One Rate, Different Creditors: Sovereign Debt Ownership and Monetary Transmission*

Sandra Gomes† Ante Šterc‡

October 23, 2026

Abstract

We study who pays for a common monetary tightening in a monetary union whose members differ in how much sovereign debt they owe and in who holds it. In a two-country heterogeneous-agent model calibrated to euro-area debt, investor holdings and household balance sheets, governments pass higher debt service on to households as lower transfers. After a 25 basis-point tightening the poorest fifth of households cuts consumption by 2 to 2.5 per cent on impact, the richest fifth by under 0.2 per cent, and the higher-debt country contracts more. Per household, the gap between the two countries is concentrated at the bottom; in the aggregate, about two fifths of it comes from wealthy hand-to-mouth households. Cross-border sovereign holdings within the union widen the gap, almost twice as much at observed debt ratios as at equal ones. Their effect reaches poorer households through transfers and richer ones through asset returns.

JEL Codes: E21, E52, E62, F45
Keywords: monetary union; sovereign debt ownership; monetary transmission; household heterogeneity

*We thank António Antunes, João Amador, Nuno Alves, João Guerreiro, and Juan Pablo Nicolini for valuable comments. The views expressed here are those of the authors and do not necessarily represent the views of Banco de Portugal, or the Eurosystem.

†Banco de Portugal. Email: sgomes@bportugal.pt.
‡Banco de Portugal and Nova School of Business and Economics. Email: asterc@bportugal.pt.

1 Introduction

Who pays for a common monetary tightening? In a monetary union the policy rate is common, but each government finances its own debt, and its creditors need not live in the issuing country. A tightening lowers inflation and raises real interest rates, and both raise what a government pays, in real terms, on the debt it has issued. Its creditors receive that return, and households pay for it as taxes and forgone transfers. If the debt is held at home, what the government pays returns to domestic households; if it is held in another member, the tax is paid at home and the return crosses the border. What a household loses therefore depends on how much its government owes, who holds that debt, and whether the household has liquid wealth to absorb a cut in its income.

We answer the question in a two-country heterogeneous-agent model of a monetary union, each country a group of euro-area members, calibrated to euro-area debt, investor holdings and household balance sheets. We follow groups of households through a common tightening. Households hold deposits, which they can adjust every period, and an illiquid claim on a domestic intermediary, which they can adjust only occasionally. The intermediary holds domestic equity and sovereign bonds: its own government's and, in the creditor country, part of the other's. Governments differ in their debt, in who holds it, and in their taxes, spending and transfers. Transfers absorb changes in debt service.

A 25 basis-point annualised tightening contracts consumption by 0.37 per cent on impact in the lower-debt country, which we call H, and by 0.48 per cent in the higher-debt country, F. Within each country the contraction falls hardest on the poorest households. The poorest fifth by net worth cuts consumption by 2.0 per cent on impact in H and 2.5 per cent in F, the richest fifth by less than 0.2 per cent, so a group with about 7 per cent of consumption accounts for close to two fifths of the fall. These households have almost no liquid wealth and live off earnings and transfers, and about 85 per cent of their fall comes through the transfer cut that passes on the higher debt service.

Which households pay depends on how governments balance their budgets; over five years, how much more the higher-debt country pays depends on it far less. When the labour tax adjusts and transfers are held at their steady-state value, the poorest fifth's share of the impact fall drops from 40 and 36 per cent to 16 and 14, while the five-year gap between the countries and the part of it that ownership accounts for change little.

Per household, the difference between the two countries is also concentrated at the bottom: half a percentage point on impact for the poorest fifth, three hundredths of a point for the richest. The aggregate gap, 0.11 percentage points on impact and 0.44 percentage-point-quarters over twenty quarters, weights each group by what it consumes. On that measure about two fifths of it comes from the wealthy hand-to-mouth, households with illiquid wealth but no liquid buffer, who are almost twice as common in F. Holding deposits costs F's households more of the return on their wealth, so more of them run their buffers down. That they also hold more wealth and deposits relative to their income cushions them, but does not close the gap.

The gap arises through the government budget: F owes more, and part of its debt service goes to households in H, whose intermediaries hold more of F's debt than F's hold of H's. Asset returns offset part of it: F's households hold more illiquid wealth relative to what they consume, and the higher return that follows the tightening supports their spending.

Where the debt is held is the margin this paper adds. We compare four calibrated unions, with and without intra-union holdings, at observed debt ratios and at a common union average. Removing the holdings lowers the cumulative gap by 0.18 percentage-point-

quarters at observed ratios and by 0.09 at the common one. This ownership increment is present even at equal debt, and unequal debt almost doubles it. Ownership works in two ways, both scaling with the stock of debt. Through the budget it moves debt service across the border, which reaches the households who live off transfers. Through asset returns it tilts the creditor's intermediary toward sovereign bonds and the debtor's toward equity, and over the tightening the book tilted toward bonds pays more; the richest fifth receives three quarters of its increment this way. The increment is therefore spread far more evenly than the gap, and the richer two fifths carry more than half of it in the aggregate.

A Shapley decomposition over four groups of primitives attributes more than the whole gap to the sovereign block, within which the debt level does most of the work, as Bellifemine et al. (2025) would predict; household differences add to the gap, and fiscal structure and country size offset part of it.

Ownership changes the exposures through which the shock reaches households: the cross-border flow of debt service, and the split of each intermediary's book between bonds and equity. The comparisons are between recalibrated economies that hit the same household targets (Section 5.3).

The closest work studies fiscal heterogeneity in monetary unions. Bellifemine et al. (2025) show that differences in legacy debt make common monetary policy transmit unevenly through debt service and fiscal space. We find the same level channel dominant and ask how the location of the creditors changes it and who pays within each country. Their members trade a single union-wide bond, which does not separately trace where a country's creditors live, the margin studied here. Chen et al. (2025) study fiscal consolidation in a euro-area heterogeneous-agent model with the same country grouping, and Bayer et al. (2024) the shifting of burdens across and within members. Both have one-period debt and no ownership margin. Foreign holdings of public debt matter for fiscal multipliers (Broner et al., 2022; Priftis and Zimic, 2021). We study the related incidence question for a common monetary shock. The sovereign-bank doom loop (Brunnermeier et al., 2016; Bocola, 2016; Farhi and Tirole, 2018) and the rise of domestic holdings in crises (Broner et al., 2014) concern banks and sovereign risk. Our holder is the household sector, through an intermediary without a balance-sheet constraint, and our shock is a common policy rate. In closed economies a larger stock of privately held public debt weakens monetary transmission (Caramp and Feilich, 2026) and lowers fiscal multipliers (Grancini, 2025), because households hold more of it. Here household wealth is held at its survey values, so a larger stock raises debt service rather than households' liquidity.

The incidence results relate to work on how monetary policy reaches households, through their earnings, their balance sheets and the government budget (Kaplan et al., 2018; Auclert, 2019; Holm et al., 2021), including in the euro area (Slacalek et al., 2020; Kase et al., 2025). Almgren et al. (2022) show that euro-area transmission varies across countries with the share of hand-to-mouth households, a fact that anchors our household calibration; here the sovereign block creates most of the asymmetry, and household liquidity decides whose spending carries it. What we add is the government's debt service and the location of its creditors. Corsetti et al. (2022) and Hintermaier and Koeniger (2024) study transmission through housing finance and tenure, which our model does not have, and Cohen (2026) the integration of private capital markets and risk sharing, which can change independently of who holds the government's debt.

Section 2 documents the facts the model is built on, Sections 3 and 4 set out the model and its calibration, and Section 5 reports the common tightening, who pays for it and the ownership comparison. The model is stated in full in Appendix A and the channel accounting in Appendix B; the local projections, the asymmetries measured one at a time,

and further results and robustness are in Appendices C, D and E.

2 Evidence

Four differences across euro-area countries discipline the model: how much sovereign debt each government owes and at what cost, who holds that debt, what household balance sheets look like, and how large and how progressive the tax-and-transfer system is. This section documents each at the country level. We group the countries only at the end, in Section 2.5, because which country belongs on which side is a grouping this paper makes on observed averages rather than a label taken from the literature.

2.1 Sovereign debt and borrowing costs

Figure 1 shows the two sovereign facts. Panel (a) plots general government debt relative to output. The dispersion is wide and persistent: it opens during the crisis years, and although the ordering moves at the margins it does not reverse. Germany sits near the bottom of the panel throughout and Italy and Greece near the top. On the narrower debt securities perimeter the model uses, Germany owes 51.9 per cent of annual output at the end of 2021 against Italy's 121.3.

Panels (b) and (c) plot ten-year yields relative to Germany. The crisis dispersion is large enough that showing it on one scale would hide everything since, which is why the recent window is drawn separately. Spreads widened sharply in 2010–2012, compressed afterwards, and have not closed. The model takes a modest input from this: borrowing costs in the higher-debt countries respond to the fiscal position while those in the lower-debt countries do not. Section 3.6 embeds that as a debt-elastic spread which is zero in steady state, so the calibration is undistorted and the spread operates only along the transition.

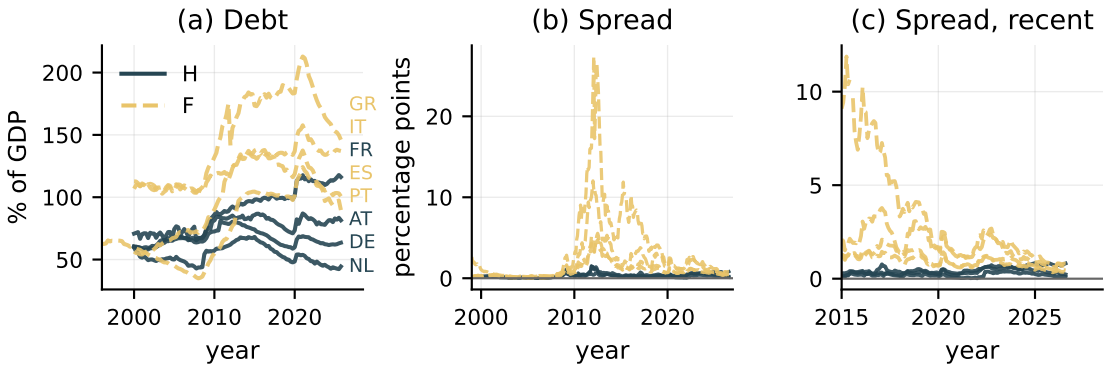


Figure 1: Sovereign debt and borrowing costs. Panel (a) is general government debt as a per cent of GDP. Panel (b) is the ten-year government bond yield minus the German yield, in percentage points, over the full sample. Panel (c) is the same spread from 2015 onwards, drawn at its own scale because the 2012 peak would otherwise compress every recent difference to invisibility. Countries are labelled at the end of their line in panel (a); panels (b) and (c) carry the same eight series, with Germany omitted from the spread by construction. Sources: Eurostat, ECB.

2.2 Who holds the debt

Figure 2 plots the share of each government's debt held by foreign investors and by its own central bank, over two decades. The point of the long window is that the cross-country differences are structural rather than an artefact of the date at which the model is calibrated: the ordering is visible before the crisis, through it, and after the asset purchase programmes changed the level for everyone.

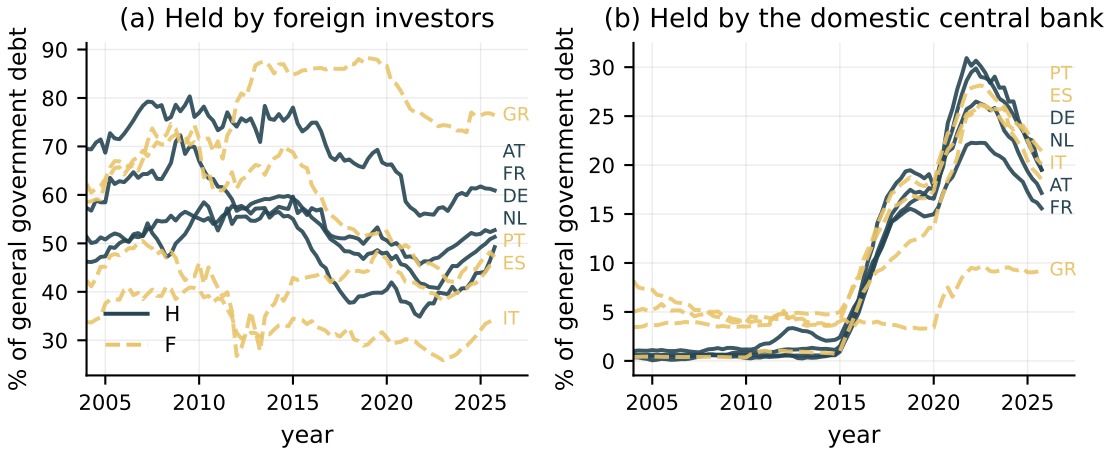


Figure 2: Sovereign debt ownership. Panel (a): share held by foreign investors. Panel (b): share held by the domestic central bank. Per cent of general government debt. Source: IMF sovereign investor base database.

The calibration uses a finer split than the figure. The investor-base data divide holders into domestic and foreign on the total debt stock, as does the cross-country evidence of Fang et al. (2025), whereas the model needs to know whether a foreign holder is inside the union or outside it, since only the first returns debt service to a household the model contains. Section 4 therefore builds the holder shares from ECB securities holdings statistics with the IMF bilateral portfolio survey as the key, on the debt securities perimeter. The two sources are consistent in ordering but not interchangeable in level, and the figure should be read as evidence that the difference is persistent rather than as the calibration input itself.

2.3 Household balance sheets

We measure household balance sheets with Wave 4 of the Eurosystem Household Finance and Consumption Survey. Deposits and sight accounts map onto the model's liquid account b. The illiquid account a, the one subject to the rebalancing friction, is net worth less deposits, which is what the calibration of Section 4 targets; directly held mutual funds, bonds and listed shares are its most liquid component, and the figures below use that financial split to show how portfolios differ.[1] The mapping is about adjustment, not returns: deposits are freely adjustable, while securities are rebalanced infrequently. In the model deposits earn the realised return on the domestic government bond net of a wedge ξ, the illiquid account is the residual claimant on the intermediary's book, and the two accounts differ in how the book's return reaches them, not in a risk premium.

[1] Country moments are computed with household survey weights and averaged across the five implicates. Group moments are weighted by 2021 GDP.

Figure 3 shows the fact that matters most for the model, the share of households living hand to mouth, under two definitions. Panel (a) counts securities as liquid, which is the usual convention. Panel (b) counts only deposits, which is what the model's liquid account literally contains. The shares are decomposed into households with no wealth to speak of and households with illiquid wealth but no liquid buffer, and the split moves a good deal between the two definitions, which is why both are shown.

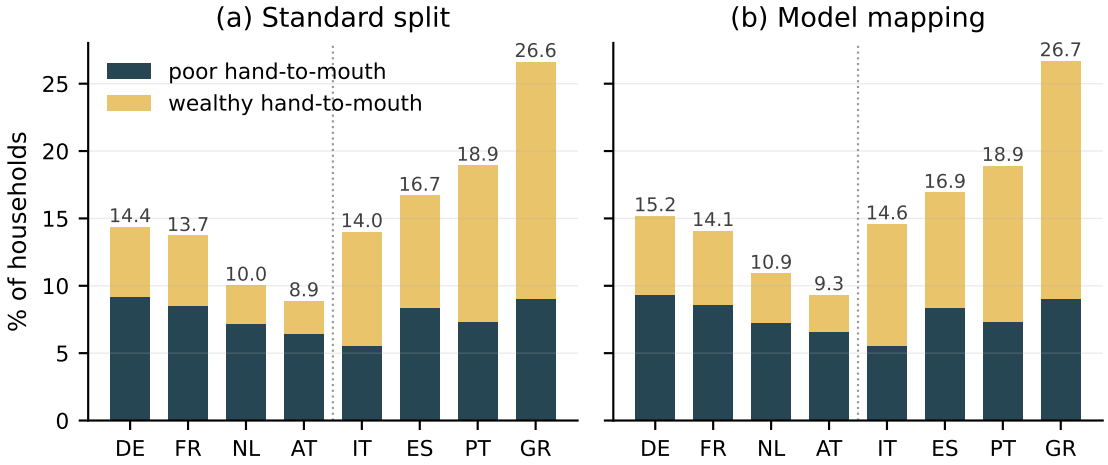


Figure 3: Hand-to-mouth shares, by country and by definition. Panel (a) uses the standard split, in which liquid wealth is deposits plus securities. Panel (b) uses the model's own mapping, in which the liquid account contains deposits only. Each bar splits into poor hand-to-mouth households, those with no wealth to speak of, and wealthy hand-to-mouth households, those with illiquid wealth but no liquid buffer. The total is printed above each bar. Source: HFCS Wave 4.

Three differences stand out. Households in the higher-debt countries hold more wealth relative to income, participate in securities markets less often, and are hand to mouth more frequently. But the higher hand-to-mouth incidence is *entirely* a wealthy hand-to-mouth phenomenon: on the model's mapping the wealthy share is 9.5 per cent against 5.3, while the poor share runs the other way, 6.9 against 8.7. These are households with substantial illiquid wealth and no liquid buffer, not households without assets.

In the portfolio evidence of Figure 4, securities participation and average holdings move in opposite directions. Securities participation is lower in the higher-debt countries, 17.1 per cent against 22.7. But securities holdings are *larger*: averaged over all households, holders and non-holders alike, they are 0.55 of income against 0.46, and the deposit share of financial assets is slightly lower, at 66.3 per cent against 68.3. Fewer households hold securities, and since a higher average is spread over fewer holders, those who do hold considerably more. Together with thin liquid buffers, this is the wealthy hand-to-mouth pattern that the liquidity wedge of Section 4.2 is calibrated to.

2.4 Fiscal structure

The fourth set of differences is fiscal. The lower-debt countries spend more through government consumption, 22.4 per cent of output against 19.2, and pay slightly more in non-pension cash benefits. Their labour-tax schedules are markedly more progressive, with an HSV progressivity parameter of 0.092 against 0.032, while their average implicit tax rate on labour is *lower*, 36.6 per cent against 40.4. Consumption and capital tax rates



Figure 4: Financial portfolio composition. Panel (a) is the split of total financial assets between deposits and directly held securities, in per cent. Panel (b) is each component as a ratio to annual net income. The dotted rule separates the two groups. Source: HFCS Wave 4.

differ by little enough that Section 5 finds no role for the first and a modest one for the second.

2.5 Two groups

The model has two countries, so the eight must be grouped. We group them on the averages documented above, principally sovereign debt and the household moments, which puts Germany, France, the Netherlands and Austria on one side and Italy, Spain, Portugal and Greece on the other. This literature sometimes calls these groups the core and the periphery. We do not use those terms, because they carry an argument about which country belongs where that is separate from the one this paper is making, and because the grouping here is a decision about observed averages rather than about geography. We label the two groups H and F, following the model's own notation, and use those labels throughout.

Table 1 reports the group averages. They are the inputs to the calibration of Section 4, and the differences in the final column are what Section 5 decomposes.

The sovereign differences in the table are large, and they run in the same direction: F owes more, and a larger part of its debt is held by private investors in the other group. Netting the rest of the world and the Eurosystem from each stock, the other group's investors hold 12.5 per cent of F's debt against 3.3 per cent of H's, so more of what F pays in debt service accrues to households the model places in H. The rest-of-world share runs the other way, with H the more exposed, which matters for a different reason (Section 5.3). The household differences are smaller and do not all point the same way, and the one that matters for transmission is the wealthy hand-to-mouth gap rather than the wealth level or the participation rate. Section 5.3 shows that the first set accounts for essentially all of the asymmetry in transmission and that the second adds to it.

	H	F	difference
Sovereign			
Debt securities, % of annual GDP	67.4	106.8	+39.4
Own-resident share of private holdings	0.906	0.766	−0.140
Rest-of-world share	0.300	0.107	−0.193
Eurosystem share	0.345	0.360	+0.015
Households			
Net wealth / annual income	8.36	10.37	+2.01
Deposits / annual income	0.98	1.02	+0.05
Securities / annual income	0.46	0.55	+0.09
Securities participation, %	22.7	17.1	−5.6
Hand-to-mouth, % (model mapping)	14.0	16.3	+2.3
of which wealthy	5.3	9.5	+4.2
of which poor	8.7	6.9	−1.8
Fiscal			
Government consumption / GDP, %	22.4	19.2	−3.2
Cash benefits / GDP, %	5.2	4.6	−0.6
Consumption tax	0.192	0.180	−0.012
Labour-tax progressivity	0.092	0.032	−0.060
Implicit labour tax rate	0.366	0.404	+0.038
Capital tax	0.356	0.298	−0.058
Share of the two groups' GDP, %	68.4	31.6	

Table 1: Group averages: the differences the model is calibrated to. H = DE, FR, NL, AT; F = IT, ES, PT, GR. Group values are 2021 GDP-weighted country averages. Sovereign holder shares are 2021Q4 on the debt securities perimeter. Household moments are HFCS Wave 4, with the hand-to-mouth share on the model's mapping, in which the liquid account is deposits only, and the asset ratios means over all households, holders and non-holders alike. Government consumption and cash benefits are 2015–24 averages; the consumption and capital tax rates are 2023 implicit tax rates, and progressivity is an estimate on labour income.

3 A Two-Country HANK Model of a Monetary Union

Time is discrete and quarterly. The union consists of two countries $j \in \{H, F\}$ that share a nominal interest rate, trade in goods, and hold claims on each other's governments. Each country is populated by a continuum of heterogeneous households, intermediate and final goods firms, capital producers, unions, a financial intermediary, and a fiscal authority; monetary policy is common. H has population share s and F $1 - s$. Every country quantity is measured per capita of its own country, and output per capita is normalised to one in steady state, so s is also H's share of union GDP. Every country-level primitive (technology, nominal rigidities, the tax system, debt level and maturity, household wealth and its composition, access to securities markets) carries a j index and may differ across members. The model is written once, at the level of the union; the one-country economy is the special case in which the two blocks are identical. Appendix A states the household, firm, wage-setting and intermediary problems and the steady state in full.

Notation and conventions

Throughout, $j \in \{H, F\}$ indexes the two member countries, grouped as in Section 2.5. Time is quarterly. Within each country, quantities produced by firms are measured in units of the domestic good; quantities held or consumed by households, the government's balance sheet, and the intermediary's book are measured in units of the domestic consumption basket. The relative price $p^c_{j,t}$ (the CPI in units of the domestic good, derived from the consumption basket in Section 3.8) converts between the two, and every equation states its units. Four conversions recur. With $w_{j,t}$ the producer-unit real wage, $d_{j,t}$ producer-unit dividends, $\pi_{j,t}$ producer-price inflation, and $r_{j,t}$ the CPI-real ex-ante rate,

$$w^c_{j,t} = \frac{w_{j,t}}{p^c_{j,t}}, \qquad d^c_{j,t} = \frac{d_{j,t}}{p^c_{j,t}}, \qquad 1 + \pi^c_{j,t} = (1 + \pi_{j,t}) \frac{p^c_{j,t}}{p^c_{j,t-1}}, \qquad 1 + r^p_{j,t} = (1 + r_{j,t}) \frac{p^c_{j,t+1}}{p^c_{j,t}}, \tag{1}$$

define the CPI-real wage and dividend, CPI inflation, and the producer real rate. When the two blocks are identical, $p^c_{j,t} = 1$ and every conversion is the identity. The symbols are defined where they are introduced.

3.1 Households

A continuum of households $i \in [0, 1]$ in country j chooses consumption of the domestic basket $c_{i,t}$, liquid holdings $b_{i,t} \geq 0$ (deposits), and illiquid holdings $a_{i,t} \geq 0$ (a claim on the intermediary's securities book, Section 3.4). Preferences are

$$u(c) = \frac{c^{1-1/\sigma}}{1 - 1/\sigma}, \qquad v(n) = \bar{\varphi}_j \frac{n^{1+1/\varphi}}{1 + 1/\varphi}, \tag{2}$$

with σ the elasticity of intertemporal substitution, φ the Frisch elasticity, and $\bar{\varphi}_j$ a country-specific scale on the disutility of work, set in steady state so that hours are normalised. *The baseline sets* $\sigma = 1$: under log utility the precautionary motive is weaker, more wealth sits near the liquid constraint, and the calibrated marginal propensity to consume roughly doubles relative to $\sigma = 0.5$ at the same wealth targets. Labour supply is delegated to unions (Section 3.3), so hours are not a household choice.

Portfolio adjustment. Following Auclert et al. (2024), rebalancing between the two accounts is subject to a Calvo lottery: household i can rebalance with i.i.d. probability χ_j

each period. A non-adjuster's illiquid account rolls over mechanically, $a_{i,t} = (1+r^a_{j,t})\, a_{i,t-1}$, and only the liquid account is active:

$$(1+\tau^c_{j,t})\, c_{i,t} + b_{i,t} = e_{i,t} + (1+r^b_{j,t})\, b_{i,t-1} + \mathrm{Tra}_{j,t}. \tag{3}$$

An adjuster chooses both accounts freely:

$$(1+\tau^c_{j,t})\, c_{i,t} + a_{i,t} + b_{i,t} = e_{i,t} + (1+r^a_{j,t})\, a_{i,t-1} + (1+r^b_{j,t})\, b_{i,t-1} + \mathrm{Tra}_{j,t}. \tag{4}$$

Consumption is taxed at $\tau^c_{j,t}$, $\mathrm{Tra}_{j,t}$ is a uniform lump-sum transfer, and after-tax labour earnings follow the progressive schedule of Heathcote et al. (2017),

$$e_{i,t} = \lambda_{j,t} \left(w^c_{j,t}\, N_{j,t}\, z_{i,t}\right)^{1-\vartheta_j}, \tag{5}$$

with $\lambda_{j,t}$ the level of the schedule, $\vartheta_j \in [0,1)$ its progressivity, $w^c_{j,t}$ the CPI-real wage of (1), and $z_{i,t}$ idiosyncratic productivity. The flat-tax case is nested at $\vartheta_j = 0$, $\lambda_{j,t} = 1-\tau_{j,t}$. Productivity follows an AR(1) in logs with persistence $\rho_{e,j}$ and stationary standard deviation $\sigma_{e,j}$, discretised by Rouwenhorst (1995) with the stationary mean normalised to one.

Return timing. $r^a_{j,t}$ and $r^b_{j,t}$ are the returns realised at t on positions carried into t. Deposits are a nominal claim on the intermediary, remunerated at the realised real return on the domestic government bond net of the intermediation cost:

$$r^b_{j,t} = r^g_{j,t} - \xi_j, \tag{6}$$

where $r^g_{j,t}$ is the realised holding return on government j's debt, defined in (20). An inflation surprise is therefore borne at once, as for any claim written in euros. With one-period debt, $1+r^g_{j,t} = (1+i_{t-1})/(1+\pi^c_{j,t})$, and a policy-rate change reaches deposit income from the following quarter. With the long debt of the calibrated union (Section 4), the realised return also carries the bond's price change, so deposits share on impact in the revaluation of the sovereign book that funds them and earn the higher return thereafter. Appendix E.4 reports the results with deposits written on the short nominal rate instead, which loads the whole revaluation of long debt onto the illiquid account. The illiquid return $r^a_{j,t}$ is determined residually by the intermediary's budget (Section 3.4).

Discount factor. Households share a single discount factor β_j, calibrated jointly with the rebalancing probability χ_j and the liquidity wedge ξ_j to the country's net worth, deposits and wealthy hand-to-mouth share (Section 4.2). The average marginal propensity to consume is then an output of the calibration rather than a target, reported in Section 4.3.

Aggregation. Let $\mathcal{D}_{j,t}$ denote the distribution over adjustment status, productivity, and the two asset holdings. The block supplies aggregate consumption $C_{j,t}$, the account stocks $A^{\mathrm{liq}}_{j,t}$ and $A^{\mathrm{illiq}}_{j,t}$ (total financial wealth $A_{j,t} = A^{\mathrm{liq}}_{j,t} + A^{\mathrm{illiq}}_{j,t}$), and the progressivity-weighted marginal utility that wage setting requires,

$$\mathrm{UCE}_{j,t} = \int z_{i,t}^{1-\vartheta_j}\, c_{i,t}^{-1/\sigma}\, \mathrm{d}\mathcal{D}_{j,t}. \tag{7}$$

It also reports hand-to-mouth shares: a household is hand-to-mouth at the liquid floor $b_{i,t} = 0$, *poor* hand-to-mouth if additionally at the illiquid floor, *wealthy* hand-to-mouth otherwise. It reports, finally, the *participation* share (households holding the illiquid claim,

$a_{i,t} > 0$), the model counterpart of the share of households with positive net worth net of deposits in the HFCS; it is not securities participation, since the account maps to all non-deposit wealth. Participation is never a calibration target: χ_j is the frequency at which the illiquid account can be rebalanced, not the share of households that hold it, and the model has no separate participation margin; the share is an outcome of the portfolio problem and an overidentifying check in Section 4, like the total hand-to-mouth share.

3.2 Firms

The production side is entirely in units of the domestic good; all discounting uses the producer real rate $r^p_{j,t}$ of (1).

Production and marginal cost. Intermediate firms operate a Cobb–Douglas technology in capital and labour, with capital share α_j and total factor productivity Z_j,

$$Y_{j,t} = Z_j \, K_{j,t-1}^{\alpha_j} \, N_{j,t}^{1-\alpha_j}, \qquad mc_{j,t} = \frac{w_{j,t} \, N_{j,t}}{(1-\alpha_j) \, Y_{j,t}}. \tag{8}$$

Investment. Capital depreciates at rate δ and accumulates subject to a quadratic adjustment cost whose curvature is ε^I, equivalently the elasticity of investment to Tobin's Q. Total investment outlay is net accumulation plus the adjustment cost,

$$I^{\text{tot}}_{j,t} = I_{j,t} + \psi^K_{j,t}, \qquad I_{j,t} = K_{j,t} - (1-\delta) K_{j,t-1}, \qquad \psi^K_{j,t} = \frac{K_{j,t-1}}{2\delta\varepsilon^I} \left(\frac{K_{j,t}}{K_{j,t-1}} - 1 \right)^2, \tag{9}$$

and the capital producer's first-order condition defines Tobin's Q as the shadow price of installed capital,

$$Q_{j,t} = 1 + \frac{1}{\delta\varepsilon^I} \left(\frac{K_{j,t}}{K_{j,t-1}} - 1 \right). \tag{10}$$

The intermediate firm's capital choice yields the capital Euler equation

$$(1 + r^p_{j,t}) \, Q_{j,t} = \alpha_j \, \frac{Y_{j,t+1}}{K_{j,t}} \, mc_{j,t+1} \; - \; \frac{I^{\text{tot}}_{j,t+1}}{K_{j,t}} \; + \; \frac{K_{j,t+1}}{K_{j,t}} \, Q_{j,t+1}, \tag{11}$$

the marginal product of capital, minus next period's investment outlay per unit of installed capital, plus the resale value. Here $r^p_{j,t}$ is the rate between t and $t+1$ (Section 3.7), and the same dating applies to price setting and equity below. At steady state, $Q = 1$ and $r = \alpha \, mc \, Y/K - \delta$.

Dividends. Gross dividends are output net of labour costs, net investment, the capital adjustment cost, and the Rotemberg (1982) price adjustment cost; profits are taxed at the country's capital income tax rate τ^K_j:

$$d^{\text{full}}_{j,t} = Y_{j,t} - w_{j,t} N_{j,t} - I_{j,t} - \psi^K_{j,t} - \psi^P_{j,t}, \qquad d_{j,t} = (1 - \tau^K_j) \, d^{\text{full}}_{j,t}, \tag{12}$$

with capital-tax revenue $\tau^K_j \, d^{\text{full}}_{j,t}$ accruing to the government. The price adjustment cost is quadratic in *log* price changes, with μ^p_j the gross price markup and κ^p the slope of the price Phillips curve it implies,

$$\psi^P_{j,t} = \frac{\mu^p_j}{\mu^p_j - 1} \, \frac{1}{2\kappa^p} \left[\ln(1 + \pi_{j,t}) \right]^2 Y_{j,t}, \tag{13}$$

which makes the resulting Phillips curve exact in logs.

Price Phillips curve. In symmetric equilibrium the firm's first-order condition gives the exact nonlinear New Keynesian Phillips curve in producer-price inflation,

$$\ln(1+\pi_{j,t}) = \kappa^p \left(mc_{j,t} - \frac{1}{\mu_j^p} \right) + \frac{1}{1+r_{j,t}^p} \frac{Y_{j,t+1}}{Y_{j,t}} \ln(1+\pi_{j,t+1}). \tag{14}$$

At steady state $\pi_j = 0$ and $mc_j = 1/\mu_j^p$.

3.3 Wage setting

Labour supply is delegated to unions, which differentiate raw labour into varieties and ration hours equally across households, so that earnings scale with idiosyncratic productivity (Auclert et al., 2024; Hagedorn et al., 2019). This removes the individual labour-supply margin: it decouples marginal propensities to consume from marginal propensities to earn (with flexible wages the two cannot simultaneously match the micro evidence, Auclert et al., 2023) and avoids the countercyclical profits that sticky-price/flexible-wage HANK models generate (Broer et al., 2023). The cost is that households cannot substitute between consumption and leisure; the labour margin is entirely passive.

Unions set nominal wages subject to adjustment costs quadratic in log wage changes, where nominal wage inflation is

$$1 + \pi_{j,t}^w = (1+\pi_{j,t}) \frac{w_{j,t}}{w_{j,t-1}}. \tag{15}$$

The resulting exact nonlinear wage Phillips curve, with slope κ^w, is

$$\ln(1+\pi_{j,t}^w) = \kappa^w \left[\bar{\varphi}_j \, N_{j,t}^{1+1/\varphi} - \frac{1-\vartheta_j}{1+\tau_{j,t}^c} \frac{\lambda_{j,t} \left(w_{j,t}^c N_{j,t} \right)^{1-\vartheta_j} \mathcal{U}_{j,t}}{\mu^w} \right] + \beta_j \ln(1+\pi_{j,t+1}^w), \tag{16}$$

in which the wage markup μ^w divides the income term, as in Auclert et al. (2024). The bracket is marginal disutility minus the marginal after-tax value of work: progressivity enters through the marginal retention rate $1-\vartheta_j$, the consumption tax through $1/(1+\tau_{j,t}^c)$, and income is valued at the CPI-real wage. *The baseline sets* $\mathcal{U}_{j,t} = C_{j,t}^{-1/\sigma}$, the marginal utility of aggregate consumption, the representative-agent convention of Auclert et al. (2025). The alternative, the union maximising average worker welfare, replaces it with $\mathrm{UCE}_{j,t}$ of (7) (Auclert et al., 2024). The choice matters: UCE inherits individual-level curvature and is dominated by households near the liquid constraint, so its elasticity to aggregate consumption is several times the representative-agent value, transmitting demand contractions into large real-wage movements. We adopt the aggregate convention. At steady state with $\pi_j^w = 0$, the disutility scale is

$$\bar{\varphi}_j = \frac{1-\vartheta_j}{1+\tau_j^c} \frac{\lambda_j \left(w_j^c N_j \right)^{1-\vartheta_j} \mathcal{U}_j}{\mu^w \, N_j^{1+1/\varphi}}. \tag{17}$$

3.4 Financial intermediary

A representative, risk-neutral, perfectly competitive intermediary in each country holds a bond portfolio $B_{j,t}^{\text{book}}$ and the domestic equity claim, whose ex-dividend value is $J_{j,t}$, and issues the two household claims: liquid deposits, at a real resource cost ξ_j per unit, and the illiquid claim. From the household's perspective the illiquid account is a single composite

buffer; the underlying bond/equity split is the intermediary's. The bond portfolio is the part of the union's sovereign debt that country-j residents hold: a set of legs $B_{jk,t}$, one per issuer k, with $B^{\text{book}}_{j,t} = \sum_k B_{jk,t}$. Under full home ownership it is the domestic sovereign stock, $B^{\text{book}}_{j,t} = B^g_{j,t}$; its composition under cross-border and outside holdings is set out in Section 3.9. All objects in this subsection are in CPI units.

Zero profit and the residual claimant. The two account returns are paid out of one book, so they cannot both be posited: perfect competition drives profits to zero period by period,

$$\sum_k (1 + r^g_{jk,t})\, B_{jk,t-1} + d^c_{j,t} + J_{j,t} = (1 + r^b_{j,t})\, A^{\text{liq}}_{j,t-1} + (1 + r^a_{j,t})\, A^{\text{illiq}}_{j,t-1} + \xi_j\, A^{\text{liq}}_{j,t-1}, \quad (18)$$

where $r^g_{jk,t}$ is the realised return on issuer k's debt in country j's CPI units, net of any holding cost the holder bears (Section 3.6), and $r^g_{jj,t} = r^g_{j,t}$. Every euro the book earns is thus paid to a household or consumed as the intermediation cost, so one market-clearing condition is redundant by Walras' law, and the condition determines exactly one return. We posit the deposit rate, (6), and let the illiquid account be the *residual claimant*: writing the book return on actual lagged market values,

$$r^{\text{book}}_{j,t} = \frac{\sum_k (1 + r^g_{jk,t})\, B_{jk,t-1} + d^c_{j,t} + J_{j,t}}{B^{\text{book}}_{j,t-1} + J_{j,t-1}} - 1, \qquad r^a_{j,t} = \frac{r^{\text{book}}_{j,t} - \varsigma_j\, r^g_{j,t}}{1 - \varsigma_j}, \quad (19)$$

where ς_j is the steady-state liquid share of household wealth.[2] Since the deposit leg plus the intermediation cost equals $(1 + r^g_{j,t})A^{\text{liq}}_{j,t-1}$ under (6), the residual is taken against r^g, and at steady state $r^{\text{book}}_j = r^a_j = r_j$. The factor $1/(1 - \varsigma_j)$ is the leverage of a claim junior to deposits: the illiquid account absorbs the entire revaluation of the book net of what the deposit leg is contractually owed. The no-arbitrage identity $r^a_{j,t} - r^b_{j,t} = \xi_j$ holds at all $t \geq 1$.

Government bonds. Government j issues nominal bonds whose coupons decay at the country-specific rate ρ^b_j (Woodford, 2001), giving a Macaulay duration of $(1+r)/(1+r-\rho^b_j)$ quarters at the steady-state rate r, which is $1/(1 - \rho^b_j)$ in the limit of a zero rate. The nominal price and the realised real holding return are

$$P^g_{j,t} = \frac{1 + \rho^b_j\, P^g_{j,t+1}}{1 + i_t}, \qquad 1 + r^g_{j,t} = \frac{1 + \rho^b_j\, P^g_{j,t}}{P^g_{j,t-1}} \frac{1}{1 + \pi^c_{j,t}}, \quad (20)$$

the return deflated by the *holder's* CPI. *The symmetric benchmark of Appendix D.1 sets* $\rho^b_j = 0$: one-period debt, $P^g_{j,t} = 1/(1 + i_t)$ and $1 + r^g_{j,t} = (1 + i_{t-1})/(1 + \pi^c_{j,t})$, with the duration channel switched off. The calibrated union of Section 4 gives both countries a common duration of 8.75 years, so there the realised return carries the bond's price change; country-specific maturity is retained as a parameter. Section 3.6 generalises the pricing equation to a sovereign spread.

[2] Holding ς_j at its steady-state value, rather than at the realised liquid share, changes the solution only through a product of two deviations from steady state, which is zero at first order, the order at which the model is solved.

Equity. The ex-dividend equity value satisfies

$$J_{j,t} = \frac{d^c_{j,t+1} + J_{j,t+1}}{1 + r_{j,t}}, \tag{21}$$

discounted at the country's ex-ante CPI-real rate: dated t, the rate between t and $t+1$, consistent with the dating of (11).

Gross legs. Each leg of the book earns its issuer's payoff deflated by its holder's CPI, net of any holding cost the holder bears. With one-period bonds and no spread every leg pays the same nominal rate i_{t-1}, so $r^g_{jk,t} = r^g_{j,t}$ for every issuer and the book earns a single return; country-specific duration or a spread separates the legs, which is why the book is kept by holder and issuer.

3.5 Government

Each government levies the progressive labour income tax of (5), a consumption tax $\tau^c_{j,t}$, and the capital income tax τ^K_j, and runs spending $G_{j,t}$ (in domestic goods) and the lump-sum transfer $\mathrm{Tra}_{j,t}$. The tax *structure* $(\vartheta_j, \tau^c_j, \tau^K_j)$ is a set of country parameters, so economies with different tax mixes can be specified and compared directly, following Bouabdallah et al. (2025), who build the same fiscal block into a two-asset heterogeneous-agent model and show that the financing instrument, the debt level and progressivity each change the size of fiscal multipliers. The labour-tax *level* λ_j is the calibrated outcome that closes the steady-state budget given the steady-state transfer, which is set to the data's cash benefits in the calibrated union (Section 4) and to zero in the symmetric benchmark of Appendix D.1, the standard normalisation for the tax function (5); in dynamics the transfer is the residual instrument, absorbing budget fluctuations around its steady-state value under the non-fiscal closures below.[3]

Aggregate labour-tax revenue is the wedge between the wage bill and after-tax earnings; since after-tax income (5) depends on the household only through $z_{i,t}$, it aggregates in closed form,

$$T^L_{j,t} = w^c_{j,t} N_{j,t} - \lambda_{j,t} \left(w^c_{j,t} N_{j,t}\right)^{1-\vartheta_j} \mathbb{E}\left[z^{1-\vartheta_j}\right], \tag{22}$$

with the moment taken under the stationary distribution of z. Writing $G^c_{j,t} = G_{j,t}/p^c_{j,t}$ for the CPI-unit cost of spending, the flow budget in CPI units is

$$B^g_{j,t} = (1 + r^g_{j,t} + \zeta_{j,t-1})\, B^g_{j,t-1} + G^c_{j,t} + \mathrm{Tra}_{j,t} - T^L_{j,t} - \tau^c_{j,t}\, C_{j,t} - \frac{\tau^K_j\, d^c_{j,t}}{1 - \tau^K_j}. \tag{23}$$

Debt service is written on the realised return path of (20), so issuer and holder revalue together, and the issuer additionally bears the spread wedge $\zeta_{j,t-1}$, which the holder does *not* receive: the issuer pays gross, the holder nets, and the difference is the deadweight cost of Section 3.6 (zero for H, and zero everywhere at the baseline). The budget holds under both closures below; what differs is the object it pins down.

[3] With $\tau^c_j \neq 0$ consumption-tax revenue depends on consumption, so the budget-closing level is solved jointly with the household problem, and in dynamics the transfer is an equilibrium unknown with the budget as an equilibrium condition.

Default closure. Sovereign debt is held constant at its steady-state real market value, $B^g_{j,t} = \bar{B}^g_j$, with the face quantity floating against the bond price, $B^{\text{face}}_{j,t} = \bar{B}^g_j / P^g_{j,t}$; the tax parameters are held at their calibrated values and $\text{Tra}_{j,t}$ absorbs the residual, the heterogeneous-agent analogue of Ricardian lump-sum financing. Because the real market value of debt is held fixed, a change in the realised return on outstanding debt, revaluation included, changes the transfer in the same quarter. A simple benchmark clarifies the mechanism. With a single household in each country, higher debt service on the debt held at home would lower the household's transfer and raise its income on the intermediary's book by the same amount. At given paths of earnings, dividends and real rates this direct effect would cancel, and the debt would reach consumption only through the country's net sovereign position against the other member and the rest of the world and through the spread. With heterogeneous households the two flows have different consumption effects, because they reach households that differ in wealth, liquidity and the ability to adjust spending: the transfer is the same for every household, and those with little liquid wealth spend a change in it at once, while the income on the book accrues to the holders of deposits and illiquid wealth, who spend a smaller part of it. A two-agent version makes this exact. If a share of households spends its income and the rest spread their receipts over time, the cancellation fails on impact in proportion to that share and to the part of the payment that savers do not spend at once, and moving the creditors across the border shifts the savers' offset between countries without changing the union's total through this channel (Appendix B). This is the closure under which all non-fiscal shocks are run.

Fiscal closure. For government-spending experiments, debt accumulates according to (23) and a feedback rule on lagged debt restores solvency, adjusting either the labour-tax level or the transfer:

$$\lambda_{j,t} = \lambda_j - \phi^B \, \frac{B^{\text{face}}_{j,t-1} - B^{\text{face}}_j}{\left(w^c_{j,t} N_{j,t}\right)^{1-\vartheta_j} \mathbb{E}\left[z^{1-\vartheta_j}\right]}, \qquad \text{or} \qquad \text{Tra}_{j,t} = \text{Tra}_j - \phi^B \left(B^{\text{face}}_{j,t-1} - B^{\text{face}}_j\right). \tag{24}$$

The rules respond to the *face* quantity outstanding, $B^{\text{face}}_{j,t} = B^g_{j,t}/P^g_{j,t}$, not to market value: responding to market value would add a fiscal response to valuation changes. The tax rule's denominator normalises it so that, via (22), ϕ^B is the *primary-surplus response per unit of debt gap* (identical to the transfer rule's by construction), so the tax-versus-transfer contrast is a like-for-like comparison of instruments, not of rule strengths. The fiscal rule fires only in fiscal experiments; it is never combined with a non-fiscal shock.

3.6 Sovereign spreads

Motivated by Section 2.1, F's sovereign borrowing costs respond to the fiscal position while H's do not. Writing $\tilde{Y}_{j,t} = Y_{j,t}/p^c_{j,t}$ for CPI-valued GDP, the spread is debt-elastic and zero in steady state, with $\psi^\zeta \geq 0$ the elasticity of the spread to the debt-to-GDP gap,

$$\zeta_{F,t} = \psi^\zeta \left(\frac{B^g_{F,t}}{\tilde{Y}_{F,t}} - \frac{\bar{B}^g_F}{\bar{\tilde{Y}}_F} \right), \qquad \psi^\zeta \geq 0, \quad \zeta_{H,t} = 0, \tag{25}$$

and enters the pricing of F debt: the government borrows at the spread-inclusive yield,

$$P^g_{F,t} = \frac{1 + \rho^b_F \, P^g_{F,t+1}}{1 + i_t + \zeta_{F,t}}. \tag{26}$$

The wedge is a *deadweight cost* in the tradition of debt-elastic premia: the holder pays a proportional cost $\zeta_{F,t-1}$ per unit of F debt carried into t (an intermediation or expected-loss cost, a real resource entering goods-market clearing alongside the intermediation cost ξ), so the holder's *net* realised return is

$$1 + r^g_{F,t} = \frac{1 + \rho^b_F \, P^g_{F,t}}{P^g_{F,t-1}} \frac{1}{1 + \pi^c_{F,t}} \; - \; \zeta_{F,t-1}, \tag{27}$$

whose expectation equals the union rate. No asset dominates any other, the no-arbitrage identity $r^a_j - r^b_j = \xi_j$ of Section 3.4 is preserved, and the calibration is untouched since $\zeta_F = 0$ in steady state. The spread is thus a purely *fiscal* cost to F: it operates through the rollover yield. In a fiscal expansion, a higher debt ratio widens the spread, raising the yield at which maturing debt is refinanced, which raises debt service and feeds back into the debt ratio: a loop the fiscal rule (24) must be responsive enough to stabilise, and which turns over faster the shorter F's maturity. The mark-to-market revaluation of *outstanding* long bonds moves the other way on impact (a spread rise lowers $P^g_{F,t}$, a valuation gain to the issuer), so the amplification is a flow, rollover-cost force. An alternative specification, in which the wedge accrues to holders as compensation for a convex cost on cross-border positions and F debt dynamics thereby feed foreign holders' returns, is left as an extension.

3.7 Monetary policy

A single central bank sets the union nominal rate on size-weighted union CPI inflation, with r^*_t the intercept of the rule, the point at which the monetary shock enters, and ϕ_π the response coefficient,

$$i_t = r^*_t + \phi_\pi \left[\, s \, \pi^c_{H,t} + (1 - s) \, \pi^c_{F,t} \, \right], \tag{28}$$

and a union monetary shock is an AR(1) innovation to r^*_t; the nominal rate then follows from the rule. Note that at equal weights the rule responds only to the union average, so the cross-country *difference* in inflation receives no direct policy response; asymmetric transmission is disciplined by relative prices and trade, not by the policy rule.

The nominal rate is the only common price. Real rates are country-specific, because each country deflates by its own consumption basket: the ex-ante Fisher equation holds per country,

$$1 + i_t = (1 + r_{j,t}) \, (1 + \pi^c_{j,t+1}), \qquad j \in \{H, F\}. \tag{29}$$

A single common real rate would be consistent only with a constant terms of trade. Taking the ratio of the two Fisher equations gives the union analogue of uncovered interest parity, in which $\mathcal{Q}_t$ is the real exchange rate, the relative price of the F consumption basket in units of the H basket, constructed in (33) below,

$$\frac{1 + r_{H,t}}{1 + r_{F,t}} = \frac{1 + \pi^c_{F,t+1}}{1 + \pi^c_{H,t+1}} = \frac{\mathcal{Q}_{t+1}}{\mathcal{Q}_t}, \tag{30}$$

which follows from the two Fisher equations rather than being imposed.

3.8 Trade and the external position

Consumption baskets and relative prices. Households in country j consume a CES basket of home and foreign goods,

$$c_{j,t} = \left[\gamma_j^{1/\eta} \left(c^{\text{dom}}_{j,t} \right)^{\frac{\eta-1}{\eta}} + (1 - \gamma_j)^{1/\eta} \left(c^{\text{imp}}_{j,t} \right)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \tag{31}$$

with home bias $\gamma_j > 1/2$ and trade elasticity η. Home bias is tied to country size through a common degree of openness θ,

$$\gamma_H = 1 - (1-s)\,\theta, \qquad \gamma_F = 1 - s\,\theta, \tag{32}$$

so each country's import share is proportional to its partner's size (Sutherland, 2005; De Paoli, 2009): the smaller country is the more open, trade is balanced at a symmetric steady state of any size, and at equal size $\gamma_j = 1-\theta/2$. The nominal exchange rate is fixed at one and the law of one price holds for every variety, so real-exchange-rate movements arise purely from basket composition. With $S_t = P_t^F/P_t^H$ the terms of trade, the CPIs in units of the domestic good and the real exchange rate are

$$p^c_{H,t} = \left[\gamma_H + (1-\gamma_H)\,S_t^{1-\eta}\right]^{\frac{1}{1-\eta}}, \quad p^c_{F,t} = \left[\gamma_F + (1-\gamma_F)\,S_t^{-(1-\eta)}\right]^{\frac{1}{1-\eta}}, \quad \mathcal{Q}_t = \frac{S_t\,p^c_{F,t}}{p^c_{H,t}}. \tag{33}$$

Because S_t is a relative price *level*, its growth rate is the producer-price inflation differential by construction,

$$\frac{S_t}{S_{t-1}} = \frac{1+\pi_{F,t}}{1+\pi_{H,t}}, \tag{34}$$

an accounting identity imposed as an equilibrium condition; the allocative content of trade lies in the demand curves and goods-market clearing below.

Demand and goods-market clearing. CES demands for the home and imported good are, in each country's own quantity units,

$$c^{\text{dom}}_{H,t} = \gamma_H\,(p^c_{H,t})^{\eta}\,C_{H,t}, \qquad c^{\text{imp}}_{H,t} = (1-\gamma_H)\left(\tfrac{S_t}{p^c_{H,t}}\right)^{-\eta} C_{H,t}, \tag{35}$$

and symmetrically for F. Each country's output is absorbed by home and foreign demand for its good, investment, government spending, the adjustment and intermediation costs, and exports to the rest of the world (all in domestic-good units, per capita of the producing country):

$$\begin{aligned} Y_{H,t} = c^{\text{dom}}_{H,t} + \frac{1-s}{s}\,c^{\text{imp}}_{F,t} + I_{H,t} + \psi^K_{H,t} + \psi^P_{H,t} + G_{H,t} \\ + p^c_{H,t}\Big(\xi_H\,A^{\text{liq}}_{H,t-1} + \sum_k \zeta_{k,t-1}\,B_{Hk,t-1} + X^R_{H,t}\Big), \end{aligned} \tag{36}$$

with the analogous condition for F, where H's per-capita imports enter scaled by $s/(1-s)$. Foreign demand is converted from the importer's to the exporter's population by the size ratio. The sum is the deadweight cost of the sovereign spread (Section 3.6), a real resource borne by *holders*: each union holder's market carries the cost of its own legs of the book, and the rest of the world bears the cost on its holdings abroad. $X^R_{j,t}$ is the net export of country j's good to the rest of the world, the counterpart of the payments on the part of j's debt held outside the union (Section 3.9). Net exports of H to F, per H capita in H-good units, are $NX_{H,t} = \frac{1-s}{s}\,c^{\text{imp}}_{F,t} - S_t\,c^{\text{imp}}_{H,t}$.

3.9 Who holds the sovereign debt

Each government's debt is held by three groups: the issuer's own intermediary, the other member's intermediary, and the rest of the world. Write $B_{jk,t}$ for the country-j intermediary's holding of issuer-k debt, per capita of j in j's CPI units, and $B^R_{k,t}$ for the rest of the

world's holding of issuer-k debt, per capita of k in k's CPI units. The model carries one cross-border leg inside the union, H's holding of F debt, $B_{HF,t} = nfa_{H,t}$ (H's net foreign asset position against F), so the issuer adding-up conditions are

$$B^g_{H,t} = B_{HH,t} + B^R_{H,t}, \qquad B^g_{F,t} = B_{FF,t} + \frac{s}{1-s}\,\frac{B_{HF,t}}{\mathcal{Q}_t} + B^R_{F,t}, \tag{37}$$

and the two books are

$$B^{\text{book}}_{H,t} = B^g_{H,t} - B^R_{H,t} + nfa_{H,t}, \qquad B^{\text{book}}_{F,t} = B^g_{F,t} - B^R_{F,t} - \frac{s}{1-s}\,\frac{nfa_{H,t}}{\mathcal{Q}_t}, \tag{38}$$

each in its holder's CPI units; H's claim is converted to a per-capita position in F using the population ratio. The domestic legs are the residuals: under the fiscal closure a new euro of issuer-k debt is bought on impact by k's own intermediary except for the fraction $\phi^R\omega^R_k$ that the rest of the world, which holds the share ω^R_k of issuer k's debt and closes the gap to its target position at rate ϕ^R, takes within the quarter; the cross-border position moves only through the balance of payments, and the rest of the world closes the remaining gap to its target only gradually (below). The domestic intermediary is thus the residual buyer by convention; a general marginal-absorption rule is left as an extension. H's F leg earns the F payoff converted into H units, net of the deadweight cost H bears on it, $r^g_{HF,t} = r^{g\times}_t - \zeta_{F,t-1}$, with the gross cross return $r^{g\times}_t$ defined in (41).

The Eurosystem is part of the public sector. National central banks remit their income to their treasuries and the ECB distributes its profits by capital key, so the coupons paid on Eurosystem-held bonds do not leave the union's public sector; and the central banks financed their purchases by issuing reserves, public liabilities held by the union's banks. We therefore consolidate the Eurosystem with the governments: the consolidated public sector owes the gross stock of securities, part of it in the form of reserves, and B^g_k is kept at the gross level. Dropping the Eurosystem's bonds would understate the public debt that private portfolios hold. In the baseline the reserves are held by the issuer's own intermediary, so the Eurosystem share ω^E_k of issuer k's debt is part of B_{kk}; the shorter duration of the consolidated liability is a robustness row of Appendix E.4.

Steady-state shares. In steady state the rest of the world holds the share ω^R_k of issuer k's debt; the rest is held within the union, and κ_k is the share of its *privately held* union part that issuer k's own residents hold: the statistic the holdings data measure. With one cross-border leg, the book carries the *net* cross-border claim,

$$\overline{nfa}_H = \frac{1-s}{s}\,(1-\kappa_F)\big(1-\omega^R_F-\omega^E_F\big)\bar B^g_F\,\bar{\mathcal{Q}} - (1-\kappa_H)\big(1-\omega^R_H-\omega^E_H\big)\bar B^g_H, \qquad \bar B^R_k = \omega^R_k\,\bar B^g_k, \tag{39}$$

per H capita, which reproduces each country's total bond holdings and the net external position exactly; netting the two gross legs changes only the exposure to issuer-specific returns, and a second gross leg is left as an extension.

The rest of the world. The rest of the world is a buy-and-hold investor with target position $\omega^R_k B^g_{k,t}$. After receiving the realised payoff on its holding and consuming its steady-state yield on the target, its excess over the current target is

$$\Xi_{k,t} = \big(1 + r^g_{k,t} + \zeta_{k,t-1}\big)B^R_{k,t-1} - \bar r\,\omega^R_k B^g_{k,t-1} - \omega^R_k B^g_{k,t};$$

it consumes the fraction ϕ^R of the excess and reinvests the rest,

$$B^R_{k,t} = \omega^R_k B^g_{k,t} + \big(1-\phi^R\big)\,\Xi_{k,t}, \qquad X^R_{k,t} = \big(1 + r^g_{k,t} + \zeta_{k,t-1}\big) B^R_{k,t-1} - B^R_{k,t}. \tag{40}$$

The second equation is the union's balance of payments with the rest of the world, issuer by issuer: the payments on the foreign-held debt (at the issuer's gross return, since the foreign holder bears the deadweight cost abroad) are either reinvested or settled by net exports $X^R_{k,t}$ of the issuer's good, which enter (36). In steady state $\bar{X}^R_k = \bar{r}\,\omega^R_k \bar{B}^g_k$: the coupon on foreign-held debt is a net factor payment, and the issuer runs the matching trade surplus with the rest of the world. Away from it, the surprises on foreign-held debt (inflation revaluations of nominal claims, capital gains on long bonds, rate changes) accrue to the foreign position and return to the union's goods markets only gradually: to first order, at constant debt and no spread, $\mathrm{d}B^R_{k,t} = (1-\phi^R)\big[(1+\bar{r})\,\mathrm{d}B^R_{k,t-1} + \bar{B}^R_k\,\mathrm{d}r^g_{k,t}\big]$, stationary for $\phi^R > \bar{r}/(1+\bar{r})$; the baseline $\phi^R = 0.05$ gives a half-life of four and a half years. For the union's households the outside holder is therefore a leakage: the government services its whole stock, so under the default closure the transfer that absorbs a rise in debt service falls by the foreign holders' share as well, while their receipts leave the union's household sector. At $\phi^R = 1$ the rule collapses to a passive investor that holds $\omega^R_k B^g_{k,t}$ every period and takes its entire realised income, revaluations included, as exports within the quarter; that special case treats foreign bondholders as spending their euro income on union goods one-for-one, which is why the baseline does not adopt it. Under the default closure the target is constant; under the fiscal closure the rest of the world absorbs the share ω^R_k of new issuance gradually, the foreign-financing margin of Broner et al. (2022).

The balance of payments inside the union. The cross-border position evolves through the intra-union balance of payments, credited at the return on the claim it actually is, the F payoff, gross of the deadweight burn (which is domestic absorption, not a cross-border flow), converted into H units:

$$nfa_{H,t} = \big(1 + r^{g\times}_t\big)\, nfa_{H,t-1} + \frac{NX_{H,t}}{p^c_{H,t}}, \qquad 1 + r^{g\times}_t = \big(1 + r^g_{F,t} + \zeta_{F,t-1}\big)\,\frac{\mathcal{Q}_t}{\mathcal{Q}_{t-1}}. \tag{41}$$

The claim is nominal, so crediting a predetermined real rate would make it implicitly indexed; and at the baseline (one-period bonds, no spread) $r^{g\times}_t = r^g_{H,t}$ exactly, since the real-exchange-rate growth factor is the CPI-inflation ratio. A cross-border position is also what lets trade imbalances arise with sticky prices: under financial autarky, goods-market clearing would reduce to balanced trade while the terms of trade is already pinned by (34), and the model would be over-determined. The holder set also disciplines the steady state: the part of each stock that a country's households hold through their own intermediary is bounded by their wealth (Appendix A.5), and in the data roughly a quarter of F debt and a third of H debt sit with the other member's investors or the rest of the world (Section 4).

3.10 Equilibrium

Given paths for the shocks $\{r^*_t, Z_{j,t}, G_{j,t}\}$, an equilibrium is a set of aggregate paths such that households solve (3)–(5), firms and unions satisfy their optimality conditions, the government budgets (23) hold under the active closure, and goods and asset markets clear. The model is solved in the sequence space (Auclert et al., 2021), to first order in aggregate shocks and nonlinearly in households' idiosyncratic risk. A shock is an innovation that is unanticipated at date 0; from then on all agents have perfect foresight over aggregate paths, and by certainty equivalence these paths are the impulse responses of the first-order

solution with aggregate risk (Auclert et al., 2024, 2021). Under the default closure the system has ten unknowns and ten targets,

$$\begin{aligned} &\big\{Y_H, Y_F, w_H, w_F, S, r_H, r_F, nfa_H, \mathrm{Tra}_H, \mathrm{Tra}_F\big\} \\ \longleftrightarrow\ &\big\{(36)\times 2,\ (16)\times 2,\ (34),\ (29)\times 2,\ (41),\ (23)\times 2\big\}, \end{aligned} \tag{42}$$

with the price Phillips curves, production, capital, and pricing blocks solved out. The fiscal closure replaces the two transfers by the two debt stocks $\{B_H^g, B_F^g\}$ as unknowns, with the fiscal rule setting the instrument. The asset-market conditions

$$A_{j,t} = B_{j,t}^{\text{book}} + J_{j,t}, \qquad j \in \{H, F\}, \tag{43}$$

are redundant by Walras' law given (18) and the household and government budgets, and hold in every solution.

4 Baseline Calibration

This section calibrates the two country blocks to the groups H and F defined in Section 2. The guiding principle is that everything observable is imposed as a primitive. Tax rates, spending and transfer ratios, debt stocks, holder shares, country size and trade openness all come from data, and the parameters common to both countries take standard values. What the calibration solves for is the small set of household parameters that no dataset reports directly, together with the terms of trade and each government's consumption-tax revenue; the production parameters follow from the steady state.

4.1 Why the calibration is joint

In a monetary union whose members hold each other's sovereign debt, the two countries cannot be calibrated one at a time. The terms of trade determine the relative price at which each country's wealth target is expressed; the external position determines how much of one country's debt service accrues to the other; and each country's discount factor responds to both. Solving country H first and country F second would impose a terms of trade that the second solve then contradicts. The calibration is therefore a single system: nine unknowns and nine targets, solved simultaneously.

Table 2 collects the primitives. Two features of the sovereign block do most of the work in Section 5. The first is the debt level, where F carries roughly sixty per cent more than H relative to output. The second is ownership: κ_j is the share of the privately held, intra-union part of country j's debt that j's own residents hold, and the other member's residents hold the rest of that part. It is the margin that decides whether a rise in debt service returns to domestic households as income or leaves the country. Outside that part, a rest-of-world holder takes income out of the union altogether, and a consolidated Eurosystem position is netted against the issuer's own book.

Common parameters. Parameters that are not specific to a country take standard values and are the same in both. Utility is logarithmic, following Kase et al. (2025), and the Frisch elasticity is 0.5 (Chetty et al., 2011). The slopes of the price and wage Phillips curves are both 0.03 (Christiano et al., 2011; Hagedorn et al., 2019), and the wage markup is 1.1. Capital depreciates at 2 per cent a quarter, the elasticity of investment to Tobin's Q is 4 (Auclert et al., 2024), and the capital stock is 2.57 times annual GDP. Log productivity has a quarterly persistence of 0.978, about 0.91 at an annual frequency, and

		H	F
Sovereign			
Debt securities / annual GDP	B_j^g	0.674	1.068
Own-resident share of private holdings	κ_j	0.906	0.766
Rest-of-world share	ω_j^R	0.300	0.107
Eurosystem share (consolidated)	ω_j^E	0.345	0.360
Duration (years, common)	D	8.75	
Fiscal			
Government consumption / GDP	G_j/Y_j	0.224	0.192
Non-pension cash benefits / GDP	T_j^{ra}/Y_j	0.052	0.046
Consumption tax	τ_j^c	0.192	0.180
Labour-tax progressivity	ϑ_j	0.092	0.032
Capital tax	τ_j^K	0.356	0.298
Households (HFCS)			
Net worth / quarterly GDP	A_j	12.05	16.14
Deposits / quarterly GDP	A_j^ℓ	1.42	1.59
Wealthy hand-to-mouth share	wHtM_j	0.053	0.095
Union			
Country size (H share of GDP)	s	0.684	
Trade openness	θ	0.25	

Table 2: Calibration primitives, by country. Sovereign holder shares are 2021Q4, face value, on the debt securities perimeter, from ECB securities holdings statistics with the IMF CPIS bilateral key. The duration is set to the average residual maturity of government debt securities (ECB), which exceeds the Macaulay duration of a coupon bond; Appendix E.4 shortens it to 5.69 years. Government consumption and cash benefits are 2015–24 averages; the consumption and capital tax rates are 2023 implicit tax rates and progressivity an estimate on labour income. Household moments are GDP-weighted HFCS wave 4. Home bias is derived from openness and country size as $\gamma_H = 1 - (1 - s)\theta$ and $\gamma_F = 1 - s\theta$, so that a change in relative size moves both countries' openness consistently.

a stationary standard deviation of 0.92, on seven grid points. The policy rule responds to union inflation with $\phi_\pi = 1.5$ (Taylor, 1993), the steady-state real rate is 5 per cent a year, the trade elasticity is $\eta = 1.5$, and the fiscal rules of the spending experiments respond to debt with $\phi^B = 0.1$. F's spread rises by about 5 basis points a year per percentage point of debt securities to annual GDP ($\psi^\zeta = 0.0032$). This is 4 basis points per point of total debt, rescaled to the securities perimeter: above the linear estimate of about 3 basis points for euro-area spreads and below the nearly 6 estimated where debt exceeds 90 per cent of GDP (Pamies et al., 2021), and close to the 3 to 4 basis points that Laubach (2009) finds for US Treasury yields. Given these values and the targets of Table 2, the steady state implies each country's capital share, price markup and productivity level: the markups are 1.07 in H and 1.11 in F, and the capital shares 0.36 and 0.37 (Appendix A.5).

4.2 The household block

Three parameters per country are free, calibrated jointly to three moments. The discount factor β_j mainly sets the level of wealth, matched to net worth. The rebalancing probability χ_j governs how often a household can move resources between the two accounts, and so mainly how much wealth is held in liquid form, matched to deposits. The third is a liquidity wedge ξ_j, the gap between the return on the illiquid account and the return on deposits.

The wedge replaces a more common device. A standard way to match hand-to-mouth shares is to give households within a country different discount factors, so that the impatient ones hold nothing. That device does not suit this calibration: with a single discount factor the model already puts more households at the constraint in H than the data do, and Kaplan and Violante (2022) argue against reading hand-to-mouth behaviour as impatience. The wealthy hand-to-mouth are households with substantial illiquid wealth and no liquid buffer, and what puts them there is the return they give up by holding deposits rather than assets, not a low discount factor.

That is the margin ξ_j governs. Raising the return gap makes it more costly to hold deposits, which pushes households with illiquid wealth to run their liquid balances down to the constraint. It moves mainly the *wealthy* hand-to-mouth margin and leaves the poor margin, which is about having no wealth at all, almost unchanged. We therefore target the wealthy share and let the total hand-to-mouth share emerge as an output. The calibrated wedges are 1.1 per cent a year in H and 2.8 per cent in F, implying real deposit returns of 3.9 and 2.2 per cent.

The three remaining unknowns are the terms of trade and, one per country, the consumption-tax revenue that enters each government's budget. Table 3 reports the solution and its fit.

4.3 The model fit

The targeted moments are matched exactly. Table 3 also reports five moments the calibration was not asked to match.

The total hand-to-mouth share is the first. It is an output of the calibration, and it comes out at 15.7 per cent in H and 17.4 in F. Both the level and, more importantly, the ordering are reasonable, and the ordering matters because the exercises in Section 5 turn on F households being more exposed to income and transfer shocks than H households. Since the wealthy share is targeted, the untargeted content is in the poor share, which the wedge barely moves and which is close to its survey counterpart in both countries.

	model		data	
	H	F	H	F
Targeted				
Net worth / quarterly GDP	12.05	16.14	12.05	16.14
Deposits / quarterly GDP	1.42	1.59	1.42	1.59
Wealthy hand-to-mouth	0.053	0.095	0.053	0.095
Transfers / GDP	0.052	0.046	0.052	0.046
Untargeted				
Hand-to-mouth share	0.157	0.174	see note	
Poor hand-to-mouth	0.104	0.079	0.087	0.069
MPC, quarterly	0.16	0.17	0.15–0.25	
MPC, annualised	0.50	0.52	0.48, 0.52	
Gini, liquid account	0.685	0.713	0.691	0.773
Gini, illiquid account	0.674	0.668	0.759	0.683
Illiquid-account participation	0.879	0.911	0.671	0.811
Labour share	0.601	0.568	0.644	0.594
Debt service / GDP	0.034	0.054	0.012	0.032
Solved parameters				
Discount factor	0.9836 / 0.9840			
Rebalancing probability	0.338 / 0.285			
Liquidity wedge (annual)	1.1% / 2.8%			

Table 3: Calibration: targeted moments, and moments the model is not asked to match. The four targeted moments are matched to solver tolerance (5×10^{-11}). The hand-to-mouth share is reported on the model's own measure, the mass of households at the liquid constraint. The survey counterpart is a band around zero, reported in Table 1. The ordering, F above H, is the same on both measures and is not targeted. The quarterly MPC is the share of a one-off lump-sum transfer spent within the quarter, consumption tax included; its data range is the one Kaplan and Violante (2022) report for the empirical literature. The annualised MPC is $1 - (1 - \text{MPC}_q)^4$, the convention of Hagedorn et al. (2019), and its data entries are the estimates of Jappelli and Pistaferri (2014) and Fagereng et al. (2021). Participation is the share of households with positive net worth net of deposits, the model's illiquid account, in both model and data.

The marginal propensity to consume is the second. Households spend 16 per cent of a one-off transfer within the quarter in H and 17 per cent in F, inside the 15 to 25 per cent range that Kaplan and Violante (2022) report for the empirical literature. Annualised as in Hagedorn et al. (2019), these are 0.50 and 0.52, close to the 0.48 that Jappelli and Pistaferri (2014) find in Italian survey data and the 0.52 that Fagereng et al. (2021) estimate from Norwegian lottery prizes.

The wealth distribution is the third. The model reproduces the Gini of the liquid account closely in H and less well in F, and understates concentration in the illiquid account. Part of this is mechanical: the model's illiquid account cannot be negative, whereas in the survey it is negative for 12.6 per cent of households in H and 6.3 per cent in F, which is most of the missing concentration.

Participation is the fourth. The model overstates participation in the illiquid account, the share of households with positive net worth net of deposits, in both countries and by a wide margin in H: in the data a third of H households and a fifth of F households hold nothing beyond their deposits, whereas in the model almost every household holds some illiquid wealth and the only question is how much. Securities participation proper, 22.7 and 17.1 per cent in Table 1, is a narrower object with no model counterpart, since the illiquid account is all non-deposit wealth. Matching either rate would require a fixed cost of entry, which the model does not have, and the overstatement should be read as the price of that omission.

Debt service is the fifth. In steady state it is roughly three times its empirical counterpart in H and 1.7 times in F. The cause is the assumed steady-state real rate of five per cent a year, which is well above what euro-area sovereigns actually paid over the sample. This matters for the experiments because it scales the steady-state fiscal flows, and Appendix E.4 reports the headline gap and the ownership increment at a rate of 2.5 per cent: the gap is slightly larger and the increment about a third smaller.

5 Cross-Country Asymmetries

This section asks what a common monetary tightening does in the calibrated union, who pays for it, and why the two countries respond differently. We begin with the aggregate result and the channels that produce it, and then follow groups of households to see where the contraction falls. We then decompose the gap into the asymmetries responsible for it, beginning with the comparison of debt levels and ownership. Appendix D measures each asymmetry on its own against a symmetric benchmark, and robustness is collected in Appendix E.4. Appendix E.1 reports the responses to productivity and spending shocks in each country.

5.1 A common tightening in the calibrated union

A 25 basis-point annualised innovation to the union policy rule raises the real rate by about 11 basis points on impact. Annualised inflation falls by about 22 basis points, so the policy rate, which responds to it, is 9 basis points lower. The lower path of nominal rates raises the price of outstanding long bonds, and lower inflation raises the real value of their payments, so the real return on the debt already outstanding rises by about 0.2 per cent in the impact quarter, most of it through the bond price. Governments hold the real market value of their debt constant, so this return is part of the debt service they pay that quarter: debt service rises by 0.58 per cent of quarterly GDP in H and 0.82 in F, more than half of what it adds over five years, and the transfer that balances each budget

passes it on to households at once (Table 4).

The tightening contracts consumption in both countries, and by more in F: -0.37 per cent against -0.48 per cent on impact. The difference is 0.11 percentage points on impact and, summing the quarterly differences, 0.44 percentage-point-quarters over twenty quarters; the countries' own cumulated responses are -2.34 and -2.79. Figure 5 shows both. The per-quarter difference is largest on impact and decays steadily, so the cumulative figure is the sum of a shrinking series rather than a gap that widens with the horizon.

The impact difference is not a calibration target, and it is close to what euro-area data show. Local projections of real consumption per capita on the monetary policy shocks of Jarociński and Karadi (2020), scaled to a shock that raises the one-year rate by as much as the model's real rate rises over its first year, put the impact difference between the two groups at 0.14 percentage points under either decomposition of the shocks; the interval of the more precise one runs from 0.03 to 0.29 and contains the model's 0.11 (Appendix C). The cumulated difference is not identified in the data.

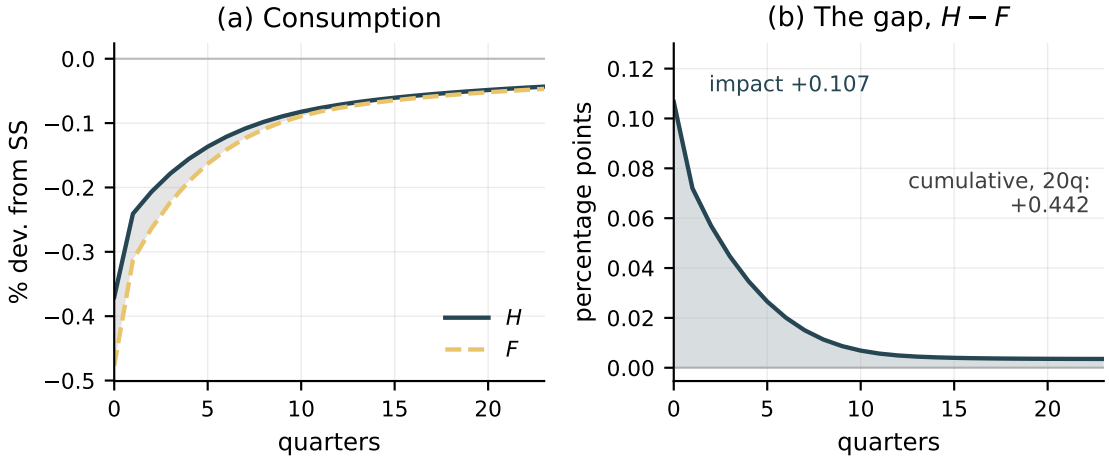


Figure 5: A common monetary tightening in the calibrated union. Panel (a) is the consumption response of each country, with the difference between them shaded. Panel (b) is that difference on its own, H minus F, in percentage points; the annotation gives the impact value and what the series cumulates to over twenty quarters. The shock is 25 basis points annualised on r^*, AR(0.85). A positive gap means F contracts by more.

What makes the gap interesting is where it comes from. The exact channel decomposition of Appendix B splits each country's consumption response into four parts, and the split is reported in Table 4.

The transfer channel dominates in both countries and is the channel on which they differ. It carries 55 per cent of the impact response in H and 62 per cent in F, and the difference persists: cumulatively, 64.5 per cent against 57.1. The tightening raises the cost of servicing government debt, and in this model that cost reaches households directly through the government budget, as lower transfers. F carries roughly sixty per cent more debt relative to output than H, so a given rise in the return on debt moves more euros there; and a larger part of F's debt service is paid to households in H, since H's investors hold more of F's debt than F's hold of H's. The tax is paid at home and part of the debt service crosses the border, and H households collect the other side of that flow. The holder outside the union, which is larger for H, is a separate leakage that Section 5.3 measures.

	H	F
Rates and the government budget		
Policy rate, impact (bp)	-8.8	
CPI inflation, impact (bp)	-22.2	-23.3
Real rate, impact (bp)	$+11.0$	$+11.7$
Return on government debt, impact (%)	$+0.22$	$+0.19$
Debt service, impact / cumulated 20q	$+0.58$ / $+0.98$	$+0.82$ / $+1.59$
Transfers, impact / cumulated 20q	-0.67 / -1.49	-0.92 / -2.11
Consumption, impact (%)	-0.37	-0.48
Consumption, cumulated 20q (%-quarters)	-2.34	-2.79
Gap $H - F$, impact / cumulated	$+0.11$ / $+0.44$	
Channel shares, impact		
Earnings	21.7%	21.7%
Transfers	55.0%	62.0%
Deposit return	5.8%	4.7%
Illiquid return	17.4%	11.6%
Channel shares, cumulated 20q		
Earnings	47.9%	44.3%
Transfers	57.1%	64.5%
Deposit return	-22.6%	-18.5%
Illiquid return	17.5%	9.7%
Contributions to the gap $H - F$, impact / cumulated 20q		
Earnings	$+0.023$ / $+0.111$	
Transfers	$+0.092$ / $+0.458$	
Deposit return	$+0.001$ / $+0.014$	
Illiquid return	-0.009 / -0.141	

Table 4: A common monetary tightening in the calibrated union: rates, the budget, responses and channel shares. The shock is 25 basis points annualised, AR(0.85), on r^*. Rates and inflation are deviations from the steady state in basis points at an annual rate; the return on government debt is the realised real holding return in the impact quarter, not annualised. Debt service and transfers are in per cent of quarterly GDP, and their cumulated values sum the first twenty quarters. Shares are of each country's own consumption response and sum to one hundred per cent by the exact decomposition of Appendix B; the deposit-return share turns negative cumulatively because higher deposit income offsets part of the contraction. The contributions are each channel's response in H minus that in F, in percentage points of each country's consumption, and add up to the gap, which is H minus F, so a positive number means F contracts by more.

The illiquid-return channel runs the other way and is the reason the gap is not larger. It carries 17.4 per cent of the impact response in H against 11.6 in F. The illiquid return falls on impact, by about 0.07 percentage points in both countries, as equity prices fall, and then stays above its steady-state value with the real rate. The returns of the two countries move almost identically; their households respond differently. F's households hold more illiquid wealth relative to what they consume, 25 quarters of consumption against 19 in H, so over five years the higher return that follows the impact quarter adds to their spending, while in H the incentive to save that it creates weighs more. In signed terms the offset is small on impact, about 0.01 percentage points against the transfer channel's 0.09, and larger cumulatively, 0.14 percentage-point-quarters against 0.46. The earnings channel carries the same share on impact in both countries, 21.7 per cent, but of different totals, so it too contributes to the gap, about 0.02 percentage points on impact and 0.11 cumulated: hours and earnings fall by more in F, the country whose consumption falls more.

The deposit-return channel lowers consumption on impact, by 5.8 per cent of the response in H and 4.7 in F. Deposits are written on the realised bond return, so they earn more from the first quarter, and that income raises consumption; households with liquid wealth also save more in anticipation of the higher returns ahead, and on impact the saving is the larger effect. Cumulatively the channel's share turns negative in both countries: the higher deposit income offsets part of the contraction, and the offset is the larger share in H, where deposits are a larger fraction of household wealth.

5.2 Who pays

The aggregate response hides where the contraction falls. Figure 6 follows households grouped at the moment the shock hits, by net worth and by liquidity, and compares each group's consumption with its own path had the shock not occurred; Appendix E.2 gives the method and the tables. Measured against what each group consumes, the fall is concentrated at the bottom of the wealth distribution. The poorest fifth of households cuts consumption by 2.0 per cent on impact in H and 2.5 per cent in F, the richest fifth by less than 0.2 per cent in both, so the poorest fifth, which accounts for about 7 per cent of consumption, accounts for 40 per cent of the fall on impact in H and 36 per cent in F. Most of these households have almost no liquid wealth and live off earnings and transfers, so the transfer cut through which the budget passes on the higher debt service reaches their spending one for one: it accounts for about 85 per cent of their fall. The richest fifth loses much less, and a larger part of what it loses comes through its portfolio; over five years the higher deposit income that follows the tightening offsets much of what it loses in earnings and transfers.

Per household, the gap between the two countries also sits at the bottom. The poorest fifth in F contracts by half a percentage point more on impact than the poorest fifth in H, the richest fifths by three hundredths of a point, and over five years the richest fifths of the two countries respond alike. For the poor hand-to-mouth, with neither liquid nor illiquid wealth, the difference is over two points, 6.1 per cent in F against 3.9 in H, because F's larger debt service means a larger transfer cut relative to what these households consume. Which households produce the aggregate gap is a different question, because it weights each group by its consumption and by how common it is in each country (Table 11). On that measure the five quintiles contribute fairly evenly on impact, and the poor hand-to-mouth contribute nothing: they respond more in F, but there are fewer of them there. The group that stands out is the wealthy hand-to-mouth, households with illiquid wealth and no liquid buffer, who are almost twice as common in F and account for about two

fifths of the aggregate gap. This is the liquidity wedge of Appendix D.5 at work.

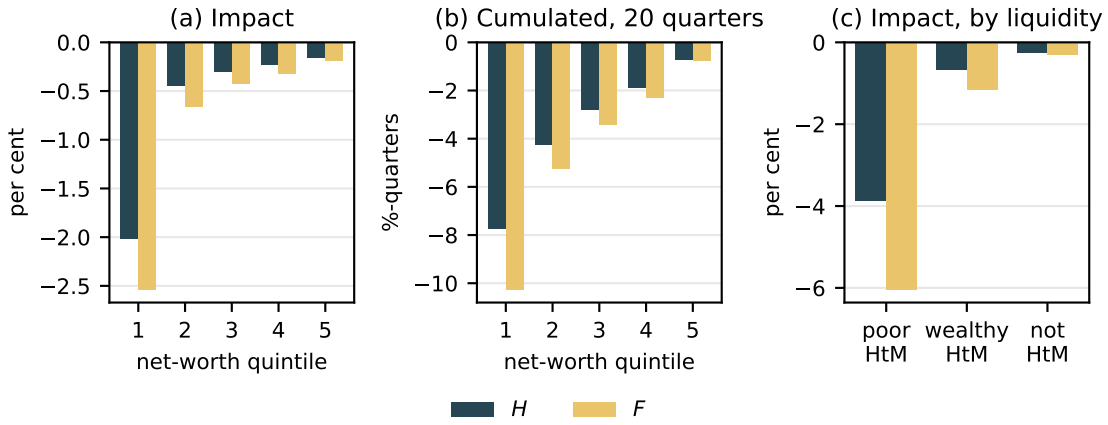


Figure 6: Who pays for a common tightening. Consumption response of household groups fixed at the moment of the shock, in per cent of each group's own consumption along its path without the shock. Panel (a): on impact, by fifth of each country's net-worth distribution. Panel (b): the same, summed over twenty quarters. Panel (c): on impact, by liquidity; poor hand-to-mouth households have neither liquid nor illiquid wealth, wealthy hand-to-mouth households have illiquid wealth but no liquid buffer. The shock is 25 basis points annualised on r^*, AR(0.85), and the responses are first order.

The effect on inequality is immediate and short-lived (Figure 7). The consumption Gini rises by 0.20 points in H and 0.27 in F on impact, on a scale from 0 to 100, halves within a quarter and leaves a small residual; the poorest fifth's response is front-loaded in the same way, while the richest fifth barely moves at any horizon. The share of households at the liquid constraint rises by half a percentage point in H and a full point in F in the quarter after the shock, most of it in F among the wealthy hand-to-mouth, as the transfer cut exhausts the buffers of households near the constraint.

How the budget is balanced decides how the cost is shared within each country. When the level of the labour tax adjusts and transfers stay at their steady-state value, the poorest fifth cuts consumption by 0.6 per cent on impact in H and 0.7 in F, and the richest fifth by 0.2, so the poorest fifth's share of the fall on impact is 16 and 14 per cent, against 40 and 36 when transfers adjust. A tax on earnings takes more, in euros, from households who earn more, and these households smooth their consumption. Inflation and the policy rate also fall by less under the tax, so the impact rise in debt service is about half as large, 0.30 and 0.32 per cent of quarterly GDP against 0.58 and 0.82. The tax spreads the contraction over time: consumption falls by less on impact and by more over five years, -2.58 and -3.01 cumulated over twenty quarters against -2.34 and -2.79. The gap between the countries over five years is almost unchanged, 0.43 percentage-point-quarters against 0.44, while the impact gap falls from 0.11 to 0.05, and the wealthy hand-to-mouth still account for about two fifths of it on impact (Table 14).

5.3 What accounts for the gap

The calibrated union differs from a symmetric one along many dimensions at once, and they do not enter additively. We therefore proceed in two steps. The first isolates the two sovereign margins the paper is about, the debt level and intra-union ownership, by

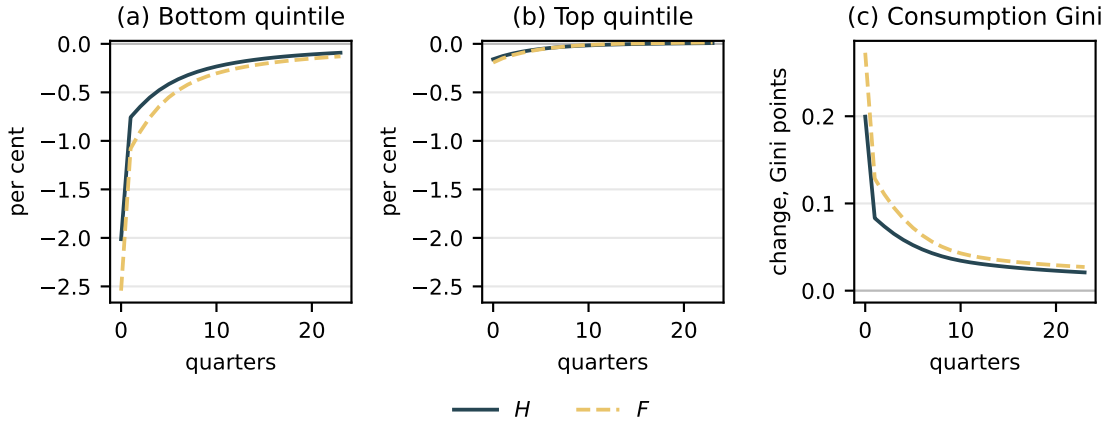


Figure 7: The dynamics of who pays. Panels (a) and (b): consumption response of the bottom and the top fifth of the net-worth distribution, in per cent of the group's own consumption, on a common scale. Panel (c): change in the consumption Gini, in points on a scale from 0 to 100. Same shock and groups as Figure 6.

comparing four calibrated unions. The second places them among every other difference between the two countries with a Shapley decomposition.

Debt levels and ownership. Table 5 compares four unions that differ in whether intra-union cross-holdings are present and in whether the two debt ratios take their observed values or the GDP-weighted union mean. Every union is a separate joint recalibration, so each row is an economy that hits the same household targets, and the differences between rows are comparisons across calibrated economies rather than the effect of moving one parameter at fixed preferences and technology.

	impact	cumulative (20q)
The consumption gap in each union		
Observed debt, ownership on	+0.107	+0.442
Observed debt, ownership off	+0.091	+0.261
Common debt, ownership on	+0.008	+0.045
Common debt, ownership off	−0.001	−0.049
Differences between unions		
Ownership increment at observed debt	+0.016	+0.181
Ownership increment at common debt	+0.009	+0.094
Debt-by-ownership interaction	+0.007	+0.087

Table 5: Debt levels and intra-union ownership: four calibrated unions. The gap is H minus F, in percentage points on impact and percentage-point-quarters cumulated over quarters 0 to 19; a positive value means F contracts by more. Ownership off removes intra-union cross-holdings, $\kappa_H = \kappa_F = 1$, and leaves the outside holder in place; common debt sets both debt ratios to the GDP-weighted union mean. Each union is a separate joint recalibration that passes the steady-state and accounting invariants. The interaction is the observed-debt increment minus the common-debt increment. Entries are rounded independently.

Within the sovereign block, the debt level does most of the work. Setting both coun-

tries' debt to the union mean, holding ownership at its observed values, collapses the cumulative gap from +0.442 to +0.045: ninety per cent of the gap goes with the fact that F owes more. That is not an additive share, since the ownership increment at observed debt is itself forty per cent of the gap and the two overlap; debt levels account for the larger difference.

Ownership matters conditional on those levels. Switching intra-union cross-holdings off at observed debt levels removes +0.181 of the cumulative gap. Doing the same at equalised debt levels removes only +0.094. Ownership is worth about twice as much at the debt levels actually observed, because both of the ways it works scale with the stock. Through the budget, the cross-border flow of debt service is the product of the share held across the border and the debt stock, so a larger stock makes a given ownership gap matter more. Through asset returns, the cross-border holding shifts the composition of the two intermediaries' books by an amount that is itself proportional to the stock. The two margins cannot be reported separately without discarding this interaction, which is why Table 5 reports the gap in all four unions and the increment at both debt levels. The difference between the two increments, +0.087 cumulated, is the interaction itself. The increment is a smaller part of the gap on impact, about a seventh, than cumulatively, about forty per cent: the level's effect arrives with the first quarter's debt service, whereas the cross-border annuity persists.

The household groups of Section 5.2 show how the ownership increment reaches households (Figure 8), comparing each group with the same-ranked group in the union without cross-holdings. Cross-holdings raise the gap in every group: every group in H contracts by less, and every group in F by more, than without them. Per household the increment is largest at the bottom, about twice what it is elsewhere over five years, but it is spread far more evenly than the gap itself, which over five years is 2.55 percentage-point-quarters for the poorest fifth and 0.04 for the richest. The reason is that ownership works through two channels, at the two ends of the distribution. At the bottom it works through the budget, as above: F pays part of its debt service to households in H, and the transfer carries the difference to the households who live off it. At the top it works through asset returns. With cross-holdings, H's intermediary holds more sovereign bonds and less domestic equity, and F's the reverse, which is the composition tilt of Proposition 1. Over the tightening a book tilted toward sovereign bonds delivers its holders more than one tilted toward equity, and the richest fifth receives three quarters of its increment this way. Because the upper two fifths consume more than three fifths of the total, they account for more than half of the aggregate increment (Table 11). The wealthy hand-to-mouth, who account for about two fifths of the gap itself, account for about a tenth of the increment: what ownership adds reaches mainly households with a liquid buffer.

For the aggregate, the illiquid return is the largest single channel of the increment, 0.10 of the 0.18 cumulated, and it adds to the budget channel rather than offsetting it: it narrows the gap between the two countries, as Section 5.1 showed, but it widens the part of the gap that ownership accounts for. The product of the share held across the border and the stock therefore describes the part of the ownership result that runs through the budget and earnings, a little over half of it; the rest runs through the composition of the two intermediaries' books (Table 13). The increment is much the same when the labour tax balances the budget and when deposits are written on the short rate, 0.17 and 0.18 cumulated, with an illiquid-return contribution of 0.10 in both (Appendix E.4).

The rest-of-world holder runs against the others. Removing it *raises* the cumulative gap by 0.018 percentage-point-quarters and leaves the impact gap within a thousandth of a point. H has the larger holder outside the union, 30 per cent of its debt against 11 per

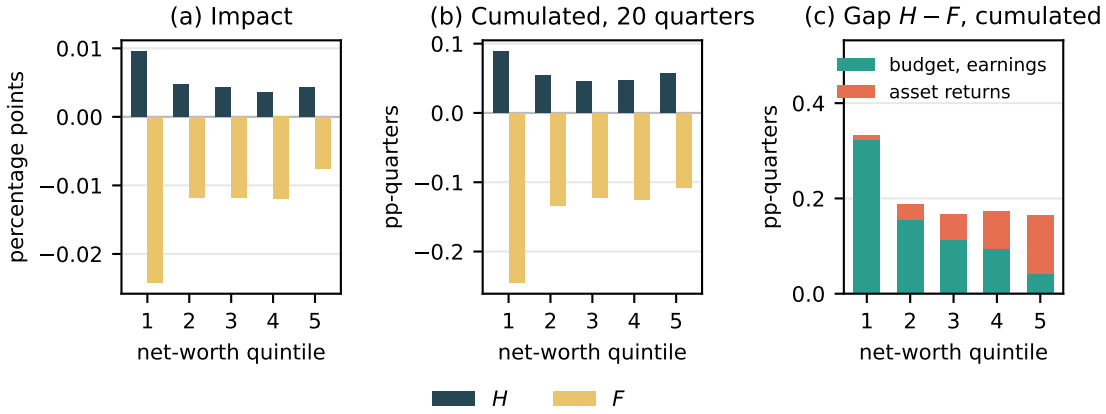


Figure 8: Who bears the ownership increment. Each group's consumption response in the calibrated union minus the same-ranked group's response in the union without intra-union holdings of Table 5, groups as in Figure 6. Panels (a) and (b): each country, on impact and summed over twenty quarters, in percentage points of the group's own consumption. Panel (c): the increment of the $H - F$ gap summed over twenty quarters, split into what reaches households through the budget and earnings and what reaches them through asset returns.

cent for F, so a larger share of H's own debt service leaks out of the euro area entirely. That leakage hurts H, and its absence would leave H better off relative to F than it is. The effect is small.

The broader sources of asymmetry. To place the two sovereign margins among everything else that differs, we compute Shapley values over four groups of primitives: the sovereign block (debt levels, ownership, and the holder structure), the household block (wealth, portfolios, and the liquidity wedge), the fiscal block (taxes, spending, and transfers), and country size. Each group's value is its average marginal contribution across all orders in which the groups could be switched on, so the four exhaust the total gap by construction.

	impact	cumulative (20q)
Shapley values, by group		
Sovereign	+0.097	+0.479
Households	+0.022	+0.107
Fiscal	−0.006	−0.071
Country size	−0.007	−0.073
Total	+0.107	+0.442

Table 6: Attribution of the $H - F$ consumption gap under a common tightening. Impact values are percentage points, cumulative values percentage-point-quarters. The four Shapley values sum to the total exactly, which is an algebraic property of the construction. Each of the sixteen coalitions is a separate joint recalibration of the steady state. The sovereign block bundles the debt levels, the intra-union ownership shares, the outside holder and the spread; Table 5 separates the first two.

The sovereign block accounts for essentially all of the gap. Its Shapley value, +0.479

cumulated, slightly exceeds the total of +0.442, because the fiscal and size blocks contribute negatively. Households add +0.107. What makes common monetary policy land unevenly across this union is above all the structure of sovereign debt, with household balance sheets second and the size and composition of government working the other way. The ranking is specific to this calibration and to these groups.

One asymmetry at a time. Appendix D moves each difference between the two countries on its own to its observed value, in an otherwise symmetric union. Together, the debt level and ownership move the gap far more than any other difference; maturity, the spread and the tax mix move it little, and among the remaining margins the liquidity wedge and the size of the tax-and-transfer system matter most. The wedge is the margin behind the household block's Shapley value, and the larger wealth and deposits of F's households offset part of it.

Who replaces the foreign creditors? Investors outside the union hold 11 per cent of F's debt. Removing them, with every parameter held fixed, leaves the union's households to absorb that debt, and the model decides which households do (Appendix E.3). Finance settles where the debt goes: H's savers end up holding almost three fifths of it and F's the rest, while the union's real rate rises by about one basis point. Trade settles what this does to the common tightening, through whose goods the departing investors bought. If they bought F's goods, their exit lowers F's terms of trade and widens the cumulated gap by close to a fifth; if their spending was spread across the union, the gap narrows slightly. Either way the burden within F keeps its distribution.

6 Conclusion

The poorest households pay most for a common monetary tightening in a monetary union, and households in the member that owes more pay more than those in the other. We have shown this in a two-country heterogeneous-agent model calibrated to euro-area data. A tightening raises debt service, that cost reaches households as taxes and forgone transfers, and the households who live off earnings and transfers have almost no liquid wealth with which to absorb it. Per household, the difference between the two countries is concentrated at the bottom of the wealth distribution. In the aggregate, about two fifths of it comes from the wealthy hand-to-mouth, who are almost twice as common in the country that owes more and turn the same change in transfers into a larger change in consumption, though the larger wealth and deposits that country's households hold cushion part of that response.

The mechanism of the gap between the countries is the government budget. The country that owes more, and pays a larger part of its debt service to the other member's households, absorbs more of the rise in debt service. Working against this, households in that country hold more illiquid wealth relative to what they consume, and the higher return that follows a tightening supports their spending. The observed asymmetry is the net of the two, and a model with only one of them would misstate it in a direction that depends on which one it kept.

Where the debt is held shapes how the burden is shared. The debt level is the larger margin, and ownership changes how much a given level matters: removing intra-union holdings takes almost twice as much out of the consumption gap at observed debt levels as it does when the two countries are given the same debt, because both of the ways ownership works scale with the stock. Through the budget it moves debt service across

the border, which reaches the households who live off transfers; through asset returns it tilts the creditor's intermediary toward bonds and the debtor's toward equity, which reaches the households who hold the illiquid wealth. For ownership the budget and asset returns point the same way, and what it adds to the gap is spread far more evenly across households than the gap itself. Measuring ownership at equalised debt understates it, and in a model whose members trade a single pooled bond the margin is absent by construction; reporting either margin on its own understates the pair.

The results carry over to a lower steady-state real rate, where the gap is slightly larger and the ownership increment remains positive and of the same order (Appendix E.4). When the labour tax balances the budget, the five-year gap and the ownership increment are almost unchanged, and the burden within each country is spread more evenly (Section 5.2). The framework also lends itself to extensions: default risk with spreads that accrue to bondholders, which would tie ownership to sovereign risk; housing and mortgage debt, which would add a second household balance-sheet channel; and a welfare analysis of how the burden of a common policy is shared.

More broadly, in a monetary union the transmission of the common policy is bound up with national fiscal positions. A tightening reaches households largely through national budgets, so how much each government owes, who holds its debt and how liquid its households are together decide who pays, within countries as well as between them. These are slow-moving features of a union that policy shapes, through fiscal rules, debt management and the integration of capital markets, and they deserve a place in any assessment of how evenly the common policy is transmitted.

References

Almgren, M., Gallegos, J.-E., Kramer, J., and Lima, R. (2022). Monetary policy and liquidity constraints: Evidence from the euro area. *American Economic Journal: Macroeconomics*, 14(4):309–340.

Altavilla, C., Brugnolini, L., Gürkaynak, R. S., Motto, R., and Ragusa, G. (2019). Measuring euro area monetary policy. *Journal of Monetary Economics*, 108:162–179.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–2367.

Auclert, A., Bardóczy, B., and Rognlie, M. (2023). MPCs, MPEs, and Multipliers: A Trilemma for New Keynesian Models. *Review of Economics and Statistics*, 105(3):700–712.

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

Auclert, A., Rognlie, M., and Straub, L. (2024). The Intertemporal Keynesian Cross. *Journal of Political Economy*, 132(12):4068–4121.

Auclert, A., Rognlie, M., and Straub, L. (2025). Fiscal and monetary policy with heterogeneous agents. *Annual Review of Economics*, 17.

Bayer, C., Kriwoluzky, A., Müller, G. J., and Seyrich, F. (2024). A HANK2 Model of Monetary Unions. *Journal of Monetary Economics*, page 103579.

Bellifemine, M., Couturier, A., and Jamilov, R. (2025). Monetary unions with heterogeneous fiscal space. *Journal of International Economics*, 156:104092.

Bilbiie, F. O. (2008). Limited asset markets participation, monetary policy and (inverted) aggregate demand logic. *Journal of Economic Theory*, 140(1):162–196.

Bocola, L. (2016). The Pass-Through of Sovereign Risk. *Journal of Political Economy*, 124(4):879–926.

Bouabdallah, O., Jacquinot, P., and Šterc, A. (2025). Fiscal Multipliers, Tax Structures, and Public Debt in a HANK Model. SSRN working paper 5119261; also Banco de Portugal Working Paper 2025-08.

Broer, T., Krusell, P., and Öberg, E. (2023). Fiscal Multipliers: A Heterogenous-Agent Perspective. *Quantitative Economics*, 14(3):799–816.

Broner, F., Clancy, D., Erce, A., and Martin, A. (2022). Fiscal Multipliers and Foreign Holdings of Public Debt. *Review of Economic Studies*, 89(3):1155–1204.

Broner, F., Erce, A., Martin, A., and Ventura, J. (2014). Sovereign Debt Markets in Turbulent Times: Creditor Discrimination and Crowding-out Effects. *Journal of Monetary Economics*, 61:114–142.

Brunnermeier, M. K., Garicano, L., Lane, P. R., Pagano, M., Reis, R., Santos, T., Thesmar, D., Van Nieuwerburgh, S., and Vayanos, D. (2016). The Sovereign-Bank Diabolic Loop and ESBies. *American Economic Review*, 106(5):508–512.

Caramp, N. and Feilich, E. (2026). Monetary Policy and Government Debt. *Journal of Money, Credit and Banking*, 58(2):389–420.

Chen, X., Lazarakis, S., and Varthalitis, P. (2025). Debt Targets and Fiscal Consolidation in a Euro Area HANK Model. *Quantitative Economics*, 16:1321–1359.

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2011). Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins. *American Economic Review*, 101(3):471–475.

Christiano, L., Eichenbaum, M., and Rebelo, S. (2011). When is the Government Spending Multiplier Large? *Journal of Political Economy*, 119(1):78–121.

Cohen, N. (2026). Who Benefits from Capital Market Integration in a Monetary Union? Job market paper, Sciences Po.

Corsetti, G., Duarte, J. B., and Mann, S. (2022). One money, many markets: Monetary transmission and housing financing in the euro area. *Journal of the European Economic Association*, 20(1):513–548.

De Paoli, B. (2009). Monetary Policy and Welfare in a Small Open Economy. *Journal of International Economics*, 77(1):11–22.

Devereux, M. B. and Sutherland, A. (2011). Country Portfolios in Open Economy Macro-Models. *Journal of the European Economic Association*, 9(2):337–369.

Fagereng, A., Holm, M. B., and Natvik, G. J. (2021). MPC Heterogeneity and Household Balance Sheets. *American Economic Journal: Macroeconomics*, 13(4):1–54.

Fang, X., Hardy, B., and Lewis, K. K. (2025). Who Holds Sovereign Debt and Why It Matters. *Review of Financial Studies*, 38(8):2326–2361.

Farhi, E. and Tirole, J. (2018). Deadly Embrace: Sovereign and Financial Balance Sheets Doom Loops. *Review of Economic Studies*, 85(3):1781–1823.

Grancini, S. (2025). Public Debt, iMPCs and Fiscal Policy Transmission. Working Paper 3106, European Central Bank.

Hagedorn, M., Manovskii, I., and Mitman, K. (2019). The Fiscal Multiplier. Working Paper 25571, National Bureau of Economic Research.

Heathcote, J., Storesletten, K., and Violante, G. L. (2017). Optimal Tax Progressivity: An Analytical Framework. *The Quarterly Journal of Economics*, 132(4):1693–1754.

Hintermaier, T. and Koeniger, W. (2024). Differences in Euro-Area Household Finances and their Relevance for Monetary-Policy Transmission. *Quantitative Economics*, 15(4):1249–1301.

Holm, M. B., Paul, P., and Tischbirek, A. (2021). The transmission of monetary policy under the microscope. *Journal of Political Economy*, 129(10):2861–2904.

Jappelli, T. and Pistaferri, L. (2014). Fiscal Policy and MPC Heterogeneity. *American Economic Journal: Macroeconomics*, 6(4):107–136.

Jarociński, M. and Karadi, P. (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

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

Kaplan, G. and Violante, G. L. (2022). The Marginal Propensity to Consume in Heterogeneous Agent Models. *Annual Review of Economics*, 14:747–775.

Kase, H., Rigato, R. D., et al. (2025). Beyond averages: heterogeneous effects of monetary policy in a hank model for the euro area. ECB Working Paper Series 3086, European Central Bank.

Laubach, T. (2009). New Evidence on the Interest Rate Effects of Budget Deficits and Debt. *Journal of the European Economic Association*, 7(4):858–885.

Obstfeld, M. and Rogoff, K. (1996). *Foundations of International Macroeconomics*. MIT Press, Cambridge, MA.

Pamies, S., Carnot, N., and Pătărău, A. (2021). Do Fundamentals Explain Differences between Euro Area Sovereign Interest Rates? European Economy Discussion Paper 141, European Commission, Directorate-General for Economic and Financial Affairs.

Priftis, R. and Zimic, S. (2021). Sources of Borrowing and Fiscal Multipliers. *Economic Journal*, 131(633):498–519.

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

Rouwenhorst, K. G. (1995). Asset pricing implications of equilibrium business cycle models. In Cooley, T. F., editor, *Frontiers of Business Cycle Research*, pages 294–330. Princeton University Press, Princeton, NJ.

Slacalek, J., Tristani, O., and Violante, G. L. (2020). Household balance sheet channels of monetary policy: A back of the envelope calculation for the euro area. *Journal of Economic Dynamics and Control*, 115:103879.

Sutherland, A. (2005). Incomplete Pass-Through and the Welfare Effects of Exchange Rate Variability. *Journal of International Economics*, 65(2):375–399.

Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, 39:195–214.

Tille, C. and van Wincoop, E. (2010). International Capital Flows. *Journal of International Economics*, 80(2):157–175.

Woodford, M. (2001). Fiscal requirements for price stability. *Journal of Money, Credit & Banking*, 33(3):669–669.

Young, E. R. (2010). Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell–Smith Algorithm and Non-stochastic Simulations. *Journal of Economic Dynamics and Control*, 34(1):36–41.

Appendix A The model in full

This appendix derives the conditions of the model of Section 3: the household, firm, wage-setting and intermediary problems, and the steady state.

A.1 Households

We derive the household optimality conditions under the Calvo-lottery portfolio adjustment of Auclert et al. (2024), for country j and in the general tax case: consumption tax $\tau^c_{j,t}$, progressive labour income (5), lump-sum transfer $\mathrm{Tra}_{j,t}$. Constraint multipliers are written $\nu^b, \nu^a \geq 0$; the country index is suppressed on household-level objects. Cash-on-hand from labour income and transfers is $y_{i,t} = e_{i,t} + \mathrm{Tra}_{j,t}$.

Value function. Each period the household draws an i.i.d. adjustment status: adjuster with probability χ_j, non-adjuster with probability $1-\chi_j$. The value of entering t with $(z_{i,t}, b_{i,t-1}, a_{i,t-1})$ is

$$V_t(z_{i,t}, b_{i,t-1}, a_{i,t-1}) = \chi_j\, V^A_t(z_{i,t}, b_{i,t-1}, a_{i,t-1}) + (1-\chi_j)\, V^{NA}_t(z_{i,t}, b_{i,t-1}, a_{i,t-1}). \tag{44}$$

Non-adjuster. Illiquid holdings roll over mechanically, $a_{i,t} = (1 + r^a_{j,t})\, a_{i,t-1}$, and only $b_{i,t}$ is chosen:

$$V^{NA}_t = \max_{b_{i,t}\geq 0}\; u\big(c^{NA}_{i,t}\big) - v(N_{j,t}) + \beta_j\, \mathbb{E}_t\, V_{t+1}\big(z_{i,t+1},\, b_{i,t},\, (1 + r^a_{j,t})\, a_{i,t-1}\big), \tag{45}$$

subject to $(1+\tau^c_{j,t})\, c^{NA}_{i,t} = y_{i,t} + (1 + r^b_{j,t})\, b_{i,t-1} - b_{i,t}$; the non-adjuster's cash-on-hand excludes the illiquid return. The first-order condition for $b_{i,t}$ is

$$\frac{u'(c^{NA}_{i,t})}{1+\tau^c_{j,t}} = \nu^{b,NA}_{i,t} + \beta_j\, \mathbb{E}_t\, \frac{\partial V_{t+1}}{\partial b}\big(z_{i,t+1},\, b_{i,t},\, (1 + r^a_{j,t})\, a_{i,t-1}\big), \tag{46}$$

with complementary slackness $\nu^{b,NA}_{i,t}\, b_{i,t} = 0$.

Adjuster. Both accounts are chosen freely:

$$V^A_t = \max_{a_{i,t}\geq 0,\; b_{i,t}\geq 0}\; u\big(c^A_{i,t}\big) - v(N_{j,t}) + \beta_j\, \mathbb{E}_t\, V_{t+1}\big(z_{i,t+1},\, b_{i,t},\, a_{i,t}\big), \tag{47}$$

subject to $(1+\tau^c_{j,t})\, c^A_{i,t} = y_{i,t} + (1 + r^a_{j,t})\, a_{i,t-1} + (1 + r^b_{j,t})\, b_{i,t-1} - a_{i,t} - b_{i,t}$, with first-order conditions

$$\frac{u'(c^A_{i,t})}{1+\tau^c_{j,t}} = \nu^{b,A}_{i,t} + \beta_j\, \mathbb{E}_t\, \frac{\partial V_{t+1}}{\partial b}, \qquad \frac{u'(c^A_{i,t})}{1+\tau^c_{j,t}} = \nu^{a,A}_{i,t} + \beta_j\, \mathbb{E}_t\, \frac{\partial V_{t+1}}{\partial a}, \tag{48}$$

both continuation derivatives evaluated at $(z_{i,t+1}, b_{i,t}, a_{i,t})$: an unconstrained adjuster equates the marginal value of the two accounts.

Envelope conditions. Differentiating (44) using (45)–(47),

$$\frac{\partial V_t}{\partial b_{t-1}} = \frac{1 + r^b_{j,t}}{1+\tau^c_{j,t}}\Big[\chi_j\, u'(c^A_{i,t}) + (1-\chi_j)\, u'(c^{NA}_{i,t})\Big], \tag{49}$$

$$\frac{\partial V_t}{\partial a_{t-1}} = (1 + r^a_{j,t})\Big[\frac{\chi_j\, u'(c^A_{i,t})}{1+\tau^c_{j,t}} + (1-\chi_j)\, \beta_j\, \mathbb{E}_t\, \frac{\partial V_{t+1}}{\partial a}\big(z_{i,t+1},\, b_{i,t},\, (1 + r^a_{j,t})a_{i,t-1}\big)\Big]. \tag{50}$$

Returns are dated at the period in which they are realised (Section 3.1). In (49) the liquid return multiplies the probability-weighted marginal utility because both types can spend liquid wealth; in (50) the non-adjuster's marginal value of illiquid wealth is the continuation value (the account is not a choice for them), while the adjuster's is current marginal utility, deflated by the consumption tax through which any wealth reaches consumption.

Numerical solution and aggregation. The problem is solved by the endogenous grid method on the state (s, z, a, b), where s is the adjustment status; because the Calvo lottery is i.i.d., the continuation value is common to adjusters and non-adjusters, which is what makes the backward iteration a standard one-asset EGM step for each type. The distribution is iterated forward following Young (2010), with exogenous transition $\Pi^{\text{Calvo}} \otimes \Pi^{z}$. Aggregation over $\mathcal{D}_{j,t}$ delivers $C_{j,t}$, $A^{\text{liq}}_{j,t}$, $A^{\text{illiq}}_{j,t}$, the hand-to-mouth indicators, and $\text{UCE}_{j,t}$ of (7), whose weight $z^{1-\vartheta_j}$ reflects that under the HSV schedule the marginal after-tax income of an efficiency unit carries the curvature $1-\vartheta_j$ (Appendix A.3).

A.2 Firms and price setting

All objects are in units of the country-j good; discounting is at the producer real rate $r^{p}_{j,t}$. The variety index is f; the country index is suppressed where unambiguous.

Intermediate producers and factor demands. Firm f produces

$$y_{f,t} = Z_j \, k_{f,t-1}^{\alpha_j} \, n_{f,t}^{1-\alpha_j}.$$

Cost minimisation gives the factor demands

$$r^{K}_{j,t} = \alpha_j \, mc_{j,t} \, \frac{Y_{j,t}}{K_{j,t-1}}, \qquad w_{j,t} = (1-\alpha_j) \, mc_{j,t} \, \frac{Y_{j,t}}{N_{j,t}}, \tag{51}$$

with r^K the rental rate of capital; the second equation is (8)'s marginal cost.

Price setting. Demand for variety f is $y_{f,t} = (P_{f,t}/P_t)^{-\varepsilon^p} Y_{j,t}$ with $\varepsilon^p = \mu^p_j/(\mu^p_j - 1)$. The firm chooses $P_{f,t}$ to maximise

$$\mathcal{J}(P_{f,t}, P_{f,t-1}) = \max_{P_{f,t}} \left\{ \frac{P_{f,t}}{P_t} \, y_{f,t} - mc_{j,t} \, y_{f,t} - \psi_{f,t} + \frac{1}{1+r^{p}_{j,t}} \, \mathcal{J}(P_{f,t+1}, P_{f,t}) \right\}, \tag{52}$$

with the adjustment cost quadratic in log price changes,

$$\psi_{f,t} = \frac{\mu^p_j}{\mu^p_j - 1} \, \frac{1}{2\kappa^p} \left[\ln(P_{f,t}/P_{f,t-1}) \right]^2 Y_{j,t}.$$

Differentiating operating profit

$$(P_{f,t}/P_t)^{1-\varepsilon^p} Y_{j,t} - mc_{j,t} (P_{f,t}/P_t)^{-\varepsilon^p} Y_{j,t}$$

and combining terms, the revenue and cost derivatives collapse to $\varepsilon^p \, (y_{f,t}/P_{f,t}) \, (mc_{j,t} - 1/\mu^p_j)$, so the first-order condition is

$$\begin{aligned} \frac{\varepsilon^p}{P_{f,t}} \left(mc_{j,t} - \frac{1}{\mu^p_j} \right) y_{f,t} = \frac{\mu^p_j}{\mu^p_j - 1} \, \frac{1}{\kappa^p \, P_{f,t}} \ln \frac{P_{f,t}}{P_{f,t-1}} \, Y_{j,t} \\ - \frac{1}{1+r^{p}_{j,t}} \, \frac{\mu^p_j}{\mu^p_j - 1} \, \frac{1}{\kappa^p \, P_{f,t}} \ln \frac{P_{f,t+1}}{P_{f,t}} \, Y_{j,t+1}. \end{aligned} \tag{53}$$

Imposing symmetric equilibrium ($P_{f,t} = P_t$, $y_{f,t} = Y_{j,t}$, $1 + \pi_{j,t} = P_t/P_{t-1}$) and using $\varepsilon^p = \mu_j^p/(\mu_j^p - 1)$ recovers (14) exactly: the log form follows from the log-quadratic cost, with no approximation. The level-quadratic Rotemberg cost delivers the same curve to first order; the log form nests the log-linearised NKPC exactly and is better behaved nonlinearly. The resource cost entering (36) is $\psi_{j,t}^P$ of (13).

Capital producers. Competitive capital producers buy old capital $(1-\delta)K_{j,t-1}$ at price $Q_{j,t}$, transform the gross goods outlay $I_{j,t}^{\text{tot}}$ into new capital subject to the adjustment cost, and sell installed capital at $Q_{j,t}$:

$$K_{j,t} = (1-\delta)K_{j,t-1} + I_{j,t}^{\text{tot}} - \frac{1}{2\delta\varepsilon^I}\left(\frac{K_{j,t}}{K_{j,t-1}} - 1\right)^2 K_{j,t-1}, \tag{54}$$

which rearranges to (9). Two objects must be kept apart: the gross outlay $I^{\text{tot}} = I + \psi^K$, which appears in the capital Euler equation, and net investment I, which enters dividends and goods-market clearing separately from ψ^K. The producer maximises $Q_{j,t}K_{j,t} - Q_{j,t}(1-\delta)K_{j,t-1} - I_{j,t}^{\text{tot}}$ taking $Q_{j,t}$ as given; the first-order condition for $K_{j,t}$ equates $Q_{j,t}$ to the marginal goods cost of installed capital, $\partial I_{j,t}^{\text{tot}}/\partial K_{j,t}$, which is (10). The intermediate firm's choice of $K_{j,t}$, discounted at $r_{j,t}^p$, yields the capital Euler equation (11); the capital block solves $\{Q_{j,t}, K_{j,t}\}$ jointly against (10) and (11), so $Q_{j,t+1}$ is carried explicitly and must not be substituted out.

A.3 Wage setting

A continuum of unions indexed by k set nominal wages $W_{k,t}$ for differentiated labour types in country j; a competitive packer aggregates them with elasticity ε^w, so demand for type k is $N_{k,t} = (W_{k,t}/W_t)^{-\varepsilon^w} N_{j,t}$, with $\mu^w = \varepsilon^w/(\varepsilon^w - 1)$. Hours are rationed uniformly across members, $n_{i,t} = N_{k,t}$, while earnings scale with idiosyncratic productivity: member i's after-tax income follows the HSV schedule on the union's own CPI-real wage, $e_{i,t} = \lambda_{j,t}\big(w_{k,t}^c N_{k,t}\, z_{i,t}\big)^{1-\vartheta_j}$, which equals (5) in symmetric equilibrium, and is spent on consumption taxed at $\tau_{j,t}^c$, so a unit of after-tax income is worth $u'(c_{i,t})/(1+\tau_{j,t}^c)$ in utility. The union maximises

$$\sum_{s=0}^{\infty} \beta_j^s \left\{ \int_i \Big[u(c_{i,t+s}) - \bar{\varphi}_j \frac{N_{k,t+s}^{1+1/\varphi}}{1+1/\varphi} \Big] \mathrm{d}i - \frac{1}{\mu^w - 1}\frac{1}{2\kappa_0^w}\left[\ln(1+\pi_{j,t+s}^w)\right]^2 \right\}, \tag{55}$$

where κ_0^w is the primitive slope parameter (re-parameterised below).

First-order condition. Raising $W_{k,t}$ raises income per hour but cuts hours by more: with $\partial N_{k,t}/\partial W_{k,t} = -\varepsilon^w N_{k,t}/W_{k,t}$, the real wage *bill* responds with elasticity $1-\varepsilon^w < 0$, while the disutility side carries $-\varepsilon^w$ (only hours enter it). The marginal benefit of a higher wage is the income response weighted by marginal utility, which carries the HSV curvature $1-\vartheta_j$ in both the income level and the productivity weight:

$$\int_i \frac{\partial e_{i,t}}{\partial W_{k,t}} \frac{u'(c_{i,t})}{1+\tau_{j,t}^c}\,\mathrm{d}i = (1-\vartheta_j)(1-\varepsilon^w)\,\frac{\lambda_{j,t}\big(w_{k,t}^c N_{k,t}\big)^{1-\vartheta_j}}{(1+\tau_{j,t}^c)\,W_{k,t}}\,\mathrm{UCE}_{j,t}, \tag{56}$$

with UCE as in (7). The marginal cost is the disutility response $\varepsilon^w \bar{\varphi}_j N_{k,t}^{1+1/\varphi}/W_{k,t}$ plus the two Rotemberg terms $\frac{1}{\mu^w-1}\frac{1}{\kappa_0^w}\big[\ln(1+\pi_{j,t}^w) - \beta_j \ln(1+\pi_{j,t+1}^w)\big]/W_{k,t}$.

Symmetric equilibrium and normalisation. Setting the sum of the three to zero, imposing $W_{k,t} = W_t$, $N_{k,t} = N_{j,t}$, multiplying by W_t, and using $1/(\mu^w - 1) = \varepsilon^w - 1$, division by ε^w gives

$$\ln(1+\pi^w_{j,t}) = \mu^w \kappa^w_0 \left[\bar{\varphi}_j N_{j,t}^{1+1/\varphi} - \frac{1-\vartheta_j}{1+\tau^c_{j,t}} \frac{\lambda_{j,t}\big(w^c_{j,t} N_{j,t}\big)^{1-\vartheta_j} \mathrm{UCE}_{j,t}}{\mu^w} \right] + \beta_j \ln(1+\pi^w_{j,t+1}), \tag{57}$$

which is (16) with $\mathcal{U} = \mathrm{UCE}$ under the re-parameterisation $\kappa^w \equiv \mu^w \kappa^w_0$: a pure relabelling of the free slope parameter, adopted so that the wage markup divides the income term, the convention of Auclert et al. (2024) and the canonical sequence-space implementation. The derivation is exact: the log form follows from the log-quadratic cost with no approximation. The baseline's representative-agent convention replaces $\mathrm{UCE}_{j,t}$ with $C_{j,t}^{-1/\sigma}$ (Section 3.3); the derivation is otherwise unchanged. At steady state the bracket must vanish, which is (17).

A.4 Intermediary and asset returns

The country-j intermediary's balance sheet, in CPI-unit market values, is the stock restriction

$$J_{j,t} + B^{\mathrm{book}}_{j,t} = A^{\mathrm{illiq}}_{j,t} + A^{\mathrm{liq}}_{j,t}, \tag{58}$$

which is asset-market clearing (43). It is not sufficient: the intermediary also faces a flow restriction (its budget), and the two together pin down the account returns.

Why one return is derived, not posited. Suppose both r^a and r^b were written down directly. The book's income over $[t-1, t]$ is $\sum_k (1+r^g_{jk,t}) B_{jk,t-1} + d^c_{j,t} + J_{j,t}$; its outlays are the two household legs plus the liquidity cost $\xi_j A^{\mathrm{liq}}_{j,t-1}$. Nothing forces these to be equal, so generically they are not: the intermediary earns a pure profit that accrues to nobody, and the discrepancy surfaces as a resource-constraint violation. Zero profit under perfect competition, (18), is therefore an equation, not a normalisation, and it determines one of the two returns.

Which return to posit. We posit the deposit rate. Deposits are a nominal claim remunerated at the realised return on the domestic sovereign leg net of the intermediation cost, $r^b_{j,t} = r^g_{j,t} - \xi_j$ per (6), so their return is pinned to one observable rate whatever the rest of the book does; at the calibrated maturity that rate includes the bond's price change, so the deposit leg shares in the sovereign revaluation. The illiquid account absorbs whatever the book delivers beyond it: a levered intermediary whose senior claim is written on its bond return passes all residual risk, on the equity leg and on the cross-border leg, to the junior claim.

Deriving r^a. Under (6) the deposit leg and the intermediation cost combine, $(1 + r^b_{j,t}) A^{\mathrm{liq}}_{j,t-1} + \xi_j A^{\mathrm{liq}}_{j,t-1} = (1 + r^g_{j,t}) A^{\mathrm{liq}}_{j,t-1}$, so the wedge cancels from the adding-up condition, which becomes

$$(1 + r^g_{j,t})\, A^{\mathrm{liq}}_{j,t-1} + (1 + r^a_{j,t})\, A^{\mathrm{illiq}}_{j,t-1} = (1 + r^{\mathrm{book}}_{j,t})\big(B^{\mathrm{book}}_{j,t-1} + J_{j,t-1}\big), \tag{59}$$

with r^{book} as in (19). Using (58) at $t-1$ and writing $\varsigma_{t-1} = A^{\mathrm{liq}}_{j,t-1} / A_{j,t-1}$ for the realised liquid share,

$$(1 + r^g_{j,t})\, \varsigma_{t-1} + (1 + r^a_{j,t})(1 - \varsigma_{t-1}) = 1 + r^{\mathrm{book}}_{j,t} \quad \Longrightarrow \quad r^a_{j,t} = \frac{r^{\mathrm{book}}_{j,t} - \varsigma_{t-1}\, r^g_{j,t}}{1 - \varsigma_{t-1}}, \tag{60}$$

which is (19) with ς_{t-1} replaced by its steady-state value ς_j. The replacement keeps the system acyclic ($A^{\text{liq}}_{j,t-1}$ is an output of a household problem that takes $r^a_{j,t}$ as an input), and the induced error, $(\varsigma_{t-1}-\varsigma_j)(r^{\text{book}}_{j,t}-r^g_{j,t})/(1-\varsigma_j)$, is a product of two deviations from steady state, hence identically zero at first order. Two properties follow. At steady state $r^{\text{book}}_j = r^a_j = r_j$, so the calibration is unaffected. Away from it, the illiquid return is levered: a one-point surprise in the book return moves r^a by $1/(1-\varsigma_j)$ points, as for any claim junior to deposits. The steady-state book composition $\bar{J}_j/(\bar{J}_j+\bar{B}^{\text{book}}_j)$ is not a parameter: it follows from the targets, and at first order it enters the returns only through the weights of (65).

The union book. Under cross-border and outside holdings the book is the set of legs of Section 3.9: H's book contains F paper, and each book excludes the part of its own government's debt held by the rest of the world. The derivation above goes through unchanged with the leg sum $\sum_k(1+r^g_{jk,t})B_{jk,t-1}$ as the book's bond income: every leg is credited at its own realised return in the holder's units, and the residual is still taken against the domestic bond return $r^g_{j,t}$, on which the deposit leg is written. Where the legs coincide (one-period bonds, no spread), the book earns a single return, because a claim's real return in the holder's units is the common nominal rate deflated by the holder's CPI, which is exactly $r^g_{j,t}$.

Bonds at steady state. With $\pi^c_j = 0$ and $i = r_j$, the pricing recursion (20) gives $\bar{P}^g_j = 1/(1+r_j-\rho^b_j)$ and $\bar{r}^g_j = r_j$; the Macaulay duration in quarters is $(1+r_j)/(1+r_j-\rho^b_j)$, so at the steady-state rate of Section 4 the common duration of 8.75 years corresponds to $\rho^b_j = 0.984$. The symmetric benchmark sets $\rho^b_j = 0$, one-period debt; the calibrated union sets the common duration of Section 4.

A.5 Steady state

At a steady state all aggregates are constant and producer and CPI inflation are zero in both countries; the terms of trade $\bar{S}$ is constant, so by (29) the common nominal rate and both real rates coincide, $i = r_H = r_F = r$, while the price-index constants $\bar{p}^c_j$ follow from (33). Per country, the real side reduces to

$$mc_j = \frac{1}{\mu^p_j}, \qquad K_j = \frac{\alpha_j\, mc_j\, Y_j}{r+\delta}, \qquad w_j = (1-\alpha_j)\, mc_j\, \frac{Y_j}{N_j}, \qquad Q_j = 1, \quad I_j = \delta K_j, \tag{61}$$

with $Y_j = Z_j K_j^{\alpha_j} N_j^{1-\alpha_j}$. Dividends and asset prices satisfy

$$d^{\text{full}}_j = Y_j - w_j N_j - \delta K_j, \qquad J_j = \frac{d^c_j}{r}, \qquad \bar{P}^g_j = \frac{1}{1+r-\rho^b_j}, \qquad r^g_j = r^a_j = r, \quad r^b_j = r-\xi_j, \tag{62}$$

the sovereign spread is zero by construction of (25), and the government budget (23) at constant debt links the steady-state transfer and the labour-tax level,

$$\overline{\text{Tra}}_j = T^L_j + \tau^c_j\, C_j + \frac{\tau^K_j\, d^c_j}{1-\tau^K_j} - G^c_j - r\, B^g_j. \tag{63}$$

The symmetric benchmark sets $\overline{\text{Tra}}_j = 0$, the standard normalisation for the tax function (5), and the calibrated union sets $\overline{\text{Tra}}_j$ to the data's non-pension cash benefits as a share

of output (Table 2); in both, (63) determines λ_j (closed form for any ϑ_j, since T_j^L in (22) is linear in λ_j). The rest of the world holds $\bar{B}_j^R = \omega_j^R \bar{B}_j^g$ and receives $\bar{X}_j^R = r\,\omega_j^R \bar{B}_j^g$ in exports of the issuer's good, the net factor payment on foreign-held debt, and the steady-state intra-union trade balance is $\overline{NX}_H = -\bar{p}_H^c\, r\, \overline{nfa}_H$. Asset and goods markets clear per (43) and (36), and the disutility scale $\bar{\varphi}_j$ follows from (17).

Calibration steps (symmetric baseline). The tax structure is a primitive, not a residual: given

$$\{r, \delta, \sigma, \varphi, \varepsilon^I, \kappa^p, \kappa^w, \mu^w, \xi_j, \rho_j^b, \rho_{e,j}, \sigma_{e,j}, \vartheta_j, \tau_j^c, \tau_j^K,$$
$$G_j/Y_j, B_j^g/Y_j, K_j/Y_j, \kappa_j, \omega_j^R, \omega_j^E, \phi^R, \theta, \eta, s, \phi_\pi, \phi^B\}$$

and the two wealth targets $\{A_j, A_j^{\text{liq}}\}$, the steady state is a recursive sequence of assignments followed by one two-dimensional household solve; the labour-tax level λ_j joins the outcomes:

1. Normalise $Y_j = N_j = 1$; set K_j and B_j^g from their ratios.
2. Book composition from (38) at the steady-state shares (39): the rest of the world's share and the net cross-border claim are netted from the domestic stock; with neither, as in the symmetric baseline, $B_j^{\text{book}} = B_j^g$.
3. Equity from asset-market clearing: $J_j = A_j - B_j^{\text{book}}$; after-tax dividends $d_j^c = rJ_j$ and gross dividends $d_j^c/(1-\tau_j^K)$.
4. Marginal cost from the dividend identity, using *gross* dividends (after-tax dividends in their place would deliver $mc_j > 1$):
$$mc_j = 1 - \frac{r}{Y_j}\left(\frac{J_j}{1-\tau_j^K} - K_j\right), \qquad \mu_j^p = 1/mc_j.$$
5. Capital share, wage, and technology: $\alpha_j = (r+\delta)K_j/(mc_j Y_j)$, then w_j and Z_j from (61).
6. Bond price and face quantity: $\bar{P}_j^g = 1/(1+r-\rho_j^b)$, $B_j^{\text{face}} = B_j^g/\bar{P}_j^g$.
7. Labour-tax level from the budget: with $\overline{\text{Tra}}_j$ at its assigned value, zero here and the data ratio in the calibrated union, (63) solved for λ_j (closed form for any ϑ_j, since T_j^L is linear in λ_j). With $\tau_j^c = 0$ every input is already determined and the step is recursive; with $\tau_j^c \neq 0$ the consumption-tax revenue makes λ_j a fixed point through C_j, solved jointly with step 8 (Section 3.5): in the implementation, one revenue unknown per country against the consistency condition that the revenue the budget uses equals $\tau_j^c C_j$ from the solved household problem.
8. Household solve: two instruments $\{\beta_j, \chi_j\}$ against the two targets $\{A_j, A_j^{\text{liq}}\}$ in the symmetric benchmark; in the calibrated union the liquidity wedge ξ_j joins as a third instrument against the wealthy hand-to-mouth share (Section 4.2). The average MPC, the total and poor hand-to-mouth shares, and participation are *outputs*, reported as validation. The frozen liquid share is $\varsigma_j = A_j^{\text{liq}}/A_j$.
9. $\mathcal{U}_j$ from the household distribution (or $C_j^{-1/\sigma}$ under the representative-agent convention), then $\bar{\varphi}_j$ from (17).
10. Verify asset-market clearing and, as a diagnostic, goods-market clearing.

Feasibility. Step 4 requires $J_j > (1-\tau_j^K)\,K_j$, i.e. $A_j - B_j^{\text{book}} > (1-\tau_j^K)\,K_j$: household-held debt beyond that ceiling drives μ_j^p below one and silently flips the sign of the Rotemberg cost prefactor $\mu^p/(\mu^p-1)$. The constraint is the holder-set restriction of Section 3.9 in steady-state form: it binds when a country's households must hold, through their own intermediary, a sovereign stock large relative to their wealth (at the wealth of the symmetric benchmark, an F stock of 135% of annual GDP held entirely at home is infeasible), and it is relaxed by the parts of each stock held by the other member and by the rest of the world.

Asymmetric steady state. When country primitives differ, the steady state is no longer recursive country by country: the terms of trade $\bar{S}$ is an unknown, every conversion in (1) is active, and the cross-border position $\overline{nfa}_H$ of (39) carries a steady-state income flow $r\,\overline{nfa}_H$ between the two budgets, and each country's trade surplus with the rest of the world, $r\,\omega_j^R \bar{B}_j^g$, enters its goods-market condition. Country size on its own leaves the per-capita steady state unchanged under (32) (it reweights demand, holdings, and policy), but it scales the per-capita cross-border position (39) and hence its income flow. The calibration then solves $\{\beta_H, \chi_H, \beta_F, \chi_F, \bar{S}\}$ jointly against the four wealth targets and one goods-market-clearing condition (the other redundant by Walras' law), with steps 2–7 evaluated inside the solver at the current $\bar{S}$; with $\tau_j^c \neq 0$ the two consumption-tax revenues join the unknowns against their consistency conditions (step 7), and in the calibrated union the two wedges $\{\xi_H, \xi_F\}$ join against the two wealthy hand-to-mouth shares, nine unknowns in all. The steady-state fiscal externality of debt ownership (the F servicing part of its debt to H holders, raising its tax burden and lowering after-tax income relative to H) is an output of this solve.

Appendix B Channel accounting: an exact first-order decomposition

This appendix decomposes the first-order consumption response along the equilibrium paths into earnings, transfers and the returns on the two accounts (Kaplan et al., 2018; Auclert et al., 2021), and shows how sovereign ownership changes the portfolio exposures that enter these returns.

The household response to its inputs. The country-j household block maps four aggregate inputs into consumption: the common after-tax earnings factor, the transfer, and the two returns. Individual after-tax income (5) factors as $e_{i,t} = \hat{e}_{j,t}\, z_{i,t}^{1-\vartheta_j}$ with $\hat{e}_{j,t} \equiv \lambda_{j,t}\big(w_{j,t}^c N_{j,t}\big)^{1-\vartheta_j}$, so the wage bill and the tax level $\lambda_{j,t}$ reach every household through the single factor $\hat{e}$, which keeps the decomposition exact under the labour-tax rule (24), where λ moves. To first order around the steady state, the consumption response to the *equilibrium* paths of these inputs is, exactly,

$$\mathrm{d}C_j = \mathcal{M}_j^{e}\,\mathrm{d}\hat{e}_j + \mathcal{M}_j^{\mathrm{Tra}}\,\mathrm{dTra}_j + \mathcal{M}_j^{r^b}\,\mathrm{d}r_j^b + \mathcal{M}_j^{r^a}\,\mathrm{d}r_j^a, \tag{64}$$

(consumption-tax experiments add a fifth kernel $\mathcal{M}_j^{\tau^c}$; the baseline holds τ_j^c fixed), where each $\mathcal{M}_j^x$ is the sequence-space Jacobian of the household block with respect to input x (a matrix whose (t,s) entry is the date-t consumption response to a date-s innovation in x), computed from the actual two-account Calvo problem of Appendix A.1. $\mathcal{M}_j^{\mathrm{Tra}}$ is the intertemporal-MPC matrix of Auclert et al. (2024). Three properties of (64) matter for what follows. First, the Jacobians are *input-specific*: a deposit-return, an illiquid-return,

and an earnings innovation induce different responses, because they reach different points of the wealth distribution with different timing. Second, each Jacobian already embeds the joint distribution of marginal propensities and asset positions (the covariance between who owns an asset and who consumes out of it), which products of averages omit. Third, the rebalancing probability χ_j governs when households can reach the illiquid account, and that access is already embedded in $\mathcal{M}_j^{r^a}$, together with the response of households who cannot yet rebalance but adjust their saving in anticipation; no separate factor χ_j multiplies the formula, which would count limited access twice.

The return paths. The inputs in (64) are themselves determined by the block structure of Section 3. The deposit return moves one-for-one with the realised bond return, $\mathrm{d}r_j^b = \mathrm{d}r_j^g$; the closed forms for $\mathrm{d}r^g$ are collected in Lemma 1 below. The illiquid return follows from the residual-claimant condition (19): using $r^{\text{book}} = r$ at steady state and writing $\mathrm{d}r_{j,t}^{eq}$ for the equity-leg return (dividend news plus revaluation),

$$\begin{aligned} \mathrm{d}r_{j,t}^{a} = \frac{\mathrm{d}r_{j,t}^{\text{book}} - \varsigma_j \, \mathrm{d}r_{j,t}^{g}}{1-\varsigma_j}, \qquad \mathrm{d}r_{j,t}^{\text{book}} = \frac{\bar{B}_j^{\text{book}}}{\bar{A}_j}\,\mathrm{d}r_{j,t}^{g} + \frac{\bar{J}_j}{\bar{A}_j}\,\mathrm{d}r_{j,t}^{eq}, \\ \mathrm{d}r_{j,t}^{eq} \equiv \frac{\mathrm{d}d_{j,t}^{c} + \mathrm{d}J_{j,t} - (1+r)\,\mathrm{d}J_{j,t-1}}{\bar{J}_j} : \end{aligned} \tag{65}$$

the book return is the holdings-weighted average of the bond and the equity return. The expressions are written for a book whose legs all earn the domestic bond return; with a spread or different durations the bond term is the sum over legs, $\sum_k (\bar{B}_{jk}/\bar{A}_j)\,\mathrm{d}r_{jk,t}^{g}$. Fluctuations in the bond *quantity* carry no first-order term: a marginal one-period bond earns the book's steady-state return, so $\mathrm{d}B_{j,t-1}^{\text{book}}$ cancels between the book's income and its base. This is the first-order portfolio-quantity irrelevance of Devereux and Sutherland (2011) and Tille and van Wincoop (2010), applied here to a cross-border sovereign book: portfolio *composition* matters at first order through the steady-state weights, portfolio *movements* do not. Equations (64)–(65) deliver the four channels of transmission with their timing: the *earnings* channel ($\mathcal{M}^e$, contemporaneous, largest where propensities to consume are high); the *deposit-return* channel ($\mathcal{M}^{r^b}$, whose input with one-period debt carries a policy-rate change from the following quarter, and with long debt also the bond revaluation in the impact quarter, since deposits are written on the realised bond return; anticipated returns move consumption from the impact quarter in either case); the *illiquid-return* channel ($\mathcal{M}^{r^a}$, whose input is front-loaded in asset prices); and the *transfer* channel ($\mathcal{M}^{\text{Tra}}$, the incidence of debt service and the residual instrument).

Closed-form exposures. The return side of the decomposition is exact algebra; only its consumption consequence passes through a kernel.

Lemma 1 (Return exposures)**.** *To first order around the steady state:*

(i) At $\rho_j^b = 0$, $\mathrm{d}r_{j,t}^{g} = \mathrm{d}i_{t-1} - (1+r)\,\mathrm{d}\pi_{j,t}^{c}$: the deposit return carries the policy rate with a one-period lag and the inflation surprise contemporaneously.

(ii) For $\rho_j^b \in [0,1)$, the bond price is the discounted sum

$$\frac{\mathrm{d}P_{j,t}^{g}}{\bar{P}_j^{g}} = -\frac{1}{1+r}\sum_{s\geq 0}\left(\frac{\rho_j^b}{1+r}\right)^{s}\mathrm{d}i_{t+s}, \qquad \mathrm{d}r_{j,t}^{g} = \rho_j^b\,\frac{\mathrm{d}P_{j,t}^{g}}{\bar{P}_j^{g}} - (1+r)\left[\frac{\mathrm{d}P_{j,t-1}^{g}}{\bar{P}_j^{g}} + \mathrm{d}\pi_{j,t}^{c}\right],$$

so, holding inflation fixed, a permanent unit rate path moves the impact return by $-\rho_j^b/(1+r-\rho_j^b) = -(D_j - 1)$*, with* $D_j = (1+r)/(1+r-\rho_j^b)$ *the Macaulay duration in quarters.*

(iii) On impact, with stocks predetermined,

$$\mathrm{d}r_{j,0}^a = \frac{1}{1-\varsigma_j}\left[\left(\frac{\bar{B}_j^{\mathrm{book}}}{\bar{A}_j} - \varsigma_j\right)\mathrm{d}r_{j,0}^g + \frac{\bar{J}_j}{\bar{A}_j}\,\mathrm{d}r_{j,0}^{eq}\right],$$

the equity revaluation reaching the illiquid account with the leverage $1/(1-\varsigma_j)$ *of a claim junior to deposits.*

Proof. Differentiate (6), (20), and (65) around the steady state, using $1+\rho_j^b \bar{P}_j^g = (1+r)\bar{P}_j^g$; part (iii) sets $\mathrm{d}J_{j,-1} = 0$. □

Sovereign debt ownership. At the baseline every claim is a one-period nominal instrument paying the common rate, so ownership creates no asset-specific return. Its content is pinned down by the following.

Proposition 1 (Ownership at the baseline)**.** *Assume one-period bonds, no spread, and the default closure.*

(i) (Annuity.) *At any steady state, the balances of payments* (41) *and* (40) *imply*

$$\overline{NX}_H = -\,r\,\overline{nfa}_H\,\bar{p}_H^c, \qquad \bar{X}_j^R = r\,\omega_j^R\,\bar{B}_j^g, \tag{66}$$

with $\overline{nfa}_H$ *from* (39)*: H runs a trade deficit with F equal to the annuity* $r\,\overline{nfa}_H$*, F the mirror image, and each issuer runs a trade surplus with the rest of the world equal to the coupon on its foreign-held debt. The identity itself is the standard steady-state balance-of-payments relation (Obstfeld and Rogoff, 1996); what is specific to this model is that* $\overline{nfa}_H$ *is pinned by* κ_F *and the debt stock, so the annuity is a function of the ownership primitive, and that its incidence is set by the fiscal instrument.*

(ii) (Conditional redundancy.) *Conditional on the complete calibrated steady state, including the household distributions and parameters, the book weights* $(\bar{B}_j^{\mathrm{book}}, \bar{J}_j)$*, household wealth* $(\bar{A}_j, \varsigma_j)$*, and the income and transfer levels, the intra-union ownership parameters* κ_j *appear nowhere in the first-order dynamic system: each asset's return process is unaffected by who holds it, the government budgets* (23) *contain* B_j^g *only, and fluctuations of the position drop out of the book return* (65)*. The outside share* ω_j^R*, by contrast, is a first-order parameter: it scales the rest of the world's claim on realised returns,* $\bar{B}_j^R\,\mathrm{d}r_{j,t}^g$*, which leaves the union's household sector and returns to its goods markets only through the rebalancing rule* (40)*.*

(iii) (Composition tilt.) *Between two calibrated steady states that differ by a reallocation* $\Delta\overline{nfa}$ *of F debt to H's book, the partial-equilibrium difference in the illiquid returns, holding aggregate paths fixed, is*

$$\begin{aligned}\Delta\,\mathrm{d}r_{H,t}^a &= +\frac{\Delta\overline{nfa}}{(1-\varsigma_H)\,\bar{A}_H}\left[\mathrm{d}r_{H,t}^g - \mathrm{d}r_{H,t}^{eq}\right],\\ \Delta\,\mathrm{d}r_{F,t}^a &= -\frac{s}{1-s}\,\frac{\Delta\overline{nfa}/\bar{\mathcal{Q}}}{(1-\varsigma_F)\,\bar{A}_F}\left[\mathrm{d}r_{F,t}^g - \mathrm{d}r_{F,t}^{eq}\right]:\end{aligned} \tag{67}$$

ownership tilts H's illiquid return toward the bond return and F's toward the equity return, in proportion to the bond–equity excess-return path, with H's per-capita claim spread over F's population by the size ratio. The statistic is an exposure measure in the sense of Auclert (2019), here on the cross-border sovereign book rather than on the household's own balance sheet.

Proof. (i) Evaluate (41) and (40) at the steady state and use (39). (ii) B_j^g and nfa_H enter the blocks exactly as stated, and ω_j^R enters (40) through the target and $\bar{B}_j^R$; the bond-quantity cancellation is shown below (65). (iii) Differentiate (65) with respect to the weights, using asset-market clearing at fixed $\bar{A}_j$, $\Delta\bar{B}_j^{\text{book}} + \Delta\bar{J}_j = 0$, with $\Delta\bar{B}_H^{\text{book}} = \Delta\overline{nfa}$ and $\Delta\bar{B}_F^{\text{book}} = -\frac{s}{1-s}\,\Delta\overline{nfa}/\bar{\mathcal{Q}}$ from (38). □

Corollary 1 (Composition irrelevance)**.** *If the two legs' return responses coincide,* $\mathrm{d}r_j^{eq} = \mathrm{d}r_j^g$*, the weights drop out of* (65) *and transmission is invariant to ownership conditional on the steady-state levels. At a symmetric steady state of two countries of equal size, moreover, the direct asset-return effect of a marginal reallocation on union aggregate consumption cancels at first order: the two tilts in* (67) *are equal with opposite sign and pass through identical kernels* $\mathcal{M}^{r^a}$*.*

At the baseline the two legs differ exactly by the equity revaluation, so the dynamic content of ownership is the same bond-versus-equity composition margin that a change in the domestic debt level moves: exchanging identical sovereign claims between issuers leaves a holder's return unchanged at a given total bond exposure, so ownership creates no return process of its own. It locates the steady-state part of the mechanism in the annuity (66), whose household incidence is set by the fiscal instrument at calibration: uniform under the transfer, earnings-proportional under the tax level (Appendix A.5). Ownership therefore acts through the annuity and through the split of each book between bonds and equity, and, when issuer returns differ (country-specific duration, or a spread), through those returns as well, which the gross holder-by-issuer book of Section 3.9 accommodates and Appendix D.3 measures. Holding by the rest of the world is different in kind: it is a first-order leakage even at the baseline, part (ii).

A two-agent benchmark. The decomposition also separates what household heterogeneity contributes, and a two-agent benchmark in the spirit of Bilbiie (2008) makes the separation analytical. It illustrates the mechanism; it is not a version of the model. Hold prices, rates, earnings, dividends and government spending at their steady-state values, and let a share h_j of country j's households be hand-to-mouth, spending their income, while the others are savers who discount at the steady-state rate r, own the intermediary's book and hold no deposits. Transfers are the same for every household. Let $\mathcal{T}_{j,t}$ be an additional payment by government j, financed by an equal cut in aggregate transfers, and $\mathcal{R}_{j,t}$ the additional receipts of country j's savers on its intermediary's book, both country totals in a common unit, and write $\mathcal{A}[x] \equiv \frac{r}{1+r}\sum_{t\geq 0}(1+r)^{-t}x_t$ for the annuity value of a path.

Proposition 2 (Fiscal burden and creditor income)**.** *The impact response of country j's consumption is*

$$\mathrm{d}C_{j,0} = -h_j\big(\mathcal{T}_{j,0} - \mathcal{A}[\mathcal{T}_j]\big) + \mathcal{A}[\mathcal{R}_j - \mathcal{T}_j]. \tag{68}$$

Proof. Hand-to-mouth households lose $h_j\mathcal{T}_{j,t}$ in every period and spend it. Savers lose $(1-h_j)\mathcal{T}_{j,t}$ and receive $\mathcal{R}_{j,t}$; since they discount at the rate they earn, their consumption change is the same in every period, and their intertemporal budget sets it to $\mathcal{A}[\mathcal{R}_j - (1-$

$h_j)\mathcal{T}_j]$. Adding the two responses and writing $\mathcal{R}_j - (1-h_j)\mathcal{T}_j = (\mathcal{R}_j - \mathcal{T}_j) + h_j\mathcal{T}_j$ gives (68). □

The first term moves income from constrained households to savers, who spread it over time; the second is the annuity of the country's net creditor receipts.

Corollary 2 (Domestic and cross-border creditors)**.** *(i) If all of the payment returns to domestic savers, $\mathcal{R}_j = \mathcal{T}_j$, and it decays geometrically, $\mathcal{T}_{j,t} = \varrho^t \mathcal{T}_{j,0}$, then $\mathrm{d}C_{j,0} = -h_j\,[1 - m(\varrho)]\,\mathcal{T}_{j,0}$ with $m(\varrho) = r/(1+r-\varrho)$. The effect vanishes with a single household, $h_j = 0$, for any path, and on impact when the payment is permanent.*

(ii) If savers in country j receive the share ϖ_{jk} of issuer k's payment and holders outside the union the share ϖ_{Ok}, with $\varpi_{Hk} + \varpi_{Fk} + \varpi_{Ok} = 1$ and a common ϱ, then

$$\begin{aligned} \mathrm{d}C_{j,0} &= -\big[h_j + (1-h_j)\,m(\varrho)\big]\mathcal{T}_{j,0} + m(\varrho)\sum_k \varpi_{jk}\,\mathcal{T}_{k,0}, \\ \mathrm{d}C_{H,0} + \mathrm{d}C_{F,0} &= -\sum_k \big[h_k\,(1-m(\varrho)) + m(\varrho)\,\varpi_{Ok}\big]\mathcal{T}_{k,0}. \end{aligned} \tag{69}$$

Moving a share ε of F's payment from F's savers to H's raises H's impact consumption and lowers F's by the same amount, $m(\varrho)\,\varepsilon\,\mathcal{T}_{F,0}$, and leaves the union's total unchanged.

Proof. Substitute $\mathcal{R}_j = \mathcal{T}_j$, or $\mathcal{R}_{j,t} = \sum_k \varpi_{jk}\mathcal{T}_{k,t}$, into (68) and sum the geometric series, $\mathcal{A}[\mathcal{T}_j] = m(\varrho)\,\mathcal{T}_{j,0}$; the union's total uses $\sum_j \varpi_{jk} = 1 - \varpi_{Ok}$. □

With a single household the cancellation holds exactly in the country's budget, along any path: the coupon and revaluation on the domestically held stock are paid once through the transfer and received once through the book, so the debt reaches consumption only through the market value of the net external sovereign position and through the spread, which is paid on the whole stock and is a deadweight loss on the part held at home, as in the representative-agent benchmark of Auclert et al. (2024). The fiscal burden follows the issuer and the creditor income follows the holder: through this channel ownership moves the burden between countries and leaves the union's total unchanged, as Corollary 1 finds for the asset-return channel. The benchmark isolates the cash flows of debt service; the ownership comparisons of Section 5.3 also move the composition of the intermediaries' books and the steady state, which the quantitative model accounts for.

In (64), a single household responds in the same way to a transfer and to the same income credited to its wealth, so the part of the transfer cut that services the debt held at home and the income that debt pays on the book offset. With heterogeneous households they do not: the transfer is the same for every household, the bond income accrues to deposits one for one and to the illiquid account with the leverage of (19), and the transfer channel carries the first while the two return channels carry the second, together with the response to anticipated returns and to equity prices. Evaluated on the model's equilibrium transfer path, with h_j the share of households at the liquid floor, the benchmark's transfer leg, $-h_j\mathcal{T}_{j,t} - (1-h_j)\mathcal{A}[\mathcal{T}_j]$, reproduces the transfer channel of Table 4 on impact within 8 per cent and its contribution to the gap within 7 per cent. It attributes that response almost entirely to the households at the floor, whereas in the model households above the floor carry two fifths to a half of it on impact and about three quarters of its five-year total. The realised payment is front-loaded, so savers spend only about 3 per cent of it in each quarter, consistent with an ownership increment that is small on impact and builds up over time (Table 5).

Relation to general equilibrium. Equation (64) is an accounting of the household block along the *equilibrium* paths; it is not a solution method. The equilibrium paths themselves are obtained from the full sequence-space system of Section 3.10, in which trade and expenditure switching, investment, both Phillips curves and the fiscal rules all feed back. The decomposition is validated column by column against the model's household Jacobians, so the channel shares it reports are exact, and any statement about their relative size at a given horizon is a statement about the solved model, not about an approximation to it.

Appendix C Local projections

This appendix compares the model's consumption gap with local projections (Jordà, 2005) on euro-area data. For each horizon $h = 0, \dots, 8$ we regress $100\,[\log C_{i,t+h} - \log C_{i,t-1}]$, where $C_{i,t}$ is real consumption per capita in country i, on the monetary policy shock of quarter t interacted with an indicator for the group F. The regressions also include the central-bank information shock interacted with the same indicator, country and quarter fixed effects, four lags of consumption growth, and one lag each of inflation, unemployment and debt to GDP. The quarter fixed effects absorb the common response, so the coefficient is the difference between the two groups' responses. The eight countries and their grouping are those of Section 2.5, each country weighted equally. The shocks are the updated euro-area monetary policy and information shocks of Jarociński and Karadi (2020), built from the announcement-window surprises of Altavilla et al. (2019) and summed within quarters; the sample runs from 2000Q2 to 2023Q3. We report both of their decompositions: the median rotation, which lets both shocks occur at one announcement, and the simpler sign restriction that assigns each announcement to one of them, the poor man's. Intervals come from a circular block bootstrap with blocks of nine quarters drawn jointly for all countries.

The model's shock is an innovation to the policy rule, and its policy rate falls on impact because inflation falls by more than the rule's intercept rises (Section 5.1), so the size of the shock has no direct counterpart in an observed rate surprise. The real rate does: it rises by 8.4 basis points on average over the model's first four quarters. Within the announcement window expected inflation barely moves, so a surprise in the one-year overnight index swap rate is close to a surprise in the one-year real rate. We therefore scale each empirical shock to the size that raises the one-year swap rate by 8.4 basis points, using its loading on that rate estimated across announcements. Scaling instead by the three-month rate and the model's impact real rate of 11.2 basis points raises every estimate by about two thirds.

Table 7 reports the results in the convention of the paper, H minus F, so a positive number means that F's consumption falls by more. On impact the difference is 0.14 percentage points under either decomposition, against 0.11 in the model. The poor man's estimate is the more precise; it differs from zero at the 10 per cent level or better in every sample in the table, and the model's value lies inside every interval. Under the poor man's decomposition the estimate rises with later starting dates, with wider intervals. Summed over the first nine quarters the estimates are positive, as in the model, but their intervals are wide and include zero: the data do not identify the cumulated difference. The comparison concerns the difference between the two groups as a whole; it does not isolate the ownership channel, which accounts for 0.016 of the model's impact difference (Table 5).

	N	T	Median rotation	Poor man's	Model
Impact					
Full sample, 2000Q2–2023Q3	743	94	+0.136	+0.140**	+0.107
			[−0.049, +0.290]	[+0.027, +0.290]	
Excluding 2020–21 shock dates	679	86	+0.101	+0.148**	
			[−0.049, +0.218]	[+0.024, +0.241]	
Excluding 2020–21 response windows	671	85	+0.108	+0.151**	
			[−0.043, +0.222]	[+0.027, +0.243]	
From 2005	591	75	+0.186	+0.225***	
			[−0.053, +0.389]	[+0.094, +0.425]	
From 2010	431	55	+0.058	+0.261**	
			[−0.267, +0.409]	[+0.040, +0.577]	
From 2015	280	35	+0.144	+0.353*	
			[−0.363, +0.296]	[−0.062, +0.678]	
Summed over quarters 0 *to* 8					
Full sample, 2000Q2–2023Q3	735	94	+2.517	+1.833	+0.388
			[−2.705, +6.333]	[−2.152, +6.578]	
Excluding 2020–21 shock dates	671	86	+1.414	+1.198	
			[−3.452, +6.073]	[−2.715, +5.027]	
Excluding 2020–21 response windows	599	77	+1.166	+0.898	
			[−4.142, +5.838]	[−3.150, +4.623]	

Table 7: Consumption responses to a common monetary tightening, H minus F: local projections and the model. Impact values are percentage points; summed values are percentage-point-quarters over quarters 0 to 8. Each empirical shock is scaled to raise the one-year overnight index swap rate in the announcement window by 8.4 basis points, the model's average real rate over its first four quarters: 1.63 quarterly standard deviations (2000Q2–2016Q4) of the median-rotation shock and 1.36 of the poor man's. The response-window exclusion drops observations whose window from $t-1$ to $t+h$ meets 2020–21. Brackets are pointwise 95 per cent intervals from 1,999 circular block-bootstrap draws with blocks of nine quarters; ***, ** and * mark 99, 95 and 90 per cent percentile intervals that exclude zero. N counts country-quarters and T quarters. The model's groups are GDP-weighted and its shock arrives at the start of the quarter; in the data countries are weighted equally and announcements fall within the quarter.

Appendix D Asymmetries one at a time

This appendix measures each difference between the two countries on its own, against a symmetric benchmark in which the two countries are exact copies, and so identifies the mechanism behind each part of the attribution of Section 5.3.

D.1 The symmetric benchmark

Both country blocks carry the single-economy euro-area calibration: log utility, a single discount factor $\beta = 0.9825$ and rebalancing probability $\chi = 0.215$ calibrated to HFCS financial wealth of 11.40 and liquid wealth of 1.80 (quarterly GDP units), with the quarterly MPC (0.14) and the hand-to-mouth share (11.5 per cent) untargeted outputs. The sovereign stock is 65 per cent of annual GDP ($B^g = 2.60$ quarterly), a realistic euro-area level, in both countries, at one-period maturity, fully domestically held, with no spread. The labour-tax level clears the budget ($\bar{\tau} = 0.30$, $\lambda = 0.70$) with zero steady-state transfers, the λ-as-outcome normalisation of Appendix A.5, here at the flat-tax point $\vartheta = \tau^c = 0$. Trade carries a home bias of $\gamma = 0.875$ in both countries, the value implied by an openness of $\theta = 0.25$ when the two are equal in size, and an elasticity $\eta = 1.5$. Monetary policy sets $\phi_\pi = 1.5$ on union CPI inflation, and the two fiscal rules are normalised to the same surplus response ($\phi^B = 0.1$), so that the tax-versus-transfer contrast compares instruments rather than rule strengths. At this benchmark the two countries are exact copies and the union replicates the one-country economy; every experiment below is a controlled deviation from it.

D.2 Design of the experiments

Each asymmetry is introduced *one at a time*, at two sizes: a small perturbation, which checks how the effect scales with the underlying difference, and the empirical pole gap, which measures its magnitude. Asymmetries that move the steady state, ownership among them, are re-calibrated jointly, so that the terms of trade, the external position and each country's (β_j, χ_j) solve together rather than being imposed on an unchanged steady state. Every configuration is subjected to a common monetary shock, a 25 basis-point annualised innovation to r^* with AR(0.85), and to *country-specific* productivity and government-spending shocks (-1 per cent TFP and $+1$ per cent of GDP, AR(0.9), the latter under both fiscal rules). The common shock measures asymmetric transmission proper, while the country-specific shocks identify cross-border spillovers, which a common shock cannot separate from domestic propagation. Consumption responses are decomposed into the four channels of Appendix B, namely earnings, transfers, the deposit return and the illiquid return. Gaps are reported as H minus F, so a positive number means F contracts by more.

D.3 Sovereign debt: maturity, spreads, level, and ownership

Table 8 collects the sovereign experiments under the common monetary shock, against a benchmark impact response of about -0.26 per cent of consumption in both countries.

Maturity. The duration gap shields the shorter-maturity country. The tightening lowers inflation and the path of nominal rates, which raises the real value of long debt by more the longer its maturity, so the shorter-maturity government absorbs a smaller rise in debt service on impact. The effect is an order of magnitude or more below the ownership

	impact	cumulative (20q)
Debt maturity (9.7y vs 7.8y)	−0.004	−0.008
Sovereign spread (F, debt-elastic)	+0.001	+0.011
Debt ownership ($\kappa_F = 0.75$, equal debt)	+0.022	+0.282
Debt level and ownership ($B^g_F = 135\%$, $\kappa_F = 0.75$)	+0.119	+0.874

Table 8: Asymmetric transmission of a common monetary tightening (consumption gap, $H - F$: percentage points on impact, percentage-point-quarters cumulated). Each row is one experiment at its empirical pole gap, all other primitives symmetric, at the common home bias of Appendix D.1 and the 25 basis-point annualised shock. Positive means F contracts by more. The single-margin effects are close to linear in the underlying difference: at one tenth of its gap each is a tenth as large, within four per cent.

channels and scales linearly with the gap. A two-year duration difference on a base of eight moves the consumption gap by well under one basis point, and it moves it in the direction that makes H, the longer-maturity country, contract by more. On its own the maturity gap is not a quantitatively important source of asymmetric monetary transmission, which is why the calibration of Section 4 gives both countries a common duration.

Sovereign spreads. Under the deadweight specification the spread is a fiscal cost, and the experiments confirm that it transmits as one. A common monetary tightening moves F's spread only through the debt ratio's denominator, and the induced asymmetry is negligible. The spread's action is in debt-financed fiscal episodes rather than in monetary transmission: when F expands spending under the labour-tax rule, the rollover cost of its debt-elastic spread adds cumulative consumption divergence per episode, and that cost roughly doubles once a quarter of F's debt is held by H, because part of the higher debt service then leaves the country.

Ownership and the debt level, separated. Holding sovereign stocks equal and moving only the domestically held share ($\kappa_F = 0.75$) isolates the ownership margin. In steady state ownership is an annuity: H absorbs $r \cdot \overline{nfa}$ of F's output each period, a trade deficit of 0.8 per cent of quarterly GDP at the pole, and after-tax incomes diverge by twice that flow (Proposition 1). In transmission, ownership contributes about one fifth of the combined level-and-ownership impact gap but close to a third of the cumulative gap. The annuity persists, so ownership matters more the longer the horizon. Its cumulative-to-impact ratio is among the highest of any margin here, though not uniquely so: the government-spending margin is comparably back-loaded. The remainder is the debt-level margin, which operates through the size of the bond books and the fiscal burden. The channel decomposition shows the composition tilt of Proposition 1(iii) at work: F's equity-tilted book makes its consumption response load more on the illiquid return, 14.1 per cent against 11.1 on impact, while H's bond-tilted book loads more on earnings, 37.1 per cent against 35.6.

Interactions. The channels are not additive, and the pattern is informative. Ownership roughly doubles the spread channel, which is the fiscal externality at work: F's debt dynamics reach H's budget through the cross-border book. The maturity gap behaves differently. Adding it to the ownership configuration raises the cumulative gap from +0.282 to +0.509, but almost all of that comes from the common *level* of duration rather than from the gap between the two countries: at a realistic common maturity of 8.75 years the revaluations are larger, and ownership has more to reallocate across the border. For

calibration this means the common duration must be realistic even where the cross-country duration gap is dropped, which is what Section 4 does.

Spillovers. The country-specific shocks deliver a sign reversal that belongs to fiscal policy rather than to any balance-sheet asymmetry. The country that expands spending contracts *less* than its neighbour on impact, since it receives the spending income while both face the common monetary response, and *more* cumulatively as its consolidation burden arrives.

The pattern is not universal, and the exception is instructive. It holds in every configuration built on one-period debt, but it fails in all three configurations with long bonds: the maturity experiments and, importantly, the calibrated union of Section 5.1, where the cumulative gap after an H spending expansion stays positive at +1.98 rather than turning negative. With long-duration debt, the bonds issued to finance the expansion are themselves revalued, and those mark-to-market movements offset enough of the later consolidation that the burden never dominates the stimulus. The reversal is therefore a short-maturity phenomenon.

D.4 Fiscal structure

With the tax structure a set of country primitives (Section 3.5), we measure four fiscal-structure asymmetries in isolation, each under the protocol of Appendix D.2: the *progressivity* gap ($\vartheta_H = 0.092$ vs. $\vartheta_F = 0.032$, the GDP-weighted bloc aggregates of the labour-income HSV estimates of Qiu and Russo), the *consumption-tax* gap ($\tau^c_H = 0.192$ vs. $\tau^c_F = 0.180$, implicit tax rates on consumption), the *capital-tax* gap ($\tau^K = 0.36$ vs. 0.30, an illustrative ±3pp split around the common 0.33), and the size of the *tax-and-transfer system*: targeted steady-state transfers financed by the labour-tax level, sized so that each country's implied average labour-tax rate matches its own implicit rate on labour in the data ($\overline{\text{Tra}}_H = 0.042$ vs. $\overline{\text{Tra}}_F = 0.067$ in units of quarterly GDP, so F carries the larger system). For the three tax gaps the *common* level is held at the data mean in both sizes and only the gap scales; the fiscal-size experiment scales from zero. In every case the steady-state budget equalises after-tax labour income across countries by construction, since the labour-tax level absorbs the revenue difference, so the experiments isolate how the *composition* of taxation shapes transmission rather than mechanical income-level differences. Table 9 collects the four under the common monetary shock.

instrument	impact	cum. 20q	contracts by more
Size of the tax-and-transfer system	+0.014	+0.086	larger system (F)
Capital tax	+0.008	+0.042	lower τ^K (F)
Progressivity	−0.008	−0.034	more progressive (H)
Consumption tax	+0.001	+0.004	neither

Table 9: Fiscal-structure asymmetries under a common monetary tightening (consumption gap, $H - F$: percentage points on impact, percentage-point-quarters cumulated). For reference, the sovereign poles of Table 8 are debt level +0.097, ownership +0.022, maturity −0.004 and spread +0.001 on impact. The fiscal-size channel is approximately linear: scaling the asymmetry tenfold scales the response by 10.5 at impact and 11.6 cumulatively.

The one-at-a-time exercises yield three findings.

First, the size of the tax-and-transfer system is the dominant fiscal-structure asymmetry, though it is not the equal of the sovereign margins. At impact it is about two

thirds of the ownership channel and one seventh of the debt level. The mechanism is fiscal incidence: the country with the larger transfer system exposes more household income to the residual transfer that absorbs debt service under the default closure, so a rise in debt service reaches its households more directly.

Second, progressivity and the capital tax are second-tier but genuine channels, both about twice the maturity gap. The more progressive country responds *more* to a common tightening, because compressed net earnings raise the hand-to-mouth mass and the transfer cut therefore bites harder; Bouabdallah et al. (2025) find that progressivity likewise shapes the response to transfer-financed fiscal expansions. The capital-tax asymmetry leaves households' steady state untouched, since the budget equalises after-tax labour income and the wealth targets pin the rest, and it operates entirely through the firm side via the implied markup. Note that this experiment uses an illustrative symmetric split around a common rate rather than the per-country rates of Section 4, so it measures the channel rather than the observed asymmetry.

Third, the consumption-tax gap has a negligible effect at the observed difference in rates, a little over one percentage point on a common rate of nineteen: the resulting consumption gap is at the resolution of the experiment. The exercise is silent on the level of the tax.

One result concerns spillovers rather than common shocks. Fiscal size is the only fiscal-structure margin that reshapes them: the impact gap following a country-specific productivity shock changes sign between the small and the empirical asymmetry, from -0.037 to $+0.028$ for a shock originating in H. With a large transfer system, households in the country that was not shocked are partly insulated on impact through their transfer income. The three tax gaps leave spillover asymmetries entirely to the shock itself.

D.5 Household balance sheets

The household block enters the calibrated gap through the Shapley value of Section 5.3, $+0.107$ cumulated, second to the sovereign block and about one and a half times the size of the fiscal and size blocks, which work the other way. This subsection isolates the margin responsible for it.

The relevant asymmetry is not wealth or portfolio composition as such but the liquidity wedge of Section 4.2, the gap between the return on the illiquid account and the return on deposits. Moving it alone, from a common value to the calibrated poles of 1.1 and 2.8 per cent a year, produces a consumption gap of $+0.044$ on impact and $+0.204$ cumulated. That makes it the largest single non-sovereign channel in the paper, three times the fiscal-size margin at impact and comfortably above every tax asymmetry.

A larger wedge makes holding deposits more costly, so households with illiquid wealth run their liquid balances down and more of them sit at the constraint. F's wedge is two and a half times H's, so more F households are liquidity-constrained and a given movement in income or transfers translates more directly into consumption. This is why the untargeted hand-to-mouth shares come out ordered as they do in Section 4, and why the common-wedge robustness row of Appendix E.4 removes so much of the gap: replacing the two wedges with a single value does not merely switch the asymmetry off, it reverses the hand-to-mouth ordering.

The horizon profile is the same as the aggregate's. The wedge gap is largest on impact and decays steadily, reaching zero around the thirteenth quarter and leaving a negative tail worth about seven per cent of the impact value. There is no impact-versus-cumulative reversal here, in contrast to the country-specific fiscal experiments of Appendix D.3, where the country that expands spending contracts less on impact and more cumulatively. Under

a *common* shock the household margin simply front-loads: the constrained households in F are hit hardest when the shock arrives and the difference erodes as balances are rebuilt.

One margin does change sign under a common shock, and it is not a household one. Moving the rest-of-world holder alone to its empirical asymmetry leaves an impact gap of $+0.0009$, which is nil, against a cumulative gap of -0.040: the per-quarter difference turns negative in the second quarter and settles near -0.003. The outside holder therefore acts with a delay rather than on impact, which is consistent with its negative increment in the comparison of Section 5.3 and with the accounting: debt service leaking out of the union is a flow that accumulates, not a revaluation that lands at once.

The wedge is not the whole household block, and the rest of it works the other way. Moving the wealth and deposit targets to their observed values while holding the wedge common, which isolates what is left once the wedge is taken out, produces a gap of -0.017 on impact and -0.019 cumulated. The sign is the point. Households in F hold more wealth and more deposits relative to output, 16.1 against 12.1 and 1.59 against 1.42 in quarterly GDP units, so they meet a common shock better buffered and adjust consumption less. That pushes H to contract by more, against the wedge, which pushes the other way by leaving F's households with less liquidity relative to the wealth they hold.

The offset is front-loaded. On impact it is about two fifths the size of the wedge's gap and cumulatively about a tenth, because the composition effect changes sign around the eighth quarter, whereas the wedge's gap stays positive for about three years. This is consistent with the household block contributing less in the attribution of Section 5.3, $+0.107$, than the wedge on its own, $+0.204$; the two exercises hold the common wealth levels at different points, so their sizes do not add up. The cumulative effect is also not proportional to the size of the difference: at one tenth of the observed differences in wealth and deposits it is $+0.005$, of the opposite sign to the -0.019 at the full difference.

Appendix E Further results and robustness

This appendix collects the responses to country-specific shocks, the distributional responses behind Section 5.2, the exit of foreign creditors from F's debt, and the robustness of the gap.

E.1 Impulse responses at the calibrated model

This section collects the impulse responses of the calibrated union to the country-specific shocks. The common monetary tightening, which is the paper's main experiment, is in the text as Figure 5 for the consumption gap and Table 4 for the rates, the budget and the channel decomposition.

Each figure overlays two mirror experiments, the same shock in H and in F, and reports ten variables for both countries: output, consumption, investment, the real wage, hours, producer-price inflation, the real rate, the household return on the illiquid account, and the two account stocks; the spending experiments add government debt. Dark lines are H and gold lines F. Solid lines are the response of the country that is shocked and dashed lines that of the other country, so in a symmetric union the two solid lines would coincide, as would the two dashed ones, and the distance between them is the asymmetry. Responses are percentage deviations from the steady state, except inflation and the two rates, which are in percentage points. All shocks are AR(0.9) and country-specific, which is what identifies cross-border spillovers: a common shock cannot separate them from domestic propagation.

Figure 9 reports the productivity shocks, and Figures 10 and 11 the spending expansions under the labour-tax rule and under the transfer rule.

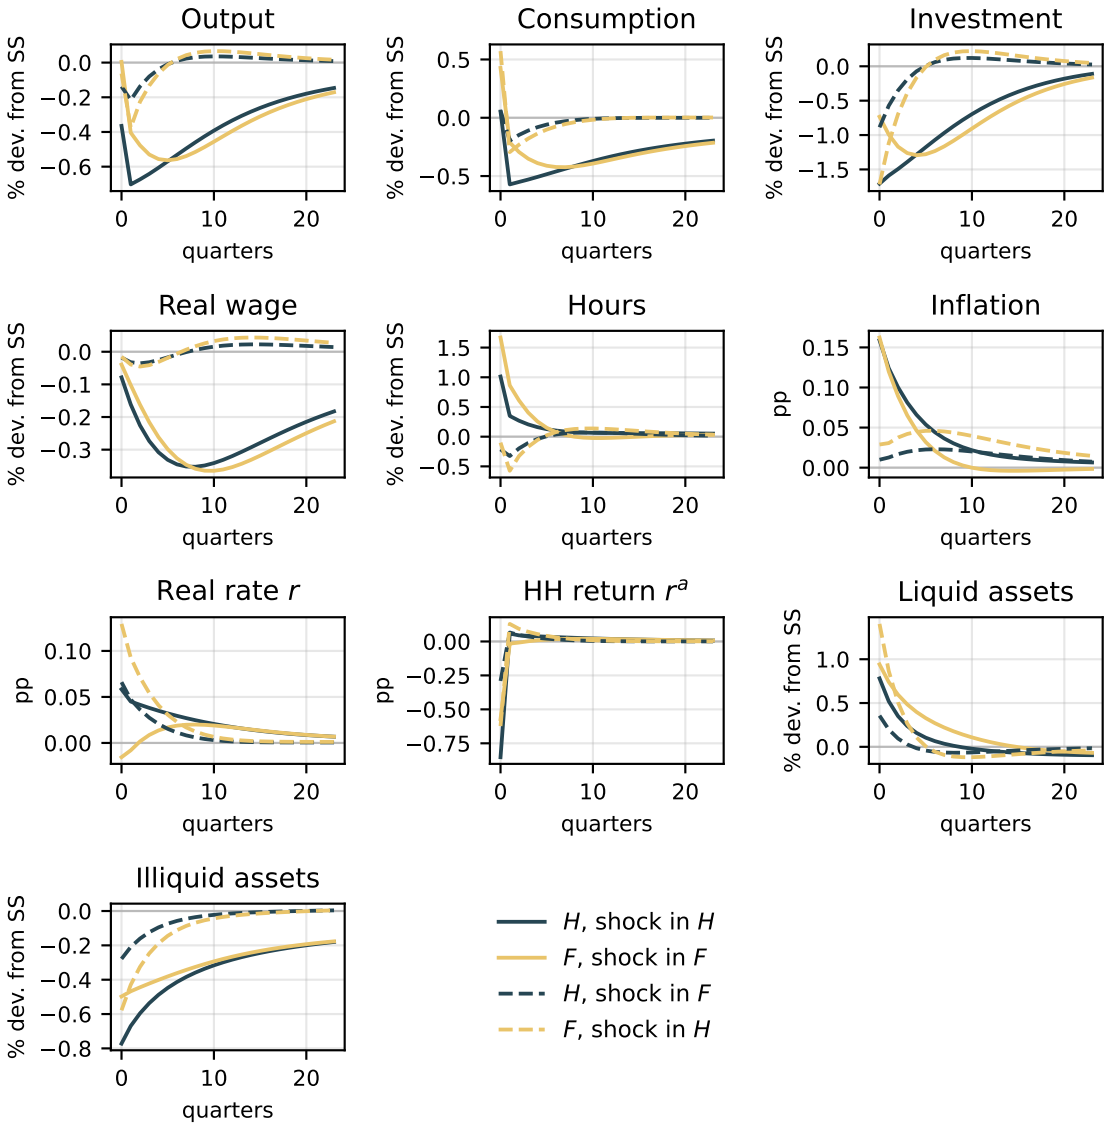


Figure 9: Productivity falls in each country. A -1 per cent AR(0.9) innovation to total factor productivity in H and, separately, in F. Dark lines are H and gold lines F; solid lines are the response of the country that is shocked and dashed lines that of the other country. From the second quarter on, output and the real wage fall by most in the shocked country, while investment falls on impact by as much or more in the other, which imports the union's monetary response, since the common rate rises with the shocked country's inflation, without the underlying shock. Consumption rises on impact in both countries after either shock, and by more in F: the rise in the common rate lowers the value of outstanding long debt, which under the constant-market-value closure lowers debt service on impact and raises the transfer, by more where the stock is larger, and that transfer response outweighs households' own revaluation losses in the first quarter. The cumulated consumption gap is -6.6 percentage-point-quarters after the shock in H and $+5.4$ after the shock in F: the sign belongs to which country is shocked, and the difference in size to the asymmetries between them.

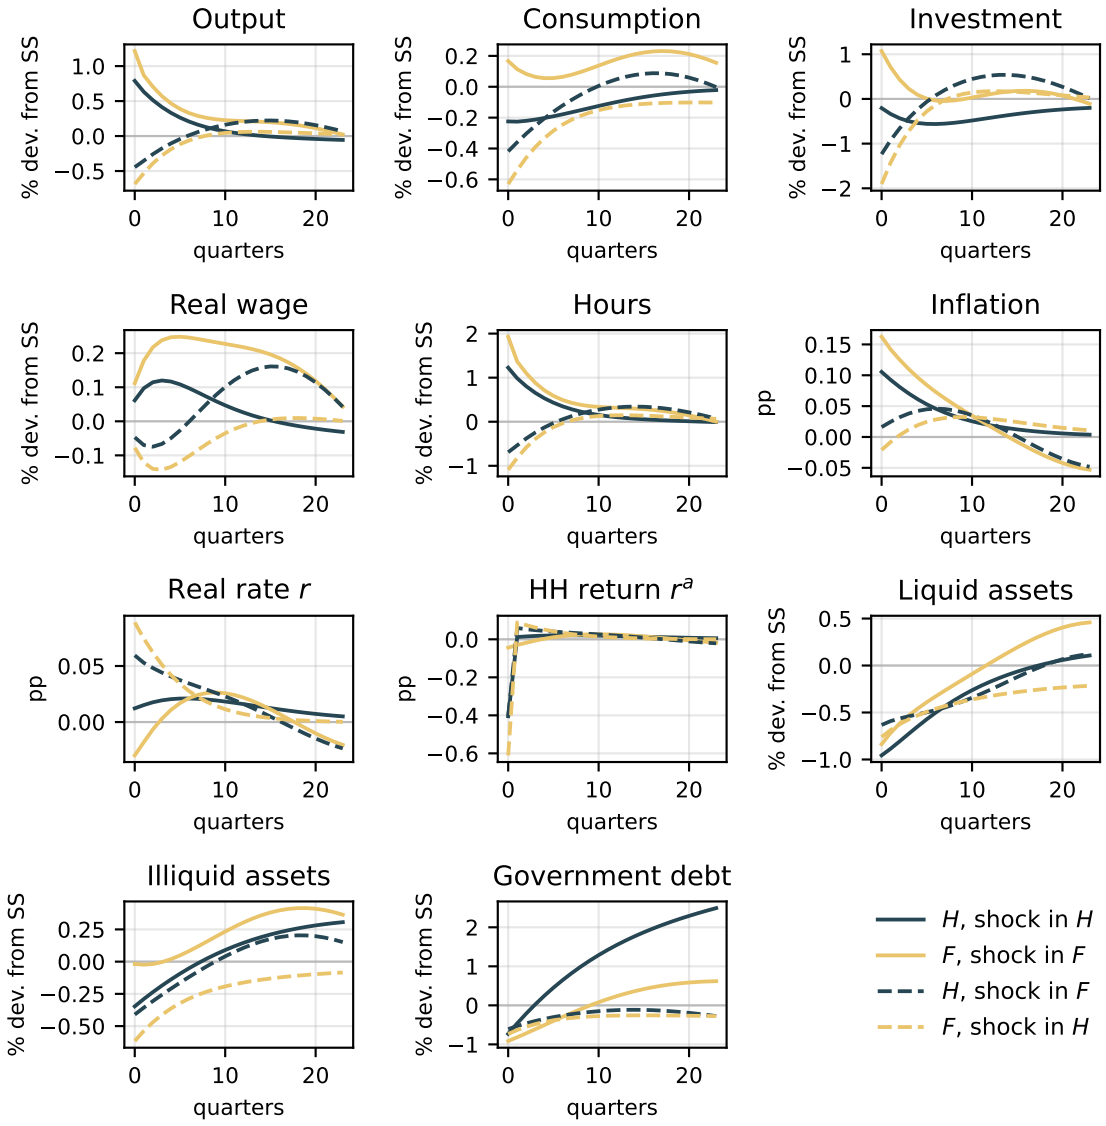


Figure 10: Government spending rises in each country, labour-tax rule. A +1 per cent of GDP AR(0.9) increase in government consumption in H and, separately, in F, consolidated later through the labour-tax level. Lines as in Figure 9. The shocked country receives the spending income while both face the common monetary response; after the shock in F, F's debt-elastic spread also operates on its own borrowing.

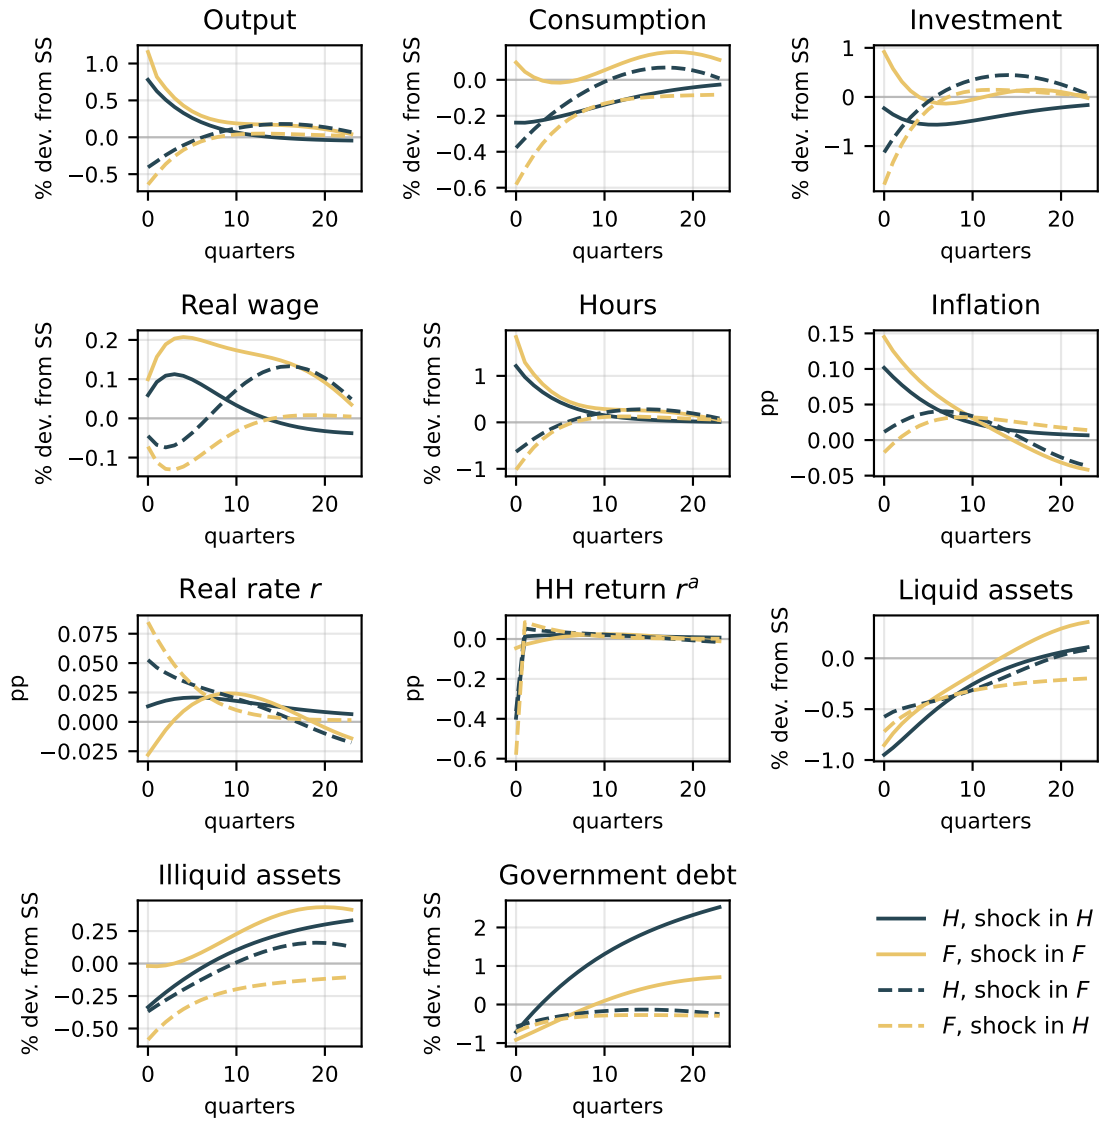


Figure 11: Government spending rises in each country, transfer rule. As in Figure 10, with the lump-sum transfer rather than the labour-tax level as the consolidating instrument; lines as in Figure 9. The consolidation now reaches hand-to-mouth households directly through their transfer income, so after the shock in H, H's consumption falls by slightly more on impact than under the tax rule, -0.24 against -0.22 per cent, while F's falls by less, -0.58 against -0.63. After the shock in F the pattern is the same: F's consumption rises by less on impact than under the tax rule, $+0.10$ against $+0.17$ per cent, and H's falls by less, -0.38 against -0.42.

E.2 Distributional responses

This section gives the method behind Figures 6, 7 and 8, and the numbers the figures do not show.

Method. Households are grouped at the moment the shock hits: by net worth, the sum of the two accounts, into fifths of each country's own distribution, with households at a common value split in proportion so that each fifth holds exactly a fifth of households; and by liquidity, into the poor hand-to-mouth (both accounts at their floors), the wealthy hand-to-mouth (the liquid account at its floor, illiquid wealth above it) and the rest, which are the model's own hand-to-mouth indicators. Each group is then followed through time. Its slice of the steady-state distribution is carried forward with the transitions implied by the households' decisions along the equilibrium paths of the household inputs, which each country's household block takes from the general-equilibrium solution. A group's response is its consumption along that path minus its consumption along the path without the shock, which is the same households' counterfactual, since the poorest fifth drifts up the distribution even without a shock, expressed in per cent of the latter. The responses are first order, like everything else in the paper: the household problem is solved along a thousandth of the equilibrium input paths and the response scaled back. The channels are those of Appendix B: earnings (the wage, hours and the tax level), transfers, the deposit return and the illiquid return.

Three checks tie the groups to the rest of the paper. They add up to the aggregate responses of Table 4 within five parts in a million; their channel split reproduces that table's shares to within 0.05 percentage points; and the hand-to-mouth masses and account Ginis at the shock are those of Table 3. For the ownership increment the same exercise is run in the union without intra-union holdings of Table 5. Each group is ranked within its own economy, and the groups are nearly identical across the two: the poorest fifth's mean net worth in H is 0.079 of quarterly GDP in one and 0.080 in the other.

Table 10 describes the groups and their share of each country's consumption fall. Table 11 gives their contribution to the aggregate gap and to its ownership increment, weighting each group by its share of consumption. Table 12 splits each group's response into the four channels, and Table 13 does the same for the ownership increment. Table 14 repeats the exercise in the economy in which the level of the labour tax balances each budget and transfers stay at their steady-state value.

Inequality measures. The consumption Gini is the measure that moves: its path is in Figure 7. The share of households at the liquid constraint rises in the quarter after the shock, by 0.45 percentage points in H and 0.99 in F, and is back to about a tenth of a point above its steady state after five years; in F most of the rise is among the wealthy hand-to-mouth. The Ginis of the two accounts barely move, by at most 0.18 points for the liquid account and 0.05 for the illiquid account on the same scale: a revaluation of the illiquid account is close to proportional across its holders, and the liquid account moves only through the households near its floor.

Nonlinearity. Along the equilibrium input paths of the full 25 basis points, the households' own decisions are not linear: the liquid floor makes the aggregate impact response 0.4 per cent larger in H and 1.8 per cent larger in F than at first order, and F's wealthy hand-to-mouth respond 9 per cent more, because a finite shock pushes households near the constraint onto it. This points to amplification of larger tightenings through the group

	net worth	hand-to-mouth	share of consumption	share of the fall	
				impact	20 quarters
Panel A: country H					
Bottom fifth	0.08	63	7	40	26
Second fifth	1.39	8	12	15	23
Middle fifth	5.10	4	17	14	20
Fourth fifth	12.72	2	23	15	19
Top fifth	40.97	1	40	17	13
Poor hand-to-mouth	0.00	100	3	30	13
Wealthy hand-to-mouth	3.87	100	4	8	7
Not hand-to-mouth	14.05	0	93	62	80
Panel B: country F					
Bottom fifth	0.17	61	7	36	27
Second fifth	2.20	13	12	16	23
Middle fifth	7.02	6	16	15	20
Fourth fifth	16.60	5	23	16	19
Top fifth	54.70	2	42	17	12
Poor hand-to-mouth	0.00	100	2	24	11
Wealthy hand-to-mouth	5.24	100	6	15	13
Not hand-to-mouth	18.93	0	92	61	77

Table 10: Household groups at the shock and their share of the consumption fall. Groups are fixed at the moment of the shock. Net worth is the group's mean holding of the two accounts, in units of quarterly GDP. Hand-to-mouth is the per cent of the group at the liquid floor, and the share of consumption the group's per cent of the country's consumption. The share of the fall is the group's per cent of the country's aggregate consumption response, on impact and summed over quarters 0 to 19. The five fifths add up to 100, as do the three liquidity groups; entries are rounded independently.

	gap $H - F$ (%)		ownership increment (%)	
	impact	20 quarters	impact	20 quarters
Bottom fifth	25	32	16	14
Second fifth	22	21	13	13
Middle fifth	18	19	16	15
Fourth fifth	19	21	23	22
Top fifth	16	7	32	37
Poor hand-to-mouth	0	−2	4	4
Wealthy hand-to-mouth	43	41	10	7
Not hand-to-mouth	58	61	86	89

Table 11: Contributions to the aggregate gap and to its ownership increment. A group's contribution is its consumption response weighted by its share of its country's consumption along its path without the shock, H minus F, as a per cent of the aggregate gap of Table 4 (+0.107 on impact, +0.442 summed over twenty quarters) or of the ownership increment of Table 5 (+0.016 and +0.181). The five fifths add up to 100, as do the three liquidity groups; entries are rounded independently.

	total	earnings	transfers	deposit return	illiquid return
Country H, impact, per cent					
Bottom fifth	−2.01	−0.26	−1.70	−0.02	−0.04
Second fifth	−0.44	−0.11	−0.23	−0.03	−0.07
Middle fifth	−0.30	−0.08	−0.12	−0.03	−0.07
Fourth fifth	−0.23	−0.07	−0.07	−0.02	−0.07
Top fifth	−0.16	−0.05	−0.04	−0.01	−0.06
Poor hand-to-mouth	−3.88	−0.38	−3.49	0.00	−0.01
Wealthy hand-to-mouth	−0.68	−0.16	−0.42	−0.03	−0.07
Not hand-to-mouth	−0.25	−0.07	−0.09	−0.02	−0.07
Country H, summed over 20 quarters, per cent-quarters					
Bottom fifth	−7.74	−2.28	−5.30	−0.01	−0.16
Second fifth	−4.27	−1.66	−2.40	+0.18	−0.39
Middle fifth	−2.78	−1.27	−1.38	+0.35	−0.48
Fourth fifth	−1.88	−1.03	−0.89	+0.54	−0.50
Top fifth	−0.75	−0.72	−0.48	+0.83	−0.38
Poor hand-to-mouth	−9.84	−2.45	−7.30	−0.02	−0.07
Wealthy hand-to-mouth	−3.84	−1.73	−1.83	+0.05	−0.33
Not hand-to-mouth	−2.02	−1.05	−1.12	+0.57	−0.43
Country F, impact, per cent					
Bottom fifth	−2.54	−0.29	−2.18	−0.02	−0.04
Second fifth	−0.66	−0.15	−0.41	−0.03	−0.06
Middle fifth	−0.43	−0.11	−0.22	−0.03	−0.07
Fourth fifth	−0.33	−0.10	−0.14	−0.03	−0.06
Top fifth	−0.19	−0.06	−0.07	−0.01	−0.05
Poor hand-to-mouth	−6.06	−0.49	−5.56	0.00	−0.01
Wealthy hand-to-mouth	−1.15	−0.23	−0.83	−0.03	−0.06
Not hand-to-mouth	−0.32	−0.09	−0.15	−0.02	−0.06
Country F, summed over 20 quarters, per cent-quarters					
Bottom fifth	−10.29	−2.55	−7.54	0.00	−0.19
Second fifth	−5.24	−1.84	−3.23	+0.17	−0.35
Middle fifth	−3.41	−1.44	−1.93	+0.34	−0.39
Fourth fifth	−2.30	−1.20	−1.26	+0.47	−0.32
Top fifth	−0.78	−0.76	−0.64	+0.80	−0.18
Poor hand-to-mouth	−14.98	−2.88	−12.01	−0.02	−0.07
Wealthy hand-to-mouth	−5.05	−1.95	−2.90	+0.07	−0.27
Not hand-to-mouth	−2.35	−1.14	−1.50	+0.56	−0.27

Table 12: Consumption response by household group and channel. Per cent of the group's own consumption along its path without the shock. The four channels are those of Appendix B and add up to the total, up to rounding.

	increment		20 quarters, by channel			
	impact	20 quarters	earnings	transfers	deposit	illiquid
Bottom fifth	+0.034	+0.332	+0.080	+0.243	−0.002	+0.012
Second fifth	+0.017	+0.187	+0.053	+0.102	−0.011	+0.043
Middle fifth	+0.016	+0.167	+0.043	+0.070	−0.021	+0.075
Fourth fifth	+0.016	+0.172	+0.045	+0.050	−0.011	+0.089
Top fifth	+0.012	+0.165	+0.019	+0.024	−0.026	+0.148
Poor hand-to-mouth	+0.109	+0.513	+0.096	+0.415	0.000	+0.002
Wealthy hand-to-mouth	+0.034	+0.220	+0.064	+0.108	−0.007	+0.054
Not hand-to-mouth	+0.014	+0.175	+0.035	+0.055	−0.019	+0.105
All households	+0.016	+0.181	+0.037	+0.063	−0.018	+0.099

Table 13: The ownership increment of the gap, by household group and channel. The group's $H-F$ gap in the calibrated union minus the same in the union without intra-union holdings, in percentage points of the group's own consumption on impact and percentage-point-quarters summed over quarters 0 to 19. The four channels add up to the summed increment within 0.0004.

that carries the largest part of the gap; the first-order equilibrium does not quantify it, since prices, wages and transfers are not re-solved.

	transfers adjust		labour tax adjusts		share of fall	
	impact	20q	impact	20q	transfer	tax
Panel A: country H, per cent of the group's consumption						
Bottom fifth	−2.01	−7.74	−0.63	−4.63	40	16
Second fifth	−0.44	−4.27	−0.36	−3.78	15	15
Middle fifth	−0.30	−2.78	−0.31	−3.08	14	18
Fourth fifth	−0.23	−1.88	−0.27	−2.53	15	22
Top fifth	−0.16	−0.75	−0.22	−1.59	17	30
Poor hand-to-mouth	−3.88	−9.84	−0.90	−4.79	30	9
Wealthy hand-to-mouth	−0.68	−3.84	−0.47	−3.97	8	7
Not hand-to-mouth	−0.25	−2.02	−0.27	−2.43	62	85
All households	−0.37	−2.34	−0.30	−2.58	100	100
Panel B: country F, per cent of the group's consumption						
Bottom fifth	−2.54	−10.29	−0.69	−5.65	36	14
Second fifth	−0.66	−5.24	−0.44	−4.49	16	15
Middle fifth	−0.43	−3.41	−0.37	−3.68	15	18
Fourth fifth	−0.33	−2.30	−0.33	−3.07	16	23
Top fifth	−0.19	−0.78	−0.25	−1.81	17	31
Poor hand-to-mouth	−6.06	−14.98	−1.03	−6.03	24	6
Wealthy hand-to-mouth	−1.15	−5.05	−0.60	−4.71	15	11
Not hand-to-mouth	−0.32	−2.35	−0.31	−2.81	61	83
All households	−0.48	−2.79	−0.34	−3.01	100	100
Panel C: contribution to the gap H − F, per cent						
Bottom fifth	25	32	1	11		
Second fifth	22	21	16	15		
Middle fifth	18	19	19	18		
Fourth fifth	19	21	29	27		
Top fifth	16	7	36	28		
Poor hand-to-mouth	0	−2	−16	−7		
Wealthy hand-to-mouth	43	41	42	35		
Not hand-to-mouth	58	61	73	71		

Table 14: Who pays under the two budget instruments. Groups are fixed at the moment of the shock, as in Table 10. When transfers adjust, the lump-sum transfer balances each budget, as in the rest of the paper; when the labour tax adjusts, its level balances each budget and transfers stay at their steady-state value. Panels A and B give each group's consumption response in per cent of its own consumption along its path without the shock, on impact and summed over quarters 0 to 19, and its per cent of its country's aggregate fall on impact. Panel C weights each group as in Table 11; the aggregate gap is +0.107 on impact and +0.442 summed over twenty quarters when transfers adjust, and +0.046 and +0.435 when the labour tax adjusts. The five fifths add up to 100, as do the three liquidity groups; entries are rounded independently.

E.3 Who replaces the foreign creditors?

Investors outside the union hold 10.7 per cent of F's debt in the calibrated union. This section removes them, taking ω_F^R from 0.107 to 0, and compares the two steady states and the response to the same common tightening in each. The two economies share every primitive: preferences, the rebalancing friction and the liquidity wedge, technology and markups, the disutility of work, the tax schedules, the bond contract, the spread rule, home bias and country size, the rest of the world's share of H's debt, and the real values of government debt and spending. What adjusts is everything an equilibrium determines: the real rate, hours, output, wages, dividends and equity values, the wealth distribution in each country, H's net claim on F, the terms of trade, and the transfers that balance each budget at the new debt service. The intercept of the policy rule moves with the natural rate, so steady-state inflation is zero in both. Because the parameters are held fixed, who ends up holding the debt is an outcome of household saving and market clearing.

The experiment is run under two rules for where the rest of the world spends its income from the union's debt: on the goods of the issuer whose debt it holds, as in the baseline, and across the union's goods in proportion to output. The second rule is a separate joint calibration with the same targets (Appendix E.4). The holdings data do not distinguish between the two, so both are reported.

Where the debt goes. Table 15 compares the steady states. The statistic of interest is the change in each group's equilibrium holdings of F's debt, from the adding-up condition (37), per unit of the outside holdings withdrawn; it compares steady states and does not trace purchases. Under the issuer rule H's savers end up holding 0.59 of the withdrawn debt and F's savers 0.41; under the pooled rule the split is 0.58 and 0.42. The union's real rate rises by about one basis point under either rule: near the steady-state rate the households' demand for assets is elastic enough that the union absorbs the exit through portfolios, with almost no change in the interest rate. With a single household in each country and a common discount factor, steady-state asset demand would be flat at the policy rate, the rate could not move and the absorber would not be determined. Heterogeneous households determine it, and most of the debt moves to the creditor country.

	issuer rule		pooled rule	
	calibrated	no outside	calibrated	no outside
Real rate, per cent a year	5.000	5.011	5.000	5.011
Terms of trade $\bar{S}$	0.957	0.939	0.972	0.975
F's debt held by				
F's intermediary	3.481	3.670	3.479	3.673
H's intermediary	0.330	0.600	0.333	0.597
the rest of the world	0.459	0	0.459	0
Share of the withdrawn holdings absorbed by				
H's savers	0.587		0.576	
F's savers	0.413		0.424	

Table 15: Steady states with and without outside holders of F's debt. Holdings are per resident of F in F's CPI units; F's debt is 4.270 in every column. $\bar{S}$ is the price of F's good relative to H's.

What it does to the common tightening. Table 16 reports the common 25 basis point tightening in each steady state, with the channel accounting of Appendix B. On

impact the gap barely moves under either rule. Cumulated, the two rules diverge, and the difference runs through the goods market. Under the issuer rule the departing investors take with them the export demand that their coupons financed, all of it for F's goods. F's terms of trade fall by 1.8 per cent, the tightening weighs more on F's earnings and transfers relative to H's, and the cumulated gap widens by 0.081, close to a fifth. When the rest of the world's spending is spread across the union in proportion to output, the lost demand falls on both countries, the terms of trade barely move, and the cumulated gap narrows by 0.030, again mostly through earnings and transfers. An external investor base shapes the incidence of the common policy less through whose books hold the debt than through whose goods its investors buy.

	issuer rule			pooled rule		
	calibrated	no outside	change	calibrated	no outside	change
Gap, impact	+0.107	+0.118	+0.012	+0.096	+0.093	−0.003
Gap, cum. 20q	+0.442	+0.522	+0.081	+0.361	+0.331	−0.030
Change in the channel contributions, cum. 20q						
Earnings			+0.060			−0.012
Transfers			+0.047			−0.013
Deposit return			−0.020			−0.003
Illiquid return			−0.006			−0.002

Table 16: The $H - F$ consumption gap under a common 25 basis point tightening, with and without outside holders of F's debt. Impact values are percentage points, cumulative values percentage-point-quarters. The four channel changes add up to the change in the cumulated gap.

Who pays. By household group (Appendix E.2), the change is broad-based under either rule. About four fifths of the change in the aggregate gap comes from households that are not hand-to-mouth, and F's distribution of its own burden is unchanged: each group's share of F's impact fall moves by less than half a percentage point. The external investor base moves how much F pays relative to H, not who pays within F.

E.4 Robustness

Each row of Table 17 is the calibrated union with one element changed and everything else held at its baseline value, so each can be read directly against the headline of Section 5.3. Every row satisfies the steady-state and accounting invariants. The rows that change the steady state are separate joint recalibrations; the last two change only how the economy responds to the shock and keep the baseline steady state.

	gap		difference from baseline	
	impact	cum. 20q	impact	cum. 20q
Calibrated union	+0.107	+0.442		
Total debt including loans	+0.147	+0.632	+0.041	+0.190
Eurosystem as an outside holder	+0.099	+0.503	−0.008	+0.062
Rest of the world spends across the union	+0.096	+0.361	−0.011	−0.080
Ownership at its conservative bound	+0.105	+0.424	−0.002	−0.018
Consolidated duration, 5.69 years	+0.105	+0.437	−0.002	−0.004
Common liquidity wedge	+0.027	+0.235	−0.080	−0.207
Real rate 2.5 per cent a year	+0.125	+0.467	+0.019	+0.026
Labour tax balances the budget	+0.046	+0.435	−0.061	−0.007
Deposits on the short rate	+0.112	+0.458	+0.006	+0.016

Table 17: Robustness of the $H - F$ consumption gap under a common tightening. Impact values are percentage points, cumulative values percentage-point-quarters. The common-wedge row replaces the per-country wedge of Section 4.2 with a single value in both countries. It does not merely switch the household asymmetry off: it reverses the hand-to-mouth ordering to 0.189 in H against 0.152 in F, against the data, which is why its effect exceeds the household Shapley value of Table 6. The real-rate row is a full recalibration at 2.5 per cent a year with every target unchanged, and it also re-solves the union without intra-union holdings, so the text below reports the ownership increment at that rate too. The last two rows keep the baseline steady state, and the text below reports the ownership increment under each.

The debt perimeter. Measuring sovereign debt inclusive of loans rather than as debt securities widens the debt gap between the two countries by forty per cent and widens the consumption gap by forty-three. That the two move almost one for one is the cleanest evidence that the level channel enters the gap close to proportionally. On the broader perimeter, the consumption gap is larger.

The consolidation convention. Whether Eurosystem holdings are netted against the issuer's own book or treated as an outside holder is the most consequential mapping choice in the holder data. Treating them as outside leaves the cumulative gap fourteen per cent larger. The headline gap therefore survives the convention, and if anything the baseline understates it; the row speaks to the total gap. Impact and cumulative responses move in opposite directions here, which is what one expects if the additional leakage bites immediately while the larger effective tradeable stock dominates at longer horizons.

Where the rest of the world spends. In the baseline the rest of the world spends its income from each issuer's debt on that issuer's goods. Spreading it across the union's

goods in proportion to output lowers the gap by a tenth on impact and by close to a fifth cumulated, and leaves the ordering of the two countries unchanged.

Ownership bounds and duration. Two rows move the gap little. Taking the own-country ownership shares to their most conservative reading moves the cumulative gap by four per cent, because it lowers ownership in both countries and leaves the *difference* between them close to its baseline value, and the difference is what the channel operates on. Shortening the common duration from 8.75 to 5.69 years moves it by one per cent. That is a numerical finding at this calibration rather than a cancellation: a common duration revalues each book in proportion to what it holds, and Appendix D.3 shows that raising the common duration from one period to 8.75 years nearly doubles the ownership gap in the symmetric benchmark. Around the calibrated duration the gap is flat in it.

The steady-state real rate. At five per cent a year, steady-state debt service is above what euro-area sovereigns have paid (Section 4.3). The real-rate row halves the rate to 2.5 per cent a year, which brings debt service to 1.7 and 2.7 per cent of GDP against 1.2 and 3.2 in the data, and recalibrates everything else to the same targets. The headline gap is slightly larger, 0.125 on impact against 0.107 and 0.467 cumulated against 0.442. A common change in the policy rate moves debt service by the change in the rate times the stock, which the level of the rate does not alter, while the recalibrated economy at the lower rate has more households at the liquid constraint, 16.8 and 18.5 per cent against 15.7 and 17.4, so the same fiscal flow moves consumption by more and the transfer channel's share of the impact response rises to 57 and 64 per cent. The ownership increment, re-measured at the lower rate by solving the union without intra-union holdings again, falls by about a third, to 0.010 on impact and 0.113 cumulated against 0.016 and 0.181. The fall is less than proportional to the rate, and the comparison moves several things at once: the steady-state annuity from debtor to creditor is smaller at the lower rate, the bond price and its response to the policy rate change with it, and the household side is recalibrated. What the row establishes is that the ownership increment remains positive and of the same order in a lower-rate economy. At the common debt level the increment at the lower rate is 0.053 cumulated, so unequal debt about doubles it, as at the baseline rate, and the interaction between the two is 0.060 against 0.087. The lower rate fits debt service better and the production side worse: with the capital-output ratio and the wealth targets held, the implied markups fall to 1.03 and 1.05 and the labour shares rise to 0.70 and 0.68, above the data's 0.64 and 0.59, which is why five per cent remains the baseline.

The budget instrument and the deposit contract. In the baseline each government balances its budget with the lump-sum transfer. Balancing it with the level of the labour tax, with transfers at their steady-state value, leaves the cumulative gap at 0.435 against 0.442 and lowers the impact gap from 0.107 to 0.046: a tax on earnings falls more on households who smooth their consumption, so the contraction is spread over time (Section 5.2). The ownership increment, re-measured in the union without intra-union holdings under the same instrument, is 0.012 on impact and 0.166 cumulated, against 0.016 and 0.181, and its illiquid-return contribution is 0.10 under either instrument. The second row writes deposits on the short nominal rate set in the previous quarter, so the illiquid account absorbs the whole revaluation of the sovereign book in the impact quarter; from the second quarter the two contracts pay the same, since under perfect foresight the long bond's holding return equals the short rate. The gap is 0.112 on impact and 0.458 cumulated, and the ownership increment 0.016 and 0.183.