Aging, Secular Stagnation and the Business Cycle*

Callum Jones†

September 2021

Abstract

By the end of 2019, US output was 14% below the level predicted by its pre-2008 trend. To understand why, I develop and estimate a model of the US with demographics, real and monetary shocks, and the occasionally binding zero lower bound on nominal rates. Demographic shocks generate slow-moving trends in interest rates, employment, and productivity. Demographics alone can explain about 40% of the gap between log output per capita and its linear trend by 2019. By lowering interest rates, demographic changes caused the zero lower bound to bind after the Great Recession, contributing to the slow recovery.

Keywords: Great Recession, Demographics, Zero Lower Bound, Forward Guidance.

JEL classifications: E3, E4, E5, J11.

*I thank the editor Olivier Coibion and anonymous referees for valuable feedback. I am particularly thankful to Virgiliu Midrigan for advice and guidance. I thank Saki Bigio, Jaroslav Borovicka, Katarina Borovickova, Tom Cooley, Tim Cogley, Mark Gertler, Dan Greenwald, Mariano Kulish, Julian Kozlowski, Steven Pennings, and Thomas Philippon for comments and discussions, and participants of a number of seminars for suggestions. All remaining errors are mine alone. The Online Appendix is available at https://callumjones.github.io. The views expressed are those of the authors and not necessarily those of the Federal Reserve Board or the Federal Reserve System.

†Federal Reserve Board, callum.j.jones@frb.gov

1 Introduction

Three significant trends over recent decades characterize the US macroeconomic landscape. First, average real growth during expansions has slowed markedly, falling from 3.2% between 1991 and 2001, to 2.8% between 2002 and 2007, and to 1.9% between 2009 to 2019. As a result, by 2019Q4 output was 14% lower than what is predicted by its long-run pre-Great Recession trend. Second, real and nominal interest rates have fallen, with the Fed Funds rate at its lower bound between 2009 and 2015. Third, the employment-population ratio has fallen significantly from its peak in the 2000s (Figure 1).

Against this macroeconomic background, the US population has aged. The average age of the population has risen from 28 to 38 in the period from 1970 to 2018. A number of papers have shown that an aging population can explain the trends in growth and interest rates observed. Eggertsson et al. (2019) and Summers (2014) show that an older population can explain why real interest rates are low and therefore why the zero lower bound (ZLB) binds – because savings behavior changes with age, an economy with a higher fraction of older people has more savings and lower interest rates (see also Carvalho et al., 2016; Gagnon et al., 2021). Feyrer (2007) and Aaronson et al. (2014) show that an aging population can rationalize why productivity growth and the employment-population ratio are low, since younger workers face a steeper human capital profile and older cohorts work fewer hours.

In this paper, I embed demographic changes into a New Keynesian model and estimate it using Bayesian methods. With the estimated model, I study how demographic changes impact the drivers of business cycles, and decompose the observed evolution of aggregate variables into demographic and business cycle shocks. By taking this approach, the key contribution of my paper relative to the existing literature is to quantitatively explore the interactions that arise between aging, monetary policy, and business cycle shocks. The existing literature, by contrast, largely studies the impact of demographics by feeding demographic shocks into large-scale overlapping generations models.

Demographic changes generate significant trends in real and nominal variables. I find that demographic shocks alone are responsible for about 40% of the decline in output relative to trend between 2008 and the end of 2019. Demographics also account for about a 1 percentage point decline in the real interest rate and a 2 percentage point decline in the nominal interest rate between 1990 and 2019. I find that the interactions between demographic changes and business cycle shocks are quantitatively important such that, had demographics not changed between 1986 and 2019, the ZLB would have been a much less significant constraint on monetary policy. In this counterfactual, the Federal Reserve would have been able to lower the real interest rate by an additional 1 to 2 percentage points in 2009 and 2010.

Estimating the model with demographic shocks also affects the posterior estimates of the model's structural parameters and the interpretation of which shocks drive business cycles. Accounting for slow-moving demographic trends lowers the importance of productivity shocks and raises the role of markup shocks for driving output and consumption, illustrating how the propagation of shocks can change under demographic trends.[1] Furthermore, the decline in interest rates caused by demographics brings the policy rate closer to the ZLB. In simulations of the estimated model, I find that the frequency of ZLB episodes increases over time. If demographics are fixed at their 1990 level, the ZLB binds for 1% of the time. If instead demographics are fixed at their 2020 level, the ZLB binds for 12% of the time. Forward guidance, or lower-for-longer policy, would extend these lower bound episodes.

The nonlinearities associated with the ZLB and forward guidance policy mean that it is a nontrivial problem to assess which shocks, other than demographics, were responsible for the decline in output following the Great Recession. To shed light on this, I feed the estimated structural shocks into a version of the model that abstracts from the ZLB, and find that the contribution of aggregate shocks that capture financial distress depressed output, but were partly offset by positive government spending shocks between 2009 and 2012. I also find that forward guidance policy was able to mitigate the impact of the ZLB and kept output from falling by a further 2 to 5 percentage points between 2011 and mid-2013.

[1] This relates to studies of how structural changes can affect the impact of shocks, e.g. Kulish and Pagan (2016); Canova et al. (2015); Wong (2015); Jaimovich and Siu (2012); Fernández-Villaverde et al. (2007).

To make it feasible to conduct a full information Bayesian estimation to parameterize the model, I need a fast and efficient solution method. Methodologically, I show how to parsimoniously model anticipated demographic shocks alongside the ZLB and a standard set of business cycle shocks. In particular, I show that my model is well approximated by an aggregate representation with time-varying and anticipated changes to the parameters of the model that are exogenous functions of demographic variables. I exploit the resulting computational advantages to filter quarterly data for the model's structural shocks, accounting for the trends caused by demographic shocks, the ZLB, and forward guidance.[2]

I find that demographic shocks reproduce well the long-run trends observed in the US economy, consistent with the work of Gagnon et al. (2021), Eggertsson et al. (2019), Aksoy et al. (2019), and Antolin-Diaz et al. (2017). Between 1990 and 2015, changes in the composition of the workforce due to an aging population explain a decline of around 2 percentage points in the labor force participation rate, with a further decline of 4 percentage points expected by 2035. This contraction of the labor force and a decline in savings and investment as fewer workers save for retirement lowers growth, consistent with the trend observed.

2 Model

This section outlines a New Keynesian model with demographic and business cycle shocks. The model features individuals of different ages, monopolistically competitive firms that produce with capital and labor and face price adjustment costs, aggregate shocks, monetary policy with nominal rates subject to the ZLB, and fiscal policy.

2.1 Households

Demographics Individuals are members of overlapping cohorts. Each cohort lives for a maximum of T periods, so that the age range of an individual is 0 to $T-1$. A cohort is of

[2] The implementation of the ZLB and forward guidance follows Kulish et al. (2017), Guerrieri and Iacoviello (2015), and Jones (2017).

mass n_t^s and comprised of a continuum of identical members of age s, measured at the start of period t. The total size of the population at time t is:

$$n_t = \sum_{s=0}^{T-1} n_t^s. \tag{1}$$

I abstract from trend population growth, and normalize the initial population size to 1.

Each period, a fraction of each cohort dies with exogenous age-specific mortality rate γ_t^s:

$$n_{t+1}^{s+1} = (1 - \gamma_t^s) n_t^s. \tag{2}$$

These mortality rates are time-varying. For example, permanent decreases in mortality rates imply increases in longevity. A maximum lifespan of T implies $\gamma_t^{T-1} = 1$.

Household problem An individual of age s has the period utility function $u(c_t^s, \ell_t^s)$ and chooses consumption c_t^s, labor supply ℓ_t^s, one-period risk-free bonds b_t^s, and loans to capital producing firms $b_{k,t}^s$, to maximize lifetime utility. The value function of an individual of age s at period t is, in recursive form:

$$V_t^s = \max_{\left\{c_t^s,\ \ell_t^s,\ b_t^s,\ b_{k,t}^s\right\}} \left\{ \chi_t u(c_t^s, \ell_t^s) + \beta(1 - \gamma_t^s) \mathbb{E}_t V_{t+1}^{s+1} \right\}, \tag{3}$$

where the expectation is taken with respect to the aggregate shocks, β is the discount factor, and χ_t is an aggregate autoregressive process subject to iid normal shocks:

$$\ln \chi_t = (1 - \rho_\chi) \ln \chi + \rho_\chi \ln \chi_{t-1} + \sigma_\chi \varepsilon_{\chi,t}. \tag{4}$$

The unintentional bequests made by those who die between periods are aggregated and redistributed to the remaining living households of the same cohort. Individuals have age-specific productivities z^s, receive a transfer from the government ξ_t^s, earn a return R_t on last period's bond holdings b_{t-1}^{s-1}, earn the return $R_{k,t}$ on last period's loans to capital-producers

$b_{k,t-1}^{s-1}$, receive $\frac{1}{n_t^s}\frac{d_t}{p_t}$ dividends from firms, and ψ_t^s for the redistributed unintentional bequest.[3] These features imply that the period budget constraint of an individual of age s is:

$$c_t^s + \frac{b_{k,t}^s}{p_t R_{k,t}} + \frac{b_t^s}{p_t R_t} \leq z^s w_t \ell_t^s (1-\tau_t^w) + \frac{b_{t-1}^{s-1}}{p_t} + \frac{b_{k,t-1}^{s-1}}{p_t} + \tau_t^s. \tag{5}$$

where $\tau_t^s = \xi_t^s + \psi_t^s + \frac{1}{n_t^s}\frac{d_t}{p_t} - T_t^g$ collects various transfers and a lump-sum tax T_t^g, w_t is the nominal wage rate, τ_t^w is a labor income tax, and p_t is the price level. In the last period of life, the budget constraint is:

$$c_t^T \leq \frac{b_{k,t-1}^{T-1}}{p_t} + \frac{b_{t-1}^{T-1}}{p_t} + \tau_t^T. \tag{6}$$

By assumption, an individual retires fully from the labor market in her last period of life. Individuals are born with zero wealth, so that $b_{k,t}^0 = 0$ and $b_t^0 = 0$ for all t, and nominal bonds are in net zero supply $b_t = \sum_s n_t^s b_t^s = 0$. Substituting in for the unintentional bequests $\psi_t^s = \frac{n_{t-1}^{s-1}}{n_t^s}\left[\frac{b_{t-1}^{s-1}}{p_t} + \frac{b_{k,t-1}^{s-1}}{p_t}\right]$, and denoting the marginal utility of wealth of an individual of age s in time t by λ_t^s, the optimal choice of risk-free bonds implies a standard Euler equation:

$$\mathbb{E}_t\left[\frac{\lambda_{t+1}^{s+1}}{\lambda_t^s}\frac{1}{\Pi_{t+1}}\right] = \frac{1}{\beta}\frac{1}{R_t}, \tag{7}$$

where $\Pi_t = p_t/p_{t-1}$ is the rate of inflation.

2.2 Firms

There are two types of firms in the economy, intermediate goods-producing firms, and capital goods-producing firms. A continuum of intermediate goods-producing firms hire capital and labor from households to supply a substitutable good $y_t(i)$ at price $p_t(i)$ to final goods producers. Final goods producers, in turn, use a CES technology with elasticity of substitution ξ_t to aggregate the intermediate goods into a final good which it sells to consumers at the

[3] The redistribution scheme for unintentional bequests scales the return on savings by $1/(1-\gamma_{t-1}^{s-1})$.

price p_t. Capital k_{t-1} is hired from capital producers, while aggregate labor hired by the firm is in efficiency units of labor $\ell_t = \sum_s z^s n_t^s \ell_t^s$. The production function of firm i is:

$$y_t(i) = \mu_t^{1-\alpha} \left(k_{t-1}(i)\right)^{\alpha} \left(Z_t \ell_t(i)\right)^{1-\alpha}, \tag{8}$$

where $\frac{Z_t}{Z_{t-1}} = z$ generates trend growth, and μ_t is an aggregate autoregressive TFP process:

$$\ln \mu_t = (1-\rho_\mu) \ln \mu + \rho_\mu \ln \mu_{t-1} + \sigma_\mu \varepsilon_{\mu,t}, \tag{9}$$

where $\varepsilon_{\mu,t}$ is a standard normal innovation. Denoting $\mathrm{mc}_t(i)$ as the Lagrange multiplier on the firm's cost minimization problem $\min_{\{k_{t-1}(i),\ \ell_t(i)\}} r_t k_{t-1}(i) + w_t \ell_t(i)$, the rental rate on capital is $r_t = \alpha \mathrm{mc}_t(i) \frac{y_t(i)}{k_{t-1}(i)}$, and the wage is $w_t = (1-\alpha)\mathrm{mc}_t(i) \frac{y_t(i)}{\ell_t(i)}$. Intermediate goods-producing firms also face a Rotemberg quadratic cost of adjusting prices, parameterized by ϕ_p. The problem of the firm i is to choose its price $p_t(i)$ to maximize firm value:

$$\max_{p_t(i)} \mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \lambda_t \left(\frac{d_t(i)}{p_t}\right), \tag{10}$$

where λ_t is a weighted average of individuals' Lagrange multipliers and real dividends are:

$$\frac{d_t(i)}{p_t} = \left(\frac{p_t(i)}{p_t}\right)^{1-\xi_t} y_t - \mathrm{mc}_t(i) \left(\frac{p_t(i)}{p_t}\right)^{-\xi_t} y_t - \frac{\phi_p}{2} \left[\frac{1}{\Pi^*} \frac{p_t(i)}{p_{t-1}(i)} - 1\right]^2 y_t, \tag{11}$$

where Π^* is the inflation target and ξ_t is the elasticity of substitution between intermediate goods which is subject to stochastic shocks:

$$\ln \xi_t = (1-\rho_\xi) \ln \xi + \rho_\xi \ln \xi_{t-1} + \sigma_\xi \varepsilon_{\xi,t}. \tag{12}$$

The first order condition for the optimal choice of price resetting is:

$$\beta \phi_p \mathbb{E}_t \tfrac{\lambda_{t+1}}{\lambda_t} \tfrac{y_{t+1}}{y_t} \left[\tfrac{\Pi_{t+1}}{\Pi^*} - 1\right] \left[\tfrac{\Pi_{t+1}}{\Pi^*}\right] = \xi_t - 1 - \xi_t \mathrm{mc}_t + \phi_p \left[\tfrac{\Pi_t}{\Pi^*} - 1\right] \left[\tfrac{\Pi_t}{\Pi^*}\right], \tag{13}$$

Log-linearizing (13) yields a standard forward-looking New Keynesian Phillips curve. I denote the slope of the log-linearized Phillips curve by ϵ_p, which is a function of the steady-state elasticity of substitution ξ and the price adjustment cost parameter ϕ_p.

The problem of the capital-producing firms is standard: they borrow from households and accumulate capital to maximize market value. Capital adjustment incurs a quadratic cost parameterized by ϕ_k. These adjustment costs are subject to an aggregate, exogenous autoregressive process κ_t, which captures changes in the efficiency of investment adjustment:

$$\ln \kappa_t = (1 - \rho_\kappa) \ln \kappa + \rho_\kappa \ln \kappa_{t-1} + \sigma_\kappa \varepsilon_{\kappa,t}. \tag{14}$$

This problem gives rise to a Tobin's Q equation, linking the investment of capital to its shadow price and its relative return.

2.3 Monetary Policy

Monetary policy operates in one of two possible regimes. In the first regime, the nominal interest rate is set according to a Taylor rule, as in Smets and Wouters (2007):

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R}\right)^{\phi_r} \left(\frac{\Pi_t}{\Pi^*}\right)^{(1-\phi_r)\phi_\pi} \left(\frac{y_t}{y_t^{\mathrm{F}}}\right)^{(1-\phi_r)\phi_y} \left(\frac{y_t/y_{t-1}}{y_t^{\mathrm{F}}/y_{t-1}^{\mathrm{F}}}\right)^{\phi_g} \exp(\sigma_R \varepsilon_{R,t}). \tag{15}$$

The nominal interest rate responds to its own lag with weight ϕ_r, deviations in inflation from a target rate Π^* with weight ϕ_π, deviations in output from its flexible-price level y_t^F with weight ϕ_y, and the growth rate of output relative to the growth rate of potential output with weight ϕ_g, and is subject to stochastic shocks $\varepsilon_{R,t}$.

In the second regime, the nominal interest rate is at the ZLB:

$$\log(R_t) = 0. \tag{16}$$

Monetary policy can be in the ZLB regime in two ways: first, if the Taylor rule calls for negative nominal interest rates – that is, $\log(R_t) = \max(0, \text{Taylor Rule}_t)$ – and second, if the Fed has announced, or has previously announced, an extension of the ZLB beyond that implied by the constraint and the Taylor rule. I assume that the Fed can manipulate expectations of how the path of interest rates evolves when it is at zero, as in Eggertsson and Woodford (2003) and Werning (2012). In estimation, I use survey data from 2009 to 2015 to discipline the expected durations of the zero interest rate regime.

2.4 Government

The government taxes labor income at the rate τ_t^w to fund a pay-as-you-go social security system. The transfer paid to individuals above eligibility age T^* depends on the accumulated pre-tax labor income of the worker, and a parameter ω governing the replacement rate of past earnings. Denote by W_t^s accumulated gross lifetime earnings, defined recursively as:

$$W_t^s = \begin{cases} w_t z^s \ell_t^s + W_{t-1}^{s-1}, & \text{if } s < T^* \\ W_{t-1}^{s-1}, & \text{if } s \geq T^*. \end{cases} \tag{17}$$

The amount ξ_t^s redistributed to an agent of age $s \geq T^*$ depends on W_t^s:

$$\xi_t^s = \omega \frac{W_t^s}{(T^* - 1)}, \tag{18}$$

where the denominator reflects the amount of time that W_t^s is accumulated over. For those younger than the eligibility age T^*, the transfer $\xi_t^s = 0$.[4] The budget constraint of the social security system is:

$$\sum_s n_t^s \xi_t^s = \sum_s n_t^s z^s w_t \ell_t^s \tau_t^w. \tag{19}$$

The tax rate τ_t^w adjusts to equalize social security outlays and tax revenues.

Finally, the government levies a lump-sum tax on households to pay for government

[4] I abstract from issues of the sustainability of pension systems in an aging society and do not allow pension funds to accumulate assets or liabilities (see Attanasio et al., 2007, for an analysis of these issues).

expenditures g_t, which are assumed to be autoregressive and subject to stochastic shocks:

$$\ln g_t = (1-\rho_g)\ln g + \rho_g \ln g_{t-1} + \sigma_g \varepsilon_{g,t}, \tag{20}$$

where the budget constraint of exogenous government expenditures is $g_t = n_t T_t^g$.

3 Model Approximation and Solution Method

To address the computational challenges arising from this model's rich sources of heterogeneity, persistence and aggregate shocks, I argue in this section that the model can be approximated very well by a representative agent framework with time-varying parameters that are functions of exogenous demographic variables. Furthermore, because demographic changes are assumed to be perfectly foreseen, the path of these time-varying parameters are also assumed to be fully anticipated. I show how this anticipated path of time-varying parameters gives rise to a VAR representation that makes a Bayesian estimation feasible.

3.1 Derivation

The approximation is derived in two steps. First, the restriction that individuals can trade assets only when alive is relaxed. In the second step, we will show that when this timing assumption is relaxed, the model has a tractable aggregate representation.

Timing Assumption In the model's overlapping generations setup, individuals are born with no wealth and make period-by-period asset trades. Assume, instead, that each generation is alive at $t=0$ and can trade claims to future consumption, and write the preferences of an individual of age s as:

$$\sum_{t=0}^{\infty} \beta^t \left[\prod_{j=0}^{t} (1-\gamma_{t+j-1}^{s+j-1})\right] \phi_t^{s+t} \sum_{\sigma^t} \Pr\left[\sigma^t|\sigma^{t-1}\right] u\left[c_t^{s+t}(\sigma^t), \ell_t^{s+t}(\sigma^t)\right], \tag{21}$$

where the term $\phi_t^{\bar{s}} = 1$ for ages $0 \leq \bar{s} \leq T-1$, and $\phi_t^{\bar{s}} = 0$ otherwise, indicating that individuals value utility only in the periods when they are between the ages of 0 and $T-1$. The term $\Pr\left[\sigma^t|\sigma^{t-1}\right]$ denotes the transition probability from state σ^{t-1} to σ^t.

A key assumption for the approximation is that unintentional bequests by individuals who die in each period are redistributed to members of the same generation. This means that individuals are insured against their only source of idiosyncratic uncertainty: that associated with mortality risk γ_t^s. The Euler equation arising from the choice of savings in the problem where individual's maximize (21) subject to the lifetime budget constraint is then:

$$\lambda_t^s(\sigma^t) = \beta \sum_{\sigma^{t+1}} \Pr\left[\sigma^{t+1}|\sigma^t\right] \lambda_{t+1}^{s+1}(\sigma^{t+1}) \frac{R_t(\sigma^t)}{\Pi_{t+1}(\sigma^{t+1})}, \tag{22}$$

where $\lambda_t^s(\sigma^t)$ is the Lagrange multiplier on the individual's budget constraint. Without idiosyncratic risk, it is possible to separate each individual's marginal utility of wealth into a component that does not depend on the aggregate state, and a component that does. Because of this, between any two individuals, across two periods t and t', the ratio of the Lagrange multipliers is constant:

$$\frac{\lambda_t^s(\sigma^t)}{\lambda_t^{s'}(\sigma^t)} = \frac{\lambda_{t'}^{s+t'}(\sigma^{t'})}{\lambda_{t'}^{s'+t'}(\sigma^{t'})} = \frac{\lambda^{s'}}{\lambda^s}, \tag{23}$$

where $\lambda^s = \frac{\lambda_t(\sigma^t)}{\lambda_t^s(\sigma^t)}$. The condition in Equation (23) is the same as that which arises when there are complete asset markets.

Aggregate Representation Under Equation (23), the economy's equilibrium can be found by solving the problem of a social planner that maximizes a weighted sum of individuals' utility functions. In the social planner's problem, the planner first determines how to allocate, within periods, aggregate consumption and aggregate labor supply between individuals. Given the optimal allocation, the planner then solves its intertemporal problem and maximizes aggregate consumption, capital, and labor supply subject to the economy's resource constraint. This approach is based on the aggregation arguments made in Constantinides (1982) and Maliar and Maliar (2003).

Assume that an individual of age s has a separable period utility function over consumption c_t^s and hours ℓ_t^s of the type $\frac{(c_t^s)^{1-\sigma}}{1-\sigma} - \upsilon^s \frac{(\ell_t^s)^{1+\varphi}}{1+\varphi}$. Under these preferences, as shown in the Online Appendix, there is a representative agent with preferences over aggregate consumption c_t and aggregate units of labor ℓ_t that take the form:

$$U(c_t, \ell_t) = \phi_t \frac{c_t^{1-\sigma}}{1-\sigma} - \upsilon_t \frac{\ell_t^{1+\varphi}}{1+\varphi}. \tag{24}$$

The representative agent's problem is to maximize (24) over time by choosing c_t, ℓ_t, and aggregate capital k_t subject to the economy's resource constraint and its production function $y_t = \theta_t^{1-\alpha} k_t^{\alpha} \ell_t^{1-\alpha}$. The relationship between units of labor and aggregate hours is $\ell_t = A_t h_t$.

In the Online Appendix, I show that θ_t and A_t are the following time-varying parameters:

$$\theta_t = \sum_s n_t^s z^s, \qquad \text{and} \qquad A_t = \frac{\sum_s n_t^s (\hat{z}^s)^{1+1/\varphi} (\upsilon^s \lambda^s)^{-1/\varphi}}{\sum_s n_t^s (\hat{z}^s)^{1/\varphi} (\upsilon^s \lambda^s)^{-1/\varphi}}, \tag{25}$$

where the value $\hat{z}^s = z^s/\theta_t$ denotes individual s's relative productivity and the λ^s parameters are the Pareto weights attached to an individual of age s. The shock θ_t encodes changes in output caused by the size of the workforce and its distribution over productivity levels. For example, populations that have a larger fraction of more productive workers have relatively higher aggregate productivity and θ_t. The time-varying parameter A_t affects the hours needed to obtain an effective unit of labor, and is a population-weighted average of relative productivity and the disutility of providing labor. If labor supply were inelastic, θ_t and A_t affect the labor input by $\frac{\sum_s n_t^s z^s}{\sum_s n_t^s}$: the labor input reflects only the population composition.

The term ϕ_t inversely affects the marginal utility of consumption, and has a simple expression mapping to the size of the population at each point in time:

$$\phi_t = \left[\sum_s n_t^s \left(\lambda^s\right)^{\frac{1}{\sigma}} \right]^{\sigma}. \tag{26}$$

The term v_t is a time-varying parameter affecting the marginal disutility of labor:

$$v_t = \left[\sum_s n_t^s (\hat{z}^s)^{\frac{1}{\varphi}+1} (v^s \lambda^s)^{-\frac{1}{\varphi}} \right]^{-\varphi}, \tag{27}$$

so that v_t is a population-weighted average of age-specific labor disutilities. The greater the relative size of the population with high disutilities of providing labor, the higher is v_t. Equating the marginal utility of consumption and the marginal disutility of labor, and substituting in for hours worked gives the labor wedge as a function of demographics ${}^{w_t}/\left({}^{\ell_t^{\varphi}}/{}_{c_t^{-\sigma}}\right) = {}^{v_t}/{}_{\phi_t}$.

Two additional trends are needed in the computations to ensure the aggregate representation will approximate the aggregate dynamics of the full lifecycle solution. The first is a gradual trend in the discount factor to account for the reduction in the average mortality rate over time. This term is constructed by mapping the change in average life expectancy to the change in the average mortality rate, and then multiplying the discount factor by the change in that average mortality.[5] The second trend is to proportional taxes that are used to finance the social security system. I take the path of labor income taxes from the non-stochastic perfect foresight path of the overlapping generations setup.

In the approximation of the model, demographics therefore affects the aggregate economy through time-varying parameters which are functions of observable population dynamics (n_t^s), the age-specific parameters of the model (z^s and v^s), and the Pareto weights that the planner attaches to each generation (λ^s). Assuming the planner equally weights each generation, these trends are straightforward to compute and do not depend on endogenous variables. In the Online Appendix, I verify that the aggregate approximation recovers closely the paths of the aggregate variables due to demographic shocks by comparing them to the paths of the decentralized lifecycle model under perfect foresight and demographic shocks.

[5] The contribution of this term to the decline in real interest rates over the estimation period is relatively small. The rise in average life expectancy maps into a very slight increase in the discount factor by 0.05% between 1986Q1 and 2019Q4 (the sample period for the Bayesian estimation). Of the 1.2 percentage point decline in the real interest rate caused by demographics, this term by itself causes the annual real interest rate to fall by only about 19 basis points.

I also show in the Online Appendix that, in a second-order approximation of a calibrated version of the full lifecycle model subject to stochastic aggregate shocks, the decision rules across individuals are approximately linear. The approach therefore shares the intuition that makes the Krusell and Smith (1998) algorithm successful: when the decision rules are close to linear in the state variables, we have a near-Gorman-aggregation setup. In this case, for any aggregate wealth distribution, the Engel curves have a similar slope for each individual, and it does not matter for the response of aggregate consumption how any additional income is allocated across agents. These results are consistent with the findings of Ríos-Rull (1996), who reports that the business cycle properties of large-scale, stochastic, overlapping-generations economies are similar to the properties of representative agent real business cycle models. The contribution in this paper is to additionally describe an approximation that makes likelihood estimation computationally feasible.

3.2 Solution Method

In this section, I describe the methodology used to solve for the path of the aggregated model under the anticipated path of time-varying demographic parameters, and describe how the methodology implements the occasionally binding ZLB.

Time-Varying Demographic Trends Let x_t be the vector of model variables, and ε_t a vector that collects the exogenous unanticipated shocks. The linearized rational-expectations approximation of the model with time-varying parameters is:

$$\mathbf{A}_t x_t = \mathbf{C}_t + \mathbf{B}_t x_{t-1} + \mathbf{D}_t \mathbb{E}_t x_{t+1} + \mathbf{F}_t \varepsilon_t, \tag{28}$$

where $\mathbf{A}_t$, $\mathbf{B}_t$, $\mathbf{C}_t$, $\mathbf{D}_t$, and $\mathbf{F}_t$ are time-varying matrices that encode the structural equations of the model linearized at each point in time around the steady-state corresponding to the time t structural parameters.[6] A solution to the problem with anticipated time-varying parameters exists if agents expect the structural matrices to be fixed at some point in the future at values which are consistent with a time-invariant equilibrium (Kulish and Pagan, 2016). In this case, the solution has a time-varying VAR representation:

$$x_t = \mathbf{J}_t + \mathbf{Q}_t x_{t-1} + \mathbf{G}_t \varepsilon_t, \tag{29}$$

where $\mathbf{J}_t$, $\mathbf{Q}_t$, and $\mathbf{G}_t$ are conformable matrices which are functions of the evolution of beliefs about the time-varying structural matrices $\mathbf{A}_t$, $\mathbf{B}_t$, $\mathbf{C}_t$, $\mathbf{D}_t$, and $\mathbf{F}_t$

$$\begin{aligned} \mathbf{Q}_t &= \left[\mathbf{A}_t - \mathbf{D}_t \mathbf{Q}_{t+1}\right]^{-1} \mathbf{B}_t \\ \mathbf{J}_t &= \left[\mathbf{A}_t - \mathbf{D}_t \mathbf{Q}_{t+1}\right]^{-1} \left(\mathbf{C}_t + \mathbf{D}_t \mathbf{J}_{t+1}\right) \\ \mathbf{G}_t &= \left[\mathbf{A}_t - \mathbf{D}_t \mathbf{Q}_{t+1}\right]^{-1} \mathbf{E}_t. \end{aligned} \tag{30}$$

This iteration is obtained by noting that, from (29), $\mathbb{E}_t x_{t+1} = \mathbf{J}_{t+1} + \mathbf{Q}_{t+1} x_t$, which is substituted into (28) and rearranged for x_t. The law of motion for the model's variables at a time period t therefore depends on the full anticipated path of the structural matrices. The final structure of the economy needed for the iteration (30) is the one that arises at the expected end of the demographic transition and under Taylor-rule policy, that is, the solution matrices $\{\mathbf{J}_T, \mathbf{Q}_T\}$ associated with the structural equations $\{\mathbf{A}_T, \mathbf{B}_T, \mathbf{C}_T, \mathbf{D}_T, \mathbf{F}_T\}$ in the final period T. Under my calibration, this final demographic structure applies from the year 2060 onwards.[7] Demographic trends in the model are thus perfectly foreseen from 1960 through to 2060. This assumption helps to make the Markov Chain Monte Carlo-based estimation computationally feasible, as standard methods to compute the solution

[6] One can instead linearize the model around its original steady-state, the steady-state associated with the time-varying system's final structure, or the steady-state implied by the structure at each point in time. Given the somewhat large movements in the steady-state induced by demographic changes, I use the latter approach, linearizing each set of structural matrices around the steady-state implied by that structure.

[7] This approach resembles that of Fernández-Villaverde et al. (2007) but instead of innovations driving parameter drift, the time-varying parameters are perfectly foreseen functions of exogenous demographics.

matrices $\{\mathbf{J}_T, \mathbf{Q}_T\}$ need to be used only once. The sequence $\{\mathbf{J}_t, \mathbf{G}_t, \mathbf{Q}_t\}$ obtained from (30) is computationally fast. In the Robustness section, I report how the trends of the model are similar if demographic changes were instead unexpected every period. In this case, to form the time-varying VAR solution, we solve for the solution associated with (28) in every period and use the sequence of those solutions to form (29) instead of the iteration (30).

Zero Lower Bound To implement the occasionally-binding ZLB in the solution (29), we follow Guerrieri and Iacoviello (2015) and Jones (2017) and define two regimes in (28) for each period, one for when the ZLB does not bind, and one for when the ZLB binds. If the ZLB binds, we assume that agents believe no shocks will occur in the future and iterate backwards through our model's equilibrium conditions from the date that the ZLB is conjectured to stop binding. We then iterate on the periods that the interest rate is conjectured to be in effect until it converges, after which the solution is (29).[8]

4 Estimation

This section discusses how the parameters of the model are set, including the calibration of the lifecycle parameters of the model, the demographic shocks which drive the trends in the model, and the estimation of the shocks that govern the business cycle.

4.1 Assigned and Calibrated Parameters

Before estimation, a subset of the parameters and the demographic shocks are calibrated.

[8]See also Canova et al. (2015) and Kulish and Pagan (2016). Jones et al. (2021) discusses how forward guidance is an extension of the ZLB regime beyond the duration implied by the structural shocks and the constraint $\log(R_t) = \max(0, \text{Taylor Rule}_t)$.

4.1.1 Lifecycle Parameters

The model is quarterly. Individuals begin life at 16 years of age and live for at most 80 more years, up to age 95. Full retirement is only imposed in the last period of life.[9]

I calibrate the disutility of providing labor v^s with a scaled cumulative density function of a normal distribution, so that v^s increasing in s (see Kulish et al., 2010). This specification is motivated by studies which link the disutility of work to deteriorating health. The parameters of the function for v^s are chosen so that the labor force participation rates by age broadly match those observed in 2000. For the social security system, I set the replacement ratio of accumulated earnings λ to 46.7%, the same value that is used in Attanasio et al. (2007). Retirement benefits are received from age 65 on ($T^* = 49$).

I calibrate the age-productivity parameters z^s to the age-experience earnings profile. I follow Elsby and Shapiro (2012) in constructing the log experience-earnings profile using deflated data on full-time, full-year workers. The data is decennial Census data from 1960 to 2000, and annual American Community Survey data from 2001 to 2007.[10] To minimize cohort effects, I pool, across years, high school dropouts, high school graduates, those with some college education, and those who have completed college or higher education.[11] Panel A of Figure 2 plots the earnings-profile over age. The estimates imply a peak increase in earnings of about 134% at age 45, before gradually declining around the age of 50: in line with the estimate of Guvenen et al. (2015) who find an increase in the earnings of the mean worker of 127%. There is less reliable data on the earnings of older workers, so after age 65, I calibrate the productivity of workers to decay by 20% a year.

I assess the calibration of the lifecycle parameters in Figure 2 by plotting, in Panel B, the labor force participation rate by age, in 2000, in the model and the data, and in Panel C,

[9] Given low labor supply at older ages, this choice is not too important.

[10] Computed off IPUMS-USA extracts. A full description is given in the Online Appendix.

[11] In robustness exercises reported in the Online Appendix, I distinguish between education groups and analyze how anticipated changes in the earnings profile map to labor supply decisions. The results indicate that the patterns of the aggregate variables are similar, suggesting that age-compositional changes in the population primarily drive aggregate dynamics.

Figure 1: Output, Interest Rates, and the Employment-to-Population Ratio

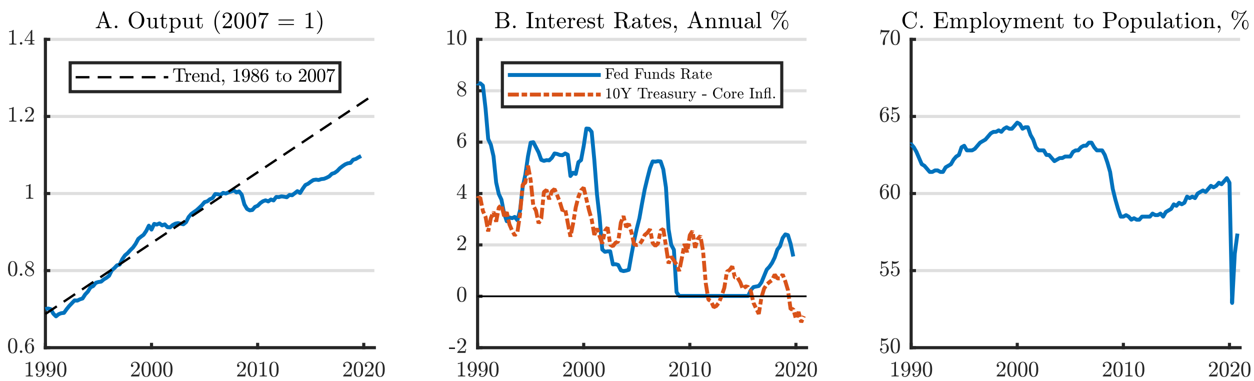

Figure 2: Parameterization of Lifecycle Parameters and Implications

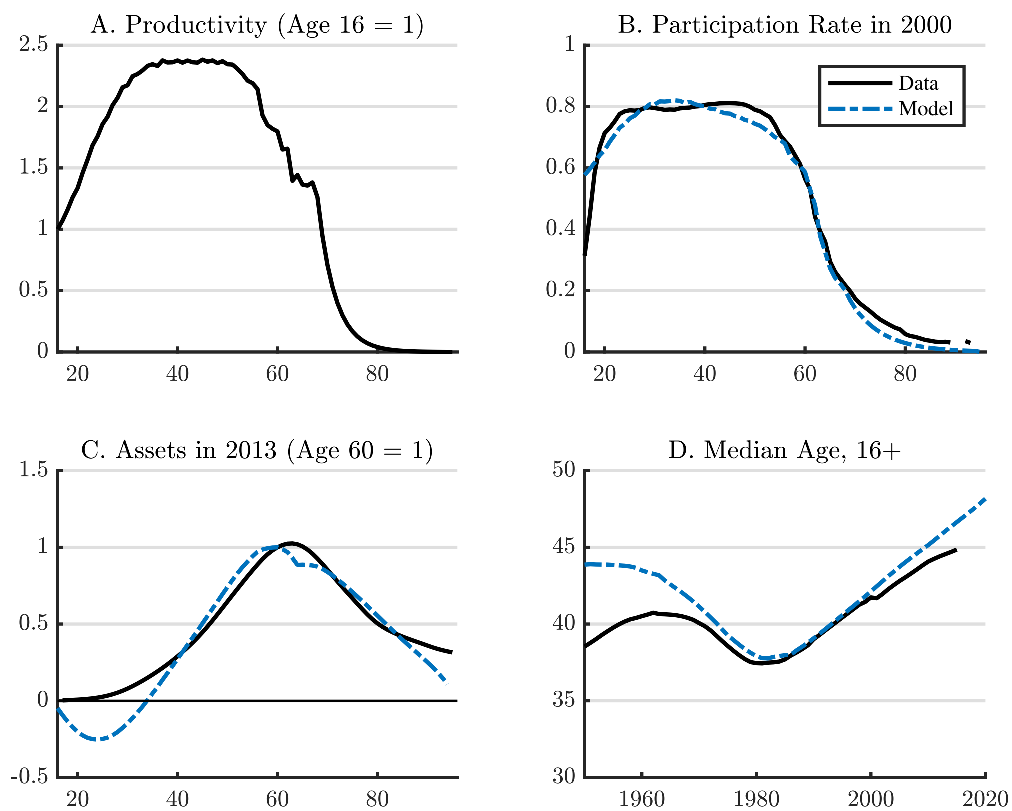

Notes: The data for the asset profile is taken from the Survey of Consumer Finances for 2013, HP-filtered, and normalized to the values of assets when 60. The productivity profile is computed from pooled Census and American Community Survey datasets. Details are in the Online Appendix.

the age-profile of assets normalized to holdings at age 60, in 2013, in the model and in the Survey of Consumer Finances. The calibration implies labor force participation rates that rise when young, flatten out during an individual's prime working life, and decline rapidly around retirement. The lifecycle asset-profile is hump-shaped and peaks around 60 years. Individuals borrow when young in the model and begin accumulating assets around age 35.

4.1.2 Mortality Profiles

Next, I calibrate the mortality probabilities of each generation during the 80 years they could possibly live, γ_t^s, to the actuarial probabilities reported by the Social Security Administration.[12] By calibrating to these probabilities, I also match changes in the life expectancy of each generation over time, conditional on an individual reaching 16 years of age. The values used are the cohort-specific survival rates computed for the cohort year of birth. These profiles include both observed survival rates of cohorts up to their current age, and extrapolated survival rates based on the Social Security Administration's forecasts of life expectancy. I assume that all changes to these actuarial probabilities are exogenous and perfectly foreseen. For the initial γ_t^s profile, I use the survival probabilities reported for those born in 1900 onwards. For those cohorts born before 1900 but who are alive in 1940, I use extrapolated values of the survival probabilities.[13] Under the calibration, between 1950 and 2020, life expectancy for 16 year olds increases from about 77 years to about 85 years.

4.1.3 Incoming Cohort Size Shocks

I choose anticipated shocks to the size of the incoming cohort so that the change in the observed cohort share is the same as the change in the model cohort share.[14] This ensures

[12] These probabilities were sourced from Table 7 from the Cohort Life Tables for the Social Security Area by Calendar Year, in Actuarial Study No. 120 by Felicitie C. Bell and Michael L. Miller, available at: https://www.ssa.gov/oact/STATS/table4c6.html. A full description is given in the Online Appendix.

[13] Because the survival probabilities are low for those years, the results are robust to alternative specifications and are not important for the model outcomes beyond 1970.

[14] Choosing initial population shocks to matching the *changes* is necessary because the model is initialized at the 1940 steady-state and matching the actual cohort sizes would imply very large and counterfactual initial population shocks.

that the model captures the wave of the baby-boomer generation and changes in the population distribution due to, for example, immigration. I assume that changes to the incoming population beyond 2015 decay to zero, so that the population distribution converges to the steady-state implied by the mortality profile that is constant from 2060.

I plot in Panel D of Figure 2 the median age of the population above 16 years of age implied by the calibrated mortality profiles and the incoming cohort size shocks. The profile tracks well the corresponding median age of those above 16 years of age in the data, declining from around the 1960s to around the 1980s to about 37 years of age, before steadily increasing as the baby-boomer population ages and longevity continues to rise.

4.1.4 Preference and Nominal Parameters

I calibrate the remaining parameters to values which imply steady-state capital-output ratios that align with those in the Bureau of Labor Studies' Multifactor Productivity (BLS-MFP) program (see Fernald, 2015).[15] I set capital to depreciate by $\delta = 8\%$ a year. The capital share α is set to 1/3, the average of the capital share reported by Fernald (2015) over 1948 to 2015. The intertemporal elasticity of substitution σ is set to 2, and the inverse Frisch elasticity of labor supply φ is set to 2, in line with the estimates of Reichling and Whalen (2012) and with the analysis of Rios-Rull et al. (2012). Both values are also in the range considered by Auerbach and Kotlikoff (1987) in a computational overlapping generations model. The quarterly discount factor β is set to $0.9935^{1/4}$. The trend growth rate is set to a value of $z = 1.0034$, the mean of output growth over the sample. Together, these parameters imply a capital-output ratio in 2000 of about 2.7, which is the value of the capital-output ratio in 2000 reported the BLS-MFP, and a real interest rate in 2000 of 1.5 percent, slightly higher than the value reported by Gagnon et al. (2021) in their baseline specification, and within the range reported in the time-series estimates of Johannsen and Mertens (2021).

I calibrate a set of the parameters describing the frictions and the nominal side of the

[15] These observed capital-output ratios vary between 2 and 2.7 over the period 1950 to 2013. A description of the BLS-MFP dataset used is given in the Online Appendix.

economy to values commonly used in the literature. The steady-state value of ξ is set to a standard value of 5, which implies a steady-state markup over marginal costs $\xi/(\xi-1)$ of 25%. The parameter governing the quadratic cost of capital adjustment ϕ_k is set to 40. The slope of the Phillips curve ϵ_p, which maps to the quadratic cost of price adjustment ϕ_p, is set to 0.01. This value translates into a Calvo probability of price adjustment every quarter of about 10%, consistent with the estimates in Del Negro et al. (2015). The annual inflation target Π^* is set to 2%. Finally, as in Jones et al. (2018), I use the values that Justiniano et al. (2011) estimate for the Taylor rule parameters, which at their posterior mode are $\phi_r = 0.86$, $\phi_\pi = 1.71$, $\phi_y = 0.05$, and $\phi_g = 0.21$.

4.2 Bayesian Estimation

I next use Bayesian likelihood techniques to estimate the parameters of the model's shocks that drive business cycle fluctuations around the demographic trend.

4.2.1 Quarterly Data

The solution (29) can be written in state-space form, allowing the use of Bayesian likelihood methods to estimate the remaining parameters of the model. The quarterly data used are:

$$\text{Data} = \left\{ \log\left(\frac{y_t}{y_{t-1}}\right),\ \log\left(\frac{c_t}{c_{t-1}}\right), \log\left(\frac{i_t}{i_{t-1}}\right), \log \Pi_t,\ \log \tilde{R}_t,\ T_t \right\}, \tag{31}$$

over the time period 1986Q1 to 2019Q4. I use, as observables, the growth rate of output per capita, of consumption per capita, of investment per capita, the GDP deflator, and the Fed Funds rate, and follow Smets and Wouters (2007) in constructing these series.[16] The nominal interest rate is removed from the set of observables when the ZLB binds between 2009Q1 and 2015Q3. I implement this with a time-varying observation equation in the state-space representation of the model (see Kulish, Morley and Robinson, 2017). The sequence

[16] The Online Appendix provides more details of the data series used in estimation, and additional details of the econometric approach.

of expected durations of the ZLB, T_t, between 2009Q1 and 2015Q3 are taken from the Blue Chip Financial Forecasts survey from 2009 to 2010 and the New York Federal Reserve's Survey of Primary Dealers from 2011 to 2015.

4.2.2 Parameter Estimates and Variance Decompositions

Table 1 reports moments of the prior and posterior distributions for the parameters governing the shock processes. The priors are diffuse, with uniform priors used for the standard deviations of the shock processes. I use a Markov Chain Monte Carlo algorithm to characterize the parameters' posterior distributions, computing two independent chains of 200,000 draws with the first 100,000 draws discarded as a burn-in. The Online Appendix provides the full details of the estimation and an analysis showing that the two chains converge to the same posterior distributions.

The estimated persistence and size of the shock processes are, by themselves, difficult to interpret. I thus report, in Table 2, the forecast error variance decompositions of the observable variables at the 2-quarter horizon (in Panel A) and at the infinite (unconditional) horizon (in Panel B).[17] These decompositions reveal how important each shock is in driving the observable variables around the demographic trend.

Inflation is largely determined by disturbances to markups at the short and long-horizons. Markup shocks also explain about 21 percent of the variance of output growth at the infinite horizon. As reported in the Online Appendix, markup shocks account for large fractions of the variance of the long-run levels of output (64 percent), consumption (29 percent), and investment (44 percent). The contribution of markup shocks to the level of output is close to the long-run variance decomposition of output caused by the combination of wage and price markup shocks in the estimated model of Smets and Wouters (2007) (about 55 percent), which they argue captures supply-side fluctuations which dominate output fluctuations in the long-run. Household preference shocks explain about 25 percent of the variance of the

[17]The Online Appendix plots the impulse responses of the observable variables to each shock with parameters set to the mode of the posterior distributions.

Fed Funds rate at the short horizon, and 38 percent at the long horizon. As discussed in the next section using counterfactual simulations, these shocks are largely accomodated by monetary policy, but can be highly contractionary when monetary policy is constrained (see also Jones, Midrigan and Philippon, 2018).

Monetary policy shocks explain between 5% and 14% of the growth rates of output, consumption, and investment at the half-year horizon. Consistent with the results in Justiniano et al. (2011), investment shocks are very important for explaining business cycle and long-run movements in all the observable variables, and account for 48% and 72% of the forecast error variance of output and investment at the half-year horizon, with similar fractions at the infinite horizon. These investment shocks capture disruptions in the intermediation between savings and investment and directly affect the shadow value of installed capital – Tobin's Q – in the reduced-form of the model. I show in the Online Appendix that these shocks correlate strongly with the Corporate Baa to 10 year Treasury yield spread, a measure of financial stress used in the literature (Del Negro, Giannoni and Schorfheide, 2015). Consistent with the large literature that has studied and found financial shocks to be important for explaining macroeconomic fluctuations around the Great Recession, I show in Section 5.4 that these shocks can explain much of the sharp drop in output around 2008-09 and its gradual recovery. Finally, technology shocks explain 20 percent of the variance of the Fed Funds rate at the 2 quarter horizon, and small fractions of the other variables at both the short and longer horizons.

Impact of Demographic Trends for Estimation I assess the impact of including demographic shocks by estimating a version of the model where demographics are held constant at their values in the year 2000.[18]

When demographics are not explicitly accounted for, the estimation assigns a more important role for technology shocks in explaining the growth rates and levels of real output and consumption. As reported in Panel C of Table 2, technology shocks explain almost 10 percent

[18] The Online Appendix reports the posterior distributions for this estimation.

of the unconditional variance of output growth (compared to 3 percent without demographics), and 14 percent of the unconditional variance of consumption growth (compared to 4 percent without demographics). As reported in the Online Appendix, without demographic trends, technology shocks account for about 19 percent of the long-run level of output, 14 percent of the level of consumption, and 8 percent of the level of investment, compared to being essentially zero for all three variables when demographic trends are included.

We can inspect the impulse responses of these variables to the exogenous shocks to understand these observations.[19] When demographic trends are not included, technology shocks have a much more persistent effect on the levels of real variables. Mechanically, this reflects the estimates – when demographics are not included in the model, the estimated persistence of technology shocks increases from 0.82 to 0.97 at the posterior mode. Economically, this is because demographic trends manifest themselves in the same way as persistent technology shocks over the estimation sample period, in that they simultaneously raise output, consumption, and investment, and lower inflation and the Fed Funds rate. This in part reflects the lifecycle accumulation of human capital by workers in the model, which at the aggregate level generates productivity dynamics that match those observed, as reported in more detail in the Online Appendix. Omitting demographic trends means that those dynamics are captured instead with persistent technology shocks. With more persistent technology shocks, less of the long-term response of real variables is explained by other shocks, reducing their importance in the variance decomposition.[20]

5 Demographics and the Business Cycle

I next study, using the estimated model, the role that demographic changes have had in explaining the decline in log output relative to its pre-crisis trend. Demographics cause a

[19]The Online Appendix plots the impulse responses under the parameters at the posterior modes from the estimations with and without demographic trends.

[20]This can also be seen in the impulse response functions – the responses of real variables to technology shocks are larger and more persistent in the model that omits demographics trends. On the other hand, markup shocks are larger and more persistent in the model with demographic trends.

direct effect on output and an indirect effect through constrained monetary policy. For the direct effect, I first show that demographic shocks alone are responsible for over one-third of the decline in output since 2007. For the indirect effect, I show that demographic changes were substantially responsible for the ZLB binding between 2009 and 2015, generating an additional non-linear effect on output. I show this by computing a counterfactual holding demographics constant from 1986. After removing the forward guidance response of the Fed to a binding ZLB, I find that output would have fallen by an additional 2% to 5% relative to trend between 2009 and 2015.

5.1 Trends Due to Demographic Shocks

First, I discuss how demographic shocks alone affect the economy's key variables, plotted in Figure 3. Starting in 1986, I turn off all shocks except those to fertility and mortality rates. Panel A plots an index of log output and illustrates how demographic changes cause a slowdown in output growth relative to log output's 1986 to 2007 trend. Panel B shows that, between 1990 and 2019, the Fed Funds rate declines by about 2.5 percentage points. In Panel C of Figure 3, I show that the real interest rate $R_t - \mathbb{E}_t \Pi_{t+1}$ is expected to fall by about 1.2 percentage points, driven by changes in the capital-output ratio. The difference between the decline in the Fed Funds rate and the real interest rate reflects an approximate 1 percentage point decline in the rate of inflation driven by demographic changes and the downward pressure that an aging population places on expected future marginal costs.

Regarding the implications of demographics alone for growth, from 1980 to 2019, my model predicts annualized output growth falls by about 1.68 percentage points, which is slightly above the decline predicted by the analysis of Gagnon et al. (2021). There are three main channels through which output growth can change over time because of changing demographics. Workers can supply more hours, affecting both output and aggregate labor. There are also changes in physical capital, as individuals save and consume out of accumulated savings in retirement. Third, the *quality* of labor can change; namely, changes in the distribution of workers resulting from demographic changes alters the average skill-level of the workforce, which shows up in a decomposition of productivity growth as fluctuations in the average productivity of labor (Fernald, 2015).

I decompose the model's predictions for output growth and labor productivity growth into their component parts and show that accelerating capital accumulation increases the growth rate of both labor productivity and total output up to 1995, after which the growth rate starts to decline.[21] The change in labor supply has a large negative effect on productivity growth, but a positive effect on total growth, when the baby boomer cohorts enter the labor force around 1960. A key component of both labor productivity and total growth is the change in the average skill level of the workforce caused by the interaction of a changing composition of the workforce with the age-productivity profile. The decomposition implies that the contribution of the change in average labor quality to the growth rate of output and output per worker peaks around 1990, adding roughly 0.3 percentage points to total growth and productivity growth. The contribution of labor quality becomes a drag on productivity growth in 2000 as a large fraction of workers reach the peak of the age-productivity profile, exhausting the potential for further growth in average human capital across the workforce. This force is forecast to depress productivity growth until 2030. In total, I find that demographics will be a drag on output growth through to 2070.

Figure 4 explores the trends generated by demographics alone in more detail by plotting, in Panel A, the employment-to-population ratio in the model and the data. Demographics alone capture the dynamics of the aggregate employment-to-population ratio well under the calibration of the lifecycle parameters, which generated age-specific labor force participation rates that are consistent with those observed. The employment-to-population ratio declines in the model at a pace that is roughly as fast as that observed and is predicted to continue to fall by a further 4 percentage points from 2020 to 2040. This result is driven by the compositional changes in the workforce towards workers with lower participation rates (Panel

[21] This decomposition is presented in the Online Appendix.

Table 1: Prior and Posterior Distributions of Estimated Parameters

	Prior				Posterior			
Parameter	Dist	Median	5%	95%	Mode	Median	5%	95%
ρ_χ	B	0.5	0.3	0.7	0.96	0.96	0.94	0.97
ρ_μ	B	0.5	0.3	0.7	0.82	0.82	0.72	0.90
ρ_ξ	B	0.5	0.3	0.7	0.94	0.94	0.92	0.95
ρ_g	B	0.5	0.3	0.7	0.98	0.98	0.97	0.99
ρ_κ	B	0.5	0.3	0.7	0.93	0.93	0.91	0.95
$100 \times \sigma_\chi$	U	2.0	0.2	3.8	1.90	1.93	1.68	2.30
$100 \times \sigma_\mu$	U	2.0	0.2	3.8	0.67	0.65	0.45	0.82
$100 \times \sigma_\xi$	U	2.0	0.2	3.8	0.05	0.05	0.05	0.06
$100 \times \sigma_R$	U	2.0	0.2	3.8	0.10	0.10	0.08	0.13
$100 \times \sigma_g$	U	2.0	0.2	3.8	2.49	2.53	2.30	2.81
$100 \times \sigma_\kappa$	U	2.0	0.2	3.8	1.16	1.18	0.90	1.54

Notes: χ_t is the household preference shock, μ_t is the technology shock, ξ_t is the price markup shock, R_t is the nominal interest rate, g_t is the exogenous government spending shock, and κ_t is the shock to investment-specific adjustment costs. 'B': beta distribution, 'U': uniform distribution.

Figure 3: Output and Interest Rates, Demographics Only

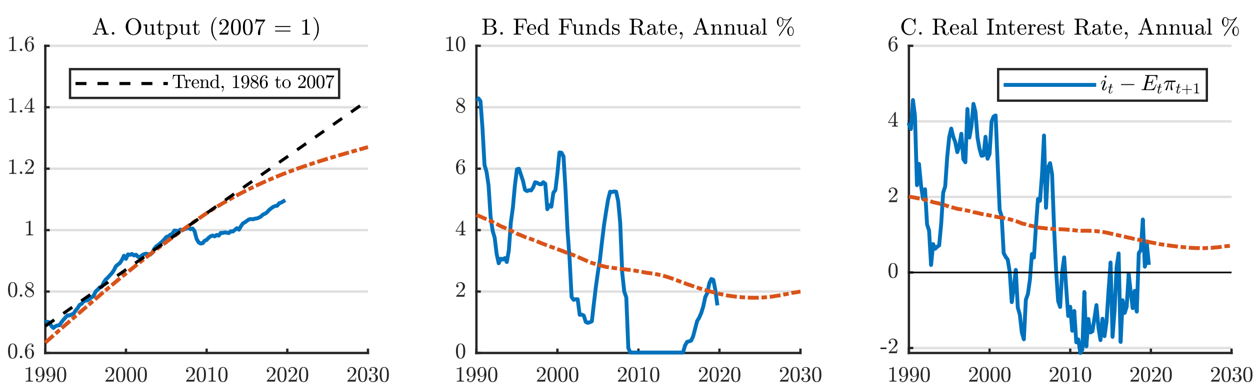

Notes: The solid blue line plots the data and the red dot-dashed line plots the counterfactual.

Table 2: Variance Decompositions, %

Variable \ Shock	Preference	Technology	Markup	Policy	Government	Investment
A. Conditional, 2 Quarter Ahead						
Fed Funds Rate	25	20	3	2	3	46
Inflation	11	4	58	8	2	17
Δ Output	1	3	20	14	14	48
Δ Consumption	39	4	2	8	31	16
Δ Investment	5	1	17	5	0	72
B. Unconditional						
Fed Funds Rate	38	7	10	1	11	34
Inflation	18	3	50	7	5	17
Δ Output	1	3	21	15	13	46
Δ Consumption	38	4	4	8	30	16
Δ Investment	5	1	18	5	0	72
C. Unconditional, Estimated Without Demographic Changes						
Fed Funds Rate	24	8	8	1	11	47
Inflation	11	3	51	10	4	21
Δ Output	2	9	13	14	8	53
Δ Consumption	27	14	3	11	33	13
Δ Investment	2	2	11	5	0	80

B of Figure 2), a point supported empirically by Aaronson et al. (2014).[22] To understand the contribution of demographics-induced changes in the composition of the workforce, a simple shift-share analysis using observed age-specific participation rates and population data reveals that about half of the total decline in the total labor force participation rate since 1996 can be explained by changes in the age-composition alone.[23]

Demographic trends have important implications for the capital-to-output ratio, plotted in the second panel of Figure 4. As life expectancy rises and mortality rates fall, aggregate savings increases to finance longer expected retirements. As a result, the capital-to-output ratio increases and the marginal product of capital and real interest rate fall. In addition, the aging of the baby boomer cohorts generates an increase, and then decrease, in the path for the real interest rate around the secular decline implied by increasing longevity. The oscillation is driven by changes in the relative size and composition of the workforce. The workforce is relatively young as the baby-boomers enter the labor market in the 1960s to 1980s, so that aggregate hours supplied is high relative to capital, thus increasing the marginal return to capital. As the baby-boomer cohort ages and accumulate savings for retirement, the marginal return to capital and the real interest rate decline. This decline is then reinforced by the withdrawal of the baby-boomer cohort from the labor market, depressing the marginal return to capital, which stays low beyond 2030 (see also Carvalho, Ferrero and Nechio, 2016).

Finally, the model does a good job at matching the aggregate net savings rate observed, as shown in Panel C of Figure 4 and consistent with its ability to match the evolution of the capital-output ratio.[24] The influx of the large baby boomer cohorts leads to an increase in

[22] Furthermore, to decompose the contribution of the mechanical effect of an aging population to the decline in the labor input, we can compare (i) the labor input predicted by the model in our baseline exercise where labor endogenously responds to wages, to (ii) the labor input predicted by the model when labor supply is inelastic ($\varphi \to \infty$). In this comparison, almost all of the forecasted decline in the labor input is due to the mechanical effect of demographic changes; the endogenous response of labor in my model to demographic changes mitigates the 8 percent decline in the labor input by 2030 (relative to 2015) by only 2 percentage points.

[23] This full shift-share analysis is presented in the Online Appendix.

[24] The net saving series in the data is from the US Bureau of Economic Analysis, net saving as a percentage of gross national income (under FRED code W207RC1Q156SBEA). I add historical and projected population growth rates to the model series (this adds roughly 1% to the net saving rate series historically, and only about 0.4% in the long-run).

the net saving rate from about 5% in 1980 to almost 7% by the mid-1990s as those cohorts save for expected retirement. As the baby boomers start to leave the workforce and dissave, the net saving rate gradually declines and turns negative in 2031, and stays negative until about 2075. Demographic changes in my model are thus projected to depress the savings rate well into the future. These forces are fundamentally the same as those driving the real interest rate lower. In contrast, however, the real rate falls from the mid-1980s onwards, which is a little earlier than the decline in the net savings rate. This reflects weaker labor supply growth relative to the growth in the capital stock starting in the mid-1980s.

These results on the implications of demographics for macroeconomic trends are consistent with the findings of other studies. Using a calibrated overlapping generations model, Gagnon et al. (2021) find that demographic trends generate a decline in the growth rate from around 2.3 percentage points in 1980 to just under 0.5 percentage points by 2030. They find, similar to what is predicted by my model, that much of the decline is due to declining fertility and the associated exit of the baby boomer generations from the labor force.[25] In an empirical study, Aksoy et al. (2019) predict an average decline in annual output growth rates across OECD countries of 1¼ percentage points between 2010 and 2030. Regarding real interest rates, Gagnon et al. (2021) find a peak in the real rate between 1975 and 1985 of around 1.7%, and a decline of about 1.5 percentage points by 2030, close to my model's predictions. In an empirical study, Johannsen and Mertens (2021) find a decline in the real rate of about ¾ percentage points between 1985 and 2015, while Eggertsson et al. (2019) forecast a larger decline in their model implied interest rate due to demographics of around 3 percentage points between the 1980s and 2030.

5.2 Demographics, Business Cycle Shocks and the ZLB

As shown in the previous section, over the estimation period 1986 to 2019, demographic changes imply quantitatively relevant declines in growth rates and real and nominal interest

[25] Decompositions of these contributing factors are presented in the Online Appendix.

rates. In this section, I study the interaction between demographic changes, business cycle shocks, and the ZLB using the estimated model.

5.2.1 Holding Demographics Constant from 1986

Here, I examine whether demographic trends were responsible for the Fed Funds rate hitting the ZLB between 2009 and 2015. To answer this, I hold the demographic profile constant at its 1986 state, which was the first year that quarterly data is used in the Bayesian estimation, and construct a counterfactual using the estimated structural shocks. As illustrated in Panel B of Figure 5, had the population not aged between 1986 and 2015, the significant recessionary shocks still push the Fed Funds rate to the ZLB in the immediate aftermath of the Great Recession. However, the period at the lower bound would have lasted about a year, with liftoff occurring by the start of 2010, and the Fed Funds rate rapidly increasing thereafter. Absent aging, there would have been greater space for conventional monetary stimulus afforded by the higher level of the Fed Funds rate. Because of this, the real interest rate, plotted in Panel C, would have declined by more than observed over 2009.

Panel A shows the effect for output of fixing demographics from 1986. Interestingly, the counterfactual response shows that output growth would have been lower up to 2007. The reason for this is that the composition of the workforce in 1986 is skewed towards younger workers and so demographics were therefore favorable to productivity growth between 1986 and the mid-2000s, as younger cohorts accumulate human capital as they move up the age-productivity profile. However, from 2011, demographic changes become a drag on total output growth, and output growth in the counterfactual with demographics fixed at their 1986 state is higher than that observed. As discussed above, demographic trends are forecast to weigh on growth relative to the observed rate well into the future.

Figure 4: Macroeconomic Trends Due to Demographics

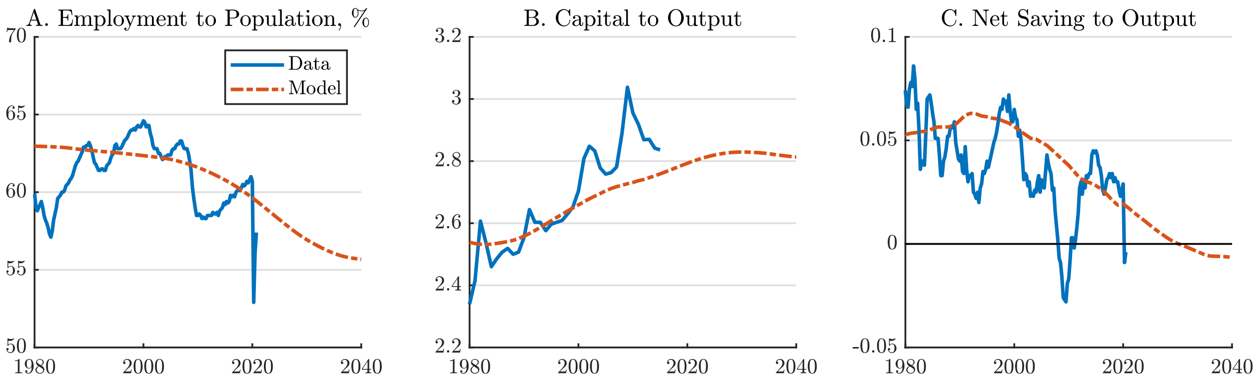

Notes: The data for the capital-output ratio is extracted from the BLS's Multifactor Productivity data published in 2016. The series is the inverse of the 'output per unit of capital services' analytic ratio for the private non-farm business sector (excluding government enterprises).

Figure 5: Output and Interest Rates, Demographics Fixed from 1986

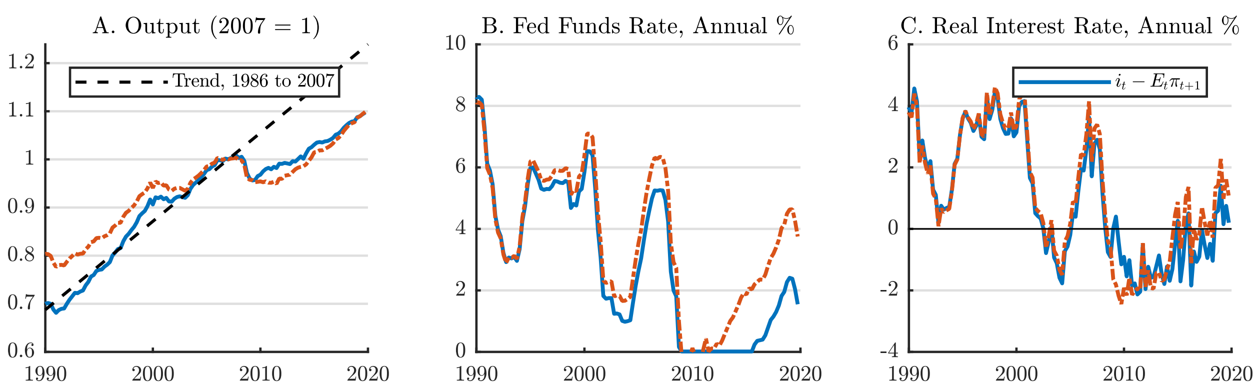

Notes: The solid blue line plots the data and the red dot-dashed line plots the counterfactual.

5.2.2 Impact of Demographic Trends for the Propagation of Shocks

Next, I examine the nonlinear interactions between the ZLB and the decline in interest rates induced by demographic changes. In Figure 6, I plot the response of the Fed Funds rate and output to a large negative investment shock for two demographic cases: one associated with the structure of the economy that arose in 1990, and the other for the structure of the economy that arose in 2008 just prior to the period of the binding ZLB. In conducting this impulse response exercise, I use the reduced-form matrices of the transition equation (29) corresponding to the years 1990 and 2008. That is, I use $\{\mathbf{J}_k, \mathbf{Q}_k, \mathbf{G}_k\}_{k=\{1990,2008\}}$ to compute the path of variables following the shock, and initialize the variables at the steady-state associated with $\{\mathbf{J}_k, \mathbf{Q}_k, \mathbf{G}_k\}_{k=\{1990,2008\}}$ by computing $x_{k=\{1990,2008\}} = (I - \mathbf{Q}_k)^{-1}\mathbf{J}_k$ and setting the initial vector of variables to x_k.

The initial value of the Fed Funds rate is higher for the demographics structure of the economy arising in 1990, when the population is younger. For the structure of the economy associated with demographics in 2008, the shock is large enough to cause the ZLB to bind for about two years, with output falling by an additional 2 percentage points on impact. The two responses illustrate how the same shock can have very different implications for the economy by bringing real and nominal interest rates closer to the ZLB.[26]

5.2.3 Frequency of ZLB Episodes

The declining Fed Funds rate caused by demographic changes implies that, going forward, the ZLB is likely to be visited more frequently. To quantitatively assess this, I simulate the model 500 times for 500 periods under three demographic profiles: first, when they are fixed at their 1990 profile; second, when they are fixed at their 2008 profile; and third, when they are fixed to their 2020 profile. In these simulations, the Fed Funds rate is set according to $\log(R_t) = \max(0, \text{Taylor Rule}_t)$, that is, without forward guidance policy. In the first case,

[26]Due to the time-varying structural changes induced by demographics, there are also small differences in how shocks propagate which is unrelated to the binding lower bound. I present additional results of such differences in the Online Appendix.

when the population is younger and the steady-state interest rate is higher, the ZLB binds 1.4% of the time, while it binds for 5.2% of the time under the 2008 demographic profile, and 11.7% of the time under the 2020 demographic profile. The duration that the ZLB binds in these simulations can be longer due to forward guidance policy, discussed in more detail in the next section. For example, in the data, the Fed Funds rate lifted off in 2015Q4 while in the counterfactual without forward guidance, the Fed Funds rate was above zero for two quarters in 2011, and lifted off zero from 2013Q2 onwards, giving a counterfactual lower bound episode of 14 quarters compared to 27 quarters in the data.

One potential policy response would be to raise the inflation target, thereby raising the steady-state nominal interest rate. In simulations, I find that the annual inflation target would need to be raised to about 3.5% to obtain approximately the same ex-ante distribution of expected ZLB episodes in 2008 as would arise under an inflation target of 2% and when demographics are set to their 1990 level.

5.3 Effect of Forward Guidance

The previous section showed how demographic trends over the past 25 years led to a decline in real and nominal interest rates, and the ZLB to be a constraint on monetary policy. I explore in this section how much of a constraint the ZLB was quantitatively by constructing a counterfactual simulation of the economy in which the Fed acts passively in response to shocks that drive the policy rate to the ZLB. In this simulation, the expected ZLB duration adjusts in response to the shock only, so that the policy interest rate is determined by $\log(R_t) = \max(0, \text{Taylor Rule}_t)$. By contrast, in the estimation, I fix the expected ZLB durations to those observed in survey data. Fixing the durations in this way allows for the possibility that the Fed extended the ZLB duration beyond the duration implied by the shocks themselves (see also Campbell et al., 2012; Jones, 2017). The counterfactual simulation therefore provides a measure of the degree to which the ZLB was a binding constraint, absent explicit forward guidance policies.

Figure 7 illustrates the counterfactual path of the economy without forward guidance. As Panel B shows, when forward guidance is removed, the structural shocks imply that the Fed Funds rate remains at the ZLB, but lifts off briefly in 2011 and then again in 2013Q2, earlier than observed, supporting the notion that the Fed Funds rate was held at zero 'lower for longer'. When there is no ZLB, the Fed Funds rate would have been lowered to about -1% from 2009 to 2013, after which it increases and lift-offs the bound by 2014Q3. Panel A plots the path of output. Absent forward guidance, output would have fallen by, at most, an additional 5 percentage points by 2011. Panel C shows that the real interest rate would have been higher by about 2 percentage points between 2009 and 2015, reflecting lower expected inflation without the extra monetary stimulus.

Next, I compute the ZLB durations implied by the shocks alone, which provides a measure of how stimulatory forward guidance is.[27] I find some degree of forward guidance stimulus every quarter between 2009 and 2015, with the strongest stimulus between 2011Q3 and 2013Q2, when the forward guidance component of the total duration is estimated to be between 8 and 11 quarters at the mode of the posterior distribution. This period corresponds to low yields on long-term Treasuries and the explicit calendar-based targets announced by the Fed. These results are also consistent with the findings in Swanson and Williams (2014), who show that between 2009 and 2011, long-term yields were relatively unconstrained, and that after 2011, long-term yields tightened significantly towards their lower bounds; consistent with the Fed announcing expansive unconventional monetary policies. In particular, in mid-2011, the Fed announced its "to mid-2013" guidance announcement, the first of many subsequent calendar-based extensions of the lower bound regime.[28]

[27] I calculate the ZLB durations implied by the structural shocks, using a method described in Jones (2017). The difference between those computed endogenous durations and the durations used in the estimation is the contribution of forward guidance, or the extension of the ZLB regime that, together with the structural shocks, will generate the observed series, as discussed in Jones et al. (2021). The decomposed ZLB durations are plotted in the Online Appendix.

[28] For example, in the FOMC press release, August 9 2011, the FOMC announced: "The Committee currently anticipates that economic conditions – including low rates of resource utilization and a subdued outlook for inflation over the medium run – are likely to warrant exceptionally low levels for the federal funds rate at least through mid-2013."

Figure 6: Impulse Response, Shock to the Efficiency of Investment

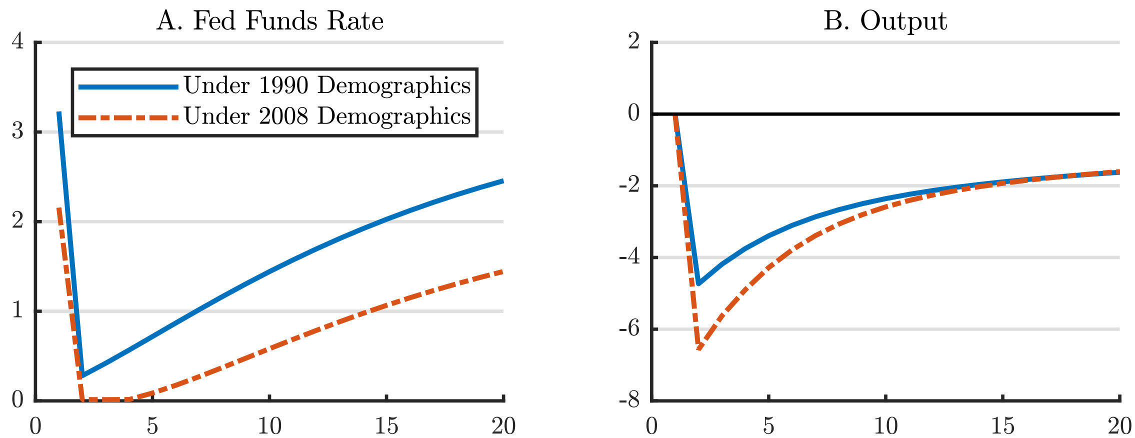

Figure 7: Output and Interest Rates, No Forward Guidance

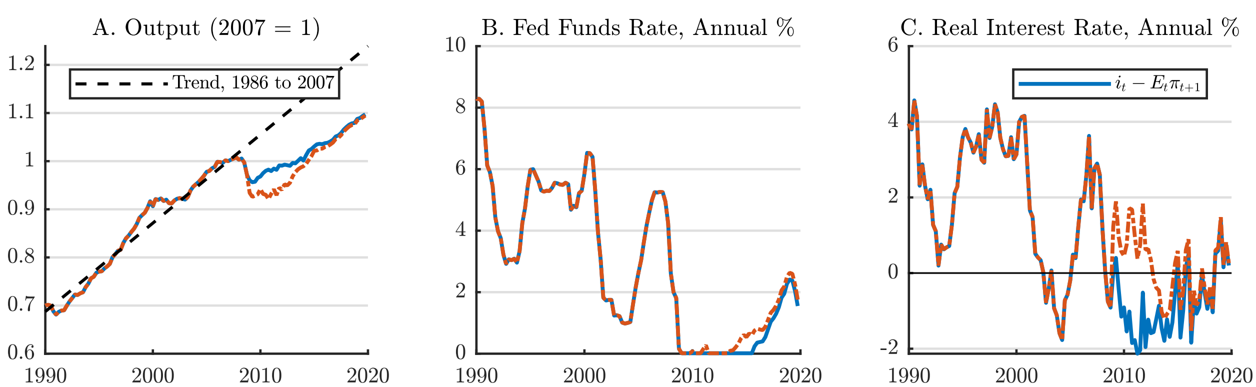

Notes: The solid blue line plots the data and the red dot-dashed line plots the counterfactual.

5.4 Decomposition of Output Since the Great Recession

In this section, I put the results from the previous sections together to study how the model decomposes the decline in log output relative to its long-run linear trend. The left panel of Figure 8 plots, in the solid line, the difference between output and its long-run trend relative to 2007. The long-run trend is computed using the average growth rate from 1986 to 2007. The data show the severity of the slowdown in output, with the gap between the data and its trend widening between 2007 and 2019 to 14%. Demographics alone account for two-fifths, or 6 percentage points, of the gap by the end of 2019. Without expansionary forward guidance policy, the gap would have been up to 5 percentage points larger, primarily between 2011 and 2013.

The nonlinearities associated with the ZLB and forward guidance policy means that there is no linear decomposition of the remaining gap between output and its demographic path into the contribution of individual shocks.[29] For this reason, I explore how each estimated shock affects output when the ZLB constraint is turned off. The right panel of Figure 8 plots the path of output relative to 2007 in counterfactuals that are computed in the following way. First, I compute counterfactual paths of output when each shock is set to zero, respectively, and the Fed Funds rate responds endogenously and can fall below zero. I then evaluate the contribution of each shock as the difference between the observed output series and the counterfactual output series relative to 2007.

The results suggest that investment efficiency shocks explain why output collapsed in 2009 and recovers slowly through to 2015. These shocks, as emphasized by Justiniano et al. (2011), capture disruptions in the intermediation between savings and investment. They therefore capture the impact that financial factors have on aggregate variables which, as shown by Jermann and Quadrini (2012) and Del Negro et al. (2015), are key to explaining the decline in output around 2008-09.[30] The results also show that government spending

[29]The time-series of shocks are shown in the Online Appendix.

[30]As discussed above, the shocks to the efficiency of investment – shocks to κ_t in (5) – correlate with the spread between the yields of Baa Corporate debt and 10 year Treasury bonds, and spike at the same time

shocks kept output from falling even further between mid-2009 and 2012. Discount factor shocks have a small effect on output when monetary policy is unconstrained and able to respond. This observation is consistent with the variance decompositions of Table 2, where preference shocks explain more than a third of the unconditional forecast error variance of the Fed Funds rate, but small fractions of the variance of output, consumption, and investment.

Finally, the right panel of Figure 8 shows that price markup shocks also imply a fall in output in years of the Great Recession and can explain the remaining persistent drop in output through to 2020. As discussed above in section 4.2.2, shocks to markups can capture supply-side factors or other unmodeled trends that depress output. These unmodeled forces could include, for example: firm entry dynamics where, as Gutierrez et al. (2021) show, positive shocks to firm entry costs since the 2000s can explain depressed investment, higher concentration, and lower real interest rates; or a rise in idiosyncratic income risk which could bring down interest rates and output through a stronger precautionary demand for savings when interest rates are near the ZLB (Auclert and Rognlie, 2018; Aladangady et al., 2021).

5.5 Robustness

I conduct a number of experiments to verify that the aggregate trend predictions from the model under expected demographic changes are robust to alternative specifications. The results of each experiment are presented in the Online Appendix but discussed briefly here.

Borrowing Constraints I first check that the model's predictions hold when individuals face a constraint restricting their borrowing early in life. With borrowing constraints, there are more savings, pushing up the capital-output ratio. As a consequence, the real interest rate is lower than in the baseline model. The magnitude of the fluctuations of the real interest rate, the participation rate and output growth are very similar to the baseline model.

as the spike in the spread in 2008, as I show in the Online Appendix.

Time-Varying Productivity Profile The second robustness check is to adopt time-varying productivity profiles to account for a possible flattening of productivity profiles over time. Such a flattening can affect the accumulation of human capital and can impact aggregate productivity measures in two ways: first, by a growth effect, by lowering the potential for new workers to accumulate human capital, and second, by a level effect, by affecting the productivity level that individuals enter the workforce on. I calibrate the age-productivity profiles by recomputing for each cross-sectional sample, the profile and then interpolating between those points in time. The overall pattern of aggregate labor productivity is much the same as the baseline model, although the magnitude of the amplitude of the change in labor productivity growth is smaller, with demographics contributing the most to labor productivity growth in 1980 rather than in 1990 (as is the case in the baseline results).

Female Labor Force Participation and Multiple Skill Types From 1985 on, the baseline predictions for the participation rate, aggregate labor productivity growth and the real interest rate are largely unaffected when the age-productivity and labor disutility profiles are calibrated to match female age-earnings profiles and female labor force participation rates from the 1940s to 1990s, after which female labor force participation is roughly constant. As a final point of comparison, I verify that the directions of the aggregate predictions are robust to a calibration where an additional source of heterogeneity is modeled–where there are two types of workers, low or high skilled, with low skill workers calibrated to those with less than college education.

Unanticipated Demographic Changes Demographic changes are assumed to be perfectly foreseen in the baseline analysis, as is common in the literature. As explained in Section 3.2, the assumption of perfect foresight also significantly aids the construction of the model's likelihood function, making Bayesian estimation feasible. We can also construct the demographics-induced trends in the model under the assumption that demographic changes arrive as a surprise every quarter. I plot in the Online Appendix the real interest rate and the growth rate of output under this assumption. Compared to the baseline results, the paths of the real interest rate and output are similar, with somewhat smaller fluctuations in the real interest rate. The anticipation of demographic changes generates larger swings in labor supply and savings behavior which influences the path of the interest rate.

6 Conclusions

This paper studies why the level of US output remains significantly below its pre-crisis trend after the Great recession. I use a New Keynesian model with demographic shocks and the ZLB to show that declining mortality rates and changes to the age population composition can generate long-run trends that match the low frequency movement of output growth, productivity, the real interest rate, and the employment-population ratio.

I estimate the model using Bayesian likelihood methods. With the estimated model, I find that the ZLB would not have been a binding constraint between 2010 and 2015 had demographics not changed from 1986. I find that demographic shocks alone are responsible for about 40% of the decline in output relative to its pre-crisis trend by 2019. Furthermore, my results suggest that absent any forward guidance policy used by the Fed, the ZLB would have caused output to fall by an additional 2 to 5 percentage points between 2011 and 2013.

I also assess the contribution of each of the estimated shocks to the decline in output since the Great Recession, and find an important role for investment shocks in causing output to fall. These shocks proxy for financial disturbances as they capture disruptions in the intermediation between savings and investment.

The results illustrate the importance of demographics as a major driver of macroeconomic trends over time. Further research could focus on how demographic trends interact with the housing market or with the efficacy of fiscal policy. It would also be interesting to model a more detailed financial sector to study the interactions between demographic changes and financial frictions, particularly as borrowing constraints may bind differently across cohorts.

References

Aaronson, Stephanie, Thomaz Cajner, Bruce Fallick, Felix Galbis-Reig, Christopher Smith, and William Wascher, "Labor Force Participation: Recent Developments and Future Prospects," *Brookings Papers on Economic Activity*, 2014, *Fall*, 197–255.

Aksoy, Yunus, Henrique S. Basso, Tobias Grasl, and Ron P. Smith, "Demographic Structure and Macroeconomic Trends," *American Economic Journal: Macroeconomics*, 2019, *11* (1), 193–222.

Aladangady, Aditya, Etienne Gagnon, Benjamin K. Johannsen, and William B. Peterman, "Macroeconomic Implications of Inequality and Income Risk," 2021.

Antolin-Diaz, Juan, Thomas Drechsel, and Ivan Petrella, "Tracking the Slowdown in Long-Run GDP Growth," *The Review of Economics and Statistics*, 05 2017, *99* (2), 343–356.

Attanasio, Orazio, Sagiri Kitao, and Giovanni L. Violante, "Global Demographic Trends and Social Security Reform," *Journal of Monetary Economics*, 2007, *54*, 144–198.

Auclert, Adrien and Matthew Rognlie, "Inequality and Aggregate Demand," 2018.

Auerbach, A. J. and L. J. Kotlikoff, *Dynamic Fiscal Policy*, Cambridge University Press, 1987.

Campbell, Jeffrey R., Charles L. Evans, Jonas D.M. Fisher, Alejandro Justiniano, Charles W. Calomiris, and Michael Woodford, "Macroeconomic Effects of Federal Reserve Forward Guidance," *Brookings Papers on Economic Activity*, 2012, *Spring*, 1–80.

Canova, Fabio, Filippo Ferroni, and Christian Matthes, "Approximating Time Varying Structural Models with Time Invariant Structures," 2015.

Carvalho, Carlos, Andrea Ferrero, and Fernanda Nechio, "Demographics and Real Interest Rates: Inspecting the Mechanism," *European Economic Review*, September 2016, *88*, 208–226.

Constantinides, George M., "Intertemporal Asset Pricing with Heterogeneous Consumers and Without Demand Aggregation," *The Journal of Business*, 1982, *55* (2), 253–267.

Eggertsson, Gauti B and Michael Woodford, "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 2003, *1*, 139–233.

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

Elsby, Michael W. L. and Matthew D. Shapiro, "Why Does Trend Growth Affect Equilibrium Employment: A New Explanation of an Old Puzzle," *American Economic Review*, 2012, *102* (4), 1378–1413.

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

Fernández-Villaverde, Jesús, Grey Gordon, Pablo A. Guerrón-Quintana, and Juan Rubio-Ramírez, "How Structural Are Structural Parameters?," *NBER Macroeconomics Annual*, June 2007, *22*, 83–132.

Feyrer, James, "Demographics and Productivity," *The Review of Economics and Statistics*, February 2007, *89* (1), 100–109.

Gagnon, Etienne, Benjamin K. Johannsen, and David Lopez-Salido, "Understanding the New Normal: The Role of Demographics," *IMF Economic Review*, 2021.

Guerrieri, Luca and Matteo Iacoviello, "Occbin: A Toolkit to Solve Models with Occasionally Binding Constraints Easily," *Journal of Monetary Economics*, March 2015, *70*, 22–38.

Gutierrez, German, Callum Jones, and Thomas Philippon, "Entry Costs and Aggregate Dynamics," 2021.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?," Working Paper 20913, National Bureau of Economic Research January 2015.

Jaimovich, Nir and Henry Siu, "The Trend is the Cycle: Job Polarization and Jobless Recoveries," 2012.

Jermann, Urban and Vincenzo Quadrini, "Macroeconomic Effects of Financial Shocks," *American Economic Review*, 2012, *102* (1), 238–71.

Johannsen, Benjamin K. and Elmar Mertens, "A Time Series Model of Interest Rates With the Effective Lower Bound," *Journal of Money, Credit and Banking*, 2021.

Jones, Callum, "Unanticipated Shocks and Forward Guidance at the Zero Lower Bound," 2017.

_ , Mariano Kulish, and Daniel Rees, "The International Spillovers of Unconventional Monetary Policy," *Journal of Applied Econometrics*, 2021. forthcoming.

_ , Virgiliu Midrigan, and Thomas Philippon, "Household Leverage and the Recession," 2018. NYU.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti, "Investment Shocks and the Relative Price of Investment," *Review of Economic Dynamics*, 2011, *14*, 102–121.

Krusell, Per and Anthony A Smith Jr, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of political Economy*, 1998, *106* (5), 867–896.

Kulish, Mariano and Adrian Pagan, "Estimation and Solution of Models with Expectations and Structural Changes," *Journal of Applied Econometrics*, 2016.

_ , Christopher Kent, and Kathryn Smith, "Aging, Retirement, and Savings: A General Equilibrium Analysis," *The BE Journal of Macroeconomics*, 2010, *10* (1).

_ , James Morley, and Tim Robinson, "Estimating DSGE Models with Zero Interest Rate Policy," *Journal of Monetary Economics*, 2017, *88*, 35–49.

Maliar, Lilia and Serguei Maliar, "The Representative Consumer in the Neoclassical Growth Model with Idiosyncratic Shocks," *Review of Economic Dynamics*, 2003, *6*, 362–380.

Negro, Marco Del, Marc Giannoni, and Frank Schorfheide, "Inflation in the Great Recession and New Keynesian Models," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 168–196.

Reichling, Felix and Charles Whalen, "Review of Estimates of the Frisch Elasticity of Labor Supply," Working Paper 13, Congressional Budget Office 2012.

Ríos-Rull, José-Víctor, "Life-Cycle Economies and Aggregate Fluctuations," *Review of Economic Studies*, 1996, *63*, 465–489.

Rios-Rull, Jose-Victor, Frank Schorfheide, Cristina Fuentes-Albero, Maxym Kryshko, and Raul Santaeulalia-Llopis, "Methods Versus Substance: Measuring the Effects of Technology Shocks," 2012, *59* (8), 826–846.

Smets, Frank and Rafael Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 2007, *96* (3), 586–606.

Summers, Lawrence H., "U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound," *Business Economics*, 2014, *49* (2), 65–73.

Swanson, Eric and John Williams, "Measuring the Effect of the Zero Lower Bound on Medium- and Long-term Interest Rates," *American Economic Review*, 2014, *104* (10), 3154–3185.

Werning, Iván, "Managing a Liquidity Trap: Monetary and Fiscal Policy," 2012.

Wong, Arlene, "Population Aging and the Transmission of Monetary Policy to Consumption," 2015.

Figure 8: Log Output Relative to Trend (2007=0)