Understanding the New Normal: The Role of Demographics*

Etienne Gagnon   Benjamin K. Johannsen   David López-Salido

February 26, 2018

Abstract

We calibrate an overlapping generation model with a rich demographic structure to observed and projected changes in U.S. population, family composition, life expectancy, and labor market activity. The model accounts for a $1\frac{1}{4}$–percentage-point decline in both GDP growth and real interest rates since 1980—essentially all of the permanent declines in those variables according to some estimates. Model-implied declines were especially pronounced over the past decade or so because of demographic factors associated with the post-war baby boom. Our model predicts GDP growth and interest rates will remain low by historical standards, consistent with a "new normal."

JEL Codes: E17, E21, J11.

Keywords: new normal, equilibrium interest rate, GDP growth, demographics.

*The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Board or any other person associated with the Federal Reserve System. We are grateful to Carter Bryson and Kathryn Holston for excellent research assistance, and to seminar participants at the Bank of Canada, the Federal Reserve Board, the 7th Joint Bank of Canada and European Central Bank Conference, the Centro de Estudios Monetarios y Financieros (CEMFI), the Einaudi Institute for Economics and Finance, the Federal Reserve Banks of Chicago, Kansas City, and Minneapolis, and the "Declining Natural Rate of Interest" session organized by the Washington Center for Equitable Growth at the 2018 annual meetings of the Allied Social Science Association. Comments can be directed to `benjamin.k.johannsen@frb.gov`.

1 Introduction

To support the economic recovery from the Great Recession, the Federal Reserve held the federal funds rate near zero for over seven years and acquired large holdings of longer-term securities. Despite these extraordinary measures, real GDP has grown at only a modest pace during the recovery. Meanwhile, long-term interest rates have declined whereas measures of longer-term inflation expectations have been relatively stable, resulting in lower real long-term interest rates, as shown in Figure 1. Some observers, such as Rogoff (2015), trace these developments to persistent, but ultimately transitory, debt deleveraging and borrowing headwinds in the wake of the global financial crisis. Others, like Summers (2014) and Eggertsson et al. (2017), see these developments as symptomatic of "secular stagnation," by which a confluence of changes in the structure or functioning of the economy is leading to weak GDP growth and low interest rates of an enduring rather than cyclical nature.[1] Commentators and policymakers have described this combination of low growth and low interest rates as a "new normal" for the U.S. economy.[2] This paper seeks to understand how much of the new normal can be explained by demographic factors in the United States.

The United States, like other advanced economies, is undergoing a dramatic demographic transition related to the unfolding of the post-war baby boom.[3] As a consequence, the growth rate of the labor force has declined and should remain low for the foreseeable future.[4] In this paper, we investigate the extent to which demographic changes, especially

[1]The expression "secular stagnation" was popularized during the Great Depression by Hansen (1939), who wrote: "[this] is the essence of secular stagnation—sick recoveries which die in their infancy and depressions which feed on themselves and leave a hard and seemingly immovable core of unemployment." One potential contributor to secular stagnation in the 21st century is a step down in technology growth, with authors such as Fernald (2015) and Gordon (2016) arguing that the high productivity gains enjoyed before the early 1970s and during the information technology boom of the mid-1990s to early 2000s are unlikely to be repeated.

[2]For an early use of the expression "new normal," see El-Erian (2010).

[3]The baby-boom generation consists of the cohorts born after World War II through the mid-1960s; see Hogan et al. (2008) for a statistical definition.

[4]Colby and Ortma (2015) describe the salient features of the U.S. demographic transition. Aaronson et al. (2014) discuss the effects of the demographic transition on recent and future movements in the aggregate labor supply.

those related to the baby boom, can explain the currently low levels of real interest rates and GDP growth. To this end, we build an overlapping generation (OG) model that is consistent with observed and projected changes in fertility, labor supply, life expectancy, family composition, and international migration. Our study focuses on the United States and traces the extent to which demographic changes can account for both the size and timing of historical changes in interest rates and real GDP growth rates. In an earlier contribution, Krueger and Ludwig (2007) use an OG model to explore the consequences of demographic transition in advanced economies for capital accumulation and rates of return from 2005 to 2080. We also investigate which historical demographic changes have been most important for recent and predicted future changes in interest rates and GDP growth.

We find that demographic factors alone account for a $1\frac{1}{4}$–percentage-point decline in the equilibrium real interest rate in the model since 1980—much, if not all, of the permanent decline in real interest rates over that period according to some recent time-series estimates, such as Johannsen and Mertens (2016b) and Holston et al. (2016). The model is also consistent with demographics having lowered real GDP growth $1\frac{1}{4}$ percentage points since 1980, primarily through lower growth in the labor supply; this decline is in line with changes in estimates of the trend of GDP growth over that period. Interestingly, the model also implies that these declines have been most pronounced since the early 2000s, so that downward pressures on interest rates and GDP growth due to demographics could be easily misinterpreted as persistent influences of the global financial crisis. When we enrich our model with total factor productivity (TFP) growth consistent with the observed historical trend, the predicted declines in GDP growth and the equilibrium real rate between 1980 and the present are essentially unchanged but the pace at which the economy has been transitioning to the new normal over the past decade is more dramatic.

In the model, the dynamics of interest rates and GDP growth are most directly connected to the consequences of the post-war baby boom. As the baby-boom generation reached working age in the 1960s, the aggregate labor supply, GDP growth, and interest rates all increased. The abundance of labor was accentuated by the fact that the baby boomers had far fewer children than their parents did, leading the United States to reap a "demographic dividend" by which the number of workers relative to the number of dependents climbed to historically high levels by the turn of the 21st century. The demographic situation stimulated aggregate capital formation as members of an unusually large cohort with few dependents simultaneously supplied labor and aimed to save ahead of retirement. The low fertility rates of the baby-boom generation also supported the accumulation of capital by facilitating greater labor-force participation of women and by freeing resources that families would have otherwise allocated to consumption by children. Furthermore, steady gains in health and life expectancy have increased the amount of time households expect to spend in retirement, in turn boosting their desire to save. The baby-boom generation has since begun to retire and the growth rates of the aggregate labor supply and aggregate output have accordingly slowed. From our model's perspective, these factors have led to a current abundance of capital relative to labor, depressing the return on capital and also causing aggregate investment rates to decline, a phenomenon that we see as consistent with puzzlingly low rates of capital investment in the recovery from the global financial crisis. Going forward, the model predicts that the capital-labor ratio will remain elevated despite low rates of aggregate investment in capital because the growth rate of the labor supply will also be low, so that both real interest rates and GDP growth will linger near their current low levels.

We use the model to show that the most consequential demographic changes for interest rates and GDP growth today took place before the 1980s. Had demographic variables held constant at their 1960 levels, both real interest rates and GDP growth today would be similar to their levels in 1980. In sharp contrast, had demographic variables held constant at their 1980 levels, the declines in the equilibrium real rate and real GDP growth would have been largely as predicted under the baseline calibration of the demographic variables. We also use these simulations to parse out the relative contributions to those declines of each of the demographic factors. In the case of real GDP growth, the sharp drop in fertility rates in the 1960s and 1970s, which ultimately led to a slowing in the growth rate of the labor supply, is the predominant contributor to the decline since 1980. In the case of real interest rates, we find that falling fertility rates and increased employment rates in the 1960s and 1970s each have contributed $\frac{1}{2}$ percentage point to the decline in the equilibrium real rate between 1980 and the present, with longer life expectancies accounting for the residual.

Our results also complement those of Ikeda and Saito (2014) for Japan and Carvalho et al. (2016) for the OECD. These studies focus on the distinction between working-age adults and retired adults, and show that a rise in the proportion of retired adults lowers real interest rates. Our model is richer, taking into account differences in demographic characteristics over the life time and across birth cohorts. We argue that the consideration of a rich model matters for obtaining reliable empirical estimates of the size and timing of the macroeconomic effects of demographics. In particular, the model of Carvalho et al. (2016) predicts a sharp decline in real interest rates in the 1990s driven by increased life expectancy. We show, within the context of our larger model, that life expectancy changes over the past several decades have had relatively little effect on interest rates and that the coming retirement of the baby boom generation had its most dramatic effects on interest rates only after the turn of the century. Our model is also related to contemporaneous work by Eggertsson et al. (2017), who find dramatic declines in interest rates from a variety of sources. Our focus on demographics and our rich demographic data allow us to quantify the effects of demographic transition in the United States.

The paper is organized as follows. Section 2 presents our model economy. Section 3 details the demographic data and other parameters of our calibration, as well as our approach to computing the model's dynamic general equilibrium. Section 4 shows the main results under our baseline calibration and gauges the relative importance of demographic factors through the consideration of a number of alternative scenarios. Section 5 offers a series of robustness checks and observations pertaining to our choice of a technology growth process, the elasticity of substitution in production, the role of dependents in household preferences, and open-economy considerations. Finally, section 6 concludes.

2 Model economy

Our OG model borrows from the vast literature exploring the long-run fiscal and macroeconomic implications of demographic transitions and social security programs. See, among many related contributions, Yaari (1965), Fehr et al. (2003), or, more recently, Reichling and Smetters (2015). The economy is comprised of households and representative firms. Households are headed by adults who are representative of their birth cohort in terms of life-cycle characteristics. Each period, these adults inelastically supply their labor endowment, which varies over the life cycle and across birth cohorts. Additionally, these adults decide how much they and their dependent children (that is, children who have yet to reach adulthood) consume, and how much capital to save for retirement. Representative firms operate a constant returns to scale production technology that uses capital and labor inputs. Given our focus on low-frequency dynamics, we abstract from business cycle shocks and aggregate uncertainty. Instead, we assume that adults and firms have perfect foresight with respect to aggregate and demographic variables, and that they face no idiosyncratic risk besides the uninsurable risk of death. The timing of events is as follows: At the beginning of each period, children are born with no capital holdings according to fertility rates that depend on the age and birth cohort of the adults. After children are born, some household members die, with probabilities conditional on age and birth cohort, and orphans are redistributed among surviving adults belonging to the same cohort as the deceased parent. The capital held by dead adults is then uniformly redistributed to surviving adults. Finally, surviving adults make their consumption/saving decisions and firms produce.

2.1 Households

In our benchmark model, we assume that households derive utility from adult consumption only.[5] Consider a cohort-representative adult aged a in period t, with a utility function given by

$$\sum_{s=0}^{A-a} \beta^s \Gamma_{a,t,t+s} \frac{(C_{a+s,t+s})^{1-\nu}}{1-\nu}, \tag{1}$$

where $C_{a,t}$ is the adults' consumption. Our specification implies that adults derive utility only in periods in which they are alive and that they have no intentional bequest motive. The probability that an adult aged a in period t survives through the end of the next s periods is given by the cohort-specific function $\Gamma_{a,t,t+s}$. Adults take fertility and mortality rates as exogenous, so the sequence of values for $\Gamma_{a,t,t+s}$ are taken as given. We assume that no adult survives past some finite length of life, A. Households face the following sequence of budget constraints

$$C_{a+s,t+s} + K_{a+s+1,t+s+1} = (K_{a+s,t+s} + \Xi_{t+s}) \left(R^K_{t+s} + 1 - \delta\right) + e_{a+s,t+s} W_{t+s}. \tag{2}$$

Here, $K_{a,t}$ is the age a adult's capital holdings at the beginning of the period, $e_{a,t}$ is the age a adult's labor endowment, R^K_t is the real rental rate of capital, and W_t is the real wage rate. We assume that the capital held by adults who die, Ξ_t, at the beginning of the period is redistributed among surviving adults. In short, the budget constraint states that consumption and net investment (in capital units) cannot exceed income from renting her capital units (both previously saved and inherited) and supplying labor.

We assume that adults receive a labor endowment in each period so the total amount of labor supplied is exogenous in the model. We could, of course, make labor supply

[5] In our robustness exercises we consider adding children's consumption to the utility function, as in Curtis et al. (2015a) and Curtis et al. (2015b)

endogenous by adding leisure time to the utility function. However, our model is not rich enough to capture the societal changes in the latter part of the 20th century that resulted in more women in the workplace than in the early part of the 20th century. To match those changes, we likely would need time-variation in the preference parameters of the cohort representative adults or in home production technology. Calibrating these preference parameters to match historical data would be akin to making labor supply exogenous. Given that we abstract from business cycle fluctuations, look at long horizons, and focus on capital accumulation dynamics, we make labor supply exogenous for simplicity.

The inter-temporal Euler equation can be written as

$$\left(\frac{C_{a+1,t+1}}{C_{a,t}}\right)^{\nu} = \beta\Gamma_{a,t,t+1}\left(1 + R^{K}_{t+1} - \delta\right). \tag{3}$$

The effective discount factor of adults depends on the usual discount factor, β, adjusted by the probability of surviving to enjoy consumption in the next period, $\Gamma_{a,t,t+1}$. In the absence of a bequest motive, adults therefore become increasingly impatient as they approach the end of their lives and the probability of dying rises. To the extent that the probability of surviving is close to 1 in the early decades of adulthood, consumption growth is fairly steady because the young are relatively patient.[6]

2.2 Firms

The representative firm rents capital and labor inputs in perfectly competitive factor markets. It uses those factors in a constant-elasticity-of-substitution (CES) production function to produce final output, Y_t. The functional form of the production technology is given by

$$Y_t = K^{\alpha}_{t}\left(Z_t L_t\right)^{1-\alpha}, \tag{4}$$

[6] Consistent with this prediction, Browning and Ejrnæs (2009) provide evidence that two-adult households without children—for which household consumption roughly equals adult consumption—have fairly flat consumption profiles over their midlife.

where Z_t is the level of labor-augmenting technology. Firm are atomistic and rent capital and hire labor in perfectly competitive factor markets so as to maximize per-period profits. In our robustness exercises, we consider other specifications of the CES production function.

2.3 International migration, aggregation, and market clearing

In addition to births and deaths, we allow the population to change through international migration, a feature that helps us to ensure that, in our simulations, the size and composition of the population matches official estimates and projections of the U.S. population. At the beginning of each period, a measure $m_{a,t}$ of individuals aged a enters the economy on net. For simplicity, and similarly to Fehr et al. (2003), we assume that these immigrants (or emigrants, if $m_{a,t}$ is negative) are identical in terms of asset holdings, labor endowments, and demographic characteristics to the resident population of age a.

Let $\mu_{a,t}$ be the number of individuals aged a at the beginning of period t after accounting for net migration and births, but before deaths have occurred. Given our assumptions, this variable evolves according to $\mu_{a,t} = \mu_{a-1,t-1}\Gamma_{a-1,t-1} + m_{a,t}$ for all $a > 0$, and $\mu_{0,t} = births_t + m_{0,t}$ for newborns. Equilibrium in the labor market requires that

$$L_t = \sum_{a=A_c+1}^{A} \mu_{a,t}\Gamma_{a,t,t+1}e_{a,t}, \tag{5}$$

that is, the economy-wide labor supply is equal to the total labor endowment of surviving adults. Equilibrium in capital markets requires

$$K_t = \sum_{a=A_c+1}^{A} \mu_{a,t}\Gamma_{a,t,t+1}K_{a,t} + \phi \sum_{a=A_c+1}^{A} \mu_{a,t}\left(1 - \Gamma_{a,t,t+1}\right)K_{a,t}. \tag{6}$$

The aggregate capital stock is composed of two summations. The first summation is the capital that surviving adults brought into the period. The second summation is the redistributed capital from adults who die in the period. Finally, goods-market clearing requires that production be equal to the sum of household consumption and gross investment aggregated across the cohort-representative adults,

$$Y_t = \sum_{a=A_c+1}^{A} \mu_{a,t}\gamma_{a,t}C^h_{a,t} + \sum_{a=A_c+1}^{A} \mu_{a,t}\left(K_{t+1,a} - (1-\delta)\left(K_{t,a} + \phi\Xi_t\right)\right). \quad (7)$$

Importantly, the presence of immigration does not change the optimality conditions of the household as the effects of immigration will show up in prices, which the household takes as given.

2.4 Solution method

We consider a perfect-foresight equilibrium such that, in each period, (a) goods and factor markets clear, (b) factor prices correspond to their marginal products, (c) the bequests received by surviving adults sum up to the capital held by adults who passed away at the beginning of the period (net of any destroyed capital), and (d) household optimality conditions are satisfied. While uncertainty may influence household decision making, we are interested in the low frequency effects of demographic changes, which are predictable far in advance. Thus, we view the perfect foresight assumption as an innocuous abstraction for the questions addressed by this paper.

We describe our solution method to find the perfect-foresight general equilibrium in the appendix. In short, we conjecture the paths of aggregate variables over the simulation period and solve the problem of the cohort-representative adults. The solution to this problem is conditional on those paths and on the household's demographic variables (namely the mortality rates, fertility rates, and labor endowments conditional on age and birth cohort), net migration, and the initial conditions at the start of the simulation period in 1900. We next compute new paths for the aggregate objects by aggregating across the individual investment, saving, and labor supply decisions of the cohort-representative adults and by imposing the first-order conditions of the representative firms. If these paths differ from our initial guesses, then we update the paths and iterate an additional time. Once we have reached numerical convergence, our search for the equilibrium stops. Because the equilibrium is conditional on an initial distribution of capital holdings and population at the start of 1900, all aggregate objects, including the equilibrium real rate, depend on these distributions. However, 1900 is sufficiently far in the past that reasonable choices of the initial distributions have little influence on equilibrium outcomes in the post-war period. The solution method is sufficiently rapid that we solve the model at a quarterly frequency.

3 Model calibration and demographic data

We take as inputs to our model demographic data, as well as labor-market participation rates. In particular, we use data on mortality to calibrate $\Gamma_{t,t,t+a}$ for each age-specific cohort. These data pin down the effective discount rate in Equation (3) for each household aged a at every date. We use population data to ensure that the correct proportions of old and young households obtain at every point in time. As a result, the rental rate of capital and real interest rates reflect the relative supply of savers and borrowers. Finally, we use labor market data to ensure that the supply of labor by age and over time is consistent with the data.

In this section, we describe the historical demographic data that we utilize to calibrate our model as well as the projections we utilize to calculate household expectations about the future. We also discuss the calibration of the model's parameters. A key feature of our calibration is that the economy eventually converges to a balanced growth path, which facilitates the computation of a solution. The simulation period runs from 1900 to 2400; such a long period recognizes that life-cycle decisions at any point in time take into account the evolution of factor prices and demographic variables over both the past and the next several decades.

3.1 Mortality

Our individual life expectancy data are derived from mortality rate estimates and projections reported by Bell and Miller (2005) from the U.S. Social Security Administration. These estimates and projections cover every decade from 1900 to 2100 and are reported separately for men and women. We aggregate mortality rates by age and birth cohort across genders in proportion to the number of surviving males and females up to the period. We interpolate mortality rates conditional on age and cohort within inter-census periods and age groups using a cubic spline. In our applications, we assume that the mortality rates in 1900 prevailed in all earlier periods and, similarly, that the mortality rates in 2100 prevail in all subsequent periods.[7] Figure 2 shows a dramatic drop in infant mortality in the early 20th century, along with a realized and projected steady increase in adult life expectancy. The most dramatic gains in life expectancy were came early in the 20th century.

3.2 Fertility

To determine population over time, we combine our data on mortality rates with data on the number of live births, female fertility, and yearly infant population estimates. Data on live births are published by the National Center for Health Statistics (NCHS) starting in 1972.[8] Prior to that year, we construct live birth estimates by adjusting infant population estimates from the Census Bureau for mortality. We note that, in recent decades, yearly observations from the NCHS series essentially coincide with corresponding five-year data

[7]The cohort born in 1900 had a lower probability of death at all ages than persons born earlier in the 19th century. For example, the probability of dying at age 65 conditional on having survived to age 64 was 3.6 percent for the cohort of men born in 1900, whereas 4.2 percent of men who turned 65 years in 1900 died that year. The corresponding gap for women is similar, suggesting that the steady fall in mortality rates around that time was not primarily due to lingering consequences of the American Civil War.

[8]We use annual estimates of live births from NCHS. The NCHS also makes available monthly data on live births, however these data contain apparent seasonal patterns.

from the United Nations' *World Population Prospects (WPP): 2015 Revision*. We use the WPP data, which extend through 2100, to construct birth-rate estimates beyond 2014. Of note, by the end of the current century, the WPP data project that live birth rates have essentially stabilized. Using this observation, we impose that the number of live births is constant from 2100 onward. Further details regarding the construction of fertility rates are available in our appendix.

3.3 Employment rates

We calibrate the labor endowments of adults in the model using data on the employment-population ratios (EPRs) by age groups published by the BLS since 1948. We complement these data with a forecast of the EPRs for 2024 based on sectoral employment projections published by the BLS. Given that we do not model labor force participation, we interchangeably refer to the EPRs as the "employment rates," although our EPR measures do not correspond to official employment rates, which are conditional on labor force participation and thus less suited for measuring the aggregate labor supply. We assume that employment rates by age before 1948 and after 2024 correspond to the employment rates by age in those years. Cyclical fluctuations in employment rates are evident in the raw data; given our focus on explaining longer-term trends in interest rates and GDP growth, we use filtered versions of the published series that strip out these movements.[9] We set the employment rates of people aged less than 18 years to zero so as to conform with the model's treatment of dependent children.

Figure 7 shows our fitted employment rates by age. A few patterns are apparent. Most notably, employment rates have risen overall since the 1950s, a consequence of women participating in the labor force in greater numbers so as to more than offset corresponding

[9]To extract the trends in quarterly EPR series, we use the Hodrick–Prescott filter with a Lagrange multiplier of 25,000. This value is much higher than the typical value of 1,600 used in the business cycle literature. However, we found that the typical multiplier is much too low to filter out the cyclical effects of the Great Recession. See our technical appendix for further discussion.

declines in male employment. The employment rates rise steeply over an individual's late teens and early 20s, but then plateau for the next few decades. Some dynamics are also apparent around the time of retirement. In the early 1960s, workers typically remained active in the labor market until later in life than in the subsequent decades. The employment rates of older workers initially declined for a couple of decades, a tendency that subsequently reversed. Our projections are consistent with a continuation of the trend toward increased participation of older workers in the labor force in the coming decade.[10]

3.4 Other parameters

To isolate the effects of demographic factors on the aggregate objects of interest, our baseline calibration features no technology growth (that is, $Z_t/Z_{t-1} = 1$). In some simulations, we will augment the model with a TFP series matching low-frequency movements in productivity growth to refine our analysis of the transition to the new normal. Our main TFP data source is an updated version of the data from Fernald (2014), whose data sample runs from 1947:Q2 to 2016:Q1. Before 1947, we assume that TFP grew at an annual rate of 2 percent, which is in the middle of the range of estimates surveyed by Shackleton (2013). After 2016, we assume that TFP is growing at a rate of 1.1 percent, which is consistent with the projections of Gordon (2015). As with employment rates, we filter our TFP series so as to strip out cyclical fluctuations.[11]

The last parameter that we need to pin down is the household's discount factor, β. We calibrate this parameter such that the average level in the 1980s of the real short-term interest rate in the model—which we refer to as the equilibrium real rate—matches the average point estimate of the real interest rate expected in the long run calculated by Johannsen

[10]These pattern are consistent with several and at times offsetting factors, including increased access to public and private savings and pension funds, improved health and access to health insurance, and longer life expectancy. For a recent review of the relevant literature on the labor supply of older workers, see Coile et al. (2016).

[11]We again use the Hodrick–Prescott filter with a Lagrange multiplier of 25,000 on quarterly data.

and Mertens (2016b) for that decade. We compare the model's path for the short-term real interest rate to the estimate from Johannsen and Mertens (2016a) of the real interest rate expected in the long run. We see these authors' estimate as a good benchmark because, like our model, it seeks to capture low-frequency movements in real interest rates that are unrelated to business cycle developments. Moreover, the contour of the estimate accords with low-frequency movements in observed real interest rates, being low in the 1960s, somewhat higher in the 1970s and 1980s, and low again in the recent period. We will return in section 5 to the discussion of our benchmark for the equilibrium real rate. For now, we note that changing β shifts the level of the equilibrium real rate path in our simulations but otherwise has little effect on the contribution of demographic factors to observed changes in real interest rates and real GDP growth over time.

4 Main results

Our dynamic simulations track the evolution of the U.S. economy from 1900 to 2400. However, our discussion focuses on the implications of the U.S. demographic transition for the recent and next few decades, especially as they relate to a possible new normal for the U.S. economy. One reason to discount the model results for the early decades of the 20th century is that the initialization of the simulation is subject to some uncertainty. Our sample begins only 35 years after the end of the U.S. Civil War, which saw the destruction of a substantial amount of capital and the tragic death of 2.5 percent of the population—mostly men who would have been part of the labor force.[12] The first half of the 20th century also encompasses two world wars and the Great Depression, whose economic consequences are outside of our framework. The uncertainty surrounding demographic projections and fu-

[12]Moreover, the 1900 Census counts by age are rather imprecise, displaying jumps in the number of persons reporting that their age in years ends with either a 0 or a 5. The quality of Census data appears to have improved quickly thereafter, making us confident that our key findings are insensitive to reasonable deviations from the initialization of demographic data.

ture technology growth similarly grows with the time horizon. That said, knowledge of the current distribution of population, mortality rates, birth rates, and immigration rates imply that the evolution of the population over the next few decades is reasonably certain.

4.1 Dynamics of real interest rates

Figure 8 shows the model's dynamic equilibrium real rate over the 1960–2030 period under our baseline specification with dependent children. At a qualitative level, the model generates a path of the equilibrium real interest rate that resembles that of Johannsen and Mertens (2016b) over the period: The equilibrium real rate rises through most of the 1960s and 1970s, reaches a peak around 1980, and falls thereafter. Time-series estimates of the equilibrium real rate are inherently uncertain, as the band in the panel highlights. Nonetheless, we find the similarity in the timing of the rise, peak, and decline in the equilibrium real rate between these authors' estimates and our baseline estimates rather remarkable. Since 1980, Johannsen and Mertens (2016b) estimate that the long-run equilibrium real rate has fallen $\frac{3}{4}$ percentage point, whereas our model predicts a modestly larger fall, at $1\frac{1}{4}$ percentage points. Put differently, demographic factors appear to be an economically important contributor to low-frequency movements in the equilibrium real rate, possibly explaining all of the permanent declines in real interest rates in recent decades.

As the figure shows, the results are broadly similar when we assume that children's consumption does not enter enter the household utility function, a common assumption in macroeconomic models. When dependents are ignored, the fall in the equilibrium real rate since 1980 occurs somewhat more gradually. By contrast, our model with dependent children has the steepest declines in real interest rates shortly after the year 2000, and thus predicts that the equilibrium real rate was falling especially rapidly during the Great Recession and its aftermath.

Figure 9 offers some guidance to understand the demographic forces that are driving our simulation. The top panel shows that the economy-wide employment rate increased consistently as the baby-boom generation entered the labor market, peaking about the year 2000. The bottom panel shows that the number of dependent children per adult dropped precipitously as the baby-boom generation reached the age of maturity. Additionally, the number of inactive adults relative to employed adults—a measure of the importance of older dependents—similarly reached a turning point around the year 2000 and is projected to increase in the next few decades as the baby-boom generation retires.[13] In this sense, the U.S. economy has had a demographic dividend as of late, in that the number of workers relative to the total population has been historically high.

4.2 Dynamics of GDP growth and other variables

Figure 10 shows the predictions from our baseline model for real GDP growth. Our baseline model abstracts from productivity growth, yet GDP grows at different rates over time for a variety of reasons. Most importantly, movements in the total number of working-age adults and in their employment rates affect the total labor supply, and thus have direct effects on real GDP growth. Demographic factors also affect the accumulation of capital, although the corresponding contribution to movements in GDP growth is small. Our model predicts that demographic factors caused real GDP growth to rise about a percentage point from 1960 to 1980. From 1980 to 2015, the transition toward lower growth in the labor supply erased that increase: Initially, the decline was modest but picked up around the year 2000.

Figure 11 provides information about the effects of the demographic changes in our data on some economic variables. By construction, the real rental rate of capital in our model (shown in the upper-left panel) mimics movements in the equilibrium real rate. The real rental rate of capital increases as the baby-boom generation enters the labor force and capital is relatively scarce. The net saving rate follows a similar pattern as the rental rate of

[13] We calculate this ratio as $(1 - \bar{e}_t) / \bar{e}_t$, where $\bar{e}_t$ is the average employment rate of adults in the sample.

capital in that it peaks around 1980 and declines sharply after the year 2000. Changes in the net saving rate reflect the desire of the baby-boom generation to save for retirement. After these workers leave the labor force, they draw down their capital, depressing net saving rates. When we include dependent children in our model, the contour of the net saving rate is somewhat changed from 1980 through 2010, reflecting the resources that the baby-boom generation allocates toward its dependent children. However, these differences are small relative to the effect of retirement saving and spending by baby-boom generation. During the period in which the proportion of adults active in the labor market is at its highest, the capital-labor ratio is relatively low, even though net savings is high (as shown in the lower-left panel). The capital-labor ratio rises even as the net saving rate falls because the retirement of workers from the baby-boom generation puts downward pressure on the growth rate of the labor supply.

The real wage (measured in efficiency units, and shown in the bottom-right panel) remained low and falls through most of the 1960s and 1970s, reaching a trough in the 1980s, before steadily rising. The movements in the real wage are much smaller than the movements in the capital-labor ratio and reflect the relative scarcity of factors of production; that is, labor is relatively scarce in our model after workers in the baby-boom generation begin to retire. Even though the net saving rate is declining after the year 2000, the model predicts that the equilibrium real interest rate is falling because the capital-labor ratio is dramatically increasing. This observation cautions that simple dynamic analyses that ignore the rich variation in demographic factors and composition of the population might draw incorrect conclusions. Moreover, it suggests that demographic factors offer an explanation for puzzlingly weak business investment in the wake of the Great Recession, which some authors, such as Banerjee et al. (2015), have instead traced back to persistent uncertainty and limited investment opportunities.

4.3 Parsing the effects of fertility, mortality, and employment rates

To assess the respective contributions of changes in fertility, mortality, and employment rates to macroeconomic developments since 1980, we perform alternative simulations in which we freeze demographic variables from 1960 onward.[14] We use the model with dependent children as our baseline model and we keep β fixed across simulations so as to isolate the effects of the demographic data.

We present the counter-factual paths of the equilibrium real rate in the top panel of Figure 12. When we keep fertility, mortality, and employment rates all fixed at their 1960 values, we see that there would have been no net decline in the equilibrium real interest rate since 1980. Put differently, the entirety of the decline in the equilibrium real interest rate that our model finds for the recent decades is a direct consequence of the demographic changes that happened from 1960 onward.

To analyze the respective contribution of each demographic factor, we consider the effects of freezing each demographic variable in turn.[15] The two most consequential demographic changes are the drop in fertility rates and the increase in employment rates, which each contributes $\frac{1}{2}$ percentage point to the $1\frac{1}{4}$–percentage-point decline in the equilibrium real rate between 1980 and the present. Intuitively, the demographic dividend stemming from the drop in fertility rates in the 1960s and 1970s facilitated the accumulation of capital in aggregate. Moreover, the retirement of a relatively large cohort of baby boomers has further put upward pressure on the aggregate capital-labor ratio of late, leading to a fall in the equilibrium real interest rate. Similarly, employment rates were relatively low in 1960, as we showed in Figure 7. Since that time, the number of households with two adults who

[14]That is, we assume that a person aged a in any period after 1960:Q1 faces the same mortality rate, fertility rate, or employment rate (or combinations of these three variables, depending on the simulation) as someone aged a did in 1960:Q1 under our baseline calibration. The net migration flows by age and period are kept the same as under our baseline calibration.

[15]The contributions of demographic factors to the fall in the equilibrium real rate since 1980 are roughly additive in the model, a fact that we verified by solving the model under all possible combinations of frozen variables.

participate in the labor market has increased dramatically. As such, the aggregate labor supply has increased through the latter part of the 20th century, in turn facilitating aggregate capital formation and exacerbating the deceleration in the aggregate labor supply once baby boomers began to retire.

The contribution of lower mortality rates to the fall in the equilibrium real rate since 1980 is only half as large, at ¼ percentage point, than those of the fall in fertility rates and rise in employment rates. That changes in life expectancy alone have had only a small effect on the equilibrium real interest rate over this period is perhaps unsurprising. To put these changes in mortality rates in perspective, we note that the mean life expectancy at birth conditional on the mortality rates observed at the start of 1980, at 79.6 years, was only three and half years less than the corresponding statistic based on the mortality rates prevailing in 2015. By contrast, the total gain in life expectancy over the 1900–2100 period is nearly three and a half decades. Moreover, the gains in life expectancy over the past few decades reflect primarily a fall in mortality rates during an adult's years in retirement, thus having almost no effect on the aggregate supply of labor. And with younger cohorts discounting the consumption of their older selves, the incentives to accumulate more capital during the productive years are little changed either. On net, the aggregate stock of capital rises only modestly, which is consistent with a small fall in the equilibrium real rate at the margin.[16]

To narrow in on the timing of the demographic changes that were most important for the decline in the equilibrium real interest rate, we do similar experiments in which we freeze demographic variables starting in 1980. These new counter-factual paths for tbhe equilibrium real interest rate are shown in the bottom panel of Figure 12. It is clear that

[16]Our estimate of the contribution of longer life expectancy to the fall in the equilibrium real interest rate is much smaller than found by Carvalho et al. (2016). The difference likely reflects the fact that, in these authors' model, all retirees face the same mortality rates independently of their age or asset holdings. By contrast, in our model, the mortality rates are calibrated to the age and cohort of each adult; an increase in longevity primarily affects older adults whose capital holdings have already declined, so that the effect on the aggregate capital stock is more limited.

holding demographic variables fixed at their 1980 levels, either independently or jointly, has essentially no effect on the subsequent path of interest rates predicted by the model. The reason for this result is that, by 1980, fertility rates were close to levels that prevail even today (which can be seen in Figure 6), employment rates had increased markedly but then stabilized (which can be seen in Figure 7), and, as mentioned above, mortality rates have fallen only somewhat modestly. In sum, the subsequent decline in real interest rates due to demographic factors was mostly due to demographic changes that took place before 1980. The subsequent dynamics reflect the life cycle of those demographic changes, meaning the contours of the path of interest rates largely reflects the impact of the baby-boom generation.

Figure 13 shows that keeping employment rates or mortality rates fixed at either their 1960 or 1980 levels has little effect on the growth rate of real GDP overall. Instead, the bulk of the fall in real GDP growth of late is accounted for by declines in the aggregate labor supply related to the gradual retirement of the baby-boom generation. As can be seen from Figure 13, fixing fertility rates at their 1960 levels would have kept real GDP growth from falling. However, because fertility rates fell, our model predicts that real GDP growth initially increased as members of the baby-boom generation reached working age and then started falling in the 1980s. The model also predicts that the most dramatic declines in real GDP growth due to demographic changes happened shortly after the year 2000, again around the time that the baby-boom generation began to retire.

4.4 The effects of productivity growth

We add productivity growth to the model to assess its effects on the transition to the new normal for the U.S. economy. Authors such as Summers (2014), Gordon (2016), and Fernald (2016) argue that the rate of potential GDP growth in the future will be modest relative to its average over the post-war period. Changes in TFP have two direct effects in theory. First, because productivity enters the production function, GDP growth should slow if productivity growth slows. Second, real interest rates are proportional to TFP growth in a broad class of models, meaning interest rates should fall if productivity growth falls. In the standard real business-cycle (RBC) model, the real interest rate is given by

$$r = \frac{(1+\Delta z)^{\nu}}{\beta} - 1,$$

where $1 + \Delta z \equiv Z_t/Z_{t-1}$.

We perform a dynamic simulation of our model in which technology growth matches prominent estimates of its historical and projected trend. Our series captures relatively rapid productivity growth before the early 1970s, a slow down from the 1970s through the early 1990s, a temporary pickup associated with the information technology boom, and finally modest rates of technology growth from the early 2000s onward of the kind envisioned by Gordon (2016). See our appendix for details on the construction of our series. Consistent with our earlier approach, we assume that these developments are anticipated and calibrate the household's discount factor, β, such that the average real rate in the 1980s corresponds to the average point estimate of Johannsen and Mertens (2016b). Absent this adjustment, the introduction of positive technology growth in the model would raise the path of the equilibrium real rate compared with our baseline specification with no technology growth.[17]

Figure 10 shows that productivity growth shifts up the model's predicted path for GDP growth. Even though we do not use data on aggregate output, the model's predictions for real GDP growth are remarkably similar to the growth rate of the Hodrick–Prescott filtered trend in actual output, also shown in Figure 10.[18] It is noteworthy how our model captures

[17]The calibration requires a value of β that is slightly above 1. In a standard RBC framework, values of β above 1 would be inconsistent with the existence of an equilibrium with balanced growth because the incentives to save would be so strong that households would want to accumulate an ever increasing amount of (technology-adjusted) capital. By contrast, in our OG model, the mortality risk lowers the effective discount factors of households to levels that ensure the existence of an equilibrium.

[18]At a business cycle frequency, one might expect that calibrating the model to match data on the aggregate labor supply would mostly capture movements in aggregate output given that the capital stock is relatively

the size of the decline in real GDP growth from the 1980s to the present. The reason is that, aside from some years of very high productivity growth in the 1990s, data from Fernald (2014) indicate that TFP growth was low even in the 1980s. By contrast, high levels of TFP growth in the 1960s correspond to a larger difference between our baseline calibration and our calibration with TFP growth in terms of real GDP growth. Thus, even with historical movements in TFP, our model predicts that much of the decline in real GDP growth since the 1980s is associated with demographic change.

Relative to the baseline case, the boom in productivity growth initiated in the 1990s keeps the equilibrium real rate elevated through much of the first decade of the 2000s. The equilibrium real rate only begins falling in earnest around the onset of the Great Recession, proceeding to drop almost one percentage point in the span of about 10 years to a level similar to that in our baseline simulation. Importantly, the fall in the equilibrium real rate under this scenario is largely concomitant with the global financial crisis. As such, factoring in the temporarily high productivity gains of the 1990s and early 2000s suggests a material risk that observers could confuse the persistent but ultimately temporary consequences of the financial crisis with other influences originating from the demographic transition and technology growth cycles.

Looking at the whole period shown, Figure 14 points to a rather weak connection between the equilibrium real rate and the pace of technology growth over long horizons. Notably, even though technology growth was elevated by historical standards prior to the early 1970s, and should thus have been associated with a high equilibrium real rate, realized real interest rates were quite low at the time. Accordingly, the figure features a persistently large gap between the equilibrium real rate predicted by the model and the estimates of Johannsen and Mertens (2016b) in the 1960s and 1970s.

---

fixed in the near term. However, at low frequencies, the endogenous dynamics of capital accumulation is also potentially important for explaining movements in the trend of real GDP growth.

5 Discussion and robustness

Judging the empirical success of our model requires taking a stand on the low-frequency paths of real interest rates and the growth rate of real GDP that the model should seek to capture. We take the change in our Hodrick–Prescott filtered output series as a good estimate of the decline in GDP growth; other low-frequency filters suggest similar contours.[19] By contrast, estimates of permanent changes in interest rates are much more uncertain, a point demonstrated in Kiley (2015). Rachel and Smith (2015) seek to explain a 450-basis-point fall in the neutral rate of interest over the past 30 years, a movement that is in line with the fall in real long-term rates over the period. We judge their benchmark to be much too large to reflect the influence of secular factors alone. As we showed in Figure 1, the 1980s were preceded by relatively low but positive real interest rates in the 1960s and early 1970s, and then negative real interest rates in the remainder of the 1970s as inflation picked up. The historically high real interest rates of the 1980s occurred as the Federal Reserve tightened the policy stance to rein in inflation. Unsurprisingly, the estimates of Johannsen and Mertens (2016b), which net out cyclical influences, suggest a significant but much smaller drop since the 1980s. Another reason for our discomfort with the size of the drop in the neutral rate reported by Rachel and Smith (2015) is that some of the explanatory factors in their analysis were already at play prior to the 1980s; one implication of extending the data farther back in time is that the drop in the neutral rate since, say, the 1960s would be implausibly large to us. Additionally, work by Hamilton et al. (2015) presents evidence that if the equilibrium real rate has declined, then that decline is likely smaller than the magnitude reported by Rachel and Smith (2015). Therefore, we assess that a decline in the equilibrium real interest rate of less than 2 percentage points since 1980 is most consistent with the new

[19] That said, we view the typical value of the Hodrick–Prescott filter Lagrange multiplier for quarterly data, at 1,600, as much too low to eliminate all cyclical fluctuations, especially in light of the unusually prolonged recovery from the global financial crisis.

normal for the U.S. economy.[20]

Our results are subject to a number of potential caveats. A particularly important assumption we have made is that adults have no intentional bequest motive, an assumption that is standard in the OG literature. Relaxing this assumption would likely weaken the size of the decline in the equilibrium real rate since 1980 by facilitating the extent to which economic agents can engage in intertemporal substitution across generations and thus smooth the effects of the baby boom and other demographic elements. By comparison, the effects on real GDP growth would likely be modest because demographics affect the pace of output growth primarily through changes in the labor supply, which is exogenous in our environment. From an empirical perspective, the extent to which intentional bequests matter for the determination of the equilibrium level of the aggregate capital-labor ratio—and thus for equilibrium interest rates—is unclear. Intentional bequests are certainly observed in practice but often come in the form of housing, which Yang (2009) argues has different properties than capital over the life cycle due notably to large transactions costs, the impossibility of fractioning most capital held in the form of housing, and the utility that housing provides to owners. The consideration of housing is beyond the scope of our model.

Below, we consider a number of alternative scenarios to assess the sensitivity of our results to other model features. We first consider the possibility that the elasticity of substitution between factors of production is lower than we have assumed. This consideration is motivated by the central role played by the aggregate capital-labor ratio in the model for explaining movements in the equilibrium real rate, as well as by empirical evidence that the long-run elasticity of substitution is lower than implied by Cobb-Douglas production

[20] Alternative estimates of the long-term real interest rate that are popular in economic analysis include the time-series model of Laubach and Williams (2003) and the related model of Holston et al. (2016). We see these estimates as less preferable for our purposes than those of Johannsen and Mertens (2016b) in light of their apparent cyclical sensitivity and the fact that they imply real equilibrium rates in the 1960s and 1970s that are well above the realized real yields on Treasuries and other risk-free instruments during those decades. The initially elevated estimates in those papers appear to be due to the assumption of a direct relationship between movements in trend productivity growth, which was high at the time, and the equilibrium real rate. See Kiley (2015) for an exploration of identification in the Laubach and Williams (2003) model.

functions. We next consider an alternative specification of preferences that some authors have argued is especially successful at matching consumption/saving decisions over the life cycle. We conclude with a comparison of the U.S. and foreign demographic transitions.

5.1 Elasticity of substitution in production

So far, our analysis has focused on how demographic characteristics affect the supply of labor and capital by households in equilibrium. We now provide a sensitivity analysis of the elasticity of substitution between capital and labor in production, $1/(1-\rho)$, which affects the demand for factors of production in equilibrium, and thus also plays a role in determining equilibrium factor prices and quantities. In particular, the dynamic of the equilibrium real rate in the model is tightly connected to that of the (technology-adjusted) capital-labor ratio. Using Equations (4) and (**??**), we can write the equilibrium real rate as

$$r_t = R_t^K - \delta = \alpha \left(\alpha + (1-\alpha) \left(\frac{Z_t L_t}{K_t} \right)^{\rho} \right)^{\frac{1-\rho}{\rho}} - \delta.$$

In our baseline calibration, we assume that the representative firm's production function is of the Cobb-Douglas form, so that the elasticity of substitution between factors of production is one (that is, $\rho = 0$ and $r_t = \alpha(Z_t L_t / K_t)^{\alpha-1} - \delta$). The empirical relevance of this standard assumption is questionable, with the survey of Chirinko (2008) suggesting a long-run elasticity that is appreciably lower than one.[21] Inspection of the above equation indicates that, for a given capital-labor ratio, both the level of r_t and its sensitivity to movements in the capital-labor ratio fall as we lower the elasticity of substitution from 1.

[21] The evidence reviewed in Chirinko (2008) is consistent with a long-run elasticity somewhere between 0.4 and 0.6, with the preferred value of the more recent study by Chirinko and Mallick (2015) being at the bottom of that range. Under a standard CES production function like Equation (4), one cannot lower the elasticity parameter down to that range while simultaneously imposing that the share of capital in aggregate income is aligned with that in the data. In particular, for a capital share of aggregate income s_K and a balanced-growth real return on capital (net of depreciation) of r, Equation (4) requires setting α equal to $r^{\rho}/(s_K)^{\rho-1}$. If we set r such that it conforms with the estimates of Johannsen and Mertens (2016b) in the 1980s, then the value of α consistent with s_K=0.35 exceeds one whenever the elasticity of substitution between capital and labor is assumed to be below 0.72.

In this robustness check, we decrease the elasticity to 0.75 while still imposing that the parameter α is consistent with a share of capital in aggregate income of 35 percent. We also re-calibrate the value of β such that the model's average equilibrium real rate in the 1980s matches the point estimate of Johannsen and Mertens (2016b). Figure 14 shows that lowering the elasticity of substitution accentuates the extent of the fall in the equilibrium real rate since 1980 that can be attributed to demographics. The current difference with respect to the baseline case is somewhat small, however, at about 2 tenths of a percentage point.

5.2 Adult-equivalent preferences

The predicted peak in household consumption in our baseline model occurs somewhat earlier in the life cycle than it does in the data. As an alternative, we consider a variant of Browning and Ejrnæs (2009) preferences in which total household consumption increases in the age of each of the children, thus postponing the peak in household consumption to when dependent children approach adulthood. These authors' specification of the utility function is of additional interest to us because they show that it can explain the magnitude and timing of the rise and peak in life-cycle consumption without recourse to other causal mechanisms, such as liquidity constraints, precautionary motives, and non-separabilities between consumption and labor supply—features that are also absent in our environment. Household utility is given by

$$\sum_{s=0}^{A-a} \beta^s \Gamma_{a,t,t+s} \frac{\left(C^h_{a+s,t+s} e^{-f_{a+s,t+s}}\right)^{1-\nu}}{1-\nu} e^{f_{a+s,t+s}}. \tag{8}$$

The function $f_{a,t} = \delta_h \log\left(m_{a,t}\right)$ allows for economies of scale in consumption based on the number of adult-equivalent members of the household, denoted $m_{a,t}$.[22] In simulations

[22] Although the utility function is specified in terms of total household consumption, we retain the dependence of the household's effective discount factor in terms of the survival rate of adults only. Current and projected

not shown, we obtain results similar to our benchmark model when we use this specification of preferences with the same parameterization as Browning and Ejrnæs (2009).

5.3 Open-economy considerations

One limitation of our framework is that it does not incorporate international influences on the U.S. equilibrium real rate. Or, there are reasons to expect foreign developments to contribute to the determination of domestic interest rates. Notably, there is a strong correlation across real interest rates worldwide, indicating either common driving factors, strong international spillovers, or both.[23] In a similar vein, the global saving glut hypothesis of Bernanke (2005) posits that U.S. interest rates have felt downward pressure from an excess of global savings emanating from Asian and oil-producing economies; this hypothesis challenged the idea that the dynamics of the U.S. current account since the mid-1990s reflected primarily domestic developments. More generally, movements in U.S. financial variables are sometimes traced back to foreign economic and political developments, as has been evident during bouts of shifting global risk appetite in the wake of the financial crisis. While we do not aim to explain the effects of these foreign elements on U.S. real interest rates, it is nonetheless legitimate to ask whether open-economy considerations might affect our estimates of the contribution of demographic factors. A number of elements give us confidence that our finding that demographic factors are an economically important driver of the U.S. equilibrium real interest rate is robust to open-economy considerations.

One such element is that the United States and the world's other advanced economies are undergoing demographic transitions that are similar in terms of magnitude and timing. In the top panel of Figure 15, we show that the step down in U.S. working-age population growth since the 1970s echoes a similar pattern in OECD economies and in the world

mortality rates are rather low for the adult's child-rearing years, so that $\Gamma_{a,t,t+1}$ remains close to one until late in life when household consumption coincides with adult consumption.

[23] On the correlation in real interest rates across countries, see, among many contributors, Desroches and Francis (2006) and King and Low (2014).

more broadly. The 2000–2030 period features an especially sharp slowdown in population growth. In the bottom panel, we show that the U.S. mean life expectancy at birth has been rising in tandem with that in OECD economies. Similar gains in life expectancy at birth are observed when we include both developed and developing economies, although the level is pulled down by still markedly higher death rates in the less-developed countries.

Another element is the size and direction of international capital flows. In our model, demographic factors put downward pressure on the return on domestic capital. In an open-economy setting with international capital mobility, a fall in the rate of return on domestic capital would, all else equal, induce domestic investors to seek investment opportunities abroad, resulting in a net capital outflow. To the extent that, in aggregate, U.S. investors have succeeded in finding investment opportunities abroad, our closed-economy estimates of the effect of demographics on the U.S. equilibrium real rate could be overestimated. However, in contrast with the prediction of a capital outflow, the U.S. current account balance has been solidly negative since the early 1980s.

A third piece of evidence is that attempts at incorporating demographics in open-economy settings do not suggest economically large differences between closed-economy and open-economy predictions for U.S. real interest rates. In a closely-related study, Krueger and Ludwig (2007) build a multi-country OG model to explore the effects of demographics factors on the rate of return on capital. For their 2005–2080 simulation period, their model projects net declines in the U.S. real interest rate of 86 basis points under the multi-country version of their model and of 79 basis points under the closed-economy version. The difference in point estimates between versions is arguably small and, taken at face value, suggests that foreign demographic developments tend to amplify rather than to dampen the extent to which the U.S. real rate is being pressured downward by demographic developments. Our model without dependent children, which is close to the model considered by Krueger and Ludwig (2007), predicts that the equilibrium real rate will fall by 54 basis points over their sample period. By comparison, our benchmark model with dependent children predicts a net fall in the equilibrium real rate of 79 basis points.

The results of Krueger and Ludwig (2007) present a number of limitations with respect to our questions of interest. One such limitation is that Krueger and Ludwig (2007) focus on projected developments over the course of this century, in contrast with our goal of explaining movements in the equilibrium real rate and other aggregate variables over the past several decades. The initial period in their simulation period—2005—coincides with a period of historically rapid fall in the equilibrium real rate in our model. Importantly for our ability to expand our framework to a multi-country setting, foreign demographic data are generally not as detailed as the U.S. demographic data we use to calibrate our model. Moreover, the time coverage of the foreign demographic data tends to be limited, in turn restricting one's ability to study phenomena unfolding over several generations. For example, their model does not incorporate dependent children and variation in the labor supply across birth cohorts, which we found to be important for explaining the timing and magnitude of the fall in the equilibrium real rate since 1980. Similarly, their exercise leaves aside non-OECD economies for which there is a dearth of demographic data. Nonetheless, these results and the other evidence presented above, taken together, give us reasonable confidence that demographic developments have been an economically important driver of low-frequency movements in U.S. real interest rates over the past few decades.

6 Conclusion

In this paper, we used a calibrated OG model with a rich demographic structure to investigate the extent to which past and projected changes in U.S. family composition, life expectancy, and the labor supply have had an influence on real interest rates, GDP growth, and other macroeconomic variables. We found that these demographic factors alone account for a $1\frac{1}{4}$–percentage-point decline in the natural rate of interest and real GDP growth since 1980—a magnitude that we argue is in line with empirical estimates. Moreover, the model is consistent with the rate of aggregate investment having slowed appreciably over that period. Looking forward, the model suggests that low interest rates, low output growth, and low investment rates are here to stay, suggesting that the U.S. economy has entered a new normal.

We believe that our results will be of interest to academics and policymakers alike. The model implies that the transition to the new normal has been especially rapid over the past decade or so because of demographic factors most-directly associated with the post-war baby boom—a phenomenon that was fully predictable—and the passing effects of the information technology boom on productivity growth. The rapidity at which our model economy has been transitioning to a new normal suggests a risk that the permanent effects of demographic factors could be misinterpreted as persistent downward pressure on the natural rate of interest and net savings stemming from the global financial crisis. Going forward, our results have implications for assessing fiscal sustainability. Moreover, the persistence of a low equilibrium real interest rate means that the scope to use conventional monetary policy to stimulate the economy during typical cyclical downturns will be more limited than it has been the case in the past for a given inflation target.

References

Aaronson, S., T. Cajner, B. Fallick, F. Galbis-Reig, C. Smith, and W. Wascher (2014). Labor Force Participation: Recent Developments and Future Prospects. *Brookings Papers on Economic Activity 49*(2), 197–275.

Banerjee, R., J. Kearns, and M. J. Lombardi (2015). (Why) Is Investment Weak? *BIS Quarterly Review March*.

Bell, F. C. and M. L. Miller (2005). Life Tables for the United States Social Security, 1900–2100. *Social Security Administration–Office of the Chief Actuary, Actuarial Study* (120).

Bernanke, B. (2005). The Global Saving Glut and the U.S. Current Account Deficit. Speech delivered at the Sandridge Lecture of the Virginia Association of Economists, Richmond, VA, March 10.

Browning, M. and M. Ejrnæs (2009). Consumption and Children. *The Review of Economics and Statistics 91*(1), 93–111.

Carvalho, C., A. Ferrero, and F. Nechio (2016). Demographics and Real Interest Rates: Inspecting the Mechanism. *European Economic Review*.

Chirinko, R. S. (2008). σ: The Long and Short of It. *Journal of Macroeconomics 30*(2), 671–686.

Chirinko, R. S. and D. Mallick (2015). The Substitution Elasticity, Factor Shares, Long-Run Growth, and the Low-Frequency Panel Model. *CESifo Working Paper Series* (4895).

Coile, C., K. S. Milligan, and D. A. Wise (2016). Social Security and Retirement Programs around the World: The Capacity to Work at Older Ages – Introduction and Summary. *NBER Working Paper Series No. 21939*.

Colby, S. L. and J. M. Ortma (2015). Projections of the Size and Composition of the U.S. Population: 2014 to 2060. *United States Census: Current Population Reports P25-1143*.

Curtis, C. C., S. Lugauer, and N. C. Mark (2015a). Demographic Patterns and Household Saving in China. *American Economic Journal: Macroeconomics 7*(2), 58–94.

Curtis, C. C., S. Lugauer, and N. C. Mark (2015b). Demographics and Aggregate Household Saving in Japan, China, and India. *National Bureau of Economic Research Working Paper No. 21555*.

Desroches, B. and M. Francis (2006). Global Savings, Investment, and World Real Interest Rates. *Bank of Canada Review 2006*(Winter), 3–17.

Eggertsson, G. B., N. R. Mehrotra, and J. A. Robbins (2017, January). A model of secular stagnation: Theory and quantitative evaluation. *National Bureau of Economic Research Working Paper Series* (23093).

El-Erian, M. A. (2010). Navigating the New Normal in Industrial Countries. Speech, Per Jacobsson Foundation Lecture, International Monetary Fund.

Fehr, H., G. Halder, S. Jokisch, and L. J. Kotlikoff (2003). A Simulation Model for the Demographic Transition in the OECD: Data Requirements, Model Structure and Calibration. W.E.P. - Wurzburg Economic Papers 45, University of Wurzburg, Chair for Monetary Policy and International Economics.

Fernald, J. (2014). A Quarterly, Utilization-Adjusted Series on Total Factor Productivity. *Federal Reserve Bank of San Francisco working paper 19*.

Fernald, J. G. (2015). Productivity and Potential Output before, during, and after the Great Recession. *NBER Macroeconomics Annual 29*(1), 1–51.

Fernald, J. G. (2016). Reassessing Longer-Run US Growth: How Low? *SF-FRB Working Papers Series*.

Gordon, R. J. (2015). Secular Stagnation on the Supply Side: U.S. Productivity Growth in the Long Run. *Digiworld Economic Journal* (100).

Gordon, R. J. (2016). *The Rise and Fall of American Growth: The US Standard of Living since the Civil War*. Princeton University Press.

Hamilton, B. E. and C. M. Cosgrove (2010). Central Birth Rates by Live-Birth Order, Current Age, and Race of Women in Each Cohort from 1911 through 1991: United States, 1960–2005. Table 1. Hyattsville, MD: National Center for Health Statistics. Released: June 30, 2010.

Hamilton, J. D., E. S. Harris, J. Hatzius, and K. D. West (2015). The Equilibrium Real Funds Rate: Past, Present and Future. *National Bureau of Economic Research Working Paper No. 21476*.

Hogan, H., D. Perez, and W. R. Bell (2008). Who (Really) Are the First Baby Boomers? *Joint Statistical Meetings Proceedings, Social Statistics Section*, 1009–1016.

Holston, K., T. Laubach, and J. Williams (2016). Measuring the Natural Rate of Interest: International Trends and Determinants. Board of Governors of the Federal Reserve System.

Ikeda, D. and M. Saito (2014). The Effects of Demographic Changes on the Real Interest Rate in Japan. *Japan and the World Economy 32*, 37–48.

Johannsen, B. K. and E. Mertens (2016a). A Time Series Model of Interest Rates With the Effective Lower Bound. *FEDS Working Paper No. 2016–033*.

Johannsen, B. K. and E. Mertens (2016b). The Expected Real Interest Rate in the Long Run: Time Series Evidence with the Effective Lower Bound. *FEDS Notes. Washington: Board of Governors of the Federal Reserve System*.

Kiley, M. T. (2015). What Can the Data Tell Us About the Equilibrium Real Interest Rate? Finance and Economics Discussion Series 2015–77, Board of Governors of the Federal Reserve System (U.S.).

King, M. and D. Low (2014). Measuring the 'World' Real Interest Rate. *National Bureau of Economic Research: Working Paper No. 19887*.

Krueger, D. and A. Ludwig (2007). On the Consequences of Demographic Change for Rates of Returns to Capital, and the Distribution of Wealth and Welfare. *Journal of monetary Economics 54*(1), 49–87.

Laubach, T. and J. C. Williams (2003). Measuring the Natural Rate of Interest. *The Review of Economics and Statistics 85*(4), 1063–1070.

Rachel, L. and T. Smith (2015). Secular Drivers of the Global Real Interest Rate. *Bank of England Working Paper*.

Reichling, F. and K. Smetters (2015). Optimal Annuitization with Stochastic Mortality and Correlated Medical Costs. *American Economic Review 105*(11), 3273–3320.

Rogoff, K. (2015). Debt Supercycle, Not Secular Stagnation. VOX: CEPR's Policy Portal, April 22., http://www.voxeu.org/article/debt-supercycle-not-secular-stagnation.

Shackleton, R. (2013). Total Factor Productivity Growth in Historical Perspective. *Congressional Budget Office, Working paper 1*.

Summers, L. H. (2014). Reflections on the 'New Secular Stagnation Hypothesis'. *Secular stagnation: Facts, causes and cures*, 27–40.

Yaari, M. E. (1965). Uncertain Lifetime, Life Insurance, and the Theory of the Consumer. *The Review of Economic Studies 32*(2), 137–150.

Yang, F. (2009). Consumption over the Life Cycle: How Different Is Housing? *Review of Economic Dynamics 12*(3), 423–443.

Figure 1: Expected real return on 10-year Treasury security

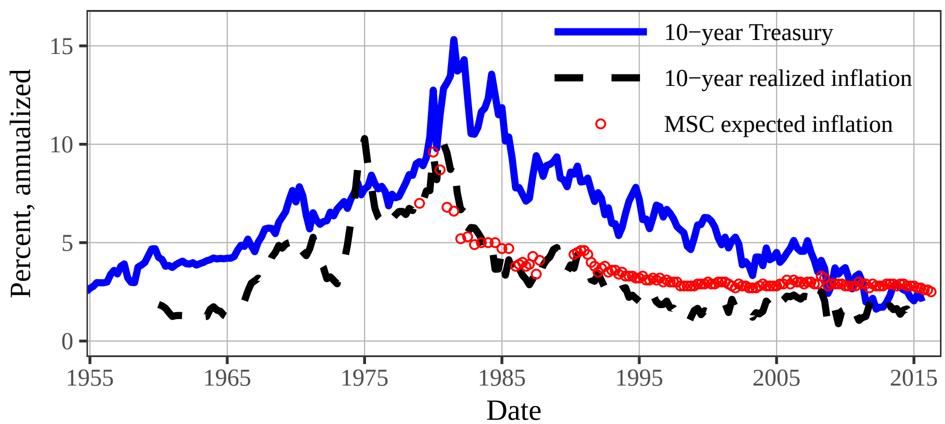

Source: Authors' calculation using data from FRED.
Notes: The "10-year Treasury" rate is the series GS10 in FRED. The "10-year realized inflation" rate is the annual average inflation rate of core PCE inflation. The "MSC expected inflation" rate is the answer to the question about the expected change in prices in general over the next 5 to 10 years in the Michigan Survey of Consumers.

Figure 2: Estimated/Projected survival function by year of birth

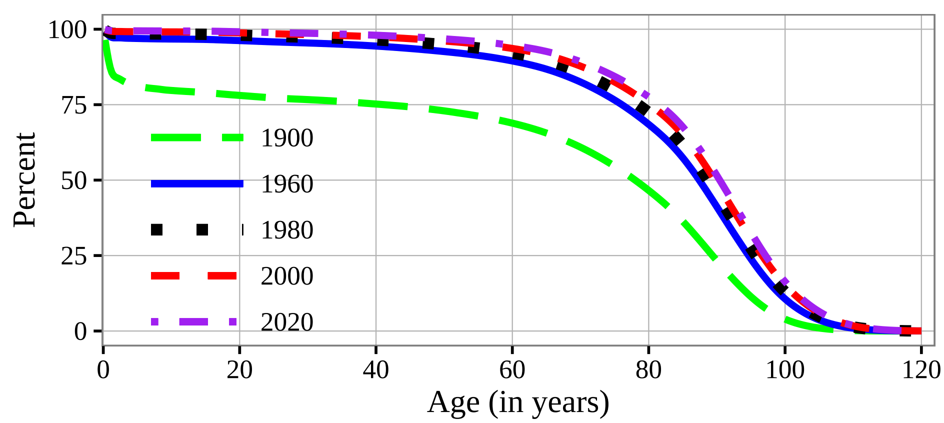

Source: Authors' interpolation of decennial mortality tables by Bell and Miller (2005).

Figure 3: Annual U.S. live birth estimates and projections: 1900 to 2030

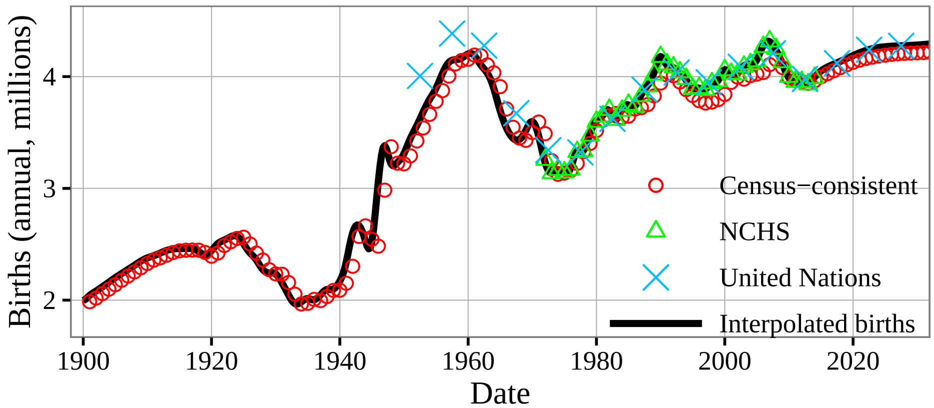

Source: Authors' calculations. Data from National Center for Health Statistics (NCHS), United Nations' *World Population Prospects: Revision 2015* (WPP), U.S. Census Bureau.
Notes: We construct a historical birth series by stitching yearly estimates from the U.S. Census Bureau for 1900–1971, yearly estimates from NCHS for 1972–2014, and five-year forecasts from the WPP estimates. The solid line shows our interpolation.

Figure 4: Median age difference between men and women at time of first marriage

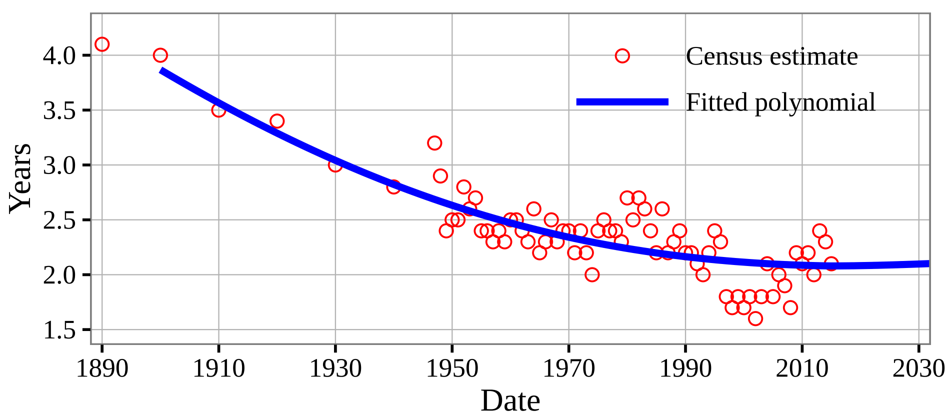

Source: Authors' calculations using data from the U.S. Census Bureau.
Notes: Our interpolation assumes that the median age difference between men and women at the time of first marriage, averaged in 2013 and 2014, corresponds to the long-run value.

Figure 5: Share of births by age of mother (in years)

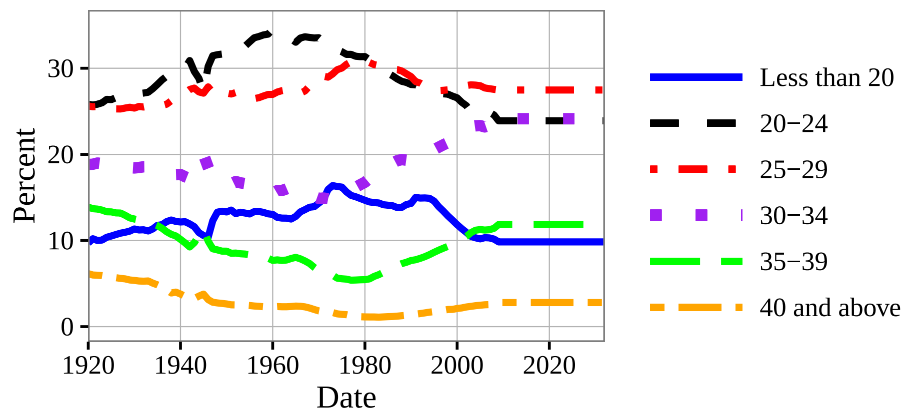

Source: Authors' calculations using U.S. data from Hamilton and Cosgrove (2010).

Figure 6: Average number of dependent children by age and cohort

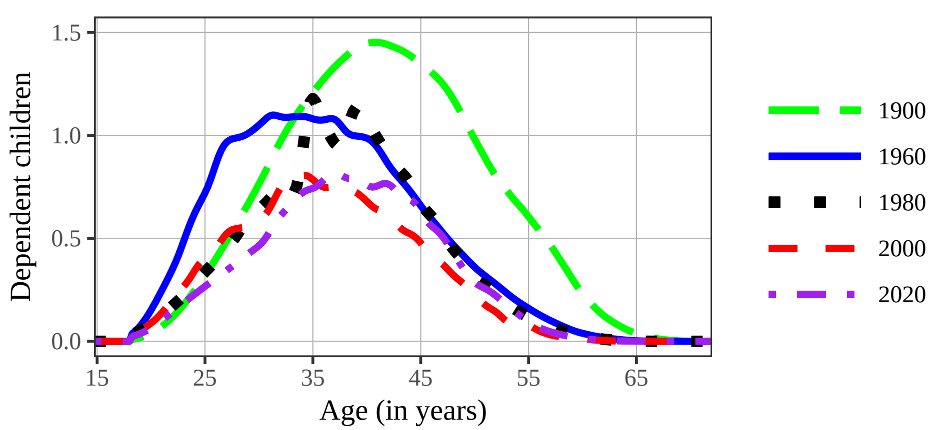

Source: Authors' calculations using data from Hamilton and Cosgrove (2010), the National Center for Health Statistics, the United Nations' *World Population Prospects: Revision 2015*, and the U.S. Census Bureau.

Figure 7: Employment rates by age and calendar year

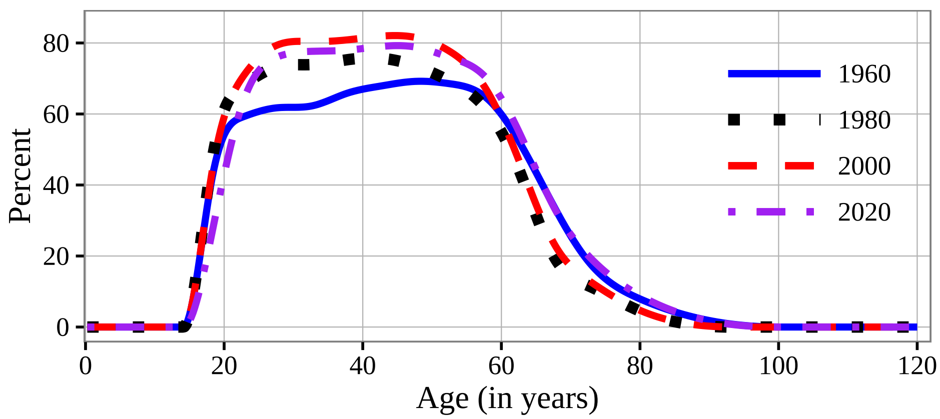

Source: Authors' interpolation of historical and projected data from the BLS.
Notes: The figure shows U.S. employment rates conditional on age. The historical yearly data begin in 1948. The data shown are an interpolation to a quarters of age performed using cubic splines.

Figure 8: Dynamic simulation under baseline demographics: equilibrium real rate

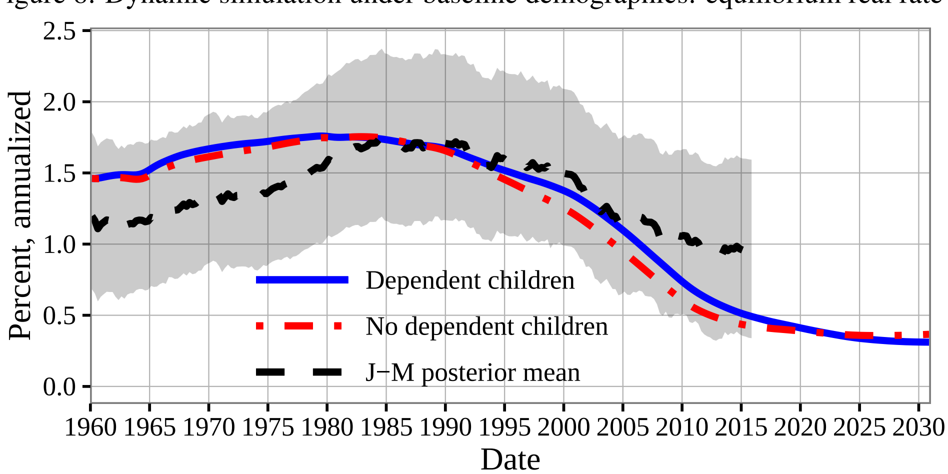

Source: Johannsen and Mertens (2016b) and authors' model simulations.
Notes: The figure shows the short-term equilibrium real rate in dynamic simulations of our model, along with the posterior mean estimate of the expected real interest rate in the long run from Johannsen and Mertens (2016b) and 50-percent uncertainty bands.

Figure 9: Dynamic simulation under baseline demographics: other variables

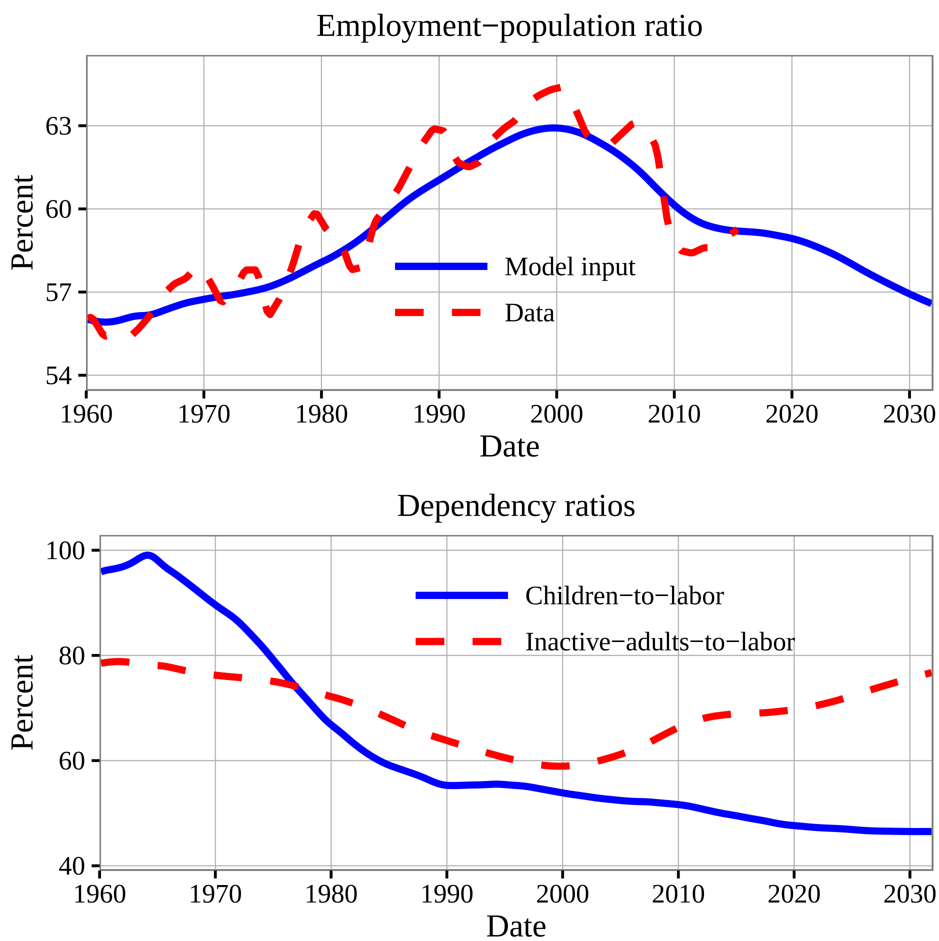

Source: Authors' calculations based on Bureau of Labor Statistics and Census Bureau data.
Notes: The top panel shows the employment-population ratio in the data and in our model after we apply the Hodrick–Prescott filter with a Lagrange multiplier of 25,000. The bottom panel shows the number of dependent children in the model scaled by the number of full-time equivalent labor endowments. The bottom panel also shows the number of full-time equivalent adults who are inactive scaled by the number of full-time equivalent labor endowments.

Figure 10: Dynamic simulation of real GDP growth

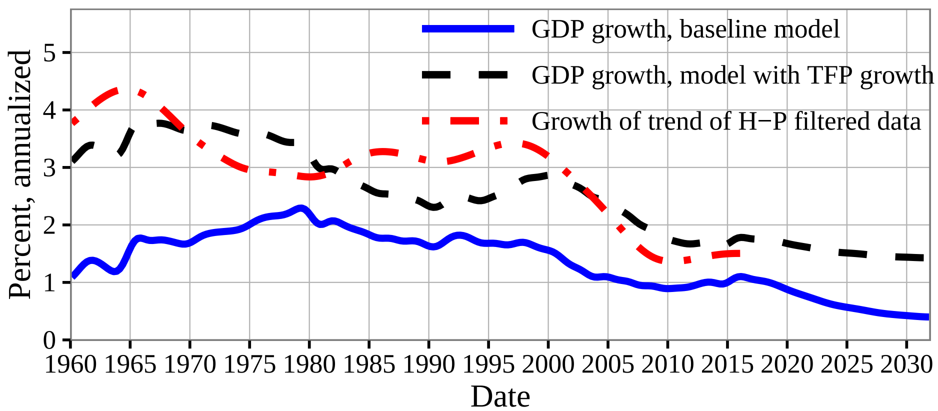

Source: Authors' calculations using data from the U.S. Bureau of Economic Analysis and authors' model simulations.

Notes: The figure shows real GDP growth under the baseline simulation, which has no time variation in technology growth, and the alternative scenario in which technology growth varies according to an estimate of its historical and projected trend. The figure also shows the growth rate of the trend in realized real GDP after applying the Hodrick–Prescott filter with a Lagrange multiplier of 25,000.

Figure 11: Dynamic simulation under baseline demographics: other variables

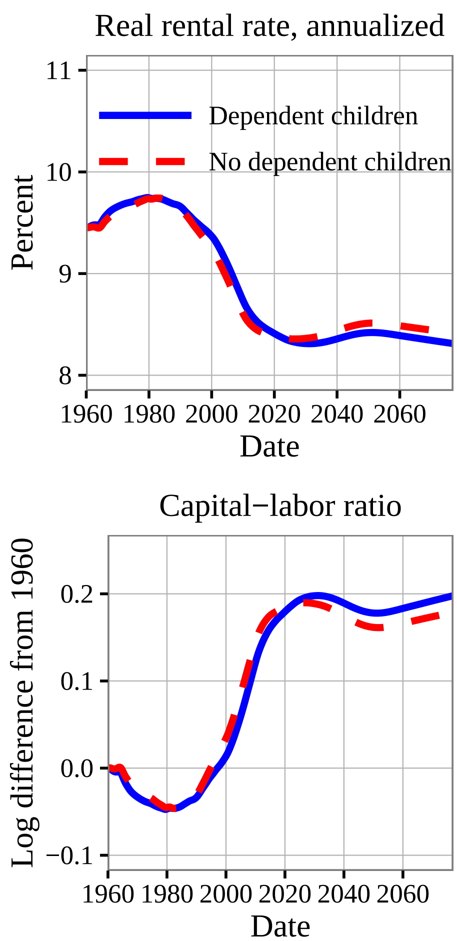

Source: Authors' model simulations.
Notes: The lines labeled "Dependent children" and "No dependent children" show simulation results under the assumption that representative adults receive utility from their children's consumption and receive no utility through an additive term, respectively. The natural logarithms of the real wage and the capital-labor ratio are normalized so that their levels are zero in 1960.

Figure 12: Equilibrium real rate with demographic variables fixed in either 1960 or 1980

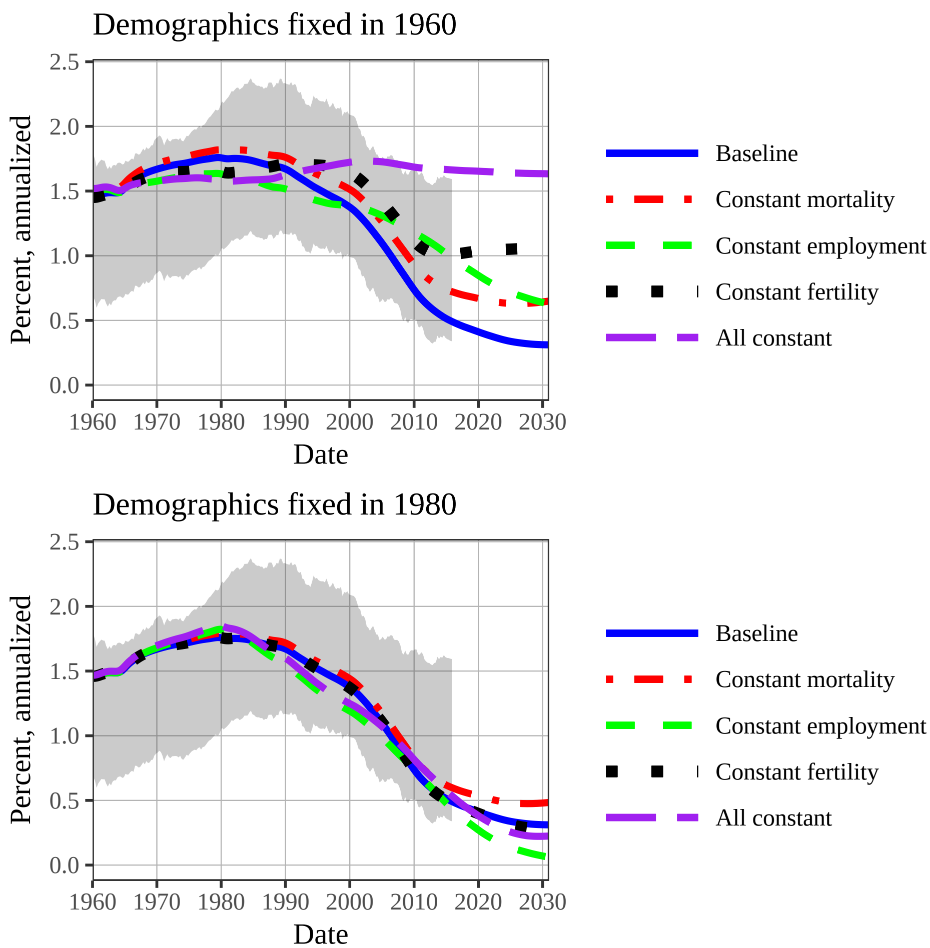

Source: Johannsen and Mertens (2016b) and authors' model simulations.
Notes: The top panel shows dynamic simulations of the model in which we hold constant, first separately and then jointly, the mortality rates, fertility rates, and employment-population ratios by age from 1960 onward. The bottom panel shows a similar set of alternative scenarios in which the demographic variables are held constant from 1980 onward. In all cases, all other parameters are held at their baseline values. The results are shown against the 50-percent uncertainty bands from Johannsen and Mertens (2016b).

Figure 13: Real GDP growth rate with demographic variables fixed in either 1960 or 1980

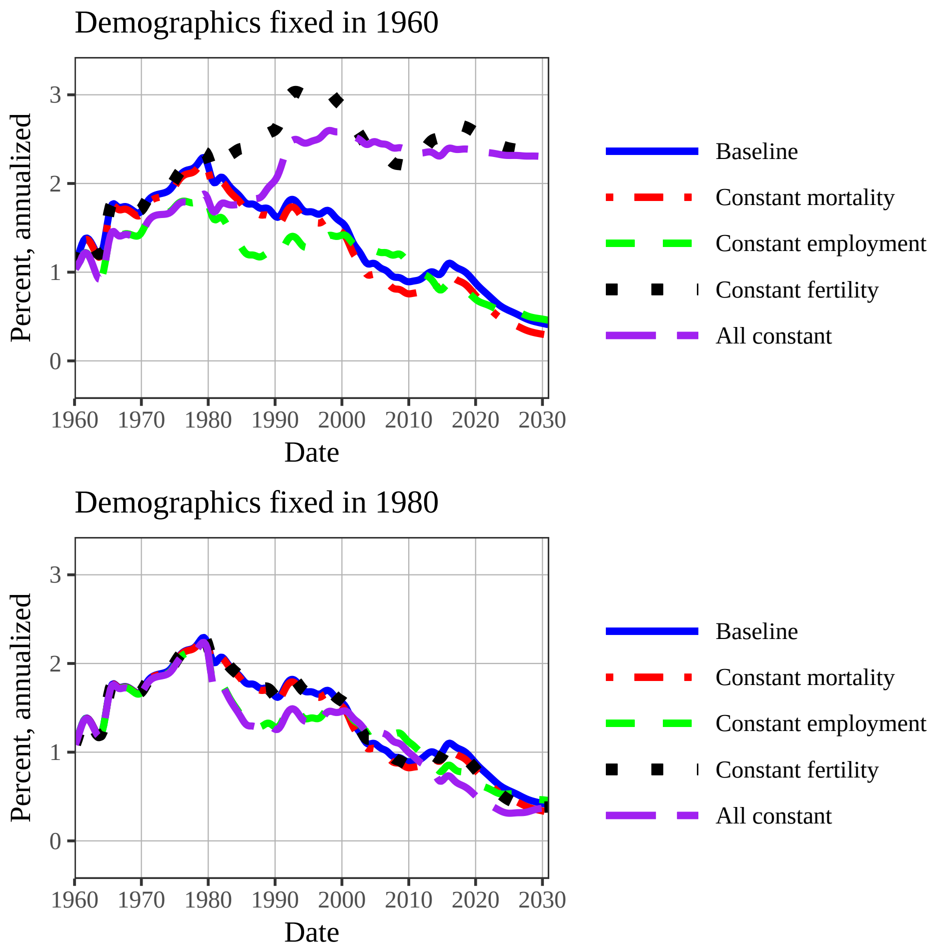

Source: Authors' model simulations.
Notes: The top panel shows dynamic simulations of the model in which we hold constant, first separately and then jointly, the mortality rates, fertility rates, and employment-population ratios by age from 1960 onward. The bottom panel shows a similar set of alternative scenarios in which the demographic variables are held constant from 1980 onward. In all cases, all other parameters are held at their baseline values.

Figure 14: Real interest rate, alternative scenarios

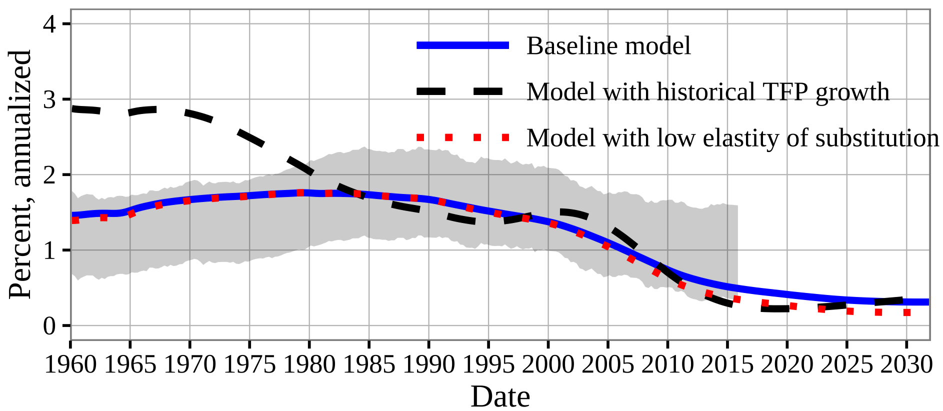

Source: Authors' model simulations.
Notes: The figure shows the equilibrium real rate under the baseline simulation, which has no time variation in technology growth, and the alternative scenario in which technology growth varies according to an estimate of its historical and projected trend. The figure also shows an alternative scenario in which the elasticity of substitution between capital and labor in the production function is set to 0.75, as opposed to 1 under the baseline simulation. The results are shown against the 50-percent uncertainty bands from Johannsen and Mertens (2016b).

Figure 15: U.S. versus global demographic changes

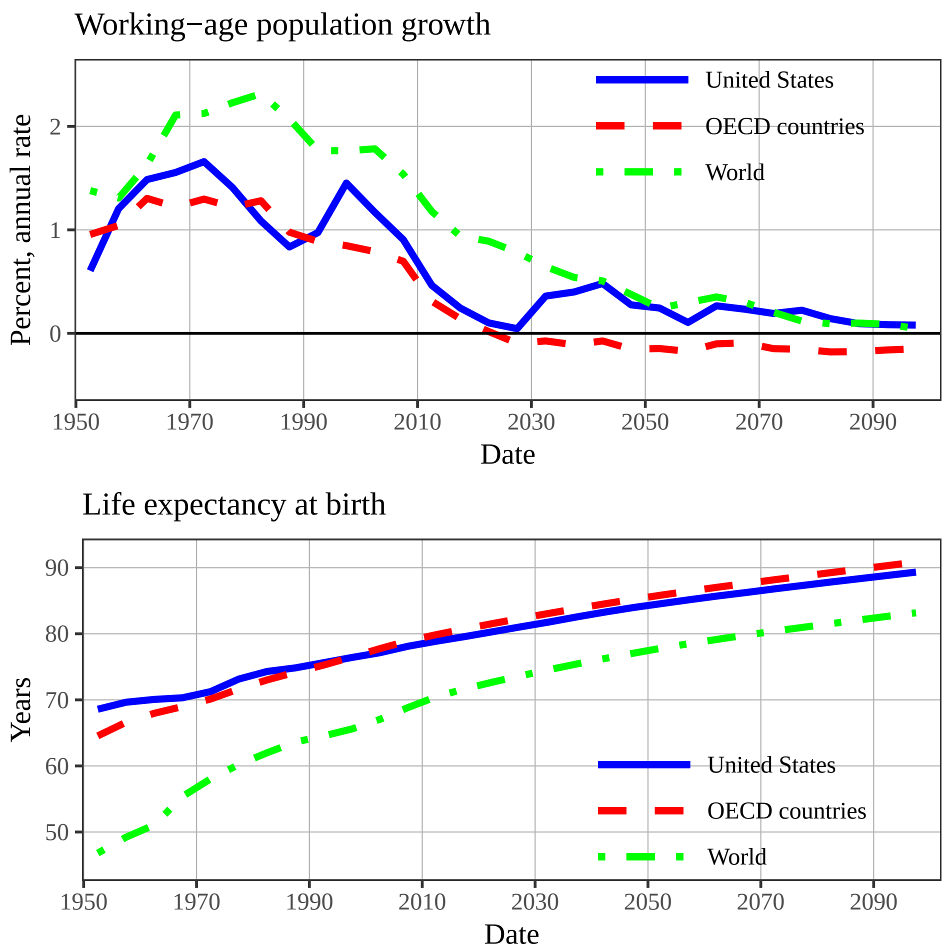

Source: Authors' calculations using data from the *World Population Prospects: Revision 2015* published by the United Nations.
Notes: The working-age population is defined as the set of individuals aged 15 to 64 years. We construct the OECD measure of mean life expectancy at birth as a population-weighted average of that statistic across OECD member countries.