *Annual review of ecology, evolution, and systematics 41*, 379–406.

Beaumont, M. A. (2019). Approximate bayesian computation. *Annual review of statistics and its application 6*, 379–403.

Bohk-Ewald, C., P. Li, and M. Myrskylä (2018). Forecast accuracy hardly improves with method complexity when completing cohort fertility. *Proceedings of the National Academy of Sciences 115*(37), 9187–9192.

Bongaarts, J. and R. E. Potter (1983). *Fertility, biology, and behavior: An analysis of the proximate determinants*. Academic Press.

Ciganda, D. and N. Todd (2021). Demographic models of the reproductive process: Past, interlude, and future. *Population Studies*.

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.

Dancik, G. M. and K. S. Dorman (2008). mlegp: statistical analysis for computer models of biological systems using r. *Bioinformatics 24*(17), 1966–1967.

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.

Gutmann, M. U. and J. Corander (2016). Bayesian optimization for likelihood-free inference of simulator-based statistical models. *The Journal of Machine Learning Research 17*(1), 4256–4302.

Hellstrand, J., J. Nisén, V. Miranda, P. Fallesen, L. Dommermuth, and M. Myrskylä (2021). Not just later, but fewer: Novel trends in cohort fertility in the nordic countries. *Demography 58*(4), 1373–1399.

Hellstrand, J., J. Nisén, and M. Myrskylä (2020). All-time low period fertility in finland: Demographic drivers, tempo effects, and cohort implications. *Population studies 74*(3), 315–329.

Hellstrand, J., J. Nisén, and M. Myrskylä (2022). Educational field, economic uncertainty, and fertility decline in finland in 2010–2019.

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.

Hiilamo, H. (2020). Why fertility has been declining in finland after the global recession?: A theoretical approach. *Finnish Yearbook of Population Research*.

HMD (2015). Human Mortality Database. University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany). Available at www.mortality.org or www.humanmortality.de (data downloaded on May 2015).

Human Fertility Database, . (2011). Human fertility database.

Leridon, H. (2004). Can assisted reproduction technology compensate for the natural decline in fertility with age? a model assessment. *Human Reproduction 19*(7), 1548–1553.

Lutz, W., V. Skirbekk, and M. R. Testa (2006). The low-fertility trap hypothesis: Forces that may lead to further postponement and fewer births in Europe. *Vienna yearbook of population research*, 167–192.

Matysiak, A., T. Sobotka, and D. Vignoli (2021). The great recession and fertility in europe: A sub-national analysis. *European Journal of Population 37*(1), 29–64.

Mode, C. J. (1985). *Stochastic processes in demography and their computer implementation*, Volume 14. Springer Science & Business Media.

Myrskylä, M., J. R. Goldstein, and Y.-h. A. Cheng (2013). New cohort fertility forecasts for the developed world: Rises, falls, and reversals. *Population and Development Review 39*(1), 31–56.

Schmertmann, C., E. Zagheni, J. R. Goldstein, and M. Myrskylä (2014). Bayesian forecasting of cohort fertility. *Journal of the American Statistical Association 109*(506), 500–513.

United Nations, D. o. E. and P. D. Social Affairs (2022). Model-based Estimates and Projections of Family Planning Indicators. Technical report, United Nations.

Vignoli, D., G. Bazzani, R. Guetto, A. Minello, and E. Pirani (2020). Uncertainty and narratives of the future: A theoretical framework for contemporary fertility. In *Analyzing contemporary fertility*, pp. 25–47. Springer.

Willekens, F. (2009). Continuous-time microsimulation in longitudinal analysis. *New frontiers in microsimulation modelling*, 353–376.

Zeigler, B. P., T. G. Kim, and H. Praehofer (2000). *Theory of modeling and simulation*. Academic press.