

The Effect of Interventions to Reduce Fertility on Economic Growth*

Quamrul H. Ashraf

Williams College

David N. Weil

Brown University and NBER

Joshua Wilde

Brown University

May 22, 2010

Abstract

We assess quantitatively the effect of exogenous reductions in fertility on output per capita. Our simulation model allows for effects that run through schooling, the size and age-structure of the population, capital accumulation, parental time input into child rearing, and crowding of fixed natural resources. The model is parameterized using a combination of microeconomic estimates, data on demographics, and natural resource income in developing countries, and standard components of quantitative macroeconomic theory. We apply the model to examine the effect of an intervention that immediately reduces TFR by 1.0, using current Nigerian vital rates as a baseline. For a base case set of parameters, we find that an immediate decline in the TFR of 1.0 will raise output per capita by approximately 13.2% at a horizon of 20 years, and by 25.4% at a horizon of 50 years.

1 Introduction

Our goal in this research is to quantitatively analyze the economic effects of interventions that reduce fertility in a developing country. Concretely, we ask how economic measures such

*We thank participants of the Fourth Annual Research Conference on Population, Reproductive Health, and Economic Development for comments. Financial support from the William and Flora Hewlett Foundation and the MacArthur Foundation is gratefully acknowledged. The usual disclaimer applies.

as GDP per capita would compare in the case where some exogenous change reduces fertility to the case where no such exogenous change takes place. The answer to this question will be very different from simply observing the natural co-evolution of fertility and economic development.

How declining fertility affects economic growth is a old question, going back at least to Malthus. The modern setting of the problem was posed by Coale and Hoover some 50 years ago. Over the last half century, the consensus view has shifted from fertility declines having strong effects, to their not being very important, and recently back toward assigning them some significance (Kelley, 2001; Sindig 2009).

For an issue of that has been studied for so long, and with such potential import, the base of evidence regarding the economic effects of population growth is rather weak. In some ways, this should not be a surprise. Population growth changes endogenously as an economy develops. Thus at the macroeconomic level it is very hard to sort out the direct effects of population growth other factors. Also, the lags at which population affects economic outcomes may be fairly long. Much of the current thinking about the effects of fertility decline relies on results from cross-country regressions, in which the dependent variable is growth of GDP per capita and the independent variables include measures of fertility and mortality, or else measures of the age structure of the population. Unfortunately, because of problems of omitted variables and reverse causation, the ability to draw inferences from the conditional correlations in these regressions is very weak (See Deaton 1999 for an critique).

While cross-country regressions suffer from severe econometric problems, they do have the advantage – if one is interested in studying the aggregate effects of fertility decline – of focusing on the right dependent variable. By contrast, a good many microeconomic studies examine the link between fertility at the household level and various outcomes for individuals in that household (for example, wages, labor force participation, education, etc.). These studies cannot directly answer the question of how fertility reduction affects the aggregate economy for three reasons. First, many of the effects of such reduction run through channels external to the household – either via externalities in the classic economic sense (for example, environmental degradation) or through changes in market prices, such as wages, land rents, and returns to capital. Second, even ignoring the issues of external effects, aggregating the different channels by which fertility affects economic outcomes is not trivial. Finally, as in the macroeconomic literature, the long time horizon over which effects of fertility change will affect the economy limits the ability of a single study to capture them.

In this research we pursue a “third way” in examining the link from fertility decline to aggregate economic growth. In particular, we build up a macro estimate starting from

microeconomic evidence on the effect of fertility decline, using economic theory to guide us in putting together the different channels by which fertility reduction works, both internal and external to the household. More specifically, we build a simulation model that takes proper account of both general equilibrium effects and the dynamic evolution of population age structure, capital accumulation, resource depletion, and so on. Throughout the research, the focus will be on giving a quantitative analysis of changes in fertility, so that we can estimate how much extra output a given intervention (for example, a drop in the TFR by 1, or a shift from the UN medium to low fertility projection) will produce over a specific time period. We hope, by showing how behavioral effects that are often studied in isolation can be integrated to answer macroeconomic questions, to reorient the academic discussion of population and development along more quantitative and practical lines.¹

There are several advantages to this approach. The simulation-based methodology allows us to take into account both general equilibrium effects and the dynamic effect of fertility reductions through channels including the evolution of the size and age-structure of the population, accumulation of physical and human capital, and resource crowding. As will be seen below, in the case of many of the channels by which fertility affects macroeconomic outcomes that we consider, it is relatively easy to calculate the steady state effect of a fertility reduction analytically. However, the transition period before an economy is well approximated by this steady state is often many, many decades. Presumably, policy makers in the real world are interested in outcomes over a shorter time horizon. The simulation approach also permits analysis of the strength of the various mechanisms at work. The simulation model that we build is general, but it has characteristics that can be tailored to the situation of particular countries. In addition to country-specific demographic characteristics (vital rates, initial age structure), the model can incorporate country-specific measures of the role of natural resources in aggregate production, the openness of the capital market, and (in future versions, we hope) other institutional characteristics.

¹In its analysis of outcomes in different demographic scenarios, our model resembles the RAPID model (see Abel, 1999). On the demographic side, the models are similar. The economic aspects, however, are completely different. In the RAPID model, the path of total GDP in a country is held fixed as different demographic scenarios are considered. Thus, for example, a halving of population relative to baseline will mechanically lead to a doubling of income per capita. This approach completely ignores the productive effect of labor, and thus leads to an unreasonably large projected impact of demographic change on the standard of living.

Young (2005) simulates the effect of the AIDS epidemic in South Africa on per capita income, using a Solow model somewhat similar to ours/ Relative to our work, however, Young is more concerned with long-run effects whereas we emphasize transition paths. Our methodological approach is also somewhat different from that of Young, in that we rely as heavily as possible on well-identified econometric estimates produced by other authors, rather than on producing our own estimates.

It is also important to note the hurdles that stand between a finding that interventions to reduce fertility raise output per capita (if that is what we find) and a conclusion that such interventions would constitute good policy. First, our analysis says nothing at all about the methods, costs, or welfare implications of such interventions. Second, GDP per capita is not necessarily the correct welfare criterion. The question of how a social planner should treat the welfare of people who may not be born as a result of some policy is notoriously difficult (Razin and Sadka, 1995; Golosov et al., 2007).

Existing literature has discussed a number of channels that lead from demographic change to economic outcomes. At the risk of some intellectual straight-jacketing, we classify these effects as follows. The most basic effect of population on output per capita is through the congestion of fixed factors, such as land. We call this the *Malthus effect*. A second channel is “capital shallowing” that results from higher growth in the labor force. We call this the *Solow effect*. Three channels run through the age structure of the population, which is a function of past fertility and mortality rates. First, a reduction in fertility leads with a lag to a higher ratio working age adults to dependents. Holding income per worker constant, this mechanically raises income per capita. We call this the *dependency effect*. Second, a concentration of population in their working years may raise national saving, feeding through to higher capital accumulation and higher output. We call this the *life-cycle saving effect*. Work by Bloom and Williamson on the “demographic dividend” has stressed a combination of the dependency and life cycle savings effects. Third, slower population growth shifts the age distribution of the working age population itself toward higher ages. In developing countries this increase in average experience would be expected to raise productivity, even though in more developed countries the shift into late middle age might lower productivity. We call this the *experience effect*. Another effect of reduced fertility is to lower the quantity of adult time that is devoted to child rearing, freeing up more time for productive labor. We call this the *labor supply effect*. Reductions in fertility are often associated with an increase in parental investment per child. We call this the *child quality effect*. Finally, an increase in the size of the population may raise productivity directly, by allowing for economies of scale, or may induce technological or institutional change that raises income per capita. We call this the *Boserup effect*. In this paper we attempt to quantify the first six of these effects (Malthus, Solow, dependency, life cycle saving, experience, labor supply, and child quality).

The rest of this paper is structured as follows. Section Two presents our simple demographic model, discusses the fertility interventions we model, and shows the dynamic paths of population size and age structure in response to the intervention. Section Three presents the economic model and discusses our choice of base case parameters. Section Four presents

simulation results for the base case model, and then discusses the sensitivity of results to altering our parameter assumptions. Section Five looks more deeply at different choices regarding the investment rate and how they interact with demographic change. Section Six similarly goes into greater depth regarding assumptions about the role of the fixed factor in production. In Section Seven, we apply our model to consider a different fertility intervention scenario than what we analyze in the rest of the paper: specifically, we look at the difference in economic outcomes resulting from a shift from the UN medium fertility scenario to the low fertility scenario in Nigeria. Section Eight concludes.

2 Demographic Model and Fertility Intervention

The demographic part of the model takes age-specific mortality and fertility schedules as inputs to project the population over time.² In practice, population is divided into 5-year age groups, and each time period in our model corresponds to five years.³

As discussed above, our analysis is focused on considering interventions that alter the path of fertility from what would occur along some baseline. Our model can be easily tailored to consider different baseline and interventions scenarios. For most of this paper, we examine simple baseline and intervention constructed using demographic data from Nigeria. This simple approach allows us to better understand the timing by which different demographic-economic channels operate. In Section Seven of the paper we consider a richer demographic scenario, based on alternative UN population projections for Nigeria.

Our simple baseline and alternative scenarios are build up using current vital rates from Nigeria. Figure 1 shows the survivorship function implied by the female life table for Nigeria in 2006.⁴ The age-specific fertility profiles for the baseline case of no intervention,

²For simplicity, our demographic projections are performed on a closed, female-only population. Considering a population of both males and females, however, would not qualitatively alter the results of our model as long as the sex-ratio-at-birth remains fixed over time.

³Formally, a population composed of n age-groups is represented by an n -dimensional vector, N_t that evolves according to:

$$N_{t+1} = \begin{cases} P^b \cdot N_t & \text{if } t < T \\ P^a \cdot N_t & \text{otherwise,} \end{cases}$$

where P^b and P^a are the $n \times n$ projection matrices before and after the shock, $N_0 > 0$ is given, and the shock period, T , is determined to occur after the pre-shock population has attained a stable age structure and rate of growth. A population projection matrix is composed of age-specific net maternity rates along the first row and age-specific survivorship rates along the sub-diagonal. The stable population growth rate implied by a projection matrix is given by its largest, real eigenvalue, and the stable age-structure by the corresponding eigenvector.

⁴This data is obtained from the WHO's *Life Tables for WHO Member States*, an online repository accessible at: http://apps.who.int/whosis/database/life_tables/life_tables.cfm




Figure 1: Female Life Table Survivorship Function, Nigeria 2006

along with the cases of 1.0 and 0.5 reductions in TFR are shown in Figure 2.⁵ We use the current fertility and mortality schedules to construct a stable population, and in the baseline scenario we assume that fertility and mortality will be constant going forward. The “base case” intervention that we consider is an instantaneous reduction of 1.0 in the TFR. In addition to this base case, we consider two alternatives: an instantaneous reduction of 0.5 in the TFR, and phased-in reduction in the TFR of 1.0 that takes place in a geometric fashion over 25 years.

Figure 3 shows the sizes of the overall population in the baseline case of no intervention, along with the case of an immediate intervention and the case of a phased-in intervention. Figure 4 depicts the population paths in the aforementioned intervention scenarios relative to the no-intervention baseline.

Figures 5, 6, and 7 show the working age (15-64), young (under 15), and elderly (65+) fractions of the population, respectively, prior to and following an immediate drop in fertility. In the pre-shock steady state, 55.46 percent of the population is working age; in the post-shock steady state, this fraction rises to 59.19 percent. In addition, as discussed by

⁵The age-specific fertility data for the baseline case of no intervention corresponds the “medium variant” profile reported by the UN in *World Population Prospects: The 2008 Revision* for Nigeria in 2005. In the case of an intervention, this profile is scaled down (proportionally across age groups) according to the size of the reduction in the TFR.




Figure 2: Age-Specific Fertility Profiles Under Different Intervention Scenarios




Figure 3: Population Size Under Different Intervention Scenarios




Figure 4: Population Size Relative to the No-Intervention Baseline

Bloom and Williamson (1998), a reduction in fertility leads to a period of several decades in which the working age-fraction is above its steady state level.

3 Economic Model and its Parameterization

3.1 Production

In our base case model, aggregate production is given by a standard Cobb-Douglas production function. The factor inputs are land (which we use as a shorthand for all fixed factors of production), capital, and human capital, so that aggregate output in period t , Y_t , is:

$$Y_t = A_t K_t^\alpha H_t^\beta X^{1-\alpha-\beta}$$

where $\alpha + \beta \leq 1$, X is a fixed arbitrary stock of land and A_t is productivity.

We assume fairly standard values for factor shares: we set $\alpha = 0.3$ and $\beta = 0.6$, meaning that the implied share of land is 10 percent. In a later section we revisit the role of fixed factors of production. We consider the sensitivity of our results to both the share of land in national income and the elasticity of substitution between land and other factors of production. We also examine data on natural resource shares of national income.




Figure 5: Working-Age Fraction of the Population in the Base Case Intervention




Figure 6: Under-15 Fraction of the Population in the Base Case Intervention




Figure 7: Over-65 Fraction of the Population in the Base Case Intervention

Productivity grows at an exogenous rate that does not respond to any of the changes in the model. For convenience, the growth rate is set to equal the stable population growth rate in the pre-shock period times the share of land, so that income per capita is constant in the pre-shock steady state. Because all of our results entail a comparison of income in the case of a fertility intervention to the case where no intervention takes place, the underlying rate of technological change is of very little importance.

3.2 Physical Capital Accumulation

In our base case, we handle capital accumulation extremely simply, by making the Solovian assumption that a fixed share of national income is saved in each period. Young (2005) makes the same assumption in his analysis of HIV/AIDS in South Africa. Accordingly, the stock of capital in period t , K_t , evolves over time according to:

$$K_{t+1} = sY_t + (1 - \delta)K_t,$$

where s and δ are the fixed saving and depreciation rates, respectively. We assume that the annual savings rate is 8.55 percent, which corresponds to the investment share of real GDP reported by the Penn World Table (version 6.3) for Nigeria in 2005. We assign a standard value to the depreciation rate of 5 percent.

In Section Five, we consider two alternative models of investment. First, we allow for variable age-specific savings rates, with workers in their prime earning years having a high saving rate. This introduces an additional channel through which demographic change affects growth.⁶ Second, we consider the case of an economy that is fully open to international capital flows. This shuts off the “Solow” channel by which slower growth of the labor force raises the level of capital per worker.

3.3 Human Capital

We model an individual’s human capital as a function of his or her schooling and experience. We assume that human capital inputs of individuals with different characteristics are perfectly substitutable. Thus the stock of human capital in period t , H_t , is:

$$H_t = \sum_{15 \leq i < 65} (h_{i,t}^s \times h_{i,t}^e) N_{i,t},$$

where $N_{i,t}$ is the number of individuals of age i in the population in period t . We assume that children enter the labor force at 15 and workers leave the labor force at 65.

Our treatment of schooling and experience is standard. Years of schooling are aggregated into human capital from schooling using the piecewise log-linear specification:

$$h_{i,t}^s = \begin{cases} \exp[\theta_1 S] & \text{if } S \leq 4 \\ \exp[4\theta_1 + \theta_2(S - 4)] & \text{if } 4 < S \leq 8 \\ \exp[4\theta_1 + 4\theta_2 + \theta_3(S - 8)] & \text{if } S > 8 \end{cases}$$

where we use values of $\theta_1 = 0.134$, $\theta_2 = 0.101$, and $\theta_3 = 0.068$, based on Hall and Jones (1999). The return to schooling will be relevant for the exercises we conduct because reductions in fertility will raise the average level of schooling.

Human capital from on-the-job experience for a worker of age i in any period t , $h_{i,t}^e$, is computed as:

$$h_{i,t}^e = \exp[\phi(i - 15) + \psi(i - 15)^2]$$

where, based on Bils and Klenow (2000), who provide an estimate of the average return to experience in a sample of 48 countries, we use a ϕ value of 0.0495 and a value of -0.0007 for ψ . Experience will play a role in our simulations because declines in mortality and fertility will lead to a population with higher average age and thus higher average experience.

⁶There is considerable controversy about the applicability of such models to developing countries. See, e.g., Lee et al. (2001) and Deaton (1999).

We also expect that lower fertility will raise the average level of schooling. Models of the fertility transition stress the movement of households along a “quality-quantity” frontier, in which investment per child in health and education rises and the number of children falls. It does not follow from this observation, however, that the change in schooling that would result from an exogenous change in fertility is the same as the change that would accompany declining fertility when both measures are evolving endogenously.

To assess the fertility intervention in which we are interested, we use results from Joshi and Schultz (2007), who analyzed a randomized intervention in Matlab, Bangladesh. Joshi and Schultz found that a 15 percent reduction in TFR resulting from the intervention led to an increase of 0.52 years of schooling for males aged 9-14.⁷ Since our base case intervention (TFR reduced by 1.0) corresponds to a percentage reduction of 18.8 percent in the TFR for Nigeria in 2005, the relevant increase in schooling is $0.52 \times \frac{18.8}{15.0} = 0.65$ years of schooling.

In our simulations we thus increase schooling by 0.65 years for cohorts in school at the time of the decrease in fertility. The effect of this rise in schooling on the average level of human capital will depend on the initial level of schooling because, as discussed above, the percentage return to schooling falls with the number of years of schooling. In our base case simulation, we consider the case where initial schooling is between 5 and 8 years, so that the return to schooling is 10.1 percent per year.

3.4 Labor Supply Effects

Raising of children requires a good deal of labor. That labor is spread over many years, and divided among many individuals, but the largest piece generally comes from the child’s mother. Reduced fertility should thus potentially increase the labor supply of women. Assigning a quantitative magnitude to this effect is difficult for several reasons.

- There are obviously strong economies of scale in child rearing. Thus, the time cost of a first child is far higher than the marginal time cost of subsequent children. For example, Tiefenthaler (1997), examining data from Cebu, Philippines, finds that 14 months after the birth, female labor market hours had declined by 39% in the case of first births but by only 10% if there were already children aged 0-5 in the household. If there were both children aged 0-5 and children aged 6-17 in the household, female labor market hours were actually slightly increased 14 months following a birth.

⁷This coefficient of 0.52 is derived from Table 9, Column 2 in their paper. They report a standardized beta of 0.54 to which we apply the standard deviation for years of schooling of 0.95 from their summary statistics.

- Not all time spent on children is subtracted from production. A good part of time devoted to child rearing may be at the expense of leisure or, in the case of siblings, human capital investment (which is not counted as part of national income).
- Child-rearing is often combined with productive activities, especially in developing countries. For example, a woman may carry a baby in a sling or watch children out of the corner of her eye while she works at a productive task. The “cost” of child-rearing in terms of productive labor in this case would only be the decrement in productivity that results from such multitasking.

Despite these caveats, the time cost of child-rearing may still be an significant component in the economic response to fertility decline. We measure the effect of fertility change on labor supply by specifying a parameter we call the labor market time cost of a marginal child. Summarizing all these considerations in a single parameter is obviously too simplistic – but in doing so we at least have a concrete measure which can be implemented in our model. Specifying the time cost of the marginal child might also be considered problematic because the marginal cost would be expected to fall with the number of children. However, for the experiments we are considering, the TFR remains above two, often by a good amount. Thus we are only considering the marginal cost of third and higher parity children, where we would not expect the decline in marginal cost to be so extreme. This matches the findings of Holmes and Tiefenthaler (1997), who conclude that the marginal time cost of children is roughly constant for the third and higher children.⁸

Mechanically, we implement the labor supply effect by increasing female labor force participation in each year by the hypothesized change in age-specific fertility multiplied by the labor market time cost (in years) of a marginal child. For example, if in our experiment age specific fertility of women aged 25-29 drops from 0.26 to 0.21 (as it does in our base case scenario of a decline in the TFR of 1.0 in Nigeria), and if the labor market time cost of a marginal child is one year, then labor force participation for women in this age group would rise by 5 percentage points.⁹

⁸Because of heterogeneity in completed fertility, a reduction in the TFR from three to two will not mean that all children not born would have been parity three. Instead, some would have been higher parity, and some would have been first or second children. Thus our method will understate the increase in labor input that results from such a reduction in fertility.

⁹Although it might seem problematic to “charge” the entire time cost of a child to the mother in the year of the child’s birth, we do not view this as too distortionary of reality for two reasons. First, time costs of child-rearing are indeed concentrated in the first years of life. Second, because we are considering an age specific fertility schedule that assigns a fractional number of births per year to each woman, the pattern of labor force increase that is generated by our method will look similar to what would result if each birth reduced labor force participation over a longer period of time. It is true, however, that our method

There only remains the question of choosing the base case parameter value for the time cost of children. In the Cebu data used by Tiefenthaler, weekly labor market hours fall from 10.4 prenatally to 5.0 at two months, 6.6 at six months, and 9.5 at 14 months for women who have other children aged 0-5 in the household; and from 13.1 prenatally to 7.6 at two months, 11.3 at six months, and 13.8 at 14 months for those with both children aged 0-5 and 6-17 in the household. Crudely interpolating these data, and allowing for an almost total cessation of labor market activity in the first month after delivery, hours averaged over the first year are reduced roughly 5 per week in the first group and 3 per week for the second group. Weekly labor market hours for men in the same households do not change much in response to a birth, and are equal to roughly 40. So in this data, women in these two groups lose, say .125 or .075 years of full-time equivalent labor market input in the first year after the birth of a marginal child. The complete or nearly complete recovery of labor hours by 14 months after delivery suggests that the decrement in subsequent years should be very small. On the other hand, there are a good number of these years. Further, we have data on neither the efficiency loss by women with small children who are working, nor any long-term health consequences of multiple pregnancies that might impede labor input for many years. As a rough guess we specify a labor market time cost of 1/2 year per marginal child.¹⁰

Table 1 shows the age-specific labor force participation rate (LFPR) for Nigerian women in 2005 and the implied levels of LFPR if the TFR were reduced by one, and also if the TFR were equal to 2.41 (which is the UN median forecast for 2050), using our baseline value of the time cost of a marginal child. For comparison, we also show the age specific LFPR for men.¹¹

3.4.1 Other Channels Not Covered

A simulation study such as ours is useful only to the extent that it covers all of the quantitatively important channels by which a change in fertility affects the macroeconomy. We

may slightly front-load the effect of lower fertility on labor force participation, both because we ignore child-rearing costs in later years and also because we apply our marginal rate to all births, whereas higher order births are concentrated at older ages.

¹⁰Bloom et al. (2009) examine the effects of fertility decline in female labor force participation in cross country data, using changes in abortion laws as an instrument for fertility. They estimate the change in the age-specific female labor force participation which results from an decrease in the TFR of one. Taking the weighted average by female population age structure, such a decline produces an increase of 13.51% in total female labor force participation. Their estimates imply an average labor market cost per marginal child of 4.4, which is far higher than the figure we use. However, the estimates in the Bloom et al. study are identified by variation in high income countries, where baseline fertility levels are far lower and where separation between home and workplace generally means that child-care and labor market input are mutually exclusive.

¹¹Source: ILO LABORSTA database for 2005.

Age	2005 Nigeria Male	2005 Nigeria Female	Reduce TFR by 1	Increase Factor	Reduce TFR to 2.41	Increase Factor
15-19	25.9%	11.8%	13.0%	1.101	15.3%	1.293
20-24	59.9%	28.1%	30.2%	1.074	34.1%	1.215
25-29	90.3%	39.4%	41.8%	1.062	46.5%	1.179
30-34	97.7%	41.6%	43.8%	1.052	47.9%	1.150
35-39	98.9%	51.6%	53.1%	1.028	55.8%	1.082
40-44	98.8%	57.1%	57.7%	1.010	58.7%	1.029
45-49	99.2%	67.0%	67.1%	1.002	67.4%	1.006

Table 1: The Labor Supply Response

have tried to keep our framework transparent and open, so that we (or someone else) can add other channels if there is an appropriate basis. Here we discuss some potential channels that we have not included, either because we think that they are of secondary quantitative importance or because we did not have a basis for quantifying them.

3.4.2 Boserup Effects

There are several channels through which higher population density could positively affect the level of income per capita. Boserup’s work on agriculture stressed that as population rose, farmers were induced to switch to more intensive methods, which meant that the land constraint did not end up lowering income per worker. A more generalized version of this effect would be that higher population would induce technological progress more generally, that is outside of agriculture, either out of “necessity,” or because more people raises the likelihood of someone having a productive idea (Jones, 1995). A completely different channel by which population growth could raise output would be by allowing for better trade and economies of scale in production. In the African context, it is often noted that long distances and poor roads lead to extremely high cost of trade.

We do not include any of these channels in our analysis, for several reasons. Regarding agriculture, some of the possibilities for intensification and substitution of other inputs for land are already included in our production function approach. In particular, Section Six (and the literature on which it draws) discusses evidence on the substitutability of inputs for land. The intensity of cultivation varies enormously in sub-Saharan Africa, but the Boserupian description in which fertile land can easily be shifted from fallow to continuous cultivation seems inappropriate for most countries. Data show that over the last decades cultivation in Africa has increasingly expanded onto marginally suitable land (Weil, 2008).

Regarding gains from scale as population rises, we are of two minds. On the one hand, we agree that costly transport raises trade costs and leads to an inefficient scale of production in many African countries (Gollin and Rogerson, 2009). However, it is not clear that population growth over the next several decades will lead to increases in rural population density that would facilitate trade. Rather, Africa is rapidly urbanizing, and so much of the growth of population in the next decades will end up in already large cities (Weil, 2008). It is hard to believe that mega cities such as Lagos do not already have sufficient size to achieve economies of scale in production.

On a more prosaic level, we were not able to find quantitative estimates of the size of Boserup effects that we could incorporate into our model.

3.4.3 Effects Through Health Improvements

Another channel through which fertility declines could possibly affect output is through improvements in health. These could result from the same quality-quantity shift that we model in the case of education. Ashraf, Lester, and Weil (2008) discuss how improvements in health can be translated quantitatively into productive human capital. However, Joshi and Schultz find no effect of the fertility intervention in Matlab, Bangladesh on child health.

4 Basic Results

Figure 8 shows the paths of physical capital per worker, human capital per worker, labor input per worker, income per worker, and income per capita in our simulation, using the base case parameters discussed above. As in all the figures that follow, we show results relative to a baseline in which no fertility intervention takes place.

The path of output per worker reflects the dynamics of human and physical capital per worker, as well as land per worker (which we do not show, but which can be inferred from Figure 8).

Figure 8 shows that in our base case the long-run effect of an exogenous reduction of 1.0 in the TFR is to raise output per capita by 25.4% percent relative to the baseline of no fertility intervention at a horizon of 50 years. At a 20 year horizon, the increase in income per capita is 13.2% percent. As discussed below, the source of income gains varies with the time horizon considered. In the early years of the simulation, the biggest gains come through changes in the dependency ratio. Over the long run, the biggest effects come through reduced pressure on fixed resources such as land. Because the intervention that we consider is a permanent reduction in fertility, and because in the baseline fertility remains




Figure 8: The Base Case Intervention Scenario

far above replacement indefinitely, the scenario implies that the income gain from reduced fertility continues to grow indefinitely.

Figure 9 compares our base case with two alternative demographic scenarios. First, an immediate reduction in TFR of 0.5, rather than the reduction of 1.0 in the base case. Second, the phased-in reduction of TFR by 1.0 discussed above. Comparing the two different scenarios with an immediate reduction in fertility shows that the effect of fertility is roughly linear, at least in the early stages of the simulation. A reduction in TFR of 1.0 yields increases in income per capita that are slightly more than twice as big as those yielded by a reduction of 0.5. For example at a horizon of 50 years, the increase in income per capita is 11.7 percent in the case of a reduction in TFR of 0.5, vs. 25.4 percent in the base case of a reduction of 1.0. The gap grows at longer time horizons, as the difference in population growth rates in the two scenarios compounds. Comparing the base case scenario of an immediate reduction in fertility to a phased-in reduction, the largest differences are not surprisingly in the early years. After 20 years, income per capita is 6.9 percent higher in the case of a phased-in reduction in fertility, vs. 13.2 percent higher in the case of an immediate reduction.

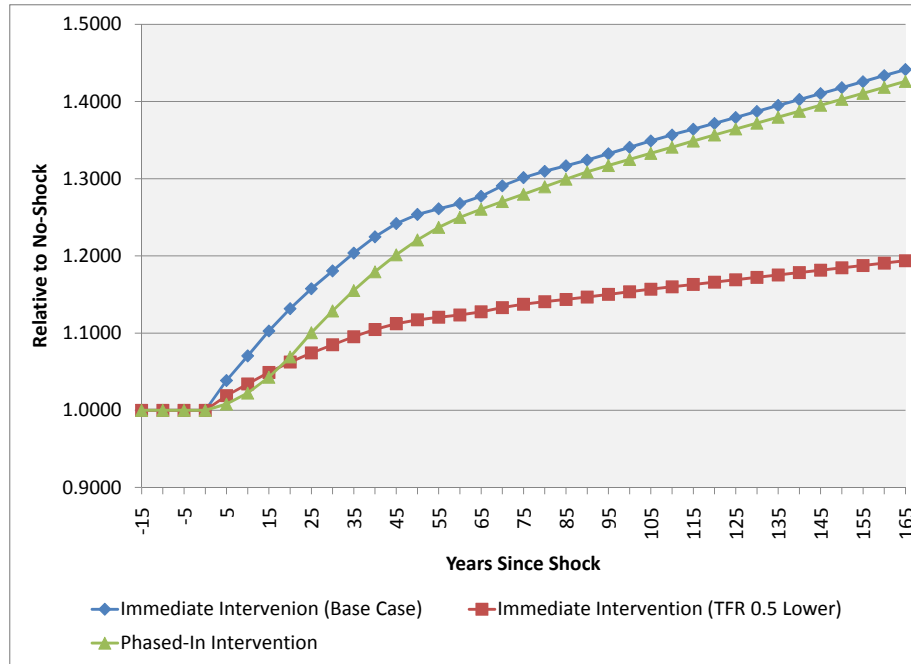


Figure 9: Income Per Capita Under Different Intervention Scenarios

4.1 Component Channels

As discussed above, demographic change affects economic outcomes through a number of channels which may operate at different relative intensities at different time horizons. It is of interest to decompose the effect of a fertility intervention into the parts that run through these different channels. Some caution is necessary, however, because there are clearly interactions among the different effects. In particular, the effect of fertility through any one channel will depend on which other channels are operative. For example, the effect of increased labor force participation of working age adults will be larger or smaller, depending on the fraction of the population made up of such adults. To address this problem, we do all of our analysis of the effects of fertility through different channels under the assumption that all the other channels are operative – that is, we consider the results in our full simulation relative to the case where one channel of effect is deleted (an alternative would be to assume that no other channels were operative).

We begin by looking at several channels individually. This allows us to perform an analysis of the sensitivity of our results to assumptions about key parameters. We then do a full decomposition showing the relative importance of different channels at different time horizons.

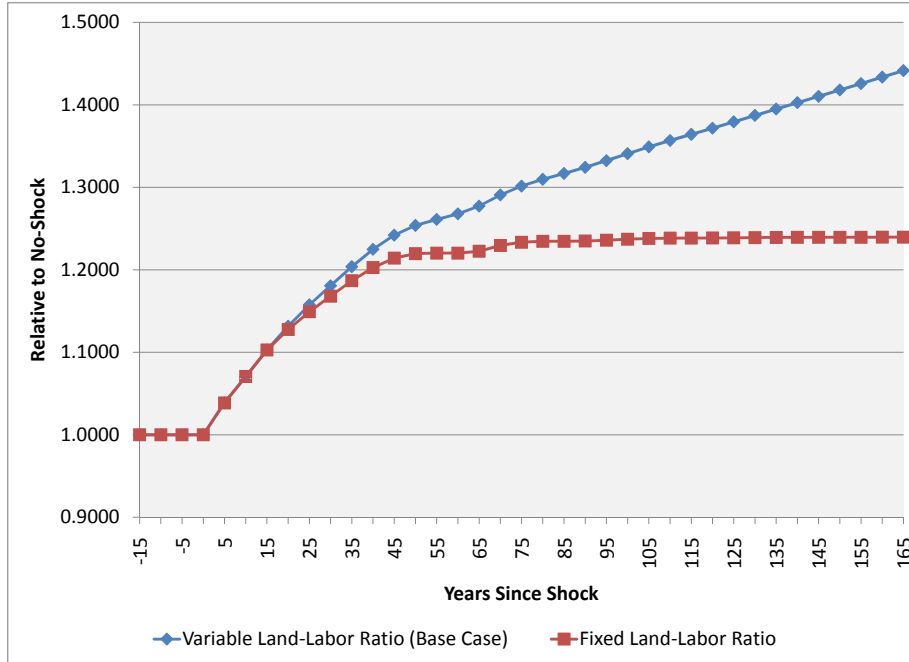


Figure 10: The Effect of the Land-Labor Ratio on Income Per Capita

To gauge the importance of non-reproducible factors, we conduct a simulation in which the level of land per worker follows the same path as it does in the baseline scenario. In other words, we ignore the effect of lower fertility in preventing the land/labor ratio from falling, while allowing for all the other economic effects of fertility decline. Figure 10 compares the path of output per capita in this scenario to the base case. The figure illustrates the extent to which the classic Malthusian channel operates only over relatively long time horizons. For the first 35 years following the shock to fertility, the path of income per capita with the Malthus effect is suppressed looks only slightly different than when the effect is present (for the first 15 of these years, this is mechanically true because the ratio of land to labor has not changed, except for the small effect from the labor supply of mothers discussed below). On the other hand, by about 70 years following the shock to fertility, all other sources of growth in income per capita (relative to the baseline) have petered out, and it is *only* the Malthus effect that produces continuing growth (recall that in our alternative scenario the growth rate of population is permanently lower than in the baseline).

We conduct a similar exercise to look at the importance of the capital-shallowing channel. Specifically, we construct a scenario in which the level of physical capital per worker is the same following a shock to fertility as it is in the baseline case. The result is show in Figure 11. As with the Malthus effect, for the first 15 years following the shock, suppressing the Solow effect makes no difference to the level of income per capita because

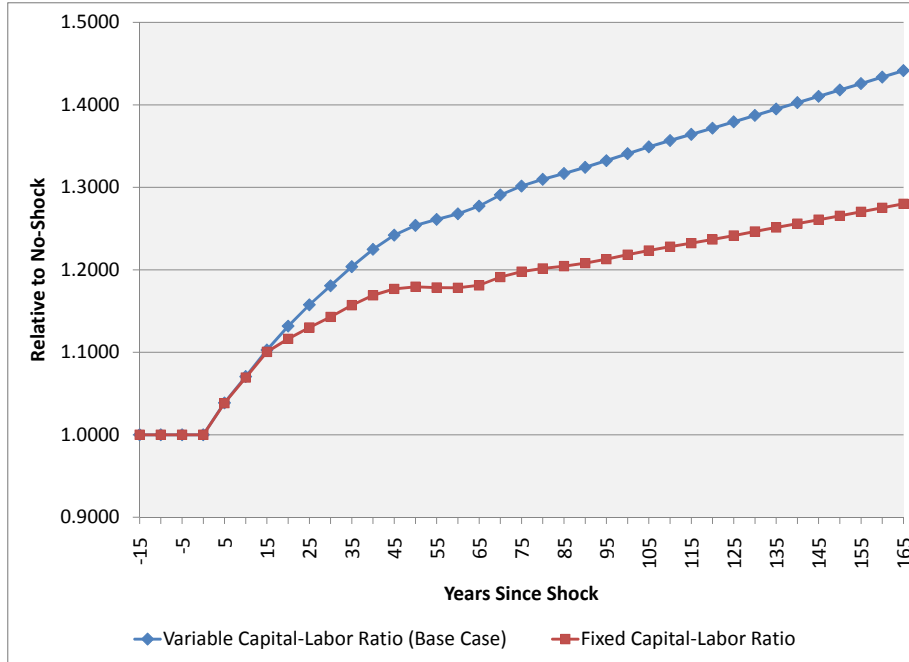


Figure 11: The Effect of the Capital-Labor Ratio on Income Per Capita

the number of workers does not differ between the baseline and alternative cases. Between year 20 and year 65, the Solow channel is at its strongest. At year 65, output per capita is 18.1% percent above baseline when the Solow effect is suppressed, compared to 27.7% percent above baseline when the Solow effect is present.

Figure 12 shows the dependency effect channel, comparing the baseline to the case where the dependency ratio remains fixed at its pre-shock steady-state level over time. Here, not surprisingly, the phase in of the effect is almost immediate. Fifteen years after the reduction in fertility, income per capita is only 1.8% percent above the baseline when the dependency effect is suppressed, vs. 10.3% percent above baseline when the effect is present. From this point onward, the difference between the two paths is roughly constant.

Figure 13 looks at the experience channel. We compare the baseline to the case where the experience effect is shut down (by assuming that the return to experience is zero). As the figure shows, the experience effect is not of great import. For example, at a 50 year horizon, the increase in output per capita is 1.4 percentage points lower when the experience effect is suppressed than in the base case. The reason that this effect is so small may be that, in our simulation, even with a reduction in TFR the labor force remains incredibly young. Simulations that looked at more slowly growing populations might find a bigger effect. It is also interesting to note that in the periods immediately following the shock to fertility, output per capita is actually slightly *higher* in the case where there is no experience




Figure 12: The Effect of the Dependency Ratio on Income Per Capita

effect than in the base case. This is because the increase in labor supply that results from a reduction in fertility is primarily among young, inexperienced workers.

In Figure 14 we vary the amount of extra human capital that is produced by an additional year of schooling. As discussed above, our base case assumption is that the return to schooling is 10.1 percent per year of schooling, which is consistent with standard estimates if initial schooling is between 5 and 8 years. In the figure we show alternative paths for two different levels of returns to education (13.4 percent and 6.8 percent, which are consistent with schooling less than five years and 9+ years, using the estimates of Hall and Jones) as well as a return of zero, which shuts down this channel completely. The figure shows that schooling plays an appreciable role in determining the economic effects from reduced fertility. At a horizon of 50 years, for example, output per capita is 19.2 percent above baseline in the scenario where the return to schooling is zero (which is the same as if there were no increase in schooling), vs. 25.4 percent higher in the base case scenario. Thus, roughly speaking, schooling accounts for one quarter of the income gain from reduced fertility at this time horizon.

As would be expected, the effect of higher schooling due to reduced fertility phases in as the cohorts which received the additional schooling enter the labor force and replace those which did not. Thus for the first 15 years after the shock to fertility, this channel contributes little to higher income.

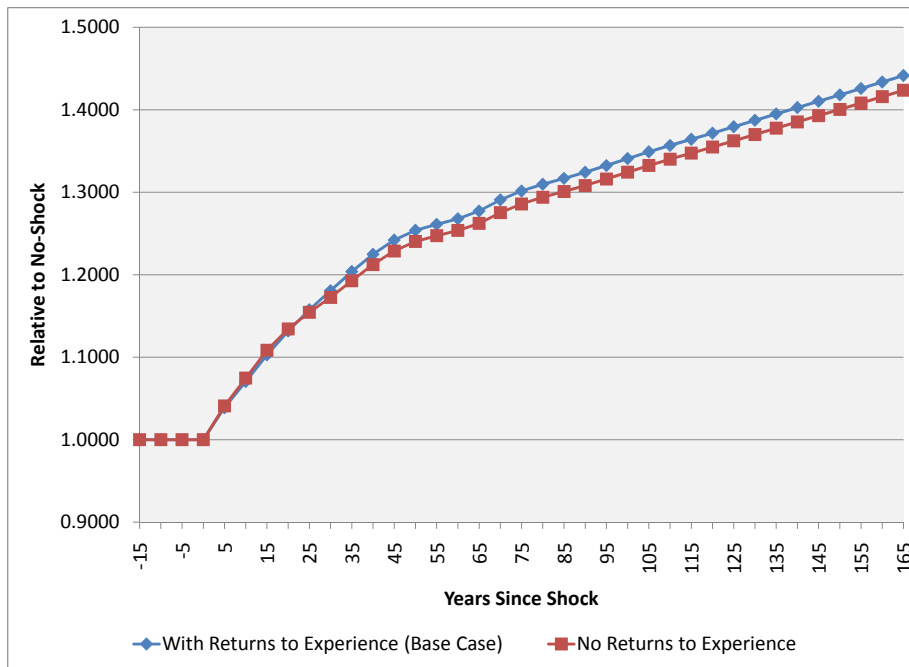


Figure 13: The Effect of Returns to Experience on Income Per Capita




Figure 14: The Effect of Returns to Schooling on Income Per Capita

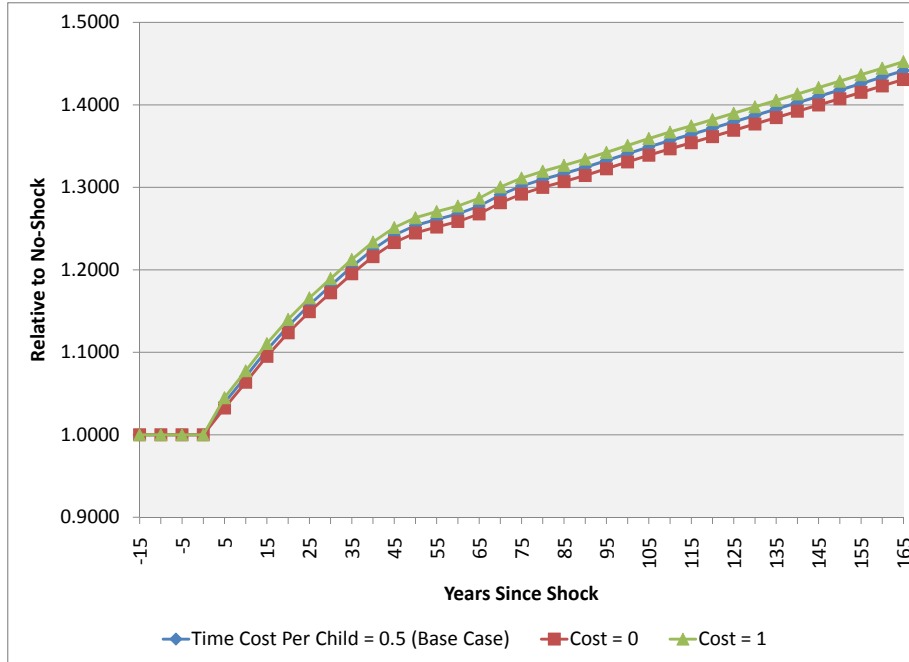


Figure 15: The Effect of the Labor Force Participation Response on Income Per Capita

Finally, Figure 15 examines the labor supply channel. We consider variations in the parameter reflecting the time cost of a marginal child. Specifically, we compare the path of income per capita using our baseline assumption of a cost per marginal child of 0.5 years (of labor market input) to alternatives of 1.0 and zero. The channel turns out to be surprisingly weak. At a horizon of 15 years, for example, income per capita in the base case is 10.3 percent above baseline when the marginal cost of a child is at its baseline value of 0.5, vs. 11.1 percent above baseline when the marginal cost of a child is 1.0, and 9.5 percent above baseline when the marginal cost of a child is zero.

4.2 Decomposition of Channels

Figure 16 presents a full decomposition of the fraction of the gain in income per capita at each point in time that is due to the different channels we study. Because of interactions among the different channels, the individual channel effects that we study above do not sum to the total effect of a decline in fertility on income per capita. To do a decomposition of the fraction of the gain in income per capita that is due to each channel, we thus proceed as follows. As above, we calculate the importance of individual channels by comparing the level of income per capita in each year in our baseline case to the level of income per capita when the channel is suppressed. We then add up these individual effects to get a proxy for

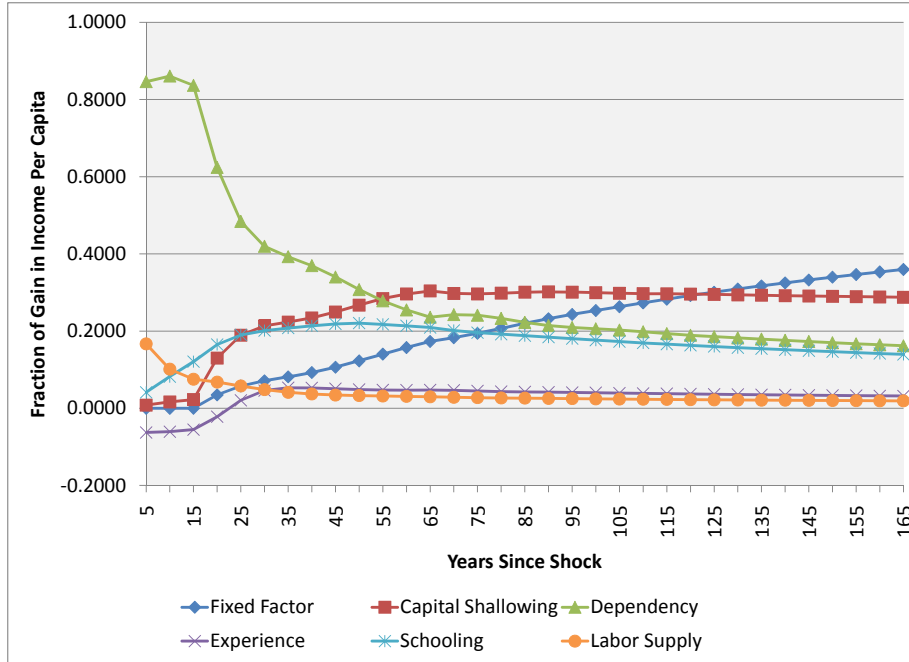


Figure 16: Decomposition of the Gain in Income Per Capita by Channel

the total effect that ignores interactions. We then divide the individual effects by the proxy for the total, to produce a share of the total income gain due to each effect at each time horizon.

The figure shows that the dependency effect is by far the dominant channel in the short run, explaining more than 80% of the income gain in the first 15 years of the simulation and only falling below 40% of the total after 35 years. The labor supply effect, which is conceptually very similar to the dependency effect, has a somewhat similar trajectory, although at a much lower level. At a horizon of 50 years, the four dominant effects are dependency (30.8 percent of the total), Solow (26.7 percent of the total), schooling (22.1 percent of the total) and Malthus (12.3 percent). At a horizon of 100 years, the same four effects are dominant, but in a different order: Solow (29.9 percent), Malthus (25.3 percent), dependency (20.6 percent), and schooling (17.7 percent).

5 Alternative Models of Investment

5.1 Age Specific Saving Rates

Discussions of the “demographic dividend” from reduced population growth often stress the benefits to national saving from having a large fraction of the population in its working

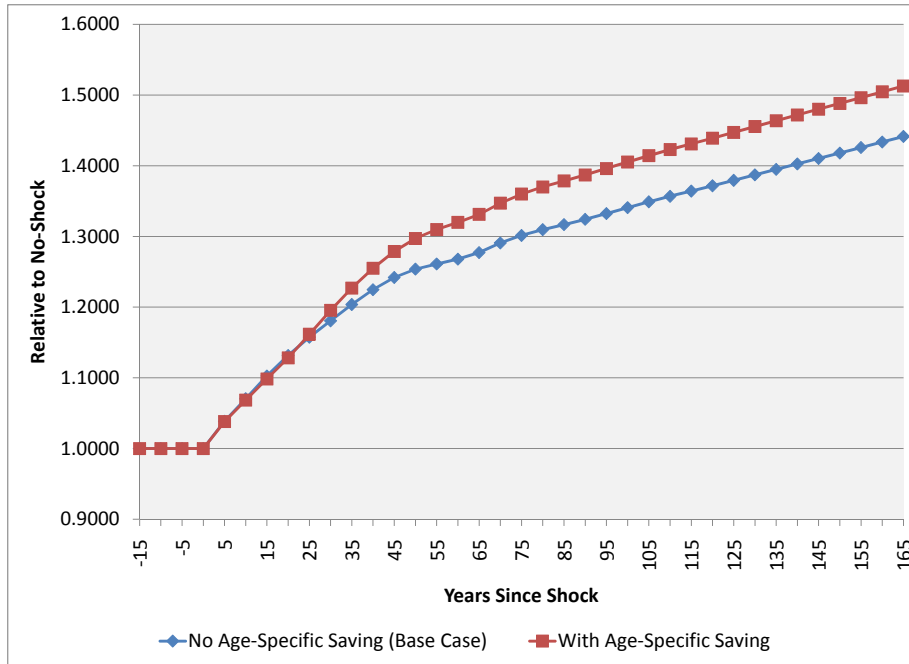


Figure 17: The Effect of Age-Specific Saving on Income Per Capita

years. In our model, we would ideally use data on age specific saving rates for Nigeria to evaluate the importance of this channel. However, such data is generally not available for the developing world, and the data that does exist shows little to no evidence of life cycle saving behavior in poor countries. For example, Deaton (1992) calculates household consumption and income over the life cycle in Côte d’Ivoire, and finds no clear relationship between age and saving, consumption, or even income. Thus, the assumption of a low constant saving rate for all ages may be more realistic than life cycle saving for developing countries.

In addition, inferring a national saving rate from household surveys is problematic. For example, Weil (1994) points out that data on household saving does not take into account the externality effects of saving across generations via bequests or other transfers. As the age structure of the population changes, so do the structure of bequests and incentives for saving among both the old and the young. It is therefore difficult to properly disentangle the general equilibrium effect of the age structure on the national saving rate.

Given the methodological and data problems with calculating age specific saving rates, we did not include this channel in our decomposition exercise, but simply assumed that the saving rate was low and constant over the life cycle according to the evidence we have. We are still trying to come up with a better way of implementing such a channel.

However, it may still be useful to determine an upper bound of the dynamic effect of life cycle saving on income per capita through changes in the age structure. To do so, we

use data on household age specific saving rates for the U.S. from Poterba (1994) to calibrate age specific saving rates for Nigeria. Specifically, we scale the magnitude of the life cycle saving profile in the U.S. to match the national saving rate in Nigeria from the Penn World Tables, conditional on the age structure of the Nigerian population. Such a methodology yields incredibly large swings in household saving rates over the life cycle – especially for a developing country such as Nigeria – from about -30% in early stages of life to about 40% in later stages. Nevertheless, this is useful for an upper bound exercise, since the true life cycle saving profile for Nigeria is undoubtedly much less pronounced.

Using these abnormally large changes in saving rates over the life cycle, Figure 17 compares the base case scenario to the alternative that incorporates age specific saving rates, showing that the path with life cycle saving does not differ much from the base case. At a horizon of 50 years, income per capita is only about 4.3 percentage points higher on the alternative path than in the base case, and only 6.5 percent higher after 100 years.

5.2 International Capital Flows

An important part of our results is driven by the assumption of Solovian saving. It is possible to adjust this assumption in a straightforward way even without building a life-cycle savings model, simply by assuming that the economy is open to international capital flows that equalize the return to capital around the world, at least up to a country fixed effect.¹²

Figures 18 (capital per worker) and 19 (income per capita) show that allowing for capital flows (assuming a fixed world interest rate) has relatively little effect on the behavior of income, although it does have a larger effect on the capital stock. With an open economy, capital stock rises rapidly in response to a fertility reduction because of the increase in effective labor supply that produces a nascent rise in the marginal product of capital; in contrast, in the base case, capital accumulates only slowly as income rises. Thus capital per worker is higher in the open economy case than in the base case for the first 25 years. In the long run, capital per worker is higher in the base case than in the open economy case because slower labor force growth drives down the marginal product of capital, which in the open economy leads to a capital outflow.¹³ The difference in output per capita between the

¹²Caselli and Feyrer (2007) make a strong case that marginal products of capital are almost completely equalized around the world.

¹³We simulate international capital flows in the following manner. Prior to the shock, capital accumulates in the usual closed-economy Solovian fashion. Note that this is equivalent to assuming that the economy is open to international capital flows but has a domestic savings rate such that there is no inflow or outflow in the pre-shock steady state. In other words, the marginal product of domestic capital in the pre-shock steady state is equal to the fixed world interest rate. Once the shock is applied, however, capital accumulates in such a fashion as to maintain its pre-shock steady-state marginal product over time.




Figure 18: The Effect of International Capital Flows on Capital Per Worker

base case of Solow savings and the case with perfectly open capital markets is surprisingly small – never more than a couple of percentage points in the first 50 years of the simulation.

6 The Role of Land in the Production Function

Our base case treatment of land involved assuming both a particular functional form (Cobb-Douglas, in other words unit elasticity of substitution) and a particular exponent on land in the production function. In this section we relax both of these assumptions. We adopt a CES production function in which we can specify an elasticity of substitution between a capital-labor-technology composite factor, on the one hand, and the fixed factor on the other:

$$Y_t = \left[(1-a) (A_t K_t^\alpha H_t^{1-\alpha})^{\frac{\sigma-1}{\sigma}} + a X^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}$$

where σ is the elasticity of substitution. If the fixed factor is paid its marginal product then its share of national income at time t , ϕ_t , will be:

$$\phi_t = \frac{a X^{\frac{\sigma-1}{\sigma}}}{(1-a) (A_t K_t^\alpha H_t^{1-\alpha})^{\frac{\sigma-1}{\sigma}} + a X^{\frac{\sigma-1}{\sigma}}}$$




Figure 19: The Effect of International Capital Flows on Income Per Capita

If the elasticity of substitution is not unity, the fixed factor’s share of national income will vary as capital and human capital are accumulated, population grows, and technology improves. For example if $\sigma > 1$, so that other factors can substitute for the fixed factor, then the fixed factor’s share of income will decline over time. Thus, one should be able to learn about the elasticity of substitution, at least in a gross sense, by observing how the income share of the fixed factor changes over time, as A , K , and H accumulate.

Figure 20 shows data for doing such an analysis. The horizontal axis measures output per worker. The data on the vertical axis is an estimate of the income share of non-reproducible factors of production, from Caselli and Feyrer (2007).¹⁴ The Caselli and Feyrer estimates are in turn built on data from the World Bank (2005) on the values of physical capital, crop land, pasture land, and subsoil resources. As the figure shows, the share of natural resources in total income in many developing countries often in excess of 30 percent.

Using data like this, Weil and Wilde (2009) present estimates of the elasticity of substitution between natural resources and other inputs into production. Their estimate is in the neighborhood of two, which is consistent with earlier estimate of Nordhaus and Tobin (1972), looking at US data. However, all of these estimates are fairly imprecise. In

¹⁴Specifically, we use $\alpha_w - \alpha_k$, where the former is the income share of all non-human factors and the latter is the share of reproducible capital. Weil and Wilde (2009) provide an alternative measure of the natural resource share in national income that is very similar.




Figure 20: Fixed-Factor Income Share and Output Per Worker Across Countries

addition, the elasticity of substitution will vary with the resource in question. For example, in developing countries with a large natural resource sector that is effectively detached from the rest of the economy (for example, offshore oil wells), the elasticity of substitution between natural resources and other inputs is infinity.

The production function can be re-written to show how total output compares at two points in time, as the quantities of physical and human capital along with the level of productivity change.

$$\frac{Y_s}{Y_t} = \left[(1 - \phi_t) \left(\frac{A_s K_s^\alpha H_s^{1-\alpha}}{A_t K_t^\alpha H_t^{1-\alpha}} \right)^{\frac{\sigma-1}{\sigma}} + \phi_t \right]^{\frac{\sigma}{\sigma-1}}$$

To do this comparison one does not need to know the quantity of the fixed factor X or the parameter a , but only the income share of the fixed factor at a point in time, the elasticity of substitution, and the growth of the inputs into production, all of which we were already measuring. We use a value of $\alpha = \frac{1}{3}$, which is consistent with our earlier parameterization of giving capital a 0.3 exponent when the land share is 10 percent.

Figure 21 shows how the results of the model are altered when the income share of land is increased from 10 percent to 20 and 30 percent. There are significant differences between the three simulations in GDP per capita following the shock. In the long run, as




Figure 21: The Effect of Land Share on Income Per Capita

expected, a higher land share means that the gains from reduced population growth are larger. Surprisingly however, for the first 50 years of the simulation there is little difference in the income gain in simulations using different land shares. The explanation is that there is an offsetting factor: the larger is the land share, the less valuable are the initial gains in labor supply and human capital that result from slower population growth.

Figure 22 shows how varying the elasticity parameter σ influences our findings by comparing our base case scenario with results obtained under $\sigma = 0.75$, where land is more complementary than in the Cobb-Douglas case, and under $\sigma = 2$, where land is more substitutable. Intuitively, the greater this substitutability, the less severe will be the consequences of increased population pressure on the fixed factor, and thus the smaller will be the income gains from lower fertility. At very long time horizons, this prediction is borne out, but as in the case of varying the land share, differences among the scenarios in the first 50 years of the simulation are minimal.

7 Applying the Model to UN Population Projections

All of the analysis above was constructed using a very stylized demographic scenario, modeled loosely on Nigeria but departing from reality in a number of important ways. In particular, we took current Nigerian fertility and mortality rates as a starting point, assuming in our

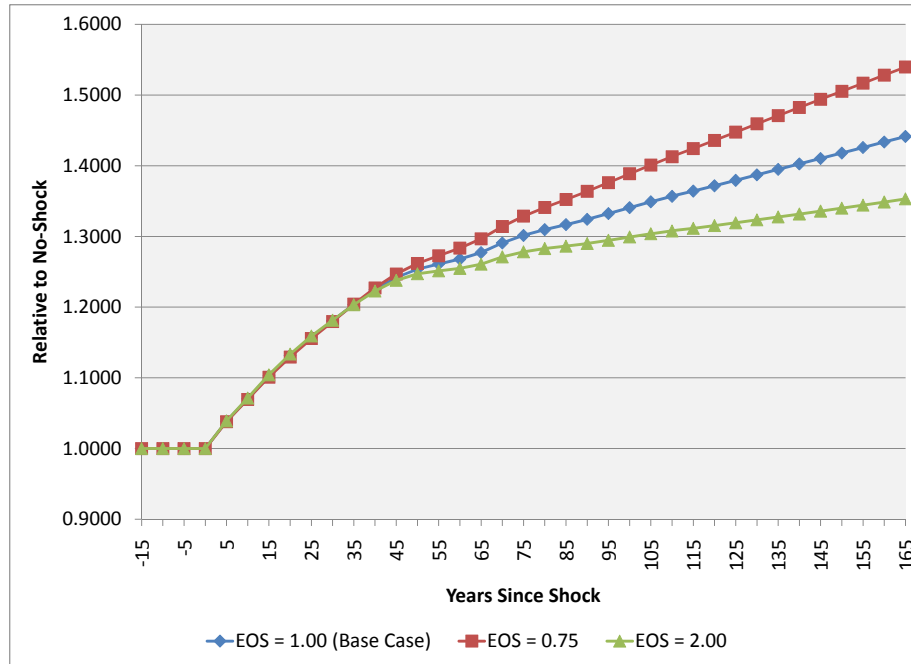


Figure 22: The Effect of Land Substitutability on Income Per Capita

baseline projection that these vital rates would stay constant indefinitely. This constancy of fertility is at variance with all projections that we know of, which say that fertility will decline. We also assumed that in our intervention scenario the TFR dropped instantly by 1.0 and then stayed constant thereafter — a somewhat unrealistic path as well. Finally, as a starting age structure of the population, we used the stable population implied by current vital rates, which is somewhat different than the actual age structure of the Nigerian population.

The benefit of this high degree of stylization was that it allowed us to more precisely analyze the different channels by which a fertility intervention affected the economy, and in particular the timing by which different channels were operative. Further, because we were interested in the effect of an intervention relative to a baseline path, our belief was that the slightly unrealistic nature of the baseline path would not have a first-order effect on our conclusions.

In this section, we examine a more realistic baseline and intervention scenario. Specifically, we take as our baseline the UN medium fertility demographic projection, and as our alternative the UN low fertility demographic projection. Figure 23 shows the paths of the total fertility rate in the two scenarios. The medium variant has the TFR declining rapidly at first, and then with some slowdown, from 5.32 in 2005-10 to 3.27 in 2025-30, and 2.41 in 2045-50. The fertility in the low variant is the same as in the medium variant, except for




Figure 23: Total Fertility Rate in the UN Demographic Projections

a fixed difference in TFR. More specifically, it has the same fertility as the medium variant in 2005-10, and then differs from the medium variant by a TFR of 0.25 in 2010-15, then by 0.4 in 2015-20, and by a fixed TFR of 0.5 thereafter (e.g. the TFR in the low variant for 2045-2050 is 1.91, compared with 2.41 for the medium variant). The two scenarios feature the same future paths of age-specific mortality. In addition, we take as our starting point the current age distribution of the population.

Figure 24 shows the paths of total population in the two scenarios. Total population in the low variant is 4.9 percent lower in the medium variant at a horizon of 2030, and 12.0 percent lower in 2050.

Figure 25, analogous to Figure 8 above, shows the paths of income per capita, income per worker, physical capital per worker, and human capital per capita (and also labor input per adult) in the alternative case (low fertility variant) relative to the base case (medium fertility variant). In the year 2055, which is 40 years after the two fertility paths diverged, income per capita is 12.1% higher in the lower fertility scenario than in the medium fertility scenario. That figure rises to 19.0% by the year 2100, assuming that the gap in fertility remains constant.

Finally, Figure 26 shows a decomposition of the different channels through which income per capita in the alternative case differs from the base case. As in our more stylized analysis above, the dependency effect, and to a lesser extent the labor supply effect, are the

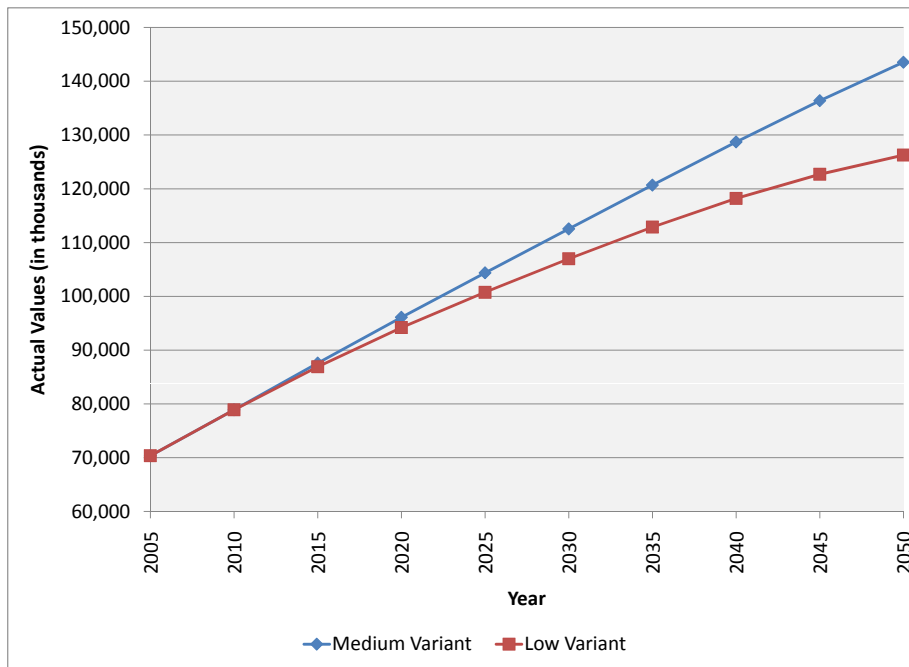


Figure 24: Population Size in the UN Demographic Projections




Figure 25: The Base Case Scenario for the UN Exercise




Figure 26: Decomposition of the Income Gain by Channel for the UN Exercise

dominant channels in the early years of the scenario. Indeed, because the decline in fertility is spread out over several years, rather than happening instantly as in the earlier analysis, the dependency effect remains the dominant channel for a longer period of time. It only falls below half of the total effect on income per capita in the year 2050.

Having begun this section by pointing out the extent to which our stylized example in the rest of the paper differed from the messy reality of actual changes in fertility, it is worth pointing out that our analysis of a more realistic demographic scenario leads to results that do not differ much from the stylized case. The UN scenario is of a gap in TFR of 0.5, phased in over a period of 15 years, relative to a baseline of changing fertility, while our stylized example was of an instant change in the TFR of 1.0 relative to a baseline of constant fertility. Roughly speaking, the results from the UN exercise look similar to those from the stylized exercise, allowing for the fertility reduction being half as big and phased in. For example, in the stable population exercise with an immediate decline in the TFR of 0.5, income per capita is 11.7 percent higher than baseline after 50 years, whereas it is 14.2 percent higher in the UN scenario at the same time horizon (2060).

8 Conclusion

Using a simulation model, we explore the economic effects of an exogenous change in population fertility. The model allows for effects that run through congestion of fixed factors (Malthus effect), capital shallowing (Solow effect), changes in the ratio of workers to dependents (dependency rate effect), change in the average experience of the population (experience effect), change in the saving rate (life cycle saving effect), change in the labor supply of working age adults (labor supply effect), and changes in the average level of schooling (child quality effect). Our simulation model is parameterized using a combination of microeconomic estimates of the effect of fertility changes on labor supply and schooling, data on demographics in developing countries, aggregate measures of the natural resource share in national income, and standard components of quantitative macroeconomic theory. The paper discusses how variations in the parameterization of the economic environment affect our results.

For a base case set of parameters, we find that a immediate decline in the TFR of 1.0 will raise output per capita by approximately 13.2% at a horizon of 20 years, and by 25.4% at a horizon of 50 years. The dependency effect (and to a lesser extent the labor supply effect) are the dominant channels by which reduced fertility affects income per capita in the short run (in the first quarter century after a fertility reduction). At a horizon of 50 years, the four dominant effects are dependency (30.8 percent of the total), Solow (26.7 percent of the total), schooling (22.1 percent of the total) and Malthus (12.3 percent). In the very long run, the Malthus effect is dominant, followed in importance by the Solow effect.

References

- [1] Ashraf, Q. H., Lester, A., and Weil, D. N. (2008). When Does Improving Health Raise GDP? In D. Acemoglu, K. Rogoff, and M. Woodford (Eds.), *NBER Macroeconomics Annual 2008, vol. 23*, Chicago, IL.: The University of Chicago Press.
- [2] Bils, M., and Klenow, P. J. (2000). Does Schooling Cause Growth? *American Economic Review*, 90(5), 1160–1183.
- [3] Bloom, D. E., Canning, D., Fink, G., and Finlay, J. E. (2009). Fertility, Female Labor Force Participation, and the Demographic Dividend. *Journal of Economic Growth*, 14(2), 79–101.

- [4] Bloom, D. E., and Williamson, J. G. (1998). Demographic Transitions and Economic Miracles in Emerging Asia. *World Bank Economic Review*, 12(3), 419–455.
- [5] Caselli, F., and Feyrer, J. (2007). The Marginal Product of Capital. *Quarterly Journal of Economics*, 122(2), 535–568.
- [6] Deaton, A. (1992). Saving and Income Smoothing in Cote d’Ivoire. *Journal of African Economies*, 1(1), 1–24.
- [7] Deaton, A. (1999). Saving and Growth. In K. Schmidt-Hebbel and L. Servén (Eds.), *The Economics of Saving and Growth: Theory, Evidence, and Implications for Policy*, 33–70, New York, NY: Cambridge University Press.
- [8] Gollin, D., and Rogerson, R. (2009). The Greatest of All Improvements: Roads, Agriculture, and Economic Development in Africa. Department of Economics, Williams College. Mimeo.
- [9] Golosov, M., Jones, L., and Tertilt, M. (2007). Efficiency with Endogenous Population Growth. *Econometrica*, 75(4), 1039–1071.
- [10] Hall, R. E., and Jones, C. I. (1999). Why Do Some Countries Produce So Much More Output Per Worker Than Others? *Quarterly Journal of Economics*, 114 (1), 83–116.
- [11] Holmes, J., and Tiefenthaler, J. (1997). Cheaper by the Dozen? The Marginal Time Costs of Children in the Philippines. *Population Research and Policy Review*, 16, 561–578.
- [12] Jones, C. I. (1995). R&D-Based Models of Economic Growth. *Journal of Political Economy*, 103(4), 759–784.
- [13] Joshi, S., and Schultz, T. P. (2007). Family Planning as an Investment in Development: Evaluation of a Program’s Consequences in Matlab, Bangladesh. Economic Growth Center, Yale University. Center Discussion Paper 951.
- [14] Lee, R. D., Mason, A., and Miller, T. (2001). Saving, Wealth and Population. In N. Birdsall, A. C. Kelley, and S. W. Sindig (Eds.), *Population Matters: Demographic Change, Economic Growth, and Poverty in the Developing World*, 137–164, New York, NY: Oxford University Press.

- [15] Nordhaus, W. D., and Tobin, J. (1972). Is Growth Obsolete? In *Economic Growth, Economic Research: Retrospect and Prospect, Fiftieth Anniversary Colloquium, Vol. V, General Series No. 96*, New York, NY: Colombia University Press (for NBER).
- [16] Poterba, J. M. (1994). *International Comparisons in Household Savings*. Chicago, IL.: The University of Chicago Press.
- [17] Razin, A., and Sadka, E. (1995). *Population Economics*. Cambridge, MA: The MIT Press.
- [18] Sindig, S. W. (2009). Population, Poverty, and Economic Development. *Philosophical Transactions of the Royal Society B*, 364(1532), 3023–3030.
- [19] Tiefenthaler, J. (1997). Fertility and Family Time Allocation in the Philippines. *Population and Development Review*, 23(2), 377–397.
- [20] Weil, D. N. (1994). The Saving of the Elderly in Micro and Macro Data. *Quarterly Journal of Economics*, 109(1), 55–81.
- [21] Weil, D. N. (2008). Rapid Population Growth Raises the Stakes for African Governments. In *Africa's Turn?* Boston Review, May 2008.
- [22] Weil, D. N., and Wilde, J. (2009). How Relevant Is Malthus for Economic Development Today? *American Economic Review Papers and Proceedings*, 99(2), 255–260.
- [23] World Bank (2005). *Where is the Wealth of Nations? Measuring Capital for the 21st Century*. Washington, DC: World Bank.
- [24] Young, A. (2005). The Gift of the Dying: The Tragedy of AIDS and the Welfare of Future African Generations. *Quarterly Journal of Economics*, 120 (2), 423–466.