The Limits to Fertility Recuperation

Daniel Ciganda, Nicolas Todd

Max Planck institute for Demographic Research

Abstract

The recuperation of period fertility rates in the first decade of the 21st century triggered a reinterpretation of the link between development and fertility. Insight into how advances in development can initially depress and later increase fertility rates have been searched primarily on the sociocultural determinants of the demand for children. But the story of post-baby boom fertility is also the story of the contraceptive transition and the move towards increased control over the reproductive process. We argue here that taking into account this dimension opens up a new interpretation of the link between development and fertility. In the explanation we propose, advances in development can, in a first stage, lead to fertility decline by reducing fertility due to poor contraception; once that first stage is completed, however, further advances in development can lead to a recuperation of fertility by reducing the obstacles couples face when trying to achieve their desired family size. We test this hypothesis by using an individual-level computational model of the reproductive process from which we simulate the reproductive trajectories of post-baby boom cohorts and forecast the potential development of aggregate fertility indicators. We estimate the parameters of our model with the aid of Gaussian process emulator using information from France, Ireland and Spain, three countries with widely different fertility trajectories. We find that, even in the presence of a strong positive effect of increasing educational attainment on the demand for children, there are limits to the recuperation scenario imposed by further reductions in exposure and the end of the transition to higher education.

1 Introduction

The main concern of most theories of fertility change developed throughout the 20th century was understanding the determinants of fertility decline. An important exception was the framework developed by Richard Easterlin, who tried to explain the *baby boom* after the Second World War and the subsequent *baby bust* as a result of a process of cyclical change in the relative socioeconomic standing of successive generations related to variations in cohort size (Easterlin, 1978). However, the baby boom *echo* predicted by Easterlin was never observed and his hypothesis largely abandoned. In fact, the years that followed the baby boom were characterized by further declines in period fertility rates. The fact that a booming economy like West Germany was first to reach low-fertility levels in the early 1980s contributed to re-strengthen the narratives of the modernization (or post-modernization) kind, which associated economic and social development to reducing fertility rates (Caldwell and Schindlmayr, 2003).

Perceptions of the relationship between fertility and development started to change in the 1990s, when the Total Fertility Rate (TFR) stabilized in countries like France, the Netherlands or the United Kingdom and even increased slightly in Scandinavian countries while it was reaching low and lowest-low levels in the south of Europe. Since then, a number of studies have documented the *reversal* of the long-term negative correlation between fertility rates and a series of social-and-economic development indicators at the macro level, such as female labor force participation (Rindfuss and Brewster, 1996; Ahn and Mira, 2002), Human Development Index (Myrskylä et al., 2009) or GDP per-capita (Luci-Greulich and Thévenon, 2014).

The idea that advances in development can contribute to increase fertility levels received additional support after the sustained recuperation of period fertility rates that consolidated around the 2000s in most European countries (Goldstein et al., 2009). Since then, modernization frameworks have been losing ground in favor of gender equality narratives which predict a return to *more family* as the new dual-earner family becomes the norm (Goldscheider et al., 2015; Esping-Andersen and Billari, 2015).

The recuperation trend also prompted an upward revision of the level at which the TFR was supposed to converge according to the United Nation's population projections (United Nations, 2010). In fact, the assumption that fertility will eventually recover to near replacement levels informed official projections for decades and was only recently relaxed in response to criticism based on the fertility trajectories of some high-income countries in East Asia (Raftery et al., 2014). Other critics have argued that a significant part of the recovery can be attributed to the diminishing pace of the postponement transition (Bongaarts and Sobotka, 2012).

Both critics are well-founded, the recuperation of TFR is not universal and it can be partially explained by timing changes, but even if this is the case, the question remains of whether this trend could *also* be indicating a real recovery process in some countries and to what extent development can contribute to increase fertility rates in the near future.

The answer to this question, we believe, lies in the linked dynamics of the three transitions that have driven post baby boom fertility in high income countries: The transition from the bread-winner to the dual earner model, the transition to higher education and the contraceptive transition.

Figure 1 illustrates the extent of the changes experienced by post-baby boom cohorts in these three domains. In parallel to a rapid increase in the proportion of women with tertiary education and the proportion of women in the labor market, there has been an equally dramatic decrease in the proportion of unplanned births. It is helpful to think of these transformations as *transitions* in the demographic sense, i.e. a process of change that represents the passage between two different regimes (or equilibriums). This notion highlights not only the magnitude of the changes involved, but the fact that these changes will eventually come to and end.

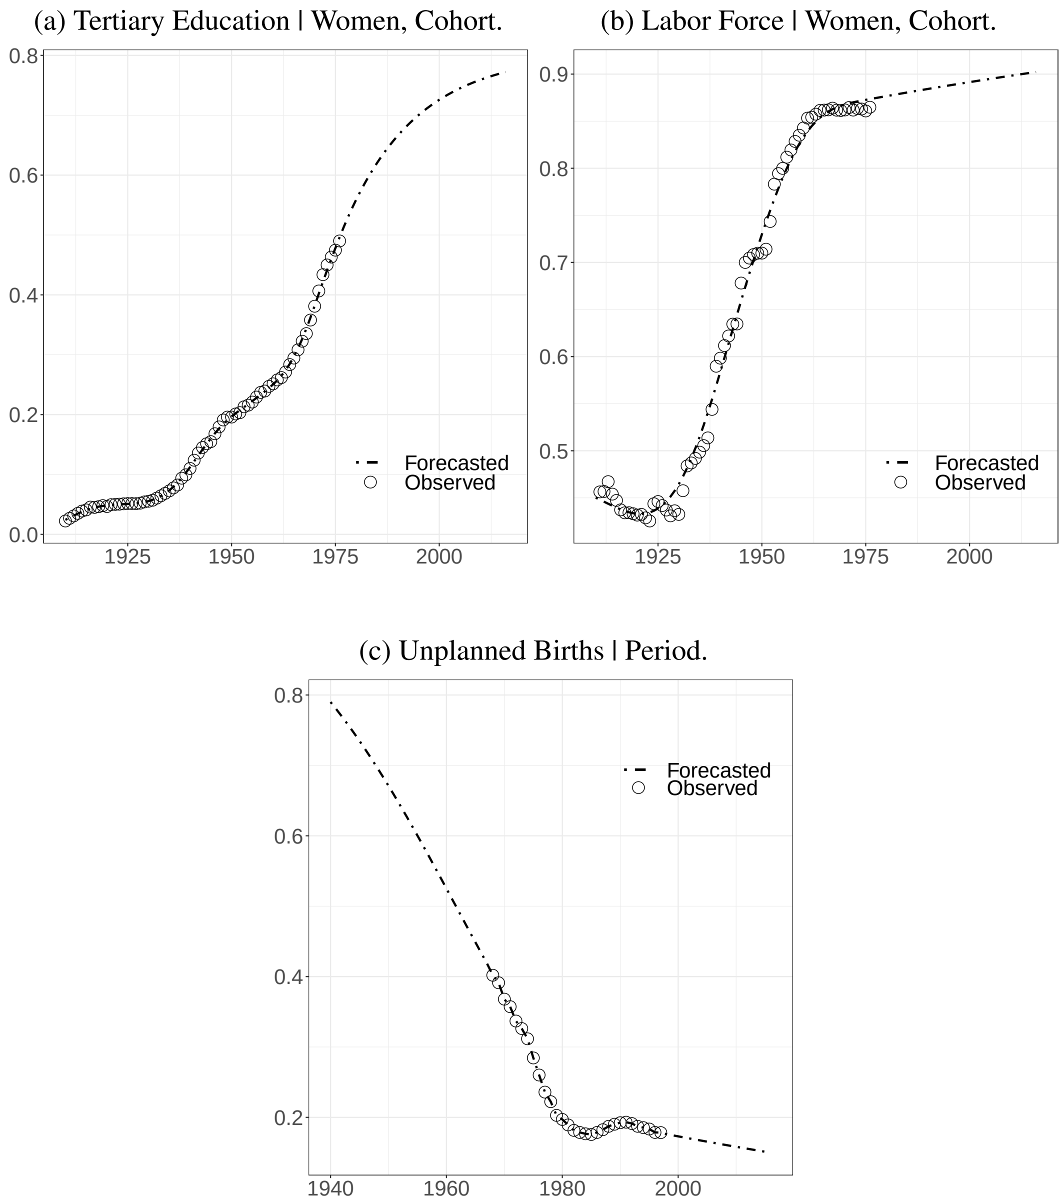

Figure 1: Transitions in Education, Contraception, Labor Force Participation | France.

Most analyses of post baby-boom fertility have focused exclusively on the transitions in the education and labor force participation domains. As a consequence, post baby-boom fertility has been perceived as a process entirely determined by individual choices.

The dismissal of the contraceptive transition is perhaps related to the fact that the methodology for measuring unplanned fertility, and the corresponding data collection efforts, were developed out of concerns for unmet family planning needs in developing countries Klijzing (2000), which also explains why the overwhelming majority of demographic analyses of unplanned fertility have focused on the regions covered by the Demographic and Health Surveys Program (Bongaarts, 1997; Casterline and El-Zeini, 2007, among others).

The few studies available in Europe, however, show that the role played by the reduction of unplanned births during the final stage of the fertility transition in Europe is larger than what is usually believed. Murphy (1993), estimated the impact of the introduction of the contraceptive pill in the fertility decline of the 1960s and 1970s in Britain, and claimed that most, if not all, of the change in fertility rates would have still occurred in the absence of the economic changes that were observed during this period. Similarly Régnier-Loilier et al. (2007) attributed half of the decline of fertility in the 1970s in France to the drop in the proportion of unplanned births and showed that this proportion has remained constant at about 20% since the mid 1980s.

These findings suggest that many fertility analysis, specially those including pre-modern-contraception birth cohorts, will benefit from thinking about fertility not as a single outcome, but as an aggregate of two outcomes: one that results from an *intentional* process (planned fertility) and one that results from an *unintentional* process (unplanned fertility).

Once this perspective is adopted is easier to see how education can have counteracting effects on fertility outcomes. While the expansion of education can depress fertility through the reduction of the risk of unplanned pregnancies or the postponement of the family formation process, it can also have a positive influence on *planned* fertility by increasing the resources that contribute to reduce the gap between desired and achieved family size.

Research on fertility intentions has indeed shown how highly educated women tend to have higher fertility intentions than women with lower education (Heiland et al., 2005; Testa, 2014; Berrington and Pattaro, 2014; De Wachter and Neels, 2011). This result has been associated with the increased bargaining power of educated women, in addition to an increased probability to meet partners who are willing to assume a larger share of unpaid work (Mills et al., 2008; Mencarini and Tanturri, 2004; Sullivan et al., 2014; Esping-Andersen, 2009). A higher level of education also usually implies an increased ability to outsource childcare activities, reduce economic uncertainty and mobilize personal and familial resources to strike a better balance between work and family.

We argue that the non-linear nature of the effects of education on fertility outcomes and the *transient* nature of the changes that drove fertility change throughout the second half of the 20$^{\text{th}}$ century are the keys to understand both the decline and recuperation of post baby-boom fertility.

More specifically, we believe that the transition to higher education, the transition to a dual-earner model and the transition to a controlled reproductive process have triggered a series of positive and negative influences on fertility outcomes. These counteracting influences have been competing with each other with varying relative strengths as the transitions moved forward. In a first, pre-modern contraception stage, the expansion of higher education contributes mostly to fertility decline, by reducing unplanned fertility, but also by promoting the participation of women in the labor market. During this first stage, the negative effects *outpace* any potential gains, but once the contraceptive transition and the transition into the labor market have been completed, the continuing expansion of education might result in the *recuperation* of fertility rates as individuals are able to mobilize an increased number of resources in order to achieve their desired family size.

This hypothesis is appealing for a number of reasons: it provides the basis for modeling the decline *and* recuperation of fertility in one unified framework; it also provides a solution to the puzzle presented by the positive educational gradient of fertility intentions while still matching the cross sectional relationship between fertility and education/resources; it takes into account both the sociocultural and biological dimensions of fertility; and, crucially, it provides insights into the *future* of fertility rates.

Given the complex nature of the mechanisms involved, we tested our hypothesis using a computational, or *algorithmic*, approach to model the reproductive process. We also adopt an individual-level approach, which allow us to derive aggregate fertility patterns from simulated reproductive trajectories, establishing a direct link between the proposed mechanisms at the individual-level and the emerging macro trends.

We fit our model to observed age-specific fertility rates from the 1960s to the present and to data on fertility goals and unintended births (when available) in three countries with widely different fertility transitions, France, Ireland and Spain.

In addition to the information obtained during the estimation of the model, we obtain insight into our question by exploring a number of scenarios which include forecasting aggregate fertility trends for the next twenty years.

The remainder of the article is structured as follows: the next section provides description of the operation of the model, followed by an explanation of the procedure used for estimation. The results and conclusions, presented in Section 5 and 6 can be fully replicated by accessing the code and data available <link to GitHub repository>.

2 Model Description

The model we introduce here can be described in a number of different ways: With respect to the formalism used, it is a *discrete event simulation*. The level at which quantities are primarily modeled also make it an *individual-level* model, and if we look at how transitions between states are modeled it is part *microsimulation*, part *agent-based model*. It is also a *computational* or *numerical* model with respect to the way in which output is obtained.

Additionally, it is probably also useful to think about the model through the lens of *natural algorithms* as proposed by Chazelle (2012).

As in other discrete event simulations time advances with the successive realization of events (Zeigler et al., 2000). The events that make the clock run in our model are the following: starting a cohabiting union, deciding to have a child, trying to have a child, having a child (an additional child), dying.

The pseudo-code presented in Appendix A details the state changes related to the experience of each of the events listed before. At each iteration the algorithm finds the event with the shortest waiting time from a list of all possible events for the entire population of simulated individuals. After the realization of the next event, the state of the system and the clock are updated and the simulation can continue to the next run. The output obtained is a set of individual trajectories from which any standard fertility measures can be computed.

The model requires only three inputs: the distribution of educational attainment, the distribution of labor force participation by educational attainment and the set of observed mortality rates; all by birth cohort. This implies that the level of education of each woman born in the population, whether she will be working or not and her time of death are determined by the observed distributions corresponding to her birth cohort.

We represent the reproductive process as starting with *union formation* and being driven by a *desired family size*. The likelihood and timing of births are determined by an individual's *intention*, but also by the *risk of an unplanned birth* they are exposed to. We detailed each of these components in the following sections.

2.1 Desired

Our definition of desired family size attempts to capture the movement described by most analyses of post baby-boom fertility in which the reduction of the demand for children is generally interpreted as a consequence of the erosion of the male bread-winner paradigm following the rapid increase of women's participation in the labor market (Murphy, 1993). Additionally, we also try to capture the influence of existing *social norms* at the time couples form their ideas about the appropriate size of a family.

In the model, the desired family size D_i of woman i is formed immediately after union formation and it is given by a truncated Gamma distribution:

$$D_i \sim \Gamma_{a \leq D_i \leq b}(\alpha, \beta_i) \tag{1}$$

where a and b are [1] the minimum and maximum number of desired children. The shape

[1]The choice of Gamma was based on computations we performed on data from the Integrated Value Sur-

parameter β_i is given by:

$$\beta_i = \overline{D}_{t[i]} \cdot (1 \pm \left(\frac{\delta}{1 + e^{\pm r \cdot (\overline{D}_{t[i]} - D^*)}}\right) \cdot p_{t[i]})/\alpha \tag{2}$$

where $t[i]$ is the time the union woman i is in was formed and $\overline{D}_{t[i]}$ is the average desired family size for the most recent cohorts of married women (25 years or below at time $t[i]$), which reflects the existing norm regarding the ideal family size. Couples consider this ideal when forming their preferences. They will tend to prefer a smaller family size than the existing norm in case the woman in the couple works, and a larger family size in case she does not. The amount of adjustment with respect to the norm is given by the second term in the equation. The maximum amount of adjustment is given by δ. Our dual-earner couple subtracts δ from $\overline{D}_{t[i]}$ if $\overline{D}_{t[i]}$ is high. The value of the correction, however, will be progressively reduced if $\overline{D}$ approaches D^*. This implies that there is less resistance to move from an ideal of six to an ideal of five than to move from an ideal of two to an ideal of one. This term is also weighted by $p_{t[i]}$, which represents the proportion of women in the population which do not share the labor force participation status of woman i. This implies that when most of population shares her characteristics she will take the existing norm as it is.

2.2 Intended

If the desired family size D_i is fixed at union formation and determines the attempted final parity, the intention I_i varies throughout the life course and determines the timing and likelihood of each birth. The *strength* of the intention depends on individual characteristics (labor force participation and education) but also on the time elapsed since the previous birth. At time t, the intention of woman i is given by:

$$I_{i,t} = \rho_i - e^{-\lambda \cdot d_{i,t}} - \omega_i \tag{3}$$

where ρ_i is a baseline which represents the intention, free of constraints, to have a child if so desired by the agent ($D_i > 0$).

The second term models how the intention is affected by $d_{i,t}$, the time elapsed (duration) since the last birth. Parameter λ is to be estimated. This formulation implies a penalty on the intention to have a child immediately after the experience of a birth, which decays exponentially as $d_{i,t}$ increases.

veys (not shown), which confirmed our expectation that the distribution of family size preferences is positively skewed, with mode in the 1-3 range, depending on country and period, and a right tail formed by a decreasing number of responses associated with larger family sizes

The last term in 3 expresses the positive effect of educational attainment on the intention to have a child among working women. It is defined as:

$$\omega_i = \eta/(1 + exp(\epsilon \cdot (y_i - \tau))) \tag{4}$$

which describes a logistic function that goes from η to 0 as the number of years of education y_i approaches the threshold τ. This formulation implies that there is a penalty on the intention when women work. At the aggregate level, the increased education attainment of working women will translate into a reduced penalty on fertility rates.

To connect intentions to events we assume that waiting times to the *attempt*[2] to conceive are exponentially distributed with a rate proportional to the intention I_i, therefore associating higher intentions to shorter waiting times.

At union formation, we draw a waiting time to the attempt to conceive for those couples that desire to do so. If the resulting waiting time is shorter than a one year period, at the end of that period our simulated couple will attempt to have a child. If it is longer, they will update their intention after one year and reconsider whether or not to have a child within the next twelve months.

If the couple fails to conceive, they will try again in a period of one to three months, until they reach a predefined limit for failed attempts. If they succeed, a waiting time equal 270 days is created which represents the duration of the pregnancy. Following the birth, there will be a period of 6 months that corresponds to the period of postpartum amenorrhea after which, and in case that the couple has not still reached D, they will go back to updating their intention/evaluating whether to try and have a new child.

2.3 Unplanned

As described earlier, fertility outcomes are often depicted as resulting exclusively from a controlled process, especially in the context of high income countries. Correctly accounting for the forces that escape human control is key, however, to understand fertility change, in particular when we are interested in changes that span over many years and include cohorts that went through their reproductive years before the contraceptive transition was completed.

Here we represent the reproductive process as composed of an *intentional* sub-process, governed by preferences and intentions as described above, and a *unintentional* sub-process, governed by the risk of experiencing an unplanned birth, which we define as:

[2]We define the event as the *attempt* to conceive in order to be able to represent the decline in the probability to experience a pregnancy with age. The timing and strength of the decline are modeled following the approximations obtained in Leridon (2004)

$$U_{i,t} = \frac{\phi / y_i^{\kappa}}{1 + e^{\theta \cdot (t - (\psi - y_i))}} + \upsilon \tag{5}$$

The evolution of the risk of having an unplanned births is closely related to the adoption of contraception, which is usually depicted as a diffusion process initially led by more educated women, but progressively reaching women and couples across strata (Rosero-Bixby and Casterline, 1993). In our model the inflection point of this process is represented by ψ, while θ controls its speed, and ϕ and υ its upper and lower bound respectively. We also introduce the years of education y in the formula to represent the idea that adoption is faster among more educated women; the extent of the differences among educational attainment groups if controlled by κ.

2.4 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 great importance. We follow previous work that have successfully approached empirical distribution in the age at marriage using the Log-normal distribution (Mode, 1985).

$$\ln(A_i) \sim \mathcal{N}(\mu_i, \sigma^2) \tag{6}$$

with $\mu_i = y_i + \xi$, where y_i is the numbers of years of education of woman i, and parameter ξ is the average waiting time (in years) to the formation of a cohabiting union after the end of the schooling period.

2.5 Death

The waiting time to death is sampled using the inverse distribution function method (for a description of the method see: Willekens, 2009) where the distribution of waiting times to death is reconstructed using age-specific cohort mortality rates from the Human Mortality Database. For the ages/years where information is not available we use the latest available figures. Missing data corresponds essentially to the most recent decades, when mortality of females under 50 years of age is very low.

2.6 Example Trajectory

After having described the central elements of our model we provide now a basic description of its operation.

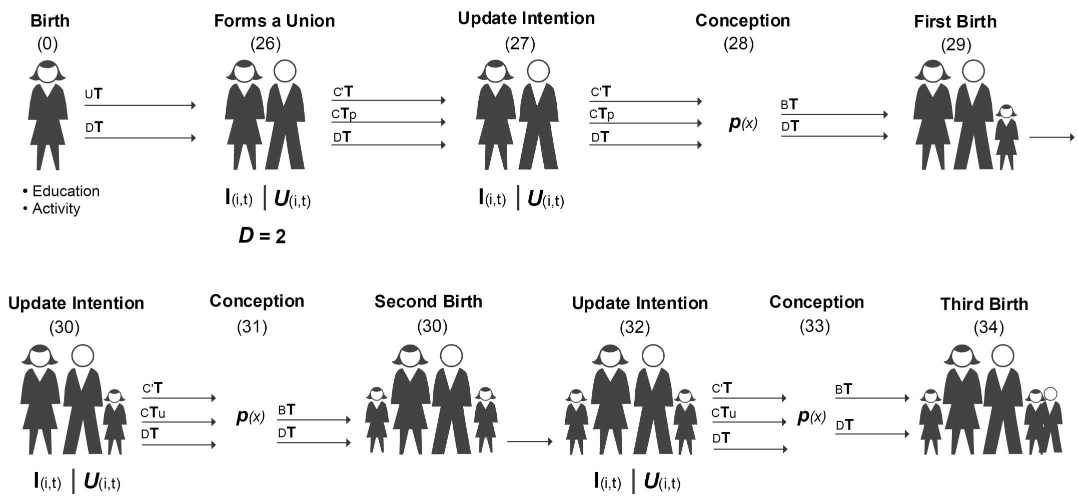

Figure 3: A Simulated Trajectory

Figure 3 depicts a hypothetical life trajectory of a woman that is born during the simulation. Based on the observed cohort trends discussed in section 2, at birth we know the level of education that she will achieve, the number of years she will be enrolled in education and whether she will work or not. A waiting time to the formation of a union ($_UT$) is computed at this moment using equation 6, which competes with her waiting time to death ($_DT$). The event with the shortest waiting time is realized, in our example, she forms a union at age 28. Two waiting times are computed at this point: $_CT_p$, the waiting time to the attempt of a planned conception, and $_CT_u$, the waiting time to the risk of an unplanned conception. The shortest of these waiting times competes both with death and $_{C'}T$, the waiting time to the update of the intention and has a fixed length of twelve months.

In our example $_{C'}T < {_C}T$, which means the next event registered is the update of the intention. At this moment $_CT_p$ and $_CT_u$ are recomputed and the shortest one competes again with $_{C'}T$. This couple decides to conceive when the woman is age 28. Their probability to succeed, p depends on her age. In this case the attempt leads to a waiting time to birth (pregnancy) $_BT$ and a first birth when she is age 29. After a period of six months they go back to being exposed to the risk of a pregnancy, which in this case leads to an unplanned birth when she is 30 years of age. At age 34 they have a third child, which is both unplanned (determined by U instead of I) and unwanted, as the couple had already achieved their desired family size.

3 Data & Tools

Data on educational attainment and labor force participation for birth cohorts until 1980 comes from pooled the available census data at IPUMS (Minnesota Population Center, 2018).

To obtain the labor market participation figures we selected women aged 30 to 55 by birth cohort in the pooled set and obtain the proportions of these women that were active (employed + unemployed) and inactive, observations of the same women in more than one census were treated as independent.

Data for the birth cohorts after 1980 is obtained by forecasting from the original IPUMS series using an Exponential Smoothing State Space model as implemented in the R *Forecast* package (Hyndman and Khandakar, 2008). The results of this procedure can be found in Appendix B.

For the estimation of the model's parameters (see Section 4) we rely on data on Age-Specific Fertility Rates, the evolution of the proportion of unplanned/unwanted births and the evolution of the average desired family size.

The Age-Specific Fertility Rates come from the Human Fertility Database (2011), as well as most of the other fertility indicators used in the article (Total Fertility Rates, Cohort Completed Fertility, mean ages at birth and at different parities).

Data on the proportions of unplanned/unwanted births was published by Régnier-Loilier et al. (2007) based on information from INED's 1988 Birth Control Survey (ERN), 1994 Family Situations and Employment Survey (ESFE) and the Fertility Intentions Survey (1998 - 2003).

The information on average desired family sizes was obtained from Leridon and Toulemon (1999) who computed these averages from information on several fertility surveys conducted by INED since the 1950s. In fact, these computations are based on a question about the ideal number of children *en situation*, i.e. the ideal number after considering the respondent's available resources, which much better approximates our concept of desired family size than question regarding the general ideal family size. In the case of Ireland and Spain, however, where we do not have information on the desired family size we approximate it with the distribution of ideal family sizes (see Appendix C) found in the supplementary material provided by Sobotka and Beaujouan (2014).

The model is implemented in R (R Core Team, 2015). For the stages when the demand for computing resources were high we rely on the MPCDF and the GWDG computing clusters. The data and scripts needed to replicate the results obtained here, including the scripts to perform computations in the cluster, are available <link to GitHub repository>.

4 Estimation Procedure

Our strategy consist in searching for the combination of parameter values that reduces the distance between the model output and a set of data points that describe the evolution of the level and timing of fertility (Age-Specific Fertility Rates); the change in the relative weight of the intentional/unintentional forces in the reproductive process (proportions of unplanned and unwanted births); the evolution of social norms regarding family size (average desired number of children).

We measure the distance between observed and simulated data by computing their Mean Squared Error (MSE) assigning equal weight to each observed data point. Due to the high computational cost of mapping the parameter space to the MSE via a full factorial design, we perform this mapping with the aid of a Gaussian Process emulator (for the use of this approach in the context of demographic simulations see Hilton and Bijak, 2017).

We train our emulator using a Latin Hypercube Sample of the parameter space and asses its accuracy with a test set of design points outside the training sample.

Given that the data needed to identify all parameter's values is not available for all countries, we proceed by stages. We start by estimating our model to the French case, which provides all the necessary data. In this first stage we estimate the 12 parameters listed in the first columns of Table 1.

Figure 4 shows a strong of correlation between the predictions of the meta-model (emulator) and the true MSE associated with each combination of parameter values in the training and test sets. This indicates the emulator provides a very accurate representation of the underling model structure.

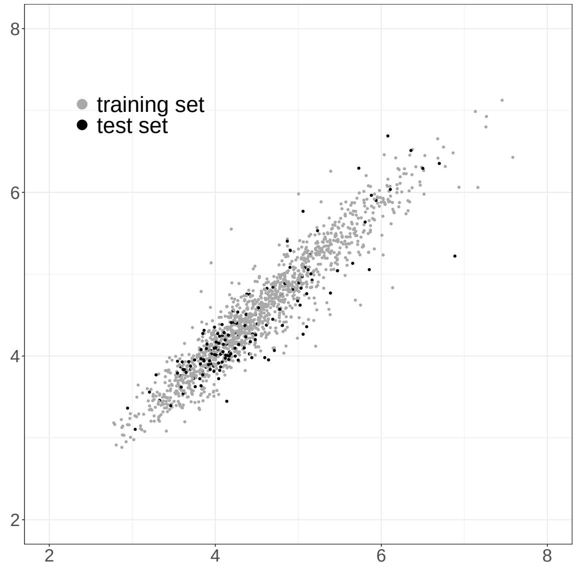

Figure 4: Real vs. Predicted Mean Squared Error | Training + Holdout Set.

The next step consist in finding the global minima in the hyperplane created by the Gaussian Process which gives the combination of parameters with better fit to the data. This combination is reported in column of 5 of Table 1 and provides the fit presented in Figure 5.

5 Results

Table 1: Model Parameters

Name	Domain	Eq.	Description	Value		
				FR	ES	IR
$\overline{D}_{t0}$	desired fam. size	2	initial value of $\overline{D}$	2.8	2.9	4.5
D^*	desired fam. size	2	inflection point resistance D	2.73	1	2.65
δ	desired fam. size	2	effect on D of activity/inactivity	0.09	0.08	0.21
λ	intention	3	effect of duration since previous birth	7.8e-08	7.8e-08	7.8e-08
η	intention	4	penalty on intention when working	0.63	0.79	0.59
τ	intention	4	years of edu. after which η is reduced	19	18.5	16.74
ϵ	intention	4	speed at which η is reduced	0.52	0.54	0.47
ϕ	unplanned births	5	max value risk of unplanned births	0.68	0.79	0.78
υ	unplanned births	5	min value risk of unplanned births	0.08	0.08	0.08
ψ	unplanned births	5	inflection year contraceptive transition	1974	1979	1979
θ	unplanned births	5	speed of the contraceptive transition	0.27	0.18	0.23
ξ	age union form.	6	time to union after schooling	4.2	6	5.9

The model very closely traces the evolution of the age schedule of period fertility rates (Figures 6a to 6d), which over time moves downward and shifts to the left, indicating the decline and postponement of fertility. The contraceptive transition seems to be also accurately captured (Figures 6e and 6f). Both proportions (of unplanned and unwanted births) experience similar trajectories marked by a strong reduction until the mid 1980s and then stabilization.

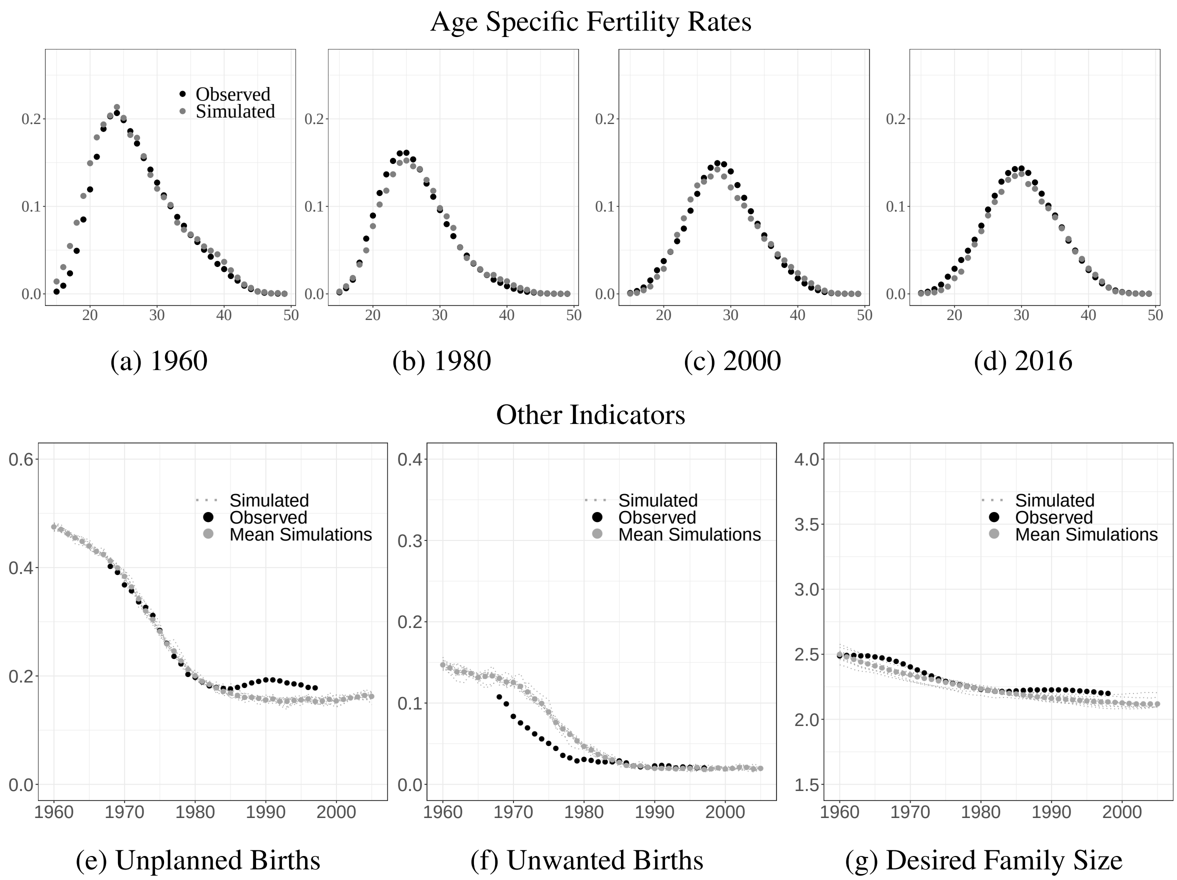

Figure 5: Observed and Simulated Fertility Indicators | France.

The degree of correspondence between observed and simulated data is also good in the case of the average desired family size (Figure 6g), both follow a similar slowly declining and finally stabilizing trajectory.

In the case of Spain and Ireland, where we can only rely on Age Specific Fertility Rates, we fix the values of a set of parameters through calibration as explained in Appendix C and estimate the remaining set. The obtained fit is similar than the previously observed for France until the late 1990s, when the simulated schedule more clearly departs from the observed one. In both countries this issue is related to an early bulge in the schedule, which is observed in some European countries with large immigration flows (Burkimsher, 2017).

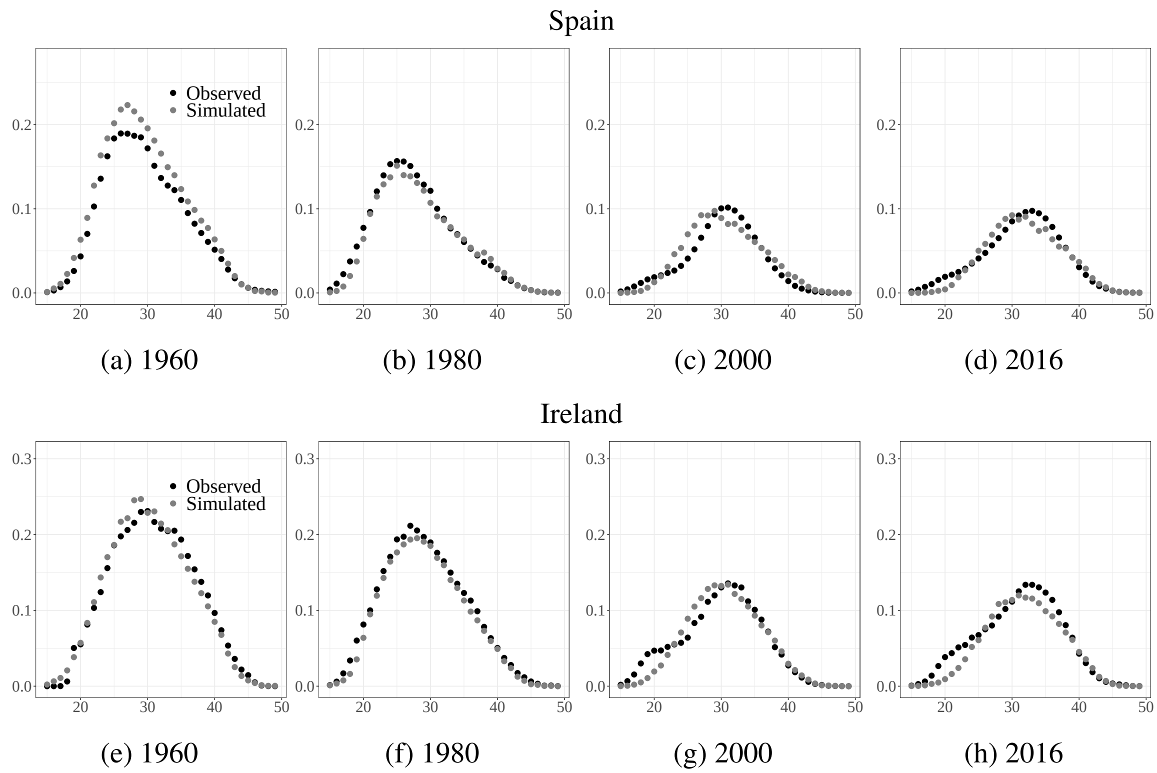

Figure 7: Observed and Simulated Fertility Indicators | Ireland, Spain.

It is also interesting to look at the model fit with respect to the Total Fertility Rate[3](Figure 9). While the model captures systematic variation in the timing of fertility associated with cohort processes, the TFR it produces is unaffected by period effects, i.e. shorter term influences associated for example with the economic cycle. This is why there are visible discrepancies at certain periods between model and data, these periods coincide with marked timing changes associated in most cases with exceptional situations in the labor market[4].

[3]The construction of the TFR rests on the notion of a *synthetic cohort* (treating women of different ages in a population at a given calendar time as if they would belong to a birth cohort), which allows to express the information on the fertility age-schedule in terms of the level of fertility achieved by this cohort. The approximation to the level of fertility of a real cohort is imperfect when the timing of fertility is changing, this is know in the literature as *tempo effects*.

[4]This dynamics are probably better understood considering also the evolution of the Mean Age at Birth (see Appendix D)

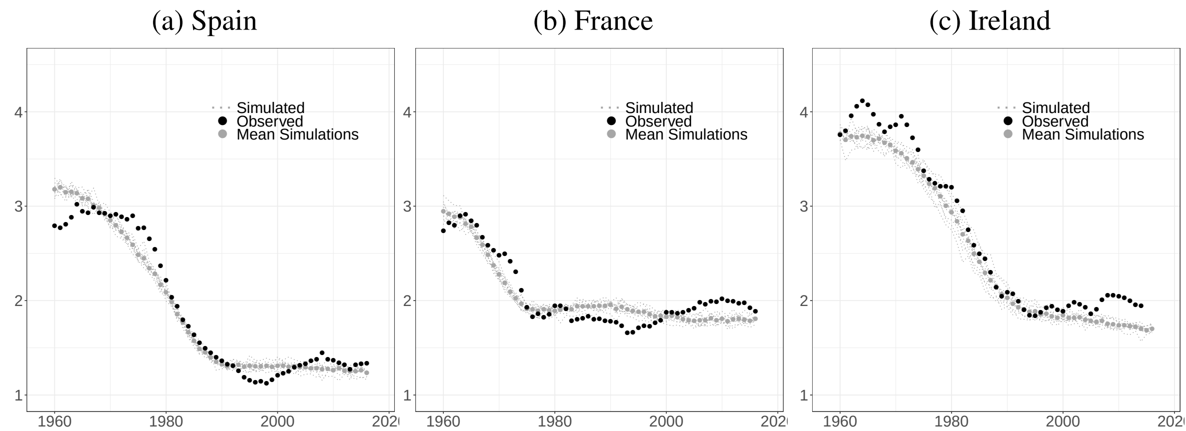

Figure 9: Observed vs Simulated Total Fertility Rate

5.1 Indirect Estimation of the Contraceptive Transition

We now illustrate another potential use of our model, which also serves as a robustness check. The lighter line in Figure 12a corresponds to the proportions of unplanned births in France (simulated using the combination of parameters that minimized the distance between model output and data). As already shown in Figure 6e, there is a strong fit between the simulated and observed proportions, which is expected as the observed proportions were involved in the estimation of the model parameters. In the case of Ireland and Spain it is not possible to assess the reliability of the model estimates because we do not count on information relative to the contraceptive transition in these countries. We can, however, assess the capacity of the model to estimate the parameters related to the contraceptive transition *in the absence of information*, i.e., by creating a scenario in which we assume the information is also not available for France. In practical term this means assigning zero weight to the proportions of unplanned/unwanted births in the estimation procedure. The simulated proportions of unplanned births obtained from this exercise are represented by the dark line in Figure 12a. Although the fit is not perfect, it provides strong indication that the estimated values for Ireland and Spain (Figures 12b and 12c) are close to the true values.

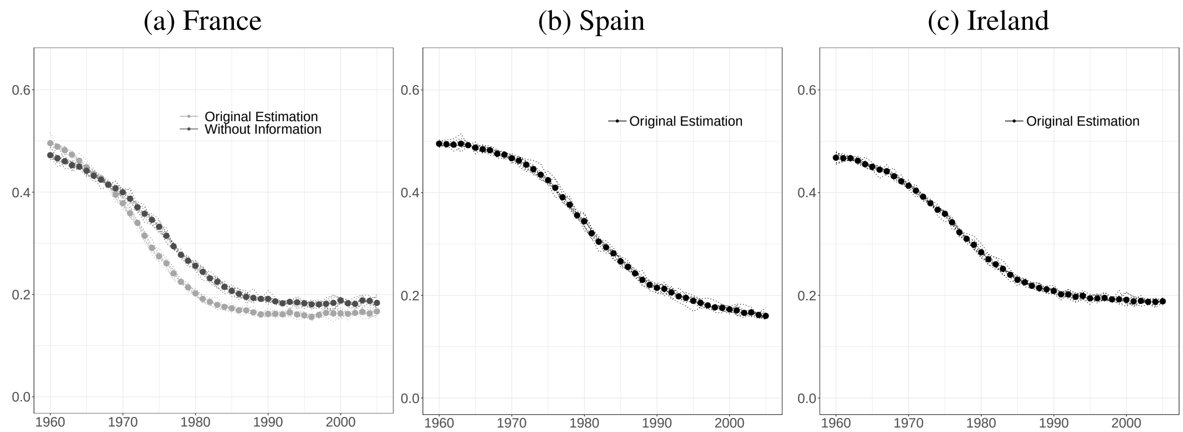

Figure 11: Direct vs Indirect Estimates of the Proportions of Unplanned Births

5.2 Forecasted Scenarios

Figure 13 presents the simulated and forecasted Total Fertility Rates, Cohort Completed Fertility and Mean Ages at Births for the three countries studied and for three different scenarios. The darker lines corresponds to a scenario in which the positive effect of educational attainment on the intention is strong (scenario 3). We obtain these results by setting $tau = 12$ while keeping the rest of parameters in the values estimated in the base sceanario. This implies that having 12 years of education already significantly contributes to reduce the penalty on fertility intentions imposed by the participation in the labor market. At the aggregate level this implies a stronger positive push on fertility rates as the number of individuals in the population with higher education increases. The second line in the scale of greys, which sits below the other two in the three plots, corresponds to a scenario in which the increase of resources associated with increased education does not have any effect on fertility intentions (scenario 2). The remaining line corresponds to the base scenario (scenario 1).

The distance between scenarios 1 and 2 shows that the model estimates a relatively strong positive effect of educational expansion on fertility rates. Without these positive effect the model would have predicted the level of fertility to be much lower than observed.

Another interesting result, is that even in the presence of a strong effect of education on fertility intentions, there is no marked upward trend on the TFRs. The reason for this result is that in all three countries there is still room for systematic increases of the Mean age at Birth, which, by construction, exerts a downward pressure on the TFR, but also has a *real* negative effect on the level of fertility via furthers reductions in exposure.

The slight upward trend towards the end of the projection in France suggest that the TFR might experience a recuperation as the expansion of education comes to an end, but this also means the potential benefit to fertility rates from a more resourceful population will also come to an end.

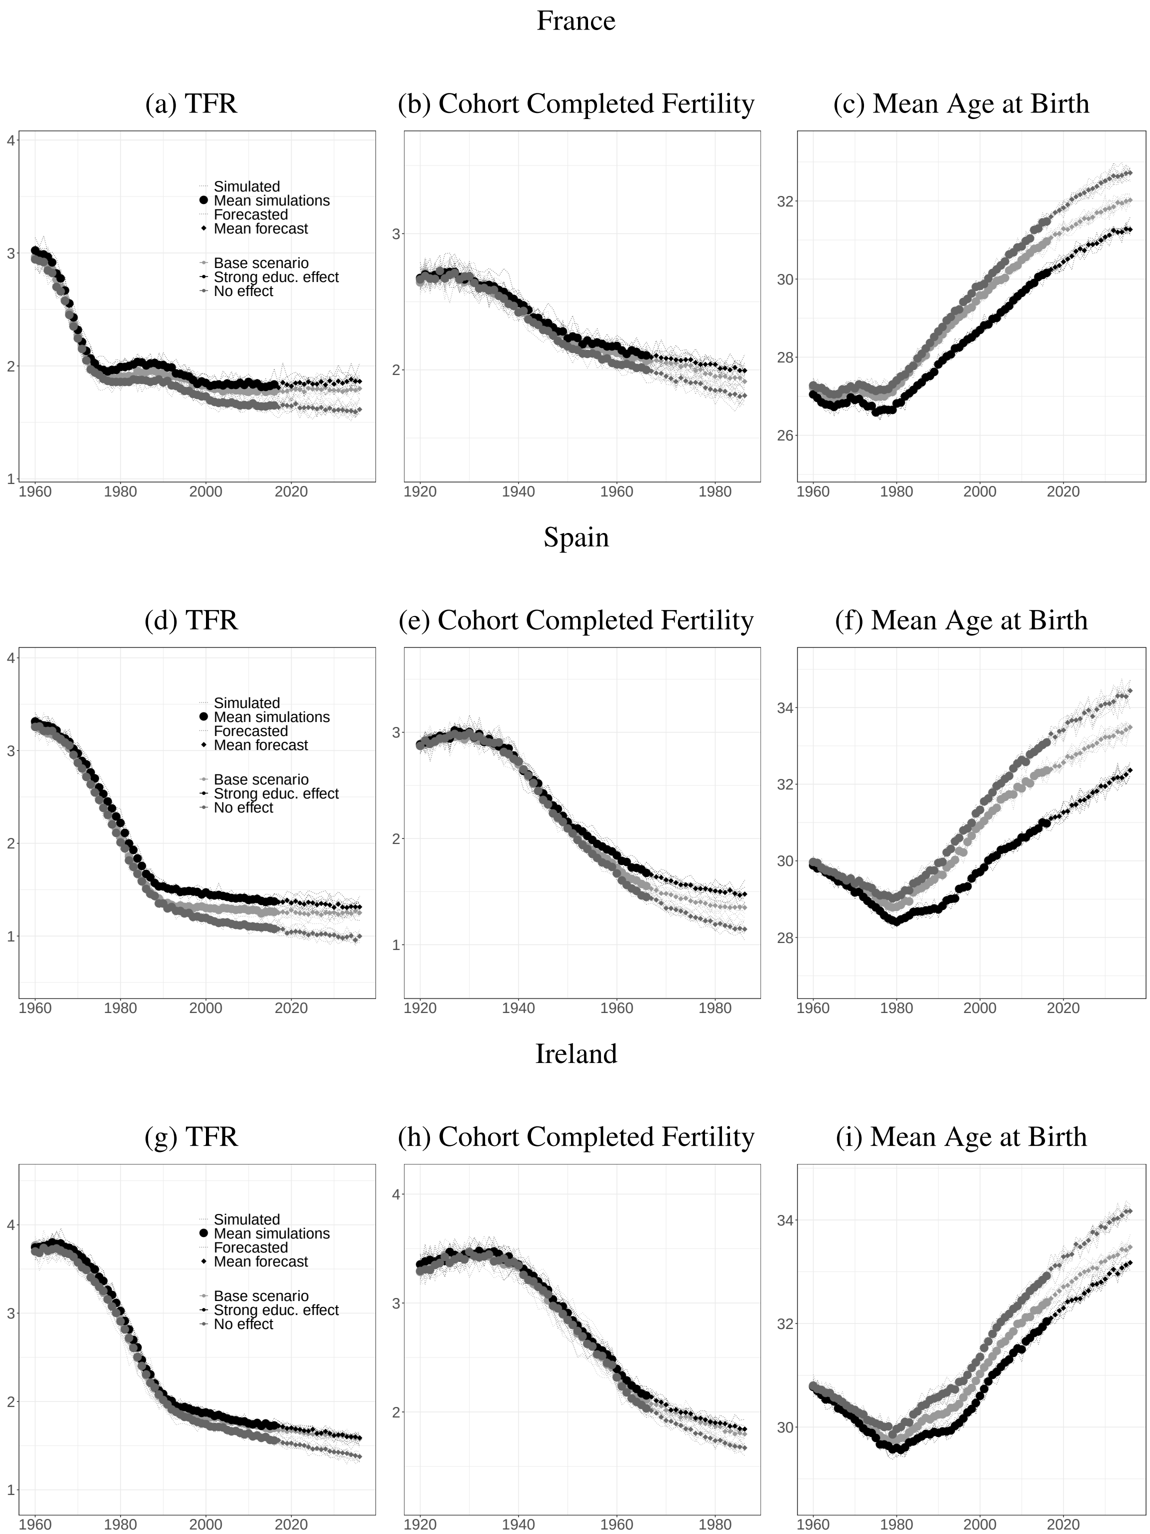

Figure 13: Forecasted Fertility Indicators | Scenarios.

6 Discussion

Systematic evidence showing that the relationship between development and fertility was more complex than originally depicted has been available for at least a decade now. During this time new evidence from a variety of contexts, time periods and levels of analysis has accumulated, but less progress has been made with respect to the *mechanisms* behind this complex relationship. Perhaps the slower pace in the search for explanations is related to our tendency to think of the decline and recuperation of fertility as two separate phases, driven by a different set of forces. This tendency is not surprising; understanding the recuperation of fertility as a new process, disconnected from earlier trends, allows us to retain the analytical notions we use to interpret the past.

The hypothesis we proposed and tested here implies taking a different road and critically examine deeply rooted ideas in the light of new evidence. What the estimation of our model shows is, in fact, that while it is correct to argue that the process of socioeconomic development has been contributing to fertility decline, it is also necessary to understand that the same set of forces was, at the same time, pushing in the opposite direction, and that the real question is which of these competing forces exerts a stronger influence at each stage of the fertility transition.

By fitting our model to data from three European countries and creating a series of counterfactual scenarios we showed that without the increasing ability and resources associated with the expansion of higher education, fertility rates will be lower than observed, specially in most recent years. But while increases in education have been key to maintain fertility levels, the expected additional increases in educational attainment will not necessarily drive an upward trend in period fertility rates. Part of this result is explained by the sensitivity of the TFR to changes in the age at birth, which in all our scenarios continues to increase as the total time spent in the education system among recent cohorts continues to increase. Apart from the mechanical effect on the TFR, this process has a *real* influence on the level of fertility express through the reduction of the time individuals in younger cohorts spent at risk of having a child.

More fundamentally, there is a harder limit to the recuperation of fertility rates in the not so distant future imposed by the deceleration and eventual end of the transition to higher education. In fact, as the three transitions that drove fertility change after the baby-boom come to an end, the space for systematic variation of fertility levels will be significantly reduced. Period rates will continue to be influenced by period effects associated to migration flows, labor market changes, etc. Other more sustainable changes might be observed related to policy efforts or cultural transformations affecting the social norm regarding the ideal family size. But we showed that the scenario of a TFR oscillating around a relatively stable level in the near future is not unlikely when we analyze the expected trajectories of the fundamental drivers of fertility levels since the second half of the 20th century.

A Appendix: Comfert Algorithm

read input data

initialization
generate wts to birth in the first year

simulation of events

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

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

if *n_event* $=$ *birth* **then**
    if *girl* **then**
        add to pop
        create attributes / wts:
        education
        activity
        wt to union
        wt to death
    end
    update indicators of mother:
    age at birth
    nr. children
    wt update intention
end
if *n_event* $=$ *union formation* **then**
    update:
    age at union
    desired family size
    birth intention
    risk unplanned
    wt conception
    wt update intention
end
if *n.event* $=$ *update intention* **then**
    update:
    intention
    risk unplanned
    wt conception
    wt update intention
end
if *n_event* $=$ *conception* **then**
    if *if random uniform* $< p_x$ **then**
        update:
        pregnant
        wt birth
        wt conception
    else
        failed attempts + 1
        wt update intention
        wt conception
end
if *n_event* $=$ *death* **then**
    remove from the population
end

if *year change* **then**
    remove those with age > *max_age*
    generate wts to birth for next year
    compute indicators
    reset indicator variables
end
save output
asfrs
unplanned births
desired family size
end

B Appendix: Forecasted Input Distributions

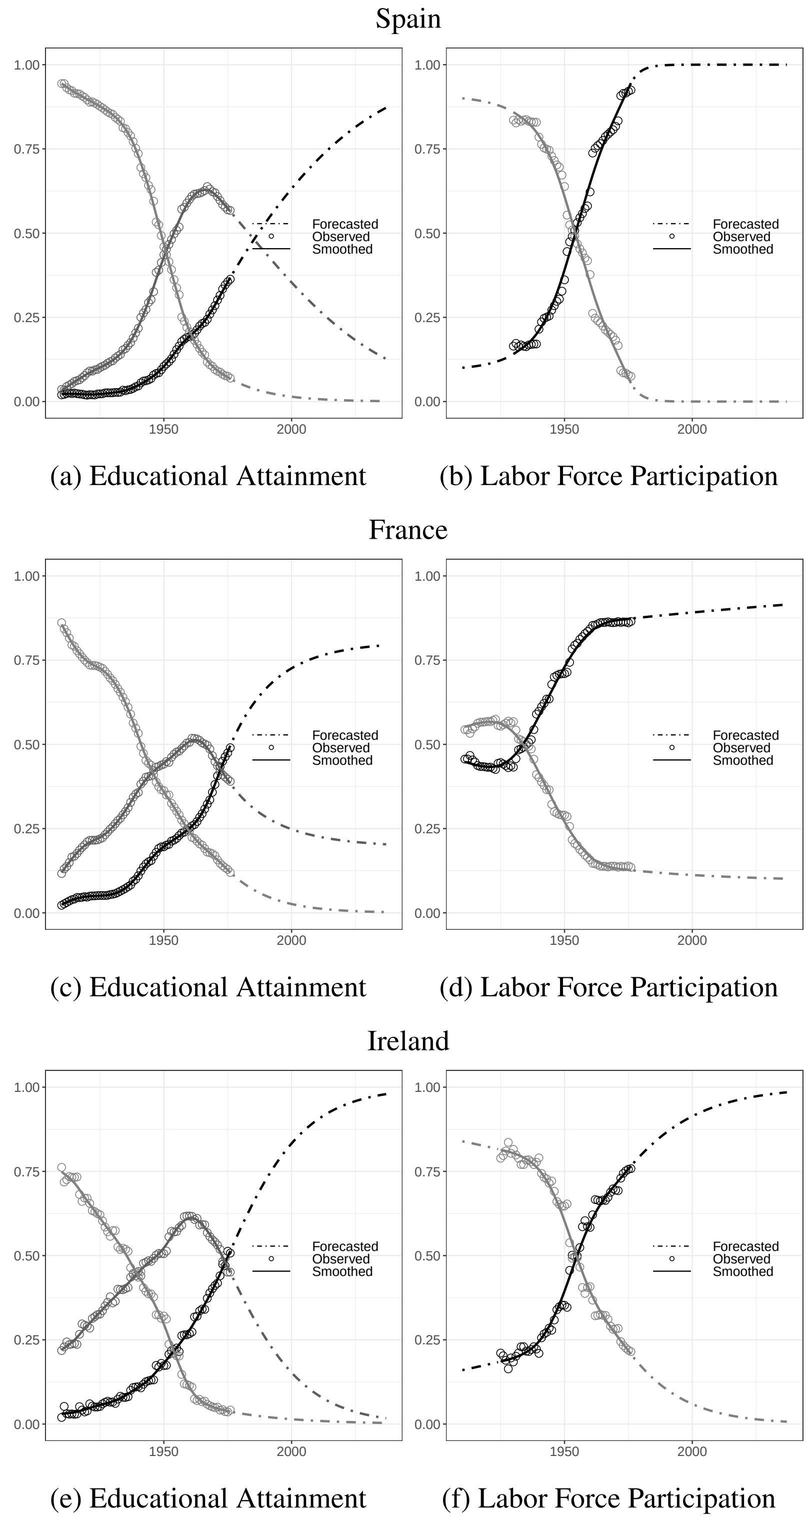

Figure 15: Observed and Forecasted Distributions of Educ. Attainment and Labor Force Participation.

C Appendix: Estimation of Parameters with Partial Data

As mention in Section 3, in the case of Ireland and Spain we cannot rely on information on average *desired* family sizes, but we can exploit the fact that there is a systematic difference between the distributions of *desired* and *ideal* family sizes (Figure 17) [5] and that we have information on the latter distribution for all countries.

The strategy, in the case of Ireland and Spain, consists in finding values for the parameters associated to D ($\overline{D}_{t0}$, D^* and δ) that results in a series of average desired family size which keeps the same distance to the general ideal as in the case of France.

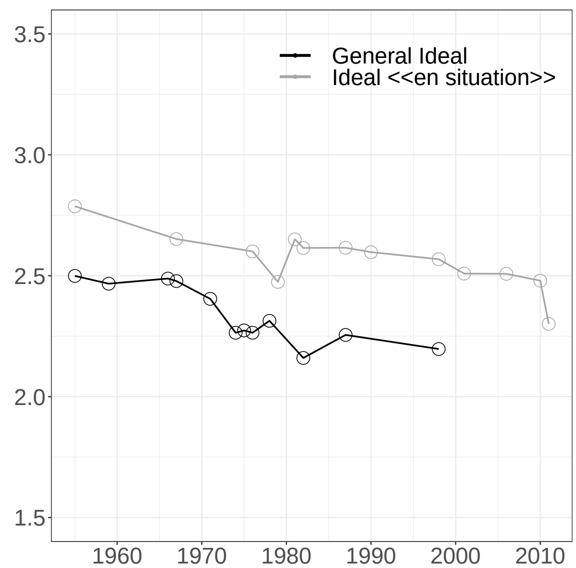

Figure 17: Ideal vs Desired Family Size | France.

Additionally, assuming that the proportions of unplanned births post-contraceptive-transition are not widely different across European countries, we use the estimated value of υ for France also for Spain and Ireland. Similarly, we do not estimate in the three cases the value of λ which is mostly determined by biological factors.

[5] This distance is only illustrative as the two series come from different sources and more importantly they are computed for different age-groups: 25-34 for the ideal *en situation* and 15-49 for the general ideal.

D Appendix: Additional Model Fit

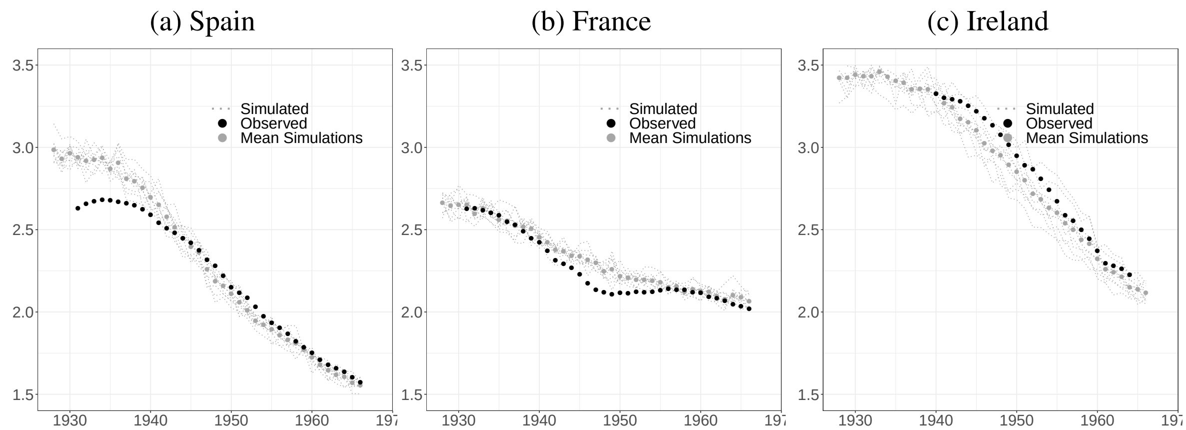

Figure 18: Observed vs Simulated Cohort Completed Fertility

References

Ahn, N. and P. Mira (2002). A note on the changing relationship between fertility and female employment rates in developed countries. *Journal of population Economics 15*(4), 667–682.

Berrington, A. and S. Pattaro (2014). Educational differences in fertility desires, intentions and behaviour: A life course perspective. *Advances in life course research 21*, 10–27.

Bongaarts, J. (1997). Trends in unwanted childbearing in the developing world. *Studies in Family planning*, 267–277.

Bongaarts, J. and T. Sobotka (2012). A Demographic Explanation for the Recent Rise in European Fertility. *Population and Development Review 38*(1), 83–120.

Burkimsher, M. (2017). Evolution of the shape of the fertility curve: Why might some countries develop a bimodal curve? *Demographic Research 37*, 295–324.

Caldwell, J. C. and T. Schindlmayr (2003). Explanations of the fertility crisis in modern societies: A search for commonalities. *Population studies 57*(3), 241–263.

Casterline, J. B. and L. O. El-Zeini (2007). The estimation of unwanted fertility. *Demography 44*(4), 729–745.

Chazelle, B. (2012). Natural algorithms and influence systems. *Communications of the ACM 55*(12), 101–110.

De Wachter, D. and K. Neels (2011). Educational differentials in fertility intentions and outcomes: family formation in flanders in the early 1990s. *Vienna Yearbook of Population Research*, 227–258.

Easterlin, R. A. (1978). What will 1984 be like? socioeconomic implications of recent twists in age structure. *Demography 15*(4), 397–432.

Esping-Andersen, G. (2009). *Incomplete revolution: Adapting welfare states to women's new roles*. Polity.

Esping-Andersen, G. and F. C. Billari (2015). Re-theorizing family demographics. *Population and Development Review 41*(1), 1–31.

Goldscheider, F., E. Bernhardt, and T. Lappegård (2015). The gender revolution: A framework for understanding changing family and demographic behavior. *Population and Development Review 41*(2), 207–239.

Goldstein, J. R., T. Sobotka, and A. Jasilioniene (2009). The End of "Lowest-Low" Fertility? *Population and Development Review 35*(4), pp. 663–699.

Heiland, F., A. Prskawetz, and W. C. Sanderson (2005). Do the more-educated prefere smaller families? Technical report, Vienna Institute of Demography Working Papers.

Hilton, J. and J. Bijak (2017). Design and analysis of demographic simulations. In *Agent-based modelling in population studies*, pp. 211–235. Springer.

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

Hyndman, R. J. and Y. Khandakar (2008). Automatic time series forecasting: the forecast package for R. *Journal of Statistical Software 26*(3), 1–22.

Klijzing, E. (2000). Are there unmet family planning needs in europe? *Family planning perspectives*, 74–88.

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.

Leridon, H. and L. Toulemon (1999). La famille idéale: combien d'enfants, à quel âge? *Insee première 652*.

Luci-Greulich, A. and O. Thévenon (2014). Does economic advancement 'Cause'a re-increase in fertility? An empirical analysis for OECD countries (1960–2007). *European Journal of Population 30*(2), 187–221.

Mencarini, L. and M. L. Tanturri (2004). Time use, family role-set and childbearing among italian working women. *Genus*, 111–137.

Mills, M., L. Mencarini, M. L. Tanturri, and K. Begall (2008). Gender equity and fertility intentions in italy and the netherlands. *Demographic research 18*, 1–26.

Minnesota Population Center, . (2018). Integrated public use microdata series, international: Version 7.0.

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

Murphy, M. (1993). The contraceptive pill and women's employment as factors in fertility change in britain 1963–1980: A challenge to the conventional view. *Population studies 47*(2), 221–243.

Myrskylä, M., H.-P. Kohler, and F. C. Billari (2009). Advances in development reverse fertility declines. *Nature 460*(7256), 741–743.

R Core Team (2015). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Raftery, A. E., L. Alkema, and P. Gerland (2014). Bayesian population projections for the united nations. *Statistical science: a review journal of the Institute of Mathematical Statistics 29*(1), 58.

Régnier-Loilier, A., H. Leridon, F. Cahen, et al. (2007). Four decades of legalized contraception in france: an unfinished revolution? Technical report, Institut National d'Études Démographiques (INED).

Rindfuss, R. R. and K. L. Brewster (1996). Childrearing and fertility. *Population and Development Review 22*, 258–289.

Rosero-Bixby, L. and J. B. Casterline (1993). Modelling diffusion effects in fertility transition. *Population Studies 47*(1), 147–167.

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.

Sullivan, O., F. C. Billari, and E. Altintas (2014). Fathers' changing contributions to child care and domestic work in very low–fertility countries: The effect of education. *Journal of Family Issues 35*(8), 1048–1065.

Testa, M. R. (2014). On the positive correlation between education and fertility intentions in europe: Individual-and country-level evidence. *Advances in life course research 21*, 28–42.

United Nations, P. D. (2010). *World population prospects: The 2010 revision*. Department of Economic and Social Affairs.

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

Willekens, F., J. Bijak, A. Klabunde, and A. Prskawetz (2017). The science of choice: an introduction.

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