American Economic Journal: Macroeconomics 2020, 12(3): 77–109
https://doi.org/10.1257/mac.20180296

The Fall in Global Fertility: A Quantitative Model[†]

By Tiloka de Silva and Silvana Tenreyro*

Over the past six decades, fertility rates have fallen dramatically in most middle- and low-income countries. To analyze these developments, we study a quantitative model of endogenous human capital and fertility choice, augmented to allow for social norms over family size. We parametrize the model using data on socioeconomic variables and information on funding for population-control policies aimed at affecting social norms and improving access to contraceptives. We simulate the implementation of population-control policies to gauge their contribution to the decline in fertility. We find that policies aimed at altering family-size norms accelerated and strengthened the decline in fertility, which would have otherwise taken place much more gradually. (*JEL* J10, J13, J18, J24, O15, Z13)

Over the past six decades, most developing countries have experienced remarkable declines in total fertility rates (TFR). The world's average TFR declined steadily, falling from an average of 5 children per woman in 1960 to an average of 2.5 in 2015. This decline in fertility is not skewed by the experience of a few countries. In 1960, more than half of the countries in the world recorded average fertility rates greater than 6 children per woman. By 2015, the median TFR was 2.2 children per woman.[1]

These large declines in fertility took place in most regions of the world despite widely varying levels of development (see Figure 1). More specifically, the relationship between fertility and development (as measured by GDP per capita) has shifted downward and become flatter. The size of the downward shift has amounted to an average of 2 children per woman, implying that today a typical woman has 2 fewer children than a woman living in a country at the same level of development in 1960 (de Silva and Tenreyro 2017). The time series of average fertility and income for developing countries further highlights the rapid transition—by 2014, the average fertility was much lower than would have been predicted by the average incomes in the cross-sections of 1960. That is, developing countries have, on average, reached

* de Silva: Department of Decision Sciences, University of Moratuwa, Moratuwa 10400 (email: tilokad@uom.lk); Tenreyro: London School of Economics, Houghton Street, London WC2A 2AE, and Bank of England, CfM, CEPR (email: S.Tenreyro@lse.ac.uk). Virgiliu Midrigan was coeditor for this article. For helpful conversations we thank Charlie Bean, Robin Burgess, Francesco Caselli, Nicola Fuchs-Schundeln, Nezih Guner, Steve Machin, Laura Castillo, Per Krusell, and participants of the CfM Macro Seminars and CEPR Annual Growth meeting. Tenreyro acknowledges financial support from ERC consolidator grant 681664–MACROTRADE. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

† Go to https://doi.org/10.1257/mac.20180296 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

[1] See online Appendix for data on the change in fertility between 1960 and 2015 by country.

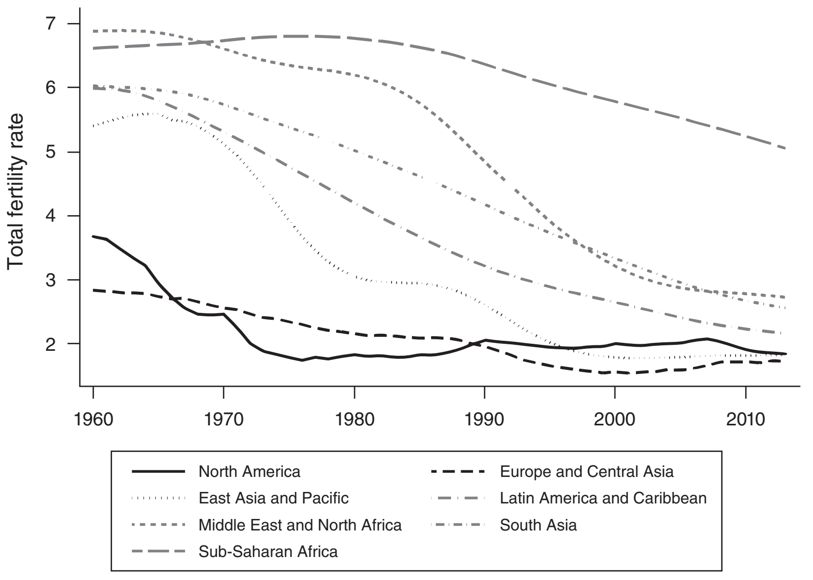

Figure 1. Fertility Trends by Region

Source: de Silva and Tenreyro (2017) using data from the World Bank's World Development Indicators database.

fertility levels similar to that of developed countries at much lower average income levels (see Figure 2).[2]

de Silva and Tenreyro (2017) have argued that while socioeconomic factors played an important role in the worldwide fertility decline, the timing and speed of the fall over the past decades suggest that the population control policies implemented in many developing countries over this period played a significant role in accelerating the process.[3] The design of population-control programs consisted of two main parts. The first was the diffusion of contraceptive supply and information. The second was the implementation of public campaigns aimed at reversing pronatalist attitudes and establishing a new small-family norm. One of the inferences drawn from our study is that the second strategy of employing public campaigns to reduce desired levels of fertility was critical in complementing contraceptive provision. The exact size of the effect, however, is not easy to gauge from the empirical evidence, as endogeneity impedes a clean causal inference.

In this paper, we study a model of endogenous fertility and human capital accumulation, building on the Barro-Becker framework of fertility choice, incorporating

[2] Recent work by Delventhal, Fernández-Villaverde, and Guner (2019) study the demographic transitions in 188 countries and find that transitions have, indeed, grown faster over time, starting from higher birth rates and lower levels of income.

[3] A number of socioeconomic factors have been cited as possible causes for low fertility, including higher income, lower mortality, increasing investments in education, and rising female labor force participation (see Manuelli and Seshadri 2009; Jones, Schoonbroodt, and Tertilt 2010; and Albanesi and Olivetti 2016 for some recent examples). Focusing on the fertility decline in developing countries, de Silva and Tenreyro (2017) report that different measures of the intensity of family planning programs are strongly and positively associated with fertility declines, even after controlling for changes in a wide range of such covariates.

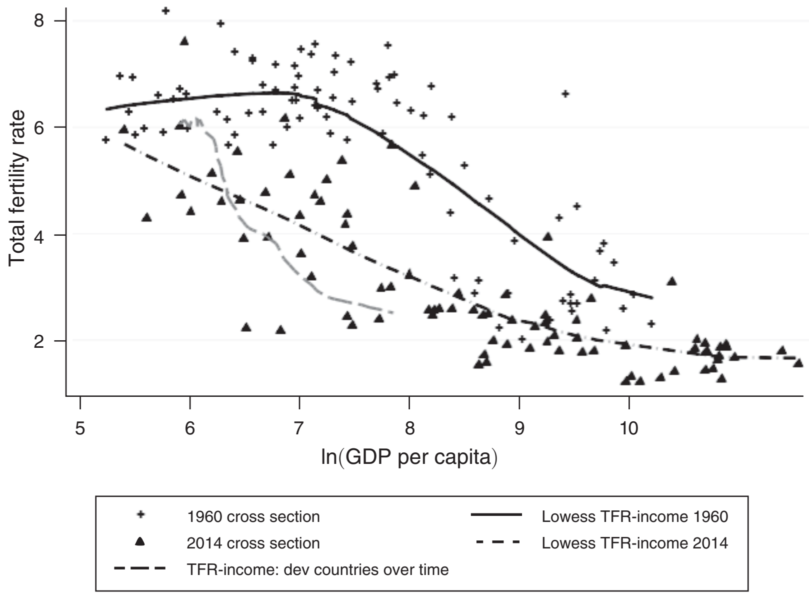

Figure 2. Fertility-Income Relationship

Notes: The graph plots the TFR-GDP per capita relationship for a cross-section of countries from 1960 and 2014 (the Lowess smoothed functions are given by the solid and dash-dot lines) and the time series of average fertility and GDP per capita for developing countries from 1960 to 2014 (dashed line).

Source: Updated from de Silva and Tenreyro (2017) using data from the World Bank's World Development Indicators database.

human capital investment (see Barro and Becker 1989, Galor and Weil 2000, Galor and Moav 2002, Moav 2005). We augment the model to include a role for endogenously evolving social norms on family size. The model allows us to analyze the factors underpinning the fertility decline observed in developing countries and quantify their causal contribution, circumventing the challenges faced by reduced-form estimations.

In the model, individuals derive utility from both the quantity and "quality" of children and dislike deviating from the social norm on the number of children.[4] Our modeling of adherence to social norms borrows from the literature on social distance and conformity (Jones 1984, Akerlof 1997) so that individuals derive disutility from a function of the distance between their realized fertility and the social norm.[5] The definition of the family-size norm builds on the sociology and demography literature, where the norm is influenced by the size of the family of origin or relevant reference groups.[6] As such, the norm is portrayed in the model as an

[4] We follow the literature's jargon, where "quality" relates to the level of human capital of the individual.

[5] We deviate from the existing work on the impact of social norms on fertility in how we model social norms. Spolaore and Wacziarg (2014), Munshi and Myaux (2006), Manski and Mayshar (2003), Palivos (2001), and Bhattacharya and Chakraborty (2012) model norms as the outcome of strategic decision-making and interaction. We take a simpler specification that is more amenable to quantification and in line with the literature on external habits or reference dependence.

[6] See, for example, Clay and Zuiches (1980) for a discussion on the importance of reference groups in forming fertility norms; Thornton (1980), Murphy (1999), and Kolk (2014), who explore the impact of parental fertility on

evolving weighted average between the fertility of the previous generation and the long-term replacement level of fertility, which we set to be equal to 2 children per woman. We choose the replacement level as the second term in the average as this was the fertility level advocated and promoted by most population-control programs in their public campaigns.

We calibrate the model's structural parameters and initial conditions to match key moments of the data for developing countries in 1960 and use it to simulate the transition to the steady-state levels of fertility and human capital. We show that the baseline model, where the only mechanism by which fertility is driven down is the accumulation of human capital, can endogenously generate only a small decline in fertility rates. Incorporating social norms into the model generates a faster and larger decline than that yielded by the model without norms, though that alone is not sufficient to match the sharp decline observed in the data.

We simulate the effect of population-control policies on family-size norms using information on funding for family planning programs. In particular, given that the majority of the programs advocated having two children, we allow the weight placed on the replacement level of fertility to increase with the intensity of these programs and shift the social norm on family size downwards.[7] The simulation shows that the introduction of policies aimed at altering family-size norms significantly accelerates and strengthens the decline in fertility that would otherwise take place much more gradually as economies move to higher levels of human capital.

We then consider several alternative mechanisms that might explain the fertility decline, with the model allowing us to gauge quantitatively the role played by the different channels. The first extension explores the role played by the fall in mortality rates and finds that, in a setting in which there is child mortality and uncertainty about how many children survive to adulthood, the decline in mortality alone is not sufficient to explain the fall in fertility observed over the past few decades.[8] The second extension of the model considers the case in which households cannot fully control fertility rates (contraception technologies are either not available or imperfect). In that setting, we study the role played by increased access to contraception

fertility outcomes; and Fernández and Fogli (2009), who find higher fertility among women whose ancestry is from high TFR countries.

[7] The main version of the model assumes that households count with the technology to control fertility. While the data on family planning funds do not allow a breakdown of funding used for increasing contraceptive access and funding used for promoting a smaller family size, what is clear in the data is that family planning funding per capita is strongly and negatively correlated with "wanted fertility" rates (as defined in Demographic and Health Surveys), which are likely to reflect preferences, but uncorrelated with "unwanted fertility" rates, which are more closely (negatively) related to access to contraception (see Appendix A1). This evidence suggests that the effect of family planning programs operated through a preference-changing channel rather than through the access to contraception channel.

[8] This point was previously made by Doepke (2005), Fernández-Villaverde (2001), and Becker and Barro (1988). The Becker and Barro (1988) model predicts that when mortality rates decrease, the total fertility rate falls, but the number of surviving children remains the same. (In other words, if people's preferences for the surviving number of children do not change, fertility falls only insofar as is necessary to achieve the same final target.) In survey data, however, we observe a decline not only in fertility rates, but also in the desired number of children, that is, a change in the final target. Cervellati and Sunde (2015) overcome this problem by introducing differential fertility across education groups, which interacts with increasing longevity, to drive down both total and net fertility. However, the authors note that while their model captures well the transition of the European economies, it does not fully capture the acceleration experienced by many developing countries after 1960. It is precisely this acceleration that our paper seeks to explain.

(the other main component of population-control policies) and find that the changing fertility norms have a much larger effect on the fertility decline than increased access to contraception, consistent with the fact that many of the family planning programs supplemented their supply-side strategies of increasing access to contraception with large-scale mass media campaigns to promote smaller family sizes.

We do not explicitly model the possibility that children provide their parents with transfers in their old age, but our modeling choices can be recast in those terms, as parents care about their children's future earning capacity.[9] We also abstract from the analysis of child labor and compulsory schooling policies, such as that in Doepke and Zilibotti (2005), leaving the joint analysis of these policies together with population-control policies for future work. In what follows, we describe the model in more detail, specifying technologies and preferences.

The rest of the paper is organized as follows. Section I describes the model. Section II explains the calibration strategy and describes the data used in the analysis. Section III presents the main results of the paper and Section IV studies various extensions of the model. Section V presents concluding remarks.

I. The Model

We consider an overlapping-generations economy in which individuals live for two periods: childhood and adulthood. In each period, the economy produces a single consumption good using as inputs the productive capacity of all working adults and a fixed factor. The human capital stock is determined by the fertility and educational choices of individuals. There is also a government, which levies taxes from households and spends all revenues on education.

A. *Technology*

Production occurs according to a constant returns to scale technology. Using the specification in Galor and Weil (2000), output at time t, Y_t, is

$$(1)\qquad Y_t = \left[\left(\bar{H} + H_t\right)L_t\right]^{\rho}(A_t X)^{1-\rho}, \quad 0 < \rho < 1,$$

where $\bar{H} + H_t$ is the productive capacity of a worker, L_t is the working age population, X is the fixed factor, and A_t is the technology at time t, with $A_t X$ referring to "effective resources." The term $\bar{H}$ is a physical labor endowment all individuals are born with and H_t is human capital produced with investments in schooling.

Output per worker at time t, y_t, is

$$(2)\qquad y_t = \left(\bar{H} + H_t\right)^{\rho} x_t^{1-\rho},$$

where $x_t = A_t X / L_t$ is the effective resources per worker at time t.

[9] There is a growing literature that addresses these intergenerational transfers explicitly (see for example Boldrin and Jones 2002; Coeurdacier, Guibaud, and Jin 2014; Choukhmane, Coeurdacier, and Jin 2017).

As in Galor and Weil (2000), we assume that the return to the fixed factor is zero. This assumption helps to keep the model simple so that the only source of earnings for households is labor income, which is a reasonable description of households' funding in developing countries. The factor X can then be interpreted as a productive public good (e.g., a natural resource) that does not yield private returns to the citizens. Galor and Weil's (2000) interpretation is that there are no property rights over this resource in the country.

The return to productive labor, w_t, is then given by its average product:

$$w_t = \left(\frac{x_t}{\bar{H} + H_t}\right)^{1-\rho}. \tag{3}$$

B. *Households*

Each household has a single decision maker, the working adult. Individuals within a generation are identical. Children consume a fraction of their parents' time. Working adults supply labor inelastically, decide on their consumption, the number of children, and their education in period t.

Parents are motivated by altruism towards their children but are conscious of the social norm on the number of children that a family should have. As such, while parents derive utility from their children (both the quantity and the quality), they derive disutility from deviating from the social norm. The utility function for a working age individual of generation t can be expressed as

$$U_t = u(C_t; n_t; q_{t+1}) - \varphi g(n_t, \hat{n}_t), \tag{4}$$

where u is a standard utility function over three goods: C_t, which denotes consumption at time t; n_t, which denotes the number of children; and q_{t+1}, which indicates the quality of children as measured by their future earning potential. Following Galor and Weil (2000) and Moav (2005), we assume $q_{t+1} = w_{t+1}(\bar{H} + H_{t+1})$, where w_{t+1} is the future wage per unit of productive labor of a child, and $\bar{H} + H_{t+1}$ is the productive capacity of a child. The factor $\varphi > 0$ governs the disutility from deviating from the social norm and $g(n_t, \hat{n}_t)$ is a function of the deviation of the chosen number of children, n_t, from the social norm on family size, $\hat{n}_t$, where $g_{11}(n_t, \hat{n}_t) > 0$; $g_{12}(n_t, \hat{n}_t) < 0$. The first condition implies that movements further away from the norm involve heavier penalties, while the second implies that the marginal cost of the additional child is decreasing in the social norm. We model the social norm on family size as a weighted average between the previous generation's fertility, n_{t-1}, and the replacement level of fertility, n^*, so that $\hat{n}_t$ can be expressed as

$$\hat{n}_t = \phi n^* + (1 - \phi) n_{t-1}, \quad 0 \leq \phi \leq 1. \tag{5}$$

The individual's choice of desired number of children and optimal education investment for each child is subject to a standard budget constraint. While parental income is given by $w_t(\bar{H} + H_t)$, we assume that a fixed fraction of income, τ_0, is spent on each child regardless of education, and a discretionary education cost for each child, $\tau_1 h_t$, which is increasing in the level of education, h_t, is chosen by the parents.[10] Households also pay a fraction, τ_g, of their income as tax. The remaining income is spent on consumption. The budget constraint at time t is therefore,

$$C_t = \left[1 - \tau_g - (\tau_0 + \tau_1 h_t) n_t\right] w_t(\bar{H} + H_t). \tag{6}$$

The cost of a year of schooling for a child is a fraction, τ_h, of income, which is met through household and government spending on education. This means

$$\tau_h n_t h_t w_t(\bar{H} + H_t) = G_t + \tau_1 n_t h_t w_t(\bar{H} + H_t), \tag{7}$$

where G_t is the government expenditure on education in period t.

The government spends all its tax revenue on education and, in equilibrium, runs a balanced budget, so

$$\tau_g w_t(\bar{H} + H_t) = G_t. \tag{8}$$

Together with equation (7), this gives

$$\tau_1 = \tau_h - \frac{\tau_g}{n_t h_t}. \tag{9}$$

We assume, for simplicity, that households internalize the government budget. The household budget constraint can thus be written as

$$C_t = \left[1 - (\tau_0 + \tau_h h_t) n_t\right] w_t(\bar{H} + H_t). \tag{10}$$

Following Becker, Murphy, and Tamura (1990) and Ehrlich and Kim (2005), we specify the human capital production function as

$$H_{t+1} = z_t(\bar{H} + H_t) h_t, \tag{11}$$

where $\bar{H} + H_t$ is the productive capacity of the parent, h_t is the educational investment (or schooling) in each child, and z_t is the human capital production technology. This specification of productive capacity prevents perfect intergenerational transmission of human capital, allowing for positive levels of human capital even for children whose parents have no schooling ($H_t = 0$).

C. *Equilibrium*

In a competitive equilibrium, agents and firms optimally solve their constrained maximization problems and all markets clear. Let $(v) = (\bar{H} + H_t, n_{t-1})$. A

[10] While changes in the gender wage gap could also have an effect on the opportunity cost of child rearing by altering the woman's bargaining position in the household (see, for example, Doepke and Kindermann 2016), modeling non-cooperative solutions is beyond the scope of this paper. However, to the extent that female labor force participation reflects some of the female power in a society, in the cross section of countries, we find no systematic relation between female labor force participation and fertility rates.

competitive equilibrium for this economy consists of a collection of policy functions for households $\{C_t(v), n_t(v), h_t(v)\}$, and prices w_t such that:

(i) policy functions $C_t(v)$, $n_t(v)$, and $h_t(v)$ maximize

$$u(C_t; n_t; q_{t+1}) - \varphi g(n_t, \hat{n}_t)$$

subject to the budget constraint (6), human capital production function (11), the law of motion for norms (5), and $(C_t, n_t, h_t) \geq 0$;

(ii) w_t satisfies equation (3);

(iii) the government runs a balanced budget, satisfying equation (8); and

(iv) the market for the final consumption good clears such that

$$C_t = \left[1 - \tau g - (\tau_0 + \tau_1 h_t) n_t\right] y_t.$$

II. Calibration

In the policy experiments that we carry out, we examine the transition of the economy from a given initial condition to a steady-state level of fertility and human capital investment. Our calibration strategy consists of choosing structural parameters and initial conditions so that the outcomes of the model in the first period match the appropriate moments for consumption, income, fertility, years of schooling, spending on education, and population in developing countries in 1960.[11] Since the economic agent in this model is an individual, the fertility rate in the model is one-half of the total fertility rate in the data. We interpret the units of investment in human capital per child, h_t, as years of education. (We ignore integer constraints in the model and treat both fertility and years of education as continuous variables; the empirical counterparts are also not integers, as they are given by the average fertility and average number of years of education.)[12,13] A period in the model corresponds to the length of one generation, which we set to be 25 years.

The data on household consumption, per capita GDP, government spending on education as a fraction of GDP, population, and fertility are obtained from the World

[11] We refer to all countries which were not classified as OECD countries prior to 1970 as developing countries in the starting period. 1960 is the first year for which cross-country data on fertility, income, and consumption are available.

[12] The data we use for education is the expected years of schooling of children of school entrance age obtained from UN Development Programme (2013).

[13] The data on expected years of schooling starts from 1980. Therefore, to obtain the average years of schooling for 1960, we compare the series with years of schooling for the adult population taken from Barro and Lee (2013). The average years of schooling for the adult population (aged 25+) in our sample of countries in 1985 is 3.67. Since this measure is likely to understate the level of education of the younger cohorts, we set expected years of schooling for children born in 1960 to be 5.

Bank's World Development Indicators (WDI) dataset, while the data on expected years of schooling are taken from the UN Development Programme (2013).

A. *Technology*

Estimates of total factor productivity in East Asian countries over the 1966–1990 period by Young (1995) indicate that, on average, annual TFP growth over the period ranged from −0.003 in Singapore to 0.024 in Taiwan. As such, we will assume a constant annual TFP growth rate of 0.018 which is compounded to obtain the TFP growth rate between generations, g_A. We set the Cobb-Douglas coefficient on labor, ρ, to 0.66.[14] Finally, we assume that there is no growth in the technology used in human capital production, z_t.

B. *Cost of Child-Rearing*

We use data on the fraction of household expenditure allocated to education reported in household surveys and government expenditure on education to calibrate the values of τ_0 and τ_h. (See Appendix A2 for a detailed description of data and sources.) In our model, the fraction of household expenditure allocated to education is represented by $\tau_1 n_t h_t$. The value for τ_1, calculated using corresponding values for n_t and h_t from the data, ranges from 0.1 percent in Latin America to 0.6 percent in Singapore. We therefore set τ_1 to its mean value, 0.003.

Government expenditure on education as a share of output is represented in our model by τ_g. We combine the average government expenditure on education as a fraction of GDP in developing countries with the calibrated value for τ_1 using equation (9) to back out the value for τ_h.[15]

We then use the household budget constraint to back out the value for τ_0, the non-discretionary component of the cost of child-rearing, given the initial levels of income, consumption, fertility, and education.

C. *Preferences*

Following the literature, we assume utility is additively log linear in consumption, the number of children, the quality of children, and social norms:

$$U_t = \ln C_t + \alpha \ln n_t + \theta \ln\left[w_{t+1}\left(\bar{H} + H_{t+1}\right)\right] - \varphi g\left(n_t, \hat{n}_t\right), \tag{12}$$

$\alpha > 0$ reflects preferences for children, $\theta > 0$ for child quality. As noted in Akerlof (1997), the use of the absolute value of the difference between individual

[14] Our specification of utility implies that the values of g_A and ρ affect the simulations only through the initial value for the human capital stock and the calibrated value of θ as wages do not have an effect on fertility or human capital investment decisions.

[15] Our specification of a fixed value for τ_h is based on the assumption that household spending on education is high when government spending is low. While there is insufficient data to check this empirically for developing countries, it is possible to use data for 39 OECD and partner countries to show that, once income differences have been controlled for, there is a negative relationship between private and public education expenditure. See Appendix A3.

fertility and the social norm gives rise to multiple equilibria. We use a more tractable functional form given by

$$g(n_t, \hat{n}_t) = (n_t - \hat{n}_t)^2,$$

where individuals derive disutility from deviating both from above as well as below the social norm and deviations in either direction are penalized symmetrically. In Section IV, we consider a different functional form which treats upward and downward deviations asymmetrically and find that the results are very similar.

Given these preferences, the first-order condition for n_t is given by

$$\text{(13)} \qquad \frac{\alpha}{n_t} = \frac{(\tau_0 + \tau_h h_t)}{1 - (\tau_0 + \tau_h h_t)\, n_t} + 2\varphi(n_t - \hat{n}_t).$$

The first-order condition equates the marginal benefit of having children with the marginal cost. The first term on the right-hand side is the marginal cost in terms of foregone consumption while the second term will be a cost if the additional child pushes the total number of children over the social norm.

The first-order condition for h_t is

$$\text{(14)} \qquad \frac{\theta z_t(\bar{H} + H_t)}{(\bar{H} + H_{t+1})} = \frac{\tau_h n_t}{(1 - (\tau_0 + \tau_h h_t)\, n_t)},$$

where the left-hand side is the marginal utility to the parent from giving her child an additional unit of education and the right-hand side is the marginal cost in terms of foregone consumption.

Our specification of utility leaves us with three preference parameters (α, θ, and φ) to be calibrated. We also require initial values for H_t and z_t. We start by calibrating a baseline model in which individuals do not care about norms ($\varphi = 0$) and pin down α from the first-order condition for n_t, using the cross-country macro data for developing countries for 1960. We use the per capita output growth in the economy to pin down $(\bar{H} + H_{t+1})/(\bar{H} + H_t)$ (which we will refer to as g_H, hereafter). We choose the value of z_t, the technology converting schooling to human capital, to match the empirical estimates of the returns to schooling. Finally, we use the first-order condition for h_t and the human capital production function to obtain values for θ, the preference for child quality, and H_1, the level of human capital of parents in the initial period.[16]

[16] Rearranging the human capital production function gives:

$$H_t = \left(\frac{1}{g_H - z_t h_t} - 1\right)\bar{H},$$

where $g_H = (\bar{H} + H_{t+1})/(\bar{H} + H_t)$. In order to obtain $H_t > 0$, it is required that $(g_H - 1)/h_t < z_t \leq g_H/h_t$. Using values for g_H and h_t from the data, we can obtain an upper and lower bound for z_t.

The Mincerian return to schooling is given by $\rho z_t/g_H$ in our model. The value for z_t we obtain for a Mincerian return of 0.11, $\rho = 0.66$ and the calibrated value of g_H falls within the upper and lower bounds of z_t. Therefore, we set z_t to 0.47.

D. *Norms*

We use the first-order condition for fertility from the full model (equation (13)) to obtain a value for ϕ (the weight placed on the replacement fertility rate in the determination of the norm), for given values of φ and n_{t-1}.[17] We do not have enough moments in the data to back out the coefficient of disutility from deviating from norms, φ, and, to the best of our knowledge, there are no empirical estimates of this parameter. Therefore, we set $\varphi = 0.1$ and conduct sensitivity tests using a range of values for this parameter. While data on fertility rates in developing countries prior to 1960 is scarce, we set n_0 to 3.5 (meaning 7 children per woman—recall that in the model n_t is fertility per single-person household) based on estimates of fertility for several non-European countries in the early twentieth century provided by Therborn (2004). Finally, the replacement level of fertility, n^*, is set to 1, reflecting a replacement level fertility rate of 2.

Table 1 summarizes the results of the calibration exercise.

E. *Estimating the Change in* ϕ

We model the role of population-control policies in changing the social norms on family size by an increase in the weight on the replacement level of fertility, ϕ. In order to gauge the value of ϕ in subsequent periods, we estimate by ordinary least squares the first-order condition for fertility using data for 2010, holding all other parameters values (other than ϕ) constant. In other words, only the weight placed on the replacement rate of fertility is allowed to change. We model ϕ as a function of the intensity of family planning programs. Specifically, we set $\phi = \phi_1 P$, where P is family planning program intensity, measured by the logarithm of per capita funds for family planning, with the data on family planning funds compiled from Nortman and Hofstatter (1978); Nortman (1982); and Ross, Mauldin, and Miller (1993). This gives rise to the following estimable equation:

$$\text{(15)} \qquad \frac{\alpha}{n_t} - \frac{(\tau_0 + \tau_h h_t)}{1 - (\tau_0 + \tau_h h_t)\, n_t} - 2\varphi(n_t - n_{t-1}) \;=\; 2\varphi\,\phi_1 P(n_{t-1} - n^*).$$

We estimate the equation using data on fertility and expected years of schooling for 2010, and the average value of per capita funds for family planning over the 1970–2000 period. Ideally, P would be the total spending per capita on family planning programs over this period. However, given that for many countries we have data only for one or two years, we use the average per capita funding over the period 1970–2000. Note that this exercise is an attempt to recover a numerical estimate for ϕ which can be used in the quantitative analysis, rather than to establish a causal link between the family planning programs and fertility.

The estimation of equation (15) provides us with a value for ϕ_1. We find that the estimated coefficient (corresponding to $2\varphi\,\phi_1$) is significantly different from zero

[17] This calibration of ϕ is based on the assumption that preferences for children are not affected by preferences about adhering to a social norm on fertility.

Table 1—Calibration of Structural Parameters

	Value	Description/Source
Parameters		
ρ	0.66	Productive labor share of output
g_A	1.56	TFP growth (Young 1995)
τ_1	0.003	Household education spending per child as a fraction of expenditure
g_H	2.886	Targeted to match per capita output growth and population growth
τ_0	0.025	Targeted to match household education expenditure
τ_h	0.006	Targeted to match public expenditure on education in 1960
α	0.1987	Targeted to match household consumption-income ratio in 1960
θ	0.1312	Targeted to match expected years of schooling in 1960
ϕ	0.204	Targeted to match TFR in 1960
φ	0.1	Disutility from deviating from social norm on fertility
n^*	1	Replacement rate of fertility
Initial conditions		
$\bar{H}$	1	Labor endowment
n_0	3.5	Fertility rates in developing countries in early twentieth century (Therborn 2004)
z	0.474	Targeted to match Mincerian return of 0.11
H_0	0.935	Obtained from human capital production function, given g_H

Note: The table reports the calibrated parameter values and initial conditions and the sources from which they are obtained.

and that the obtained value has the expected sign and magnitude (see Table 2).[18] We calculate ϕ at the sample average of total spending, P, to obtain a value of 0.62, which implies that the weight on n^* has tripled over the past fifty years.

III. Results

The dynamics of fertility and human capital accumulation in the economy are governed by equations (5), (11), (13), and (14).[19] We use the calibrated model to investigate how the two channels in our model—human capital accumulation and the presence of social norms on fertility—contribute to the fertility decline. We begin from an initial level of human capital stock and fertility and examine the transition to a steady state.

We start by analyzing a baseline model in which individuals do not care about social norms $(\varphi = 0)$ and the only mechanism by which fertility falls is the faster accumulation of human capital. We compare this model with our extended model of fertility and social norms. We consider two cases: in the first case, ϕ remains unchanged over time; in the second case, ϕ rises to the value estimated in the previous section (referred to as the model with policy changes). Since the estimated values are for 2010, we set ϕ in 1985 to be in between the values of the initial calibration for 1960 and the estimated value for 2010. We do not impose any changes to the parameters after the third period.

Figure 3 shows the model's predicted path of fertility (multiplied by 2 to make it comparable with the data) and investment in education (measured in years of education) under the different versions outlined above. The corresponding values

[18] This also indicates that our choice of 0.1 for φ is not unreasonable.

[19] Note that since neither first-order condition depends on w_t, the production side of the economy doesn't affect the dynamics of fertility and human capital.

TABLE 2—ESTIMATION OF ϕ

Parameter	Value
ϕ_1	0.167
	(0.000)
$\phi\ (= \phi_1 \bar{P})$	0.626
Observations	53
R^2	0.699

Notes: The table reports the results from estimating equation (15). The estimation is carried out using data on fertility and years of schooling for 2010, and the average annual per capita spending on family planning over the 1970–2000 period. ϕ is calculated as $\phi = \phi_1 \bar{P}$, where $\bar{P}$ is the sample average of per capita spending on family planning. The value in parentheses is the *p*-value of the regression coefficient from which the value for ϕ_1 is backed out and is based on robust standard errors.

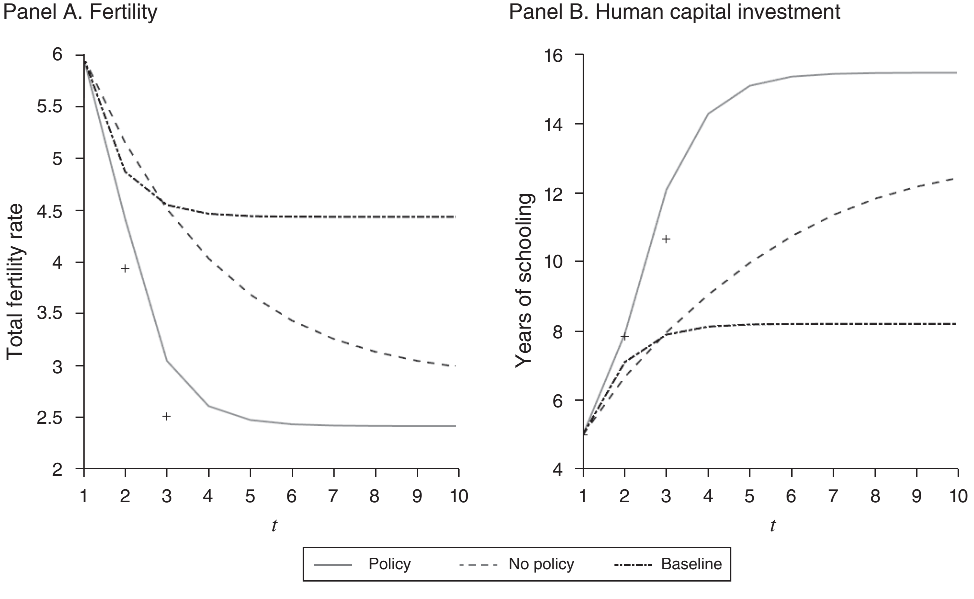

FIGURE 3. TRANSITION TO STEADY STATE

Notes: The figure plots the path of fertility and investment in education for the different versions of the model. The dash-dot line corresponds to the baseline model where $\varphi = 0$. The dashed line represents the case where ϕ and φ remain unchanged over time, while the solid line represents changes in ϕ to 0.4 and 0.62 at $t = 2$ and $t = 3$, respectively. The points marked by "+" refer to the values observed in the data where $t = 2$ is 1985 and $t = 3$ is 2010.

in the data (only available for the first three periods for fertility and education) are marked by crosses.

The baseline model (given by the dash-dot line), in which individuals do not care about norms, generates a small decline in fertility. TFR falls to 4.8 in $t = 2$ and reaches a steady-state value of around 4.4 children per woman while investment in education rises to 7 years of schooling in $t = 2$ and reaches a steady state of roughly

8.2. The inclusion of social norms on fertility generates a larger decline in fertility in the long term, even when ϕ remains unchanged, though this decline occurs at slower pace. In this case, TFR falls from 6 children per woman to 3.4 within six generations and a steady state of 2.9 is reached after approximately twelve periods. At the same time, human capital investment reaches a steady state of around 13 years of schooling. The existence of endogenously evolving social norms on fertility is enough to generate a decline in fertility that is much larger than the decline generated by the baseline model.

We next consider the effect of the population control policies (given by the solid line), which we interpret as an increase in ϕ. As can be expected, the increase in ϕ (a larger weight placed on the replacement level of fertility) generates a much larger decline in fertility, increase in education, and a quicker convergence to the steady state. We allow ϕ to rise from 0.2 in $t = 1$ to 0.4 and then 0.62 in the two subsequent periods, which corresponds to a change in the norm on number of children from 6 children in the initial period to around 3 by $t = 3$. Accordingly, the model predicts a decline in TFR to 3 at $t = 3$ and fertility reaches a steady state of around 2.4 after eight periods. At the same time, years of schooling rises from 5 to 12 in just three generations.

Comparing the results of the model with the data indicates that the inclusion of social norms with an increase in ϕ over time improves the predictions of fertility and the number of years of schooling considerably. The model is able to replicate the patterns for both fertility and years of schooling in the second period very well; in the third period, both model-generated variables fall just slightly above the data. In particular, the predicted steady-state level of fertility is very close to the currently observed level of fertility. Note that we do not allow ϕ to change after $t = 3$. If we allowed ϕ to increase continuously over time, convergence to a steady state fertility rate of 2 children per woman would be even faster.

The changes in ϕ which would be required to exactly match the data would be an increase to 0.5 in $t = 2$ and then to 0.85 by $t = 3$. While we estimate the change in ϕ captured by spending on family planning programs, it is likely that when taking into account other factors such as increased access to mass media and modernization, the actual increase in ϕ is larger than that estimated in this paper.

To summarize, this quantitative exercise points to the importance of changing social norms on family size for the decline in fertility observed in developing countries over the past few decades. We use data on family planning program funds to capture the change in social norms brought about by these programs which were widely adopted in developing countries during this period. The results suggest that the change in social norms brought about by these programs considerably accelerated the fertility decline. This is consistent with empirical studies that find evidence of the effectiveness of public persuasion measures in reducing fertility (La Ferrara, Chong, and Duryea 2012; and Bandiera et al. 2014).

A. *Individual Country Simulations*

As an additional test, we now apply the model to individual countries for which sufficient data is available. Focusing on countries with at least 10 data points for spending on family planning programs (the exceptions are the sub-Saharan countries for which fewer data points are available) as well as data on the other macroeconomic variables required for calibration, we use a sample of 15 countries. The spending on family planning programs in these countries range from zero in Benin to constant 2005 US$0.93 in Indonesia to US$3.14 in El Salvador, while the decreases in fertility between 1960 and 2010 range from 1.2 births per woman in Benin to 4.9 in Tunisia and Korea.

We then recalibrate the model, country by country, using country-specific data on the required macroeconomic variables and country-specific estimates of the Mincerian coefficient compiled by Psacharopoulos and Patrinos (2004).[20] The only parameter values that we use from the original calibration are the labor income share (ρ), the cost of educating a child for a household, τ_1, and the technology growth rate (g_A). Since we can no longer use a regression to obtain the individual values to which ϕ rises, we simply back it out from equation (15) using data on fertility and years of schooling for each of the 15 countries for 2010.[21]

Using the reparametrized model, we simulate the path of fertility for each of the 15 countries. The plots in Figure 4 show the model's predictions together with the data. As the figure illustrates, the model does reasonably well at predicting a significant part of the fertility decline in most countries, with three key exceptions: Thailand, Indonesia, and Tunisia. In these cases, the model under-predicts the decline in fertility indicating that the change in ϕ (the policy parameter) was not sufficiently large.[22]

It is worth noting that the deviations from the model's predictions are not related to the level of spending on family planning programs. For instance, Tunisia and Thailand were countries in which strong government-led family planning programs were implemented. But so were South Korea and India, where the model performs remarkably well. The model also does reasonably well at predicting fertility transitions in the sub-Saharan African countries, where family planning programs were introduced much later. It is more likely that when the model deviates substantially from the data, it does so because spending on family planning programs is an inadequate proxy for the effectiveness of the program. Or, in other words, because the spending data we have does not cover the full effort put on fertility reduction.

It is possible, of course, that there are other confounding factors that played a role in Tunisia, Thailand, and Indonesia (e.g., media spillovers from neighboring countries) which we have not been able to pin down with our model. However, in all, given the data limitations, the model matches the fertility transitions in most economies quite well.

[20] The results are not so different when we use a flat rate of 11 percent for the Mincerian coefficient for all countries.

[21] In doing so, three countries (Singapore, South Korea, and Thailand) record values of ϕ greater than one or less than zero. As such, we impose an upper bound of one and a lower bound of zero for the calibrated value of ϕ. The full set of recalibrated parameters for each country is available in Appendix A4.

[22] In fact, for Thailand, the calibration strategy results in a decrease in ϕ rather than an increase.

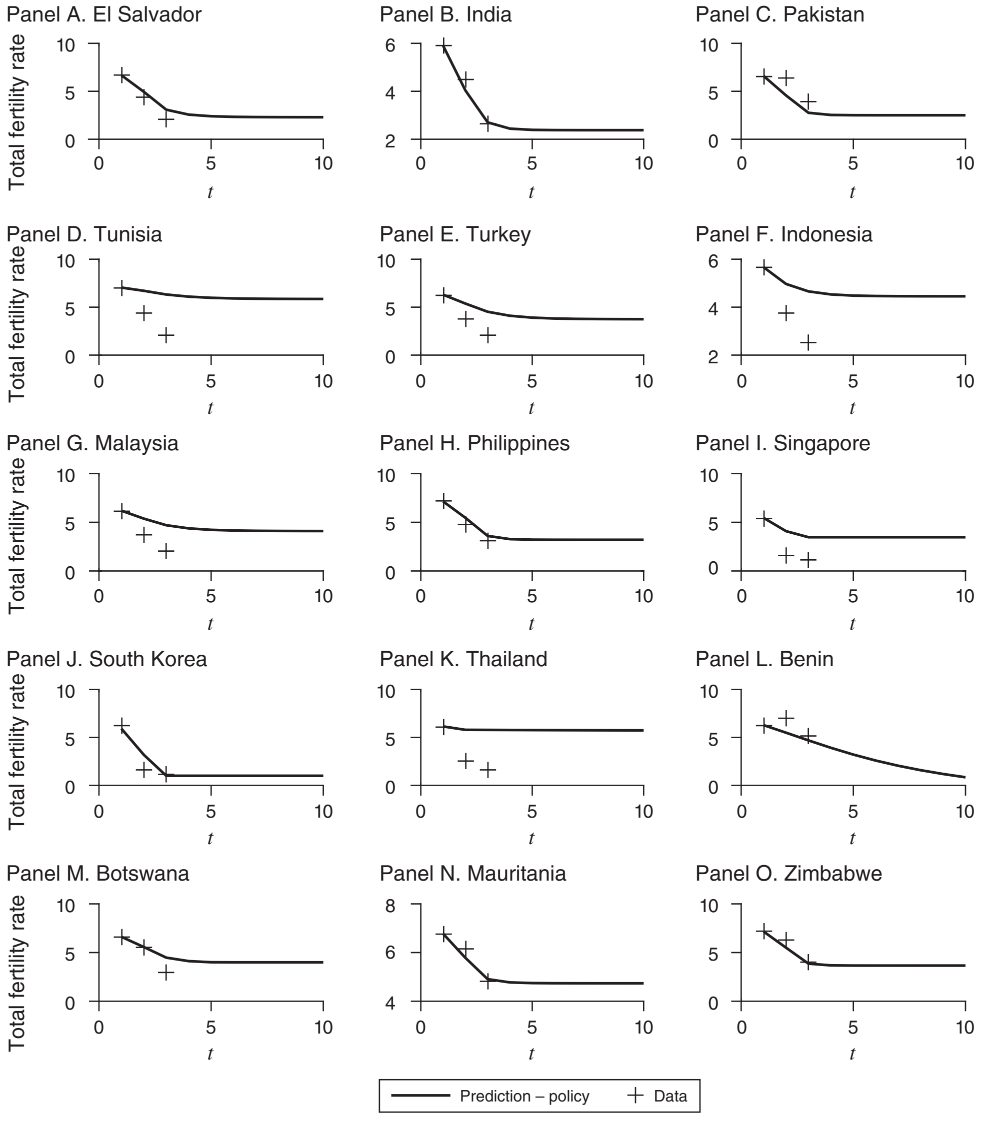

Figure 4. Fertility Transitions

Note: The figure plots the simulated fertility transition for each country against the data (the points marked by "+" refer to the values observed in the data where $t = 2$ is 1985 and $t = 3$ is 2010).

B. *Out-of-Sample Fit*

As another means of validating the model's use in measuring the impact of population policy interventions, we consider the model's predictions in a completely different setting—the fertility transition of the advanced Western economies. Given the absence of population policy interventions in these economies in that period, we compare the predictions of the model with norms but no policy intervention with the fertility rates observed in advanced European and North American economies between 1850 and 1950.[23]

We start by recalibrating the model to match the average fertility and years of schooling in these countries in 1850.[24] Historical data on fertility is obtained from Gapminder.org, data on GDP per capita and population is obtained from the Maddison Project database (2018), and data on years of schooling is taken from the Lee and Lee Long-Run Education dataset (2016). Given that data limitations do not permit the recalibration of all parameters, we keep the labor income share (ρ), technology growth rate (g_A), and the fraction of income required to raise a child (τ_0) unchanged from our main exercise. Given the very low level of expected years of schooling (average years of schooling in 1870 is less than 2), we set public spending on education to zero (which means $\tau_g = 0$) and the Mincerian coefficient to 0.2. The initial conditions and parameter values used for this exercise are given in Table 3.

We then simulate the fertility transition of the advanced economies between 1850 and 1950 using the model with norms but no policy intervention. Figure 5 plots the predictions of the model against the data.

We find that the predictions of the model, calibrated to match the initial conditions in advanced Western economies in 1850, fits the data well. The predictions are in line with the slow decline in fertility that took place in these countries during this period, in which fertility-related policy intervention was minimal. The only part of the transition the model does not pick up is the rapid decline in fertility observed during the first half of the twentieth century but given the occurrence of World War I and World War II, it is unsurprising that the model cannot replicate fertility trends during this period. However, the model performs very well for the second half of the nineteenth century, and its TFR prediction for 1950 is back in line with the data. The model's fit with this out-of-sample data confirms, to us, the credibility of using the model to measure the effect of the population-control policies implemented in the developing economies.

IV. Extensions and Robustness Checks

In this section we discuss a number of extensions of the model. First, we extend the model to introduce a role for declining infant and child mortality in the fertility fall. Next, we incorporate imperfect control over fertility to study the role of improvements in contraceptive technologies. Finally, we carry out two robustness checks that consider the effect of changing the value of the coefficient of disutility from norm deviation, φ, and the effect of changing the specification of disutility from deviating from the norm, allowing upward and downward deviations to be treated asymmetrically.

[23] We limit the comparison to this period given the developments in the technology of modern contraceptives such as the oral contraceptive pill during the 1960s and the establishment of the Population Council and the International Planned Parenthood Federation in the 1950s.

[24] See Appendix A5 for the list of countries used for this analysis.

Table 3—Calibration of Structural Parameters for Out-of-Sample Exercise

	Value	Description/Source
Parameters		
ρ	0.66	Same as original model
g_A	1.56	Same as original model
τ_0	0.025	Same as original model
g_H	1.446	Targeted to match per capita output growth and population growth between 1850 and 1875
τ_1	0.032	Targeted to match consumption-income ratio of 0.8 in 1850
α	0.2554	Targeted to match TFR of 4.8 in 1850
θ	0.3218	Targeted to match 2 years of schooling in 1850
ϕ	0.035	Targeted to match fertility rate of 4.9 in 1825
φ	0.1	Disutility from deviating from fertility norm
n^*	1	Replacement rate of fertility
Initial conditions		
$\bar{H}$	1	Labor endowment
n_0	2.45	Average fertility rate of 4.9 in 1825
z	0.4315	Targeted to match Mincerian return of 0.2
H_0	0.7166	Obtained from human capital production function, given g_H

Note: The table reports the calibrated parameter values and initial conditions and the sources from which they are obtained.

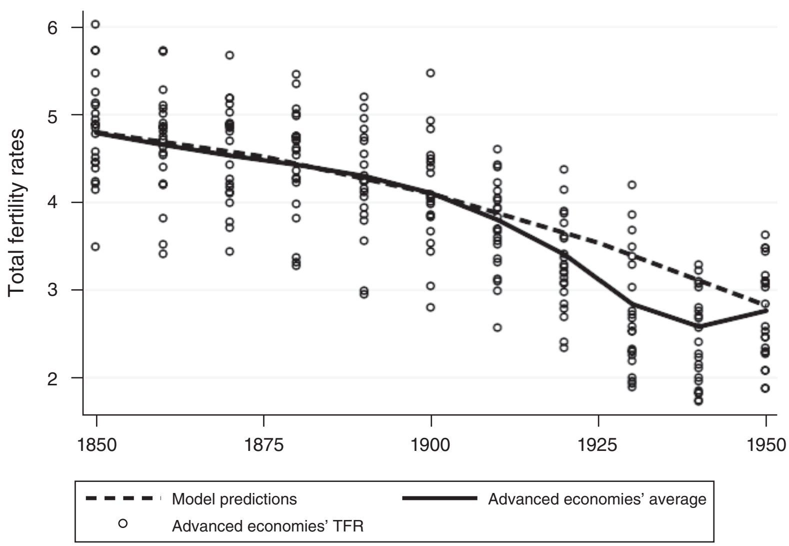

Figure 5. Fertility Transition in Advanced Economies

Notes: The figure plots fertility rates in advanced economies from 1850 and 1950 and the predictions of the model with norms but no policy intervention. Fertility data on the 22 European and North American countries between 1850 and 1950 are obtained from Gapminder.org.

A. *Including Mortality*

The model presented in the previous section did not take into account the mortality decline observed in developing countries during this period. In this section, we extend our model to include uncertainty regarding the number of children who survive to adulthood. We then investigate the impact of an increase in survival rates on fertility and human capital investment. We follow Kalemli-Ozcan (2003) in how we incorporate mortality into the model.[25]

Parents choose the number of children, n_t, but only N_t of the infants survive to childhood and all children survive to adulthood. Parents spend on rearing and educating their surviving children and derive utility from the quantity and quality of these children.[26] In addition, parents care about how the number of their surviving children compares with the social norm on family size. The utility function for an adult of generation t can then be written as

$$(16)\qquad E_t U_t = E_t\left\{\ln C_t + \alpha \ln N_t + \theta \ln\left[w_{t+1}\left(\bar{H} + H_{t+1}\right)\right] - \varphi\left(N_t - \hat{N}_t\right)^2\right\},$$

where $\hat{N}_t = \phi n^* + (1-\phi) N_{t-1}$ is the norm on family size.

Expected utility is maximized subject to

$$(17)\qquad C_t = \left[1 - \tau_g - (\tau_0 + \tau_1 h_t) N_t\right] w_t\left(\bar{H} + H_t\right),$$

and the human capital production function (11).

As in Kalemli-Ozcan (2003), N_t is a random variable drawn from a binomial distribution, with $s_t \in [0,1]$ the survival probability of each infant. We use a second-order approximation of the expected utility function around the mean value of N_t (i.e., $n_t s_t$). The approximated expected utility function is given by

$$(18)\qquad E_t U_t = E_t\Bigg\{ \ln\left[\left(1 - (\tau_0 + \tau_h h_t)\, n_t s_t\right) w_t\left(\bar{H} + H_t\right)\right] + \alpha \ln\left(n_t s_t\right) + \theta \ln\left[w_{t+1}\left(\bar{H} + H_{t+1}\right)\right] - \varphi\left(n_t s_t - \hat{N}_t\right)^2 - \frac{n_t s_t (1 - s_t)}{2}\left[\left(\frac{(\tau_0 + \tau_h h_t)}{(\tau_0 + \tau_h h_t)\, n_t s_t}\right)^2 + \frac{\alpha}{(n_t s_t)^2} + 2\varphi\right]\Bigg\},$$

which incorporates the budget constraint (17). The last three terms represent the disutility arising from uncertainty in the number of infants who survive to adulthood.

[25] In the original Barro-Becker (1989) framework, child mortality is modeled as an explicit cost of child-rearing. Doepke (2005) studies three variations of this model: a baseline model where fertility choice is continuous and there is no uncertainty over the number of surviving children, which is contrasted with an extension involving discrete fertility choice and stochastic mortality and another with sequential fertility choice. He finds that while the total fertility rate falls as child mortality declines in each model, the number of surviving children increases, and concludes that factors other than declining infant and child mortality were responsible for the fertility transition observed in industrialized countries.

[26] This is a slight deviation from Kalemli-Ozcan (2003) where education is provided before the uncertainty is realized.

The first-order conditions for fertility and human capital investment become:

$$(19)\quad \frac{\alpha}{n_t}\left(1+\frac{(1-s_t)}{2n_ts_t}\right)$$
$$= 2\varphi s_t\left(n_ts_t-\hat{N}_t\right)+\varphi s_t(1-s_t)$$
$$+\frac{(\tau_0+\tau_h h_t)s_t}{1-(\tau_0+\tau_h h_t)n_ts_t}\left[1+\frac{1+(\tau_0+\tau_h h_t)n_ts_t}{2\left(1-(\tau_0+\tau_h h_t)n_ts_t\right)}\frac{(\tau_0+\tau_h h_t)(1-s_t)}{\left(1-(\tau_0+\tau_h h_t)n_ts_t\right)}\right],$$

$$(20)\quad \frac{\theta z_t\left(\bar{H}+H_t\right)}{\left(\bar{H}+H_{t+1}\right)} = \frac{\tau_h n_t s_t}{\left(1-(\tau_0+\tau_h h_t)\,n_ts_t\right)}\left[1+\frac{(\tau_0+\tau_h h_t)(1-s_t)}{\left(1-(\tau_0+\tau_h h_t)\,n_ts_t\right)^2}\right].$$

The key difference between this setup and that in Section I is that there is now an additional term in the marginal cost of both fertility and schooling, which reflects the cost of uncertainty.

Calibration and Results.—The calibration exercise is carried out in the same way as before: we start from a model with mortality and no norms to back out all the parameters except ϕ and then use the extended model with norms and mortality to get an initial value for ϕ. We use the mortality rate for children below 5 years of age (measured as the number of deaths of children below 5 years of age per 1,000 live births) for developing countries in 1960 (from the WDI database) as a measure of $1-s_t$. The recalibration causes τ_0, τ_h, and θ to change slightly (to 0.021, 0.007, and 0.1504 respectively) while ϕ changes substantially to 0.02, much lower than 0.2 in the model without mortality.

To identify the change in ϕ over the past two periods, we carry out the same estimation exercise as before, again setting $\phi = \phi_1\bar{P}$ but now using equation (19). We see a much larger increase in the value of ϕ, in both absolute and relative terms, than in the model without mortality. Table 4 shows the values of the parameters obtained from the estimation.

We then plot the transition paths of fertility and human capital to their steady states for three cases: the baseline model with no norms or mortality (given by the dash-dot line), the model with falling mortality rates and no norms (given by the dashed line), and the extended model of mortality and social norms (given by the solid line). We allow s_t to rise over time from 0.77 in $t = 1$ to 0.90 and 0.96 in $t = 2$ and $t = 3$ as seen in the data. As before, since the estimation of ϕ is for 2010, the value of ϕ for 1985 is set to be in between the values of the initial calibration for 1960 and the estimate for 2010 and do not change after the third period.

As Figure 6 shows, the incorporation of mortality into the baseline model generates a slightly larger decline in fertility than the baseline model which only includes human capital accumulation with TFR converging to around 3.5 births per woman rather than 3.8. On the other hand, the incorporation of mortality into the baseline model results in a smaller predicted increase in years of schooling than the baseline. This is because the decline in the number of surviving children—

TABLE 4—ESTIMATION OF ϕ WITH MORTALITY

Parameter	Value
ϕ_1	0.141 (0.000)
$\phi\ (= \phi_1 \bar{P})$	0.558
Observations	52
R^2	0.576

Notes: The table reports the results from estimating equation (19). The estimation is carried out using data on fertility, child mortality rates, and years of schooling for 2010, and the average annual per capita spending on family planning over the 1970–2000 period. ϕ is calculated as $\phi = \phi_1 \bar{P}$, where $\bar{P}$ is the sample average of per capita spending on family planning. The value in parentheses is the p-value of the regression coefficient from which the value for ϕ_1 is backed out and is based on robust standard errors.

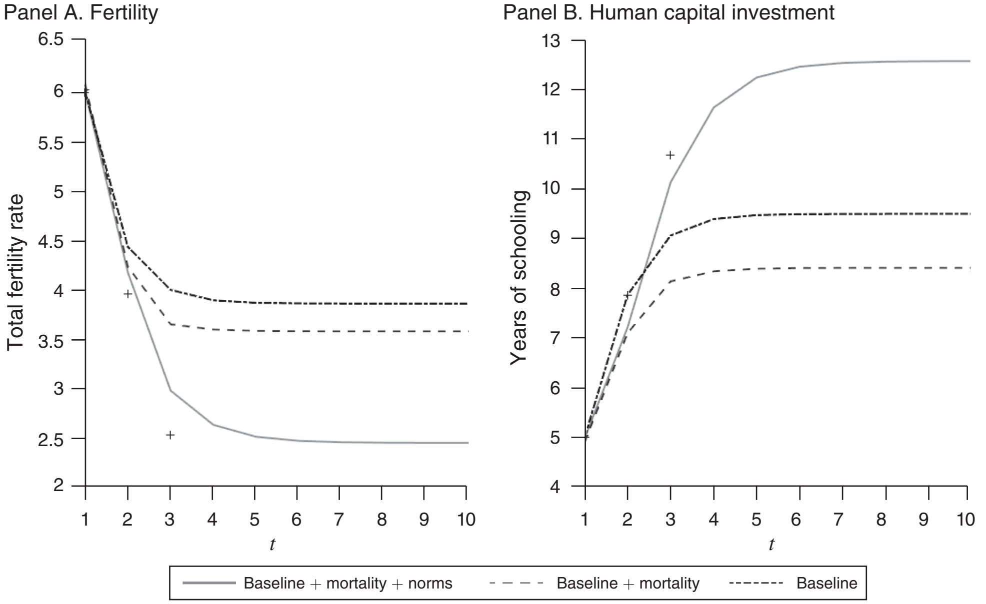

FIGURE 6. INCORPORATING MORTALITY

Notes: The figure plots the path of fertility and investment in education in the three versions of the model. The dash-dot line represents the baseline model with no mortality or social norms while the dashed line represents the baseline model augmented to include mortality where s_t rises to 0.9 at $t = 2$, and to 0.96 at $t = 3$, where it remains in all successive periods. The solid line represents the model with mortality and social norms. Here, s_t rises as described earlier while ϕ rises to 0.3 and 0.55 in the second and third periods, respectively. The points marked "+" refer to the values observed in the data.

which is the value on which years of schooling is determined—is actually larger in the baseline model, even though the level of fertility is higher (see Figure 7). In the baseline model that incorporates the mortality decline, the number of surviving children drops from 4.7 to just 3.4 compared to the decline from 5.9 to 3.8

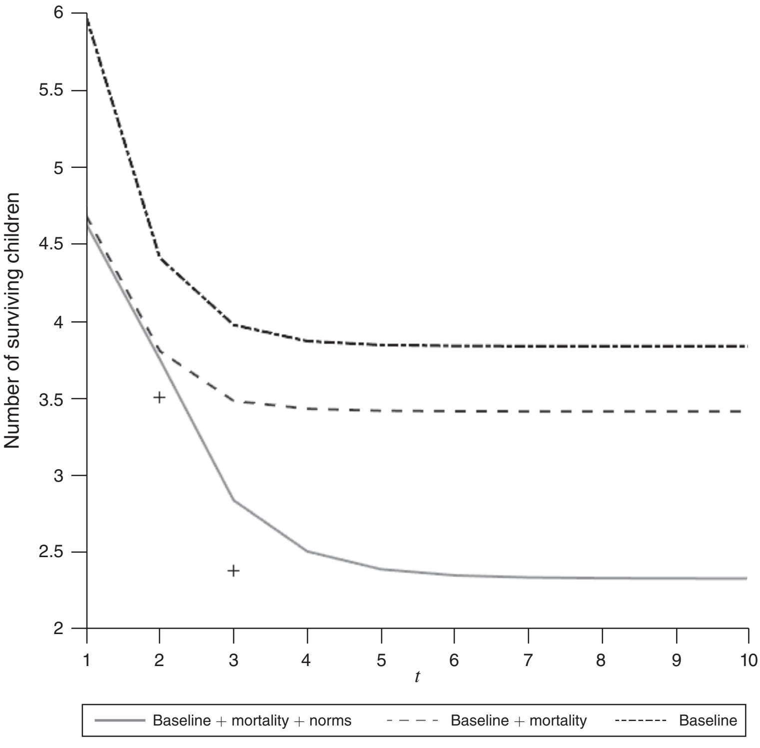

Figure 7. Number of Surviving Children

Notes: The figure plots the number of surviving children predicted by the three versions of the model. The dash-dot line represents the baseline model with no mortality or social norms while the dashed line represents the baseline model augmented to include mortality where s_t rises to 0.9 at $t = 2$ and to 0.96 at $t = 3$, where it remains in all successive periods. The solid line represents the model with mortality and social norms, where ϕ rises to 0.3 and 0.55 in the second and third periods, respectively.

in the baseline model without mortality. By contrast, including a social norm that falls over time generates a large decline in the number of surviving children—a drop from 4.6 to 2.3. Given that the investment in schooling is made for surviving children, a smaller decline in surviving children leads to a smaller increase in the years of schooling.

Our modeling of mortality, which is based on Kalemli-Ozcan (2003), generates a hoarding effect, where the risk of child mortality results in a precautionary demand for children. The decline in fertility generated by the decline in social norms is marginally smaller than that in the model described in the previous section because of this effect. However, the simulations clearly show that it is the presence of changing social norms that significantly accelerates the fertility fall, indicating that the mortality transition cannot rule out the role of the population control policies in the fertility decline. Taken as a whole, we would argue that while the decline in mortality rates did play an important role in triggering the introduction of population-control policies, its role in precipitating the fast fall in fertility through individual responses, without the policy intervention, is less clear.

B. *Incorporating Unwanted Fertility*

So far we have simulated the effect of population control policies on the fertility decline by focusing on their role in changing the norm on family size. We now extend the model such that individuals do not perfectly control fertility. In other words, we allow the lack of contraceptive technologies to cause a discrepancy between the desired and actual number of children.[27] This allows us to examine the impact of a reduction in unwanted fertility caused by the introduction of widespread modern contraceptives, which was the second main component of the population control policies.

We do not explicitly model the choice of contraceptive usage (see, for example, Cavalcanti, Kocharkov, and Santos 2017) but consider individuals' ability to control fertility to be exogenously determined. So while the production side of the model is the same as before, we now assume that parents' inability to perfectly control their fertility leads to a distinction between the desired or chosen number of children, n_t^d, and the actual number of children, n_t^a. Specifically,

$$n_t^a = n_t^d + \varepsilon_t,$$

where ε_t is a stochastic error term causing the desired number of children, n_t^d, to differ from the actual number of children, n_t^a.

Individuals now have to maximize expected utility, which, for an adult of generation t, is given by

$$(21) \qquad E_t U_t = E_t\Big[\ln C_t + \alpha \ln n_t^a + \theta \ln\big[w_{t+1}\big(\bar{H} + H_{t+1}\big)\big] - \varphi\left(n_t^a - \hat{n}_t\right)^2\Big],$$

where E_t denotes expectations as of time t.

Individuals maximize expected utility with respect to the human capital production function (same as before) and the budget constraint, which is now changed slightly to

$$(22) \qquad C_t = \left[1 - \tau_g - \left(\tau_0 + \tau_1 h_t\right) n_t^a\right] w_t\big(\bar{H} + H_t\big).$$

The formulation of the expected utility function requires some distributional assumptions about unwanted fertility, ε_t. The data on wanted fertility rates in developing countries (obtained from Demographic and Health Surveys) indicates that ε_t is usually positive and has a positively skewed distribution. We assume that ε_t follows a Poisson distribution with mean λ. Thus, a reduction in λ translates to a reduction in uncertainty as well as average unwanted fertility. We then carry out a second-order approximation of the expected utility around the mean of unwanted

[27] The key difference between this and the mortality extension is that now individuals face the risk of overshooting their desired number of children whereas in the case of uncertainty about mortality, individuals faced the risk of ending up with fewer children than they wanted.

fertility. Substituting in the budget constraint and human capital production function, the household problem can be rewritten as

$$\{n_t^d, h_t\} = \operatorname{argmax} \Bigg\{ \ln\Big[\Big(1 - (\tau_0 + \tau_h h_t)\big(n_t^d + \lambda\big)\Big) w_t\big(\bar{H} + H_t\big)\Big] + \theta \ln\Big[W_{t+1}\big(\bar{H} + z_t\big(\bar{H} + H_t\big) h_t\big)\Big] + \alpha \ln\big[n_t^d + \lambda\big] - \varphi\big(n_t^d + \lambda - \hat{n}_t\big)^2 - \frac{\lambda}{2}\left[\frac{(\tau_0 + \tau_h h_t)^2}{\Big(1 - (\tau_0 + \tau_h h_t)\big(n_t^d + \lambda\big)\Big)^2} + 2\varphi + \frac{\alpha}{\big(n_t^d + \lambda\big)^2}\right] \Bigg\} \tag{23}$$

subject to: $\big(n_t^d, h_t\big) \geq 0$.

The first-order conditions for n_t^d and h_t are given by

$$\frac{\alpha}{n_t^d + \lambda} = \frac{(\tau_0 + \tau_h h_t)}{\Big(1 - (\tau_0 + \tau_h h_t)\big(n_t^d + \lambda\big)\Big)} + 2\varphi\big(n_t^d + \lambda - \hat{n}_t\big) + \lambda\left[\frac{(\tau_0 + \tau_h h_t)^3}{\Big(1 - (\tau_0 + \tau_h h_t)\big(n_t^d + \lambda\big)\Big)^3} - \frac{\alpha}{\big(n_t^d + \lambda\big)^3}\right], \tag{24}$$

$$\frac{\theta z_t\big(\bar{H} + H_t\big)}{\big(\bar{H} + H_{t+1}\big)} = \frac{\tau_h\big(n_t^d + \lambda\big)}{\Big(1 - (\tau_0 + \tau_h h_t)\big(n_t^d + \lambda\big)\Big)} + \lambda\left[\frac{\tau_1(\tau_0 + \tau_h h_t)}{\Big(1 - (\tau_0 + \tau_h h_t)\big(n_t^d + \lambda\big)\Big)^3}\right], \tag{25}$$

where the last term on the right-hand side in equation (25) reflects the cost of uncertainty. Since parents derive utility from all children (unwanted or not), the second line in equation (24) reflects the cost of uncertainty adjusted for the gain in utility caused by having an extra child.

Calibration and Results.—The calibration strategy follows the same procedure as the main model, leaving parameters α, θ, τ_0, τ_1, g_H, ρ, and n^* and the initial conditions unchanged. However, ϕ needs to be recalibrated using equation (24) for given values of φ and λ. The parameter λ is chosen using data on wanted fertility rates obtained from Demographic and Health Surveys that start in the late 1980s. Unwanted fertility (calculated as the difference between TFR and wanted fertility rate) is around 1 birth, on average, in the 1980s. Since this is well after the introduction of the oral contraceptive pill and the implementation of many family planning programs worldwide, we set initial λ to 1 (reflecting an average of 2 unwanted births). We then use equation (24) to obtain the value of ϕ, with φ set to 0.1 as before. This gives us $\phi = 0.22$, which is very close to the value obtained in the main model. As such, we allow ϕ to rise to the same levels estimated in Section IIE.

We then consider two policy experiments using this model: one in which social norms on fertility and unwanted fertility both change and one in which only social norms change. In other words, we allow the weight on the replacement level of fertility, ϕ, to rise in both versions but allow unwanted fertility, λ, to fall only in one. The fall in λ reflects the increased contraceptive prevalence over the past few decades. Using the data on wanted fertility we allow λ to fall from 1 in the first period to 0.55 in the second, 0.31 in the third and then remain at 0.31 in all successive periods. Figure 8 plots the two transition paths.

As seen in Figure 8, both channels play a role in the fertility decline. However, it appears that a large portion of the decline can be explained by the change in social norms alone. The simulations indicate that the change in norms brings down fertility from 6 children per woman to 3.4, which is more than 85 percent of the decline predicted by the model in which both unwanted fertility and social norms change. The change in social norms accounts for less of the increase in years of schooling but still accounts for 75 percent of the total increase predicted by the model in which both parameters change.

The comparison between the two models indicates that changing the norms on fertility has a much larger effect on fertility decisions than merely increasing access to contraception. This is consistent with the fact that many of the family planning programs supplemented their supply-side strategies of increasing access to contraception with large-scale mass media campaigns to promote smaller family sizes. This point was made by demographers Enke (1960) and Davis (1967) at early stages of the global population control movement, and later by Becker (1992), who argued that family planning programs focused on increasing contraceptive usage are effective only when the value of having children is lowered. The result is also consistent with Cavalcanti, Kocharkov, and Santos (2017), who find that aggregate fertility is unresponsive to improved contraceptive access even though there are significant compositional differences between education groups.

C. *Sensitivity to Choice of φ*

In all of the simulations we carried out, the value of φ, which measures how much individuals dislike deviating from family-size norms, was set to 0.1 given the lack of sufficient moments in the data. We now consider the sensitivity of our results to this choice of φ by redoing the computations for $\varphi = 0.05$ and for $\varphi = 0.5$. For each case, we re-estimate the value of ϕ for subsequent periods using equation (15).[28]

Table 5 presents the new estimates for ϕ for the different values of φ. The regression results show that the change in φ is compensated by the change in ϕ, though the estimated changes are small. For instance, the change in ϕ when $\varphi = 0.5$ is slightly smaller than the change when $\varphi = 0.05$. As before, we use the re-estimated values of ϕ for the third period and set ϕ in the second period to an in-between value.

[28] Using the values of ϕ estimated in Section IIE rather than these re-estimated values has hardly any effect on the results.

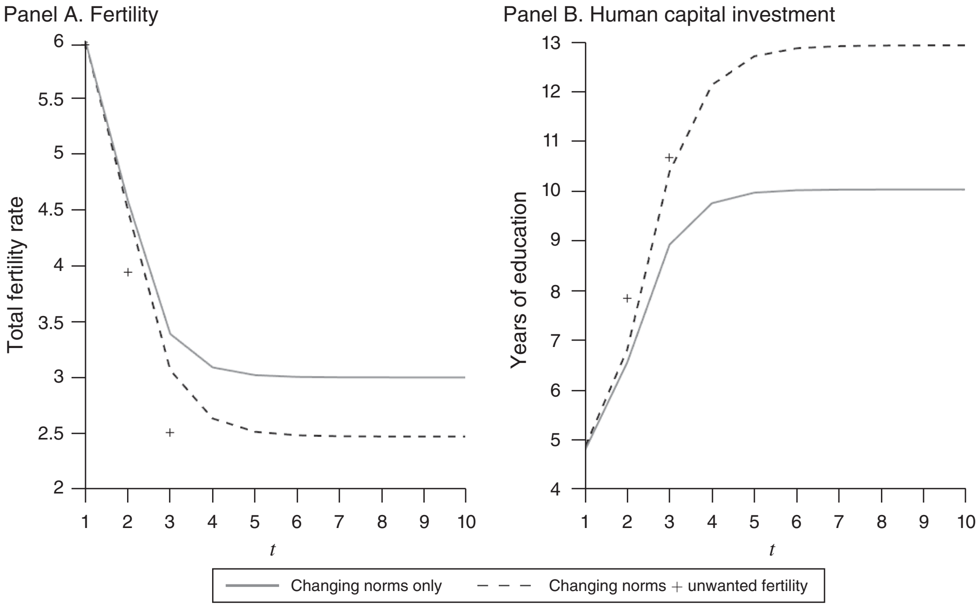

Figure 8. Incorporating Unwanted Fertility

Notes: The figure plots the path of fertility and investment in education in two models. The solid line represents the model where only the social norm changes while the dashed line represents the model where both social norms and unwanted fertility change. In both models ϕ rises to 0.4 and then 0.62 in the second and third periods. In the model where unwanted fertility also changes, λ falls from 1 in the first period to 0.55 in the second, and 0.31 in the third, where it remains in all successive periods. The points marked by "+" refer to the values observed in the data.

Table 5—Estimation of ϕ for Different Values of φ

	Value	
Parameter	$\varphi = 0.05$	$\varphi = 0.5$
ϕ_1	0.18	0.16
	(0.000)	(0.000)
$\phi\ (= \phi_1 \bar{P})$	0.66	0.59
Observations	53	53
R^2	0.532	0.837

Notes: The table reports the results from estimating equation (15) for different values of φ. The estimation is carried out using data on fertility and years of schooling for 2010, and the average annual per capita spending on family planning over the 1970–2000 period. ϕ is calculated as $\phi = \phi_1 \bar{P}$, where $\bar{P}$ is the sample average of per capita spending on family planning. Values in parentheses are *p*-values of the regression coefficients from which the values for ϕ_1 are backed out and are based on robust standard errors.

The transition paths of fertility and investment in human capital to their steady-state values under the alternative values for φ are plotted in Figure 9. The figure shows that the results do not vary much in response to φ; the transition path for the first three periods is virtually the same under all three scenarios. The key difference is in the steady-state values to which fertility and schooling converge. The higher the coefficient of disutility from deviating from social norms on fertility, the lower the steady-state level of fertility and the higher the steady-state investment in human capital. However, moving from $\varphi = 0.05$ to $\varphi = 0.5$, a tenfold increase, results in

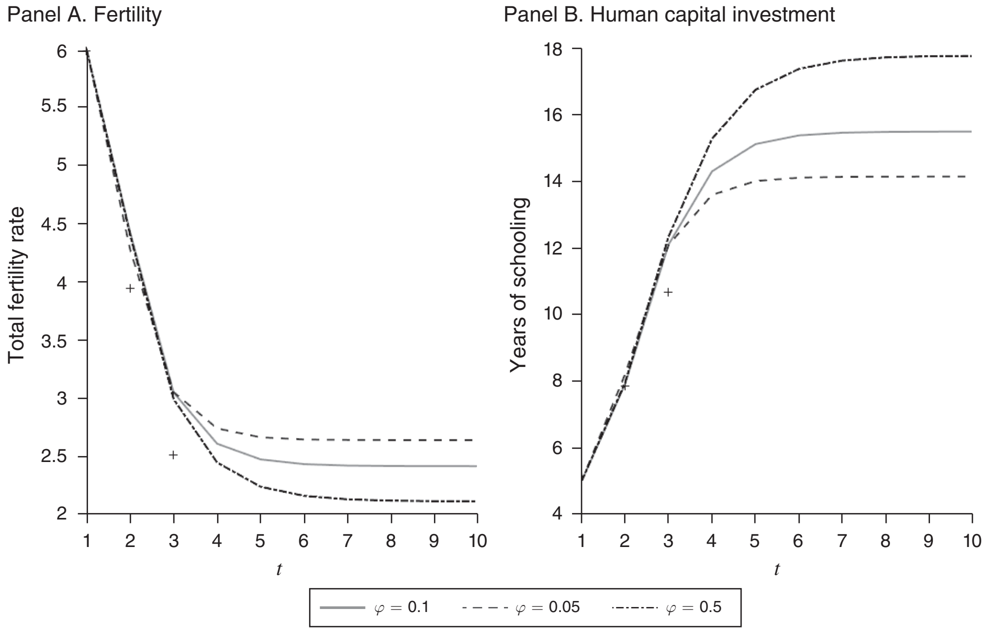

FIGURE 9. ALTERNATIVE VALUES FOR φ

Notes: The figure plots the path of fertility and investment in education in the full model under different values of φ: 0.05, 0.1, and 0.5, corresponding to the dashed, solid, and dash-dot lines, respectively. For each variation, ϕ starts from the same initial value but follows a different path. When $\varphi = 0.1$, ϕ rises to 0.4 and then 0.62 in the second and third periods. When $\varphi = 0.05$, ϕ rises to 0.4 and 0.66, and when $\varphi = 0.5$, ϕ rises to 0.4 and 0.59. The points marked by "+" refer to the values observed in the data.

a reduction of the steady state fertility level of less than 0.5. Furthermore, the fertility decline generated is still much larger than in the baseline model with no norms. Carrying out the simulations for the different values of φ under the same path for ϕ (as estimated in Section IIE) shows a nearly identical picture. It does not appear, therefore, that our results are too reliant on the assumed value of φ.

D. *Functional Form of Disutility from Deviation from the Norm*

We now consider the robustness of our results to an alternative specification for the disutility from deviating from the norm. In particular, we now use a functional form that treats upward and downward deviations from the norm asymmetrically with deviations below the norm being penalized more heavily than deviations above. This would be consistent with societal norms in developing countries where not having children is considered taboo. For this purpose, we set:

$$g(n_t, \hat{n}) = \left[\ln(n_t/\hat{n}_t)\right]^2.$$

The first-order condition for fertility changes to the following:

$$\frac{\alpha}{n_t} = \frac{(\tau_0 + \tau_h h_t)}{\left(1 - (\tau_0 + \tau_h h_t)\, n_t\right)} + 2\varphi \frac{1}{n_t} \ln(n_t/\hat{n}_t), \tag{26}$$

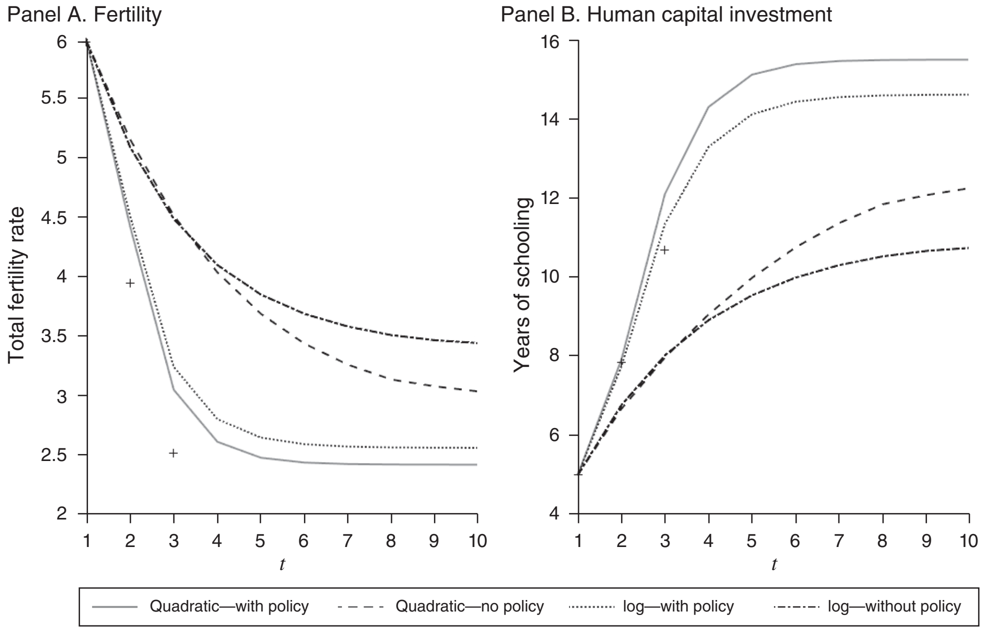

Figure 10. Comparing Functional Forms

Notes: The figure plots the path of fertility and investment in education in the full model under two functional forms: quadratic disutility from norm deviation (main analysis) and log disutility from norm deviation. For each functional form we consider two experiments: one where ϕ rises (to the levels estimated in Section IIE) and the other where it remains unchanged. The solid and dashed lines correspond to quadratic disutility with and without policy changes, respectively. The dotted and dash-dot lines correspond to log disutility with and without policy changes. The points marked by "+" refer to the values observed in the data.

while the first-order condition for human capital investment remains unchanged.

Under the same parameter and initial condition values as in the previous section, we plot the transition paths of fertility and investment in human capital to their steady-state values. We consider two experiments: one in which ϕ increases and the other in which ϕ remains unchanged over time. We compare the results of this model with the results of the main model with quadratic disutility from deviating from the norm.

The results show that the two functional forms yield results that are very similar. The decline in fertility is slightly smaller in the log disutility version (corresponding to the dotted line) with and without the policy change, reflecting the larger penalties for deviating below the norm. Given the slightly lower predicted fertility rates, the years of schooling predicted by the log disutility version are marginally closer to the data than those predicted by the quadratic disutility model.

V. Conclusion

We develop a tractable framework that allows us to quantitatively assess the role played by different mechanisms in the large decline in fertility rates experienced by developing countries over the past decades. Our framework explicitly models the influence of population-control policies aimed at affecting social norms and fostering contraceptive technologies. Population-control policies were put in place by most countries in the world to lower fertility rates by affecting social norms and increasing contraceptive use. These policies, however, were often left out of the analysis in standard macroeconomic models of fertility and development. Our model seeks to bring those policies into the standard framework and analyze their role together with those of other determinants of fertility. To do so, we build on the Barro-Becker framework of endogenous fertility choice, incorporating human capital accumulation and social norms over the number of children. Using data on a number of socioeconomic variables as well as information on funding for family planning programs to parametrize the model, we simulate the implementation of population-control policies. We consider several extensions of the model to assess the robustness of the results. The model suggests that, while a decline in fertility would have gradually taken place as economies moved to higher levels of human capital and lower levels of infant and child mortality, policies aimed at altering the norms on family size played a significant role in accelerating and strengthening the decline.

Appendix

Table A1—Family Planning Funds and Wanted and Unwanted Fertility

	Wanted fertility	Unwanted fertility
ln(average funds per capita)	−0.447	0.005
	[0.140]	[0.059]
ln(GDP per capita)	−0.513	0.141
	[0.282]	[0.105]
IMR	0.029	0.007
	[0.009]	[0.003]
Urban population, percent of total	0.010	−0.006
	[0.012]	[0.005]
Years of schooling of adults	0.168	0.010
	[0.089]	[0.040]
Observations	37	37
R^2	0.637	0.150

Notes: The table reports the results of regressing wanted and unwanted fertility (the latter is defined as the difference between total and wanted fertility rates) on the logged real value of average per capita funds for family planning for the 1970s, 1980s, and 1990s; logged GDP per capita; infant mortality rate; proportion of urban population; and years of schooling of the population aged 25 and more. Data on wanted fertility, which comes from Demographic and Health Surveys, covers different countries in different years, so for each country, we use data from the latest year for which wanted fertility is available (the earliest observation is from 1987 but more than 80 percent of the observations are from after 2000). Since years of schooling is available at five-year intervals, we replace missing values with data from the closest year for which data is published. All regressions include a constant. Per capita funds for family planning are converted to US$ in 2005 before averaging. The values in parentheses are robust standard errors.

Sources: Authors' calculations. Data on total fertility rate, wanted fertility rate, urban population, per capita GDP, and infant mortality rate are from the World Development Indicators. Data on years of schooling are from Barro and Lee (2013). Data on funds for family planning are from Nortman and Hofstatter (1978); Nortman (1982); and Ross, Mauldin, and Miller (1993).

Table A2—Spending on Education

Country	$\tau_1 n_t h_t$	n_t	h_t	τ_1	Year	Source
India	0.026	1.4	10.5	0.002	2007/2008	Tilak 2009
Singapore	0.055	0.6	15.4	0.006	2012/2013	Singapore Dept. of Statistics 2014
Sub-Saharan Africa	0.042	2.75	8.72	0.002	2001–2008	Foko, Tiyab, and Husson 2012
Sri Lanka	0.039	1.71	10	0.002	1980/1981	Department of Census and Statistics
Sri Lanka	0.056	1.22	13.6	0.003	2012/2013	of Sri Lanka 2015
Latin America and the Caribbean	0.019	1.1	13.9	0.001	2010	Regional Bureau of Education for Latin America and the Caribbean 2013
South Korea[a]	0.039	0.61	17	0.004	2012	OECD 2016a, OECD 2016b
Chile[a]	0.037	0.929	15.1	0.003	2012	OECD 2016a, OECD 2016b
Indonesia[a]	0.007	1.22	12.7	0.0005	2012	OECD 2016a, OECD 2016b
Egypt[b]	0.028	1.6	13	0.001	2010	Rizk and Abou-Ali 2016
Jordan[b]	0.068	1.85	13.4	0.003	2010	Rizk and Abou-Ali 2016
Sudan[b]	0.05	2.45	7.3	0.003	2009	Rizk and Abou-Ali 2016

Notes: The table reports the fraction of household expenditure spent on education by households and the backed out value for τ_1, which is the fraction of household expenditure incurred per child per year of education using data for different countries and years. The sources for data on household expenditure on education are given in the last column, while data for the corresponding years on fertility and years of schooling are obtained from the World Development Indicators and Barro-Lee datasets. Given that years of education are published at five-year intervals, we choose data from the closest year for backing out τ_1.

[a] $\tau_1 n_t h_t$ calculated using private spending as a percentage of GDP and household expenditure as a percentage of GDP. Private spending on education excludes expenditure outside educational institutions such as textbooks purchased by families, private tutoring for students, and student living costs so possibly underestimates household spending on education.

[b] Household spending on education is obtained as a fraction of household income rather than expenditure and therefore the obtained values of τ_1 are likely to be understated.

Table A3—Public and Private Spending on Education

	Private education expenditure
Public education expenditure (percent of GDP)	−0.12
	[0.067]
ln(GDP per capita)	−0.63
	[0.35]
Observations	113
Number of countries	39
Country and year fixed effects	Yes

Notes: The table reports the results from regressing private education expenditure (as a percentage of GDP) against public education expenditure (as a percentage of GDP), controlling for GDP per capita. Values in brackets are standard errors. The data covers 39 countries over four years. Therefore, the estimated model is a panel regression with country and year fixed effects.

Sources: Authors' calculations. Data on education expenditure is from the OECD's Education at a Glance (OECDa 2016) and data on GDP per capita is from the WDI.

Table A4—Parameters from Country-Specific Calibration

	τ_0	τ_h	α	θ	ϕ_{60}	ϕ_{2010}	z	H_0
El Salvador	0.038	0.005	0.252	0.175	0.935	0.267	0.094	1.663
India	0.042	0.006	0.242	0.147	0.781	0.231	0.292	0.337
Pakistan	0.028	0.005	0.178	0.085	0.920	0.157	0.899	0.226
Tunisia	0.132	0.006	1.186	0.276	1.000	0.904	0.397	0.004
Turkey	0.065	0.006	0.430	0.154	0.861	0.644	0.465	2.108
Indonesia	0.140	0.007	0.938	0.231	0.733	0.731	0.558	0.005
Malaysia	0.077	0.005	0.515	0.155	0.839	0.669	0.586	9.772
Philippines	0.047	0.005	0.455	0.212	1.000	0.253	0.201	128.393
Singapore	0.120	0.005	0.697	0.144	0.691	0.000	1.279	151.597
South Korea	0.000	0.005	0.080	0.150	0.831	0.000	1.021	102.601
Thailand	0.048	0.003	0.233	0.066	0.829	1.000	0.824	0.041
Benin	0.035	0.010	0.226	0.227	0.856	0.856	0.265	0.022
Botswana	0.075	0.008	0.661	0.247	0.923	0.494	3.694	186.758
Mauritania	0.163	0.006	1.644	0.337	0.955	0.279	0.289	0.256
Zimbabwe	0.082	0.004	0.587	0.136	1.000	0.178	0.313	3.753

Note: The table shows the calibrated values for the parameters and initial conditions for each of the 15 countries.

Table A5—Advanced Economies Sample

Australia	Italy
Austria	Luxembourg
Belgium	Netherlands
Canada	New Zealand
Denmark	Norway
Finland	Portugal
France	Spain
Germany	Sweden
Greece	Switzerland
Iceland	United Kingdom
Ireland	United States

REFERENCES

Akerlof, George A. 1997. "Social Distance and Social Decisions." *Econometrica* 65 (5): 1005–27.

Albanesi, Stefania, and Claudia Olivetti. 2016. "Gender Roles and Medical Progress." *Journal of Political Economy* 124 (3): 650–95.

Bandiera, Oriana, Niklas Buehren, Robin Burgess, Markus Goldstein, Selim Gulesci, Imran Rasul, and Munshi Sulaiman. 2014. "Women's Empowerment in Action: Evidence From A Randomised Control Trial In Africa." STICERD – Economic Organisation and Public Policy Discussion Papers Series 50.

Barro, Robert J., and Gary S. Becker. 1989. "Fertility Choice in a Model of Economic Growth." *Econometrica* 57 (2): 481–501.

Barro, Robert J., and Jong Wha Lee. 2013. "A New Data Set of Educational Attainment in the World, 1950–2010." *Journal of Development Economics* 104: 184–98.

Becker, Gary S. 1960. "An Economic Analysis of Fertility." In *Demographic and Economic Change in Developed Countries*, edited by Universities-National Bureau, 209–40. New York: Columbia University Press.

Becker, Gary S. 1992. "Fertility and the Economy." *Journal of Population Economics* 5 (3): 185–201.

Becker, Gary S., and Robert J. Barro. 1988. "A Reformulation of the Economic Theory of Fertility." *Quarterly Journal of Economics* 103 (1): 1–25.

Becker, Gary S., Kevin M. Murphy, and Robert Tamura. 1990. "Human Capital, Fertility, and Economic Growth." *Journal of Political Economy* 98 (5): S12–37.

Bhattacharya, Joydeep, and Shankha Chakraborty. 2012. "Fertility Choice under Child Mortality and Social Norms." *Economics Letters* 115 (3): 338–41.

Boldrin, Michele, and Larry E. Jones. 2002. "Mortality, Fertility, and Saving in a Malthusian Economy." *Review of Economic Dynamics* 5 (4): 775–814.

Cavalcanti, Tiago, Georgi Kocharkov, and Cezar Santos. 2017. "Family Planning and Development: Aggregate Effects of Contraceptive Use." Rede de Economia Aplicada Working Paper 105.

Cervellati, Matteo, and Uwe Sunde. 2015. "The Economic and Demographic Transition, Mortality, and Comparative Development." *American Economic Journal: Macroeconomics* 7 (3): 189–225.

Choukhmane, Taha, Nicolas Coeurdacier, and Keyu Jin. 2017. "The One-Child Policy and Household Savings." https://personal.lse.ac.uk/jink/pdf/onechildpolicy_ccj.pdf.

Clay, Daniel C., and James J. Zuiches. 1980. "Reference Groups and Family Size Norms." *Population and Environment* 3 (3–4): 262–79.

Coeurdacier, Nicolas, Stéphane Guibaud, and Keyu Jin. 2014. "Fertility Policies and Social Security Reforms in China." *IMF Economic Review* 62 (3): 371–408.

Davis, K. 1967. "Population Policy: Will Current Programs Succeed?" *Science* 158 (3802): 730–39.

Delventhal, Matthew J., Jesús Fernández-Villaverde, and Nezih Guner. 2019. "Demographic Transitions across Time and Space." https://www.sas.upenn.edu/~jesusfv/Demographic_Transitions.pdf.

Department of Census and Statistics of Sri Lanka. 2015. *Household Income and Expenditure Survey 2012/13: Final Report*. Colombo: Department of Census and Statistics of Sri Lanka. http://www.statistics.gov.lk/HIES/HIES2012_13FinalReport.pdf.

de Silva, Tiloka, and Silvana Tenreyro. 2017. "Population Control Policies and Fertility Convergence." *Journal of Economic Perspectives* 31 (4): 205–28.

Doepke, Matthias. 2005. "Child Mortality and Fertility Decline: Does the Barro-Becker Model Fit the Facts?" *Journal of Population Economics* 18 (2): 337–66.

Doepke, Matthias, and Fabrizio Zilibotti. 2005. "The Macroeconomics of Child Labor Regulation." *American Economic Review* 95 (5): 1492–524.

Doepke, Matthias, and Fabian Kindermann. 2016. "Bargaining over Babies: Theory, Evidence, and Policy Implications." NBER Working Paper 22072.

Enke, Stephen. 1960. "The Economics of Government Payments to Limit Population." *Economic Development and Cultural Change* 8 (4): 339–48.

Ehrlich, Isaac, and Jinyoung Kim. 2005. "Endogenous Fertility, Mortality and Economic Growth: Can a Malthusian Framework Account for the Conflicting Historical Trends in Population?" *Journal of Asian Economics* 16 (5): 789–806.

Fernández, Raquel, and Alessandra Fogli. 2009. "Culture: An Empirical Investigation of Beliefs, Work, and Fertility." *American Economic Journal: Macroeconomics* 1 (1): 146–77.

Fernández-Villaverde, Jesús. 2001. "Was Malthus Right? Economic Growth and Population Dynamics." Penn Institute for Economic Research Working Paper 01-046.

Foko, Borel, Beifith Kouak Tiyab, and Guillaume Husson. 2012. "Household Education Spending: An Analytical and Comparative Perspective for 15 African Countries." https://unesdoc.unesco.org/ark:/48223/pf0000216719.

Galor, Oded, and Omer Moav. 2002. "Natural Selection and the Origin of Economic Growth." *Quarterly Journal of Economics* 117 (4): 1133–91.

Galor, Oded, and David N. Weil. 2000. "Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and Beyond." *American Economic Review* 90 (4): 806–28.

Gapminder. 2019. "Data." www.gapminder.org/data/ (accessed on January 30, 2019).

Jones, Larry E., Alice Schoonbroodt, and Michèle Tertilt. 2010. "Fertility Theories: Can They Explain the Negative Fertility-Income Relationship?" In *Demography and the Economy*, edited by John B. Shoven, 43–100. Chicago: University of Chicago Press.

Jones, Stephen R.G. 1984. *The Economics of Conformism*. Oxford: Basil Blackwell.

Kalemli-Ozcan, Sebnem. 2003. "A Stochastic Model of Mortality, Fertility, and Human Capital Investment." *Journal of Development Economics* 70 (1): 103–18.

Kolk, Martin. 2014. "Multigenerational Transmission of Family Size in Contemporary Sweden." *Population Studies* 68 (1): 111–29.

La Ferrara, Eliana, Alberto Chong, and Suzanne Duryea. 2012. "Soap Operas and Fertility: Evidence from Brazil." *American Economic Journal: Applied Economics* 4 (4): 1–31.

Lee, Jong-Wha, and Hanol Lee. 2016. "Human Capital in the Long Run." *Journal of Development Economics* 122: 147–69.

Manski, Charles F., and Joram Mayshar. 2003. "Private Incentives and Social Interactions: Fertility Puzzles in Israel." *Journal of the European Economic Association* 1 (1): 181–211.

Manuelli, Rodolfo E., and Ananth Seshadri. 2009. "Explaining International Fertility Differences." *Quarterly Journal of Economics* 124 (2): 771–807.

Moav, Omer. 2005. "Cheap Children and the Persistence of Poverty." *Economic Journal* 115 (500): 88–110.

Munshi, Kaivan, and Jacques Myaux. 2006. "Social Norms and the Fertility Transition." *Journal of Development Economics* 80 (1): 1–38.

Murphy, M. 1999. "Is the Relationship between Fertility of Parents and Children Really Weak?" *Social Biology* 46 (1–2): 122–45.

Nortman, Dorothy. 1982. *Population and Family Planning Programs: A Compendium of Data through 1981*. New York: Population Council.

Nortman, Dorothy L., and Ellen Hofstatter. 1978. *Population and Family Planning Programs*. 9th ed. New York: Population Council.

OECD. 2016a. *Education at a Glance 2016: OECD Indicators*. Paris: OECD.

OECD. 2016b. "Household Spending." https://data.oecd.org/hha/household-spending.htm (accessed on January 30, 2017).

Palivos, Theodore. 2001. "Social Norms, Fertility and Economic Development." *Journal of Economic Dynamics and Control* 25 (12): 1919–34.

Psacharopoulos, George, and Harry Anthony Patrinos. 2004. "Returns to Investment in Education: A Further Update." *Education Economics* 12 (2): 111–34.

Regional Bureau of Education for Latin America and the Caribbean. 2013. *The State of Education in Latin America and the Caribbean: Towards a Quality Education for All*. Santiago: UNESCO.

Rizk, Reham, and Hala Abou-Ali. 2016. "Out of Pocket Health Expenditure and Household Budget: Evidence from Arab Countries." Economic Research Forum Working Paper 1065.

Ross, John A., W. Parker Maudlin, and Vincent C. Miller. 1993. *Family Planning and Population: A Compendium of International Statistics*. New York: Population Council.

Singapore Department of Statistics. 2014. *Report on the Household Expenditure Survey 2012/13*. Singapore: Ministry of Trade and Industry. https://www.singstat.gov.sg/-/media/files/publications/households/hes1213.pdf.

Spolaore, Enrico, and Romain Wacziarg. 2014. "Fertility and Modernity." https://www.brown.edu/academics/economics/sites/brown.edu.academics.economics/files/uploads/Enrico%20Spolaore_FertilityModernity_Paper.pdf.

STAT Compiler. 2015. "The Demographic and Health Surveys Program." www.statcompiler.com (accessed on June 18, 2015).

Therborn, Göran. 2004. *Between Sex and Power: Family in the World 1900–2000*. London: Routledge.

Thornton, Arland. 1980. "The Influence of First Generation Fertility and Economic Status on Second Generation Fertility." *Population and Environment* 3 (1): 51–72.

Tilak, Jandhyala B.G. 2009. "Household Expenditure on Education and Implications for Redefining the Poverty Line in India." http://planningcommission.gov.in/reports/genrep/tilak.pdf.

United Nations Development Programme. 2013. Expected Years of Schooling. http://hdr.undp.org/en/content/expected-years-schooling-children-years (accessed on January 15, 2018).

World Bank. 2018. "World Bank Group: Development Indicator." http://data.worldbank.org/data-catalog/world-development-indicators (accessed on March 3, 2018).

Young, Alwyn. 1995. "The Tyranny of Numbers: Confronting the Statistical Realities of the East Asian Growth Experience." *Quarterly Journal of Economics* 110 (3): 641–80.