Is Sub-Replacement Fertility Bad for the Economy?

David N. Weil
Brown University and NBER
September 2025

In 2023, the total fertility rate (TFR) in the United States was 1.62 births per woman.[1] That figure was not only the lowest in the country's history, but also well below the replacement rate of 2.08 births per woman that would be required to maintain constant population size in the long run with current mortality rates and absent any net migration. The United States previously experienced a bout of sub-replacement fertility between 1972 and 1989, but that dip was not severe, and came after the Baby Boom, during which the TFR had been well above 3. From 1990 until 2007, the TFR remained quite close to replacement. The current episode of low fertility, which started with a trend break at the time of the Great Recession, is expected by many forecasters to persist into the foreseeable future. This expectation is based largely on the ubiquity and durability of sub-replacement fertility among developed countries worldwide -- a situation in which the United States had long been the most notable outlier (Kearney, Levine, and Pardue, 2022).

Sub-replacement fertility is frequently discussed as a crisis. Commentators worry that the burden of supporting the dependent elderly will crush the economic hopes of working age adults, that smaller populations will be unable to service the debt accumulated by earlier generations, and that innovation will dry up. Some of the concern transcends economics. In 2021 the billionaire Elon Musk warned that "if people don't have more children, civilization is going to crumble, mark my words." The next year, Pope Francis commented that "this denial of fatherhood or motherhood diminishes us, it takes away our humanity. And in this way civilization becomes aged and without humanity, because it loses the richness of fatherhood and motherhood. And our homeland suffers, as it does not have children."[2]

Corresponding to this sense of alarm has been a search for explanations as well as a discussion of the feasibility of deploying various pronatalist policies. While both explaining the roots of current low fertility and exploring the efficacy of pro-natalist policies are interesting projects, I leave them for other research. In this paper, I instead focus narrowly on the question of how sub-replacement fertility affects the economy. For the most part, I will take as my

[1] https://www.cdc.gov/nchs/data/nvsr/nvsr74/nvsr74-1.pdf
[2] Hoffower, Hillary and Marguerite Ward, 2021, "Elon Musk says civilization will crumble if more people don't have more children — and his comments shine light on a heated demographic debate," *Business Insider*, Dec 7. Pope Francis, General Audience, 5 January 2022.

outcome measure the level of consumption per capita, suitably adjusted for the consumption needs of people of different ages.

This issue of focusing on per capita outcomes is worth emphasizing. There is a natural tendency to think of an economy where total output is growing at 3% per year and population is growing at 1% per year as being more "robust" than one in which total output is growing at 1% per year and population is shrinking at 1% per year. But that tendency should be resisted. In both cases, income per capita is growing at the same rate of 2% per year. Income per capita (or consumption per capita, which moves similarly) is, in turn, our most common measure of economic success.

Because sub-replacement fertility is such a recent phenomenon, and because its effects take a very long time to manifest in economic outcomes, my approach is primarily to use off-the-shelf economic and demographic models to study its impact. My conclusion is that much of the worry about sub-replacement fertility is overstated. Sometimes, commentators are simply confused about aggregate vs. per capita measures. There are certainly channels by which low fertility will lower the standard of living, most importantly by raising the flow of transfers to the elderly that are funded via distortionary taxation. However there are also offsetting effects, such as reduced need for investment to supply new workers with capital.

Quantitatively, the net effect of sub-replacement fertility on the standard of living is relatively small. Comparing scenarios for the United States in which the TFR is either two or one (an enormous difference in demographic terms), and focusing on the most easily quantified channels, the steady-state difference in consumption per capita would be about 8.7%. Phased in over many decades, this would come to less than a third of a percent per year. Much of the fiscal cost of future population aging is not due to sub-replacement fertility at all, but rather to variations in fertility that took place before the Great Recession as well as to ongoing declines in mortality. Finally, it is important to note that any attempt to fix the economic problems brought about by low fertility by raising the birth rate will entail a period of higher overall dependency in the decades that it takes the resulting children to become productive adults.

Demographic Effects of Sub-Replacement Fertility

The economic impacts of sub-replacement fertility run through a country's demography, that is, the numbers of people at different ages. It is pretty obvious that a country with a lower fertility rate will, over time, have a lower rate of population growth and a smaller population than it would have had otherwise. With a little more thinking it is also easy to see that lower fertility will produce a population that is older, because it will have a higher ratio of people born long ago (the elderly) to people born recently (the young).

The demography of a country is often visualized through population pyramids. Figure 1 shows such pyramids for the United States for the years 1980, 2025, and projected for 2075.[3] In 1980,

[3] https://www.census.gov/data-tools/demo/idb/#/dashboard

the population was 227.2 million, and was growing at a rate of 1.0% per year; in 2025 the figures are 338.0 million and 0.4% per year. In the projection, which assumes that the TFR will remain at approximately 1.6 children per woman for the next 50 years, the population in 2075 is 369.0 million and is growing at a rate of almost exactly zero.

While lower fertility is reflected in the evolution of these pyramids, its effect is complicated for two reasons. First, changes in fertility affect a country's population demography with an extremely long lag. This lag is due in part to the slow pace at which humans move through their life cycles, but also results from the so-called demographic momentum that is built into the relative sizes of different birth cohorts. For example, in Japan, the cohort of girls aged 0-4 is only 70% as large as the cohort of women aged 30-34, meaning that it would take a significant rise in the birthrate to prevent the annual number of births from falling over the next several decades.

The second complication is that changing fertility is not the only thing that affects these demographic measures. Rates of age-specific mortality have also been changing over time, and will continue to do so in the future. Similarly, the pace of net migration, as well as the ages of these migrants, are reflected in the pyramids. For these reasons, even though low fertility is one of the things that will affect the country's demographic future, it is definitely not the only one. In what follows, I use several different approaches to isolate the effect of low fertility from other drivers of demographic change.

Age Structure, Dependency, and Transfers

The most obvious impact of low fertility is in changing the age structure of the population, that is, the relative numbers of people in different age groups. These changes in the age structure are relevant to economic outcomes because both labor input and consumption have strong life-cycle patterns. In advanced economies, children and old people on average consume more than they produce, with people in the middle of their lives doing the opposite. Thus demographic change as shown in Figure 1 is reflected in the consumption possibilities of the economy.

Before going farther we can demonstrate the effect just described in the simplest possible economic-demographic model. Consider an economy where the population is divided into three groups: young (aged 0-19), working-age (20-64), and old (65+). We will assume that output is produced solely with the labor of working age people (all of whom work), and that consumption in each period is equal to output. Further, we will assume that people of all ages have equal consumption. We will put aside for now any discussion of how claims on consumption are transferred from working age people to those in dependent age groups. Finally, we take the level of income per working age person as exogenous and unrelated to the age structure of the population.

In this simple setting, GDP per capita (which is equal to consumption per capita) is just equal to output per worker multiplied by the share of the population that is working age. Further, the

relationship between the growth rates of GDP per capita, GDP per worker, and the fraction of the population that is working age is given by the equation:

$$g_{GDP\ per\ capita} = g_{GDP\ per\ worker} + g_{workers\ per\ capita}$$

In the United States, the fraction of the population aged 20-64 peaked in the year 2011. In the decades prior to that, the growth rate of working-age people per capita was positive, allowing the growth of GDP per capita to exceed the growth rate of GDP per worker. For example, over the period 1965-85, this tailwind contributed an average of 0.70% per year to the growth of GDP per capita. After 2011, of course, this effect turned negative. Between 2025 and 2050, using the UN median population scenario, the working age fraction of the population is expected to fall from 0.579 to 0.566, implying that $g_{workers\ per\ capita} \approx -0.17\%$. In other words, GDP per capita will grow 0.17 percentage points more slowly than GDP per worker.

There is, of course, a lot that is left out of this simple model. One deficiency of the simple model is its description of all working-age people producing the same amount, with no one in other age groups producing anything else, and everyone having the same level of consumption. To address this point, I use data from the National Transfer Accounts project (Lee *et al.*, 2014).[4] Figure 2 shows the average consumption and labor income by age for the US in 2011. Consumption includes private expenditures as well as government services imputed by age; labor income includes wages as well as imputed labor income of the self-employed and unpaid family workers. Not surprisingly, Figure 2 shows that average consumption among old people is higher than among working age, which in turn is higher than among children. This will tend to raise the economic burden of population aging that results from low fertility.

Using the data in Figure 2 along with data and projections of the number of people at each age by year, I construct a measure of the *adjusted support ratio*. This is the weighted sum of population by age group, where the weights are average labor income by age, divided by a similar sum where the weights are average consumption by age. Both labor income and consumption are normalized relative to the average for working-age adults. Conceptually, the adjusted support ratio functions just as the working-age share of the population in the simple calculations above: The larger its value, the higher will consumption will be relative to the earnings of working-age adults.

In the United States, the adjusted support ratio reached a local maximum of 0.614 in 2006. Over the period 2025-2075, it is projected to fall from 0.568 to 0.501, a pace of 0.25% per year. Just as in the more simplified calculation above, one can think of this as a headwind that will mean that growth of consumption per capita (adjusted for age-specific needs) will grow one quarter of a percent per year more slowly than income per worker.

A second deficiency of the simple model is that, contrary to the model, in real life not all of the gap between consumption and labor income in old age is funded by transfers from the working

[4] https://www.ntaccounts.org/web/nta/show/

age population. The other important channel is people's own accumulated assets. In a world where all old-age consumption was funded out of savings, the fraction of the population made up of non-working elderly would be relevant to individuals' welfare only to the extent that it affected the return to capital. In practice, both transfers and assets are important contributors to the consumption of the elderly in most countries, and heterogeneity in their relative weights may have important implications for how costly low fertility is. In the United States, 37% of the gap between consumption and labor income among the elderly is funded by government transfers, with the rest of the gap being filled by capital income and dissaving. In Japan, the share of the gap funded by government transfers is 58%, while in Italy it is 80% (National Transfer Accounts, 2016).

The calculation just performed is also problematic as a means of capturing the effect of low fertility on economic outcomes because not all of the demographic change, and in particular population aging, that is projected to take place over the next several decades is due to the current episode of sub-replacement fertility. As mentioned above, the demography of a country evolves only slowly in response to changes in fertility, and also is affected by other factors, specifically mortality and immigration, that themselves evolve over time.

One way to demonstrate this point is by looking at alternative projection scenarios that vary fertility while holding constant mortality and immigration. For this purpose, I compare the median and "instant replacement" projections from the United Nations (2024). The former holds the TFR between 1.6 and 1.65 for the remainder of the century. The latter, as the name implies, has the TFR jump immediately to the replacement level. Figure 3 shows the adjusted support ratio under these two fertility scenarios. The key finding here is that the adjusted support ratio is significantly *lower* in the instant replacement scenario than in the scenario where fertility remains below replacement. The gap between the two scenarios is largest 20 years after the two scenarios diverge, reflecting the fact that up to that point the high-fertility scenario features more children but the same numbers of working-age and old people as the low fertility scenario. Eventually, higher fertility leads to a higher adjusted support ratio by lowering the share of the population made up of old people, but that doesn't happen until 2071.

The exercise just conducted is one way to isolate the effect of low fertility, but it is still not ideal because in the projections presented, future changes in fertility are mixed together with ongoing demographic adjustment from past fertility changes as well as changes in mortality and immigration. Another approach to this issue is to take advantage of a tool from demography, the *stable population*. A stable population is the theoretical distribution of people across ages that would result if both age-specific fertility and age-specific mortality rates remained constant, and there was no migration. It can be thought of as somewhat analogous to a steady state in a macroeconomic model. Comparing stable populations with different fertility rates -- assuming the same pattern of mortality -- allows for a *ceteris paribus* comparison between different fertility scenarios, abstracting from the short run dynamics as observed in actual populations.

Table 1 shows data for a set of stable populations constructed using age- and sex-specific mortality rates for the United States for 2019, along with a number of different total fertility

rates.[5] For each TFR, the table shows the associated stable-population growth rates of population, the shares in different age groups, and the adjusted support ratio. One can think of these stable populations as representing a menu of different possible steady-state age structures.

The TFR that maximizes the adjusted support ratio in Table 1 is 2.24, with an associated support ratio of 0.571. This finding suggests that from the perspective of maximizing the long-run support ratio, fertility below the replacement level is indeed a bad thing. There are two caveats, however. First, we will see in the next section that there are economic effects of demography beyond population age structure that will somewhat alter the calculus of what TFR maximizes consumption. Second, besides telling us an optimum, Table 1 is also useful for showing the magnitude of the dependency cost of being away from the optimum.

To pursue this latter question, I will focus on a somewhat stark comparison: a TFR of 2 vs. a TFR of 1. The former, as the table shows, is only slightly below the level of fertility that maximizes adjusted consumption. It is also, of course, just below the replacement rate. The latter is by most accounts an extremely low level -- far below the current TFR in the United States, although still above South Korea, which in 2022 had a TFR of 0.78.

Calculating the magnitude of the effect of going from a TFR of 2 to a TFR of 1 is a simple matter of looking at the ratio of the adjusted support ratios in the two cases. The relevant values are .570 for a TFR of 2 and .490 for a TFR of 1. The ratio is 0.86, implying that consumption in the low-fertility scenario is 14.0% lower than when the TFR is just below replacement.

Before going further it is worth noting that because of the concavity of the adjusted support ratio, deviations from optimal fertility that are smaller than the one just analyzed have much less significant effects. For example, consumption in a scenario with TFR of 1.5 is only 4% lower than when TFR is 2.

The Demographic Dividend and the Demographic Debit

A striking point that comes from comparing Figure 3 and Table 1 is that the peak of the adjusted support ratio in actual data is well above the value associated with any stable population. Specifically, the maximum of the adjusted support ratio was 0.614, achieved in 2016, while the highest value of the adjusted support ratio in a stable population is 0.571, which occurs when the TFR is 2.24.

The explanation for this phenomenon is quite intuitive: when fertility falls, the immediate effect is that there are fewer children. The ratio of old people to those who are working age is not affected for several decades, and only after an additional several decades does the ratio of elderly to working age populations rise by enough to undo the effect of lower youth dependency.

[5] Age-specific fertility rates for these scenarios are constructed by starting with age-specific fertility in the United States in 2019 and scaling proportionally to achieve the desired TFR.

Bloom, Canning, and Sevilla (2003) call this transitory period of low dependency that arises following a decline in fertility the *demographic dividend*. The rise in dependency (and decline in the adjusted support ratio) that is currently taking place in the United States is in large part simply the fading away of this demographic dividend, which was inevitable (Weil, 1997).

This analysis of the dynamics of the support ratio in response to fertility changes has an unfortunate implication for pro-natalist policies. Just as a decline in fertility produces a transitory demographic dividend, a rise in fertility produces a transitory decline in the support ratio. The immediate effect of a rise in fertility is to increase the number of dependent children without impacting the ratio of dependent elderly to working age adults. Only after two decades does the size of the working age population rise above what it would have been absent the increase in fertility, and it takes an additional two decades for the ratio of working age to elderly rise by enough to undo the effect of higher youth dependency. This phenomenon does not yet have a recognized name, but I propose calling it the *demographic debit*.[6]

Figure 4 demonstrates the transitory effects of both the demographic dividend and the demographic debit using a simple simulation model (based on age-specific mortality rates from the United States in 2021 and assuming zero migration). The solid line shows the support ratio in a country that starts with a stable population and a TFR of 2. Beginning in year zero, the TFR jumps to 1. As the figure shows (and consistently with Table 1), the long-run value of the support ratio at the new TFR is lower than at the initial value. However, the figure also shows that there is a long period in which the support ratio is higher than its initial value. The dashed line shows the opposite experiment: a rise in the TFR from 1 to 2. In this case, there is a long run increase in the support ratio, but a temporary decline. More specifically, in the case of a rise in fertility, it takes 41 years from the onset of the fertility increase for the adjusted support ratio to rise above its initial value. This observation suggests that policies designed to raise fertility would not be appealing (at least on economic grounds) to elderly or near elderly voters, and that such policies would similarly be a bad choice for politicians whose planning horizon rarely extends so far into the future.

Fiscal Effects of Transfers

The simple analysis of dependency presented above does not address the issue of how income is transferred from working age to dependent groups. In most developed economies, the family is the most important channel by which resources are transferred to children, while the government dominates transfers to the elderly. Thus population aging due to low fertility is a more serious issue when viewed through the lens of government transfers than when one assesses the issue of dependency more broadly. This section uses the US Social Security system to illustrate the fiscal impact of low fertility.

[6] Bloom and Kotschy (2023) use the term “demographic drag” to describe the reduction in the growth of output per capita relative to output per worker that results when old-age dependency rises in response to low fertility several decades earlier.

To carry out this analysis, I take advantage of tables from actuaries report (Social Security Administration, 2025) that calculate the 75-year actuarial balance of Social Security under a baseline set of assumptions and then under a series of scenarios that vary one or another assumption. Under baseline assumptions, the system's 75-year actuarial balance is -3.82%. This actuarial balance indicates the approximate amount that the payroll tax rate would have to immediately rise above its current level of 12.4% in order to keep the system solvent for the next three-quarters of a century (the exact amount of the required rise is slightly smaller, for complicated reasons).

The baseline scenario assumes that the TFR will gradually adjust from its 2023 level of 1.62 to a value of 1.9 in 2045, after which it will remain constant. Were the TFR instead to remain constant (specifically, 1.6 from 2045 onward), the 75-year actuarial balance would be -4.49%, and were the TFR instead to adjust to a higher level of 2.1 in 2045 and onwards, the actuarial balance would -3.40%.[7]

The high and low TFR scenarios just described, that is, TFR of 2.1 vs. 1.6, certainly do not encompass all of the uncertainty regarding how fertility will evolve in coming decades. Many developed countries have TFRs significantly below 1.6, while we know that the US had fertility well above 2.1 in the not-too-distant past. Still this range can usefully serve as a benchmark to illustrate the extent to which low fertility is contributing to the problems of financing support for the dependent elderly. The answer is that fertility is not the dominant factor: between the low and high fertility scenarios, the difference in the actuarial balances is 1.09% of taxable payrolls, which is only 29% of the current actuarial balance. Thus, low fertility going forward is not the major explanation for Social Security's financial problems. Put differently, reversing low fertility would not on its own fix Social Security, although it could certainly be part of the picture.

One can go further and compare the effect of changing fertility to other adjustments that similarly would affect the actuarial balance. One comparison is with immigration, which is often seen as a natural counterbalance to low fertility. Immigrants tend to arrive young and ready to work -- in this sense, they are far more efficient at countering the burden of dependency than are newborn babies. Adding an additional 440 thousand migrants on top of the 1.25 million annual flow (both legal and illegal) assumed in the baseline of the actuaries would improve the actuarial balance by 0.42 percentage points. Lowering the flow by 420 thousand would worsen the balance by 0.46 percentage points.[8] Combining these estimates for the effects of fertility and migration, and assuming linearity for small changes around the reported projections, one can easily calculate the change in immigration that would be equivalent, in terms of the Social Security actuarial balance, to a given change in fertility. The answer is that the difference between a long-run TFR of 2.1 and a TFR of 1.6 has roughly the same impact as a change in the annual flow of migrants of 1.07 million people.

This analysis of the fiscal effects of reduced fertility can be extended to include the cost of distortions associated with raising revenue. In the Benefit-Cost Ratio approach of Heckman *et*

[7] Tables V.A1 and VI.D1.
[8] Table VI.D3.

al. (2010), if we assume that there is no social benefit to the extra Social Security spending beyond the value to the recipient, then the cost of one dollar of extra Social Security benefits is just $(1 + \phi)$ dollars, where ϕ is marginal deadweight loss, taken to be 0.5. The part of this cost that is not deadweight loss was already incorporated into the calculations above. To give a feel for the deadweight loss: As noted above, the difference in the actuarial balance of Social Security between a future with TFR of 2.1 and a future with TFR of 1.6 is 1.09% of taxable payrolls. Over the 75 year projection horizon used by Social Security, taxable payrolls average 35% of GDP.[9] Thus the additional deadweight loss associated with fertility following a lower trajectory would be $1.09\% \times 0.5 \times 0.35 = 0.19\%$ of GDP per year. This is a relatively small number, because of two points noted earlier: in the United States, only a relatively small fraction of consumption in old age is funded by public transfers, and only a small part of the actuarial imbalance in Social Security is due to sub-replacement fertility. In many developed countries outside the US, where the share of support for the elderly that runs through the government is higher, projected fertility is lower, and tax rates are higher (meaning that the marginal deadweight loss is higher as well), this effect would be of greater magnitude.

Low or Negative Population Growth

Investment and Capital Accumulation

In 2023, gross domestic investment in the US (inclusive of government investment) came to 21.5% of GDP.[10] Investment goes to do three things: replacing depreciated capital, increasing the capital/labor ratio, and, if the labor force is growing, supplying new workers with capital. If the labor force grows more slowly, and even more so if it shrinks, the need for this last category of investment is reduced, which expands the consumption possibilities of the economy.

In general, tracing the effect of slower population growth on investment is complex. In the neoclassical growth model analyzed by Cutler *et al*. (1990), slowing population growth induces the social planner to raise the capital/labor ratio, since the investment cost of maintaining this higher ratio is reduced. In an economy with decentralized saving and investment decisions, the impact of demographic change on investment will depend on, among other things, the extent to which the capital market is open to the world and the model determining individual saving.

Rather than delve into these complexities, I pursue a quantitative shortcut that captures the essence of the impact of slower population growth. First, I will assume that the growth rates of the total population and the labor force are equal, as holds in the case of a stable population. Consider a path for investment expenditures relative to GDP (I/Y) that, for a given growth rate of the population, would yield a particular path of the capital/output ratio (K/Y). If we then ask what would happen to required I/Y in order to maintain that same path of K/Y if there were a different growth rate of population, the answer is simple: I/Y would have to change by an amount equal to the desired K/Y ratio multiplied by the change in population growth.

---

[9] Table VI.G5.

[10] Council of Economic Advisers (2025), Table B19.

To give a concrete example, consider the comparison of a stable population with a TFR of 1 vs. a stable population with a TFR of 2. From Table 1, the former has a growth rate of population that is 2.3% per year lower than the latter. The K/Y ratio in the United States in 2019 was approximately 3.3.[11] Thus the version of the United States with a TFR of 1 would have investment required to maintain capital per worker that would be $3.3 \times 2.3\% = 7.6\%$ GDP lower than the version of the United States with TFR of 2. This is a large effect -- although it should be noted that the comparison of a TFR of 1 vs. a TFR of 2 is also somewhat extreme.

In addition to the change in required investment, demographic change due to lower fertility may have other macroeconomic impacts that are of interest. One that has received attention in the literature is a reduction in the real interest rate, via both decreased investment demand and the shifting of population into ages with higher average wealth holdings (Eggertsson, Mehrotra, and Robbins, 2019). Tracing the welfare effect of such a reduction is difficult, but it is certainly the case that lower interest rates do not qualify as being bad for the economy in a *prima facie* sense.

Servicing Government Debt

My analysis of debt sustainability parallels that of capital accumulation, although the direction of the effect is reversed. In an economy with a growing population, government debt is diluted over time. This means that fewer resources have to be devoted to maintaining a given debt/GDP ratio; more concretely, an economy in this situation is required to run a smaller primary surplus (or may run a larger primary deficit) in order to achieve a given path of the debt/GDP ratio. When population growth becomes negative, this process goes into reverse: the primary surplus has to be increased to maintain a given debt/GDP ratio.

Of course in the current context of US fiscal policy, it is a bit fanciful to discuss stabilization of the debt/GDP ratio, but the same logic regarding the role of population growth holds even when considering paths of debt/GDP that are not flat. Consider a path of the primary budget surplus relative to GDP that, for a given growth rate of the population, would yield a particular path of the debt/GDP ratio. If we then ask what would happen to the primary surplus that would be required to achieve this path with some different population growth rate, the answer is that the primary surplus would have to grow by the debt/GDP ratio multiplied by the *negative* of the change in the population growth rate.

Applying the same comparison of a version of the United States with a TFR of 1 vs a version of the country with a TFR of 2 that was discussed above, we would just multiply the difference in population growth rates (2.3%) by the debt/GDP ratio. In the US, that ratio (specifically using debt held by the public) is almost exactly 1.0, and so the math is particularly easy: in the low-TFR scenario, the primary surplus would have to be 2.3% of GDP larger than in the high fertility scenario in order to achieve the same path of the debt/GDP ratio.

[11] Penn World Tables version 10.01, Feenstra, Inklaar, and Timmer (2015).

Combining Age Structure and Population Growth Rate Effects

So far, I have discussed three channels through which low fertility affects the standard of living: age structure effects on transfers, and the effect of population growth on both required capital investment and servicing government debt. These are definitely not the only economic impacts from low fertility, but they are the three that are most frequently discussed, and, as shown above, they are all relatively easy to quantify. Here I combine the earlier results to produce a more comprehensive assessment of the economic impact of low fertility. For simplicity, I will continue to examine the case of stable populations; more concretely, I will continue to compare a stable population with a TFR of 1 to a stable population with a TFR of 2.

The combined effect of these three channels is constructed by simply adding their individual effects. Comparing TFR of 1 to TFR of 2, the dependency effect reduces consumption by 14.0%; the investment effect raises consumption by 7.6%; and the debt sustainability effect reduces consumption by 2.3% (with these latter two effects being scaled to the current K/Y and debt/GDP ratios of the United States). The combined effect is to reduce consumption by 8.7%.

This strikes me as a relatively small impact. First, note that the fertility rates that I am comparing differ by an enormous amount. They were chosen for illustrative purposes rather than realism. For example, in the probabilistic projections of the United Nations (2024), a TFR of 1 in the United States in the year 2100 is at the very bottom of the 95% probability interval, while a TFR of 2 is at the top of the 80% probability interval.

Even taking this extreme comparison, a reduction in consumption by 8.7% is not all that large. It is significantly smaller than the difference between consumption per capita in Canada and the United States. Another way to look at it is that it is slightly smaller than five-year's growth of personal consumption per capita in the US.[12] While such a reduction in consumption would be unfortunate, it does not seem to rise to the level of being a crisis, as it is frequently depicted. Further, for the reasons discussed above, the reduction in consumption resulting from a decline in fertility of this magnitude would be phased in over many decades, and during this phase-in period consumption would actually rise due to the demographic dividend.

A Smaller Population Size

The growth rate of the population impacts the size of the population, so it might seem like the growth rate and size of the population are the same issue. But this is not completely true. For example, if we imagine that the growth rate of the population will someday reach zero and stay there (for some unspecified reason), there could be different possible sizes of the resulting constant population depending on what growth rates were observed in the period leading up to that stabilization. The population pessimists who coined the slogan "zero population growth" actually thought that the population of the earth at the time that they wrote was already too high.

[12] Real personal consumption expenditures per capita from the FRED database over the period Q3 1990 to Q3 2019 grew at an annual rate of 1.76%.

What they wanted was negative population growth for a while, to get the level down, followed by zero after that. By contrast, Spears and Geruso (2025) argue that while declining world population is now almost inevitable, welfare would be improved by stabilizing world population quickly, rather than letting it decline for several generations before such a stabilization.

Natural Resources and the Environment

The interaction of population size with natural resources is the issue on which the field of population economics cut its teeth, starting with Malthus (1798). In a constant-returns production function with land or another fixed factor as an input, a smaller population size leads directly to a higher standard of living. It was this issue of the ratio of population to resources that energized population pessimists during the second half of the 20th century (e.g. Ehrlich, 1968; Hardin, 1995). While these scholars focused particularly on developing countries, they were also quite clear in their pronouncements that richer countries like the United States would have a higher standard of living if they had fewer people.

There were a number of reasons why the dire predictions of the population pessimists did not pan out. These included the possibilities of substitution among resource inputs in production, technological progress that allowed for new productive techniques, and the ability of highly populated but resource-poor countries such as Japan to use trade as a substitute for locally available inputs. Further, an analysis of how countries' populations compare to their available resources suggests that the United States is certainly far from "full," as President Donald Trump proclaimed in 2019. The US ratio of people to land area, with the latter adjusted to reflect characteristics that support economic activity, is about 40% of the world average (Henderson, Storeygard, and Weil, 2020).

To the extent that the issue of population size relative to available resources is still relevant today, it is in two contexts. The first is among countries that are still largely autarkic. This includes most of the poorest countries in the world, where costs of or other barriers to trade are high, most of the population is dependent on small-scale farming or herding, and land degradation from over-use is a serious problem. While these countries would benefit economically from a smaller population, they almost all still have fertility rates that are at or above the replacement level. By contrast, almost all countries with fertility well below replacement level are open to trade.

The population/resource ratio is also relevant for the world as a whole. While population pessimists of the 20th century worried about the ratio of people to farmland and subsoil resources, most current analyses of the interaction of population and the natural environment focus on the emissions of greenhouse gasses (GHGs) as drivers of climate change. It is clear that holding other things constant smaller populations will emit fewer GHGs. Thus from a climate change perspective, lower fertility in highly developed countries, which have the highest levels of GHG emissions per capita, is a good thing. Whether changes in fertility from the present moment onward would have a significant impact on cumulative GHG emissions depends on the path of decarbonization going forward. As pointed out by Budolfson *et al*.

(2025), many projections have the property that GHG emissions per unit of output will fall very quickly over the next 75 years in order to avoid the worst effects of climate change. In such a scenario, conceivable variations in fertility in the wealthiest countries going forward will have only a minor additional impact on cumulative emissions as of the end of the century. Of course, if rapid decarbonization in per-capita terms did not happen, then a declining population would mitigate the resulting climate disaster.

Technological Progress

In just about every model that economists construct, the only driver of economic growth over the very long run is technological progress. New technologies have the property that they are non-rival, so that benefits from an invention are scaled by the population that has access to it. In turn, creation of new technologies requires resources in the form of labor, human capital, and physical capital. An economy with more people will *ceteris paribus* have more people working on creating new technologies, and thus faster economic growth. Jones (2022) argues that for this reason, zero or negative population growth will have a negative impact on the long-run standard of living.

While this is a reasonable argument, there are several caveats that suggest that it may not be of great quantitative importance, at least for horizons of less than a century. First, since technological progress is shared among the set of countries that are at or near the technology frontier, the speed of technological progress experienced by any one country (unless it is very large) is invariant to that country's own rate of population growth. (In 2021, the United States accounted for 19% of the world's science and technology researchers.[13])

Second, for the next several decades, the stock of researchers in the world will be growing despite contraction in the labor forces of many developed countries. This is both because of new countries joining the group that is at the technological cutting edge, and because of rising human capital of the labor force, and thus growth in the number of potential researchers, in populous countries such as China and India. We can put some quantitative flesh on this point by looking not at R&D workers directly (for whom there are no projections), but at the number of people with post-secondary education. These are the type of people who are going to come up with almost all of the innovations that are relevant to the standard of living in the United States. According to the Wittgenstein Centre for Demography and Global Human Capital (KC *et al.*, 2024), the combined population of the regions where the vast bulk of current R&D takes place (Europe, North America, East Asia, Australia, and New Zealand) is projected to decline by 22% between 2050 and 2100, while the population of the world as a whole is projected to increase by only 3.1% over this same period. However, over this same period, the number of people with post-secondary education worldwide is projected to increase by 78%.[14]

---

[13] Full time equivalents. OECD Main Science and Technology Indicators database https://www.oecd.org/en/data/datasets/main-science-and-technology-indicators.html .
[14] Medium (SSP2) scenario. https://dataexplorer.wittgensteincentre.org/wcde-v3/

Third, a driving mechanism in the Jones model is that the only lever for influencing the speed of technological progress is the size of the population. The ratio of R&D workers to the population is generally taken as fixed. In practice, however, producing more people is a relatively inefficient way to produce more scientists. It is far cheaper and faster to simply improve access to high quality STEM education; and similarly far simpler to use government resources for direct expenditures on R&D rather than raising the size of the population. While the fraction of the labor force doing R&D obviously can't rise forever because it would eventually be greater than 100%, the current fraction is small enough that such a limit would not bind for a long time.

Finally, even taking the Jones model fully seriously, the transition to zero or negative population growth affects the speed of technological progress with an extremely long lag. It is true that in this model technological progress will asymptote to zero if the size of population stabilizes or declines, while if population continues to grow, technology will grow without limit. Eventually, the difference in living standards between these different scenarios becomes infinite. However, given that these differences play out only at a horizon of a century or more, it is hard to reconcile this perspective with the idea that currently low fertility represents a crisis.

Conclusion

It is only over the last half-century that we started to observe peaceful and prosperous countries with birth rates below the replacement level, and yet today more than half of the world's population lives in countries where this is the case. The United States, where fertility remained near replacement until the Great Recession, was long an outlier among rich countries.

In this paper, I assess the effect of sub-replacement fertility on economic wellbeing. My finding is that this impact, while likely to be negative, is relatively small. To give a feel for the magnitude, consider the comparison of a country where the total fertility rate is two children per woman with an otherwise similar country where the TFR is one. This is a dramatic difference in fertility, given that the former country would have a population growth rate very close to zero (-0.1% per year), while in the latter the population growth rate would be -2.4% per year. Combining the three most easily quantifiable channels by which fertility affects economic outcomes -- the change in the sizes of dependent age groups, investment required to maintain the capital stock, and taxation required to stabilize government debt -- I find that the country with low fertility would have consumption per capita 8.7% lower than the country where fertility was high. Further, this result only obtains after a long adjustment period. When fertility declines there is a positive transitory effect on consumption that lasts for several decades.

Throughout the article my focus has been on the impact of low fertility on consumption per capita. This is, after all, the most frequently used measure of economic success. When people say that something is "bad for the economy" they are usually referring to consumption per capita. Nonetheless, in closing it is certainly worth mentioning other social welfare criteria that might be brought to bear, and how their consideration might alter one's view about whether sub-replacement fertility was a problem.

One reason to favor a bigger population is that in such a setting there are more potential friends or mates to choose among, more people who might share one's interest, and so on. Mill (1848) raised this issue, but dismissed it, writing that "The density of population necessary to enable mankind to obtain, in the greatest degree, all the advantages both of co-operation and of social intercourse, has, in all the most populous countries, been attained" at a time when the population of the United Kingdom was 40% of its current level. A related critique of the per-capita approach, goes back to the utilitarian imperative to achieve "the greatest good for the greatest number." The philosophical approach known as totalist utilitarianism aims to maximize the sum of total utility, rather than the per-capita level. Most implementations of this approach, taking into account the curvature of the utility function and the productivity of extra people, as well as the utility that people get from being alive rather than just consuming, point to an optimal population size far higher than the current one, and thus suggest that sub-replacement fertility leads to large welfare losses by moving population size in the wrong direction (Adhami *et al*., 2024).

Yet another philosophical approach, known as long-termism, expands the totalist utilitarian framework by considering the welfare of potential people who could live in the future, without any time discounting. This approach rapidly leads to the conclusion that the optimal fertility rate is the one that maximizes the probability of the survival of the human species, given that if the species does survive, the vast majority of human lives will take place far in the future (MacAskill, 2022). Whether that survival maximizing fertility rate is higher or lower than the currently observed level is subject to debate. Many environmentalists argue that in the present setting, low fertility is most conducive to humanity's survival.

References

Adhami, Mohamad, Mark Bils, Charles I. Jones, and Peter J. Klenow, "Population and Welfare: Measuring Growth when Life is Worth Living," NBER Working Paper 31999, August 2024.

Bloom, David, David Canning, and Jaypee Sevilla. 2003. *The demographic dividend: A new perspective on the economic consequences of population change*, Rand Corporation.

Bloom, David, and Rainer Kotschy, "Population Aging and Economic Growth: From Demographic Dividend to Demographic Drag?" NBER working paper 31585, August 2023.

Budolfson, Mark, Michael Geruso, Kevin Kuruc, Dean Spears, and Sangita Vyas "Is Less Really More? Comparing the Climate and Productivity Impacts of a Shrinking Population" working paper, 2025.

Council of Economic Advisers, *Economic Report of the President*, Washington, DC: Government Publishing Office, 2025.

Cutler, David M, J.M. Poterba, L.M. Sheiner, L.H. Summers, “Aging Society: Opportunity Challenge?” *Brookings Papers on Economic Activity*, 1 (1990), pp. 1-73

Eggertsson, Gauti B., Neil R. Mehrotra, and Jacob A. Robbins. 2019. "A Model of Secular Stagnation: Theory and Quantitative Evaluation." *American Economic Journal: Macroeconomics* 11 (1): 1–48.

Ehrlich, Paul, *The Population Bomb*, New York: Ballantine Books, 1968.

Feenstra, Robert C., Robert Inklaar and Marcel P. Timmer (2015), "The Next Generation of the Penn World Table" American Economic Review, 105(10), 3150-3182, available for download at www.ggdc.net/pwt

Hardin, Garrett, *Living Within Limits: Ecology, Economics, and Population Taboos*, Oxford: Oxford University Press, 1995.

Heckman, James J., Seong Hyeok Moon, Rodrigo Pinto, Peter A. Savelyev, and Adam Yavitz, “The Rate of Return to the High Scope Perry Preschool Program,” *Journal of Public Economics*, 94 (2010), 114–128.

Henderson, J. Vernon, Adam Storeygard, and David N. Weil, “Land Quality,” NBER Working Paper 28070, 2020.

Jones, Charles I., 2022, “The End of Economic Growth? Unintended Consequences of a Declining Population,” *American Economic Review* 112(11): 3489–3527.

KC, Samir, *et al.*, “Updating the Shared Socioeconomic Pathways (SSPs) Global Population and Human Capital Projections,” Working Paper, International Institute for Applied Systems Analysis, 2024. Link.

Kearney, Melissa S., Phillip B. Levine, and Luke Pardue. 2022. "The Puzzle of Falling US Birth Rates since the Great Recession." *Journal of Economic Perspectives* 36 (1): 151–76.
MacAskill, William, *What We Owe the Future,* New York: Basic Books, 2022.

Lee, Ronald *et al*., “Is low fertility really a problem? Population aging, dependency, and consumption,” *Science* 346, 229-234 (2014).

Malthus, T. R. (1798). *An Essay on the Principle of Population*. London: Oxford University Press.

Mill, John Stuart, *Principles of Political Economy*, London: John W. Parker, 1848.

National Transfer Accounts, NTA Data Sheet, 2016.
https://ntaccounts.org/doc/repository/NTA%20Data%20Sheet%202016.pdf

Social Security Administration, *The 2025 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds,* Washington, DC, 2025.

Spears, Dean, and Michael Geruso, *After the Spike Population, Progress, and the Case for People*, New York: Simon and Schuster, 2025.

United Nations, Department of Economic and Social Affairs, World Population Prospects 2024.

Weil, David N. ‘‘The Economics of Population Aging,’’ in Mark R. Rosenzweig and Oded Stark, eds., *Handbook of population and family economics*. Amsterdam: NorthHolland, 1997, pp. 967–1014.

Weil, David N., “Replacement Fertility is Neither Natural nor Optimal nor Likely,” working paper, 2024.

Table 1: Stable Populations for the United States

TFR	n (%)	Share 0−19 (%)	Share 20−64 (%)	Share 65+ (%)	Adj. Support Ratio
1.00	−2.4	10.0	47.9	42.2	0.490
1.25	−1.7	13.6	51.0	35.4	0.526
1.50	−1.1	17.3	52.8	30.0	0.549
1.75	−0.6	20.8	53.6	25.6	0.563
2.00	−0.1	24.1	53.9	22.1	0.570
2.25	0.3	27.2	53.7	19.2	0.571
2.50	0.6	30.1	53.2	16.8	0.570
2.75	0.9	32.8	52.5	14.8	0.565
3.00	1.2	35.3	51.6	13.1	0.559

Stable populations are constructed using US age- and sex-specific mortality from 2019 and proportionally scaling age-specific fertility from 2019.

Figure 1: Population Pyramids for the United States

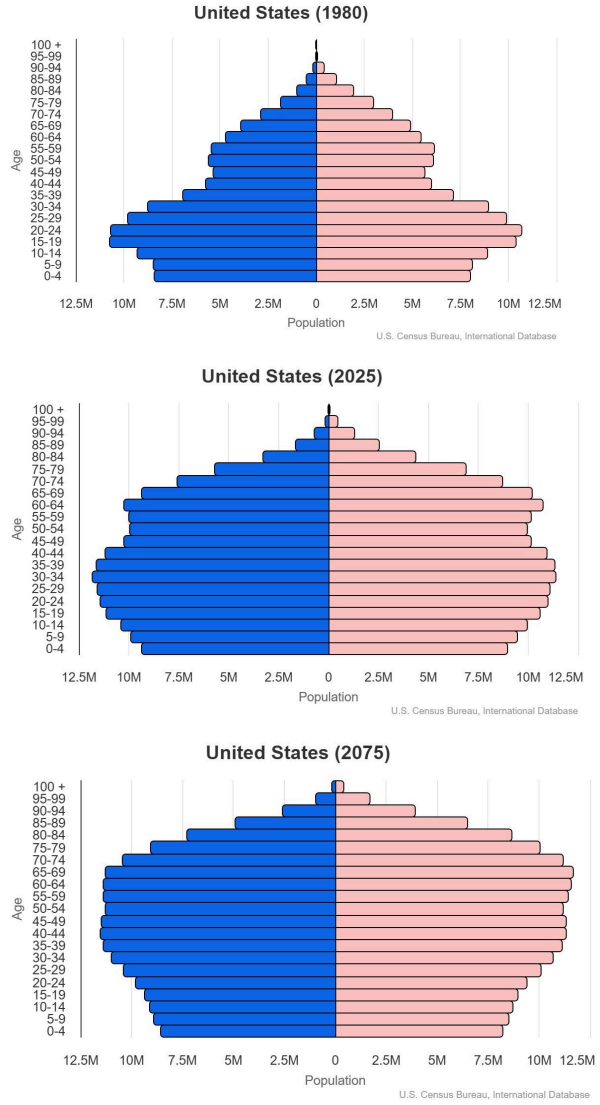

Figure 2: Age Profiles of Consumption and Labor Income

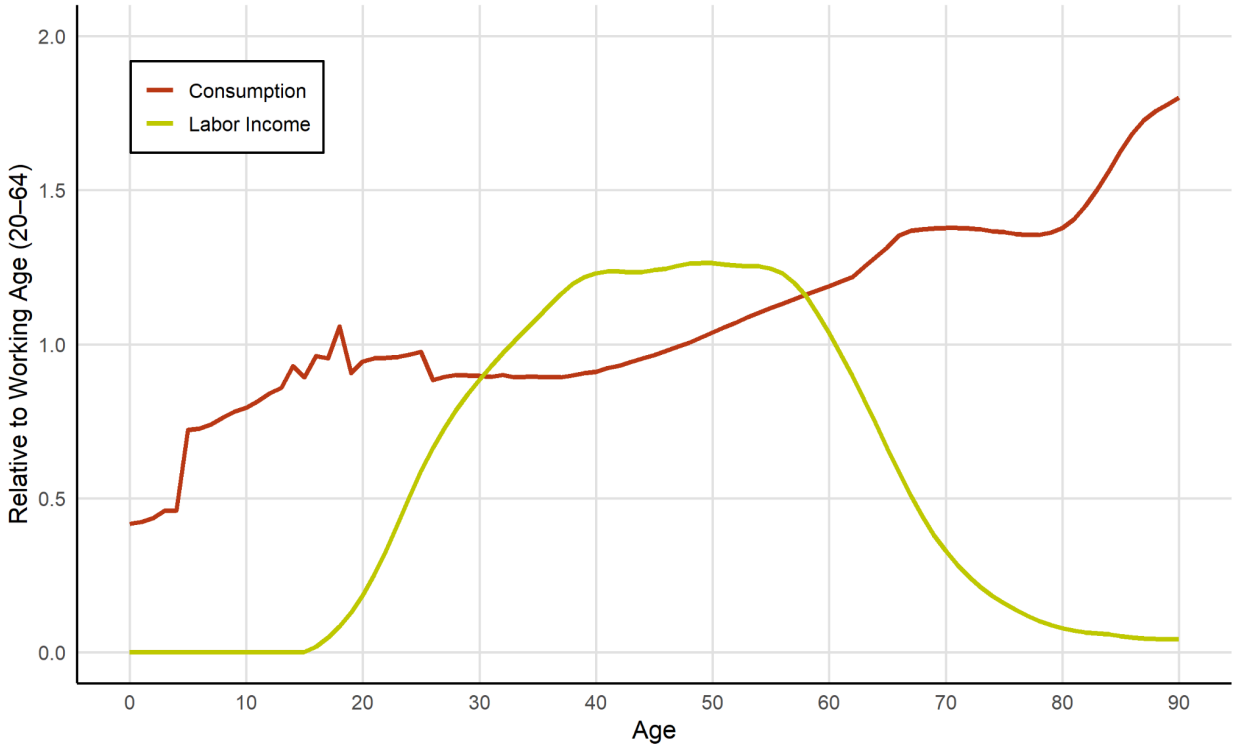

Source: National Transfer Accounts (2011). Consumption includes private expenditures as well as government services imputed by age; labor income includes wages as well as imputed labor income of the self-employed and unpaid family workers.

Figure 3: Adjusted Support Ratio by UN Fertility Scenario

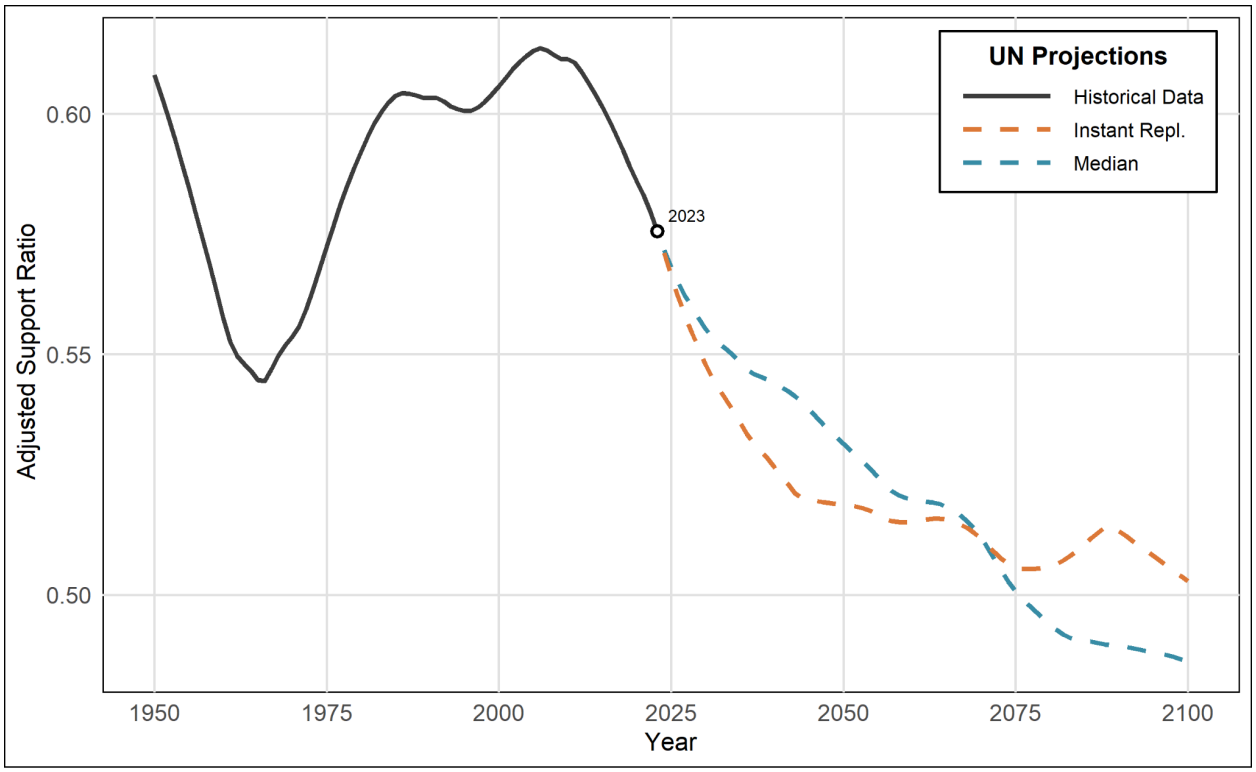

Figure 4: The Demographic Dividend and the Demographic Debit