EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Charles de Beauffort, Ansgar Rannenberg

Fiscal Policy and sectoral spillovers in open-economy HANK

No 3279

Disclaimer: This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Abstract

Government spending falls disproportionately on non-tradable services. We show empirically that government spending shocks stimulate private consumption along with sizable spillovers to the goods sector and a relative decline in goods prices. We rationalize these findings with a two-sector open-economy HANK model. Uninsurable income risk and precautionary savings lead to a persistent income-driven expansion in private consumption. In the tradable sector, import intensity and limited labor reallocation dampen wage pass-through to prices, matching observed co-movements. The resulting expenditure switching produces a positive tradable output response despite deteriorating net exports. Household heterogeneity and trade openness jointly shape sectoral fiscal multipliers.

JEL Codes: E62, F41, E21, C11, C32

Keywords: Fiscal policy, Heterogeneous agents, Open economy, Sectoral spillovers, Government spending, Consumption, Trade, SVAR

Non-technical summary

Why does government spending on services also boost the goods sector?

When governments spend money, they mostly buy services like healthcare, education, or public administration. Yet this kind of spending seems to stimulate the production and consumption of goods as well. This paper asks why, and whether standard economic models can explain it.

What data shows

Looking at US data from 1954 to just before COVID-19, the authors find four consistent patterns following a government spending increase: output rises in both the services and goods sectors; private households consume more of both, with the increase in goods consumption being especially pronounced; goods become relatively cheaper compared to services; and the trade balance deteriorates.

The puzzle is that standard models, where a single "representative" household rationally plans ahead, cannot reproduce these patterns. In those models, households respond to higher anticipated real interest rates and future taxes by cutting private consumption, so the positive spillovers to the goods sector simply do not emerge.

The key ingredients

The paper builds a richer model with three departures from the standard framework. First, households are heterogeneous. Many face borrowing constraints and imperfect insurance against income fluctuations. When government spending raises wages and employment, constrained households spend a significant share of the additional income immediately rather than saving it for the future. Precautionary motives mean this elevated spending persists for years, not just a few quarters.

Second, workers do not move freely between sectors. When government spending concentrates in services, labor market pressure is higher there, pushing up services wages and prices more than in the goods sector. Goods therefore become relatively cheaper, encouraging households to substitute toward them, which explains the disproportionate rise in goods consumption.

Third, the model accounts for international trade realistically. Import prices partly reflect domestic distribution and retail costs, so a rising exchange rate does not fully erode competitiveness. This limits how much the goods sector loses to foreign competition, allowing domestic goods output to rise despite a deterioration of the trade balance.

What the model delivers

Combining these ingredients, the model successfully replicates all four empirical patterns. The representative-agent version of the same model fails on every count. This confirms that household heterogeneity is not just a detail; it is central to understanding how fiscal policy is transmitted across sectors.

What this means for policy

Governments should not assume that spending concentrated in services remains confined there. The income gains extend to goods demand as well, but how much depends on the economy's openness. In a more open economy, a larger share of that demand leaks abroad through imports and an appreciated exchange rate, reducing the domestic output gains. Policies that support the competitiveness of domestic firms help ensure that the income generated by fiscal expansions stays within the home economy.

1 Introduction

Government spending is a key macroeconomic stabilization tool and has been crucial to mitigate the adverse effects of both the Global Financial Crisis and the COVID-19 shock. In this paper, we extend existing evidence on its sectoral effects in an open economy and study the underlying transmission mechanisms through the lens of a New Keynesian model with rich household heterogeneity.

Using a Structural Bayesian Vector Auto-Regressive (S-BVAR) model estimated on US data, we identify the macroeconomic effects of an orthogonal increase in government spending. We document four main stylized facts: (i) a sizable increase in goods production, despite government spending being largely concentrated in services; (ii) a joint rise in private consumption of goods and services, with a substantially larger response in goods; (iii) a relative decline in the price of consumption goods, and; (iv) a deterioration of the home economy's net exports.

We develop an open economy Heterogeneous-Agent New-Keynesian (HANK) framework to account for the observed sectoral co-movements. The model features two production sectors—tradable and non-tradable—capturing goods and services in the data. Labor is imperfectly mobile across sectors and government spending is concentrated in the non-tradable sector, in line with empirical evidence. A central mechanism is that deficit-financed fiscal expansion increases the expected disposable income of heterogeneous households by raising labor demand with positive effects on wages and hours worked. Since many households face borrowing constraints and incomplete insurance against income fluctuations, a sizable share of this additional income is consumed on impact, while precautionary motives imply that the consumption response persists beyond the initial shock.

Imperfect labor mobility causes wages to rise more in the non-tradable sector, where government spending is concentrated, leading to higher inflation there. By contrast, tradable prices increase less, owing also to their foreign component. The resulting decline in the relative price of tradable goods matches the empirical evidence and triggers expenditure switching towards cheaper tradable goods. The flip side of the strong demand response is an appreciation of the real exchange rate, which dampens the competitiveness of domestic goods both at home and abroad. A distribution sector partially offsets these competitiveness losses, as the non-tradable cost component embedded in import prices rises alongside domestic inflation, cushioning the deterioration of net exports. We show that for a relatively closed economy like the US, the rise in domestic tradable consumption more than offsets the fall in exports, so that output rises sensibly in the tradable sector. The HANK model thus succeeds in replicating all the stylized facts described above.

We compare these results with those of an otherwise identical model with a representative agent. As shown in previous work, this framework generates several puzzles regarding sectoral co-movements of quantities and prices (see e.g. Monacelli and Perotti, 2008). Consistently with those results, our model predicts that consumption declines in both sectors because households anticipate higher real interest rates and higher future taxes. As a result, the sectoral composition of the spending shock becomes the unique driver of the positive output response. When government spending is fully allocated to the non-tradable sector, output in the tradable sector contracts, generating counterfactual negative co-movements between sectoral outputs.

This paper provides new evidence on the role of household heterogeneity in shaping the sectoral transmission of fiscal shocks in an open economy. It shows that a fiscal shock tilted toward services can generate substantial spillovers to the goods sector through the private consumption response to higher income. The magnitude of these spillovers depends on the international repercussions of the associated increase in domestic prices and the real exchange rate. Consequently, a more open and less specialized economy is likely to experience smaller output gains in its tradable sectors. From a policy perspective, this suggests that governments should not only consider which domestic industries receive fiscal support, but also adopt measures that strengthen the competitive and innovative edge of domestic firms to ensure that as much of the additional income as possible is spent within the home economy.

Our paper relates to the literature on the transmission of fiscal policy in representative-agent models. Monacelli and Perotti (2010) and Ravn et al. (2012), among others, show that the standard framework cannot account for the positive aggregate consumption response observed in the data. Monacelli and Perotti (2008) further study the sectoral dimension of this puzzle and document several shortcomings of the representative-agent model. First, consumption declines in both sectors: sectoral co-movements are therefore positive, as in the data, but of the wrong sign in each sector. Second, while trade openness can rationalize the observed price co-movements, the model struggles to replicate the observed output co-movements—especially when government spending is concentrated in services, in which case sectoral output co-movements are unambiguously negative. Relative to their empirical approach, we incorporate net exports as an additional indicator of the trade effects of fiscal shocks. We also extend the sample up to just before the COVID-19 pandemic, thereby encompassing the Global Financial Crisis, and estimate the model using Bayesian methods, which provide greater flexibility in handling a large parameter space. Building on this analysis, we then develop a full HANK model that resolves the puzzles uncovered from their representative-agent framework.

Beyond this literature, our paper connects to previous work on fiscal policy in heterogeneous-agent models. Galí et al. (2007) provide a resolution to the consumption puzzle in a closed economy using reduced-form household heterogeneity. Their Two-Agent New Keynesian (TANK) model, combined with sticky prices, generates a positive consumption response to additional income from government spending. Erceg et al. (2005) obtain similar results in an open-economy setting, but favor specifications with a small share of liquidity-constrained households to match evidence of limited effects of government spending on the trade balance. Compared with those models, we study an open-economy model with sectoral spillovers under full household heterogeneity. This richer structure allows for a more granular characterization of the transmission mechanisms and delivers a substantially tighter quantitative fit, simultaneously matching the observed responses of consumption and trade dynamics. The household block of our model builds on Auclert et al. (2024), who show that the intertemporal New Keynesian cross mechanism delivers larger output multipliers of government spending than those found in representative-agent models. More recently, Druedahl et al. (2024) extend this framework to a single-sector open economy and demonstrate that the trade channel offsets part of these effects, thereby dampening the multiplier. In a related contribution, Aggarwal et al. (2023) study the international transmission of fiscal transfers in a multi-country HANK model and show that high household intertemporal marginal propensities to consume out of income amplify the aggregate demand response and generate persistent current account deficits. While our model shares these general amplification mechanisms with both papers, the key distinction is our two-sector structure with imperfect labor mobility, which generates endogenous relative price dynamics between tradables and non-tradables. This allows us to address sectoral co-movements and terms of trade responses that are outside the scope of both frameworks. All of these forces operate simultaneously in our two-sector environment and help reconcile theoretical predictions with empirical evidence.

Our extension incorporating imperfect labor mobility is inspired by Cardi et al. (2020), who provide empirical evidence on the effects of fiscal shocks on sectoral labor shares. They find that government consumption is more intensive in non-tradable goods and that labor reallocation across sectors is associated with a relative increase in non-tradable wages which is consistent with costly labor adjustments.

Finally, the identification strategy in our S-BVAR follows the seminal contribution of Blanchard and Perotti (2002). Many subsequent studies have adopted variants of this approach, including Monacelli and Perotti (2008), Monacelli and Perotti (2010), Ravn et al. (2012), and Cardi et al. (2020).

The remainder of the paper is organized as follows. Section 2 describes the S-BVAR model

and highlights the main empirical insights that we draw from it. Section 3 develops the open economy New-Keynesian model with heterogeneous households. Section 4 contains the main analysis. After calibrating the model, Section 4.2 provides the baseline simulation results and compares them with the empirical evidence to give a structural interpretation of the observed dynamics. Section 4.3 investigates the channels at work by testing the sensitivity of the quantitative results to key parameters regarding household expectations, sectoral allocation of consumption and international trade. Section 5 concludes.

2 The empirical model

We estimate the orthogonal effect of a government spending shock using the seminal Blanchard and Perotti (2002) identification strategy within a Structural Bayesian Vector Auto-Regressive model (S-BVAR). This strategy relies on a Cholesky decomposition in which government spending is ordered first.[1] Our sample comprises quarterly data from 1954Q1 to 2019Q4 that we obtain for nine variables from the US Bureau of Economic Analysis.[2] The S-BVAR model is formally defined as

$$A \begin{bmatrix} \hat{g}_t \\ \hat{y}_t \\ \hat{y}_{T,t} \\ \hat{y}_{NT,t} \\ \hat{c}_t \\ \hat{c}_{T,t} \\ \hat{c}_{NT,t} \\ \widehat{nxy}_t \\ \hat{p}_{S,t} \end{bmatrix} = B(L) \begin{bmatrix} \hat{g}_{t-1} \\ \hat{y}_{t-1} \\ \hat{y}_{T,t-1} \\ \hat{y}_{NT,t-1} \\ \hat{c}_{t-1} \\ \hat{c}_{T,t-1} \\ \hat{c}_{NT,t-1} \\ \widehat{nxy}_{t-1} \\ \hat{p}_{S,t-1} \end{bmatrix} + \epsilon_t$$

[1] Another strand of the literature uses a narrative approach—typically based on military spending episodes—to identify fiscal shocks (e.g., Ramey and Shapiro, 1998; Ramey, 2011). Two considerations make the narrative approach unsuitable here. First, our analysis focuses on government consumption shocks biased toward non-tradable sectors, which are predominantly non-military in nature; military spending, by contrast, is relatively tradable-intensive, making Ramey-type series a poor proxy for the shocks we study. Second, as noted by Ravn et al. (2007), narrative series based on military spending announcements capture anticipation effects of future expenditures, whereas the HANK model we use as a theoretical benchmark features fully unanticipated shocks—rendering a direct comparison with estimates that embed anticipation effects conceptually inconsistent. An alternative would be to purge anticipation effects using analysts' forecast revisions (e.g., Auerbach and Gorodnichenko, 2012), but even under the most favorable data availability—relying on Greenbook forecasts for government spending, which run from 1966 to 2004—our sample would be truncated at both ends relative to our current 1954–2019 baseline, at the cost of losing key low-frequency variation. Crucially, even if some degree of fiscal foresight were present in our baseline estimates, Mertens and Ravn (2011) show, using a SVAR estimator explicitly designed to accommodate anticipated fiscal shocks, that anticipation effects do not overturn the standard findings from the fiscal SVAR literature. Our results are therefore unlikely to be materially affected by this concern.

[2] The quarterly frequency is important because a lack of discretionary fiscal response in the same period would be much less plausible at lower frequency. We start in 1954 to exclude the Korean War and stop before the COVID-19 period to avoid dealing with associated structural breaks and unusually large government spending packages. See Galí et al. (2007) for a discussion about the effects of including the Korean war on results from a structural VAR.

where $\hat{g}_t$ denotes federal government consumption expenditures, $\hat{y}_t$ denotes real total GDP, $\hat{y}_{T,t}$ denotes real GDP of the goods producing sector, $\hat{y}_{NT,t}$ denotes real GDP for services, $\hat{c}_t$ denotes total private consumption, $\hat{c}_{T,t}$ denotes private consumption of goods, $\hat{c}_{NT,t}$ denotes private consumption of services, $\widehat{nxy}_t$ denotes the net export-to-GDP ratio and $\hat{p}_{S,t}$ denotes the goods over services price ratio (deflators in the data). Hatted variables stand for log-deviation from trend, except for $\widehat{nxy}_t$ which is in absolute deviations. All variables are seasonally adjusted by the data provider, and detrended with a linear and quadratic trend. The vector of disturbances ϵ_t is mean-zero and serially uncorrelated. In estimation, we include dummy variables to account for the oil crisis in 1975-Q2 and the Global Financial crisis from 2008-Q2 to 2009-Q2. We include four lags so that $B(L) \equiv B_0 + B_1 L + B_2 L^2 + B_3 L^3$.

As for inference, we adopt a Bayesian approach to deal more efficiently with the increase in variance stemming from the large number of parameters involved with nine variables and four lags. We apply a Normal-Inverse Wishart prior that shrinks coefficients toward zero at a rate that exponentially increases with more distant lags. This specification favors shorter lag structures and smooth impulse responses. Since all our variables are stationary, we set to 0.8 the mean of the autoregressive coefficient corresponding to the first lag.

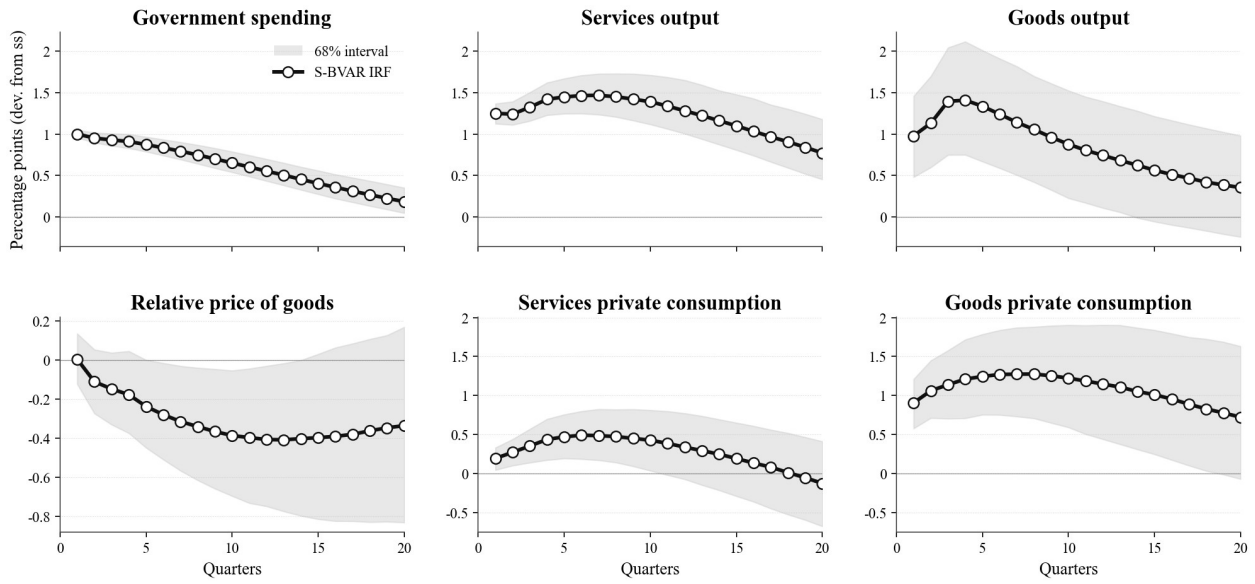


Figure 1: **Empirical impulse responses to an orthogonal increase in government spending**
Notes: Impulse responses to an orthogonal 1% increase in government spending as a share of output identified by the S-BVAR. Solid line with circles is the estimated response and grey area corresponds to 68% confidence bands.

Some key stylized facts from the empirical model appear in Figure 1. First, the output response on impact is close to 1% in both sectors and is hump-shaped for subsequent periods. Second,

consumption also increases in both sectors. This increase is much larger for tradable goods. Third, tradable goods become relatively cheaper when compared to non-tradable goods. A fourth stylized fact not shown in the figure (but investigated below) is a fall in the home economy's net exports.[3] In the next section, we present a HANK model which can shed light on the transmission channels behind these empirical responses.

3 The HANK model

Our framework is a Heterogeneous-Agent New Keynesian (HANK) model with international trade. Time is discrete and the horizon is infinite. There are two production sectors, tradable and non-tradable, respectively denoted by $S \in \{T, NT\}$. Tradable goods can be produced domestically or abroad.

We consider perfect-foresight impulse responses to a government spending shock starting from a steady state without aggregate uncertainty ("MIT shocks"). We use the sequence-space Jacobian method of Auclert et al. (2021) and linearize with respect to these shocks.[4] Government consumption is composed of imperfectly substitutable non-tradable and tradable consumption.

3.1 Private sector

Heterogeneous households We borrow from Auclert et al. (2024) to model household heterogeneity. The economy is populated by a unit mass of households indexed with i facing idiosyncratic income uncertainty captured by time varying labor productivity $e_t(i)$.

The sectoral effective labor supply of household i is given by $e_t(i)n_{S,t}(i)$, with $n_{S,t}(i)$ denoting household i's labor effort so that aggregate sectoral hours are given by $N_{S,t} = \int e_t(i)n_{S,t}(i)di$. Heterogeneous households delegate their labor decisions $n_{S,t}(i)$ to sectoral labor unions that sluggishly set the nominal wage per unit of effective labor $W_{S,t}$. Labor hours are allocated proportionally across households via the rule

$$n_{S,t}(i) = N_{S,t} \tag{1}$$

[3] In Appendix A, we show that those stylized facts are robust across different specifications of government spending in the BVAR (excluding military spending and excluding gross capital formation) and to a reduced sample period excluding the GFC. The latter sample is also directly comparable to Mertens and Ravn (2014); as shown in Figure A.3, our identified shock series correlates at 0.76 with theirs, lending further confidence to the consistency of our identification across specifications.

[4] Bianchi and Kaplan (2026) show that linearization can overstate consumption responses to shocks with positive income effects. This concern is mitigated in our setting for two reasons. First, we normalize the shock to a one-percent increase in the government spending share of GDP. For this magnitude, the authors' own quantitative comparisons show modest differences between linear and nonlinear solutions. Second, unlike the transfer experiment in Bianchi and Kaplan (2026), government spending raises household income only indirectly, through general equilibrium effects on wages and employment, limiting the scope for linearization bias.

which implies normalization $\int e_t(i)di = 1$. The productivity shock follows an AR(1) process with persistence ρ_e and standard deviation σ_e.

Taxes are assumed to be levied proportionally to household i's productivity. Post-tax labor income is therefore given by $z_t(i) = (w_{T,t}n_{T,t}(i) + w_{NT,t}n_{NT,t}(i) - T_t)e_t(i)$ where $w_{S,t}$ is the real wage and T_t denotes aggregate lump-sum tax collections. The term T_t should be understood as net taxes, i.e. tax payments net of transfer receipts. This specification implies that the pre- and post tax distributions of labor income are identical. Consequently, variations in total tax revenue are neutral with respect to the income distribution.

The household budget constraint can then be written as

$$a_t(i) + c_t(i) = (1 + r_t)a_{t-1}(i) + z_t(i) \tag{2}$$

where $a_t(i)$ denotes real asset holdings and r_t the ex-post real interest rate. In addition, households face the borrowing constraint

$$a_t(i) \geq 0 \tag{3}$$

Each household maximizes

$$\sum_{t=0}^{\infty} \beta^t \left(\frac{c_t(i)^{1-\sigma}}{1-\sigma} - \chi \frac{n_t(i)^{1+\varphi}}{1+\varphi} \right) \tag{4}$$

subject to (2) and (3).

Households consume a composite of tradable and non-tradable goods. The respective demand schedules derived from the CES functions read

$$C_{T,t} = \phi_T P_{T,t}^{-\lambda_T} C_t \tag{5}$$

$$C_{NT,t} = (1 - \phi_T) P_{NT,t}^{-\lambda_T} C_t \tag{6}$$

where $P_t = 1$ since all prices are expressed in consumption units of the final good.

The tradable good, in turn, is a composite of a domestic tradable good and a foreign tradable good with respective demand schedules

$$C_{D,t} = \phi_D \left(\frac{P_{D,t}}{P_{T,t}} \right)^{-\lambda_D} C_{T,t} \tag{7}$$

$$C_{F,t} = (1 - \phi_D) \left(\frac{P_{F,t}}{P_{T,t}} \right)^{-\lambda_D} C_{T,t} \tag{8}$$

Wage setting Previous studies have shown that imperfect re-allocation of labor across sectors is an important friction to explain the sectoral transmission of shocks in open economy models (see e.g. Cardi and Restout, 2015; Cardi et al., 2020; Bertinelli et al., 2022). Following their approach, we assume that aggregate labor is a CES basket of the sectoral labor types

$$N_t = \left[\phi_H^{-\frac{1}{\lambda_H}} (N_{NT,t})^{\frac{1+\lambda_H}{\lambda_H}} + (1-\phi_H)^{-\frac{1}{\lambda_H}} (N_{T,t})^{\frac{1+\lambda_H}{\lambda_H}} \right]^{\frac{\lambda_H}{1+\lambda_H}} \tag{9}$$

The sectoral labor input $N_{S,t}$ employed by firms, in turn, is a CES basket of labor varieties. Hence, the demand for variety j is given by

$$N_{S,t}(j) = \left(\frac{W_{S,t}(j)}{W_{S,t}} \right)^{-\theta_w} N_{S,t} \tag{10}$$

We model wage stickiness as in Hagedorn et al. (2019); Auclert et al. (2023) and assume that nominal wages paid for individual varieties $W_{S,t}(j)$ are set by a labor union that maximizes the utility of an average household subject to quadratic adjustment costs. This approach has the advantage of remaining neutral with respect to the household block of the model. We introduce partial indexation to past consumer price inflation into the wage adjustment cost. The cost of adjusting the wage of variety j in sector S becomes

$$\frac{\psi_w}{2} \left(\frac{W_{S,t}(j)}{W_{S,t-1}(j)\, \pi_{t-1}^{\iota_w}} - 1 \right)^2 N_{S,t}(j)\, W_{S,t}(j), \qquad \iota_w \in [0,1] \tag{11}$$

where $\pi_{t-1} \equiv P_{t-1}/P_{t-2}$ is gross CPI inflation and ι_w governs the degree of indexation. The case $\iota_w = 0$ recovers the baseline, while $\iota_w = 1$ indexes fully to past inflation. The union's objective reads

$$\begin{aligned} \sum_{t=0}^{\infty} \beta^t \Bigg(& -\chi \frac{\left\{ \left[\phi_H^{-\frac{1}{\lambda_H}} \left(\left[\frac{W_{NT,t}(j)}{W_{NT,t}} \right]^{-\theta_w} N_{NT,t} \right)^{\frac{1+\lambda_H}{\lambda_H}} + (1-\phi_H)^{-\frac{1}{\lambda_H}} \left(\left[\frac{W_{T,t}(j)}{W_{T,t}} \right]^{-\theta_w} N_{T,t} \right)^{\frac{1+\lambda_H}{\lambda_H}} \right]^{\frac{1}{1+\lambda_H}} \right\}^{1+\varphi}}{1+\varphi} \\ & + C_t^{-\sigma} \Bigg\{ W_{NT,t}(j) \left(\left[\frac{W_{NT,t}(j)}{W_{NT,t}} \right]^{-\theta_w} N_{NT,t} \right) \left(1 - \frac{\psi_w}{2} \left(\frac{W_{NT,t}(j)}{W_{NT,t-1}(j)\, \pi_{t-1}^{\iota_w}} - 1 \right)^2 \right) \\ & + W_{T,t}(j) \left(\left[\frac{W_{T,t}(j)}{W_{T,t}} \right]^{-\theta_w} N_{T,t} \right) \left(1 - \frac{\psi_w}{2} \left(\frac{W_{T,t}(j)}{W_{T,t-1}(j)\, \pi_{t-1}^{\iota_w}} - 1 \right)^2 \right) \Bigg\} \Bigg) \end{aligned} \tag{12}$$

Define the indexed wage inflation gap as

$$\tilde{\pi}_{w,S,t} \equiv \frac{\pi_{w,S,t}}{\pi_{t-1}^{\iota_w}} - 1, \qquad \pi_{w,S,t} \equiv \frac{W_{S,t}}{W_{S,t-1}} \tag{13}$$

After imposing a symmetric equilibrium, the first order conditions of this program yield the sectoral non-linear wage Phillips curves. For the non-tradable sector, it is given by

$$\begin{aligned}\frac{MRS_{NT,t}}{W_{NT,t}} = \frac{\theta_w - 1}{\theta_w}\left(1 - \frac{\psi_w}{2}\tilde{\pi}_{w,NT,t}^2\right) \\ + \frac{\psi_w}{\theta_w}\tilde{\pi}_{w,NT,t}\pi_{w,NT,t} - \beta\mathbb{E}_t\frac{\psi_w}{\theta_w}\tilde{\pi}_{w,NT,t+1}\pi_{w,NT,t+1}^2\frac{C_{t+1}^{-\sigma}}{C_t^{-\sigma}}\frac{N_{NT,t+1}}{N_{NT,t}}\end{aligned} \tag{14}$$

where

$$MRS_{NT,t} = \chi N_t^{\varphi - \frac{1}{\lambda_H}} \phi_H^{-\frac{1}{\lambda_H}} N_{NT,t}^{\frac{1}{\lambda_H}} C_t^{\sigma}$$

The wage Phillips curve for the tradable sector is symmetric, only now:

$$MRS_{T,t} = \chi N_t^{\varphi - \frac{1}{\lambda_H}} (1 - \phi_H)^{-\frac{1}{\lambda_H}} N_{T,t}^{\frac{1}{\lambda_H}} C_t^{\sigma}$$

As in Hagedorn et al. (2019), the wage adjustment process does not involve actual costs, but optimal choices are made as-if those costs were actually present. This assumption prevents wage adjustment costs from becoming a non-trivial fraction of output following a government spending shock.

Household expectations We assume that households hold sticky expectations as in Mankiw and Reis (2002); Auclert et al. (2020). In contrast to the rational expectations paradigm, in which agents perfectly foresee the entire path of a shock, households update their information sets only sluggishly and gradually incorporate new information into their decisions.

$$\mathbf{E} = \begin{pmatrix} 1 & 1-\vartheta & 1-\vartheta & \cdots \\ 1 & 1 & 1-\vartheta^2 & \cdots \\ 1 & 1 & 1 & \ddots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \tag{15}$$

Formally, let $\vartheta \in (0,1)$ denote the fraction of households that do not update their expectations in a given period. Expectations then evolve according to the updating matrix (15).[5] The first row corresponds to expectations in period 1. If the shock occurs after this period, households learn

[5] In the context of the solution method of our model, we recover the modified Jacobian with sticky information from

about it with probability $1-\vartheta$. The second row represents expectations in period 2. In this case, if the shock occurs after period 2, households have had two opportunities to update, implying that a fraction $1-\vartheta^2$ has incorporated the shock. This logic extends recursively to subsequent periods. As a result, shocks that materialize sufficiently far in the future are effectively fully internalized by households.

Production Firms operate under perfect competition and set prices flexibly to maximize profits subject to linear technology $Y_{S,t}(k) = N_{S,t}(k)$.[6] So, in equilibrium, they equalize their marginal costs and marginal profits

$$P_{Z,t} = W_{S,t} \tag{16}$$

where $P_{Z,t} \in \{P_{NT,t}, P_{D,t}\}$ is the domestic sectoral price.[7]

In the tradable sector, firms apply producer currency pricing so that the price of imports is

$$P_{M,t} = P^*_{DT,t} Q_t \tag{17}$$

where the star superscript denotes the foreign economy and Q_t is the real exchange rate defined as how many units of domestic goods can be exchanged for one unit of foreign good (an increase in Q_t thus represents a depreciation of the exchange rate).

3.2 Public sector

There are two public authorities, a government and a central bank. The government levies lump-sum taxes to stabilize debt following government spending shocks and the central bank sets the short-term nominal interest rate to stabilize inflation.

the full-information Jacobian using the following mapping

$$\mathcal{J} = \begin{pmatrix} \mathcal{J}_{00} & (1-\vartheta)\mathcal{J}_{01} & (1-\vartheta)\mathcal{J}_{02} & \cdots \\ \mathcal{J}_{10} & (1-\vartheta)\mathcal{J}_{11} + \vartheta\mathcal{J}_{00} & (1-\vartheta)\mathcal{J}_{12} + \vartheta(1-\vartheta)\mathcal{J}_{01} & \cdots \\ \mathcal{J}_{20} & (1-\vartheta)\mathcal{J}_{21} + \vartheta\mathcal{J}_{10} & \cdots & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

where $\mathcal{J}s\tau$ denotes the Jacobian for the group of households that learns about the shock to input sequence s at date τ. This procedure allows to build the modified Jacobian with sticky information at almost no extra computational costs (see Section 4.3 of Auclert et al. (2020) for more details).

[6] Assuming only sticky wages is standard in the HANK literature (Auclert et al., 2023, 2024) and is sufficient to introduce nominal rigidities while avoiding countercyclical firm profits. Nevertheless, we assess the robustness of our results to the introduction of sticky prices in Appendix C.

[7] Here, we have to make a distinction between S and Z because the domestic tradable price does not include contribution of foreign tradable prices which only enters in the total tradable price $P_{T,t}$.

Fiscal policy Government spending is exogenous $G_t = G_t^{exo}$. The shock G_t^{exo} follows the path of its empirical counterpart from the S-BVAR model.[8] As for the private sector, public spending is allocated to production sectors according to a CES function. The demand schedules are

$$G_{T,t} = \phi_{GT} P_{T,t}^{-\lambda_T} G_t \tag{18}$$

$$G_{NT,t} = (1 - \phi_{GT}) P_{NT,t}^{-\lambda_T} G_t \tag{19}$$

where the tradable public good is a composite of a domestic and a foreign good with respective demand schedules

$$G_{D,t} = \phi_D^G \left(\frac{P_{D,t}}{P_{T,t}} \right)^{-\lambda_D} G_{T,t} \tag{20}$$

$$G_{F,t} = (1 - \phi_D^G) \left(\frac{P_{F,t}}{P_{T,t}} \right)^{-\lambda_D} G_{T,t} \tag{21}$$

The government budget constraint reads

$$B_t = (1 + r_{t-1}) B_{t-1} + P_{G,t} G_t - T_t \tag{22}$$

Lump-sum taxes respond to inherited debt according to the rule

$$T_t = T^{SS} + \kappa_b (B_{t-1} - B^{SS}) \tag{23}$$

where κ_b captures a very gradual return to steady state debt after a deficit-financed spending shock.

Monetary policy The nominal interest rate is set according to a simple Taylor rule

$$i_t = \zeta i_{t-1} + (1 - \zeta) \left(i^{SS} + \kappa_\pi (\pi_t - 1) + \kappa_y \frac{\Delta Y_t}{Y_{t-1}} \right) \tag{24}$$

Parameters κ_π and κ_y respectively govern the reaction of the central bank to inflation and output growth. Parameter ζ captures its interest rate smoothing motive.

[8] This experiment captures the average transmission of an exogenous government spending shock. State-dependent effects — such as those emerging in permanent-income or credit-tightening crises — are beyond the scope of our analysis. We refer the reader to Guntin et al. (2023) for a treatment of consumption multipliers in such environments.

3.3 Financial sector

As in Auclert et al. (2023), a mutual fund issues claims to households with aggregate value A_t and invests in financial assets at home and abroad. Profit maximization implies that the expected returns on those two assets are equal. The UIP is then given by

$$(1 + i_t) = (1 + i^*)\frac{\mathcal{E}_{t+1}}{\mathcal{E}_t}\Gamma_t \tag{25}$$

where $\mathcal{E}_t$ is the nominal exchange rate and $\Gamma_t = \exp(-10^{-9} nfa_t)$ creates a wedge between domestic and foreign rates of return on bonds. In the RANK model, this wedge or an equivalent assumption is necessary in order to stabilize net foreign assets in the long run (Schmitt-Grohé and Uribe (2003)). By contrast, in the HANK model the precautionary savings motive stabilizes net foreign assets, but we keep the wedge in the HANK model for better comparability of results with the RANK model. Define the ex-ante real interest rate and the real exchange rate as

$$(1 + r_t) \equiv (1 + i_t)\frac{P_t}{P_{t+1}} \tag{26}$$

$$Q_t \equiv \frac{P_t^* \mathcal{E}_t}{P_t} \tag{27}$$

Then, (25) can be written as

$$(1 + r_t) = (1 + i^*)\frac{Q_{t+1}}{Q_t}\Gamma_t \tag{28}$$

The net foreign asset position evolves according to

$$nfa_t = (1 + i^*)\frac{Q_t}{Q_{t-1}} nfa_{t-1} + NX_t \tag{29}$$

where $NX_t = P_{D,t}X_t - P^*_{D,t}Q_t M_t$ are the net exports.

3.4 International trade

The tradable sector is subject to distribution costs, following Corsetti and Dedola (2005). Importers at home and abroad combine local distribution services with imports to produce the distributed foreign consumption good using the following Leontief technology:

$$Y_{F,t} = \min\left\{\frac{1 + \mu_F}{\mu_F} Y^d_{F,t}, (1 + \mu_F) M_t\right\} \tag{30}$$

Similarly, final producers of the distributed domestic tradable good use technology:

$$Y_{D,t} = \min\left\{\frac{1+\mu_D}{\mu_D}Y^d_{D,t}, (1+\mu_D)Y_{T,t}\right\} \tag{31}$$

In equilibrium, $Y_{F,t} = C_{F,t}$ and $Y_{D,t} = C_{D,t}$. The price indices then become:

$$P_{F,t} = \frac{\mu_F}{1+\mu_F}P_{NT,t} + \frac{1}{1+\mu_F}P_{M,t} \tag{32}$$

$$P_{D,t} = \frac{\mu_D}{1+\mu_D}P_{NT,t} + \frac{1}{1+\mu_D}P_{DT,t} \tag{33}$$

Thus, distribution costs generate lower pass-through of real exchange rate movements to the price of imported goods, as the total tradable price now incorporates a non-tradable component. Exports and imports are respectively given by:

$$X_t = \frac{(1-\phi^*_D)(P^*_{F,t})^{-\lambda^*_D}C^*_t}{1+\mu^*_F} \tag{34}$$

$$M_t = \frac{C_{F,t}+G_{F,t}}{1+\mu_F} \tag{35}$$

Real tradable and non-tradable output are:

$$Y_{NT,t} = C_{NT,t} + G_{NT,t} + \frac{\mu_D}{1+\mu_D}C_{D,t} + \frac{\mu_F}{1+\mu_F}C_{F,t} \tag{36}$$

$$Y_{T,t} = \frac{1}{1+\mu_D}(C_{D,t}+G_{D,t}) + X_t \tag{37}$$

3.5 Market clearing

The asset and goods market clearing conditions read

$$A_t = B_t + nfa_t \tag{38}$$

$$Y_t = C_t + P_{G,t}G_t + NX_t \tag{39}$$

4 Sectoral transmission of fiscal shocks

We begin by examining whether the impulse responses to the government shock in our HANK model are consistent with the predictions from the empirical S-BVAR model.

4.1 Calibration

The model is calibrated on the US economy before the COVID-19 period. A time period represents one quarter of a year. We follow the calibration of Auclert et al. (2024) for the heterogeneous household block. The intertemporal elasticity of consumption and the inverse of the Frisch elasticity are equal to 1. The original values of the AR(1) coefficient and the innovation standard deviation of the idiosyncratic income process in Auclert et al. (2024) are for annual data. We find our quarterly values by simulating quarterly data for a given pair of values, aggregating to annual, and then estimating an AR(1) on the annual data. We choose the quarterly persistence and standard deviation of innovation to minimize $(\hat{\rho}_a-\rho_a)^2+(\hat{\sigma}_a-\sigma_a)^2$, where $\hat{\rho}_a/\hat{\sigma}_a$ are the persistence/standard deviation estimated from the annual data and ρ_a and σ_a are the values of Auclert et al. (2024). We choose the household discount factor β consistent with a steady-state real interest rate equals 2%, and, following Auclert et al. (2024), set steady-state value of government bonds such that the first year marginal propensity to consume out of a transitory income shock equals the estimate of Fagereng et al. (2021).[9]

Table 1: Key moments: annual MPCs and steady-state wealth distribution

Intertemporal MPCs		*Wealth shares*	
Year	MPC	Quintile	Wealth share
Year 1	0.5067	Q1 (bottom 20%)	0.0000
Year 2	0.1375	Q2	0.0000
Year 3	0.0955	Q3	0.0047
Year 4	0.0698	Q4	0.0861
Year 5	0.0526	Q5 (top 20%)	0.9092

Notes: Annual MPCs are sums of quarterly iMPCs from the income Jacobian $\mathcal{J}^{C,Z}$, in response to a period-0 unit income shock. Wealth shares are computed from the steady-state asset distribution, with households sorted into quintiles by asset holdings.

Table 1 shows the key moments from our model at steady state.[10] For the representative agent benchmark, β targets a real interest rate of 2% annually. The sticky expectations parameter ϑ is 0.935 as estimated by Auclert et al. (2020) and captures the very gradual movements in consumption observed in the data.

The parameters of the wage Phillips curves follow Lindé et al. (2016) which update estimates

[9] As in Auclert et al. (2024), this approach results in a government debt-to-GDP-to-annual-ratio of substantially below empirical values, namely 19%. A model with liquid and an illiquid asset as in Auclert et al. (2024) might be able to match a more realistic level of liquid assets.

[10] To gauge the fit of the model with evidence from Fagereng et al. (2021), one has to compare the annual MPCs in the table with the first panel of Figure 2 in the reference paper. As is well known from the HANK literature, the one-asset model overstates wealth concentration at the top when calibrated to match those iMPCs. While this is inconsequential for our aggregate dynamics results, we note that it would be possible to match both moments simultaneously by extending the model to include a second illiquid asset, as in Kaplan et al. (2018), while keeping the same iMPCs.

from Smets and Wouters (2007) by extending the data set up to the COVID-19 period. The elasticity of substitution between labor services θ_w corresponds to a wage markup over the marginal cost of 50%. Given the value of θ_w, the parameter of the wage adjustment cost, ψ_w, is set to target a wage Phillips curve slope of 0.0044 (in both sectors). We assume that wages are not indexed in the baseline ($\iota_w = 0$) and turn indexation on when studying the response of inflation and monetary policy in Section 4.3.

Table 2: Baseline calibration

Symbol	Description	Value	Source/target
Households and firms			
β	Subjective discount factor (RA - HA)	0.995 - 0.986	Annual rate of 2%
θ_w	Elasticity of substitution labor services	3	Market power of 50%
$\frac{\theta_w - 1}{\psi_w}$	Wage PC slope	0.0044	Lindé et al. (2016)
ι_w	Wage indexation	0	Non-indexed wages
σ	Intertemporal elasticity of C	1	Auclert et al. (2024)
φ	Intertemporal elasticity of H	1	Auclert et al. (2024)
σ_e	Volatility idiosyncratic shock	0.58	Auclert et al. (2024)
ρ_e	Persistence idiosyncratic shock	0.968	Auclert et al. (2024)
ϑ	Expectation stickiness	0.935	Auclert et al. (2020)
Trade			
ϕ_T	Share of tradable consumption	0.623	Adjusted data average
λ_T	Sectoral elasticity of substitution	1	Cashin and McDermott (2003) Cardi et al. (2020)
ϕ_D	Home bias	0.57	Import share of 12%
λ_D	Trade elasticity	1	Standard
λ_H	Elasticity of substitution sectoral labor	1	Horvath (2000)
μ_F	Share of local distribution services in C_F	0.754	Hale et al. (2019)
μ_D	Share of local distribution services in C_D	0.266	Yuskavage (2006)
Public policy			
ζ	Interest rate smoothing	0.9	Standard
κ_π	Inflation response	1.5	Standard
κ_y	Output growth response	0.125	Standard
κ_b	Lump-sum tax response to debt	0.027	Bi et al. (2016)
ϕ_{GT}	Tradable content of public consumption	0.25	OECD COFOG
ϕ_D^G	Domestic content of public consumption	1	Only domestic
Determined by normalizations of steady-state values			
C_F^*	Foreign demand for domestic goods	0.21	$Q = 1$
ϕ_H	Non-tradable sector labor bias	0.6	$P_M^* = 1$
χ	Labor disutility weight	0.51	$Y = 1$

Notes: See text for rationale behind the parameter values.

The steady-state share of tradable consumption in total consumption, ϕ_T, is set to 62.3%, reflecting the sample mean of our constructed time series.[11] The home bias ϕ_D is disciplined by targeting

[11]The raw National Accounts tradable consumption share is 43.35%. Since distribution services account for 15% of US GDP (Yuskavage, 2006), equivalent to 18.9% of steady-state consumption, and are classified as non-tradable in

a 12% share of imports in GDP.

Estimates of the elasticity of substitution between tradable and non-tradable goods are scarce for the US. Two notable exceptions are Cardi et al. (2020) and Cashin and McDermott (2003). The former uses a regression of relative expenditures on relative prices estimated with fully modified OLS (FMOLS) and finds a value of 0.78, while the latter combines a cointegration-based estimate of preference parameters with generalized method of moments (GMM)-based estimation of Euler equations and finds a value of 3.5 for the US, though this estimate is imprecise.[12] We use the Cobb-Douglas aggregator as a baseline which corresponds to a unit elasticity of substitution ($\lambda_T = 1$) and therefore does not impose strong complementarity or substitutability between tradable and non-tradable goods. This provides a neutral benchmark given the uncertainty surrounding the empirical estimates and the possibility that US-specific values are higher than found in cross-country panel studies. For the international trade elasticity, we use $\lambda_D = 1$ as in Obstfeld and Rogoff, 2002; Gali and Monacelli, 2005 and conduct robustness with $\lambda_D = 1.5$, another widely used value in the literature. The elasticity of substitution between sectoral labor types λ_H is also set to 1, following the conventional value for the United States established by Horvath (2000). Robustness exercises around those elasticities are presented in Section 4.3.

Distribution costs are calibrated separately for imported and domestically produced tradable goods. For imports, we target a share of distribution costs in foreign final goods C_F of 43%, as estimated by Hale et al. (2019), which yields $\mu_F = 0.754$. This share implies that import-related distribution services account for 9% of total output. For domestically produced tradable goods, we calculate the share of distribution services using the estimate from Yuskavage (2006) that the distribution sector represents approximately 15% of GDP, which implies that distribution services constitute 21% of the value of domestically produced and distributed tradable goods. This share corresponds to $\mu_D = 0.266$.

The monetary response to inflation κ_π and output growth κ_y correspond to 1.5 and 0.125, respectively. Interest rate smoothing ζ equals 0.9. For the lump-sum tax response to government debt, we set $\kappa_b = 0.027$, following Bi et al. (2016) who adapt the estimate from Leeper et al. (2010) for the US to a rule in levels consistent with ours. The baseline calibration assumes that

the National Accounts despite forming part of tradable consumption in our framework, we add this margin to obtain $\phi_T = 0.4335 + 0.189 = 0.623$. This implies a share of distribution services in tradable consumption of $0.189/0.623 = 30\%$, consistent with the evidence in Burstein et al. (2003).

[12] The upper end of this range is nonetheless corroborated by the precisely estimated Canadian counterpart in Cashin and McDermott (2003), which yields a value of 3.4 — Canada being a large industrialized economy closely comparable to the US in terms of income level and consumption structure. Note that cross-country panel estimates typically find lower values ranging from 0.44 (Stockman and Tesar, 1995) to 0.74 (Mendoza, 1991). However, country-specific estimates for large industrialized economies such as the US tend to exceed cross-country panel averages (Akinci, 2011).

$\phi_{GT} = 0.25$, so that government spending is mainly non-tradable. We obtain this value from the sectoral composition of government spending in the US using OECD COFOG data.[13] Our baseline assumes that all tradable public spending is domestic. Section 4.3 explores the sensitivity of the results to a positive import share of public tradable spending, $(1 - \phi_D^G) = 0.35$.

Finally, three parameters, the steady-state level of foreign demand for domestic goods C_F^*, the labor disutility weight χ and the weight on non-tradable labor in the labor aggregate ϕ_H are set to normalize $Y = Q = P_M^* = 1$. Table 2 summarizes the baseline parameter values.

4.2 Explaining sectoral co-movements in consumption and output

We now examine the extent to which our theoretical model can replicate the empirical findings documented in Section 2. An important transmission channel of the model stems from the presence of uninsurable idiosyncratic income risk. Since there is limited liquidity in the economy, many households have a short but non-zero planning horizon. When those households receive additional income, they consume some of it and use the rest to build savings. This precautionary behavior generates intertemporal marginal propensities to consume (iMPCs) out of income shocks which are positive on average and consistent with empirical micro evidence in terms of magnitude (see Auclert et al., 2024).

To illustrate this channel, Figure 2 shows the partial equilibrium response of consumption to a 1% increase in aggregate income with persistence 0.9. Under sticky expectations, the dashed blue line shows that consumption jumps on impact and initially continues to rise as more households update their information set and reduce their savings in anticipation of higher income, producing a hump-shaped response. Importantly, the consumption response is more persistent than income. The dash-dotted orange line captures the consumption response in the standard RANK model. In this model, the response is guided by the permanent income hypothesis and is completely flat due to consumption smoothing. The rise on impact is therefore negligible. The dotted green line shows the partial response in a two-agent New Keynesian (TANK) model as an intermediate case. In the TANK model, a fraction μ of households are hand-to-mouth (HtM), with consumption given by $C_t^{HtM} = Z_t$, while the remaining $(1 - \mu)$ households smooth consumption according to their Euler equation. For the sake of the comparison, we set $\mu = 0.497$ so as to match the fraction of liquidity constrained

[13] There is no clear agreement around the exact tradable share of government spending. Monacelli and Perotti (2008) consider scenarios ranging from 0% to 50% (when excluding compensation of public employees). Bertinelli et al. (2022) estimate this share to be 11% also using COFOG data. However, they assume no tradable content of defense spending, which is implausible when including gross capital formation. We choose to consider the category "Economic Affairs" as purely tradable and attach a 50% weight to categories "Health" and "Defense". Our value is then an average over the 1970-2019 period.

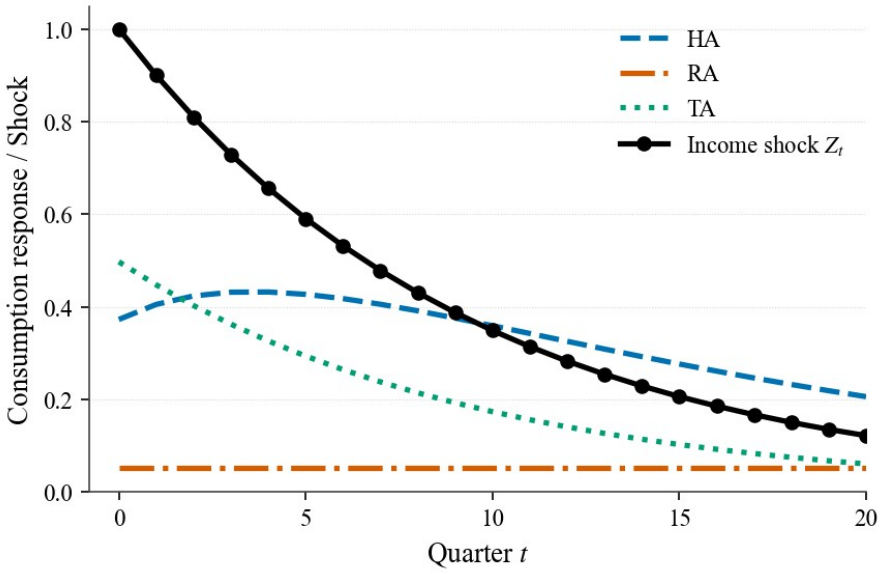


Figure 2: **Partial equilibrium consumption response to an income shock**

Notes: Response of aggregate consumption under sticky expectations to a 1% increase in disposable income with persistence 0.9 in the HANK model (dashed blue line), in the RANK model (dash-dotted orange line), and in the TANK model (dotted green line). The black line with circles shows the persistence of the income shock itself.

agents in the steady state of our HANK model.[14] The initial jump in consumption, stemming from HtM households' unit MPC, is above the HANK counterpart. However, consumption declines much more rapidly thereafter. This reflects the absence of precautionary savings behavior in the TANK which implies that consumption closely tracks the shock process, up to the share of HtM households, and falls to the same level as the RANK counterpart after 20 quarters.

Turning to the government spending shock, the increase in labor demand raises hours worked and pushes up the equilibrium wage in both sectors through union wage bargaining. Since the fiscal expansion is deficit-financed, these effects translate into higher disposable income for households. Figure 3 compares the impulse responses from the S-BVAR with those from our open-economy HANK model and a representative-agent benchmark (RANK). In the baseline HANK specification (solid blue line), the rise in disposable income generates an increase in household consumption, driven by positive micro-founded iMPCs. The gradual build-up induced by sticky expectations aligns closely with the empirical response (upper left panel). As suggested by the partial equilibrium analysis in Figure 2, a positive consumption response would also emerge qualitatively in the TANK model. The key difference, however, lies in the absence of endogenous persistence, which leaves consumption dynamics largely muted over the medium term and unable to produce the gradual build-up observed in the data (see Appendix B for details).

[14] Our model generates a HtM share of 49.7%, of which approximately 34.3% are in low income states (poor HtM) and 15.4% are in middle income states, consistent with the wealthy HtM households identified by Kaplan et al. (2018). As already noted, a two-asset extension would likely redistribute these households into the wealthy HtM category without materially affecting the aggregate MPC.

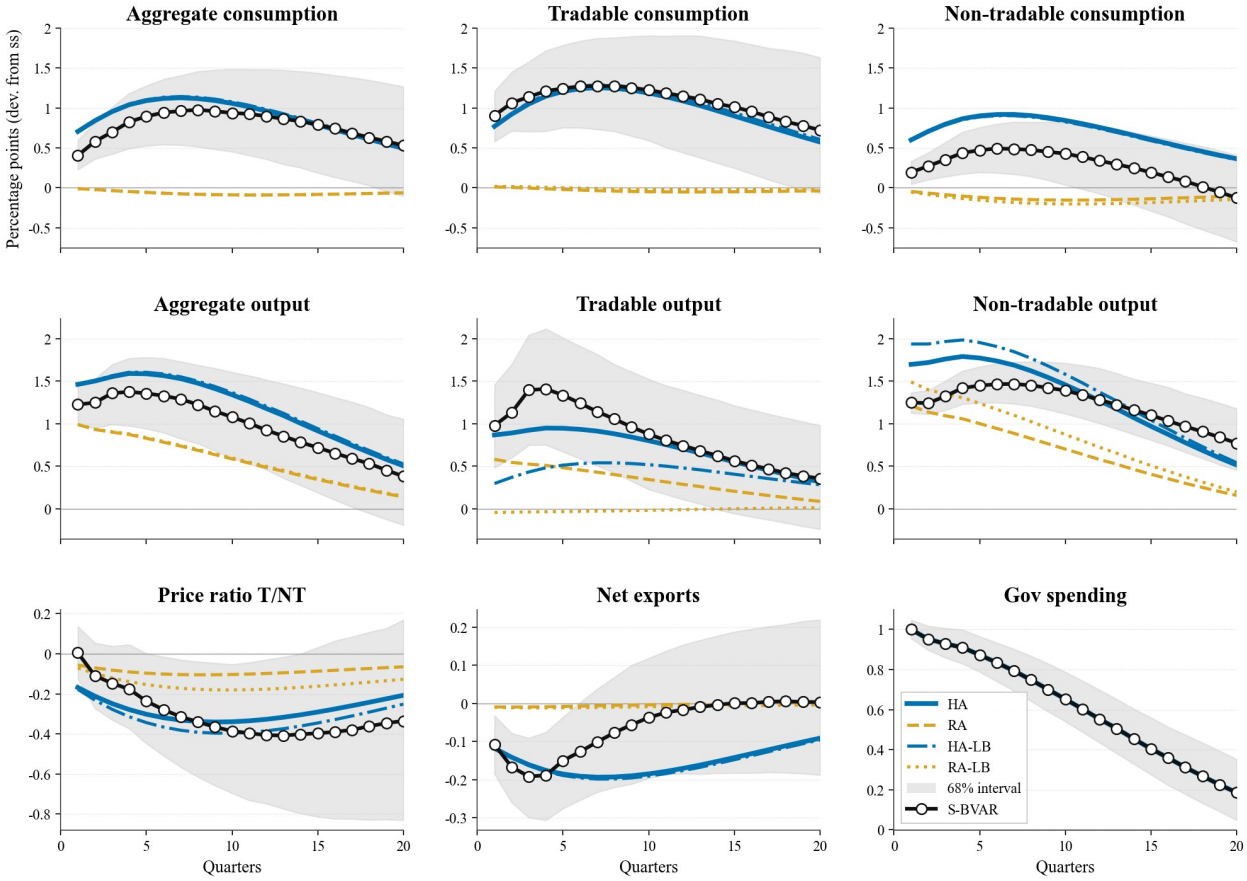


Figure 3: **Impulse responses to a government spending shock: empirical versus theoretical**
Notes: Impulse responses to an orthogonal 1% increase in government spending as a share of output. All variables are in percentage deviation from steady state except for net exports which are in absolute deviations. The solid line with circles is the estimated response from the S-BVAR and the grey area corresponds to 68% confidence bands. The solid and dash-dotted blue lines are the responses from the HANK model under the baseline calibration and when assuming the lower bound case (HA-LB) of no tradable content in government spending, respectively. The dashed and dotted yellow lines are the responses from the RANK model under the same model variations as for the HANK.

At the sectoral level, consumption increases for both tradable and non-tradable goods (top middle and right panels), as households demand for both is proportional to total consumption. However, there is an additional price effect that introduces an asymmetry in the sectoral consumption responses. Following the government spending shock, tradable goods become relatively cheaper when compared to non-tradable goods (bottom left panel). This occurs in the model for two reasons. First, the price of tradable goods is a composite of domestic prices and foreign prices through the import share of tradable consumption. Since the shock is purely domestic, foreign prices do not change and dampen variations in tradable good prices. Second, government spending is more intensive in non-tradable goods and labor market pressures are thus relatively higher in this sector. Under imperfect labor mobility, this translates into asymmetric sectoral wage increases that lower the relative price of tradable goods. Those price effects encourage substitution of private

domestic consumption towards tradable goods and provide a rationale behind the larger increase in tradable consumption observed in the empirical responses.[15]

The theoretical model thus makes explicit how consumption dynamics underlie the positive empirical co-movements in sectoral output. Even though the direct output effect of the government spending shock is smaller in the tradable sector, this is compensated by a larger gain in private consumption. The extent to which the model can replicate the sizable observed rise in tradable output, however, hinges on the strength of the international transmission of the shock, as illustrated in Figure 4.

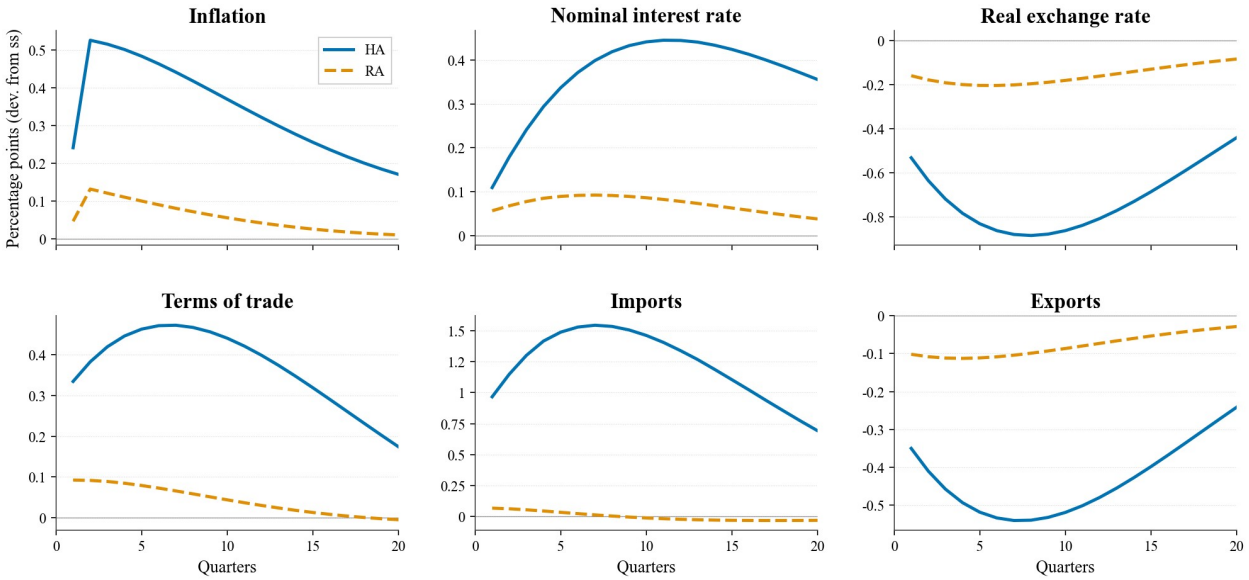


Figure 4: **International transmission of the fiscal shock**

Notes: All variables are in percentage deviation from steady state; net exports and the nominal interest rate are in absolute deviations. Inflation and the nominal interest rate are annualized. Terms of trade are defined as $ToT_t = \frac{P_{DT,t}}{P_{F,t}}$. The blue line is the baseline HANK model and the dashed yellow line is the benchmark RANK model.

In the HANK framework, the joint expansion of public and private demand exerts upward pressure on domestic inflation (upper left panel) and triggers a tightening of monetary policy (upper middle panel). The amplified inflation and demand responses relative to the RANK model are consistent with the general finding in Druedahl et al. (2024) and Aggarwal et al. (2023) that household heterogeneity and high MPCs amplify the aggregate effects of fiscal expansions. Those effects lead to an appreciation of the real exchange rate (upper right panel) and an improvement of the terms of trade (bottom left panel). As Aggarwal et al. (2023) document in a multi-country

[15]The HANK model somewhat undershoots the magnitude of expenditure switching under baseline calibration. Expenditure switching is increasing in the elasticity of substitution between tradable and non-tradable goods and decreasing in the elasticity of substitution between labor types. Raising λ_T toward the upper bound of Cashin and McDermott (2003) of 3.5 brings the model closer to the empirical tradable versus non-tradable consumption responses.

setting, this appreciation feeds into a persistent deterioration of the current account — a mechanism also at work in our model (see Figure 3, final row). However, unlike either of these frameworks, the terms of trade response in our model reflects an endogenous relative price wedge between tradable and non-tradable inflation, driven by imperfect labor mobility across sectors and the concentration of government spending in non-tradables. As imports become relatively cheaper compared to exports, the fiscal shock stimulates domestic demand for imports (bottom middle panel) and weakens foreign demand for exports (bottom right panel), both of which reduce net exports in the bottom middle panel of Figure 3 and thereby attenuate the increase in tradable output. This dampening of the tradable output multiplier through the trade channel is consistent with Druedahl et al. (2024), but our two-sector structure further reveals that its magnitude depends on expenditure switching from private consumers, with imports rising strongly on the back of higher domestic demand and exports falling more moderately as the real exchange rate appreciation erodes competitiveness abroad.

The solid yellow line in Figure 3 shows the impulse responses of the RANK benchmark under the baseline calibration. In this flexible exchange rate RANK setup, a government spending shock raises domestic interest rates, inducing households to postpone consumption, while the resulting real exchange rate appreciation slightly crowds out net exports. Combined with the negative wealth effect from expected future taxes, these channels cause private consumption to fall in response to the fiscal expansion (see Monacelli and Perotti (2008) for the original discussion of this puzzle and its sectoral implications). This stands in sharp contrast to the HANK framework. Consequently, the RANK model struggles to replicate the empirical dynamics. Output does co-move across sectors, but it is driven almost exclusively by the sectoral composition of the shock and is quantitatively too small compared to the data. The international transmission of the shock (see Figure 4) is similarly negligible, driven by the muted inflation response. The mismatch becomes even more pronounced when government spending is assumed to be entirely non-tradable (see Figure 3, dotted yellow line). In this scenario, the flat response in tradable consumption maps directly into a flat response in tradable output, eliminating any cross-sectoral co-movement. The HANK model, by contrast, successfully replicates positive output co-movements even in the absence of a tradable component in government spending — though in this case sectoral differences are somewhat overstated (dash-dotted blue lines).

4.3 Inspecting the transmission channels

In this section, we investigate the contribution of four key channels of the HANK model in matching the sizable increase in tradable output: the sticky expectations, the monetary policy response, the cross-sectoral substitution of consumption and the output leakage from international trade.

Household expectations In our baseline specification, we assume that households form sticky expectations and therefore update their information sets only gradually. This informational friction enables the model to generate hump-shaped impulse responses while maintaining consistency with empirically calibrated and micro-founded intertemporal marginal propensities to consume (iMPCs). For a detailed discussion of this mechanism, see Auclert et al. (2020).

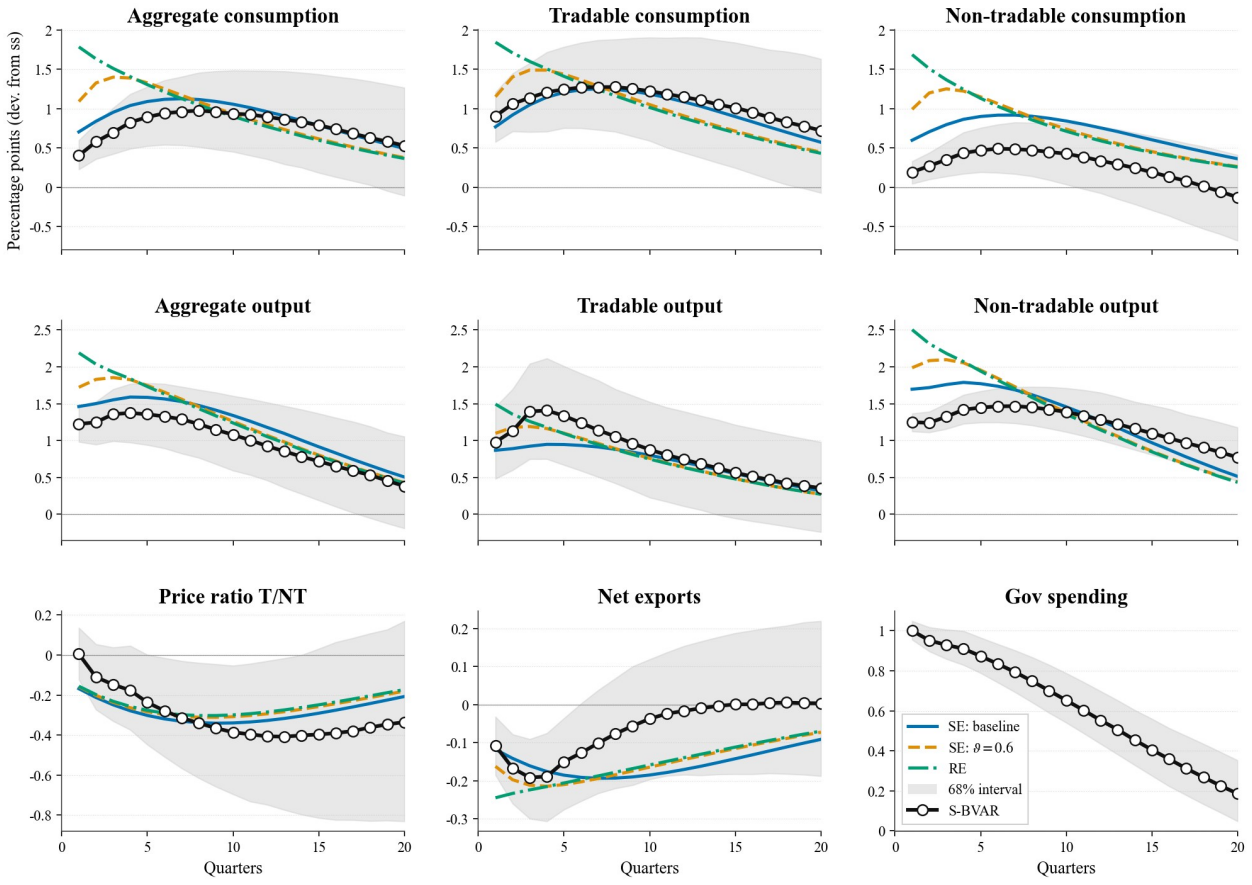


Figure 5: **Comparison of sticky expectations and rational expectations**

Notes: The figure compares impulse responses under sticky expectations (solid blue line, baseline), fully rational expectations (dash-dotted green line) and a faster expectation-updating pace than our baseline with $\vartheta = 0.6$ (dashed yellow line). The solid black line represents the empirical response from the S-BVAR model, with gray shaded areas indicating 68% confidence bands. All responses are normalized to a 1% increase in government spending relative to GDP.

Figure 5 contrasts these results with those obtained under the standard rational expectations paradigm (solid blue line versus dash-dotted green line). Under rational expectations, households

perfectly anticipate the entire path of the fiscal shock from the outset and consequently increase consumption more strongly on impact. This front-loading places the consumption response outside the empirical confidence bands in the initial periods, at odds with the data, where both consumption and output display pronounced hump-shaped dynamics. Since the fiscal shock itself does not follow a hump-shaped path, the failure to generate a consumption hump under rational expectations mechanically carries over to output, which similarly lacks hump-shaped dynamics. This comparison suggests that the hump-shaped output dynamics observed in both the model and the data originate primarily from the consumption response. The dashed yellow line corresponds to a faster expectation-updating pace than our baseline. While responses in this case display some degree of hump-shape, the initial jump remains too strong, which provides further support for the view that the updating pace estimated by Auclert et al. (2020) offers a good fit for our framework.

More broadly, this exercise underscores the importance of household heterogeneity in matching not only the magnitude but also the dynamic profile of empirical impulse responses. In a representative-agent framework, where the consumption response is nearly flat, there is no endogenous mechanism that can generate hump-shaped output dynamics. It is important to note, however, that sticky expectations are not essential for the paper's qualitative conclusions. The main results continue to hold under rational expectations, albeit with different dynamic profiles.

Inflation and monetary policy response Given that monetary policy matters for the consumption response to the government spending shock, we investigate whether the predictions of our HANK model for inflation and the policy rate are consistent with the empirical evidence. To this end, we re-estimate the empirical model, augmenting the variable set with the log change in the GDP deflator and the 3-month Treasury bill secondary market rate. As shown by Perotti (2005), inflation may affect government spending within the quarter through its mechanical effect on the real value of the public wage bill. We therefore adapt our identification strategy by imposing a negative sign restriction on the first period response of government spending to inflation.[16]

Figure 6 compares the results across the two models. In the HANK model, we assume partial wage indexation, setting $\iota_w = 0.69$ in line with the empirical estimate from Lindé et al. (2016), which generates the gradual inflation build-up observed in the data. The blue line confirms that the positive consumption response also obtains in this version of the B-SVAR with unrestricted response of consumption. Inflation is negative on impact but recovers and turns positive after six

[16]Since this sign restriction is less strict than the zero restrictions in our baseline model, additional discipline is required. To avoid imposing our main results from the outset, we restrict only the impact responses of variables for which there is clear consensus in the literature: output and non-tradable output are constrained to respond positively and net exports negatively.

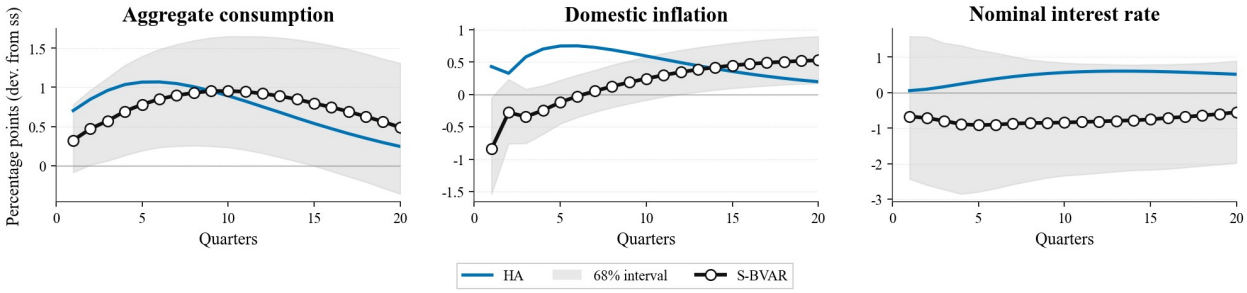


Figure 6: **Inflation and monetary policy response**

Notes: The figure compares impulse responses under partial wage indexation (solid blue line, baseline) and the empirical response from the S-BVAR model estimated under sign restrictions (solid black line, with gray shaded areas indicating 68% confidence bands). All responses are normalized to a 1% increase in government spending relative to GDP.

quarters, overlapping with the model in the longer run. Consistent with the initial inflation decline, the interest rate falls on impact and remains negative throughout the 20-quarter horizon. Although the HANK model implies a positive interest rate response, it remains within the wide empirical confidence bands. We interpret these results as indicative of an accommodative monetary policy response to the fiscal shock despite uncertainty around the sign. This further supports the use of the inertial interest rate rule in the HANK model, as both models imply a decline in the real interest rate over the short to medium run.

International trade pass-through In our baseline HANK model, distribution costs in the tradable sector attenuate exchange-rate pass-through—which is typically found to be too strong in standard open-economy models with producer-currency pricing—and dampen the fall in net exports as follows.[17] The fact that non-tradable goods have to be used to distribute imported tradable goods means that the price of imported goods paid by households $P_{F,t}$ declines less in response to a drop (appreciation) of the real exchange rate (see equation (32)), the more so the higher μ_F. Therefore, the substitution of foreign for domestic tradable goods by domestic and foreign households is attenuated. This effect is counteracted by the positive effect of the aforementioned increase in the price of non-tradable goods $P_{NT,t}$ on the price of domestic tradable goods $P_{D,t}$ (see equation (33)), the more so the higher μ_D, which tends to lower the demand for domestic tradable goods. However, under our calibration this compensation remains partial, as μ_F is much larger than μ_D, implying the observed attenuation of the decline of net exports.

Shutting down this channel ties export prices directly to the exchange rate, amplifying the

[17] Local-currency pricing and border price stickiness are another commonly used mechanism serving a similar purpose (Itskhoki and Mukhin, 2021; de Walque et al., 2023).

competitiveness losses that arise from the exchange rate appreciation. The dashed yellow line shows that net exports consequently fall more sharply relative to the baseline (solid blue line) and now lie outside the empirical confidence bands. As a result, tradable output remains slightly subdued despite a stronger rise in tradable consumption, which itself stems from a larger decline in relative tradable prices that encourages expenditure switching.

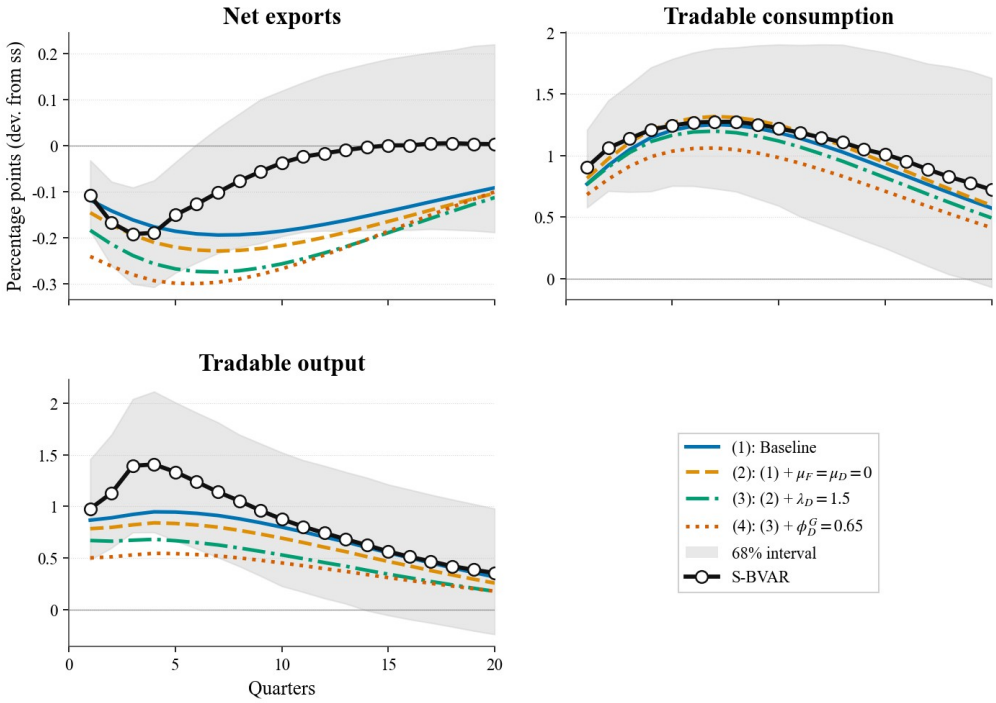


Figure 7: **Sensitivity to international trade**

Notes: Impulse responses to an orthogonal 1% increase in government spending as a share of output. All variables are in percentage deviation from steady state except for net exports which are in absolute deviations. The solid line with circles is the estimated response from the S-BVAR and the grey area corresponds to 68% confidence bands. The solid blue line is the baseline calibration. The dashed yellow line shuts down the distribution sector channel, the dash-dotted green line additionally increases trade elasticity to $\lambda_D = 1.5$ and the dotted red line keeps those modifications and decreases the home bias in government spending to $\phi_D^G = 0.65$.

The key parameters governing the international transmission of the shock are the elasticity of substitution between domestic and foreign tradable goods (λ_D) and the home bias in public tradable consumption (ϕ_D^G). Figure 7 thus also reports a sensitivity analysis around these parameters. The dash-dotted green line increases the trade elasticity to $\lambda_D = 1.5$, another commonly used value in the trade literature (Backus et al., 1994; Chari et al., 2002), while keeping the distribution sector turned off. As domestic and foreign tradables become closer substitutes, competitiveness losses generate larger international leakages: net exports (bottom left panel) deteriorate further, and the resulting second-round effects, via lower domestic activity, reduce households' disposable income and tradable consumption, and impose an additional drag on tradable output. A similar transmis-

sion mechanism operates when we further reduce the home bias in public tradable consumption so that international leakages arise not only through private demand but also directly through public spending. The dotted red line illustrates the case with $\phi_D^G = 0.65$, the same home bias as for private consumption (in the absence of distribution services). When all these effects are combined, the rise in tradable output is significantly curtailed and even falls outside the confidence bands for a short period.[18]

Cross-sectoral substitution of consumption For the HANK model to reproduce the relative increase in private consumption observed in the data, it requires both a sizable decline in the relative price of tradable goods—consistent with the empirical evidence—and a sufficiently high elasticity of substitution between tradable and non-tradable consumption. Figure 8 examines the role of these parameters in shaping consumption co-movements. The blue line corresponds to the baseline calibration.

The yellow line increases the elasticity of substitution between sectoral labor types to $\lambda_H = 100$, effectively imposing fully flexible labor markets. Under this calibration, labor reallocates freely across sectors, and wages equalize to clear the labor market. This eliminates the asymmetric wage response that would otherwise arise when government spending is tilted toward non-tradable goods. As a result, the only remaining source of asymmetry comes from the foreign component of the tradable price, which dampens price increases in that sector. The upper-left panel illustrates that relative prices fall by less in this case, generating weaker substitution of consumption toward tradable goods.

When we additionally reduce the elasticity of substitution between sectoral goods to $\lambda_T = 0.44$—as in Stockman and Tesar, 1995—cross-sectoral substitution becomes even more limited (dash-dotted green line). Consequently, non-tradable consumption remains above the empirical 68% confidence bands throughout the horizon. This provides further evidence that, provided the HANK model is a good description of the empirical economy, tradable and non-tradable goods are substitutes rather than complements, at least for a large developed country such as the US.

[18] Another key factor is the home bias in private consumption, which captures the degree of openness of the economy. Since this parameter is pinned down by the country's import share, its value is uniquely identified for the US. Nevertheless, we investigate the consequences of greater trade openness in Appendix D. A more open economy suffers from stronger output leakages following the real exchange rate appreciation, echoing the lower fiscal multipliers found by Druedahl et al. (2024) in small open economies with heterogeneous households.

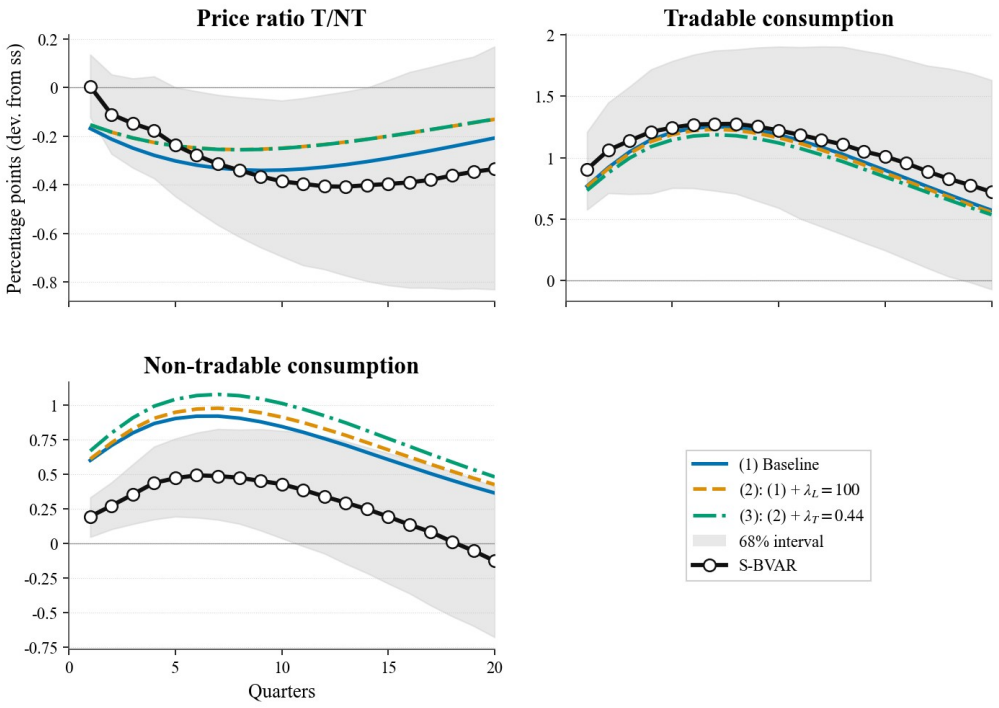


Figure 8: **Sensitivity to cross-sectoral substitution**

Notes: Impulse responses to an orthogonal 1% increase in government spending as a share of output. All variables are in percentage deviation from steady state except for net exports which are in absolute deviations. The solid line with circles is the estimated response from the S-BVAR and the grey area corresponds to 68% confidence bands. The solid blue line is the baseline calibration. The dashed yellow line increases the elasticity of substitution between sectoral labor types to $\lambda_H = 100$, effectively imposing fully flexible labor markets. The dash-dotted green line additionally reduces the elasticity of substitution between sectoral goods to $\lambda_T = 0.44$.

5 Conclusion

This paper shows that household heterogeneity is essential for understanding the sectoral transmission of government spending in an open economy. Using a S-BVAR estimated on US data up to just before COVID-19, we document that a fiscal expansion concentrated in services generates strong spillovers to the goods sector, raises private consumption across the board, and lowers the relative price of tradable goods.

Our open-economy HANK model replicates all these facts. Borrowing constraints, incomplete insurance, and precautionary saving motives amplify the consumption response to deficit-financed spending, while imperfect labor mobility and sector-specific price adjustments generate the observed shifts in relative prices and consumption patterns. In contrast, a representative-agent version fails to produce these co-movements, underscoring the central role of heterogeneity.

These findings imply that fiscal policy can generate sizable cross-sector spillovers, but that their strength depends on openness and the real exchange rate. Policies that support domestic competitiveness can help ensure that the income gains from fiscal expansions translate into stronger

domestic production.

References

Aggarwal, R., A. Auclert, M. Rognlie, and L. Straub (2023): "Excess savings and twin deficits: The transmission of fiscal stimulus in open economies," *NBER Macroeconomics Annual*, 37, 325–412.

Akinci, O. (2011): "A note on the estimation of the atemporal elasticity of substitution between tradable and nontradable goods," *manuscript, Columbia University*.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021): "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models," *Econometrica*, 89, 2375–2408.

Auclert, A., H. Monnery, M. Rognlie, and L. Straub (2023): "Managing an energy shock: Fiscal and monetary policy," Tech. rep., National Bureau of Economic Research.

Auclert, A., M. Rognlie, and L. Straub (2020): "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model," Tech. rep., National Bureau of Economic Research.

——— (2024): "The intertemporal keynesian cross," *Journal of Political Economy*, 132, 4068–4121.

Auerbach, A. J. and Y. Gorodnichenko (2012): "Measuring the output responses to fiscal policy," *American Economic Journal: Economic Policy*, 4, 1–27.

Backus, D. K., P. J. Kehoe, and F. E. Kydland (1994): "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?" *The American Economic Review*, 84, 84–103.

Bertinelli, L., O. Cardi, and R. Restout (2022): "Labor market effects of technology shocks biased toward the traded sector," *Journal of International Economics*, 138, 103645.

Bi, H., W. Shen, and S.-C. S. Yang (2016): "Debt-dependent effects of fiscal expansions," *European Economic Review*, 88, 142–157.

Bianchi, J. and G. Kaplan (2026): "How Small is Small? Non-linearities in Heterogeneous Agent Models," *Working Manuscript*.

Bilbiie, F. O. (2025): "Monetary policy and heterogeneity: An analytical framework," *Review of Economic Studies*, 92, 2398–2436.

Blanchard, O. and R. Perotti (2002): "An empirical characterization of the dynamic effects of changes in government spending and taxes on output," *the Quarterly Journal of economics*, 117, 1329–1368.

Burstein, A. T., J. C. Neves, and S. Rebelo (2003): "Distribution costs and real exchange rate dynamics during exchange-rate-based stabilizations," *Journal of monetary Economics*, 50, 1189–1214.

Cardi, O. and R. Restout (2015): "Imperfect mobility of labor across sectors: a reappraisal of the Balassa–Samuelson effect," *Journal of International Economics*, 97, 249–265.

Cardi, O., R. Restout, and P. Claeys (2020): "Imperfect mobility of labor across sectors and fiscal transmission," *Journal of Economic Dynamics and Control*, 111, 103815.

Cashin, P. and C. J. McDermott (2003): "Intertemporal Substitution and Terms-of-Trade Shocks," *Review of International Economics*, 11, 604–618.

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2002): "Can sticky price models generate volatile and persistent real exchange rates?" *The review of economic studies*, 69, 533–563.

Corsetti, G. and L. Dedola (2005): "A macroeconomic model of international price discrimination," *Journal of International Economics*, 67, 129–155.

de Walque, G., T. Lejeune, A. Rannenberg, and R. Wouters (2023): "Low pass-through and international synchronization in general equilibrium: Reassessing vertical integration," *Journal of International Economics*, 140, 103710.

Debortoli, D. and J. Galí (2025): "Heterogeneity and aggregate fluctuations: Insights from tank models," *NBER Macroeconomics Annual*, 39, 307–357.

Druedahl, J., S. H. Ravn, L. Sunder-Plassmann, J. M. Sundram, and N. Waldstrøm (2024): "Fiscal Multipliers in Small Open Economies With Heterogeneous Households," .

Erceg, C. J., L. Guerrieri, and C. Gust (2005): "Expansionary fiscal shocks and the US trade deficit," *International Finance*, 8, 363–397.

Fagereng, A., M. B. Holm, and G. J. Natvik (2021): "MPC heterogeneity and household balance sheets," *American Economic Journal: Macroeconomics*, 13, 1–54.

Galí, J., J. D. López-Salido, and J. Vallés (2007): "Understanding the effects of government spending on consumption," *Journal of the european economic association*, 5, 227–270.

Gali, J. and T. Monacelli (2005): "Monetary policy and exchange rate volatility in a small open economy," *The Review of Economic Studies*, 72, 707–734.

Guntin, R., P. Ottonello, and D. J. Perez (2023): "The micro anatomy of macro consumption adjustments," *American Economic Review*, 113, 2201–2231.

Hagedorn, M., J. Luo, I. Manovskii, and K. Mitman (2019): "Forward guidance," *Journal of Monetary Economics*, 102, 1–23.

Hale, G., B. Hobijn, F. Nechio, D. Wilson, et al. (2019): "How much do we spend on imports?" *FRBSF Economic Letter*, 1.

Horvath, M. (2000): "Sectoral shocks and aggregate fluctuations," *Journal of Monetary Economics*, 45, 69–106.

Itskhoki, O. and D. Mukhin (2021): "Exchange rate disconnect in general equilibrium," *Journal of Political Economy*, 129, 2183–2232.

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

Leeper, E. M., M. Plante, and N. Traum (2010): "Dynamics of fiscal financing in the United States," *Journal of Econometrics*, 156, 304–321.

Lindé, J., F. Smets, and R. Wouters (2016): "Challenges for central banks' macro models," in *Handbook of macroeconomics*, Elsevier, vol. 2, 2185–2262.

Mankiw, N. G. and R. Reis (2002): "Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve," *The Quarterly Journal of Economics*, 117, 1295–1328.

Mendoza, E. G. (1991): "Real business cycles in a small open economy," *The American Economic Review*, 797–818.

Mertens, K. and M. O. Ravn (2011): "Understanding the aggregate effects of anticipated and unanticipated tax policy shocks," *Review of Economic dynamics*, 14, 27–54.

——— (2014): "A reconciliation of SVAR and narrative estimates of tax multipliers," *Journal of Monetary Economics*, 68, S1–S19.

Monacelli, T. and R. Perotti (2008): "Openness and the sectoral effects of fiscal policy," *Journal of the European Economic Association*, 6, 395–403.

——— (2010): "Fiscal policy, the real exchange rate and traded goods," *The Economic Journal*, 120, 437–461.

Obstfeld, M. and K. Rogoff (2002): "Global implications of self-oriented national monetary rules," *The Quarterly journal of economics*, 117, 503–535.

Perotti, R. (2005): "Estimating the effects of fiscal policy in OECD countries," *Available at SSRN 717561*.

Ramey, V. A. (2011): "Identifying government spending shocks: It's all in the timing," *The quarterly journal of economics*, 126, 1–50.

Ramey, V. A. and M. D. Shapiro (1998): "Costly capital reallocation and the effects of government spending," in *Carnegie-Rochester conference series on public policy*, Elsevier, vol. 48, 145–194.

Ravn, M. O., S. Schmitt-Grohé, and M. Uribe (2007): "Explaining the effects of government spending shocks on consumption and the real exchange rate," .

——— (2012): "Consumption, government spending, and the real exchange rate," *Journal of Monetary Economics*, 59, 215–234.

Rotemberg, J. J. (1982): "Sticky prices in the United States," *Journal of Political Economy*, 90, 1187–1211.

Schmitt-Grohé, S. and M. Uribe (2003): "Closing small open economy models," *Journal of international Economics*, 61, 163–185.

Smets, F. and R. Wouters (2007): "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American economic review*, 97, 586–606.

Stockman, A. and L. Tesar (1995): "Tastes and Technology in a Two-Country Model of the Business Cycle: Explaining International Comovements," *American Economic Review*, 85, 168–85.

Yuskavage, R. E. (2006): "Distributive Services in the US Economic Accounts," in *Conference on Research in Income and Wealth*, vol. 2006.

Appendix

A Robustness of empirical results

We first compare our baseline empirical results with those of Monacelli and Perotti (2008), the benchmark paper that first documented the puzzle surrounding the sectoral transmission of fiscal shocks in an open-economy RANK model. Their sample runs from 1954Q1 to 2006Q2, and the paper is relatively silent on implementation details such as the number of lags, detrending choices, and the treatment of constants. For comparability, we re-estimate our SVAR over the same sample and with the same variable set as Monacelli and Perotti (2008): government spending on goods and services, net taxes (revenues less transfers), GDP, output of goods, output of services, private investment, the three-month Treasury bill rate, and the relative price of goods to services. We then compare results obtained under our Bayesian procedure with those from an OLS estimation using four lags and both linear and quadratic trends, following as closely as possible the original specification. A specific concern raised by the use of a Bayesian prior is that it may introduce spurious persistence in the impulse responses, which could artificially favour our model with sticky expectations. Figure A.1 addresses this concern directly. The OLS estimates (dashed blue line) successfully replicate the stylized facts reported in Figure 1 of Monacelli and Perotti (2008), confirming that our re-estimation recovers their original findings. Comparing the two lines, the Bayesian prior does not meaningfully increase persistence relative to OLS. Its main effect is to smooth out the erratic impact responses that arise under OLS: in particular, aggregate and tradable outputs no longer exhibit the sharp initial drop followed by a delayed increase that characterises the OLS estimates, instead displaying the hump-shaped responses more typical of standard macroeconomic models. The key qualitative features of the impulse responses—including the sectoral transmission pattern—are therefore not an artifact of the Bayesian prior, but a genuine feature of the data already present under OLS.

Having established that the Bayesian prior does not drive our results, we now assess the robustness of our findings to the two dimensions in which our baseline specification departs most substantively from Monacelli and Perotti (2008): the definition of government spending and the sample period. Our baseline S-BVAR defines government spending as general government consumption expenditures and gross investment over the period 1954Q1 to 2019Q3. In what follows, we examine whether the key stylized facts are preserved when we exclude military expenditures or gross capital formation from the spending aggregate, and when we restrict the sample to exclude the Global Financial Crisis.

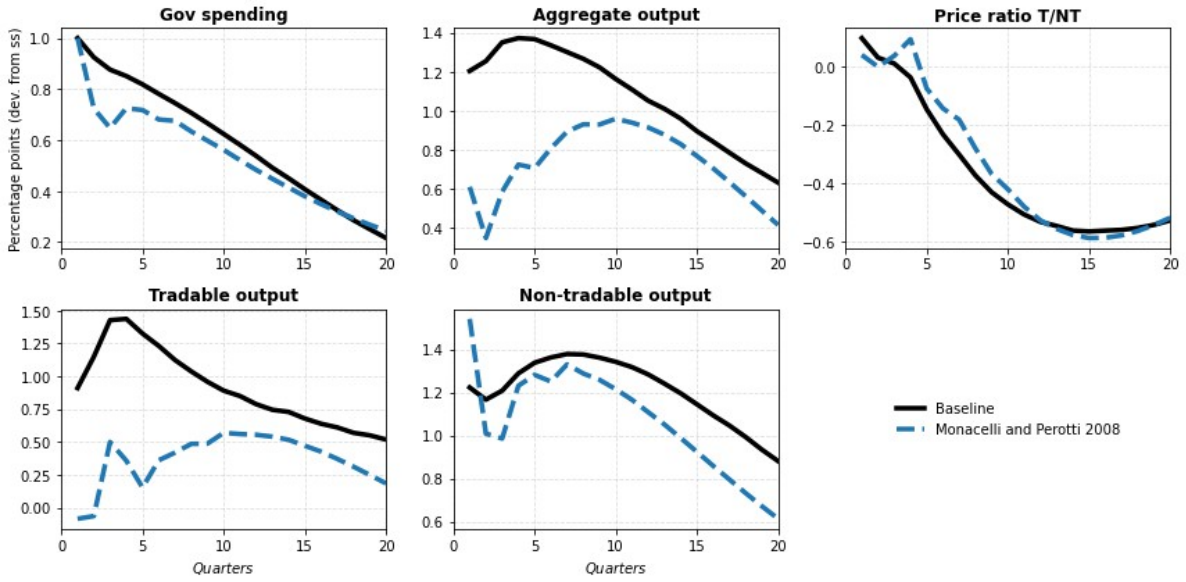


Figure A.1: **Baseline BVAR versus Monacelli and Perotti (2008): impulse responses to a government spending shock**

Notes: The figure compares impulse responses to a government spending shock identified under two estimation procedures, over the sample 1954Q1–2006Q2 and using the variable set of Monacelli and Perotti (2008). The solid black line corresponds to our baseline Bayesian VAR. The dashed blue line corresponds to OLS estimation under four lags with linear and quadratic trends, replicating the original specification of Monacelli and Perotti (2008). All responses are normalised to a 1% increase in government spending on impact.

Figure A.2 presents the empirical impulse responses under three alternative specifications:

1. **Exclusion of military expenditures** (dashed blue line). This specification addresses potential concerns that military spending may disproportionately drive the expansion in tradable output, as defense expenditures represent potentially large shocks with higher tradable goods intensity. Furthermore, the anticipated component of military spending is typically larger, which may represent a challenge for our identification strategy.

2. **Exclusion of gross capital formation** (dotted orange line). This robustness check follows similar reasoning to the first specification, as investment expenditures may exhibit different sectoral compositions and anticipation effects compared to consumption expenditures.

3. **Exclusion of the Global Financial Crisis** (dash-dotted green line). This specification restricts the sample period to exclude the financial crisis, thereby alleviating concerns that this exceptional episode may bias the results. It also facilitates comparison with earlier literature on the sectoral transmission of fiscal shocks, notably Monacelli and Perotti (2008).

As Figure A.2 demonstrates, the qualitative patterns remain consistent across all specifications. The primary differences concern the relative persistence of macroeconomic variables. In the specification excluding military expenditures, the underlying shock process differs substantially, which

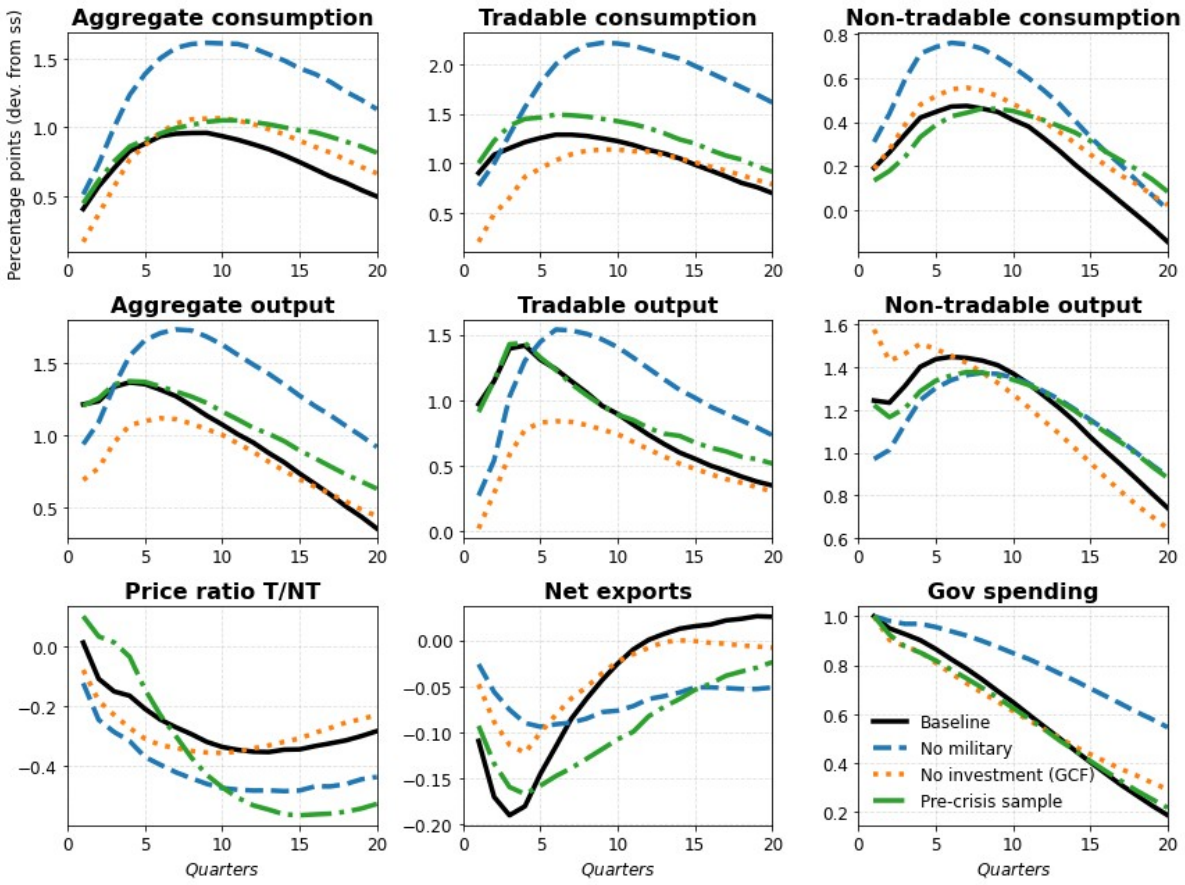


Figure A.2: **Robustness checks for baseline empirical specification**

Notes: The figure displays impulse responses to a government spending shock under alternative specifications. The solid black line represents the baseline specification (general government consumption and investment, 1954Q1–2019Q3). The dashed blue line excludes military expenditures from government spending. The dotted orange line excludes gross capital formation. The dash-dotted green line excludes the Global Financial Crisis period from the sample. Gray shaded areas represent 68% confidence bands for the baseline specification. All responses are normalized to a 1% increase in government spending relative to GDP.

accounts for the apparently larger responses of other variables.

As an additional robustness check, we compare our identified structural government spending shock series with that of Mertens and Ravn (2014), whose sample overlaps with our reduced sample excluding the Global Financial Crisis. Figure A.3 plots both series over their common sample period. The two series track each other closely, with a correlation of 0.76, despite relying on independent identification strategies. This close correspondence provides further evidence that our empirical model recovers government spending shocks that are consistent with those identified in prominent contributions to the literature.

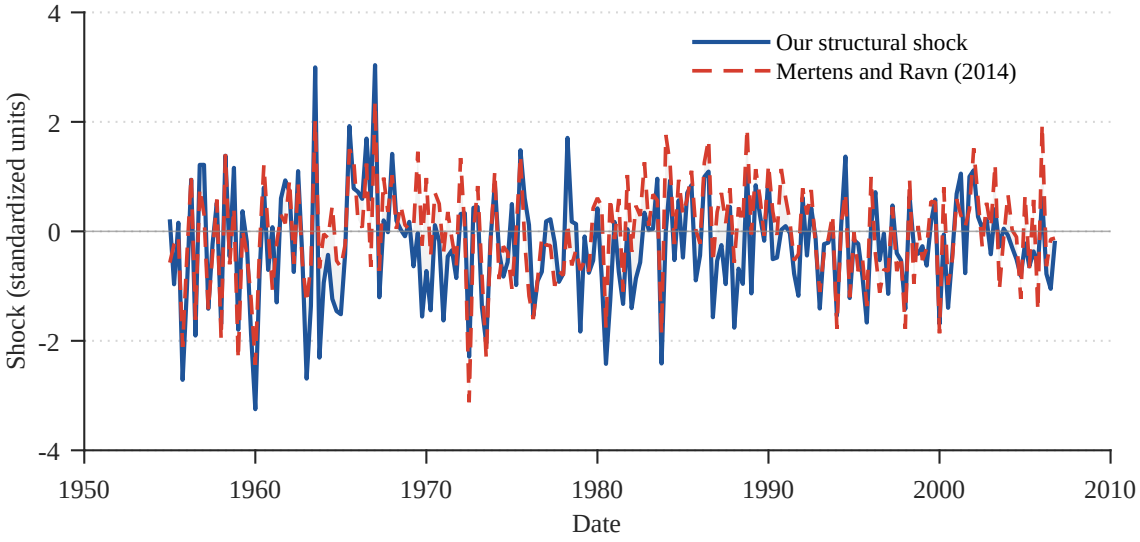


Figure A.3: **Comparison of identified government spending shocks**

Notes: The figure plots the structural government spending shock series identified by our baseline SVAR specification against that of Mertens and Ravn (2014), over their common sample period. Both series are identified using Blanchard and Perotti (2002)'s methodology. The correlation between the two series is 0.76.

B Comparison with Two-Agent model

A substantial literature captures household heterogeneity in a reduced-form manner by distinguishing between two types of agents: liquidity-constrained households that consume their entire current income each period, and Ricardian households that smooth consumption intertemporally by satisfying an Euler equation, as in the standard representative-agent New Keynesian (RANK) framework. Several contributions (e.g., Debortoli and Galí, 2025) argue that this two-agent New Keynesian (TANK) model provides a tractable and empirically plausible alternative to fully specified HANK models.

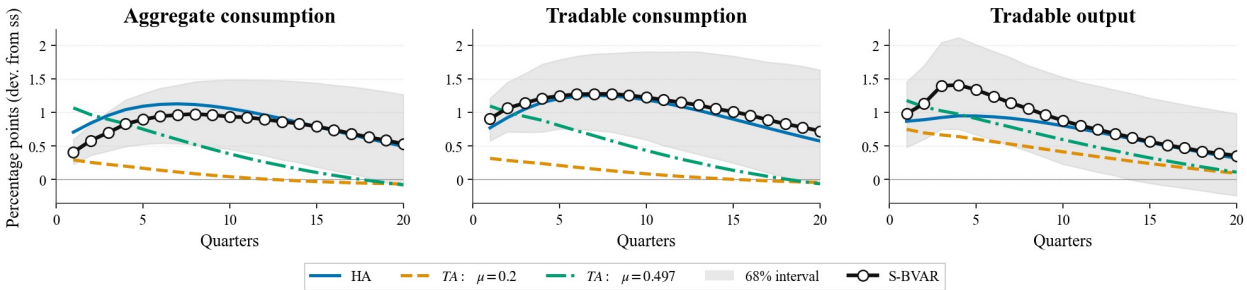


Figure B.1: **Comparison of HANK and TANK models**

Notes: The figure compares impulse responses from the HANK model (solid blue line, baseline) with two TANK specifications: $\mu = 0.2$ (dashed yellow line) and $\mu = 0.497$ (dash-dotted green line), where μ represents the fraction of hand-to-mouth (HtM) households. The solid black line represents the empirical response from the S-BVAR model, with gray shaded areas indicating 68% confidence bands.

Figure B.1 compares the impulse responses generated by our HANK model with those obtained from a simple TANK specification. In the TANK model, a fraction μ of households are hand-to-mouth (HtM), with consumption given by $C_t^{HtM} = Z_t$, while the remaining $(1 - \mu)$ households smooth consumption according to their Euler equation. The dashed yellow line corresponds to $\mu = 0.2$, a calibration commonly used in the literature. Under this parameterization, aggregate consumption responds only weakly, and the tradable output response remains close to the RANK benchmark. Matching the magnitude of the initial consumption response therefore requires a substantially larger share of HtM households.

The dash-dotted green line increases this share to $\mu = 0.497$ so as to match the fraction of liquidity constrained agents in the steady state of our HANK model. While this calibration generates a sizable impact response of consumption, it fails to replicate the observed persistence: consumption declines too rapidly relative to the empirical response. The underlying reason is the absence of a precautionary savings motive in the TANK framework. Consequently, intertemporal marginal propensities to consume (iMPCs) are effectively zero beyond the initial period of the income shock, and aggregate consumption dynamics closely track current income, which in our model is tightly linked to the government spending shock.

This comparison highlights a key advantage of the HANK framework in matching empirical impulse responses in both magnitude and dynamic shape. In particular, the hump-shaped consumption response emerges from the interaction between precautionary savings—which endogenously generates persistence in HANK models—and sticky expectations.[19]

C Sticky prices

Our baseline model features sticky wages and flexible prices. This configuration is standard in HANK models (e.g. Auclert et al., 2024), as it is sufficient in principle to introduce nominal rigidities while avoiding countercyclical profits. Nevertheless, it is informative to assess the robustness of our results to the introduction of price stickiness. To this end, we extend the model by incorporating nominal rigidities in prices.

Specifically, we assume that the intermediate production sector consists of a continuum of monopolistically competitive firms that maximize profits subject to linear production technology $Y_{S,t}(k) = N_{S,t}(k)$ and quadratic price adjustment costs à la Rotemberg (1982). From the profit

[19] A related literature seeks to combine a precautionary savings motive with the tractability of TANK models (e.g. Bilbiie, 2025). Such approaches provide valid alternatives when the primary objective is to match empirical impulse responses, as in the present paper.

maximization problem of the final producer, the demand function for the generic firm producing variety k is given by:

$$Y_{S,t}(k) = \left(\frac{P_{Z,t}(k)}{P_{Z,t}}\right)^{-\theta_\pi} Y_{S,t} \tag{C.1}$$

where θ_π denotes the elasticity of substitution between varieties and $P_{Z,t} \in \{P_{DNT,t}, P_{DT,t}\}$ represents the domestic sectoral price. The optimization problem of firm k is:

$$\max_{P_{Z,t}(k)} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t C_t^{-\sigma} Y_{S,t} \left[\left(\frac{P_{Z,t}(k)}{P_{Z,t}}\right)^{-\theta_\pi} \left(\frac{P_{Z,t}(k)}{P_{Z,t}} - W_{S,t}\right) - \frac{\psi_{\pi,S}}{2} \left(\frac{P_{Z,t}(k)}{P_{Z,t-1}(k)} - 1\right)^2 \right] \tag{C.2}$$

where $\psi_{\pi,S}$ represents the price adjustment cost parameter. The first-order condition of this problem yields the following sectoral price Phillips curves:

$$\frac{W_{S,t}}{P_{Z,t}} = \frac{\theta_\pi - 1}{\theta_\pi} + \frac{\psi_{\pi,S}}{\theta_\pi} \left(\pi_{p,Z,t}(\pi_{p,Z,t} - 1) - \beta \mathbb{E}_t \left[\frac{C_t^\sigma}{C_{t+1}^\sigma} \frac{Y_{S,t+1}}{Y_{S,t}} \pi_{p,Z,t+1}(\pi_{p,Z,t+1} - 1) \right] \right) \tag{C.3}$$

where $\pi_{p,Z,t} \equiv P_{Z,t}/P_{Z,t-1}$.

For the calibration, we set $\theta_\pi = 6$ and choose $\psi_{\pi,S}$ to target a price Phillips curve slope of 0.0066, as estimated by Lindé et al. (2016)—the same study we use for the calibration of the wage Phillips curve.

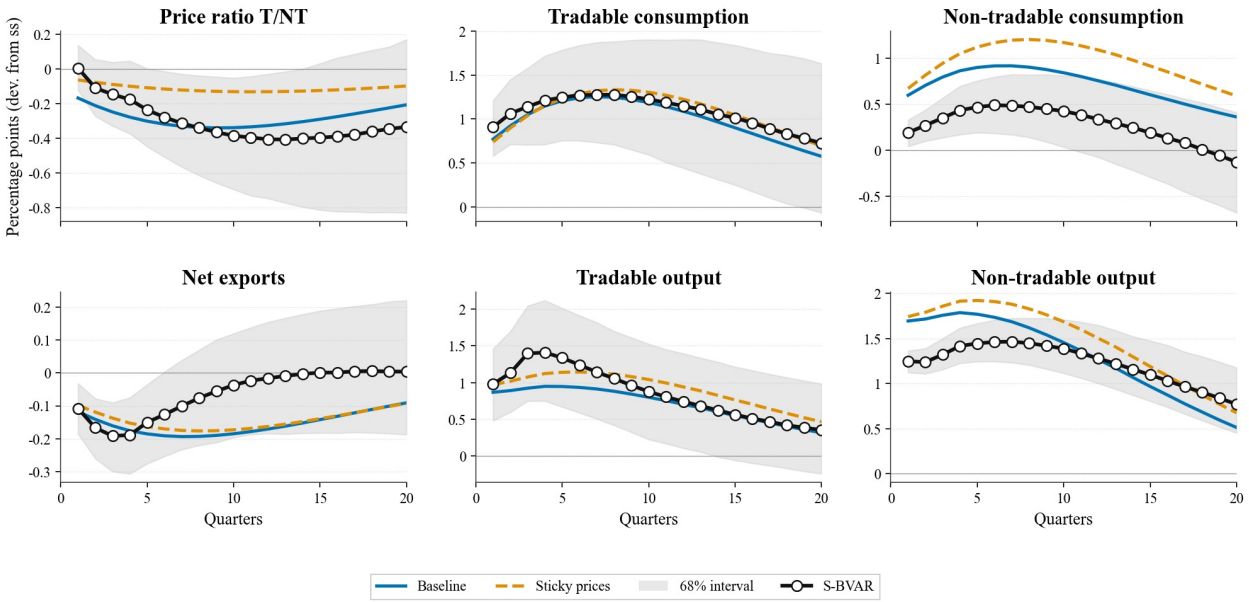


Figure C.1: **Effects of price stickiness**

Notes: The figure compares impulse responses under flexible prices (solid blue line, baseline) and sticky prices (dashed yellow line). The solid black line represents the empirical response from the S-BVAR model, with gray shaded areas indicating 68% confidence bands. The sticky price specification incorporates Rotemberg (1982) quadratic price adjustment costs with a Phillips curve slope of 0.0066.

Figure C.1 compares the baseline economy with flexible prices to an economy featuring both wage and price stickiness (solid blue line versus dashed yellow line). The introduction of price rigidities amplifies the increase in real wages following the public spending shock and attenuates the loss of international competitiveness. Both effects raise household income and, in turn, support consumption across all sectors. Price stickiness also dampens relative sectoral price adjustments, thereby reducing the degree of cross-sectoral substitution. As a result, consumption rises by similar magnitudes in both the tradable and non-tradable sectors.

D Import share

The United States is a relatively closed economy and is thus less exposed to international leakages following real exchange rate appreciation. The dashed yellow line in Figure D.1 explores the implications of the fiscal shock for a more open economy. To this end, we increase the import share from 12% to 25%.

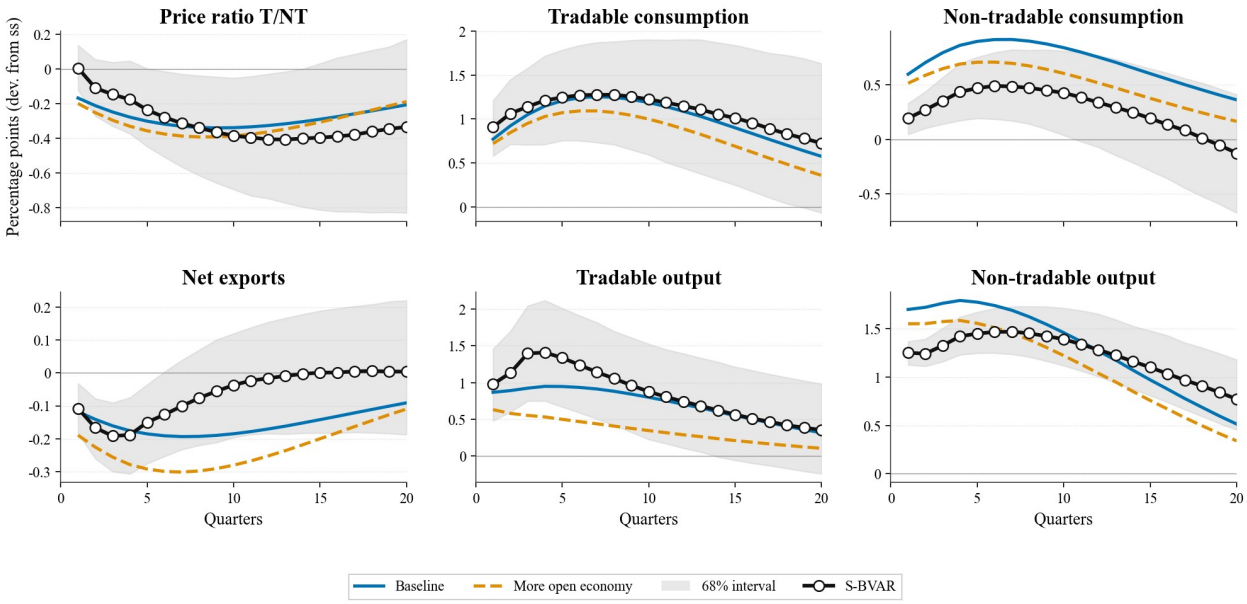


Figure D.1: **Effect of economy openness**

Notes: The figure compares impulse responses from the baseline (solid blue line) to a more open economy (dashed yellow line). The solid black line represents the empirical response from the S-BVAR model, with gray shaded areas indicating 68% confidence bands.

In this scenario, the economy experiences a larger decline in net exports (bottom left panel), which curtails the rise in tradable output. Moreover, there are second-round effects on aggregate output and consumption as the income gains are reduced. Overall, tradable output significantly undershoots its counterpart in the US calibration and thus falls short of the empirical response.

Acknowledgements

We thank Christian Wolf, Greg Kaplan, Raf Wouters, Jolan Mohimont and one anonymous referee for useful comments. We also benefited from helpful discussions during presentations at the National Bank of Belgium, UCLouvain, 30th International Conference on Macroeconomic Analysis and International Finance, 32nd Conference on Computing in Economics and Finance, Annual Royal Economic Society Conference and the 41st meeting of the European Economic Association. We thank Adrien Auclert for providing guidance on his Sequence-space Jacobian method for solving heterogeneous-agent models.
The views expressed in this paper are those of the authors and should not be attributed to the National Bank of Belgium, the European Central Bank or the Eurosystem. All remaining errors are ours.

Charles de Beauffort (corresponding author)
Nationale Bank van België/Banque nationale de Belgique, Brussels, Belgium; email: charles.debeauffort@nbb.be

Ansgar Rannenberg
European Central Bank, Frankfurt am Main, Germany; email: ansgar-romero.rannenberg@ecb.europa.eu

© European Central Bank, 2026

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the Social Science Research Network electronic library or from RePEc: Research Papers in Economics. Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website.

PDF ISBN 978-92-899-7993-1 ISSN 1725-2806 doi:10.2866/1987088 QB-01-26-207-EN-N