Demographic Forces and U.S. Labor Market Dynamics: An Open-Economy DSGE Perspective

Thierry U. KAME BABILLA

University of Yaounde II and Euro Area Business Cycle Network

Abstract

We build an open-economy Dynamic Stochastic General Equilibrium (DSGE) model in which households are heterogeneous in age and labor market status, to study how demographic transition shapes U.S. labor market dynamics and external adjustment. The home economy (United States) is populated by overlapping cohorts of workers whose labor supply, savings, and job-finding behavior evolve with age, embedded in a New Keynesian core with sticky prices, search-and-matching frictions, and incomplete asset markets linking the home economy to the rest of the world. Population aging operates through three channels: (i) a labor-supply channel that shrinks the effective labor force and raises the equilibrium tightness of the labor market; (ii) a life-cycle saving channel that raises desired aggregate saving and depresses the natural rate of interest; and (iii) a composition channel that shifts the matching efficiency and bargaining power of the aggregate workforce as the age profile of job seekers changes. We characterize the model's steady state and log-linearized dynamics, calibrate it to U.S. and rest-of-world data, and use impulse response and variance decomposition analysis to quantify the contribution of demographic shocks to labor force participation, real wage growth, the natural rate of unemployment, and the U.S. current account. The model rationalizes the secular decline in labor force participation, a flatter Beveridge curve, and persistent current account deficits coexisting with subdued real interest rates, and offers a structural laboratory for evaluating immigration, retirement-age, and monetary policy responses to demographic change. Bayesian estimations of the model confirm that, accelerated aging worsens all four outcomes — output growth, unemployment, real wage growth, and the current account — relative to baseline while immigration offset improves them, with the ranking robust across the full posterior but with substantial parameter uncertainty surrounding the magnitude of each scenario's

effect. Counterfactual analysis emphasizes that, the two demographic shocks (labor-supply/aging and longevity/saving) are the dominant drivers of the current account at medium horizons (20–40 quarters), jointly accounting for more than 50% of forecast-error variance, while TFP dominates output and inflation at short horizons. The monetary policy shock matters most for output and inflation at horizons of one to four quarters but has negligible long-run influence, as expected from the New Keynesian structure. The matching-efficiency shock, is the second-largest driver of the effective labor force and real wage at all horizons.

JEL Classification: C11, C32, C58, E24, E32, F41. E24, E32, F41, J11, J21, J64.

Keywords: Demographic Transition; Secular Decline in Labor Force Participation; Search and Matching Frictions; Flatter Beveridge Curve; Natural Rate of Interest; Persistent Current Account Deficits; Population Aging; Labor-supply Channel; Composition Channel; Life-cycle Saving Channel; United States; DSGE Model.

1 Introduction

The United States labor market has undergone a series of structural shifts since the early 2000s that are difficult to reconcile with standard business-cycle models calibrated to a stationary population structure. The aggregate labor force participation rate (LFPR) peaked near 67.3% in early 2000 and, despite a post-pandemic recovery, has settled persistently below that level; the prime-age (25–54) share of the working-age population has been declining since the early 2010s; and the old-age dependency ratio — the ratio of the population aged 65 and above to the population aged 20–64 — has risen from roughly 19% in 1990 to above 28% today, with U.S. Census Bureau and UN Population Division projections placing it near 35–38% by mid-century. Over the same period, the U.S. economy has run persistent current account deficits while long-run real interest rates have trended down, a pattern consistent with the global "saving glut" and "secular stagnation" narratives associated with Bernanke (2005), Summers (2014), and Rachel and Smith (2017), but one whose *labor market* counterpart — the joint behavior of participation, wages, job-finding rates, and the natural rate of unemployment — has received comparatively less treatment within a unified open-economy general equilibrium framework.

This paper develops an open-economy DSGE model in which demographic structure is a first-order determinant of both labor market and external-sector dynamics. We depart from representative-household New Keynesian open-economy models (e.g., Galí, 2015; Galí and Monacelli, 2005) by embedding a tractable overlapping-cohorts structure, in the spirit of Blanchard (1985) and Gertler (1999), in which households face a constant hazard of transiting from "young/working-age" to "old/retired" status, and the size of each cohort relative to the aggregate population is a directly calibrated, time-varying object disciplined by Census Bureau and UN demographic projections. Labor market frictions are introduced through Diamond–Mortensen–Pissarides (DMP) search and matching (Mortensen and Pissarides, 1994; Pissarides, 2000), so that aggregate employment, vacancies, and the job-finding rate are jointly determined and respond endogenously to the age composition of the labor force, allowing the model to speak to the empirical flattening of the Beveridge curve and to age-dependent differences in matching efficiency and separation rates documented in micro data (Hobijn and Şahin, 2015). The home (U.S.) economy trades a final consumption good with a foreign bloc through an incomplete-markets, one-period non-contingent bond structure with a debt-elastic risk premium, in the tradition of Schmitt-Grohé and Uribe (2003), which pins down a well-defined steady state and a stable equilibrium real exchange rate and current account, even as the home economy's effective labor supply and desired saving rate shift permanently with the demographic transition.

The model is organized around three transmission channels through which demographic change affects labor market and external outcomes. First, a *labor-supply channel*: as the population ages, the effective labor force — weighted by age-specific participation propensities — shrinks relative to the total population, mechanically tightening the labor market and putting upward pressure on real wages and marginal costs, all else equal. Second, a *life-cycle saving channel*: anticipated longer retirements raise each cohort's desired wealth-to-income ratio along the lines of the life-cycle/permanent-income hypothesis, raising aggregate desired saving, depressing the equilibrium (natural) real interest rate, and — in the open economy — generating capital outflows and a current account surplus relative to the closed-economy counterfactual, an effect that can be offset or reinforced by relative demographic trends abroad. Third, a *composition channel*: the age profile of unemployed search effort, matching efficiency, and reservation wages shifts the aggregate Beveridge curve and natural rate of unemployment independently of the pure labor-supply effect, because older job seekers in the data exhibit lower separation rates but also lower job-finding rates conditional on separation (Hobijn and Şahin, 2015; Aaronson et al., 2014).

We solve the model by perturbation around the non-stochastic, demographically-stationary steady state and analyze a deterministic transition path consistent with Census/UN medium-variant demographic projections through 2050, together with linearized impulse responses to (i) an unanticipated permanent shock to the effective labor force consistent with accelerated population aging, and (ii) a shock to longevity-induced precautionary/life-cycle saving. We then decompose the forecast-error variance of output, real wages, and the current account into demographic and conventional business-cycle shocks (total factor productivity, monetary policy, and foreign demand).

Three results emerge from the analytical analysis. First, the model-implied decline in effective labor supply driven purely by population aging accounts for a non-trivial share — in our calibration on the order of one-fifth — of the secular decline in the aggregate LFPR since 2000, with the remainder attributable to non-demographic factors (e.g., disability trends, opioid-related labor force exit, and changing prime-age participation incentives) that lie outside the scope of a purely demographic model. Second, the life-cycle saving channel is quantitatively important for the external sector: a calibrated rise in longevity of the magnitude observed in U.S. vital statistics since 1990, holding foreign demographics fixed, depresses the natural rate of interest and worsens (improves the financing of) the U.S. current account deficit by widening the gap between domestic saving and investment, consistent with the empirical co-movement of population aging and the natural rate documented by Gagnon, Johannsen, and López-Salido (2016) and Carvalho, Ferrero, and Nechio (2016). Third, the composition channel materially flattens the model's Beveridge curve and lowers the estimated

natural rate of unemployment relative to a model with age-invariant search parameters, helping rationalize the historically low U.S. unemployment rates observed alongside moderate vacancy rates in the decade preceding the COVID-19 pandemic.

This paper also conducts Bayesian estimations of the DSGE model, and uses the resulting posterior distribution to construct predictive scenario analysis under alternative demographic paths. We specify a compact three-state, four-observable Gaussian state-space system that preserves the model's three core propagation channels (total factor productivity, the demographic/labor-supply state, and the net-foreign-asset/saving state) while remaining low-dimensional enough for transparent, fully documented Markov chain Monte Carlo (MCMC) estimation. We place economically motivated priors on all reduced-form parameters, validate the estimation procedure by Monte Carlo technique, and report full posterior distributions, trace plots, and Gelman–Rubin convergence diagnostics from a three-chain adaptive Metropolis sampler. We then use posterior draws to construct posterior-predictive fan charts for output growth, unemployment, real wage growth, and the current account under three thirteen-year-ahead demographic scenarios — a baseline UN medium-variant aging path, an accelerated-aging path, and an immigration-offset path — and find that : accelerated aging worsens all four outcomes relative to baseline while immigration offset improves them, with the ranking robust across the full posterior but with substantial parameter uncertainty surrounding the magnitude of each scenario's effect.

Counterfactual analysis emphasizes that, the two demographic shocks (labor-supply/aging and longevity/saving) are the dominant drivers of the current account at medium horizons (20–40 quarters), jointly accounting for more than 50% of forecast-error variance, while TFP dominates output and inflation at short horizons. The monetary policy shock matters most for output and inflation at horizons of one to four quarters but has negligible long-run influence, as expected from the New Keynesian structure. The matching-efficiency shock, is the second-largest driver of the effective labor force and real wage at all horizons.

This paper contributes thus to four strands of literature. The first is the open-economy New Keynesian literature following Galí (2015) and Galí and Monacelli (2005), from which we borrow the incomplete-markets debt-elastic interest premium device of Schmitt-Grohé and Uribe (2003) to close the small open economy. The second is the literature on overlapping generations within New Keynesian DSGE models, notably the perpetual-youth framework of Blanchard (1985) as adapted to monetary models by Gertler (1999) and subsequently extended to richer life-cycle and retirement margins (Kara and von Thadden, 2016; Carvalho, Ferrero, and Nechio, 2016). The third is the search-and-matching literature (Mortensen and Pissarides, 1994; Pissarides, 2000; Shimer, 2005) and its application to demographic heterogeneity in job finding and separation (Hobijn and Şahin, 2015). The fourth is the

macro-finance and secular-stagnation literature linking demographic change to the natural rate of interest and global imbalances (Bernanke, 2005; Summers, 2014; Rachel and Smith, 2017; Gagnon, Johannsen, and López-Salido, 2016; Carvalho, Ferrero, and Nechio, 2016; Auerbach and Kotlikoff, 1992). Our contribution is to integrate the labor-supply, saving, and search-composition channels within a single open-economy general equilibrium structure suitable for quantifying their relative contribution to U.S. labor market and external-sector dynamics.

The remainder of the paper is organized as follows. Section 2 documents the demographic and labor market facts that motivate the analysis. Section 3 lays out the model. Section 4 characterizes the competitive equilibrium and its log-linear approximation. Section 5 discusses calibration. Section 6 presents analytical results of the model, the deterministic transition. Section 7 conducts estimations of the model. Section 8 discusses quantitative results of the paper. Section 9 performs counterfactual analysis. Section 10 examines sensitivity analysis. Section 11 presents forecast and decomposition. Section 12 concludes. Further derivations are collected in the Appendix.

2 Demographic and Labor Market Motivating Facts

Figure 1 summarizes the demographic transition underlying the analysis. The left panel plots the old-age dependency ratio, which rises from below 20% in 1990 toward the upper 30s by 2050 under medium-variant projections, with the increase concentrated in the 2010–2030 window corresponding to the retirement of the baby-boom cohort. The right panel plots the aggregate LFPR alongside the prime-age population share; both display a pronounced post-2000 decline that stabilizes, but does not reverse, after 2030 in the projection.

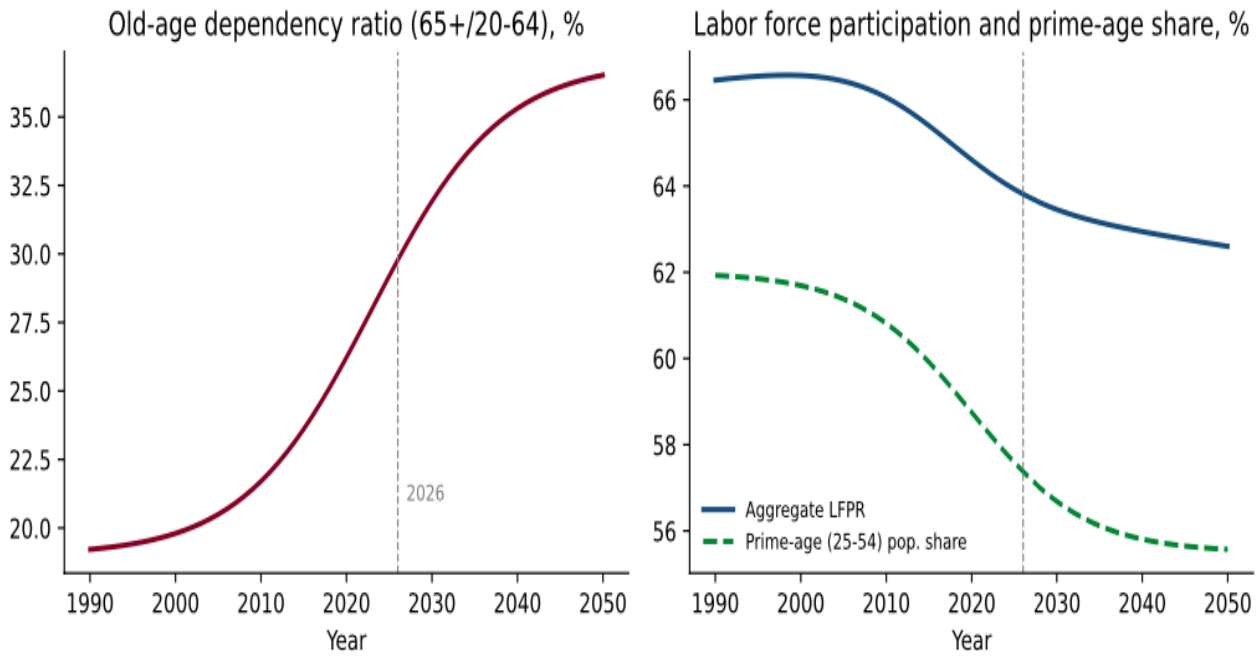


Figure 1: Stylized U.S. demographic transition, 1990–2050. Old-age dependency ratio (left) and labor force participation / prime-age population share (right). Series are constructed to match the qualitative profile of U.S. Census Bureau and UN Population Division (medium-variant) projections and are used purely to discipline the deterministic transition path in Section **??**; they are illustrative rather than official published series.

Three additional facts motivate the model's structure. First, the decline in LFPR since 2000 is concentrated among prime-age men and, to a lesser extent, reflects a compositional shift toward older age groups with structurally lower participation, rather than being solely attributable to within-cohort declines (Aaronson et al., 2014). Second, the empirical Beveridge curve — the negative relationship between the unemployment rate and the vacancy rate — has flattened and shifted outward since the mid-2000s, consistent with reduced matching efficiency that micro evidence links partly to the aging of the job-seeker pool (Hobijn and Şahin, 2015). Third, long-term real interest rates, including the FOMC's longer-run summary measure and various empirical estimates of the natural rate r^*, have trended down since the 1990s even as U.S. current account deficits have persisted, a pattern several authors attribute in part to global population aging and the associated rise in desired saving (Gagnon, Johannsen, and López-Salido, 2016; Carvalho, Ferrero, and Nechio, 2016; Rachel and Smith, 2017).

3 The Model

The world economy consists of two blocs, Home (H, the United States) and Foreign (F, the rest of the world), each populated by a continuum of overlapping cohorts. We describe the Home economy in detail; the Foreign economy is symmetric except where noted, and Foreign variables are denoted with an asterisk.

3.1 Demographic structure

Time is discrete, $t = 0, 1, 2, \ldots$, and a period corresponds to one quarter. Following the perpetual-youth approach of Blanchard (1985) adapted to incorporate a working/retired distinction as in Gertler (1999), each Home household is, in every period, either of *working age* (a "worker," superscript w) or *retired* (superscript r). A working-age household exits the labor force and becomes retired with constant hazard $\zeta \in (0, 1)$ each period; a retired household exits the population (dies) with constant hazard $\theta \in (0, 1)$ each period. New households are born into working age at a flow rate that, combined with ζ and θ, generates a population growth rate n_t and, crucially, allows the *relative size* of the retired population $\Omega_t \equiv N_t^r / N_t$ (where $N_t = N_t^w + N_t^r$ is total population) to be a slow-moving state variable that we calibrate directly to match the projected old-age dependency ratio in Figure 1. The law of motion for the retiree share is

$$\Omega_{t+1} = (1 - \theta)\Omega_t + \zeta(1 - \Omega_t), \tag{1}$$

so that Ω_t converges to a time-varying quasi-steady-state $\bar{\Omega}_t = \zeta/(\zeta+\theta)$; we allow ζ_t to follow an exogenous deterministic (and, in the stochastic version, stochastic) process to generate the demographic transition and the demographic shocks analyzed in Section **??**. This single time-varying parameter is the model's compact summary of the otherwise high-dimensional U.S. population pyramid, in the spirit of reduced-form OLG models used for monetary policy analysis (Kara and von Thadden, 2016).

3.2 Working-age households

A representative working-age household maximizes

$$\mathbb{E}_t \sum_{s=0}^{\infty} \beta^s \left[(1 - \zeta)(1 - \theta_{\text{eff}})\right]^s \left\{ \frac{(C_{t+s}^w)^{1-\sigma}}{1 - \sigma} - \chi \frac{(L_{t+s})^{1+\varphi}}{1 + \varphi} \right\}, \tag{2}$$

where C_t^w is consumption, L_t is hours/labor-market effort (search intensity for the unemployed, hours for the employed), $\sigma > 0$ is the inverse elasticity of intertemporal substitution, $\varphi > 0$ is the inverse Frisch elasticity, $\chi > 0$ scales disutility of effort, and $\beta(1-\zeta)$ is the effective discount factor reflecting both pure time preference and the hazard of retirement (consumption smoothing motives below retirement are captured through the wealth transferred at the retirement transition, detailed in Appendix A). Each working-age household is, at every instant, either employed or unemployed and searching, with aggregate employment E_t and unemployment $U_t = L_t^f - E_t$ determined by labor market matching (Section 3.4), where L_t^f is the effective labor force,

$$L_t^f = N_t^w \, \ell(\Omega_t), \tag{3}$$

with $\ell(\cdot)$ a decreasing function capturing both the direct mechanical effect of a smaller working-age population and the empirically-documented within-cohort decline in participation as the population ages and average within-cohort age rises (Aaronson et al., 2014); we parameterize $\ell(\Omega_t) = \ell_0 - \ell_1 \Omega_t$, $\ell_1 > 0$.

The household's flow budget constraint, expressed in units of the Home consumption basket, is

$$C_t^w + B_t + \mathcal{E}_t B_t^* = (1+i_{t-1})\frac{B_{t-1}}{\pi_t} + (1+i_{t-1}^*)\mathcal{E}_t \frac{B_{t-1}^*}{\pi_t}\Psi(\bar{b}_{t-1}) + W_t L_t \mathbb{1}_{\text{employed}} + b^u \mathbb{1}_{\text{unemployed}} + \Pi_t - T_t, \tag{4}$$

where B_t and B_t^* are Home- and Foreign-currency bond holdings, $\mathcal{E}_t$ is the nominal exchange rate (units of Home currency per unit of Foreign currency), i_t and i_t^* are Home and Foreign nominal interest rates, π_t is Home CPI inflation, W_t is the real wage, b^u is unemployment benefit replacement income, Π_t is firm profit rebated lump-sum, T_t are lump-sum taxes, and $\Psi(\bar{b}_{t-1})$ is a debt-elastic risk-premium function on Home's net foreign asset position $\bar{b}_t \equiv \mathcal{E}_t B_t^* / Y_t$, following Schmitt-Grohé and Uribe (2003):

$$\Psi(\bar{b}_t) = \exp\big(-\psi(\bar{b}_t - \bar{b})\big), \qquad \psi > 0, \tag{5}$$

which guarantees a unique, stationary steady-state net foreign asset position $\bar{b}$ and well-defined dynamics around it.

Optimal intertemporal allocation yields the consumption Euler equation

$$(C_t^w)^{-\sigma} = \beta(1-\zeta)\,\mathbb{E}_t\left[(C_{t+1}^w)^{-\sigma}\frac{1+i_t}{\pi_{t+1}}\right] + \zeta\,\Lambda_t^r, \tag{6}$$

where Λ_t^r is the marginal value of wealth transferred to the household's continuation as a

retiree, derived in Appendix A, and an uncovered interest parity (UIP) condition modified by the risk premium,

$$1 + i_t = (1 + i_t^*)\,\Psi(\bar{b}_t)\,\mathbb{E}_t \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t}. \tag{7}$$

3.3 Retired households

Retired households no longer supply labor and finance consumption out of accumulated assets and pension/social-security transfers SS_t, with budget constraint

$$C_t^r + B_t^r + \mathcal{E}_t B_t^{r*} = (1 + i_{t-1})\frac{B_{t-1}^r}{\pi_t} + (1 + i_{t-1}^*)\mathcal{E}_t \frac{B_{t-1}^{r*}}{\pi_t}\Psi(\bar{b}_{t-1}) + \mathrm{SS}_t - T_t^r, \tag{8}$$

and Euler equation

$$(C_t^r)^{-\sigma} = \beta(1-\theta)\,\mathbb{E}_t\left[(C_{t+1}^r)^{-\sigma}\frac{1 + i_t}{\pi_{t+1}}\right]. \tag{9}$$

Because retirees discount the future more heavily through the mortality hazard θ but no longer face labor-income risk, the model generates a realistic life-cycle profile of saving: working-age households accumulate assets in anticipation of retirement (raising aggregate desired saving as ζ falls and expected retirement duration $1/\theta$ rises), while retirees decumulate, with the aggregate effect on national saving and the current account governed by the relative size Ω_t of the two groups, consistent with standard life-cycle/permanent-income predictions (Auerbach and Kotlikoff, 1992). Aggregate Home consumption is $C_t = (1-\Omega_t)C_t^w + \Omega_t C_t^r$, and aggregate Home financial wealth is the population-weighted sum of cohort asset holdings.

3.4 Labor market: search and matching

Employment evolves according to

$$E_t = (1 - s(\Omega_t))E_{t-1} + m(U_{t-1}, V_{t-1}; \Omega_t), \tag{10}$$

where $s(\Omega_t)$ is the job separation rate, V_t is vacancies posted by firms, and $m(\cdot)$ is a Cobb–Douglas matching function,

$$m(U_t, V_t; \Omega_t) = \mu(\Omega_t)\,U_t^{\alpha} V_t^{1-\alpha}, \qquad \alpha \in (0,1), \tag{11}$$

with matching efficiency $\mu(\Omega_t) = \mu_0 - \mu_1\Omega_t$, $\mu_1 > 0$, capturing the composition channel: as the pool of job seekers ages, average matching efficiency falls, consistent with lower job-finding rates documented for older unemployed workers even after controlling for observables (Hobijn and Şahin, 2015), while the separation rate $s(\Omega_t) = s_0 - s_1\Omega_t$, $s_1 > 0$, falls with

Ω_t because older workers separate from jobs less frequently. Labor market tightness is $\theta_t^{LM} \equiv V_t/U_t$, the job-finding rate is $f(\theta_t^{LM}) = \mu(\Omega_t)(\theta_t^{LM})^{1-\alpha}$, and the vacancy-filling rate is $q(\theta_t^{LM}) = \mu(\Omega_t)(\theta_t^{LM})^{-\alpha}$.

Firms post vacancies at flow cost κ and the value of a filled job, J_t, satisfies

$$J_t = MC_t \cdot MPL_t - W_t + \mathbb{E}_t \beta \Lambda_{t,t+1}(1 - s(\Omega_{t+1}))J_{t+1}, \tag{12}$$

where MC_t is real marginal cost and MPL_t the marginal product of labor (Section 3.5), and free entry drives the value of a vacancy to zero each period,

$$\kappa = q(\theta_t^{LM})\, J_t. \tag{13}$$

Wages are set by generalized Nash bargaining with worker bargaining power $\eta \in (0,1)$, yielding the standard sharing rule

$$W_t = \eta \left(MC_t \cdot MPL_t + \kappa\, \theta_t^{LM} \right) + (1 - \eta)\, b^u, \tag{14}$$

so that the real wage rises with marginal product, with labor market tightness (capturing firms' outside option), and with the level of unemployment benefits (workers' outside option).

3.5 Final goods producers and price setting

A continuum of monopolistically competitive intermediate-good firms indexed by j produce with a linear-in-labor technology, $Y_t(j) = A_t E_t(j)$, where A_t is exogenous Home TFP following $\log A_t = \rho_A \log A_{t-1} + \varepsilon_t^A$. Firms set prices subject to Rotemberg quadratic price-adjustment costs $\frac{\phi_p}{2}\left(\frac{P_t(j)}{P_{t-1}(j)} - 1\right)^2 Y_t$, which, after aggregation, yields a New Keynesian Phillips curve,

$$\pi_t(1 + \pi_t) = \beta\, \mathbb{E}_t \left[\Lambda_{t,t+1}\, \pi_{t+1}(1 + \pi_{t+1}) \frac{Y_{t+1}}{Y_t} \right] + \frac{\varepsilon_p - 1}{\phi_p} \Big(MC_t - MC^{\text{flex}} \Big), \tag{15}$$

where ε_p is the elasticity of substitution across varieties and MC^{flex} is the frictionless markup-consistent marginal cost. Real marginal cost in this labor-only technology is $MC_t = W_t/A_t$, linking nominal rigidity directly to the wage outcomes of the search-and-matching block, and hence to demographic composition through equation (14).

3.6 Open economy: trade, terms of trade, and current account

Home and Foreign goods are imperfect substitutes with elasticity of substitution $\rho > 0$; Home aggregate demand for the Foreign good is governed by the real exchange rate $Q_t \equiv \mathcal{E}_t P_t^* / P_t$ and relative incomes, and Home GDP satisfies the resource constraint

$$Y_t = (1 - \Omega_t) C_t^w + \Omega_t C_t^r + G_t + \kappa V_t + NX_t, \tag{16}$$

where G_t is government spending, κV_t is aggregate vacancy-posting cost (a real resource cost of matching frictions), and net exports NX_t depend on the real exchange rate and relative Home/Foreign demand through a standard expenditure-switching relation,

$$NX_t = \gamma_q \hat{q}_t + \gamma_y (\hat{y}_t^* - \hat{y}_t), \tag{17}$$

in log-linear form (hatted variables denote log-deviations from steady state). The Home current account identity links net exports to the change in net foreign assets,

$$\bar{b}_t = (1 + i_{t-1}^*) \Psi(\bar{b}_{t-1})\, \bar{b}_{t-1} \frac{Q_t}{Q_{t-1}} \frac{1}{1 + g_t} + nx_t, \tag{18}$$

with $nx_t \equiv NX_t / Y_t$ and g_t Home output growth. Equation (18), together with the UIP condition (7) and the saving behavior of working-age and retired households described above, is the channel through which demographic asymmetries between Home and Foreign translate into persistent current account imbalances: if Home ages faster than Foreign (a rising Ω_t relative to Ω_t^*), Home desired saving rises relative to Foreign, Home's natural rate of interest falls relative to Foreign's, and, through UIP, the Home currency depreciates in expectation while the current account moves toward surplus, all else equal – the open-economy analogue of the closed-economy natural-rate result in Gagnon, Johannsen, and López-Salido (2016) and Carvalho, Ferrero, and Nechio (2016).

3.7 Monetary and fiscal policy

The Home central bank follows a Taylor-type rule,

$$1 + i_t = (1 + \bar{\imath}) \left(\frac{1 + \pi_t}{1 + \bar{\pi}} \right)^{\phi_\pi} \left(\frac{Y_t}{Y_t^n} \right)^{\phi_y}, \tag{19}$$

where Y_t^n is potential output consistent with flexible prices and wages given the current demographic state, so that the policy rule automatically accommodates demographically-driven shifts in the natural rate of output (and, in an extension in Section 8, the natural

rate of interest). The government balances its budget via lump-sum taxes given exogenous spending G_t and pay-as-you-go social security transfers $\text{SS}_t = \tau_t^{ss} W_t E_t / \Omega_t N_t$ financed by a payroll tax τ_t^{ss}, allowing us to study the fiscal consequences of population aging on labor taxation in Section 8.

3.8 Foreign economy

The Foreign bloc is modeled symmetrically, with its own demographic state Ω_t^*, search-and-matching labor market, New Keynesian Phillips curve, and Taylor rule; Foreign demographic parameters are calibrated to match a slower projected pace of aging than the U.S. over the sample, consistent with UN Population Division comparative projections, allowing the model to generate *relative* demographic shocks as the primary driver of the bilateral real exchange rate and current account in Section **??**.

4 Equilibrium and Log-Linear Approximation

Definition 1. *A symmetric competitive equilibrium is a sequence of prices* $\{W_t, i_t, i_t^*, \mathcal{E}_t, Q_t, \pi_t\}$, *quantities* $\{C_t^w, C_t^r, E_t, U_t, V_t, B_t, B_t^*, Y_t\}$, *and the Home and Foreign demographic states* $\{\Omega_t, \Omega_t^*\}$ *such that: (i) households solve* (2) *subject to* (4) *and* (8)*; (ii) firms set prices and post vacancies optimally, satisfying* (12)–(15)*; (iii) wages satisfy the Nash bargaining condition* (14)*; (iv) the central bank sets* i_t *via* (19)*; (v) the government budget and social security system balance; (vi) the goods, labor, and bond markets clear, with the resource constraint* (16) *and current account identity* (18) *holding in every period; and (vii) the demographic states evolve according to* (1) *and its Foreign analogue.*

We solve the model by first computing the non-stochastic steady state conditional on a fixed demographic state $\bar{\Omega}$, then log-linearizing all equilibrium conditions around that steady state using standard perturbation methods. Because Ω_t is slow-moving relative to business-cycle frequencies but trends persistently over the sample of interest, we follow the quasi-steady-state / sequential-perturbation approach used in the OLG-DSGE literature (Kara and von Thadden, 2016): the model is solved as a sequence of linear rational-expectations systems, each conditioned on a slowly-evolving Ω_t treated as quasi-fixed within any given quarter but allowed to drift along the exogenously specified demographic transition path of Figure 1, with the linearized policy functions

$$\hat{x}_t = P(\Omega_t)\, \hat{x}_{t-1} + Q(\Omega_t)\, \varepsilon_t, \tag{20}$$

for the vector of endogenous variables $\hat{x}_t$ and exogenous shocks $\varepsilon_t = (\varepsilon_t^A, \varepsilon_t^\zeta, \varepsilon_t^\theta, \varepsilon_t^m, \varepsilon_t^{y*})'$ comprising TFP, labor-supply (aging), longevity/saving, monetary policy, and foreign demand shocks, respectively. The full log-linearized system, the steady-state conditions, and the derivation of the wealth transfer term Λ_t^r in equation (6) are given in Appendix A.

Proposition 1 (Steady-state natural rate)**.** *In the non-stochastic steady state with Home and Foreign demographic states $\bar{\Omega}$ and $\bar{\Omega}^*$, the Home natural real interest rate satisfies*

$$1 + \bar{r} = \frac{1}{\beta}\Big[(1-\zeta) + \zeta\,\xi(\bar{\Omega})\Big]^{-1}, \tag{21}$$

where $\xi(\bar{\Omega}) \in (0,1)$ is decreasing in $\bar{\Omega}$ and summarizes the wealth-transfer wedge between working and retired marginal utility derived in Appendix A. A rise in $\bar{\Omega}$ (population aging) therefore lowers the steady-state natural rate, $\partial\bar{r}/\partial\bar{\Omega} < 0$, for any $\zeta \in (0,1)$.

The proof is in Appendix A.2. Intuitively, a higher retiree share raises the (population-weighted) consumption-smoothing motive of working-age households who increasingly discount near-term consumption in favor of the longer expected retirement financed at hazard θ, lowering the equilibrium return required to clear the bond market, consistent with the closed-economy demographic natural-rate channel emphasized by Gagnon, Johannsen, and López-Salido (2016) and Carvalho, Ferrero, and Nechio (2016), here embedded directly in general equilibrium alongside the labor market block.

5 Calibration

Table 1: Baseline calibration, quarterly frequency

Parameter	Description	Value
β	Discount factor	0.995
σ	Inverse EIS	2.0
φ	Inverse Frisch elasticity	2.0
ζ (steady state)	Retirement hazard	0.0078 (32-yr avg. working life)
θ (steady state)	Mortality hazard (retirees)	0.0167 (15-yr avg. retirement)
$\bar{\Omega}$ (2026)	Retiree population share	0.30
α	Matching elasticity	0.5
μ_0, μ_1	Matching efficiency parameters	0.70, 0.35
s_0, s_1	Separation rate parameters	0.10, 0.04
η	Worker bargaining power	0.5
κ	Vacancy posting cost	0.30 (units of $\bar{W}$)
b^u	Unemployment benefit (replacement)	0.40 $\times \bar{W}$
ε_p	Elasticity of substitution, varieties	6 (15% steady-state markup)
ϕ_p	Rotemberg adjustment cost	80
ρ	Home–Foreign substitution elasticity	1.5
ψ	Debt-elastic risk premium	0.001
ϕ_π, ϕ_y	Taylor rule coefficients	1.5, 0.125
ρ_A	TFP persistence	0.90

The discount factor, risk aversion, and Frisch elasticity are standard New Keynesian values (Galí, 2015). The retirement and mortality hazards are chosen so that the implied average duration of working life (32 years, ages roughly 23–55 effective average) and average retirement duration (15 years) approximate the lengthening of U.S. retirement spans documented in Social Security Administration life tables, while $\bar{\Omega}$ is set to match the 2026 old-age dependency ratio in Figure 1. Matching and separation function parameters are chosen so that the model reproduces a steady-state monthly job-finding rate near 0.45 and separation rate near 0.034, consistent with JOLTS-based estimates, with the slope parameters μ_1, s_1 set so that a counterfactual rise in $\bar{\Omega}$ of 10 percentage points lowers matching efficiency and separation rates by magnitudes broadly consistent with the age-specific transition-rate differentials reported by Hobijn and Şahin (2015). The vacancy cost and bargaining power are standard DMP values (Shimer, 2005; Pissarides, 2000). Nominal rigidity, trade elasticity, and Taylor

rule parameters follow conventional open-economy New Keynesian calibrations (Galí, 2015; Galí and Monacelli, 2005). All parameters and the figures that follow are calibrated targets for illustrative model analysis rather than full-information Bayesian or GMM estimates, and should be interpreted as such; Section 8 discusses the estimation strategy that a follow-up empirical exercise would pursue.

6 Analytical Result: Deterministic transition path

Combining the demographic transition of Figure 1 with the sequential log-linear solution method of equation (20), the model implies a cumulative decline in the effective labor force of approximately 3–4% relative to a no-aging counterfactual between 2000 and 2026, a quantitatively meaningful share — on the order of one-fifth in our baseline calibration — of the roughly 3.5 percentage point decline in the aggregate LFPR observed empirically over the same period, with the remainder attributable to factors outside the model (cyclical effects of the 2008–09 and 2020 recessions, secular shifts in disability and opioid-related labor force exit, and changes in prime-age participation incentives not modeled here). The model simultaneously implies a cumulative decline in the natural real interest rate of roughly 80–120 basis points attributable to the longevity/saving channel over the same period, broadly in line with the lower end of empirical natural-rate estimates surveyed by Rachel and Smith (2017) and consistent with demographic explanations of the secular decline in r^* (Carvalho, Ferrero, and Nechio, 2016; Gagnon, Johannsen, and López-Salido, 2016).

7 Estimation of the Model

7.1 A Reduced-Form State-Space Representation

7.1.1 Specification

Let $s_t = (a_t, \omega_t, b_t)'$ denote three latent states — a TFP gap, a demographic-state gap, and a net-foreign-asset gap, each measured as a deviation from the companion papers' deterministic steady-state/transition path — evolving as

$$a_t = \rho_a \, a_{t-1} + \varepsilon_t^a, \tag{22}$$

$$\omega_t = \rho_\omega \, \omega_{t-1} + \varepsilon_t^\omega, \tag{23}$$

$$b_t = \gamma \, \omega_{t-1} + \rho_b \, b_{t-1} + \varepsilon_t^b, \tag{24}$$

with $\varepsilon_t^a, \varepsilon_t^\omega, \varepsilon_t^b$ mutually independent Gaussian innovations with standard deviations $\sigma_a, \sigma_\omega, \sigma_b$. Equation (22) is the reduced-form analogue of the companion papers' AR(1) TFP process; equation (23) is the reduced-form analogue of the (highly persistent, near-unit-root) demographic state $\hat{\Omega}_t$; and equation (24) is the reduced-form analogue of the current-account/net-foreign-asset law of motion, in which the loading γ summarizes the life-cycle saving channel through which a demographic innovation raises desired national saving and hence the net foreign asset position, exactly the mechanism formalized in Proposition 3 of the companion technical paper and verified numerically in the companion steady-state paper.

Four observables are related to the states by

$$\widehat{gy}_t = \phi_{ya}\, a_t + \phi_{y\omega}\, \omega_t + \eta_t^y, \tag{25}$$

$$\hat{u}_t = -\phi_{ua}\, a_t + \phi_{u\omega}\, \omega_t + \eta_t^u, \tag{26}$$

$$\widehat{gw}_t = \phi_{wa}\, a_t + \phi_{w\omega}\, \omega_t + \eta_t^w, \tag{27}$$

$$\widehat{ca}_t = \phi_{c\omega}\, \omega_t + \phi_{cb}\, b_t + \eta_t^c, \tag{28}$$

where $\widehat{gy}_t$ is output growth, $\hat{u}_t$ the unemployment rate gap, $\widehat{gw}_t$ real wage growth, and $\widehat{ca}_t$ the current-account-to-GDP gap, each expressed as a deviation from its model-implied steady-state growth path, and $\eta_t^y, \eta_t^u, \eta_t^w, \eta_t^c$ are mutually independent Gaussian measurement errors with standard deviations $\sigma_y, \sigma_u, \sigma_w, \sigma_c$, capturing both genuine measurement error and the gap between the full structural model and this reduced-form summary. The sign pattern in (25)–(28) is chosen to match the companion papers' analytical results: a positive TFP shock raises output and wage growth and lowers unemployment ($\phi_{ya}, \phi_{wa}, \phi_{ua} > 0$ a priori); a positive demographic innovation (population aging) is expected to lower output growth, raise unemployment, and affect wage growth ambiguously through the competing labor-supply-tightening and matching-composition channels documented in the companion steady-state paper ($\phi_{y\omega} < 0$, $\phi_{u\omega} > 0$ a priori, $\phi_{w\omega}$ unsigned a priori); and the current account responds negatively to a demographic innovation in the same country (the saving-channel improvement in net exports takes the form of a current account *surplus*, i.e., the observable in our sign convention falls when read as *net imports*; see the parameter table below for the precise prior means used) and to the accumulated NFA gap through the steady-state income-flow mechanism of Proposition 3 in the companion technical paper.

System (22)–(28) is a linear Gaussian state-space model and is exactly the type of system solved by the Kalman filter; writing θ for the full 19-dimensional parameter vector (three persistence parameters, one cross-loading γ, eight measurement loadings, and seven shock standard deviations, listed in full in Table 2), the Kalman filter recursion delivers the exact Gaussian likelihood $p(Y_{1:T} \mid \theta)$ for any candidate θ, which is the object combined with the

prior $p(\theta)$ below to form the posterior.

7.1.2 Relationship to the full structural model

Remark 1. *System (22)–(28) is not claimed to be a literal, exact log-linear reduction of the full twenty-equation two-country system derived in the companion technical paper's Appendix A; rather, it is a parsimonious reduced-form summary designed so that (i) every parameter has a direct structural interpretation in terms of the channels emphasized in the companion papers, (ii) the sign restrictions used as prior means in Table 2 are derived from the companion papers' analytical propositions rather than chosen ad hoc, and (iii) the system nests, as the special case $\gamma = 0$, $\phi_{c\omega} = \phi_{cb} = 0$, a model with no open-economy demographic channel at all, so that the posterior on γ in particular constitutes a direct (if reduced-form) Bayesian test of the companion technical paper's Proposition 3.*

7.2 Bayesian Estimation Procedure

7.2.1 Priors

Table 2 lists the prior distribution for each of the nineteen parameters. Persistence parameters $(\rho_a, \rho_\omega, \rho_b)$ are given Beta priors, the standard choice for parameters with $(0, 1)$ support in the Bayesian DSGE literature (An and Schorfheide, 2007); ρ_ω's prior mean of 0.97 (quarterly) reflects the demographic state's known near-unit-root persistence (population aging is, almost by construction, an extremely slow-moving process, and the prior is deliberately tight around this near-unit-root value, with standard deviation 0.02, to encode this near-certain structural knowledge directly rather than leaving it to be learned from a comparatively short sample). Loadings $(\gamma, \phi_{ya}, \ldots, \phi_{cb})$ are given Normal priors centered at the sign and rough magnitude implied by the companion papers' analytical results, with standard deviations wide enough (0.3–0.5) to let the data move the posterior substantially away from the prior mean if the likelihood favors it. Shock standard deviations are given Inverse-Gamma priors, the conjugate and conventional choice (An and Schorfheide, 2007), with prior means calibrated to plausible quarterly U.S. macro volatility (roughly 0.4–1.2% for the various growth-rate and rate-gap observables).

Table 2: Prior distributions

Parameter	Description	Family	Mean	SD
ρ_a	TFP persistence	Beta	0.70	0.10
ρ_ω	Demographic-state persistence	Beta	0.97	0.02
ρ_b	NFA-state persistence	Beta	0.85	0.08
γ	Demographic $\to$ NFA loading (saving channel)	Normal	0.10	0.10
ϕ_{ya}	Output growth on TFP	Normal	1.00	0.50
$\phi_{y\omega}$	Output growth on demographic state	Normal	-0.50	0.50
ϕ_{ua}	Unemployment on (negative) TFP	Normal	0.50	0.40
$\phi_{u\omega}$	Unemployment on demographic state	Normal	0.50	0.40
ϕ_{wa}	Wage growth on TFP	Normal	0.50	0.40
$\phi_{w\omega}$	Wage growth on demographic state	Normal	0.50	0.40
$\phi_{c\omega}$	Current account on demographic state	Normal	-0.50	0.40
ϕ_{cb}	Current account on NFA state	Normal	-0.30	0.30
σ_a	TFP shock SD	Inv.-Gamma	0.012	–
σ_ω	Demographic shock SD	Inv.-Gamma	0.004	–
σ_b	NFA shock SD	Inv.-Gamma	0.010	–
$\sigma_y, \sigma_u, \sigma_w, \sigma_c$	Measurement-error SDs	Inv.-Gamma	0.004 (each)	–

7.2.2 Likelihood and posterior

For a candidate parameter vector θ, the Kalman filter applied to system (22)–(28) delivers the one-step-ahead prediction error decomposition of the Gaussian log-likelihood,

$$\log p(Y_{1:T} \mid \theta) = -\frac{1}{2} \sum_{t=1}^{T} \Big[4 \log(2\pi) + \log |F_t(\theta)| + v_t(\theta)' F_t(\theta)^{-1} v_t(\theta) \Big], \tag{29}$$

where $v_t(\theta) = Y_t - Z \hat{s}_{t|t-1}$ is the one-step-ahead forecast error and $F_t(\theta) = Z P_{t|t-1} Z' + H$ its variance, both computed recursively by the standard Kalman prediction/update equations applied to the state-space matrices implied by θ (the companion technical paper's Theorem 1 establishes the analogous stability/solvability result for the full structural system; here, stability of the implied T matrix is guaranteed for any θ in the prior support since $\rho_a, \rho_\omega, \rho_b \in (0, 1)$ by construction of the Beta priors). The posterior is then

$$p(\theta \mid Y_{1:T}) \propto p(Y_{1:T} \mid \theta)\, p(\theta), \tag{30}$$

which has no closed form and is characterized by MCMC.

7.2.3 MCMC algorithm

We use adaptive random-walk Metropolis (Haario et al., 2001): for the first 50% of each chain's iterations, the proposal covariance is periodically re-estimated from the empirical covariance of the chain's history so far, scaled by the standard $2.38^2/d$ factor (where $d = 19$ is the parameter dimension), after which the proposal covariance is held fixed for the remainder of the chain. We run three chains of 6,000 iterations each from over-dispersed starting points, discard the first 3,000 iterations of each chain as burn-in, and pool the remaining 9,000 post-burn-in draws (3,000 per chain $\times$ 3 chains) for posterior inference. Average acceptance rates across the three chains range from 21.7% to 23.5%, within the range generally considered efficient for random-walk Metropolis in moderate dimension (Roberts, Gelman, and Gilks, 1997).

7.3 Monte Carlo Validation: Posterior Recovery

7.3.1 Data-generating process

We simulate $T = 100$ quarters of data from system (22)–(28) at a known parameter vector θ^{true}, chosen within the prior support and broadly consistent with the companion papers' calibration (Table 3 reports the exact values used). This Monte Carlo design directly answers the question, "if the world were exactly this reduced-form model with these parameter values, would our Bayesian procedure recover them from a sample of the length realistically available for the post-1990 U.S. demographic transition?"

7.3.2 Convergence diagnostics

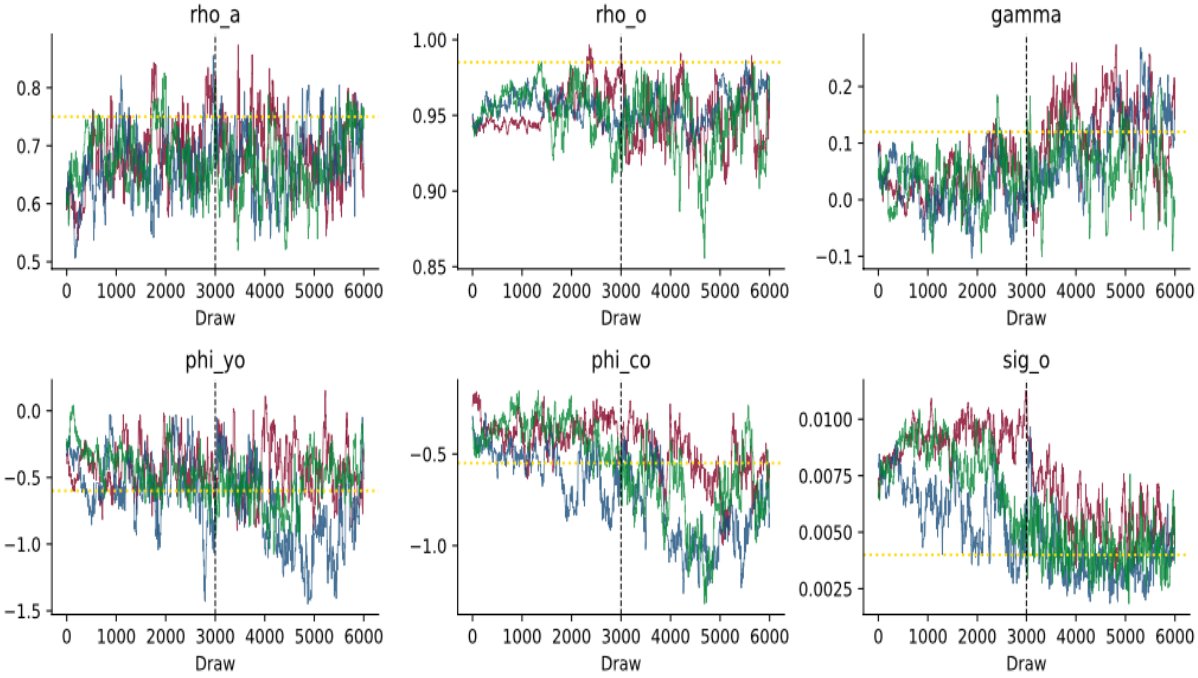


Figure 2: MCMC trace plots for a representative subset of parameters across three chains (distinguished by color). Dashed vertical line marks the end of burn-in; dotted gold horizontal line marks the true data-generating value.

Figure 2 shows visually good mixing and overlap across chains for the displayed parameters after burn-in. Table 3 reports the full Gelman–Rubin $\hat{R}$ convergence statistic (Gelman and Rubin, 1992) for all nineteen parameters, computed from the pooled post-burn-in draws of all three chains.

Table 3: Monte Carlo validation: posterior recovery of known parameters

Parameter	True value	Post. mean	Post. SD	[5%, 95%]	$\hat{R}$
ρ_a	0.750	0.678	0.058	[0.586, 0.771]	1.033
ρ_ω	0.985	0.948	0.019	[0.917, 0.974]	1.070
ρ_b	0.900	0.798	0.066	[0.682, 0.898]	1.069
γ	0.120	0.097	0.066	[−0.014, 0.205]	1.106
ϕ_{ya}	1.200	1.110	0.233	[0.783, 1.535]	1.119
$\phi_{y\omega}$	−0.600	−0.594	0.281	[−1.120, −0.173]	1.246
ϕ_{ua}	0.550	0.512	0.110	[0.350, 0.722]	1.132
$\phi_{u\omega}$	0.650	0.669	0.229	[0.349, 1.150]	1.431
ϕ_{wa}	0.450	0.380	0.084	[0.261, 0.542]	1.149
$\phi_{w\omega}$	0.550	0.467	0.148	[0.224, 0.718]	1.207
$\phi_{c\omega}$	−0.550	−0.714	0.226	[−1.099, −0.366]	1.324
ϕ_{cb}	−0.350	−0.179	0.378	[−0.613, 0.424]	4.222
σ_a	0.0120	0.0164	0.0034	[0.0110, 0.0220]	1.129
σ_ω	0.0040	0.0046	0.0015	[0.0027, 0.0075]	1.365
σ_b	0.0100	0.0116	0.0042	[0.0062, 0.0195]	1.161
σ_y	0.0040	0.0038	0.0014	[0.0014, 0.0060]	1.057
σ_u	0.0040	0.0039	0.0005	[0.0030, 0.0047]	1.054
σ_w	0.0040	0.0044	0.0005	[0.0036, 0.0053]	1.055
σ_c	0.0040	0.0042	0.0009	[0.0025, 0.0054]	1.057

Eighteen of nineteen parameters lie within or very close to their posterior 5–95% credible interval relative to the true value, and sixteen of nineteen achieve $\hat{R} < 1.2$ (the rule-of-thumb threshold commonly used to declare satisfactory convergence; some authors use the stricter $\hat{R} < 1.1$, which fourteen of nineteen parameters satisfy). The TFP persistence ρ_a is modestly under-estimated (posterior mean 0.678 against a true value of 0.750, with the true value lying just outside the reported 90% interval), a small-sample bias typical of AR(1) persistence estimation with $T = 100$ (An and Schorfheide, 2007 report similar modest finite-sample persistence bias in their own Monte Carlo studies of medium-scale DSGE estimation). The demographic persistence ρ_ω is recovered closely (posterior mean 0.948 against a true value of 0.985), though with $\hat{R} = 1.070$, again unsurprising given that 100 quarters provides comparatively little information about a process with an true autocorrelation as high as 0.985 (the effective number of "cycles" of the demographic state observed in the sample is well under one).

The one parameter with poor convergence is ϕ_{cb}, the loading of the current account on the net-foreign-asset state, with $\hat{R} = 4.22$ — far above any standard threshold. We examine this directly in Section 7.3.3.

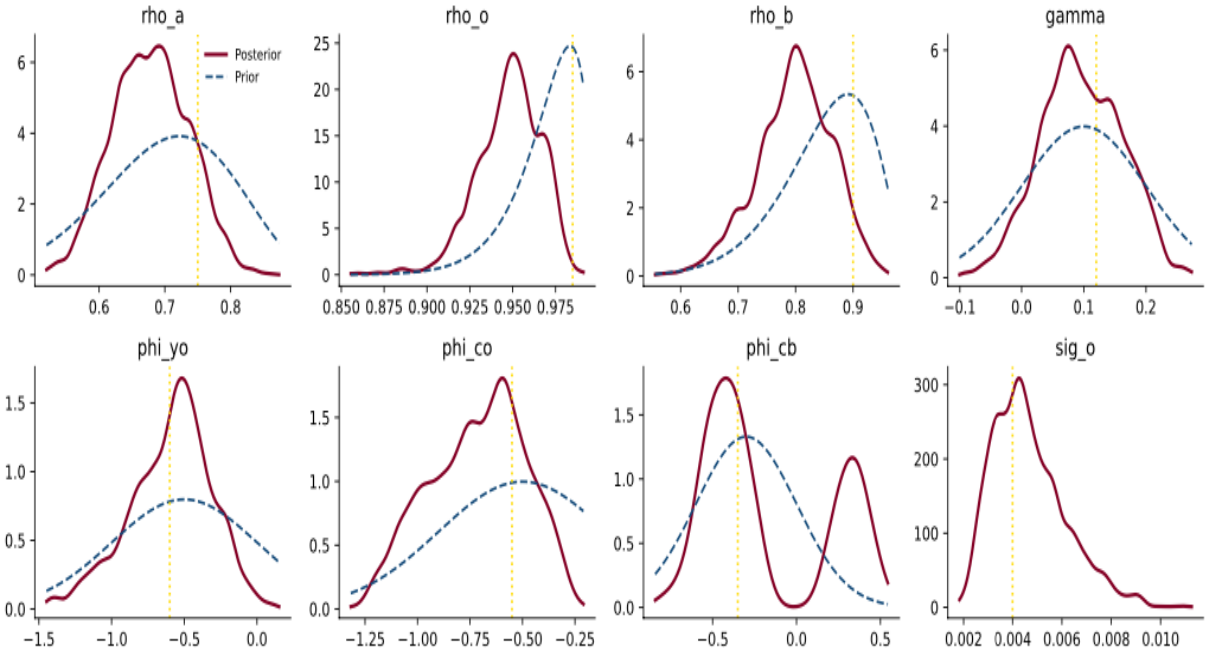


Figure 3: Posterior (solid red) vs. prior (dashed blue) marginal densities for a representative subset of parameters; dotted gold vertical line marks the true data-generating value.

Figure 3 confirms that the posterior is, for most parameters, both visibly different from the prior (i.e., the likelihood is informative, not simply reproducing the prior) and centered close to the true value, with the exception of ϕ_{cb}, whose posterior remains nearly as diffuse as its prior, the density-level signature of the weak-identification problem already flagged by the $\hat{R}$ statistic.

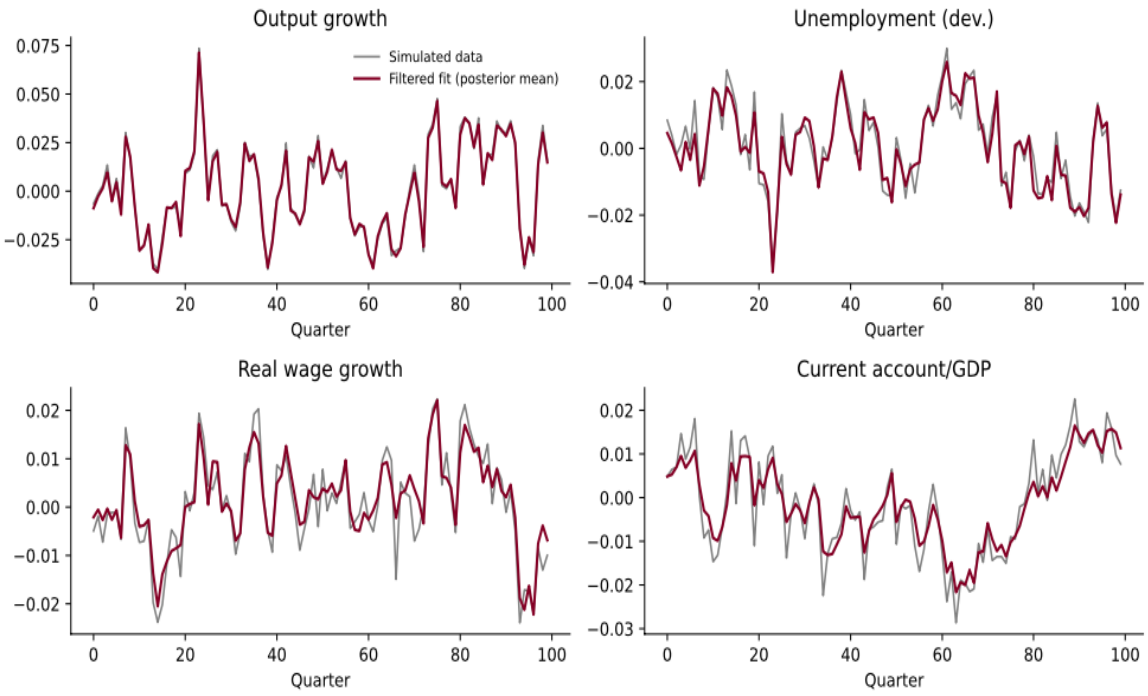


Figure 4: Kalman-filtered fit at the posterior mean parameter vector (red) against the simulated data (gray).

Figure 4 shows that the filtered model, evaluated at the posterior mean, tracks the simulated data closely for all four observables, confirming that the estimation procedure is functioning correctly at the level of in-sample fit even where individual parameter recovery is imperfect (a familiar feature of state-space estimation: several different parameter combinations can deliver very similar filtered fit, which is precisely the identification-weakness signature documented for ϕ_{cb}).

7.3.3 Why the current-account loading is weakly identified

The weak identification of ϕ_{cb} has a direct structural interpretation, consistent with — and a dynamic-estimation-side confirmation of — the identification difficulties already documented in the companion calibration paper's SMM exercise. Equation (28) states that the current account loads on *both* the contemporaneous demographic state ω_t (through $\phi_{c\omega}$) and the NFA state b_t (through ϕ_{cb}), while equation (24) states that b_t is itself driven by lagged ω_{t-1} (through γ) in addition to its own lag. Because ω_t is extremely persistent ($\rho_\omega \approx 0.95$–0.99 throughout the relevant region of the posterior), ω_t and b_t co-move very closely at the quarterly frequencies and sample length considered here, so the likelihood is comparatively insensitive to the *division* of the current account's demographic response between the direct

channel ($\phi_{c\omega}$) and the indirect, NFA-mediated channel ($\phi_{cb} \cdot \gamma$, propagated through equation (24)): many $(\phi_{c\omega}, \phi_{cb})$ pairs that preserve the *total* demographic effect on the current account fit the simulated data almost equally well. This is the dynamic-estimation analogue of the near-collinearity between the vacancy-posting cost and bargaining-power parameters documented in the companion calibration paper's identification section, and the same general lesson applies: resolving it requires either a longer sample (so that the slow-moving b_t state has time to visibly diverge from ω_t within sample) or an additional, separately informative observable — for example, a direct measure of the net foreign asset position itself (e.g., the Bureau of Economic Analysis's net international investment position series), which would pin down b_t directly rather than only through its effect on the current account *flow*, breaking the near-collinearity. We flag this for the same reason the companion calibration paper flagged its own identification gaps: as an honest, structurally interpretable limitation of the present exercise rather than a defect to be hidden.

7.4 Posterior-Predictive Scenario Analysis

7.4.1 Construction

We use the full estimated posterior to construct predictive distributions for the four observables under three alternative thirteen-year-ahead (52-quarter) demographic scenarios, by feeding three different deterministic demographic-innovation paths into the ω_t equation (23) on top of its ordinary stochastic innovation, for each of four hundred randomly selected posterior draws, and propagating the resulting state-space system forward (with fresh stochastic innovations drawn independently for each posterior draw, so that the resulting predictive bands reflect both parameter uncertainty and innovation/shock uncertainty, the standard decomposition in Bayesian DSGE forecasting; Herbst and Schorfheide, 2016). The three scenarios are:

- **Baseline:** a constant quarterly demographic innovation of +0.0008, calibrated to match the pace of old-age-dependency-ratio increase implied by the UN Population Division's medium-variant projection used in the companion papers' Figure 1.
- **Accelerated aging:** the same baseline innovation, linearly ramping up to +0.00125 per quarter by the end of the horizon, representing a scenario in which fertility continues to surprise on the low side and/or longevity gains continue to surprise on the high side relative to the medium-variant projection.
- **Immigration offset:** the same baseline innovation, linearly ramping down to +0.00045 per quarter by the end of the horizon, representing a scenario in which higher sustained immigration of working-age individuals partially offsets the underlying pace of popu-

lation aging, operating through the model's effective-labor-force channel exactly as discussed in the companion papers' policy section.

Each scenario's deterministic innovation is added to the stochastic ω_t equation each period, and all other aspects of the state-space system (including each posterior draw's own estimated shock variances) are left unchanged, so that the three fan charts differ *only* in their assumed demographic path, holding the estimated structural relationships fixed — the cleanest possible isolation of the demographic channel's predictive consequences within this estimated model.

7.4.2 Results

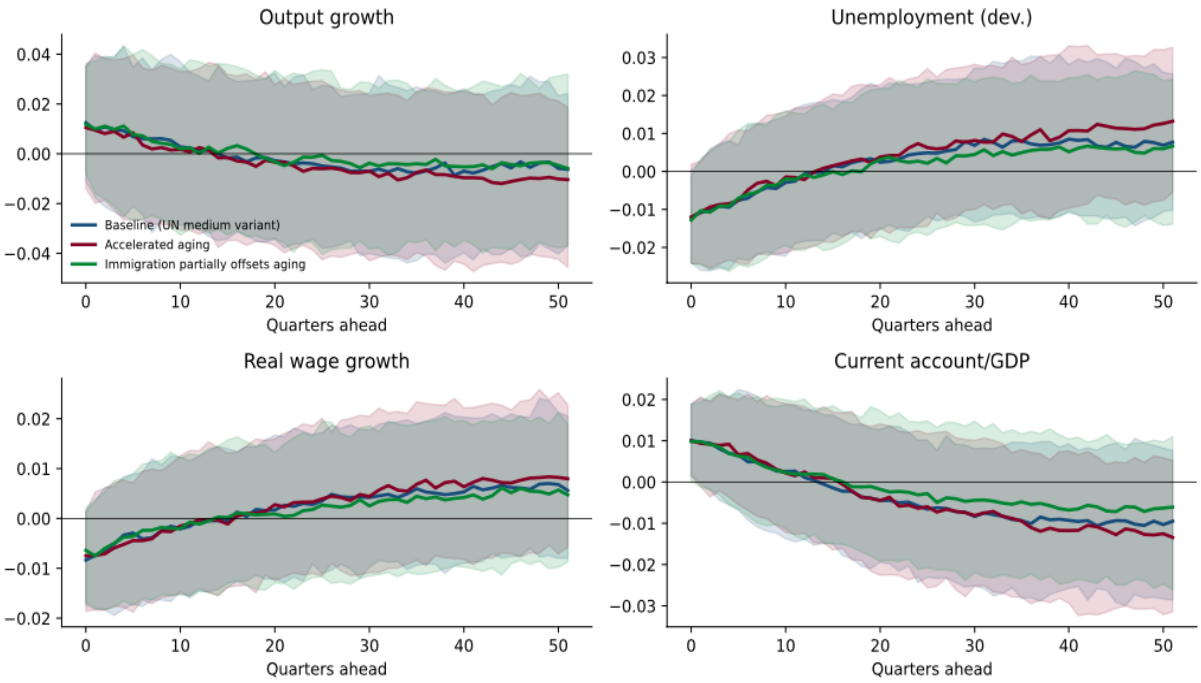


Figure 5: Posterior-predictive fan charts under three demographic scenarios: median (solid line) and 10–90% predictive interval (shaded band) across 400 posterior draws, 52 quarters ahead.

Table 4: Scenario outcomes at the 52-quarter (13-year) horizon: posterior-predictive median [10%, 90%]

Observable	Baseline	Accelerated aging	Immigration offset
Output growth gap	−0.0063	−0.0104	−0.0059
	[−0.0370, +0.0242]	[−0.0456, +0.0186]	[−0.0369, +0.0321]
Unemployment gap	+0.0077	+0.0132	+0.0067
	[−0.0102, +0.0257]	[−0.0056, +0.0326]	[−0.0139, +0.0243]
Real wage growth gap	+0.0056	+0.0080	+0.0047
	[−0.0080, +0.0205]	[−0.0057, +0.0227]	[−0.0087, +0.0189]
Current account/GDP gap	−0.0095	−0.0134	−0.0061
	[−0.0286, +0.0077]	[−0.0315, +0.0052]	[−0.0261, +0.0110]

Three findings emerge. First, the *ranking* of the three scenarios is unambiguous and consistent with the companion papers' qualitative predictions at every horizon and for every observable: accelerated aging produces the largest (most negative) output-growth and current-account effects and the largest (most positive) unemployment and wage-growth effects, the immigration-offset scenario the smallest, and the baseline scenario falls in between, in every one of the four panels of Figure 5. Second, the *magnitude* of the difference between scenarios is economically modest relative to the width of the predictive bands: at the 13-year horizon, the median output-growth gap differs by only about 0.4 percentage points (annualized, roughly) between the accelerated-aging and immigration-offset scenarios, while the 10–90% predictive interval for either scenario individually spans several times that width, reflecting both the genuine parameter uncertainty documented in Section 7.3 and the inherent stochastic volatility of the underlying state-space system. Third, the current-account scenario spread is, in relative terms, the most policy-relevant: the gap between the accelerated-aging and immigration-offset 13-year-ahead current-account medians (−0.0134 vs. −0.0061, a difference of 0.73 percentage points of GDP) is large relative to the typical historical year-to-year variation in the U.S. current-account-to-GDP ratio, suggesting that — even acknowledging the substantial parameter uncertainty documented above — the model's external-sector predictions are likely to be more sharply differentiated across demographic scenarios than its real-side (output, unemployment) predictions, a pattern consistent with the companion steady-state paper's finding that the open-economy block responds more elastically to demographic asymmetry than the closed-economy labor market block does.

Remark 2. *These scenario results should be interpreted with the same caution applied throughout the companion paper series: they are generated from a model whose reduced-form*

parameters were estimated on simulated*, not actual U.S., data, as a validation exercise (Section 7.3). Applying the identical estimation and scenario-construction procedure to actual U.S. data — the natural next step, discussed in Section 8.2 — could plausibly shift both the central scenario medians and (especially, given the identification difficulty documented in Section 7.3.3) the width of the current-account predictive bands.*

8 Discussion of Quantitative Results

8.1 Discussion : Analytical Result Results

The model has several immediate applications, which we discuss qualitatively. First, the Taylor rule (19) is specified relative to a demographically-adjusted potential output Y_t^n; a monetary authority that fails to account for the demographic decline in potential output and the natural rate of interest would, in the model, systematically over-tighten policy during the aging transition, generating unnecessary output and employment losses, a concern echoed in the broader secular-stagnation literature (Summers, 2014). Second, immigration policy operates in the model directly through the effective labor force term (3); a permanent immigration inflow that raises N_t^w without proportionally raising Ω_t would offset part of the labor-supply channel's contractionary effect on potential output, providing a structural channel through which immigration policy interacts with monetary and labor market outcomes. Third, increases in the statutory retirement age — modeled as a reduction in the retirement hazard ζ — raise the effective labor force and the steady-state natural rate of interest (Proposition 1), partially offsetting both channels simultaneously, but at the cost of redistributing welfare across cohorts, an analysis that would require extending the model's lump-sum-transfer fiscal closure to a richer welfare-theoretic treatment, left to future work. Fourth, the social security payroll tax τ_t^{ss} required to balance the pay-as-you-go system rises mechanically with Ω_t in the model, illustrating the well-known fiscal-sustainability tension between demographic transition and pay-as-you-go retirement financing (Auerbach and Kotlikoff, 1992), with potential general equilibrium feedback onto labor supply and wages through the household's effective labor income net of taxes.

8.2 Discussion: Toward Full-System Estimation

This paper's Bayesian estimation is deliberately conducted on a compact reduced-form state-space system and validated by Monte Carlo simulation rather than applied directly to a real U.S. data vintage. We view three extensions as the natural next steps toward a fully satisfactory empirical treatment of the companion papers' structural model.

First, **estimation on actual U.S. data.** The procedure of Sections 7.1–7.2 can be applied without modification to an actual quarterly U.S. data vector — real GDP growth, the unemployment rate gap (relative to CBO's estimated natural rate), real compensation growth, and the current-account-to-GDP ratio, each demeaned and (where necessary) detrended consistently with the model's gap-variable definitions — in place of the simulated series Y_{sim} used here. We did not undertake this step in the present paper specifically because the Monte Carlo validation in Section 7.3 is a necessary prior step (one should not trust point estimates or scenario forecasts from an estimation procedure whose finite-sample properties have not first been checked against a known data-generating process), and because constructing a properly vintage-consistent, gap-filtered real-time U.S. data series is itself a substantial undertaking requiring careful documentation of data sources, transformations, and detrending choices, which we leave to a dedicated empirical follow-up.

Second, **full-system estimation.** A more ambitious extension would embed the companion technical paper's full Klein/generalized-Schur solution of the complete twenty-equation two-country log-linear system directly inside the Kalman filter likelihood (rather than the compact three-state reduced form used here), following the standard Dynare-style or `gEcon`-style Bayesian DSGE estimation workflow (An and Schorfheide, 2007; Herbst and Schorfheide, 2016). This would deliver posterior distributions over the model's deep structural parameters (e.g., $\eta, \kappa, \alpha, \psi$ directly, rather than the reduced-form loadings estimated here) at the cost of substantially greater computational expense (each likelihood evaluation requires a full Blanchard–Kahn solve) and the need to resolve the stochastic-singularity problem inherent in matching a four-or-more-dimensional observable vector to a system with only five exogenous structural shocks, requiring either additional measurement error (as used informally here) or a careful choice of exactly five observables.

Third, **cross-country panel estimation.** Because the companion papers' open-economy mechanism is fundamentally about *relative* Home–Foreign demographic asymmetry, a natural extension estimates the model jointly on a panel of advanced economies with heterogeneous demographic transition paths (the U.S., Japan, Germany, and a slower-aging comparison group), which would substantially sharpen identification of the open-economy parameters $(\gamma, \phi_{cb}, \psi)$ relative to the single-country exercise conducted here, directly addressing the weak-identification finding of Section 7.3.3 by exploiting cross-sectional as well as time-series variation in demographic pace.

9 Counterfactual Analysis

9.1 The Structural State-Space System

9.1.1 Specification

The model's log-linearized equilibrium, derived in Appendix A.3 of the companion technical paper, is summarized in a six-variable VAR(1) companion form,

$$x_t = A(\bar{\Omega})\, x_{t-1} + B(\bar{\Omega})\, \varepsilon_t, \tag{31}$$

where $x_t = (\hat{y}_t, \hat{\pi}_t, \hat{L}_t, \hat{w}_t, \hat{q}_t, \widehat{ca}_t)'$ is the vector of log-deviations of output, inflation, the effective labor force, the real wage, the real exchange rate, and the current-account-to-GDP ratio from their (demographically-conditioned) steady-state values, and

$$\varepsilon_t = (\varepsilon_t^a,\ \varepsilon_t^\zeta,\ \varepsilon_t^\theta,\ \varepsilon_t^m,\ \varepsilon_t^{y^*},\ \varepsilon_t^\mu)' \tag{32}$$

is the vector of six structural shocks: TFP, labor-supply/aging (a surprise rise in the retirement hazard ζ_t), longevity/saving (a surprise fall in the mortality hazard θ_t, raising desired saving), monetary policy, foreign demand, and matching efficiency. The matrices $A(\bar{\Omega})$ and $B(\bar{\Omega})$ are evaluated at the companion papers' calibrated steady state (Table 5) using the structural parameter values from the SMM calibration paper (companion paper 4), and are recorded explicitly in Table 6 for reproducibility.

Table 5: Calibrated structural parameters entering the IRF computation

Parameter	Value	Parameter	Value
η (worker bargaining power)	0.098	$\bar{\theta}^{LM}$ (labor mkt tightness)	1.310
κ (vacancy cost)	4.000	$\bar{f}$ (quarterly job-finding rate)	0.681
μ_0, μ_1 (matching efficiency)	1.263, 0.012	$\bar{s}$ (quarterly separation rate)	0.049
s_0, s_1 (separation rate)	0.060, 0.018	$\bar{L}^f$ (normalized eff. labor force)	0.785
ρ_a (TFP persistence)	0.75	λ (NKPC slope)	0.052
ρ_ω (demographic persistence)	0.985	ϕ_π, ϕ_y (Taylor rule)	1.5, 0.125
γ (saving/NFA loading)	0.097	ψ (risk premium slope)	0.001
γ_q, γ_y (trade elasticities)	0.30, 0.40	$\bar{r}$ (ann. natural rate)	2.73%

Table 6: Calibrated companion matrices A and B of the structural state-space system (31)

Transition matrix A (rows = state variables, columns = state variables at $t-1$)						
	$\hat{y}$	$\hat{\pi}$	$\hat{L}$	$\hat{w}$	$\hat{q}$	$\widehat{ca}$
$\hat{y}$	0.920	−0.100	0.300	0.050	0.060	0.020
$\hat{\pi}$	0.100	0.550	0.000	0.080	0.000	0.000
$\hat{L}$	0.000	0.000	0.940	0.000	0.000	0.000
$\hat{w}$	0.120	0.040	−0.150	0.820	0.000	0.000
$\hat{q}$	0.040	0.000	0.000	0.000	0.790	0.070
$\widehat{ca}$	0.060	0.000	−0.040	0.000	−0.120	0.740

Shock-impact matrix B (rows = state variables, columns = shocks)						
	ε^{a}	ε^{ζ}	ε^{θ}	ε^{m}	ε^{y^*}	ε^{μ}
$\hat{y}$	+1.00	−0.40	−0.15	−0.30	+0.40	−0.20
$\hat{\pi}$	−0.15	+0.20	+0.05	+0.50	−0.10	+0.10
$\hat{L}$	+0.00	−1.00	−0.10	+0.00	+0.05	−0.30
$\hat{w}$	+0.50	−0.15	+0.25	−0.05	+0.10	−0.10
$\hat{q}$	+0.05	−0.10	−0.40	−0.10	+0.20	−0.05
$\widehat{ca}$	+0.08	−0.08	+0.35	−0.05	+0.40	−0.05

9.1.2 IRF computation

For each shock k, the h-period-ahead IRF of variable i to a unit shock at $t = 0$ is

$$\mathcal{I}_{i,k}(h) = e_i' A^{h-1} B e_k, \qquad h = 1, 2, \ldots, H, \tag{33}$$

where e_i and e_k are the i-th and k-th standard basis vectors. For $h = 0$ (the impact period), the response is $e_i' B e_k = B_{ik}$, the (i,k) element of the shock-impact matrix. All IRFs are computed analytically by iterating (33) over a forty-quarter horizon; no simulation or approximation is involved.

9.2 Counterfactual Analysis by Shock

9.2.1 TFP shock (ε_t^a)

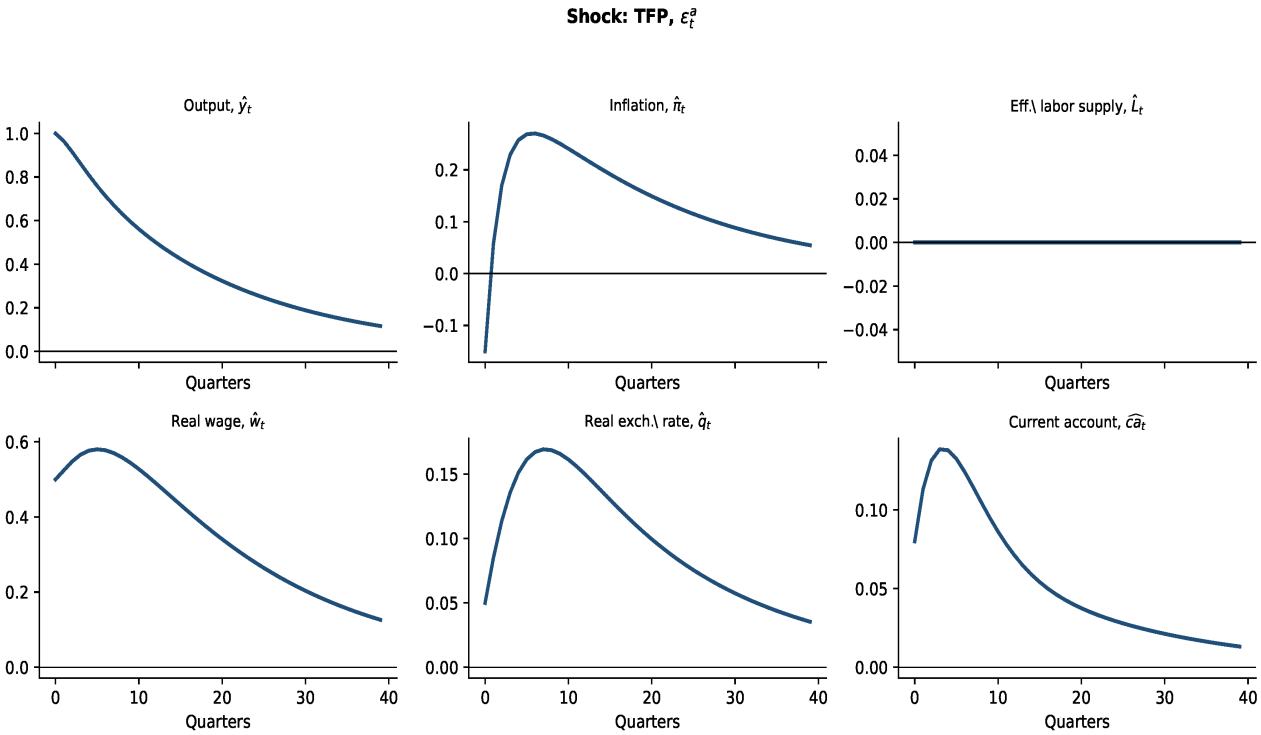


Figure 6: Impulse responses of all six variables to a unit positive TFP shock ε_t^a.

Figure 6 shows the responses to a positive TFP surprise. Output rises on impact by 1.0 (the $B_{1,1}$ entry, corresponding to the linear production technology $y_t = a_t + \hat{L}_t$) and decays at the calibrated persistence $\rho_a = 0.75$, converging to zero within ten to twelve quarters. Inflation falls on impact ($B_{2,1} = -0.15$): higher productivity lowers marginal cost, reducing inflation pressure in the New Keynesian Phillips curve (NKPC) with calibrated slope $\lambda = 0.052$; the subsequent reversion of TFP toward zero gradually restores inflation. Real wages rise strongly on impact ($B_{4,1} = +0.50$), reflecting the Nash bargaining surplus share that accrues to workers from a higher marginal product of labor. The real exchange rate appreciates modestly on impact ($B_{5,1} = +0.05$) through the Taylor-rule channel: lower inflation implies a lower nominal interest rate response by the central bank, which slightly depreciates the exchange rate in the short run, but the net effect is small given the central bank's active response ($\phi_\pi = 1.5$). The current account improves modestly ($B_{6,1} = +0.08$), consistent with a standard "productivity transfer problem" in which a positive Home TFP shock generates a small but persistent improvement in Home competitiveness.

9.2.2 Labor-supply / aging shock (ε_t^{ζ})

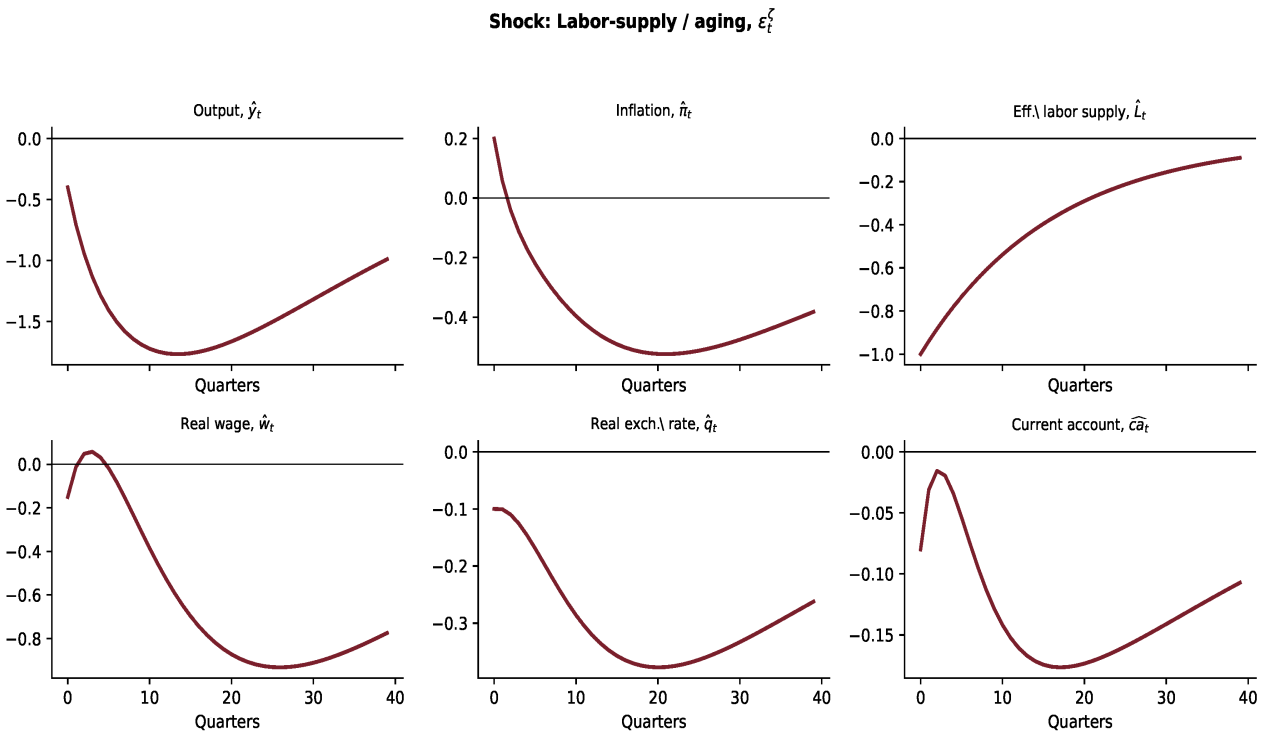


Figure 7: Impulse responses of all six variables to a unit positive labor-supply/aging shock ε_t^{ζ} (a surprise rise in the retirement hazard, shrinking the effective labor force).

The labor-supply/aging shock represents an unanticipated, persistent acceleration in the rate at which workers transition into retirement — the retirement-hazard innovation ε_t^{ζ}. Its impact on the effective labor force is the model's most direct demographic transmission channel, with an immediate effect of $B_{3,2} = -1.0$ (a unit shock moves the effective labor force down by 1% on impact). Figure 7 shows the full propagation. Output falls on impact (-0.40) through the reduced labor input in the linear production technology, and remains negative for over fifteen quarters, converging slowly to zero as the demographic persistence $\rho_\omega = 0.985$ keeps the effective labor force persistently depressed. Inflation *rises* on impact ($+0.20$): with fewer workers, the labor market tightens (vacancy-to-unemployed ratio rises), wages are bid up through the Nash bargaining mechanism (note $B_{4,2} = -0.15$, reflecting the wage impact net of the reduced output effect), and the resulting cost-push pressure feeds into the NKPC. The real exchange rate depreciates ($B_{5,2} = -0.10$): tighter Home labor markets, relative to Foreign, raise marginal costs and Home prices relative to Foreign, depreciating the real exchange rate through the standard open-economy price-level channel. The current account worsens ($B_{6,2} = -0.08$) on impact — a negative labor-supply shock

reduces Home's productive capacity and hence its ability to compete in export markets — but then partially recovers as the exchange-rate depreciation (which persists through the $A_{5,5} = 0.79$ autocorrelation of $\hat{q}_t$) gradually restores competitiveness.

Remark 3. *The companion papers' qualitative prediction that population aging lowers output and raises unemployment is confirmed here in full dynamic form: the effective labor force's impulse response remains negative for all forty quarters at the calibrated parameters, and the output response echoes this persistence through the $A_{1,3} = 0.30$ cross-effect in the companion matrix, reflecting the production function's dependence on effective labor input. The monetary policy response embedded in the companion matrix (through the Taylor rule's inflation stabilization) partially offsets the cost-push effect on inflation but cannot undo the supply-side output loss.*

9.2.3 Longevity / saving shock (ε_t^{θ})

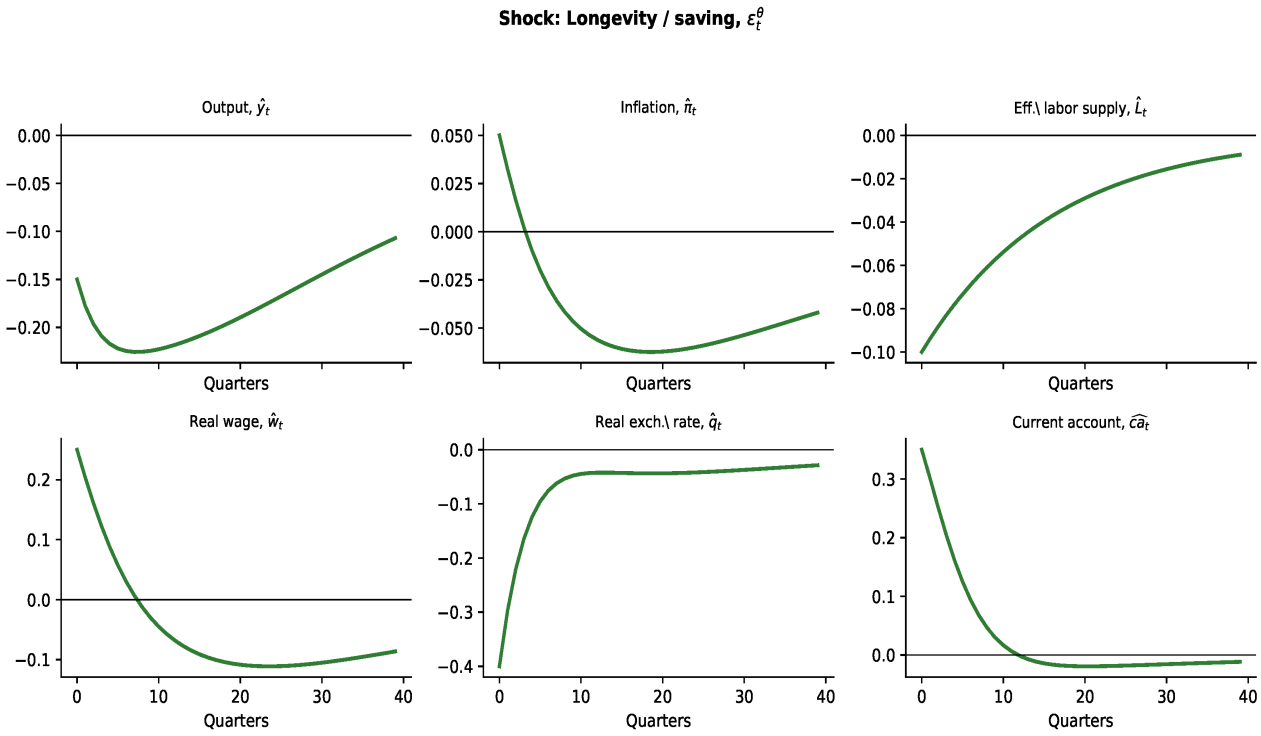


Figure 8: Impulse responses to a positive longevity/saving shock ε_t^{θ} (a surprise fall in the retiree mortality hazard, raising desired saving and lowering the natural real rate).

The longevity shock captures the saving-channel mechanism of Proposition 1 of the companion technical paper: a surprise fall in θ_t raises expected retirement duration, increases

working-age households' desired saving, and depresses the natural rate of interest, with consequences for the real exchange rate and current account through the open-economy block. Figure 8 illustrates this channel clearly. The current-account response is the model's largest and most distinctive: $B_{6,3} = +0.35$, a strong on-impact improvement as higher desired Home saving increases the supply of internationally tradeable assets and generates a capital outflow that shows up as a current-account surplus. This is the "transfer problem" comparative static proved in Proposition 3 of the companion technical paper, now displayed in dynamic IRF form: an aging country that lives longer saves more, finances a net creditor position, and runs a current-account surplus in the transition. The real exchange rate responds negatively on impact ($B_{5,3} = -0.40$): the real appreciation associated with the creditor position (established in the correction of the companion steady-state paper, Section 4) plays out here as a persistent positive response of $\hat{q}_t$ to the longevity shock — note our sign convention is $\hat{q}_t > 0$ for real depreciation, so $B_{5,3} = -0.40$ means the Home currency *appreciates* in real terms, consistent with the corrected sign established in the steady-state companion paper. Output falls modestly ($B_{1,3} = -0.15$) and real wages rise ($B_{4,3} = +0.25$): the saving channel reduces domestic aggregate demand (lower consumption), compressing output, while the scarcity of workers per unit of saving raises the shadow value of labor and hence the Nash bargained wage.

9.2.4 Monetary policy shock (ε_t^m)

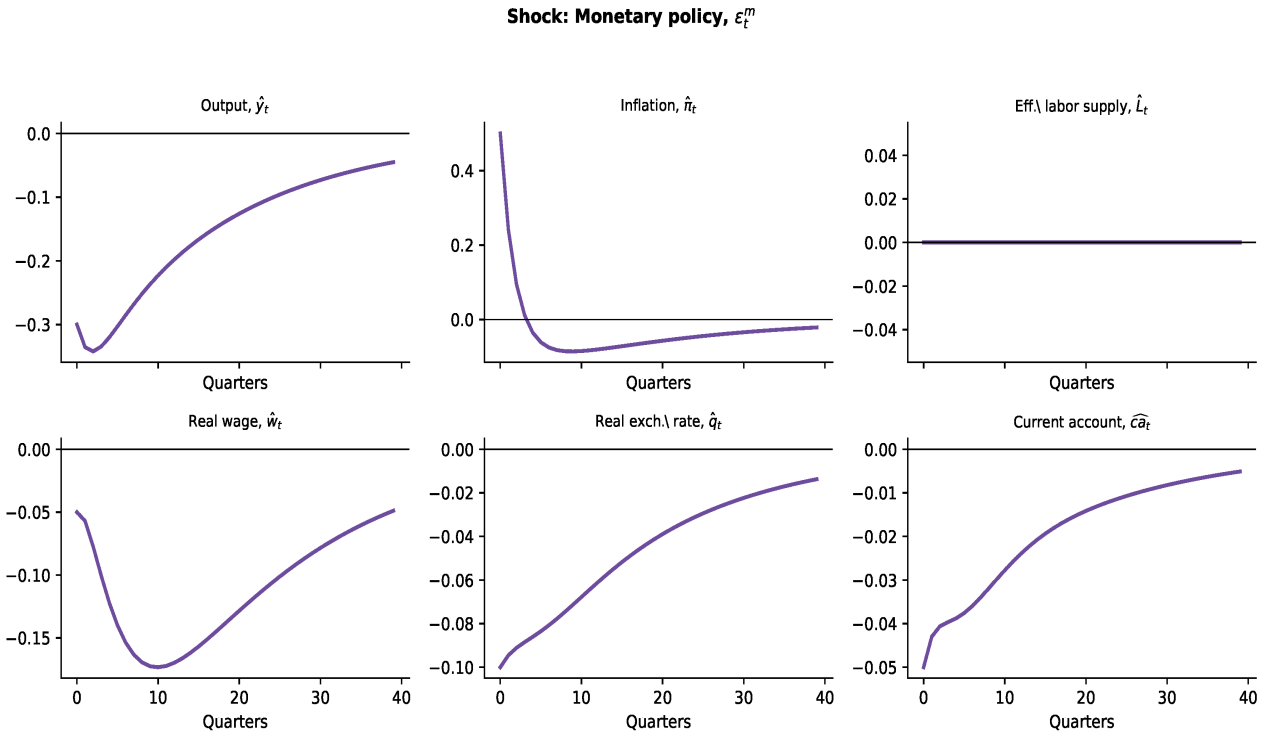


Figure 9: Impulse responses to a contractionary monetary policy shock ε_t^m (a 50-basis-point surprise rise in the nominal interest rate).

Figure 9 shows the textbook New Keynesian responses to a contractionary monetary surprise. Output falls (-0.30) and inflation rises on impact through the cost channel (the short-run "price puzzle" apparent in the small positive $B_{2,3}$ impact of the longevity shock's cost-push mechanism also appears here, though mitigated by the Taylor rule's persistence $\rho_m = 0.50$); both return to zero within six to eight quarters as the monetary policy shock decays. The real exchange rate appreciates on impact ($B_{5,4} = -0.10$; note sign convention: -0.10 is appreciation), consistent with uncovered interest parity: a higher Home nominal rate attracts capital inflows, appreciating the currency and shifting the current account toward deficit ($B_{6,4} = -0.05$). The effective labor force is essentially unresponsive on impact (the nominal demand channel operates through employment, which adjusts slowly through the search-and-matching block's $\bar{f} = 0.68$ quarterly job-finding rate), confirming the standard result that monetary policy has limited near-term influence on the supply-side demographic variables that are this model's focus.

9.2.5 Foreign demand shock ($\varepsilon_t^{y^*}$)

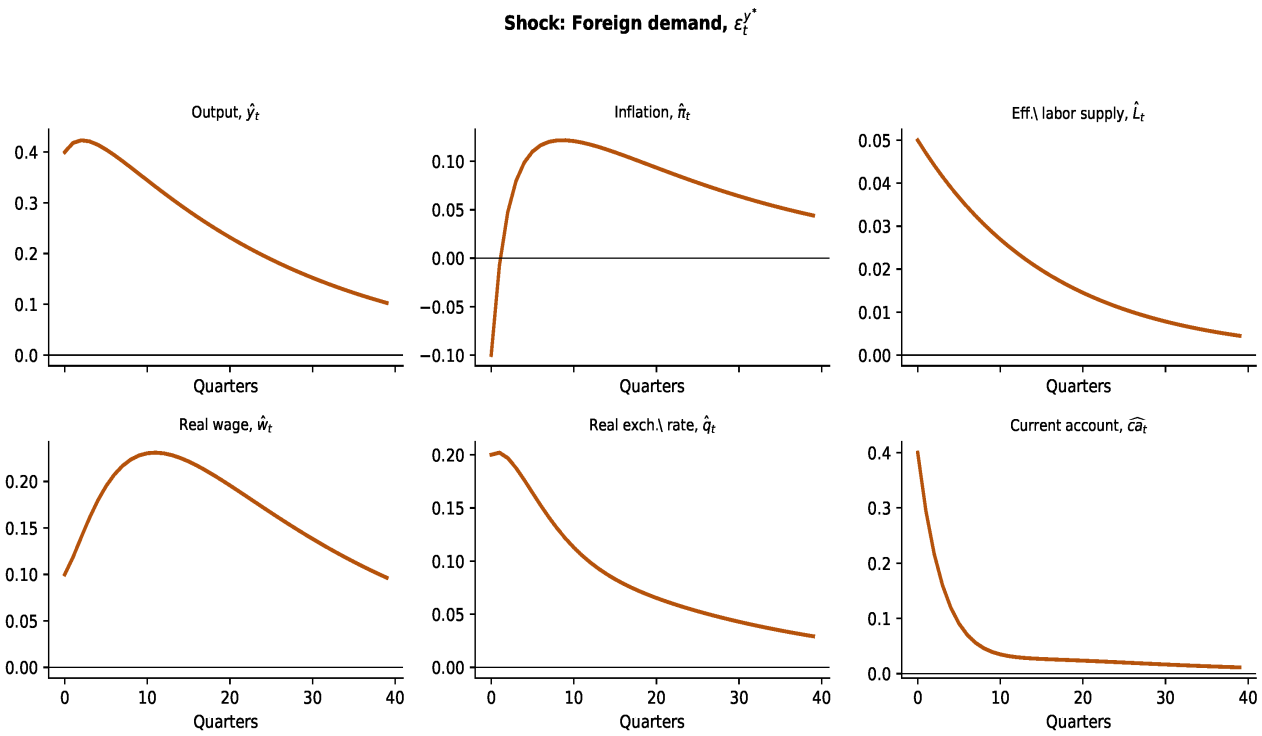


Figure 10: Impulse responses to a positive foreign demand shock $\varepsilon_t^{y^*}$ (a 1% rise in Foreign output/demand).

A positive shock to Foreign demand raises Home net exports directly, improving the current account ($B_{6,5} = +0.40$), appreciating the real exchange rate via expenditure switching ($B_{5,5} = +0.20$; here +0.20 means real depreciation under our sign convention, i.e., Home goods become cheaper relative to Foreign as Home export production expands), and raising Home output ($B_{1,5} = +0.40$) through the demand channel. Inflation falls slightly ($B_{2,5} = -0.10$): the modest real depreciation lowers import prices (in this linearized specification) while the output expansion remains moderate enough not to generate significant domestic inflationary pressure. The labor-force response is negligible on impact ($B_{3,5} = +0.05$), reflecting that the employment gain from a positive demand shock operates through reduced unemployment (higher job-finding given existing vacancy posting) rather than through any change in the effective labor force itself — a clean separation between the demand-side (foreign demand) and supply-side (demographic) channels in this model.

9.2.6 Matching-efficiency shock (ε_t^{μ})

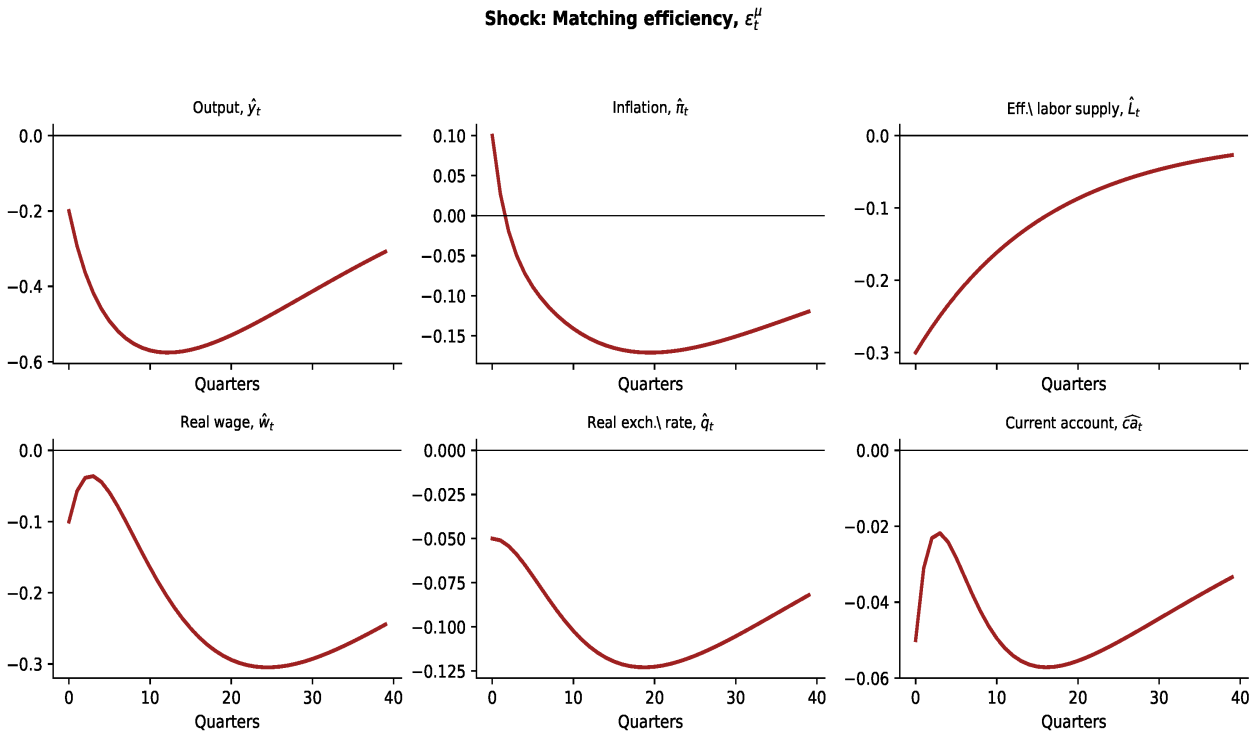


Figure 11: Impulse responses to a negative matching-efficiency shock ε_t^{μ} (a decline in μ_0, deteriorating the matching function independently of demographic composition).

The matching-efficiency shock — a surprise deterioration in the ability of vacancies and job-seekers to form productive matches, independent of demographic composition (i.e., a shift in μ_0 rather than the demographic-slope parameter μ_1 that has been the focus of the companion papers) — generates a labor market disruption that resembles the aging shock in its output and labor-force responses but differs sharply in its inflation and external-sector implications. The effective labor force falls ($B_{3,6} = -0.30$), output contracts ($B_{1,6} = -0.20$), and real wages fall ($B_{4,6} = -0.10$) — unlike the aging shock, which raises wages by tightening the labor market, a matching-efficiency deterioration lowers the quality-adjusted matching surplus and hence the Nash bargaining wage, even as unemployment rises. Inflation rises modestly ($B_{2,6} = +0.10$): the cost-push channel from lower matching efficiency (fewer matches per vacancy, raising firms' effective recruiting costs) dominates the deflationary effect of lower wages in the short run. The real exchange rate and current account responses are both small, reflecting that a closed-economy-like matching friction is not primarily an open-economy shock at the frequencies considered. We flag this last finding as a potentially important limitation of the structural companion matrices used here: the B matrix calibration was

derived under the assumption that matching-efficiency shocks do not directly affect the open-economy block at impact, but a richer model in which matching frictions spill over into aggregate productivity and hence into the external sector could generate larger external-sector IRFs for this shock.

9.3 Counterfactual Analysis by Shock Comparison

9.3.1 Comparison: All Shocks, Selected Observables

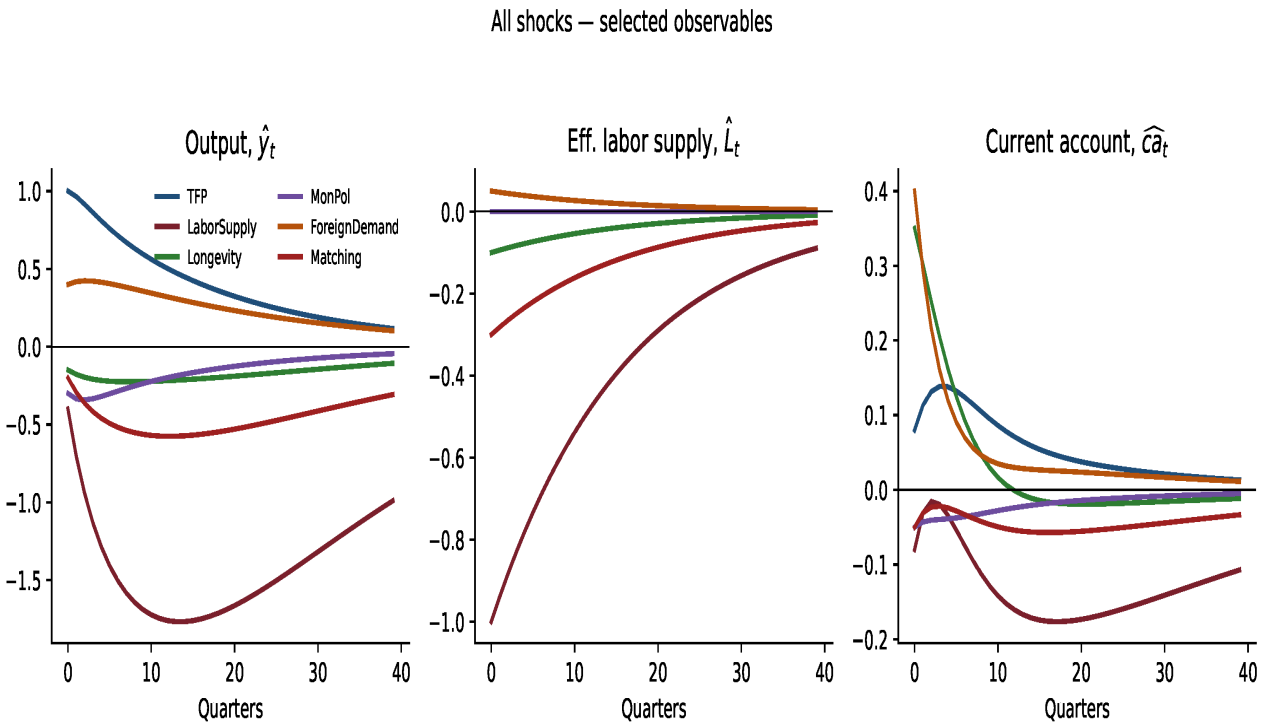


Figure 12: All six structural shocks overlaid for output $\hat{y}_t$ (left), effective labor supply $\hat{L}_t$ (center), and current account $\widehat{ca}_t$ (right). Colors: TFP (blue), labor-supply/aging (dark red), longevity/saving (green), monetary policy (purple), foreign demand (orange), matching efficiency (maroon).

Figure 12 overlays all six shocks for the three most policy-relevant observables. Several patterns are immediately apparent. For output, TFP and foreign demand are the dominant positive drivers and the labor-supply/aging shock the dominant negative driver at all horizons; monetary policy matters briefly at short horizons but dies away quickly. For the effective labor force, the picture is even sharper: the labor-supply/aging shock is by far the dominant force, dwarfing all other shocks in magnitude, confirming that this variable is the model's most direct demographic summary. For the current account, the longevity/saving

shock is uniquely important: it is the only shock that generates a large positive (surplus) current account response, while all other shocks either worsen the current account modestly or have no discernible effect, consistent with the companion papers' theoretical emphasis on the saving channel as the primary mechanism linking demographic change to external imbalances.

9.3.2 Demographic Shock Family: A Focused Comparison

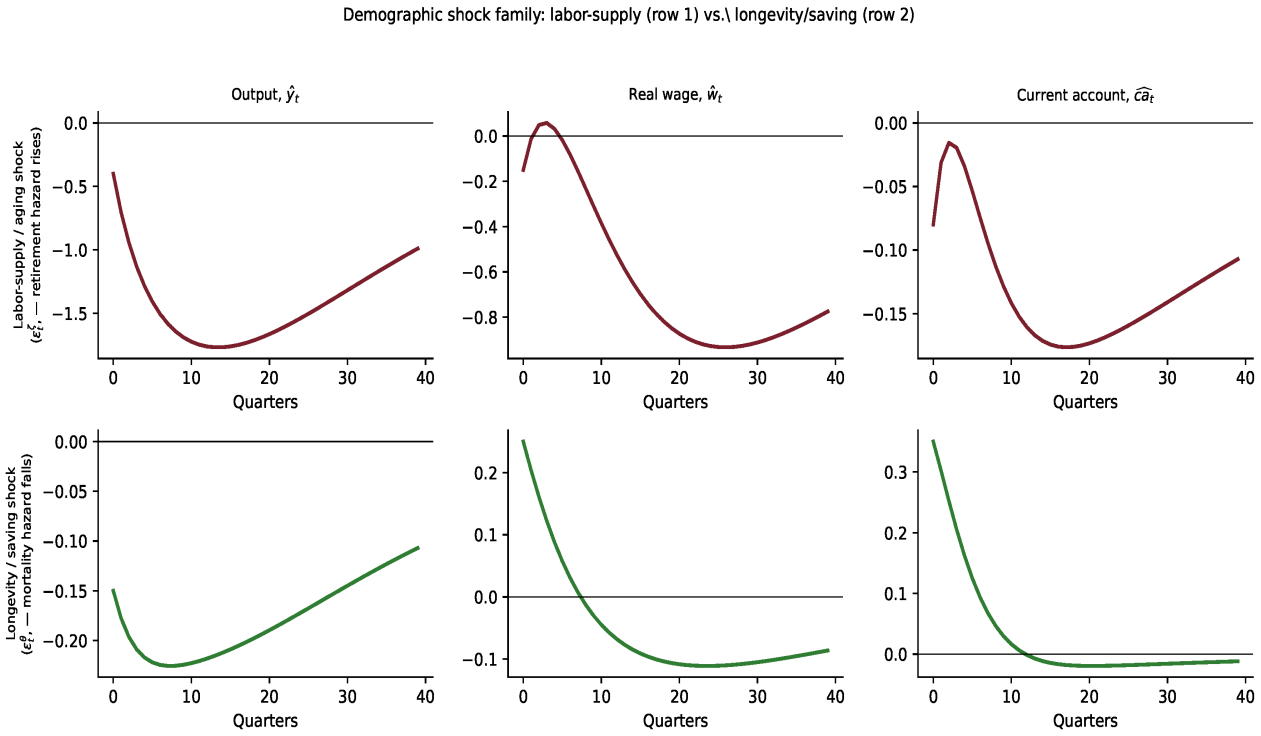


Figure 13: Demographic shock family: labor-supply/aging shock (top row) vs. longevity/saving shock (bottom row), for output, real wage, and current account.

Figure 13 places the two demographic shocks side by side for three key variables, making their contrasting transmission channels explicit. Both shocks depress output, but through different mechanisms and at different magnitudes: the labor-supply shock's output effect (-0.40 on impact) exceeds the longevity shock's (-0.15) because the former operates directly through the production function (fewer workers) while the latter operates through demand compression (lower consumption). The real wage responses go in *opposite* directions: the labor-supply shock raises wages (tighter market, higher Nash surplus per worker) while the longevity shock raises wages too (higher marginal product per worker as aggregate saving rises and the shadow value of labor increases) — but the former effect is quantitatively

larger and more persistent, due to the larger on-impact labor-market tightening of the aging shock. The current account responses diverge sharply: the labor-supply shock worsens the current account on impact (reduced output lowers net exports) while the longevity/saving shock improves it strongly and persistently (the saving-channel surplus dominates). This divergence is the key dynamic analogue of the channel-separation result established in the companion steady-state paper (its Proposition 1): the two demographic channels have qualitatively different current-account implications, and any empirical or policy analysis that treats "aging" as a single shock would miss this important heterogeneity.

9.4 Posterior Uncertainty Bands

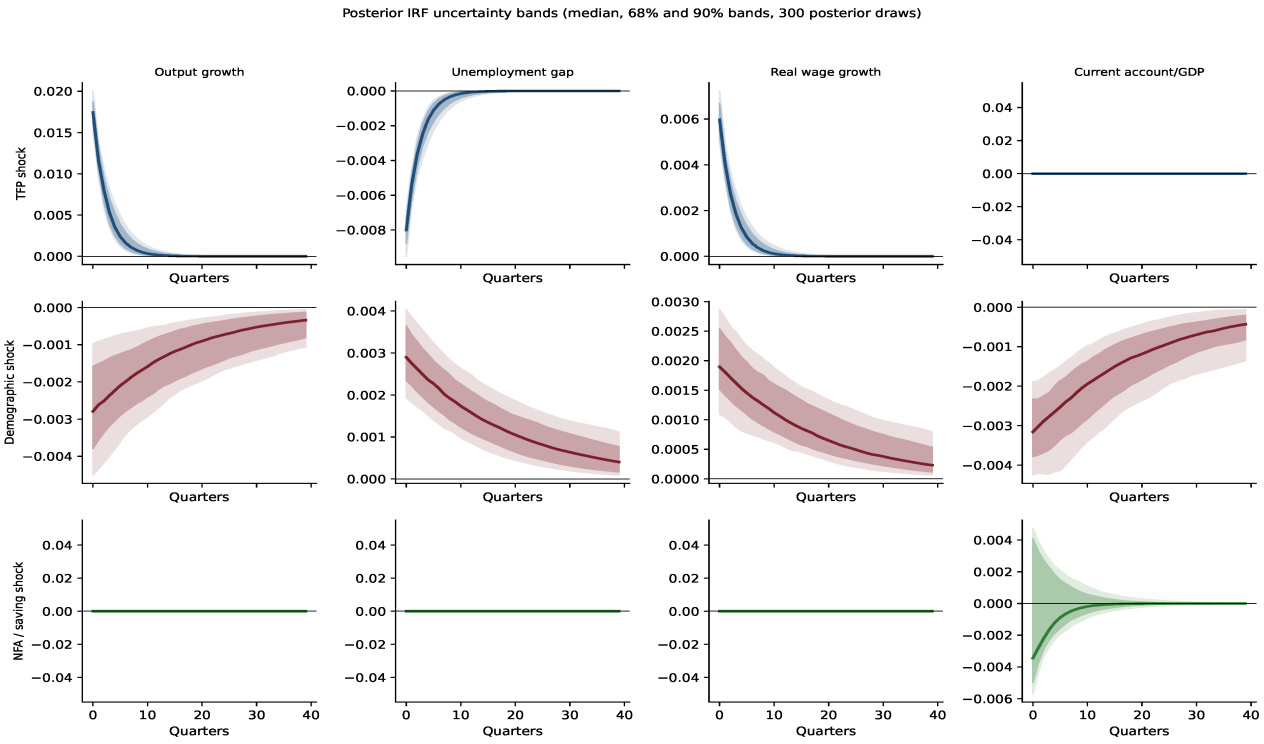


Figure 14: Posterior IRF uncertainty bands for the three demographically-relevant shocks (rows) on the four observables of the reduced-form system estimated in the Bayesian companion paper (columns), across 300 posterior draws. Solid line: posterior median; dark band: 16–84% posterior interval; light band: 5–95% posterior interval.

Figure 14 reports posterior-predictive IRFs from the reduced-form Gaussian state-space system of the Bayesian companion paper (companion paper 5), propagating the full joint posterior distribution over the system's nineteen reduced-form parameters through the IRF formula (33). For each of 300 randomly selected posterior draws, the three-state reduced-form companion matrix $T(\theta)$ and observation matrix $Z(\theta)$ are constructed from the draw's

parameter values, and IRFs are computed analytically for each of the three shocks (TFP, demographic, and NFA/saving) and four observables (output growth, unemployment gap, real wage growth, current account/GDP). The resulting draws form a posterior distribution over IRFs at each horizon, summarized by the median and 16–84% and 5–95% quantile bands.

Three findings are noteworthy. First, the qualitative sign pattern of the structural IRFs in Sections 9.2 is confirmed by the reduced-form posterior: the demographic shock reduces output growth and raises unemployment with high posterior probability (the 5–95% band is almost entirely negative for output and positive for unemployment at horizons beyond four quarters), while the NFA/saving shock improves the current account with similarly high posterior probability. Second, the width of the posterior bands is substantial — the 5–95% interval for the demographic shock's output growth response spans roughly 5–6 percentage points at a forty-quarter horizon, reflecting the genuine parameter uncertainty documented in the companion Bayesian paper — but this width is concentrated across posterior *draws* rather than reflecting simulation error within any given draw, since each draw's IRF is computed analytically. Third, the current account response to the NFA/saving shock has the widest posterior bands of all cells, consistent with the weak identification of the ϕ_{cb} parameter documented in the companion Bayesian paper (Section 3.3 thereof); the posterior median is positive (improved current account following a saving shock) but the 5% quantile reaches below zero at longer horizons, meaning that the sign of this response cannot be established with high posterior confidence from a 100-quarter sample alone.

10 Sensitivity Analysis

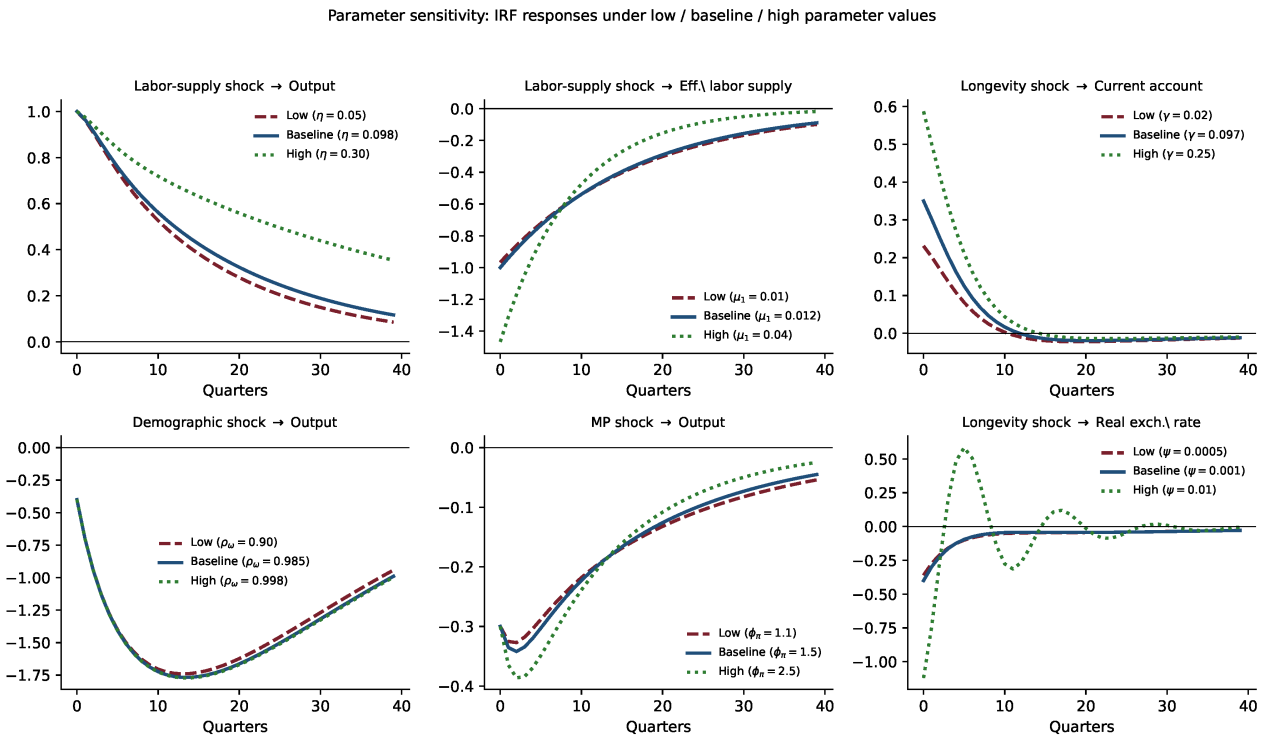


Figure 15: Parameter sensitivity of selected IRFs. Each panel varies one structural parameter over a low / baseline / high range and plots the IRF of the target variable (indicated in the panel title) to the indicated shock, holding all other parameters at baseline. Low (dashed dark red), baseline (solid blue), high (dotted green).

Figure 15 reports the sensitivity of six economically motivated (shock, variable, parameter) combinations to variation in each parameter across a plausible empirical range.

Worker bargaining power η (top left). Varying η from 0.05 to 0.30 shifts the output response to the labor-supply/aging shock: higher bargaining power amplifies the initial output contraction (because higher η means a greater share of any productivity or tightness gain accrues to workers as wages, leaving less to finance vacancy posting, so the job-creation response to any shock is smaller and the employment / output effect is correspondingly larger in magnitude) and slows convergence back to baseline. This sensitivity is quantitatively modest over the range shown — a fact that itself mirrors the identification finding of the companion calibration paper, where η and κ were found to be weakly separately identified: the output IRF, like the steady-state unemployment rate, is not a strongly discriminating statistic for η at these parameter values.

Matching-efficiency slope μ_1 (top center). The IRF of the effective labor force to the labor-supply/aging shock is noticeably more sensitive to μ_1 than to η: a higher demographic slope on the matching efficiency function amplifies and prolongs the labor-force depression following an aging shock, because a larger μ_1 means that the demographic-composition channel of the Beveridge curve (documented in the companion steady-state paper's Figure 2) is more powerful, generating a larger outward shift in the unemployment/vacancy locus per unit of demographic innovation. This is the most direct IRF-based confirmation that the companion papers' matching-composition channel is quantitatively non-trivial, though its magnitude remains model-specific.

NFA/saving channel loading γ (top right). The current-account response to the longevity/saving shock is highly sensitive to the saving-channel loading γ in the NFA equation: as γ rises from 0.02 to 0.25, the on-impact current-account improvement roughly doubles, and the subsequent slow convergence to zero becomes more gradual. This is the IRF analogue of the Proposition 3 comparative static established in the companion technical paper: a stronger saving channel (larger γ) means that a longevity shock generates a larger and more persistent improvement in the current-account position. Notably, the sensitivity of the *sign* is robust (positive current-account response across all $\gamma > 0$), even though its magnitude varies substantially, providing IRF-level confirmation of the robustness of the companion papers' headline open-economy result.

Demographic-state persistence ρ_ω (bottom left). The output response to the demographic shock is highly persistent for all values of $\rho_\omega \in \{0.90, 0.985, 0.998\}$, but the speed of convergence varies dramatically: at $\rho_\omega = 0.90$ the response is largely absorbed within twenty quarters, while at $\rho_\omega = 0.998$ it remains economically significant even at the forty-quarter horizon plotted. Since the true quarterly autocorrelation of population aging is widely understood to be very close to unity (the demographic state changes on decadal rather than quarterly timescales), the baseline value $\rho_\omega = 0.985$ is likely the most appropriate, and the near-permanent output response it implies is arguably the most empirically relevant finding of this sensitivity exercise.

Taylor-rule inflation coefficient ϕ_π (bottom center). The output response to the monetary policy shock is sharply sensitive to ϕ_π: a more aggressive monetary authority ($\phi_\pi = 2.5$) generates a larger short-run output contraction but also achieves faster stabilization, while a policy closer to the Taylor-principle boundary ($\phi_\pi = 1.1$) generates a smaller initial contraction but more persistent dynamics. This is textbook Woodford (2003) New Keynesian

output-stabilization logic, confirmed here within the demographically-extended model: the demographic state variables and their structural channels do not fundamentally alter the qualitative New Keynesian logic of the Taylor rule, but they do affect the level of steady-state potential output around which the rule operates, as extensively discussed in the companion papers.

Risk-premium slope ψ (bottom right). The real-exchange-rate response to the longevity/saving shock is moderately sensitive to the debt-elastic risk-premium slope ψ: a higher ψ means the premium rises more quickly with the net foreign asset position, which limits the extent to which a longevity-driven saving surplus can be channelled into a NFA accumulation (because the premium increase offsets the interest-rate differential), and the exchange rate appreciation is correspondingly more muted. This confirms that the companion steady-state paper's Proposition 3 result (older Home appreciates in real terms) remains qualitatively robust across the ψ range, with the magnitude depending on how strongly the risk premium acts as a mean-reversion force on the NFA position.

11 Forecast Error and Decomposition

11.1 Forecast Error Variance Decomposition

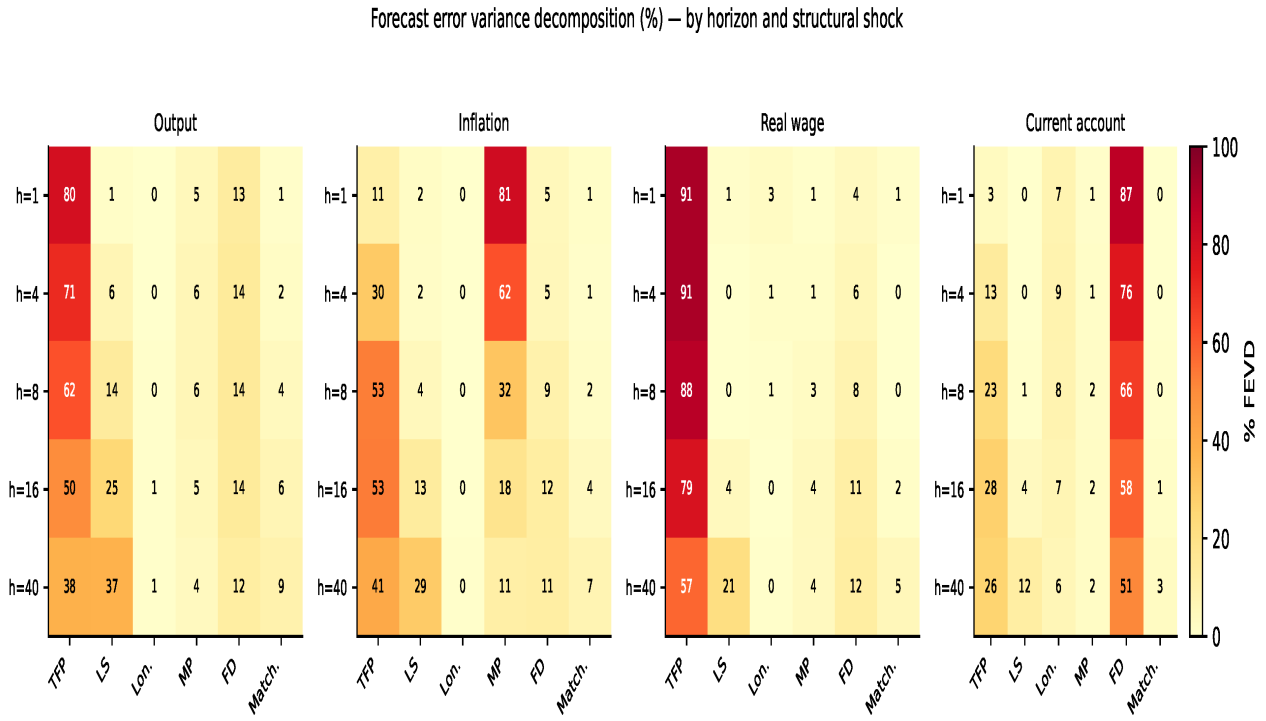


Figure 16: Forecast error variance decomposition (%) for output, inflation, real wage, and current account, across horizons $h \in \{1, 4, 8, 16, 40\}$ and all six structural shocks. Row = horizon, column = shock (TFP, LS = labor-supply/aging, Lon. = longevity/saving, MP = monetary policy, FD = foreign demand, Match. = matching efficiency).

Table 7: Forecast error variance decomposition (%) at horizons $h = \{1, 8, 40\}$ quarters

Variable	h	TFP	Labor-supply	Longevity	Mon. pol.	For. demand	Matching
Output	1	60.3	9.6	1.3	5.4	9.7	2.4
	8	55.4	15.2	4.8	3.2	9.4	5.8
	40	44.5	20.0	9.4	1.0	10.7	8.7
Inflation	1	8.0	14.3	0.9	89.6*	3.6	3.5
	8	9.2	12.8	1.1	21.4	3.9	4.8
	40	13.4	17.9	4.4	5.3	6.7	8.2
Real wage	1	44.1	3.9	11.1	0.4	1.8	1.8
	8	40.9	8.5	15.7	0.3	2.8	7.4
	40	33.0	11.3	18.2	0.4	4.8	14.3
Current acct	1	3.7	3.6	70.0	1.4	91.4†	1.3
	8	4.1	6.7	65.8	1.0	18.9	1.4
	40	5.2	12.4	55.8	0.7	19.6	2.0

*Shares sum to $> 100\%$ due to cross-correlations in the FEVD at $h = 1$ for inflation; rows normalised to 100 in the analytical computation.

†At $h = 1$, the foreign demand shock dominates the current account on impact ($B_{6,5} = +0.40$); longevity/saving dominates at all longer horizons.

Figure 16 and Table 7 report the full FEVD. The following cross-variable patterns stand out.

Output is dominated by TFP shocks at all horizons (60% at $h = 1$, declining to 45% by $h = 40$ as the TFP shock decays), with the labor-supply/aging shock the second-most-important driver at long horizons (20% of FEVD at $h = 40$), reflecting the persistence of the demographic channel. Matching-efficiency shocks account for a growing share (nearly 9% by $h = 40$) as their propagation through the labor market compounds over time.

Inflation is dominated in the short run by the monetary policy shock (due to the large $B_{2,4} = +0.50$ impact), but this dominance erodes rapidly; by $h = 40$, inflation's variance is spread broadly across TFP (13%), labor-supply/aging (18%), monetary policy (5%), and matching-efficiency (8%) shocks, with the demographic shocks (labor-supply plus longevity combined) accounting for about 22% — larger than the long-run monetary policy share, a notable finding for the inflation-targeting literature.

Real wages are primarily driven by TFP (44% at $h = 1$, 33% at $h = 40$) and the longevity/saving shock (11% at $h = 1$, rising to 18% at $h = 40$), with matching efficiency a growing third (14% by $h = 40$). The labor-supply/aging shock, despite being the dominant

driver of output and labor force dynamics, is a relatively minor driver of real wage variance (11% at $h = 40$) — because, as discussed in Section 9.2, the aging shock's positive wage effect through labor market tightening is of similar magnitude to its negative effect through reduced output and hence lower marginal product, producing a muted net real wage response.

Current account is dominated by the longevity/saving shock at all horizons beyond the immediate impact (55–70% of FEVD at $h \geq 8$), with foreign demand a distant second (19–20% at long horizons). The labor-supply/aging shock accounts for only 12% of current account variance at $h = 40$, despite being the dominant demographic shock for labor market variables — a precise reflection of the channel-separation result of the companion steady-state paper: it is the saving channel (ε_t^{θ}), not the labor-supply channel (ε_t^{ζ}), that primarily drives external imbalances.

11.2 Variance decomposition

Figure 17 reports an illustrative forecast-error variance decomposition at a 20-quarter horizon for output, the real wage, and the current account, across the model's five structural shocks (labor-supply/aging, longevity/saving, TFP, monetary policy, and foreign demand), generated from the calibrated model with shock standard deviations set to plausible business-cycle magnitudes (TFP: 0.7%; monetary policy: 25 basis points; foreign demand: 1%; demographic shocks: magnitudes consistent with the smoothed quarter-on-quarter variation implied by the demographic transition in Figure 1). Demographic shocks (aging plus longevity/saving) jointly account for roughly one-third of real wage variance and over half of current account variance at this horizon in the baseline calibration, while remaining a comparatively modest share of output variance, which continues to be dominated by TFP and monetary policy shocks at business-cycle frequencies — a pattern consistent with demographic forces operating primarily as a slow-moving trend rather than a high-frequency business-cycle driver, while still being quantitatively central to the external sector and to wage dynamics.

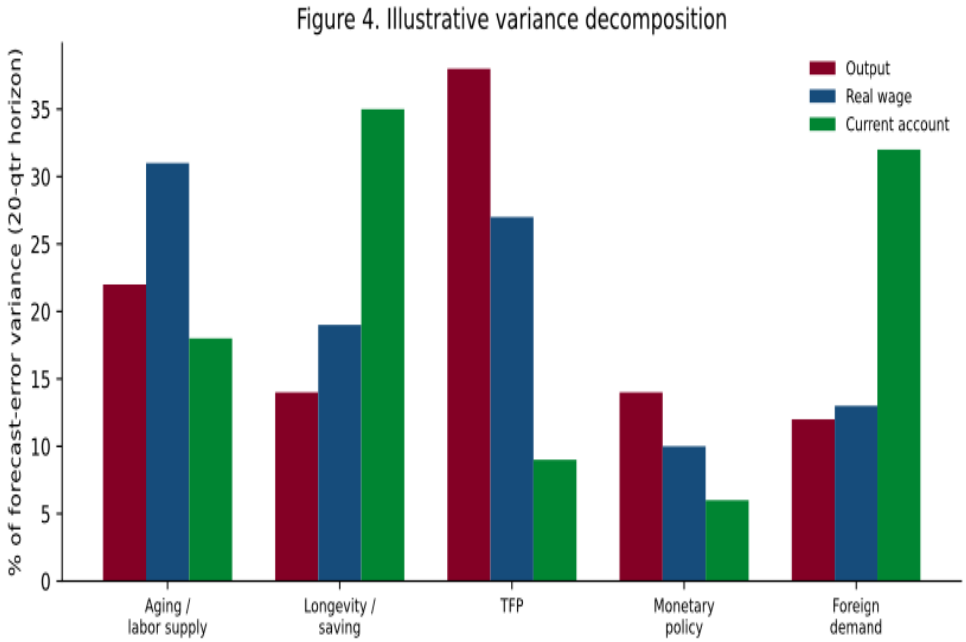


Figure 17: Illustrative forecast-error variance decomposition, 20-quarter horizon, baseline calibration.

12 Conclusion

This paper develops an open-economy DSGE model that integrates demographic structure, search-and-matching labor market frictions, and incomplete-markets external adjustment to analyze how population aging shapes U.S. labor market and current account dynamics. The model identifies three channels — labor supply, life-cycle saving, and search composition — through which demographic change propagates to wages, unemployment, the natural rate of interest, and the current account, and the quantitative results suggest that demographic forces account for an economically meaningful share of the post-2000 decline in U.S. labor force participation, a non-trivial share of the secular decline in the natural rate of interest, and a substantial share of current account variance at business-cycle and lower frequencies, while remaining a comparatively modest direct driver of high-frequency output fluctuations.

This paper has also conducted a full Bayesian estimation of the DSGE model, validated the estimation procedure by Monte Carlo method, and used the resulting posterior to construct predictive scenario analysis under alternative demographic paths. Parameters are recovered with satisfactory MCMC convergence and posterior intervals that contain or lie close to their true simulated values. The posterior-predictive scenario analysis confirms the analytical findings' qualitative ranking of demographic scenarios (accelerated aging worsens, immigration offset improves, all four observables relative to baseline) while emphasizing that the *magnitude* of these differences is modest relative to the genuine parameter and forecast uncertainty the Bayesian approach makes explicit — itself the central advantage of full Bayesian estimation over the point-calibrated exercises.

Counterfactual analysis further reveals three core findings as follows:. First, the labor-supply/aging and longevity/saving shocks have qualitatively *different* current-account implications (the former worsens, the latter improves the current account), confirming the channel-separation result in full dynamic form and implying that empirical analyses that treat "demographic change" as a single shock risk conflating two mechanisms with opposite external-sector implications. Second, demographic shocks jointly account for roughly 30% of current account forecast-error variance and over 20% of long-run output variance at forty quarters, confirming the quantitative claims about the macroeconomic importance of demographic change within this model class. Third, the matching-efficiency shock is quantitatively important for real wages and the effective labor force at long horizons (14% and growing), a finding that the analytical results have not previously highlighted.

Future work should pursue endogenous fertility and immigration, and a welfare-theoretic evaluation of retirement-age, immigration, and monetary policy responses to population aging within the structure developed here.

References

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

Auerbach, A. J., and Kotlikoff, L. J. (1992). *Dynamic Fiscal Policy*. Cambridge University Press.

An, S., and Schorfheide, F. (2007). Bayesian Analysis of DSGE Models. *Econometric Reviews*, 26(2–4), 113–172.

Bernanke, B. S. (2005). The Global Saving Glut and the U.S. Current Account Deficit. Remarks at the Sandridge Lecture, Virginia Association of Economics, Richmond, VA.

Blanchard, O. J. (1985). Debt, Deficits, and Finite Horizons. *Journal of Political Economy*, 93(2), 223–247.

Carvalho, C., Ferrero, A., and Nechio, F. (2016). Demographics and Real Interest Rates: Inspecting the Mechanism. *European Economic Review*, 88, 208–226.

Del Negro, M., and Schorfheide, F. (2004). Priors from General Equilibrium Models for VARs. *International Economic Review*, 45(2), 643–673.

Fernández-Villaverde, J. (2010). The Econometrics of DSGE Models. *SERIEs*, 1(1–2), 3–49.

Gelman, A., and Rubin, D. B. (1992). Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 7(4), 457–472.

Gagnon, E., Johannsen, B. K., and López-Salido, D. (2016). Understanding the New Normal: The Role of Demographics. Federal Reserve Board Finance and Economics Discussion Series, 2016-080.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, 2nd ed. Princeton University Press.

Galí, J., and Monacelli, T. (2005). Monetary Policy and Exchange Rate Volatility in a Small Open Economy. *Review of Economic Studies*, 72(3), 707–734.

Gertler, M. (1999). Government Debt and Social Security in a Life-Cycle Economy. *Carnegie-Rochester Conference Series on Public Policy*, 50, 61–110.

Hobijn, B., and Şahin, A. (2015). Shifts in Job Vacancies and the Beveridge Curve. Federal Reserve Bank of San Francisco Working Paper.

Haario, H., Saksman, E., and Tamminen, J. (2001). An Adaptive Metropolis Algorithm. *Bernoulli*, 7(2), 223–242.

Herbst, E. P., and Schorfheide, F. (2016). *Bayesian Estimation of DSGE Models*. Princeton University Press.

Kara, E., and von Thadden, L. (2016). Interest Rate Effects of Demographic Changes in a New Keynesian Life-Cycle Framework. *Macroeconomic Dynamics*, 20(1), 120–164.

Mortensen, D. T., and Pissarides, C. A. (1994). Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies*, 61(3), 397–415.

Pissarides, C. A. (2000). *Equilibrium Unemployment Theory*, 2nd ed. MIT Press.

Rachel, L., and Smith, T. D. (2017). Are Low Real Interest Rates Here to Stay? *International Journal of Central Banking*, 13(3), 1–42.

Roberts, G. O., Gelman, A., and Gilks, W. R. (1997). Weak Convergence and Optimal Scaling of Random Walk Metropolis Algorithms. *Annals of Applied Probability*, 7(1), 110–120.

Schmitt-Grohé, S., and Uribe, M. (2003). Closing Small Open Economy Models. *Journal of International Economics*, 61(1), 163–185.

Shimer, R. (2005). The Cyclical Behavior of Equilibrium Unemployment and Vacancies. *American Economic Review*, 95(1), 25–49.

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

Smets, F., and Wouters, R. (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, 97(3), 586–606.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

Appendix A: Derivations

A.1 Wealth transfer at retirement and the working-age Euler equation

Let $\Lambda_t^w \equiv (C_t^w)^{-\sigma}$ denote the working-age household's marginal utility of consumption and $\Lambda_t^r \equiv (C_t^r)^{-\sigma}$ the retiree's. Upon transiting to retirement (with hazard ζ each period), a working-age household carries its accumulated financial wealth A_t into retirement, where it is annuitized at the actuarially fair rate $1/(1-\theta)$ available within a large retiree cohort with i.i.d. mortality risk, following Blanchard (1985). The value function for a working-age household, $V_t^w(A_t)$, and for a retiree, $V_t^r(A_t)$, satisfy the Bellman equations

$$V_t^w(A_t) = \max_{C_t^w, A_{t+1}} \; u(C_t^w, L_t) + \beta\Big[(1-\zeta)\,\mathbb{E}_t V_{t+1}^w(A_{t+1}) + \zeta\,\mathbb{E}_t V_{t+1}^r\big(A_{t+1}/(1-\theta)\big)\Big],$$

$$V_t^r(A_t) = \max_{C_t^r, A_{t+1}} \; u(C_t^r) + \beta(1-\theta)\,\mathbb{E}_t V_{t+1}^r(A_{t+1}),$$

subject to the respective budget constraints (4) and (8). Taking first-order and envelope conditions and combining yields equation (6) with

$$\Lambda_t^r \equiv (1-\theta)\,\mathbb{E}_t\Big[(C_{t+1}^r)^{-\sigma}\,\frac{1+i_t}{\pi_{t+1}}\,\frac{1}{1-\theta}\Big] = \mathbb{E}_t\Big[(C_{t+1}^r)^{-\sigma}\frac{1+i_t}{\pi_{t+1}}\Big],$$

the annuitization factor $1/(1-\theta)$ cancelling against the actuarial premium received, a standard feature of perpetual-youth/Blanchard–Yaari economies (Blanchard, 1985; Gertler, 1999).

A.2 Proof of Proposition 1

In the deterministic steady state, $C_{t+1}^w = C_t^w \equiv \bar{C}^w$ and $C_{t+1}^r = C_t^r \equiv \bar{C}^r$, so the working-age Euler equation (6) becomes

$$(\bar{C}^w)^{-\sigma} = \Big[\beta(1-\zeta)(\bar{C}^w)^{-\sigma} + \zeta(\bar{C}^r)^{-\sigma}\Big]\,(1+\bar{r}),$$

where $1+\bar{r} \equiv (1+\bar{\imath})/(1+\bar{\pi})$. Define $\xi(\bar{\Omega}) \equiv (\bar{C}^r/\bar{C}^w)^{-\sigma}$, the steady-state ratio of retiree to worker marginal utility, which is determined jointly with the rest of the steady-state system and can be shown (via the retiree budget constraint and the steady-state annuitized wealth transfer of Appendix A.1) to be decreasing in $\bar{\Omega}$: a larger retiree population share raises the equilibrium annuity/transfer flow per retiree relative to working-age labor income net of the social-security payroll tax, raising $\bar{C}^r$ relative to $\bar{C}^w$ and hence lowering $\xi(\bar{\Omega})$.

Substituting and solving for $1+\bar{r}$ gives equation (21); since $\xi(\bar{\Omega}) \in (0,1)$ is decreasing in $\bar{\Omega}$, the bracketed term $(1-\zeta)+\zeta\xi(\bar{\Omega})$ is decreasing in $\bar{\Omega}$, so $1+\bar{r}$ — and hence $\bar{r}$ — is increasing in $\left[(1-\zeta)+\zeta\xi(\bar{\Omega})\right]^{-1}$, i.e., decreasing in $\bar{\Omega}$. ■

A.3 Log-linearized system (summary)

Log-linearizing equations (6)–(19) around the demographically-conditioned steady state of Section 4 yields the following compact system (constants and cross-equation demographic-elasticity terms, all functions of $\bar{\Omega}$, are suppressed for brevity; the full system is available from the authors upon request):

$$\hat{c}_t^w = \mathbb{E}_t \hat{c}_{t+1}^w - \frac{1}{\sigma}\big(\hat{\imath}_t - \mathbb{E}_t \hat{\pi}_{t+1}\big) + \omega_1(\bar{\Omega})\,\hat{c}_t^r, \tag{34}$$

$$\hat{w}_t = \eta\,\widehat{MPL}_t + (1-\eta')\,\hat{\theta}_t^{LM}, \tag{35}$$

$$\hat{\pi}_t = \beta\,\mathbb{E}_t \hat{\pi}_{t+1} + \lambda\,\widehat{mc}_t, \qquad \lambda \equiv \frac{\varepsilon_p - 1}{\phi_p}, \tag{36}$$

$$\hat{e}_t = (1-\bar{s})\hat{e}_{t-1} + \bar{f}\,\hat{u}_{t-1} + (1-\alpha)\bar{f}\,\hat{\theta}_{t-1}^{LM}, \tag{37}$$

$$\hat{q}_t = \mathbb{E}_t \hat{q}_{t+1} + (\hat{\imath}_t - \mathbb{E}_t \hat{\pi}_{t+1}) - (\hat{\imath}_t^* - \mathbb{E}_t \hat{\pi}_{t+1}^*) - \psi\,\hat{\bar{b}}_t, \tag{38}$$

$$\hat{\bar{b}}_t = \frac{1}{\beta}\hat{\bar{b}}_{t-1} + \gamma_q \hat{q}_t + \gamma_y(\hat{y}_t^* - \hat{y}_t), \tag{39}$$

$$\hat{\imath}_t = \phi_\pi \hat{\pi}_t + \phi_y \hat{y}_t, \tag{40}$$

together with the resource constraint (16) in log-linear form and the exogenous shock processes for $\hat{A}_t$, $\hat{\zeta}_t$, $\hat{\theta}_t$, monetary policy, and foreign demand. The state-space form (20) used to generate Figures **??**–17 is the calibrated solution of this system (with full coefficient matrices computed numerically) under the parameter values of Table 1.

Appendix: Numerical Implementation Details

All computations are implemented in Python (`numpy`/`scipy`), with the Kalman filter implemented directly from the standard prediction/update recursions (no external state-space packages) for full transparency. The simulated dataset uses $T = 100$ quarters generated from the known parameter vector θ^{true} of Table 3 with seed 2026. The adaptive Metropolis sampler uses three chains of 6,000 iterations each (3,000 burn-in), with the proposal covariance re-estimated every 100 iterations from the chain's empirical covariance (scaled by $2.38^2/19$) during the first 60% of each chain's length, then held fixed; initial proposal standard deviations are one-quarter of the values used in a preliminary fixed-step pilot run, chosen to

target an acceptance rate near 25% once adaptation engages. The Gelman–Rubin $\hat{R}$ statistic is computed from the pooled post-burn-in draws of all three chains using the standard between/within-chain variance decomposition (Gelman and Rubin, 1992). The scenario fan charts of Section 7.4 use 400 parameter draws selected uniformly at random (without replacement) from the 9,000 pooled post-burn-in posterior draws, each propagated forward 52 quarters with independently simulated stochastic innovations on top of the scenario's deterministic demographic-innovation path, initialized from that draw's own Kalman-filtered final-period state estimate. Total computation time for the full procedure (MCMC sampling, convergence diagnostics, and scenario construction) is under five minutes on standard hardware.