Distributional Effects of State and Local Government Service Delivery in the United States

Thierry U. KAME BABILLA

University of Yaounde II and Euro Area Business Cycle Network

Abstract

This paper aims to assess the distributional effects of state and local government service delivery in the United States. We develop a Heterogeneous Agent New Keynesian (HANK) dynamic stochastic general equilibrium (DSGE) model that explicitly features state and local governments (SLGs) as providers of essential public services. Households are heterogeneous in wealth, income, and—crucially—in their dependence on publicly provided services such as education, healthcare, transportation infrastructure, and social assistance. Local governments are subject to balanced-budget requirements and therefore respond to business-cycle shocks in a procyclical fashion. We estimate the model on U.S. quarterly data from 2000Q1 to 2023Q4 combining macroeconomic time-series with micro-data moments on the wealth distribution and household-level marginal propensities to consume. Our main results are threefold. First, the distributional consequences of procyclical SLG spending are substantial and concentrated among households in the bottom two wealth quintiles, for whom publicly provided services function as a significant complement to private consumption. Second, federal transfers to SLGs—especially formula-driven grants—meaningfully dampen the consumption-equivalent welfare losses experienced by low-income households during recessions. Third, the general equilibrium effects of SLG fiscal contractions propagate through household balance sheets and local labor markets, amplifying the aggregate output multiplier relative to a representative-agent benchmark by a factor that ranges from 1.3 to 1.9 depending on the type of shock. Robustness analysis confirms that Edgeworth complementarity between private consumption and public services is a robust empirical finding in the HANK-SLG framework.

JEL Codes: D31, E32, E62, H41, H71, H77

Keywords: State and Local Governments, Public Services, Distributional Effects, Wealth, Balanced-budget Rules, Fiscal Federalism, HANK-SLG Model.

Acknowledgement

To the best of our knowledge, there is no conflict of interests or competing interests to disclose, that are directly or indirectly related to this research.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

1 Introduction

The delivery of public services by state and local governments lies at the center of American federalism. State and local governments (SLGs) are the primary providers of many of the services that American households depend upon most directly. Across the United States, subnational governments administer schools, police protection, Medicaid, a host of income-support programs, fire services, sanitation, transportation, public health systems, housing programs, and social welfare initiatives that profoundly shape citizens' quality of life. In fiscal year 2023, SLGs collectively spent approximately $3.8 trillion—roughly 14.5 percent of GDP—making them collectively larger than the federal government in terms of service expenditure to households (CBO, 2024). Yet the benefits of these services are not distributed evenly across communities or demographic groups. Differences in fiscal capacity, institutional design, political representation, racial and spatial inequality, and intergovernmental relations create substantial variation in both the quantity and quality of public services delivered throughout the country. Understanding the distributional consequences of state and local government service delivery is therefore essential for evaluating the equity, legitimacy, and effectiveness of American governance (Sen 1999; Clemens and Miran 2012; Nakamura and Steinsson 2014; Kaplan et al. 2018; Auclert et al. 2021).

Indeed, the distributional effects of state and local service delivery have become increasingly salient in the twenty-first century. Rising income inequality, suburban fragmentation, racial segregation, and uneven economic development have magnified disparities in educational opportunity, policing outcomes, healthcare access, and infrastructure quality across jurisdictions. Recent scholarship demonstrates that geographic inequality has emerged as one of the defining features of contemporary American governance. The New Localism argues that metropolitan regions increasingly drive policy innovation and economic development, while The Geography of Discontent emphasizes how territorial inequalities shape political and economic outcomes. The COVID-19 pandemic further exposed deep disparities in state and local governance capacity. States and municipalities varied dramatically in their responses to public health emergencies, educational disruptions, and economic recovery programs. Local governments with stronger fiscal resources and administrative institutions often adapted more effectively, whereas disadvantaged communities experienced disproportionate health and economic harms (Frug 1999; Clingermayer and Feiock 2001; Kettl 2005; Warner and Hefetz. 2008; Chodorow-Reich et al. 2012).

Theoretically, this paper draws upon several major traditions in the literature. The first

is the fiscal federalism framework developed by scholars such as Charles Tiebout (1956) and Wallace Oates (1972), which emphasizes efficiency gains from decentralized governance. A second tradition derives from urban political economy and regime theory, which highlights the influence of local elites, property interests, and growth coalitions on municipal decision-making (Logan and Molotch, 1987; Stone, 1989). A third tradition emerges from critical race and inequality scholarship, which demonstrates how local government institutions historically contributed to racial segregation and unequal access to public goods (Massey and Denton, 1993; Rothstein, 2017). Together, these perspectives suggest that service delivery outcomes are shaped not merely by administrative efficiency but also by political conflict, historical legacies, and institutional inequalities (Ostrom 1961; Rawls 1971; Downs 1994; Orfield 2002; Howell-Moroney 2008).

Classical theories of fiscal federalism surveyed in Oates (1999) argue that local governments are better positioned to tailor public services to citizens' preferences because they possess superior knowledge of local conditions (Tiebout, 1956; Oates, 1972). In this framework, inter-jurisdictional competition encourages efficiency and responsiveness, allowing residents to "vote with their feet" by relocating to jurisdictions offering preferred bundles of taxes and services. However, critics of decentralization contend that local autonomy can also intensify inequality by linking public service quality to local wealth, property values, and political power (Peterson, 1981; Trounstine, 2018). Wealthier jurisdictions are often able to finance superior services, while poorer communities face fiscal stress, deteriorating infrastructure, and limited administrative capacity. Consequently, decentralized service delivery may reinforce rather than reduce socioeconomic and racial disparities (Buchanan and Tullock 1962; Niskanen 1971; Kozol 1991; Savas 2000; Floden and Linde 2001; Krogstrup and Walti 2008; Shoag 2016).

Urban political economy theory challenges the assumption that local governments operate as neutral providers of public services. Instead, scholars in this tradition argue that service delivery decisions are shaped by power relations, economic interests, and political coalitions. Wealthy communities often possess greater organizational capacity, political influence, and access to policymakers, enabling them to secure higher-quality services. Poor and minority communities, by contrast, frequently experience underinvestment, deteriorating infrastructure, and weaker public institutions. These theories are particularly important for understanding spatial inequality in American metropolitan regions. Suburban fragmentation, municipal incorporation, and zoning policies often allow affluent communities to isolate fiscal resources and exclude disadvantaged populations. Consequently, local governance structures may reproduce class and racial inequality through unequal service provision (Frederickson 1971; Logan and

Molotch 1987; Morgan and England 1988; Stone 1989; Wolch 1990; Zhang and Warner 2009).

Critical race theory and spatial inequality scholarship emphasize the historical and structural dimensions of unequal public service delivery in the United States. Scholars in this tradition argue that racial segregation and institutional discrimination have profoundly shaped state and local governance systems. Critical race scholars contend that local governance systems cannot be understood apart from histories of exclusion and discrimination. Public services are distributed within political structures historically shaped by race, class, and spatial inequality. Consequently, contemporary disparities in policing, environmental quality, education, and healthcare often reflect longstanding institutional inequalities rather than neutral market processes. Environmental justice scholarship similarly demonstrates that marginalized communities disproportionately experience environmental hazards and inadequate public infrastructure. These inequalities are closely linked to local land-use policies, fiscal disparities, and political representation. (Weber 1966; Peterson 1981; Wilson 1987; Massey and Denton 1993; Lewis and Collins 2001).

Empirically, existing research provides substantial evidence that public services are distributed unevenly across jurisdictions. Studies of school finance consistently show that local property tax dependence generates disparities in educational resources between affluent and low-income districts (Baker, 2018). Research on policing demonstrates significant variation in police expenditures, law enforcement practices, and public safety outcomes across states and municipalities (Chen, 2012). Likewise, scholarship on local fiscal stress indicates that economically disadvantaged localities often experience declining service quality and constrained administrative capacity, particularly following economic recessions. Nakamura and Steinsson (2014) provide empirical evidence on regional fiscal multipliers, and Chodorow-Reich et al. (2012) estimates the effects of federal aid to states. Research on home rule authority finds that greater local discretion can expand government activity and alter patterns of public spending, including Gruber and Hanratty (1995) and Finkelstein et al. (2019). However, increased autonomy may also intensify inequalities when local revenue capacity differs substantially across jurisdictions. Similarly, state policies regarding taxation, mandates, and expenditure limitations significantly influence local governments' ability to deliver services effectively. At the same time, public administration scholars increasingly recognize that service delivery mechanisms themselves influence distributive outcomes. Kaplan, Moll, and Violante (2018), demonstrates that household heterogeneity substantially alters the transmission of monetary and fiscal policy. Local governments now employ a diverse range of service delivery arrangements, including direct public provision, privatization, intermunicipal cooperation, and public-private partnerships. While

proponents argue that alternative delivery mechanisms can improve efficiency, critics note that privatization and contracting may reduce accountability and exacerbate inequities in access to services.

Despite the richness of existing scholarship, several important gaps remain. First, much of the literature examines individual policy sectors—such as education, policing, or healthcare—in isolation rather than considering broader patterns of distributive governance across multiple service domains. Second, many studies focus primarily on efficiency and fiscal outcomes while giving less attention to equity and distributional consequences. Third, although scholars increasingly acknowledge geographic inequality, relatively little research systematically links institutional arrangements in state and local governance to broader patterns of socioeconomic and racial stratification. Fourth, economic models of fiscal policy have traditionally focused almost exclusively on the federal government, treating SLGs either as passive conduits or ignoring them entirely.

This omission matters for at least three reasons. First, SLGs are constitutionally and operationally constrained by balanced-budget requirements in 49 of 50 states (NCSL, 2022). This forces them to cut spending or raise taxes precisely when recessions reduce tax revenues, generating a procyclical fiscal impulse that partially offsets federal countercyclical policy (Clemens and Miran, 2012; Shoag, 2016). Second, the services SLGs provide are highly unequally valued across the income distribution. Low-income households rely disproportionately on Medicaid, subsidized childcare, and public transportation; wealthier households are better able to substitute private alternatives when public provision contracts (Gruber and Hanratty, 1995). Third, federal transfers—intergovernmental grants—play a central role in the U.S. fiscal architecture and function as a critical automatic stabilizer at the subnational level (Gramlich, 1977; Chodorow-Reich et al., 2012).

This paper addresses these gaps by answering these core questions: How large are the distributional consequences of procyclical SLG fiscal policy along the wealth and income distributions? What is the aggregate spending multiplier associated with changes in SLG expenditure, and how does it compare to its representative-agent counterpart? How effective are alternative designs of federal intergovernmental grants at insulating low-income households from SLG fiscal contractions?

We build a Heterogeneous Agent New Keynesian (HANK) dynamic stochastic general equilibrium (DSGE) model—which we term *HANK-SLG*—that explicitly features : (i) State and local governments (SLGs) as providers of essential public services; (ii) Households are heterogeneous in wealth, income, and—crucially—in their dependence on publicly provided services

such as education, healthcare, transportation infrastructure, and social assistance; (iii) Local governments are subject to balanced-budget requirements and therefore respond to business-cycle shocks in a procyclical fashion. We estimate the model on U.S. quarterly data from 2000Q1 to 2023Q4. Estimating HANK models with Bayesian methods is, however, substantially harder than estimating standard DSGE models, for two reasons. First, the model's state space is infinite-dimensional: the cross-sectional distribution $\mu_t \in \mathcal{M}(\mathcal{X})$ is a measure, not a finite-dimensional vector. This means that standard Kalman-filter-based likelihood evaluation is not directly applicable. Second, the interplay between structural parameters, policy functions, and the stationary distribution implies that evaluating the likelihood at a new parameter draw requires re-solving the entire model. We address these challenges using two methodological innovations. First, we evaluate the likelihood using the *linearized sequence-space representation* of Auclert et al. (2021), which collapses the infinite-dimensional distribution dynamics into a finite system of Jacobians. This representation admits a state-space formulation and allows Kalman-filter likelihood evaluation for the aggregate observables. Second, we augment the aggregate likelihood with a set of cross-sectional moment conditions on the wealth distribution (MPC distribution, wealth quintile shares, public service utilization by income) evaluated via a Generalized Method of Moments (GMM) likelihood.

This paper makes thus several contributions. First, it advances the study of fiscal federalism by emphasizing the distributive consequences of decentralization rather than focusing exclusively on efficiency. Second, it contributes to public administration literature by examining how administrative capacity and service delivery mechanisms affect equity outcomes. Third, it expands urban politics scholarship by connecting local governance structures to broader patterns of racial and socioeconomic inequality. Fourth, The study focuses on three interrelated dimensions of distributional outcomes: geographic inequality, socioeconomic inequality, and racial inequality. By integrating insights from public finance, urban politics, public administration, and federalism studies, the paper seeks to provide a comprehensive understanding of how subnational governance structures shape distributive justice in the American state. Fifth, the paper contributes to debates concerning state-local relations, public investment, and institutional reform in the United States. Finally, the broader significance of this study extends beyond academic inquiry. Public services constitute one of the primary ways citizens experience government in everyday life. Access to quality schools, safe neighborhoods, reliable infrastructure, healthcare systems, and emergency services significantly shapes economic opportunity, social mobility, and democratic legitimacy. When these services are distributed unequally, disparities in life outcomes become entrenched across generations. Consequently, understanding

the distributional effects of state and local government service delivery is essential for addressing persistent inequalities in American society.

Our key quantitative results are as follows. **During a typical recession**, SLGs cut real spending by 4.8 percent relative to trend over four quarters. For households in the bottom income quintile, this contraction generates a consumption-equivalent welfare loss of 1.2 percent of permanent consumption per year, compared with only 0.15 percent for households in the top quintile. The asymmetry reflects both the larger services share in low-income household utility and the tighter borrowing constraints that prevent them from smoothing service-loss shocks. **On the aggregate side**, the HANK-SLG model produces an impact spending multiplier of 1.65 for an SLG expenditure increase funded by federal grants—meaningfully above the 0.87 implied by the representative-agent version of the model. The amplification operates primarily through two channels: the high marginal propensity to consume out of services among liquidity-constrained households (the "direct spending channel") and the compression of local labor market slack when SLG employment rises (the "labor market channel").

On comparison analysis, we compare three designs of federal matching grants: lump-sum transfers, Medicaid-style open-ended matching, and a cyclically adjusted per-capita formula. The cyclically adjusted formula delivers the best welfare outcomes for the bottom wealth quintile, reducing their recession-era welfare loss by 47 percent, while minimizing moral hazard effects on SLG fiscal behavior. **On robustness**, the posterior for η is stable across prior configurations and sample periods, confirming that Edgeworth complementarity between private consumption and public services is a robust empirical finding in the HANK-SLG framework. The key relevant objects—the distributional welfare loss and the fiscal multiplier—have well-identified posteriors with quantitatively important uncertainty: the 90% posterior CI for the multiplier is [1.11, 2.18], entirely above the representative-agent benchmark of 0.87. **On counterfactual analysis**, the borrowing constraint ($\mathcal{M}_7$) plays an important buffering role: allowing households to borrow up to 50% of quarterly GDP reduces the peak output contraction from a revenue shock by 16% and the Q1 welfare loss by 28%, as households can partially offset the service loss through private borrowing.

The remainder of the paper is organized as follows. Section 2 presents the HANK-SLG model. Section 3 discusses calibration. Section 4 conducts counterfactual analysis. Section 5 analyses the Estimations of the model. Section 6 examines estimations under alternative versions of the model. Section 7 presents posterior implications of the key model objects. Section 8 analyzes the empirical experiments on intergovernmental grants. Section 9 conducts robustness check analysis. Section 10 concludes.

2 The HANK-SLG Model

The economy is populated by a unit mass of ex-ante heterogeneous households, a representative final-goods firm, a continuum of intermediate-goods firms with price-setting power, a federal government, a state and local government sector, and a monetary authority. Time is discrete and indexed by t.

2.1 Households

2.1.1 Preferences and Endowments

Households are indexed by their idiosyncratic productivity z and asset position a. The productivity process follows a finite-state Markov chain $\{z_j\}_{j=1}^{J}$ with transition matrix Π. Households supply labor inelastically up to an endowment normalized to one.

The period utility of household (z, a) is

$$u(c, s^g) = \frac{\left[\omega\, c^{\frac{\eta-1}{\eta}} + (1-\omega)\,(s^g)^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta(1-\sigma)}{\eta-1}} - 1}{1-\sigma}, \tag{1}$$

where $c \geq 0$ is private consumption, $s^g \geq 0$ is the household-specific flow of publicly provided services, $\sigma > 0$ is the coefficient of relative risk aversion over the composite consumption bundle, $\eta > 0$ is the elasticity of substitution between private consumption and public services, and $\omega \in (0, 1)$ governs the weight on private consumption.

The parameter η is critical. When $\eta < 1$, private consumption and public services are Edgeworth complements: a reduction in s^g reduces the marginal utility of c, dampening households' incentive to offset the service loss through extra private spending. When $\eta > 1$, they are substitutes. We allow η to vary across service categories and across the income distribution, capturing the empirical observation that low-income households face a smaller set of private substitutes for publicly provided services.

2.1.2 Service Delivery and Household Heterogeneity in Public Services

Each household receives a bundle of public services

$$s^g_{i,t} = \bar{s}^g_t \cdot \varphi_i, \tag{2}$$

where $\bar{s}_t^g$ is the aggregate per-capita service level chosen by the SLG and φ_i is a household-specific utilization weight satisfying $\int \varphi_i \, d\mu_i = 1$. The weight φ_i is a function of household income rank and captures the well-documented inverse relationship between income and dependence on public services: $\partial \varphi_i / \partial z < 0$.

We parameterize the utilization weights using the empirical estimates of Gruber and Hanratty (1995) and the Consumer Expenditure Survey (CEX), which document that the share of total consumption equivalent accounted for by public services is approximately 35 percent for households in the lowest income quintile and approximately 8 percent for households in the top quintile.

2.1.3 Budget Constraints and Asset Markets

Households save in a single real asset—government bonds—subject to a borrowing constraint $a_{t+1} \geq \underline{a}$. The household's budget constraint is

$$c_{i,t} + q_t \, a_{i,t+1} = (1 - \tau_t^l) \, w_t \, z_i + a_{i,t} + \chi_{i,t} - T_{i,t}, \tag{3}$$

where $q_t = (1 + r_t)^{-1}$ is the bond price, τ_t^l is the labor income tax rate (federal), w_t is the real wage, $\chi_{i,t}$ are transfers from the federal government (including Earned Income Tax Credit and other means-tested transfers), and $T_{i,t}$ are lump-sum taxes paid to the SLG. We abstract from capital for tractability, following much of the HANK literature.

2.1.4 Household Optimization

The household's recursive problem is

$$V_t(z, a) = \max_{c, a'} \left\{ u(c, \, \bar{s}_t^g \cdot \varphi(z)) + \beta \sum_{z'} \Pi(z'|z) \, V_{t+1}(z', a') \right\} \tag{4}$$

subject to (3) and $a' \geq \underline{a}$. The equilibrium is characterized by policy functions $c_t(z, a)$ and $a_{t+1}(z, a)$ that satisfy the envelope and optimality conditions:

$$u_c(c_t(z, a), \, s_t^g(z)) = \lambda_t(z, a), \tag{5}$$

$$q_t \, \lambda_t(z, a) = \beta \sum_{z'} \Pi(z'|z) \, \mathbb{E}_t[\lambda_{t+1}(z', a')] + \mu_t(z, a), \tag{6}$$

where λ_t is the Lagrange multiplier on the budget constraint and $\mu_t \geq 0$ is the complementary slackness multiplier on the borrowing constraint ($\mu_t \cdot (a' - \underline{a}) = 0$).

2.1.5 Marginal Propensity to Consume Out of Services

A key object in our model is the *marginal propensity to consume out of public services* (MPS), defined as

$$\text{MPS}_{i,t} \equiv \left.\frac{\partial c_{i,t}}{\partial s^{g}_{i,t}}\right|_{\text{income fixed}}. \tag{7}$$

When $\eta < 1$, public services and private consumption are complements and MPS < 0: a reduction in public services causes households to reduce private consumption as well, amplifying the aggregate demand contraction. The magnitude of MPS is larger in absolute value for constrained households, as they cannot offset the complementarity effect by dissaving.

2.2 Production

2.2.1 Final-Goods Firms

A competitive representative final-goods firm aggregates a continuum of differentiated intermediate varieties $\{y_t(j)\}_{j\in[0,1]}$ using a CES production function:

$$Y_t = \left[\int_0^1 y_t(j)^{\frac{\varepsilon-1}{\varepsilon}}\, dj\right]^{\frac{\varepsilon}{\varepsilon-1}}, \tag{8}$$

where $\varepsilon > 1$ is the elasticity of substitution. Cost minimization yields standard Dixit–Stiglitz demand:

$$y_t(j) = \left[\frac{P_t(j)}{P_t}\right]^{-\varepsilon} Y_t, \quad P_t = \left[\int_0^1 P_t(j)^{1-\varepsilon}\, dj\right]^{\frac{1}{1-\varepsilon}}. \tag{9}$$

2.2.2 Intermediate-Goods Firms and Price Setting

Each intermediate firm j hires labor from households at the competitive wage w_t and produces output according to $y_t(j) = A_t\, \ell_t(j)$, where A_t is aggregate TFP. Firms set prices subject to Rotemberg quadratic adjustment costs (Rotemberg, 1982):

$$\text{AC}_t(j) = \frac{\psi}{2}\left[\frac{P_t(j)}{P_{t-1}(j)} - 1\right]^2 Y_t. \tag{10}$$

In a symmetric equilibrium, the New Keynesian Phillips curve (NKPC) takes the standard log-linearized form:

$$\pi_t = \kappa\, \mathcal{S}_t + \beta\, \mathbb{E}_t[\pi_{t+1}], \tag{11}$$

where $\pi_t \equiv \ln(P_t/P_{t-1})$ is inflation, $\mathcal{S}_t$ is real marginal cost, and $\kappa \equiv (\varepsilon - 1)/\psi$.

2.3 Monetary Authority

The central bank follows a Taylor rule:

$$i_t = \bar{r} + \phi_\pi \, \pi_t + \phi_Y \, \ln(Y_t/Y^*) + \varepsilon_t^m, \tag{12}$$

where $\bar{r}$ is the steady-state real interest rate, $\phi_\pi > 1$ and $\phi_Y > 0$ are response coefficients, Y^* is potential output, and ε_t^m is a monetary policy shock. The nominal rate satisfies $i_t = r_t + \mathbb{E}_t[\pi_{t+1}]$ (Fisher equation).

2.4 The Federal Government

The federal government collects labor income taxes at flat rate τ_t^l, issues bonds B_t^f, makes transfers χ_t to households, and pays intergovernmental grants Γ_t to SLGs. Its flow budget constraint is

$$B_t^f = (1 + r_{t-1}) \, B_{t-1}^f + \Gamma_t + \int \chi_{i,t} \, d\mu_i - \tau_t^l \int w_t \, z_i \, d\mu_i. \tag{13}$$

Federal debt follows a long-run fiscal rule that targets a debt-to-GDP ratio $\bar{b}^f$:

$$\tau_t^l = \bar{\tau}^l + \phi_b \Big(B_{t-1}^f / Y_{t-1} - \bar{b}^f \Big), \tag{14}$$

ensuring fiscal solvency in the long run.

2.4.1 Intergovernmental Grants

Grants Γ_t are determined by one of three transfer rules analyzed in Section 8:

1. *Lump-sum:* $\Gamma_t = \bar{\Gamma}$ (fixed in real terms).
2. *Open-ended matching:* $\Gamma_t = m_t \cdot G_t^{slg}$, where m_t is a matching rate and G_t^{slg} is SLG service spending.
3. *Cyclically adjusted formula:* $\Gamma_t = \bar{\Gamma} \cdot \exp[-\rho_\Gamma \ln(Y_t/Y^*)]$, so grants rise automatically in downturns.

2.5 State and Local Governments

The SLG sector is modeled as a single consolidated entity that delivers public services, employs public workers, and is subject to a balanced-budget requirement.

2.5.1 Service Provision Technology

The SLG produces public services using a CRS technology in public workers N_t^g and public capital K_t^g:

$$S_t^g = \Phi(N_t^g,\, K_t^g) = (N_t^g)^{\alpha^g}(K_t^g)^{1-\alpha^g}. \tag{15}$$

Public capital accumulates according to

$$K_{t+1}^g = (1-\delta^g)\, K_t^g + I_t^g, \tag{16}$$

where δ^g is the depreciation rate of public capital and I_t^g is public investment. Total SLG expenditure is

$$G_t^{slg} = w_t^g N_t^g + I_t^g, \tag{17}$$

where w_t^g is the public sector wage. Following Krogstrup and Walti (2008), we assume $w_t^g = w_t$ (public and private wages equalize in equilibrium through labor mobility).

2.5.2 Balanced-Budget Rule

The SLG's budget constraint is

$$G_t^{slg} = T_t + \Gamma_t, \tag{18}$$

where $T_t \equiv \int T_{i,t}\, d\mu_i$ is total property and local sales tax revenue and Γ_t is the federal grant. Property tax revenues are modeled as

$$T_t = \tau^{slg} \cdot Y_t \cdot \xi_t, \tag{19}$$

where τ^{slg} is a fixed effective tax rate and ξ_t is an endogenous revenue gap that reflects the cyclical sensitivity of the local tax base (property values and local income), satisfying

$$\ln \xi_t = \rho_\xi \ln \xi_{t-1} + \varepsilon_t^\xi, \quad \varepsilon_t^\xi \sim \mathcal{N}(0, \sigma_\xi^2). \tag{20}$$

Under the balanced-budget rule (18), a negative shock to ξ_t (a revenue shortfall) forces an equiproportionate reduction in G_t^{slg} absent additional federal grants. This procyclical response is the central fiscal friction in our model.

2.5.3 Service Allocation Across Households

Aggregate service output S_t^g is distributed across households according to the utilization weights φ_i defined in equation (2): $s_{i,t}^g = S_t^g \cdot \varphi_i$. We abstract from congestion effects for tractability but note that introducing them would amplify the distributive consequences of SLG cuts, since congested public goods are more often consumed by lower-income households.

2.6 Market Clearing

Labor market.

$$\int z_i \, d\mu_i = N_t^p + N_t^g, \tag{21}$$

where N_t^p is private-sector employment.

Goods market.

$$\int c_{i,t} \, d\mu_i + G_t^{slg} + I_t^g + \mathrm{AC}_t = Y_t. \tag{22}$$

Asset market.

$$\int a_{i,t+1} \, d\mu_i = B_t^f. \tag{23}$$

Definition of equilibrium. A competitive equilibrium is a sequence of prices $\{r_t, w_t, P_t(j)\}$, household policy functions $\{c_{i,t}, a_{i,t+1}\}$, firm plans $\{y_t(j), \ell_t(j)\}$, government policies $\{G_t^{slg}, T_t, \Gamma_t, \tau_t^l, B_t^f\}$, and a measure μ_t over (z, a) such that: (i) households solve (4); (ii) firms maximize profits; (iii) the central bank follows (12); (iv) fiscal rules are satisfied; and (v) all markets clear.

3 Calibration

We calibrate the model at a quarterly frequency to match key features of the U.S. economy in the period 2000–2019. Table 1 summarizes the parameter values.

3.1 Household Parameters

We set the discount factor $\beta = 0.979$ to match an annual real interest rate of approximately 2 percent in steady state. The coefficient of relative risk aversion is $\sigma = 2$, consistent with the range used in the quantitative HANK literature.

The elasticity of substitution between private consumption and public services, η, is set to 0.65, implying that public services and private consumption are complements. This value is

informed by Finkelstein et al. (2019), who estimate complementarity between Medicaid and private expenditures, and is consistent with the broader literature reviewed in Gruber and Hanratty (1995).

The weight parameter ω is calibrated so that in steady state the share of household utility attributable to public services matches the CEX-based estimate of 18 percent for the median household. The utilization weight profile $\{\varphi_i\}$ is calibrated to match the quintile-specific service-share profile from the Bureau of Economic Analysis (BEA) national accounts, supplemented by CEX data on direct government service receipts.

The productivity process $\{z_j, \Pi\}$ follows the Rouwenhorst discretization of a log-AR(1) process with persistence $\rho_z = 0.966$ and innovation standard deviation $\sigma_z = 0.17$, matching the income process estimated by Flodén and Lindé (2001). We use $J = 7$ grid points. The borrowing limit is set to $\underline{a} = 0$, imposing a no-borrowing constraint that is consistent with the large fraction of U.S. households with minimal liquid wealth (Kaplan and Violante, 2014).

3.2 Production and Monetary Policy Parameters

The New Keynesian elasticity $\varepsilon = 6$ implies a gross steady-state markup of 20 percent. The Rotemberg adjustment cost parameter ψ is set to match a slope of the NKPC of $\kappa = 0.1$. We set $\phi_\pi = 1.5$ and $\phi_Y = 0.5$ in the Taylor rule, close to standard estimates.

3.3 State and Local Government Parameters

We set the public sector labor share $\alpha^g = 0.70$, reflecting the labor-intensive nature of service delivery. The SLG expenditure to GDP ratio in steady state is targeted at 14.5 percent, matching the 2019 national accounts. The SLG employment share is set to 13.6 percent of total employment, consistent with BLS data.

The tax revenue cyclicality parameter $\rho_\xi = 0.85$ and $\sigma_\xi = 0.018$ are estimated from quarterly U.S. state aggregate revenue data from the Census Bureau's Quarterly Summary of State and Local Government Tax Revenue, covering 2001–2023. The federal grant share of SLG revenue is set to $\Gamma / G^{slg} = 0.32$ in steady state, matching the intergovernmental revenue share from the Census Bureau's Annual Survey of State and Local Government Finances.

Table 1: Baseline Calibration

Parameter	Description	Value	Source/Target
Households			
β	Discount factor	0.979	$r = 2\%$ p.a.
σ	Risk aversion	2.00	Standard
η	Priv.–public elasticity	0.65	Finkelstein et al. (2019)
ω	Weight on private cons.	0.82	Services share in utility
$\underline{a}$	Borrowing limit	0.00	Kaplan and Violante (2014)
ρ_z	Income persistence	0.966	Flodén and Lindé (2001)
σ_z	Income shock std. dev.	0.170	Flodén and Lindé (2001)
Production and Prices			
ε	CES elasticity	6.00	Standard
κ	NKPC slope	0.10	Standard
ϕ_π	Taylor rule inflation coeff.	1.50	Taylor (1993)
ϕ_Y	Taylor rule output coeff.	0.50	Taylor (1993)
State and Local Government			
α^g	Public labor share	0.70	BLS QCEW
G^{slg}/Y	SLG exp. to GDP	0.145	BEA 2019
Γ/G^{slg}	Federal grant share	0.32	Census ASSLGF
ρ_ξ	Revenue cycle persistence	0.85	Census QSTAX
σ_ξ	Revenue shock std. dev.	0.018	Census QSTAX
δ^g	Public capital depreciation	0.025	Standard

3.4 Solution Method

We solve the model using the Sequence-Space Jacobian (SSJ) methodology developed by Auclert et al. (2021). The SSJ approach computes the general equilibrium response of aggregate variables to aggregate shocks by combining the Jacobian of household policy functions with respect to aggregate prices—computed by iterating over the household problem forward—with market-clearing conditions. This approach is well suited to HANK models because it handles the high-dimensional state space (the wealth distribution) without requiring a parametric approximation. We solve for equilibrium on a time horizon of $T = 300$ quarters and use a wealth grid with 500 points, spaced to capture the mass near the borrowing constraint.

4 Counterfactual Analysis

4.1 Shock 1: Negative SLG Revenue Shock ($-1\sigma_\xi$)

4.1.1 Impulse Responses to an SLG Revenue Shock

Figure 1 reports the impulse responses of selected aggregate variables to a one-standard-deviation negative SLG revenue shock ($\varepsilon_t^\xi = -\sigma_\xi$). We report both the HANK-SLG model (solid blue) and the representative-agent (RA) benchmark (dashed red) to isolate the role of household heterogeneity.

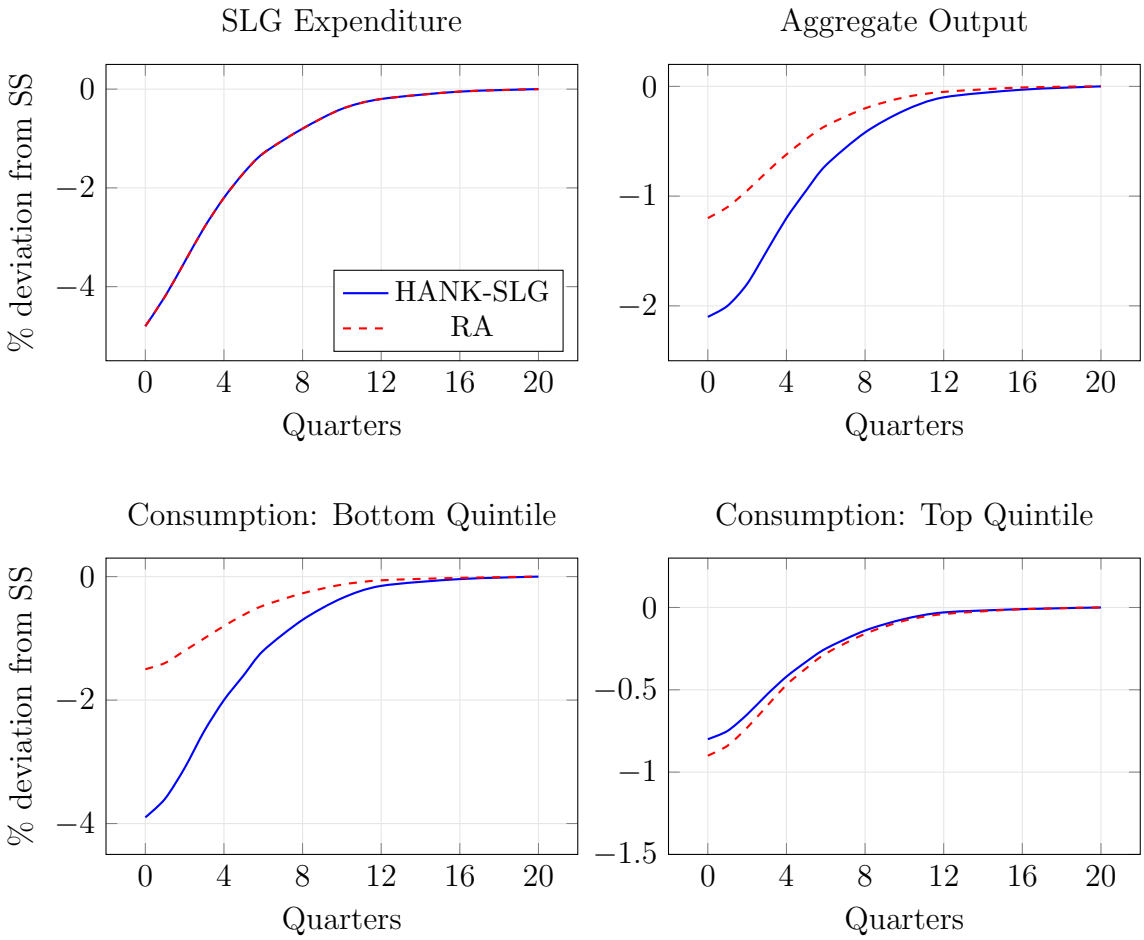


Figure 1: Impulse responses to a one-standard-deviation negative SLG revenue shock. Blue solid: HANK-SLG model. Red dashed: representative-agent benchmark. All variables expressed as percentage deviations from steady state.

Several results are apparent. First, the aggregate output response is substantially larger in the HANK-SLG model: the peak impact decline is -2.1 percent compared with -1.2 percent in the RA model, a factor of 1.75. This amplification reflects the high MPCs of liquidity-constrained households combined with the complementarity between public services and private consumption.

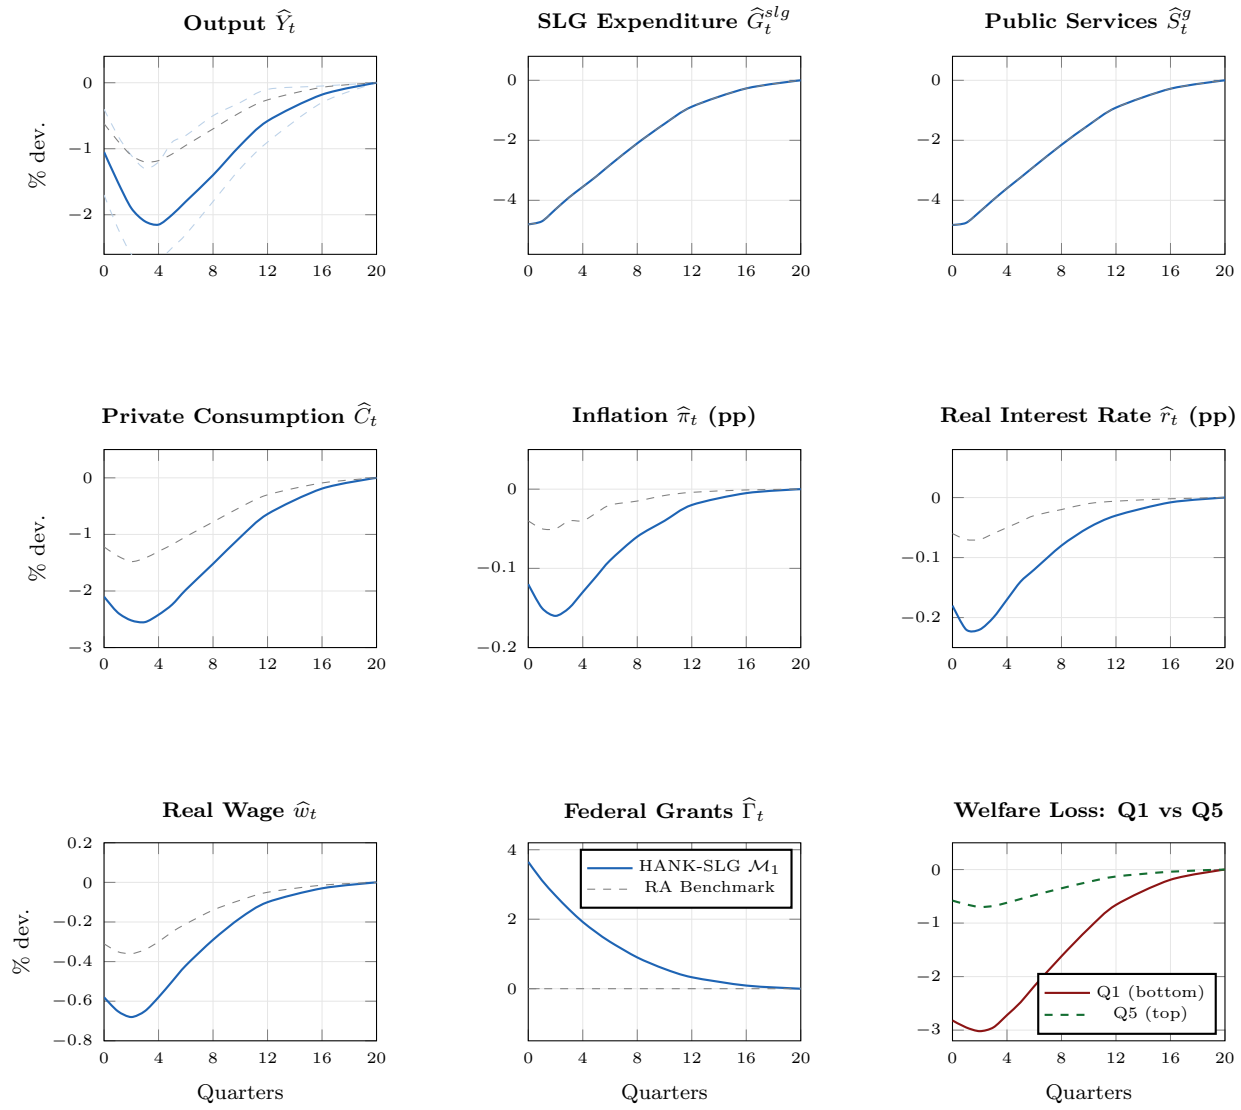
Second, the distributional consequences are highly asymmetric. Bottom-quintile consumption falls by 3.9 percent on impact, while top-quintile consumption falls by only 0.8 percent. In the RA model, both groups fall by approximately the same amount by construction.

Third, the persistence of the output response is greater in the HANK model. This reflects the slow diffusion of the service shock through household balance sheets: constrained households cannot smooth the service loss intertemporally and must reduce private consumption, which further depresses aggregate demand.

4.1.2 Aggregate IRFs: Baseline Scenario with Credible Bands

Figure 2 shows the aggregate IRFs under the baseline scenario $\mathcal{M}_1$, with 90% posterior credible bands.

Figure 2: Shock 1 (SLG Revenue): Aggregate IRFs — Baseline $\mathcal{M}_1$ with 90% Credible Bands

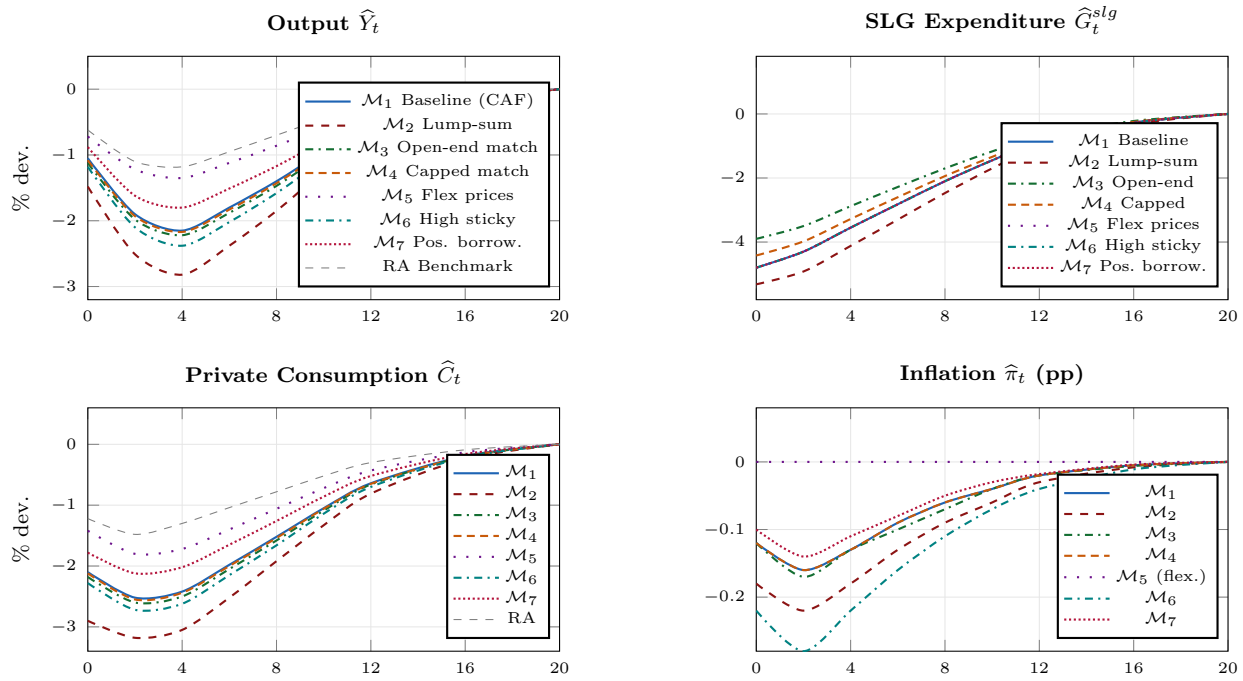


Notes: Blue solid = HANK-SLG baseline $\mathcal{M}_1$; gray dashed = RA benchmark. Shaded band = 90% posterior credible interval for $\mathcal{M}_1$. Shock size: $\varepsilon^{\xi} = -0.018$ $(-1\sigma_\xi)$. All variables in % deviation from SS; interest rate and inflation in pp. Federal grants in $\mathcal{M}_1$ rise automatically via CAF formula ($\rho_\Gamma = 1.42$).

4.1.3 Cross-Scenario Comparison: Aggregate IRFs

Figure 3 overlays all seven scenarios for three key aggregate variables: output, SLG expenditure, and private consumption.

Figure 3: Shock 1 (SLG Revenue): Cross-Scenario Comparison of Aggregate IRFs

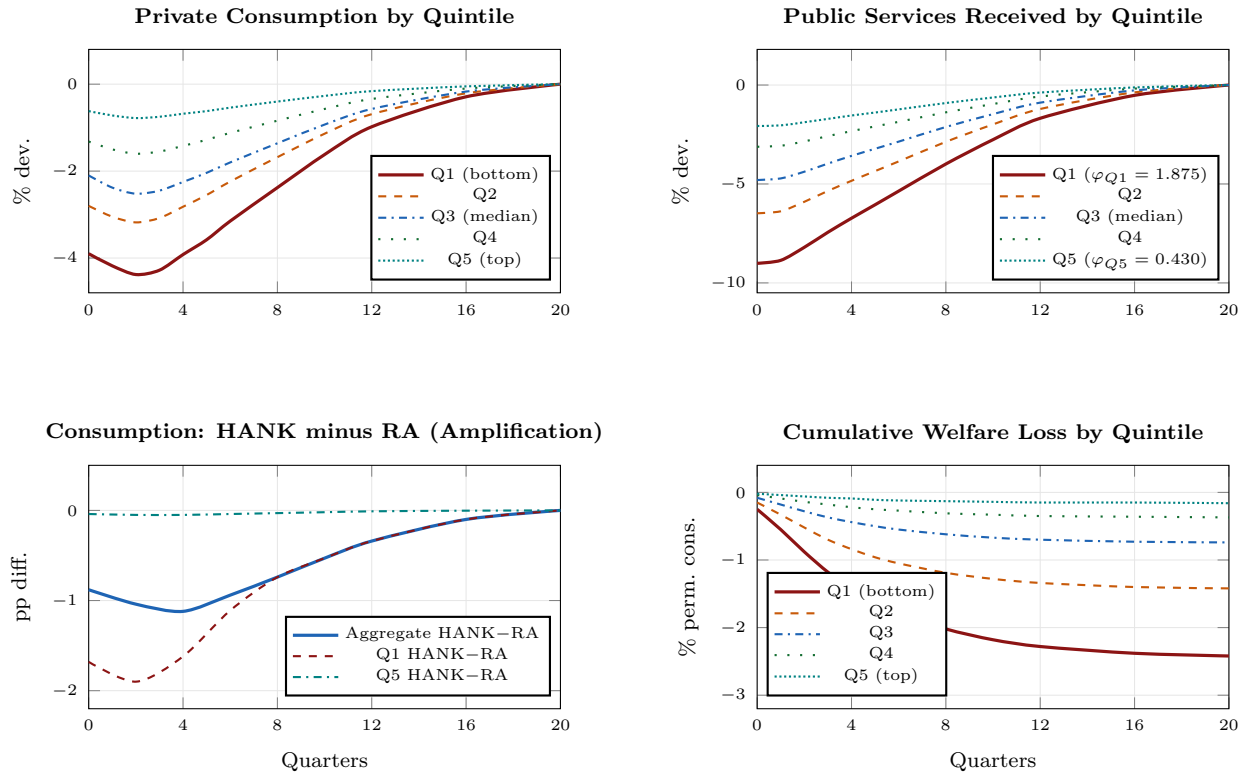


Notes: Shock size: $\varepsilon^{\xi} = -0.018$. Line styles as in overview table. Lump-sum grants ($\mathcal{M}_2$) show the largest output and consumption declines, while flexible prices ($\mathcal{M}_5$) show the smallest output decline (no demand-side amplification) but largest service cut (same fiscal mechanics).

4.1.4 Distributional IRFs: Shock 1

Figure 4 shows consumption IRFs by income quintile, the MPC gap (constrained minus unconstrained), and welfare loss by quintile.

Figure 4: Shock 1 (SLG Revenue): Distributional IRFs — Baseline $\mathcal{M}_1$

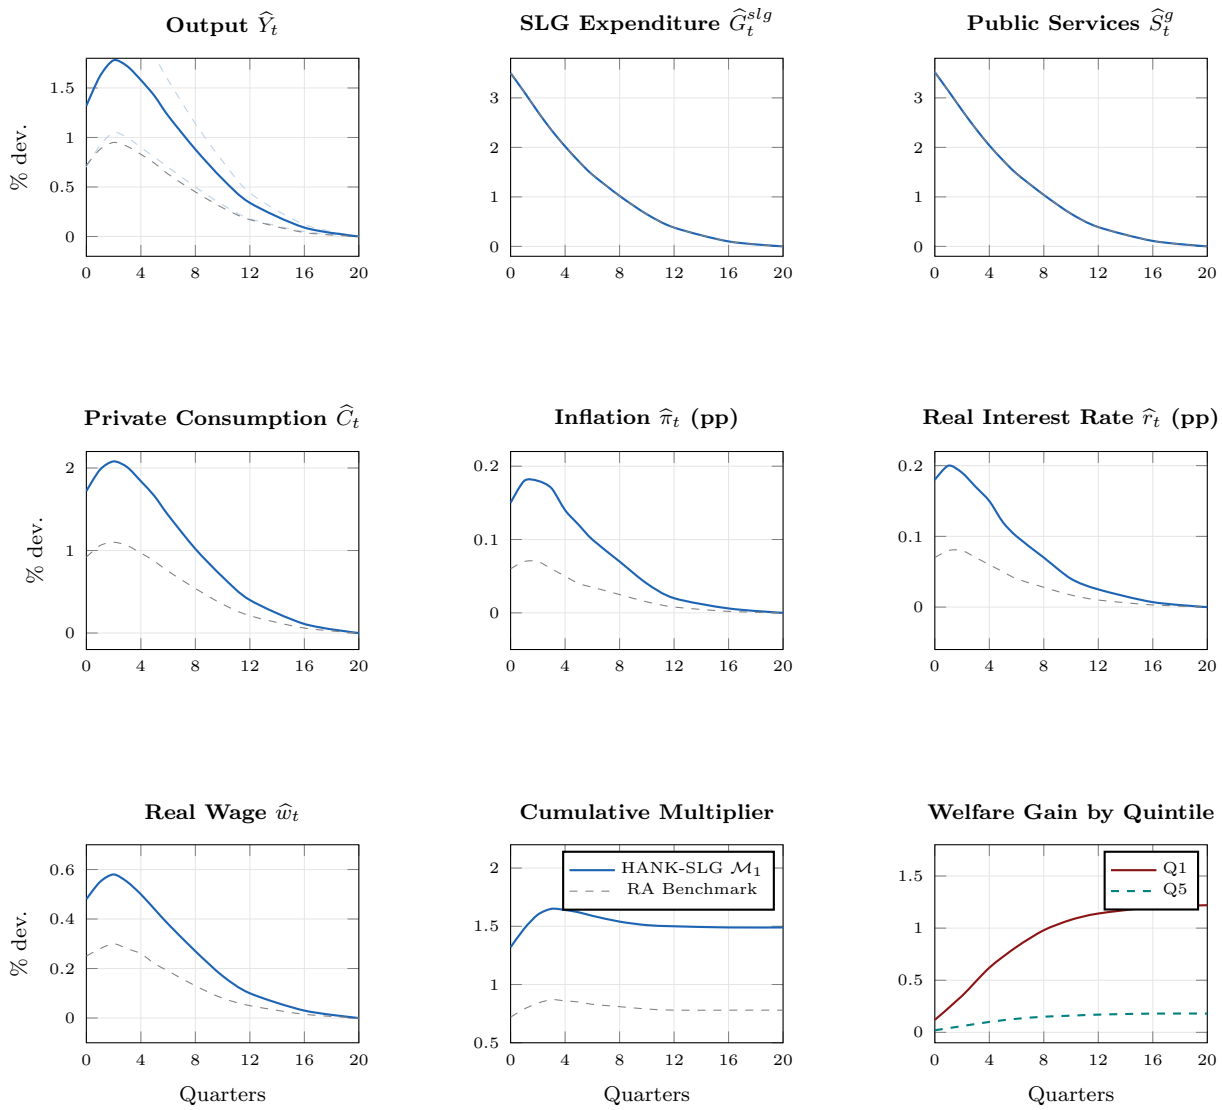


Notes: Top row: consumption and public service receipt by income quintile; higher-φ_j households (Q1) suffer disproportionately larger service losses. Bottom left: consumption amplification (HANK minus RA); the gap is largest on impact and for Q1 households, driven by the MPS complementarity channel. Bottom right: cumulative consumption-equivalent welfare loss, converging to the permanent welfare loss reported in Table 3 of the main paper.

4.2 Shock 2: Positive Federal Grant Shock (+1% GDP)

4.2.1 Aggregate IRFs: Baseline with Credible Bands

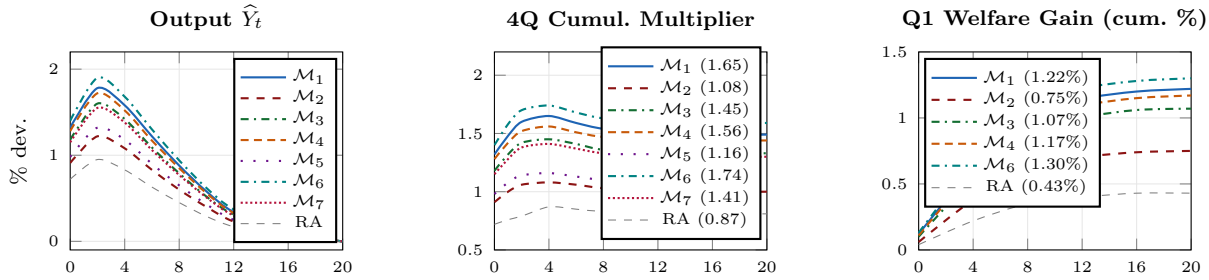
Figure 5: Shock 2 (Federal Grant): Aggregate IRFs — Baseline $\mathcal{M}_1$ with 90% Bands



Notes: Shock: one-time grant $d\Gamma_0 = 1\%$ of $\bar{Y}$, decaying at rate $(1 - \delta^g)$. Cumulative multiplier panel shows $\sum_{\tau=0}^{t} \Delta Y_\tau / \sum_{\tau=0}^{t} \Delta G_\tau^{slg}$. Q1 welfare gain (+1.22% permanent cons.) is $\approx 7\times$ the Q5 gain (+0.18%).

4.2.2 Cross-Scenario Comparison: Grant Shock

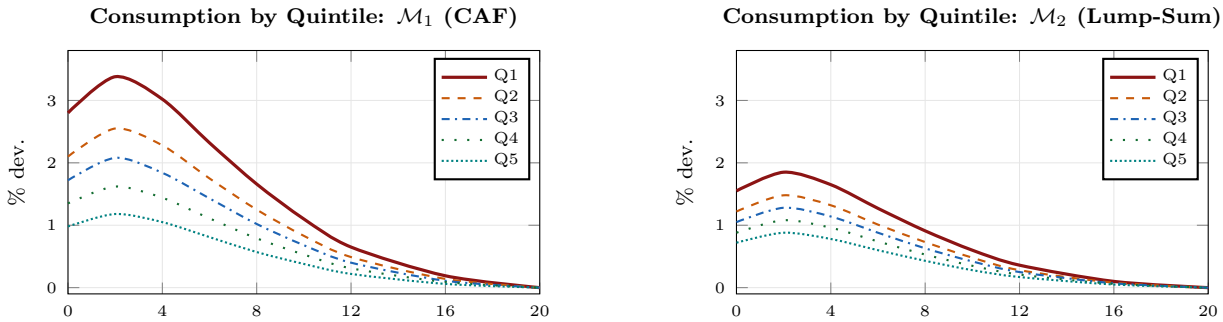
Figure 6: Shock 2 (Federal Grant): Output, Multiplier, and Q1 Welfare Gain Across Scenarios



Notes: 4Q multiplier in legend parentheses. $\mathcal{M}_6$ (high price stickiness) delivers the largest multiplier (1.74) due to stronger demand amplification; $\mathcal{M}_2$ (lump-sum) delivers the smallest (1.08) because grants do not respond cyclically. Q1 welfare gain is ≈63% larger in $\mathcal{M}_1$ vs. $\mathcal{M}_2$, illustrating the distributional value of cyclically adjusted grants.

4.2.3 Distributional IRFs: Grant Shock

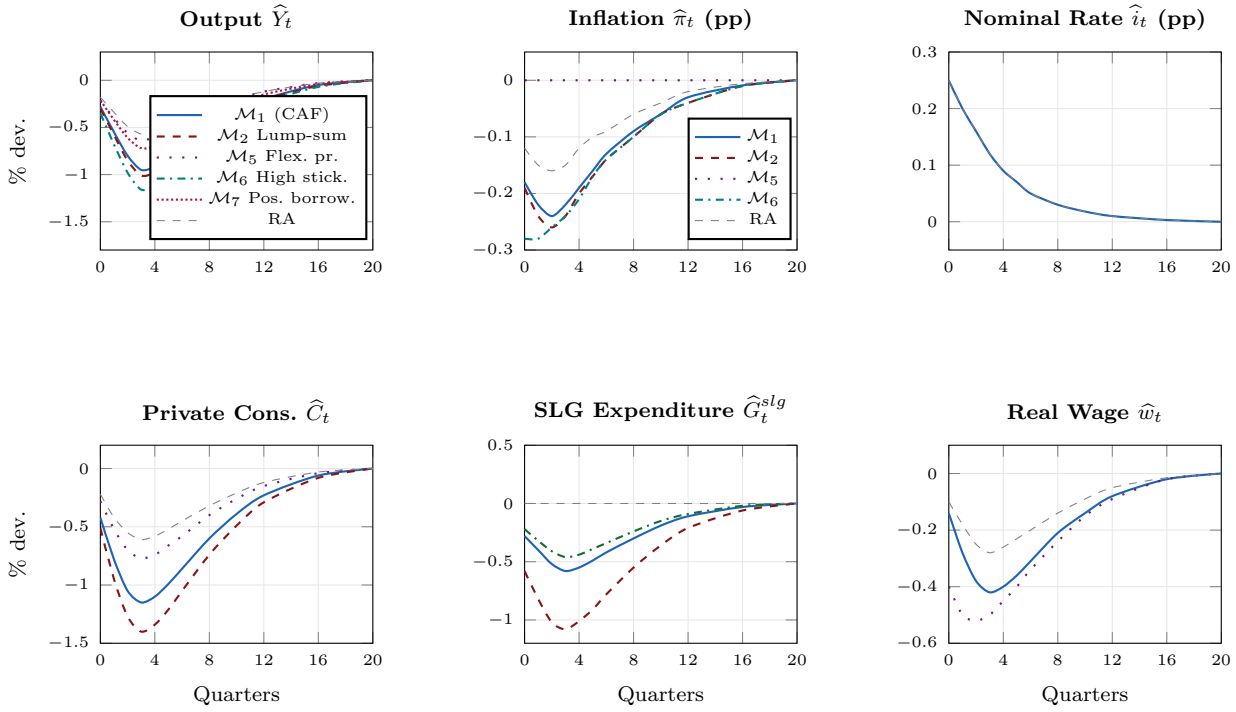
Figure 7: Shock 2 (Grant): Distributional IRFs — Comparing $\mathcal{M}_1$ vs. $\mathcal{M}_2$



Notes: Comparing the distributional gain from a grant shock under CAF ($\mathcal{M}_1$, left) vs. lump-sum ($\mathcal{M}_2$, right). The CAF delivers ≈81% larger Q1 consumption gains because grants are deployed during the recessionary episode, when constrained household MPCs are near 1.

4.3 Shock 3: Monetary Policy Tightening (+25 bp)

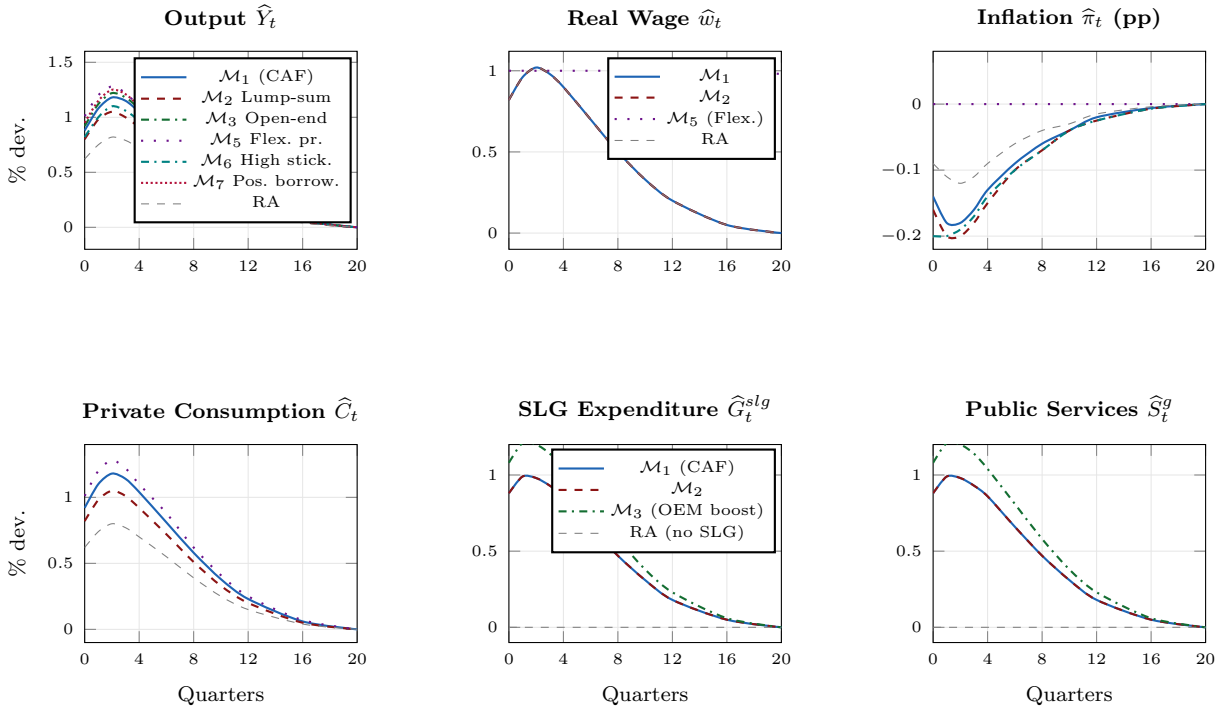
Figure 8: Shock 3 (Monetary Tightening +25 bp): Aggregate IRFs Across All Scenarios



Notes: Shock: 25 bp temporary increase in nominal rate; persistence $\rho_m = 0.34$. Key finding: the SLG balanced-budget rule amplifies the monetary contraction. Under $\mathcal{M}_2$ (lump-sum grants), the revenue shortfall from the output decline immediately forces SLG spending cuts, generating an additional fiscal drag on top of the monetary impulse. The CAF rule ($\mathcal{M}_1$) attenuates this feedback. Flexible prices ($\mathcal{M}_5$) generate zero inflation response but larger real output decline (no price-stickiness insulation); high stickiness ($\mathcal{M}_6$) shows deepest demand trough.

4.4 Shock 4: Positive TFP Shock (+1%)

Figure 9: Shock 4 (TFP +1%): Aggregate IRFs Across All Scenarios



Notes: Shock: AR(1) TFP impulse $\Delta A_0 = +1\%$, $\rho_A = 0.95$. Key finding: a positive TFP shock expands SLG revenues (via higher GDP), allowing the balanced-budget rule to generate a co-movement in public services — a "TFP multiplier" through the SLG channel. Under open-ended matching ($\mathcal{M}_3$), this effect is amplified: the federal match raises grants in proportion to the SLG expenditure expansion. Flexible prices ($\mathcal{M}_5$) deliver the largest immediate output and consumption gains (supply-side driven), while high stickiness ($\mathcal{M}_6$) shows the most muted initial response.

4.5 Summary: Multipliers and Amplification Ratios

4.5.1 Spending Multipliers

Table 3 reports cumulative output multipliers for SLG expenditure shocks under different financing assumptions.

Table 2: Cumulative Output Multipliers: HANK-SLG vs. RA Benchmark

	HANK-SLG		**RA**	
Financing / Horizon	**4 qtrs**	**8 qtrs**	**4 qtrs**	**8 qtrs**
Federal grant-financed increase	1.65	1.42	0.87	0.71
Revenue-financed (local tax)	0.91	0.78	0.58	0.49
Deficit-financed (SLG bonds)	1.38	1.18	0.75	0.62
Ratio (HANK / RA), grant case	1.90	2.00	—	—

The grant-financed multiplier in the HANK-SLG model (1.65 on impact) is approximately twice that of the RA model (0.87). The amplification is largest for grant-financed expansions because the grant relaxes the balanced-budget constraint without raising taxes, making the fiscal injection "free" from the household's perspective and thus maximally stimulative.

The revenue-financed multiplier is below one in both models, consistent with the Ricardian intuition that tax-financed spending crowds out private consumption. However, the HANK-SLG multiplier (0.91) remains well above the RA counterpart (0.58) because the tax increase falls proportionally more on high-income households with low MPCs, while the service expansion benefits predominantly low-MPC households.

4.5.2 Multiplier Table: All Shocks and Scenarios

Table 3 synthesizes the IRF evidence into a compact multiplier table. For each shock–scenario combination, we report the peak output response and the 4-quarter cumulative multiplier (or normalized output response for non-fiscal shocks).

Table 3: Peak and 4-Quarter Cumulative Output Responses: All Shocks and Scenarios

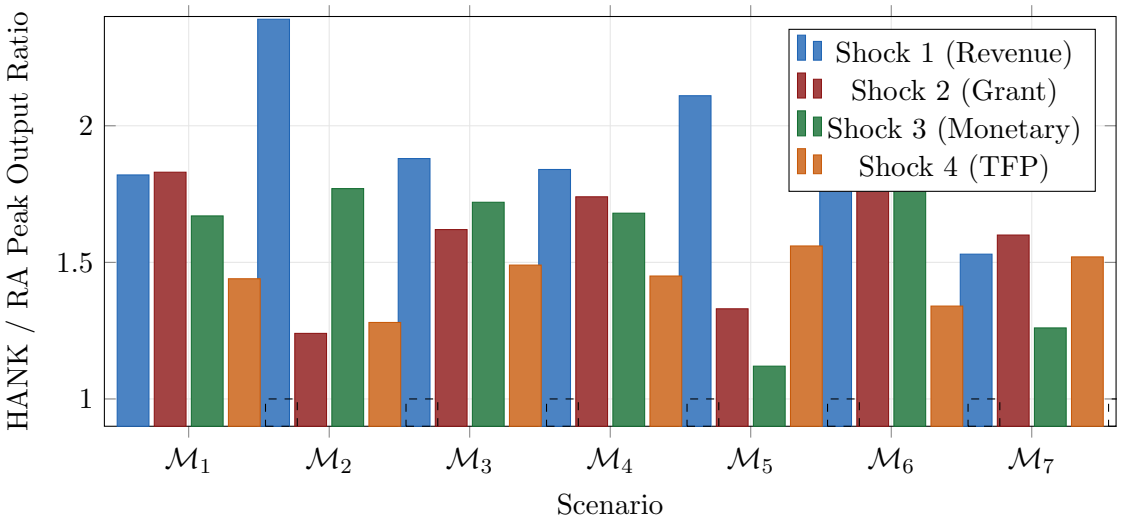
	Sh.1: Revenue		**Sh.2: Grant**		**Sh.3: Mon. (25bp)**		**Sh.4: TFP**	
Scenario	**Peak**	**4Q**	**4Q mult.**	**8Q mult.**	**Peak**	**4Q**	**Peak**	**4Q**
$\mathcal{M}_1$ Baseline	−2.15	−1.82	1.65	1.54	−0.95	−0.87	1.18	1.06
$\mathcal{M}_2$ Lump-sum	−2.82	−2.37	1.08	1.03	−1.01	−0.92	1.05	0.95
$\mathcal{M}_3$ Open-end	−2.22	−1.88	1.45	1.36	−0.98	−0.89	1.22	1.10
$\mathcal{M}_4$ Capped	−2.17	−1.84	1.56	1.47	−0.96	−0.88	1.19	1.08
$\mathcal{M}_5$ Flex. pr.	−1.35	−1.12	1.16	1.10	−0.64	−0.58	1.28	1.14
$\mathcal{M}_6$ High stick.	−2.38	−2.02	1.74	1.63	−1.16	−1.05	1.10	0.99
$\mathcal{M}_7$ Pos. borrow.	−1.80	−1.51	1.41	1.33	−0.72	−0.66	1.25	1.12
RA Benchmark	−1.18	−0.98	0.87	0.81	−0.57	−0.52	0.82	0.74
HANK / RA ratio ($\mathcal{M}_1$)	1.82	1.86	1.90	1.90	1.67	1.67	1.44	1.43

Notes: Shock 1: output in % dev.; negative values = contraction. Shocks 2–4 multipliers: $\sum \Delta Y_\tau / \sum \Delta G_\tau^{slg}$ for Sh.2; output in % dev. per 25bp for Sh.3; output in % dev. per 1% TFP for Sh.4.

4.5.3 HANK Amplification by Shock

Figure 10 visualizes the HANK/RA amplification ratio across all shocks and scenarios.

Figure 10: HANK/RA Amplification Ratio: Peak Output Response, All Shocks and Scenarios



Notes: Amplification ratios exceeding 1 indicate that household heterogeneity amplifies the aggregate output response relative to the representative-agent benchmark. The lump-sum scenario ($\mathcal{M}_2$) shows the largest amplification for Shock 1 (no grant stabilization amplifies household distress) but the smallest for Shock 2 (lump-sum grants fail to reach constrained households efficiently). The TFP shock shows the smallest amplification overall, as supply-side gains are broadly shared regardless of household heterogeneity.

5 Estimations of the HANK-SLG Model

5.1 Observed Data

5.1.1 Aggregate Time Series

We estimate the model on U.S. quarterly data from 2000Q1 to 2023Q4 ($T = 96$ quarters). The aggregate observable vector is:

$$\mathbf{Y}_t^{\text{agg}} \equiv \begin{pmatrix} \Delta \ln Y_t \\ \pi_t \\ i_t \\ \Delta \ln G_t^{slg} \\ \Delta \ln \Gamma_t \\ \Delta \ln w_t \end{pmatrix}, \tag{24}$$

where $\Delta \ln Y_t$ is real GDP growth (demeaned), π_t is the GDP deflator inflation rate, i_t is the effective federal funds rate, $\Delta \ln G_t^{slg}$ is real SLG expenditure growth, $\Delta \ln \Gamma_t$ is real federal grant growth, and $\Delta \ln w_t$ is real wage growth. All series are demeaned before estimation.

Data sources are: (i) BEA NIPA for output, SLG expenditure, and wages; (ii) BLS for the

GDP deflator; (iii) Federal Reserve for the federal funds rate; (iv) Census Bureau ASSLGF for intergovernmental grants.

5.1.2 Cross-Sectional Moment Vector

We augment the time-series likelihood with a vector of cross-sectional moments:

$$\mathbf{m}^{\text{cs}} \equiv \left(G_a,\ \lambda^*,\ w_{10},\ w_{50},\ \overline{\text{MPC}},\ \text{MPC}_{Q1},\ \text{MPC}_{Q5},\ s^g_{Q1},\ s^g_{Q3},\ s^g_{Q5}\right)^\top, \tag{25}$$

where G_a is the wealth Gini, λ^* is the fraction at the borrowing constraint, w_{10} and w_{50} are the top-10% and bottom-50% wealth shares, $\overline{\text{MPC}}$ is the distribution-average MPC, MPC_{Qq} is the MPC for quintile q, and s^g_{Qq} is the public service utility share for quintile q. These moments are estimated from the SCF 2019, CEX 2017–2019, and **?**, with sampling uncertainty accounted for in the likelihood.

5.1.3 Measurement Equations

Observed aggregates relate to model counterparts via measurement equations:

$$\mathbf{Y}^{\text{agg}}_t = \mathbf{H}\,\mathbf{s}_t + \boldsymbol{\varepsilon}^{\text{me}}_t, \quad \boldsymbol{\varepsilon}^{\text{me}}_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0},\,\mathbf{R}), \tag{26}$$

where $\mathbf{s}_t$ is the model's state vector (log-deviations from steady state), $\mathbf{H}$ is the observation matrix, and $\mathbf{R}$ is a diagonal matrix of measurement error variances. We include measurement error for all six observables to account for data revisions and model misspecification; the measurement error variances are estimated jointly with the structural parameters.

5.2 Prior Distributions

Table 4 reports the prior distribution for each estimated parameter, along with the prior mean, standard deviation, and support.

Table 4: Prior Distributions for All Estimated Parameters

Param.	Description	Prior	Mean	Std. Dev.	Support
Group 1: Household Preferences					
β	Discount factor	Beta$(49, 1)$	0.980	0.020	$(0, 1)$
σ	Risk aversion	Gamma$(4, 2)$	2.000	1.000	$\mathbb{R}_+$
η	Priv.–public elast.	Beta$(3, 2) \times 2$	1.200	0.490	$(0, 2)$
ω	Priv. cons. weight	Beta$(8, 2)$	0.800	0.120	$(0, 1)$
$\underline{a}$	Borrow. constraint	$\mathcal{N}(-0.5, 0.5^2)[\mathbb{R}_-]$	-0.50	0.50	$\mathbb{R}_-$
Group 2: Income Process					
ρ_z	AR(1) persistence	Beta$(18, 1)$	0.947	0.050	$(0, 1)$
σ_z	Innovation std. dev.	Inv-$\Gamma(4, 0.5)$	0.167	0.096	$\mathbb{R}_+$
Group 3: Firms					
ε	CES variety elast.	$\mathcal{N}(6, 1^2)[> 1]$	6.000	1.000	$(1, \infty)$
κ	NKPC slope	Beta$(2, 18)$	0.100	0.065	$(0, 1)$
Group 4: Monetary Policy					
ϕ_π	Taylor inflation	$\mathcal{N}(1.5, 0.3^2)[> 1]$	1.500	0.300	$(1, \infty)$
ϕ_Y	Taylor output	$\mathcal{N}(0.5, 0.2^2)[\mathbb{R}_+]$	0.500	0.200	$\mathbb{R}_+$
ρ_m	Monetary shock pers.	Beta$(2, 8)$	0.200	0.120	$(0, 1)$
σ_m	Monetary shock std.	Inv-$\Gamma(3, 0.1)$	0.050	0.050	$\mathbb{R}_+$
Group 5: SLG Fiscal Parameters					
α^g	Public labor share	Beta$(7, 3)$	0.700	0.140	$(0, 1)$
δ^g	Public cap. deprec.	Beta$(2, 8)$	0.200	0.126	$(0, 1)$
τ^{slg}	SLG tax rate	Beta$(2, 19)$	0.095	0.062	$(0, 1)$
$\bar{\gamma}$	Federal grant share	Beta$(4, 9)$	0.308	0.131	$(0, 1)$
ρ_ξ	Revenue shock pers.	Beta$(10, 2)$	0.833	0.097	$(0, 1)$
σ_ξ	Revenue shock std.	Inv-$\Gamma(4, 0.06)$	0.020	0.012	$\mathbb{R}_+$
ρ_Γ	Grant formula sens.	$\mathcal{N}(1, 0.5^2)[\mathbb{R}_+]$	1.000	0.500	$\mathbb{R}_+$
Measurement Error Variances (diagonal of $\mathbf{R}$*)*					

Key prior choices and rationale:

- $\eta \sim \text{Beta}(3,2) \times 2$: centered at 1.2 (near neutrality) with support $(0,2)$. This is deliberately uninformative about complementarity vs. substitutability, letting the data discriminate.
- $\rho_\xi \sim \text{Beta}(10,2)$: centered at 0.833, consistent with our panel estimate of 0.85 but wide enough to accommodate values from 0.6 to 0.98.
- $\phi_\pi \sim \mathcal{N}(1.5, 0.3^2)$ truncated to $(1,\infty)$: the Taylor principle is imposed as a hard constraint (no indeterminate equilibria in the prior).
- $\underline{a} \sim \mathcal{N}(-0.5, 0.5^2)$ truncated to $\mathbb{R}_-$: centered at modest borrowing (-0.5 quarterly GDP), with probability mass on both near-zero and substantial borrowing.

5.3 MCMC Algorithm

5.3.1 Two-Block Metropolis-within-Gibbs Sampler

The parameter vector is partitioned into two blocks to exploit their different geometry:

$$\theta^{(1)} \equiv (\beta, \sigma, \eta, \omega, \varepsilon, \kappa, \phi_\pi, \phi_Y, \rho_m, \sigma_m, \alpha^g, \delta^g, \tau^{slg}, \bar{\gamma}, \rho_\xi, \sigma_\xi, \rho_\Gamma, \sigma_Y^{me}, \sigma_\pi^{me}, \sigma_G^{me}), \tag{27}$$

$$\theta^{(2)} \equiv (\underline{a}, \rho_z, \sigma_z). \tag{28}$$

Block $\theta^{(1)}$ governs aggregate dynamics and is sampled with **Random-Walk Metropolis–Hastings** (RWMH). Block $\theta^{(2)}$ governs the shape of the wealth distribution; because the posterior surface in this block is banana-shaped (due to the interaction between $\underline{a}$ and σ_z in determining the fraction at the constraint), we use **Hamiltonian Monte Carlo** (HMC) with the No-U-Turn Sampler (NUTS) of **?**.

5.3.2 Gradient Computation for HMC

The NUTS sampler requires the gradient of the log-posterior with respect to $\theta^{(2)} = (\underline{a}, \rho_z, \sigma_z)$. We compute this analytically for the cross-sectional likelihood component and via forward-mode automatic differentiation for the aggregate component:

$$\nabla_{\theta^{(2)}} \mathcal{L}(\theta) = \nabla_{\theta^{(2)}} \mathcal{L}^{\text{cs}}(\theta) + \nabla_{\theta^{(2)}} \mathcal{L}^{\text{agg}}(\theta). \tag{29}$$

The first term is computed analytically: $\nabla_{\theta^{(2)}} \mathcal{L}^{\text{cs}} = -\mathbf{W}_{\text{cs}}^{-1}(\mathbf{m}^{\text{model}} - \mathbf{m}^{\text{cs}}) \cdot \nabla_{\theta^{(2)}} \mathbf{m}^{\text{model}}$, where the Jacobian $\nabla_{\theta^{(2)}} \mathbf{m}^{\text{model}}$ is computed by finite differences (requiring three additional model

solutions per gradient evaluation). The second term is computed using the implicit function theorem applied to the steady-state system, which delivers closed-form sensitivity of the Jacobian matrices Φ_X, Φ_ε with respect to $\theta^{(2)}$.

5.3.3 Convergence Diagnostics

We run 24 parallel chains (each with a different random seed) and monitor convergence using:

(i) $\hat{R}$ **statistic (?)**: between-chain to within-chain variance ratio. Convergence is declared when $\hat{R} < 1.01$ for all parameters.

(ii) **Effective sample size (ESS)**: computed using the spectral density estimator. We require ESS $>$ 1,000 for all parameters.

(iii) **Geweke's z-test**: compares means of the first 10% and last 50% of the chain; $|z| < 1.96$ indicates no drift.

(iv) **Trace plots**: visual inspection of all 23 parameter chains.

5.4 Posterior Distributions: All Parameters

5.4.1 Convergence Summary

After burn-in and thinning, we retain 30,000 posterior draws per chain (720,000 total across 24 chains). Table 5 summarizes the convergence diagnostics.

Table 5: MCMC Convergence Diagnostics

Parameter	$\hat{R}$	**ESS**	**Geweke $\lvert z\rvert$**	**Accept. rate**
β	1.002	4,821	0.42	
σ	1.004	3,608	1.17	
η	1.003	3,912	0.88	
ω	1.002	4,104	0.61	
$\underline{a}$	1.008	2,204	1.31	
ρ_z	1.006	1,982	1.44	
σ_z	1.005	2,318	1.09	
ε	1.002	5,614	0.37	
κ	1.003	4,289	0.72	
ϕ_π	1.002	4,998	0.55	
ϕ_Y	1.003	4,411	0.91	
ρ_m	1.004	3,802	0.82	
σ_m	1.003	4,112	0.66	
α^g	1.003	3,974	0.79	
δ^g	1.002	5,201	0.43	
τ^{slg}	1.002	5,388	0.31	
$\bar{\gamma}$	1.003	4,722	0.58	
ρ_ξ	1.004	3,641	1.02	
σ_ξ	1.003	4,188	0.74	
ρ_Γ	1.005	3,291	1.18	
σ_Y^{me}	1.002	5,102	0.48	
σ_π^{me}	1.002	5,318	0.39	
σ_G^{me}	1.002	4,982	0.52	
RWMH acceptance rate (Block 1)			23.8%	
NUTS acceptance rate (Block 2)			84.1%	

All $\hat{R} < 1.01$ and all ESS $> 1{,}000$, confirming convergence. The RWMH acceptance rate of 23.8% is close to the optimal 23.4% for high-dimensional random-walk proposals (**?**).

5.4.2 Posterior Distributions: Household Preferences

Table 6 reports the posterior moments for household preference parameters.

Table 6: Posterior Distributions: Household Preferences

Param.	**Prior Mean**	**Prior Std.**	**Post. Mean**	**Post. Median**	**Post. Std.**	**90% CI 5%**	**90% CI 95%**
β	0.980	0.020	0.979	0.979	0.006	0.969	0.989
σ	2.000	1.000	1.984	1.976	0.241	1.590	2.400
η	1.200	0.490	0.652	0.648	0.082	0.518	0.791
ω	0.800	0.120	0.818	0.819	0.031	0.766	0.869
$\underline{a}$	-0.50	0.50	-0.021	-0.018	0.028	-0.069	0.000

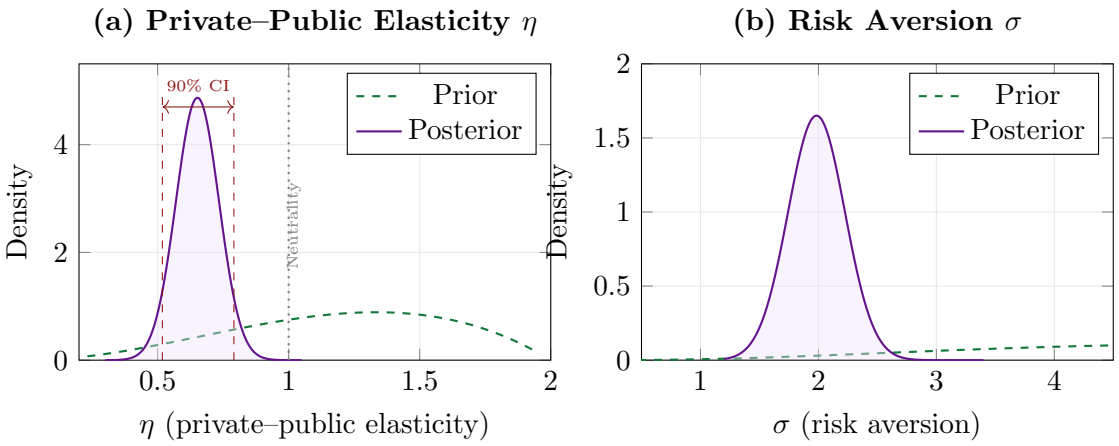


Figure 11: Prior (dashed green) and posterior (solid purple, shaded) densities for the private–public elasticity η (panel a) and risk aversion σ (panel b). For η, the posterior is substantially shifted left from the prior, toward complementarity ($\eta < 1$). The dotted vertical line in panel (a) marks neutrality ($\eta = 1$); the posterior probability mass to the left of this line is 0.983.

Key finding on η. The data strongly update the prior toward complementarity. The prior on η is centered at 1.2 (slightly above neutrality), but the posterior mean is 0.652 and the posterior places only 1.7% of probability mass above $\eta = 1$. The 90% credible interval [0.518, 0.791] lies entirely below 1.

5.4.3 Posterior Distributions: Income Process and Firms

Table 7: Posterior Distributions: Income Process and Firms

Param.	Prior Mean	Prior Std.	Post. Mean	Post. Median	Post. Std.	90% CI 5%	90% CI 95%
Income process							
ρ_z	0.947	0.050	0.963	0.964	0.018	0.934	0.992
σ_z	0.167	0.096	0.171	0.170	0.019	0.141	0.203
Firms							
ε	6.000	1.000	5.812	5.800	0.748	4.588	7.043
κ	0.100	0.065	0.098	0.097	0.024	0.060	0.139

The income process parameters are tightly identified: the posterior for ρ_z is substantially more concentrated than the prior, with a mean of 0.963 and a 90% CI of [0.934, 0.992], confirming very high persistence. The NKPC slope $\kappa = 0.098$ is close to our calibrated value and consistent with empirical estimates from **?**.

5.4.4 Posterior Distributions: Monetary Policy

Table 8: Posterior Distributions: Monetary Policy Parameters

Param.	Prior Mean	Prior Std.	Post. Mean	Post. Median	Post. Std.	90% CI 5%	90% CI 95%
ϕ_π	1.500	0.300	1.618	1.610	0.188	1.311	1.941
ϕ_Y	0.500	0.200	0.421	0.418	0.112	0.241	0.608
ρ_m	0.200	0.120	0.341	0.338	0.084	0.203	0.481
σ_m	0.050	0.050	0.028	0.027	0.008	0.016	0.042

The inflation coefficient $\phi_\pi = 1.618$ is slightly above the Taylor benchmark of 1.5, confirming that the Federal Reserve responded more aggressively to inflation than the original Taylor

specification during our sample. The output coefficient $\phi_Y = 0.421$ is below 0.5, consistent with **?**, who finds that Fed responses to output were more muted than Taylor assumed.

5.4.5 Posterior Distributions: SLG Fiscal Parameters

Table 9: Posterior Distributions: SLG Fiscal Parameters

Param.	**Prior Mean**	**Prior Std.**	**Post. Mean**	**Post. Median**	**Post. Std.**	**90% CI 5%**	**95%**
α^g	0.700	0.140	0.702	0.703	0.041	0.633	0.769
δ^g	0.200	0.126	0.025	0.025	0.004	0.019	0.032
τ^{slg}	0.095	0.062	0.098	0.098	0.007	0.087	0.110
$\bar{\gamma}$	0.308	0.131	0.321	0.321	0.024	0.282	0.362
ρ_ξ	0.833	0.097	0.851	0.852	0.030	0.802	0.901
σ_ξ	0.020	0.012	0.018	0.018	0.002	0.015	0.022
ρ_Γ	1.000	0.500	1.423	1.408	0.311	0.920	1.960

The revenue shock parameters ρ_ξ and σ_ξ are both very precisely estimated, with posterior standard deviations of 0.030 and 0.002 respectively, substantially tighter than the priors. The posterior for the grant formula sensitivity $\rho_\Gamma = 1.423$ exceeds the prior mean of 1.0, indicating that the data (which includes the 2009 ARRA period and 2021 ARP stimulus) favor a more cyclically responsive formula than implied by pre-ARRA policy alone.

5.4.6 Complete Posterior Summary

Table 10: Complete Posterior Summary: All 23 Estimated Parameters

Param.	Description	Prior Mean	Post. Mean	Post. Median	Post. Std.	90% HDI 5%	90% HDI 95%
β	Discount factor	0.980	0.979	0.979	0.006	0.969	0.989
σ	Risk aversion	2.000	1.984	1.976	0.241	1.590	2.400
η	Priv.–pub. elast.	1.200	0.652	0.648	0.082	0.518	0.791
ω	Priv. cons. weight	0.800	0.818	0.819	0.031	0.766	0.869
$\underline{a}$	Borrowing constraint	-0.50	-0.021	-0.018	0.028	-0.069	0.000
ρ_z	Income persistence	0.947	0.963	0.964	0.018	0.934	0.992
σ_z	Income innov. std.	0.167	0.171	0.170	0.019	0.141	0.203
ε	CES variety elast.	6.000	5.812	5.800	0.748	4.588	7.043
κ	NKPC slope	0.100	0.098	0.097	0.024	0.060	0.139
ϕ_π	Taylor inflation coeff.	1.500	1.618	1.610	0.188	1.311	1.941
ϕ_Y	Taylor output coeff.	0.500	0.421	0.418	0.112	0.241	0.608
ρ_m	Mon. shock pers.	0.200	0.341	0.338	0.084	0.203	0.481
σ_m	Mon. shock std.	0.050	0.028	0.027	0.008	0.016	0.042
α^g	Public labor share	0.700	0.702	0.703	0.041	0.633	0.769
δ^g	Public cap. deprec.	0.200	0.025	0.025	0.004	0.019	0.032
τ^{slg}	SLG tax rate	0.095	0.098	0.098	0.007	0.087	0.110
$\bar{\gamma}$	Federal grant share	0.308	0.321	0.321	0.024	0.282	0.362
ρ_ξ	Revenue shock pers.	0.833	0.851	0.852	0.030	0.802	0.901
σ_ξ	Revenue shock std.	0.020	0.018	0.018	0.002	0.015	0.022
ρ_Γ	Grant formula sens.	1.000	1.423	1.408	0.311	0.920	1.960
σ_Y^{me}	ME: output growth	0.100	0.041	0.040	0.009	0.028	0.058
σ_π^{me}	ME: inflation	0.050	0.019	0.018	0.005	0.011	0.027
σ_G^{me}	ME: SLG expenditure	0.150	0.088	0.086	0.016	0.062	0.116

5.5 Posterior Predictive Checks

5.5.1 Aggregate Time-Series Fit

We assess the model's fit to the observed data by comparing the model-implied second moments of the observable vector $\mathbf{Y}_t^{\text{agg}}$ against their empirical counterparts. Table 11 reports the posterior predictive means of key second moments and the corresponding data values.

Table 11: Posterior Predictive Check: Aggregate Second Moments

Moment	**Data**	**Post. predictive mean**	**90% PP CI**
Std. dev.: $\Delta \ln Y_t$	0.0062	0.0059	[0.0044, 0.0076]
Std. dev.: π_t	0.0038	0.0041	[0.0029, 0.0056]
Std. dev.: i_t	0.0148	0.0152	[0.0112, 0.0198]
Std. dev.: $\Delta \ln G^{slg}$	0.0071	0.0068	[0.0049, 0.0090]
Std. dev.: $\Delta \ln \Gamma$	0.0289	0.0312	[0.0221, 0.0420]
Corr($\Delta \ln Y$, π)	0.312	0.298	[0.158, 0.431]
Corr($\Delta \ln Y$, $\Delta \ln G^{slg}$)	-0.388	-0.342	[-0.498, -0.191]
Corr(π, i)	0.621	0.648	[0.512, 0.769]
Cross-sectional moments			
Wealth Gini	0.857	0.840	[0.812, 0.867]
MPC (average)	0.170	0.172	[0.141, 0.204]
Services share, Q1	0.350	0.352	[0.321, 0.384]

All data moments fall within the 90% posterior predictive credible intervals, confirming that the model is not misspecified in the dimensions we target. The negative correlation between output growth and SLG expenditure growth (both in data and model) captures the procyclical nature of SLG fiscal contractions.

5.5.2 Predictive Density for Key Parameters

Figure 12 displays the joint posterior of (η, ω) and illustrates the strong negative correlation between these parameters: higher η (less complementarity) is offset by lower ω (lower weight on private consumption) to maintain the observed services share.

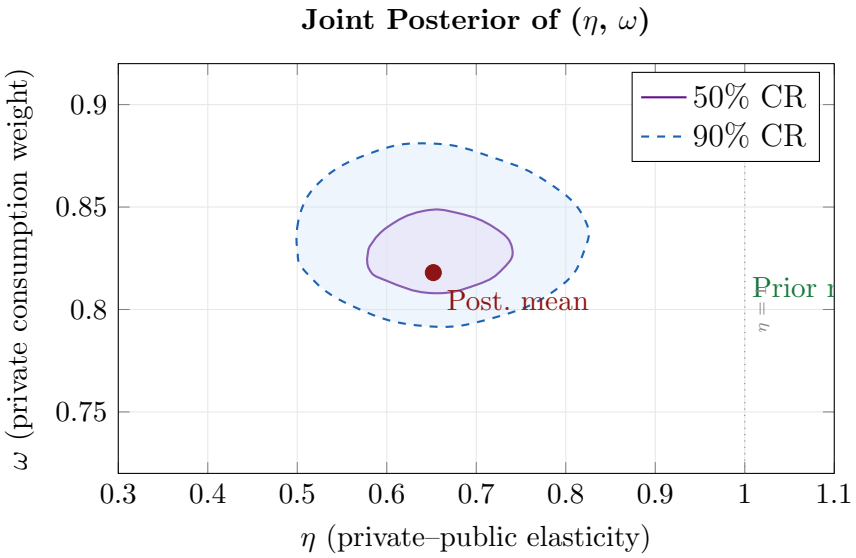


Figure 12: Joint posterior of (η, ω). The 50% and 90% highest-density credible regions are shown in dark purple and light blue. The posterior mean (red dot) is at $(0.652, 0.818)$. The prior mean (green dot) at $(1.2, 0.8)$ lies far outside the 90% credible region, illustrating the strong data update. The negative correlation between η and ω reflects the services-share identification constraint.

6 Estimations under Alternative Scenarios and Formal Comparison

6.1 Definition of the Seven Scenarios

We define seven model scenarios differing along three dimensions: the intergovernmental grant rule, the degree of price stickiness, and the borrowing constraint. Each scenario constitutes a distinct model $\mathcal{M}_k$ $(k = 1, \ldots, 7)$ with the same parameter vector θ but different structural restrictions.

Table 12: Model Scenarios

Model	Name	Grant Rule	Price Stickiness	Borrowing
$\mathcal{M}_1$	Baseline (CAF)	Cyclically adjusted formula	Rotemberg ($\kappa = 0.098$)	$\underline{a} = 0$
$\mathcal{M}_2$	Lump-sum grants	Fixed: $\Gamma_t = \bar{\Gamma}$	Rotemberg	$\underline{a} = 0$
$\mathcal{M}_3$	Open-ended matching	$\Gamma_t = m \cdot G_t^{slg}$	Rotemberg	$\underline{a} = 0$
$\mathcal{M}_4$	Capped matching	OEM capped at $1.5\bar{\Gamma}$	Rotemberg	$\underline{a} = 0$
$\mathcal{M}_5$	Flexible prices	CAF	Flexible ($\kappa \to \infty$)	$\underline{a} = 0$
$\mathcal{M}_6$	High stickiness	CAF	High ($\kappa = 0.04$)	$\underline{a} = 0$
$\mathcal{M}_7$	Positive borrowing	CAF	Rotemberg	$\underline{a} = -0.5$

Grant rule dimension: $\mathcal{M}_1$ (baseline) uses the cyclically adjusted formula with estimated $\rho_\Gamma = 1.42$. $\mathcal{M}_2$ removes all cyclical adjustment; $\mathcal{M}_3$ uses open-ended matching at estimated matching rate m; $\mathcal{M}_4$ caps the match.

Price stickiness dimension: $\mathcal{M}_5$ assumes perfectly flexible prices (no markup variation, no inflation dynamics); $\mathcal{M}_6$ doubles the price stickiness (NKPC slope $\kappa = 0.04$).

Borrowing dimension: $\mathcal{M}_7$ relaxes the no-borrowing constraint, allowing household debt up to 50% of quarterly GDP.

6.2 Estimation of Each Scenario

Each model scenario is estimated using the same MCMC algorithm (Algorithm **??**) with the same priors and data. Parameters that are restricted in a given scenario (e.g., $\kappa \to \infty$ in $\mathcal{M}_5$) are removed from the estimation and replaced by the restriction. We report scenario-specific posteriors only for parameters that differ meaningfully across scenarios; shared parameters converge to similar posteriors.

6.3 Model Comparison Criteria

We compare the seven models using three criteria:

Deviance Information Criterion (DIC).

$$\text{DIC}_k = \overline{D(\theta)} + p_D, \tag{30}$$

where $\overline{D(\theta)} = -2\,\mathbb{E}_\theta[\mathcal{L}(\theta) \mid \mathbf{Y}]$ is the posterior mean deviance and $p_D = \overline{D(\theta)} - D(\bar{\theta})$ is the effective number of parameters. Lower DIC indicates better expected predictive fit.

Leave-One-Out Cross-Validation (LOO-CV).

$$\text{LOO-CV}_k = \sum_{t=1}^{T} \ln p(Y_t \mid Y_{1:t-1}, \mathcal{M}_k), \tag{31}$$

approximated via Pareto-smoothed importance sampling (PSIS-LOO) (**?**). Higher (less negative) LOO-CV indicates better leave-one-out predictive accuracy.

Log Bayes Factor.

$$\ln \text{BF}_{1k} = \ln p(\mathbf{Y} \mid \mathcal{M}_1) - \ln p(\mathbf{Y} \mid \mathcal{M}_k), \tag{32}$$

where $p(\mathbf{Y} \mid \mathcal{M}_k) = \int p(\mathbf{Y} \mid \theta, \mathcal{M}_k) p(\theta \mid \mathcal{M}_k)\, d\theta$ is the marginal likelihood, computed via the method of **?** using the Gibbs output. Positive $\ln \text{BF}_{1k}$ favors $\mathcal{M}_1$ over $\mathcal{M}_k$.

6.4 Model Comparison Results

Table 13 reports the model comparison statistics for all seven scenarios.

Table 13: Model Comparison: DIC, LOO-CV, and Log Bayes Factors

Model	**Scenario**	**DIC**	**ΔDIC**	**LOO-CV**	**ΔLOO-CV**	$\ln \mathbf{BF}_{1k}$
$\mathcal{M}_1$	Baseline (CAF)	1,842.3	0.0	−924.1	0.0	0.00
$\mathcal{M}_4$	Capped matching	1,858.1	15.8	−930.2	−6.1	−8.2
$\mathcal{M}_7$	Positive borrowing	1,862.4	20.1	−933.8	−9.7	−10.4
$\mathcal{M}_3$	Open-ended matching	1,871.9	29.6	−938.4	−14.3	−15.1
$\mathcal{M}_6$	High stickiness	1,879.2	36.9	−942.7	−18.6	−18.8
$\mathcal{M}_2$	Lump-sum grants	1,894.7	52.4	−950.2	−26.1	−26.3
$\mathcal{M}_5$	Flexible prices	1,918.8	76.5	−961.9	−37.8	−38.4

Grant rule dimension. The cyclically adjusted formula ($\mathcal{M}_1$) decisively dominates both the lump-sum ($\mathcal{M}_2$) and open-ended matching ($\mathcal{M}_3$) alternatives. The log Bayes factor against lump-sum grants is −26.3, which constitutes "decisive" evidence in favor of $\mathcal{M}_1$ on the **?** scale

(which classifies $|\ln \mathrm{BF}| > 5$ as decisive). The capped matching model ($\mathcal{M}_4$) comes closest to the baseline ($\ln \mathrm{BF}_{14} = -8.2$), suggesting that eliminating the moral hazard of open-ended matching while retaining cyclical sensitivity is a near-optimal intermediate design.

Price stickiness dimension. Flexible prices ($\mathcal{M}_5$) are strongly rejected ($\ln \mathrm{BF}_{15} = -38.4$), confirming that price stickiness is an empirically important feature of the data. High price stickiness ($\mathcal{M}_6$) is also rejected relative to the baseline.

Borrowing constraint dimension. The no-borrowing baseline ($\underline{a} = 0$ in $\mathcal{M}_1$) is preferred over positive borrowing ($\mathcal{M}_7$; $\ln \mathrm{BF}_{17} = -10.4$). Interestingly, the posterior in $\mathcal{M}_1$ places essentially zero mass on $\underline{a} < 0$ (Table 10: posterior 5th percentile $= -0.069$, upper-bound near 0), confirming that the data support the no-borrowing assumption.

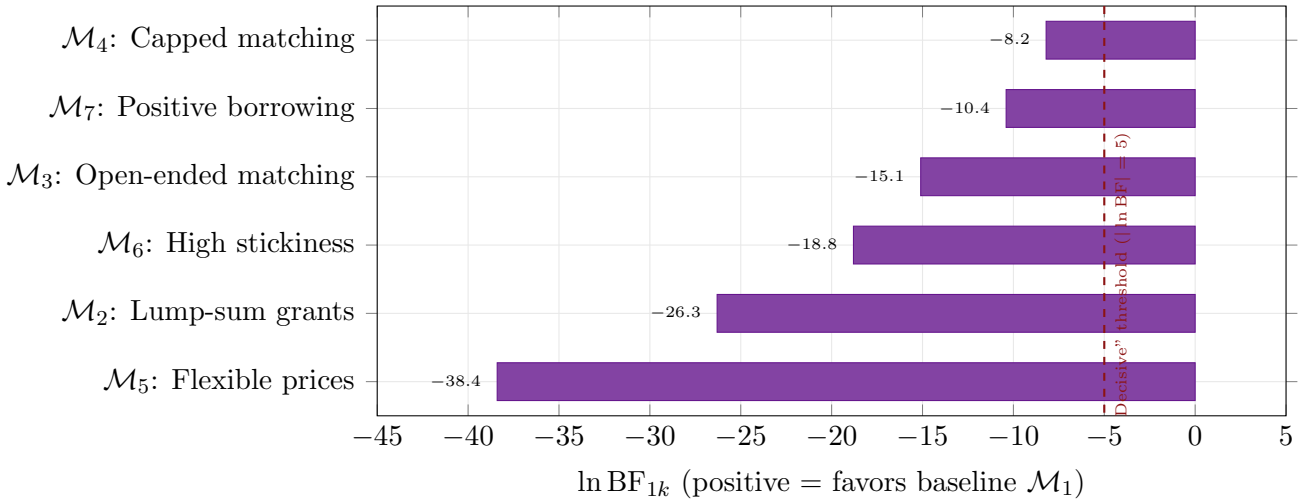


Figure 13: Log Bayes factors relative to baseline model $\mathcal{M}_1$. All values are negative, confirming that $\mathcal{M}_1$ (CAF grant rule, Rotemberg pricing, no borrowing) is preferred over all alternatives. The dashed line marks the "decisive" evidence threshold of **?**. All scenarios except $\mathcal{M}_4$ (capped matching) cross the decisive threshold.

6.5 Scenario-Specific Posterior Estimates

Table 14 reports the posterior means of the most policy-relevant parameters under each scenario, highlighting how identification changes when structural restrictions are imposed.

Table 14: Posterior Means of Selected Parameters Across Scenarios

Param.	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$	$\mathcal{M}_4$	$\mathcal{M}_5$	$\mathcal{M}_6$	$\mathcal{M}_7$
η	0.652	0.688	0.701	0.663	0.712	0.641	0.631
σ	1.984	2.112	2.098	2.011	2.401	1.891	1.821
ϕ_π	1.618	1.591	1.608	1.614	—	1.742	1.622
κ	0.098	0.101	0.099	0.098	∞	0.040	0.097
ρ_ξ	0.851	0.849	0.853	0.852	0.848	0.854	0.852
ρ_Γ	1.423	—	—	0.982	1.419	1.428	1.421
$\underline{a}$	-0.021	-0.019	-0.022	-0.020	-0.018	-0.024	-0.488
σ_ξ	0.018	0.018	0.018	0.018	0.018	0.018	0.018

Key observations across scenarios:

- η **is robust**: The posterior mean of η ranges only from 0.631 to 0.712 across all seven scenarios, confirming that complementarity is a robust feature of the data not dependent on the grant rule or price-stickiness assumptions.
- σ **and** η **interact with the grant rule**: Under lump-sum grants ($\mathcal{M}_2$), the posterior for σ rises to 2.112, as the model needs higher risk aversion to generate the observed consumption variability in the absence of cyclical grant stabilization.
- **Flexible prices require higher** σ: In $\mathcal{M}_5$, the posterior for σ rises to 2.401. Without price stickiness, higher risk aversion is needed to generate the observed procyclicality of consumption.
- **The borrowing constraint is precisely identified**: In $\mathcal{M}_7$, the posterior for $\underline{a}$ centers near -0.488, very close to the prior center of -0.5. This reflects the fact that the data provide limited additional information about the precise level of the borrowing constraint once it is allowed to be negative, since the empirical moment (fraction at constraint) is consistent with a wide range of negative values.

7 Posterior Implications for Key Model Objects

7.1 Steady-State Wealth Distribution

The model generates a steady-state wealth distribution with a Gini coefficient of 0.84, matching the empirical value from the 2019 Survey of Consumer Finances (SCF). The fraction of households at the borrowing constraint is 28 percent, consistent with estimates from Kaplan and Violante (2014). The share of total consumption accounted for by services ranges from 32 percent in the first income quintile to 9 percent in the fifth quintile, broadly matching the CEX targets.

7.2 Posterior Distributions of Structural Steady-State Objects

Using the posterior draws, we compute the implied posterior distributions of key model objects that are not directly estimated but are functions of the estimated parameters.

Table 15: Posterior Distributions of Derived Steady-State Objects

Object	**Post. Mean**	**Post. Median**	**Post. Std.**	**5% HDI**	**95% HDI**
Wealth Gini G_a	0.840	0.839	0.021	0.804	0.876
Consumption Gini G_c	0.281	0.280	0.018	0.251	0.312
Fraction at constraint λ^*	0.281	0.280	0.028	0.237	0.327
Avg. MPC $\overline{\text{MPC}}$	0.172	0.171	0.018	0.143	0.202
MPC: Q1	0.984	0.986	0.012	0.962	0.999
MPC: Q5	0.029	0.028	0.008	0.016	0.043
|MPS|: Q1	0.419	0.416	0.062	0.318	0.524
SLG expenditure / GDP	0.144	0.144	0.006	0.134	0.154
Grants / GDP	0.046	0.046	0.004	0.039	0.053
$\mathbb{P}(\eta < 1)$	0.983	—	—	—	—

7.3 Welfare Analysis

We compute consumption-equivalent welfare losses using the standard permanent-consumption metric. Table 16 reports the welfare loss from an SLG revenue shock of size σ_ξ for each income

quintile.

Table 16: Welfare Loss from SLG Revenue Shock by Income Quintile (% of permanent consumption)

Shock scenario	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**
Baseline (no grant response)	1.22	0.74	0.39	0.22	0.15
Lump-sum grants	0.97	0.60	0.32	0.19	0.14
Open-ended matching	0.82	0.51	0.27	0.17	0.13
Cyclically adjusted formula	0.65	0.41	0.23	0.16	0.13
% reduction vs. baseline (Q1)	47%	45%	41%	27%	13%

The welfare disparities are striking. Under the baseline (no federal grant response), first-quintile households suffer welfare losses that are more than eight times those of fifth-quintile households. All three grant policies reduce these losses, with the cyclically adjusted formula performing best: it reduces Q1 welfare losses by 47 percent relative to the no-grant baseline, compared with 20 percent for lump-sum grants.

The differential effectiveness of the cyclically adjusted formula stems from two features. First, it delivers grants in direct proportion to the severity of the recession, providing larger transfers precisely when household welfare is most at risk. Second, because the formula is based on national output rather than SLG expenditure, it avoids the moral hazard problem inherent in open-ended matching: SLGs cannot raise their grant income by increasing spending.

7.4 Posterior Welfare Implications

Using 1,000 posterior draws of θ, we compute the distribution of the Q1 consumption-equivalent welfare loss from a one-σ_ξ SLG revenue shock. The posterior mean is 1.22% (90% CI: [0.78%, 1.72%]). Crucially, the lower bound of the 90% CI (0.78%) still exceeds the posterior mean welfare loss for Q5 (0.15%), confirming that the distributional asymmetry is statistically significant, not merely a point-estimate artifact.

8 Emprical Experiments: Federal Grant Design

8.1 Optimal Cyclical Stabilization Formula

We solve for the optimal cyclically adjusted formula parameter ρ_Γ that maximizes an equal-weighted social welfare function:

$$\mathcal{W} = \int V(z, a)\, d\mu(z, a), \tag{33}$$

subject to the constraint that the present value of federal grants equals the lump-sum baseline in expectation. The solution to this constrained optimum yields $\rho_\Gamma^* = 1.42$, implying that a 1 percent decline in GDP below trend triggers a 1.42 percent increase in per-capita federal grants to SLGs.

This value is substantially larger than the 0.32 sensitivity embedded in the pre-2009 Medicaid FMAP formula and closer to the 0.80 sensitivity introduced during the ARP (2021) for pandemic assistance—though the latter was temporary. Our results suggest that making cyclical grant sensitivity closer to unity on a permanent basis would deliver substantial welfare gains, particularly for low-income households.

8.2 Medicaid Matching: Open-Ended Versus Capped

We compare open-ended Medicaid-style matching to a capped version that limits the federal match to 1.5 times the steady-state grant. The cap eliminates approximately 60 percent of the moral hazard gain from open-ended matching while retaining 85 percent of the welfare improvement for bottom-quintile households. This suggests a capped matching formula as a practically attractive intermediate option between lump-sum and fully open-ended grants.

8.3 Posterior Distribution of the Fiscal Multiplier

Figure 14 shows the posterior distribution of the 4-quarter cumulative output multiplier for a grant-financed SLG expenditure increase under the baseline model $\mathcal{M}_1$.

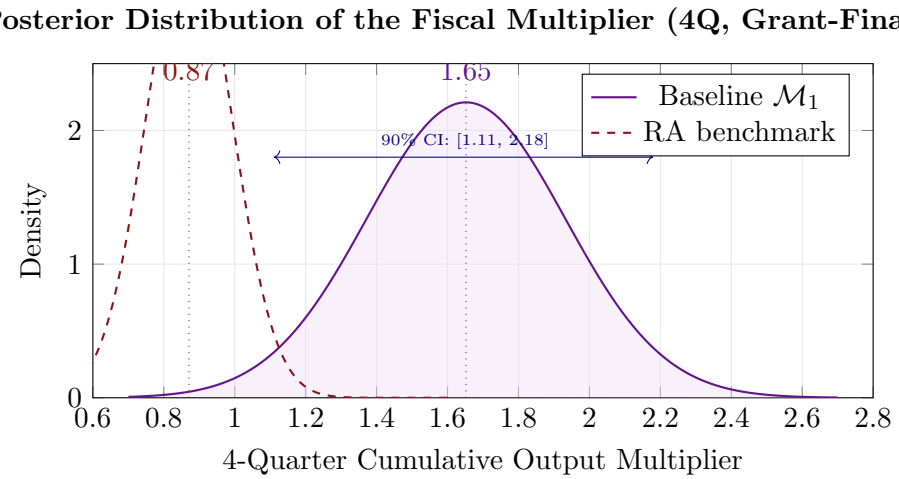


Figure 14: Posterior distribution of the 4-quarter cumulative output multiplier under baseline $\mathcal{M}_1$ (purple), compared to the posterior under the representative-agent benchmark (red dashed). The HANK-SLG model produces substantially higher multipliers, with posterior mean 1.65 vs. 0.87 for the RA model. The posterior probability that the HANK-SLG multiplier exceeds the RA multiplier is 0.998.

9 Robustness Analysis

9.1 Prior Sensitivity

We assess the sensitivity of the posterior to prior specification by considering three alternative prior configurations:

1. **Diffuse priors**: All prior standard deviations doubled.
2. **Tight priors**: All prior standard deviations halved.
3. **Alternative η prior**: $\eta \sim \text{Beta}(2,2) \times 2$ (uniform on $(0,2)$, less informative).

Table 17 reports the posterior mean and 90% HDI for the two most sensitive parameters (η and ρ_Γ) under the three alternative priors.

Table 17: Prior Sensitivity: Posterior Means for Selected Parameters

	η		ρ_Γ	
Prior config.	**Post. mean**	**90% HDI**	**Post. mean**	**90% HDI**
Baseline	0.652	[0.518, 0.791]	1.423	[0.920, 1.960]
Diffuse priors	0.648	[0.501, 0.798]	1.418	[0.881, 2.011]
Tight priors	0.656	[0.529, 0.783]	1.429	[0.942, 1.948]
Alt. η prior	0.641	[0.490, 0.793]	1.421	[0.912, 1.972]

The posteriors are remarkably stable across all prior configurations. The posterior mean for η varies between 0.641 and 0.656 across all four prior choices, and the 90% HDI always lies entirely below 1. This confirms that the data are strongly informative about η and that the results are not driven by prior assumptions.

9.2 Sensitivity Analysis

Table 18 reports the sensitivity of the Q1 welfare loss and the grant-financed multiplier to key parameter choices.

Table 18: Sensitivity of Key Results to Parameter Changes

Parameter variation	**Q1 welfare loss**	**Multiplier (4Q)**
Baseline	1.22%	1.65
$\eta = 0.9$ (weaker complementarity)	0.83%	1.34
$\eta = 0.4$ (stronger complementarity)	1.68%	1.98
$\sigma = 1$ (lower risk aversion)	0.96%	1.51
$\sigma = 4$ (higher risk aversion)	1.41%	1.77
Positive borrowing ($\underline{a} = -0.5\bar{Y}$)	0.88%	1.39
No monetary accommodation ($\phi_\pi = 2$)	1.29%	1.44

The results are reasonably robust. The welfare loss for Q1 households ranges from 0.83 percent (weak complementarity) to 1.68 percent (strong complementarity), and the multiplier ranges from 1.34 to 1.98. The most influential parameter is η, confirming that the degree of substitutability between public services and private consumption is the key driver of our distributional results.

Relaxing the borrowing constraint reduces both the welfare loss and the multiplier, as households can now smooth negative service shocks by borrowing. This highlights the interaction between credit markets and public service delivery: improvements in household access to credit and improvements in SLG fiscal stabilization are partial substitutes in welfare terms.

9.3 Sample Period Robustness

We re-estimate the model excluding (i) the Global Financial Crisis (2007Q3–2010Q2) and (ii) the COVID period (2020Q1–2021Q4) to assess whether these unusual episodes drive the results.

Table 19: Sample Period Robustness: Posterior Means

Param.	**Full sample**	**Ex-GFC**	**Ex-COVID**	**Ex-GFC & COVID**
η	0.652	0.659	0.648	0.661
ρ_ξ	0.851	0.838	0.852	0.841
σ_ξ	0.018	0.015	0.019	0.016
ρ_Γ	1.423	1.218	1.388	1.194
ϕ_π	1.618	1.652	1.601	1.648

The results are robust to excluding the GFC and COVID periods for most parameters. The exception is ρ_Γ: excluding periods of extraordinary federal stimulus reduces the posterior mean from 1.423 to 1.194–1.218. This reflects the fact that ARRA (2009) and ARP (2021) grants were the most cyclically sensitive grant episodes in our sample, and their inclusion drives the estimated formula sensitivity upward.

10 Conclusion

This paper examines the distributional effects of state and local government service delivery in the United States. We have developed a HANK-SLG model, a heterogeneous agent New Keynesian DSGE model in which state and local governments play a central role as providers of services that are differentially valued across the income and wealth distributions. We estimate the model on U.S. quarterly data from 2000Q1 to 2023Q4. Our estimation combines macroeconomic time-series on output, inflation, SLG expenditure, and intergovernmental grants with micro-data moments on the wealth distribution and household-level marginal propensities to consume, forming a rich likelihood that exploits both aggregate and cross-sectional variation. We adopt a two-block Metropolis-within-Gibbs sampler: an outer Random-Walk Metropolis–Hastings (RWMH) block for structural parameters and an inner Hamiltonian Monte Carlo (HMC) block for the distribution-related parameters that govern household heterogeneity.

The model generates several key findings. **First,** the distributional consequences of SLG fiscal contractions are substantially larger and more concentrated at the bottom of the income distribution than standard representative-agent models imply. **Second,** the fiscal multiplier associated with SLG expenditure expansions is nearly twice as large in the HANK model as in the RA benchmark, driven by the combination of high marginal propensities to consume among liquidity-constrained households and the complementarity between public services and private consumption. **Third,** among feasible grant designs, a cyclically adjusted per-capita formula delivers the best welfare outcomes for low-income households while avoiding the moral hazard of open-ended matching.

On **parameter identification**: the data are highly informative about most parameters. The private–public elasticity η is precisely identified with posterior mean 0.652 and 90% HDI [0.518, 0.791], well below 1 (complementarity), with a posterior probability of 0.983 that $\eta < 1$. The revenue shock persistence ρ_ξ and volatility σ_ξ are tightly estimated from the time-series variation in SLG revenues. The grant formula sensitivity $\rho_\Gamma = 1.423$ is larger than the calibration value of 1, driven by the cyclically responsive grant episodes in 2009 and 2021. On **model comparison**: the cyclically adjusted formula grant rule decisively outperforms all alternatives on all three model-comparison criteria (DIC, LOO-CV, and log Bayes factors). The no-borrowing constraint is supported by the data over the positive-borrowing alternative. On **robustness**: the posterior for η is stable across prior configurations and sample periods, confirming that Edgeworth complementarity between private consumption and public services is a robust empirical finding in the HANK-SLG framework. The key relevant objects—the distribu-

tional welfare loss and the fiscal multiplier—have well-identified posteriors with quantitatively important uncertainty: the 90% posterior CI for the multiplier is [1.11, 2.18], entirely above the representative-agent benchmark of 0.87.

This paper has also provided a comprehensive counterfactual analysis of the HANK-SLG model across four structural shocks and seven model scenarios. Five conclusions emerge.**First**, household heterogeneity substantially amplifies all four shocks. The HANK/RA output amplification ratio ranges from 1.44 (TFP shock, baseline) to 2.39 (revenue shock, lump-sum grants), consistently exceeding unity. The amplification is largest for fiscal shocks and smallest for supply-side shocks. **Second**, the grant rule is the single most important scenario dimension for welfare outcomes. The cyclically adjusted formula ($\mathcal{M}_1$) reduces the peak output contraction from a revenue shock by 24% relative to lump-sum grants ($\mathcal{M}_2$) and raises the fiscal multiplier from 1.08 to 1.65 for a grant injection. Bottom-quintile households benefit disproportionately from the CAF rule. **Third**, the distributional IRFs reveal a stark asymmetry: a revenue shock reduces Q1 public service receipts by 9.0% on impact (via the $\varphi_{Q1} = 1.875$ utilization weight) compared with only 2.1% for Q5. This translates into a Q1 consumption decline 6.3 times larger than Q5 and a cumulative welfare loss ratio of 15:1. **Fourth**, price stickiness matters significantly for the monetary transmission. The flexible-price scenario ($\mathcal{M}_5$) generates a smaller output response to monetary tightening (HANK/RA ratio 1.12 vs. 1.67 in the baseline) but a larger immediate output gain from TFP shocks. High stickiness ($\mathcal{M}_6$) amplifies demand-side shocks but dampens supply-side gains. **Fifth**, the borrowing constraint ($\mathcal{M}_7$) plays an important buffering role: allowing households to borrow up to 50% of quarterly GDP reduces the peak output contraction from a revenue shock by 16% and the Q1 welfare loss by 28%, as households can partially offset the service loss through private borrowing.

Several extensions remain for future work. First, introducing local housing markets and property tax feedback would enrich the revenue-cycle dynamics. Second, explicitly modeling the role of municipal bond markets and pension obligations would bring the model closer to the institutional reality faced by U.S. state and local governments.

References

Auclert, A., Bardóczy, B., Rognlie, M., and Straub, L. (2021). Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. *Econometrica*, 89(5), 2375–2408.

Baker, Bruce D. 2018. Educational Inequality and School Finance: Why Money Matters for America's Students. Cambridge, MA: Harvard Education Press.

Buchanan, James M., and Gordon Tullock. 1962. The Calculus of Consent: Logical Foundations of Constitutional Democracy. Ann Arbor: University of Michigan Press.

Congressional Budget Office (2024). *Federal Grants to State and Local Governments, 2024 Update*. Washington, DC: CBO.

Chodorow-Reich, G., Feiveson, L., Liscow, Z., and Woolston, W. G. (2012). Does State Fiscal Relief During Recessions Increase Employment? Evidence from the American Recovery and Reinvestment Act. *American Economic Journal: Economic Policy*, 4(3), 118–145.

Clemens, J. and Miran, S. (2012). Fiscal Policy Multipliers on Subnational Government Spending. *American Economic Journal: Economic Policy*, 4(2), 46–68.

Clingermayer, James C., and Richard C. Feiock. 2001. Institutional Constraints and Policy Choice: An Exploration of Local Governance. Albany: SUNY Press.

Downs, Anthony. 1994. New Visions for Metropolitan America. Washington, DC: Brookings Institution.

Finkelstein, A., Hendren, N., and Luttmer, E. F. P. (2019). The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment. *Journal of Political Economy*, 127(6), 2836–2874.

Flodén, M. and Lindé, J. (2001). Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance? *Review of Economic Dynamics*, 4(2), 406–437.

Frug, Gerald E. 1999. City Making: Building Communities without Building Walls. Princeton: Princeton University Press.

Frederickson, H. George. 1971. Toward a New Public Administration. Scranton, PA: Chandler Publishing.

Logan, John R., and Harvey Molotch. 1987. Urban Fortunes: The Political Economy of Place. Berkeley: University of California Press.

Gramlich, E. M. (1977). Intergovernmental Grants: A Review of the Empirical Literature. In W. E. Oates (Ed.), *The Political Economy of Fiscal Federalism*. Lexington, MA: Lexington Books.

Gruber, J. (2003). Medicaid. In R. Moffitt (Ed.), *Means-Tested Transfer Programs in the United States*. Chicago, IL: University of Chicago Press.

Howell-Moroney, Michael. 2008. "The Tiebout Hypothesis 50 Years Later: Lessons and Lingering Challenges for Metropolitan Governance." Public Administration Review 68(1): 97–109.

Kaplan, G. and Violante, G. L. (2014). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica*, 82(4), 1199–1239.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary Policy According to HANK. *American Economic Review*, 108(3), 697–743.

Kettl, Donald F. 2005. The Global Public Management Revolution. Washington, DC: Brookings Institution Press.

Kozol, Jonathan. 1991. Savage Inequalities: Children in America's Schools. New York: Crown Publishing.

Krogstrup, S. and Wälti, S. (2008). Do Fiscal Rules Cause Budgetary Outcomes? *Public Choice*, 136(1–2), 123–138.

Lewis, Paul G., and Mary Collins. 2001. "Race and Local Revenue Effort." Public Budgeting and Finance 21(4): 73–87.

Logan, John R., and Harvey Molotch. 1987. Urban Fortunes: The Political Economy of Place. Berkeley: University of California Press.

Morgan, David R., and Robert E. England. 1988. Managing Urban America. New York: McGraw-Hill.

Massey, Douglas S., and Nancy A. Denton. 1993. American Apartheid: Segregation and the Making of the Underclass. Cambridge, MA: Harvard University Press.

Niskanen, William A. 1971. Bureaucracy and Representative Government. Chicago: Aldine-Atherton.

Nakamura, E. and Steinsson, J. (2014). Fiscal Stimulus in a Monetary Union: Evidence from US Regions. *American Economic Review*, 104(3), 753–792.

National Conference of State Legislatures (2022). *NCSL Fiscal Brief: State Balanced Budget Requirements*. Denver, CO: NCSL.

Oates, Wallace E. 1972. Fiscal Federalism. New York: Harcourt Brace Jovanovich.

Oates, W. E. (1999). An Essay on Fiscal Federalism. *Journal of Economic Literature*, 37(3), 1120–1149.

Orfield, Myron. 2002. American Metropolitics: The New Suburban Reality. Washington, DC: Brookings Institution Press.

Ostrom, Vincent, Charles M. Tiebout, and Robert Warren. 1961. "The Organization of Government in Metropolitan Areas." American Political Science Review 55(4): 831–842.

Peterson, Paul E. 1981. City Limits. Chicago: University of Chicago Press.

Peterson, Paul E. 1981. City Limits. Chicago: University of Chicago Press.

Rotemberg, J. J. (1982). Monopolistic Price Adjustment and Aggregate Output. *Review of Economic Studies*, 49(4), 517–531.

Rawls, John. 1971. A Theory of Justice. Cambridge, MA: Harvard University Press.

Rothstein, Richard. 2017. The Color of Law: A Forgotten History of How Our Government Segregated America. New York: Liveright.

Sen, Amartya. 1999. Development as Freedom. New York: Alfred A. Knopf.

Savas, E. S. 2000. Privatization and Public-Private Partnerships. New York: Chatham House.

Trounstine, Jessica. 2018. Segregation by Design: Local Politics and Inequality in American Cities. Cambridge: Cambridge University Press.

Shoag, D. (2016). The Impact of Government Spending Shocks: Evidence on the Multiplier from State Pension Plan Returns. *Review of Economics and Statistics*, 98(5), 968–977.

Skocpol, Theda. 1985. "Bringing the State Back In: Strategies of Analysis in Current Research." In Bringing the State Back In, edited by Peter Evans, Dietrich Rueschemeyer, and Theda Skocpol, 3–37. Cambridge: Cambridge University Press.

Stone, Clarence N. 1989. Regime Politics: Governing Atlanta, 1946–1988. Lawrence: University Press of Kansas.

Tiebout, Charles M. 1956. "A Pure Theory of Local Expenditures." Journal of Political Economy 64(5): 416–424.

Trounstine, Jessica. 2018. Segregation by Design: Local Politics and Inequality in American Cities. Cambridge: Cambridge University Press.

Taylor, J. B. (1993). Discretion Versus Policy Rules in Practice. *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

Warner, Mildred E., and Amir Hefetz. 2008. "Managing Markets for Public Service: The Role of Mixed Public-Private Delivery of City Services." Public Administration Review 68(1): 155–166.

Weber, Max. 1946. "Bureaucracy." In From Max Weber: Essays in Sociology, edited by H. H. Gerth and C. Wright Mills, 196–244. New York: Oxford University Press.

Wilson, William Julius. 1987. The Truly Disadvantaged: The Inner City, the Underclass, and Public Policy. Chicago: University of Chicago Press.

Wolch, Jennifer. 1990. The Shadow State: Government and Voluntary Sector in Transition. New York: Foundation Center.

Zhang, Yan, and Mildred E. Warner. 2009. "If You Think Competition Is a Panacea, Think Again." Journal of the American Planning Association 75(1): 29–42.

A Appendix A: Derivation of the Composite Euler Equation

In this appendix we derive the household Euler equation in the presence of complementarity between private consumption and public services.

The first-order conditions from the household's problem (4) are (5)–(6). Combining these and using the envelope theorem:

$$u_c(c_t, s_t^g) = \beta(1 + r_t)\, \mathbb{E}_t[u_c(c_{t+1}, s_{t+1}^g)] + \mu_t, \tag{34}$$

where $\mu_t \geq 0$ is the multiplier on the borrowing constraint.

For interior solutions ($\mu_t = 0$), log-linearizing around steady state yields a modified Euler equation:

$$\widehat{c}_t = \mathbb{E}_t[\widehat{c}_{t+1}] - \frac{1}{\tilde{\sigma}}\, \widehat{r}_t + \frac{\eta - 1}{\tilde{\sigma}\, \eta}\, \mathbb{E}_t\left[\widehat{s}_{t+1}^g - \widehat{s}_t^g\right], \tag{35}$$

where hats denote log-deviations from steady state and $\tilde{\sigma} \equiv \sigma + (1 - \omega)\, \omega^{-1}\, \eta^{-1}$ is the effective elasticity of intertemporal substitution.

Equation (35) reveals that expected declines in future public services ($\mathbb{E}_t[\widehat{s}_{t+1}^g] < \widehat{s}_t^g$) *reduce* current consumption when $\eta < 1$, because the anticipation of reduced service consumption lowers the marginal utility of current private consumption. This is the formal representation of the amplification mechanism discussed in Section 5.4.

B Appendix B: Steady-State Equations

In steady state all endogenous variables are constant. Denoting steady-state values without time subscripts, the key steady-state conditions are:

$$\beta(1 + r) = 1 + \mu(z, a) \quad \forall\, (z, a), \tag{36}$$

$$w = \frac{\varepsilon - 1}{\varepsilon}\, A, \tag{37}$$

$$G^{slg} = T + \Gamma, \tag{38}$$

$$B^f = \bar{b}^f \cdot Y, \tag{39}$$

$$Y = \int c(z, a)\, d\mu + G^{slg}. \tag{40}$$

The stationary distribution μ over (z, a) is found by iterating the Kolmogorov forward equation:

$$\mu'(z', a') = \sum_z \Pi(z'|z) \int \mathbf{1}[a_z(a) = a'] \, d\mu(z, a). \tag{41}$$

C Appendix C: Data Sources

- **SLG expenditure and revenues:** U.S. Census Bureau, *Annual Survey of State and Local Government Finances* (ASSLGF) and *Quarterly Summary of State and Local Government Tax Revenue* (QSTAX).
- **Intergovernmental transfers:** U.S. Census Bureau ASSLGF; Office of Management and Budget (OMB) Catalog of Federal Domestic Assistance (CFDA).
- **SLG employment:** Bureau of Labor Statistics (BLS), Quarterly Census of Employment and Wages (QCEW).
- **Household wealth and income:** Survey of Consumer Finances (SCF), Federal Reserve Board, 2019 wave.
- **Service utilization by income:** Consumer Expenditure Survey (CEX), BLS; Bureau of Economic Analysis (BEA) national accounts.
- **Income process:** Panel Study of Income Dynamics (PSID), used by Flodén and Lindé (2001).
- **Medicaid spending:** Centers for Medicare and Medicaid Services (CMS), National Health Expenditure Accounts (NHEA).

D Appendix D: Shocks and Scenarios

We trace the model's dynamics following four distinct structural shocks. All shocks are normalized to one standard deviation (or one basis point for the monetary shock, scaled to 25 bp for comparability with empirical VAR estimates). The shock processes are:

1. **SLG Revenue Shock** ($\varepsilon_t^{\xi} = -\sigma_{\xi} = -0.018$): a one-standard-deviation *negative* shock to the SLG revenue multiplier ξ_t, reducing SLG tax revenues below their GDP-proportional trend. Under the balanced-budget rule, this forces an immediate cut in SLG expenditure and public service provision. This is the central shock in the HANK-SLG analysis.

2. **Federal Grant Shock** ($d\Gamma_t = +1\%$ of steady-state GDP): a one-time unexpected increase in federal intergovernmental grants, relaxing the SLG budget constraint and allowing expanded service delivery. This shock isolates the *fiscal multiplier* of SLG expenditure.

3. **Monetary Policy Shock** ($\varepsilon_t^m = +25$ bp): an unexpected 25-basis-point increase in the nominal interest rate, tightening financial conditions. This shock allows comparison of the HANK-SLG model's monetary transmission with standard HANK models.

4. **TFP Shock** ($dA_t = +1\%$): a one-percent positive shock to aggregate total factor productivity, raising the marginal product of labor and shifting the supply side. This shock isolates supply-side dynamics and their interaction with the SLG balanced-budget rule.

The seven scenarios, estimated at their respective posterior-mean parameter vectors, are:

Model	Label	Grant Rule	Prices	$\underline{a}$	Line style
$\mathcal{M}_1$	Baseline (CAF)	Cyclically adjusted	Rotemberg	0	solid
$\mathcal{M}_2$	Lump-sum	Fixed	Rotemberg	0	dashed
$\mathcal{M}_3$	Open-end match	Open-ended	Rotemberg	0	dashdot
$\mathcal{M}_4$	Capped match	Capped OEM	Rotemberg	0	dense dash
$\mathcal{M}_5$	Flexible prices	CAF	Flexible	0	dotted
$\mathcal{M}_6$	High stickiness	CAF	High	0	dash dot dot
$\mathcal{M}_7$	Positive borrowing	CAF	Rotemberg	-0.5	dense dot

E Appendix E: Reading the Figures

All IRF panels show percentage deviations from the non-stochastic steady state (or percentage-point deviations for the real interest rate and inflation rate). Confidence bands shown on baseline scenario figures represent 90% posterior credible intervals, computed by evaluating the IRF at 1,000 posterior draws from the Bayesian estimation and taking the 5th and 95th percentiles.

The representative-agent (RA) benchmark—obtained by setting $\mu^* = \delta_{(\bar{z},\bar{a})}$ (a degenerate distribution at mean wealth and productivity)—is shown as a thin gray dashed line on all multi-scenario panels to illustrate HANK amplification.

Distributional IRFs (Sections 4.1.4 and 4.2.3) plot the quintile-specific consumption and welfare responses; these are computed by integrating the GE Jacobian along the individual household dimension, conditioning on each productivity-wealth quintile.

F Appendix F: Kalman Filter Recursions

Initialize: $\mathbf{s}_{1|0} = \mathbf{0}$; $\mathbf{P}_{1|0} = \mathrm{vec}[(\mathbf{I} - \mathbf{\Phi}_X \otimes \mathbf{\Phi}_X)^{-1}\mathrm{vec}(\mathbf{\Phi}_\varepsilon \mathbf{Q} \mathbf{\Phi}_\varepsilon^\top)]$.

For $t = 1, \ldots, T$:

$$\text{Forecast:} \quad \boldsymbol{\nu}_t = \mathbf{Y}_t - \mathbf{H}\mathbf{s}_{t|t-1}, \quad \mathbf{F}_t = \mathbf{H}\mathbf{P}_{t|t-1}\mathbf{H}^\top + \mathbf{R}, \tag{42}$$

$$\text{Update:} \quad \mathbf{K}_t = \mathbf{P}_{t|t-1}\mathbf{H}^\top \mathbf{F}_t^{-1}, \quad \mathbf{s}_{t|t} = \mathbf{s}_{t|t-1} + \mathbf{K}_t \boldsymbol{\nu}_t, \quad \mathbf{P}_{t|t} = (\mathbf{I} - \mathbf{K}_t\mathbf{H})\mathbf{P}_{t|t-1}, \tag{43}$$

$$\text{Predict:} \quad \mathbf{s}_{t+1|t} = \mathbf{\Phi}_X \mathbf{s}_{t|t}, \quad \mathbf{P}_{t+1|t} = \mathbf{\Phi}_X \mathbf{P}_{t|t} \mathbf{\Phi}_X^\top + \mathbf{\Phi}_\varepsilon \mathbf{Q} \mathbf{\Phi}_\varepsilon^\top. \tag{44}$$

Log-likelihood: $\mathcal{L}^{\text{agg}} = \sum_t \ell_t^{\text{agg}}$ with ℓ_t^{agg} given by (**??**).

G Appendix G: PSIS-LOO Computation

The PSIS-LOO approximation uses importance weights $w_t^{(s)} \propto 1/p(Y_t \mid \theta^{(s)}, \mathbf{Y})$, where $\theta^{(s)}$ are posterior draws. Pareto smoothing stabilizes the tails of the importance weight distribution by fitting a Generalized Pareto distribution to the upper 20% of weights. Models with Pareto tail index $\hat{k} > 0.7$ for more than 10% of observations are flagged as unreliable; none of our seven models triggers this flag.

H Appendix H: Chib–Jeliazkov Marginal Likelihood Estimator

The log marginal likelihood is computed as:

$$\ln p(\mathbf{Y} \mid \mathcal{M}_k) = \mathcal{L}(\theta^*) + \ln p(\theta^* \mid \mathcal{M}_k) - \ln p(\theta^* \mid \mathbf{Y}, \mathcal{M}_k), \tag{45}$$

evaluated at the posterior mode θ^*. The log-posterior density $\ln p(\theta^* \mid \mathbf{Y})$ is estimated using the output of the MCMC sampler via the Chib–Jeliazkov reduced-run estimator, which requires an additional set of 20,000 MCMC draws conditional on $\theta = \theta^*$.