Taxing the Many to Subsidize the Few: The Political Economy of Industrial Policy*

Jesús Fernández-Villaverde† Yang Yu‡ Francesco Zanetti♭

August 4, 2026

Abstract

This paper studies the conditions for the democratic enactment of industrial policy and how they compare with the welfare criterion of a social planner. We develop a two-sector model that integrates learning-by-doing, inter-sectoral spillovers, technological synergies, and labor-market monopsony. The set of policies supported by the taxed majority is strictly smaller than the set of welfare-improving policies, opening the possibility that welfare-improving policies may fail democratically. Enactment requires strong inter-sectoral learning spillovers, low synergy costs, and a small subsidized sector; monopsony power broadens support by channeling rents to shareholders through diversified equity holdings. Using U.S. firm-level customer-supplier data and a novel IV strategy, we identify a small set of industries—transportation equipment, computer and electronic products, primary metals, and chemicals—for which industrial policy is politically feasible, aligning with the historical U.S. policy targets over the past seven decades.

Keywords: Industrial policy, political economy, learning spillovers, synergies, monopsony power, voting.

JEL Classification: D72, L52, O25, O38.

*We thank Jing Cai and participants at SED Athens, Renmin University of China, and Wuhan University for helpful comments and suggestions. PRELIMINARY DRAFT, please do not circulate without permission from the authors.

†Jesús Fernández-Villaverde: University of Pennsylvania, US jesusfv@econ.upenn.edu; ‡Yang Yu: Shanghai Jiao Tong University, China. yu.yang.econ@sjtu.edu.cn; ♭Francesco Zanetti: University of Oxford, U.K. francesco.zanetti@economics.ox.ac.uk.

1 Introduction

In the United States, industrial policy has repeatedly returned to a remarkably stable set of industries. Transportation equipment has been supported by Department of Defense and NASA procurement since the 1950s, and remains central to interventions ranging from the 2009 auto bailouts to the electric-vehicle credits in the 2022 Inflation Reduction Act. Semiconductors were developed in part through military contracts authorized under the Defense Production Act of 1950, and are now the explicit target of the $52.7 billion CHIPS and Science Act of 2022. Primary metals have long been protected through the Buy American Act of 1933, Section 232 national-security tariffs, and successive executive orders, most recently in 2025. Across Democratic and Republican administrations alike, the United States has repeatedly intervened in the same narrow slate of manufacturing sectors while leaving most others largely untouched. Why are these industries politically viable targets of industrial policy when most others are not? And are the resulting interventions welfare improving?

These questions are difficult because industrial policy is inherently redistributive and it requires democratic support from voters in non-beneficiary sectors who finance subsidies to beneficiary sectors (Acemoglu and Robinson, 2008). Voters in the non-beneficiary sectors bear the costs directly, while benefiting only indirectly from any productivity gains that the subsidies may eventually generate. Whether welfare-improving industrial policy secures democratic support therefore depends on how such gains propagate through the economy—via learning-by-doing in subsidized sectors, technological spillovers across sectors, and the allocation of rents in imperfectly competitive markets.

This paper develops a unified framework for analyzing these issues and disciplines it using new data. We build a two-sector, two-period model that incorporates four mechanisms central to the welfare and political economy of industrial policy: within-sector learning-by-doing, intersectoral learning spillovers, technological synergies that depend on relative productivity, and labor-market monopsony power. Industrial policy is modeled as a budget-balanced transfer to the beneficiary sector, financed by a tax on the non-beneficiary sector. We derive the conditions under which workers in each sector and a utilitarian social planner support the policy. We then estimate the structural parameters governing these mechanisms using newly assembled firm-level customer-supplier data from Compustat and FactSet, covering more than 360,000 customer-supplier pairs and 18,000 U.S. firms from 1976 to 2020. To identify the synergy and spillover parameters, we introduce two instruments—pre-relationship flows of senior managers between firms and pre-relationship size-distance between trading partners—that exploit exogenous variation in the initial productivity gap at the formation of the trading relationship.

Three findings stand out and together reframe the debate around industrial policy. First, from a welfare perspective, we show that industrial policy is justified when the beneficiary sector generates sufficiently strong inter-sectoral learning spillovers, exhibits low synergy costs from productivity gaps that impair technological synergies, and is small relative to the rest of the economy. However, the political support condition for workers in the non-beneficiary sector is more demanding as it requires resulting spillovers to be large enough to offset not only the cost of the sectoral productivity gap but also the direct tax burden on the non-beneficiary sector. A non-trivial range of industrial policies is therefore welfare-improving yet democratically infeasible, and a separate range is politically attractive yet socially suboptimal.

Second, labor-market monopsony power in the subsidized sector—typically viewed as a distortion to be corrected—acts in our framework as a political lubricant: by channeling rents to diversified shareholders, it partially offsets the tax burden borne by workers outside the sector and broadens the democratic support.[1] This result stands in sharp contrast to the standard result that treats monopsony and other similar market powers as a market failure to be eliminated.

Third, when we map our industry-level estimates to three-digit NAICS sectors, the small set of industrial policies that enhance welfare and simultaneously garner democratic support—Fabricated Metals, Transportation Equipment, Machinery, Computer and Electronic Products, Primary Metals, and Chemicals—coincides closely with the historical targets of U.S. industrial policy over the past seven decades. The historical pattern of U.S. industrial policy, in other words, closely resembles the set of sectors our model identifies as both welfare improving and politically viable.

Our identification strategy exploits firm-level heterogeneity within narrowly defined sectors. Within a given three-digit NAICS industry, firms can be ranked by productivity and matched with customers and suppliers of different productivity ranks. The central identification challenge is that the productivity gap between a firm and its trading partners is endogenous to unobserved firm characteristics—organizational capital, technological sophistication—that also influence performance. We address this with two instruments that predict the productivity gap at the time the trading relationship was first formed. Pre-relationship flows of senior managers between

[1]We use labor market monopsony as a tractable stand-in for market power more broadly: any force that channels a share of profits distributed to shareholders—whether it originates in the labor market or, alternatively, in product-market markups—generates the same political-economy mechanism we study below.

firms increase the likelihood of matches between firms of similar productivity, providing a plausibly exogenous source of assortative matching. Pre-relationship differences in firm size capture a structural trade-off in which large, established firms accept less productive partners in exchange for scale. Standard diagnostics verify the validity of our IV strategy.

Our study makes one contribution—a quantitative political-economy theory of industrial policy that connects to data and explain history—built on three pillars. The first is a formal characterization of when democratic voting yields welfare-improving industrial policy. We embed heterogeneous, sector-specific workers in an explicit voting framework, derive a closed-form condition under which the taxed majority supports the policy, prove that this condition is strictly stronger than the social planner's welfare criterion, and show how monopsony rents systematically widen the political-feasibility region. The second pillar is an identification strategy that uses firm-to-firm trading-relationship data to estimate the spillover and synergy parameters that govern the welfare and political analysis. The combination of customer-supplier panel data, and pre-relationship instruments yields causal estimates of both quantities at the firm and industry level. The third pillar is the mapping from theory to empirical evidence: by computing the conditions for democratic support at the three-digit NAICS level, we identify the small set of welfare-improving and politically feasible interventions and show that this set tracks remarkably well seven decades of U.S. industrial policy. The three pillars together show not only when industrial policy is welfare-improving in principle, but under what conditions democratic majorities enact it in practice, and show that those conditions have been met in U.S. history.

The paper builds on and is distinguished from three influential recent contributions on industrial policy. Bartelme et al. (2025) provide the textbook welfare case for industrial policy in a quantitative trade model and estimate sector-specific learning rates from international trade variation. We share their focus on the structural foundations of industrial policy but ask a complementary question: under what conditions is a welfare-improving industrial policy politically feasible in a democracy that must finance subsidies through visible taxation? The gap between welfare-improving and politically feasible policies is the central object of our analysis. Liu (2019) characterizes optimal sectoral subsidies in a production-network model under a benevolent planner with lump-sum taxes and treats market power as a distortion to be corrected. We retain the production-network intuition but model industrial policy as a budget-balanced tax-and-subsidy scheme that imposes a visible cost on the taxed majority and requires democratic support. We show that labor-market monopsony generates rents whose redistribution garners democratic support that makes the policy possible. Juhász and Lane (2024) provide a comprehensive assessment of the political economy of industrial policy and emphasize state capacity and time-inconsistency as central frictions; we provide a formal, quantitative framework that derives the voting condition, identifies the structural parameters from firm-level data, and maps the resulting support sets to observed U.S. industries. Our work also links to the seminal contribution of Fernandez and Rodrik (1991) on resistance to reform: where they show that ex-ante uncertainty over individual outcomes can block welfare-improving change, we show that even under full information and transparent redistribution, the political feasibility is more demanding than the welfare criterion, with implications that can be quantified industry by industry.

We also contribute to the empirical literature on learning-by-doing and productivity spillovers and we contextualize the importance of these forces for industrial policy. Building on seminal work documenting within-industry learning (Irwin and Klenow, 1994; Bartelme et al., 2025) and firm-level spillovers (Barrios and Strobl, 2004), we provide novel estimates of inter-sectoral learning using firm-to-firm trading relationships. Critically, we argue that within-sector learning alone cannot justify industrial policy since subsidizing one sector inevitably impedes learning opportunities in taxed sectors. The potential for inter-sectoral spillovers becomes the decisive economic rationale, and our instrumental variables strategy—using manager transitions and size-distance trade-offs—provides new and powerful identification of these spillover effects.

Finally, we formalize the economic case for industrial policy rooted in learning externalities and productivity spillovers, a theme closely linked to development economics (Chenery, 1960; Pack and Westphal, 1994). Like Stiglitz et al. (2013), we emphasize spillovers as a core market failure justifying intervention. Our contribution is to embed empirically-grounded estimates of productivity transmission (Amiti et al., 2023; Irwin and Klenow, 1994; Kong and Dhyne, 2025) within a formal political economy model that we use to interpret industrial policy in the U.S, demonstrating how these spillovers interact with voter preferences to determine both welfare implications and political viability.

The remainder of the paper proceeds as follows. Section 2 develops the theoretical model. Section 2.5 derives the voting conditions and the social planner's welfare criterion and characterizes the parameter space in which industrial policy is welfare-improving, politically feasible, or both. Section 5 presents the empirical strategy, the customer-supplier data, and the instrumental-variables approach to estimating synergies and inter-sectoral learning. Section 6 estimates industry-specific parameters and uses the framework to study the enactment of industrial policy in U.S. history. Section 7 concludes.

2 Model

We develop a two-sector, two-period model to analyze the welfare and political economy implications of industrial policy. The model incorporates four mechanisms central to both the economic rationale for, and the democratic feasibility of, such interventions: within-sector learning-by-doing, inter-sectoral learning spillovers, technological synergies that depend on relative sectoral productivity, and labor-market monopsony power. Industrial policy is embedded in an explicit voting framework in which households in each sector decide whether to authorize a subsidy to one sector, financed by a tax on the other sector under a balanced-budget constraint.

The timing of the model is as follows. In period 1, the government proposes an industrial policy consisting of a sectoral tax and a sectoral subsidy, which is put to a vote by households in both sectors. If ratified (i.e., supported by the majority of households), the tax and subsidy rates are set, production takes place. Capital input in period 1 determines sectoral productivity in period 2 through the learning mechanisms described below. Period 2 has no taxes or subsidies; firms produce using the productivity levels that result from period-1, and all output is consumed.

2.1 Production Technology

The economy operates over two periods, $t \in \{1, 2\}$, and has two sectors, $i \in \{A, B\}$. Each sector comprises a continuum of homogeneous firms indexed by $j \in [0, N_i]$, where N_i is the measure of firms in sector i.

Intermediate goods production. The output of firm j in sector i follows a Cobb-Douglas production function:

$$y_{i,j,t} = z_{i,t} l_{i,j,t}^{1-\alpha} k_{i,j,t}^{\alpha}, \tag{1}$$

where $z_{i,t}$ is sectoral productivity (common to all firms in sector i), $l_{i,j,t}$ is the labor input of firm j at time t, $k_{i,j,t}$ its capital input, and $\alpha \in (0, 1)$ is the capital share. Both sectors have equal productivity in the first period: $z_{A,1} = z_{B,1} = z_1$. Productivity in period 2, $z_{i,2}$, accumulates through the learning mechanisms described below.

Intermediate goods firm's problem and monopsony power. An intermediate goods firm in sector i chooses labor and capital to maximize after-tax profit:

$$\max_{l_{i,j,1}, k_{i,j,1}} (1 - \tau_i) p_{i,1} y_{i,j,1} - l_{i,j,1} w_{i,j,1} - k_{i,j,1},$$

subject to the production function (1). The firm exercises monopsony power in the labor market, which we capture in reduced form through the wage-setting rule:

$$w_{i,j,1} = (1-\kappa)(1-\tau_i)p_{i,1} \cdot \mathrm{MPL}_{i,j,1},$$

where $\mathrm{MPL}_{i,j,1} = (1-\alpha)z_1 k_{i,j,1}^{\alpha} l_{i,j,1}^{-\alpha}$ denotes the marginal product of labor. The parameter $\kappa \in [0,1]$ measures the degree of monopsony power, interpreted as the wage markdown below the marginal revenue product of labor. When $\kappa = 0$, the labor market is competitive and workers are paid their full marginal revenue product; when $\kappa > 0$, the firm pays a fraction $(1-\kappa)$ of the marginal revenue product and retains the residual κ as monopsony rents, which constitute firm profit. A higher κ thus corresponds to a lower labor income share and a higher profit share. Since aggregate labor supply is fixed in the baseline model, this monopsony power is purely redistributive rather than distortionary at the aggregate level. In Appendix E, we relax this assumption by allowing labor supply to be elastic, which introduces a deadweight loss from monopsony power; the core political economy insight of the paper is unchanged.

Labor is immobile across sectors, and each worker supplies one unit of labor inelastically. The population of sector i is N_i, which coincides with the measure of firms.[2] In equilibrium, each firm hires exactly one worker: $l_{i,j,1} = 1$. Firms purchase capital from the competitive final goods market, where capital goods are converted from the final good one-for-one. The final good is the numeraire, so capital good price is one. Capital constitute a cost to firms and does not accrue to households as rental income.[3] The first-order condition for capital is:

$$k_{i,j,1} = \alpha(1-\tau_i)p_{i,1}y_{i,j,1}, \tag{2}$$

which shows that a higher tax rate reduces the firm's capital input.

Final goods production. A competitive final goods sector combines the outputs of the two intermediate sectors:

$$y_t = e^{x-\gamma|\ln(z_{A,t})-\ln(z_{B,t})|} y_{A,t}^{\chi_A} y_{B,t}^{\chi_B}, \tag{3}$$

or in log form:

$$\ln(y_t) = x - \gamma|\ln(z_{A,t}) - \ln(z_{B,t})| + \chi_A \ln(y_{A,t}) + \chi_B \ln(y_{B,t}), \tag{4}$$

where χ_A and χ_B are the factor shares of the intermediate goods, with $\chi_A + \chi_B = 1$, x is the

[2]Throughout the paper, N_i denotes both the measure of firms in sector i and the sector's population. Below, we additionally impose $N_i = \chi_i$, where χ_i is the sector's factor share. We discuss this normalization when it is introduced.

[3]Alternatively, one can think of capital supplied through a competitive rental market by agents outside the two sectors who hold the economy's capital stock.

exogenous component of aggregate TFP (assumed constant), and $\gamma \in [0, 1]$ captures the degree of synergies between sectors.[4] As mentioned, the final good serves as the numeraire and its price p_t is one.

The synergy parameter γ measures the extent to which productivity dispersion between sectors reduces aggregate final output. When $\gamma = 0$, sectors are fully independent and productivity gaps impose no aggregate cost. When $\gamma > 0$, the term $-\gamma|\ln(z_{A,t}) - \ln(z_{B,t})|$ penalizes disparities in sectoral productivity, capturing complementarities or coordination frictions that arise when one sector's technology lags substantially behind the other's. This reduced-form specification is motivated by a growing body of evidence that productivity mismatches between sectors generate efficiency costs at the aggregate level.[5]

Final goods firm's problem. Final goods producers operate in perfectly competitive markets and take prices as given. Their profit maximization, given the Cobb-Douglas production function (3), yields the standard factor share condition:

$$\chi_i y_t = y_{i,t} p_{i,t}. \tag{5}$$

where $p_{i,t}$ denotes the price of sector i's output. This condition states that expenditure on each intermediate input equals its factor share times total revenue.

Sectoral aggregation. All firms within sector i face a common sectoral productivity $z_{i,t}$, output price $p_{i,t}$, and tax rate τ_i, and each hires one unit of labor in equilibrium ($l_{i,j,t} = 1$). Given these common objects, the symmetric first-order condition for capital (2) and the production function (1) yield a unique solution of $(k_{i,j,t}, y_{i,j,t})$, implying that all firms within the sector choose identical capital inputs and produce identical output. Sectoral output is obtained by aggregating across firms:

$$y_{i,t} = \int_0^{N_i} y_{i,j,t}\, dj = z_{i,t} l_{i,t}^{1-\alpha} k_{i,t}^{\alpha}, \tag{6}$$

where $l_{i,t} = \int_0^{N_i} l_{i,j,t}\, dj = N_i$ and $k_{i,t} = \int_0^{N_i} k_{i,j,t}\, dj = N_i k_{i,j,t}$.

For tractability, we impose the normalization $N_i = \chi_i$, equating each sector's population—the number of firms and workers—to its income share in final goods production. This normalization

[4] An alternative specification would nest sectoral productivities in a CES aggregator, with the elasticity of substitution governing the degree of complementarity. We instead adopt the reduced-form penalty $e^{-\gamma|\ln(z_{A,t})-\ln(z_{B,t})|}$ because it maps directly onto our empirical strategy (Section 5 and Appendix D). A CES aggregator would instead require structurally estimating an elasticity of substitution between sector-level productivities, a more demanding exercise given our firm-to-firm trading data.

[5] See Fernández-Villaverde et al. (2021, 2025a,b) and Jones (2011) for theoretical and empirical analyses of inter-sectoral synergies. See Acemoglu et al. (2024); Jones and Tonetti (2025) for a recent discussion of how inter-sectoral synergies may constrain aggregate productivity growth.

ensures that, in the absence of taxes, both sectors attain identical per capita income in the decentralized equilibrium, providing a natural reference point in which workers have no incentive to move across sectors even if mobility were permitted. A formal proof of this result is provided in Appendix A. Without loss of generality, we designate sector A as the smaller sector ($\chi_A < \chi_B$). Sector A represents a nascent or targeted industry that may warrant industrial policy support, while sector B represents the rest of the economy, which would bear the tax burden of funding such a policy.

From the capital optimality condition (2) and the factor share condition (5), sectoral capital input is: $k_{i,t} = N_i k_{i,j,t} = N_i \alpha(1-\tau_i) p_{i,t} y_{i,j,t}$. Using $y_{i,t} = N_i y_{i,j,t}$, $N_i = \chi_i$, and $p_{i,t} y_{i,t} = \chi_i y_t$ from equation (5), this simplifies to:

$$k_{i,t} = \alpha(1-\tau_i)\chi_i y_t. \tag{7}$$

Substituting equation (7) into equation (6) and using $l_{i,t} = \chi_i$ gives sectoral output as a function of aggregate output and sectoral productivity:

$$y_{i,t} = z_{i,t} \left[\alpha(1-\tau_i) y_t\right]^{\alpha} \chi_i. \tag{8}$$

2.2 Learning by Doing and Spillovers

Productivity in period 2 evolves through two channels, both of which operate as sector-level externalities: within-sector learning-by-doing and inter-sectoral learning spillovers.

Within-Sector Learning-by-Doing. Higher per capita capital in period 1 raises productivity in period 2 through learning-by-doing:

$$\ln(\widehat{z}_{i,2}) = \ln(z_1) + \psi \ln(k_{i,1}/N_i), \tag{9}$$

where $\psi \geq 0$ is the elasticity of productivity with respect to per capita capital—the within-sector learning rate. The variable $\widehat{z}_{i,2}$ denotes the productivity that sector i would attain through its own learning-by-doing in the absence of spillovers from the other sector. This mechanism formalizes the classic infant-industry argument that production experience raises productivity over time.

For parsimony and because sector-specific learning rates cannot be reliably estimated with available data, our baseline model imposes symmetric learning-by-doing across the two sectors ($\psi_A = \psi_B$). In Appendix B, we relax this restriction and analyze how sector-specific rates ($\psi_A \neq \psi_B$) affect our results. Allowing the subsidized sector to learn faster renders the conditions for democratic support less restrictive, as the taxed majority anticipates larger future spillovers. This quantitative adjustment preserves the qualitative logic of our findings.

Inter-Sectoral Learning Spillovers. A sector can raise its productivity by learning from the other sector when that sector is more productive:

$$\ln(z_{i,2}) = \ln(\widehat{z}_{i,2}) + \phi \max\{\ln(\widehat{z}_{-i,2}) - \ln(\widehat{z}_{i,2}), 0\}, \tag{10}$$

where $\phi \in [0, 1]$ is the inter-sectoral learning rate and $z_{i,2}$ denotes the final productivity of sector i in period 2, net of spillovers.[6] When the other sector is more productive ($\widehat{z}_{-i,2} > \widehat{z}_{i,2}$), sector i closes a fraction ϕ of the productivity gap through knowledge spillovers. When $\phi = 0$, spillovers are absent; when $\phi = 1$, sector i fully converges to the productivity level of the more advanced sector. This specification is consistent with a well-documented empirical pattern: firms learn from more productive trading partners through multiple channels, including technology transfer, knowledge diffusion, exposure to superior management practices, and the stronger innovation incentives induced by stricter technical and quality standards imposed by high-productivity partners.

Because productivity is determined at the sectoral level while individual firms are atomistic, both learning mechanisms operate as sector-level externalities that lie beyond the control of any single firm. This externality provides the potential efficiency rationale for industrial policy.

2.3 Households

Household income consists of two components: a sector-specific wage and an equal per capita share of aggregate corporate profits. Each household holds a fully diversified portfolio of equity in all intermediate goods firms in both sectors, so dividend income is identical across households irrespective of their sector of employment.[7] Final goods firms earn zero profits under perfect competition, so their ownership is irrelevant.

The utility function of sector i's representative household is

$$U_{i,t} = \ln(c_{i,t}),$$

where $c_{i,t}$ is per capita consumption in sector i at time t. We abstract from labor disutility because labor supply is fixed and inelastic. Households maximize lifetime utility over the two periods, $U_i = u(c_{i,1}) + \beta u(c_{i,2})$, where $\beta \in (0, 1)$ is the discount factor.

To obtain closed-form voting conditions, we abstract from household saving or borrowing decisions; households consume their period income each period. In Appendix C, we show

[6] Following the literature, we assume that the rate of technology spillover is independent of sector size.

[7] We abstract from heterogeneity in equity holdings. This assumption is consistent with our political-economy interpretation: the households most influential and most responsive to industrial policy are those with financial stakes in its outcomes (Leighley and Nagler, 2013). Our representative household is thus best understood as the politically influential segment of each sector.

that allowing for consumption smoothing does not alter our main results: when log-linearized around the no-policy equilibrium, the conditions in the consumption-smoothing case coincide with those in the benchmark model.[8] The consumption of a representative household in sector i is:

$$c_{i,t} = w_{i,t} + \frac{1}{N_A + N_B} \left(\sum_{j \in A} \pi_{A,j,t} + \sum_{j \in B} \pi_{B,j,t} \right), \tag{11}$$

where $w_{i,t}$ is the sectoral wage and $\pi_{i,j,t}$ is the period-t profit of firm j in sector i. The term in parentheses is total dividend income, which is divided equally across the total population $N_A + N_B$. Labor income is sector-specific, but dividend income is pooled and equally distributed across all households.

2.4 Government and Industrial Policy

The government levies an output tax on one sector and uses the proceeds to finance an output subsidy to the other. Sector i's period-1 output is taxed (or subsidized, if $\tau_i < 0$) at rate τ_i. The government operates under a balanced-budget constraint, so that total tax revenue equals total subsidy expenditure in period 1:

$$\tau_A y_{A,1} p_{A,1} + \tau_B y_{B,1} p_{B,1} = 0, \tag{12}$$

where $p_{i,1}$ is the price of sector i's output in period 1. No taxes or subsidies apply in period 2: the intervention is temporary and aims solely at jumpstarting learning in sector A.

Equation (12) makes transparent the redistributive nature of industrial policy in our framework: every dollar of subsidy to one sector must be financed by a dollar of tax revenue from the other. We focus on a policy that subsidizes the smaller sector A by taxing the larger sector B:

$$\tau_A < 0, \quad \tau_B > 0.$$

The subsidy to sector A stimulates capital usage and production in that sector in period 1, raising its within-sector learning-by-doing and potentially generating spillovers to sector B in period 2. At the same time, the policy may widen the productivity gap between sectors, which reduces final output through the synergy mechanism described above. The central questions are whether the learning gains dominate the synergy costs of the induced productivity dispersion, and whether such a policy would secure democratic support from workers in both sectors.

[8] Intuitively, in the no-policy equilibrium, per capita income is identical across sectors, so the saving/borrowing margin is inactive. Around this point, the voting conditions with and without saving are locally identical.

2.5 Aggregate Output, Consumption and Lifetime Utility

We now derive the model's equilibrium allocations. We proceed in three steps: first, we solve for period-1 aggregate output and consumption; second, we determine period-2 productivity, output, and consumption; third, we compute lifetime utilities, which form the foundation for the voting analysis in Section 4.

2.5.1 Period 1: Equilibrium Allocations

Aggregate output. Substituting the sectoral output expressions (8) into the final goods production function (4) yields aggregate output in period 1:

$$\ln(y_1) = \frac{x + \ln(z_1) + \alpha \ln(\alpha) + \sum_i \chi_i \ln(\chi_i)}{1-\alpha}, \tag{13}$$

where the tax terms cancel out due to the government's balanced budget constraint. Hence, industrial policy has no effect on the level of aggregate period-1 output; it affects only its sectoral composition.

Firm profits. From the monopsony wage-setting rule, the profit of firm j in sector i is the fraction κ of the labor revenue product:

$$\pi_{i,j,1} = \kappa(1-\alpha)(1-\tau_i)p_{i,1}y_{i,j,1}.$$

Using the factor share condition $\chi_i y_1 = y_{i,1}p_{i,1}$ and $y_{i,1} = N_i y_{i,j,1} = \chi_i y_{i,j,1}$, per capita profit in sector i simplifies to:

$$\pi_{i,1} = \pi_{i,j,1} = \kappa(1-\alpha)(1-\tau_i)y_1. \tag{14}$$

Dividend income. Since households hold a fully diversified portfolio, dividend income per capita is identical across sectors. Aggregating profits using equation (14) yields:

$$\sum_i \chi_i \pi_{i,1} = \kappa(1-\alpha)y_1 \sum_i \chi_i(1-\tau_i) = \kappa(1-\alpha)y_1, \tag{15}$$

where the final equality follows from the balanced budget constraint. Thus, every household receives the same dividend income regardless of its sector of employment.

Labor income. The labor income per worker equals the fraction $(1-\kappa)$ of the labor revenue product:

$$l_{i,j,1}w_{i,j,1} = (1-\tau_i)(1-\kappa)(1-\alpha)y_1. \tag{16}$$

Consumption. Consumption per capita in sector i is the sum of labor income (equation (16)) and dividend income (equation (15)):

$$c_{i,1} = (1-\tau_i)(1-\kappa)(1-\alpha)y_1 + \kappa(1-\alpha)y_1 = (1-\tau_i+\tau_i\kappa)(1-\alpha)y_1.$$

Taking logs and applying the normalization $x = 0$ and $z_1 = 1$ from equation (13) yields:

$$\ln(c_{i,1}) = \ln(1-\alpha) + \ln(1-\tau_i+\tau_i\kappa) + \frac{\alpha\ln(\alpha) + \sum_i \chi_i \ln(\chi_i)}{1-\alpha}. \tag{17}$$

Equation (17) encapsulates a central tension in the political economy of industrial policy. A tax on sector i ($\tau_i > 0$) reduces per capita consumption in that sector through the term $(1-\tau_i+\tau_i\kappa)$. However, higher monopsony power κ partially offsets this reduction: the rents extracted from the taxed sector's labor market accrue to all households as dividend income via their diversified equity portfolios. As we will show, monopsony power, conventionally viewed as a distortion to be addressed by industrial policy, renders industrial policy more palatable to the taxed majority.

2.5.2 Period 2: Productivity, Output, and Consumption

In period 2, taxes and subsidies are set to zero. Sectoral output is determined by the productivity levels attained through the learning mechanisms described in Section 2.2.

Within-sector learning. Substituting the per capita capital expression from equation (7) into the within-sector learning equation (9) yields:

$$\ln(\widehat{z}_{i,2}) = \psi\left[\ln(1-\tau_i) + \frac{\ln(\alpha) + \sum_i \chi_i \ln(\chi_i)}{1-\alpha}\right]. \tag{18}$$

This shows that a higher tax rate reduces second-period productivity by dampening capital input and learning-by-doing, while subsidies increase productivity.

Inter-sectoral learning spillovers. Since sector A receives a subsidy ($\tau_A < 0$) and sector B is taxed ($\tau_B > 0$), equation (18) implies $\widehat{z}_{A,2} > \widehat{z}_{B,2}$. Therefore, only sector B learns from sector A:

$$\ln(z_{A,2}) = \ln(\widehat{z}_{A,2}), \tag{19}$$

$$\begin{aligned}\ln(z_{B,2}) &= (1-\phi)\ln(\widehat{z}_{B,2}) + \phi\ln(\widehat{z}_{A,2}) \\ &= \psi\left[(1-\phi)\ln(1-\tau_B) + \phi\ln(1-\tau_A) + \frac{\ln(\alpha) + \sum_i \chi_i \ln(\chi_i)}{1-\alpha}\right].\end{aligned} \tag{20}$$

Aggregate output. Substituting the productivity equations (19) and (20) into the final goods production function equation (4) yields aggregate output in period 2:

$$
\begin{aligned}
(1-\alpha)\ln(y_2) = &\left(1+\frac{\psi}{1-\alpha}\right)\left[\ln(\alpha)+\sum_i \chi_i \ln(\chi_i)\right] \\
&+\psi\left[\chi_A+\phi\chi_B-(1-\phi)\gamma\right]\ln(1-\tau_A) \\
&+\psi(1-\phi)(\chi_B+\gamma)\ln(1-\tau_B).
\end{aligned} \tag{21}
$$

The first term captures the direct effect of period-1 capital on productivity in the absence of policy. The second and third terms capture the additional effects of the subsidy and tax, operating through the synergy and spillover channels.

Consumption. With no taxes in period 2, household income per capita consists of labor income, $(1-\kappa)(1-\alpha)y_2$, and dividend income, $\kappa(1-\alpha)y_2$. Consumption per capita is therefore the sum of these two terms:

$$
c_{i,2} = (1-\alpha)y_2,
$$

which is common across sectors. In log form:

$$
\ln(c_{i,2}) = \ln(1-\alpha) + \ln(y_2). \tag{22}
$$

2.5.3 Lifetime Utility

Combining period-1 consumption equation (17) and period-2 consumption equation (22) with the output expression equation (21), the lifetime utility of a household in sector i is:

$$
\begin{aligned}
U_i = &\ln(1-\tau_i+\tau_i\kappa) + \frac{\beta\psi}{1-\alpha}\left[\chi_A+\phi\chi_B-(1-\phi)\gamma\right]\ln(1-\tau_A) \\
&+\frac{\beta\psi}{1-\alpha}(1-\phi)(\chi_B+\gamma)\ln(1-\tau_B)+\zeta,
\end{aligned} \tag{23}
$$

where ζ is a constant that does not depend on policy:

$$
\zeta = (1+\beta)\ln(1-\alpha) + \left[\frac{1}{1-\alpha}+\beta\frac{1-\alpha+\psi}{(1-\alpha)^2}\right]\left[\ln(\alpha)+\sum_i \chi_i \ln(\chi_i)\right]. \tag{24}
$$

In the absence of industrial policy ($\tau_A = \tau_B = 0$), both sectors achieve identical lifetime utility, $U_A = U_B = \zeta$. This ζ term serves as the foundation for the voting analysis that follows.

3 Voting Conditions

We now derive the conditions under which workers in each sector support industrial policy, together with the welfare criterion of a utilitarian social planner.

3.1 Voting by Sector B Workers

Households in sector B support industrial policy if it increases their lifetime utility relative to the no-policy baseline, i.e., if $U_B(\tau_A < 0, \tau_B > 0) > U_B(0,0) = \zeta$. Using the utility expression equation (23), this condition is:

$$\ln(1-\tau_B+\kappa\tau_B) + \frac{\beta\psi}{1-\alpha}(1-\phi)(\chi_B+\gamma)\ln(1-\tau_B) \\ + \frac{\beta\psi}{1-\alpha}\left[\chi_A+\phi\chi_B-(1-\phi)\gamma\right]\ln(1-\tau_A) > 0. \tag{25}$$

Applying the approximations $\ln(1-\tau_i) \approx -\tau_i$ (valid for small $|\tau_i|$) and the government budget constraint $\chi_A\tau_A + \chi_B\tau_B = 0$, which implies $\tau_B = -(\chi_A/\chi_B)\tau_A$, condition (25) simplifies to:

$$\frac{(1-\alpha)(1-\kappa)\chi_A}{\beta\psi} + (1-\phi)\gamma < \phi\chi_B. \tag{26}$$

The first term on the LHS of condition (26) is the tax burden on sector B workers. It increases with the size of sector A (χ_A) and is partially offset by monopsony power (κ). The second term is the synergy cost arising from the productivity gap, which is partially mitigated by learning spillovers. The RHS captures the benefit to sector B workers from inter-sectoral learning spillovers ($\phi\chi_B$).

Interpretation. Condition (26) reveals the forces that make sector B workers more likely to support industrial policy:

- χ_A is low: A small subsidized sector minimizes the tax burden on sector B.
- γ is low: Weak synergies reduce the cost of the productivity gap created by the subsidy.
- ϕ is high: Strong inter-sectoral learning generates substantial spillover benefits to sector B.
- κ is high: Greater monopsony power partially offsets the tax burden on sector B workers through increased dividend income.
- α is high: A higher capital share attenuates the tax cost on sector B via reduced labor income.
- β is high: Patient households place greater weight on future productivity gains.
- ψ is high: Strong within-sector learning amplifies the productivity effects of the subsidy.

The role of monopsony power deserves particular emphasis. When κ is high, taxing sector B reduces its workers' labor income but generates profits that flow back to them (and all other households) as dividends. This redistribution channel makes industrial policy more appealing to the taxed majority.

3.2 Voting by Sector A Workers

Households in sector A support industrial policy if $U_A(\tau_A < 0, \tau_B > 0) > U_A(0,0) = \zeta$. Following the same steps as above, this condition can be derived as:

$$(1-\phi)\gamma < \phi\chi_B + \frac{(1-\alpha)(1-\kappa)\chi_B}{\beta\psi}. \tag{27}$$

The first term on the left-hand side and the first term on the right-hand side of condition (27) are the synergy cost and the learning benefit, respectively. These are the same as in condition (26) because workers in both sectors are proportionally affected by the policy's impact on the aggregate economy, owing to the fixed factor share in the final goods production function. The second term on the right-hand side is the direct subsidy benefit to sector A workers. It increases with the size of the taxed sector (χ_B) and is partially offset by monopsony power (κ).

Interpretation. Condition (27) is strictly weaker than condition (26) for sector B workers, indicating that sector A households are more likely to support industrial policy than sector B households. This is intuitive: sector A workers receive direct subsidies, while sector B workers bear tax burdens.

Condition (27) is more likely to hold when:

- χ_B is high: A large taxed sector provides a broader tax base to finance the subsidy.
- γ is low: Synergies are weak, so the productivity gap created by the subsidy imposes lower cost on aggregate output.
- ϕ is high: Inter-sectoral learning is strong, generating substantial productivity spillovers that benefit sector A workers through higher future aggregate income.
- κ is low: Monopsony power is low, so the subsidy flows primarily to workers as higher labor income rather than being captured as profits.
- α is low: The labor income share is high, so sector A workers capture a larger fraction of the subsidy through wages.

- β is low: Impatient households place greater weight on the immediate consumption gains from the subsidy than on future productivity benefits that are partially shared with sector B.

- ψ is low: The immediate consumption gains from the subsidy becomes more important than the productivity benefits that are shared with sector B.

The last two conditions may appear counterintuitive but have a clear economic interpretation. The benefits of subsidies accrue to sector A workers primarily in period 1 through higher labor income, while the period-2 benefits from stimulated learning-by-doing are shared with sector B through spillovers. When sector A workers are impatient (low β) or learning is weak (low ψ), they care more about immediate consumption gains than future productivity gains.

3.3 Social Planner

We now study the conditions under which industrial policy improves aggregate social welfare from the perspective of a utilitarian social planner who maximizes the population-weighted sum of utilities across sectors:

$$U = \chi_A U_A + \chi_B U_B, \tag{28}$$

where we have imposed the normalization $N_i = \chi_i$.

An industrial policy raises U relative to the no-policy baseline whenever:

$$\chi_A \left[\phi\chi_B + \frac{\chi_B}{\beta\psi} - (1-\phi)\gamma \right] + \chi_B \left[\phi\chi_B - \frac{\chi_A}{\beta\psi} - (1-\phi)\gamma \right] > 0. \tag{29}$$

The $\chi_A\chi_B/(\beta\psi)$ terms with opposite signs cancel, yielding the compact condition:

$$(1-\phi)\gamma < \phi\chi_B. \tag{30}$$

Interpretation. The social planner's condition (30) depends on only three parameters. The inequality is satisfied when the subsidized sector is small (χ_A low, equivalently χ_B high), synergies are weak (γ low), and inter-sectoral learning is strong (ϕ high).

The social planner's condition is independent of monopsony power κ. In the baseline model with fixed labor supply, κ is purely redistributive in nature and therefore does not affect aggregate welfare. In Appendix E, we relax the assumption of inelastic labor supply and show that monopsony power introduces a deadweight loss that renders the social planner's condition more restrictive; the fundamental political economy tension, however, is unchanged.

A comparison of the three conditions establishes their relative restrictiveness. The social planner's condition (30) occupies an intermediate position: it is weaker than the condition for sector B workers (condition (26)) but stronger than the condition for sector A workers (condition (27)). This ordering yields two central implications. First, there exists a non-empty region of the parameter space in which industrial policy is welfare-improving (satisfying conditon (30)) yet fails to secure political support from sector B workers (violating condition (26)). Second, sector A workers support industrial policy over a broader range of parameter values than is socially optimal, reflecting their direct receipt of subsidies while bearing only a fraction of the aggregate cost.

3.4 Numerical Illustration

To build intuition, we present a quantitative illustration of the theoretical results calibrated to empirically grounded parameter values. Although the key parameters are estimated from microdata in Section 5, the illustration serves to crystallize the trade-offs between political feasibility and social welfare, demonstrating that, under plausible parameterizations, industrial policy can be welfare-improving yet politically unattainable, or politically viable yet socially inefficient.

Calibration. We calibrate the key parameters as follows. The within-sector learning elasticity is set to $\psi = 0.25$, following Bartelme et al. (2025), who estimate the elasticity of productivity with respect to production scale from international trade shocks. The inter-sectoral learning rate is set to $\phi = 0.21$, based on the empirical estimates reported in Section 5.

The monopsony power parameter κ is calibrated to capture the share of the surplus that accrues as pure profit to shareholders, after accounting for the cost of capital. In our model, the firm's period-1 profit is $\pi_{i,j,1} = \kappa(1-\alpha)(1-\tau_i)p_{i,1}y_{i,j,1}$ and labor income is $w_{i,j,1}l_{i,j,1} = (1-\tau_i)(1-\kappa)(1-\alpha)p_{i,1}y_{i,j,1}$, so κ represents the fraction of the non-capital residual that is captured as shareholder surplus rather than paid to workers. Empirically, we measure κ for each industry i using U.S. national accounts as:

$$\kappa_i = \frac{\text{Operating Surplus}_i - (r+\delta_i)K_i}{\text{Employee Compensation}_i + \text{Operating Surplus}_i - (r+\delta_i)K_i}, \tag{31}$$

where the numerator is profit –the excess of operating surplus over the full cost of capital services– and the denominator is the sum of profit and employee compensation. Operating surplus is the gross return to capital, measured as value added less employee compensation. Capital cost is the user cost of capital, $(r+\delta_i)K_i$, where K_i is the capital stock, r is the real interest rate, and δ_i is the depreciation rate.[9] For our numerical illustration below, we estimate the

[9]Data are from the BEA 2002 benchmark Input-Output Accounts and Fixed Assets Accounts. Specifically, value added equals operating surplus plus employee compensation; operating surplus equals capital cost plus profit;

aggregate monopsony power as the average value of κ_i weighted by industry value added, which is approximately 0.21.

The subjective discount factor is set to $\beta = 0.974$, which implies a steady-state real interest rate of approximately 2.6 percent, close to the sample average of 2.56 percent. Finally, the synergy parameter is set to $\gamma = 0.25$, based on the empirical estimates reported in Section 5.

As discussed in Section 5, these parameters exhibit substantial heterogeneity across industries, which can affect the political feasibility of industrial policy in different sectors. We exploit this heterogeneity in Section 6 to evaluate industrial policy proposals for the U.S. manufacturing sector. The numerical illustration in the present section serves to isolate the key trade-offs generated by the theoretical framework.

Parameter space analysis. Figures 1, 2, and 3 illustrate the region of the parameter space in which each of the three decision-makers—sector B workers, sector A workers, and the social planner—supports industrial policy.

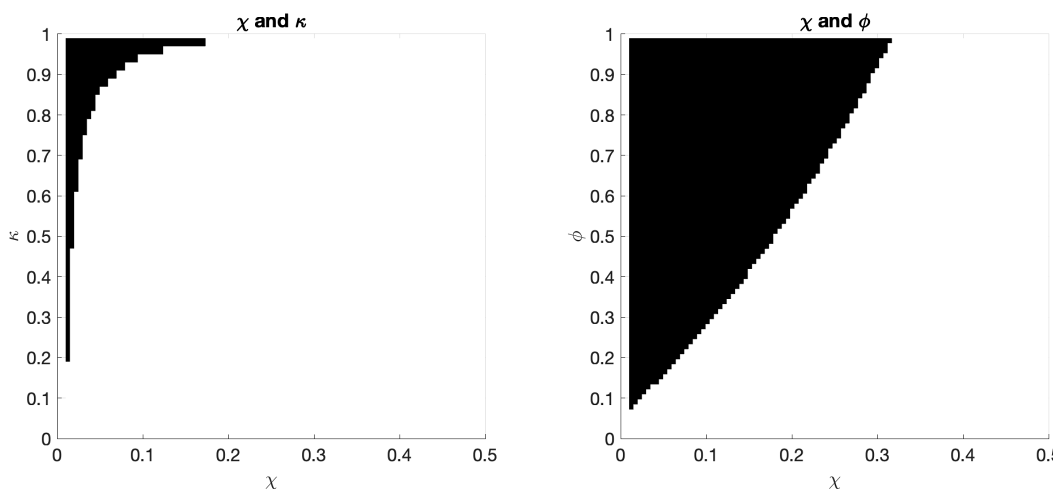

Figure 1: Parameter Range for Sector B's Support (Equation 26)

Sector B support (Figure 1). The figure illustrates the parameter combinations satisfying the voting condition for sector B to support industrial policy in sector A (equation 26), which is the *necessary condition for the policy to pass democratically*. The shaded regions represent parameter combinations where sector B would vote in favor. As predicted by our theory, support for industrial policy increases when:

and capital cost is proxied by $(r + \delta)K$. The capital stock is the net stock of private fixed assets; the depreciation rate is the ratio of consumption of fixed capital to the net capital stock; and the real interest rate of 2.56% is the average of the Federal Reserve Bank of Cleveland's real interest rate series over 1982–2020. We use 2002 to maintain consistency across all BEA sources.

- Inter-sectoral learning (ϕ) is strong, allowing sector B to benefit substantially from knowledge spillovers

- Sector A is small (χ_A is low), minimizing the tax burden on sector B

- Monopsony power is high (κ is high), increasing sector B's dividend income from holding sector A's equity

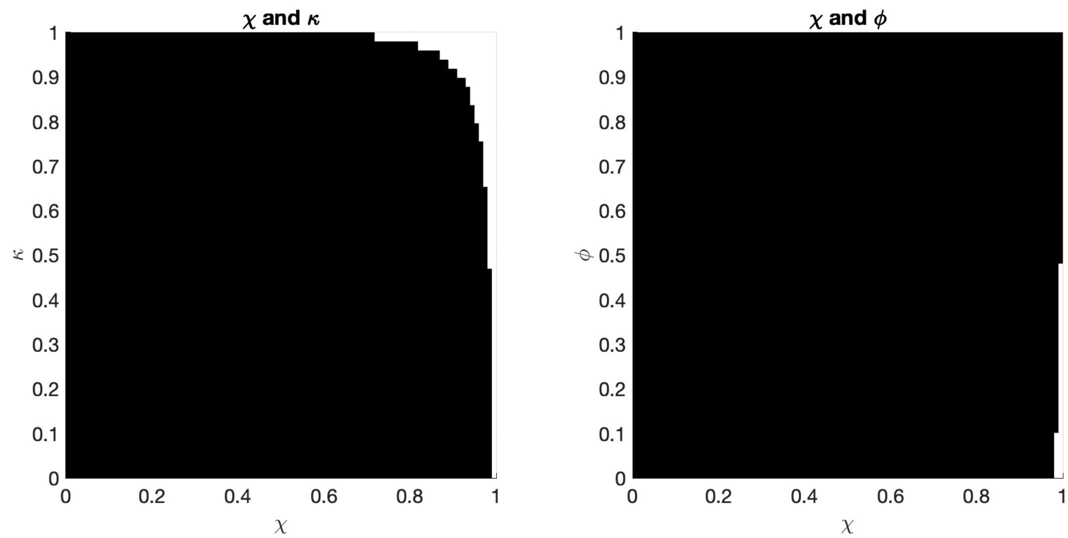

Figure 2: Parameter Range for Sector A's Support (Equation 27)

Sector A support (Figure 2). This figure shows parameter combinations where sector A supports industrial policy in its own sector (satisfying inequality 27), which is the *necessary condition for the policy to be proposed*. The shaded regions are substantially larger than in Figure 1, confirming that sector A's support is much broader. This reflects the fact that sector A receives direct subsidies while bearing only a fraction χ_A of the aggregate tax burden.

At the estimated synergy parameter of $\gamma = 0.25$, sector A supports the policy across a wide range of parameter combinations, although higher values of γ would shrink the support region. Opposition from sector A's own workers arises only in corner cases: when monopsony power is very high, so that the bulk of the subsidy is captured as rents by shareholders rather than flowing to workers, or when the sector is extremely large, making the aggregate inefficiency of the subsidization too costly even for its direct beneficiaries.

Figure 2 reveals a strong inherent bias in favor of sector-specific industrial policy: workers in any given sector have a powerful incentive to seek subsidies for their own industry. That such proposals are not ubiquitous in practice is consistent with the strategic anticipation that workers in other sectors would refuse to bear the associated tax burden.

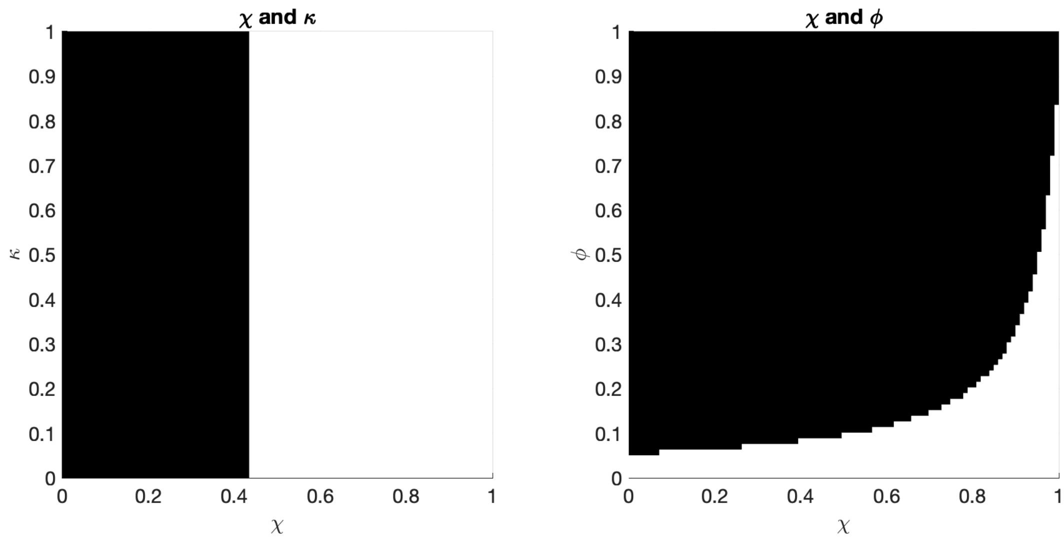

Figure 3: Parameter Range for Social Planner's Support (Equation 30)

Social planner support (Figure 3). This figure displays parameter combinations satisfying the social planner's condition (equation 30), which is the *necessary condition for industrial policy to be welfare-improving*. The social planner's condition is intermediate in restrictiveness between sector A's and sector B's conditions. Support increases when:

- Inter-sectoral learning (ϕ) is strong
- Sector A is small (χ_A is low), amplifying the per capita effect of subsidies on learning-by-doing

Comparison and discussion. A joint reading of the three figures yields several insights:

First, as the direct beneficiary of subsidies, sector A supports industrial policy across a wide range of parameter values. Condition (27) is strictly weaker than condition (26), reflecting sector A's structurally advantageous position as the recipient of transfers.

Second, as the taxed sector, sector B imposes the most stringent conditions for supporting industrial policy. The parameter region generating sector B's support (Figure 1) is substantially narrower than that generating sector A's support (Figure 2). Sector B workers require clear evidence of economy-wide spillover benefits to justify the tax burden they must bear.

Third, the social planner supports industrial policy when aggregate benefits exceed aggregate costs, a condition that is more demanding than sector A's support condition but less demanding than sector B's. A comparison of Figures 1 and 3 reveals a non-trivial region of the parameter space in which industrial policy is welfare-improving yet fails to pass democratically, underscoring a fundamental tension between economic efficiency and political incentives.

The central political-economy tension—beneficiaries strongly favor subsidies while the taxed majority demands tangible spillover benefits—implies that many welfare-improving industrial policies are likely to be politically infeasible under majority voting, while many policies proposed by beneficiaries are likely to be welfare-reducing.

4 Equilibrium

For analytical completeness, we now characterize the equilibrium of the model.

Definition 1 (Decentralized equilibrium). *Given a policy (τ_A, τ_B) satisfying the government budget constraint (12), a decentralized equilibrium is a collection of prices $\{p_{i,t}\}_{i,t}$, wages $\{w_{i,j,t}\}_{i,j,t}$, firm allocations $\{l_{i,j,t}, k_{i,j,t}, y_{i,j,t}\}_{i,j,t}$, and household consumption $\{c_{i,t}\}_{i,t}$ such that:*

(i) ***Intermediate goods firm optimization.*** *Taking $(p_{i,t}, \tau_i)$ as given, each firm chooses $l_{i,j,t}$ and $k_{i,j,t}$ to solve the profit-maximization problem of Section 2, setting wages according to the monopsony rule.*

(ii) ***Final goods firm optimization.*** *Taking $\{p_{i,t}\}$ as given, the representative final goods producer chooses $\{y_{i,t}\}$ to maximize profit subject to production technology equation (3).*

(iii) ***Labor market clearing.*** *Each sector's labor market clears at $l_{i,t} = N_i$, consistent with inelastic unit labor supply.*

(iv) ***Household consumption.*** *Each household in sector i consumes the total income, the sum of labor income and dividend income.*

(v) ***Goods market clearing.*** *The final good market clears, $y_t = N_A c_{A,t} + N_B c_{B,t} + \alpha y_t$, where the last term is the capital input.*

Sectoral productivity $\{z_{i,2}\}_i$ evolves across periods according to the learning-by-doing and spillover equations (9)–(10), taking period-1 equilibrium capital $k_{i,1}$ as given.

5 Empirical Strategy: Estimating Synergies and Inter-Sectoral Learning

This section presents our empirical approach for estimating the synergy parameter γ and the inter-sectoral learning rate ϕ that are key parameters of our theoretical model. We first discuss the core intuition behind our identification strategy, then describe the new data constructed by merging different datasets, and finally present the estimation results. While the estimation in this section is focused on the aggregate economy, Section 6 applies the same empirical methodology to estimate these parameters separately for each three-digit Naics industry, enabling us to evaluate which industrial policies would receive support from different stakeholder groups.

5.1 Identification Strategy

Estimating synergies and inter-sectoral learning at the sectoral level is challenging (Griffith et al., 2004; Bernstein and Nadiri, 1988; Keller, 2002; Bloom et al., 2013; Fernández-Villaverde et al., 2025b). The direct comparisons of productivity levels across sectors are problematic because sectoral productivity measures may reflect fundamental differences in technology, capital intensity, or market structure rather than the presence of synergies or learning potential. For instance, while the finance sector typically exhibits higher measured productivity than manufacturing, this does not necessarily imply that (i) finance employs more advanced technologies from which manufacturing could learn, (ii) manufacturing acts as a bottleneck in joint production with finance, or (iii) productivity gaps between these sectors generate meaningful synergy costs.

To overcome these challenges, we exploit firm-level heterogeneity *within* sector pairs. Our identification strategy relies on the premise that within a given industry, firms can be ranked by their productivity. By focusing on changes in productivity rankings, which are independent of sectoral factors, we can identify and quantify the distinct channels of synergies and inter-sectoral learning. We then examine how the productivity *rankings* of trading partners is linked to firm performance and productivity growth. This approach allows us to isolate the effects of productivity gaps and learning opportunities from other sector-specific factors.

Consider two illustrative trading relationships: a high-productivity R&D firm supplying a low-productivity manufacturer, versus the same R&D firm supplying a high-productivity manufacturer. If synergies matter, the productivity gap in the first relationship should create bottlenecks and reduce output more severely. If inter-sectoral learning occurs, the low-productivity manufacturer should grow faster, narrowing the gap with its high-productivity partner relative to peers facing the same industry shocks. Testing these two predictions with firm-level trading data lets us estimate the magnitude of synergies (γ) and the rate of inter-sectoral learning (ϕ)—the key parameters of our theoretical model. Appendix D derives the formal mapping from these firm-level coefficients to the sector-level parameters γ and ϕ in the aggregate production function and the inter-sectoral learning equation. The remaining concern is that mismatch formation may be endogenous to unobserved firm characteristics; We address this with an instrumental variables strategy based on exogenous predictors of mismatch.

5.2 Data

Our empirical analysis integrates three distinct datasets to develop comprehensive information on inter-firm trading relationships and firm performance in the U.S. The first is the *Compustat Segment* data, which documents customer-supplier relationships for all publicly listed U.S. firms. Since disclosure regulations require firms to report trading partners accounting for more than 10% of their annual sales, this dataset identifies 72,694 distinct partner relationships among approximately 18,000 firms over the period 1976–2020. The second is the *FactSet Supply Chain Relationships* data, which compiles trading relationships from public sources including SEC 10-K filings, investor presentations, and press releases. For the period 2003–2020, this dataset provides details on 289,239 distinct partner-supplier relationships.[10]

We merge both datasets with *Compustat Fundamentals Annual* data, which provides firm-level information on output (sales), employment, capital, and financial positions. Our integrated dataset enables us to link trading relationships with firm performance and productivity measures.

5.3 Synergies and productivity gaps

This subsection studies how synergies and mismatches in productivity between firms and their trading partners are linked to firm performance. Our theoretical framework shows that synergies arise when firms with similar productivity levels collaborate, and large productivity gaps create bottlenecks that reduce output. We test these hypotheses using firm-level data on customer-supplier relationships.

5.3.1 Empirical specification

We examine productivity gaps from both the customer and supplier perspectives using two symmetric empirical specifications. For the customer side, we estimate:

$$y_{j,t} = \widehat{\gamma}^{cus}\overline{\Delta}_{j,t}^{cus} + \beta z_{j,t} + \eta \times \overline{decile}\left(z_{j,t}^{cus}\right) + firm_j + \chi_t + \epsilon_{j,t}, \tag{32}$$

where $y_{j,t}$ and $z_{j,t}$ denote firm j's log value-added and log productivity in year t, respectively. $\overline{\Delta}_{j,t}^{cus}$ measures the average productivity gap between firm j and its customers, and $\overline{decile}\left(z_{j,t}^{cus}\right)$ controls for customers' average productivity level.

[10]See Xu et al. (2025) for a detailed description of the FactSet Supply Chain Relationships data.

For the supplier side, we estimate an analogous specification:

$$y_{j,t} = \widehat{\gamma}^{sup}\overline{\Delta}_{j,t}^{sup} + \beta z_{j,t} + \eta \times \overline{decile}\left(z_{j,t}^{sup}\right) + firm_j + \chi_t + \epsilon_{j,t}, \tag{33}$$

where $\overline{\Delta}_{j,t}^{sup}$ measures the average productivity gap between firm j and its suppliers, and $\overline{decile}\left(z_{j,t}^{sup}\right)$ controls for suppliers' average productivity level.

In both specifications, productivity is measured in the following way. For goods-producing industries (e.g., construction, manufacturing), we use total factor productivity (TFP) estimated via the enhanced Wooldridge-Levinsohn-Petrin method, which corrects for simultaneity bias and functional dependence in production function estimation (Levinsohn and Petrin, 2003; Wooldridge, 2009). For service-providing industries (e.g., professional and business services), we use labor productivity (LP), defined as value added per employee, given the greater difficulty in accurately measuring capital inputs in these sectors.

The key explanatory variables $\overline{\Delta}_{j,t}^{cus}$ and $\overline{\Delta}_{j,t}^{sup}$ capture the average productivity gap between firm j and its customers or suppliers, respectively. For each partner k, the mismatch is defined as $\Delta_{j,k,t} = |decile\left(z_{j,t}\right) - decile\left(z_{j,k,t}\right)|$, which ranges from 0 (perfect alignment) to 9 (maximum mismatch). Here, $decile\left(z_{j,t}\right)$ and $decile\left(z_{j,k,t}\right)$ represent the decile of productivity ranking **within** the firm's three-digit Naics industry. We compute firm-level averages $\overline{\Delta}_{j,t}^{cus}$ and $\overline{\Delta}_{j,t}^{sup}$ across all of firm j's customers or suppliers.

We use decile ranks rather than log-productivity levels because deciles are more comparable across industries with different productivity distributions. Appendix D formalizes the mapping between decile mismatches and the log-productivity gap γ that enters the theoretical model, and shows that under our estimated scaling factor the two are approximately equivalent, so the regression coefficients below can be interpreted directly as (the negative of) the structural synergy parameter.

We control for the average productivity decile of partners, $\overline{decile}\left(z_{j,t}^{cus}\right)$ or $\overline{decile}\left(z_{j,t}^{sup}\right)$, to account for the possibility that firms systematically match with higher- or lower-productivity partners for reasons unrelated to synergies. Firm fixed effects ($firm_j$) absorb time-invariant firm characteristics, while year fixed effects (χ_t) capture common macroeconomic shocks. The coefficients of interest, $\widehat{\gamma}^{cus}$ and $\widehat{\gamma}^{sup}$, measure the effect of productivity gap on firm output from customer and supplier perspectives, respectively. Consistent with the synergy hypothesis, we expect $\widehat{\gamma}^{cus} < 0$ and $\widehat{\gamma}^{sup} < 0$: larger mismatches should reduce firm output.

5.3.2 Instrumental variables strategy

A key identification challenge is that productivity gap may be endogenous. Firms might select trading partners based on unobserved characteristics that also affect output, leading to biased estimates of $\widehat{\gamma}^{cus}$ and $\widehat{\gamma}^{sup}$. For example, better-managed firms are both more productive and more adept at forming well-aligned partnerships, creating a spurious correlation between mismatch and performance. To address this concern, we implement an instrumental variables (IV) approach using two novel instruments that plausibly influence productivity gap but are exogenous to unobserved determinants of firm output.

The first instrument is **prior inter-firm manager transitions**, which counts senior managers who have worked at both firms prior to the formation of the trading relationship. We define senior managers as workers in seniority levels 5–7 under Revelio Labs' 7-tier taxonomy, which standardizes job titles into an integer scale from 1 (most junior) to 7 (most senior). Levels 5–7 correspond to director-, executive-, and senior executive-level roles (e.g., Head of Legal, Managing Director, COO/CEO). The taxonomy is constructed using Revelio's dataset, which aggregates career histories from over 33 million U.S. LinkedIn profiles. Since managers tend to move between firms of similar productivity levels, this variable facilitates more assortative matching through enhanced mutual understanding. Consequently, we expect the *initial productivity gap* at the formation of trading relationsip to be negatively correlated with the pre-relationship flow of senior managers. Because mismatch is highly persistent over time, pre-relationship manager flows also predict current productivity gap, making our IV relevant.

The second instrument is **size-distance**, defined as the absolute difference in sales deciles between the firm and its partner *in the year before the match formed*: $|decile\left(y_{j,t_0-1}\right) - decile\left(y_{j,k,t_0-1}\right)|$, where t_0 denotes the initial year of trading relationship. The intuition is that highly productive firms may accept less productive partners if those partners are substantially larger, creating a trade-off between productivity alignment and the opportunity to work with large, established firms (which presumably offer greater business stability conditional on productivity).

Both instruments are averaged across all partners to construct firm-level measures. Table 1 reports the first-stage results, confirming that our instruments are strong predictors of mismatch.

A central feature of our identification strategy is to exploit the high persistence of productivity gaps, i.e., the productivity gap observed at relationship formation (Δ_{t_0}) strongly predicts productivity gap several years later (Δ_t). Because our instruments affect Δ_{t_0}, they therefore affect Δ_t through this persistence channel. In the language of dynamic panel data models, we are effectively using $\widehat{\Delta}_{t_0}$ (the fitted initial mismatch) as an instrument for Δ_t. This logic is analogous to using lagged endogenous variables (Δ_{t_0}) as instruments, though our instruments yield a more credible exclusion restriction.

Table 1: First-stage regression results for instrumental variables

	(1)	(2)
Dependent variable	Productivity gap ($\Delta_{j,k,t}$)	
Pre-relationship senior manager transition	-0.0056***	
	(0.0014)	
Pre-relationship size-distance		0.0469***
		(0.0085)
Time fixed effect	Yes	Yes
Firm fixed effect	Yes	Yes
Adjusted R^2	0.448	0.445
Observations	61,368	58,118

Note: Sample: 1976-2020. The table reports first-stage regressions for the instrumental variables used in equations (32) and (33). Column (1) shows the regression of productivity gap on the number of senior managers in both firms. Column (2) shows the regression of productivity gap on size-distance mismatch. The number of senior managers in both firms is the count of senior managers (seniority level higher than 4 out of 7) who have worked at both firms prior to the trading relationship, using data from Revelio Labs. Size-distance mismatch is the absolute difference in sales deciles between the firm and its partner in the year before the match formed. Robust standard errors in parentheses. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.3.3 Estimation results

Table 2 presents the IV estimation results for both customer and supplier specifications. Column (1) shows the results for supplier firms, where the dependent variable is supplier log sales and the key explanatory variable is $\overline{\Delta}^{cus}_{j,t}$ (customer productivity gap). Column (2) shows the results for customer firms, where the dependent variable is customer log sales and the key explanatory variable is $\overline{\Delta}^{sup}_{j,t}$ (supplier productivity gap).

For supplier firms in column (1), we find a highly significant negative effect: $\widehat{\gamma}^{cus} = -0.413$ (significant at the 1% level). This indicates that a one-unit increase in the average customer productivity gap reduces supplier sales by approximately 0.413 log points. The economic magnitude is substantial: moving from perfect alignment (mismatch = 0) to the maximum mismatch of 9 would reduce supplier output by about 3.7%. For our numerical illustration in the prior section, we therefore used 0.25, the negative of the two columns' average, as the calibrated synergies parameter.

Table 2: Output decline with productivity gap

	(1)	(2)
Dependent variable	Supplier firm's log sales	Customer firm's log sales
$\overline{\Delta}_{j,t}^{cus}$	-0.413***	
	(0.073)	
$\overline{\Delta}_{j,t}^{sup}$		-0.086
		(0.279)
$\overline{decile}\left(z_{j,t}^{cus}\right)$	-0.110***	
	(0.021)	
$\overline{decile}\left(z_{j,t}^{sup}\right)$		0.007
		(0.014)
$z_{j,t}$	0.305***	0.561*
	(0.046)	(0.327)
Time FE	Yes	Yes
Firm FE	Yes	Yes
Observations	16,205	14,440
Identification tests for IV specifications:		
Underid. test (KP rk LM)	69.518***	5.329*
	(0.000)	(0.070)
Weak id. test (CD Wald F)	59.100	2.296
Weak id. test (KP Wald rk F)	35.025	2.078
Overid. test (Hansen J)	0.562	1.054
	(0.453)	(0.305)

The data span 1976–2020. The dependent variable is log sales for the firm receiving the productivity gap effect: supplier firms in column (1) and customer firms in column (2). Both columns report IV estimates. Robust standard errors are in parentheses. For identification tests, p-values are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

For customer firms in column (2), the estimated effect is smaller and statistically insignificant: $\widehat{\gamma}^{sup} = -0.086$. This suggests that supplier productivity gaps have a less pronounced effect on customer performance compared to customer mismatches on supplier performance.

The identification tests support the validity of our instruments for the supplier regression. The Kleibergen-Paap rk LM statistic of 69.518 rejects underidentification ($p < 0.001$), and the Kleibergen-Paap Wald F-statistic of 35.025 exceeds the Stock-Yogo critical value for 10% maximal IV size (19.93). The Hansen J test does not reject instrument validity ($p = 0.453$). For the customer regression, however, the identification is weaker, with a Kleibergen-Paap LM statistic of 5.329 ($p = 0.070$) and a Wald F-statistic of 2.078, which falls below conventional critical values.

The negative coefficients on both $\widehat{\gamma}^{cus}$ and $\widehat{\gamma}^{sup}$ provide empirical support for the synergy mechanism in our theoretical model. When firms collaborate with partners of similar productivity levels, they avoid bottlenecks and coordination frictions, leading to higher output. The asymmetry in the magnitude and significance of the effects suggests that customer mismatches may be more consequential for firm performance than supplier mismatches, potentially because customers often drive production requirements and quality standards in supply relationships.

5.4 Inter-sectoral learning from customers and suppliers

This subsection studies whether firms can enhance their productivity by learning from higher-productivity trading partners, a central mechanism in our theoretical model. We test this hypothesis using two complementary approaches: learning from customers and learning from suppliers.

5.4.1 Empirical specification

We examine learning from both customers and suppliers using symmetric empirical specifications. For learning from customers, we estimate:

$$decile\left(z^{sup}_{j,k,t+5}\right) - decile\left(z^{sup}_{j,k,t}\right) = \phi^{cus}\widetilde{\Delta}^{cus}_{j,k,t} + firm_{j,k} + \chi_t + \epsilon_{j,k,t}, \tag{34}$$

where the dependent variable represents five-year productivity decile growth for firm j's kth supplier. The learning potential for this supplier (from customer j) is defined as

$$\widetilde{\Delta}^{cus}_{j,k,t} = \max\left[0, decile\left(z_{j,t}\right) - decile\left(z^{sup}_{j,k,t}\right)\right],$$

which ranges from zero (when firm j's productivity ranking equals or exceeds its customer's) to nine (maximum learning potential).

For learning from suppliers, we estimate an analogous specification:

$$decile\left(z^{cus}_{j,k,t+5}\right) - decile\left(z^{cus}_{j,k,t}\right) = \phi^{sup}\widetilde{\Delta}^{sup}_{j,k,t} + firm_{j,k} + \chi_t + \epsilon_{j,k,t}, \tag{35}$$

where the dependent variable is the five-year productivity decile growth for firm j's kth customer. $\widetilde{\Delta}^{sup}_{j,k,t}$ measures the learning potential of this customer from supplier j, defined analogously:

$$\widetilde{\Delta}^{sup}_{j,k,t} = \max\left[0, decile\left(z_{j,t}\right) - decile\left(z^{cus}_{j,k,t}\right)\right]$$

In both specifications, we control for the firm fixed effects for both firms ($firm_{j,k}$) and year fixed effects (χ_t). The coefficients of interest, $\widehat{\phi}^{cus}$ and $\widehat{\phi}^{sup}$, measure the effect of learning potential from customers and suppliers on productivity growth, respectively. We expect $\widehat{\phi}^{cus} > 0$ and $\widehat{\phi}^{sup} > 0$: firms with higher-productivity partners should experience faster productivity growth due to learning opportunities.

The measures of learning potential may be endogenous if unobserved firm characteristics influence both partner selection and productivity growth. For instance, more innovative firms might both attract higher-productivity partners and experience faster productivity growth for reasons unrelated to learning. To address this concern, we employ an instrumental variables (IV) approach with two instruments similar to those used in the synergy analysis. The first instrument is prior interfirm manager transitions (described in Section 5.3.1), which counts senior managers who have worked at both firms prior to the trading relationship. The second instrument is adjusted size-distance, defined as $\max\left[0, decile\left(y_{j,t}\right) - decile\left(y_{j,k,t}\right)\right]$. These instruments address the endogeneity concern by providing exogenous variation in learning potential that is driven by factors related to the formation of trading relationships rather than firms' inherent growth trajectories.

5.4.2 Estimation results

Table 3 presents the IV estimation results for both learning specifications. Column (1) shows the results for learning from customers (supplier productivity growth), while column (2) shows the results for learning from suppliers (customer productivity growth).

Table 3: Productivity increases with partner's productivity

	(1)	(2)
Dependent variable	Supplier's prod. growth	Customer's prod. growth
$\widetilde{\Delta}^{cus}_{j,k,t}$	0.216^{**}	
	(0.099)	
$\widetilde{\Delta}^{sup}_{j,k,t}$		0.203^{*}
		(0.110)
Time FE	Yes	Yes
Firm FE (both sides)	Yes	Yes
Observations	43,793	51,692
Identification tests for IV specifications:		
Underid. test (KP rk LM)	146.353^{***}	86.049^{***}
	(0.0000)	(0.0000)
Weak id. test (CD Wald F)	131.195	56.683
Weak id. test (KP Wald rk F)	66.554	38.735
Overid. test (Hansen J)	1.160	0.454
	(0.2815)	(0.5006)

The data span 1976-2020. Both columns report IV estimates. For identification tests, p-values are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

For learning from suppliers, the IV estimate in column (1) yields $\widehat{\phi}^{cus} = 0.216$ (significant at the 5% level). This indicates that a one-unit increase in learning potential from suppliers raises customers' five-year productivity growth by approximately 0.22 percentage points. For learning from customers, the IV estimate in column (2) yields $\widehat{\phi}^{sup} = 0.203$ (significant at the 10% level), indicating that a one-unit increase in learning potential from customers raises suppliers' five-year productivity growth by approximately 0.20 percentage points.

The identification tests support the validity of our instruments in both specifications. For learning from suppliers, the Kleibergen-Paap rk LM statistic of 146.353 rejects underidentification ($p < 0.001$), and the Kleibergen-Paap Wald F-statistic of 66.554 exceeds the Stock-Yogo critical value for 10% maximal IV size (19.93). The Hansen J test does not reject instrument validity ($p = 0.282$). For learning from customers, the instruments are also strong, with a Kleibergen-Paap LM statistic of 86.049 ($p < 0.001$) and a Wald F-statistic of 38.735, which also exceeds conventional critical values. The Hansen J test shows no evidence of overidentification ($p = 0.501$).

The economic magnitude of these effects is meaningful. Moving from no learning potential to the maximum learning potential of 9 would increase five-year productivity growth by approximately 1.9 percentage points for suppliers learning from customers and 1.8 percentage points for customers learning from suppliers. These effects translate to annual productivity growth improvements of approximately 0.38 and 0.36 percentage points, respectively.

The positive coefficients on both $\widehat{\phi}^{cus}$ and $\widehat{\phi}^{sup}$ provide empirical support for the inter-sectoral learning mechanism in our theoretical model. Firms indeed learn from higher-productivity trading partners, whether they are customers or suppliers. The similar magnitudes for customer and supplier learning suggest that knowledge flows operate symmetrically in both directions within buyer-supplier relationships. Suppliers may learn from customers about market requirements, quality standards, and product specifications, while customers may learn from suppliers about production processes, technological innovations, and efficiency improvements. These bidirectional learning effects provide empirical proxies for the inter-sectoral learning parameter ϕ in our theoretical framework.

6 Electoral support for industrial policies

In this section, we apply our theoretical framework and empirical estimates to assess which industrial policies would garner political support. We consider each three-digit NAICS industry as a potential subsidized sector (industry A) and evaluate whether subsidizing that industry would: (i) improve aggregate welfare from a social planner's perspective, (ii) receive support from workers in the subsidized industry, and (iii) receive support from workers in the rest of the economy (industry B).

Our theoretical analysis in Section 2.5 identifies distinct determinants of support for each group. For the social planner, support depends primarily on the balance between inter-sectoral learning spillovers and synergy costs. The social planner condition $(1-\phi_p)\gamma_p < \phi_p\chi_{-p}$ (equation 30) indicates that industries with **high inter-sectoral learning** rates (ϕ_p) and **low synergy costs** (γ_p) are more likely to receive support, as they generate substantial productivity spillovers to the rest of the economy while minimizing coordination frictions. Appendix D provides the theoretical derivation of how the firm-level estimates $\widehat{\gamma}_p^{cus}$, $\widehat{\gamma}_p^{sup}$, $\widehat{\phi}_p^{cus}$, and $\widehat{\phi}_p^{sup}$ map to the sector-level parameters γ_p and ϕ_p used in these conditions.

For workers in the rest of the economy (sector B), support requires an additional consideration: the immediate tax burden. The sector B condition $\frac{(1-\kappa_p)\chi_p}{\beta\psi} + (1-\phi_p)\gamma_p < \phi_p\chi_{-p}$ highlights that **high monopsony power (κ_p)** in the subsidized industry makes sector B more likely to support industrial policy. This occurs because sector B workers, as shareholders of diversified firms, benefit from the redistribution of profits when monopsony power is high. The tax burden on sector B is partially offset by increased dividend income from sector A's profits.

Workers in the subsidized industry (sector A) has the loosest support condition for industrial policy targeting their own sector, as they receive direct subsidies while bearing only a fraction of the costs. The sector A condition is $(1-\phi_p)\gamma_p < \phi_p\chi_{-p} + \frac{(1-\kappa_p)\chi_{-p}}{\beta\psi}$.

Our empirical evaluation proceeds in two steps. First, we estimate industry-specific parameters using U.S. data: the synergy parameter γ_p, the inter-sectoral learning rate ϕ_p, industry size $\chi_{A,p}$, and monopsony power κ_p, and calibrate the conventional parameters that apply uniformly to all industries. Second, we apply these estimates to the three support conditions derived in Section 2.5 to identify which industries would receive support from each stakeholder group.

Estimating industry-level synergies and learning. We estimate industry-specific parameters by applying the methodology from Section 4 separately to each three-digit NAICS industry. For each industry p, we estimate the following regressions using only firms (indexed by j) in that industry:

$$y_{j,t} = \widehat{\gamma}_p^{cus}\overline{\Delta}_{j,t}^{cus} + \beta_p z_{j,t} + \eta_p \times \overline{decile}\left(z_{j,t}^{cus}\right) + firm_j + \chi_t + \epsilon_{j,t},$$

$$y_{j,t} = \widehat{\gamma}_p^{sup}\overline{\Delta}_{j,t}^{sup} + \beta_p z_{j,t} + \eta_p \times \overline{decile}\left(z_{j,t}^{sup}\right) + firm_j + \chi_t + \epsilon_{j,t},$$

$$z_{j,t+5} - z_{j,t} = \widehat{\phi}_p^{cus}\widetilde{\Delta}_{j,t}^{cus} + \beta_p z_{j,t} + firm_j + \chi_t + \epsilon_{j,t},$$

$$z_{j,t+5} - z_{j,t} = \widehat{\phi}_p^{sup}\widetilde{\Delta}_{j,t}^{sup} + \beta_p z_{j,t} + firm_j + \chi_t + \epsilon_{j,t}.$$

Variables and instrumental variables are defined as in Section 5. Coefficients that are not statistically significant at the 10% level are set to zero.

For each industry, we construct aggregate parameters by combining customer- and supplier-side estimates. The overall synergy parameter is $\gamma_p = -[\omega_p \widehat{\gamma}_p^{cus} + (1-\omega_p)\widehat{\gamma}_p^{sup}]$ and the overall learning rate is $\phi_p = \omega_p \widehat{\phi}_p^{cus} + (1-\omega_p)\widehat{\phi}_p^{sup}$, where the weight ω_p represents industry p's downstream orientation in the supply chain, measured as the fraction of inter-sectoral sales to downstream industries using U.S. input-output tables.

Calibrating industry size and monopsony power. Industry size χ_p is measured as the industry's share of total value added in the U.S. economy, averaged over 2010–2020 using data from the BEA. The size of the rest of the economy is $\chi_{-p} = 1 - \chi_p$. Monopsony power κ_p is measured as the industry p's profitability using BEA data, as in equation (31).

Evaluating support conditions. Using the estimated parameters for each industry, we evaluate the three conditions from Section 3:

1. **Social planner support:** Equation (30) requires $(1-\phi_p)\gamma_p < \phi_p \chi_{-p}$.
2. **Sector B (rest of economy) support:** Equation (26) requires $\frac{(1-\kappa_p)\chi_p}{\beta\psi} + (1-\phi_p)\gamma_p < \phi_p \chi_{-p}$.
3. **Sector A (subsidized industry) support:** Equation (27) requires $(1-\phi_p)\gamma_p < \phi_p \chi_{-p} + \frac{(1-\kappa_p)\chi_{-p}}{\beta\psi}$.

6.1 Results

We present our main results in three figures that visualize the support conditions across all three-digit NAICS industries. Each figure plots the difference between the right-hand side and left-hand side of the corresponding support condition (equations 30, 26, and 27). A positive value indicates support for industrial policy, a negative value indicates opposition, and zero represents indifference.

Figure 4 presents social planner support across all three-digit NAICS industries. Only a minority of industries yield positive net benefits from the social planner's perspective. The supported industries include: Fabricated Metal Product Manufacturing (332), Transportation Equipment Manufacturing (336), Machinery Manufacturing (333), Computer and Electronic Product Manufacturing (334), Primary Metal Manufacturing (331), Support Activities for Mining (213), Petroleum and Coal Products Manufacturing (324), Chemical Manufacturing (325), and Miscellaneous Manufacturing (339).

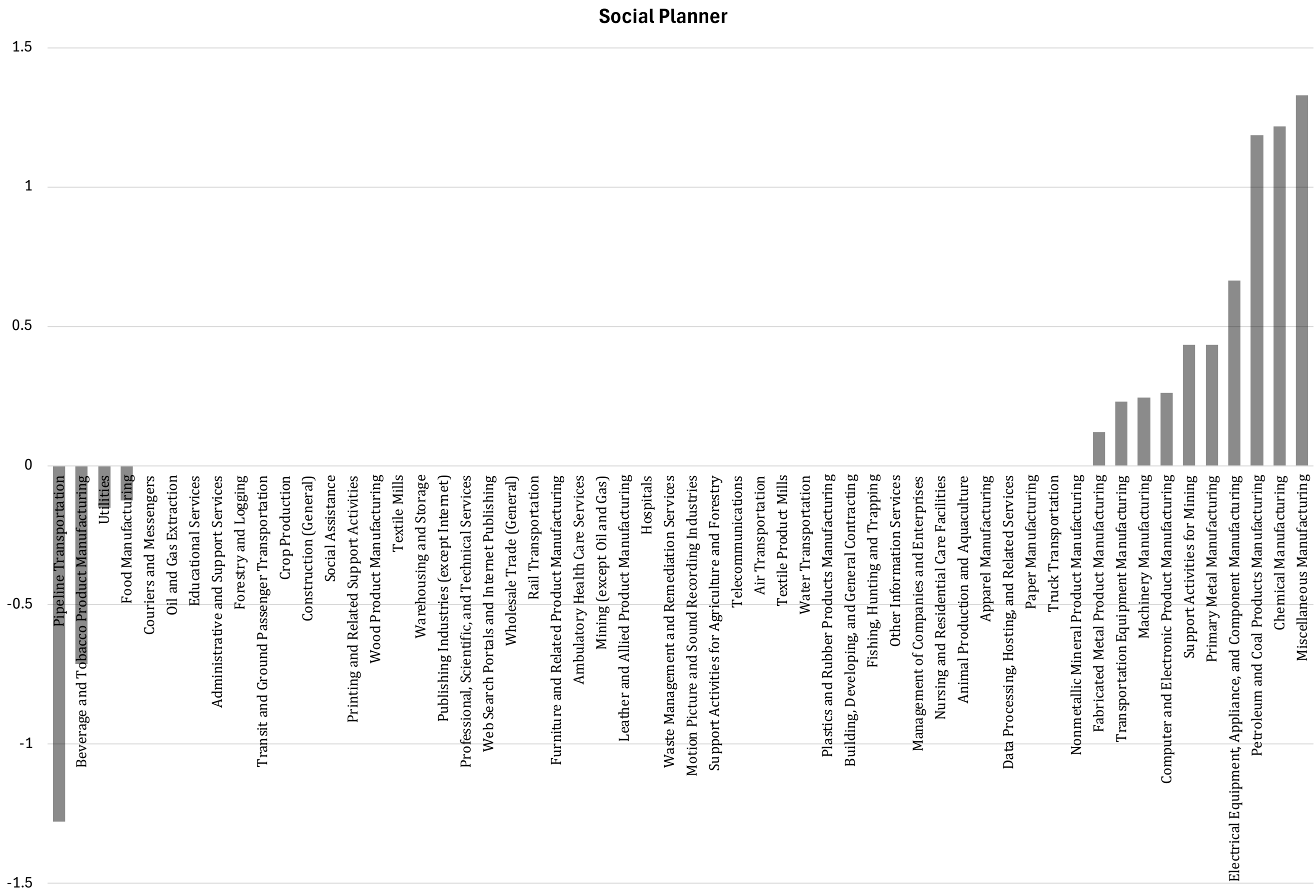

Figure 4: Social planner support for industrial policy across three-digit NAICS industries. The vertical axis shows the net social benefit $\phi_p \chi_{-p} - (1 - \phi_p)\gamma_p$ from subsidizing each industry (equation 30). Positive values indicate industries that would receive support from a welfare-maximizing social planner, as inter-sectoral learning spillovers exceed synergy costs.

These industries share common characteristics: they exhibit strong inter-sectoral learning spillovers (ϕ_p is high) that generate positive externalities for the rest of the economy, while their synergy costs (γ_p) are sufficiently low.

Figure 5 displays support from workers in the rest of the economy (sector B). The pattern of support closely mirrors the social planner's preferences but is more restrictive. This is because these industries that receive social planner's support are predominantly those that are relatively small (χ_p is low), which minimizes the immediate tax burden on sector B workers and receive their support. Hypothetically, for much larger industries, even when subsidies would be welfare-improving, sector B workers oppose them due to the substantial tax costs they would bear. Notably, there exists a broad range of industries for which the social planner is indifferent (net benefit near zero) because estimated synergies or inter-sectoral learning effects

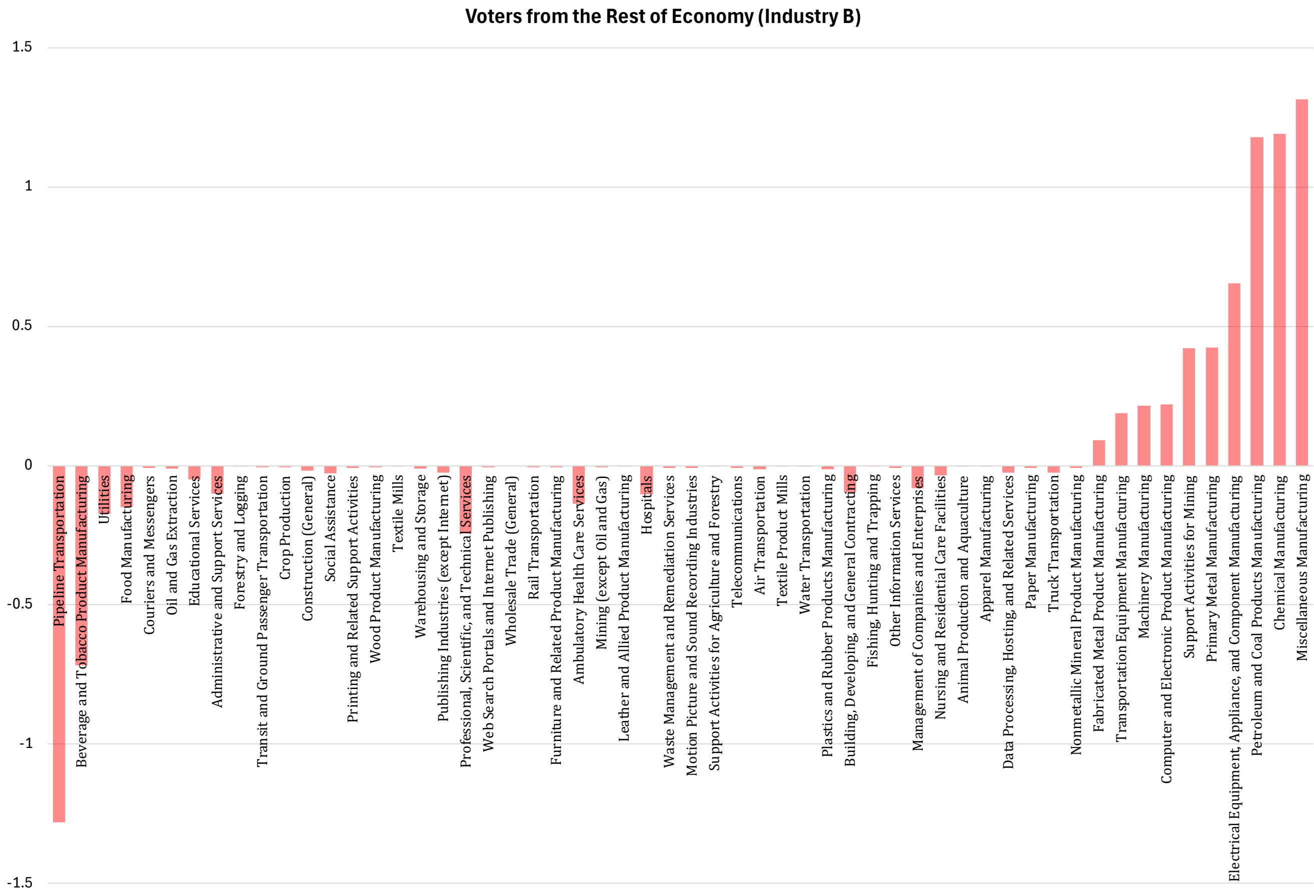

Figure 5: Support for industrial policy from workers in the rest of the economy (Sector B) across three-digit NAICS industries. The vertical axis shows the net benefit $\phi_p \chi_{-p} - \left[\frac{(1-\kappa_p)\chi_p}{\beta\psi} + (1-\phi_p)\gamma_p\right]$ to workers in the rest of the economy (equation 26). Positive values indicate industries that would receive sector B support.

are not statistically significant. Despite the absence of significant social benefits or costs, sector B workers oppose subsidies to these industries, reflecting their aversion to bearing taxes without clear compensatory spillovers.

Our finding that sector B workers would support industrial policies targeting Transportation Equipment, Computer and Electronic Product Manufacturing, and Primary Metals is consistent with the historical trajectory of U.S. industrial policy. The Transportation Equipment sector has received sustained government support for over seven decades. The aerospace industry benefited from Department of Defense procurement contracts beginning in the 1950s and NASA's establishment in 1958, while the automotive sector received emergency support through the Emergency Economic Stabilization Act of 2008 and the American Recovery and Reinvestment Act of 2009 (Klier and Rubenstein, 2012). More recently, the Inflation Reduction Act of 2022 provided substantial tax credits for electric vehicle production and adoption (IRA, U.S. Congress, 2022).

The Computer and Electronic Product Manufacturing sector, particularly semiconductors, similarly emerged from government intervention. Military procurement contracts under the Defense Production Act of 1950 and Department of Defense research funding in the 1960s–1970s cultivated the early semiconductor industry (Flamm, 1996; Mowery, 1999). Contemporary policy has reinforced this pattern: the CHIPS and Science Act of 2022 allocated $52.7 billion for semiconductor manufacturing subsidies and R&D (CHIPS, U.S. Congress, 2022).

Primary Metal Manufacturing has received protection through Section 232 national security tariffs imposed in 2018 under the Trade Expansion Act of 1962, which levied 25% tariffs on steel and 10% on aluminum imports (Bown, 2021). "Buy American" provisions, codified in the Buy American Act of 1933 and strengthened through subsequent executive orders, have further sheltered this sector (McGuinness and Manuel, 2017). Chemical Manufacturing benefits from strong intellectual property protection under the Bayh-Dole Act of 1980 and receives indirect support through defense and energy programs (Mowery et al., 2004). Even Petroleum and Coal Products have received support through tax provisions in the Energy Policy Act of 2005 and the Strategic Petroleum Reserve established in 1975 (Bordoff and Losz, 2020).

This historical pattern demonstrates that the sectors our model identifies as generating substantial inter-sectoral learning spillovers while imposing limited synergy costs are precisely those that have attracted sustained policy intervention. This empirical consistency validates our theoretical framework's ability to rationalize observed industrial policy choices through the political economy lens of voter preferences and welfare effects.

Figure 6 illustrates support from workers in potentially subsidized industries (sector A). As anticipated, sector A workers exhibit near-universal support for industrial policy targeting their own industries. The only exceptions occur in rare cases where synergy costs are exceptionally high and inter-sectoral learning is negative or negligible. The additional term $\frac{(1-\kappa_p)\chi_{-p}}{\beta\psi}$ in equation (27) represents the indirect benefit to sector A workers from reduced monopsony power in the rest of the economy, further bolstering their support.

The divergence between sector A's support (near-universal) and sector B's support (highly selective) reveals a fundamental tension in the political economy of industrial policy: direct beneficiaries strongly favor subsidies regardless of aggregate welfare effects, while the taxed majority requires demonstrable economy-wide spillovers to lend support. This tension implies that welfare-improving industrial policies may nevertheless fail to achieve democratic majority sup-

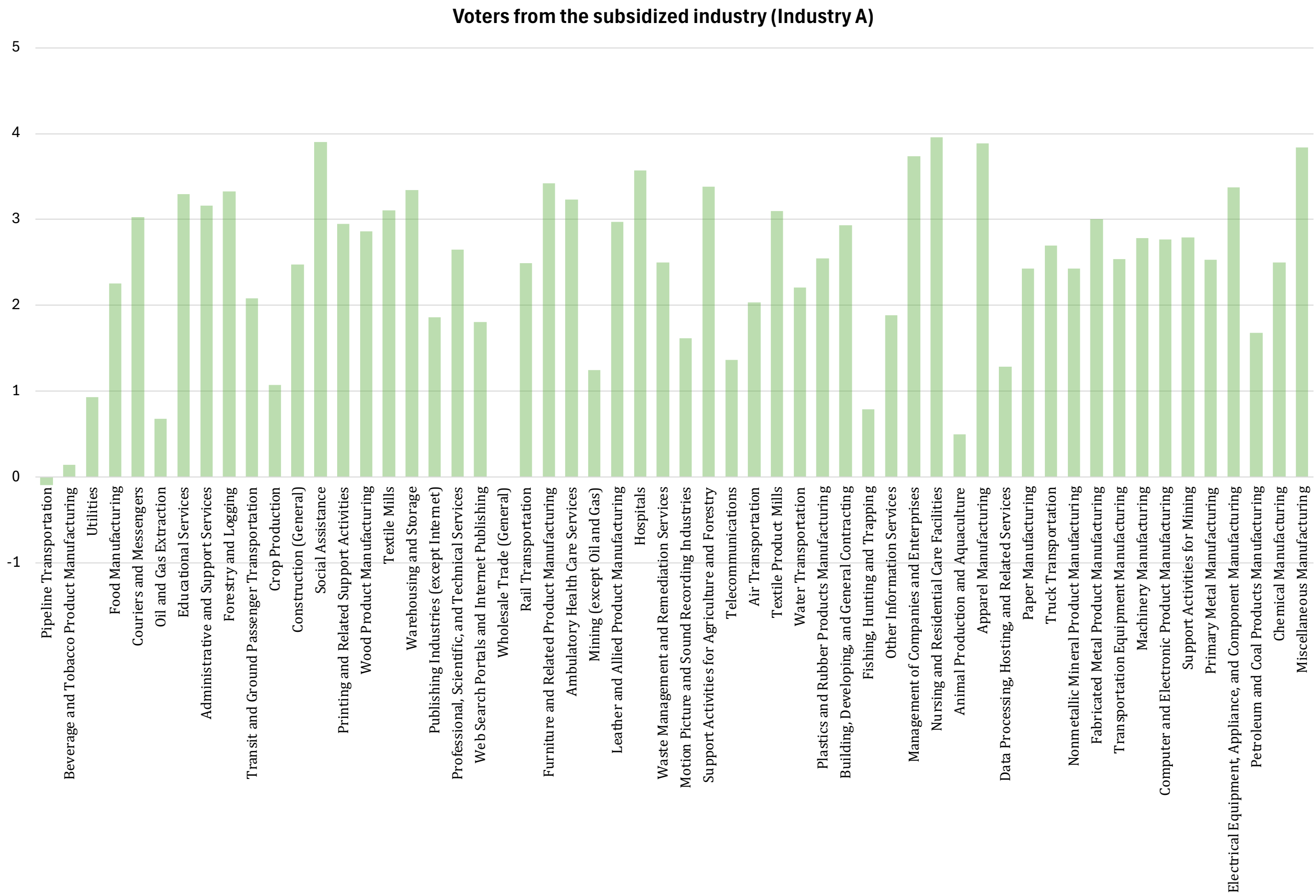

Figure 6: Sector A worker support for industrial policy across three-digit NAICS industries. The vertical axis shows the net benefit $\phi_p \chi_{-p} + \frac{(1-\kappa_p)\chi_{-p}}{\beta\psi} - (1-\phi_p)\gamma_p$ to workers in the subsidized industry (equation 27). Positive values indicate support.

port, and that politically viable policies must generate—and clearly communicate—substantial spillover benefits.

7 Conclusion

References

Acemoglu, D., Autor, D., and Patterson, C. (2024). Bottlenecks: Sectoral imbalances and the us productivity slowdown. *NBER Macroeconomics Annual*, 38(1):153–207.

Acemoglu, D. and Robinson, J. A. (2008). Persistence of power, elites, and institutions. *American Economic Review*, 98(1):267–293.

Amiti, M., Duprez, C., Konings, J., and Van Reenen, J. (2023). Fdi and superstar spillovers: Evidence from firm-to-firm transactions. Working Paper 31128, National Bureau of Economic Research.

Barrios, S. and Strobl, E. (2004). Learning by doing and spillovers: Evidence from firm-level panel data. *Review of Industrial Organization*, 25(2):175–203.

Bartelme, D., Costinot, A., Donaldson, D., and Rodriguez-Clare, A. (2025). The textbook case for industrial policy: Theory meets data. *Journal of Political Economy*, 133(5):1527–1573.

Bernstein, J. I. and Nadiri, M. I. (1988). Interindustry r&d spillovers, rates of return, and production in high-tech industries. *American Economic Review*, 78(2):429–434.

Bloom, N., Schankerman, M., and Van Reenen, J. (2013). Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4):1347–1393.

Bordoff, J. and Losz, A. (2020). The strategic case for the us strategic petroleum reserve. *Columbia SIPA Center on Global Energy Policy*.

Bown, C. P. (2021). Us-china phase one tracker: China's purchases of us goods. *Peterson Institute for International Economics*.

Chenery, H. B. (1960). Patterns of industrial growth. *American Economic Review*, 50(4):624–654.

CHIPS, U.S. Congress (2022). CHIPS and science act of 2022. Public Law 117-167. 117th Congress.

Fernandez, R. and Rodrik, D. (1991). Resistance to reform: Status quo bias in the presence of individual-specific uncertainty. *American Economic Review*, 81(5):1146–1155.

Fernández-Villaverde, J., Mandelman, F., Yu, Y., and Zanetti, F. (2025a). Search complementarities, aggregate fluctuations, and fiscal policy. *Review of Economic Studies*, 92(4):2502–2536.

Fernández-Villaverde, J., Mandelman, F., and Zanetti, F. (2021). The "matthew effect" and market concentration: Search complementarities and monopsony power. *Journal of Monetary Economics*, 121:62–90.

Fernández-Villaverde, J., Yu, Y., and Zanetti, F. (2025b). Technological synergies, heterogeneous firms, and uncertainty shocks. Nber working paper, National Bureau of Economic Research. NBER Working Paper.

Flamm, K. (1996). *Mismanaged trade? Strategic policy and the semiconductor industry*. Brookings Institution Press.

Griffith, R., Redding, S., and Van Reenen, J. (2004). Mapping the two faces of r&d: Productivity growth in a panel of oecd industries. *Review of Economics and Statistics*, 86(4):883–895.

IRA, U.S. Congress (2022). Inflation reduction act of 2022. Public Law 117-169, United States Government.

Irwin, D. A. and Klenow, P. J. (1994). Learning-by-doing spillovers in the semiconductor industry. *Journal of Political Economy*, 102(6):1200–1227.

Jones, C. I. (2011). Intermediate goods and weak links in the theory of economic development. *American Economic Journal: Macroeconomics*, 3(2):1–28.

Jones, C. I. and Tonetti, C. (2025). Past automation and future ai: How weak links tame the growth explosion. Technical report, Working paper, Stanford GSB and NBER.

Juhász, R. and Lane, N. (2024). The political economy of industrial policy. *Journal of Economic Perspectives*, 38(4):27–54.

Keller, W. (2002). Geographic localization of international technology diffusion. *American Economic Review*, 92(1):120–142.

Klier, T. and Rubenstein, J. (2012). Did the cash for clunkers program help the economy? *Economic Perspectives*, 36(4):40–51.

Kong, X. and Dhyne, E. (2025). Learning in firm-to-firm trade.

Leighley, J. E. and Nagler, J. (2013). *Who votes now? Demographics, issues, inequality, and turnout in the United States*. Princeton University Press.

Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *The review of economic studies*, 70(2):317–341.

Liu, E. (2019). Industrial policies in production networks. *The Quarterly journal of economics*, 134(4):1883–1948.

McGuinness, K. M. and Manuel, K. M. (2017). Buy american laws: A legislative and legal overview. *Congressional Research Service Report*.

Mowery, D. C. (1999). *US industry in 2000: Studies in competitive performance*. National Academies Press.

Mowery, D. C., Nelson, R. R., Sampat, B. N., and Ziedonis, A. A. (2004). *Ivory tower and industrial innovation: University-industry technology transfer before and after the Bayh-Dole Act*. Stanford University Press.

Pack, H. and Westphal, L. E. (1994). Industrial policy: Growth elixir or poison? In *The World Bank Research Observer*, volume 9, pages 59–84. Oxford University Press.

Stiglitz, J. E., Lin, J. Y., and Patel, E. (2013). Industrial policy, learning, and development. In Stiglitz, J. E. and Lin, J. Y., editors, *The Practice of Industrial Policy: Government–Business Coordination in Africa and East Asia*, pages 23–39. Oxford University Press, Oxford.

Wooldridge, J. M. (2009). On estimating firm-level production functions using proxy variables to control for unobservables. *Economics letters*, 104(3):112–114.

Xu, L., Yu, Y., and Zanetti, F. (2025). The adoption and termination of suppliers over the business cycle. *Journal of Monetary Economics*, 151:103730.

Appendix

A Sectoral Symmetry Under the Population-Income Share Normalization

In the main text, we adopt the normalization $N_i = \chi_i$, equating each sector's population—the number of firms and workers—to its income share in final goods production. This appendix formally demonstrates that this normalization ensures identical per capita outcomes across sectors in both periods in the absence of industrial policy ($\tau_A = \tau_B = 0$). This result justifies using the no-policy equilibrium as a symmetric benchmark in which workers would have no incentive to move across sectors even if labor mobility were permitted.

A.1 Period 1

Sectoral output is given by the aggregation of firm-level output:

$$y_{i,1} = z_1 l_{i,1}^{1-\alpha} k_{i,1}^{\alpha},$$

where $l_{i,1} = N_i$ and $k_{i,1} = N_i k_{i,j,1}$. Substituting these expressions yields:

$$y_{i,1} = z_1 N_i^{1-\alpha} (N_i k_{i,j,1})^{\alpha} = z_1 N_i k_{i,j,1}^{\alpha}. \tag{A-1}$$

Firm-level capital is $k_{i,j,1} = \alpha(1-\tau_i)\chi_i y_1 / N_i$. With no taxes ($\tau_i = 0$) and $\chi_i = N_i$, this becomes $k_{i,j,1} = \alpha y_1$. Sectoral output is therefore: $y_{i,1} = z_1 N_i (\alpha y_1)^{\alpha}$. Sectoral output per capita is:

$$\frac{y_{i,1}}{N_i} = z_1 (\alpha y_1)^{\alpha}. \tag{A-2}$$

As already shown, capital per capita in sector i is:

$$\frac{k_{i,1}}{N_i} = k_{i,j,1} = \alpha y_1. \tag{A-3}$$

For consumption per capita, labor income per capita in sector i is $(1-\tau_i)(1-\kappa)(1-\alpha)\chi_i y_1 / N_i$. With no taxes and $N_i = \chi_i$, this becomes $(1-\kappa)(1-\alpha)y_1$. Dividend income per capita is $\kappa(1-\alpha)y_1 \sum_s \chi_s (1-\tau_s)$, which simplifies to $\kappa(1-\alpha)y_1$. Consumption per capita is therefore:

$$\frac{c_{i,1}}{N_i} = (1-\kappa)(1-\alpha)y_1 + \kappa(1-\alpha)y_1 = (1-\alpha)y_1. \tag{A-4}$$

Equation (A-2)-(A-4) show that output, capital, and consumption per capita are identical across sectors in period 1.

A.2 Period 2

Given the identical per capita capital across sectors in period 1 ($k_{A,1}/N_A = k_{B,1}/N_B$), the within-sector learning equation (9) yields identical within-sector productivity in period 2:

$$z_{A,2} = z_{B,2}. \tag{A-5}$$

With the same derivation as in period 1, we can show that output, capital, and consumption per capita are identical across sectors in period 2.

The normalization $N_i = \chi_i$ delivers a clean symmetric benchmark. In the absence of industrial policy ($\tau_A = \tau_B = 0$), per capita capital, output, and consumption are identical across sectors in period 1. Because per capita capital is identical in period 1, within-sector learning yields identical productivity across sectors in period 2. With no productivity gap, inter-sectoral learning spillovers are absent, and final productivity is identical, resulting in equal per capita consumption across sectors. This symmetry implies that workers in both sectors achieve the same lifetime utility in the absence of policy, providing a natural reference point from which to evaluate the welfare effects of industrial policy.

B Sector-specific rate of learning by doing

We now extend our baseline model to allow for sector-specific within-sector learning-by-doing parameters, ψ_A and ψ_B. This captures the empirical reality that different industries may have different capacities for learning from their own production experience.

B.1 Modified Productivity Dynamics

Under this extension, the within-sector learning equations become sector-specific:

$$\ln(\widehat{z}_{A,2}) = \psi_A \left[\ln(1-\tau_A) + \frac{\ln(\alpha) + \sum_i \chi_i \ln(\chi_i)}{1-\alpha} \right], \tag{A-6}$$

$$\ln(\widehat{z}_{B,2}) = \psi_B \left[\ln(1-\tau_B) + \frac{\ln(\alpha) + \sum_i \chi_i \ln(\chi_i)}{1-\alpha} \right]. \tag{A-7}$$

The inter-sectoral spillover equation remains unchanged:

$$\ln(z_{i,2}) = \ln(\widehat{z}_{i,2}) + \phi \max\{\ln(\widehat{z}_{-i,2}) - \ln(\widehat{z}_{i,2}), 0\}. \tag{A-8}$$

Given that $\tau_A < 0$ and $\tau_B > 0$, we still have $\widehat{z}_{A,2} > \widehat{z}_{B,2}$, so sector B learns from sector A:

$$\ln(z_{A,2}) = \ln(\widehat{z}_{A,2}) = \psi_A \left[\ln(1-\tau_A) + C\right], \tag{A-9}$$

$$\begin{aligned} \ln(z_{B,2}) &= (1-\phi)\ln(\widehat{z}_{B,2}) + \phi \ln(\widehat{z}_{A,2}) \\ &= (1-\phi)\psi_B \left[\ln(1-\tau_B) + C\right] + \phi\psi_A \left[\ln(1-\tau_A) + C\right], \end{aligned} \tag{A-10}$$

where $C = \frac{\ln(\alpha)+\sum_i \chi_i \ln(\chi_i)}{1-\alpha}$.

B.2 Modified Aggregate Output and Voting Conditions

The aggregate output in period 2 is:

$$
\begin{aligned}
\ln(y_2) &= -\gamma\left[\ln(z_{A,2}) - \ln(z_{B,2})\right] + \chi_A \ln(z_{A,2}) + \chi_B \ln(z_{B,2}) \\
&= -\gamma(1-\phi)\left[\psi_A \ln(1-\tau_A) - \psi_B \ln(1-\tau_B)\right] \\
&\quad + \chi_A \psi_A \ln(1-\tau_A) + \chi_B\left[(1-\phi)\psi_B \ln(1-\tau_B) + \phi\psi_A \ln(1-\tau_A)\right] + \text{constant}. \qquad \text{(A-11)}
\end{aligned}
$$

The lifetime utilities for the two sectors are modified accordingly. Using the approximations $\ln(1-\tau_i) \approx -\tau_i$ and $\ln(1-\tau_i+\tau_i\kappa) \approx -\tau_i(1-\kappa)$ for small τ_i, and the government budget constraint $\chi_A\tau_A + \chi_B\tau_B = 0$, we obtain the following modified voting conditions:

B.2.1 Support Conditions with Heterogeneous Learning Rates

The voting conditions for sector A (subsidized) and sector B (taxed), along with the social planner's welfare criterion, can be expressed in comparable forms that highlight how relative learning rates modify the political calculus. When learning rates differ ($\psi_A \neq \psi_B$), the conditions become:

Sector B (taxed majority) supports policy if:

$$\frac{(1-\kappa)\chi_A}{\beta\psi_B} + (1-\phi)\gamma \ < \ \phi\chi_B \cdot \frac{\psi_A}{\psi_B}. \qquad \text{(A-12)}$$

Sector A (subsidized sector) supports policy if:

$$(1-\phi)\gamma \ < \ \phi\chi_B \cdot \frac{\psi_A}{\psi_B} \ + \ \frac{(1-\kappa)\chi_B}{\beta\psi_B}. \qquad \text{(A-13)}$$

Social planner supports policy if:

$$\psi_B(1-\phi)(\chi_B+\gamma)\chi_A \ < \ \psi_A\chi_B\left[\chi_A + \phi\chi_B - (1-\phi)\gamma\right]. \qquad \text{(A-14)}$$

These conditions reveal how relative learning capacity ψ_A/ψ_B enters the political economy of industrial policy.

First, spillover benefits are scaled by relative learning rates. The right-hand side of inequalities (A-12) and (A-13) includes the term $\phi\chi_B \cdot (\psi_A/\psi_B)$. When the subsidized sector learns faster ($\psi_A/\psi_B > 1$), the perceived spillover benefit grows, making support more likely. When it learns slower ($\psi_A/\psi_B < 1$), the benefit shrinks.

Second, immediate costs are weighted by the taxed sector's learning rate. The tax-burden term $(1-\kappa)\chi_A/(\beta\psi_B)$ in condition (A-12) is divided by ψ_B. A higher ψ_B reduces this cost term, reflecting the fact that taxing a fast-learning sector imposes a larger opportunity cost in forgone learning.

Third, dividend benefits are also weighted by the taxed sector's learning rate. The additional benefit term in condition (A-13) is divided by ψ_B. This captures how the taxed sector's learning capacity influences the dividend channel.

Fourth, the social planner's condition is more complex, incorporating both learning rates through the product terms on both sides. When $\psi_A = \psi_B = \psi$, condition (A-14) simplifies to $(1-\phi)\gamma < \phi\chi_B$, which matches the main text.

When $\psi_A = \psi_B = \psi$, the conditions reduce exactly to their symmetric counterparts in the main text:

$$\frac{(1-\kappa)\chi_A}{\beta\psi} + (1-\phi)\gamma < \phi\chi_B, \qquad (1-\phi)\gamma < \phi\chi_B + \frac{(1-\kappa)\chi_B}{\beta\psi}, \qquad (1-\phi)\gamma < \phi\chi_B.$$

The ordering of restrictiveness remains: sector A workers support industrial policy over the broadest parameter range, the social planner's condition is intermediate, and sector B workers require the strongest evidence of net benefits. However, the exact thresholds now depend on the ratio ψ_A/ψ_B, which can significantly expand or contract the politically feasible parameter space.

B.3 Implications

The introduction of sector-specific learning rates modifies the analysis of industrial policy in several meaningful ways.

First, the relative magnitudes of learning rates now shape the conditions for political support. The decision to subsidize a sector no longer depends solely on its absolute learning capacity, but rather on how its learning rate compares to that of the taxed sector. A sector with a high ψ_p relative to ψ_{-p} is more likely to gain support from both the social planner and workers in the taxed sector, because the productivity gains from subsidizing it are larger relative to the opportunity cost of taxing a sector that could otherwise have learned rapidly on its own.

Second, changes in learning rates have asymmetric effects on political support. An increase in ψ_p makes industrial policy more appealing to workers in the taxed sector—the right-hand side of condition (A-12) rises—but it also makes the policy slightly less attractive to workers in the subsidized sector, because the right-hand side of (A-13) rises as well, tightening the inequality.

In practice, however, the direct beneficiaries in sector p almost always favor subsidies, so this second effect is usually minor. Conversely, an increase in ψ_{-p} reduces the appeal of industrial policy to taxed workers—the left-hand side of (A-12) grows—while making it more attractive to subsidized workers, as the left-hand side of (A-13) falls.

Third, applying this extension empirically would require estimating within-sector learning rates industry by industry. One could, for example, regress productivity growth on lagged measures of production scale or capital input at the industry level, using instrumental variables to address endogeneity concerns. Such estimates would help identify which sectors possess high intrinsic learning capacity.

Finally, the policy implications become more nuanced. Industries that combine strong within-sector learning (ψ_p) with strong inter-sectoral spillovers (ϕ) are especially compelling targets for industrial policy, both in terms of economic efficiency and political feasibility. However, even a sector with high ψ_p may struggle to gain political support if the rest of the economy learns even faster ($\psi_{-p} > \psi_p$), because the opportunity cost of taxation is then prohibitively high.

B.4 Numerical Illustration of Heterogeneous Learning Rates

To illustrate how heterogeneous learning rates affect the political feasibility of industrial policy, we simulate the support condition for sector B—the taxed majority—across three scenarios: 1. $\psi_p = 0.4 > \psi_{-p} = 0.2$ (the subsidized sector learns faster); 2. $\psi_p = \psi_{-p} = 0.25$ (the symmetric baseline, matching our main analysis). 3. $\psi_p = 0.2 < \psi_{-p} = 0.4$ (the taxed sector learns faster). We hold other parameters at their baseline values as in the original calibration.

Figure B.1 displays the parameter space for sector B support in the first scenario, where the subsidized sector has a higher learning rate. The left panel varies the subsidized sector's size (χ) and monopsony power (κ), holding inter-sectoral learning fixed at baseline value; the right panel varies χ and ϕ, holding κ fixed at baseline value. Black regions indicate combinations where sector B supports the policy; white regions indicate opposition.

Three patterns stand out. First, sector size matters less than in the symmetric benchmark: the black region extends further to the right along the χ-axis. When the subsidized sector learns quickly, the taxed majority is willing to accept a larger sector because the spillover benefits compensate for the higher tax burden. Second, greater monopsony power (κ) expands the

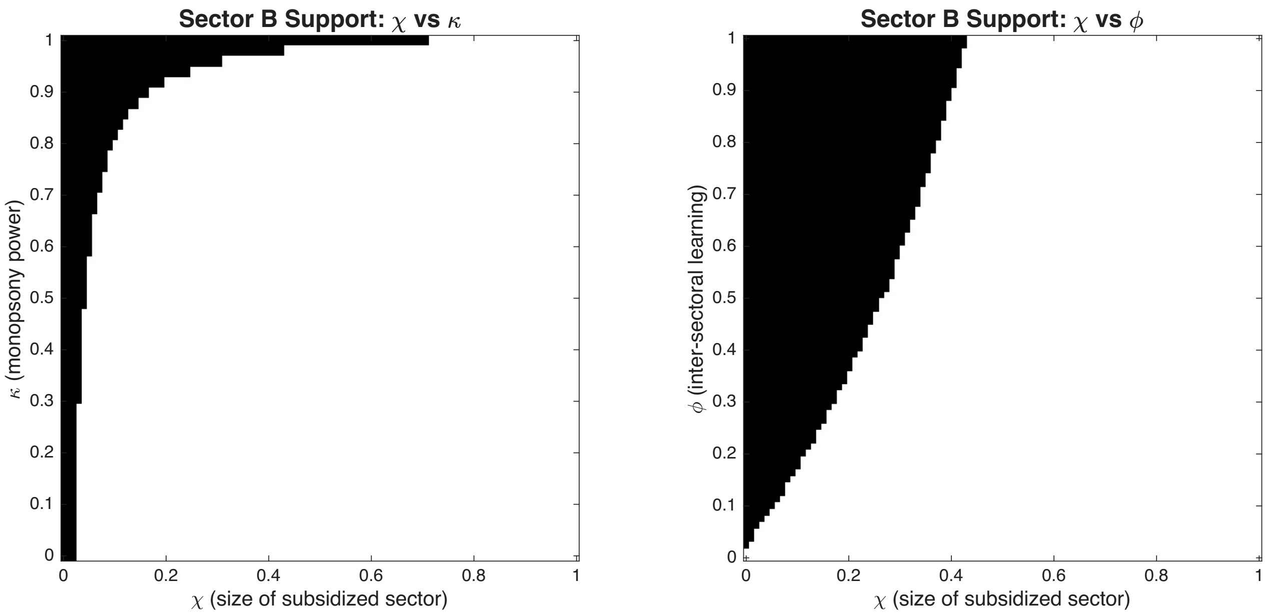

Figure B.1: Parameter space for sector B support when $\psi_p = 0.4 > \psi_{-p} = 0.2$. Black regions indicate support; white regions indicate opposition. Left: Variation in χ (subsidized sector size) and κ (monopsony power). Right: Variation in χ and ϕ (inter-sectoral learning).

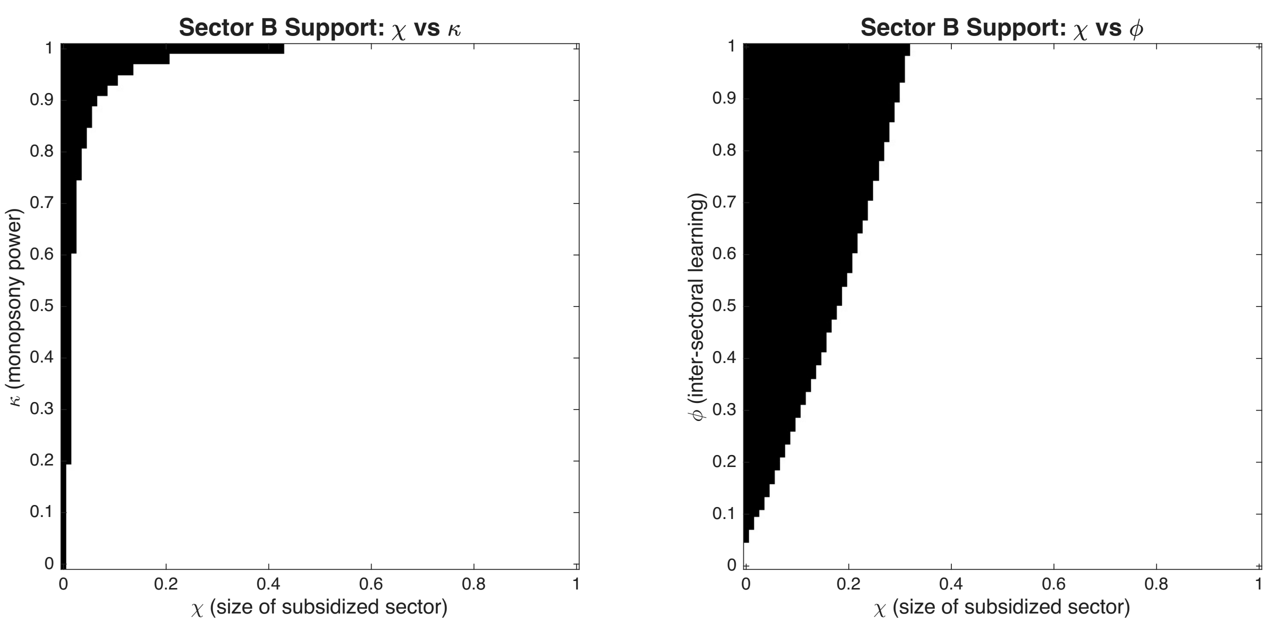

Figure B.2: Parameter space for sector B support in the symmetric case ($\psi_p = \psi_{-p} = 0.25$), matching our baseline analysis.

support region, as before, because taxed workers benefit from dividend income. Third, inter-sectoral learning (ϕ) remains important—higher ϕ enlarges the support region—but even at moderate values ($\phi \approx$ 0.3–0.4) substantial support exists when $\psi_p > \psi_{-p}$.

Figure B.2 shows the symmetric baseline, which reproduces exactly the parameter space from the main analysis (Figure 1). This provides a benchmark for the asymmetric cases.

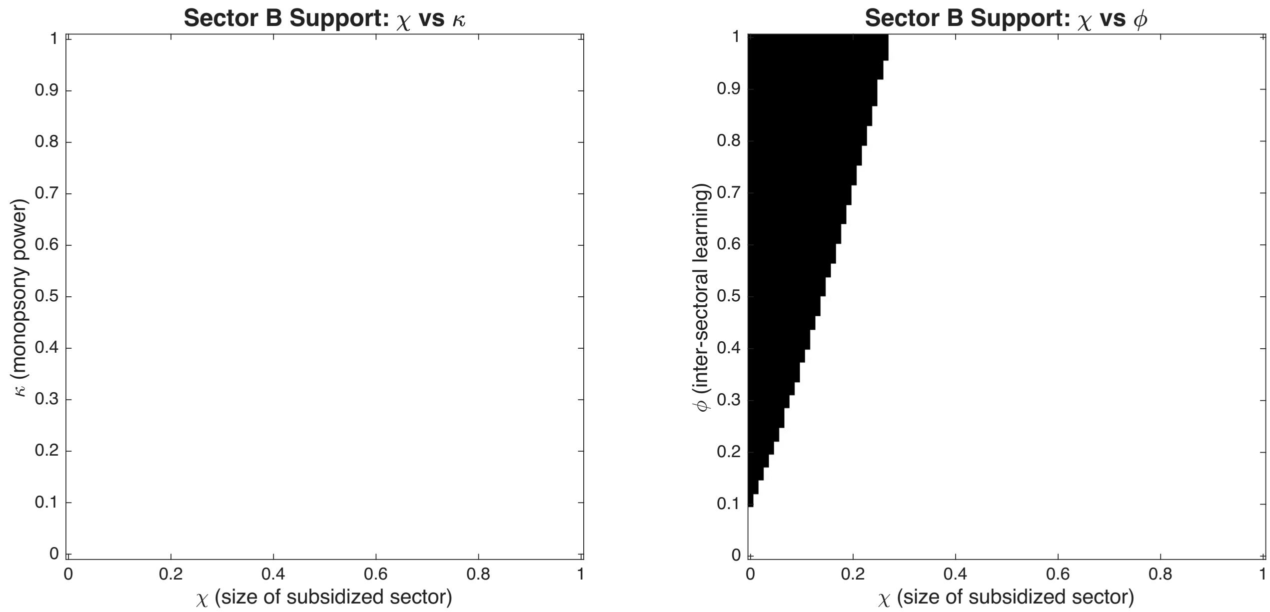

Figure B.3: Parameter space for sector B support when $\psi_p = 0.2 < \psi_{-p} = 0.4$. Support (dark region) is substantially more restrictive than in Scenarios 1 and 2.

Figure B.3 presents the third scenario, where the taxed sector learns faster. The support region shrinks dramatically. Sector B now backs the policy only for very small subsidized sectors, because taxing a fast-learning sector entails a high opportunity cost. Monopsony power still helps, but its effect is muted; the dominant factor is the loss of own-sector learning. Moreover, support emerges only at very high inter-sectoral learning rates, implying that nearly perfect knowledge transfer is needed to justify taxing a rapidly learning sector.

Together, these numerical results underscore the central role of relative learning rates. When the subsidized sector learns faster ($\psi_p > \psi_{-p}$), taxed workers view the policy as an investment in future productivity: the forgone learning in their own sector is outweighed by direct spillovers, indirect gains from a more productive economy, and dividend income. When the taxed sector learns faster ($\psi_p < \psi_{-p}$), industrial policy becomes a hard political sell, because taxed workers bear a triple burden—immediate taxes, lost learning opportunities, and limited spillovers from a weaker sector.

While allowing for sector-specific learning rates enriches the theoretical framework and better captures real-world heterogeneity, the core political-economy tension remains: workers in the subsidized sector almost always favor subsidies, while workers in the taxed sector require clear evidence of substantial spillover benefits. The threshold for those benefits, however, now depends critically on the relative learning capacities of the two sectors. Future empirical work that estimates industry-specific ψ parameters could therefore yield more nuanced policy recommendations.

C Allowing for Household Saving and Borrowing

In the main text, we abstracted from household saving and borrowing decisions to focus on the redistributive and productivity channels of industrial policy. In this appendix, we relax this assumption and allow households to smooth consumption across periods by saving or borrowing at a common real interest rate r.

C.1 Household Problem

The household in sector i solves:

$$\max_{c_{i,1},c_{i,2}} \ln(c_{i,1}) + \beta \ln(c_{i,2})$$

subject to the intertemporal budget constraint:

$$c_{i,1} + \frac{c_{i,2}}{1+r} = y_{i,1} + \frac{y_{i,2}}{1+r}.$$

The first-order conditions yield the Euler equation:

$$\frac{c_{i,2}}{c_{i,1}} = \beta(1+r).$$

From the budget constraint and the Euler equation, period-1 consumption is:

$$c_{i,1} = \frac{1}{1+\beta}\left(y_{i,1} + \frac{y_{i,2}}{1+r}\right).$$

Substituting the income expressions $y_{i,1} = (1 - \tau_i + \tau_i\kappa)(1-\alpha)y_1$ and $y_{i,2} = (1-\alpha)y_2$ from the main text gives:

$$c_{i,1} = \frac{1}{1+\beta}\left[(1 - \tau_i + \tau_i\kappa)(1-\alpha)y_1 + \frac{(1-\alpha)y_2}{1+r}\right]. \tag{A-15}$$

Period-2 consumption follows from the Euler equation:

$$c_{i,2} = \beta(1+r)c_{i,1}.$$

C.2 Modified Voting Conditions

With saving and borrowing, lifetime utility can be derived as:

$$U_i = (1+\beta)\ln(c_{i,1}) + \beta \ln(\beta(1+r)),$$

where period-1 consumption is given by equation (A-15).

Households in sector B support industrial policy if $U_B(\tau_A < 0, \tau_B > 0) > U_B(0,0)$. Following the same log-linearization procedure as in the main text, we approximate $\ln(c_{B,1})$ around the no-policy equilibrium ($\tau_A = \tau_B = 0$). Let

$$X_B(\tau) \equiv (1-\tau_B+\tau_B\kappa)(1-\alpha)y_1 + \frac{(1-\alpha)y_2(\tau)}{1+r}$$

denote the term inside the logarithm in the expression for $c_{B,1}$. Log-linearizing $\ln(X_B)$ around $\tau = 0$ gives:

$$\ln(X_B(\tau)) - \ln(X_B(0)) \approx \frac{1}{X_B(0)}\left(\frac{\partial X_B}{\partial \tau_A}\tau_A + \frac{\partial X_B}{\partial \tau_B}\tau_B\right).$$

Using the derivatives of y_2 from the main text, the government budget constraint $\chi_A\tau_A + \chi_B\tau_B = 0$, and the approximation $\ln(1-\tau_i) \approx -\tau_i$, this simplifies to:

$$\frac{(1-\alpha)(1-\kappa)\chi_A}{\psi}\frac{1+r}{g} + (1-\phi)\gamma < \phi\chi_B,$$

where $g \equiv y_2^0/y_1$ is the growth factor of aggregate output in the absence of policy.

To compare this with the benchmark model, note that in the no-policy equilibrium, the Euler equation implies $g = c_{i,2}/c_{i,1} = \beta(1+r)$. Substituting $\beta(1+r) = g$ into the condition above yields:

$$\frac{(1-\alpha)(1-\kappa)\chi_A}{\beta\psi} + (1-\phi)\gamma < \phi\chi_B,$$

which is identical to the benchmark condition (26) from the main text.

Similarly, sector A workers support the policy if:

$$(1-\phi)\gamma < \phi\chi_B + \frac{(1-\alpha)(1-\kappa)\chi_B}{\beta\psi},$$

which is identical to the benchmark condition (27).

The social planner's condition remains unchanged from the main text:

$$(1-\phi)\gamma < \phi\chi_B.$$

Thus, when log-linearized around the no-policy equilibrium, the voting conditions with and without saving are identical. The intuition is straightforward: in the no-policy equilibrium, per capita income is identical across sectors, so households have no incentive to borrow or lend across sectors. The saving/borrowing margin is therefore inactive at the point around which we log-linearize, and the local voting conditions coincide.

D Mapping Firm-Level Estimates to Sector-Level Parameters

The empirical analysis in Section 5 estimates four firm-level coefficients for each three-digit NAICS industry p: synergy parameters $\widehat{\gamma}_p^{cus}$ and $\widehat{\gamma}_p^{sup}$, and learning parameters $\widehat{\phi}_p^{cus}$ and $\widehat{\phi}_p^{sup}$. This appendix provides the theoretical bridge between these firm-level estimates and the sector-level parameters γ and ϕ that enter the aggregate production function (4) and the inter-sectoral learning equation (10).

We start with a set of common assumptions that underpin both mappings.

Shared Assumptions

(A1) **Common distributional shape.** For any two sectors i and j, the log-productivity distributions are identical up to a location shift with common dispersion σ:

$$F_i(\ln z) = G\left(\frac{\ln z - \mu_i}{\sigma}\right), \qquad i = A, B,$$

where G is a baseline CDF with density $g = G'$ and μ_i is the geometric mean (so $\mu_i = \ln Z_i$). This assumption is plausible for mature industries within the same country and period, and it is testable in the data.

(A2) **Local linearity of decile ranks.** For any firm j in sector i, its decile rank satisfies

$$\text{decile}(z_j) \approx 10\, F_i(\ln z_j) \approx 10\, G(0) + \delta\,(\ln z_j - \mu_i),$$

where $\delta \equiv 10\, g(0)/\sigma$. Consequently, for any customer-supplier pair (j, k),

$$\text{decile}(z_j) - \text{decile}(z_k) \;\approx\; \delta\,(\ln z_j - \ln z_k) + \alpha_{pq},$$

where α_{pq} is a pair-specific intercept that absorbs any additive differences in baseline decile levels across industry pairs (p, q).

These two assumptions are sufficient to map the decile-based regressors used in the data to the log-productivity gaps that appear in the theoretical model. The scaling factor δ is estimated directly from the data via the fixed-effects regression described below; we find $\widehat{\delta} \approx 1$, so no further rescaling is required for the point estimates.

D.1 From Firm-Level Sector-Level Synergy Costs

Consider a continuum of customer-supplier pairs (j, k), with $j \in A$, $k \in B$. In the absence of synergy frictions, pair output is $y_{j,k}$; with a productivity-gap penalty, actual output is

$$\ln y^{\text{actual}}_{j,k} = \ln y_{j,k} - \gamma_{\text{firm}}\, |\ln z_j - \ln z_k|, \qquad \gamma_{\text{firm}} \geq 0.$$

Final goods are produced by a log-linear aggregator over all pairs, so

$$\ln Y = \iint \ln y_{j,k}\, d\mu(j,k) \;-\; \gamma_{\text{firm}} \iint |\ln z_j - \ln z_k|\, d\mu(j,k).$$

The first integral aggregates to sectoral outputs:

$$\iint \ln y_{j,k}\, d\mu(j,k) = \chi_A \ln Y_A + \chi_B \ln Y_B + \text{constant}.$$

Using Assumption 2, the absolute log gap is approximately proportional to the absolute decile gap:

$$|\ln z_j - \ln z_k| \;\approx\; \frac{1}{\delta}\, |\text{decile}(z_j) - \text{decile}(z_k)|.$$

Hence,

$$\ln Y = \text{constant} + \chi_A \ln Y_A + \chi_B \ln Y_B - \frac{\gamma_{\text{firm}}}{\delta}\, \mathbb{E}\big[|\text{decile}(z_j) - \text{decile}(z_k)|\big].$$

The empirical regression in Section 5 estimates the effect of the average absolute decile mismatch on $\ln Y$; for the customer side,

$$\widehat{\gamma}^{cus}_p = \frac{\partial \ln Y}{\partial \overline{\Delta}^{cus}} = -\frac{\gamma_{\text{firm}}}{\delta}.$$

Therefore the structural synergy parameter is

$$\gamma_{\text{firm}} = -\delta\, \widehat{\gamma}^{cus}_p = -\delta\, \widehat{\gamma}^{sup}_p,$$

and the industry-level parameter is the weighted average of the two directional estimates (using the downstream orientation weight ω_p):

$$\gamma_p = \omega_p\,(-\widehat{\gamma}_p^{cus}) + (1-\omega_p)\,(-\widehat{\gamma}_p^{sup}).$$

Since $\widehat{\delta} \approx 1$, γ_p is directly the negative of the estimated mismatch coefficient.

D.2 From Firm-Level to Sector-Level Learning Spillover

The inter-sectoral learning equation states that sector B closes a fraction ϕ of the productivity gap with sector A:

$$\ln Z_{B,t+1} = \ln \widehat{Z}_{B,t+1} + \phi\ \max\{\ln Z_{A,t+1} - \ln Z_{B,t+1},\, 0\}.$$

At the firm level, we estimate regressions of the form

$$\Delta\text{decile}(z_k) = \phi^{cus}\ \max\{0,\ \text{decile}(z_j) - \text{decile}(z_k)\} + \text{controls}.$$

To connect ϕ^{cus} to ϕ, we make one additional, mild assumption:

(A3) **Dominance of the sector-level gap.** The sector-level productivity gap $\Delta \equiv \ln Z_A - \ln Z_B$ is the main driver of the decile gap between a customer and a supplier; idiosyncratic deviations are mean-zero and uncorrelated with Δ. Equivalently,

$$\max\{0,\ \text{decile}(z_j) - \text{decile}(z_k)\}\ \approx\ \delta\ \max\{0,\ \ln Z_A - \ln Z_B\} + \varepsilon_{jk},$$

with $\mathbb{E}[\varepsilon_{jk} \mid \Delta] = 0$.

Under this assumption, the firm-level regression identifies

$$\phi^{cus} = \frac{\text{Cov}\big(\Delta\text{decile}(z_k),\ \max\{0, \text{decile gap}\}\big)}{\text{Var}\big(\max\{0, \text{decile gap}\}\big)}.$$

For the average firm in the lagging sector, the change in its decile rank is approximately δ times the change in the sector's log mean, and the positive part of the decile gap is approximately δ

times the positive part of the sector-level gap. Therefore,

$$\phi^{cus} = \phi.$$

Likewise, $\phi^{sup} = \phi$. The industry-level learning parameter is then the weighted average of the two directional estimates:

$$\phi_p = \omega_p \, \widehat{\phi}_p^{cus} + (1 - \omega_p) \, \widehat{\phi}_p^{sup},$$

where ω_p is the fraction of pairs in which industry p acts as a supplier (so that it is the learner in the customer-side regression). This weighting improves precision under the null that both coefficients estimate the same ϕ.

D.3 Estimation of the Scaling Factor

We estimate δ directly from the firm-level data using the following panel regression with industry-pair fixed effects α_{pq}:

$$\Delta \text{decile}_{jk} = \alpha_{pq} + \delta \, \Delta \ln z_{jk} + \varepsilon_{jk}, \qquad (j \in p, \ k \in q).$$

The fixed effects absorb any additive differences across industry pairs, so δ is identified from the within-pair variation in the log-productivity difference. This is precisely the variation that corresponds to the local linearity in Assumption 2. Table D.1 reports the result.

	(1)
Dependent variable	Current decile mismatch (Δ_t)
Log-productivity gap ($\Delta \ln z_{jk}$)	1.000***
	(0.018)
Year FE	Yes
Industry-pair FE	Yes
Observations	26,187
R^2	0.204
Notes: Robust standard errors in parentheses. $^{***}p < 0.01$.	

Table D.1: Direct estimate of the scaling factor δ.

The estimated coefficient is indistinguishable from unity, implying an approximately one-to-one mapping between decile ranks and log-productivity differences in our sample. Consequently, the estimated mismatch coefficients $\widehat{\gamma}^{cus}$ and $\widehat{\gamma}^{sup}$ are directly interpretable as structural synergy parameters in log-productivity units, and no further rescaling is needed for the learning coefficients either.

Robustness to unequal dispersions. If the two sectors have different dispersions ($\sigma_A \neq \sigma_B$), the scaling factor becomes sector-specific: $\delta_i = 10g(0)/\sigma_i$. The synergy mapping then involves a harmonic mean of δ_A and δ_B, which biases the estimated γ by a factor that depends on the dispersion ratio. For learning, the firm-level regression would identify $\phi^{cus} = \phi \cdot 2\sigma_A/(\sigma_A + \sigma_B)$. However, because each industry trades with many partners and we average the two directional estimates, these biases partially cancel across partners and directions. In the empirically relevant case where dispersions are similar, the bias is negligible; our main results are robust to allowing for unequal dispersions, as discussed in the online appendix.

Summary of the Mapping

Under the shared assumptions (common shape, local linearity, and—for learning—sector-gap dominance), the firm-level estimates $\widehat{\gamma}_p^{cus}, \widehat{\gamma}_p^{sup}, \widehat{\phi}_p^{cus}, \widehat{\phi}_p^{sup}$ map directly to sector-level parameters:

$$\gamma_p = -\omega_p \, \widehat{\gamma}_p^{cus} - (1 - \omega_p) \, \widehat{\gamma}_p^{sup}, \qquad \phi_p = \omega_p \, \widehat{\phi}_p^{cus} + (1 - \omega_p) \, \widehat{\phi}_p^{sup}.$$

These are the parameters used in the quantitative exercise.

E Monopsony Power with Elastic Labor Supply

In the main text, we assumed a perfectly inelastic labor supply to highlight the redistributive role of monopsony power. Here, we relax that assumption and adopt the standard textbook monopsony model in which each firm faces an upward-sloping labor supply curve. This introduces a deadweight loss from monopsony, which alters the welfare and political support conditions for industrial policy.

E.1 Model Modification

Each firm j in sector i faces a labor supply schedule

$$w_{i,j} = \eta_i \, l_{i,j}^{\theta}, \qquad \theta > 0,$$

where η_i is a sector-specific supply shifter and θ is the inverse elasticity of labor supply (common across sectors for simplicity). The firm's problem in period 1 becomes

$$\max_{l_{i,j,1},k_{i,j,1}} (1-\tau_i)p_{i,1}y_{i,j,1} - w_{i,j,1}l_{i,j,1} - k_{i,j,1},$$

subject to $y_{i,j,1} = z_1 k_{i,j,1}^{\alpha} l_{i,j,1}^{1-\alpha}$ and $w_{i,j,1} = \eta_i l_{i,j,1}^{\theta}$. The first-order conditions are:

$$(1-\tau_i)p_{i,1}(1-\alpha)\frac{y_{i,j,1}}{l_{i,j,1}} = (1+\theta)w_{i,j,1}, \qquad k_{i,j,1} = \alpha(1-\tau_i)p_{i,1}y_{i,j,1}.$$

The term $(1+\theta) > 1$ represents the monopsony wedge: the marginal revenue product of labor exceeds the wage. The profit of the firm is

$$\pi_{i,j,1} = (1-\tau_i)p_{i,1}y_{i,j,1} - w_{i,j,1}l_{i,j,1} - k_{i,j,1} = \left[1 - \frac{1-\alpha}{1+\theta} - \alpha\right](1-\tau_i)p_{i,1}y_{i,j,1}.$$

Define $\kappa_i \equiv 1 - \frac{1-\alpha}{1+\theta} - \alpha = \frac{\alpha\theta+\theta}{1+\theta}$. Then the profit share of output is κ_i, identical in form to the reduced-form monopsony parameter in the main model. However, the labor supply elasticity θ now also affects the employment level.

E.2 Equilibrium with Elastic Supply

Because labor supply is elastic, aggregate employment in period 1 is no longer fixed. Solving the model yields:

$$l_{i,1} = \left(\frac{(1-\alpha)(1-\tau_i)}{(1+\theta)\eta_i}\right)^{\frac{1}{\theta}} \left(\alpha^{\alpha}(1-\alpha)^{1-\alpha}z_1\right)^{\frac{1-\alpha}{\theta}} (1-\tau_i)^{\frac{\alpha(1-\alpha)}{\theta}} + \text{terms in } y_1.$$

Crucially, a higher tax τ_i now reduces employment, creating a deadweight loss beyond the pure redistribution channel. The aggregate output in period 1 becomes

$$\ln y_1 = \frac{x + \ln z_1 + \alpha\ln\alpha + (1-\alpha)\ln(1-\alpha) + \sum_i \chi_i \left[\frac{1-\alpha}{\theta}\ln(1-\tau_i) + \ln\chi_i\right]}{1-\alpha-\frac{(1-\alpha)^2}{\theta}}.$$

Taxation now lowers y_1 directly because it reduces labor supply.

E.3 Modified Voting Conditions

Using the same log-linearization as in the main text, the condition for sector B workers to support subsidies to sector A becomes:

$$\frac{(1-\tilde{\kappa})\chi_A}{\beta\psi} \;+\; (1-\phi)\gamma \;+\; \frac{(1-\alpha)(1-\tilde{\kappa})\chi_A}{\beta\psi\theta} \;<\; \phi\chi_B,$$

where $\tilde{\kappa} = \frac{\alpha\theta+\theta}{1+\theta}$ is the equilibrium profit share (which differs from the reduced-form κ used earlier because α now also appears). The extra term $\frac{(1-\alpha)(1-\tilde{\kappa})\chi_A}{\beta\psi\theta}$ captures the additional deadweight loss from taxing sector B: it reduces employment and output, making sector B even less willing to support the policy. Conditional on the previous force, however, a higher monopsony power would make workers from sector B more likely to support industrial policy to A.

The condition for sector A workers becomes:

$$(1-\phi)\gamma \;<\; \phi\chi_B \;+\; \frac{(1-\tilde{\kappa})\chi_B}{\beta\psi} \;+\; \frac{(1-\alpha)(1-\tilde{\kappa})\chi_B}{\beta\psi\theta}.$$

The deadweight loss term now appears on the right-hand side, which might appear to make sector A *more* supportive. Closer inspection reveals that the deadweight loss from taxing sector B reduces the tax burden that sector A would otherwise have to bear indirectly (because the government collects less revenue from a shrunken sector B).

The social planner's condition becomes:

$$(1-\phi)\gamma \;<\; \phi\chi_B \;-\; \frac{(1-\alpha)(1-\tilde{\kappa})}{\beta\psi\theta}\chi_A.$$

Relative to the baseline condition $(1-\phi)\gamma < \phi\chi_B$, the planner now faces an additional cost term that makes intervention less attractive. This reflects the fact that taxing sector B creates a deadweight loss even before considering synergies.