Modeling and assessing low-fertility traps: Intergenerational feedback effects under multigenerational optimization

Michael Geruso* Nicholas Lawson† Dean Spears‡

paper for Wittgenstein Centre Conference: November 2021

Conference question:
"How to include potential demographic, social, economic, and environmental feedback effects into models of depopulation?"

Abstract

The low-fertility trap hypothesis is one of the most important theories about depopulation. A low-fertility trap is an intergenerational stable equilibrium, where low fertility in one generation causes low fertility in the next through strong intergenerational forces. An equilibrium is a trap if such intergenerational forces overpower intragenerational determinants. We construct a model of optimizing parents who live in a succession of cohorts. As a reduced-form proxy for intergenerational forces, we model a role for dispersion within a cohort's fertility outcomes. Because fertility is bounded below at zero, as average fertility decreases, variance decreases, too. Various mechanisms may discourage outlier fertility. Yet, we find that it is difficult for such a model to generate a low-fertility trap: intergenerational forces must be quite strong, relative to other factors that influence fertility decisions. This result is consistent with the empirical fact that fertility in low-fertility populations has not settled into stationary equilibria. [150 words]

*Population Research Center and Economics Department, University of Texas at Austin; NBER.

†Université du Québec à Montréal.

‡**Corresponding author**: dspears@utexas.edu. Population Research Center and Economics Department, University of Texas at Austin; r.i.c.e.; IZA.

1 Introduction

The hypothesis of a low-fertility trap is one of the most important theories of depopulation (Lutz et al., 2006). As Lutz and Skirbekk (2005) explained, introducing the hypothesis: "Social inertia and self-reinforcing processes may make it difficult to return to higher levels once fertility has been very low for some time, creating a possible 'low-fertility trap.'" Because of the importance of the possibility of low-fertility traps to understanding depopulation and the prospects for policy responses, it is essential that this idea be formalized into a model that can be used to test theories and inform policy.

Here, we apply the tools of economics to the task of modeling intergenerational feedback effects in a low-fertility trap. Although we present a formal model below in the spirit of Barro and Becker (1989), we begin with an informal analogy to the large economic literature about "poverty traps." Just as a low fertility trap would be a self-reinforcing mechanism whereby populations that reach low fertility are caused to remain at low fertility, a poverty trap is a hypothesized self-reinforcing mechanism whereby countries that start poor remain poor (Kraay and McKenzie, 2014; Barrett et al., 2016). Being in a poverty trap is distinct from being poor, or even from being enduringly poor: poverty can be stationary without being trapped. A poverty trap is a *stable equilibrium in an intertemporal process*, where past poverty causes future poverty to persist, even if some conditions change. Poverty traps are important because they have policy implications: modest capital infusions or other policy interventions that might work outside of a trap will have no effect if a country is stuck in an equilibrium poverty trap, but if policy makes a "big push," the economy might switch to a prosperous equilibrium.

Figure 1 applies these ideas to a low-fertility trap. Panel (a) repurposes the classic diagram of a poverty trap to illustrate a fertility trap. The illustration makes clear that a low-fertility trap must be an *intergenerational* processes. The curved line plots that process, mapping from the fertility of one generation to the statistically expected fertility of the next. If the fertility of

Figure 1: How intergenerational forces might or might not create a trap

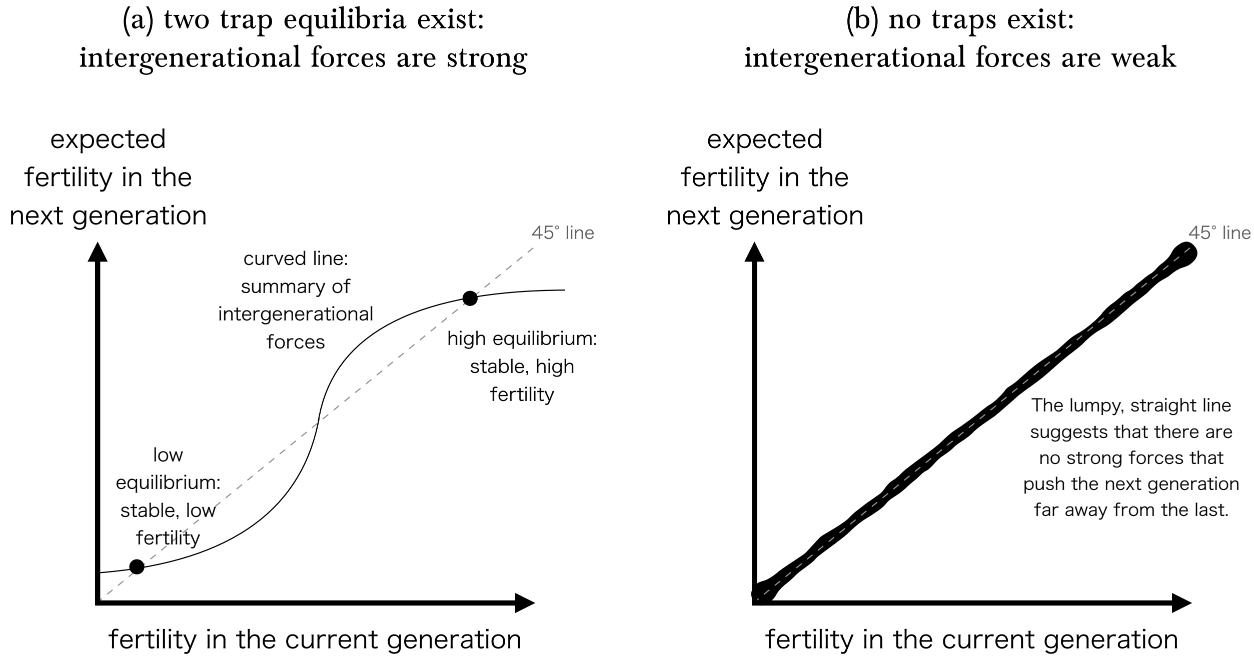

one generation is low, this intergenerational effect pushes the fertility of the next generation towards a stable low equilibrium. What makes the low equilibrium a low-fertility trap is that it is stable. If policy, social change, or another factor were to push fertility above baseline fertility for a generation, then intergenerational forces would pull fertility back towards the equilibrium trap.

In contrast, Panel (b) of Figure 1 illustrates the possibility that there is no fertility trap. In Panel (b), this is because there are no strong intergenerational forces that move the expectation for the next generation's fertility away from the level of the current generation's fertility. Instead, the next generation will probably be like the last, on average, unless some other factor changes. In Panel (b), any point can be an unstable equilibrium. There are no stable equilibria: if cultural or economic factors change the number of children parents want to have, then fertility levels can enduringly change. The critical difference between Panels (a) and (b) is the strength of *intergenerational* forces which, in Panel (a), can overwhelm changes in *intragenerational* determinants of fertility.

Whether actual fertility dynamics resemble Panel (a), Panel (b), or something else is an empirical question. Is there evidence of a low fertility trap? Figure 2 considers the simplest intergenerational dynamics: how does past fertility in low-fertility populations predict future fertility? Once fertility becomes low, does it remain in a stable equilibrium? The axes of Panels (a) and (c) of Figure 2 resemble the axes in Figure 1: the horizontal axis is the fertility of an earlier generation and the vertical axis is the fertility of a later generation. Panels (b) and (d) subtract the difference between the generations, to clarify how past conditions predict future outcomes.

We see little evidence in Figure 2 that, so far, low-fertility countries have become trapped in a stable equilibrium. We do not see, for example, that countries that started out with lower fertility were more likely subsequently to be stationary or declining. If anything, countries that started with lower fertility tended to show fertility increases, such as Russia. Countries that share a starting point show divergent trajectories. In 1995, Slovenia, Portugal, Austria, and Estonia, for example, all had tempo-adjusted TFR within the range 1.71 to 1.77. Portugal and Austria fell towards 1.6, Slovenia stayed near 1.7, and Estonia increased to 2.3. To be sure, the literature on convergence dynamics in international development economics teaches us that convergence to equilibria is difficult to assess with statistics (De Long, 1988). Each country, for one possibility, may be moving on its own path towards its own unique low-fertility trap. But no low-fertility trap is apparent in Figure 2.

In our formal contribution, we consider why an intergenerational low-fertility trap may be hard to find. Recalling the shape of the curve in Figure 1, if we are to find a low fertility trap, we must find an mechanism by which, if fertility in one generation is low but above the equilibrium level in the trap, then parents in the next generation would choose to have even fewer children. We construct an intertemporal model of optimizing agents — women who may become mothers — who live in a succession of cohorts. Each cohort chooses its fertility in light of its goals, the costs, and the benefits — and under the influence of the fertility behavior of the cohorts who

Figure 2: Trap or motion?: Empirical dynamics of total fertility rates (TFR)

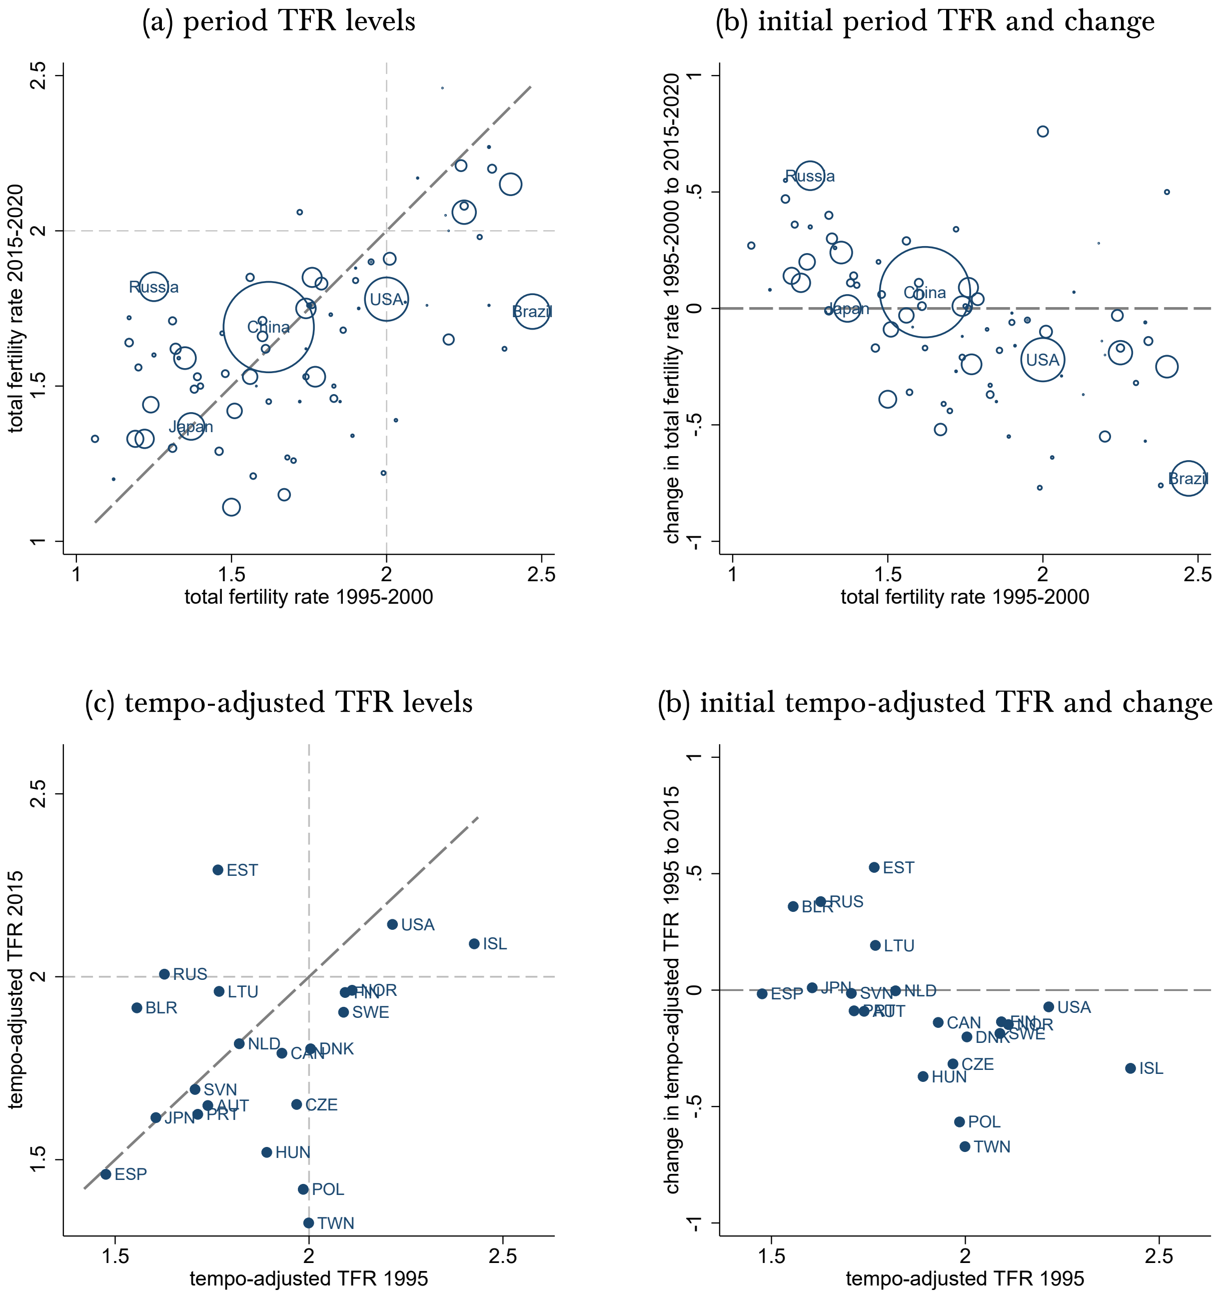

Note: Observations are countries. Panels (a) and (b) use the same observations and Panels (c) and (d) use the same observations. Data in Panels (a) and (b) are from the 2019 UN World Population Prospects and circles are weighted by 2015-2020 population size. Tempo-adjusted TFR attempts to adjust for changes in the age-distribution of fertility, to permit comparisons of the quantum of fertility. Data in Panels (c) and (d) are from the Human Fertility Database. (Max Planck Institute for Demographic Research (Germany) and Vienna Institute of Demography (Austria). Available at www.humanfertility.org (data downloaded on 8 November 2021).)

have come before. Such intergenerational influence could come in various forms, for example: economic, as employment, educational, and child care offerings are endogenously optimized to suit prevailing fertility trends; social, as the distribution of one generation's fertility outcomes influences the ideal or expected fertility of the next; and policy-mediated, as governments design programs and policies for average members of the population. Our stylized model abstracts away from the details of how intergenerational forces do their work.

We particularly closely explore a possible role of *dispersion* within a population's fertility outcomes. We intend this as a reduced-form representation that could encompass a wide range of actual intergenerational social, economic, and political forces. Because fertility is bounded below at zero, as average fertility decreases, the variance of fertility is likely to decrease, as well.[1] For example, consider a case where fertility in a cohort has a poisson distribution, which is a plausible discrete distribution with a variance equal to its mean. In a cohort with a Total Fertility Rate of 1 and a poisson distribution, only 8% of women would have three or more children, and only 2% would have four or more. The consequence is that moderately above-average levels of fertility would be high percentiles in the tail of the distribution, in tightly-distributed, low-fertility populations. Various mechanisms may discourage outlier levels of fertility. Importantly, this low-variance mechanism would generate an asymmetric *low* fertility trap, because low levels of fertility in a low-fertility cohort would not be outliers but high levels of fertility would be: fertility is bounded below at zero but not (for practical purposes in present-day populations) biologically bounded above.

Ultimately, we find that it is difficult for such a model to generate a low-fertility trap: intergenerational forces must be quite strong, relative to other factors that influence fertility decisions. This result is consistent with the facts of Figure 2: fertility levels in low-fertility populations do not appear to have settled into stationary equilibria.

[1]A version of this statistical fact has been previously documented by Hruschka and Burger (2016), who do not model its consequences.

Figure 3: Across place and time, cohorts with lower average fertility have lower fertility variance

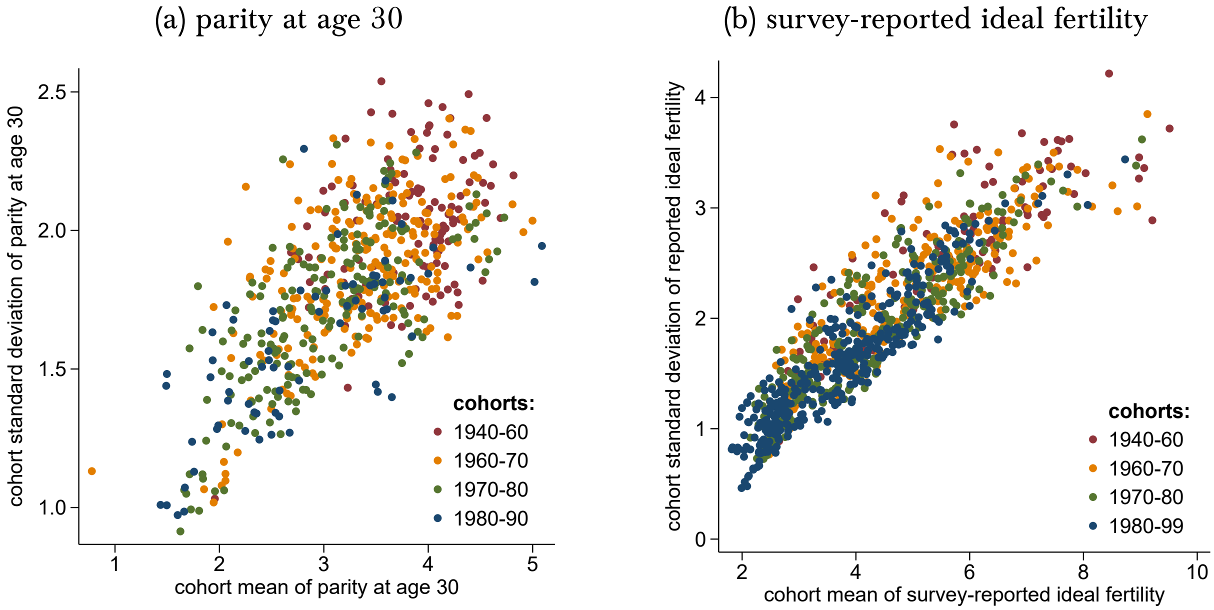

Note: Points in panel (a) are 615 five-year birth cohorts of women from 131 Demographic and Health Survey rounds; the mean point reflects 1,848 women. Points in panel (b) are 972 five-year birth cohorts of women from 130 Demographic and Health Survey rounds; the mean point reflects 2,324 women.

2 Background:

Low fertility averages create low fertility variances

This section presents preliminary empirical information about the reduced-form mechanism that our model uses to stand in for a broad set of intergenerational forces: the fact that low fertility levels create low fertility variance. Figure 3 plots actual and survey-reported ideal fertility from cohorts in the Demographic and Health Surveys (DHS) of developing countries. Each dot is a five-year birth cohort in one country (more specifically, in one DHS survey wave in one country). The horizontal axes plot average fertility in the cohort and the vertical axes plot the standard deviation of fertility within the cohort. Our purpose here is to document the clear positive slope: the variance of fertility is low in cohorts where the average is low.

Figure 4: Case study: Declining variance in actual and ideal cohort fertility in India

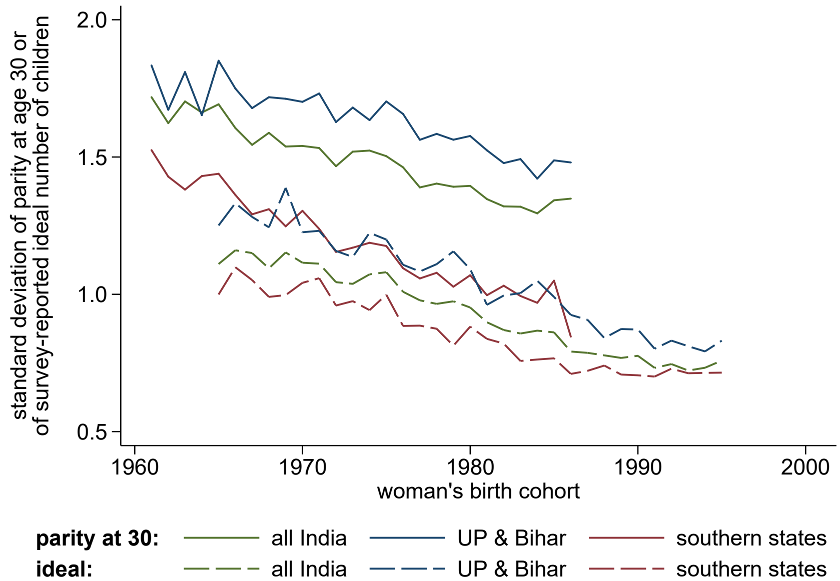

Note: Data are from the women's recodes of the 2005-6 and the 2015-16 rounds of the Demographic and Health Survey in India.

Figure 4 presents a striking case study: India, which had high fertility rates in the mid-20th century that were central to global fertility policy debates (Connelly, 2008), but which has subsequently had a demographic transition to low fertility (Visaria and Visaria, 2003). The vertical axis plots the standard deviation[2] of cohort fertility for birth cohorts, computed at one-year intervals, for India as a whole and for two large regions. The message of the figure is the downward trend: in India, as in other populations, fertility is not only becoming lower, it is also becoming *more homogenous* across women in a cohort.

In isolation, the fact that the variance of fertility is higher when the mean is higher is hardly startling, given the nature of this variable. However, the regularity of this fact is striking and worth documenting. The more important reason to consider this fact here is that our model

[2]In this way, our analysis is distinguished from complementary work in the prior literature that present visually-similar plots of how the *mean*, rather than the variance, of cohort fertility has changed over time (*e.g.*, Frejka and Calot, 2001; Lesthaeghe, 2010; Zeman et al., 2018).

will attempt to use this empirical regularity to formalize the *consequences* of this pattern for intergenerational feedback in our optimizing model.

In the rest of our paper, we demonstrate theoretically and numerically to what extent it is difficult to generate a low-fertility trap, and why that is the case. There are two main avenues for generating a low-fertility trap: convergence of the distribution of fertility to a degenerate distribution at a single point, or multiple equilibria.

In Section 3, we show that it is possible for the distribution of fertility to converge to a single point if the variance drops below a certain threshold, which would only tend to happen if the average level of fertility is low. We present this result in a setting in which fertility can take any continuous value greater than or equal to zero, for simplicity. However, we show that in any setting, this argument relies on a force of social learning that becomes infinitely strong, such that all variation in fertility disappears, which seems likely to be counterfactual in the real world.

In Section 4, we consider the possibility that the fertility distribution may exhibit multiple equilibria. We model this in a setting with discrete fertility, where the number of children is restricted to be a whole number, to make the analysis tractable. When average fertility is low, the equilibrium distribution may be one in which a force of social learning causes significant concentration at a given number of children, and it is possible that if fertility preferences subsequently increase, the equilibrium will remain one with high concentration at the same low level of fertility, which is not true of an alternate equilibrium that could have existed if fertility preferences had always been high. However, we shall see that it is not easy to generate such a trap, and that in any case, it is not clear that such a "low-fertility trap" will actually involve a particularly low level of average fertility.

3 Continuous Fertility Model with Degenerate Distribution

We begin with the possibility of fertility converging to a single value, using a continuous model of fertility for simplicity. Let us assume that there is a succession of non-overlapping generations of parents who each live in one period denoted by t. Each parent has a type defined by their value of b, which is a personal utility from fertility, and which is distributed according to a Gamma distribution with shape parameter k and scale parameter θ;[3] as discussed later, these parameters of the distribution may change over time. We assume that the parent can choose any non-negative number n of children, subject to a convex utility cost from fertility (representing time and money costs of raising children). There is also an intergenerational social learning term that is a function of the difference between a parent's fertility and average fertility in the previous period, divided by the variance of fertility in the previous period; basically, we assume that parents learn what is "normal" from the example of their parents' generation.

For individual i in period t with type b_{it} and fertility choice n_{it}, utility is given by:

$$U_{it}(n_{it}) = b_{it}n_{it} - \frac{1}{2}n_{it}^2 - \alpha\frac{(n_{it} - \mathrm{E}_{t-1}(n))^2}{\mathrm{Var}_{t-1}(n)} \tag{3.1}$$

where α represents the strength of the intergenerational force of social learning, and $\mathrm{E}_{t-1}(n)$ and $\mathrm{Var}_{t-1}(n)$ are the mean and variance of fertility in the previous period. The optimal value of n_{it} is given by:

$$\begin{aligned} U'_{it}(n_{it}) &= b_{it} - n_{it} - \frac{2\alpha}{\mathrm{Var}_{t-1}(n)}(n_{it} - \mathrm{E}_{t-1}(n)) = 0 \\ \rightarrow n_{it} &= \frac{\mathrm{Var}_{t-1}(n)b_{it} + 2\alpha\mathrm{E}_{t-1}(n)}{\mathrm{Var}_{t-1}(n) + 2\alpha} \end{aligned}$$

which we can see is just a weighted average of the privately-optimal fertility (in the absence of

[3] Such a distribution has a mean equal to $k\theta$ and a variance of $k\theta^2$. The chi-squared distribution with v degrees of freedom is a special case of the Gamma distribution with $k = v/2$ and $\theta = 2$.

social learning) b_{it} and the average fertility in the previous period $\mathrm{E}_{t-1}(n)$.

This gives us the following difference equations for the mean and variance of fertility:

$$\mathrm{E}_t(n) = \frac{\mathrm{Var}_{t-1}(n)\mathrm{E}_t(b) + 2\alpha\mathrm{E}_{t-1}(n)}{\mathrm{Var}_{t-1}(n) + 2\alpha} \tag{3.2}$$

$$\mathrm{Var}_t(n) = \mathrm{Var}_t(b)\left(\frac{\mathrm{Var}_{t-1}(n)}{\mathrm{Var}_{t-1}(n) + 2\alpha}\right)^2. \tag{3.3}$$

In a steady-state in which k and θ are constant, the mean must satisfy the following equation:

$$\mathrm{E}(n) = \frac{\mathrm{Var}(n)\mathrm{E}(b) + 2\alpha\mathrm{E}(n)}{\mathrm{Var}(n) + 2\alpha} \tag{3.4}$$

and if the variance is greater than zero, we can solve this to find that $\mathrm{E}(n) = \mathrm{E}(b) \equiv k\theta$. If, however, $\mathrm{Var}(n) = 0$ in steady-state, then the equation degenerates to $\mathrm{E}(n) = \mathrm{E}(n)$ and any value is hypothetically possible for the mean.

Meanwhile, the steady-state variance must satisfy the following equation:

$$\mathrm{Var}(n) = \mathrm{Var}(b)\left(\frac{\mathrm{Var}(n)}{\mathrm{Var}(n) + 2\alpha}\right)^2. \tag{3.5}$$

and the first thing to notice here is that a variance of zero is an equilibrium satisfying this equation. However, if the variance is positive, we can solve for two additional roots, using $\beta \equiv \mathrm{Var}(b) \equiv k\theta^2$ and $V \equiv \mathrm{Var}(n)$ for notational simplicity:

$$V^2 + (4\alpha - \beta)V + 4\alpha^2 = 0$$

$$\to V = \frac{\beta - 4\alpha \pm \sqrt{(4\alpha-\beta)^2 - 16\alpha^2}}{2} = \frac{\beta}{2} - 2\alpha \pm \sqrt{\frac{1}{4}\beta^2 - 2\alpha\beta}.$$

We can see that if $\beta < 8\alpha$, this equation does not give any real solution. Therefore, if the parameters are such that $\mathrm{Var}(b) < 8\alpha$, the only possible equilibrium value of steady-state variance of n_{it} is zero. When $\beta \geq 8\alpha$, we can evaluate which of the roots are stable by returning to equation (3.5), and taking the derivative of each side with respect to $V \equiv \text{Var}(n)$: if the derivative of the right-hand side is smaller than that of the left-hand side, then the root is stable (because an increase in V raises the right-hand side by less than the increase in V, meaning that any deviation from the root must collapse back to the root). The derivative of the left-hand side is always 1, whereas the derivative of the right-hand side is:

$$\frac{\partial RHS}{\partial V} = 2\beta\left(\frac{V}{V+2\alpha}\right)\left(\frac{2\alpha}{(V+2\alpha)^2}\right) = \frac{4\beta\alpha V}{(V+2\alpha)^3}. \tag{3.6}$$

We can clearly see that $V = 0$ is a stable equilibrium; and for $V > 0$, we can use the fact that $1 = \frac{\beta V}{(V+2\alpha)^2}$, as implied by equation (3.5), to simplify the above equation to $\frac{\partial RHS}{\partial V} = \frac{4\alpha}{V+2\alpha}$. At the largest root, where $V = \frac{\beta}{2} - 2\alpha + \sqrt{\frac{1}{4}\beta^2 - 2\alpha\beta}$, we find that:

$$\frac{\partial RHS}{\partial V} = \frac{4\alpha}{V+2\alpha} = \frac{4\alpha}{\frac{\beta}{2} + \sqrt{\frac{1}{4}\beta^2 - 2\alpha\beta}} < \frac{8\alpha}{\beta} < 1. \tag{3.7}$$

Therefore there are two stable equilibrium values of $\text{Var}(n)$ given by zero and $\frac{\beta}{2} - 2\alpha + \sqrt{\frac{1}{4}\beta^2 - 2\alpha\beta}$, and by continuity we know that the middle value given by $\frac{\beta}{2} - 2\alpha - \sqrt{\frac{1}{4}\beta^2 - 2\alpha\beta}$ is an unstable equilibrium. As a result, we can interpret the unstable equilibrium as a critical value for variance: if the variance of fertility drops below $\hat{V} = \frac{\beta}{2} - 2\alpha - \sqrt{\frac{1}{4}\beta^2 - 2\alpha\beta}$, it will converge towards zero, and will not rise again unless the parameters of the model change.

It is also apparent that the variance of fertility is a function of the variance of b, which is equal to $k\theta^2$. Suppose that k exogenously declines in the population, leading to lower fertility preferences (the average of b is given by $k\theta$); the variance will also decline, and if that variance drops below $\hat{V}$, it will converge to zero. Even if k subsequently increases, leading to higher fertility preferences, if the variance has dropped enough that $\text{Var}_{t-1}(n) < \hat{V}$ (recognizing that $\hat{V}$ is a function of the variance of b), it will be insufficient to escape the low-fertility trap. Since the variance of b tends to vary in the same direction as the mean of b, it is clear that this trap is indeed likely to be triggered by sustained low levels of fertility.

We have simulated this model to provide an illustration of the low-fertility trap. Suppose that we are initially in an equilibrium in which $k = 1.6$ and $\theta = 1.25$, so that $\mathrm{E}(b) = 2$ and $\mathrm{Var}(b) = 2.5$, and that in period $t = 1$ parents are hit by a shock that lowers k to 0.8. Fertility preferences remain low for 10 periods, after which k jumps back up to 1.6. We simulate the mean and variance of fertility for 20 periods in total, for 4 different scenarios given by $\alpha = \{0.12, 0.15, 0.18, 0.21\}$, and the results can be found in Figure 5.

With $k = 0.8$ and $\theta = 1.25$, the critical value of α above which only the equilibrium at zero variance exists is 0.15625; at values of intergenerational social force α below that critical value, a stable positive value of fertility variance exists. When k returns to 1.6, the critical value of variance required to escape the trap depends on α, varying from approximately 0.0289 when $\alpha = 0.12$ to about 0.1141 when $\alpha = 0.21$. Accordingly, at low values of α, the variance remains high when k drops to 0.8, and average fertility increases rapidly when k jumps back up to 1.6. When $\alpha = 0.18$, the only equilibrium value of fertility variance is zero once k drops to 0.8, but the convergence to zero is slow enough that variance remains above the critical value when k jumps up to 1.6, and the population escapes the low-fertility trap. At $\alpha = 0.21$, the variance of fertility drops to zero more quickly, and even a doubling of fertility preferences cannot escape the low-fertility trap.

This model demonstrates that a low level of fertility, which tends to generate low variance of fertility as well, could generate a low-fertility trap if fertility remains sufficiently low for a sustained period of time; even a subsequence increase in fertility preferences could be insufficient to escape the trap. However, this version of the trap depends on the variance of fertility converging to zero, which is a radical scenario that does not seem likely to occur in the real world.

The logic of this result could be applied to generalizations of this model, including a version

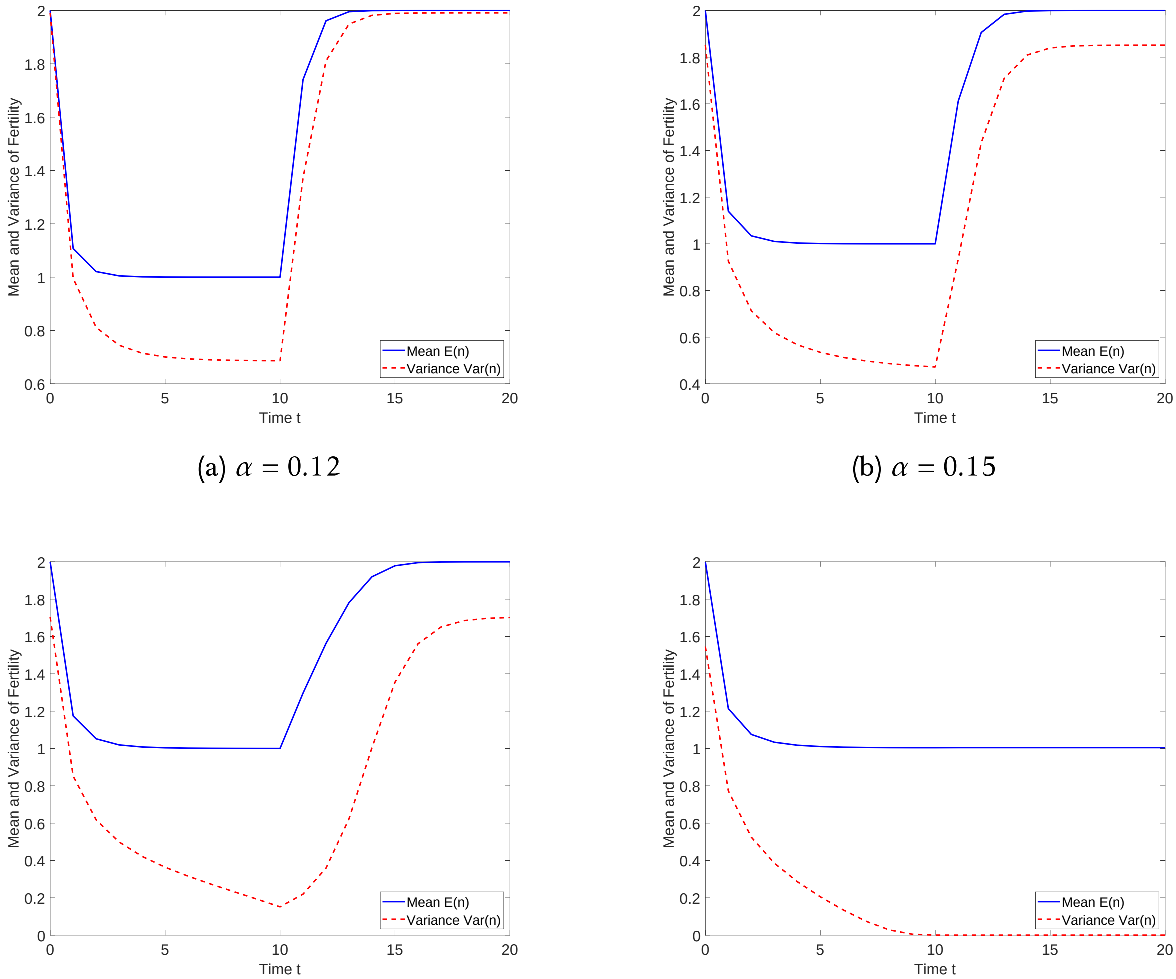

(a) $\alpha = 0.12$

(b) $\alpha = 0.15$

(c) $\alpha = 0.18$

(d) $\alpha = 0.21$

Figure 5: Illustration of the Degenerate Low-Fertility Trap

in which fertility is restricted to whole numbers, as the underlying logic is that the distribution of fertility will always gradually converge back to its equilibrium unless the force of intergenerational social learning holding it in place is infinitely strong. As a result, in the absence of multiple equilibria or the possibility of the variance of fertility converging to zero, a low-fertility trap could not exist. To understand this, suppose that fertility preferences were low and fertility was distributed around values close to zero, but that subsequently fertility preference k increased in period t (due to changing attitudes or to government policies to reduce the net cost of children to parents). Due to intergenerational social learning, each parent in period t would choose a value of n_{it} that is in between their personal preference b_{it} and average fertility in the prior generation $\mathrm{E}_{t-1}(n)$, and that value of n_{it} will be higher than the value they would have chosen had k not increased; as a result, average fertility will increase slightly from period $t-1$ to t, and in the subsequent period that increase in average fertility will further raise each parent's chosen fertility. Over time, in the absence of an infinitely strong intergenerational social force, average fertility will gradually converge back to $\mathrm{E}(n) = k\theta$, and the distribution of fertility will converge to one that is appropriate to the new value of k.[4]

4 Discrete Fertility Model with Multiple Equilibria

We now consider the possibility of multiple fertility equilibria, in the context of a model with discrete choice over fertility. The possibility of multiple equilibria is much easier to model in a setting in which fertility is restricted to whole numbers, because it is necessary to specify an intergenerational social learning term that depends on the percentage of the population choosing each possible level of fertility.

To keep the model as similar as possible to that from the previous section, we use the following utility function:

$$U_{it}(n_{it}) = b_{it}n_{it} - \frac{1}{2}n_{it}^2 + \alpha \frac{p_{t-1}(n_{it})^\gamma}{p_{t-1}(n_{it})^\gamma + (1 - p_{t-1}(n_{it}))^\gamma} \tag{4.1}$$

where $\gamma > 0$, and where $p_{t-1}(n_{it})$ is the fraction of parents in period $t-1$ that chose the fertility level given by n_{it}. The social learning term now takes the form of an S-shaped curve in the fraction of parents choosing a given level of fertility: when a large fraction of parents chose n_{it}

[4] In a model with discrete fertility, in which n is limited to whole numbers, some fraction of the population will react to an increase in k by raising their fertility by 1 unit, which will raise average fertility; the adjustment over time would be in the form of a fraction of the population raising fertility by 1 unit each period rather than all of the population raising fertility by some fraction of a unit.

in the previous period, it increases a parent's desire to choose that level of fertility.

It is apparent that this model specification could generate multiple equilibria, in the sense that both low percentages and high percentages of parents at a given level of fertility could be self-sustaining. To see this algebraically in a simplified version of the model, let us consider a binary and static version of the model in which parents can choose $n = \{0, 1\}$; this does not literally have to be understood as 0 or 1 child, but rather as "low" and "high" levels of fertility. Suppose that the relative utility of choosing $n = 1$ is given by $U_1(p; b) = b - \frac{1}{2}(1+\alpha) + \alpha \frac{p^\gamma}{p^\gamma + (1-p)^\gamma}$, where p is the fraction of parents choosing $n = 1$, and suppose that b is distributed uniformly between 0 and 1. In this case, the equilibrium p is equal to 1 minus the critical value of b at which parents are indifferent, and therefore we must have in equilibrium:

$$p = \frac{1}{2}(1-\alpha) + \alpha \frac{p^\gamma}{p^\gamma + (1-p)^\gamma}. \tag{4.2}$$

Clearly, $p = 0.5$ is always an equilibrium, but if α and γ are large enough, this equilibrium may be unstable and there may be additional equilibria at higher and lower values of p. This will be the case if the derivative of the right-hand side of (4.2) with respect to p is larger than 1 at $p = 0.5$, which will be true if $\alpha\gamma > 1$. In that case, the only stable equilibria involve values that are different from 0.5, and potentially quite far from 0.5. However, this example of multiple equilibria is too simple for our purposes, as the binary setting does not permit any asymmetry regarding low- versus high-fertility traps.

Returning to the broader model represented in (4.1), the logic of a low-fertility trap would be as follows: the social learning term provides an incentive for parents to concentrate at numbers of children that were popular in previous generations, and particularly as the fraction of the previous generation that chose a particular n rises in the vicinity of 0.5.[5] When average fertility is high, the distribution of fertility levels is likely to be spread out sufficiently that this social

[5]It is possible to generalize the model to place the inflection point of the social learning term somewhere other than 0.5, without significantly changing the model's implications.

learning term only has a modest effect; but if fertility preferences and thus average fertility decrease, the distribution is likely to become more concentrated around, for example, $n = 1$ or $n = 2$, making it possible for the social learning term to dominate and lead to an even higher level of concentration. If fertility preferences subsequently increase again, the distribution may remain stuck in an equilibrium in which the distribution is highly concentrated at a particular level of fertility.

An example of this mechanism is provided by the following simulation, which assumes $\theta = 1$, $\alpha = 1.75$, and $\gamma = 8$, and starts with $k = 2$. We start at an initial equilibrium to which the model converges from the allocation in the absence of social learning, at which average fertility is approximately equal to 2. From this initial equilibrium, we again assume that in period $t = 1$ parents are hit by a shock that lowers k, this time to 1. Fertility preferences remain low for 10 periods, after which k jumps back up to 2. The results can be found in Figure 6, which plots average fertility over time, and in Figure 7, which presents the fertility distribution at each of the three equilibria.

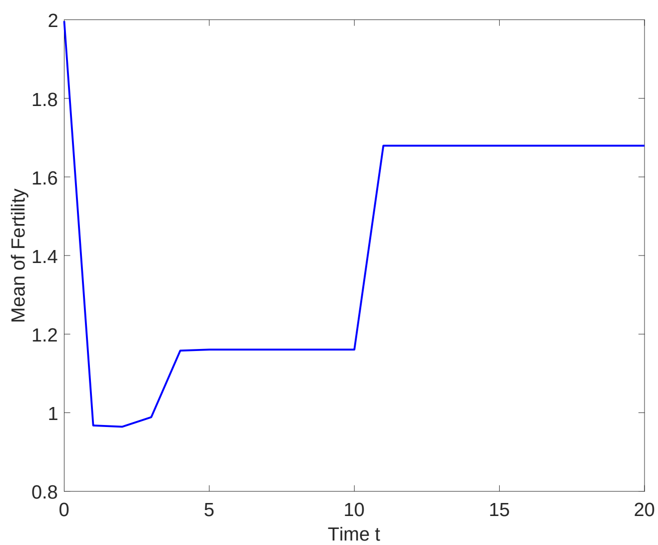

Figure 6: Average Fertility

Figure 7 shows that when fertility preferences drop, the distribution of fertility becomes

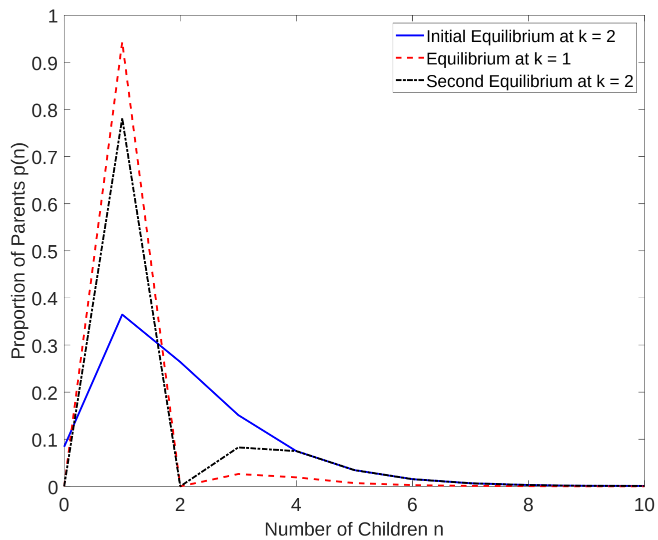

Figure 7: Equilibrium Distributions of Fertility

concentrated around $n = 1$, and when k subsequently increases, the proportion of parents choosing $n = 1$ remains much higher than in the initial equilibrium. However, this corresponds to an average level of fertility that is only moderately lower than in the original equilibrium, as Figure 6 shows that average fertility rises to about 1.7 instead of 2. This is because, while the new equilibrium involves parents who previously chose to have 2 or 3 children now choosing $n = 1$, parents who would have chosen $n = 0$ are also now choosing $n = 1$. Furthermore, the long right tail of the distribution is unaffected, which seems logical: parents who were planning to have 6 children probably have unusually strong preferences over children which should not be significantly affected by knowing the distribution of fertility around 0, 1, 2, or 3 children.

It is also worth noting that it is not easy to generate this sort of low-fertility trap: there is no difference between the first and last equilibria if β is reduced to 6, for example, and the gap between the average fertility levels gradually declines when α declines, vanishing around $\alpha = 1.15$. As a result, what we learn from this model is that it is difficult to generate a low-fertility trap with multiple equilibria, and that the logic explained above means that it is hard for this trap to actually involve very low fertility. An alternative model may exist that would generate a deeper and lower form of a low-fertility trap, but it would require significant modification from the simple and intuitive model presented here.

5 Conclusion

All leading long-term global population projections agree on projecting continued fertility decline, in individual countries and for the human population as a whole (KC and Lutz, 2017; United Nations, 2019; Vollset et al., 2020). We see no reason to disagree. But sustained, low fertility is not sufficient to conclude that there is a low-fertility trap. This is because, rather than low fertility persisting because it is trapped low due to intergenerational forces, it may be that many successive generations choose low fertility for their own reasons. After all, when conditions changed in the 20th century, intergenerational forces did not trap parents at high fertility. In both our simple, macro-level empirical description of fertility and in our stylized micro-level model of intergenerational feedback effects, a low-fertility trap—understood as a stable intergenerational equilibrium—proves difficult to find. If the implication is that future generations have their fertility options open to the choices that they may find fit, then (depending upon what they choose and its consequences) this conclusion could be good news.

References

Barrett, Christopher B, Teevrat Garg, and Linden McBride (2016) "Well-being dynamics and poverty traps," *Annual Review of Resource Economics*, Vol. 8, pp. 303–327.

Barro, Robert J and Gary S Becker (1989) "Fertility choice in a model of economic growth," *Econometrica*, pp. 481–501.

Connelly, Matthew (2008) *Fatal misconception: The struggle to control world population*: Harvard University Press.

De Long, J Bradford (1988) "Productivity growth, convergence, and welfare: comment," *The American Economic Review*, Vol. 78, pp. 1138–1154.

Frejka, Tomas and Gérard Calot (2001) "Cohort reproductive patterns in low-fertility countries," *Population and Development Review*, Vol. 27, pp. 103–132.

Hruschka, Daniel J and Oskar Burger (2016) "How does variance in fertility change over the demographic transition?" *Philosophical Transactions of the Royal Society B: Biological Sciences*, Vol. 371, p. 20150155.

KC, Samir and Wolfgang Lutz (2017) "The human core of the shared socioeconomic pathways: Population scenarios by age, sex and level of education for all countries to 2100," *Global Environmental Change*, Vol. 42, pp. 181–192.

Kraay, Aart and David McKenzie (2014) "Do poverty traps exist? Assessing the evidence," *Journal of Economic Perspectives*, Vol. 28, pp. 127–48.

Lesthaeghe, Ron (2010) "The unfolding story of the second demographic transition," *Population and Development Review*, Vol. 36, pp. 211–251.

Lutz, Wolfgang and Vegard Skirbekk (2005) "Policies addressing the tempo effect in low-fertility countries," *Population and Development Review*, Vol. 31, pp. 699–720.

Lutz, Wolfgang, Vegard Skirbekk, and Maria Rita Testa (2006) "The low-fertility trap hypothesis: Forces that may lead to further postponement and fewer births in Europe," *Vienna Yearbook of Population Research*, pp. 167–192.

United Nations (2019) *World Population Prospects 2019: Highlights*: Department of Economic and Social Affairs, Population Division.

Visaria, Leela and Pravin Visaria (2003) "Long-term population projections for major states, 1991-2101," *Economic and Political Weekly*, pp. 4763–4775.

Vollset, Stein Emil, Emily Goren, Chun-Wei Yuan, Jackie Cao, Amanda E. Smith, Thomas Hsiao, Catherine Bisignano, Gulrez S. Azhar, Emma Castro, Julian Chalek, Andrew J. Dolgert, Tahvi Frank, Kai Fukutaki, Simon I. Hay, Rafael Lozano, Ali H. Mokdad, Vishnu Nandakumar, Maxwell Pierce, Martin Pletcher, Toshana Robalik, Krista M. Steuben, Han Yong Wunrow, Bianca S. Zlavog, and Christopher J. L. Murray (2020) "Fertility, mortality, migration, and population scenarios for 195 countries and territories from 2017 to 2100: a forecasting analysis for the Global Burden of Disease Study," *The Lancet*, Vol. 396, pp. 1285–1306.

Zeman, Kryštof, Éva Beaujouan, Zuzanna Brzozowska, and Tomáš Sobotka (2018) "Cohort fertility decline in low fertility countries: Decomposition using parity progression ratios," *Demographic Research*, Vol. 38, pp. 651–690.