

The Network Origins of Inflation Stances in a Currency Union

ZHIHAO XU*

THIS VERSION: AUG 19, 2026

ABSTRACT

This paper proposes a production-network mechanism to explain divergent inflation stances within a currency union. I develop a heterogeneous-agent input–output (HAIO) model and show that a country’s inflation stance reflects redistributive incentives to manipulate the terms of trade, with its magnitude and direction determined by production networks, ownership structures, and nominal rigidities. Applied to the euro area, a country’s stance is inversely related to its upstreamness, and the policy-alignment loss increases quadratically with its stance deviation from the union-wide consensus, implying that countries at either extreme of the production network, whether highly upstream or highly downstream, incur the largest welfare costs.

KEYWORDS: production networks, optimal monetary policy, currency union, heterogeneous agents, inflation stances

*Emory University. E-mail: owenzxu@gmail.com. I am grateful to Tao Zha for invaluable and continued guidance. I thank Kaiji Chen, Alireza Tahbaz-Salehi, Le Xu and Vivian Yue for constructive conversations. I also thank Yan Bai, Zhen Huo, Emi Nakamura, Vincenzo Quadrini, Shengxing Zhang, and the participants at the NASM and ESEM for helpful discussions.

1 Introduction

A salient feature of monetary policymaking in the euro area is the persistent divergence in inflation stances across member states, often discussed as a North–South divide (see, e.g., [Iversen et al., 2016](#)). Countries such as Germany and the Netherlands are commonly associated with a stronger emphasis on price stability, while others, including Italy, Greece, and Portugal, are typically viewed as more supportive of accommodative policies. This persistent disagreement reflects a fundamental tension at the heart of any currency union: a single monetary instrument must serve member states with structurally different preferred inflation outcomes. Why do these preferences diverge so systematically across countries, and what are the welfare consequences of imposing a common policy on economies that would each, on their own, choose differently?

This paper proposes a production-network explanation for divergent inflation stances within a currency union. It departs from the traditional New Open Economy Macroeconomics (NOEM) literature, which emphasizes cross-country asymmetries in shocks, preferences, and nominal rigidities as the primary sources of policy divergence (see, e.g., [Benigno, 2004](#)), by highlighting the role of input–output linkages. I argue that asymmetries in production networks and ownership structures can endogenously generate divergent policy preferences through their impact on the propagation and incidence of shocks across member states.

Building on [Baqae and Farhi \(2024\)](#), I construct a multisector general equilibrium model with heterogeneous agents, input–output linkages, and nominal rigidities, in which countries differ in labor supply, consumption baskets, and firm ownership structures. Within this framework, monetary policy induces first-order changes in country-level allocative efficiency, which decompose into two distinct effects: a direct-incidence effect and a factoral terms-of-trade effect. In a setting where there are no input–output linkages and firms are owned domestically, the direct-incidence effect coincides with the conventional notion of commodity terms of trade. More generally, it captures network-adjusted commodity terms of trade, which combine with factoral terms of trade to determine a country’s overall terms of trade—the exclusive source of first-order changes in allocative efficiency in the absence of initial distortions.

How do production networks reshape the conventional commodity terms of trade? By further decomposing the direct-incidence effect into an input–output multiplier channel, a home-bias channel, and a sectoral-heterogeneity channel, I show that production networks generate terms-of-trade redistribution even when markup responses are uniform across

sectors: countries with above-average input–output multipliers experience a deterioration in their commodity terms of trade under monetary expansion, a mechanism entirely absent in standard economies without input–output linkages.¹

I then derive a second-order approximation to the utilitarian welfare loss function and characterize the Ramsey-optimal monetary policy under arbitrary Pareto weights. When Pareto weights coincide with countries’ income shares, the optimal policy is a pure price-stabilization rule that places greater weight on industries that are larger, exhibit greater nominal rigidity, and belong to countries whose labor supply is more sensitive to monetary expansions. By contrast, when Pareto weights deviate from income shares—reflecting a misalignment between political influence and economic size—monetary policy faces a fundamental trade-off between aggregate stabilization and first-order redistribution. This tension generates a redistributive inflation bias, as the central bank has incentives to tilt the overall terms of trade in favor of certain member states.

Building on this result, the analysis turns to policy evaluation from the perspective of an individual country. When Pareto weights are unilateral—that is, when full weight is assigned to a single country—the corresponding unilateral optimal policy minimizes that country’s welfare loss. While this policy implements the second-best allocation, it presumes that a single country has full control over the policy instruments—an arrangement ruled out by the union’s common price-targeting regime. Once the price-index regime is taken as given and inflation must be non-state-contingent, the country faces a third-best problem. The unilateral inflation stance is therefore defined as the union-wide inflation rate that minimizes the country’s expected welfare loss under this institutional constraint. The analysis shows that this stance is proportional to the inflation bias under the unilateral optimal policy and summarizes how the country’s overall terms of trade respond to monetary expansions: a positive value reflects an improvement, while a negative value reflects a deterioration.

How does divergence in inflation stances translate into the welfare incidence of a common monetary policy? To address this question, I introduce the policy-alignment loss (PAL), defined as the welfare loss a country incurs when the union-wide policy differs from the policy that would be optimal from that country’s perspective.² I show that a country’s expected PAL decomposes into an inflation-misalignment component, which is quadratic in its inflation stance deviation from the union-wide consensus, and a price-index-misalignment component that vanishes when its preferred price index coincides with the

¹The country-level input–output multiplier is the analog of the economy-wide multiplier of [Baqae and Farhi \(2019\)](#): it measures a country’s gross domestic sales relative to its domestic final demand.

²The PAL is a within-union measure of policy misalignment: it compares alternative uses of the same common monetary instrument, rather than alternative monetary regimes.

union's.

I then apply the theoretical framework to a quantitative analysis of the euro area. Using the World Input–Output Database (WIOD), I calibrate the full inter- and intra-country production network and match it with industry-level data on nominal rigidities. I compute the centralized optimal monetary policy and compare its resulting welfare losses to those under alternative price-stabilization rules, both at the union-wide and country-specific levels. Next, I compute each country's unilateral optimal policy and identify its unilateral inflation stance and policy-alignment loss. The results demonstrate that policy-alignment losses are driven primarily by inflation misalignment, so that a country's stance deviation serves as a sufficient statistic for its policy-alignment loss. Quantitatively, a one-percentage-point deviation predicts a loss of about 0.29 percentage points of steady-state consumption; a two-percentage-point deviation predicts roughly four times that amount.

How does a country's inflation stance relate to its position in the union's production network? The quantitative analysis reveals a robust negative relationship between inflation stance and production upstreamness.³ Among major euro-area economies, the Netherlands and Germany exhibit lower preferred inflation rates than France, Italy, and Spain, closely mirroring the widely discussed North–South divide.

The mechanism operates primarily through the input–output multiplier channel, which by itself accounts for most of the cross-country variation in inflation stances. Because upstreamness is strongly correlated with the input–output multiplier, upstream countries typically have above-average multipliers and tend to experience an overall terms-of-trade deterioration under monetary expansion, leading them to prefer tighter policy. By contrast, downstream countries tend to experience an improvement in their terms of trade under monetary expansion and therefore prefer more accommodative policy. Combined with the quadratic relationship between stance deviation and policy-alignment loss, this implies that countries at either extreme of the production chain—those with highly upstream or highly downstream structures—tend to exhibit more extreme inflation stances and incur larger welfare losses under a common monetary policy.

Related literature. This paper is primarily related to the NOEM literature on currency-union policy design (see e.g., [Benigno, 2004](#); [Galí, 2008](#); [Ferrero, 2009](#); [Farhi and Werning, 2017](#)). Foundational work in this field typically focuses on the optimal policy without

³Country upstreamness, following [Antràs et al. \(2012\)](#), captures how upstream a country is in the union's production network, measured as the value-added-weighted average distance of its industries from final consumption.

redistribution concerns (under income-share Pareto weights). In contrast, paralleling recent progress in the HANK literature (e.g., [Bhandari et al., 2021](#); [Acharya et al., 2023](#); [Dávila and Schaab, 2023](#); [Le Grand et al., 2022](#); [La’O and Morrison, 2024](#); [Nuño and Thomas, 2022](#)), this paper demonstrates that, in a heterogeneous-agent economy, monetary policy faces a fundamental trade-off between aggregate stabilization and first-order redistribution whenever Pareto weights deviate from income shares. The novelty here is that the redistribution motive operates through asymmetry in production-network exposure rather than through wealth heterogeneity. This stabilization–redistribution trade-off results in an inflation bias that reflects a policymaker’s incentive to tilt the terms of trade between members—a mechanism with a long tradition in the NOEM literature on terms-of-trade manipulation. One of the first to formalize this mechanism is [Corsetti and Pesenti \(2001, 2005\)](#), who show that the central bank of a country with monopoly power in trade has an incentive to appreciate its currency to improve the terms of trade, resulting in a deflationary bias. Subsequent research shows that the direction of this bias—whether inflationary or deflationary—depends on the elasticity of substitution within and across countries, which determines the direction of expenditure switching (e.g., [Tille, 2001](#); [De Paoli, 2009](#)), and on the price-setting regime, which influences the degree of exchange-rate pass-through (e.g., [Devereux and Engel, 2003](#)).⁴ Building on this insight, this paper contributes to the literature by demonstrating that a country’s unilateral inflation stance operates through the input–output multiplier that governs its network-adjusted terms-of-trade exposure, and thus relates to its position in the supply chain.

This contribution also places the paper within the burgeoning literature on production networks, which studies how microeconomic shocks propagate through input–output linkages. For example, building on [Long and Plosser \(1983\)](#), a strand of work, including [Acemoglu et al. \(2012\)](#), [Acemoglu et al. \(2013, 2017\)](#); [Baqae and Farhi \(2019\)](#); [Dew-Becker \(2023\)](#) and [Taschereau-Dumouchel \(2026\)](#), examines how input–output linkages transmit micro-level productivity shocks into macroeconomic outcomes and shape the aggregate output distribution.⁵ In contrast, [Jones \(2013\)](#); [Liu \(2019\)](#); [Baqae and Farhi \(2020\)](#) and [Bigio and La’O \(2020\)](#) focus on the macroeconomic consequences of micro-level distortions. Extending this line, [Baqae and Farhi \(2024\)](#) develop a general heterogeneous-agent framework, which serves as a foundation for this paper.⁶ Relative to their general setup,

⁴Recent studies include [Bergin and Corsetti \(2023\)](#) and [Bianchi and Coulibaly \(2025\)](#).

⁵Other papers in this line of work include [Gabaix \(2011\)](#); [Foerster et al. \(2011\)](#); [Atalay \(2017\)](#); [Baqae \(2018\)](#); [Carvalho et al. \(2021\)](#) and [Acemoglu and Tahbaz-Salehi \(2025\)](#), among others. Also see [Carvalho \(2014\)](#); [Carvalho and Tahbaz-Salehi \(2019\)](#) and [Baqae and Rubbo \(2023\)](#) for surveys.

⁶This framework has also been extended in subsequent work; see, for example, [Baqae and Burstein \(2025\)](#).

this paper considers a more tractable case with Cobb–Douglas consumption and production functions. This simplification enables explicit welfare analysis at the individual level, showing that cross-sector misallocation captures second-order changes in both the labor wedge and the factorial terms of trade.

This paper also contributes to the literature on monetary transmission in multi-sector economies with input–output linkages, including [Pastén et al. \(2020, 2024\)](#); [Wei and Xie \(2020\)](#); [Ghassibe \(2021\)](#); [Ferrari and Ghassibe \(2024\)](#); [La’O and Tahbaz-Salehi \(2022, 2025\)](#); [Rubbo \(2023, 2025\)](#); [Afrouzi and Bhattarai \(2023\)](#); [Baqaei et al. \(2024\)](#); [Schaab and Tan \(2023\)](#); [Xu and Yu \(2025\)](#) and [Qiu et al. \(2026\)](#). Among these, a small but growing strand of the literature studies optimal monetary policy in production networks. [La’O and Tahbaz-Salehi \(2022\)](#) show that monetary policy cannot implement the first-best allocation in a multisector production network economy. Instead, the optimal policy should stabilize a price index with greater weights assigned to larger, stickier, and more upstream industries. In a parallel work, [Rubbo \(2023\)](#) emphasizes how nominal rigidities accumulate through production chains and flatten both the sectoral and aggregate Phillips curves, and derives a divine coincidence index that outperforms the consumer price index in Phillips curve regressions. Following this line of research, [Xu and Yu \(2025\)](#) characterize optimal monetary policy in production networks with steady-state distortions, while [Qiu et al. \(2026\)](#) study optimal monetary policy in a small open economy with domestic and cross-border input–output linkages. These studies primarily characterize optimal monetary policy from the perspective of a centralized planner. This paper complements the literature by extending the optimal-policy characterization to a heterogeneous-agent economy with arbitrary Pareto weights in welfare aggregation. This generalization changes the object of analysis: beyond the planner’s preferred price index, the paper characterizes countries’ divergent preferences over the common monetary instrument, and identifies how production-network positions generate systematic disagreement over monetary policy.

The closest antecedent to this paper is [Rubbo \(2025\)](#), which also studies monetary policy in a HAIIO economy. While both papers move beyond standard HANK models by emphasizing the role of production structures in shaping the incidence of monetary policy on labor income, this paper extends the analysis from first-order incidence to second-order welfare analysis, thereby facilitating a full characterization of the optimal policy. This extension is essential for moving from incidence to policy choice: it allows the paper to characterize not only how monetary policy affects different agents, but also which policy each country would prefer for the common monetary instrument.

Outline. The rest of the paper is organized as follows. Section 2 sets up the NK–HAIO environment and defines both the sticky-price and flexible-price equilibria. Section 3 log-linearizes the model around the efficient equilibrium and characterizes monetary transmission. Section 4 derives a closed-form characterization of the Ramsey-optimal monetary policy under arbitrary Pareto weights and introduces the concepts of unilateral inflation stance and policy-alignment loss. Section 5 presents the quantitative analysis. Section 6 discusses extensions of the basic framework to a nested CES economy and to a global economy under dominant currency pricing. All proofs, derivations, and additional extensions are included in the Online Appendix.

2 Framework

This section sets up a New Keynesian model with monopolistic competition and nominal rigidities in a currency union composed of C member countries and N industries. Each country $c \in \{1, \dots, C\}$ is represented by a representative household that supplies a distinct type of labor and differs across countries in labor supply elasticity, sectoral labor input shares, ownership structures, and consumption baskets.

The model is formulated in a general environment that does not require geographic segmentation of labor inputs or ownership claims. In the quantitative analysis, such segmentation is imposed as a calibration baseline, under which each country c is associated with a subset of industries N_c , in order to map the model to a currency-union setting and to discipline the magnitude of cross-country asymmetries.

2.1 Firms

In each industry $i \in \{1, \dots, N\}$, a continuum of monopolistically competitive firms (indexed by $k \in [0, 1]$) produces differentiated varieties using the same constant-returns-to-scale production function,

$$y_{ik} = z_i F_i(\{x_{ij,k}\}_{j=1}^N, \{L_{ic}\}_{c=1}^C) = z_i \varsigma_i \prod_{j=1}^N x_{ij,k}^{\omega_{ij}} \prod_{c=1}^C L_{ic,k}^{\alpha_{ic}}$$

where $\alpha_{ic} \geq 0$ denotes the share of labor supplied by country c used by the firm; $\omega_{ij} \geq 0$ denotes the share of intermediate input j ; and ς_i is a normalization constant independent of the shocks. Sectoral productivity shocks $z_i > 0$ are jointly log-normally distributed, with $\log z = (\log z_1, \dots, \log z_N)' \sim N(\mathbf{0}, \Sigma_z)$.

Differentiated varieties from all the firms k within an industry i are aggregated into an industry-level output using a CES aggregator:

$$y_i = \left(\int_0^1 y_{ik}^{\frac{\theta_i-1}{\theta_i}} dk \right)^{\frac{\theta_i}{\theta_i-1}}, \quad (1)$$

where $\theta_i > 1$ denotes the elasticity of substitution between varieties within industry i .

The implied sectoral price index is thus given by

$$p_i = \left(\int_0^1 p_{ik}^{1-\theta_i} dk \right)^{\frac{1}{1-\theta_i}}. \quad (2)$$

All firms in industry i solve the following cost-minimization problem:

$$mc_i = \min_{\{x_{ijk}\}_{j=1}^N, \{L_{ic,k}\}_{c=1}^C} \sum_{c=1}^C w_c L_{ic,k} + \sum_{j=1}^N p_j x_{ij,k}, \quad s.t. \ y_{ik} = 1$$

where w_c is the wage rate for labor supplied by agent c . Under constant returns to scale, all firms within an industry face the same marginal cost and employ inputs in the same proportions.

Nominal rigidities. Nominal rigidities are introduced by assuming that, in each sector, only a fraction $\delta_i \in (0, 1]$ of firms can adjust their prices after observing the realizations of productivity shocks and the money supply, while the remaining firms charge their steady-state prices. In the static environment considered here, this formulation is equivalent to a sticky-information model (Mankiw and Reis, 2002), in which a fraction δ_i of firms receive updated information about the state of the economy, while the remaining fraction $1 - \delta_i$ set prices based on their priors. The parameter δ_i captures the degree of sectoral price flexibility, with lower values indicating greater nominal rigidity. Given industry level output y_i , price p_i , and marginal cost mc_i , the optimal reset price p_i^* maximizes profits:

$$\Pi_{ik} = \max_{p_{ik}} [p_{ik} - (1 - \tau_i)mc_i] \left(\frac{p_{ik}}{p_i} \right)^{-\theta_i} y_i, \quad (3)$$

where τ_i is an input subsidy provided by the government. Throughout this paper, the subsidies τ_i are non-state-contingent and set to eliminate the distortions that arise under the CES demand structure, resulting in the profit-maximizing price being equal to pre-subsidy

marginal costs:

$$1 - \tau_i = \frac{\theta_i - 1}{\theta_i}.$$

This assumption is standard in the New Keynesian literature (see, e.g., [Woodford, 2003](#); [Galí, 2008](#)) and eliminates the incentive for monetary policy to use nominal adjustments to substitute for missing tax instruments, thereby isolating the welfare effects of nominal rigidities.

After productivity and monetary shocks are realized, firms within each industry charge different prices due to price rigidities. This implies that, to a first-order approximation, the vector of sectoral inflation rates is given by

$$\log p = \delta \log mc$$

where $\log p$ denotes the vector of sectoral inflation rates, $\log mc$ is the vector of changes in sectoral marginal costs, and $\delta = \text{diag}(\delta_1, \dots, \delta_N)$ is a diagonal matrix of industry-specific price adjustment probabilities. This then implies that sectoral markups are related to sectoral inflation through

$$\log \mu = \log(p/mc) = -(\delta^{-1} - I) \log p. \quad (4)$$

2.2 Households

Each country is populated by a representative household that earns income from labor supply and from sectoral profits. Each household retains exclusive ownership of its labor income, so labor earnings are not shared across countries. Profit income, by contrast, is distributed according to an ownership structure that governs the allocation of firm profits across countries. Following [Baqaee and Farhi \(2024\)](#), I introduce an $N \times C$ ownership matrix Φ , where each entry Φ_{ic} denotes the share of profits generated in industry i that accrues to the representative household in country c . A block-diagonal structure of Φ (i.e., $\Phi_{ic} = \mathbf{1}\{i \in N_c\}$) corresponds to geographically segmented ownership, whereby the representative household in country c receives profits only from industries operating domestically. More general ownership structures allow profits to be distributed independently of production location, thereby altering the cross-country incidence of income and the redistribution effects induced by monetary policy.

The representative household in country c has log balanced-growth preferences,⁷

$$U_c(C_c, L_c) = \log C_c - \psi_c \frac{L_c^{1+1/\eta_c}}{1+1/\eta_c},$$

where L_c denotes labor supply, η_c is the Frisch elasticity, and ψ_c is calibrated to match steady-state labor supply. Consumption is summarized by the Cobb–Douglas aggregator

$$C_c = \prod_{i=1}^N \left(\frac{c_{ci}}{\beta_{ci}} \right)^{\beta_{ci}},$$

where c_{ci} is consumption of good i by the household in country c , and $\beta_{ci} \geq 0$ denotes the expenditure share of good i in country c 's consumption basket.

Each household maximizes utility subject to the budget constraint

$$\sum_{i=1}^N p_i c_{ci} \leq w_c L_c + \sum_{i=1}^N \Phi_{ic} \Pi_i - T_c, \quad (5)$$

where $\Pi_i = \int_0^1 \Pi_{i,k} dk$ denotes total profits generated in industry i , and T_c is a lump-sum tax. The budget constraint equates country c 's nominal gross national expenditure (GNE_c) with its nominal gross national income (GNI_c), given by labor and profit income accruing to domestic residents, net of international taxes and transfers.

2.3 Policy Instruments

The economy also features a government comprising fiscal and monetary authorities, each responsible for implementing its respective policy instruments.

To ensure budget balance, fiscal policy finances subsidies to firms via lump-sum taxes imposed on households in proportion to their ownership shares. In particular, the tax paid by the representative household in country c is given by

$$T_c = \sum_{i=1}^N \Phi_{ic} \tau_i m c_i \int_0^1 y_{ik} dk \quad (6)$$

where $\tau_i m c_i \int_0^1 y_{ik} dk$ denotes the total subsidy allocated to sector i .

This tax scheme renders the fiscal system neutral: in steady state, each country's nominal

⁷This specification corresponds to the log-separable case of KPR preferences (King et al., 1988).

expenditure is fully financed by labor income, and trade is balanced across countries. For example, under geographically segmented ownership, equation (6) implies that subsidies to domestic firms are financed entirely by taxes raised domestically.

On the monetary side, I adopt the cash-in-advance setup of [Pastén et al. \(2020\)](#), in which the money supply m directly determines nominal aggregate demand:

$$\text{GNE} = \sum_{c=1}^C \sum_{i=1}^N p_i c_{ci} = m.$$

This formulation abstracts from monetary micro-foundations and treats m as an intermediary policy instrument.

2.4 Equilibrium

I now formalize equilibrium in the model economy. The environment features nominal rigidities faced by firms, and clearing in both goods and factor markets. I define two equilibrium concepts: one with nominal rigidities and one with fully flexible prices. These serve as the basis for analyzing how monetary policy transmits through the economy and for quantifying the inefficiencies introduced by price rigidity.

Before presenting the formal definitions, note that the market-clearing conditions for goods and labor markets are given by

$$y_i = \sum_{c=1}^C c_{ci} + \sum_{j=1}^N \int_0^1 x_{jik} dk,$$

$$L_c = \sum_{i=1}^N \int_0^1 L_{ic,k} dk,$$

for all industries $i \in N$ and countries $c \in C$. The first condition requires that total output in each industry equals its use for final consumption across countries and as intermediate inputs across firms. The second condition equates total labor supplied by country c to the sum of labor services employed across industries.

Definition 1. For any realization of productivity shocks $\{z_i\}_{i=1}^N$ and a monetary shock m , a sticky-price equilibrium consists of: a vector of prices $\{p_i\}_{i=1}^N$, a vector of sectoral output $\{y_i\}_{i=1}^N$, a vector of nominal wages $\{w_c\}_{c=1}^C$, a vector of labor supply $\{L_c\}_{c=1}^C$, a matrix of intermediate inputs $\{x_{ij}\}_{i,j \in N}$, a matrix of primary inputs $\{L_{ic}\}_{i \in N, c \in C}$ and a matrix of final use $\{c_{ci}\}_{c \in C, i \in N}$, such that: (i) firms optimally choose intermediate inputs and labor demand to

minimize their costs, and optimally reset their prices when they have a chance to adjust; (ii) each representative household maximizes utility subject to the budget constraint; (iii) the government budget constraint is satisfied; (iv) all markets for goods and labor clear.

To isolate the effects of nominal rigidities, I define a benchmark equilibrium in which all prices are fully flexible.

Definition 2. The flexible-price equilibrium is defined by the same conditions as the sticky-price equilibrium, except that all firms are assumed to reset their prices optimally in response to realized shocks.

3 Log-Linearized Model

In this section, I approximate the model around the efficient equilibrium and characterize some important results of the log-linearized economy. I state these results in terms of changes in ex post markups.

3.1 Definitions and Notations

Throughout this paper, I express the log deviation of the variable x from the flexible-price equilibrium x^* as:

$$\hat{x} = \log x - \log x^*.$$

For example, the variable $\hat{c}_c = \log(C_c/C_c^*)$ denotes the consumption gap, representing the percentage deviation of country c 's consumption in the sticky-price equilibrium relative to the first-best flexible-price benchmark.

Before presenting results for the log-linearized economy with input–output linkages, I introduce notation to streamline the expressions. Table 1 defines key input–output parameters, including the Leontief inverse and Domar weights.

Input–output matrices. The entry ω_{ij} in the input–output matrix represents the direct elasticity of sector i 's marginal cost with respect to the price of sector j . In contrast, the element Ψ_{ij} of the associated Leontief inverse captures the total exposure of sector i to sector j , incorporating both direct and indirect effects—namely, how price changes in sector j transmit through input–output linkages to affect the marginal cost of sector i .⁸ By analogy, the element Ψ_{if} in the Leontief inverse of factor inputs captures the total exposure of sector

⁸This follows from the Neumann series expansion of the Leontief inverse: $\Psi = (I - \Omega)^{-1} = I + \Omega + \Omega^2 + \dots$, which captures successive rounds of input propagation across sectors.

Table 1: input–output definitions.

Income share of country c	$\chi \in \mathbb{R}^C, \chi_c = \frac{\text{GNE}_c}{\text{GNE}}$
Consumption basket of country c	$\beta_c \in \mathbb{R}^N, \beta_{ci} = \frac{p_i c_{ci}}{\text{GNE}_c}$
Union-wide consumption shares	$b \in \mathbb{R}^N, b_i = \sum_{c=1}^C \chi_c \beta_{ci}$
Input–output matrix	$\Omega \in \mathbb{R}^{N \times N}, \omega_{ij} = \frac{p_j x_{ij}}{m c_i y_i}$
Labor input matrix	$\alpha \in \mathbb{R}^{N \times C}, \alpha_{ic} = \frac{w_c L_{ic}}{m c_i y_i}$
Leontief inverse	$\Psi = (I - \Omega)^{-1}$
Leontief inverse of factors	$\Psi_{(f)} = \Psi \alpha_{(f)}$
Pass-throughs of nominal wages into prices	$\varrho^w = (I - \delta \Omega)^{-1} \delta \alpha, \varrho_{ic}^w = \frac{d \log p_i}{d \log w_c}$
Exposures of country c to sectors	$(\lambda^c)' = \beta_c' \Psi$
Domar weights	$\lambda' = b' \Psi, \lambda_i = \frac{p_i y_i}{\text{GNE}}$
Exposures of country c to factors	$(\Lambda^c)' = \beta_c' \Psi \alpha$
Labor income shares	$\Lambda' = b' \Psi \alpha, \Lambda_c = \frac{w_c L_c}{\text{GNE}}$

i 's marginal cost to country f 's nominal wage, accounting for both direct and indirect effects. Since value added in each sector is entirely composed of labor inputs from different countries, it follows that $\sum_{f=1}^C \Psi_{if} = 1$ holds for all sectors.

Domar weights. The Domar weight (or sales share) of sector i , denoted by $\lambda_i = p_i y_i / \text{GNE}$, captures the economy's total exposure to sector i , directly and indirectly. Formally, it aggregates the i -th column of the Leontief inverse Ψ using union-wide consumption shares b , i.e., $\lambda_i = \sum_{j=1}^N b_j \Psi_{ji}$. Similarly, the labor income share of country c , denoted by $\Lambda_c = w_c L_c / \text{GNE}$, measures country c 's labor income as a share of total nominal output. It reflects the economy's total exposure to labor from country c , and relates to the Leontief structure via $\Lambda_c = \sum_{j=1}^N b_j \Psi_{jc}$. Note that the ratio $\Lambda_c / \chi_c = \frac{w_c L_c}{\text{GNE}_c}$ defines country c 's labor wedge, measuring the share of labor income in country c 's total expenditure.⁹ In the flexible-price (efficient) equilibrium, the fiscal setup implies $\Lambda_c = \chi_c$, so the labor wedge equals one in all countries.

Analogously, country-specific Domar weights are constructed using country c 's own consumption shares β_c . The weight $\lambda_i^c = \sum_{j=1}^N \beta_{cj} \Psi_{ji}$ measures country c 's consumption cost exposure to sector i , while $\Lambda_f^c = \sum_{j=1}^N \beta_{cj} \Psi_{jf}$ measures its exposure to labor from country f . These country-level weights provide a decomposition of aggregate Domar weights across consumers, satisfying the identities: $\lambda_i = \sum_{c=1}^C \chi_c \lambda_i^c$ and $\Lambda_f = \sum_{c=1}^C \chi_c \Lambda_f^c$.

⁹Given that $\Lambda_c / \chi_c = -\frac{\partial U_c / \partial L_c}{\partial U_c / \partial C_c} \frac{L_c}{C_c}$, the labor wedge defined here closely follows the standard intratemporal labor-wedge measure (see, e.g., Farhi and Werning, 2017).

Pass-throughs of nominal wages into prices. Due to nominal rigidities, sectors only partially transmit changes in marginal costs along the production chain. As a result, the pass-through of nominal wages to sectoral prices is dampened relative to the frictionless benchmark and is lower than sectors' total structural exposure to labor inputs.

Formally, the matrix of pass-throughs of nominal wages to sectoral prices is given by $\varrho^w = (I - \delta\Omega)^{-1}\delta\alpha$, where each element ϱ_{ic}^w quantifies the elasticity of the price in sector i with respect to the nominal wage in country c . The production network governs how cost shocks propagate across sectors, with nominal rigidities preventing full transmission. In particular, the effective pass-through to a sector's price is bounded above by the product of that sector's own price adjustment probability and its total (direct and indirect) exposure to labor from country c : $\varrho_{ic}^w \leq \delta_i \Psi_{ic}$.

3.2 Basic Results

In the presence of nominal rigidities, monetary policy generates distortions that reallocate resources across industries and countries. This section establishes that, to a first-order approximation, country-level employment gaps, income shares, and consumption gaps are all proportional to weighted sums of sectoral ex post markups, with weights determined by production networks, ownership structures, and labor supply elasticities.

Lemma 1. To a first-order approximation, the country-level employment gaps $\{\hat{l}_c\}_{c=1}^C$ and income share changes $\{\hat{\chi}_c\}_{c=1}^C$ are linear functions of sectoral ex post markups $\{\mu_i\}_{i=1}^N$

$$\hat{l}_c = \sum_{i=1}^N \ell_{ic}^\mu \log \mu_i, \quad \ell_{ic}^\mu \doteq \frac{d \log L_c}{d \log \mu_i} = -\frac{\eta_c}{1 + \eta_c} \cdot \frac{\Phi_{ic} \lambda_i}{\chi_c} \quad (7)$$

$$\hat{\chi}_c = \sum_{i=1}^N \Gamma_{ic} \log \mu_i, \quad \Gamma_{ic} \doteq \frac{d \log \chi_c}{d \log \mu_i} = \chi_c^{-1} \lambda_i \sum_{f=1}^C Q_{cf} (\Phi_{if} - \Psi_{if}). \quad (8)$$

Here, $\ell^\mu \in \mathbb{R}^{N \times C}$ and $\Gamma \in \mathbb{R}^{N \times C}$ respectively record the elasticities of country-level labor supply and income shares with respect to sectoral markups. The matrix $Q = (I - \Lambda')|_{\mathcal{S}}^{-1}$ denotes the inverse of $I - \Lambda'$ restricted to the subspace $\mathcal{S} = \{\mathbf{x} \in \mathbb{R}^C : \mathbf{1}'\mathbf{x} = 0\}$, where $\Lambda \in \mathbb{R}^{C \times C}$ has entries $\Lambda(c, f) = \Lambda_{cf}^c$, representing country c 's consumption cost exposure to labor from country f .

The first equation in Lemma 1 shows that each sector i 's contribution to country c 's employment gap is proportional to its sales share (λ_i) and to country c 's ownership share

Φ_{ic} . This result provides a significant departure from representative-agent frameworks, such as [Rubbo \(2023\)](#), where sectoral contributions to the aggregate employment gap are governed solely by the Domar weights. The intuition rests on endogenous fluctuations in the country-specific labor wedge. An increase in sectoral markups acts as an implicit tax on production, yet its impact is mediated by income composition. Countries that hold larger ownership shares in highly distorted sectors receive a greater fraction of income in the form of profits relative to total expenditure and therefore exhibit labor wedges that are more sensitive to changes in ex post markups. As ex post markups rise and labor wedges widen, the associated wealth effects are stronger in these countries, leading to larger reductions in labor supply. Consequently, employment responses are inherently shaped by a country's position within the economy's value chain.

The second equation in Lemma 1 characterizes how country c 's income share responds to sectoral markup changes. The structure of the elasticity Γ_{ic} reflects three forces: (i) the sector's centrality in nominal expenditure (λ_i); (ii) the differential between country f 's ownership claims on sector i (Φ_{if}) and sector i 's exposure to labor from country f (Ψ_{if}); and (iii) the country-income network Q , which governs the general equilibrium propagation of income changes across countries.

Structurally, the matrix Q functions as a country-level Leontief inverse or a global income multiplier. In an economy characterized by deep interdependence—where countries consume baskets that embody the labor of others—an initial shift in sectoral profits or labor demand triggers a cascade of expenditure-income feedback. As a country's income share adjusts, its endogenous expenditure response alters the labor demand and income of all other countries in the union. The matrix Q aggregates these infinite rounds of expenditure-income feedback, mapping primitive profit-vs-factor shocks into final equilibrium adjustments.

The economic intuition behind the income redistribution mechanism is as follows. An increase in the markup of sector i simultaneously raises sectoral profits and depresses output demand. The resulting direct redistribution of income across countries is governed by the vector of differences between ownership claims and labor exposure, $\{\Phi_{if} - \Psi_{if}\}_{f=1}^C$: country f benefits from higher markups through its ownership claims on sector i 's profits, captured by Φ_{if} . At the same time, higher markups reduce demand for sector i 's output and, through sector i 's exposure to labor from country f , Ψ_{if} , lead to a contraction in labor demand and labor income in that country. These primitive incidence effects are subsequently processed through the linear operator Q , which accounts for the general-equilibrium feedback across the country network. Since both the primitive incidence and the network propagation are linear, the resulting elasticity of country c 's income share is a linear combination of

all country-level differentials $(\Phi_{if} - \Psi_{if})$, as expressed in equation (8). This expression highlights that income-share adjustments are entirely driven by asymmetries between sectoral profit ownership and sectoral labor exposure across industries and countries.

In the special case with two countries—say, countries h and r —equation (8) admits an explicit solution. The change in country h 's income share is given by

$$\hat{\chi}_h = \chi_h^{-1}(\Lambda_h^r + \Lambda_h^h)^{-1} \sum_{i=1}^N \lambda_i (\Phi_{ih} - \Psi_{ih}) \log \mu_i. \quad (9)$$

This formula makes transparent how markup changes redistribute income across countries. An increase in $\log \mu_i$ raises profits in sector i , benefiting country h 's profit income in proportion to its ownership claims Φ_{ih} . At the same time, higher markups reduce sector i 's demand for labor across countries. Since sector i 's exposure to labor from country h is given by Ψ_{ih} , the contraction in labor demand lowers the labor income accruing to country h in proportion to Ψ_{ih} . These opposing forces partially offset each other, with the net effect scaled by the sector's Domar weight λ_i . The resulting expression summarizes the total impact of markup changes on country h 's income share, incorporating both profit incidence and factor-demand effects.

Remark 1. Combining equation (7) with the relationship between sectoral markups and prices in equation (4), employment gaps can be related to sectoral price changes as follows,

$$(1 + 1/\eta_c) \chi_c \hat{\ell}_c = \lambda' (\delta^{-1} - I) \text{diag}(\Phi(:, c)) \log p.$$

This equation reveals that the transmission of monetary policy to the employment of country c breaks down under each of the following limiting cases: (i) there are no nominal rigidities across sectors, i.e., $\delta_i = 1$ for all i ; (ii) labor supply is perfectly inelastic, $\eta_c = 0$; or (iii) country c has no profit claims across sectors, $\Phi_{ic} = 0$ for all i .

Remark 2. A key implication of equation (8) is that markup shocks have no incidence on country income shares under *labor-equivalent ownership*, defined by $\Phi_{if} = \Psi_{if}$ for all sectors i and countries f . Under this condition, each country receives profits from a sector in exactly the same proportion as it supplies direct and indirect labor services to that sector, so profit-income and labor-income effects exactly offset one another. This eliminates the structural asymmetry that drives income-share redistribution, ensuring that monetary shocks—while still generating aggregate distortions—do not induce any ex post redistribution of income across countries. Consequently, $\hat{\chi}_c = 0$ for all countries.

I now turn to country-level consumption gaps and show that, in a heterogeneous-agent economy, consumption and employment gaps generally do not coincide. Proposition 1 establishes that the difference between consumption and employment gaps is proportional to a sum of sectoral ex post markups.

Proposition 1 (Country-Level Allocative Efficiency). To a first-order (log-linear) approximation, the allocative efficiency of country c , measured by the difference between its consumption and employment gaps, can be decomposed into a direct incidence effect and a [Viner's](#) factoral terms-of-trade effect,

$$\hat{c}_c - \hat{\ell}_c = \underbrace{\sum_{i=1}^N \left(\frac{\lambda_i \Phi_{ic}}{\chi_c} - \lambda_i^c \right) \log \mu_i}_{\text{direct incidence (DI)}} + \underbrace{d \log w_c - \sum_{f=1}^C \Lambda_f^c d \log w_f}_{\text{factoral terms-of-trade (FToT)}}. \quad (10)$$

Furthermore, the FToT component can be rewritten in terms of sectoral markups as

$$\Delta \text{FToT}_c = \sum_{i=1}^N \left[\left(\Gamma_{ic} + \frac{1}{\eta_c} \ell_{ic}^\mu \right) - \sum_{f=1}^C \Lambda_f^c \left(\Gamma_{if} + \frac{1}{\eta_f} \ell_{if}^\mu \right) \right] \log \mu_i. \quad (11)$$

Combining both components, country-level allocative efficiency can be written compactly as a linear exposure to sectoral markups

$$\hat{c}_c - \hat{\ell}_c = \sum_{i=1}^N \mathcal{J}_{ic} \log \mu_i = \mathcal{J}_c' \log \mu,$$

where $\mathcal{J}_c \in \mathbb{R}^N$ collects the total markup-exposure coefficients relevant for the allocative efficiency of country c . For completeness, I also construct the direct-incidence index $\mathcal{J}_c^{\text{DI}}$ and the factoral terms-of-trade index $\mathcal{J}_c^{\text{FToT}}$, which together decompose $\mathcal{J}_c$ into its direct-incidence and terms-of-trade components.¹⁰

The direct incidence (DI) effect represents a generalized terms-of-trade effect that captures how changes in sectoral markups redistribute real income across countries while holding factor prices fixed. In the classical international macroeconomics literature, a country's terms of trade improve when the price of its exports rises relative to its imports, shifting rents from foreign consumers to domestic producers. The direct incidence effect extends

¹⁰The coefficients appearing in the decomposition are $\mathcal{J}_{ic}^{\text{DI}} = \frac{\lambda_i \Phi_{ic}}{\chi_c} - \lambda_i^c$, $\mathcal{J}_{ic}^{\text{FToT}} = \Gamma_{ic} + \frac{1}{\eta_c} \ell_{ic}^\mu - \sum_{f=1}^C \Lambda_f^c \left(\Gamma_{if} + \frac{1}{\eta_f} \ell_{if}^\mu \right)$, and they sum to the total coefficient $\mathcal{J}_{ic} = \mathcal{J}_{ic}^{\text{DI}} + \mathcal{J}_{ic}^{\text{FToT}}$.

this logic to a production-network environment by recognizing each country's dual role as both a consumer of sectoral goods and a residual claimant on sectoral profits.

Specifically, a rise in the markup $\log \mu_i$ increases the cost of country c 's consumption basket by λ_i^c representing a deterioration in the country's terms of trade on the consumption side. At the same time, it increases country c 's income by $\lambda_i \Phi_{ic} / \chi_c$ through its ownership claims on sectoral profits, corresponding to a terms-of-trade improvement on the income side. The direct-incidence index $\mathcal{J}_{ic}^{\text{DI}}$ summarizes the net effect of these opposing forces. It is positive when country c is a net owner of good i , so that profit income more than offsets the increased consumption cost, and negative when country c is a net user, bearing higher user cost without sufficient profit compensation.

The factorial terms-of-trade (FToT) effect, by contrast, captures the redistribution arising from endogenous changes in relative factor prices. It measures the change in relative nominal wages between the labor supplied by country c and the labor to which its consumption basket is exposed. Specifically, it reflects the difference between the change in country c 's own wage, $d \log w_c$, and the exposure-weighted average wage change across all labor types f embodied in its consumption basket, $\sum_{f \in C} \Lambda_f^c d \log w_f$. Since the weights satisfy $\sum_{f \in C} \Lambda_f^c = 1$, this component isolates a pure relative-wage change between labor types across countries. Equation (11) then maps these wage movements to the underlying structural distortions, expressing the FToT effect in terms of sectoral markups.

The following example illustrates how, in a simplified environment, the direct-incidence effect reduces to the classical terms-of-trade channel.

Example 1 (Horizontal Economy). In a currency union without input–output linkages and with domestic ownership segmentation, the direct-incidence index captures export exposure: $\mathcal{J}_{ic}^{\text{DI}} = \sum_{f \neq c} \chi_f \beta_{fi} / \chi_c > 0$ for domestic industries and $\mathcal{J}_{ic}^{\text{DI}} = -\beta_{ci} < 0$ for foreign industries. The direct-incidence effect then maps precisely to the classical terms-of-trade decomposition:

$$\Delta \text{DI}_c = \text{Cov}(\mathcal{J}_c^{\text{DI}}, \log \mu) = \underbrace{\sum_{i \in N_c} \sum_{f \neq c} \chi_f \beta_{fi} / \chi_c \log \mu_i}_{\text{change in export price}} - \underbrace{\sum_{i \notin N_c} \beta_{ci} \log \mu_i}_{\text{change in import price}}. \quad (12)$$

In the general HAIO economy, the direct-incidence effect can therefore be understood as changes in a country's network-adjusted commodity terms of trade.

Taken together, the direct-incidence and factorial terms-of-trade effects show that, in a heterogeneous economy without steady-state distortions, country-level allocative efficiency

arises exclusively from changes in *overall terms of trade*. Monetary policy influences these relative terms of trade by shifting sectoral markups, thereby redistributing allocative efficiency across countries. However, as established in the following corollary, such redistribution is neutral from an aggregate efficiency perspective: when evaluated using income-share-weighted aggregation, these first-order gains and losses perfectly cancel out, reflecting the zero-sum nature of terms-of-trade shifts in a closed economy.

Corollary 1 (Aggregate Neutrality). Country-level allocative efficiencies, while generally nonzero, exactly offset at first order when aggregated using income shares

$$\sum_{c=1}^C \chi_c (\hat{c}_c - \hat{l}_c) = \sum_{c=1}^C \chi_c \mathcal{J}'_c \log \boldsymbol{\mu} = 0. \quad (13)$$

This aggregate neutrality reflects a macro-level envelope theorem. As shown by [Negishi \(1960\)](#), the competitive equilibrium allocation can be replicated by a Negishi planner who maximizes a weighted sum of country utilities.¹¹ As a result, the economy-wide allocative efficiency corresponds to the first-order variation in the planner's welfare function. Since the equilibrium allocation is optimal from the planner's perspective, any marginal reallocation—including that induced by a monetary shock—yields no first-order improvement in aggregate welfare. Monetary policy therefore redistributes across countries without improving union-wide allocative efficiency.

While the aggregate neutrality result holds generally, the pattern of redistribution across countries depends critically on production-network structure. In the horizontal economy of [Example 1](#), the direct-incidence effect arises only when ex post markups in export and import sectors respond differently to monetary policy; a uniform markup change leaves every country's terms of trade unchanged. Input-output linkages break this neutrality, as the following example illustrates.

Example 2 (Two-Stage Vertical Chain). Consider a vertical production chain ([Figure 1](#)) with two countries. For simplicity, I shut down the factorial terms-of-trade channel by assuming labor-equivalent ownership ($\Phi_{if} = \Psi_{if}$) and perfectly elastic labor supply in all countries. The upstream producer (sector 2) uses labor from country 2 to manufacture an intermediate input, which is then purchased by the downstream producer (sector 1). The downstream sector combines this input with labor from country 1 to produce the final good, with a labor share of α_1 . [Proposition 1](#) implies that the allocative efficiencies across countries are

¹¹Lemma 3 shows that when utility is logarithmic in consumption, the Negishi welfare weights coincide with income shares.

determined by the ex post markup of the upstream sector:

$$\hat{c}_1 - \hat{l}_1 = -(1 - \alpha_1) \log \mu_2 \quad \text{and} \quad \hat{c}_2 - \hat{l}_2 = \alpha_1 \log \mu_2.$$

In this setting, an expansionary monetary policy that reduces the upstream sector's markup improves country 1's terms of trade but deteriorates those of country 2. This is because, under labor-equivalent ownership, both countries' terms of trade are independent of changes in downstream markups.¹² Since country 2 is a net supplier of the upstream good while country 1 is a net user, a reduction in the upstream sector's markup benefits users at the expense of suppliers.




Figure 1: Two-stage vertical chain

The role of production networks in redistribution. To formalize how production-network position shapes the redistribution illustrated above, I introduce a country-level input-output multiplier following [Baqee and Farhi \(2019\)](#).

Definition 3. The *input-output multiplier* of country c is

$$\xi_c \equiv \sum_{i \in N_c} \frac{\lambda_i}{\chi_c}, \quad (14)$$

measuring the total domestic sales as a fraction of domestic final demand. The union-wide counterpart is $\mathbb{E}_\chi[\xi] = \sum_{i \in N} \lambda_i$.

When factors are sourced domestically, ξ_c exceeds one whenever country c 's domestic sectors use intermediate inputs.¹³ In particular, under geographic ownership segmentation,

¹²Given the assumed labor-equivalent ownership structure, country 2 owns sector 2 and holds a $1 - \alpha_1$ share of sector 1.

¹³To see this, note that $\chi_c = \Lambda_c = \sum_{i \in N_c} \lambda_i \alpha_{ic} \leq \sum_{i \in N_c} \lambda_i$ at the steady state.

the direct-incidence indices sum to $\sum_{i=1}^N \mathcal{J}_{ic}^{\text{DI}} = \xi_c - \sum_{i=1}^N \lambda_i^c$, motivating the following decomposition.

Corollary 2. Under ownership segmentation ($\Phi_{ic} = \mathbf{1}\{i \in N_c\}$), the direct-incidence effect in Proposition 1 admits the decomposition

$$\Delta \text{DI}_c = \underbrace{(\xi_c - \mathbb{E}_\chi[\xi]) \overline{\log \mu}}_{\text{IO-multiplier channel}} + \underbrace{(\mathbb{E}_\chi[\xi] - \sum_{i \in N} \lambda_i^c) \overline{\log \mu}}_{\text{home-bias channel}} + \underbrace{\sum_{i \in N} \mathcal{J}_{ic}^{\text{DI}} (\log \mu_i - \overline{\log \mu})}_{\text{sectoral-heterogeneity channel}}, \quad (15)$$

where $\overline{\log \mu}$ denotes the union-wide average ex post markup change.

The first two channels are evaluated at the union-wide average markup change. The IO-multiplier channel captures cross-country differences in production-network centrality: countries with above-average IO multipliers ($\xi_c > \mathbb{E}_\chi[\xi]$) experience a terms-of-trade improvement when the union-wide markup rises and a deterioration when it falls. The home-bias channel reflects heterogeneity in countries' total consumption exposure to sectoral wedges, $\sum_{i \in N} \lambda_i^c$, relative to the union-wide counterpart $\sum_{i \in N} \lambda_i = \mathbb{E}_\chi[\xi]$; it vanishes when consumption baskets are homogeneous across countries. The sectoral-heterogeneity channel captures the interaction between a country's direct-incidence exposures and the cross-sectional dispersion of sectoral markup responses.

In the horizontal economy of Example 1, the first two channels are neutralized ($\xi_c = \sum_{i \in N} \lambda_i^c = 1$ for all c), so that the direct-incidence effect arises only when ex post markups in export and import sectors respond differently to monetary policy. Input–output linkages break this neutrality. With production networks, a uniform markup response need not eliminate the direct-incidence motive, because countries differ in their IO multipliers and total consumption exposures. In particular, when the IO-multiplier channel dominates cross-country variation, countries with above-average IO multipliers experience a deterioration in their terms of trade under monetary expansion, which compresses the union-wide markup.

3.3 Monetary Transmission

I now characterize the transmission of monetary policy within the HAO framework. The following proposition derives the general-equilibrium elasticities of sectoral prices and country-level employment with respect to a monetary shock, accounting for the endogenous feedback between nominal rigidities and cross-country heterogeneity.

Proposition 2 (Monetary Transmission in HAIO). Assume that the matrix $I - \varrho^w(\Gamma + \ell^\mu \eta^{-1})'(I - \delta^{-1})$ is nonsingular. Then, in response to a change in the money supply, the elasticities of sectoral prices with respect to nominal aggregate demand are given by

$$\varrho^m \equiv \frac{d \log p}{d \log m} = [I - \varrho^w(\Gamma + \ell^\mu \eta^{-1})'(I - \delta^{-1})]^{-1} \varrho^w \mathbf{1},$$

where $\eta = \text{diag}(\eta_1, \dots, \eta_C)$ denotes the diagonal matrix of household-specific labor supply elasticities. The corresponding effects on household employment are:

$$\ell^m \equiv \frac{d \log L}{d \log m} = \ell^{\mu'}(I - \delta^{-1})\varrho^m.$$

Proposition 2 describes how a monetary expansion is transmitted through the HAIO economy. Sectoral prices respond through two mechanisms. First, a *numéraire effect* raises nominal wages uniformly across countries and, under nominal price rigidities, mechanically raises sectoral prices through marginal-cost pass-through. Second, the *redistribution effect* operates through endogenous variation in ex post markups: monetary shocks reallocate income and employment across countries via the factoral terms-of-trade channel, altering relative nominal wages, which feed back into marginal costs and either amplify or dampen sectoral price responses. The matrix inverse in ϱ^m captures this general-equilibrium feedback loop, while ℓ^m summarizes the resulting labor reallocation across countries. Monetary policy thus moves prices and employment not only through the common numéraire effect but also through endogenous redistribution across countries embedded in the production network.

The following two examples build intuition for how production-network structure shapes the sectoral price pass-through of monetary shocks. In both—a vertical supply chain and a roundabout economy—intermediate-input linkages dampen the price response to monetary shocks, echoing the amplification mechanism of [Basu \(1995\)](#).

Example 3 (General Vertical Economy). This example extends the two-stage vertical chain in Example 2 to a general setting with C countries and C sectors. I maintain the assumptions of labor-equivalent ownership ($\Phi_{if} = \Psi_{if}$) and perfectly elastic labor supply in all countries, which together neutralize the redistribution feedback from ex post markups to nominal wages and collapse the matrix inverse in Proposition 2 to the identity. As illustrated in Figure 2, the most upstream sector C uses labor from country C to produce intermediate good C . Each downstream sector k then combines intermediate input $k + 1$ from its immediate upstream sector with labor supplied by the representative household in country k to

produce good k . Only the final good produced by sector 1 is consumed by households across countries.




Figure 2: General vertical economy

Under this structure, Proposition 2 implies that sector i 's price elasticity satisfies the recursion

$$\varrho_i^m = \alpha_i \delta_i + (1 - \alpha_i) \delta_i \varrho_{i+1}^m,$$

with terminal condition $\varrho_C^m = \delta_C$. Nominal rigidities accumulate along the supply chain: since $\varrho_i^m \in (0, 1)$, a sector that relies more heavily on upstream inputs—smaller α_i —effectively imports the stickiness of its entire upstream network, causing its price to respond more sluggishly to monetary shocks.

Example 4 (Roundabout Economy). Consider a roundabout economy in which all sectors share the same input bundle, $\alpha_{ic} = \bar{\alpha}_c$ and $\omega_{ij} = \bar{\omega}_j$ for all i , while nominal rigidities are allowed to vary across sectors. Proposition 2 then implies,

$$\varrho_i^m = \delta_i / \tilde{\delta} \quad \text{and} \quad \ell_c^m = \frac{\sum_{i=1}^N \Phi_{ic} \lambda_i (1 - \delta_i)}{\tilde{\delta} \chi_c (1 + 1/\eta_c)},$$

where $\tilde{\delta} = \mathbb{E}_b(\delta) + \sum_{i=1}^N \sum_{c=1}^C \lambda_i (1 - \delta_i) \frac{\Phi_{ic}}{1 + 1/\eta_c}$ is a constant that increases in the Domar weights λ_i . As the economy becomes more roundabout—characterized by a higher intensity of intermediate input use (lower $\sum_c \bar{\alpha}_c$), the Domar weights rise, leading to a more sluggish aggregate price adjustment.¹⁴ Production complexity therefore acts as a structural amplifier of nominal rigidities.

¹⁴In this example the Domar weights satisfy $\lambda_i = b_i + \frac{\bar{\omega}_i}{\sum_c \bar{\alpha}_c}$.

4 Optimal Monetary Policy and Inflation Stance

In multi-sector economies with input–output linkages and nominal rigidities, monetary policy cannot attain the first-best allocation and instead trades off competing welfare objectives (see, e.g., [La'O and Tahbaz-Salehi, 2022](#)). This section formulates a Ramsey problem for a heterogeneous-agent input–output economy under a utilitarian Bergson–Samuelson welfare criterion with arbitrary Pareto weights. The resulting characterization of optimal policy motivates the definitions of unilateral inflation stances and policy-alignment losses, which are used to analyze the heterogeneous incidence of a common monetary policy across countries.

4.1 Welfare Loss and Policy Objective

A central question in welfare measurement is how to aggregate country utilities into an economy-wide metric. Throughout the paper, I adopt a utilitarian Bergson–Samuelson welfare function,

$$W(\{\kappa_c\}_{c=1}^C) = \sum_{c=1}^C \kappa_c U_c,$$

where $\{\kappa_c\}_{c=1}^C$ denotes an arbitrary set of non-negative Pareto weights, which reflect the political influence of each country. Let $\mathbf{e}_c$ denote the unilateral Pareto weights, where $\kappa_c = 1$ and $\kappa_{c'} = 0$ for all $c' \neq c$. This corresponds to a scenario in which the monetary authority places full weight on country c 's welfare. Under this formulation, the aggregate welfare criterion collapses to that country's utility: $W(\mathbf{e}_c) = U_c$. The welfare analysis begins with this polar case of policy weighting and then generalizes to arbitrary Pareto weights, leading to the following result.

Lemma 2 (Unilateral Welfare Loss). Up to a second-order approximation, the welfare loss under unilateral Pareto weights for country c is given by:

$$W(\mathbf{e}_c) - W^*(\mathbf{e}_c) = \mathcal{J}'_c \log \boldsymbol{\mu} - \frac{1}{2} \log \boldsymbol{\mu}' \mathcal{L}_c \log \boldsymbol{\mu}, \quad (16)$$

where the matrix $\mathcal{L}_c$ summarizes second-order welfare losses and is decomposed as

$$\mathcal{L}_c \equiv \mathcal{L}_c^{\text{e.g.}} + \mathcal{L}_c^{\text{within}} + \mathcal{L}_c^{\text{across}},$$

with the following components:

$$\begin{aligned}\mathcal{L}_c^{\text{e.g.}} &= (1 + 1/\eta_c) \ell_{(:,c)}^\mu [\ell_{(:,c)}^\mu]', \\ \mathcal{L}_c^{\text{within}} &= \text{diag}(\lambda^c \circ \theta)(\delta^{-1} - I)^{-1}, \\ \mathcal{L}_c^{\text{across}} &= \frac{\eta_c}{1 + \eta_c} \Upsilon^c + \sum_{f=1}^C \frac{1}{1 + \eta_f} \Lambda_f^c \Upsilon^f - \Xi^c + \sum_{f=1}^C \Lambda_f^c \Xi^f,\end{aligned}$$

where the matrices $\Upsilon^c, \Xi^c \in \mathbb{R}^{N \times N}$ summarize the second-order effects of sectoral wedges on equilibrium allocations. In particular, Ξ^c governs the second-order approximation of the income share χ_c , while Υ^c governs the second-order approximation of the labor wedge Λ_c/χ_c . Explicit closed-form expressions for these matrices are provided in Lemma 4 in the Online Appendix.¹⁵

Lemma 2 provides a closed-form second-order approximation of country-level welfare loss in a HAIIO economy. Specifically, equation (16) illustrates that the welfare loss function for country c consists of a first-order allocative efficiency specific to that country, along with a second-order loss comprising three distinct components.

The first two components parallel those in the standard New Keynesian model (e.g., Galí, 2008), reflecting the classical trade-off between employment-gap stabilization and price stability. The first term, $\log \mu' \mathcal{L}_c^{\text{e.g.}} \log \mu = (1 + 1/\eta_c) \hat{l}_c^2$, captures welfare loss from the volatility in the employment gap of country c . This term vanishes when labor supply is perfectly inelastic ($\eta_c = 0$).

The second term, $\log \mu' \mathcal{L}_c^{\text{within}} \log \mu = \sum_{i=1}^N \lambda_i^c \theta_i \frac{\delta_i}{1 - \delta_i} \log \mu_i^2$ measures welfare loss from price dispersion within industries. This loss is strictly increasing with the within-sector elasticity of substitution, θ_i , since greater substitutability amplifies resource misallocation arising from relative price distortions. The key distinction from La'O and Tahbaz-Salehi (2022) lies in the weighting of sectoral misallocation. In this framework, price dispersion in sector i is weighted by country c 's consumption cost exposure, λ_i^c , rather than sector i 's

¹⁵Specifically, Lemma 4 in the Online Appendix formalizes the closed-form expressions

$$\begin{aligned}\Xi_{ij}^c &= -\Gamma_{jc} \Gamma_{ic} + \chi_c^{-1} \sum_{f=1}^C Q_{cf} (\Phi_{if} - \Psi_{if}) \left(2 \sum_{g=1}^C \lambda_i^g \chi_g \Gamma_{jg} - 2 \lambda_j \Psi_{ji} + \lambda_i \iota_{ij} \right), \\ \Upsilon_{ij}^c &= \frac{\Phi_{ic}}{\chi_c} \left(2 \lambda_j \Psi_{ji} - 2 \sum_{f=1}^C \lambda_i^f \chi_f \Gamma_{jf} + 2 \lambda_i \Gamma_{jc} - \lambda_i \iota_{ij} - \lambda_i \lambda_j \frac{\Phi_{jc}}{\chi_c} \right),\end{aligned}$$

where ι_{ij} denotes the (i, j) -th entry of the identity matrix. Both matrices are normalized to be symmetric.

Domar weight, λ_i . This refinement highlights how heterogeneity in consumption baskets across countries generates unequal exposure to sectoral inefficiencies.

The third term, $\log \mu' \mathcal{L}_c^{\text{across}} \log \mu$, arises exclusively in multisector economies and characterizes the misallocation of resources across industries. Building on second-order approximations of labor wedges and income shares, it extends the cross-sector misallocation term in [Baqee and Farhi \(2024\)](#) to the country level. In a standard representative-agent production network economy, such a term typically captures the second-order approximation of the aggregate labor wedge (see e.g., [Baqee and Farhi, 2020](#); [La'O and Tahbaz-Salehi, 2022](#)). In this heterogeneous-agent framework, however, the term is broader in scope: it integrates the second-order approximation of country c 's labor wedge, Λ_c/χ_c , alongside the second-order approximation of the country's factoral terms-of-trade. Consequently, it captures not only the standard inefficiencies in sectoral labor allocation but also the distortionary effects of monetary shocks on the relative factor prices embedded in the country's consumption basket.

Building on Lemma 2, Proposition 3 establishes the general Bergson–Samuelson welfare loss by aggregating the unilateral welfare losses using arbitrary Pareto weights.

Proposition 3 (Bergson–Samuelson Welfare Loss). The Bergson–Samuelson welfare loss function with arbitrary Pareto weights $\{\kappa_c\}_{c=1}^C$ can be expressed as the weighted sum of country-level welfare losses:

$$\begin{aligned} W(\{\kappa_c\}_{c=1}^C) - W^*(\{\kappa_c\}_{c=1}^C) &= \sum_{c=1}^C \kappa_c [W(\mathbf{e}_c) - W^*(\mathbf{e}_c)] \\ &= \sum_{c=1}^C \kappa_c \mathcal{J}'_c \log \mu - \frac{1}{2} \log \mu' \mathcal{L}(\{\kappa_c\}) \log \mu, \end{aligned} \quad (17)$$

where the aggregate loss matrix $\mathcal{L}(\{\kappa_c\})$ is the Pareto-weighted average of country-level loss matrices,

$$\mathcal{L}(\{\kappa_c\}) \equiv \sum_{c=1}^C \kappa_c \mathcal{L}_c = \mathcal{L}^{\text{e.g.}}(\{\kappa_c\}) + \mathcal{L}^{\text{within}}(\{\kappa_c\}) + \mathcal{L}^{\text{across}}(\{\kappa_c\}) \quad (18)$$

with each component of $\mathcal{L}(\{\kappa_c\})$ defined as

$$\mathcal{L}^{\text{e.g.}}(\{\kappa_c\}) = \sum_{c=1}^C \kappa_c \mathcal{L}_c^{\text{e.g.}}, \quad \mathcal{L}^{\text{within}}(\{\kappa_c\}) = \sum_{c=1}^C \kappa_c \mathcal{L}_c^{\text{within}}, \quad \mathcal{L}^{\text{across}}(\{\kappa_c\}) = \sum_{c=1}^C \kappa_c \mathcal{L}_c^{\text{across}}.$$

Proposition 3 generalizes the welfare loss expression from a New Keynesian production network economy with a representative agent (e.g., La'O and Tahbaz-Salehi, 2022) to a fully heterogeneous-agent input–output (HAIO) economy. While Baqaee and Farhi (2024) similarly provide a second-order approximation of welfare in a HAIO framework, their analysis is restricted to the special case where Pareto weights are given by income shares. In contrast, Proposition 3 accommodates arbitrary welfare weights, facilitating a more flexible evaluation of the distributional incidence of a common monetary policy. This generalization is central to characterizing the policy-alignment loss (PAL) in Section 4.4, as it explicitly accounts for the potential misalignment between a country's unilateral preferences and the centralized objectives of the monetary authority.

In another comparison to Baqaee and Farhi (2024), the welfare function in equation (17) implies that monetary policy faces a fundamental trade-off between aggregate stabilization and first-order redistribution. These redistributive motives are silenced under the income-share weighting assumed in their framework, forcing the central bank to focus exclusively on aggregate efficiency. The following corollary formalizes this result, showing that the welfare objective collapses to a purely efficiency-based, quadratic measure when Pareto weights coincide with countries' income shares.

Corollary 3 (Redistribution Neutrality under Income-Share Weighting). When Pareto weights coincide with countries' income shares, such that $\kappa_c = \chi_c$ for all c , the optimal monetary policy is redistribution-neutral. In this case, the central bank cannot enhance aggregate efficiency through manipulating the countries' overall terms of trade. Consequently, the aggregate welfare loss becomes purely quadratic in ex post markups¹⁶

$$W(\{\chi_c\}_{c=1}^C) - W^*(\{\chi_c\}_{c=1}^C) = -\frac{1}{2} \log \mu' \mathcal{L}(\{\chi_c\}) \log \mu.$$

4.2 Optimal Policy

To characterize optimal monetary policy, I model the monetary authority as a Ramsey planner who chooses the policy to minimize aggregate welfare loss, subject to the competitive-equilibrium constraints of the economy. The nominal policy instrument is the money supply $\log m(z)$, which affects equilibrium prices and markups through the cash-in-advance constraint.

While policy is implemented via $\log m(z)$, welfare depends only on the induced movements in prices and markups. It is therefore without loss of generality to represent monetary

¹⁶See Corollary 4 in the Online Appendix for a full characterization of the welfare loss function under income-share weighting, which yields a substantially simplified expression.

policy by a sectoral price index target of the form

$$\zeta' \log \mathbf{p} = \pi,$$

for some $(\zeta, \pi) \in \mathcal{P} \subseteq \mathbb{R}^N \times \mathbb{R}$. This representation captures the reduced-form implications of monetary policy through its impact on sectoral prices.

Lemma 6 in the Online Appendix shows that, to a first-order approximation, any price-targeting rule of the form $\zeta' \log \mathbf{p} = \pi$ can be implemented by an appropriate choice of the money supply rule $\log m(z)$. In particular, for any admissible (ζ, π) , there exists a policy of the form $\log m(z) = \varsigma_0 + \varsigma' \log z$ that reproduces the same equilibrium price allocation. This result justifies treating (ζ, π) as reduced-form policy instruments throughout the analysis.

Building on this result, the Ramsey-optimal policy under arbitrary Pareto weights is characterized by the optimal price-targeting index ζ^* and the associated inflation bias π^* that jointly minimize the aggregate welfare loss specified in equation (17).

Theorem 1 (Optimal Policy). Given the Pareto weights $\{\kappa_c\}_{c=1}^C$, the optimal monetary policy solves for (ζ^*, π^*) that minimizes the welfare loss in equation (17). The resulting policy targets a sectoral price index of the form

$$\sum_{j=1}^N \zeta_j^*(\{\kappa_c\}) \log p_j = \pi^*(\{\kappa_c\}), \quad (19)$$

where

$$\pi^*(\{\kappa_c\}) = \sum_{c=1}^C \kappa_c \mathcal{J}'_c(I - \delta^{-1}) \varrho^m. \quad (20)$$

The optimal sectoral weight can be further broken down into three components,

$$\zeta_j^*(\{\kappa_c\}) = \zeta_j^{\text{e.g.}}(\{\kappa_c\}) + \zeta_j^{\text{within}}(\{\kappa_c\}) + \zeta_j^{\text{across}}(\{\kappa_c\}),$$

with

$$\begin{aligned}\zeta_j^{\text{e.g.}}(\{\kappa_c\}) &= (\delta_j^{-1} - 1) \lambda_j \sum_{c=1}^C \kappa_c \Phi_{jc} \ell_c^m / \chi_c, \\ \zeta_j^{\text{within}}(\{\kappa_c\}) &= (\delta_j^{-1} - 1) \theta_j \varrho_j^m \sum_{c=1}^C \kappa_c \lambda_j^c, \\ \zeta_j^{\text{across}}(\{\kappa_c\}) &= (\delta_j^{-1} - 1) \sum_{i=1}^N (\delta_i^{-1} - 1) \varrho_i^m \mathcal{L}_{ij}^{\text{across}}(\{\kappa_c\}).\end{aligned}$$

Theorem 1 characterizes the optimal policy as a function of sectoral nominal rigidities, the production network, and ownership structures. As implied by Corollary 3, the monetary authority faces a trade-off between aggregate stabilization and first-order redistribution whenever Pareto weights deviate from income shares. In such cases, an inflation bias arises on the right-hand side of equation (19), reflecting incentives to redistribute across countries by shifting the overall terms of trade in favor of certain groups. By contrast, when Pareto weights align with income shares, the optimal monetary policy reduces to a pure price-stabilization rule with zero inflation bias ($\pi^*(\{\chi_c\}) = 0$). This redistribution-neutral policy serves as the benchmark for centralized optimal monetary policy.¹⁷

The optimal industry weights in equation (19) determine which price index monetary policy should target to minimize second-order welfare losses arising from nominal rigidities. Recall that the monetary authority must balance between three distinct sources of misallocation: the aggregate volatility of employment gaps, price dispersion within sectors, and price dispersion across sectors. Accordingly, the optimal weights decompose into three components, each capturing the relative welfare importance of these distortions.

4.3 Unilateral Inflation Stance

This section studies monetary policy from the perspective of an individual member country. Taking the unilateral optimal policy as a benchmark, it characterizes the constrained unilateral inflation stance under the union's common price-index regime and examines its relationship to the production structure.

To facilitate the analysis, policy instruments are mapped into sectoral markups, the channel through which monetary policy affects real allocations. Given the realization of

¹⁷Corollary 5 in the Online Appendix shows that when cross-country asymmetries are shut down by homogeneous consumption baskets and ownership shares aligned with income shares (i.e., $\beta_{ci} = b_i$ and $\Phi_{ic} = \chi_c$ for all c and i), the centralized optimal policy coincides with the one derived in La'O and Tahbaz-Salehi (2022).

productivity shocks $\log z$, the vector of ex post markups is, to a first-order approximation, a differentiable function of the policy pair $(\zeta, \pi) \in \mathcal{P} \subseteq \mathbb{R}^N \times \mathbb{R}$:

$$\log \mu = \mathcal{M}(\zeta, \pi; z),$$

where $\mathcal{M} : \mathcal{P} \rightarrow \mathbb{R}^N$ represents the reduced-form mapping from monetary policy instruments to sectoral markups.¹⁸ For notational simplicity, the dependence on z is suppressed in the following analysis.

Substituting this mapping into country c 's unilateral welfare loss in equation (16) yields

$$\mathbb{L}_c(\zeta, \pi) \equiv -\mathcal{J}'_c \mathcal{M}(\zeta, \pi) + \frac{1}{2} \mathcal{M}(\zeta, \pi)' \mathcal{L}_c \mathcal{M}(\zeta, \pi).$$

As established in Theorem 1, assigning full weight to country c implies that the associated optimal policy solves

$$(\zeta^*(\mathbf{e}_c), \pi^*(\mathbf{e}_c)) \in \arg \min_{(\zeta, \pi) \in \mathcal{P}} \mathbb{L}_c(\zeta, \pi).$$

This policy therefore defines country c 's unilateral optimal policy and implements a second-best allocation by minimizing its welfare loss over the full set of monetary instruments.

In practice, however, such an allocation is generally unattainable, as it requires the country to exercise sole authority over those instruments—an arrangement precluded by the union's common price-index regime ζ .¹⁹ Once the price-index regime is taken as given and inflation is restricted to be non-state-contingent, the country's decision problem becomes a third-best one. The following definition formalizes the associated third-best inflation stance.

Definition 4. The unilateral inflation stance of country c is the inflation rate it selects to minimize its expected welfare loss, taking the union's target price-index regime ζ (normalized to sum to one) as given

$$\pi_c(\zeta) = \arg \min_{\pi \in \mathbb{R}} \mathbb{E}[\mathbb{L}_c(\zeta, \pi)].$$

Unlike the second-best unilateral optimal policy, which minimizes welfare loss for every realization of sectoral shocks, the inflation stance is defined in expectation. The expectation operator reflects the non-state-contingent inflation constraint: the country must choose a single union-wide inflation rate ex ante rather than adjust inflation state

¹⁸See Lemma 7 in the Online Appendix for the full expression.

¹⁹Lemma 8 shows that, under a common price-index regime, replicating the unilateral optimal policy requires state-contingent inflation.

by state. Consequently, the third-best policy $(\zeta, \pi_c(\zeta))$ weakly underperforms the second-best unilateral optimal policy $(\zeta^*(\mathbf{e}_c), \pi^*(\mathbf{e}_c))$. Nevertheless, as the following propositions demonstrate, the two objects remain closely connected.

Proposition 4. Let $(\zeta^*(\mathbf{e}_c), \pi^*(\mathbf{e}_c))$ denote country c 's unilateral optimal policy. Given the union's target price-index regime ζ (normalized to sum to one), the unilateral inflation stance of country c satisfies

$$\pi_c(\zeta) = \underbrace{\frac{\zeta' \mathbf{q}^m}{\zeta^*(\mathbf{e}_c)' \mathbf{q}^m}}_{\text{adjustment scalar due to different target price indices}} \pi^*(\mathbf{e}_c). \quad (21)$$

Equation (21) shows that the unilateral inflation stance is proportional to the inflation bias $\pi^*(\mathbf{e}_c)$ of the unilateral optimal policy, scaled by a factor that accounts for the difference between the country's preferred target price index $\zeta^*(\mathbf{e}_c)$ and the union's common index ζ . When the two indices coincide, the factor equals one and the stance replicates the unilateral optimal inflation bias. Economically, the unilateral inflation stance summarizes how a country's overall terms of trade respond to monetary expansions in the presence of nominal rigidities. A positive stance indicates that monetary expansion improves the country's terms of trade; a negative stance implies the opposite. By analogy with equations (10) and (15), the stance admits a parallel decomposition: first into a direct-incidence component and a factoral terms-of-trade component — reflecting how a country's commodity terms of trade and factoral terms of trade, respectively, respond to monetary expansion; and further, within direct incidence, into IO-multiplier, home-bias, and sectoral-heterogeneity channels.

Proposition 5. For each country c , the unilateral inflation stance $\pi_c(\zeta)$ replicates the expected first-order allocative-efficiency gain generated by country c 's unilateral optimal policy:

$$\mathbb{E}[\mathcal{J}'_c \mathcal{M}(\zeta, \pi_c(\zeta))] = \mathbb{E}[\mathcal{J}'_c \mathcal{M}(\zeta^*(\mathbf{e}_c), \pi^*(\mathbf{e}_c))] \geq 0. \quad (22)$$

By contrast, any centralized policy with zero inflation bias is redistribution-neutral to first order: $\mathbb{E}[\mathcal{J}'_c \mathcal{M}(\zeta, 0)] = 0$ for all c .

Despite the institutional constraints imposed by the union's price-index regime, Proposition 5 demonstrates that the unilateral inflation stance preserves the expected first-order effects of the unconstrained unilateral optimal policy.

4.4 Policy-Alignment Loss

To assess how divergent policy preferences translate into welfare outcomes, this section introduces the policy-alignment loss (PAL). The PAL measures the welfare loss a country incurs when the implemented common monetary policy differs from the policy that would be optimal from that country's own perspective. It therefore provides a within-union measure of policy misalignment and a benchmark for evaluating the distributional consequences of centralized monetary policy in a heterogeneous-agent economy.

Definition 5. Let $(\zeta^*(\mathbf{e}_c), \pi^*(\mathbf{e}_c))$ denote country c 's unilateral optimal monetary policy. Given a union-wide monetary policy (ζ, π) , the *policy-alignment loss* (PAL) for country c is defined as

$$\text{PAL}_c(\zeta, \pi) \equiv \mathbb{L}_c(\zeta, \pi) - \mathbb{L}_c(\zeta^*(\mathbf{e}_c), \pi^*(\mathbf{e}_c)) \geq 0, \quad (23)$$

where the inequality holds because (ζ, π) remains a feasible policy in country c 's optimization problem, but may not be optimal from its individual perspective.

A country's policy-alignment loss reflects two distinct dimensions of disagreement with the union-wide policy: the inflation rate implemented under the common instrument and the composition of the price index used to define that instrument. To separate these two dimensions, consider the intermediate policy $(\zeta, \pi_c(\zeta))$, which combines the union's target price index with country c 's unilateral inflation stance conditional on that index. Then PAL can be written as

$$\text{PAL}_c(\zeta, \pi) = \underbrace{\mathbb{L}_c(\zeta, \pi) - \mathbb{L}_c(\zeta, \pi_c(\zeta))}_{\text{inflation misalignment}} + \underbrace{\mathbb{L}_c(\zeta, \pi_c(\zeta)) - \mathbb{L}_c(\zeta^*(\mathbf{e}_c), \pi^*(\mathbf{e}_c))}_{\text{price-index misalignment}}. \quad (24)$$

The first component measures the welfare cost of implementing the inflation rate π rather than $\pi_c(\zeta)$, holding the price index fixed at ζ . The second term measures the remaining loss from targeting a price index that differs from country c 's optimal index, even when inflation is set at the country's preferred level. Proposition 6 provides closed-form expressions for each component in expectation.

Proposition 6. The expected policy-alignment loss admits the decomposition:

$$\mathbb{E}[\text{PAL}_c(\zeta, \pi)] = \underbrace{\frac{1}{2} \frac{\zeta^*(\mathbf{e}_c)' \varrho^m}{(\zeta' \varrho^m)^2} [\pi_c(\zeta) - \pi]^2}_{\text{inflation misalignment}} + \underbrace{\mathcal{I}_c(\zeta)}_{\text{price-index misalignment}}, \quad (25)$$

where $\pi_c(\zeta)$ is country c 's unilateral inflation stance defined in equation (21), and $\mathcal{I}_c(\zeta) \geq 0$ measures the variance loss from targeting a price index ζ that differs from country c 's optimal index $\zeta^*(\mathbf{e}_c)$, with $\mathcal{I}_c(\zeta^*(\mathbf{e}_c)) = 0$.²⁰

The decomposition shows that inflation misalignment is quadratic in the *inflation stance deviation*—the distance between country c 's unilateral inflation stance and the union-wide consensus. By contrast, price-index misalignment is independent of π and represents an irreducible loss from using a target price index that differs from the country's preferred index. When price-index misalignment is quantitatively negligible, the inflation stance deviation provides a sufficient statistic for its policy-alignment loss.

5 Quantitative Analysis

This section applies the theoretical framework to the euro area. I first evaluate union-wide welfare by solving for the centralized optimal policy and comparing it to alternative price-stabilization rules. I then analyze countries' unilateral welfare, computing each country's preferred inflation stance and policy-alignment loss, and relate these outcomes to positions in the union's production network.

5.1 Calibration

Given the availability of rich cross-country input–output data, I focus on the euro area as the empirical context.²¹ The model is calibrated at a quarterly frequency. To construct the input–output matrix Ω , I use data from the World Input–Output Database (WIOD) (Timmer et al., 2015), which provides annual IO tables for 44 countries, each with 56 industries, from 2000 to 2014. I extract the 20 euro area member countries and reconstruct their integrated production network, preserving all intra-euro linkages at the country–industry level. Assuming that direct factor inputs are entirely domestically sourced, I set α_{ic} equal to the value-added share of industry i when it is located in country c , and zero otherwise. For

²⁰See Online Appendix B for a closed-form expression.

²¹For example, interstate input–output linkages are not directly available in U.S. data.

each country c , I construct the vector of consumption shares, β_c , to replicate the distribution of final uses across industries within that country. Income shares (GNE shares), χ_c , are calibrated from each country's total nominal expenditure relative to the euro area aggregate.

Given the industry classification provided by the WIOD, I next calibrate nominal rigidities across sectors. Sectoral price flexibilities are inferred from the frequency of price adjustment (FPA) data of [Pastén et al. \(2020\)](#), which provide monthly FPA estimates derived from U.S. Producer Price Index (PPI) microdata. Although the model is calibrated to a European production network, U.S. PPI data are well suited for three reasons. First, the cross-country difference in price rigidity at the consumer level is modest: after excluding sales, [Gautier et al. \(2024\)](#) report that 8.5% of consumer prices adjust each month in the euro area versus 10% in the United States. Second, PPI-based FPA captures price adjustment among producers—the relevant margin for a production network model—rather than retail prices faced by final consumers. Third, the PPI-based estimates provide broader sectoral coverage, spanning 54 of the 56 WIOD sectors, compared with 46 sectors for the CPI-based estimates of [Gautier et al. \(2024\)](#). Indeed, across the 46 jointly covered sectors, the PPI-based FPA is on average considerably higher than the CPI-based FPA (correlation of 0.49), a pattern consistent with the finding of [Nakamura and Steinsson \(2008\)](#) that intermediate-goods producer prices adjust more frequently than consumer prices. Each WIOD industry is matched to the closest NAICS code with available FPA estimates, following the mapping procedure described in the Online Appendix. Nominal rigidities are assumed to be uniform across countries but vary across industries according to the matched FPA data. In addition, following standard practice in the production network literature (e.g., [La'O and Tahbaz-Salehi, 2022](#)), I incorporate nominal wage rigidities by introducing, for each country, a pseudo-industry that transforms domestic labor into labor services supplied to all other domestic industries. The degree of nominal wage rigidity is calibrated based on the empirical estimates of [Barattieri et al. \(2014\)](#) and [Beraja et al. \(2019\)](#), setting the price flexibility of the pseudo-industry to 0.30.

I then calibrate the distribution of sectoral productivity shocks using data from the WIOD Socio-Economic Accounts (2000–2014). I construct annual series of productivity growth rates through a Törnqvist decomposition based on observed inputs and outputs, and linearly interpolate them to quarterly frequency to match the model's time period. The log productivity shocks, $(\log z_1, \dots, \log z_N)$, are assumed to be jointly normally distributed, with variance–covariance matrix matched to the variance–covariance matrix of the linearly detrended quarterly productivity series.

Finally, I set the within-industry elasticity of substitution to $\theta_i = 6$ for all industries

and the Frisch elasticity of labor supply to $\eta_c = 2$ for all countries. These parameter values are consistent with standard calibrations in the New Keynesian literature. The baseline calibration also assumes ownership segmentation, under which firms are owned exclusively by their domestic countries.

5.2 Welfare Comparison of Alternative Monetary Policies

This section evaluates union-wide expected welfare losses under alternative monetary policy regimes in a currency union. The analysis adopts a welfare aggregation in which Pareto weights coincide with countries' income shares, implying that centralized monetary policy addresses a pure efficiency problem. This specification is commonly used in the literature and provides a natural benchmark for the euro area. Although the European Central Bank (ECB) does not explicitly assign welfare weights to member states, its institutional design—such as permanent membership on the Executive Board and the rotating voting system of the Governing Council—implicitly links policy influence to economic size.²²

Table 2 reports the expected welfare losses arising from nominal rigidities, expressed as a percentage of steady-state union-wide real consumption. Following the decomposition in equation (18), the total loss is partitioned into three distinct sources of misallocation. The first column presents results under the centralized optimal policy, which yields an expected welfare loss equivalent to 0.584% of quarterly consumption relative to the flexible-price equilibrium. The largest source of this welfare loss is within-industry misallocation, responsible for 0.327 percentage points of the total loss, followed by across-industry misallocation at 0.248 percentage points. In contrast, the welfare cost attributable to the aggregate volatility of employment gaps is an order of magnitude smaller.²³

Table 2 also compares the performance of the optimal policy with two alternative price-stabilization rules. The first is an employment-gap (EG) stabilization policy, which minimizes the aggregate volatility of employment gaps across countries.²⁴ This corresponds to a price-stabilization rule of the form $\sum_{i=1}^N \zeta_i^{\text{e.g.}}(\{\chi_c\}) \log p_i = 0$, where the industry weights are given by $\zeta_i^{\text{e.g.}}(\{\chi_c\}) = \lambda_i(\delta_i^{-1} - 1)(\sum_{c=1}^C \Phi_{ic} \ell_c^m)$, which simplifies under ownership segmentation to $\zeta_i^{\text{e.g.}}(\{\chi_c\}) = \lambda_i(\delta_i^{-1} - 1)\ell_c^m$ for $i \in N_c$. Intuitively, the EG stabilization policy targets a price index that assigns greater weight to industries that are larger (higher λ_i), stickier (lower δ_i),

²²See Online Appendix D, which validates this benchmark using voting-based Pareto weights.

²³The ordering between within- and across-industry misallocations reflects the model's structure: the high elasticity of substitution within sectors ($\theta_i = 6$) makes within-sector price dispersion more damaging than the across-sector misallocation which is controlled by the unit across-sector elasticity implied by the Cobb–Douglas production functions.

²⁴Formally, the EG stabilization policy minimizes $\log \mu' \mathcal{L}^{\text{e.g.}}(\{\chi_c\}) \log \mu = \sum_{c=1}^C (1 + 1/\eta_c) \chi_c \hat{p}_c^2$.

Table 2: Expected welfare losses under various policies.

	(1) Cent. Optimal	(2) EG Stabilization	(3) CPI Stabilization
Total welfare loss	0.584	0.592	0.601
Employment Gaps	0.009	0.006	0.008
Within-sector misallocation	0.327	0.339	0.346
Across-sector misallocation	0.248	0.246	0.246
Cosine similarity to optimal policy	1	0.976	0.120

and located in countries with more pronounced employment responses to monetary shocks (higher ℓ_c^m).

As shown in the second column of the table, the EG stabilization policy generates a welfare loss equivalent to a 0.592% reduction in quarterly consumption, and is just 0.008 percentage points higher than under the optimal policy. It achieves slightly better stabilization of employment gaps but at the cost of greater within-sector misallocation. The approximate optimality of the EG stabilization echoes prior results in [La'O and Tahbaz-Salehi \(2022\)](#) and [Rubbo \(2023\)](#), which show that stabilizing the output or employment gap can be approximately optimal. A key distinction here is that the heterogeneous-agent structure implies multiple employment gaps (one per country), so the EG policy must balance stabilization across member states. Importantly, the target price indices implied by the two policies are also very similar: their cosine similarity exceeds 97%. Thus, the centralized optimal policy implicitly shares the same logic as the EG stabilization, assigning greater weight to sectors that are larger, stickier, and belong to countries with more pronounced employment responses.

The third column examines a CPI stabilization rule, $\sum_{i=1}^N b_i \log p_i = 0$, which stabilizes the union-wide consumer price index. Relative to the other two policies, CPI stabilization performs substantially worse, generating larger welfare losses primarily due to within-sector misallocation. This underperformance reflects a pronounced misalignment between the CPI and the welfare-relevant price index targeted by the optimal policy, as indicated by a cosine similarity of only 12 percent. To illustrate this difference in underlying price indices, [Figure A.1](#) in the Online Appendix aggregates industry weights to the country level and plots each country's weight in the target price index against its income share under both the optimal policy and CPI stabilization. Under CPI stabilization, country weights closely track income shares, with most observations near the 45-degree line. In contrast, under the optimal policy, while weights broadly reflect income shares, five countries—Belgium, Germany,

Ireland, Luxembourg, and the Netherlands—receive systematically higher weights. These countries are often seen as occupying upstream positions in the union production network, consistent with the theoretical result that the optimal policy stabilizes a price index with greater weights assigned to more upstream industries.

5.3 Country-Level Welfare and Policy Alignment

The previous section evaluated alternative stabilization rules from a union-wide welfare perspective. This section shifts attention to the country level, examining how a common monetary policy translates into heterogeneous welfare outcomes across member states.

Figure 3 provides a first step in this direction by comparing country-level welfare losses under alternative monetary policy regimes. Panel (a) shows that CPI stabilization generates higher welfare losses than the centralized optimal policy for nearly all countries: 18 of 20 observations lie above the 45-degree line, with Ireland and Malta as the only exceptions, though the differences are economically negligible. Switching from CPI stabilization to the optimal policy would therefore yield welfare gains across member states. Panel (b) compares welfare losses under the centralized optimal policy with those under each country’s unilateral optimal policy. All observations lie below the 45-degree line, as the vertical distance to the frontier corresponds to the policy-alignment loss defined in equation (23).

In light of the uniform dominance documented in Panel (a), the remainder of the analysis adopts the centralized optimal policy as the benchmark. All subsequent figures report policy-alignment losses and unilateral inflation stances relative to this benchmark, with results under CPI stabilization reported in the Online Appendix.

To shed light on the sources of the policy-alignment losses documented above, Figure 4 provides two complementary decompositions. Panel (a) decomposes each country’s PAL into first- and second-order components. The first-order component in equation (22) isolates the welfare gain a country derives from manipulating its overall terms of trade under unilateral optimal policy. By plotting the total PAL against this first-order component, a striking pattern emerges: for 18 out of 20 countries, the first-order component exceeds the total policy-alignment loss, placing these observations above the 45-degree line. This outcome highlights a significant redistribution–stabilization trade-off under unilateral policymaking. Once a country gains unilateral control over monetary policy, it can aggressively pursue its own first-order redistributive objectives—even at the cost of larger second-order welfare losses due to price dispersion and employment volatility. By contrast, Panel (b) decomposes PAL into inflation misalignment and price-index misalignment, as




Figure 3: Welfare loss under alternative monetary policies. *Note:* Each dot represents a euro-area country, with size proportional to its income share. For each country, welfare loss is expressed as a percentage of its steady-state consumption. Panel (a) compares welfare losses under the centralized optimal policy and CPI stabilization. Panel (b) compares welfare losses under the centralized optimal policy and each country's unilateral optimal policy. The 45-degree line indicates identical welfare outcomes across policies.

in equation (25). Observations cluster tightly around the 45-degree line, confirming that policy-alignment losses are driven primarily by inflation misalignment rather than by disagreement over the composition of the target price index. This implies that, under the centralized optimal price index, setting union-wide inflation equal to a country's unilateral inflation stance would be approximately optimal from that country's perspective.

To better understand the sources of these welfare losses, I next zoom in on a representative member state and decompose its welfare losses across various policies using equation (16). Relative to the union-wide analysis in the previous section, country-level welfare loss includes an additional component arising from underutilized allocative efficiency. Using Spain as an illustrative case, Table 3 reports welfare losses, expressed as a percentage of Spain's steady-state consumption, under the unilateral optimal policy, the centralized optimal policy, CPI stabilization, as well as counterfactual policies that retain the union's price-targeting regime but implement a country's unilateral inflation stance. For completeness, Table A.1 in the Online Appendix reports the corresponding welfare losses for all 20 euro-area countries under the same set of policies.

The first column reports the unilateral optimal policy, in which the country fully internalizes monetary policy decisions at the union level and chooses both the target




Figure 4: Decomposition of policy-alignment loss. *Note:* Panels (a) and (b) plot total policy-alignment loss against its underutilized allocative-efficiency and inflation-misalignment components, respectively, under the centralized optimal policy. The 45-degree line indicates full accounting.

price index and the inflation rate to minimize its own welfare loss. The associated price-targeting rule is $\sum_{i=1}^N \zeta_i^*(\mathbf{e}_c) \log p_i = \pi^*(\mathbf{e}_c)$. The second and fourth columns correspond to the centralized optimal policy and CPI stabilization, respectively, as introduced in the previous section. The third column keeps the centralized optimal price-index regime but sets inflation equal to Spain's unilateral inflation stance, $\sum_{i=1}^N \zeta_i^*(\{\chi_c\}) \log p_i = \pi_c(\zeta^*(\{\chi_c\}))$. Similarly, the fifth column retains CPI targeting while allowing Spain to choose its preferred inflation rate under that regime, $\sum_{i=1}^N b_i \log p_i = \pi_c(b)$.

Table 3: Expected welfare loss for Spain under various policies.

	(1) Uni. Opt.	(2) Cent. Opt.	(3) Cent. Uni. π	(4) CPI Stab.	(5) CPI Uni. π
Total welfare loss	0.333	0.369	0.336	0.384	0.351
Allocative efficiency (FO)	-0.066	0	-0.066	0	-0.066
Employment gaps	0.010	0.004	0.013	0.006	0.014
Within-sector misallocation	0.271	0.247	0.271	0.261	0.284
Across-sector misallocation	0.118	0.117	0.118	0.117	0.118
Cosine similarity to unilateral optimal policy	1	0.315	0.315	0.035	0.035

Three results are worth highlighting. First, welfare losses due to underutilized allocative efficiency are quantitatively important under stabilization policies. As shown in Proposi-

tion 5, the expected first-order allocative efficiency term is invariant across policies (1), (3), and (5), and equals zero under stabilization policies (2) and (4). Allowing Spain to set its preferred inflation rate (policies (1), (3), and (5)) yields substantially lower welfare losses than pure stabilization policies (policies (2) and (4)). These gains reflect improvements in allocative efficiency driven by a favorable adjustment of the overall terms of trade, albeit at the cost of larger second-order losses arising from employment volatility and within-sector price dispersion.

Second, policy-alignment losses are largely attributable to inflation misalignment. Under both the centralized optimal and CPI regimes, aligning the union-wide inflation rate with Spain's unilateral inflation stance defined under the corresponding price-index regime substantially reduces welfare losses toward the unilateral optimum. In particular, under the centralized optimal price-index regime, policy (3) delivers welfare outcomes very close to the unilateral optimal policy, with a remaining loss of only 0.003 percentage points of consumption. This near-optimality indicates that policy disagreements within a currency union mainly reflect heterogeneity in preferred inflation rates rather than the design of the target price index.

Third, price-index misalignment is generally smaller under the centralized optimal regime than under CPI targeting. The price index targeted by the centralized optimal policy is considerably closer to that implied by the unilateral optimal policy (cosine similarity of 0.315) than is the CPI (cosine similarity of 0.035). As a result, policy (3) substantially outperforms policy (5). For the same reason, the centralized optimal policy (2) dominates CPI stabilization (4), consistent with Panel (a) of Figure 3, which shows uniformly lower welfare losses under the centralized optimal policy across countries.

5.4 Inflation Stance and Policy-Alignment Loss

When policy-alignment losses are driven primarily by inflation misalignment, as established in the previous subsection, a country's unilateral inflation stance becomes the key object predicting the welfare incidence of a common monetary policy across member states. By Proposition 6, policy-alignment loss is quadratic in the inflation stance deviation. Under income-share weighting, the centralized optimal policy $(\zeta^*(\{\chi_c\}), 0)$ is inflation-neutral, so the deviation reduces to the absolute value of the unilateral inflation stance, $|\pi_c(\zeta^*(\{\chi_c\}))|$.

Figure 5 confirms this relationship by plotting each country's policy-alignment loss under the centralized optimal monetary policy against its absolute unilateral inflation

stance. A quadratic fit with an intercept but no linear term yields an R-square of 0.98.²⁵ Quantitatively, a one-percentage-point inflation stance deviation predicts a policy-alignment loss of about 0.29 percentage points of steady-state consumption; doubling this stance deviation approximately quadruples the predicted loss. The convexity of this relationship implies that the welfare costs of policy misalignment fall disproportionately on countries at the tails of the stance distribution—whether hawkish or dovish—while countries near the union-wide consensus bear negligible costs.




Figure 5: Inflation stance deviation and policy-alignment loss. *Note:* Each dot represents a euro-area country, with size proportional to its income share. The figure relates the absolute unilateral inflation stance to the corresponding policy-alignment loss, measured in percent of steady-state consumption. The red curve shows a quadratic fit with an intercept but no linear term.

Having established that inflation-stance deviations provide a sufficient statistic for policy-alignment losses, I now turn to the cross-country variation in the stances themselves. Figure 6 decomposes unilateral inflation stances under the centralized optimal price-index regime in two complementary ways. Panel (a) plots each country's inflation stance against its direct-incidence component, decomposed analogously to equation (10); most observations lie close to the 45-degree line, indicating that the direct-incidence channel accounts for the bulk of cross-country variation in stances. Panel (b) drills further into the direct-incidence channel, plotting the inflation stance against its IO-multiplier component alone, decomposed analogously to equation (15). With Luxembourg (omitted from the panel as a far outlier) and Malta as exceptions, observations again cluster tightly around the 45-degree line, indicating

²⁵Figure A.4 in the Online Appendix confirms robustness under CPI stabilization.

that the IO-multiplier channel alone reproduces most of the cross-country variation.²⁶ Since the input–output multiplier in equation (14) is the building block of this channel, these results identify production-network structure as a central determinant of inflation stances.




Figure 6: Decomposition of inflation stance. *Note:* Panels (a) and (b) plot the unilateral inflation stance against its direct-incidence component and its IO-multiplier channel, respectively, under the centralized optimal price-index regime. The direct-incidence component corresponds to the first term in the allocative-efficiency decomposition (equation (10)), while the IO-multiplier channel corresponds to the first term in the direct-incidence decomposition (equation (15)). The red 45-degree line denotes equality.

5.5 Inflation Stance and Production Networks

How does a country’s inflation stance relate to its position in the union’s production network? To answer this question, I first construct a country-level measure of upstreamness. Following Antràs et al. (2012), I define industry-level upstreamness as the average distance of an industry’s output from final consumption, then aggregate to the country level using value-added shares as weights. The resulting measure captures how far a country’s production lies from final demand within the broader union supply chain.

Figure 7 displays a clear negative relationship between country-level upstreamness and unilateral inflation stances: more upstream countries tend to prefer lower inflation, while more downstream countries tend to prefer higher inflation.²⁷ This gradient mirrors the familiar North–South divide in the euro area, with core Northern economies occupying

²⁶Figure A.2 in the Online Appendix confirms robustness under CPI regime.

²⁷Figures A.5 and A.6 confirm robustness under CPI stabilization and alternative price-rigidity estimates.

more upstream positions and peripheral Southern economies located further downstream. In particular, Luxembourg (LUX), the Netherlands (NLD), and Ireland (IRL) combine highly upstream production structures with a stronger preference for price stabilization, while Greece (GRC), Cyprus (CYP), and Spain (ESP) are more downstream and favor more accommodative policy. Germany (DEU), a major upstream producer, also leans toward price stabilization.

This relationship is confirmed by a bivariate weighted least-squares (WLS) regression, using countries' income shares as weights and including no additional controls. The fitted red line in Figure 7 has an estimated slope of -32.7 (standard error 4.27), statistically significant at the 0.1% level, with an R-square of 0.765. Quantitatively, the estimate implies that a one-unit increase in upstreamness is associated with an average 3.1 percentage point reduction in a country's preferred inflation rate. Table A.2 in the Online Appendix reports the country-level characteristics and shows that this relationship is not driven by any single country: leave-one-out regressions yield slopes ranging from -34.4 to -30.3 , all statistically significant at the 0.1% level. The slope also remains negative and significant after jointly excluding Luxembourg, Ireland, and the Netherlands, three hub economies whose network positions may be affected by financial activity, multinational production, or re-exports.




Figure 7: Inflation stance and upstreamness. *Note:* Each dot represents a euro-area country, with size proportional to its income share. The figure plots unilateral inflation stances, evaluated under the centralized optimal price-index regime, against upstreamness in the union production network. The red line shows the fitted linear relationship.

Motivated by Figure 6's finding that the IO-multiplier channel alone reproduces most of the cross-country variation in inflation stances, Figure 8 unpacks the underlying mechanism

by introducing the local input–output multiplier as the mediating network statistic. Panel (a) shows that country-level upstreamness is strongly associated with the local input–output multiplier, with an R-square of 0.77: countries whose domestic sectors supply intermediates along longer production chains have their Domar weights amplified by the Leontief inverse, raising their local multipliers above the union-wide benchmark (dashed line).²⁸ Panel (b) shows that a larger local input–output multiplier is in turn associated with a lower unilateral inflation stance, with an R-square of 0.75. Together, the two panels trace out the channel through which production-network position shapes monetary-policy preferences: upstream countries tend to have above-average input–output multipliers, so that even a uniform markup compression under monetary expansion deteriorates their terms of trade, inducing a preference for tighter policy. Downstream countries, by contrast, tend to have below-average input–output multipliers and experience a terms-of-trade improvement under monetary expansion, giving them less reason to resist accommodative policies.




Figure 8: The input–output multiplier channel. *Note:* Panel (a) plots upstreamness against the input–output multiplier. Panel (b) plots the multiplier against the unilateral inflation stance under the centralized optimal price-index regime. Each dot represents a euro-area country, with size proportional to its income share. In each panel, the red line shows the fitted linear relationship.

Combined with the quadratic relationship in Figure 5, these results imply that countries at either end of the production chain—those with highly upstream or downstream structures—tend to exhibit more extreme inflation stances and suffer larger policy-alignment

²⁸The empirical association in Panel (a) has a clean theoretical counterpart. In the roundabout economy of Example 4, upstreamness and the input–output multiplier coincide at the union level, both equal to $\frac{1}{\sum_c \bar{\alpha}_c}$, the inverse of the aggregate labor share.

losses. Production network structure thus plays a systematic role in shaping both inflation preferences and the incidence of policy misalignment under a common monetary policy.

6 Robustness and Extensions

Online Appendices [E](#) and [F](#) examine two extensions of the baseline framework. The first considers a standard nested-CES economy and shows that elasticities of substitution affect first-order outcomes only through the income-share Jacobian, leaving the direct-incidence index unchanged. In a calibrated production-network economy, the direct-incidence channel continues to account for most of the variation in inflation stances across a wide range of elasticities, reinforcing the input–output multiplier as the key object organizing cross-country variation in stances. The second extends the analysis to a global economy with dominant-currency pricing, where the relevant benchmark is set by the dominant-currency country rather than by a union-wide planner. Policy-alignment losses therefore increase quadratically with the distance between a country’s inflation stance and that of the dominant-currency country, and economies whose production networks more closely resemble that of the dominant economy incur smaller welfare losses.

7 Conclusion

This paper studies the origins and welfare implications of divergent inflation preferences within a currency union using a heterogeneous-agent input–output framework. Monetary policy induces first-order changes in country-level allocative efficiency through shifts in overall terms of trade. When these allocative effects are not neutralized in welfare aggregation, the monetary authority faces a trade-off between aggregate stabilization and first-order redistribution. This trade-off generates an inflation bias under the unilateral optimal policy, and therefore a third-best unilateral inflation stance under the institutional constraint of the union’s common price-index regime.

Applying the model to the euro area, I find that this unilateral inflation stance is fundamentally shaped by production structure: countries with more upstream production networks prefer lower union-wide inflation. Moreover, a country’s inflation stance deviation strongly predicts its policy-alignment loss, indicating that countries at either end of the production chain bear the largest welfare costs under a common monetary policy.

This framework provides a tractable approach for analyzing the structural foundations of policy divergence and its distributional consequences in integrated monetary unions. It can

be extended to explore other dimensions of macroeconomic coordination, including fiscal transfers, institutional design, and political representation in heterogeneous economies.

Acknowledgment. This paper received no external funding. The author declares no financial relationships or other potential conflicts of interest.

References

- Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012). The network origins of aggregate fluctuations. *Econometrica* 80(5), 1977–2016.
- Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2013). The network origins of large economic downturns. NBER Working Papers 19230, National Bureau of Economic Research, Inc.
- Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2017). Microeconomic origins of macroeconomic tail risks. *American Economic Review* 107(1), 54–108.
- Acemoglu, D. and A. Tahbaz-Salehi (2025). The Macroeconomics of Supply Chain Disruptions. *Review of Economic Studies* 92(2), 656–695.
- Acharya, S., E. Challe, and K. Dogra (2023). Optimal monetary policy according to HANK. *American Economic Review* 113(7), 1741–1782.
- Afrouzi, H. and S. Bhattarai (2023). Inflation and GDP dynamics in production networks: A sufficient statistics approach. NBER Working Papers 31218, National Bureau of Economic Research, Inc.
- Antràs, P., D. Chor, T. Fally, and R. Hillberry (2012). Measuring the upstreamness of production and trade flows. *American Economic Review* 102(3), 412–416.
- Atalay, E. (2017). How important are sectoral shocks? *American Economic Journal: Macroeconomics* 9(4), 254–280.
- Baqaei, D. and A. Burstein (2025). Aggregate efficiency with heterogeneous agents. Technical report, National Bureau of Economic Research.
- Baqaei, D. and E. Rubbo (2023). Micro Propagation and Macro Aggregation. *Annual Review of Economics* 15, 91–123.
- Baqaei, D. R. (2018). Cascading failures in production networks. *Econometrica* 86(5), 1819–1838.
- Baqaei, D. R. and E. Farhi (2019). The macroeconomic impact of microeconomic shocks: Beyond Hulten’s theorem. *Econometrica* 87(4), 1155–1203.

- Baqae, D. R. and E. Farhi (2020). Productivity and misallocation in general equilibrium. *The Quarterly Journal of Economics* 135(1), 105–163.
- Baqae, D. R. and E. Farhi (2024). Networks, Barriers, and Trade. *Econometrica* 92(2), 505–541.
- Baqae, D. R., E. Farhi, and K. Sangani (2024). The Supply-Side Effects of Monetary Policy. *Journal of Political Economy* 132(4), 1065–1112.
- Barattieri, A., S. Basu, and P. Gottschalk (2014). Some evidence on the importance of sticky wages. *American Economic Journal: Macroeconomics* 6(1), 70–101.
- Basu, S. (1995). Intermediate goods and business cycles: implications for productivity and welfare. *American Economic Review* 85(3), 512–31.
- Benigno, P. (2004). Optimal monetary policy in a currency area. *Journal of International Economics* 63(2), 293–320.
- Beraja, M., E. Hurst, and J. Ospina (2019). The aggregate implications of regional business cycles. *Econometrica* 87(6), 1789–1833.
- Bergin, P. R. and G. Corsetti (2023). The macroeconomic stabilization of tariff shocks: What is the optimal monetary response? *Journal of International Economics* 143, 103758.
- Bhandari, A., D. Evans, M. Golosov, and T. J. Sargent (2021). Inequality, Business Cycles, and Monetary-Fiscal Policy. *Econometrica* 89(6), 2559–2599.
- Bianchi, J. and L. Coulibaly (2025). The optimal monetary policy response to tariffs. Technical report, National Bureau of Economic Research.
- Bigio, S. and J. La’O (2020). Distortions in production networks. *Quarterly Journal of Economics* 135, 2187–2253.
- Carvalho, V. M. (2014). From micro to macro via production networks. *Journal of Economic Perspectives* 28(4), 23–48.
- Carvalho, V. M., M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi (2021). Supply chain disruptions: Evidence from the Great East Japan earthquake. *The Quarterly Journal of Economics* 136(2), 1255–1321.
- Carvalho, V. M. and A. Tahbaz-Salehi (2019). Production networks: A primer. *Annual Review of Economics* 11(1), 635–663.
- Corsetti, G. and P. Pesenti (2001). Welfare and macroeconomic interdependence. *The Quarterly Journal of Economics* 116(2), 421–445.
- Corsetti, G. and P. Pesenti (2005). International dimensions of optimal monetary policy. *Journal of Monetary Economics* 52(2), 281–305.

- 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 C. Engel (2003). Monetary policy in the open economy revisited: Price setting and exchange-rate flexibility. *The Review of Economic Studies* 70(4), 765–783.
- Dew-Becker, I. (2023). Tail risk in production networks. *Econometrica* 91(6), 2089–2123.
- Dávila, E. and A. Schaab (2023). Optimal monetary policy with heterogeneous agents: Discretion, commitment, and timeless policy. Technical report, National Bureau of Economic Research.
- Farhi, E. and I. Werning (2017). Fiscal unions. *American Economic Review* 107(12), 3788–3834.
- Ferrari, A. and M. Ghassibe (2024). A Disaggregated Economy with Optimal Pricing Decisions. Technical report, Working Paper.
- Ferrero, A. (2009). Fiscal and monetary rules for a currency union. *Journal of International Economics* 77(1), 1–10.
- Foerster, A. T., P.-D. G. Sarte, and M. W. Watson (2011). Sectoral versus aggregate shocks: A structural factor analysis of industrial production. *Journal of Political Economy* 119(1), 1–38.
- Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica* 79(3), 733–772.
- Galí, J. (2008). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework*. Princeton University Press.
- Gautier, E., C. Conflitti, R. P. Faber, B. Fabo, L. Fadejeva, V. Jouvanceau, J.-O. Menz, T. Messner, P. Petroulas, P. Roldan-Blanco, F. Rumler, S. Santoro, E. Wieland, and H. Zimmer (2024, October). New facts on consumer price rigidity in the euro area. *American Economic Journal: Macroeconomics* 16(4), 386–431.
- Ghassibe, M. (2021). Monetary policy and production networks: An empirical investigation. *Journal of Monetary Economics* 119(C), 21–39.
- Iversen, T., D. Soskice, and D. Hope (2016). The Eurozone and political economic institutions. *Annual Review of Political Science* 19(1), 163–185.
- Jones, C. I. (2013). Misallocation, Economic Growth, and Input–Output Economics. In D. Acemoglu, M. Arellano, and E. Dekel (Eds.), *Advances in Economics and Econometrics: Tenth World Congress*, Econometric Society Monographs, pp. 419–456. Cambridge University Press.
- King, R. G., C. I. Plosser, and S. T. Rebelo (1988). Production, growth and business cycles: I. The basic neoclassical model. *Journal of Monetary Economics* 21(2-3), 195–232.

- La'O, J. and W. A. Morrison (2024). Optimal monetary policy with redistribution. Technical report, National Bureau of Economic Research.
- La'O, J. and A. Tahbaz-Salehi (2022). Optimal monetary policy in production networks. *Econometrica* 90(3), 1295–1336.
- La'O, J. and A. Tahbaz-Salehi (2025). Missing Tax Instruments: Attaining Production Efficiency in Disaggregated, Production Network Economies with Nominal Rigidities.
- Le Grand, F., A. Martin-Baillon, and X. Ragot (2022). Should monetary policy care about redistribution? Optimal monetary and fiscal policy with heterogeneous agents.
- Liu, E. (2019). Industrial policies in production networks. *The Quarterly Journal of Economics* 134(4), 1883–1948.
- Long, J. B. and C. I. Plosser (1983). Real Business Cycles. *Journal of Political Economy* 91(1), 39–69.
- Mankiw, N. G. and R. Reis (2002). Sticky information versus sticky prices: a proposal to replace the new keynesian phillips curve. *The Quarterly Journal of Economics* 117(4), 1295–1328.
- Nakamura, E. and J. Steinsson (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics* 123(4), 1415–1464.
- Negishi, T. (1960). Welfare economics and existence of an equilibrium for a competitive economy. *Metroeconomica* 12(2-3), 92–97.
- Nuño, G. and C. Thomas (2022). Optimal redistributive inflation. *Annals of Economics and Statistics* (146), 3–64.
- Pastén, E., R. Schoenle, and M. Weber (2020). The propagation of monetary policy shocks in a heterogeneous production economy. *Journal of Monetary Economics* 116(C), 1–22.
- Pastén, E., R. Schoenle, and M. Weber (2024). Sectoral Heterogeneity in Nominal Price Rigidity and the Origin of Aggregate Fluctuations. *American Economic Journal: Macroeconomics* 16(2), 318–52.
- Qiu, Z., Y. Wang, L. Xu, and F. Zanetti (2026). Monetary policy in open economies with production networks. *Journal of Monetary Economics* 159, 103918.
- Rubbo, E. (2023). Networks, phillips curves, and monetary policy. *Econometrica* 91(4), 1417–1455.
- Rubbo, E. (2025). Monetary Non-Neutrality in the Cross-Section.
- Schaab, A. and S. Y. Tan (2023). Monetary and fiscal policy according to hank-io. Technical report.
- Taschereau-Dumouchel, M. (2026). Cascades and Fluctuations in an Economy with an Endogenous Production Network. *The Review of Economic Studies* 93(2), 1354–1392.

- Tille, C. (2001). The role of consumption substitutability in the international transmission of monetary shocks. *Journal of International Economics* 53(2), 421–444.
- Timmer, M. P., E. Dietzenbacher, B. Los, R. Stehrer, and G. J. Vries (2015). An illustrated user guide to the world input–output database: The case of global automotive production. *Review of International Economics* 23(3), 575–605.
- Viner, J. (1937). *Studies in the theory of international trade*. Routledge.
- Wei, S.-J. and Y. Xie (2020). Monetary policy in an era of global supply chains. *Journal of International Economics* 124, 103299.
- Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.
- Xu, Z. and C. Yu (2025). Optimal monetary policy in production networks with distortions. *Journal of Economic Theory* 225, 105979.

Online Appendix to “*The Network Origins of Inflation Stances in a Currency Union*”

Zhihao Xu¹

A Additional Figures and Tables

Table A.1: Expected welfare losses by countries.

	Uni. Opt.	Cent. Opt.	Cent. Uni. π	CPI Stab.	CPI Uni. π
AUT	0.412	0.417	0.413	0.434	0.430
BEL	0.435	0.437	0.436	0.454	0.453
CYP	0.499	0.858	0.509	0.875	0.525
DEU	0.671	0.708	0.674	0.723	0.689
ESP	0.333	0.369	0.336	0.384	0.351
EST	3.726	3.796	3.751	3.803	3.758
FIN	0.606	0.615	0.607	0.632	0.624
FRA	0.282	0.330	0.284	0.350	0.305
GRC	2.185	2.459	2.225	2.488	2.255
HRV	1.127	1.220	1.146	1.228	1.153
IRL	2.644	3.266	2.704	3.259	2.697
ITA	0.320	0.329	0.323	0.350	0.344
LTU	3.501	3.599	3.533	3.620	3.554
LUX	-0.315	0.946	-0.307	0.964	-0.290
LVA	2.754	2.886	2.775	2.890	2.779
MLT	2.638	3.049	2.656	3.048	2.654
NLD	-0.167	0.332	-0.164	0.343	-0.153
PRT	0.111	0.234	0.112	0.253	0.131
SVK	2.277	2.337	2.317	2.374	2.354
SVN	0.528	0.549	0.537	0.557	0.545

¹Department of Economics, Emory University. Address: 1602 Fishburne Drive, Atlanta, GA 30322, USA.
E-mail: owenzxu@gmail.com.

Table A.2: Leave-one-out robustness: inflation stances and upstreamness

Dropped country	Characteristics of dropped country				WLS regression	
	Inflation stance (%)	Upstreamness	IO multiplier	Income share (%)	Slope	R ²
None	–	–	–	–	–32.72	0.765
AUT	0.13	2.238	2.980	3.26	–32.80	0.769
BEL	–0.05	2.461	3.271	3.97	–32.49	0.809
CYP	1.21	2.275	2.512	0.20	–33.12	0.776
DEU	–0.35	2.304	3.079	27.05	–34.43	0.754
ESP	0.38	2.020	2.873	10.87	–31.46	0.755
EST	0.42	2.566	3.086	0.21	–32.87	0.777
FIN	0.18	2.310	3.052	2.05	–32.91	0.777
FRA	0.44	2.039	2.785	22.33	–30.80	0.712
GRC	0.99	1.994	2.509	1.97	–33.49	0.765
HRV	0.55	2.277	2.804	0.43	–32.92	0.772
IRL	–1.38	2.643	3.655	1.55	–32.97	0.746
ITA	0.15	2.116	2.987	16.16	–32.22	0.768
LTU	0.50	2.424	2.807	0.37	–32.96	0.778
LUX	–2.04	3.027	5.544	0.38	–32.28	0.751
LVA	0.67	2.573	3.075	0.25	–33.06	0.785
MLT	1.34	2.383	3.328	0.08	–32.97	0.773
NLD	–1.25	2.677	3.531	5.91	–30.30	0.645
PRT	0.73	2.110	2.772	1.86	–33.39	0.773
SVK	0.29	2.481	3.250	0.76	–32.99	0.788
SVN	0.22	2.408	3.058	0.36	–32.79	0.771
LUX, IRL, NLD	–	–	–	–	–28.68	0.535

Note: Each row reports a weighted least-squares regression of country-level upstreamness on unilateral inflation stance after excluding the indicated country from the estimation sample. Columns 2–5 report the inflation stance, upstreamness, input–output multiplier, and income share of the excluded country. The row labeled “None” reports the full-sample specification. The slope remains negative across all leave-one-out specifications. The final row excludes Luxembourg, Ireland, and the Netherlands jointly.




Figure A.1: Country weights in the target price index under alternative monetary policies. *Note:* Panels (a) and (b) compare country income shares with their weights in the target price indices under the centralized optimal policy and CPI stabilization, respectively. The red 45-degree line indicates perfect coincidence.




Figure A.2: Decomposition of inflation stance under CPI regime. *Note:* Panels (a) and (b) plot the unilateral inflation stance against its direct-incidence component and its IO-multiplier channel, respectively, under the CPI stabilization regime. The direct-incidence component corresponds to the first term in the allocative-efficiency decomposition (equation (10)), while the IO-multiplier channel corresponds to the first term in the direct-incidence decomposition (equation (15)). The red 45-degree line denotes equality.




Figure A.3: Decomposition of policy-alignment loss under CPI stabilization. *Note:* Panel (a) plots total PAL against the component attributable to underutilized allocative efficiency. Panel (b) plots total PAL against the component attributable to inflation misalignment, as in equation (25). In each panel, the 45-degree line indicates that the component accounts for the full policy-alignment loss.




Figure A.4: Robustness to CPI stabilization. *Note:* Each dot represents a euro-area country. Dot size reflects the country's income share. The horizontal axis measures the absolute value of the country's unilateral inflation stance. The vertical axis shows the corresponding policy-alignment loss as a percentage of the country's steady-state consumption. The red curve shows a quadratic fit with an intercept but no linear term.




Figure A.5: Robustness to CPI stabilization. *Note:* The horizontal axis reports each country's unilateral inflation stance, computed using the CPI as the common price index rather than the centralized optimal price index used in the baseline. The vertical axis reports country-level upstreamness in the union production network. The red line depicts the fitted linear trend.




Figure A.6: Robustness to euro-area price rigidity. *Note:* The horizontal axis reports each country's unilateral inflation stance, computed under the centralized optimal price-index regime. The vertical axis reports country-level upstreamness in the union production network. Unlike the baseline calibration, which uses producer-price FPA estimates from [Pastén et al. \(2020\)](#), this figure calibrates sectoral price flexibilities using the euro-area consumer-price FPA estimates from [Gautier et al. \(2024\)](#). The red line depicts the fitted linear trend.

B Related Proofs

Throughout the appendix, I use a bar to denote the steady-state (initial) value of a variable, and a superscript $*$ to denote its value in the flexible-price equilibrium. For example, $\bar{L}_c$ and L_c^* denote household c 's labor supply in the steady state and under flexible prices, respectively. To simplify notation, I normalize the initial nominal aggregate demand to one ($\bar{m} = 1$), so that the monetary shock can be represented directly by $\log m$. Accordingly, I define the vector of exogenous shocks as $\xi \equiv (\log z', \log m)'$, which includes both sectoral productivity and monetary shocks.

I begin proofs by defining sectoral wedges following [La'O and Tahbaz-Salehi \(2022\)](#), a formulation that is convenient for analyzing within-sector distortions. Importantly, to a first-order approximation, the sectoral wedge is equivalent (up to a negative sign) to the sectoral markup introduced in equation (4). This equivalence is established formally in Remark 3.

Cost minimization implies that the demand of firm k in industry i for the good produced by industry j can be expressed as

$$x_{ij,k} = \omega_{ij} y_{ik} m c_i / p_j.$$

Firm k 's output in industry i is

$$y_{ik} = y_i (p_{ik} / p_i)^{-\theta_i}$$

where the sectoral price index is defined as

$$p_i = \left(\int_0^1 p_{ik}^{1-\theta_i} dk \right)^{1/(1-\theta_i)}.$$

Aggregating over the unit mass of firms in industry i , the total demand for goods from industry j equals

$$\int_0^1 x_{ij,k} dk = \omega_{ij} p_i y_i \varepsilon_i / p_j$$

where the sectoral wedge ε_i is defined as

$$\varepsilon_i = \frac{m c_i}{p_i} \int_0^1 (p_{ik} / p_i)^{-\theta_i} dk. \quad (\text{B-1})$$

In the flexible-price equilibrium, firms in industry i set identical prices and charge no markups. As a result, equation (B-1) implies $\varepsilon_i^* = 1$ for all i .

Let $\varepsilon_{ik} = mc_i/p_{ik}$ denote the firm-level wedge. Using the representation above, the sectoral wedge aggregates firm-level wedges as

$$\log \varepsilon_i = \log \int_0^1 \varepsilon_{ik}^{\theta_i} dk - \log \int_0^1 \varepsilon_{ik}^{\theta_i-1} dk$$

Expanding around the efficient allocation, I obtain

$$\varepsilon_{ik}^{\theta_i} = 1 + \theta_i \log \varepsilon_{ik} + o(\|\xi\|), \quad \text{and} \quad \varepsilon_{ik}^{\theta_i-1} = 1 + (\theta_i - 1) \log \varepsilon_{ik} + o(\|\xi\|),$$

which yields

$$\begin{aligned} \log \varepsilon_i &= \log \left(1 + \theta_i \int_0^1 \log \varepsilon_{ik} dk \right) - \log \left(1 + (\theta_i - 1) \int_0^1 \log \varepsilon_{ik} dk \right) + o(\|\xi\|) \\ &= \int_0^1 \log \varepsilon_{ik} dk + o(\|\xi\|). \end{aligned}$$

Thus, to a first-order approximation, the sectoral wedge equals the cross-sectional average of firm-level wedges.

Moreover, the sectoral wedge is directly related to the sectoral markup. Up to a first-order approximation,

$$\log \varepsilon_i = \log mc_i - \log p_i + o(\|\xi\|) = -\log \mu_i + o(\|\xi\|)$$

Remark 3. Throughout the appendix, sectoral distortions are expressed using the sectoral wedge $\log \varepsilon_i$, defined in (B-1). In the main text, I instead work with the equivalent representation in terms of sectoral markups $\log \mu_i$. Since $\log \varepsilon_i = -\log \mu_i + o(\|\xi\|)$, the two notations are interchangeable to a first-order approximation.

Proof of Lemma 1. I first show (8) of Lemma 1. I begin by aggregating profits across the unit mass of firms in industry i . Using the firm-level profit expression (3), the total profit in

industry i is

$$\begin{aligned}
\int_0^1 \Pi_{ik} dk &= \int_0^1 [p_{ik} - (1 - \tau_i)mc_i] y_{ik} dk \\
&= \int_0^1 p_{ik} y_{ik} dk - (1 - \tau_i)mc_i \int_0^1 y_{ik} dk \\
&= p_i y_i - (1 - \tau_i)mc_i \int_0^1 y_{ik} dk.
\end{aligned}$$

Substituting this expression, together with the tax scheme (6), into the household budget constraint (5), yields

$$P_c C_c = w_c L_c + \sum_{i=1}^N \Phi_{ic} \left(p_i y_i - mc_i \int_0^1 y_{ik} dk \right) = w_c L_c + \sum_{i=1}^N \Phi_{ic} p_i y_i (1 - \varepsilon_i),$$

where the second equality uses the definition of the sectoral wedge (B-1).

Applying the input–output definitions yields a convenient representation of the labor wedge for household c :

$$\frac{\Lambda_c}{\chi_c} = \frac{w_c L_c}{P_c C_c} = 1 - \sum_{i=1}^N \frac{\lambda_i \Phi_{ic}}{\chi_c} (1 - \varepsilon_i). \quad (\text{B-2})$$

This expression reveals that deviations of labor income from total consumption expenditure are driven exclusively by sectoral wedges, $1 - \varepsilon_i$, scaled by sectoral Domar weights, λ_i , and household-specific ownership exposures, Φ_{ic} . Thus, the labor wedge captures the transmission of production-network distortions into non-labor income for household c . In the undistorted flexible-price equilibrium ($\varepsilon_i^* = 1$ for all i), the wedge vanishes, implying $\Lambda_c^* = \chi_c^*$.

A first-order approximation of the labor wedge implies

$$\log(\Lambda_c / \chi_c) = \log \Lambda_c - \log \chi_c = \sum_{i=1}^N \frac{\lambda_i \Phi_{ic}}{\chi_c} \log \varepsilon_i + o(\|\xi\|). \quad (\text{B-3})$$

Market clearing for good i implies

$$\begin{aligned} p_i y_i &= p_i \sum_{c=1}^C c_{ci} + p_i \sum_{j=1}^N \int_0^1 x_{jik} dk \\ &= \sum_{c=1}^C p_i c_{ci} + \sum_{j=1}^N \omega_{ji} p_j y_j \varepsilon_j \end{aligned}$$

Using the input-output definitions, this is equivalent to

$$\lambda_i = \sum_{c=1}^C \chi_c \beta_{ci} + \sum_{j=1}^N \omega_{ji} \lambda_j \varepsilon_j$$

which in vector form gives

$$\lambda = (I - \Omega' \text{diag}(\varepsilon))^{-1} \beta \chi \quad (\text{B-4})$$

where $\beta = [\beta_1, \dots, \beta_C] \in \mathbb{R}^{N \times C}$.

Thus, to a first order, we have

$$\begin{aligned} \lambda - \lambda^* &= (I - \Omega' \text{diag}(\varepsilon))^{-1} \beta \chi - (I - \Omega')^{-1} \beta \chi^* \\ &= (I - \Omega' \text{diag}(\varepsilon))^{-1} \Omega' (\text{diag}(\varepsilon) - I) (I - \Omega')^{-1} \beta \chi + (I - \Omega')^{-1} \beta (\chi - \chi^*) \\ &= \Psi' \Omega' \text{diag}(\log \varepsilon) \Psi' \beta \chi + \Psi' \beta (\chi - \chi^*) + o(\|\xi\|) \\ &= (\Psi' - I) \text{diag}(\log \varepsilon) \lambda^* + \lambda^* (\chi - \chi^*) + o(\|\xi\|) \end{aligned}$$

or componentwise,

$$\lambda_i = (1 - \log \varepsilon_i) \lambda_i^* + \sum_{j=1}^N \Psi_{ji} \lambda_j^* \log \varepsilon_j + \sum_{c=1}^C \lambda_i^{c*} (\chi_c - \chi_c^*) + o(\|\xi\|). \quad (\text{B-5})$$

By analogy,

$$\Lambda_f = \Lambda_f^* + \sum_{j=1}^N \Psi_{jf} \lambda_j^* \log \varepsilon_j + \sum_{c=1}^C \Lambda_f^{c*} (\chi_c - \chi_c^*) + o(\|\xi\|) \quad (\text{B-6})$$

Combining (B-3) and (B-6) yields the first-order approximation for household income-

share changes in (8), as stated in Lemma 1,

$$\hat{\chi}_c = \log \chi_c - \log \chi_c^* = - \underbrace{\sum_{j=1}^N \sum_{f=1}^C (\chi_c^*)^{-1} Q_{cf} \lambda_j^* (\Phi_{jf} - \Psi_{jf})}_{\Gamma_{jc}} \log \varepsilon_j + o(\|\xi\|). \quad (\text{B-7})$$

The matrix Q denotes the inverse of $I - \Lambda'$ restricted to the subspace $\mathcal{S} = \{x \in \mathbb{R}^C : \mathbf{1}'x = 0\}$,¹ and can equivalently be written as

$$Q = Z(Z'(I - \Lambda')Z)^{-1}Z',$$

where $Z \in \mathbb{R}^{C \times (C-1)}$ is any orthonormal basis for $\mathcal{S}$.

I now turn to the first-order approximation (7) for the household employment gap in Lemma 1.

The consumption-leisure optimality condition for household c implies

$$\psi_c L_c^{1/\eta_c} = C_c^{-1} w_c / P_c$$

where ψ_c is a household-specific preference parameter and $P_c = \prod_{i=1}^N p_i^{\beta_{ci}}$ is household c 's consumer price index (CPI).

Combining this condition with the definition of the labor wedge in (B-2), I can express household consumption and labor supply directly in terms of the household's real wage and labor wedge

$$C_c = (w_c / P_c) (\Lambda_c / \chi_c)^{-\frac{1}{1+\eta_c}} \psi_c^{-\frac{\eta_c}{1+\eta_c}}, \quad (\text{B-8})$$

$$L_c = (\Lambda_c / \chi_c)^{\frac{\eta_c}{1+\eta_c}} \psi_c^{-\frac{\eta_c}{1+\eta_c}}. \quad (\text{B-9})$$

A direct implication of (B-9) is that natural (flexible-price) employment, $L_c^* = \psi_c^{-\frac{\eta_c}{1+\eta_c}}$, is invariant to both productivity and monetary shocks. Log-linearizing (B-9) yields the

¹Income shares satisfy $\sum_{c=1}^C \chi_c = 1$ in every equilibrium. Hence their level deviations must sum to zero, which means that perturbations of χ lie in the $(C-1)$ -dimensional subspace orthogonal to $\mathbf{1}$.

first-order approximation for employment gap (7) in Lemma 1:

$$\begin{aligned}
\hat{l}_c &= \log L_c - \log L_c^* = \frac{\eta_c}{1 + \eta_c} \log(\Lambda_c / \chi_c) \\
&= \frac{\eta_c}{1 + \eta_c} \sum_{i=1}^N \frac{\lambda_i \Phi_{ic}}{\chi_c} \log \varepsilon_i + o(\|\xi\|) \\
&= \sum_{i=1}^N \underbrace{-\frac{\eta_c}{1 + \eta_c} \frac{\lambda_i \Phi_{ic}}{\chi_c}}_{\ell_{ic}^\mu} \log \mu_i + o(\|\xi\|),
\end{aligned}$$

where the second line substitutes the first-order expression for the labor wedge in (B-3). $\square$

Lemma 3 (Income shares as Negishi weights). The competitive equilibrium allocation can be replicated by a Negishi planner who maximizes a weighted sum of household utilities, where the Negishi weights coincide with households' income shares.

Proof of Lemma 3. The competitive equilibrium allocation solves a social planner's problem that maximizes a weighted sum of household utilities. Letting $\{\kappa_c\}$ denote undetermined welfare weights, the planner chooses allocations $\{c_{ci}, x_{ij}, L_{ic}, L_c\}$ to solve (omitting firm subscripts as firms are identical within each industry in the absence of nominal rigidities)

$$\begin{aligned}
\max \sum_{c=1}^C \kappa_c & \left(\sum_{i=1}^N \beta_{ci} \log c_{ci} - \psi_c \frac{L_c^{1+1/\eta_c}}{1 + 1/\eta_c} \right) \\
& + \sum_{i=1}^N \tilde{p}_i \left(z_i \zeta_i \prod_{j=1}^N x_{ij}^{\omega_{ij}} \prod_{c=1}^C L_{ic}^{\alpha_{ic}} - \sum_{j=1}^N x_{ji} - \sum_{c=1}^C c_{ci} \right) + \sum_{c=1}^C \tilde{w}_c \left(L_c - \sum_{i=1}^N L_{ic} \right),
\end{aligned}$$

where $\{\tilde{p}_i\}_{i=1}^N$ and $\{\tilde{w}_c\}_{c=1}^C$ are the Lagrange multipliers on the goods-market and labor constraints, respectively.

The planner's first-order conditions with respect to intermediate inputs x_{ij} and consumption c_{ci} imply

$$\begin{aligned}
\tilde{p}_i \omega_{ij} \frac{y_i}{x_{ij}} &= \tilde{p}_j, \\
\tilde{p}_i c_{ci} &= \kappa_c \beta_{ci}.
\end{aligned}$$

In a competitive equilibrium, optimality implies

$$p_i \omega_{ij} \frac{y_i}{x_{ij}} = p_j,$$

$$p_i c_{ci} = \beta_{ci} \chi_c m,$$

where $m = \sum_c \sum_i p_i c_{ci}$ is nominal aggregate expenditure.

Comparing planner and market allocations, it follows that $\tilde{p}_i = p_i/m$, $\tilde{w}_c = w_c/m$, and $\kappa_c = \chi_c$. Thus, the welfare weights that rationalize the planner's allocation coincide with the income shares realized in the competitive equilibrium.

To verify consistency on the labor margin, the planner's first-order conditions with respect to labor inputs L_{ci} and labor supply L_c imply

$$\tilde{p}_i \alpha_{ic} \frac{y_i}{L_{ic}} = \tilde{w}_c, \quad \tilde{w}_c = \kappa_c \psi_c L_c^{1/\eta_c}.$$

Substituting $\kappa_c = \chi_c$, $\tilde{w}_c = w_c/m$, and using the competitive-equilibrium condition $w_c L_c = \chi_c m$, we obtain

$$L_c = \psi_c^{-\frac{\eta_c}{1+\eta_c}},$$

which coincides with labor supply in the flexible-price equilibrium.

Hence, the competitive allocation is exactly rationalized by a Negishi planner whose welfare weights equal households' income shares. $\square$

Lemma 4 (Second-Order Approximations). Up to a second-order approximation in the sectoral wedges $\{\varepsilon_i\}_{i=1}^N$, the income share χ_c and the labor wedge Λ_c/χ_c satisfy

$$\log \chi_c - \log \chi_c^* = - \sum_{j=1}^N \Gamma_{jc} \log \varepsilon_j + \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \Xi_{ij}^c \log \varepsilon_j \log \varepsilon_i + o(\|\xi\|^2)$$

and

$$\log(\Lambda_c/\chi_c) = \sum_{i=1}^N \frac{\lambda_i \Phi_{ic}}{\chi_c} \log \varepsilon_i + \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \Upsilon_{ij}^c \log \varepsilon_j \log \varepsilon_i + o(\|\xi\|^2)$$

respectively. The associated second-order coefficient matrices are

$$\begin{aligned}\Xi_{ij}^c &= -\Gamma_{jc}\Gamma_{ic} + \chi_c^{-1} \sum_{f=1}^C Q_{cf}(\Phi_{if} - \Psi_{if}) \left(2 \sum_{g=1}^C \lambda_i^g \chi_g \Gamma_{jg} - 2\lambda_j \Psi_{ji} + \lambda_i \iota_{ij} \right), \\ \Upsilon_{ij}^c &= 2 \frac{\lambda_j \Phi_{ic}}{\chi_c} \Psi_{ji} - 2 \sum_{f=1}^C \frac{\lambda_i^f \Phi_{ic}}{\chi_c} \chi_f \Gamma_{jf} + 2 \frac{\lambda_i \Phi_{ic}}{\chi_c} \Gamma_{jc} - \frac{\lambda_i \Phi_{ic}}{\chi_c} \iota_{ij} - \frac{\lambda_i \Phi_{ic}}{\chi_c} \frac{\lambda_j \Phi_{jc}}{\chi_c}.\end{aligned}$$

Proof of Lemma 4. To establish Lemma 4, I derive log-quadratic approximations to the labor wedges in (B-2) and the sectoral Domar weights in (B-4), expanding both expressions around the economy's steady state. The overall structure parallels the first-order analysis but extends the derivations to second order in the sectoral wedges.

I begin with the labor-wedge expression (B-2). A second-order expansion yields

$$\begin{aligned}\Lambda_c - \chi_c &= - \sum_{i=1}^N \lambda_i \Phi_{ic} (1 - \varepsilon_i) \\ &= \sum_{i=1}^N \lambda_i \Phi_{ic} (\log \varepsilon_i + \frac{1}{2} \log^2 \varepsilon_i + o(\|\xi\|^2)) \\ &= \sum_{i=1}^N \lambda_i \Phi_{ic} \log \varepsilon_i + \frac{1}{2} \sum_{i=1}^N \lambda_i \Phi_{ic} \log^2 \varepsilon_i + o(\|\xi\|^2).\end{aligned}\tag{B-10}$$

I then substitute the first-order expression for Domar weights, (B-5), into the above expansion and use $\Lambda_c^* = \chi_c^*$ to obtain

$$\begin{aligned}\Lambda_c - \Lambda_c^* - (\chi_c - \chi_c^*) &= \sum_{i=1}^N \lambda_i^* \Phi_{ic} \log \varepsilon_i + \sum_{i=1}^N \sum_{j=1}^N \lambda_j^* \Phi_{ic} \Psi_{ji} \log \varepsilon_j \log \varepsilon_i \\ &\quad + \sum_{i=1}^N \sum_{f=1}^C \lambda_i^{f*} (\chi_f - \chi_f^*) \Phi_{ic} \log \varepsilon_i - \frac{1}{2} \sum_{i=1}^N \lambda_i^* \Phi_{ic} \log^2 \varepsilon_i + o(\|\xi\|^2).\end{aligned}\tag{B-11}$$

Next, I derive a log-quadratic approximation to sectoral Domar weights using (B-4).

$$\begin{aligned}
\lambda - \lambda^* &= (I - \Omega' \text{diag}(\varepsilon))^{-1} \beta \chi - (I - \Omega')^{-1} \beta \chi^* \\
&= (I - \Omega' \text{diag}(\varepsilon))^{-1} \Omega' (\text{diag}(\varepsilon) - I) (I - \Omega')^{-1} \beta \chi + (I - \Omega')^{-1} \beta (\chi - \chi^*) \\
&= \Psi' \Omega' \text{diag}(\log \varepsilon + \frac{1}{2} \log^2 \varepsilon) \Psi' \beta \chi^* + \Psi' \Omega' \text{diag}(\log \varepsilon) \Psi' \Omega' \text{diag}(\log \varepsilon) \Psi' \beta \chi^* \\
&\quad + \Psi' \Omega' \text{diag}(\log \varepsilon) \Psi' \beta (\chi - \chi^*) + \Psi' \beta (\chi - \chi^*) + o(\|\xi\|^2) \\
&= (\Psi' - I) \text{diag}(\log \varepsilon + \frac{1}{2} \log^2 \varepsilon) \lambda^* + [(\Psi' - I) \text{diag}(\log \varepsilon)]^2 \lambda^* \\
&\quad + (\Psi' - I) \text{diag}(\log \varepsilon) \lambda^* (\chi - \chi^*) + \lambda^* (\chi - \chi^*) + o(\|\xi\|^2)
\end{aligned}$$

where $\lambda^* \equiv \Psi' \beta$ denotes the $N \times C$ matrix of country-level Domar exposures, whose (i, c) -th element is λ_i^{c*} .

Equivalently, in component form,

$$\begin{aligned}
\lambda_i - \lambda_i^* &= \sum_{c=1}^C \lambda_i^{c*} (\chi_c - \chi_c^*) - \lambda_i^* \log \varepsilon_i + \sum_{j=1}^N \lambda_j^* \Psi_{ji} \log \varepsilon_j + \sum_{j=1}^N \sum_{c=1}^C \lambda_j^{c*} (\Psi_{ji} - \iota_{ji}) (\chi_c - \chi_c^*) \log \varepsilon_j \\
&\quad - \frac{1}{2} \lambda_i^* \log^2 \varepsilon_i + \frac{1}{2} \sum_{j=1}^N \lambda_j^* \Psi_{ji} \log^2 \varepsilon_j + \sum_{j=1}^N \sum_{k=1}^N \lambda_j^* (\Psi_{jk} - \iota_{jk}) (\Psi_{ki} - \iota_{ki}) \log \varepsilon_j \log \varepsilon_k + o(\|\xi\|^2).
\end{aligned}$$

By analogy, a second-order approximation of the labor income share Λ_f is given by

$$\begin{aligned}
\Lambda_f - \Lambda_f^* &= \sum_{c=1}^C \Lambda_f^{c*} (\chi_c - \chi_c^*) + \sum_{j=1}^N \lambda_j^* \Psi_{jf} \log \varepsilon_j + \sum_{j=1}^N \sum_{c=1}^C \lambda_j^{c*} \Psi_{jf} (\chi_c - \chi_c^*) \log \varepsilon_j \\
&\quad + \frac{1}{2} \sum_{j=1}^N \lambda_j^* \Psi_{jf} \log^2 \varepsilon_j + \sum_{j=1}^N \sum_{k=1}^N \lambda_j^* (\Psi_{jk} - \iota_{jk}) \Psi_{kf} \log \varepsilon_j \log \varepsilon_k + o(\|\xi\|^2) \\
&= \sum_{c=1}^C \Lambda_f^{c*} (\chi_c - \chi_c^*) + \sum_{j=1}^N \lambda_j^* \Psi_{jf} \log \varepsilon_j + \sum_{j=1}^N \sum_{c=1}^C \lambda_j^{c*} \Psi_{jf} (\chi_c - \chi_c^*) \log \varepsilon_j \\
&\quad - \frac{1}{2} \sum_{j=1}^N \lambda_j^* \Psi_{jf} \log^2 \varepsilon_j + \sum_{j=1}^N \sum_{k=1}^N \lambda_j^* \Psi_{jk} \Psi_{kf} \log \varepsilon_j \log \varepsilon_k + o(\|\xi\|^2)
\end{aligned}$$

Combining this expression with (B-11) and eliminating $\Lambda_f - \Lambda_f^*$ results in the following

linear system

$$\begin{aligned}
\chi_f - \chi_f^* &= \sum_{c=1}^C \Lambda_f^{c*} (\chi_c - \chi_c^*) - \sum_{j=1}^N \lambda_j^* (\Phi_{jf} - \Psi_{jf}) \log \varepsilon_j \\
&+ \frac{1}{2} \sum_{j=1}^N \lambda_j^* (\Phi_{jf} - \Psi_{jf}) \log^2 \varepsilon_j + \sum_{i=1}^N \sum_{j=1}^N \lambda_j^* \Psi_{ji} (\Psi_{if} - \Phi_{if}) \log \varepsilon_j \log \varepsilon_i \\
&+ \sum_{i=1}^N \sum_{c=1}^C \lambda_i^{c*} (\chi_c - \chi_c^*) (\Psi_{if} - \Phi_{if}) \log \varepsilon_i + o(\|\xi\|^2)
\end{aligned}$$

together with the adding-up constraint $\sum_{f=1}^C (\chi_f - \chi_f^*) = 0$.

Solving this system using the restricted inverse Q of $I - \Lambda'$ and collecting terms yields

$$\begin{aligned}
\log \chi_c - \log \chi_c^* &= - \sum_{j=1}^N \Gamma_{jc} \log \varepsilon_j - \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \Gamma_{jc} \Gamma_{ic} \log \varepsilon_j \log \varepsilon_i \\
&+ \frac{1}{2} (\chi_c^*)^{-1} \sum_{j=1}^N \sum_{f=1}^C Q_{cf} \lambda_j^* (\Phi_{jf} - \Psi_{jf}) \log^2 \varepsilon_j \\
&+ (\chi_c^*)^{-1} \sum_{i=1}^N \sum_{j=1}^N \sum_{f=1}^C \sum_{g=1}^C Q_{cf} \lambda_i^{g*} \chi_g^* \Gamma_{jg} (\Phi_{if} - \Psi_{if}) \log \varepsilon_j \log \varepsilon_i \\
&+ (\chi_c^*)^{-1} \sum_{i=1}^N \sum_{j=1}^N \sum_{f=1}^C Q_{cf} \lambda_j^* \Psi_{ji} (\Psi_{if} - \Phi_{if}) \log \varepsilon_j \log \varepsilon_i + o(\|\xi\|^2) \\
&= - \sum_{j=1}^N \Gamma_{jc} \log \varepsilon_j + \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \Xi_{ij}^c \log \varepsilon_j \log \varepsilon_i + o(\|\xi\|^2)
\end{aligned}$$

where

$$\begin{aligned}
\Xi_{ij}^c &= - \Gamma_{jc} \Gamma_{ic} + 2\chi_c^{-1} \sum_{f=1}^C \sum_{g=1}^C Q_{cf} \lambda_i^g \chi_g^* \Gamma_{jg} (\Phi_{if} - \Psi_{if}) \\
&+ \chi_c^{-1} \sum_{f=1}^C Q_{cf} \lambda_j (\Phi_{jf} - \Psi_{jf}) \iota_{ij} + 2\chi_c^{-1} \sum_{f=1}^C Q_{cf} \lambda_j \Psi_{ji} (\Psi_{if} - \Phi_{if}) \\
&= - \Gamma_{jc} \Gamma_{ic} + \chi_c^{-1} \sum_{f=1}^C Q_{cf} (\Phi_{if} - \Psi_{if}) \left(2 \sum_{g=1}^C \lambda_i^g \chi_g^* \Gamma_{jg} - 2\lambda_j \Psi_{ji} + \lambda_i \iota_{ij} \right),
\end{aligned}$$

This establishes the claimed log-quadratic approximation for the income share χ_c .

To establish a second-order approximation for the labor wedge Λ_c/χ_c , I start from (B-2), which implies

$$\begin{aligned}\Lambda_c/\chi_c &= 1 - \sum_{i=1}^N \frac{\lambda_i \Phi_{ic}}{\chi_c} (1 - \varepsilon_i) \\ &= 1 + \sum_{i=1}^N \frac{\lambda_i \Phi_{ic}}{\chi_c} (\log \varepsilon_i + \frac{1}{2} \log^2 \varepsilon_i + o(\|\xi\|^2)) \\ &= 1 + \sum_{i=1}^N \frac{\lambda_i \Phi_{ic}}{\chi_c} \log \varepsilon_i + \frac{1}{2} \sum_{i=1}^N \frac{\lambda_i \Phi_{ic}}{\chi_c} \log^2 \varepsilon_i + o(\|\xi\|^2)\end{aligned}$$

which then implies

$$\log(\Lambda_c/\chi_c) = \sum_{i=1}^N \frac{\lambda_i \Phi_{ic}}{\chi_c} \log \varepsilon_i + \frac{1}{2} \sum_{i=1}^N \frac{\lambda_i \Phi_{ic}}{\chi_c} \log^2 \varepsilon_i - \frac{1}{2} \left(\sum_{i=1}^N \frac{\lambda_i \Phi_{ic}}{\chi_c} \log \varepsilon_i \right)^2 + o(\|\xi\|^2). \quad (\text{B-12})$$

Next, substitute the first-order expression for Domar weights (B-5) and the first-order expansion for the inverse income share,

$$\frac{1}{\chi_c} = \frac{1}{\chi_c^*} + \sum_{i=1}^N \frac{\Gamma_{ic}}{\chi_c^*} \log \varepsilon_i + o(\|\xi\|)$$

into (B-12) and collect terms up to second order. This yields

$$\begin{aligned}\log(\Lambda_c/\chi_c) &= \sum_{i=1}^N \frac{\lambda_i^* \Phi_{ic}}{\chi_c^*} \log \varepsilon_i + \sum_{i=1}^N \sum_{j=1}^N \frac{\lambda_j^* \Phi_{ic}}{\chi_c^*} \Psi_{ji} \log \varepsilon_j \log \varepsilon_i - \sum_{i=1}^N \sum_{j=1}^N \sum_{f=1}^C \frac{\lambda_i^{f*} \Phi_{ic}}{\chi_c^*} \chi_f^* \Gamma_{jf} \log \varepsilon_j \log \varepsilon_i \\ &\quad + \sum_{i=1}^N \sum_{j=1}^N \frac{\lambda_i^* \Phi_{ic}}{\chi_c^*} \Gamma_{jc} \log \varepsilon_j \log \varepsilon_i - \frac{1}{2} \sum_{i=1}^N \frac{\lambda_i^* \Phi_{ic}}{\chi_c^*} \log^2 \varepsilon_i - \frac{1}{2} \left(\sum_{i=1}^N \frac{\lambda_i^* \Phi_{ic}}{\chi_c^*} \log \varepsilon_i \right)^2 + o(\|\xi\|^2) \\ &= \sum_{i=1}^N \frac{\lambda_i^* \Phi_{ic}}{\chi_c^*} \log \varepsilon_i + \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \Upsilon_{ij}^c \log \varepsilon_j \log \varepsilon_i + o(\|\xi\|^2)\end{aligned}$$

where the second-order coefficients are given by

$$\Upsilon_{ij}^c = 2 \frac{\lambda_j \Phi_{ic}}{\chi_c} \Psi_{ji} - 2 \sum_{f=1}^C \frac{\lambda_i^f \Phi_{ic}}{\chi_c} \chi_f \Gamma_{jf} + 2 \frac{\lambda_i \Phi_{ic}}{\chi_c} \Gamma_{jc} - \frac{\lambda_i \Phi_{ic}}{\chi_c} \iota_{ij} - \frac{\lambda_i \Phi_{ic}}{\chi_c} \frac{\lambda_j \Phi_{jc}}{\chi_c}.$$

This establishes the claimed log-quadratic approximation for the labor wedge Λ_c/χ_c . $\square$

Proof of Proposition 2. Cost minimization by firm k in industry i implies that

$$mc_i = \frac{1}{z_i \varsigma_i} \prod_{j=1}^N \left(\frac{p_j}{\omega_{ij}} \right)^{\omega_{ij}} \prod_{c=1}^C \left(\frac{w_c}{\alpha_{ic}} \right)^{\alpha_{ic}}, \quad (\text{B-13})$$

where ς_i is a normalization constant chosen so that the steady-state price and marginal cost satisfy $\bar{p}_i = \bar{mc}_i = 1$ for all i . Since all prices are normalized in the steady state, each $\bar{w}_c$ represents the real wage of household c . The constant ς_i thus absorbs these steady-state real wages and is given by

$$\varsigma_i = \prod_{j=1}^N \omega_{ij}^{-\omega_{ij}} \prod_{c=1}^C \left(\frac{\bar{w}_c}{\alpha_{ic}} \right)^{\alpha_{ic}}.$$

Taking logs of (B-13) yields the marginal-cost system

$$\log mc_i = -\log z_i + \sum_{j=1}^N \omega_{ij} \log p_j + \sum_{c=1}^C \alpha_{ic} \Delta \log w_c \quad (\text{B-14})$$

where $\Delta \log w_c \equiv \log w_c - \log \bar{w}_c$ denotes the deviation of household c 's wage from its steady-state level.

Or in matrix notation,

$$\log mc = -\log z + \Omega \log p + \alpha \Delta \log w \quad (\text{B-15})$$

To characterize the response of nominal wages to monetary policy, start from the definition of the labor income share:

$$\log w = \log m + \log \Lambda - \log L. \quad (\text{B-16})$$

Applying the first-order approximation results from Lemma 1 yields

$$d \log w = d \log m + d \log \chi + \boldsymbol{\eta}^{-1} d \log L + o(\|\xi\|) \quad (\text{B-17})$$

$$= \log m + (\Gamma + \ell^\mu \boldsymbol{\eta}^{-1})' \log \boldsymbol{\mu} + o(\|\xi\|) \quad (\text{B-18})$$

Here, the term $\varrho^w(\Gamma + \ell^\mu \boldsymbol{\eta}^{-1})' \log \boldsymbol{\mu}$ captures the endogenous feedback of sectoral markups onto nominal wages. Note that sectoral markups are linked to sectoral prices through (4): $\log \boldsymbol{\mu} = (I - \boldsymbol{\delta}^{-1}) \log \boldsymbol{p} + o(\|\xi\|)$.

Substituting these expressions into the marginal-cost system gives

$$\begin{aligned} \log \boldsymbol{p} &= \boldsymbol{\delta}(-\log \boldsymbol{z} + \Omega \log \boldsymbol{p} + \alpha d \log w) + o(\|\xi\|) \\ &= -(I - \boldsymbol{\delta}\Omega)^{-1} \boldsymbol{\delta} \log \boldsymbol{z} + (I - \boldsymbol{\delta}\Omega)^{-1} \boldsymbol{\delta} \alpha d \log w + o(\|\xi\|) \\ &= -(I - \boldsymbol{\delta}\Omega)^{-1} \boldsymbol{\delta} \log \boldsymbol{z} + \varrho^w \mathbf{1} \cdot \log m + \varrho^w(\Gamma + \ell^\mu \boldsymbol{\eta}^{-1})'(I - \boldsymbol{\delta}^{-1}) \log \boldsymbol{p} + o(\|\xi\|) \end{aligned}$$

Rearranging terms gives the first-order equilibrium mapping from productivity and monetary shocks to sectoral prices,

$$\begin{aligned} \log \boldsymbol{p} &= -[I - \varrho^w(\Gamma + \ell^\mu \boldsymbol{\eta}^{-1})'(I - \boldsymbol{\delta}^{-1})]^{-1} (I - \boldsymbol{\delta}\Omega)^{-1} \boldsymbol{\delta} \log \boldsymbol{z} \\ &\quad + [I - \varrho^w(\Gamma + \ell^\mu \boldsymbol{\eta}^{-1})'(I - \boldsymbol{\delta}^{-1})]^{-1} \varrho^w \mathbf{1} \cdot \log m + o(\|\xi\|). \end{aligned} \quad (\text{B-19})$$

□

The next result summarizes the second-order approximation to sectoral wedges in terms of prices, nominal wages and productivities.

Lemma 5. Up to a second-order approximation around the steady-state equilibrium, the sectoral wedge ε_i satisfies

$$\log \varepsilon_i = \sum_{j=1}^N \omega_{ij} \log p_j - \log p_i + \sum_{c=1}^C \alpha_{ic} \Delta \log w_c - \log z_i + \frac{1}{2} \theta_i \vartheta_i + o(\|\xi\|^2) \quad (\text{B-20})$$

where $\Delta \log w_c \equiv \log w_c - \log \bar{w}_c$ denotes the log-deviation of the wage of household c from its steady-state level, and

$$\vartheta_i = \text{Var}(\log p_{ik}) = \int_0^1 (\log p_{ik} - \log p_i)^2 dk - \left(\int_0^1 (\log p_{ik} - \log p_i) dk \right)^2$$

is the cross-sectional dispersion of prices within sector i . With nominal rigidities, this

dispersion satisfies

$$\vartheta_i = \frac{1 - \delta_i}{\delta_i} (\log p_i)^2 + o(\|\xi\|^2) = \frac{\delta_i}{1 - \delta_i} \log^2 \varepsilon_i + o(\|\xi\|^2).$$

Proof of Lemma 5. First, consider the CES price dispersion expression $\int_0^1 (p_{ik}/p_i)^{-\theta_i} dk$. To a second-order approximation,

$$\log \int_0^1 (p_{ik}/p_i)^{-\theta_i} dk = -\theta_i \int_0^1 (\log p_{ik} - \log p_i) dk + \frac{1}{2} \theta_i^2 \vartheta_i + o(\|\xi\|^2).$$

The definition of the sectoral price index in (2) then implies

$$\begin{aligned} 0 &= \frac{1}{1 - \theta_i} \log \left[\int_0^1 (p_{ik}/p_i)^{1-\theta_i} dk \right] \\ &= (1 - \theta_i) \int_0^1 (\log p_{ik} - \log p_i) dk + \frac{1}{2} (1 - \theta_i)^2 \vartheta_i + o(\|\xi\|^2), \end{aligned}$$

which yields

$$\int_0^1 (\log p_{ik} - \log p_i) dk = \frac{1}{2} (\theta_i - 1) \vartheta_i + o(\|\xi\|^2).$$

Thus, the CES price-dispersion term simplifies to

$$\log \int_0^1 (p_{ik}/p_i)^{-\theta_i} dk = -\theta_i \left(\frac{1}{2} (\theta_i - 1) \vartheta_i \right) + \frac{1}{2} \theta_i^2 \vartheta_i + o(\|\xi\|^2) = \frac{1}{2} \theta_i \vartheta_i + o(\|\xi\|^2)$$

This equation, together with (B-14), allows for a log-quadratic approximation of the sectoral wedges in terms of prices, nominal wages, and productivities:

$$\begin{aligned} \log \varepsilon_i &= \log mc_i - \log p_i + \log \int_0^1 (p_{ik}/p_i)^{-\theta_i} dk \\ &= \sum_{c=1}^C \alpha_{ic} \Delta \log w_c - \log z_i + \sum_{j=1}^N \omega_{ij} \log p_j - \log p_i + \frac{1}{2} \theta_i \vartheta_i + o(\|\xi\|^2) \end{aligned}$$

□

Proof of Lemma 2. The household's welfare loss relative to the flexible-price equilibrium

takes the form

$$W(\mathbf{e}_c) - W^*(\mathbf{e}_c) = U_c - U_c^* = (\log C_c - \log C_c^*) - \psi_c \left(\frac{L_c^{1+1/\eta_c}}{1+1/\eta_c} - \frac{(L_c^*)^{1+1/\eta_c}}{1+1/\eta_c} \right).$$

A second-order approximation of the disutility from labor yields

$$\begin{aligned} \psi_c \left(\frac{L_c^{1+1/\eta_c}}{1+1/\eta_c} - \frac{(L_c^*)^{1+1/\eta_c}}{1+1/\eta_c} \right) &= \psi_c(L_c^*)^{1+1/\eta_c} (\hat{l}_c + \frac{1+1/\eta_c}{2} \tilde{l}_c^2) + o(\|\xi\|^2) \\ &= \hat{l}_c + \frac{1+1/\eta_c}{2} \tilde{l}_c^2 + o(\|\xi\|^2) \end{aligned}$$

where the second line uses the fact that $\psi_c(L_c^*)^{1+1/\eta_c} = \Lambda_c^*/\chi_c^* = 1$ as established in (B-9).

Substituting this approximation into the welfare loss expression gives

$$\begin{aligned} W(\mathbf{e}_c) - W^*(\mathbf{e}_c) &= \hat{c}_c - \left(\hat{l}_c + \frac{1+1/\eta_c}{2} \tilde{l}_c^2 \right) + o(\|\xi\|^2) \\ &= \hat{c}_c - \hat{l}_c - \frac{1+1/\eta_c}{2} \tilde{l}_c^2 + o(\|\xi\|^2). \end{aligned} \tag{B-21}$$

The next part of the proof is devoted to obtaining a second-order approximation of allocative efficiency, represented by the difference $\hat{c}_c - \hat{l}_c$.

First, combining equations (B-8) and (B-9), household c 's consumption-labor ratio satisfies

$$\log C_c - \log L_c = \log w_c - \log P_c - \log(\Lambda_c/\chi_c).$$

Second, applying Lemma 5 and aggregating sectoral wedges using household-level Domar weights $(\lambda_i^c)^*$ yields

$$\begin{aligned} \sum_{i=1}^N (\lambda_i^c)^* \log \varepsilon_i &= \sum_{i=1}^N \sum_{j=1}^N (\lambda_i^c)^* \omega_{ij} \log p_j - \sum_{i=1}^N (\lambda_i^c)^* \log p_i \\ &\quad + \sum_{i=1}^N \sum_{f=1}^C (\lambda_i^c)^* \alpha_{if} \Delta \log w_f - \sum_{i=1}^N (\lambda_i^c)^* \log z_i + \frac{1}{2} \sum_{i=1}^N (\lambda_i^c)^* \theta_i \vartheta_i + o(\|\xi\|^2) \\ &= - \sum_{i=1}^N \beta_{ci} \log p_i + \sum_{f=1}^C (\Lambda_f^c)^* \Delta \log w_f - \sum_{i=1}^N (\lambda_i^c)^* \log z_i + \frac{1}{2} \sum_{i=1}^N (\lambda_i^c)^* \theta_i \vartheta_i + o(\|\xi\|^2) \end{aligned}$$

where I use the identities $\sum_{i=1}^N (\lambda_i^c)^* \omega_{ij} = (\lambda_j^c)^* - \beta_{jc}$ and $\sum_{i=1}^N (\lambda_i^c)^* \alpha_{if} = (\Lambda_f^c)^*$.

Recall that the household-level price index satisfies $\log P_c = \sum_{i=1}^N \beta_{ci} \log p_i$. Substituting the expression above into the consumption-labor ratio yields

$$\begin{aligned} \log C_c - \log L_c &= \sum_{i=1}^N (\lambda_i^c)^* \log z_i + \sum_{i=1}^N (\lambda_i^c)^* \log \varepsilon_i - \frac{1}{2} \sum_{i=1}^N (\lambda_i^c)^* \theta_i \vartheta_i \\ &\quad + \log w_c - \sum_{f=1}^C (\Lambda_f^c)^* \log w_f - \log(\Lambda_c/\chi_c) + F_c^1 + o(\|\xi\|^2) \end{aligned} \quad (\text{B-22})$$

where $F_c^1 = \sum_{f=1}^C (\Lambda_f^c)^* \log \bar{w}_f$ is a constant.

Using the expression for the labor income share in equation (B-16), the factorial terms-of-trade $\log w_c - \sum_{f=1}^C (\Lambda_f^c)^* \log w_f$ can be rewritten as

$$\begin{aligned} \text{FTToT}_c &= (\log m + \log \Lambda_c - \log L_c) - \sum_{f=1}^C (\Lambda_f^c)^* (\log m + \log \Lambda_f - \log L_f) \\ &= (\log \Lambda_c - \log L_c) - \sum_{f=1}^C (\Lambda_f^c)^* (\log \Lambda_f - \log L_f) \\ &= \frac{1}{1 + \eta_c} \log \Lambda_c + \frac{\eta_c}{1 + \eta_c} \log \chi_c - \sum_{f=1}^C (\Lambda_f^c)^* \left(\frac{1}{1 + \eta_f} \log \Lambda_f + \frac{\eta_f}{1 + \eta_f} \log \chi_f \right) + F_c^2 \\ &= \frac{1}{1 + \eta_c} \log(\Lambda_c/\chi_c) - \sum_{f=1}^C \frac{1}{1 + \eta_f} (\Lambda_f^c)^* \log(\Lambda_f/\chi_f) + \log \chi_c - \sum_{f=1}^C (\Lambda_f^c)^* \log \chi_f + F_c^2 \end{aligned} \quad (\text{B-23})$$

where the second equality uses the fact that $\sum_{f=1}^C (\Lambda_f^c)^* = 1$, and the third equality substitutes $\log L_c = \frac{\eta_c}{1 + \eta_c} (\log \Lambda_c - \log \chi_c - \log \psi_c)$ with the resulting constant term $F_c^2 = \frac{\eta_c}{1 + \eta_c} \log \psi_c - \sum_{f=1}^C (\Lambda_f^c)^* \frac{\eta_f}{1 + \eta_f} \log \psi_f$.

Thus, substituting the expression for the terms-of-trade effect yields

$$\begin{aligned} \log C_c - \log L_c &= \sum_{i=1}^N (\lambda_i^c)^* \log z_i + \sum_{i=1}^N (\lambda_i^c)^* \log \varepsilon_i - \frac{1}{2} \sum_{i=1}^N (\lambda_i^c)^* \theta_i \vartheta_i - \frac{\eta_c}{1 + \eta_c} \log(\Lambda_c/\chi_c) \\ &\quad - \sum_{f=1}^C \frac{1}{1 + \eta_f} (\Lambda_f^c)^* \log(\Lambda_f/\chi_f) + \log \chi_c - \sum_{f=1}^C (\Lambda_f^c)^* \log \chi_f + F_c + o(\|\xi\|^2), \end{aligned}$$

where $F_c = F_c^1 + F_c^2$ is a constant.

In the flexible-price equilibrium, where $\varepsilon_i^* = 1$, $\vartheta_i^* = 0$, and $\Lambda_c^* = \chi_c^*$, the household-level consumption-labor ratio simplifies to

$$\log C_c^* - \log L_c^* = \sum_{i=1}^N (\lambda_i^c)^* \log z_i + \log \chi_c^* - \sum_{f=1}^C (\Lambda_f^c)^* \log \chi_f^* + F_c + o(\|\xi\|^2).$$

Taking the difference between the two expressions gives the second-order approximation of allocative efficiency:

$$\begin{aligned} \hat{c}_c - \hat{l}_c &= \log C_c - \log L_c - (\log C_c^* - \log L_c^*) \\ &= \sum_{i=1}^N (\lambda_i^c)^* \log \varepsilon_i - \frac{1}{2} \sum_{i=1}^N (\lambda_i^c)^* \theta_i \vartheta_i + \hat{\chi}_c - \sum_{f=1}^C (\Lambda_f^c)^* \hat{\chi}_f \\ &\quad - \frac{\eta_c}{1 + \eta_c} \log(\Lambda_c / \chi_c) - \sum_{f=1}^C \frac{1}{1 + \eta_f} (\Lambda_f^c)^* \log(\Lambda_f / \chi_f) + o(\|\xi\|^2) \end{aligned} \quad (\text{B-24})$$

Finally, applying Lemma 4 allows us to express the welfare loss in terms of sectoral wedges,

$$\begin{aligned} W(\mathbf{e}_c) - W^*(\mathbf{e}_c) &= \sum_{i=1}^N \left[(\lambda_i^c)^* - \frac{\eta_c}{1 + \eta_c} \frac{\lambda_i^* \Phi_{ic}}{\chi_c^*} - \sum_{f=1}^C \frac{1}{1 + \eta_f} (\Lambda_f^c)^* \frac{\lambda_i^* \Phi_{if}}{\chi_f^*} - \Gamma_{ic} + \sum_{f=1}^C (\Lambda_f^c)^* \Gamma_{if} \right] \log \varepsilon_i \\ &\quad - \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \left[\frac{\eta_c}{1 + \eta_c} \Upsilon_{ij}^c + \sum_{f=1}^C \frac{1}{1 + \eta_f} \Lambda_f^c \Upsilon_{ij}^f - \Xi_{ij}^c + \sum_{f=1}^C \Lambda_f^c \Xi_{ij}^f \right] \log \varepsilon_j \log \varepsilon_i \\ &\quad - \frac{1}{2} \sum_{i=1}^N (\lambda_i^c)^* \theta_i \vartheta_i - \frac{1 + 1/\eta_c}{2} \hat{l}_c^2 + o(\|\xi\|^2). \end{aligned} \quad (\text{B-25})$$

Note that $\vartheta_i = \frac{\delta_i}{1 - \delta_i} \log^2 \varepsilon_i + o(\|\xi\|^2)$ and $\hat{l}_c^2 = (\sum_{i=1}^N \ell_{ic}^\mu \log \varepsilon_i)^2 + o(\|\xi\|^2)$, so the last two terms are quadratic in $\log \varepsilon_i$. Grouping them accordingly completes the proof. $\square$

The second-order approximation of allocative efficiency derived in Lemma 2 serves as a key intermediate step for establishing Proposition 1, which provides a first-order characterization of household-level allocative efficiency.

Proof of Proposition 1. Starting from the second-order approximation to the consumption–labor

ratio in equation (B-22), a first-order approximation yields:

$$\begin{aligned}
\log C_c - \log L_c &= \sum_{i=1}^N (\lambda_i^c)^* \log z_i + \sum_{i=1}^N (\lambda_i^c)^* \log \varepsilon_i - \log(\Lambda_c / \chi_c) \\
&\quad + \log w_c - \sum_{f=1}^C (\Lambda_f^c)^* \log w_f + F_c^1 + o(\|\xi\|) \\
&= \sum_{i=1}^N (\lambda_i^c)^* \log z_i - \sum_{i=1}^N \left(\frac{\lambda_i^* \Phi_{ic}}{\chi_c^*} - (\lambda_i^c)^* \right) \log \varepsilon_i \\
&\quad + \log w_c - \sum_{f=1}^C (\Lambda_f^c)^* \log w_f + F_c^1 + o(\|\xi\|)
\end{aligned}$$

where the second equality uses the first-order approximation of $\log(\Lambda_c / \chi_c)$ from equation (B-3).

In the flexible-price equilibrium, where $\varepsilon_i^* = 1$, and $\log w_c = \log \bar{w}_c + \log m$ for all c , the consumption-labor ratio simplifies to

$$\log C_c^* - \log L_c^* = \sum_{i=1}^N (\lambda_i^c)^* \log z_i + \log \bar{w}_c - \sum_{f=1}^C (\Lambda_f^c)^* \log \bar{w}_f + F_c^1 + o(\|\xi\|)$$

Taking the difference between sticky-price and flexible-price outcomes yields the household-level allocative efficiency

$$\hat{c}_c - \hat{l}_c = - \sum_{i=1}^N \left(\frac{\lambda_i^* \Phi_{ic}}{\chi_c^*} - (\lambda_i^c)^* \right) \log \varepsilon_i + d \log w_c - \sum_{f=1}^C (\Lambda_f^c)^* d \log w_f + o(\|\xi\|).$$

Approximating the factorial terms-of-trade effect to a first order using equation (B-23)

and Lemma 4 yields

$$\begin{aligned}
\Delta \text{FTOT}_c &= d \log w_c - \sum_{f=1}^C (\Lambda_f^c)^* d \log w_f \\
&= \frac{1}{1 + \eta_c} d \log (\Lambda_c / \chi_c) - \sum_{f=1}^C \frac{1}{1 + \eta_f} (\Lambda_f^c)^* d \log (\Lambda_f / \chi_f) + d \log \chi_c - \sum_{f=1}^C (\Lambda_f^c)^* d \log \chi_f \\
&= - \sum_{i=1}^N \left[\left(\Gamma_{ic} - \frac{1}{1 + \eta_c} \frac{\lambda_i^* \Phi_{ic}}{\chi_c^*} \right) - \sum_{f=1}^C (\Lambda_f^c)^* \left(\Gamma_{if} - \frac{1}{1 + \eta_f} \frac{\lambda_i^* \Phi_{if}}{\chi_f^*} \right) \right] \log \varepsilon_i + o(\|\xi\|).
\end{aligned}$$

□

Proof of Proposition 3. The social welfare loss is defined as the Pareto-weighted sum of household-level losses:

$$W(\{\kappa_c\}) - W^*(\{\kappa_c\}) = \sum_{c=1}^C \kappa_c [W(\mathbf{e}_c) - W^*(\mathbf{e}_c)].$$

Applying Lemma 2 to each term, we have

$$W(\mathbf{e}_c) - W^*(\mathbf{e}_c) = \mathcal{J}'_c \log \boldsymbol{\mu} - \frac{1}{2} \log \boldsymbol{\mu}' \mathcal{L}_c \log \boldsymbol{\mu}.$$

Summing across households yields

$$W(\{\kappa_c\}) - W^*(\{\kappa_c\}) = \sum_{c=1}^C \kappa_c \mathcal{J}'_c \log \boldsymbol{\mu} - \frac{1}{2} \log \boldsymbol{\mu}' \mathcal{L}(\{\kappa_c\}) \log \boldsymbol{\mu},$$

where $\mathcal{L}(\{\kappa_c\}) = \sum_{c=1}^C \kappa_c \mathcal{L}_c$.

□

Corollary 4 (Welfare Loss under Income-Share Weights). When Pareto weights are set equal to households' income shares, i.e., $\kappa_c = \chi_c$ for all c , reallocations across households do not generate first-order welfare gains. This implies that monetary policy cannot improve aggregate efficiency through redistribution at the first order. In this case, the aggregate welfare loss simplifies to a purely quadratic form in ex post markups

$$W(\{\chi_c\}_{c=1}^C) - W^*(\{\chi_c\}_{c=1}^C) = -\frac{1}{2} \log \boldsymbol{\mu}' \mathcal{L}(\{\chi_c\}_{c=1}^C) \log \boldsymbol{\mu}.$$

The aggregate loss matrix is given by

$$\mathcal{L}(\{\chi_c\}_{c=1}^C) = \mathcal{L}^{\text{e.g.}}(\{\chi_c\}) + \mathcal{L}^{\text{within}}(\{\chi_c\}) + \mathcal{L}^{\text{across}}(\{\chi_c\}),$$

where

$$\begin{aligned}\mathcal{L}_{ij}^{\text{e.g.}}(\{\chi_c\}) &= \sum_{c=1}^C \frac{\eta_c}{1 + \eta_c} \chi_c^{-1} \lambda_i \lambda_j \Phi_{ic} \Phi_{jc}, \\ \mathcal{L}_{ij}^{\text{within}}(\{\chi_c\}) &= \lambda_i \theta_i \frac{\delta_i}{1 - \delta_i} l_{ij}, \\ \mathcal{L}_{ij}^{\text{across}}(\{\chi_c\}) &= \lambda_i \Psi_{ij} + \lambda_j \Psi_{ji} - \lambda_i l_{ij} \\ &\quad + \sum_{c=1}^C \left[(\lambda_i \Phi_{ic} - \lambda_i^c \chi_c) \Gamma_{jc} + (\lambda_j \Phi_{jc} - \lambda_j^c \chi_c) \Gamma_{ic} - \chi_c^{-1} \lambda_i \lambda_j \Phi_{ic} \Phi_{jc} \right].\end{aligned}$$

Proof of Corollary 4. From equation (B-21), the aggregate welfare loss under income-share Pareto weights $\{\chi_c\}$ is:

$$\begin{aligned}W(\{\chi_c\}_{c=1}^C) - W^*(\{\chi_c\}_{c=1}^C) &= \sum_{c=1}^C \chi_c^* [W(\mathbf{e}_c) - W^*(\mathbf{e}_c)] \\ &= \sum_{c=1}^C \chi_c^* (\hat{c}_c - \hat{l}_c) - \sum_{c=1}^C \frac{1 + 1/\eta_c}{2} \chi_c^* \hat{l}_c^2 + o(\|\xi\|^2)\end{aligned}$$

Using the second-order approximation of allocative efficiency in (B-24) and aggregating

across households with weights χ_c^* yields

$$\begin{aligned}
\sum_{c=1}^C \chi_c^* (\hat{c}_c - \hat{l}_c) &= \sum_{c=1}^C \sum_{i=1}^N \chi_c^* (\lambda_i^c)^* \log \varepsilon_i - \frac{1}{2} \sum_{c=1}^C \sum_{i=1}^N \chi_c^* (\lambda_i^c)^* \theta_i \vartheta_i + \sum_{c=1}^C \chi_c^* \hat{\chi}_c - \sum_{c=1}^C \sum_{f=1}^C \chi_c^* (\Lambda_f^c)^* \hat{\chi}_f \\
&\quad - \sum_{c=1}^C \frac{\eta_c}{1 + \eta_c} \chi_c^* \log(\Lambda_c / \chi_c) - \sum_{c=1}^C \sum_{f=1}^C \frac{1}{1 + \eta_f} \chi_c^* (\Lambda_f^c)^* \log(\Lambda_f / \chi_f) + o(\|\xi\|^2) \\
&= \sum_{i=1}^N \lambda_i^* \log \varepsilon_i - \frac{1}{2} \sum_{i=1}^N \lambda_i^* \theta_i \vartheta_i + \sum_{c=1}^C \chi_c^* \hat{\chi}_c - \sum_{f=1}^C \chi_f^* \hat{\chi}_f \\
&\quad - \sum_{c=1}^C \frac{\eta_c}{1 + \eta_c} \chi_c^* \log(\Lambda_c / \chi_c) - \sum_{f=1}^C \frac{1}{1 + \eta_f} \chi_f^* \log(\Lambda_f / \chi_f) + o(\|\xi\|^2) \\
&= \sum_{i=1}^N \lambda_i^* \log \varepsilon_i - \frac{1}{2} \sum_{i=1}^N \lambda_i^* \theta_i \vartheta_i - \sum_{c=1}^C \chi_c^* \log(\Lambda_c / \chi_c) + o(\|\xi\|^2), \tag{B-26}
\end{aligned}$$

where the second equality uses the identities: $\sum_{c=1}^C \chi_c^* (\lambda_i^c)^* = \lambda_i^*$ and $\sum_{c=1}^C \chi_c^* (\Lambda_f^c)^* = \Lambda_f^* = \chi_f^*$.

Next, aggregating the second-order labor-wedge expansion (B-12) across households using income shares gives

$$\begin{aligned}
\sum_{c=1}^C \chi_c^* \log(\Lambda_c / \chi_c) &= \sum_{i=1}^N \lambda_i^* \log \varepsilon_i + \sum_{i=1}^N \sum_{j=1}^N \lambda_j^* \Psi_{ji} \log \varepsilon_j \log \varepsilon_i - \sum_{i=1}^N \sum_{j=1}^N \sum_{f=1}^C \lambda_i^f \chi_f^* \Gamma_{jf} \log \varepsilon_j \log \varepsilon_i \\
&\quad + \sum_{i=1}^N \sum_{j=1}^N \sum_{c=1}^C \lambda_i^* \Phi_{ic} \Gamma_{jc} \log \varepsilon_j \log \varepsilon_i - \frac{1}{2} \sum_{i=1}^N \lambda_i^* \log^2 \varepsilon_i - \frac{1}{2} \sum_{c=1}^C \chi_c^* \left(\sum_{i=1}^N \frac{\lambda_i^* \Phi_{ic}}{\chi_c^*} \log \varepsilon_i \right)^2 + o(\|\xi\|^2) \\
&= \sum_{i=1}^N \lambda_i^* \log \varepsilon_i + \sum_{i=1}^N \sum_{j=1}^N \lambda_j^* \Psi_{ji} \log \varepsilon_j \log \varepsilon_i + \sum_{i=1}^N \sum_{j=1}^N \sum_{c=1}^C (\lambda_i^* \Phi_{ic} - \lambda_i^c \chi_c^*) \Gamma_{jc} \log \varepsilon_j \log \varepsilon_i \\
&\quad - \frac{1}{2} \sum_{i=1}^N \lambda_i^* \log^2 \varepsilon_i - \frac{1}{2} \sum_{c=1}^C (\chi_c^*)^{-1} \left(\sum_{i=1}^N \lambda_i^* \Phi_{ic} \log \varepsilon_i \right)^2 + o(\|\xi\|^2).
\end{aligned}$$

Substituting (B-26) into the aggregate welfare loss function and eliminating the interme-

diate term $\sum_c \chi_c^* \log(\Lambda_c/\chi_c)$ yields

$$\begin{aligned}
W(\{\chi_c\}) - W^*(\{\chi_c\}) &= \sum_{c=1}^C \chi_c^* (\hat{c}_c - \hat{l}_c) - \sum_{c=1}^C \frac{1 + 1/\eta_c}{2} \chi_c^* \hat{l}_c^2 + o(\|\xi\|^2) \\
&= -\frac{1}{2} \left[\sum_{c=1}^C \frac{\eta_c}{1 + \eta_c} (\chi_c^*)^{-1} \left(\sum_{i=1}^N \lambda_i^* \Phi_{ic} \log \varepsilon_i \right)^2 + \sum_{i=1}^N \lambda_i^* \theta_i \vartheta_i \right. \\
&\quad + 2 \sum_{i=1}^N \sum_{j=1}^N \lambda_i^* \Psi_{ij} \log \varepsilon_j \log \varepsilon_i + 2 \sum_{i=1}^N \sum_{j=1}^N \sum_{c=1}^C (\lambda_i^* \Phi_{ic} - \lambda_i^c \chi_c^*) \Gamma_{jc} \log \varepsilon_j \log \varepsilon_i \\
&\quad \left. - \sum_{i=1}^N \lambda_i^* \log^2 \varepsilon_i - \sum_{c=1}^C (\chi_c^*)^{-1} \left(\sum_{i=1}^N \lambda_i^* \Phi_{ic} \log \varepsilon_i \right)^2 \right] + o(\|\xi\|^2) \\
&= -\frac{1}{2} \log \mu' \mathcal{L}(\{\chi_c\}) \log \mu + o(\|\xi\|^2).
\end{aligned}$$

where the aggregate loss matrix $\mathcal{L}(\{\chi_c\})$ admits the decomposition stated in the corollary. $\square$

Lemma 6 (Implementability of Price-Targeting Rules). To a first-order approximation, any price-targeting rule of the form $\zeta' \log \mathbf{p} = \pi$, for some pair $(\zeta, \pi) \in \mathbb{R}^N \times \mathbb{R}$, can be implemented by a monetary policy rule for nominal aggregate demand of the form $\log m(z) = \varsigma_0 + \sum_{i=1}^N \varsigma_i \log z_i$, for a suitable choice of coefficients (ς_0, ς) .

Proof of Lemma 6. From the first-order approximation in (B-19), the vector of sectoral prices satisfies

$$\log \mathbf{p} = -\varrho^z \log z + \varrho^m \cdot \log m + o(\|\xi\|),$$

where $\varrho^z \equiv \left[I - \varrho^w (\Gamma + \ell^\mu \eta^{-1})' (I - \delta^{-1}) \right]^{-1} (I - \delta \Omega)^{-1} \delta$ characterizes the elasticity of sectoral prices with respect to productivity shocks under price stickiness.

Taking inner products with ζ and applying the price-targeting condition gives

$$\pi = \zeta' \log \mathbf{p} = -\zeta' \varrho^z \log z + \zeta' \varrho^m \cdot \log m + o(\|\xi\|).$$

Solving for $\log m$ yields

$$\log m(z) = \frac{\pi}{\zeta' \varrho^m} + \frac{\zeta' \varrho^z}{\zeta' \varrho^m} \log z + o(\|\xi\|).$$

Thus, setting

$$\varsigma_0 = \frac{\pi}{\zeta' \varrho^m}, \quad \varsigma_i = \frac{\zeta' \varrho^{z_{(:,i)}}}{\zeta' \varrho^m}, \quad i = 1, \dots, N,$$

defines a monetary policy rule that implements the price target to a first-order approximation. Note that since $\varrho^m = \varrho^z \alpha \mathbf{1}$, price-targeting rules span only a subset of admissible monetary policies to first order. $\square$

Proof of Theorem 1. The monetary authority chooses the policy instrument $\log m$ to minimize the aggregate welfare loss in (17). Since monetary policy affects welfare only through its impact on prices and markups, the first-order condition is obtained by differentiating the loss with respect to $\log m$:

$$\frac{d}{d \log m} [W(\{\kappa_c\}) - W^*(\{\kappa_c\})] = \frac{d \log \mu'}{d \log m} \left(\sum_{c=1}^C \kappa_c \mathcal{J}_c - \mathcal{L}(\{\kappa_c\}) \log \mu \right) = 0. \quad (\text{B-27})$$

With nominal rigidities, sectoral markups are related to sectoral prices through (4), $\log \mu = -(\delta^{-1} - I) \log p$. Moreover, from Proposition 2, sectoral prices respond to the policy shock according to $\frac{d \log p}{d \log m} = \varrho^m$. Combining the two expressions implies

$$\frac{d \log \mu}{d \log m} = -(\delta^{-1} - I) \varrho^m.$$

Substituting into the first-order condition (B-27) gives

$$(\varrho^m)' (I - \delta^{-1}) \mathcal{L}(\{\kappa_c\}) \log \mu = \sum_{c=1}^C \kappa_c \mathcal{J}'_c (I - \delta^{-1}) \varrho^m.$$

Using (4) once more, the condition can be re-expressed as a target criterion for a weighted sectoral price index

$$\sum_{j=1}^N \zeta_j^*(\{\kappa_c\}) \log p_j = \pi^*(\{\kappa_c\}),$$

where the optimal weights and target level are given by

$$\zeta_j^*(\{\kappa_c\}) = (\delta_j^{-1} - 1) \sum_{i=1}^N (\delta_i^{-1} - 1) \varrho_i^m \mathcal{L}_{ij}(\{\kappa_c\}),$$

and

$$\pi^*({\kappa}_c) = \sum_{c=1}^C \kappa_c \mathcal{J}'_c(I - \delta^{-1}) \varrho^m,$$

respectively.

Finally, using the decomposition

$$\mathcal{L}({\kappa}_c) = \mathcal{L}^{\text{e.g.}}({\kappa}_c) + \mathcal{L}^{\text{within}}({\kappa}_c) + \mathcal{L}^{\text{across}}({\kappa}_c),$$

the sectoral weights $\zeta_j^*({\kappa}_c)$ admit the corresponding decomposition:

$$\zeta_j^*({\kappa}_c) = \zeta_j^{\text{e.g.}}({\kappa}_c) + \zeta_j^{\text{within}}({\kappa}_c) + \zeta_j^{\text{across}}({\kappa}_c),$$

where each term reflects the contribution of employment-gap volatility, within-sector price dispersion, and cross-sector misallocation to the overall optimal sectoral weight. $\square$

Theorem 1 characterizes optimal monetary policy in the general heterogeneous-agent setting, in which Pareto weights, consumption baskets, and ownership shares can vary arbitrarily across households. To connect this characterization to the existing production-network literature, the following corollary specializes to the symmetric benchmark in which Pareto weights and ownership shares coincide with income shares, and consumption baskets are homogeneous across households. In this limit, the heterogeneity-driven cross-household terms collapse, and the optimal industry weights reduce to the closed-form characterization of [La'O and Tahbaz-Salehi \(2022\)](#).

Corollary 5. Consider the benchmark in which Pareto weights coincide with households' income shares, consumption baskets are homogeneous across households, and ownership shares are proportional to income shares: $\beta_{ci} = b_i$ and $\Phi_{ic} = \chi_c$ for all c and i . In this case, the industry weights in the optimal target price index simplify to

$$\begin{aligned} \zeta_j^{\text{e.g.}}({\chi}_c)_{c=1}^C &= (1/\delta_j - 1)\lambda_j, \\ \zeta_j^{\text{within}}({\chi}_c)_{c=1}^C &= (1/\delta_j - 1)\lambda_j\theta_j\mathcal{K}_j, \\ \zeta_j^{\text{across}}({\chi}_c)_{c=1}^C &= (1/\delta_j - 1) \sum_{i=1}^N (1/\delta_i - 1)\mathcal{K}_i(\lambda_i\Psi_{ij} + \lambda_j\Psi_{ji} - \lambda_{itij} - \lambda_{itj}), \end{aligned}$$

where $\mathcal{K}_j = \varrho_j^m / (\sum_{c=1}^C \chi_c \ell_c^m)$ represents the slope of the Phillips curve for sector j .

Lemma 7. Given a realization of productivity shocks $\log z$, the vector of ex post markups is, to a first-order approximation, a function of the policy pair (ζ, π) :

$$\log \mu = \mathcal{M}(\zeta, \pi) = \frac{\pi}{\zeta' \varrho^m} (I - \delta^{-1}) \varrho^m + (I - \delta^{-1}) \left(\frac{\varrho^m \zeta'}{\zeta' \varrho^m} - I \right) \varrho^z \log z + o(\|\xi\|). \quad (\text{B-28})$$

Proof of Lemma 7. By Lemma 6, under the policy instrument (ζ, π) , sectoral prices satisfy

$$\begin{aligned} \log p &= -\varrho^z \log z + \varrho^m \log m(z) + o(\|\xi\|), \\ &= -\varrho^z \log z + \varrho^m \left(\frac{\pi}{\zeta' \varrho^m} + \frac{\zeta' \varrho^z}{\zeta' \varrho^m} \log z \right) + o(\|\xi\|), \\ &= \frac{\pi}{\zeta' \varrho^m} \varrho^m + \left(\frac{\varrho^m \zeta'}{\zeta' \varrho^m} - I \right) \varrho^z \log z + o(\|\xi\|) \end{aligned}$$

Substituting this expression into (4) delivers the stated mapping $\log \mu = \mathcal{M}(\zeta, \pi)$. $\square$

Lemma 8. Under a given union-wide price-index regime ζ , the unilateral optimal policy can be implemented by a state-contingent inflation rule, $\tilde{\pi}_c(\zeta; z) = \arg \min_{\pi} \mathbb{L}(\zeta, \pi)$. To a first-order approximation, this rule is given by

$$\tilde{\pi}_c(\zeta; z) = \frac{\zeta' \varrho^m}{\zeta^*(\mathbf{e}_c)' \varrho^m} \pi^*(\mathbf{e}_c) + \zeta' \left[I - \frac{\varrho^m \zeta^*(\mathbf{e}_c)'}{\zeta^*(\mathbf{e}_c)' \varrho^m} \right] \varrho^z \log z + o(\|\xi\|).$$

Proof of Lemma 8. By Lemma 6, the unilateral optimal policy for country c can, to a first-order approximation, be implemented by a monetary rule of the form

$$\log m(z) = \frac{\pi^*(\mathbf{e}_c)}{\zeta^*(\mathbf{e}_c)' \varrho^m} + \frac{\zeta^*(\mathbf{e}_c)' \varrho^z}{\zeta^*(\mathbf{e}_c)' \varrho^m} \log z + o(\|\xi\|).$$

Under a given union-wide price-index regime ζ , any admissible price-targeting rule that delivers inflation π induces a monetary policy of the form

$$\log m(z) = \frac{\pi}{\zeta' \varrho^m} + \frac{\zeta' \varrho^z}{\zeta' \varrho^m} \log z + o(\|\xi\|).$$

To replicate the unilateral optimum under regime ζ , the induced monetary rule must coincide with the unilateral rule state by state, up to first order. Equating the two expressions yields

$$\frac{\pi}{\zeta' \varrho^m} + \frac{\zeta' \varrho^z}{\zeta' \varrho^m} \log z = \frac{\pi^*(\mathbf{e}_c)}{\zeta^*(\mathbf{e}_c)' \varrho^m} + \frac{\zeta^*(\mathbf{e}_c)' \varrho^z}{\zeta^*(\mathbf{e}_c)' \varrho^m} \log z + o(\|\xi\|).$$

Rearranging the equality establishes the stated inflation rule. $\square$

Proof of Proposition 4. By Lemma 7, the expected welfare loss of country c can be written as

$$\begin{aligned}\mathbb{E}[\mathbb{L}_c(\zeta, \pi)] &= \mathbb{E}\left[-\mathcal{J}'_c \mathcal{M}(\zeta, \pi) + \frac{1}{2} \mathcal{M}(\zeta, \pi)' \mathcal{L}_c \mathcal{M}(\zeta, \pi)\right] \\ &= -\frac{\pi}{\zeta' \varrho^m} \mathcal{J}'_c (I - \delta^{-1}) \varrho^m + \frac{1}{2} \left(\frac{\pi}{\zeta' \varrho^m}\right)^2 (\varrho^m)' (I - \delta^{-1}) \mathcal{L}_c (I - \delta^{-1}) \varrho^m \\ &\quad + \frac{1}{2} \text{tr}((\mu^z)' \mathcal{L}_c \mu^z \Sigma_z) \\ &= -\frac{\pi}{\zeta' \varrho^m} \pi^*(\mathbf{e}_c) + \frac{1}{2} \left(\frac{\pi}{\zeta' \varrho^m}\right)^2 \zeta^*(\mathbf{e}_c)' \varrho^m + \frac{1}{2} \text{tr}((\mu^z)' \mathcal{L}_c \mu^z \Sigma_z)\end{aligned}\quad (\text{B-29})$$

where $\mu^z(\zeta) = (I - \delta^{-1}) \left(\frac{\varrho^m \zeta'}{\zeta' \varrho^m} - I\right) \varrho^z$ and $\Sigma_z = \mathbb{E}[\log z \log z']$ denotes the variance–covariance matrix of the sectoral productivity shocks.

The last equality follows from Theorem 1, which characterizes the unilateral optimal policy

$$(\zeta^*(\mathbf{e}_c), \pi^*(\mathbf{e}_c)) = \left((I - \delta^{-1}) \mathcal{L}_c (I - \delta^{-1}) \varrho^m, \mathcal{J}'_c (I - \delta^{-1}) \varrho^m\right).$$

Note that the trace term is independent of π . Minimization with respect to π yields the unilateral inflation stance of country c :

$$\pi_c(\zeta) = \arg \min_{\pi \in \mathbb{R}} \mathbb{E}[\mathbb{L}_c(\zeta, \pi)] = \frac{\zeta' \varrho^m}{\zeta^*(\mathbf{e}_c)' \varrho^m} \pi^*(\mathbf{e}_c).$$

$\square$

Proof of Proposition 5. By Lemma 7 and using $\mathbb{E}[\log z] = \mathbf{0}$, the expected sectoral markup distortions satisfy

$$\mathbb{E}[\mathcal{M}(\zeta, \pi)] = \frac{\pi}{\zeta' \varrho^m} (I - \delta^{-1}) \varrho^m + o(\|\xi\|).$$

Evaluating at $\pi = \pi_c(\zeta) = \frac{\zeta' \varrho^m}{\zeta^*(\mathbf{e}_c)' \varrho^m} \pi^*(\mathbf{e}_c)$ yields

$$\mathbb{E}[\mathcal{M}(\zeta, \pi_c(\zeta))] = \frac{\pi^*(\mathbf{e}_c)}{\zeta^*(\mathbf{e}_c)' \varrho^m} (I - \delta^{-1}) \varrho^m + o(\|\xi\|) = \mathbb{E}[\mathcal{M}(\zeta^*(\mathbf{e}_c), \pi^*(\mathbf{e}_c))].$$

Premultiplying by $\mathcal{J}'_c$ yields the stated equality in expected allocative efficiency

$$\mathbb{E}[\mathcal{J}'_c \mathcal{M}(\zeta, \pi_c(\zeta))] = \mathbb{E}[\mathcal{J}'_c \mathcal{M}(\zeta^*(\mathbf{e}_c), \pi^*(\mathbf{e}_c))] = \frac{[\pi^*(\mathbf{e}_c)]^2}{\zeta^*(\mathbf{e}_c)' \varrho^m} \geq 0 = \mathbb{E}[\mathcal{J}'_c \mathcal{M}(\zeta, 0)].$$

□

Proof of Proposition 6. The policy-alignment loss admits the decomposition

$$\text{PAL}_c(\zeta, \pi) = \underbrace{\mathbb{L}_c(\zeta, \pi) - \mathbb{L}_c(\zeta, \pi_c(\zeta))}_{\text{inflation misalignment}} + \underbrace{\mathbb{L}_c(\zeta, \pi_c(\zeta)) - \mathbb{L}_c(\zeta^*(\mathbf{e}_c), \pi^*(\mathbf{e}_c))}_{\text{price-index misalignment}}.$$

For the first term, equations (B-29) and (21) yield

$$\mathbb{E}[\mathbb{L}_c(\zeta, \pi) - \mathbb{L}_c(\zeta, \pi_c(\zeta))] = \frac{1}{2} \frac{\zeta^*(\mathbf{e}_c)' \varrho^m}{(\zeta' \varrho^m)^2} [\pi_c(\zeta) - \pi]^2.$$

For the second term, define

$$\begin{aligned} \mathcal{I}_c(\zeta) &\equiv \mathbb{E}[\mathbb{L}_c(\zeta, \pi_c(\zeta)) - \mathbb{L}_c(\zeta^*(\mathbf{e}_c), \pi^*(\mathbf{e}_c))] \\ &= \frac{1}{2} \text{tr}(\mu^z(\zeta)' \mathcal{L}_c \mu^z(\zeta) \Sigma_z) - \frac{1}{2} \text{tr}(\mu^z(\zeta^*(\mathbf{e}_c))' \mathcal{L}_c \mu^z(\zeta^*(\mathbf{e}_c)) \Sigma_z) \\ &= \frac{1}{2} (\varrho^m)' (I - \delta^{-1}) \mathcal{L}_c (I - \delta^{-1}) \varrho^m \left(\frac{\zeta}{\zeta' \varrho^m} - \frac{\zeta^*(\mathbf{e}_c)}{\zeta^*(\mathbf{e}_c)' \varrho^m} \right)' \varrho^z \Sigma_z \varrho^{z'} \left(\frac{\zeta}{\zeta' \varrho^m} - \frac{\zeta^*(\mathbf{e}_c)}{\zeta^*(\mathbf{e}_c)' \varrho^m} \right) \geq 0, \end{aligned}$$

where non-negativity follows from the definition of PAL: if $\mathcal{I}_c(\hat{\zeta}) < 0$ for some $\hat{\zeta}$, then $\mathbb{E}[\text{PAL}_c(\hat{\zeta}, \pi_c(\hat{\zeta}))] = \mathcal{I}_c(\hat{\zeta}) < 0$, contradicting $\text{PAL}_c(\zeta, \pi) \geq 0$. □

C Data Appendix

Input–output matrix. The input–output calibration uses the World Input–Output Database (WIOD; [Timmer et al., 2015](#)), which reports annual industry-level data for 44 economies, including a composite “rest of the world,” and 56 industries over the period 2000–2014.

I extract the 20 euro area countries and reconstruct their joint production network. All intra-euro input–output linkages at the industry level are retained exactly as in the WIOD, while the remaining non-euro countries are collapsed into 56 aggregate external sectors representing bilateral trade flows with the euro area. These external sectors transmit import and export linkages but do not engage in domestic production or trade among themselves,

ensuring that the rest of the world enters the system only through its trade exposure to the euro area.

The resulting dataset combines: (i) a fully detailed euro area production network with cross-country and cross-industry linkages, and (ii) a compact representation of external trade linkages aggregated by industry.

From the WIOD tables, I collect gross output, intermediate inputs, and value added, and aggregate final demand across destinations. Country income shares are then computed in proportion to their final-use shares and normalized to sum to one across the euro area. The final data are organized as an annual panel for 2000–2014 and form the empirical basis for model calibration.

Price flexibility calibration. To calibrate the vector of price flexibilities, I use industry-level frequencies of price adjustment (FPA) from [Pastén et al. \(2020\)](#), which provide monthly FPA estimates for NAICS industries at the 2- through 4-digit level. Each WIOD sector is mapped to a NAICS code using one of three approaches. First, when a WIOD sector corresponds directly to a single NAICS industry, it is matched to that code; if the exact code lacks FPA data, the next broader (truncated) NAICS level is used. Second, when a WIOD sector spans multiple NAICS industries that together constitute an exclusive 2-digit NAICS parent (e.g., Mining and quarrying covers all of NAICS 21), it is assigned the FPA of that 2-digit sector. Third, for the remaining WIOD sectors whose constituent NAICS industries do not exhaust a common 2-digit parent, FPA values are computed as sales-weighted averages of the available constituent-level estimates, with weights derived from gross output in the 2007 BEA Benchmark input–output tables.

As a robustness check, I also construct an alternative FPA vector from the CPI-based estimates of [Gautier et al. \(2024\)](#), which report monthly frequencies of price adjustment using consumer price microdata from eleven euro area countries over 2010–2019, classified by COICOP (Classification of Individual Consumption by Purpose). The mapping to WIOD sectors exploits the natural concordance between the European CPA (Classification of Products by Activity) and the NACE Rev. 2 classification underlying WIOD. In most cases, a WIOD sector corresponds to a single CPA code, and its FPA is computed as the average across all matched COICOP items. When a WIOD sector spans multiple CPA codes sharing a common NACE parent, these are averaged; for the single case where constituent CPA codes belong to different NACE parents, an expenditure-weighted average is used. I adopt the FPA measure that excludes sales and applies the frequency filter. This matching yields CPI-based FPA values for 46 of the 56 WIOD sectors.

Sectoral productivity shocks. To estimate the variance–covariance matrix of sectoral productivity shocks, I construct annual series of sectoral total factor productivity (TFP) growth using the WIOD Socio-Economic Accounts (SEA, 2016 release) for euro-area economies over the period 2000–2014. The SEA provide industry-level data on nominal values, price indices, and real volume indices for gross output and intermediate inputs, as well as employment.

Nominal cost shares of labor and intermediate inputs in gross output are defined as

$$s_{i,t}^L \equiv \frac{VA_{i,t}}{GO_{i,t}}, \quad s_{i,t}^M \equiv \frac{\Pi_{i,t}}{GO_{i,t}},$$

where $VA_{i,t}$ and $\Pi_{i,t}$ denote value added and intermediate inputs, respectively, and $GO_{i,t}$ is gross output at current basic prices.

Let GO_QI and Π_QI denote the real volume indices for gross output and intermediate inputs (2010 = 100), and EMP denote total employment. Annual growth rates of real variables are computed as log differences:

$$\Delta \log Y_{i,t} \equiv \log GO_QI_{i,t} - \log GO_QI_{i,t-1},$$

$$\Delta \log L_{i,t} \equiv \log EMP_{i,t} - \log EMP_{i,t-1},$$

$$\Delta \log M_{i,t} \equiv \log \Pi_QI_{i,t} - \log \Pi_QI_{i,t-1}.$$

Sectoral TFP growth is then obtained using a Törnqvist decomposition,

$$\Delta \log TFP_{i,t} = \Delta \log Y_{i,t} - \frac{1}{2} \left[(s_{i,t}^L + s_{i,t-1}^L) \Delta \log L_{i,t} + (s_{i,t}^M + s_{i,t-1}^M) \Delta \log M_{i,t} \right].$$

which measures the Solow residual as output growth net of input growth weighted by average nominal cost shares. The Solow residuals are winsorized at the 1st and 99th percentiles to reduce the influence of extreme outliers.

These resulting series $\Delta \log TFP_{i,t}$ form the basis for constructing the empirical variance–covariance matrix of sectoral productivity shocks, after interpolation to quarterly frequency and detrending.

D Beyond Income-Share Weighting

The quantitative analysis in Section 5 adopts income-share weighting as the benchmark. This section exercises Theorem 1 along two alternatives: a one-dimensional path tilting the

Pareto weights toward a single country, and a politically motivated weighting derived from the ECB voting rotation system.

Tilting toward a single country. I trace the optimal policy and welfare incidence along a one-dimensional family of Pareto weights interpolating between income-share weighting and unilateral weighting on Germany,

$$\kappa_c(t) = (1 - t)\chi_c + t \mathbf{e}_{\text{DEU}}, \quad t \in [0, 1].$$

By Theorem 1, $\zeta^*(\{\kappa_c\})$ and $\pi^*(\{\kappa_c\})$ are linear in $\{\kappa_c\}$, so the optimal policy at any interior t is a convex combination of the two endpoint policies. Figure D.1 reports two consequences of this path. Panel (a) plots Germany's weight in the optimal price index against its Pareto weight: both grow linearly in t but with different slopes, so the curves cross at $t \approx 0.23$. At the income-share benchmark ($t = 0$), the optimal index assigns Germany approximately 0.30, slightly above its Pareto weight of 0.27—this 0.03 excess maps directly to Figure A.1: Germany is one of five euro-area countries whose weight in the centralized optimal price index systematically exceeds its income share, owing to upstreamness. At $t = 1$, when the planner places full weight on Germany, the optimal index still assigns approximately 0.08 to non-German sectors, because Germany's welfare is exposed to foreign sectors through input–output linkages. Panel (b) plots expected welfare losses for Germany and the income-share-weighted average for the remaining nineteen member states: as the planner shifts weight toward Germany, Germany's expected welfare loss falls from 0.71 to 0.67 percent of steady-state consumption, while the rest of the union's average loss rises from 0.54 to 0.59 percent. The pattern directly visualizes the redistribution motive identified in Theorem 1: tilting Pareto weights toward a member state improves its welfare at the direct expense of the rest of the union.

An institutional benchmark: ECB voting weights. The ECB Governing Council comprises six Executive Board members with permanent voting rights and the twenty national central bank governors, who rotate through a fixed roster of voting slots. Governors are partitioned by economic size into two groups: five large economies (Germany, France, Italy, Spain, the Netherlands) share four voting slots; the remaining fifteen share eleven. The implied long-run voting frequencies are 4/5 for Group 1 and 11/15 for Group 2. To map this to Pareto weights, I allocate the six Executive Board votes by income shares—treating them as representing the union as a whole—and the fifteen country-governor votes according




Figure D.1: Optimal price-index weights and welfare losses along a Pareto-weight path. *Note:* The path is $\kappa_c(t) = (1 - t)\chi_c + t\mathbf{e}_{\text{DEU}}$. Panel (a) plots Germany's weight in the optimal price index (solid red) against its Pareto weight (dashed blue). Panel (b) plots expected welfare losses, in percent of steady-state consumption, for Germany (red) and the income-share-weighted average for the remaining nineteen member states (blue).

to the rotation frequencies. The resulting vector κ^{vote} over-represents small downstream economies relative to income shares but less aggressively than pure rotation weights.

Figure D.2 plots each country's policy-alignment loss against its unilateral inflation stance under the voting-weight regime; the dashed vertical line marks the union-wide inflation consensus $\pi^*(\{\kappa^{\text{vote}}\})$. The central message is robustness: even under Pareto weights that depart substantially from income shares, the consensus inflation rate barely moves from zero. Panel (a) reports the blended regime with the six Executive Board votes allocated by income shares; the consensus is $\pi^*(\{\kappa^{\text{vote}}\}) \approx 0.0005$. Panel (b) reports the pure-rotation regime in which the Executive Board votes are dropped entirely; the consensus shifts only marginally, to $\pi^*(\{\kappa^{\text{vote}}\}) \approx 0.0007$. Both values are an order of magnitude smaller than the typical unilateral inflation stance across member states, indicating that the income-share-weighted policy used in the main analysis closely approximates the optimal policy even under a politically motivated weighting scheme. The cross-sectional incidence is correspondingly similar to the centralized optimal regime: tail countries—Luxembourg, Ireland, and the Netherlands on the hawkish side, Cyprus and Malta on the dovish side—bear the largest policy-alignment losses, while large economies near the consensus, notably Germany, Italy, France, and Spain, incur negligible PAL.




Figure D.2: Policy-alignment loss against unilateral inflation stance under ECB voting weights. *Note:* Each dot is a euro-area country, with size proportional to its income share. The red curve is a quadratic fit centered at the union-wide inflation consensus $\pi^*(\{k^{\text{vote}}\})$, marked by the dashed vertical line. Panel (a) reports the blended specification with the six Executive Board votes allocated by income shares; Panel (b) reports the pure-rotation specification with the Executive Board votes dropped.

E Nested CES Economies

This section extends the Cobb–Douglas benchmark to a standard nested-CES economy. I use a calibrated production-network model to show that the inflation stance continues to operate primarily through the direct-incidence channel, and therefore remains closely related to the input–output multiplier.

Preferences. The consumption aggregator for country c is generalized to a CES form with elasticity of substitution σ_c :

$$\frac{C_c}{\bar{C}_c} = \left(\sum_{i=1}^N \beta_{ci}^{\frac{1}{\sigma_c}} \left(\frac{c_{ci}}{\bar{c}_{ci}} \right)^{\frac{\sigma_c-1}{\sigma_c}} \right)^{\frac{\sigma_c}{\sigma_c-1}}. \quad (\text{E-1})$$

Technology. The production function of industry i is generalized to a CES form with elasticity of substitution σ_i (omitting the firm subscript):

$$\frac{y_i}{\bar{y}_i} = z_i \left(\sum_{j=1}^N \omega_{ij}^{\frac{1}{\sigma_i}} \left(\frac{x_{ij}}{\bar{x}_{ij}} \right)^{\frac{\sigma_i-1}{\sigma_i}} + \sum_{c=1}^C \alpha_{ic}^{\frac{1}{\sigma_i}} \left(\frac{L_{ic}}{\bar{L}_{ic}} \right)^{\frac{\sigma_i-1}{\sigma_i}} \right)^{\frac{\sigma_i}{\sigma_i-1}}. \quad (\text{E-2})$$

Lemma 9. Under nested CES preferences with country-specific elasticities $\{\sigma_c\}_{c=1}^C$ and CES production with sector-specific elasticities $\{\sigma_j\}_{j=1}^N$, the first-order changes in country-level income shares satisfy

$$\hat{\lambda}_c = \sum_{j=1}^N \tilde{\Gamma}_{jc} \log \mu_j, \quad \tilde{\Gamma}_{jc} = \sum_{f=1}^C (\lambda_c^*)^{-1} \tilde{Q}_{cf} \lambda_j^* \left(\Phi_{jf} - \Psi_{jf} + \Delta_{jf} - \sum_{g=1}^C \frac{\Delta_{gf} \Phi_{jg}}{1 + \eta_g} \right), \quad (\text{E-3})$$

where $\tilde{Q} = (I - \tilde{\Lambda}')|_{\mathcal{S}}^{-1}$ denotes the inverse of $I - \tilde{\Lambda}'$ restricted to $\mathcal{S} = \{\mathbf{x} \in \mathbb{R}^C : \mathbf{1}'\mathbf{x} = 0\}$, with $\tilde{\Lambda}(c, f) = \Lambda_f^c - \Delta_{cf}$, and the CES substitution wedge Δ_{kf} is defined by

$$\Delta_{kf} \equiv (\lambda_k^*)^{-1} \sum_{j \in N \cup C} \lambda_j^* (\sigma_j - 1) \text{Cov}_{\Omega_{(j,:)} } (\Psi_{(:,k)}, \Psi_{(:,f)}), \quad (\text{E-4})$$

where, for consumption nests, I adopt the convention $\lambda_c^* \equiv \lambda_c^*$ and $\Omega_{(c,:)} \equiv \beta'_c$.

Under Cobb–Douglas ($\sigma_j = \sigma_c = 1$ for all j and c), $\Delta_{kf} = 0$ and $\tilde{\Lambda} = \Lambda$, so equation (E-3) reduces to equation (8) of Lemma 1.

Proof. Expanding the consumer- c Domar weight λ_i^c to first order yields

$$\begin{aligned} \lambda_i^c - \lambda_i^{c*} &= - \sum_{j=1}^N \lambda_j^{c*} (\Psi_{ji} - \iota_{ji}) \log \mu_j \\ &+ \sum_{j=1}^N \lambda_j^{c*} (\sigma_j - 1) \text{Cov}_{\Omega_{(j,:)} } (\log \mathbf{p}, \Psi_{(:,i)}) + (\sigma_c - 1) \text{Cov}_{\beta_c} (\log \mathbf{p}, \Psi_{(:,i)}) + o(\|\xi\|). \end{aligned}$$

Applying the product rule to $\Lambda_f = \sum_c \chi_c \Lambda_f^c$ gives $\Lambda_f - \Lambda_f^* = \sum_c \chi_c^* (\Lambda_f^c - \Lambda_f^{c*}) + \sum_c \Lambda_f^{c*} (\chi_c - \chi_c^*) + o(\|\xi\|)$. Substituting the consumer- c expansion above at a factor index f (where $\iota_{jf} = 0$) and using $\lambda_j^* = \sum_c \chi_c^* \lambda_j^{c*}$ yields

$$\begin{aligned} \Lambda_f - \Lambda_f^* &= \sum_{c=1}^C \Lambda_f^{c*} (\chi_c - \chi_c^*) - \sum_{j=1}^N \lambda_j^* \Psi_{jf} \log \mu_j \\ &+ \sum_{j \in N \cup C} \lambda_j^* (\sigma_j - 1) \text{Cov}_{\Omega_{(j,:)} } (\log \mathbf{p}, \Psi_{(:,f)}) + o(\|\xi\|). \end{aligned} \quad (\text{E-5})$$

The cost-based price system, combined with the first-order condition for labor-supply

(B-9), gives

$$\begin{aligned}
\log p &= \log \mu + \Omega \log p + \alpha d \log w, \\
&= \sum_k \Psi_{(:,k)} \log \mu_k + \sum_g \Psi_{(:,g)} d \log w_g, \\
&= \sum_k \Psi_{(:,k)} \log \mu_k + \sum_g \Psi_{(:,g)} (\log m + d \log \chi_g + \frac{1}{\eta_g} d \log L_g), \\
&= \sum_k \Psi_{(:,k)} \log \mu_k + \sum_g \Psi_{(:,g)} (d \log \chi_g + \frac{1}{\eta_g} d \log L_g) + \log m,
\end{aligned}$$

where $\log m$ factors out of the inner sum because $\sum_g \Psi_{ig} = 1$ under constant returns to scale production functions.

Substituting into the covariance terms of (E-5) and using the definition (E-4) of Δ_{kf} :

$$\begin{aligned}
\sum_{j \in N \cup C} \lambda_j^* (\sigma_j - 1) \text{Cov}_{\Omega_{(j,:)}}(\log p, \Psi_{(:,f)}) &= \sum_{k=1}^N \lambda_k^* \Delta_{kf} \log \mu_k \\
&\quad + \sum_{g=1}^C \chi_g^* \Delta_{gf} (d \log \chi_g + \frac{1}{\eta_g} d \log L_g), \quad (\text{E-6})
\end{aligned}$$

where the $\log m$ term drops out because $\text{Cov}_{\Omega_{(j,:)}}(\cdot, \Psi_{(:,f)})$ annihilates constants.

Substituting the first-order approximation of labor supply (7) yields

$$\chi_g^* (d \log \chi_g + \frac{1}{\eta_g} d \log L_g) = (\chi_g - \chi_g^*) - \frac{1}{1 + \eta_g} \sum_{k=1}^N \lambda_k^* \Phi_{kg} \log \mu_k. \quad (\text{E-7})$$

Finally, the labor-wedge identity (B-3) gives $\Lambda_f - \Lambda_f^* = (\chi_f - \chi_f^*) - \sum_k \lambda_k^* \Phi_{kf} \log \mu_k$. Combining this with (E-5), (E-6), and (E-7) and collecting terms gives

$$\chi_f - \chi_f^* = \sum_{c=1}^C (\Lambda_f^{c*} - \Delta_{cf}) (\chi_c - \chi_c^*) + \sum_{k=1}^N \lambda_k^* \left(\Phi_{kf} - \Psi_{kf} + \Delta_{kf} - \sum_{g=1}^C \frac{\Delta_{gf} \Phi_{kg}}{1 + \eta_g} \right) \log \mu_k. \quad (\text{E-8})$$

□

Corollary 6 (Nested-CES extension of Propositions 1 and 2). Under a nested CES structure with elasticities $(\{\sigma_c\}_{c=1}^C, \{\sigma_j\}_{j=1}^N)$, Propositions 1 and 2 continue to hold after replacing the income-share Jacobian Γ with $\tilde{\Gamma}$ of Lemma 9. In particular, the direct-incidence index

$\mathcal{J}_{ic}^{\text{DI}} = \frac{\lambda_i \Phi_{ic}}{\lambda_c} - \lambda_i^c$ is unchanged.

Proof. The nested-CES extension affects the first-order system through the income-share response. Given the calibrated steady-state network, substitution elasticities enter through cost-share adjustment, summarized by the modified income-share Jacobian $\tilde{\Gamma}$. The direct-incidence index $\mathcal{J}^{\text{DI}}$ is not directly modified by these substitution terms: it is computed from steady-state Domar weights and ownership shares. $\square$

Implementation. I bring the nested-CES economy in equations (E-1)–(E-2) to WIOD data as follows. For each producer and each household, I introduce *importer pseudo-industries* that aggregate same-industry varieties across source countries through an Armington nest with elasticity σ_{arm} . These importer pseudo-industries have zero value added and flexible prices. Producers then substitute, with elasticity σ_p , between domestically sourced value added and intermediate input composites purchased from importer pseudo-industries. Households substitute across consumption composites with elasticity σ_c .

Figure E.1 verifies that the mechanism identified in the Cobb–Douglas benchmark continues to operate under nested CES. The direct-incidence component accounts for most of the cross-country variation in the inflation bias of unilateral optimal policy across a wide range of elasticities of substitution. By Proposition 4, the unilateral inflation stance is proportional to the inflation bias of unilateral optimal policy. Hence the direct-incidence channel remains the main driver of the inflation stance after Cobb–Douglas consumption and production functions are replaced by a standard nested-CES structure.

By Lemma 9 and Corollary 6, nested CES substitution modifies the income-share response through $\tilde{\Gamma}$ but leaves the direct-incidence index $\mathcal{J}^{\text{DI}}$ unchanged. Because $\mathcal{J}^{\text{DI}}$ depends only on steady-state Domar weights, consumption shares, and the ownership structure, it is invariant to the elasticities of substitution. The input–output multiplier therefore remains the key object organizing cross-country variation in inflation stances under nested CES.

F Global Economy under Dominant Currency Pricing

This section extends the analysis to a global economy under dominant currency pricing (DCP), with a quantitative focus on the optimal monetary policy of the dominant-currency country and the resulting policy-alignment losses across non-dominant-currency economies.




Figure E.1: Direct incidence dominates the allocative-efficiency response under nested CES. *Note:* Each panel plots, for the twenty euro-area members, the inflation bias of unilateral optimal policy defined in (20) (horizontal axis) against its direct-incidence component, decomposed analogously to equation (10), on the vertical axis. The four panels span the production elasticity $\sigma_p \in \{0.5, 2.0\}$ across columns and the Armington elasticity $\sigma_{arm} \in \{2, 4\}$ across rows, with the consumption elasticity fixed at $\sigma_c = 0.9$. In each panel, the 45-degree line marks coincidence.

Irrelevance of exchange rates under dominant currency pricing. Assuming that all goods are tradable, DCP implies that nominal rigidities operate exclusively in dominant-currency prices. Domestic exchange rate movements then do not affect ex post markups or relative prices, rendering non-dominant-currency economies locally equivalent to members of a single-currency area in terms of real allocations. Exchange rates are thus irrelevant for resource allocation under this assumption.

Optimal monetary policy of the dominant-currency country. Given this allocation equivalence, the optimal monetary policy problem is effectively centralized at the level of the dominant-currency country, the United States. The optimal policy of the dominant-currency country, which is given by the price-index targeting rule $\sum_{i=1}^N \zeta_i^*(\mathbf{e}_{US}) \log p_i = \pi^*(\mathbf{e}_{US})$, governs global nominal rigidities and, through them, shapes real allocations worldwide. In this sense, it serves as a de facto global benchmark, analogous to the union-wide optimal policy in a monetary union. Given this benchmark price-index regime, a country's unilateral inflation stance is

$$\pi_c = \frac{[\zeta^*(\mathbf{e}_{US})]' \varrho^m}{[\zeta^*(\mathbf{e}_c)]' \varrho^m} \pi^*(\mathbf{e}_c).$$

As in the currency-union case, Figure F.1 documents a negative relationship between unilateral inflation stances and upstreamness in the global production network. Consistent with its downstream position, the optimal monetary policy of the United States is inflationary. In addition, domestic sectors receive approximately 93.2% of the weight in the optimal price index, indicating that U.S. monetary policy primarily reflects domestic production structures while anchoring global nominal rigidities.

Policy-alignment losses across non-dominant-currency economies. Under DCP, policy-alignment loss is defined relative to the dominant-currency country's optimal monetary policy rather than a union-wide consensus. For country c ,

$$\text{PAL}_c \equiv \mathbb{L}_c(\zeta^*(\mathbf{e}_{US}), \pi^*(\mathbf{e}_{US})) - \mathbb{L}_c(\zeta^*(\mathbf{e}_c), \pi^*(\mathbf{e}_c)) \geq 0,$$

which captures the welfare cost of operating under the dominant-currency country's policy rather than the country's own unilateral optimum.

Figure F.2 shows that policy-alignment losses increase with the distance between a country's unilateral inflation stance and the dominant-currency country's optimal inflation rate. Countries such as Brazil (BRA), France (FRA), Hungary (HUN), Japan (JPN), and




Figure F.1: Unilateral inflation stance and upstreamness in the global production network. *Note:* The horizontal axis reports each country's unilateral inflation stance computed under the price-index regime implied by the dominant-currency country's optimal policy. The vertical axis measures country-level upstreamness. The red line shows a fitted linear trend.

the United Kingdom (GBR), whose production network structures are closer to that of the dominant-currency economy, experience smaller policy-alignment losses.




Figure F.2: Inflation-stance deviation and policy-alignment loss under dominant currency pricing. *Note:* The horizontal axis measures the distance between the country's unilateral inflation stance and the dominant-currency country's optimal inflation rate. The vertical axis reports the policy-alignment loss as a percentage of steady-state consumption. The red curve shows a quadratic fit with an intercept but no linear term.