

The Textbook Case for Industrial Policy: Theory Meets Data

Dominick Bartelme

Independent Researcher

Arnaud Costinot

Massachusetts Institute of Technology

Dave Donaldson

Massachusetts Institute of Technology

Andrés Rodríguez-Clare

University of California, Berkeley

The textbook case for industrial policy is well understood: sectors with relatively large external economies of scale should be subsidized at the expense of other sectors. Little is known, however, about the magnitude of the welfare gains from such interventions. We develop an empirical strategy that leverages commonly available trade data to estimate sector-level economies of scale and, in turn, to quantify the gains from optimal industrial policy in a general equilibrium environment. Our results point toward significant economies of scale across manufacturing sectors but gains from industrial policy that are hardly transformative, even among the most open economies.

We are grateful to Susanto Basu, Jonathan Eaton, Giammario Impullitti, Yoichi Sugita, and David Weinstein for discussions; Tishara Garg, Daniel Haanwinckel, and Walker Ray

Electronically published March 13, 2025

Journal of Political Economy, volume 133, number 5, May 2025.

© 2025 The University of Chicago. All rights reserved. Published by The University of Chicago Press.

<https://doi.org/10.1086/734129>

I. Introduction

Industrial policy has been practiced around the globe for centuries. It has been a source of controversy among economists for about as long. Is industrial policy a critical input into the making of growth miracles? Or even in the best of worlds, where market failures have been correctly identified and benevolent governments are solely motivated by the desire to fix them, are the potential gains from industrial policy likely to be modest instead? The goal of our paper is to make progress on these and other related questions.

We focus on the textbook case for industrial policy in economies where some sectors are subject to external economies of scale. Accordingly, private marginal costs of production are higher than social ones. This creates a rationale for Pigouvian subsidies equal to the difference between the two, with the associated welfare gains equal to the area of the Harberger triangle located between the demand and social marginal cost curves, as illustrated in figure 1. In what follows, we propose to bring these classical ideas to the data and offer a first quantification of the textbook gains from industrial policy.

The contribution of our paper is threefold. First, we offer a modern general equilibrium treatment of these classical ideas in an open economy environment. Our theoretical analysis characterizes unilaterally optimal industrial policy and its welfare consequences, up to a second-order approximation, as a function of sector-level scale elasticities. As a second step, we show how to estimate the previous elasticities using commonly available trade data as well as existing estimates of sector-level trade elasticities. Our empirical analysis points toward sizable economies of scale across manufacturing sectors, which opens up the possibility of substantial wedges between private and social costs of production. As a final step, we combine our theoretical and empirical results to quantify the gains from industrial policy. Despite significant economies of scale in manufacturing, the gains from industrial policy that we estimate are hardly transformative, reflecting the fact that little reallocation across sectors actually takes place in response to the optimal industrial policy, even in the most open economies.

Section II presents our theoretical framework. Our baseline analysis considers a Ricardian economy with multiple sectors, each subject to external economies of scale. Our focus on this environment is motivated by its long intellectual history—from the work of Marshall (1920) to Graham's (1923) famous argument for trade protection and the formal treatment of external economies of scale in Chipman (1970) and Ethier

for outstanding research assistance; and many seminar and conference participants for comments. Costinot and Donaldson thank the National Science Foundation (grant 1559015) for financial support. This paper was edited by Greg Kaplan.




FIG. 1.—Textbook case for industrial policy. Because of external economies of scale in a sector, private marginal cost MC exceeds social marginal cost SMC , and so output Q is less than the social optimum Q^* . The optimal industrial policy is a subsidy $s = MC(Q^*) - SMC(Q^*)$, which gives rise to gains equal to the area of the gray Harberger triangle.

(1982)—as well as the recent emergence of the Ricardian model as a workhorse model for quantitative work. We first show that in the presence of optimal trade policy, the optimal industrial policy takes the form of an employment subsidy whose level depends on only the elasticity of productivity with respect to sector size, or scale elasticity, consistent with Pigou's (1932) logic. We then show that, up to a second-order approximation, the welfare gains from industrial policy are equal to half the product of the scale elasticity and the proportional change in sector size that it generates, summed across all sectors and weighted by the share of each sector in gross domestic product (GDP), that is, the sum of the areas of Harberger triangles in each sector.

Section III describes our empirical strategy to estimate scale elasticities. It builds on two simple observations. First, if there are positive economies of scale, larger sectors should tend to sell their products at lower prices. Second, if prices are lower in those sectors, quantities demanded should tend to be higher. It follows that one can estimate scale elasticities by tracing out the impact of exogenous (i.e., demand driven) variation in sector size on equilibrium quantities. We operationalize this general idea by assuming that within each sector, (i) productivity is a log-linear function of total sector employment, so that we have constant scale elasticities, and (ii) the demand for output from different countries is a log-linear

function of their prices, so that we have constant trade elasticities.¹ Under these restrictions, the (log of the) export price from a country is proportional to (the log of) its sector size, with a slope given by the scale elasticity, and the (log of the) export price is proportional to (the log of) its bilateral exports, with a slope given by the inverse of the trade elasticity. Given existing estimates of sector-level trade elasticities in the literature, we can therefore estimate sector-level scale elasticities using a log-linear regression of bilateral exports, adjusted by the trade elasticity, on sector size.

Since idiosyncratic productivity differences across countries and sectors affect both sector size and bilateral exports, identification of scale elasticities requires demand-side instrumental variables (IVs) that are correlated with sector size yet uncorrelated with productivity shocks. To construct such instruments, we exploit variation in countries' population and preferences by multiplying the former with estimates of structural demand residuals in each country-sector pair. The logic of our IV strategy is that within each sector, employment should be higher in countries that are larger and/or have a stronger taste for goods from that sector. Under the assumption that population and demand residuals are uncorrelated with idiosyncratic productivity shocks sector by sector, the previous procedure provides valid (and, in practice, strong) instruments for sector size at the country level. Importantly, our identification strategy deliberately draws on cross-sectional variation alone so as to isolate the long-run notion of scale economies that animates the textbook case for industrial policy.

Section IV presents our estimates of scale elasticities. Drawing on a cross section of 61 of the world's largest countries in 2010, our results point to statistically significant scale elasticities in every two-digit manufacturing sector, with an average of 0.21. The existence of positive economies of scale reflects the fact that within each sector, we find a positive relationship between demand-driven size and bilateral trade flows, a finding that is robust across samples and specifications. Our results also suggest substantial heterogeneity across sectors, with estimates ranging from 0.11 to 0.31, mostly driven by the heterogeneity of existing estimates of sector-level trade elasticities.² In the data, the relationship between demand-driven sector size and bilateral trade flows is fairly stable across sectors. Through the lens of our model, this creates a tendency for

¹ These parametric restrictions are satisfied by the multisector gravity models analyzed in Kucheryavy, Lyn, and Rodríguez-Clare (2023), a set that includes models with perfect competition and external economies of scale, as in this paper, but also models with monopolistic competition and free entry; in this case, scale effects arise from product differentiation and love of variety within industries, as in Krugman (1980).

² Interestingly, the previous numbers are below the inverse of the trade elasticity in all sectors, implying that scale effects are weaker than those implicitly assumed in trade models with monopolistic competition à la Krugman (1980) or Melitz (2003), as discussed in Costinot and Rodríguez-Clare (2014) and Kucheryavy, Lyn, and Rodríguez-Clare (2023).

sectors with higher trade elasticities—whose sales should respond more to prices—to be those with lower scale elasticities, so that their prices respond less to sector size. Accordingly, alternative sets of trade elasticity estimates may lead to very different scale elasticity estimates and, in turn, very different recommendations about the structure of optimal industrial policies.³

Section V uses our empirical estimates to evaluate the gains from industrial policy for each country one at a time. We focus on the best-case scenario in which governments unilaterally maximize their own residents' welfare without retaliation from other governments, have full knowledge of the underlying economy, and have full access to trade policy that can neutralize the adverse terms-of-trade effects that would occur when implementing industrial policy on its own. Yet even under these optimistic assumptions, we find that gains from industrial policy in our baseline calibration range from 0.59% to 2.06% of GDP, with larger gains for more open economies, and an average gain equal to 1.08%. These modest benefits from industrial policy do not reflect modest wedges. According to our estimates of scale elasticities, if labor were to reallocate fully to the manufacturing sector with the largest scale elasticity, the average welfare gains predicted by the areas of Harberger triangles would be equal to 12.4%. Modest gains instead reflect the fact that only modest labor reallocations take place from low-wedge to high-wedge sectors because of both inelastic domestic demand (due to a low estimated elasticity of substitution between sectors) and inelastic foreign demand (due to estimates of trade elasticities that are also small). Graphically, Harberger triangles have large heights but small bases, even for the most open countries. Interestingly, while the exact structure of optimal industrial policy is very sensitive to the choice of trade elasticities, we find that the welfare gains from such policies are fairly robust across alternative choices.

Section VI extends our theoretical and empirical analysis to a general environment featuring multiple factors of production and input-output linkages across sectors. We show that for a given set of trade elasticities, the structure of optimal industrial policy is remarkably similar to that in our baseline environment, with a correlation between the optimal industrial policy estimated in the two environments equal to 0.98. In terms of the magnitude of the welfare gains though, the introduction of input-output

³ In his original discussion of optimal industrial policy, Pigou (1932, 227) noted, "Attempts to develop and expand [these theoretical results] are sometimes frowned upon on the ground that they cannot be applied to practice. For, it is argued, though we may be able to say that . . . economic welfare would be increased by granting bounties to industries falling into one category and by imposing taxes on those falling into another category, we are not able to say to which of our categories the various industries of real life belong." The extent to which our analysis addresses these concerns (or not) therefore depends on the confidence (or lack thereof) that one has in existing estimates of sector-level estimates of trade elasticities.

linkages raises the average gains from industrial policy from 1.08% to 4.06%. This derives both from a mechanical increase in the tax base, since gross output rather than value added is being subsidized, and from larger reallocations across sectors, since subsidizing a sector now also tends to lower the price of inputs in that sector.

There is a large literature studying the rationale and potential consequences of industrial policy, as reviewed in Harrison and Rodríguez-Clare (2010), Lane (2020), Juhász, Lane, and Rodrik (2024), and Juhász and Steinwender (2024). Recent work includes reduced-form empirical analysis on the consequences of the Napoleonic blockade (Juhász 2018) and South Korea's transition to a military dictatorship (Lane 2017), text-based analysis to measure various types of industrial policy (Juhász et al. 2023), as well as theoretical work on optimal industrial policy in the presence of financial frictions (Itskhoki and Moll 2019; Liu 2019). However, despite the prominence of industrial policy in accounts of growth takeoffs and the recent popularity of industrial policy initiatives in many countries' economic agendas—from Made in China 2025 to the European Green Deal Industrial Plan and from the CHIPS and Science Act to the Inflation Reduction Act in the United States—there is a dearth of work that has tried to combine both theory and data in order to evaluate the potential benefits of industrial policy.

A notable exception is Lashkaripour and Lugovskyy (2018), who study a monopolistically competitive environment à la Krugman (1980), where the elasticity of substitution between domestic varieties may differ from the elasticity of substitution between domestic and foreign varieties. In this model, the scale elasticity is indirectly determined by the elasticity of substitution between domestic varieties, whereas the trade elasticity is determined by the elasticity of substitution between domestic and foreign varieties, so estimates of these two demand elasticities—obtained from monthly exchange rate variation in Colombia—can be used to calculate the effects of optimal policy. In contrast, our empirical strategy directly identifies scale elasticities from the responses of sector-level productivity, as revealed by exports, to changes in sector size caused by long-run variation in domestic demand.

Our paper also relates to a large literature that uses gravity models for counterfactual analysis, including Kucheryavy, Lyn, and Rodríguez-Clare (2023), who develop the generalization of the Ricardian model with industry-level economies of scale used in our estimation of the gains from industrial policy. As discussed in Kucheryavy, Lyn, and Rodríguez-Clare (2023), the quantitative predictions of gravity models hinge on two key elasticities: trade elasticities and scale elasticities. While the former have received significant attention in the empirical literature (see, e.g., Head and Mayer 2013), the latter have not. Scale economies, when introduced in gravity models, are instead indirectly calibrated using information

about the elasticity of substitution across goods in monopolistically competitive environments, as emphasized by Costinot and Rodríguez-Clare (2014). One of the goals of our paper is to offer more direct and credible evidence about scale elasticities for use in quantitative multisector gravity models.

In line with the aforementioned quantitative trade literature, we focus on the estimation of scale effects that operate at the level of a country-sector pair. This matches the level at which industrial policy is most often enacted. A related literature in urban and regional economics estimates agglomeration effects at the subnational level, as surveyed in Rosenthal and Strange (2004) and Combes and Gobillon (2015). Depending on whether these agglomeration effects are region specific or region and sector specific, this creates an additional rationale for place-based policies (as in Kline and Moretti 2014) or so-called cluster policies (as in Duranton 2011).

Our empirical analysis at the country-sector level is most closely related to Caballero and Lyons (1992) and Basu and Fernald (1997), who estimate returns to scale for US manufacturing sectors at the two- or three-digit levels. A distinctive feature of our empirical strategy is that we do not rely on measures of real output, or price indexes, collected by statistical agencies. Instead, we use existing estimates of trade elasticities to infer such prices. This provides a theoretically grounded way to adjust for potential quality differences across origins within the same sector, as discussed in detail in Bartelme et al. (2019), as well as an approach that works symmetrically for a large set of countries around the world, albeit one whose credibility depends on the credibility of existing estimates of trade elasticities. In addition, our method can only be applied to estimate scale elasticities in tradable sectors. Although tradable manufacturing sectors have been the traditional focus of industrial policy—from Japan, South Korea, and Taiwan in the twentieth century to China in the twenty-first—this is a nontrivial limitation of our approach (which we will have to deal with by considering a wide range of scale elasticities outside manufacturing in our counterfactual analysis).

Finally, the general idea of using trade data to infer economies of scale bears a direct relationship to empirical tests of the home market effect (see, e.g., Costinot et al. 2019; Head and Ries 2001; Davis and Weinstein 2003). Indeed, the home market effect—that is, a positive effect of demand on exports—implies the existence of economies of scale at the sector level. Our empirical strategy is also closely related to previous work on revealed comparative advantage (see, e.g., Costinot, Donaldson, and Komunjer 2012; Levchenko and Zhang 2016). The starting point of these papers, like ours, is that trade flows contain information about relative costs of production, a point also emphasized by Antweiler and Trefler (2002).

II. Theory

A. Baseline Environment

Consider an economy with many countries, indexed by i or j , and many sectors, indexed by k . Labor is the only factor of production, with L_i the fixed labor supply in each country.

Technology.—Production is subject to local external economies of scale that depend on total employment in each sector. Holding total employment fixed, firm-level production functions exhibit constant returns to scale.⁴ If firms from an origin country i use $\ell_{ij,k}$ units of labor to produce good k for a destination country j , they can deliver $y_{ij,k}$ units of good k to consumers in that country, with

$$y_{ij,k} = A_{ij,k} E_k(L_{i,k}) \ell_{ij,k}.$$

Transportation costs, if any, are reflected in $A_{ij,k}$, while external economies of scale are reflected in $E_k(L_{i,k})$, with $L_{i,k}$ denoting total employment in sector k and country i :

$$L_{i,k} = \sum_j \ell_{ij,k}. \quad (1)$$

From now on, we simply refer to $L_{i,k}$ as sector size.

Preferences.—In each destination country j , there is a representative agent with utility,

$$u_j(c_j),$$

where $c_j \equiv \{c_{ij,k}\}_{i,k}$ is the vector of consumption and $c_{ij,k}$ denotes consumption of good k from a particular origin i in a destination country j .⁵

Prices and taxes.—There are three types of ad valorem taxes: import tariffs $\{t_{ij,k}^m\}$, export taxes $\{t_{ij,k}^x\}$, and employment subsidies $\{s_{j,k}\}$. These taxes create wedges between the prices and wages faced by consumers $\{p_{ij,k}, w_j\}$ and those faced by firms $\{q_{ij,k}, v_{j,k}\}$:

⁴ In our working paper (Bartelme et al. 2019), we have allowed each sector to comprise multiple goods, each potentially produced by heterogeneous firms. In this more general environment, the critical assumption is constant returns to scale at the level of each good but not at the level of firms. As is well understood, constant returns to scale at the good level may reflect the free entry of heterogeneous firms, each subject to decreasing returns to scale, as in Hopenhayn (1992).

⁵ For expositional purposes, we have restricted economies of scale to operate through physical productivity. As shown in our working paper (Bartelme et al. 2019), one can easily extend our analysis to allow sector size to affect quality. Formally, if we were to assume more generally that $y_{ij,k} = A_{ij,k} E_k^q(L_{i,k}) \ell_{ij,k}$ and $u_j(\{E_k^q(L_{i,k}) c_{ij,k}\})$, then external economies of scale would affect quality-adjusted productivity via $E_k(L_{i,k}) \equiv E_k^q(L_{i,k}) E_k^p(L_{i,k})$, without any further implication for our analysis.

$$p_{ij,k} = (1 + t_{ij,k}^m) \bar{p}_{ij,k}, \tag{2}$$

$$q_{ij,k} = (1 - t_{ij,k}^x) \bar{p}_{ij,k}, \tag{3}$$

$$v_{j,k} = (1 - s_{j,k}) w_j, \tag{4}$$

where $\bar{p}_{ij,k}$ denotes the world (untaxed) price of a good k produced by country i and sold in country j . Net revenues from all taxes and subsidies imposed in any country j are rebated through a lump-sum transfer, T_j , to the representative agent in that country.

B. Competitive Equilibrium

In a competitive equilibrium, firms maximize profits, consumers maximize their utility, good and labor markets clear, and governments balance their budgets.

Profit maximization.—For any origin country i , any destination country j , and any sector k , profit maximization requires

$$(\ell_{ij,k}, y_{ij,k}) \in \operatorname{argmax}_{(\ell,y)} \{q_{ij,k}y - v_{i,k}\ell \mid y = A_{ij,k}E_k(L_{i,k})\ell\}. \tag{5}$$

We let $\pi_{ij,k}(q_{ij,k}, v_{i,k}, L_{i,k})$ denote the value function associated with (5), that is, the profit functions of firms from country i selling good k in country j .

Utility maximization.—For any destination country j , utility maximization requires

$$c_j \in \operatorname{argmax}_c \left\{ u_j(c) \mid \sum_{i,k} p_{ij,k} c_{ij,k} = w_j L_j + \sum_{i,k} \pi_{ji,k} + T_j \right\}. \tag{6}$$

We let $V_j(p_j, w_j L_j + \sum_{i,k} \pi_{ji,k} + T_j)$ denote the value function associated with (6), that is, the indirect utility function in country j , with $p_j \equiv \{p_{ij,k}\}_{i,k}$ the vector of good prices faced by consumers in that country.

Market clearing.—All good and labor markets clear:

$$c_{ij,k} = y_{ij,k}, \tag{7}$$

$$\sum_{j,k} \ell_{ij,k} = L_i. \tag{8}$$

Government budget balance.—For any country i , the government’s budget is balanced:

$$T_i = - \sum_k s_{i,k} w_i L_{i,k} + \sum_{j \neq i, k} t_{ij,k}^m \bar{p}_{ji,k} c_{ji,k} + \sum_{j \neq i} t_{ij,k}^x \bar{p}_{ij,k} y_{ij,k}. \tag{9}$$

Definition.—A competitive equilibrium with employment subsidies, $\{s_{j,k}\}$, import tariffs, $\{t_{ij,k}^m\}$, and export taxes, $\{t_{ij,k}^x\}$ corresponds to an allocation, $\{c_{ij,k}, \ell_{ij,k}, y_{ij,k}\}$, with sector sizes, $\{L_{i,k}\}$, good prices, $\{p_{ij,k}, q_{ij,k}, \bar{p}_{ij,k}\}$, wages, $\{w_i, v_{i,k}\}$, and lump-sum transfers, $\{T_j\}$, such that equations (1)–(9) hold.

C. *Optimal Industrial and Trade Policy*

The government’s problem.—Let $\tau \equiv \{s_{j,k}, t_{ij,k}^m, t_{ij,k}^x\}$ denote the full vector of employment subsidies and trade taxes around the world. Each competitive equilibrium with taxes τ induces equilibrium utility levels $\{U_j(\tau)\}$. The government’s problem in any country j is to choose its policy $\tau_j \equiv \{s_{j,k}, t_{ij,k}^m, t_{ij,k}^x\}_{i \neq j,k}$ to maximize the utility of its representative agent,

$$\max_{\tau_j} U_j(\tau), \tag{10}$$

taking employment subsidies and trade taxes in other countries as given. In the rest of our analysis, an optimal industrial and trade policy $\tau_j^* \equiv \{s_{j,k}^*, t_{ij,k}^{m*}, t_{ij,k}^{x*}\}_{i \neq j,k}$ refers to an interior solution to (10).

To simplify the characterization of the optimal trade policy, we assume that country j is small in the sense that import prices, $\{\bar{p}_{ij,k}\}_{i \neq j,k}$, and other foreign equilibrium variables, $\{c_{im,k}, \ell_{im,k}, y_{im,k}, L_{i,k}, p_{im,k}, q_{im,k}, \bar{p}_{im,k}, w_i, v_{i,k}, T_m\}_{i \neq j, m \neq j,k}$, are treated as independent of τ_j , while export prices, $\{\bar{p}_{ji,k}\}_{i \neq j,k}$, are allowed to vary with its exports, $\bar{p}_{ji,k} = \tilde{p}_{ji,k}(y_{ji,k})$, reflecting the fact that goods may be differentiated by country of origin.⁶

The structure of optimal taxes.—To characterize the structure of optimal taxes in country j , it is convenient to start from the following identity,

$$U_j(\tau) = V_j(p_j(\tau), I_j(\tau)),$$

where both country j ’s consumer prices p_j and its income $I_j \equiv w_j L_j + \sum_{i,k} \pi_{ji,k} + T_j$ are expressed as functions of τ . For a small change in country j ’s employment subsidies and trade taxes $d\tau_j$ that generates changes in equilibrium prices and quantities, a standard envelope argument implies that

$$\begin{aligned} \frac{dU_j}{\lambda_j} = & -\sum_k s_{j,k} w_j dL_{j,k} + \sum_{i \neq j,k} t_{ij,k}^m \bar{p}_{ij,k} dc_{ij,k} + \sum_{i \neq j,k} t_{ij,k}^x \bar{p}_{ij,k} dy_{ji,k} \\ & + \sum_k \epsilon_{j,k}^E (1 - s_{j,k}) w_j dL_{j,k} + \sum_{i \neq j,k} \epsilon_{ji,k}^P \bar{p}_{ji,k} dy_{ji,k}, \end{aligned} \tag{11}$$

where λ_j denotes the marginal utility of income in country j , $\epsilon_{j,k}^E \equiv d \ln E_k(L_{j,k}) / d \ln L_{j,k}$ denotes the elasticity of productivity with respect to sector size, which we will simply refer to as the scale elasticity, and

⁶ The small economy assumption is an implicit restriction on the primitives—preferences, technology, and labor endowments—described in sec. II.A. It holds if preferences are weakly separable across sectors and rest-of-the-world expenditure shares on goods from country j converge to zero within each sector, as we will assume in the quantitative analysis of secs. V and VI. In the context of a simple Armington model, Costinot and Rodríguez-Clare (2014) show that the optimal import tariff is very close to the one predicted by the small economy approximation, reflecting the fact that even the largest countries account for a very small fraction of expenditure in the rest of the world.

$\epsilon_{ji,k}^p \equiv d \ln \tilde{p}_{ji,k}(y_{ji,k}) / d \ln y_{ji,k}$ denotes the elasticity of export prices with respect to export volumes.⁷

The first three terms in equation (11) represent the marginal increases in the deadweight loss from employment subsidies, import tariffs, and export taxes, respectively, whereas the final two terms represent the gains for country j of correcting the two sources of distortions: external economies of scale, as captured by $\sum_k \epsilon_{j,k}^E (1 - s_{j,k}) w_j dL_{j,k}$, and monopoly power in world markets, as captured by $\sum_{i \neq j,k} \epsilon_{ji,k}^p \bar{p}_{ji,k} dy_{ji,k}$. At the optimal policy mix, changes in utility associated with any tax variation should be zero, leading to the following necessary condition:

$$\begin{aligned} \sum_k s_{j,k}^* w_j^* dL_{j,k} - \sum_{i \neq j,k} t_{ij,k}^{m*} \bar{p}_{ij,k}^* dc_{ij,k} - \sum_{i \neq j,k} t_{ji,k}^{x*} \bar{p}_{ji,k}^* dy_{ji,k} \\ = \sum_k \epsilon_{j,k}^{E*} (1 - s_{j,k}^*) w_j^* dL_{j,k} + \sum_{i \neq j,k} \epsilon_{ji,k}^{p*} \bar{p}_{ji,k}^* dy_{ji,k}, \end{aligned} \tag{12}$$

where asterisks denote variables in the equilibrium with optimal policy in country j . Condition (12) states that the marginal cost from taxation should be equal to its marginal benefit in terms of reducing distortions. It is immediate that (12) holds for (i) $s_{j,k}^* = \epsilon_{j,k}^{E*} / (1 + \epsilon_{j,k}^{E*})$, (ii) $t_{ij,k}^{m*} = 0$, and (iii) $t_{ji,k}^{x*} = -\epsilon_{ji,k}^{p*}$. Distortions caused by external economies of scale are best targeted by employment subsidies, whereas those caused by endogenous export prices are best targeted by export taxes. Since the prices of foreign goods are not manipulable, no import tariffs are necessary.

More generally, optimal industrial and trade policies can be characterized as follows:

PROPOSITION 1. In our baseline environment, the optimal industrial and trade policy are such that (i) $s_{j,k}^* = (\bar{s}_j + \epsilon_{j,k}^{E*}) / (1 + \epsilon_{j,k}^{E*})$, (ii) $t_{ij,k}^{m*} = \bar{t}_j$, and (iii) $t_{ji,k}^{x*} = 1 - (1 + \bar{t}_j) (1 + \epsilon_{ij,k}^{p*})$ for some \bar{s}_j and \bar{t}_j .

The formal proof can be found in appendix section A.2. The two shifters, \bar{s}_j and \bar{t}_j , reflect two distinct sources of tax indeterminacy. First, since labor supply is perfectly inelastic, a uniform employment subsidy \bar{s}_j affects only the level of wages in country j but leaves the equilibrium allocation unchanged. Second, a uniform change in all trade taxes again affects the level of prices in country j but leaves the trade balance condition and the equilibrium allocation unchanged, an expression of Lerner symmetry. In the rest of our analysis, we normalize both \bar{s}_j and \bar{t}_j to zero.

It is worth noting that many assumptions imposed previously can be relaxed without affecting proposition 1. In terms of technology, our proof does not require the existence of a single factor and a linear production function. It can accommodate arbitrary convex technologies, a feature

⁷ Details are provided in app. sec. A.1 (apps. A–C are available online). Similar variational arguments can be found in Greenwald and Stiglitz (1986) and Costinot and Werning (2023), among many others.

that we will use in section VI when introducing multiple factors of production and input-output linkages. Our proof can also accommodate external economies of scale $E_{j,k}(\cdot)$ that are both country and sector specific; the restriction that $E_{j,k}(\cdot) = E_k(\cdot)$ will be used only for identification in section III. Finally, our focus on a small economy is relevant only for the structure of optimal trade policy. For a large country, the optimal trade taxes would depend on the elasticities of export and import prices with respect to the entire vector of imports and exports, as described in Dixit (1985), not just exports within a given sector and destination, as described in proposition 1. But regardless of whether country j is small, the optimal industrial policy, which is the main focus of our analysis, is given by $\epsilon_{j,k}^{E*} / (1 + \epsilon_{j,k}^{E*})$.⁸

The welfare gains from optimal taxes.—We define the welfare gains from optimal taxes, W_j^* , as the variation in country j 's income under laissez-faire that would be equivalent to moving from laissez-faire, $\tau_j = 0$, to the optimal policy mix, $\tau_j = \tau_j^*$. Omitting foreign taxes for notational convenience, W_j^* is implicitly given by the solution to

$$V_j(p_j(\tau_j^*), I_j(\tau_j^*)) = V_j(p_j(0), I_j(0) + W_j^*). \tag{13}$$

Likewise, we define the gains from optimal industrial policy, W_j^I , as the gains from introducing optimal employment subsidies $s_j^* \equiv \{s_{j,k}^*\}$ conditional on having already imposed optimal trade taxes $t_j^* \equiv \{t_{ij,h}^{m*}, t_{ji,h}^{x*}\}_{i \neq j,k}$ and, conversely, the welfare gains from optimal trade policy, W_j^T , as the gains from introducing the optimal trade taxes t_j^* conditional on having already imposed optimal employment subsidies s_j^* :

$$V_j(p_j(\tau_j^*), I_j(\tau_j^*)) = V_j(p_j(0, t_j^*), I_j(0, t_j^*) + W_j^I), \tag{14}$$

$$V_j(p_j(\tau_j^*), I_j(\tau_j^*)) = V_j(p_j(s_j^*, 0), I_j(s_j^*, 0) + W_j^T). \tag{15}$$

By definition, W_j^* , W_j^I , and W_j^T are all positive since the underlying changes in policy bring the economy toward full efficiency.⁹

⁸ This derives from the fact that the welfare impact of marginal changes in sector sizes would still be equal to $\sum_k [\epsilon_{j,k}^{E*} (1 - s_{j,k}^*) - s_{j,k}^*] w_j^* dL_{j,k}$. The critical assumption for the structure of industrial policy is that external economies of scale are country and sector specific. More generally, if productivity in a given country and sector were allowed to vary with the size of other sectors, then optimal industrial policy would depend on the elasticities of productivity with respect to the full vector of sector size, not just own sector size. This is analogous to the issue that arises for trade policy as we go from a small to a large country.

⁹ In contrast, starting from laissez-faire and imposing only the industrial or trade policy characterized in proposition 1 may reduce welfare, as policies that target only a subset of distortions may aggravate the others. As an extreme example, imagine an economy that (i) exports all of its output, so that employment subsidies and export taxes are perfect substitutes, and (ii) features $\epsilon_{ij,k}^{p*} = -\epsilon_{j,k}^{E*} / (1 + \epsilon_{j,k}^{E*})$ for all i, j , and k under laissez-faire. In this case, laissez-faire replicates the optimal allocation because the optimal employment subsidy exactly cancels out the export tax. Thus, only imposing $s_{j,k}^* = \epsilon_{j,k}^{E*} / (1 + \epsilon_{j,k}^{E*})$ in the

Building on the first-order necessary conditions used to characterize optimal policies, our next result offers a second-order approximation to these gains.

PROPOSITION 2. In our baseline environment, up to a second-order approximation, the welfare gains from optimal taxes satisfy

$$\frac{W_j^*}{I_j} \approx \frac{1}{2} \sum_k \frac{w_j L_{j,k}}{I_j} \cdot \frac{(\Delta L_{j,k})^*}{L_{j,k}} \cdot \epsilon_{j,k}^{E*} + \frac{1}{2} \sum_{i \neq j,k} \frac{\bar{p}_{i,k} y_{i,k}}{I_j} \cdot \frac{(\Delta y_{i,k})^*}{y_{i,k}} \cdot \epsilon_{j,k}^{P*}, \quad (16)$$

where $(\Delta L_{j,k})^*$ and $(\Delta y_{i,k})^*$ denote the employment and export changes associated with imposing τ_j^* and all other variables are evaluated under laissez-faire. Likewise, up to a second-order approximation, gains from industrial and trade policy satisfy

$$\frac{W_j^I}{I_j} \approx \frac{1}{2} \sum_k \frac{w_j L_{j,k}}{I_j} \cdot \frac{(\Delta L_{j,k})^I}{L_{j,k}} \cdot \epsilon_{j,k}^{E*}, \quad (17)$$

$$\frac{W_j^T}{I_j} \approx \frac{1}{2} \sum_{i \neq j,k} \frac{\bar{p}_{i,k} y_{i,k}}{I_j} \cdot \frac{(\Delta y_{i,k})^T}{y_{i,k}} \cdot \epsilon_{j,k}^{P*}, \quad (18)$$

where $(\Delta L_{j,k})^I$ denotes the employment change associated with imposing s_j^* , conditional on having already imposed optimal trade taxes, and $(\Delta y_{i,k})^T$ denotes the export change associated with imposing t_j^* , conditional on having already imposed optimal employment subsidies.

The formal proof can be found in appendix section A.3. Like the proof of proposition 1, it remains valid for arbitrary convex technologies, including those featuring multiple factors and input-output linkages as in section VI, and it uses only the restriction to a small economy in order to characterize the welfare gains from trade policy.

Proposition 2 can be viewed as the mathematical counterpart in general equilibrium of figure 1. In each of the second-order approximations displayed in equations (16)–(18), the welfare gains correspond to a weighted sum of the areas of Harberger triangles, with weights given by employment shares, $w_j L_{j,k}/I_j$, for industrial policy and export shares, $\bar{p}_{i,k} y_{i,k}/I_j$, for trade policy. The height of each triangle is equal to the magnitude of the underlying distortion, either $\epsilon_{j,k}^{E*}$ or $\epsilon_{j,k}^{P*}$, whereas the base of each triangle is equal to the changes in employment and exports generated by the policies targeting those distortions, $\Delta L_{j,k}/L_{j,k}$ and $\Delta y_{i,k}/y_{i,k}$.¹⁰

absence of trade taxes or only imposing $\tau_{j,k}^{x*} = -\epsilon_{ij,k}^{P*}$ in the absence of employment subsidies would necessarily lower welfare.

¹⁰ Although we refer to our triangles as Harberger triangles, it is worth pointing out one important conceptual distinction between proposition 2 and the standard Harberger triangle analysis (see, e.g., Hines 1999). The standard analysis fixes the primitives of the economy and introduces small exogenous taxes in order to evaluate their welfare cost. Our exercise instead varies the primitives of the economy in order to generate small endogenous

For gains from industrial policy to be large, proposition 2 highlights two conditions. First, production processes need to display large external economies of scale—such that a nation’s productivity in a given sector is increasing in the size of that sector—and scale economies that differ in strength across sectors—such that any productivity-enhancing expansion of scale in one sector does not just lead to an equal and opposite decline in productivity elsewhere in the economy.¹¹ Second, countries need to produce highly substitutable and tradable goods, such that a country can simultaneously expand employment in one sector and find useful domestic or foreign substitutes for the sector that it chooses to shrink. These are the conditions that lead to large changes in employment, $\Delta L_{j,k}/L_{j,k}$. Likewise, for gains from trade policy to be large, world prices need to be highly manipulable, as reflected in high values of $\epsilon_{ji,k}^{p*}$, and trade flows need to be responsive to changes in trade policy, as reflected in high values of $\Delta y_{ji,k}/y_{ji,k}$.

III. Empirical Strategy

We aim now to quantify the benefits that a country can be expected to enjoy if it enacted the optimal industrial and trade policy described above. Doing so relies on estimates of scale and export price elasticities, $\{\epsilon_{j,k}^{E*}, \epsilon_{ji,k}^{p*}\}_{i \neq j,k}$. The latter can be directly recovered from existing estimates of so-called gravity equations, as we discuss in detail below. In this section, we describe a procedure for obtaining estimates of $E_k(\cdot)$ and, in turn, $\{\epsilon_{j,k}^{E*}\}$.

A. General Idea

Our empirical strategy for estimating external economies of scale builds on two observations. First, the presence of external economies of scale should be reflected in prices. Everything else being equal, if there are positive external economies of scale, larger sectors should tend to sell their products at lower prices. Second, lower prices should themselves be reflected in higher quantities demanded for goods in that sector. Combining these two observations, one can therefore identify the presence of external economies by tracing out the impact of exogenous variation in sector size on quantities demanded.

Formally, the profit maximization condition (5) implies the equality of producer price and marginal cost for any good k produced by country i for destination j :

optimal taxes and evaluate their welfare gains. At a technical level, this distinction is reflected in the fact that the terms appearing in our Taylor expansion are $\epsilon_{j,k}^{E*}$ or $\epsilon_{ji,k}^{p*}$ and not $s_{j,k}^*$ or $t_{ji,k}^{ex}$.

¹¹ If $\epsilon_{j,k}^{E*} = \epsilon_j^{E*}$ for all k , then $\sum_k (w_j L_{j,k}/I_j) \cdot (\Delta L_{j,k}/L_{j,k}) \cdot \epsilon_{j,k}^{E*} = (w_j/I_j) \epsilon_j^{E*} \sum_k \Delta L_{j,k} = 0$.

$$q_{ij,k} = \frac{v_{i,k}}{A_{ij,k}E_k(L_{i,k})}. \tag{19}$$

Similarly, the utility maximization condition (6) implies the equality of relative consumer price and the marginal rate of substitution for any good k sold by two countries in destination j . That is,

$$\frac{p_{ij,k}}{p_{1j,k}} = \frac{u_{ij,k}(c_j)}{u_{1j,k}(c_j)}, \tag{20}$$

where $u_{ij,k}(c_j)$ is the partial derivative of $u_j(\cdot)$ with respect to the quantity $c_{i,k}$ evaluated at the consumption vector c_j and prices are expressed relative to those of goods from country 1. Combining these expressions, using the nonarbitrage conditions (2)–(4) relating producer and consumer prices, and averaging across destination countries implies

$$\frac{1}{J} \sum_j \text{DD}_{i_0,k_0} \{ \ln u_{ij,k}(c_j) \} = -\text{DD}_{i_0,k_0} \{ \ln E_k(L_{i,k}) \} - \frac{1}{J} \sum_j \text{DD}_{i_0,k_0} \{ \ln \alpha_{ij,k} \}, \tag{21}$$

where $\alpha_{ij,k} \equiv A_{ij,k}(1 - t_{ij,k}^x) / [(1 - s_{i,k})(1 + t_{ij,k}^m)]$ is the tax-adjusted productivity of firms from country i exporting good k to destination j ; $\text{DD}_{i_0,k_0} \{ \cdot \}$ is the double difference relative to a reference country i_0 and a reference sector k_0 , for example, $\text{DD}_{i_0,k_0} \{ \ln u_{ij,k}(c_j) \} \equiv [\ln u_{ij,k}(c_j) - \ln u_{i_0j,k}(c_j)] - [\ln u_{1j,k_0}(c_j) - \ln u_{i_0j,k_0}(c_j)]$; and J is the number of destinations.¹²

In this cross-sectional difference-in-differences, the first difference, taken between country i and country i_0 , reflects the fact that, as in equation (20), consumer choices reveal only relative prices, whereas the second difference, taken between sector k and sector k_0 , derives from a desire to eliminate the endogenous wage w_i that is part of $v_{i,k}$ in equation (19). Last, the average over destinations j on the left-hand side is introduced to focus on the fact that when buying from a given origin i and sector k , consumers in many locations j are facing potentially distinct prices and have distinct preferences, but they all face a price that is influenced by the origin sector’s size $L_{i,k}$ to the extent that external economies of scale are active.¹³

¹² It should be clear that the absence of internal decreasing returns to labor at the sector-level plays a key role in establishing eqq. (19) and (21). If such forces were active, our empirical strategy would identify only the combination of external and internal scale effects.

¹³ While eq. (21) focuses on averages across all destinations, it would be equally valid from a theoretical standpoint to use averages across any subsets of destinations with positive trade flows. For instance, our model also implies

$$\frac{1}{J} \sum_{j \neq i} \text{DD}_{i_0,k_0} \{ \ln u_{ij,k}(c_j) \} = -\text{DD}_{i_0,k_0} \{ \ln E_k(L_{i,k}) \} - \frac{1}{J} \sum_{j \neq i} \text{DD}_{i_0,k_0} \{ \ln \alpha_{ij,k} \}.$$

In practice, this leave-out specification, excluding domestic sales, leads to estimates of external economies of scale that are very similar to the baseline estimates reported in table 1.

Equation (21) is the starting point of our empirical analysis. Given existing estimates of consumer's preferences, $u_j(\cdot)$, this equation corresponds to a nonparametric regression of the form $y = h(x) + \xi$. The left-hand-side variable y can be measured using data on consumption choices $c_{i,k}$ or, as we do below, data on expenditures, $X_{ij,k} \equiv p_{ij,k} c_{ij,k}$. The first term on the right-hand side, which depends on external economies of scale, $E_k(\cdot)$, is the unknown function h to be estimated; it is evaluated at observable industry sizes $L_{i,k}$ corresponding to x . And the second term, which depends on $(1/J)\sum_j \ln \alpha_{ij,k}$, is the structural error term ξ . While $L_{i,k}$ is endogenous and hence may be correlated with $(1/J)\sum_j \ln \alpha_{ij,k}$, suitable IVs can be used to achieve nonparametric identification of $E_k(\cdot)$, as in Newey and Powell (2003).¹⁴

B. Parametric Model

In theory, one could point identify external economies of scale, $E_k(\cdot)$, by tracing out nonparametrically the impact of exogenous changes in sector sizes $\{L_{i,k}\}$ on consumption choices $\{c_j\}$. In practice, our dataset will include only 61 countries. So, estimation inevitably needs to proceed parametrically, as we now describe.

Functional form assumptions.—We assume that external economies of scale satisfy

$$E_k(L_{i,k}) = L_{i,k}^{\gamma_k} \quad (22)$$

Hence, scale elasticities are constant within each sector, $\epsilon_{i,k}^{E*} = \gamma_k$, but may vary across sectors, as is critical for industrial policy considerations, as seen in propositions 1 and 2. Similarly, we restrict preferences to be nested constant elasticity of substitution (CES):

$$u_j = \left(\sum_k \beta_{j,k}^{1/(\rho+1)} C_{j,k}^{\rho/(\rho+1)} \right)^{(\rho+1)/\rho}, \quad (23)$$

$$C_{j,k} = \left(\sum_i c_{ij,k}^{\theta_k/(1+\theta_k)} \right)^{(1+\theta_k)/\theta_k}, \quad (24)$$

with $1 + \theta_k$ the elasticity of substitution between goods from different origins, $1 + \rho$ the elasticity of substitution between goods from different sectors, and the taste shifters $\beta_{j,k}$ normalized such that $\sum_k \beta_{j,k} = 1$ for all j .¹⁵ As

¹⁴ The interested reader can find further details in our working paper (Bartelme et al. 2019).

¹⁵ The absence of bilateral taste shifters $\beta_{j,k}$ in eq. (24) is without loss of generality, as units of account can always be chosen such that it holds, with the productivity term $A_{ij,k}$ adjusted accordingly. We come back to this point briefly in sec. III.D.

is well known, such preferences give rise to a so-called gravity equation for within-sector trade flows, with θ_k the (sector-level) trade elasticity.¹⁶

Baseline specification.—As established in appendix section B.1, substituting for u_j and E_k using equations (22)–(24), we can rewrite equation (21) as

$$Y_{i,k} = \delta_i + \delta_k + \gamma_k \ln L_{i,k} + \epsilon_{i,k}, \quad (25)$$

where $Y_{i,k} \equiv ((1/J)\sum_j \ln X_{ij,k})/\theta_k$ is the log expenditure $X_{ij,k}$ on goods from country i in sector k , averaged across all destinations and adjusted by the trade elasticity θ_k , and δ_i and δ_k are country and sector fixed effects, respectively. The productivity shock, $\epsilon_{i,k} \equiv (1/J)\sum_j \ln \alpha_{ij,k} - \mathbb{E}[(1/J)\sum_j \ln \alpha_{ij,k} | i] - \mathbb{E}[(1/J)\sum_j \ln \alpha_{ij,k} | k] + \mathbb{E}[(1/J)\sum_j \ln \alpha_{ij,k}]$, is demeaned so that $\mathbb{E}[\epsilon_{i,k} | i] = 0$ for all i and $\mathbb{E}[\epsilon_{i,k} | k] = 0$ for all k , where $\mathbb{E}[\epsilon_{i,k} | i]$ refers to expectation, holding i fixed (and hence the expectation is taken across k only), and analogously for $\mathbb{E}[\epsilon_{i,k} | k]$.

To gain intuition about how equation (25) identifies scale elasticities, consider a hypothetical dataset with only two sectors, $k \in \{\text{food, textiles}\}$; three exporters, $i \in \{\text{Canada, Costa Rica, France}\}$; and one importer, $j \in \{\text{United States}\}$. Suppose that these three exporters have the same productivity in all sectors, $\epsilon_{i,k} = 0$, but that because of demand-side considerations, they have different sector sizes: $L_{\text{Fra,Food}}/L_{\text{Can,Food}} > L_{\text{Fra,Text}}/L_{\text{Can,Text}} = 1$ and $L_{\text{Cos,Text}}/L_{\text{Can,Text}} > L_{\text{Cos,Food}}/L_{\text{Can,Food}} = 1$.¹⁷ Differencing out the country and sector fixed effects in equation (25) then implies

$$\gamma_{\text{Food}} = \frac{(\ln X_{\text{FraUS,Food}} - \ln X_{\text{CanUS,Food}})/\theta_{\text{Food}} - (\ln X_{\text{FraUS,Text}} - \ln X_{\text{CanUS,Text}})/\theta_{\text{Text}}}{\ln L_{\text{Fra,Food}} - \ln L_{\text{Can,Food}}},$$

$$\gamma_{\text{Text}} = \frac{(\ln X_{\text{CosUS,Text}} - \ln X_{\text{CanUS,Text}})/\theta_{\text{Text}} - (\ln X_{\text{CosUS,Food}} - \ln X_{\text{CanUS,Food}})/\theta_{\text{Food}}}{\ln L_{\text{Cos,Text}} - \ln L_{\text{Can,Text}}}.$$

Intuitively, for a given difference in sector size, $\ln L_{\text{Fra,Food}} - \ln L_{\text{Can,Food}} > 0$, the scale elasticity in the food sector γ_{Food} is larger if (i) the country with a large food sector exports more to the United States than the country with a small food sector, that is, $\ln X_{\text{FraUS,Food}} - \ln X_{\text{CanUS,Food}}$ is larger; (ii) US imports in the food sector are less responsive to changes in prices, that is, θ_{Food} is lower; and (iii) the country with a large food sector has a higher wage, as revealed by its relative exports in the textile sector, that is,

¹⁶ While the preferences in (23) and (24) correspond to an Armington model of trade, Costinot, Donaldson, and Komunjer's (2012) multisector extension of Eaton and Kortum's (2002) Ricardian model delivers observationally equivalent gravity equations. The same microfoundations can be invoked in the presence of external economies of scale, as in Kucheryavy, Lyn, and Rodríguez-Clare (2023), with identical implications for our analysis, as shown in our working paper (Bartelme et al. 2019).

¹⁷ The assumption that France and Costa Rica have the same employment as Canada in the textile and food sectors, respectively, simplifies the algebra but is not critical for identification. All that matters is that France and Costa Rica have different distributions of employment, $L_{\text{Fra,Food}}/L_{\text{Fra,Text}} \neq L_{\text{Cos,Food}}/L_{\text{Cos,Text}}$.

$(\ln X_{\text{FraUS,Text}} - \ln X_{\text{CanUS,Text}})/\theta_{\text{Text}}$ is lower. The scale elasticity in the textiles sector γ_{Text} is identified in a similar manner by comparing the (US-bound) exports of Costa Rica with those of Canada.

This example operationalizes the general idea, described in section III.A, that one can identify scale elasticities, that is, by how much increases in employment lowers unit costs and, in turn, prices by tracing out how changes in employment affect consumption choices and, in turn, expenditures. All that is required to go from the latter to the former is an estimate of demand that maps changes in expenditures into changes prices. Under our parametric assumptions, this boils down to knowledge of the trade elasticities θ_k .

In practice, of course, bilateral trade flows may depend on productivity shocks, $\epsilon_{i,k} \neq 0$, and as noted at the end of section III.A, sector size $\ln L_{i,k}$ may respond endogenously to these shocks, $\mathbb{E}[\ln L_{i,k} \times \epsilon_{i,k} | k] \neq 0$ for all k . Estimation of the vector of supply-side parameters $\{\gamma_k\}$ thus requires at least as many demand-side IVs. We now describe a procedure for constructing such variables.

C. Instrumental Variables

In our model, sector size $L_{i,k}$ is an endogenous object determined by countries' factor supply, preferences, technology, and taxes. In addition, the adjusted productivity residual $\epsilon_{i,k}$ in our baseline specification (25) depends on technology and taxes. To construct IVs, we propose to exploit variation in countries' factor supply and preferences, which we will assume to be orthogonal to the unobserved variation in technology and taxes.

To fix ideas, consider first the special case where upper-level preferences are Cobb-Douglas ($\rho \rightarrow 0$). In the absence of trade, the level of sectoral employment $L_{i,k}$ predicted by our model would then be $\beta_{i,k} \times L_i$, where $\beta_{i,k}$ can be directly obtained as the share of expenditure by country i on goods from sector k across all origins, $x_{i,k} \equiv \sum_j X_{ji,k} / \sum_{j,l} X_{ji,l}$, and L_i can be proxied by country i 's population, \hat{L}_i . Given the magnitude of trade costs in practice, we would expect $x_{i,k} \times \hat{L}_i$ to remain a good predictor of $L_{i,k}$ in the presence of trade in this Cobb-Douglas special case.

Our IVs generalize the previous idea to the case where ρ may be different from zero. The key difference between the Cobb-Douglas and CES case is that the demand shifters $\beta_{i,k}$ used to predict sectoral employment no longer satisfy $\beta_{i,k} = x_{i,k}$. Under CES, equations (23) and (24) instead imply

$$\beta_{i,k} = \frac{x_{i,k} / (P_{i,k})^{-\rho}}{\sum_l x_{i,l} / (P_{i,l})^{-\rho}}, \quad (26)$$

where $P_{i,k} \equiv (\sum_j p_{ji,k}^{-\theta_i})^{-1/\theta_i}$ is sector k 's price index in country i . Provided that ρ is different from zero, the adjustment by $1/(P_{i,k})^{-\rho}$ is required to purge expenditure shares from prices and, in turn, the technological and tax considerations that may affect them. Section IV.B describes in detail how we obtain estimates of sector-level price indexes, $\hat{P}_{i,k}$, and the elasticity of substitution, $\hat{\rho}$, in order to recover estimates of the demand residuals $\hat{\beta}_{i,k}$ via equation (26). Given such estimates, we then construct a measure of demand-predicted sector size as $\hat{L}_{i,k} \equiv \hat{\beta}_{i,k} \times \hat{L}_i$, which will serve as the basis for constructing our IVs as the interaction of sector indicator variables and $\hat{L}_{i,k}$. This will deliver consistent estimates of the parameters $\{\gamma_k\}$ under the exclusion restriction:

$$\mathbb{E}[\ln \hat{L}_{i,k} \times \epsilon_{i,k} | k] = 0 \text{ for all } k. \quad (27)$$

This requires that there be no systematic relationship, within any sector k , between a country's demand-predicted sector size $\ln \hat{L}_{i,k}$ and the productivity shock $\epsilon_{i,k}$.¹⁸

D. Threats to the Exclusion Restriction

By construction, the (log) demand-predicted sector size variable consists of two parts, $\ln \hat{L}_{i,k} = \ln \hat{L}_i + \ln \hat{\beta}_{i,k}$. There are therefore two potential sources of violation for the above exclusion restriction:

$$\mathbb{E}[\ln \hat{L}_i \times \epsilon_{i,k} | k] \neq 0 \text{ for some } k, \quad (28)$$

$$\mathbb{E}[\ln \hat{\beta}_{i,k} \times \epsilon_{i,k} | k] \neq 0 \text{ for some } k. \quad (29)$$

Though condition (27) may, in principle, hold when both violations exactly compensate each other, it is hard to imagine such a situation being relevant in practice. We therefore discuss each of these two sources of violation separately.

The first violation (28) arises, for instance, if there are country-wide scale effects that are heterogeneous across sectors, as in $A_{ij,k} \propto (\hat{L}_i)^{\Gamma_i}$. As demonstrated in appendix section B.2, this leads to an overestimate of the extent of external economies of scale γ_k in any sector k with higher than average values of country-wide scale effects, $\Gamma_k > \mathbb{E}[\Gamma_k]$. It should be clear, however, that the source of bias in this scenario is not the existence of country-wide scale effects but their heterogeneity across sectors.

The second violation arises if countries that have a higher propensity to buy in some sectors because of preference considerations also tend to have a higher propensity to sell in those same sectors because of technological

¹⁸ Appendix sec. B.1 explicitly describes the two stages of our IV strategy using dummy variable notation.

or tax considerations. More specifically, suppose that $\hat{\beta}_{i,k} \propto (\bar{\alpha}_{i,k})^\phi$, with $\ln \bar{\alpha}_{i,k} \equiv (1/J) \sum_j \ln \alpha_{ij,k} - \mathbb{E}[(1/J) \sum_j \ln \alpha_{ij,k} | i]$. As demonstrated in appendix section B.2, if $\phi > 0$, this leads to an overestimate of the extent of economies of scale and more so in sectors where $\text{Var}(\epsilon_{i,k} | k)$ is higher. It is worth noting that the issue here is not the standard concern about supply and demand residuals being correlated because of quality considerations. Such a correlation would affect the structural interpretation of our tax-adjusted productivity measures $\alpha_{ij,k}$ but leave the rest of our analysis unchanged.¹⁹ The source of bias here is more subtle. It requires, for instance, countries with lower costs in a given sector k to develop tastes that are tilted toward goods from that sector (which may be important for cultural goods subject to habit formation, like particular food items in Atkin [2013], but is not something we expect to be prevalent among the 15 manufacturing sectors that we consider) or countries with stronger tastes for consumption in sector k to impose policies set to favor production in that sector (despite optimal policy in our model being a function of only the elasticities θ_k and γ_k , as discussed in sec. V).

The previous discussion implicitly focuses on biases arising from the structural parameters, L_i and $\beta_{i,k}$, being systematically correlated with the productivity shocks, $\epsilon_{i,k}$. Even if they are not, our exclusion restriction may also fail if the difference between the true parameters, L_i and $\beta_{i,k}$, and our proxies for these parameters, \hat{L}_i and $\hat{\beta}_{i,k}$, is itself correlated with $\epsilon_{i,k}$, perhaps because of misspecification of preferences in equations (23) and (24), bias in the estimation of ρ , or some other form of measurement error. To alleviate concerns about these potential sources of biases in our estimates of scale elasticities, section IV.D reports the robustness of our results to the inclusion of controls for various sources of productivity differences across countries and sectors. Section VI further introduces controls for cost differences deriving from multiple factors of production and tradable inputs.

IV. Empirical Results

A. Data

To estimate scale elasticities γ_k in equation (25) and construct our IVs, we need measures of bilateral trade flows $X_{ij,k}$, population \hat{L}_i , and sector size $L_{i,k}$. We discuss each of these in turn.

¹⁹ Formally, suppose that preferences take the form $u_j(\{B_{j,k} c_{ij,k}\}_{i,k})$, with $B_{j,k}$ a measure of the quality of good k sold by country i in country j . The standard concern is that $A_{ij,k}$ and $B_{j,k}$ may be positively correlated. To see that this is irrelevant for our empirical analysis, note that starting from this more general model, one can always define quality-and-tax-adjusted measures of productivity, $\tilde{\alpha}_{ij,k} \equiv A_{ij,k} B_{j,k} (1 - t_{ij,k}^x) / [(1 - s_{i,k})(1 + t_{ij,k}^m)]$, without any further implication.

Bilateral trade flows.—We obtain data on bilateral trade flows $X_{ij,k}$ from the Organization for Economic Cooperation and Development (OECD) Inter-Country Input-Output (ICIO) tables. This source documents bilateral trade among 61 major exporters i and importers j , listed in table C.1 (tables B.1–C.4 are available online). ICIO tables report all bilateral flows, including domestic sales $X_{ii,k}$, in each sector, a feature that we use below to construct aggregate measures of expenditure, $X_{j,k} \equiv \sum_i X_{ij,k}$ and $X_j \equiv \sum_{i,k} X_{ij,k}$, as well as sales, $S_{i,k} \equiv \sum_j X_{ij,k}$ and $S_i \equiv \sum_{j,k} X_{ij,k}$. We focus our empirical analysis on 15 manufacturing sectors k defined at a similar level to the two-digit Standard Industrial Classification code and listed in table C.1.²⁰ Our baseline estimates use ICIO data from 2010, but we report similar estimates from other available cross sections (1995, 2000, and 2005).

Population.—We take our measure of population \hat{L}_i from the population (POP) variable in the Penn World Tables (ver. 9.0). In practice, this variable is highly correlated with alternative measures, such as the total labor force.

Sector size.—According to the baseline model developed here, the total wage bill in a sector is equal to total sales across all destinations, $w_i L_{i,k} = S_{i,k}$. In turn, the share of total employment allocated to a given sector k is equal to the share of its sales, $S_{i,k}/S_i$. Using country i 's population \hat{L}_i as a measure of its labor supply L_i , we can therefore construct sector size $L_{i,k}$ as $(S_{i,k}/S_i)\hat{L}_i$.²¹

B. Estimates of Auxiliary Elasticities

Our estimation of scale elasticities γ_k via equation (25) requires estimates of two auxiliary parameters: (i) trade elasticities θ_k in each sector to construct the dependent variable and (ii) the elasticity of substitution between manufacturing sectors $\rho + 1$ to construct our demand-side instruments. We begin with a discussion of these auxiliary estimates.

Trade elasticities.—Nested CES preferences in (23) and (24) imply the following log-linear relationship between expenditure shares $X_{ij,k}/X_{j,k}$ and relative prices $p_{ij,k}/P_{j,k}$ within any sector k :

²⁰ We omit a sixteenth manufacturing sector—recycling and manufacturing not elsewhere classified—from the estimation, as this label covers a small amount of output in a range of highly heterogeneous activities.

²¹ Given this measure of sector size, any discrepancy between population and labor supply in efficiency units is therefore also implicitly part of the error term $\epsilon_{i,k}$ in eq. (25). That is, if $L_i = (\exp \chi_i)\hat{L}_i$, then $\epsilon_{i,k}$ also includes $\chi_i \gamma_k - \mathbb{E}[\chi_i \gamma_k | i] - \mathbb{E}[\chi_i \gamma_k | k] + \mathbb{E}[\chi_i \gamma_k]$, with our exclusion restriction (27) applying to this term as well—though, as we describe in sec. IV.C, our estimates are robust to the inclusion of controls for per capita GDP interacted with sector indicators, which lends support to this assumption.

$$\ln\left(\frac{X_{ij,k}}{X_{j,k}}\right) = -\theta_k \ln\left(\frac{p_{ij,k}}{P_{j,k}}\right). \quad (30)$$

Equation (30) is at the core of a vast gravity literature, reviewed in Head and Mayer (2013), that provides estimates of trade elasticities θ_k by finding exogenous sources of variation in $p_{ij,k}$ and using fixed effects or other strategies to control for $P_{j,k}$. Rather than reestimating trade elasticities θ_k ourselves by combining the data on bilateral trade flows $X_{ij,k}$ presented in section IV.A with additional data on prices $p_{ij,k}$ or shifters of these prices, such as import tariffs and shipping costs, we use recent estimates from Giri, Yi, and Yilmazkuday (2021). Although the structural model considered by these authors does not feature external economies of scale, the empirical procedure that they use starts from equation (30) and is fully consistent with our model. This is because productivity levels, regardless of whether they are exogenous (as in Giri, Yi, and Yilmazkuday's [2021] model) or endogenous (as in ours), are differenced out in their empirical procedure.²² As such, the trade elasticities that they estimate are fully portable into our analysis. For convenience, we report these estimates in column 6 of table 1, along with their standard errors.

Elasticity of substitution between sectors.—Nested CES preferences in (23) and (24) further imply that the share of expenditure in country j on sector k is given by

$$\ln\left(\frac{X_{j,k}}{X_j}\right) = -\rho \ln\left(\frac{P_{j,k}}{P_j}\right) + \ln\beta_{j,k}, \quad (31)$$

where $P_j = (\sum_s \beta_{j,s} P_{j,s}^{-\rho})^{-1/\rho}$ is the overall price index for country j . We estimate ρ via the following specification:

$$\ln X_{j,k} = \phi_j + \phi_k - \rho \ln \hat{P}_{j,k} + \phi_{j,k}, \quad (32)$$

where ϕ_j is treated as a country fixed effect and ϕ_k is treated as a sector fixed effect. In the absence of data on the price index $P_{j,k}$, we use a proxy obtained by averaging equation (30) across all origins, $\ln \hat{P}_{j,k} \equiv (1/J) \sum_i \ln(X_{ij,k}/X_{j,k})/\theta_k$. This means that, after controlling for ϕ_j and ϕ_k , the structural error term, $\phi_{j,k}$, comprises both the preference shock from equation (31), $\ln\beta_{j,k}$, and the average of the log productivity across all origin countries for that sector and destination, $\ln \hat{\alpha}_{j,k} \equiv (1/J) \sum_i \ln \alpha_{ij,k}$, each demeaned across countries and sectors, as further described in appendix

²² Formally, eq. (30) implies

$$\theta_k = -\frac{\sum_{i,j} [\ln(X_{ij,k}/X_{j,k}) - \ln(X_{ii,k}/X_{i,k})]}{\sum_{i,j} [\ln(p_{ij,k}/P_{j,k}) - \ln(p_{ii,k}/P_{i,k})]},$$

which is the starting point of Giri, Yi, and Yilmazkuday's (2021) simulated method of moments estimation using micro-level price data, as in Simonovska and Waugh (2012).

TABLE 1
ESTIMATES OF SECTOR-LEVEL SCALE ELASTICITIES (γ_k)

Sector	OLS (1)	IV (2)	Reduced Form (3)	First Stage		Trade Elasticity (6)
				F (4)	SW F (5)	
Food, beverages, and tobacco	.26 (.02)	.24 (.02)	.28 (.03)	192.4	1,678.9	3.57 (.27)
Textiles	.22 (.02)	.21 (.02)	.24 (.03)	117.9	2,216.8	4.43 (.31)
Wood products	.20 (.04)	.19 (.04)	.22 (.04)	218.8	1,383.6	4.32 (1.03)
Paper products	.32 (.05)	.31 (.05)	.35 (.05)	138.1	1,608.7	2.97 (.42)
Coke/petroleum products	.10 (.05)	.11 (.05)	.15 (.06)	14.8	472.8	8.94 (5.33)
Chemicals	.23 (.03)	.22 (.03)	.26 (.04)	70.1	1,192.6	3.75 (.51)
Rubber and plastics	.21 (.05)	.20 (.05)	.23 (.05)	154.9	1,595.4	4.13 (1.18)
Mineral products	.20 (.04)	.19 (.03)	.22 (.04)	301.6	1,485.0	5.14 (1.17)
Basic metals	.12 (.05)	.11 (.05)	.13 (.06)	137.1	1,342.2	8.94 (5.33)
Fabricated metals	.19 (.08)	.18 (.09)	.21 (.09)	191.1	2,321.4	5.07 (2.46)
Machinery and equipment	.28 (.05)	.26 (.04)	.31 (.05)	136.8	1,764.6	3.27 (.52)
Computers and electronics	.28 (.04)	.27 (.04)	.32 (.05)	71.8	612.5	3.27 (.52)
Electrical machinery (National Electrical Code)	.27 (.04)	.25 (.04)	.30 (.05)	109.7	1,474.2	3.27 (.52)
Motor vehicles	.20 (.04)	.19 (.04)	.26 (.04)	80.4	1,392.2	4.47 (.78)
Other transport equipment	.20 (.03)	.20 (.04)	.24 (.04)	25.7	846.8	4.47 (.78)

NOTE.—Columns 1 and 2 report the OLS and IV estimates, respectively, of γ_k in eq. (25). Column 3 reports the reduced form coefficients. The instruments are the log of (country population \times sectoral demand shifter), interacted with sector indicators. All regressions in cols. 1–3 include exporter and sector fixed effects. Columns 4 and 5 report the conventional and Sanderson-Windmeijer (SW) F -statistics, respectively, from the first-stage regression corresponding to the endogenous regressor formed by interacting the log of sector size with an indicator for the sector named in each row. Column 6 reports the estimate of θ_k from Giri, Yi, and Yilmazkuday (2021) used to construct $Y_{i,k}$ in eq. (25). Standard errors in cols. 1–3 are clustered at the exporter level and adjusted to account for the uncertainty in the estimates of θ_k , as described in app. sec. B.5. Standard errors are in parentheses.

section B.3. Our price proxy $\hat{P}_{j,k}$ captures the fact that, everything else being equal, an exporter i should tend to have a larger share of sales in destination j and sector k when the price index $P_{j,k}$ in that same destination and sector is high, with the trade elasticity θ_k controlling the steepness of this relationship, $P_{j,k} \propto (X_{i,j,k}/X_{j,k})^{1/\theta_k}$. The logic is the same as in the

formula for the welfare gains from trade in Arkolakis, Costinot, and Rodríguez-Clare (2012), with the market share of domestic producers and the trade elasticity revealing the CES price index in a destination, $P_{j,k} \propto (X_{j,k}/X_{j,k})^{1/\theta_k}$. Since the previous logic can be applied to any exporter, our proxy instead uses the simple average across all i in order to minimize the role of idiosyncratic productivity differences.²³

For standard reasons, ordinary least squares (OLS) estimates of the demand elasticity ρ based on this expression would suffer from simultaneity bias. We therefore use an IV procedure in which the instruments for $\ln \hat{P}_{j,k}$ in (32) are formed from the product of country j 's log population $\ln \hat{L}_j$ and a full set of sector indicators. The relevance of such instruments draws on the supply-side logic of our economies of scale model. In particular, we expect large countries to be relatively productive (and hence have relatively low price indexes $P_{j,k}$) in sectors with relatively large scale elasticities, γ_k ; this, in turn, means that the impact of population on prices should vary in distinct ways across sectors. The validity of this IV procedure relies on the exclusion restriction: $\mathbb{E}[\ln \hat{L}_j \times \phi_{j,k} | k] = 0$. Recalling that the error term includes both $\ln \beta_{j,k}$ and $\ln \hat{\alpha}_{j,k}$, it requires no systematic tendency for larger countries (those with higher values of $\ln \hat{L}_j$) to have stronger preferences for a given sector k (higher values of $\ln \beta_{j,k}$) or access to better technologies to source good k (higher values of $\ln \hat{\alpha}_{j,k}$).²⁴

Our estimate of ρ is reported in table B.2, with the corresponding first-stage results in table B.1. The IV estimate is $\hat{\rho} = -0.13$, implying an elasticity of substitution $\hat{\rho} + 1 = 0.87$ in the complements range, though not statistically different from the Cobb-Douglas case (i.e., $\rho = 0$). This IV estimate is lower than the OLS estimate (of $\hat{\rho} = 2.44$), as is consistent with the presence of increasing returns at the sector level. When supply curves slope downward, positive demand shocks lead to reductions in prices. Hence, the OLS estimate of the impact of prices on expenditure shares, which confounds a truly downward-sloping demand curve with the negative correlation between demand shocks and prices, will be an underestimate of $-\rho$, leading to an overestimate of ρ .

Reassuringly, our IV estimate of $\hat{\rho} = -0.13$ is in the same range as estimates seen in prior work using alternative sources of variation. Across two-digit sectors, Oberfield and Raval (2021) obtain estimates of ρ that

²³ Summing across all i also allows us to control for the impact of wages and sector sizes on expenditure shares through the sector fixed effect ϕ_k , as shown in app. sec. B.3. We come back to the choice of our price proxies in the sensitivity analysis of sec. IV.D.

²⁴ This second condition mirrors the orthogonality condition between population and productivity shocks, $\mathbb{E}[\ln \hat{L}_i \times \epsilon_{i,k} | k] = 0$, discussed in sec. III.D. There we were ruling out systematic advantages of larger countries as sellers in specific sectors, i.e., higher values of $(1/J) \sum_j \ln \alpha_{j,k}$. Here we are also ruling out systematic advantages of larger countries as buyers in those same sectors, i.e., higher values of $(1/J) \sum_i \ln \alpha_{j,k}$.

range from -0.14 to 0.27 depending on the specification. At higher levels of aggregation (i.e., across the broad categories of agriculture, manufacturing, and services), Herrendorf, Rogerson, and Valentinyi (2013) obtain -0.19 using US time series consumption data, and Comin, Lashkari, and Mestieri (2021) estimate -0.50 using cross-country panel data. At lower levels of aggregation (i.e., across four-digit US import data), Redding and Weinstein (2018) estimate $\rho = 0.36$. Sections IV.D and V.D demonstrate how our broad conclusions about the effects of industrial policy are insensitive to using a range of alternative values for ρ that spans the estimates from other work discussed here for both constructing IVs and conducting counterfactual analysis.

C. *Estimates of Scale Elasticities*

We now return to the estimation of scale elasticities γ_k . We begin by reporting OLS estimates of γ_k in column 1 of table 1.

All of these estimates imply precisely estimated economies of scale, that is, $\gamma_k > 0$, but as previously discussed, we expect OLS to deliver biased estimates of true economies of scale. For this reason, we turn to the IV procedure described in section III.C, which relies on the demand-side instruments constructed from $\hat{L}_{i,k} \equiv \hat{\beta}_{i,k} \times \hat{L}_i$, with $\hat{\beta}_{i,k} = [x_{i,k}/(\hat{P}_{i,k})^{-\rho}]/[\sum_l x_{i,l}/(\hat{P}_{i,l})^{-\rho}]$ and $\hat{\rho} = -0.13$ (see table B.2).

The IV estimates of γ_k are reported in column 2 of table 1, with summary statistics indicating the strength of the instruments in columns 4 and 5 and the first-stage coefficients themselves summarized in table B.3.²⁵ These are our preferred estimates of the strength of economies of scale within each of the 15 manufacturing sectors in our sample. The results point to substantial economies of scale—with an average scale elasticity of 0.21 —that are statistically significantly different from zero in every sector. At the same time, there is widespread heterogeneity, with estimates ranging from $\gamma_k = 0.11$ in the coke/petroleum and basic metals sectors to $\gamma_k = 0.31$ in the paper products sector.²⁶ We can reject the hypothesis of coefficient equality at the 1% level.

²⁵ For each first-stage regression, table B.3 reveals that the demand-based IV for any given sector has a strong correlation with its own sector size and a far weaker correlation with any other sector's size. Consequently, as seen in col. 4 of table 1, the conventional F -statistic from the 15 instruments in each first-stage equation is large, and the Sanderson and Windmeijer (2016) F -statistic in col. 5—which assesses the extent to which each first stage is affected by independent variation in the instruments from that in the other 14 first stages—is considerably larger. Figure B.1 (figs. B.1–C.1 are available online) provides a visualization of the first-stage relationships for each sector, and fig. B.2 does the same for the corresponding reduced-form relationships.

²⁶ This finding of significant economies of scale in manufacturing is consistent with prior estimates using alternative empirical strategies. For example, Caballero and Lyons (1992) estimate a scale elasticity for (pooled) US manufacturing sectors of 0.07 – 0.29 depending on the instrument used, Basu and Fernald (1997) estimate that the equivalent of γ_k for

To understand what moments in the data lead us to the previous conclusion, note that estimating γ_k in our baseline specification is equivalent to estimating $\xi_k \equiv \theta_k \gamma_k$ in the following alternative specification:

$$\frac{1}{J} \sum_j \ln X_{ij,k} = \theta_k \delta_i + \theta_k \delta_k + \xi_k \ln L_{i,k} + \theta_k \epsilon_{i,k},$$

where we have multiplied both the left- and the right-hand-side variables of equation (25) by the trade elasticity θ_k . Our IV estimates in basic metals and paper products are $\xi_{\text{Metals}} = 0.98$ and $\xi_{\text{Paper}} = 0.92$. Since estimates of trade elasticities in these two sectors are $\theta_{\text{Metals}} = 8.94$ and $\theta_{\text{Paper}} = 2.97$ (see col. 6 of table 1), we conclude that the scale elasticities must be significantly smaller in metals than in paper: $\gamma_{\text{Metals}} = 0.98/8.94 = 0.11 \ll \gamma_{\text{Paper}} = 0.92/2.97 = 0.31$. Intuitively, if changes in employment have similar effects on bilateral trade flows in the two sectors ($\xi_{\text{Metals}} \approx \xi_{\text{Paper}}$), but bilateral trade flows are much less sensitive to changes in costs in paper than in metals ($\theta_{\text{Paper}} \ll \theta_{\text{Metals}}$), then the impact of changes in employment on costs must be much smaller in metals ($\gamma_{\text{Metals}} \ll \gamma_{\text{Paper}}$).²⁷

Another feature of these estimates is that in all sectors but one, the OLS estimate of γ_k is larger than its corresponding IV estimate. This upward OLS bias is to be expected in an open economy in which countries specialize in sectors where they have a comparative advantage. However, since ρ is close to zero and trade volumes are low for the average country, the OLS bias is not large, ranging between 0 and 0.02 across sectors in table 1.

Finally, column 3 reports the reduced-form parameter estimates of the impact of predicted (log) domestic demand $\ln \hat{L}_{i,k}$ in each sector on the dependent variable, $Y_{i,k} \equiv ((1/J) \sum_j \ln X_{ij,k}) / \theta_k$. We see that countries with higher predicted domestic demand in a sector tend to have larger values of $Y_{i,k}$ in that sector. This is again consistent with the existence of increasing returns at the sector level, which implies that positive shocks to domestic demand cause lower prices and, in turn, greater exports. This is a manifestation of the home market effect.

(a weighted average of) US manufacturing sectors is 0.06, Antweiler and Trefler (2002) use international trade data and a Heckscher-Ohlin-Vanek approach to estimate scale elasticities ranging from 0–0.40 among the set of manufacturing sectors in which precise estimates are obtained, and Costinot et al. (2019) estimate a value of $\gamma_k = 0.25$ for the global pharmaceutical sector.

²⁷ Across sectors, our empirical results imply less variation in the elasticities ξ_k of average bilateral trade flows with respect to employment, which range from 0.82 to 0.98, than in the estimates of trade elasticities θ_k , which range from 2.97 to 8.94. As a result, the correlation between scale elasticities γ_k and trade elasticities θ_k is -0.91 .

D. Sensitivity to Alternative Samples and Specifications

As characterized by propositions 1 and 2, the estimates of scale economies γ_k in table 1 shape both the structure of and the gains from optimal industrial policy. Before turning to such implications of our γ_k estimates, we describe three exercises that probe their sensitivity to alternative samples and specifications.

Heterogeneity over time.—The estimates of γ_k in table 1 are obtained from a single cross section (2010). As a first robustness check, we reestimate the scale elasticity parameters γ_k in three additional cross sections (from 1995, 2000, and 2005) as well as a specification that pools across all years 1995–2010. These estimates, displayed in table B.4, are very similar to each other and to those reported in table 1; indeed, in all 3 years, the correlation with our 2010 baseline is 0.99. This finding highlights how our conclusions are not specific to one particular year.

Alternative instrumental variables.—To construct our instrumental variables, we use demand residuals $\hat{\beta}_{i,k} = [x_{i,k}/(\hat{P}_{i,k})^{-\hat{\rho}}]/[\sum_l x_{i,l}/(\hat{P}_{i,l})^{-\hat{\rho}}]$, whose value depends on both our estimate of the upper-level elasticity $\hat{\rho}$ as well as our price proxies, $\hat{P}_{i,k}$. Our second robustness exercise explores the sensitivity of our estimates of γ_k to using alternative demand residuals. We start by recomputing $\hat{\beta}_{i,k}$ under alternative values of $\hat{\rho} = -0.9, 0$ and 3 , a wide range that spans prior estimates from the literature. It includes the Cobb-Douglas special case, $\hat{\rho} = 0$, for which no price proxies are required. Even if there were no measurement error in our price proxies, $\hat{P}_{i,k} = P_{i,k}$, using a value of $\hat{\rho}$ that is distinct from its true value ρ creates a wedge between the estimated and the true demand residuals, $\hat{\beta}_{i,k}/\beta_{i,k} = (P_{i,k})^{\hat{\rho}-\rho}/[\sum_l \beta_{i,l} (P_{i,l})^{\hat{\rho}-\rho}]$. Thus, even if the true demand residuals satisfy the exclusion restriction, $\mathbb{E}[\ln \beta_{i,k} \times \epsilon_{i,k} | k] = 0$, this restriction may not be satisfied for the estimated demand residuals, $\mathbb{E}[\ln \hat{\beta}_{i,k} \times \epsilon_{i,k} | k] \neq 0$, leading to a bias in our IV estimates of scale elasticities γ_k whose magnitude would scale up with the difference between $\hat{\rho}$ and ρ . The estimates of γ_k displayed in columns 2–4 of table B.5 conform with that observation, though they are all close to our baseline estimates, also reported in column 1 of table B.5. These small differences are consistent with the small difference between the OLS and IV estimates of γ_k documented in table 1.

The fact that we use price proxies $\hat{P}_{i,k}$ rather than the true prices $P_{i,k}$ is another source of measurement error in the construction of the demand residuals $\hat{\beta}_{i,k}$ and a potential source of bias in the IV estimation of γ_k . Equation (30) implies that the log difference between the two prices is equal to the average of the log prices across exporters in a given sector and destination, $\ln \hat{P}_{i,k} - \ln P_{i,k} = -(1/J) \sum_j \ln p_{ji,k}$. Hence, even if we knew the true upper-level elasticity, $\hat{\rho} = \rho$, there would still be a wedge between the estimated and the true demand residuals, $\hat{\beta}_{i,k}/\beta_{i,k} = \exp(- (1/J) \sum_j \ln p_{ji,k})^\rho / [\sum_l \beta_{i,l} \exp(- (1/J) \sum_j \ln p_{jl,k})^\rho]$. This would lead to

upward bias in the estimates of γ_k if consumers in more productive countries also tend to face lower average prices (e.g., because country i 's local productivity in sector k , $A_{i,k}$, mechanically lowers domestic prices $p_{i,k}$) or downward bias if the opposite pattern prevails (e.g., because country i 's higher productivity in sector k tends to lower its neighbors' employment, $L_{j,k}$, and, in turn, to raise their export prices, $p_{j,k}$). To address this endogeneity concern, we consider an alternative price proxy $\hat{P}_{i,k}^{alt} \equiv (1/J) \sum_{j \in D(i)} \ln(X_{j,i,k}/X_{i,k})/\theta_k$ that builds on the same strategy as in section IV.B of inferring price indexes from exporters' market shares but incorporates only the exporters located in a "doughnut" $D(i)$ that excludes all countries whose distance to country i is less than half the median distance between countries in our sample. The estimates of γ_k using these alternative price proxies are reported in column 5 of table B.5. Reassuringly, they are very similar to the baseline estimates reported in column 1.

Additional controls.—As discussed above, our estimation procedure requires that each country's comparative advantage—deriving from cost differences across sectors k and exporting countries i , which are absorbed in the error term $\epsilon_{i,k}$ of equation (25)—is orthogonal to our demand-side instruments. One test of this requirement is that adding observable proxies for comparative advantage should not substantially change our estimates of γ_k . We report here several variants of this idea that are based on Ricardian comparative advantage, deriving from productivity differences alone, consistent with the model of section II. Section VI presents additional robustness checks based on other sources of cost differences (coming from multiple factors of production and tradable intermediate inputs) across countries and sectors.

A prominent source of Ricardian comparative advantage stems from heterogeneity in institutions across countries and the differential implications that those institutions have for productivity across sectors. Following the empirical literature surveyed by Nunn and Trefler (2014), we model this as an interaction between two components: (i) country-level proxies for institutional quality and (ii) sector-level characteristics. For the country component, we use both a measure of contract enforcement and a measure of financial development, thereby encompassing the sources of comparative advantage stressed in Levchenko (2007) and Nunn (2007) as well as Beck (2002) and Manova (2008).²⁸ For the sector component, we use a full set of sector indicator variables, which has the benefit of controlling for any form of systematic Ricardian comparative advantage based on contract enforcement or financial development. As an extension to this procedure, we add regressors formed from interactions between

²⁸ Following Nunn (2007), we measure contract enforcement by the rule of law variable due to Kaufmann, Kraay, and Mastruzzi (2003; in 1997–98) and financial development by the (log of the) ratio of private bank credit to GDP (in 1997).

the exporter's per capita GDP and sector indicators, which hence controls for any potential reason for relatively rich countries to be differentially productive in certain sectors.²⁹

The results from both of these exercises are reported in table B.6. In each case, we see only minor changes in the estimates of γ_k as compared with our baseline (a correlation of 0.99 in both extensions), suggesting that our IV strategy utilizes variation in sector size driven by factors that are orthogonal to observable sources of Ricardian comparative advantage. This lends credence to the view that our estimates draw on only demand-side variation in order to identify the supply-side scale economies γ_k .

Discussion.—The findings in this section show that our estimates of scale elasticities γ_k remain stable across different samples and specifications. However, before turning to our counterfactual analysis, it is important to stress again that these estimates are directly affected by the auxiliary estimates of trade elasticities θ_k used in our empirical analysis for the reasons discussed in section IV.C.³⁰ This issue is critical for the exact structure of optimal policy, as described in proposition 1, since the rankings of scale and trade elasticities determine the ranking of employment and export subsidies, respectively. It is not a priori obvious, however, that precise knowledge of which sectors have high and low trade elasticities is equally important for the aggregate welfare consequences of optimal policies, as described in proposition 2. We will come back to this important point in section V.D.

V. The Textbook Gains from Industrial Policy

A. The Calibrated Economy

We focus on a world economy that comprises the 61 countries from the ICIO dataset, the 15 manufacturing sectors for which we have estimated scale economies, as well as the 17 nonmanufacturing sectors also included in the ICIO dataset. Throughout our quantitative exercise, we maintain the parametric restrictions imposed in equations (22)–(24). Thus, scale elasticities are given by γ_k , whereas trade elasticities are given by θ_k , and the elasticity of substitution between sectors is given by $1 + \rho$.

²⁹ We obtain data on per capita GDP by dividing the real output variable (RGDPO) in the Penn World Tables by \hat{L}_i .

³⁰ Our working paper (Bartelme et al. 2019) presents estimates of scale elasticities obtained using alternative estimates of trade elasticities from Caliendo and Parro (2015) and Shapiro (2016). Given the long-run focus of our analysis, we prefer the estimates in Giri, Yi, and Yilmazkuday (2021) that use cross-sectional variation to those in Shapiro (2016) that use high-frequency variation. We also prefer not to rely on the orthogonality assumption invoked by Caliendo and Parro (2015), which precludes (the levels of) tariffs from responding systematically to (the levels of) imports (because of lobbying or other political economy considerations).

To calibrate γ_k , we use the IV estimates reported in column 2 of table 1 for all manufacturing sectors and set $\gamma_k = 0$ for all nonmanufacturing sectors. This implies that there are welfare gains from reallocating resources from nonmanufacturing to manufacturing sectors that have $\gamma_k > 0$, and so the overall gains from industrial policy will be higher than if we had set γ_k in nonmanufacturing to some positive value. To calibrate θ_k , we use the trade elasticities from Giri, Yi, and Yilmazkuday (2021) reported in column 6 of table 1 for each manufacturing sector, in line with the empirical analysis of section IV, and we set $\theta_k = 4.32$ for all nonmanufacturing sectors, which is the median trade elasticity throughout manufacturing. Finally, we set $\rho = -0.13$, in line with the IV estimate for manufacturing sectors reported in column 2 of table B.2. This implies the same elasticity of substitution between manufacturing and nonmanufacturing sectors as within the set of manufacturing sectors. We consider alternative cases in the sensitivity analysis of section V.C.

Following the theoretical analysis of section II.C, we treat each country as small. Under the assumption of nested CES preferences, as a given country j becomes arbitrarily small relative to other countries, sector-level prices and expenditures, $P_{j,k}$ and $X_{j,k}$, become independent of the prices of its exports, $p_{ji,k}$, and the price elasticities $\epsilon_{ij,k}^{p*}$ converge to $-1/(1 + \theta_k)$. Under the normalization $\bar{s}_j = \bar{t}_j = 0$, the optimal industrial and trade policy described in proposition 1 are therefore given by

$$s_{j,k}^* = \frac{\gamma_k}{1 + \gamma_k} \text{ for all } k, \quad (33)$$

$$t_{ji,k}^{x*} = \frac{1}{1 + \theta_k} \text{ for all } k \text{ and } i \neq j, \quad (34)$$

$$t_{ij,k}^{m*} = 0 \text{ for all } k \text{ and } i. \quad (35)$$

Given estimates of the optimal policies, we compute counterfactual equilibria under those policies using exact hat algebra, as in Dekle, Eaton, and Kortum (2008), under the assumption that the initial equilibrium observed in the data (corresponding to 2010, the final year in the ICIO dataset) features neither taxes nor subsidies and that trade deficits D_i correspond to transfers across countries, which we hold fixed in the counterfactual equilibria. The full nonlinear system of equations that determines these counterfactual changes can be found in appendix section C.1.

B. Baseline Results

Table 2 shows the gains from optimal trade and industrial policy for a subset of countries as well as the average across countries, both unweighted and weighted by GDP. Results for each of the 61 countries in our dataset can be found in table C.1.

TABLE 2
GAINS FROM OPTIMAL POLICIES: SELECTED COUNTRIES

Country	Fully Optimal (1)	Trade (2)	Industrial (3)
United States	.82	.51	.59
China	.91	.74	.76
Germany	1.35	1.51	1.08
Ireland	2.72	2.35	2.06
Vietnam	2.07	1.88	1.19
Average:			
Unweighted	1.67	1.46	1.08
GDP weighted	1.06	.87	.76

NOTE.—Column 1 reports the gains associated with fully optimal policies, as defined in eq. (13). Column 2 reports the gains associated with optimal trade policy, as defined in eq. (15). Column 3 reports the gains associated with optimal industrial policy, as defined in eq. (14). All gains are reported as a share (%) of initial income.

Column 1 reports the gains from the fully optimal policy, defined in equation (13) and expressed as a share of initial income. This corresponds to the welfare gains of going from laissez-faire to export taxes equal to $1/(1 + \theta_k)$ and employment subsidies equal to $\gamma_k/(1 + \gamma_k)$. The separate gains from optimal trade and industrial policy—defined in equations (14) and (15) and again expressed as a share of initial income—are reported in columns 2 and 3, respectively.

The results in table 2 reveal that the gains from fully optimal policy (col. 1) are on average 1.67%. The gains from optimal industrial policy (col. 3) are smaller than those from optimal trade policy (col. 2): focusing on the unweighted average across countries, these gains are 1.08% and 1.46%, respectively. Interestingly, the gains from fully optimal policy are sometimes lower than the gains from optimal trade policy, reflecting the fact that using only industrial policy and ignoring their terms of trade implications leads to large welfare losses in some countries.

The gains from industrial policy are closely aligned with those predicted by the areas of Harberger triangles in proposition 2, as can be seen in figure 2, which reports the exact gains from industrial policy generated by the model plotted against the second-order approximation in equation (17). A linear regression of the exact gains on our second-order approximation yields a slope equal to 0.93 and an R^2 equal to 0.97.³¹

Modest gains from industrial policy do not reflect modest wedges. According to our estimates of scale elasticities, if all labor were to reallocate to the manufacturing sector with the largest scale elasticity $\gamma_k = 0.31$, the average welfare gains predicted by the area of the Harberger triangles

³¹ Harberger triangles also provide good approximations for the gains from fully optimal policy and optimal trade policy, with corresponding slope and R^2 given by 1.49 and 0.87 for the fully optimal policy and 0.66 and 0.84 for optimal trade policy.

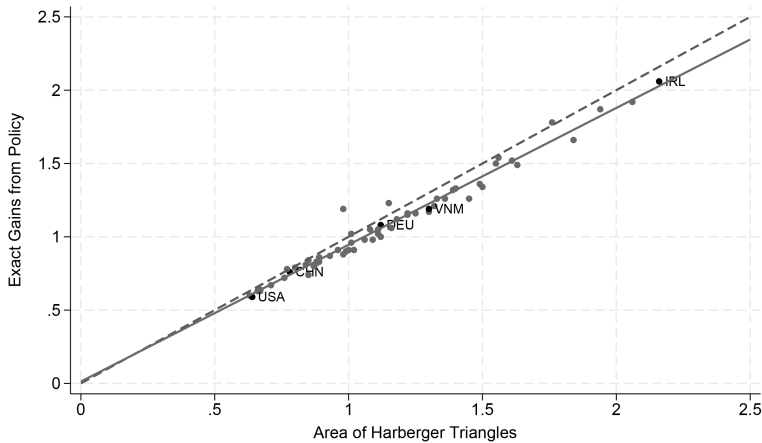


FIG. 2.—Exact policy gains versus areas of Harberger triangles. The figure reports the exact gains from industrial policy (as a share of initial income) on the y -axis, as defined in equation (14), against the second-order approximation from proposition 2 on the x -axis, as described in equation (17). Solid line = line of best fit; dashed line = 45° line.

would be equal to 12.4%.³² Modest gains instead reflect the fact that only modest labor reallocations take place from low-wedge to high-wedge sectors. This is due to a relatively low elasticity of substitution between sectors and to trade elasticities that are far from infinity. Going back to the textbook case for industrial policy illustrated in figure 1, our empirical analysis points to Harberger triangles that have nontrivial heights but small bases.

Across countries, our quantitative results nevertheless exhibit substantial heterogeneity. Smaller countries, in particular, tend to gain more from industrial policy than larger ones. This is revealed by the fact that the simple average in column 3 of table 2 is higher than the corresponding GDP-weighted average. As an example, Ireland has gains from industrial policy that are more than three times higher than those of the United States (2.06% vs. 0.59%).³³

A key difference between small and large countries is their degree of openness to trade. Smaller countries are more open, and more open countries gain more from optimal industrial policy. This is illustrated in figure 3,

³² Specifically, the sector with the largest elasticity $\gamma_k = 0.31$ is paper products. If labor were to fully reallocate to that sector, the welfare change in a given country j predicted by eq. (17) would be $(1/2) \times [(1 - L_{j,\text{Paper}}/L_j) \times 0.31 - \sum_{k \neq \text{Paper}} (L_{j,k}/L_j) \gamma_k]$. The average of the previous number across all countries in our dataset is equal to 0.124.

³³ Smaller countries also gain more from trade policy and the fully optimal policy, as can also be seen by comparing the simple averages in cols. 1 and 2 with the corresponding GDP-weighted averages. The reason why smaller countries gain more from optimal trade policy is simple: such a policy improves a country's terms of trade, and since small countries tend to trade more, they benefit more from that improvement.

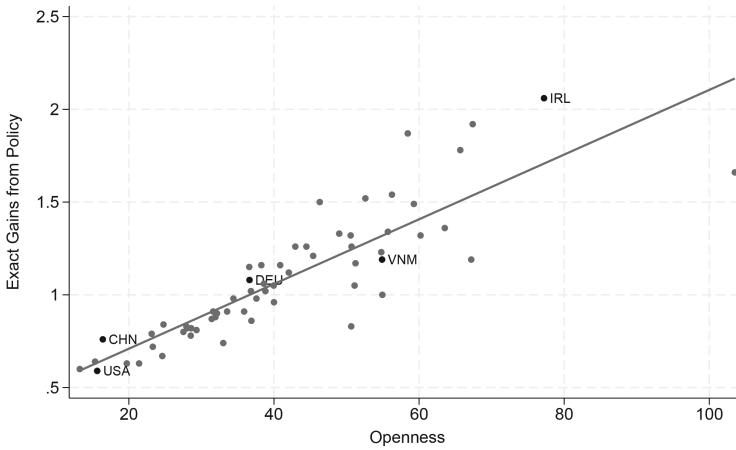


FIG. 3.—Industrial policy gains versus openness. The figure reports the exact gains from industrial policy (as a share of initial income) on the y -axis, as defined in equation (14), against openness on the x -axis, measured as exports plus imports over gross output.

which shows a scatterplot of the gains from industrial policy (y -axis) against openness, measured as exports plus imports over gross output (x -axis). Intuitively, inelastic domestic demand exerts a weaker restraint on labor reallocation in more open economies, and so their Harberger triangles have larger bases, leading to bigger gains from industrial policy. In line with the previous intuition, a cross-country regression of $\sum_k |(\Delta L_{j,k})^l / L_j|$ on openness yields a strong positive relationship, with an R^2 of 0.93.

An additional implication of the Harberger formula described in equation (17) is that for given elasticities $(\Delta L_{j,k})^l / L_{j,k}$, countries with higher employment in sectors with stronger scale economies should benefit more from industrial policy. Our quantitative results are also consistent with this observation: countries with a higher correlation between L_k/L and γ_k do indeed have higher gains from optimal industrial policy. However, this channel accounts for a smaller fraction of the cross-country variation in the gains compared with the openness channel discussed above: a regression of the gains from industrial policy on openness yields $R^2 = 0.74$, while adding the variable $\sum_k L_k/L \cdot \gamma_k$ as a predictor increases the R^2 only to 0.87.

C. Gains from Industrial Policy in the Presence of Trade Agreements

The previous quantitative results assume that countries are free to pursue their unilaterally optimal trade policies. In practice, explicit trade

agreements or implicit threats of foreign retaliation may prevent countries from doing so. How would such considerations affect the gains from industrial policy?

We address this issue under two benchmark scenarios. In the first case, countries are forced to set zero trade taxes but still face incentives to manipulate their terms of trade using industrial policy, as in Lashkaripour and Lugovskyy (2018). In the second, we assume that despite the availability of other policy instruments, trade agreements have been designed to internalize terms of trade externalities and restore global efficiency, as in Bagwell and Staiger (2001).³⁴

Under the first scenario, we numerically find the employment subsidies that maximize utility in a given country conditional on zero trade taxes. As mentioned above, these constrained optimal employment subsidies involve a compromise between the textbook Pigouvian motive of internalizing production externalities and the goal of improving the country's terms of trade.³⁵ Column 2 of table 3 reports the gains from industrial policy under this first scenario.³⁶ For convenience, column 1 reports again the gains from industrial policy when trade policy is unconstrained as well, that is, column 2 of table 2. On average, the gains from this type of constrained industrial policy are a bit lower than those from the optimal industrial policy reported in our baseline, with the unweighted world average decreasing from 1.08% to 0.97%.

Under the second scenario, we assume that employment subsidies are chosen in a Pareto-efficient manner, with lump-sum transfers between countries available if necessary. Hence, only the Pigouvian motive remains, and all countries set employment subsidies satisfying $s_{j,k} = \gamma_k / (1 + \gamma_k)$ for all j, k . The gains from industrial policy (gross of lump-sum transfers between countries, if any) in this case are reported in column 3 of table 3. The GDP-weighted average of the gains associated with this policy is 0.23% but with (gross of transfer) gains of 1.66% in Vietnam and losses of 2.14% in Ireland. Such welfare losses derive from adverse terms of trade effects: larger employment subsidies in sectors with high scale elasticities

³⁴ Another potentially interesting scenario would be to consider the gains from industrial and trade policy in an environment where all countries simultaneously impose policies noncooperatively. Although we stop short of analyzing such Nash equilibria, we expect countries to be less open when they all choose trade taxes in a noncooperative manner and, in turn, the gains from industrial policy to be smaller, consistent with the results of fig. 3. Indeed, when considering counterfactual autarkic economies, we find gains from industrial policy that are 0.41% on average.

³⁵ As an example, consider an economy that exports all its output. In this case, employment subsidies equal to $1 - (1/(1 + \gamma_i))((1 + \theta_i)/\theta_i)$ would perfectly replicate the effect of both the employment subsidies and export taxes in the unconstrained policy case. As long as there is some sector in which part of domestic production is sold at home, however, the constrained optimal employment subsidies will deviate from those and the corresponding gains will be strictly lower than those when policy is unconstrained.

³⁶ Results for all countries in our dataset can be found in table C.2.

TABLE 3
GAINS FROM CONSTRAINED AND GLOBALLY EFFICIENT INDUSTRIAL POLICIES:
SELECTED COUNTRIES

Country	Baseline (1)	Constrained (2)	Globally Efficient (3)
United States	.59	.48	.62
China	.76	.49	.02
Germany	1.08	.60	-.63
Ireland	2.06	1.64	-2.14
Vietnam	1.19	1.28	1.66
Average:			
Unweighted	1.08	.97	.35
GDP weighted	.76	.59	.23

NOTE.—Each column reports the gains—expressed as a share (%) of initial real national income—that could be achieved by each type of policy.

cause an expansion of these sectors and a deterioration of the terms of trade of countries specializing in them.³⁷

D. Sensitivity to Calibrated Parameters

We conclude this section by exploring the sensitivity of our findings with respect to the main parameters used in our quantitative exercise: trade elasticities θ_k , scale elasticities γ_k , and the elasticity of substitution between sectors ρ .

Trade elasticities.—Trade elasticities play a dual role in our analysis. First, they affect the demand functions that we use to estimate scale elasticities γ_k in equation (25) and in turn the structure of optimal industrial policy. Second, they shape the magnitude of the reallocations associated with industrial policy for any value of γ_k and in turn the associated welfare gains. Our baseline estimates used θ_k from Giri, Yi, and Yilmazkuday (2021). We now explore the sensitivity of the gains from industrial policy to alternative calibrations of θ_k that do not rely on the heterogeneity uncovered by Giri, Yi, and Yilmazkuday (2021).

For our first exercise, we vary the extent of heterogeneity in trade elasticities across sectors by setting $\theta_k = \lambda\theta_{k,\text{baseline}} + (1 - \lambda)\text{median}\{\theta_{s,\text{baseline}}\}$, where $\theta_{k,\text{baseline}}$ denotes the estimate of the trade elasticity in sector k from Giri, Yi, and Yilmazkuday (2021) and λ controls the dispersion in trade elasticities across sectors. When $\lambda = 1$, this dispersion is as estimated by Giri, Yi, and Yilmazkuday (2021). When $\lambda = 0$, there is none. For each alternative value of the vector of trade elasticities $\{\theta_k\}$, we then repeat the entire two-step estimation procedure from section IV to estimate new

³⁷ The correlation between the gains from the coordinated industrial policy considered in col. 3 of table 3 and the comparative advantage of countries in high scale elasticity sectors (which we measure simply as the correlation, within each country, between its sector-level net exports and sector scale elasticities) is equal to -0.54 .

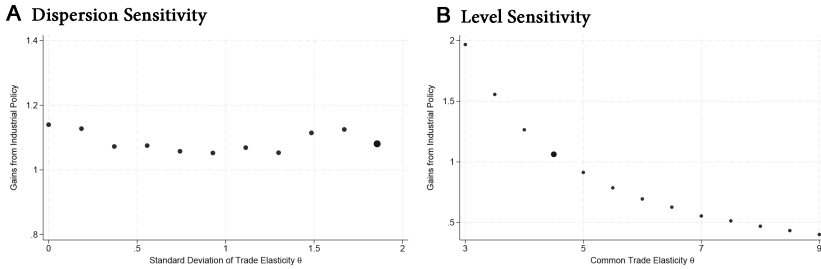


FIG. 4.—Gains from industrial policy: alternative trade elasticities. *A* reports the (global, unweighted average) gains from industrial policy, as defined in equation (14), for different values of the standard deviation of trade elasticities achieved by varying $\lambda \in [0, 1]$ in $\theta_k = \lambda\theta_{k,\text{baseline}} + (1 - \lambda)\text{median}\{\theta_{k,\text{baseline}}\}$. Large circle indicates the trade elasticities used in table 2, that is, $\lambda = 1$. *B* does the same for different values of a common trade elasticity $\theta_k = \theta$. Large circle indicates the median of the estimates from Giri, Yi, and Yilmazkuday (2021), that is, $\theta = 4.32$.

scale elasticities $\{\gamma_k\}$ before recomputing the counterfactual equilibria from section V.A. The results of this exercise are reported in figure 4A, with the standard deviation of θ_k corresponding to each λ on the x -axis. Although the extent of heterogeneity across sectors varies widely across simulations, we see that it has very little effect on the magnitude of the gains from industrial policy.

For our second exercise, we start from the case where there is no heterogeneity in trade elasticities, $\theta_k = \theta$, and vary the common level of the trade elasticity θ around the median value of 4.32 estimated by Giri, Yi, and Yilmazkuday (2021) from their minimum estimate ($\theta_k = 2.97$) to their maximum ($\theta_k = 8.94$). We then again repeat the estimation from section IV and counterfactual analysis from section V.A. Figure 4B shows that gains from industrial policy are more sensitive to the level of the (common) trade elasticity than to its standard deviation, with smaller trade elasticities associated with larger gains from industrial policy. Intuitively, there are two forces at play. On the one hand, smaller trade elasticities lead to more modest labor reallocations, which imply Harberger triangles with smaller bases. On the other hand, smaller trade elasticities also require larger scale elasticities to rationalize the response of exports to sector size observed in the data, which implies Harberger triangles with higher heights. Around our calibrated economy, the second effect dominates the first. For a common trade elasticity of $\theta = 2.97$, which is at the low end of existing estimates in the literature, gains from industrial policy are around 2% on average.³⁸

³⁸ In our working paper (Bartelme et al. 2019), we also report gains from industrial policy using alternate trade elasticities from Caliendo and Parro (2015) and Shapiro (2016). The average gains from industrial policy in the case of Shapiro (2016) remain very similar

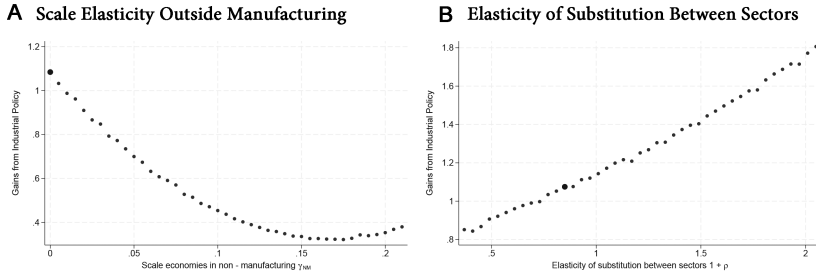


FIG. 5.—Gains from industrial policy: other parameter values. A reports the (global, unweighted average) gains from industrial policy, as defined in equation (14), for different values of the scale elasticity outside manufacturing, γ_{NM} . B does the same for different values of the elasticity of substitution between sectors, $1 + \rho$. Large circles indicate baseline parameter values used in table 2.

Scale economies in nonmanufacturing.—We next consider the implications of the value of scale economies in nonmanufacturing sectors, a parameter that we denote γ_{NM} . Our baseline assumption of $\gamma_{NM} = 0$ —in which only manufacturing sectors exhibit economies of scale—fits the traditional view behind industrial policy. But in the absence of strong evidence to suggest that industrial sectors have superior economies of scale, it is important to explore the quantitative implications of this assumption.

To shed light on this issue, figure 5A plots the average of the estimated gains from optimal industrial policy (across all countries in the world) as a function of γ_{NM} .³⁹ Not surprisingly, as we start raising γ_{NM} from zero, the relative size of manufacturing and nonmanufacturing sectors in the competitive equilibrium gets closer to its optimal value, and the gains from industrial policy fall. Indeed, the gains are minimized when γ_{NM} is around 0.17, close to the average of the manufacturing sector γ_k values of 0.21. Our baseline finding that the gains from optimal industrial policy are relatively modest therefore appears to hold across a range of reasonable values for scale elasticities outside manufacturing, in line with the fact that these modest gains derive from Harberger triangles with modestly sized bases rather than heights.

Elasticity of substitution across sectors.—Finally, we consider the role played by the elasticity of substitution $1 + \rho$. In our baseline analysis, we have used our IV estimate $1 + \hat{\rho} = 0.87$. In figure 5B, we plot the average gains from

to those in our baseline calibration but are around 3.28% on average in the case of Caliendo and Parro (2015) because of extremely low trade elasticity estimates in a small number of sectors.

³⁹ We continue to calculate the gains from optimal policy as described in eqq. (33)–(35); that is, now that $\gamma_{NM} > 0$, the optimal policy includes an employment subsidy in nonmanufacturing sectors.

optimal industrial policy for a wide range of alternative parameters. As discussed above, we expect that a higher ρ would lead to larger reallocations in response to optimal industrial policy and, in turn, larger welfare gains. Qualitatively, this is confirmed in figure 5B. Quantitatively, however, we see that even for an elasticity of substitution of $1 + \rho = 2$, much higher than the recent estimates in Cravino and Sotelo (2019) and Comin, Lashkari, and Mestieri (2021), the average gains from industrial policy are only 1.76%.

VI. Beyond Ricardian Economies

In sections II–V, we have explored the structure and welfare consequences of optimal industrial and trade policy in a Ricardian environment in the tradition of Graham (1923), Chipman (1970), and Ethier (1982). In this final section, we extend our earlier analysis to richer environments featuring multiple factors of production and input-output linkages.

A. General Environment

Consider a generalized version of the environment of section II.A. Each country is now endowed with both labor and physical capital, L_i and K_i , and all production functions now use labor, physical capital, and intermediate goods from other countries and sectors:

$$\begin{aligned} y_{ij,k} &= A_{ij,k} E_k(Z_{i,k}) z_{ij,k}, \\ z_{ij,k} &= f_k(\ell_{ij,k}, k_{ij,k}, m_{ij,k}), \end{aligned}$$

where $m_{ij,k} \equiv \{m_{oij,sk}\}$ denotes the vector of intermediate inputs, $m_{oij,sk}$ from a given origin country o and source sector s used by firms producing good k in country i and exporting it to country j ; $k_{ij,k}$ denotes the amount of physical capital that they use; and f_k is homogeneous of degree 1 in all inputs.⁴⁰ The relevant measure of sector size is now given by the total use of the composite input in country i and sector k :

$$Z_{i,k} = \sum_j z_{ij,k}. \quad (36)$$

In terms of policy, each country j 's government can still tax imports and exports at rates $t_{ij,k}^m$ and $t_{ji,k}^x$ as well as subsidize purchases of the composite input at rate $s_{j,k}$, leading to the arbitrage condition

⁴⁰ We allow for Hicks-neutral differences across countries and sectors but restrict sector-level production functions, f_k , to be constant across countries. In our theoretical analysis, this restriction can be dispensed with. In our quantitative analysis, the assumption of a common f_k across countries helps to avoid computational problems arising from the combination of economies of scale and very imbalanced input-output tables in some countries while still allowing differences in factor endowments to act as a source of comparative advantage.

$$v_{j,k} = (1 - s_{j,k})w_{j,k}, \tag{37}$$

where $w_{j,k}$ denotes the price received by firms from country j producing sector k 's composite input and $v_{j,k}$ denotes the price paid by the firms purchasing it. The rest of the equilibrium conditions can be found in appendix section A.4.

B. Revisiting the Structure of Optimal Policy

Within this environment, we can use the same strategy as in section II.C to characterize the optimal production subsidies, $s_{j,k}^*$, and trade taxes, $t_{ji,k}^{x*}$ and $t_{ij,k}^{m*}$. The same Pigouvian logic leads to the following generalization of proposition 1.

PROPOSITION 3. In the general environment of section VI.A, the optimal industrial and trade policy are such that (i) $s_{j,k}^* = \epsilon_{j,k}^{E*} / (1 + \epsilon_{j,k}^{E*})$, (ii) $t_{ij,k}^{m*} = \bar{t}_j$, and (iii) $t_{ji,k}^{x*} = 1 - (1 + \bar{t}_j)(1 + \epsilon_{ij,k}^{P*})$ for some \bar{t}_j .

The formal proof can be found in appendix section A.4.2. The only difference between propositions 1 and 3 comes from the level of industrial policy. Since the total supply of the composite input generating the externality, $\sum_k Z_{i,k}$, is elastic, the absolute level of the scale elasticity is no longer irrelevant, and there is no longer an indeterminacy in the level of the input subsidy, \bar{s}_j : optimal subsidies now require $\bar{s}_j = 0$.

To go from theory to data, we maintain the same parametric restrictions as in section III.B—that is, E_k and u_j satisfy equations (22)–(24)—and assume that f_k is nested CES:

$$f_k(\ell_{ij,k}, k_{ij,k}, m_{ij,k}) = [(\ell_{ij,k})^{a_k} (k_{ij,k})^{1-a_k}] \prod_s [M_s(\{m_{oij,sk}\}_o)]^{b_{sk}}, \tag{38}$$

$$M_s(\{m_{oij,sk}\}_o) = \left(\sum_o m_{oij,sk}^{\theta_s/(1+\theta_s)} \right)^{(1+\theta_s)/\theta_s}, \tag{39}$$

where $b_k + \sum_s b_{sk} = 1$ and $a_k \in [0, 1]$, with average shares consistent with the OECD ICIO tables. The estimation of scale elasticities γ_k can then proceed following the same steps as in section III given knowledge of the sector-level production function f_k that maps labor, physical capital, and intermediate goods into a country-and-sector-specific composite factor. The counterpart of our baseline specification (25) is

$$Y_{i,k}^{IO} = \delta_i + \delta_k + \gamma_k \ln Z_{i,k} + \epsilon_{i,k}, \tag{40}$$

where $Y_{i,k}^{IO} \equiv (1/J) \sum_j (\ln X_{ij,k} / \theta_k) + (1 - a_k) b_k \ln(r_i / w_i) + \sum_s b_{sk} \ln(P_{i,s} / w_i)$ now also controls for differences in relative factor prices, r_i / w_i , and relative intermediate good prices, $P_{i,s} / w_i$, as described in appendix section B.6. This




FIG. 6.—Optimal industrial policy beyond Ricardian economies. The y-axis reports the optimal industrial policy $s_{j,k}^* = \gamma_k / (1 + \gamma_k)$ estimated for 15 manufacturing sectors in the environment with capital and input-output linkages (as described in sec. VI.A), whereas the x-axis reports the optimal industrial policy $s_{j,k}^* = \gamma_k / (1 + \gamma_k)$ estimated in the baseline environment (as described in sec. II.A).

adjustment is necessary to isolate the impact of external economies of scale on quantities demanded.

Figure 6 plots the optimal industrial policy $s_{j,k}^* = \gamma_k / (1 + \gamma_k)$ estimated in the general environment of section VI.A against the optimal industrial policy $s_{j,k}^* = \gamma_k / (1 + \gamma_k)$ estimated in the Ricardian environment of section II.A for the same 15 manufacturing sectors. All observations are close to but below the 45° line.⁴¹ The average optimal subsidy in manufacturing sectors is about 27% lower in the presence of physical capital and input-output linkages. This reflects the fact that our baseline estimation does not control for variation in the price of intermediate goods. Since each sector and country sources a large fraction of its inputs from itself, countries with larger sectors also tend to have lower input prices. Controlling for input prices in our modified estimating equation (40) removes this source of upward bias in our baseline estimation. Remarkably, the identity of the sectors that should be subsidized remains virtually unchanged, with a correlation between the optimal policies estimated in the two models equal to 0.98.⁴²

⁴¹ All estimates of γ_k remain statistically significant at the 5% level (except basic metals), and the instruments remain strong by conventional standards (with the lowest first-stage Sanderson-Windmeijer F -statistic equal to 100.7). Full estimates are reported in table B.7.

⁴² Since none of the estimates of trade elasticities are affected by the presence of physical capital and input-output linkages, the structure of optimal trade policy is trivially unchanged.

C. *Revisiting the Gains from Optimal Policy*

While the structure of optimal policy is very similar to that in our baseline environment, the consequences of such policy are not. As can be seen from figure 7, the gains from optimal industrial policy are significantly larger in the presence of physical capital and input-output linkages. The average gains estimated here are equal to 4.06%, almost four times larger than in a Ricardian environment.⁴³ To understand why the gains from optimal policy are higher in this more general environment, it is useful to turn to the following generalization of proposition 2.

PROPOSITION 4. In the general environment of section VI.A, up to a second-order approximation, the welfare gains from optimal taxes satisfy

$$\frac{W_j^*}{I_j} \approx \frac{1}{2} \sum_k \frac{w_{j,k} Z_{j,k}}{I_j} \cdot \frac{(\Delta Z_{j,k})^*}{Z_{j,k}} \cdot \epsilon_{j,k}^{E*} + \frac{1}{2} \sum_{i \neq j,k} \frac{\bar{p}_{i,k} y_{i,k}}{I_j} \cdot \frac{(\Delta y_{i,k})^*}{y_{i,k}} \cdot \epsilon_{j,k}^{P*}, \quad (41)$$

where $(\Delta Z_{j,k})^*$ and $(\Delta y_{i,k})^*$ denote the input and export changes associated with imposing τ_j^* and all other variables are evaluated under *laissez-faire*. Likewise, up to a second-order approximation, gains from industrial and trade policy satisfy

$$\frac{W_j^I}{I_j} \approx \frac{1}{2} \sum_k \frac{w_j Z_{j,k}}{I_j} \cdot \frac{(\Delta Z_{j,k})^I}{Z_{j,k}} \cdot \epsilon_{j,k}^{E*}, \quad (42)$$

$$\frac{W_j^T}{I_j} \approx \frac{1}{2} \sum_{i \neq j,k} \frac{\bar{p}_{i,k} y_{i,k}}{I_j} \cdot \frac{(\Delta y_{i,k})^T}{y_{i,k}} \cdot \epsilon_{j,k}^{P*}, \quad (43)$$

where $(\Delta Z_{j,k})^I$ denotes the input change associated with imposing s_j^* , conditional on having already imposed optimal trade taxes, and $(\Delta y_{i,k})^T$ denotes the export change associated with imposing t_j^* , conditional on having already imposed optimal employment subsidies.

The formal proof can be found in appendix section A.4.3. The key distinction between propositions 2 and 4 comes from the shares, $w_{j,k} Z_{j,k}/I_j$ and $\bar{p}_{i,k} y_{i,k}/I_j$, appearing in equations (41)–(43). The numerators, $w_{j,k} Z_{j,k}$ and $\bar{p}_{i,k} y_{i,k}$, are gross flows, since gross output is what is being subsidized by industrial policy and taxed by trade policy. In contrast, the denominator, I_j , still measures total value added in country j . Thus, shares now add up to numbers that are strictly greater than 1, which mechanically raises the gains associated with both industrial and trade policy.

Quantitatively, the mechanical adjustment from value added to gross flows plays an important role in generating these larger gains. Starting

⁴³ The counterpart of the exact hat algebra of sec. V used for the counterfactual analysis in this section can be found in app. sec. C.4, whereas the gains from industrial policy for all countries can be found in table C.3, along with the gains from trade policy and the fully optimal policy.

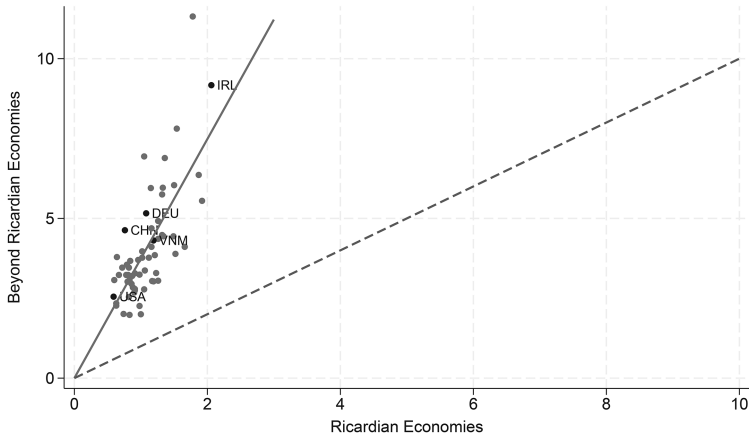


FIG. 7.—Gains from industrial policy beyond Ricardian economies. The figure reports the gains from industrial policy, as defined in equation (14), in the environment with capital and input-output linkages (as described in sec. VI.A) on the y -axis against the gains from industrial policy in the baseline environment (as described in sec. II.A) on the x -axis.

from the second-order approximations in proposition 2 and substituting the shares from proposition 4 while keeping everything else constant reduces the difference between the gains from industrial policy in the two environments by over 80%.⁴⁴

Overall, though, gains from industrial policy remain significantly smaller than the potential gains that one might have expected given the magnitude of scale elasticities that we have estimated (12.4% in the simple back-of-the-envelope calculations of sec. V.B). This reflects the fact that while the wedges that we have estimated empirically are substantial, the cross-sectoral reallocations taking place under the optimal policy are fairly modest in all our quantitative exercises.

VII. Concluding Remarks

In this paper, we have evaluated the textbook case for industrial policy in an open economy. We have shown how to use commonly available trade and production data to estimate economies of scale at the sector-level and, in turn, how to characterize the structure and consequences of optimal industrial policy.

Our estimates imply that even under the optimistic assumption that governments maximize welfare and have full knowledge of the underlying

⁴⁴ Like in the baseline Ricardian environment, our second-order formulas provide good approximations in the general environment of sec. VI.A. This can be seen from fig. C.1.

economy, gains from optimal industrial policy are hardly transformative, ranging from an average across countries of 1.08% of GDP in our baseline analysis to 4.06% in the general environment with physical capital and input-output linkages.

To put these numbers in perspective, recall that South Korea—a country often presented as an industrial policy success story—experienced gains in real GDP per capita of 6.82% per year from 1960 to 1989, as documented in Rodrik (1995). A 4.06% long-run welfare increase is nothing to spit at, but a miracle it is not.

Intuitively, for the textbook gains from industrial policy to be large, two sets of conditions should hold. First, production processes should exhibit significant external economies of scale that differ in strength across sectors, conditions that would give rise to a Harberger triangle with a large height in figure 1. Second, high elasticities of substitution and trade openness should allow significant reallocations of resources across sectors, conditions that would give rise to a Harberger triangle with a large base in figure 1. Across our quantitative exercises, given low elasticities of demand, the nontrivial gap between social and marginal costs that we estimate remains too small to generate large gains from industrial policy, even for the most open economies.

There are, of course, many other market failures that industrial policy may target. Alleviating financial distortions, as in Itskhoki and Moll (2019) and Liu (2019); solving coordination failures, as in Murphy, Shleifer, and Vishny (1989), Rodrik (1996), and Buera et al. (2021); and fostering sectors with positive spillovers to the rest of the economy, as in Greenwald and Stiglitz (2006), may all, in theory, generate transformative gains from industrial policy. Our results, however, offer little empirical support for the notion that these gains can arise from the textbook case based on external economies of scale, at least at the level of aggregation considered in our analysis. An interesting open question is whether allowing for more granular economies of scale within both more narrowly defined sectors and regions may magnify the gains from policy intervention by combining features of optimal industrial and trade policy, as emphasized in this paper, with those of optimal place-based policy, as emphasized in Rossi-Hansberg, Sarte, and Schwartzman (2019) and Fajgelbaum and Gaubert (2020). We hope that the combination of theory and empirics developed in this paper will prove useful in making progress on this and other related questions.

Data Availability

Data and code replicating the tables and figures in this article can be found in Bartelme et al. (2024) in the Harvard Dataverse, <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/JNED5H>.

References

- Antweiler, Werner, and Daniel Treffer. 2002. "Increasing Returns and All That: A View from Trade." *A.E.R.* 92 (1): 93–119.
- Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. "New Trade Models, Same Old Gains?" *A.E.R.* 102 (1): 94–130.
- Atkin, David. 2013. "Trade, Tastes, and Nutrition in India." *A.E.R.* 103 (5): 1629–63.
- Bagwell, Kyle, and Robert W. Staiger. 2001. "Domestic Policies, National Sovereignty, and International Economic Institutions." *Q.J.E.* 116 (2): 519–62.
- Bartelme, Dominick, Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare. 2019. "The Textbook Case for Industrial Policy: Theory Meets Data." Working Paper no. 26193, NBER, Cambridge, MA.
- Bartelme, Dominick, Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare. 2024. "Replication Data for: 'The Textbook Case for Industrial Policy: Theory Meets Data.'" Harvard Dataverse, <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/JNED5H>.
- Basu, S., and J. G. Fernald. 1997. "Returns to Scale in U.S. Production: Estimates and Implications." *J.P.E.* 105:249–83.
- Beck, Thorsten. 2002. "Financial Development and International Trade: Is There a Link?" *J. Internat. Econ.* 57 (1): 107–31.
- Buera, Francisco J., Hugo Hopenhayn, Yongseok Shin, and Nicholas Trachter. 2021. "Big Push in Distorted Economies." Working Paper no. 28561, NBER, Cambridge, MA.
- Caballero, Ricardo J., and Richar K. Lyons. 1992. "External Effects in U.S. Pro-cyclical Productivity." *J. Monetary Econ.* 29 (2): 209–25.
- Caliendo, Lorenzo, and Fernando Parro. 2015. "Estimates of the Trade and Welfare Effects of NAFTA." *Rev. Econ. Studies* 82 (1): 1–44.
- Chipman, John S. 1970. "External Economies of Scale and Competitive Equilibrium." *Q.J.E.* 84 (3): 347–85.
- Combes, Pierre-Philippe, and Laurent Gobillon. 2015. "The Empirics of Agglomeration Economies." In *Handbook of Regional and Urban Economics*, vol. 5, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, 247–348. New York: Elsevier.
- Comin, Diego, Danial Lashkari, and Marti Mestieri. 2021. "Structural Change with Long-Run Income and Price Effects." *Econometrica* 89 (1): 311–74.
- Costinot, Arnaud, Dave Donaldson, and Ivana Komunjer. 2012. "What Goods Do Countries Trade? A Quantitative Exploration of Ricardo's Ideas." *Rev. Econ. Studies* 79 (2): 581–608.
- Costinot, Arnaud, Dave Donaldson, Margaret Kyle, and Heidi Williams. 2019. "The More We Die, the More We Sell? A Simple Test of the Home-Market Effect." *Q.J.E.* 134 (2): 843–94.
- Costinot, Arnaud, and Andrés Rodríguez-Clare. 2014. "Trade Theory with Numbers: Quantifying the Consequences of Globalization." In *Handbook of International Economics*, vol. 4, edited by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, 197–261. New York: Elsevier.
- Costinot, Arnaud, and Iván Werning. 2023. "Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation." *Rev. Econ. Studies* 90 (5): 2261–91.
- Cravino, Javier, and Sebastian Sotelo. 2019. "Trade-Induced Structural Change and the Skill Premium." *American Econ. J. Macroeconomics* 11 (3): 289–326.
- Davis, Donald R., and David Weinstein. 2003. "Market Access, Economic Geography and Comparative Advantage: An Empirical Test." *J. International Econ.* 59:1–23.

- Dekle, Robert, Jonathan Eaton, and Samuel Kortum. 2008. "Global Rebalancing with Gravity: Measuring the Burden of Adjustment." *IMF Staff Papers* 55 (3): 511–40.
- Dixit, Avinash K. 1985. "Tax Policy in Open Economies." In *Handbook of Public Economics*, vol. 1, edited by Martin Feldstein, 313–74. New York: Elsevier.
- Durantón, Gilles. 2011. "California Dreamin': The Feeble Case for Cluster Policies." *Rev. Econ. Analysis* 3:3–45.
- Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography and Trade." *Econometrica* 70 (5): 1741–79.
- Ethier, Wilfred J. 1982. "Decreasing Costs in International Trade and Frank Graham's Argument for Protection." *Econometrica* 50 (5): 1243–68.
- Fajgelbaum, Pablo D., and Cecile Gaubert. 2020. "Optimal Spatial Policies, Geography, and Sorting." *Q.J.E.* 135 (2): 959–1036.
- Giri, Rahul, Kei-Mu Yi, and Hakan Yilmazkuday. 2021. "Gains from Trade: Does Sectoral Heterogeneity Matter?" *J. Internat. Econ.* 129:103429.
- Graham, Frank D. 1923. "Some Aspects of Protection Further Considered." *Q.J.E.* 37 (2): 199–227.
- Greenwald, Bruce, and Joseph E. Stiglitz. 2006. "Helping Infant Economies Grow: Foundations of Trade Policies for Developing Countries." *A.E.R.* 96 (2): 141–46.
- Greenwald, Bruce C., and Joseph E. Stiglitz. 1986. "Externalities in Economies with Imperfect Information and Incomplete Markets." *Q.J.E.* 101 (2): 229–64.
- Harrison, Ann E., and Andrés Rodríguez-Clare. 2010. "Trade, Foreign Investment, and Industrial Policy for Developing Countries." In *Handbook of Development Economics*, vol. 5, edited by Dani Rodrik and Mark Rosenzweig, 4039–214. New York: Elsevier.
- Head, Keith, and Thierry Mayer. 2013. "Gravity Equations: Toolkit, Cookbook, Workhorse." In *Handbook of International Economics*, vol. 4, edited by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, 131–95. New York: Elsevier.
- Head, Keith, and John Ries. 2001. "Increasing Returns versus National Product Differentiation as an Explanation for the Pattern of U.S.-Canada Trade." *A.E.R.* 91 (4): 858–76.
- Herrendorf, Berthold, Richard Rogerson, and Akos Valentinyi. 2013. "Two Perspectives on Preferences and Structural Transformation." *A.E.R.* 103 (7): 2752–89.
- Hines, James R. 1999. "Three Sides of Harberger Triangles." *J. Econ. Perspectives* 13 (2): 167–88.
- Hopenhayn, Hugo A. 1992. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica* 60 (5): 1127–50.
- Itskhoki, Oleg, and Benjamin Moll. 2019. "Optimal Development Policies with Financial Frictions." *Econometrica* 87 (1): 139–73.
- Juhász, Réka. 2018. "Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade." *A.E.R.* 108 (11): 3339–76.
- Juhász, Réka, Nathan Lane, Emily Oehlsen, and Veronica C. Perez. 2023. "The Who, What, When, and How of Industrial Policy: A Text-Based Approach." Working paper.
- Juhász, Réka, Nathan Lane, and Dani Rodrik. 2024. "The New Economics of Industrial Policy." *Ann. Rev. Econ.* 16:213–42.
- Juhász, Réka, and Claudia Steinwender. 2024. "Industrial Policy and the Great Divergence." *Ann. Rev. Econ.* 16:27–54.
- Kaufmann, Daniel, Aart Kraay, and Massimo Mastruzzi. 2003. "Governance Matters III: Governance Indicators for 1996–2002." Working Paper no. 3106, World Bank, Washington, DC.

- Kline, Patrick, and Enrico Moretti. 2014. "People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs." *Ann. Rev. Econ.* 6:629–62.
- Krugman, Paul. 1980. "Scale Economies, Product Differentiation, and the Pattern of Trade." *A.E.R.* 70 (5): 950–59.
- Kucheryavy, Konstantin, Gary Lyn, and Andrés Rodríguez-Clare. 2023. "Grounded by Gravity: A Well-Behaved Trade Model with Industry-Level Economies of Scale." *American Econ. J. Macroeconomics* 15 (2): 372–412.
- Lane, Nathan. 2017. "Manufacturing Revolutions: Industrial Policy and Networks in South Korea." Working paper.
- . 2020. "The New Empirics of Industrial Policy." *J. Industry, Competition, and Trade* 20:209–34.
- Lashkaripour, Ahmad, and Volodymyr Lugovskyy. 2023. "Profits, Scale Economies, and the Gains from Trade and Industrial Policy." *A.E.R.* 113 (10): 2759–808.
- Levchenko, Andrei A. 2007. "Institutional Quality and International Trade." *Rev. Econ. Studies* 74 (3): 791–819.
- Levchenko, Andrei A., and Jing Zhang. 2016. "The Evolution of Comparative Advantage: Measurement and Welfare Implications." *J. Monetary Econ.* 78:96–111.
- Liu, Ernest. 2019. "Industrial Policies in Production Networks." *Q.J.E.* 134 (4): 1883–948.
- Manova, Kalina. 2008. "Credit Constraints, Equity Market Liberalizations and International Trade." *J. Internat. Econ.* 76 (1): 33–47.
- Marshall, Alfred. 1920. *Principles of Economics: An Introductory Volume*. 8th ed. London: Macmillan.
- Melitz, Marc J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6): 1695–725.
- Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny. 1989. "Industrialization and the Big Push." *J.P.E.* 97 (5): 1003–26.
- Newey, W. K., and J. L. Powell. 2003. "Instrumental Variable Estimation in Nonparametric Models." *Econometrica* 71 (5): 1565–78.
- Nunn, Nathan. 2007. "Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade." *Q.J.E.* 122 (2): 569–600.
- Nunn, Nathan, and Daniel Trefler. 2014. "Domestic Institutions as a Source of Comparative Advantage." In *Handbook of International Economics*, vol. 4, edited by G. Gopinath, E. Helpman, and K. Rogoff, 263–315. New York: Elsevier.
- Oberfield, Ezra, and Devesh Raval. 2021. "Micro Data and Macro Technology." *Econometrica* 89 (2): 703–32.
- Pigou, A. C. 1932. *The Economics of Welfare*. 4th ed. London: Macmillan.
- Redding, Stephen J., and David E. Weinstein. 2024. "Accounting for Trade Patterns." *J. Internat. Econ.* 150:103910.
- Rodrik, Dani. 1995. "Getting Interventions Right: How South Korea and Taiwan Grew Rich." *Econ. Policy* 10 (20): 53–107.
- . 1996. "Coordination Failures and Government Policy: A Model with Applications to East Asia and Eastern Europe." *J. Internat. Econ.* 40 (1–2): 1–22.
- Rosenthal, Stuart S., and William C. Strange. 2004. "Evidence on the Nature and Sources of Agglomeration Economies." In *Handbook of Regional and Urban Economics*, vol. 4, edited by J. V. Henderson and J. F. Thisse, 2119–71. New York: Elsevier.
- Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Felipe Schwartzman. 2019. "Cognitive Hubs and Spatial Redistribution." Working Paper no. 26267, NBER, Cambridge, MA.

- Sanderson, Eleanor, and Frank Windmeijer. 2016. "A Weak Instrument F -Test in Linear IV Models with Multiple Endogenous Variables." *J. Econometrics* 190 (2): 212–21.
- Shapiro, Joseph. 2016. "Trade Costs, CO₂, and the Environment." *American Econ. J. Econ. Policy* 8 (4): 220–54.
- Simonovska, Ina, and Michael Waugh. 2012. "Different Trade Models, Different Trade Elasticities?" Working paper.