The Political Economy of Industrial Policy*

Jesús Fernández-Villaverde† Yang Yu‡ Francesco Zanetti♭

March 22, 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 integrating learning-by-doing, inter-sectoral spillovers, technological synergies, and labor market monopsony within an explicit democratic voting framework. The model shows that policies supported by the taxed majority is strictly smaller than the set of policies that improve aggregate welfare, opening the possibility that welfare-improving policies fail to pass democratically. The democratic enactment of welfare-improving policies requires strong inter-sectoral learning spillovers, low synergy costs, and a small subsidized sector. Monopsony power in the beneficiary sector enhances political support from workers in the rest of the economy who benefit through increased dividend income that partially offset their tax contribution to support industrial policy. Using novel U.S. firm-level data on customer-supplier relationships for the period 1976-2020, we validate our framework.

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

JEL Classification: D72, L52, O25, O38.

*We would like to thank Jing Cai for helpful comments and suggestions.
†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 modern democracies, the enactment of industrial policy –the political decision about the distribution of economic resources across sectors– is shaped by the relative political power of different groups, the institutions governing collective choice, and requires public approval through legislative votes or electoral mandates (Acemoglu and Robinson, 2008). In the U.S., for instance, these policies are typically enacted through Congressional Acts (which reflect competing group interests) or Presidential Orders (which are often limited by Congress's control over funding). This reality raises a fundamental question: why do voters support some industrial policies but reject others? And critically, are the policies that receive democratic support actually welfare-improving? What economic conditions facilitate democratic adoption of welfare-improving interventions? Understanding the alignment, or misalignment, between political feasibility and economic efficiency is essential for designing effective industrial interventions in democratic societies.[1]

This paper addresses these questions by developing a novel theoretical framework and using new data to study the political economy of industrial policy in U.S. history. Our model features a simple two-sector economy that integrates four key economic mechanisms into a explicit model of democratic voting: (i) within-sector learning-by-doing, which generates sectoral economies of scale; (ii) inter-sectoral learning spillovers, enabling high-productivity sectors to raise long-run productivity elsewhere; (iii) technological synergies between sectors due to complementarities in production; (iv) labor market monopsony power, which creates substantial profits that can be shared beyond the sector. The model derives the conditions under which voters from the beneficiary sector and the taxed sector would support subsidizing a specific industry, and allows us to compare these political support conditions with the welfare criterion of a social planner.

Our theoretical framework reveals three key insights about the political economy of industrial policy. First, from a welfare perspective, we show that industrial policy is justified when the beneficiary sector generates strong inter-sectoral learning spillovers, exhibits low synergy costs from productivity mismatches that impair technological synergies, and is small relative to the rest of the economy.

Second, workers in non-beneficiary sectors –who bear the tax burden of industrial policy– support industrial policy only when spillover benefits are large, the subsidized sector is small

[1] Some industrial policies are purely motivated by national security—such as requiring data centers to be located onshore. These fall outside the scope of this paper, which focuses on the economic dimensions of industrial policy, including the broader concept of economic security.

(minimizing their tax burden), synergy costs are low, and high profit shares in the subsidized sector generate compensating dividend income for all workers through diversified equity holdings. Central to the policy viability of industrial policy is the role of *monopsony power* in the subsidized sector. When the subsidized industry exhibits high monopsony power and firms in these sectors extracts most of the surplus generates by the policy that foster high growth and workers in the rest of the economy benefit through increased dividend income from their diversified equity holdings, offsetting their tax burden from industrial policy. This channel arising from monopsony power makes industrial policy more politically palatable to the taxed majority.

Third, workers in potentially subsidized industries exhibit near-universal support for policies targeting their own sectors, as they receive direct benefits while bearing only a small fraction of the tax burden. Comparing these conditions reveals a fundamental political economy constraint: the set of policies supported by the taxed majority is strictly smaller than the set that would improve aggregate welfare, creating the risk that welfare-improving policies fail to pass democratically.

We bring our theoretical framework to data by integrating firm-level trading relationships from Compustat and FactSet with firm performance data, covering over 360,000 customer-supplier pairs and 18,000 U.S. firms from 1976 to 2020. This rich dataset enables us to empirically estimate the key structural parameters of our model –synergies, inter-sectoral learning rates, and monopsony power– for each three-digit NAICS industry, and then use the estimates in our theoretical model to determine the key drivers of political approvals of industrial policies in the U.S.

Our identification strategy leverages firm-level heterogeneity within narrowly defined sectors, comparing firms that trade with partners of different productivity ranks. This allows us to examine how the distance in trading partners' productivity ranking affects a firm's current output and future productivity growth. Using relative productivity ranks within sectors isolates relational mismatches from aggregate sectoral trends, sharpening the identification by focusing on within-sector-pair comparison.

The central identification challenge is endogenous trading relationship formation: productivity mismatch between a firm and its partners is likely correlated with unobserved firm characteristics, such as organizational capital, or technological sophistication, that directly influence the firm's performance. To isolate exogenous variation in partner productivity, we use two novel instrumental variables. First, prior flows of senior managers between firms increase the

likelihood of matches between similar firms, providing a plausibly exogenous source of assortative matching. Second, pre-relationship differences in firm size capture a strategic trade-off where productive firms may accept a less productive partner for commercial scale. Standard diagnostics verify the validity of our IV strategy.

Using our instrumental variable procedure, we estimate industry-specific parameters for each three-digit NAICS industries and map these estimates to our theoretical framework to determine which industries garner support from workers in the beneficiary industry, workers in the non-beneficiary industries, and a welfare-maximizing planner.

Our empirical results provide strong support to the theoretical predictions in the model and yield concrete policy implications for U.S. industrial strategy. We find that only a narrow set of industries would receive support from both a social planner and the taxed majority: Fabricated Metal Products (NAICS 332), Transportation Equipment (336), Machinery (333), Computer and Electronic Products (334), Primary Metals (331), Support Activities for Mining (213), Petroleum and Coal Products (324), Chemical Manufacturing (325), and Miscellaneous Manufacturing (339). These industries share common characteristics—strong inter-sectoral learning spillovers combined with manageable synergy costs—that generate positive externalities for the broader economy while limiting the coordination frictions that arise from productivity mismatches.

Strikingly, this set of industries aligns closely with the historical targets of U.S. industrial policy over the past seven decades. Transportation Equipment has received sustained support through Department of Defense and NASA procurement beginning in the 1950s, emergency interventions such as the auto bailouts under the Emergency Economic Stabilization Act of 2008 and American Recovery and Reinvestment Act of 2009, and contemporary electric vehicle incentives under the Inflation Reduction Act of 2022. Computer and Electronic Product Manufacturing, particularly semiconductors, was cultivated through military contracts under the Defense Production Act of 1950 and Department of Defense R&D funding in the 1960s–1970s, and is now the explicit focus of the CHIPS and Science Act of 2022, which allocated $52.7 billion for domestic semiconductor production and research. Primary Metals has enjoyed protection through Section 232 national security tariffs (25% on steel, 10% on aluminum) and "Buy American" procurement rules codified in the Buy American Act of 1933. This historical consistency demonstrates that our framework captures the core economic rationale—high inter-sectoral spillovers, limited synergy costs, and favorable labor share parameters—that has repeatedly motivated actual policy decisions in democratic settings.

Our paper contributes to three strands of literature. First, we advance the political economy

literature on industrial policy by formalizing the democratic voting mechanism with heterogeneous, sector-specific workers. While existing studies models policy as the outcome of organized lobbying (Grossman and Helpman, 1994) or analysis focuses on qualitative institutional constraints (Juhász and Lane, 2024), we derive precise conditions –based on learning spillovers, synergy costs, industry size, and monopsony power– under which welfare-improving industrial policies may nevertheless fail to secure majority support. The seminal paper by Fernandez and Rodrik (1991) shows that ex-ante uncertainty about individual outcomes can lead to resistance against reforms that are ex-post welfare-improving. Our analysis builds on this insight by providing a formal model that derives explicit conditions under which this misalignment arises in the context of industrial policy, even without uncertainty.

Second, we 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.

Third, we formalize the economic case for industrial policy rooted in learning externalities and productivity spillovers, a theme with deep roots in 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; Bartelme et al., 2025; 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, incorporating learning-by-doing, inter-sectoral spillovers, synergies between sectors with different productivity levels, and monopsony in the labor market. Section 3 derives the voting conditions for workers in each sector and the social planner solution, and provides numerical illustrations of the parameter space for political support. Section 4 presents our empirical strategy for estimating synergies and inter-sectoral learning using firm-level trading data, including our novel

instrumental variables approach. Section 5 estimates industry-specific parameters and uses the framework to study the enactment of industrial policy in U.S. history. Section 6 concludes.

2 Model

We develop a two-sector, two-period model to analyze how industrial policy affects social welfare and the welfare of voters in different industries. The model incorporates four key mechanisms: within-sector learning-by-doing, inter-sectoral learning spillovers, synergies between sectors that depend on their relative productivity levels, and monopsony in the labor market. Our model embeds industrial policy choice within an explicit democratic voting framework where voters who work in each sector decide whether to vote in favour of industrial policy that subsidizes one sector and the government imposing taxes on the non-beneficiary sector to maintain a balanced budget.

2.1 Production Technology

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

Firm-level production. The output of firm j in sector i is given by a Cobb-Douglas production function in period 1:

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

where $z_{i,1}$ represents sectoral productivity (common to all firms within sector i), $l_{i,j,1}$ denotes labor input, $k_{i,j,1}$ denotes capital input, and $\alpha \in (0, 1)$ is the capital share parameter. For simplicity, we assume that the two sectors have equal productivity levels in the first period: $z_{A,1} = z_{B,1} = z_1$.

In period 2, firms produce output using only the productivity gains accumulated through learning:

$$y_{i,j,2} = z_{i,2},$$

where $z_{i,2}$ is the sector-level productivity determined by the learning mechanisms described below. This specification captures the idea that period 1 production generates knowledge capital (i.e., learning by doing and spillover as described below) that enhances period 2 productivity.

Sectoral aggregation. Sectoral outputs are obtained by aggregating across firms in the sector:

$$y_{i,1} = \int_0^1 y_{i,j,1}\, dj = z_{i,1} l_{i,1}^{1-\alpha} k_{i,1}^{\alpha}, \tag{2}$$

and

$$y_{i,2} = \int_0^1 y_{i,j,2}\, dj = z_{i,2},$$

where $l_{i,1} = \int_0^1 l_{i,j,1}\, dj$ and $k_{i,1} = \int_0^1 k_{i,j,1}\, dj$ denote aggregate labor and capital in sector i.

Labor market structure. We assume that the population in each sector is fixed at N_i, equal to the measure of firms within the sector. Workers cannot move across sectors and each worker supplies one unit of labor inelastically. This implies that each firm j employs exactly one unit of labor in equilibrium: $l_{i,j,1} = 1$. Total labor input in period 1 is $l_1 = l_{A,1} + l_{B,1}$. Since production in period 2 requires no labor input (as it depends only on accumulated productivity through knowledge capital learning), we have $l_2 = 0$.

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

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

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 satisfying $\chi_A + \chi_B = 1$, x denotes the exogenous component of aggregate TFP (assumed constant), and $\gamma \in [0, 1]$ captures the degree of synergies between sectors.

The synergy parameter γ measures the extent to which productivity differences between sectors reduce final output. When $\gamma = 0$, there are no synergies and sectoral productivity gaps have no effect on aggregate output. When $\gamma > 0$, larger productivity gaps $|\ln(z_{A,t}) - \ln(z_{B,t})|$ reduce final output, reflecting complementarities or coordination frictions between sectors. This specification captures the empirical finding that mismatches in productivity levels between trading partners impose costs on aggregate production.[2]

Sector size. For analytical tractability, we assume $N_i = \chi_i$, so that sector size (measured by the number of firms and workers) equals its income share. This normalization ensures that in the decentralized equilibrium, both sectors achieve the same per capita income. Without loss of generality, we assume sector A is smaller than sector B: $\chi_A < \chi_B$. Intuitively, sector A represents a nascent or emerging industry that may require industrial policy support, while sector B represents the rest of the economy that would bear the tax burden of such policies.

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

2.2 Learning by Doing and Spillovers

Sectoral productivity in period 2 is determined by two learning mechanisms that operate as externalities at the sector level.

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

$$\ln(\widehat{z}_{i,2}) = \ln(z_1) + \psi \ln(k_{i,1}/N_i), \tag{5}$$

where $\psi \geq 0$ represents the elasticity of productivity with respect to per capita capital, capturing the rate of within-sector learning. The variable $\widehat{z}_{i,2}$ denotes the productivity level sector i would achieve through its own learning, absent any spillovers from the other sector. This formulation follows the infant industry argument that production experience generates productivity improvements over time.

For simplicity and due to data limitations (we lack a sufficiently robust empirical approach to estimate sector-specific learning rates), our baseline model assumes that the rate of learning-by-doing is symmetric across the two sectors ($\psi_A = \psi_B$). In Appendix A, we relax this assumption and examine how sector-specific learning rates ($\psi_A \neq \psi_B$) affect our results. This extension highlights that when the subsidized sector possesses a higher learning rate than the rest of the economy, the conditions for democratic support become less restrictive, as the taxed majority anticipates larger future spillovers. Despite this quantitative adjustment, the qualitative core of our findings remains unchanged.

2. Inter-Sectoral Learning Spillovers. Sectors can enhance their productivity by learning from the other sector when it has achieved higher productivity levels:

$$\ln(z_{i,2}) = \ln(\widehat{z}_{i,2}) + \phi \max\{\ln(\widehat{z}_{-i,2}) - \ln(\widehat{z}_{i,2}), 0\}, \tag{6}$$

where $\phi \in [0,1]$ measures the rate of inter-sectoral learning and $z_{i,2}$ is the final productivity of sector i in period 2 after accounting for spillovers. When sector $-i$ (the other sector) has higher productivity ($\widehat{z}_{-i,2} > \widehat{z}_{i,2}$), sector i can close a fraction ϕ of the productivity gap through knowledge spillovers. When $\phi = 0$, there are no inter-sectoral spillovers; when $\phi = 1$, sector i completely catches up to the more productive sector. This formulation captures the well-documented phenomenon whereby firms learn from more productive partners via technology transfer, knowledge diffusion, exposure to superior management practices, and the heightened innovation incentives driven by stricter technical and quality standards.

Since productivity is determined at the sectoral level and individual firms are infinitesimally

small, both learning mechanisms operate as externalities that are beyond the control of individual firms. This market failure provides the potential economic rationale for industrial policy intervention.

2.3 Households

Workers from each sector belong to separate, sector-specific representative households. Each household owns a fully diversified portfolio of shares in all intermediate goods firms across both sectors, implying they receive identical dividend income regardless of which sector they work in.[3] This assumption captures the idea that workers can diversify their financial holdings even though they cannot switch sectors for employment.

The utility function of sector i's representative household is given by:

$$U_{i,t} = \ln(c_{i,t}), \tag{7}$$

where $c_{i,t}$ denotes per capita consumption in sector i at time t. Labor disutility is omitted from the utility function because labor supply is fixed and inelastic.

2.4 Government and Industrial Policy

The government can implement industrial policy by taxing one sector and using the revenue to subsidize the other sector. Specifically, sector i's output in period 1 is subject to a tax (or subsidy if negative) at rate τ_i. The government maintains a balanced budget that ensures total tax revenue equals total subsidies in period 1:

$$\tau_A y_{A,1} p_{A,1} + \tau_B y_{B,1} p_{B,1} = 0, \tag{8}$$

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

This formulation highlights the redistributional nature of the industrial policy we consider. In particular, we focus on a policy that subsidizes the smaller sector A by taxing the larger sector B:

$$\tau_A < 0, \quad \tau_B > 0. \tag{9}$$

The subsidy to sector A stimulates capital accumulation and production in that sector during period 1, thereby enhancing its learning-by-doing and potentially generating spillovers to sector B in period 2. At the same time, the industrial policy may introduce productivity differences between sectors that reduce final output. The key question is whether these dynamic benefits

[3] Final goods firms earn zero profits due to perfect competition, so their ownership is irrelevant.

justify the costs of synergies due to redistribution, and whether such a policy would receive democratic support from workers in both sectors.

3 Equilibrium

This section characterizes the decentralized equilibrium of the economy. We solve the model backward: first, we analyze the optimization problems of firms and households in period 1; next, we derive the resulting productivities and consumption in period 2; and finally, we compute lifetime utilities to characterize the conditions under which households in each industry A and B support industrial policy.

3.1 Period 1: Production and Consumption Decisions

3.1.1 Final Good Firms

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:

$$\chi_i y_t p_t = y_{i,t} p_{i,t}, \tag{10}$$

where p_t denotes the price of final goods and $p_{i,t}$ denotes the price of sector i output. This condition states that expenditure on each intermediate input equals its factor share times total revenue.

Normalizing the final good price to unity ($p_t = 1$), equation (10) simplifies to:

$$\chi_i y_t = y_{i,t} p_{i,t}. \tag{11}$$

This implies that each sector receives an income share equal to its factor share χ_i in the aggregate production function.

3.1.2 Intermediate Good Firms

Firms in the intermediate goods sectors use labor and capital to produce output. We introduce labor market monopsony power to capture the empirically observed wedge between wages and marginal products.[4] Intermediate good firms have *monopsony power in the labor market*, allowing them to set wages below the after-tax marginal product of labor. The resulting gap is captured by firms as profits and distributed to shareholders.

Critically, in our model the monopsony power is purely redistributive rather than distortionary, since aggregate labor supply is fixed at the population level. Monopsony power signifi-

[4]See recent studies by Kline et al. (2019); Fernández-Villaverde et al. (2021); Garin and Silvério (2024) for evidence and discussion of monoposony power.

cantly influences electoral preferences through two channels that reduce: (i) the labor income of workers in period 1, and (ii) the flow of profits to all households through their diversified equity holdings. The second channel is critical to the political economy of industrial policy: workers outside the subsidized sector benefit from higher dividend income, which partially offsets their tax burden.

An intermediate good firm in sector i solves:

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

subject to: (i) the production function $y_{i,j,1} = z_1 k_{i,j,1}^{\alpha} l_{i,j,1}^{1-\alpha}$, (ii) the wage-setting rule with monopsony power $w_{i,j,1} = (1-\kappa)(1-\tau_i)p_{i,1}\text{MPL}_{i,j,1}$, and (iii) the one-for-one conversion of capital into final goods. The parameter $\kappa \in [0,1]$ in the wage-setting rule captures labor market monopsony power in reduced form: a higher κ means a lower labor income share within the industry, with more revenue captured as profit.

The optimality conditions for the firm are:

$$l_{i,j,1}w_{i,j,1} = (1-\alpha)(1-\tau_i)(1-\kappa)p_{i,1}y_{i,j,1}, \tag{12}$$

$$k_{i,j,1} = \alpha(1-\tau_i)p_{i,1}y_{i,j,1}. \tag{13}$$

Equation (12) shows that the wage bill equals marginal revenue product of labor scaled by $(1-\kappa)$, with the fraction κ retained as profit. Equation (13) is the standard condition equating capital cost to its marginal revenue product.

Capital demand and the role of industrial policy. Using equation (10) and the aggregation $y_{i,1} = N_i y_{i,j,1}$, we obtain the capital demand for firm i:

$$k_{i,j,1} = \alpha(1-\tau_i)\chi_i y_1 / N_i = \alpha\,(1-\tau_i)\,y_1, \tag{14}$$

where the second equality uses $\chi_i = N_i$.

Equation (14) illustrates a key mechanism through which industrial policy enhances the economy: capital input decreases with the sectoral tax rate τ_i and increases with subsidies ($\tau_i < 0$). Since period 2 productivity depends on period 1 capital accumulation through learning-by-doing (equation 5), subsidizing sector A stimulates its capital investment, thereby enhancing future productivity.

Firm and sectoral output. Given the fixed labor supply $l_{i,j,1} = 1$, firm-level output in logs is:

$$\begin{aligned}\ln\left(y_{i,j,1}\right) &= \ln\left(z_1\right) + \alpha \ln\left(k_{i,j,1}\right) \\ &= \ln\left(z_1\right) + \alpha\left[\ln\left(\alpha\right) + \ln\left(1-\tau_i\right) + \ln\left(y_1\right)\right].\end{aligned}$$

Aggregating across firms yields sectoral output:

$$\begin{aligned}\ln\left(y_{i,1}\right) &= \ln\left(y_{i,j,1}\right) + \ln\left(N_i\right) \\ &= \ln\left(z_1\right) + \alpha\left[\ln\left(\alpha\right) + \ln\left(1-\tau_i\right) + \ln\left(y_1\right)\right] + \ln\left(\chi_i\right).\end{aligned} \tag{15}$$

Aggregate output. Substituting equation (15) into the final goods production function (4) yields the aggregate output in period 1:

$$\ln(y_1) = x + \ln\left(z_1\right) + \alpha \ln\left(\alpha\right) + \alpha \ln\left(y_1\right) + \sum_i \chi_i\left[\alpha \ln\left(1-\tau_i\right) + \ln\left(\chi_i\right)\right]. \tag{16}$$

The government's balanced budget constraint (equation 8) implies:

$$\sum_i \chi_i \ln\left(1-\tau_i\right) \approx -\sum_i \chi_i \tau_i = 0,$$

so the tax terms cancel out in the aggregate output (equation 16), giving:

$$\ln\left(y_1\right) = \frac{x + \ln\left(z_1\right) + \alpha \ln\left(\alpha\right) + \sum_i \chi_i \ln\left(\chi_i\right)}{1-\alpha}. \tag{17}$$

Equation (17) shows that while industrial policy redistributes resources across sectors, it does not affect aggregate period 1 output due to the balanced budget constraint.

Capital per capita. From equation (14) and using $\chi_i = N_i$:

$$\ln\left(k_{i,1}/N_i\right) = \ln(1-\tau_i) + \frac{\ln(\alpha) + x + \ln\left(z_1\right) + \sum_i \chi_i \ln\left(\chi_i\right)}{1-\alpha}. \tag{18}$$

Equation (18) shows that subsidies ($\tau_i < 0$) increase per capita capital accumulation in sector i, which is the channel through which industrial policy enhances future productivity via learning-by-doing.

3.1.3 Profits and Household Income

Firm profits. Profits for firm i derive from monopsony power in the labor market:

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

and using equation (11), sector i's profit per capita is:

$$\pi_{i,1} = \pi_{i,j,1} = \kappa(1-\alpha)(1-\tau_i)\chi_i y_1 / N_i = \kappa(1-\alpha)(1-\tau_i)y_1. \tag{19}$$

Household consumption. Workers in sector i receive income from two sources:

(i) *Labor income per capita:*

$$l_{i,j,1} w_{i,j,1} = (1-\tau_i)(1-\kappa)(1-\alpha)\chi_i y_1 / N_i = (1-\tau_i)(1-\kappa)(1-\alpha) y_1. \tag{20}$$

This shows that higher taxes reduce workers' labor income directly, while higher monopsony power κ also reduces labor income (with the difference captured as profit).

(ii) *Dividend income per capita from diversified equity holdings:*

$$\sum_i N_i \pi_{i,1} = \kappa(1-\alpha) y_1 \sum_i \chi_i (1-\tau_i) = \kappa(1-\alpha) y_1. \tag{21}$$

The second equality uses the balanced budget constraint. Critically, all households receive the same dividend income regardless of which sector they work in, because they hold fully diversified portfolios.

Total consumption per capita. Combining the two income sources in equations (20) and (21) yields total consumption per capita in sector i:

$$c_{i,1} = (1-\tau_i)(1-\kappa)(1-\alpha)y_1 + \kappa(1-\alpha)y_1 = (1-\tau_i+\tau_i\kappa)(1-\alpha)y_1, \tag{22}$$

or, equivalently, in logs:

$$\ln(c_{i,1}) = \ln(1-\alpha) + \ln(1-\tau_i+\tau_i\kappa) + \frac{x + \ln(z_1) + \alpha\ln(\alpha) + \sum_i \chi_i \ln(\chi_i)}{1-\alpha}, \tag{23}$$

where the term $(1-\tau_i+\tau_i\kappa)$ captures how taxes and monopsony power jointly affect consumption. When $\tau_i > 0$ (taxation), this term is less than 1, reducing consumption. However, if monopsony power is high (i.e., κ close to 1), the negative effect is partially offset because dividend income from the taxed sector's profits accrues to all households.

Normalization. To simplify subsequent expressions, we normalize $x = 0$ and $z_1 = 1$:

$$\ln(c_{i,1}) = \ln(1-\alpha) + \ln(1-\tau_i+\tau_i\kappa) + \frac{\alpha\ln(\alpha) + \sum_i \chi_i \ln(\chi_i)}{1-\alpha}, \tag{24}$$

$$\ln(k_{i,1}/N_i) = \ln(1-\tau_i) + \frac{\ln(\alpha) + \sum_i \chi_i \ln(\chi_i)}{1-\alpha}. \tag{25}$$

3.2 Period 2: Learning and Consumption

In period 2, there is no production using labor or capital inputs. Instead, sectoral output equals the productivity level achieved through the two learning mechanisms described in Section 2.2.

3.2.1 Within-Sector Learning-by-Doing

Substituting equation (25) into the within-sector learning equation (5):

$$\ln(\widehat{z}_{i,2}) = \psi \left[\ln(1-\tau_i) + \frac{\ln(\alpha) + \sum_i \chi_i \ln(\chi_i)}{1-\alpha} \right], \tag{26}$$

which shows how industrial policy affects second-period productivity through its impact on capital accumulation in period 1. Equation (26) shows that higher tax rates ($\tau_i > 0$) reduce second-period productivity by dampening capital accumulation and thus learning-by-doing, while subsidies ($\tau_i < 0$) enhance productivity growth.

3.2.2 Inter-Sectoral Learning Spillovers

Given that sector A receives a subsidy ($\tau_A < 0$) and sector B is taxed ($\tau_B > 0$), equation (26) implies higher productivity in sector A relative to sector B, i.e., $\widehat{z}_{A,2} > \widehat{z}_{B,2}$. Therefore:

Sector B learns from sector A: Sector B can improve productivity via learning spillovers from sector A, whose productivity has risen in response to the industrial policy subsidy. In particular, using equation (6), sector B can partially catch up by learning from the more productive sector A:

$$\begin{aligned} \ln(z_{B,2}) &= (1-\phi)\ln(\widehat{z}_{B,2}) + \phi \ln(\widehat{z}_{A,2}) \\ &= \psi \left[(1-\phi)\ln(1-\tau_B) + \phi \ln(1-\tau_A) + \frac{\ln(\alpha) + \sum_i \chi_i \ln(\chi_i)}{1-\alpha} \right], \end{aligned} \tag{27}$$

where the parameter ϕ determines the strength of inter-sectoral learning in closing the productivity gap.

Sector A does not learn from sector B: Since $\widehat{z}_{A,2} > \widehat{z}_{B,2}$, there is no learning potential from B to A, as shown by using equation (6):

$$\ln(z_{A,2}) = \ln(\widehat{z}_{A,2}) = \psi \left[\ln(1-\tau_A) + \frac{\ln(\alpha) + \sum_i \chi_i \ln(\chi_i)}{1-\alpha} \right]. \tag{28}$$

3.2.3 Aggregate Output and the Synergy Cost

The productivity gap between sectors imposes a synergy cost on final output, as captured by the penalty term in equation (4). Substituting equations (28) and (27) into (4):

$$\begin{aligned}\ln(y_2) &= -\gamma\left[\ln(z_{A,2}) - \ln(z_{B,2})\right] + \chi_A \ln(z_{A,2}) + \chi_B \ln(z_{B,2}) \\ &= \psi\left[\frac{\ln(\alpha) + \sum_i \chi_i \ln(\chi_i)}{1-\alpha}\right] + \psi\left[\chi_A + \phi\chi_B - (1-\phi)\gamma\right]\ln(1-\tau_A) \\ &\quad + \psi(1-\phi)(\chi_B + \gamma)\ln(1-\tau_B).\end{aligned} \tag{29}$$

3.2.4 Consumption in Period 2

Since production requires no inputs in period 2, and $\chi_i = N_i$, consumption per capita in sector i for period 2 is:

$$c_{i,2} = \chi_i y_2 / N_i = y_2, \tag{30}$$

implying:

$$\ln(c_{i,2}) = \ln(y_2). \tag{31}$$

Equation (31) shows that each sector i achieves equal consumption per capita in period 2. We model sectoral output as being fully paid out as labor income, which is isomorphic to firms earning profits and redistributing them entirely to households.

3.3 Lifetime Utility and Voting

We now derive the lifetime utility of households in each sector A and B and study the conditions under which sector B workers support industrial policy undertaken in sector A. We begin by combining period 1 and period 2 consumption to compute lifetime utility in each sector, then derive in turn the voting condition for sectors B and A households and study the forces that shape their voting over industrial policy.

3.3.1 Lifetime Utility

Combining period 1 consumption (equation 24) and period 2 consumption (equations 29 and 31), the lifetime utility of sector A households is:

$$\begin{aligned}U_A = u(c_{A,1}) + \beta u(c_{A,2}) &= \ln(1-\alpha) + \frac{(\alpha+\beta\psi)\ln(\alpha) + (1+\beta\psi)\sum_i \chi_i \ln(\chi_i)}{1-\alpha} \\ &\quad + \ln(1-\tau_A+\tau_A\kappa) + \beta\psi\left[\chi_A + \phi\chi_B - (1-\phi)\gamma\right]\ln(1-\tau_A) \\ &\quad + \beta\psi(1-\phi)(\chi_B+\gamma)\ln(1-\tau_B).\end{aligned} \tag{32}$$

Similarly, sector B households have lifetime utility:

$$U_B = u(c_{B,1}) + \beta u(c_{B,2}) = \ln(1-\alpha) + \frac{(\alpha+\beta\psi)\ln(\alpha) + (1+\beta\psi)\sum_i \chi_i \ln(\chi_i)}{1-\alpha} + \ln(1-\tau_B+\tau_B\kappa) + \beta\psi\left[\chi_A + \phi\chi_B - (1-\phi)\gamma\right]\ln(1-\tau_A) + \beta\psi(1-\phi)(\chi_B+\gamma)\ln(1-\tau_B). \tag{33}$$

Benchmark: No industrial policy. With $\tau_i = 0$ for all i:

$$U_A = U_B = \ln(1-\alpha) + \frac{(\alpha+\beta\psi)\ln(\alpha) + (1+\beta\psi)\sum_i \chi_i \ln(\chi_i)}{1-\alpha}, \tag{34}$$

confirming that both sectors achieve identical lifetime utility in the absence of industrial policy.

3.3.2 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 $U_B(\tau_A < 0, \tau_B > 0) > U_B(0,0)$. This condition is:

$$\beta\psi(1-\phi)(\chi_B+\gamma)\ln(1-\tau_B) + \ln(1-\tau_B+\kappa\tau_B) + \beta\psi\left[\chi_A + \phi\chi_B - (1-\phi)\gamma\right]\ln(1-\tau_A) > 0. \tag{35}$$

Using 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 (35) simplifies to:

$$\left[1 - \kappa + \beta\psi(1-\phi)(\chi_B+\gamma)\right]\chi_A < \beta\psi\left[\chi_A + \phi\chi_B - (1-\phi)\gamma\right]\chi_B. \tag{36}$$

Rearranging:

$$\frac{(1-\kappa)\chi_A}{\beta\psi} + (1-\phi)\gamma < \phi\chi_B. \tag{37}$$

Interpretation: Condition (37) is more likely to hold –meaning sector B households support industrial policy for sector A– when:

1. χ_A is low (equivalently, χ_B is high): Sector A is small, minimizing the tax burden on sector B.
2. γ is low: Synergies are weak, so the productivity gap creates minimal costs.
3. ϕ is high: Inter-sectoral learning is strong, generating substantial spillover benefits to sector B.
4. κ is high: Labor market monopsony power is high, so sector B workers benefit from increased dividend income that partially offsets their tax burden.
5. β is high: Households are patient and value future productivity gains.

6. ψ is high: Within-sector learning is strong, amplifying the productivity effects of subsidies.

The role of monopsony power deserves 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.

3.3.3 Voting by Sector A Workers

Similarly, households in sector A support industrial policy if:

$$\beta\psi\left(1-\phi\right)\left(\chi_B+\gamma\right)\chi_A<\left\{1-\kappa+\beta\psi\left[\chi_A+\phi\chi_B-\left(1-\phi\right)\gamma\right]\right\}\chi_B, \tag{38}$$

which simplifies to:

$$\left(1-\phi\right)\gamma<\phi\chi_B+\frac{(1-\kappa)\chi_B}{\beta\psi}. \tag{39}$$

Interpretation: Condition (39) is strictly weaker than condition (37), confirming that sector A households are far more likely to support industrial policy than sector B households. This is intuitive: sector A workers receive direct subsidies in period 1 while bearing only a fraction χ_A of the aggregate tax burden (since they also receive dividends from sector B's taxed profits).

Condition (39) is more likely to hold when:

1. χ_A is low (sector A is small).
2. γ is low (synergies are weak).
3. ϕ is high (inter-sectoral learning is strong).
4. κ is *low* (labor market monopsony power is low).
5. β is *low* (households are impatient).
6. ψ is *low* (within-sector learning is weak).

The last three 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 the immediate consumption gains rather than the future productivity gains from the implementation of industrial policy. When monopsony power is low (low κ), more of the subsidy flows to workers rather than being captured as profits that are shared across sector A and B households.

3.4 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 weighted sum of utilities across sectors. A utilitarian social planner maximizes aggregate welfare weighted by population shares:

$$U = N_A U_A + N_B U_B. \tag{40}$$

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

$$N_A \left[\phi \chi_B + \frac{\chi_B}{\beta \psi} - (1 - \phi)\gamma \right] + N_B \left[\phi \chi_B - \frac{\chi_A}{\beta \psi} - (1 - \phi)\gamma \right] > 0. \tag{41}$$

Using $N_i = \chi_i$, this condition simplifies dramatically to:

$$(1 - \phi)\gamma < \phi \chi_B. \tag{42}$$

Interpretation: The social planner's condition (42) is remarkably simple and depends on only three parameters. The inequality is satisfied when:

1. χ_A is low (equivalently, χ_B is high): Sector A is small, so the learning-by-doing effect is strong relative to sector size.

2. γ is low: Synergies are weak, minimizing the cost of creating productivity gaps.

3. ϕ is high: Inter-sectoral learning is strong, generating substantial spillovers to sector B.

Notably, the social planner's condition does not depend on monopsony power κ because this parameter is purely redistributive in our model with fixed labor supply. The condition is intermediate in restrictiveness: it is weaker than sector B's condition (37) but stronger than sector A's condition (39). This ordering implies:

- There exists a parameter space where industrial policy is welfare-improving (satisfies 42) but fails to gain political support from sector B workers (violates 37).

- Sector A workers support industrial policy across a broader parameter space than is socially optimal, reflecting their direct receipt of benefits while bearing only partial costs.

3.5 Numerical Illustration

To gain intuition, we develop a numerical illustration of our theoretical results calibrated using empirically grounded parameter values. While we estimate the values of the key parameters in the data in the next section (Section 4), we use numerical illustration to highlight the

distinct trade-offs between the political support for industrial policy and its implications for social welfare, demonstrating that industrial policy can be welfare-improving while nonetheless failing to garner political support, or might obtain political support yet be socially suboptimal for plausible parameterizations.

Calibration. We set values for the key parameters in the model as follows:

- *Within-sector learning* ($\psi = 0.25$): Following Bartelme et al. (2025), who estimate the elasticity of productivity with respect to production scale using international trade shocks.
- *Inter-sectoral learning* ($\phi = 0.44$): Based on our empirical estimates (that we will show in Section 4), specifically the IV estimate of productivity spillovers from high-productivity customers (0.44) from Table 3.
- *Monopsony power* ($\kappa = 0.47$): Calibrated to match the average profit share across sectors, measured as the ratio of operating surplus to value added in the U.S. 2012 input-output tables.
- *Discount factor* ($\beta = 0.99$): Standard annual discount factor in the macro literature.
- *Synergies* ($\gamma = 1.14$): Based on our empirical estimates in Section 4, specifically the IV estimate of output decline with customer productivity mismatch (-1.14) from Table 2, where $-\widehat{\gamma}^{cus}$ provides a proxy for the synergy parameter.

As we discuss in Section 4, there is significant heterogeneity in these parameters across industries, which can influence the political feasibility of industrial policy in different sectors. We will study this heterogeneity in the data (Section 4), and use it to inform our empirical analysis of industrial policy proposals in the U.S. manufacturing sector (Sector 5). Thus, the current numerical illustration serves to highlight the key trade-offs in our theoretical framework.

Parameter Space Analysis. Figure 1, 2 and 3 visualize the range of key parameters that satisfy the voting conditions for the three decision-makers: sector B workers, sector A workers, and the social planner.

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 37), 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:

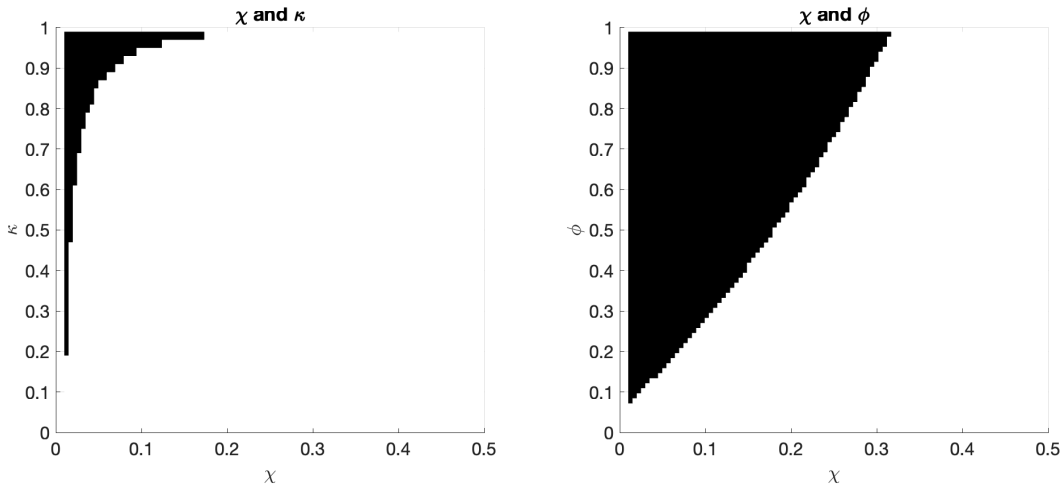


Figure 1: Parameter Range for Sector B's Support (Equation 37)

- 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

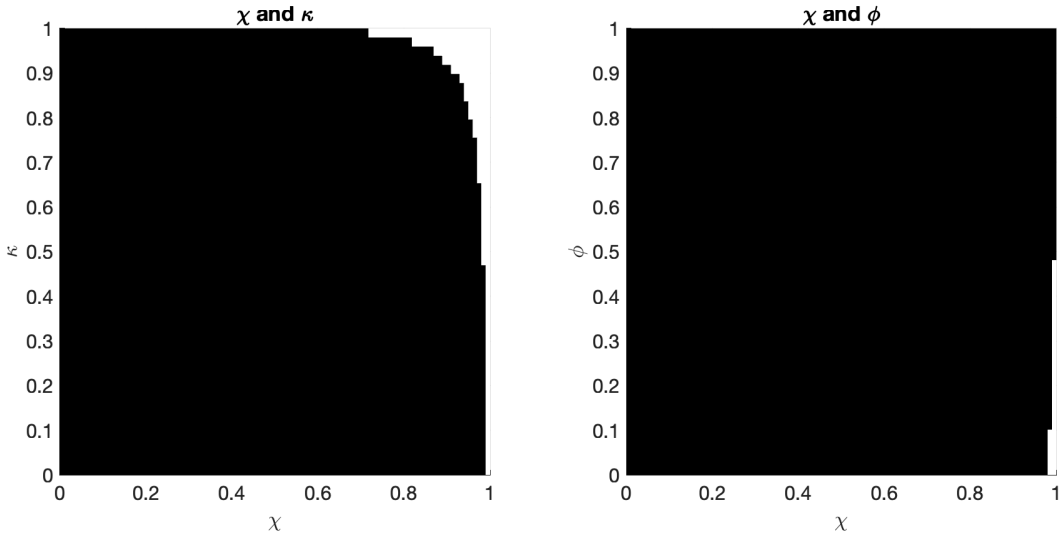


Figure 2: Parameter Range for Sector A's Support (Equation 39)

Sector A support (Figure 2). This figure shows parameter combinations where sector A supports industrial policy in its own sector (satisfying inequality 39), 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.

Given the estimated synergy parameter ($\gamma = 1.140$), sector A would vote in favor of the policy in many parameter combinations, though the higher synergy cost makes support less universal than with a lower γ. Exceptions occur under conditions: when monopsony power is very high (minimizing worker benefits relative to shareholder gains) and when sector A is extremely large (making subsidization highly inefficient at the aggregate level). In these cases, even sector A's own workers would be better off without the policy.

Figure 2 demonstrates a strong inherent incentive for workers to seek industrial policy for their own industry. However, the absence of such proposals in reality can be explained by workers' strategic anticipation that other sectors would not support them.

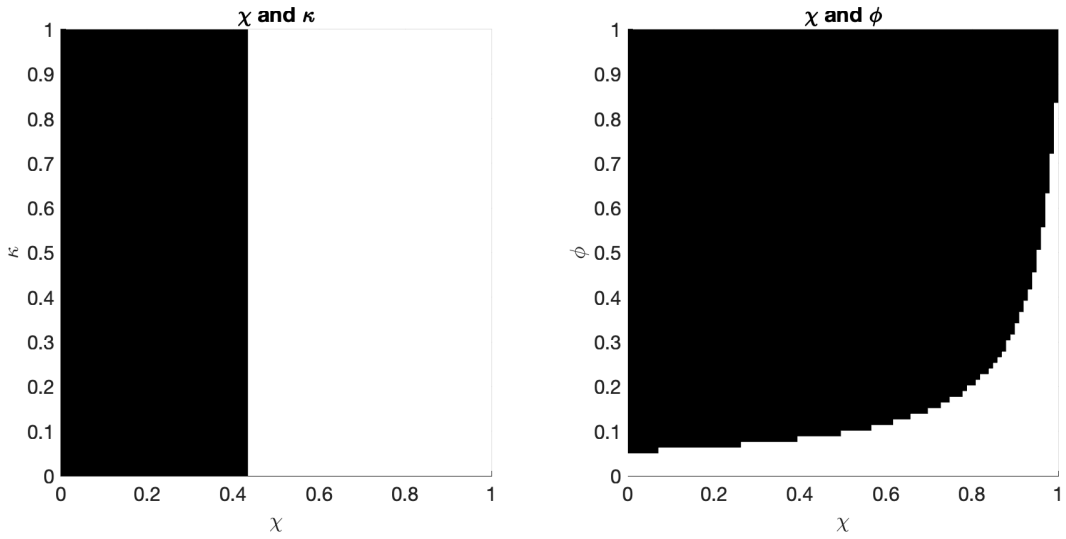


Figure 3: Parameter Range for Social Planner's Support (Equation 42)

Social planner support (Figure 3). This figure displays parameter combinations satisfying the social planner's condition (equation 42), 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

Taking stock. Comparing the three figures reveals important insights about the political economy of industrial policy:

First, as the direct beneficiary of subsidies, sector A supports industrial policy across a wide range of parameter values. The condition in equation (39) is strictly weaker than sector B's condition (37), reflecting sector A's favorable position as subsidy recipient.

Second, as the taxed sector, sector B requires the most stringent conditions to support industrial policy. The parameter space for sector B's support (Figure 1) is substantially smaller than for sector A (Figure 2). Sector B workers demand clear evidence of economy-wide spillover benefits to justify bearing the tax burden.

Third, the social planner supports industrial policy when aggregate benefits exceed aggregate costs, which requires stronger conditions than sector A's support but weaker conditions than sector B's support. Comparing Figures 1 and 3 reveals a substantial parameter space where industrial policy is welfare-improving yet fails to pass democratically, underscoring a clear conflict between economic efficiency and political incentives.

The political economy tension –beneficiaries strongly favor subsidies while the taxed majority requires tangible spillover benefits– suggests that many welfare-improving industrial policies may be politically infeasible under democratic voting.

4 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 5 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.

4.1 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 mismatches and learning opportunities from other sector-specific factors.

More concretely, consider the following two illustrative trading relationships: (i) a high-productivity R&D firm supplying a low-productivity manufacturer, and (ii) a high-productivity R&D firm supplying a high-productivity manufacturer. If synergies matter, the productivity mismatch in the first relationship should create greater bottlenecks and reduce output more severely. If inter-sectoral learning occurs, the low-productivity manufacturer in the first relationship should experience faster productivity growth by learning from its high-productivity partner, and therefore raise in ranking quickly while facing the same shocks as the other firms in the industry. By testing these hypotheses using firm-level trading data, we can estimate the magnitude of synergies (γ) and the rate of inter-sectoral learning (ϕ), the key parameters in our theoretical model. These firm-level coefficient estimates provide the empirical counterparts to the synergy and learning parameters featured in our theoretical framework. A remaining concern is the potential endogeneity of mismatch formation, which may correlate with unobserved firm characteristics. To address this, we employ an instrumental variables strategy using an exogenous predictor of mismatch, which we detail in the following subsections.

4.2 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.[5]

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.

4.3 Synergies and productivity mismatches

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 mismatches create bottlenecks that reduce output. We test these hypotheses using firm-level data on customer-supplier relationships.

4.3.1 Empirical specification

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

$$y_{j,t} = \widehat{\gamma}^{cus}\overline{\Delta}^{cus}_{j,t} + \beta z_{j,t} + \eta \times \overline{decile}\left(z^{cus}_{j,t}\right) + firm_j + \chi_t + \epsilon_{j,t}, \tag{43}$$

where $y_{j,t}$ and $z_{j,t}$ denote firm j's log value-added and log productivity in year t, respectively. For the supplier side, we estimate an analogous specification:

$$y_{j,t} = \widehat{\gamma}^{sup}\overline{\Delta}^{sup}_{j,t} + \beta z_{j,t} + \eta \times \overline{decile}\left(z^{sup}_{j,t}\right) + firm_j + \chi_t + \epsilon_{j,t}, \tag{44}$$

where $\overline{\Delta}^{sup}_{j,t}$ measures the average productivity mismatch between firm j and its suppliers, and $\overline{decile}\left(z^{sup}_{j,t}\right)$ 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.

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

The key explanatory variables $\overline{\Delta}_{j,t}^{cus}$ and $\overline{\Delta}_{j,t}^{sup}$ capture the average productivity mismatch 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 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 mismatch 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.

4.3.2 Instrumental variables strategy

A key identification challenge is that productivity mismatch 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 mismatch 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 it to be negatively correlated with the productivity gap between partners.

The second instrument is **size-distance**, defined as the absolute difference in sales deciles be-

tween the firm and its partner *in the year before the match formed*: $|decile\left(y_{j,t-1}\right) - decile\left(y_{j,k,t-1}\right)|$. 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.

Table 1: First-stage regression results for instrumental variables

	(1)	(2)
Dependent variable	Productivity mismatch ($\overline{\Delta}_{j,t}$)	
Number of senior managers in both firms	-0.0056***	
	(0.0014)	
Size-distance mismatch		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 (43) and (44). Column (1) shows the regression of productivity mismatch on the number of senior managers in both firms (N_flow). Column (2) shows the regression of productivity mismatch on size-distance mismatch (gap2). 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. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

4.3.3 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 $\overline{\Delta}_{j,t}^{cus}$ (customer productivity mismatch). Column (2) shows the results for customer firms, where the dependent variable is customer log sales and the key explanatory variable is $\overline{\Delta}_{j,t}^{sup}$ (supplier productivity mismatch).

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 mismatch reduces supplier sales by approximately 0.413 log points. The economic

Table 2: Output decline with productivity mismatch

	(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 mismatch 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. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

magnitude is substantial: moving from perfect alignment (mismatch = 0) to the maximum mismatch of 9 would reduce supplier output by about 3.7%.

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

4.4 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.

4.4.1 Empirical specification

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

$$decile\left(z^{sup}_{j,k,t+5}\right) - decile\left(z^{sup}_{j,k,t}\right) = \phi^{cus}\widetilde{\Delta}^{cus}_{j,k,t} + firm_{j,k} + \chi_t + \epsilon_{j,k,t}, \tag{45}$$

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:

$$decile\left(z^{cus}_{j,k,t+5}\right) - decile\left(z^{cus}_{j,k,t}\right) = \phi^{sup}\widetilde{\Delta}^{sup}_{j,k,t} + firm_{j,k} + \chi_t + \epsilon_{j,k,t}, \tag{46}$$

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 4.3.1), 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.

4.4.2 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	Customer's prod. growth	Supplier'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.

5 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 3 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 **??**) 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.

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) almost universally support industrial policy targeting their own sector, as they receive direct subsidies while bearing only a fraction of the costs. The sector A condition $(1-\phi_p)\gamma_p < \phi_p\chi_{-p} + \frac{(1-\kappa_p)\chi_{-p}}{\beta\psi}$ is the least restrictive, reflecting this inherent bias.

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. Second, we apply these estimates to the three support conditions derived in Section 3 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 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 4. 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.

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 Bureau of Economic Analysis. The size of the rest of the economy is $\chi_{-p} = 1 - \chi_p$. Monopsony power κ_p is measured as the average ratio of operating surplus to value added for firms in industry p using Compustat data, reflecting the portion of revenue captured as profit rather than paid to workers.

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

1. **Social planner support:** Equation (42) requires $(1-\phi_p)\gamma_p < \phi_p \chi_{-p}$.
2. **Sector B (rest of economy) support:** Equation (37) 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 (39) requires $(1-\phi_p)\gamma_p < \phi_p \chi_{-p} + \frac{(1-\kappa_p)\chi_{-p}}{\beta\psi}$.

We use the calibrated values of $\beta = 0.99$ (annual discount factor) and $\psi = 0.25$ (within-sector learning rate) from Section 3.6.

5.1 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 42, 37, and 39). 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

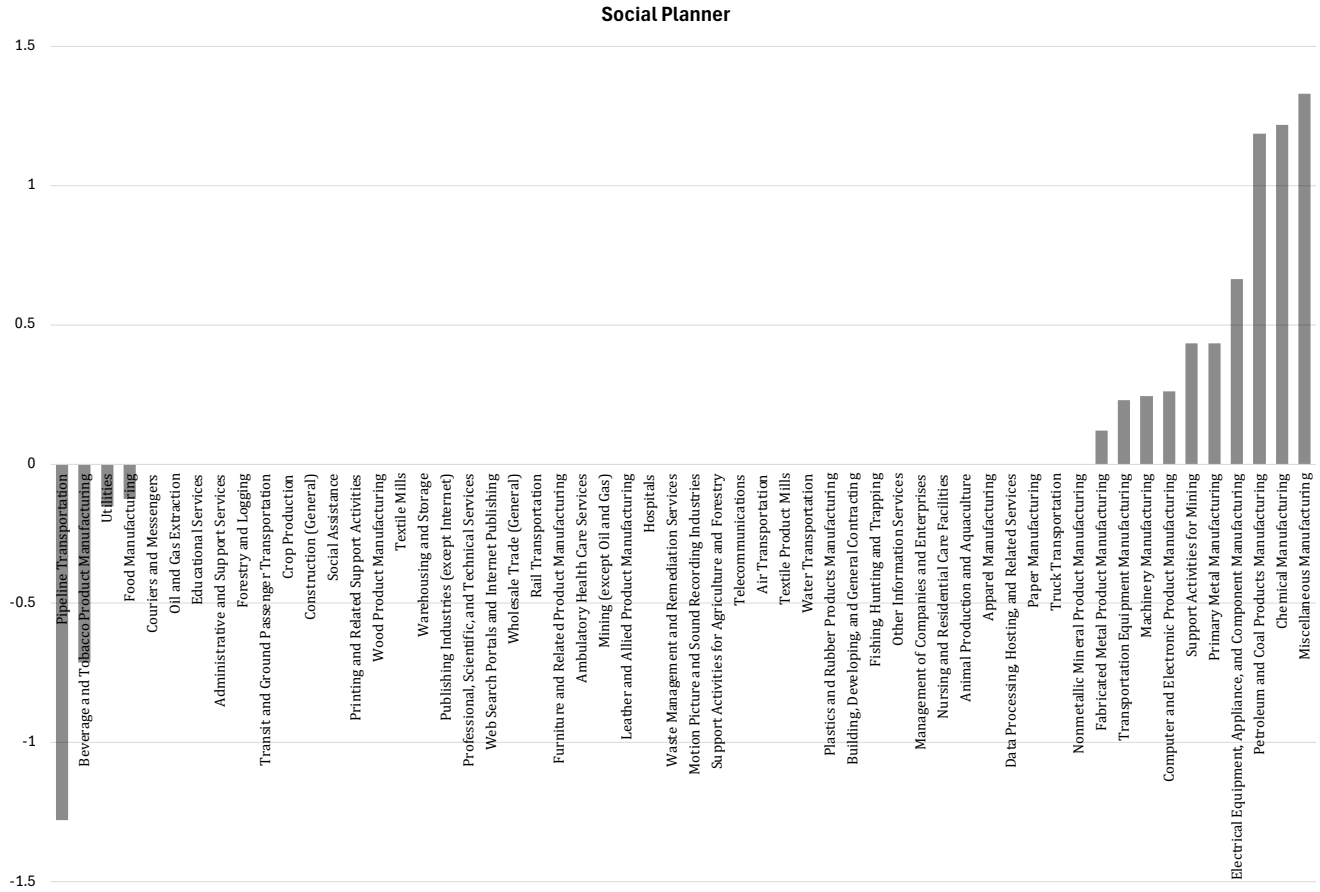


Figure 4: Social planner support for industrial policy across three-digit NAICS industries. The vertical axis shows the net social benefit $\phi_p \chi_{-p} - (1 - \phi_p)\gamma_p$ from subsidizing each industry (equation 42). Positive values indicate industries that would receive support from a welfare-maximizing social planner, as inter-sectoral learning spillovers exceed synergy costs.

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).

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



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 37). Positive values indicate industries that would receive sector B support.

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 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 (39) 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

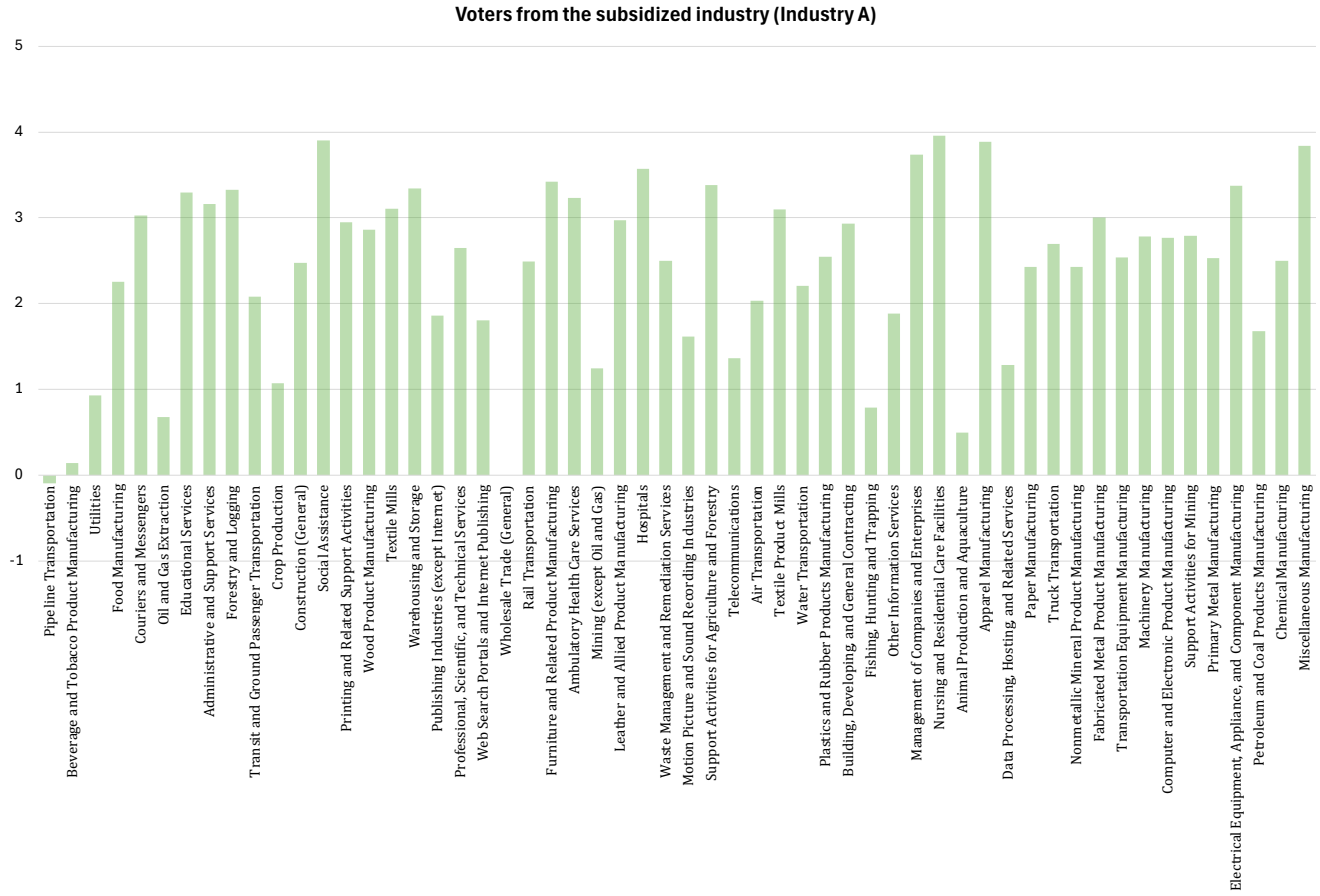


Figure 6: Sector A worker support for industrial policy across three-digit NAICS industries. The vertical axis shows the net benefit $\phi_p \chi_{-p} + \frac{(1-\kappa_p)\chi_{-p}}{\beta\psi} - (1-\phi_p)\gamma_p$ to workers in the subsidized industry (equation 39). Positive values indicate support.

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

6 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.

Garin, A. and Silvério, F. (2024). How responsive are wages to firm-specific changes in labour demand? evidence from idiosyncratic export demand shocks. *Review of Economic Studies*, 91(3):1671–1710.

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.

Grossman, G. M. and Helpman, E. (1994). Protection for sale. *American Economic Review*, 84(4):833–850.

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.

Kline, P., Petkova, N., Williams, H., and Zidar, O. (2019). Who profits from patents? rent-sharing at innovative firms. *The Quarterly Journal of Economics*, 134(3):1343–1404.

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

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

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 Sector-specific rate of learning by doing

B Model Extension: Sector-Specific Learning Rates

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:

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

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

The inter-sectoral spillover equation remains unchanged:

$$\ln(z_{i,2}) = \ln(\widehat{z}_{i,2}) + \phi \max\{\ln(\widehat{z}_{-i,2}) - \ln(\widehat{z}_{i,2}), 0\}. \tag{A-3}$$

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:

$$\ln(z_{A,2}) = \ln(\widehat{z}_{A,2}) = \psi_A \left[\ln(1-\tau_A) + C\right], \tag{A-4}$$

$$\begin{aligned}\ln(z_{B,2}) &= (1-\phi)\ln(\widehat{z}_{B,2}) + \phi \ln(\widehat{z}_{A,2}) \\ &= (1-\phi)\psi_B \left[\ln(1-\tau_B) + C\right] + \phi\psi_A \left[\ln(1-\tau_A) + C\right],\end{aligned} \tag{A-5}$$

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:

$$\begin{aligned}\ln(y_2) &= -\gamma\left[\ln(z_{A,2}) - \ln(z_{B,2})\right] + \chi_A \ln(z_{A,2}) + \chi_B \ln(z_{B,2}) \\ &= -\gamma(1-\phi)\left[\psi_A \ln(1-\tau_A) - \psi_B \ln(1-\tau_B)\right] \\ &\quad + \chi_A\psi_A \ln(1-\tau_A) + \chi_B\left[(1-\phi)\psi_B \ln(1-\tau_B) + \phi\psi_A \ln(1-\tau_A)\right] + \text{constant}. \quad \text{(A-6)}\end{aligned}$$

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:

$$\frac{(1-\kappa)\chi_A}{\beta\psi_B} + (1-\phi)\gamma \;<\; \phi\chi_B \cdot \frac{\psi_A}{\psi_B}. \quad \text{(A-7)}$$

Sector A (subsidized sector) supports policy if:

$$(1-\phi)\gamma \;<\; \phi\chi_B \cdot \frac{\psi_A}{\psi_B} \;+\; \frac{(1-\kappa)\chi_B}{\beta\psi_B}. \quad \text{(A-8)}$$

Social planner supports policy if:

$$\psi_B(1-\phi)(\chi_B+\gamma)\chi_A \;<\; \psi_A\chi_B\left[\chi_A + \phi\chi_B - (1-\phi)\gamma\right]. \quad \text{(A-9)}$$

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-7) and (A-8) 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-7) 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-8) 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-9) 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-7) rises—but it also makes the policy slightly less attractive to workers in the subsidized sector, because the right-hand side of (A-8) 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-7) grows—while making it more attractive to subsidized workers, as the left-hand side of (A-8) 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 accumulation 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: $\gamma = 0.05$, $\phi = 0.08$ (when fixed), $\kappa = 0.47$, and $\beta = 0.99$. The symmetric case uses $\psi = 0.25$, 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 $\phi = 0.08$; the right panel varies χ and ϕ, holding κ fixed at 0.47. 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

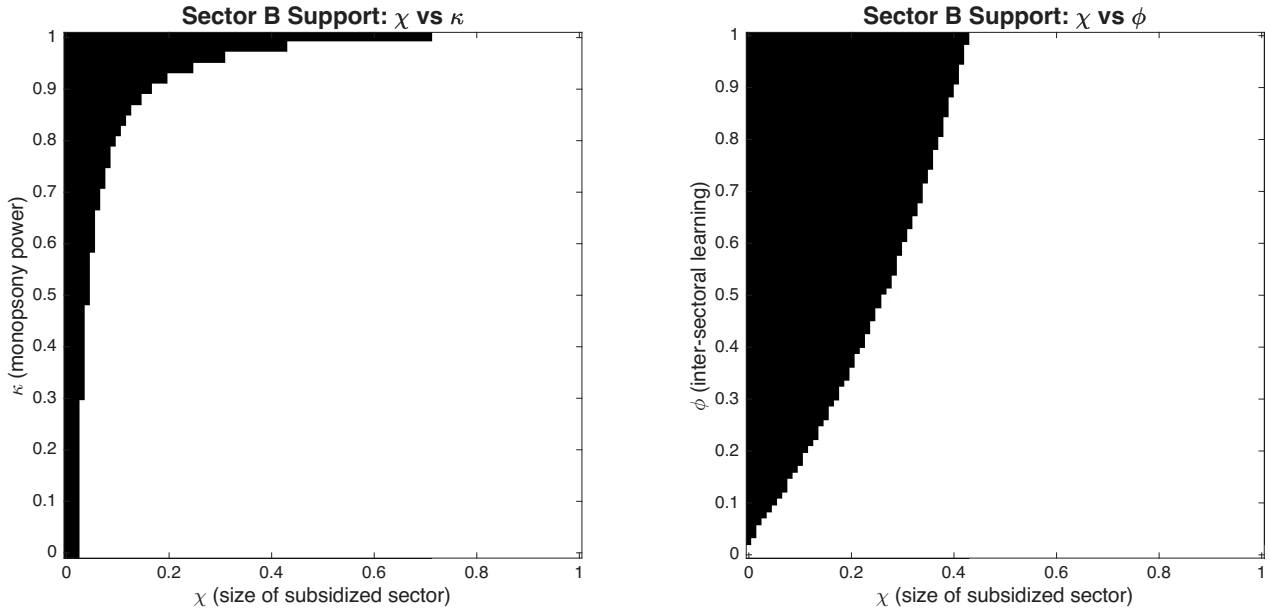


Figure B.1: Parameter space for sector B support when $\psi_p = 0.4 > \psi_{-p} = 0.2$. Black regions indicate support; white regions indicate opposition. Left: Variation in χ (subsidized sector size) and κ (monopsony power). Right: Variation in χ and ϕ (inter-sectoral learning).

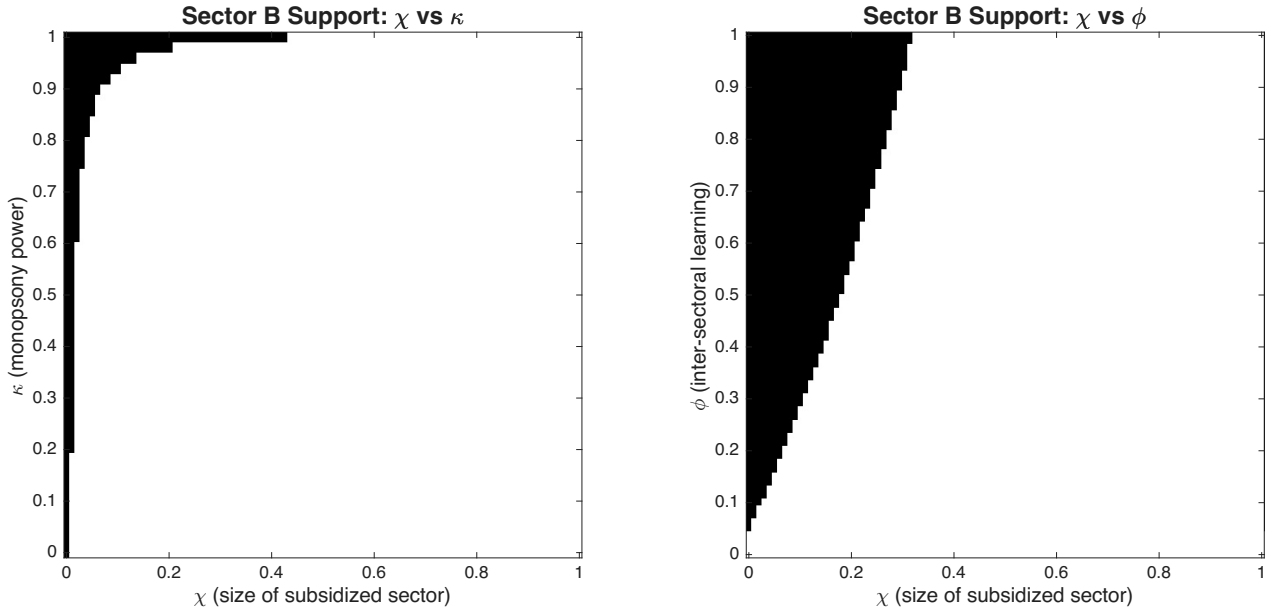


Figure B.2: Parameter space for sector B support in the symmetric case ($\psi_p = \psi_{-p} = 0.25$), matching our baseline analysis.

compensate for the higher tax burden. Second, greater monopsony power (κ) expands the 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.

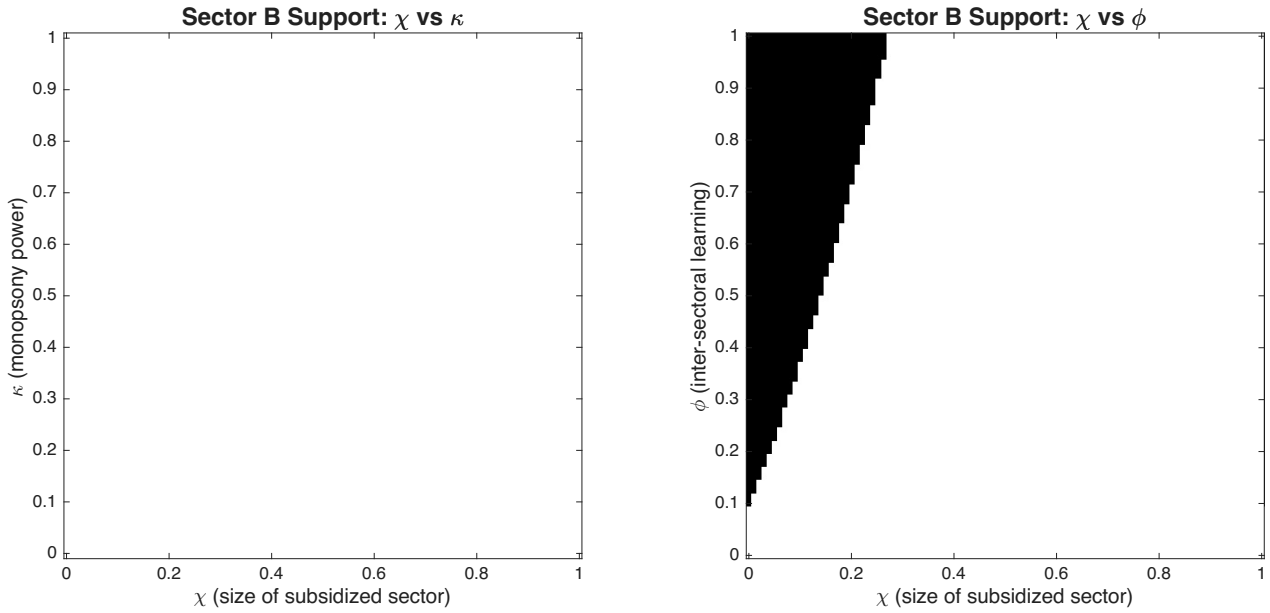


Figure B.3: Parameter space for sector B support when $\psi_p = 0.2 < \psi_{-p} = 0.4$. Support (dark region) is substantially more restrictive than in Scenarios 1 and 2.

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.