

Profits, Scale Economies, and the Gains from Trade and Industrial Policy[†]

By AHMAD LASHKARIPOUR AND VOLODYMYR LUGOVSKYY*

This paper examines the efficacy of second-best trade restrictions at correcting sectoral misallocation due to scale economies or profit-generating markups. To this end, we characterize optimal trade and industrial policies in an important class of quantitative trade models with scale effects and profits, estimating the structural parameters that govern policy outcomes. Our estimates reveal that standalone trade policy measures are remarkably ineffective at correcting misallocation, even when designed optimally. Unilateral adoption of corrective industrial policies is also ineffective due to immiserizing growth effects. But industrial policies coordinated internationally via a deep agreement are more transformative than any unilateral policy alternative. (JEL F12, F13, L52, O19, O25)

The United States will likely adopt an explicit industrial policy in the coming decade. Similar developments are well underway in other countries (Aiginger and Rodrik 2020). And with industrial policy back on the scene, we are witnessing a revival of old but questionable trade policy practices. Governments are often turning to protectionist trade policy measures to pursue their industrial policy objectives—as manifested by the United States National Trade Council’s mission or the Chinese *Made in China 2025* initiative.¹

These developments have sparked new interest in old but open questions regarding trade and industrial policy. For instance: (i) is trade policy an effective tool for correcting misallocation in the domestic economy? (ii) if not, should governments

*Lashkaripour: Indiana University (email: ahmadlp@gmail.com); Lugovskyy: Indiana University (email: vlugovsk@indiana.edu). Emi Nakamura was the coeditor for this article. We are grateful to James Anderson, Adina Ardelean, Dominick Bartelme, Kerem Cosar, Farid Farrokhi, Harald Fadinger, Fabio Ghironi, David Hummels, Kala Krishna, Konstantin Kucheryavy, Danial Lashkari, Nuno Limao, Gary Lyn, Ralph Ossa, James Rauch, Andrés Rodríguez-Clare, Kadee Russ, Peter Schott, Alexandre Skiba, Anson Soderbery, Robert Staiger, Jonathan Vogel, and conference participants at the Midwest Trade Meetings, Chicago Fed, UECE Lisbon Meetings, 2017 NBER ITI Summer Institute, 2019 WCTW, Indiana University, Purdue University, SIU, IBA, Boston College, University of Mannheim, and University of Michigan for helpful comments and suggestions. We thank Nicolas de Roux and Santiago Tabares for providing us with data on the Colombian HS10 product code changes over time. We are grateful to Fabio Gomez for research assistance. Lugovskyy thanks Indiana University SSRC for financial support. All errors are our own.

[†]Go to <https://doi.org/10.1257/aer.20210419> to visit the article page for additional materials and author disclosure statements.

¹See Bhagwati (1988) and Irwin (2017) for a historical account of trade restrictions being used by governments to promote their preferred industries. A prominent example dates back to 1791, when Alexander Hamilton approached Congress with “the Report on the Subject of Manufactures,” encouraging the implementation of protective tariffs and industrial subsidies. These policies were intended to help the US economy catch up with Britain.

undertake *unilateral* domestic policy interventions to correct misallocation? or (iii) Should they coordinate their industrial policies via a deep trade agreement?

To answer these questions, we characterize optimal trade and industrial policies in an important class of multi-industry, multicountry quantitative trade models where misallocation occurs due to scale economies or profit-generating markups. Guided by theory, we estimate the key parameters that govern the welfare consequences of trade and industrial policy in open economies. We then combine our estimated parameters with optimal policy formulas to quantify the *ex ante* gains from trade and industrial policy among 43 major countries.

Our estimation reveals that trade policy is remarkably ineffective at correcting misallocation, reflecting a deep tension between allocative efficiency and terms of trade. Unilateral adoption of corrective industrial policies can also backfire, as it often triggers *immiserizing growth*. These considerations, we argue, may have spurred a global *race to the bottom*, wherein governments either avoid corrective industrial policies or pair them with hidden trade barriers. A deep agreement can remedy this problem and deliver welfare gains that are more transformative than any unilateral policy intervention.

Section I presents our theoretical framework. Our baseline model is a generalized multi-industry Krugman (1980) model that features a nonparametric utility aggregator across industries and a nested CES utility aggregator within industries. This specification has an appealing property wherein the degrees of firm-level and country-level love for variety can diverge. We analyze both the restricted and free entry cases of the model to distinguish between the *short-run* and *long-run* consequences of policy. With a reinterpretation of parameters, our baseline framework also nests the multi-industry Melitz (2003)-Pareto model and the multi-industry Eaton and Kortum (2002) model with industry-level Marshallian externalities. We later extend our baseline model to accommodate nonparametric input-output linkages.

Section II derives sufficient statistics formulas for *first-best* and *second-best* trade and industrial policies. Unilaterally optimal policies in our framework pursue two objectives: first, they seek to improve the home country's *terms of trade* (*ToT*) vis-à-vis the rest of the world. Second, they seek to restore *allocative efficiency* in the domestic economy by reallocating resources toward high returns-to-scale or high-profit industries.

The *first-best* optimal policy consists of misallocation-blind import tariffs and export subsidies that purely maximize *ToT* gains. Allocative efficiency under the first-best is restored via domestic Pigouvian subsidies.² While these insights resonate with the targeting principle, our optimal policy formulas have other implications worth highlighting. First, even though first-best trade policies are blind to misallocation (the dispersion in scale elasticities), they depend on the overall strength of scale economies (the level of scale elasticities). Second, the optimal tariff formula is input-output-blind provided that export subsidies are assigned optimally.³

²The optimal subsidy rate in each industry equals the inverse of the industry-level scale elasticity or markup.

³Specifically, optimal import tariffs are equal to the inverse export supply elasticity regardless of input-output relationships. Consequently, in the absence of profits and scale economies, optimal tariffs become uniform when export subsidies are set optimally—revealing that the uniformity result in Costinot et al. (2015) extends to environments with input-output linkages.

Our *second-best* trade policy formulas apply to scenarios where governments are reluctant to use industrial subsidies for correcting misallocation—the root of which could be political pressures or institutional barriers.⁴ Second-best trade tax-cum-subsidies are composed of two elements: a neoclassical ToT-improving component and a misallocation-correcting component. The former aims to restrict relative exports in nationally differentiated industries. The latter seeks to restrict imports and promote exports in high returns-to-scale industries, mimicking the first-best Pigouvian subsidies.

Section III, guided by our optimal policy framework, puts forth two conjectures concerning the efficacy of trade and industrial policy: first, if industry-level trade and scale elasticities exhibit a strong negative correlation, stand-alone trade policy measures deliver limited welfare gains even when set optimally—reflecting their inability to strike a balance between ToT-improving and misallocation-correcting objectives.⁵ Second, in these conditions, unilateral scale (or markup) correction via industrial policy can cause *immiserizing growth* due to adverse ToT effects. Consequently, shallow trade agreements may be insufficient for reaching global efficiency. Once governments agree to abandon inefficient trade restrictions under a shallow agreement, they become tangled in a coordination game involving corrective industrial policies. The outcome of this game is a *race to the bottom* wherein governments are reluctant to implement corrective industrial policies without violating their commitments to free trade. A deep agreement can remedy this problem.

To test these conjectures, Section V estimates the structural parameters necessary for ex ante policy evaluation. Our optimal policy formulas reveal that the sufficient statistics for measuring policy outcomes include observables and two sets of parameters: (i) industry-level scale elasticities that govern the extent of misallocation and (ii) industry-level trade elasticities that control the scope for ToT manipulation. In our generalized Krugman (1980) model, scale elasticities reflect the degree of *love for variety* whose social benefits are not internalized by firms' entry decisions; and trade elasticities represent the degree of national product differentiation. We recover both elasticities from firm-level demand parameters, which are estimated by fitting a structural demand function to the universe of Colombian import transactions covering over 225,000 firms from 251 countries. Our identification strategy leverages high-frequency data on import transactions and exchange rates to construct a shift-share instrument that measures exposure to aggregate exchange rate fluctuations at the firm-product-year level.⁶ This estimation strategy is well suited for our end goal of policy evaluation, as it separately identifies the scale elasticities from trade elasticities, accurately pinpointing their covariance.

⁴Trade policy has been regularly used—in place of domestic industrial policy—to promote critical industries (Bhagwati 1988; Harrison and Rodríguez-Clare 2010; Irwin 2017). Relatedly, see Lane (2020) for a historical account of various industrial policy practices around the world.

⁵In some canonical cases, optimal second-best trade policies are even industry-blind—unable to beneficially correct interindustry misallocation or manipulate the ToT on an industry-by-industry basis.

⁶We develop this identification strategy to overcome challenges that are difficult to resolve with standard demand estimation techniques. Firm-level demand estimation techniques from the Industrial Organization literature leverage information on observed product characteristics, which are unavailable at the scale at which we conduct our estimation. Also, unlike traditional *country-level* import demand estimations, we cannot rely on tariff data for identification, as tariff rates do not vary sufficiently among firms from the same country.

Section VI combines our estimated scale and trade elasticity parameters, our optimal policy formulas, and macro-level data from the 2014 World Input-Output Database to quantify the (maximal) *ex ante* gains from policy among 43 major economies. Our analysis delivers three main findings.

First, we find that trade policy is remarkably ineffective at correcting misallocation in the domestic economy—even without factoring in the cost of retaliation by trading partners. Under free entry, second-best export subsidies and import taxes can raise the average country's real GDP by only 1.19 percent, which amounts to less than 4/10 of the gains attainable under the unilaterally first-best policy. Third-best import taxes are even less effective as a stand-alone policy, raising real GDP by a mere 0.63 percent. These findings corroborate the argument that trade policy has difficulty striking a balance between ToT and misallocation-correcting objectives.

Second, the unilateral adoption of corrective industrial policies triggers severe immiserizing growth in most countries. The average country's real GDP declines by 2.7 percent if they implement scale-correcting subsidies without reciprocity by trading partners. Aversion to these consequences, we argue, may have spurred a global race to the bottom in industrial policy implementation. To escape immiserizing growth, governments either avoid corrective policies or pair them with hidden trade barriers that breach shallow trade cooperation.⁷

Third, deep agreements can remedy the race to the bottom and deliver welfare gains that are more transformative than any unilateral intervention. To offer some perspective, corrective industrial policies coordinated via a deep agreement can elevate the average country's real GDP by 3.2 percent. These welfare gains rival the already-realized gains from shallow agreements for most countries. They, moreover, exceed any welfare gains achievable through unilateral trade or industrial policy interventions—even not considering that unilateralism often backfires in the form of retaliation by trading partners.

Related Literature.—Our theory relates to an emerging literature on optimal policy in *distorted* open economies. In a concurrent paper, Bartelme et al. (2019) characterize the first-best optimal policy for a small open economy in a multisector Ricardian model with Marshallian externalities. Relatedly, Haaland and Venables (2016) characterize optimal policy for a small open economy in two-by-two Krugman and Melitz models.⁸ Beyond optimal policy, Campolmi, Fadinger, and Forlati (2018) employ a two-sector Melitz-Pareto model to elucidate the trade-offs facing countries that join shallow and deep trade agreements.⁹ Our analysis of *second-best* trade policies speaks to an older literature emphasizing the firm-delocation rationale

⁷The Chinese government, for instance, pairs its domestic subsidies with hidden export taxes. These hidden barriers are applied via partial value-added tax rebates and are designed to restore China's ToT (Garred 2018).

⁸Demidova and Rodríguez-Clare (2009) and Felbermayr, Jung, and Larch (2013) characterize optimal tariffs in a *single-industry* Melitz-Pareto model. The single-industry assumption ensures that markets are efficient (Dhingra and Morrow 2019) and import and export taxes are equivalent (the Lerner symmetry). So, the unilaterally first-best can be reached with import tariffs alone. Costinot, Rodríguez-Clare, and Werning (2016) examine optimal policy in the *single-industry* Melitz-Pareto model from a different lens, characterizing optimal *firm-level* taxes.

⁹Other papers have also used new or quantitative trade models to analyze piecemeal policy reforms in distorted economies or optimal policy in *nondistorted* economies—e.g., Costinot and Rodríguez-Clare (2014); Campolmi, Fadinger, and Forlati (2014); Costinot et al. (2015); Bagwell and Lee (2018); Caliendo et al. (2015); Demidova (2017); Beshkar and Lashkaripour (2019, 2020).

for trade restrictions (e.g., Venables 1987; Ossa 2011) and supplements Bagwell and Staiger's (2001, 2004) result about the role of trade agreements in distorted economies. We also build on Kucheryavyy, Lyn, and Rodríguez-Clare (2023a) to establish isomorphism between our baseline model and other workhorse models in the literature. Our *quantitative* examination of trade and industrial policy connects to two strands of literature. First, a mature line of research measuring the ex post consequences of tariff cuts (Costinot and Rodríguez-Clare 2014; Caliendo and Parro 2015; Ossa 2014, 2016; Spearot 2016). Second, a growing literature examining the ex ante consequences of optimal policy. Ossa (2014), most notably, quantifies the consequences of cooperative and noncooperative import tariffs in a multi-industry Krugman model with restricted entry. Lastly, our work relates to a vibrant literature examining the impacts of exogenous trade shocks in distorted economies (e.g., De Blas and Russ 2015; Edmond, Midrigan, and Xu 2015; Baqaee and Farhi 2019).

We contribute to the calculus of optimal policy in open economies by developing a new dual technique for optimal policy derivation in general equilibrium quantitative trade models with many countries, increasing returns-to-scale production technologies, and input-output linkages. Our approach has applications beyond those considered in this paper. Lashkaripour (2021), for instance, adopts a special case of this technique to characterize Nash tariffs in a monopolistic competition model with restricted entry.¹⁰ Farrokhi and Lashkaripour (2021) extend this technique to analyze optimal carbon pricing under international climate externalities.

Lastly, our paper supplements the broader quantitative trade literature by developing an estimation technique that separately identifies the scale elasticity from the trade elasticity in certain environments. Our indirect approach to scale elasticity estimation complements the direct method concurrently proposed by Bartelme et al. (2019). The main limitation of our indirect approach is that it cannot detect scale externalities unrelated to love for variety, as it does *not* directly leverage scale-related moments. Despite this limitation, our approach has two useful properties for policy evaluation. First, it separately identifies the trade elasticity from the scale elasticity, provided that scale economies arise from love for variety, which is valuable since the covariance between these elasticities regulates policy outcomes. Second, our approach is robust to the presence of quasi-fixed production inputs, enabling us to isolate scale effects that impair allocative efficiency from fixed-input-driven diseconomies of scale.

I. Theoretical Framework

Our baseline model is a generalized multi-industry, multicountry Krugman model with semiparametric preferences. In Section IV we show that our theory readily applies to alternative models featuring firm selection à la Melitz–Chaney and external economies of scale à la Kucheryavyy, Lyn, and Rodríguez-Clare (2023a). We also extend our theory later to accommodate arbitrary input-output networks and political economy pressures.

¹⁰Lashkaripour (2021) examines the cost of noncooperative import restrictions when governments simultaneously apply their third-best optimal import tariffs. To expedite the computational process, Lashkaripour (2021) uses analytic formulas for Nash tariffs, which correspond to a special of our Theorem 3.

We consider a world economy consisting of multiple countries and industries. Countries are indexed by $i, j, n \in \mathbb{C}$. Industries are indexed by $g, k \in \mathbb{K}$. Industries can differ in fundamentals such as the degree of scale economies or trade elasticity. Each country $i \in \mathbb{C}$ is populated by L_i individuals who supply one unit of labor inelasticity. Labor is the sole primary factor of production in each economy. Workers cannot relocate between countries but are perfectly mobile across industries within a country and are paid a countrywide wage, w_i .

A. Preferences

Each good in our model is indexed by a triplet, which signifies its location of production (origin), its location of final consumption (destination), and the industry under which the good is classified. To give an example, good “ ji, k ” denotes a good corresponding to *origin country* j –*destination country* i –*industry* k .

Cross-Industry Demand.—The representative consumer in country $i \in \mathbb{C}$ faces a vector of industry-level consumer price indexes $\tilde{\mathbf{P}}_i = \{\tilde{P}_{i,k}\}$, where index $\tilde{P}_{i,k} \equiv \tilde{P}_{i,k}(\tilde{P}_{1i,k}, \dots, \tilde{P}_{Ni,k})$ aggregates over industry k goods sourced from various origins. The consumer chooses their demand for industry-level bundles $\mathbf{Q}_i \equiv \{Q_{i,k}\}$ to maximize a nonparametric utility function subject to a budget constraint. This choice yields an indirect utility, which is a function of the consumer’s income, Y_i , and the vector of industry-level “*consumer*” price indexes in market i , $\tilde{\mathbf{P}}_i$:

$$(1) \quad V_i(Y_i, \tilde{\mathbf{P}}_i) = \max_{\mathbf{Q}_i} U_i(\mathbf{Q}_i)$$

subject to

$$\sum_{k \in \mathbb{K}} \tilde{P}_{i,k} Q_{i,k} = Y_i.$$

Throughout this paper, the tilde notation on price is used to distinguish between “consumer” and “producer” prices. The former includes taxes, whereas the latter does not. Problem (1) yields an *industry-level* Marshallian demand function, which we denote by $Q_{i,k} = \mathcal{D}_{i,k}(Y_i, \tilde{\mathbf{P}}_i)$. This function tracks how (given prices and total income) consumers allocate their expenditure across industries. A special case of our general cross-industry demand function is the Cobb-Douglas case, wherein $U_i(\mathbf{Q}_i) = \prod_{k \in \mathbb{K}} Q_{i,k}^{e_{i,k}}$, implying that $Q_{i,k} = e_{i,k} Y_i / \tilde{P}_{i,k}$.

Within-Industry Demand.—Each industry-level bundle aggregates over various origin-specific composite varieties: $Q_{i,k} \equiv Q_{i,k}(Q_{1i,k}, \dots, Q_{Ni,k})$. Each origin-specific composite variety, itself, aggregates over multiple firm-level varieties: $Q_{ji,k} \equiv Q_{ji,k}(\mathbf{q}_{ji,k})$, where $\mathbf{q}_{ji,k} = \{q_{ji,k}(\omega)\}_{\omega \in \Omega_{j,k}}$ is a vector with each element $q_{ji,k}(\omega)$ denoting the quantity consumed of firm ω ’s output.¹¹ We assume that the

¹¹ $\Omega_{j,k}$ denotes the set of all firms operating in origin j –industry k . In our baseline model, firms do *not* incur a fixed exporting cost, so each firm in $\Omega_{j,k}$ serves market i . We relax this assumption in Section IV.

within-industry utility aggregator has a nested-CES structure, which enables us to abstract from variable markups and direct our attention to the scale-driven and profit-shifting effects of policy.

ASSUMPTION (A1): *The within-industry utility aggregator is nested-CES. In particular,*

$$Q_{i,k} = \left(\sum_{j \in \mathbb{C}} Q_{ji,k}^{\frac{\sigma_k-1}{\sigma_k}} \right)^{\frac{\sigma_k}{\sigma_k-1}}, \text{ where } Q_{ji,k} = \left[\int_{\omega \in \Omega_{j,k}} \varphi_{ji,k}(\omega)^{\frac{1}{\gamma_k}} q_{ji,k}(\omega)^{\frac{\gamma_k-1}{\gamma_k}} d\omega \right]^{\frac{\gamma_k}{\gamma_k-1}},$$

with $\gamma_k \geq \sigma_k > 1$ and $\varphi_{ji,k}(\omega) > 0$ corresponding to a constant variety-specific taste shifter.

Based on (A1), the demand for the composite *national-level* variety ji,k (origin country j –destination country i –industry k) is given by

$$(2) \quad Q_{ji,k} = \left(\tilde{P}_{ji,k} / \tilde{P}_{i,k} \right)^{-\sigma_k} Q_{i,k},$$

where $\tilde{P}_{ji,k}$ and $\tilde{P}_{i,k}$ respectively denote the origin-specific and industry-level CES price indexes.¹² Recall that $Q_{i,k}$ denotes industry-level demand, which is given by $Q_{i,k} = \mathcal{D}_{i,k}(Y_i, \tilde{\mathbf{P}}_i)$. The demand facing individual firms from country j is, accordingly, given by

$$(3) \quad q_{ji,k}(\omega) = \varphi_{ji,k}(\omega) \left[\frac{\tilde{P}_{ji,k}(\omega)}{\tilde{P}_{ji,k}} \right]^{-\gamma_k} \left(\frac{\tilde{P}_{ji,k}}{\tilde{P}_{i,k}} \right)^{-\sigma_k} \mathcal{D}_{i,k}(Y_i, \tilde{\mathbf{P}}_i).$$

Importantly, the above parameterization of demand allows for the *firm-level* and *national-level* degrees of market power to diverge. γ_k governs the degree of firm-level market power and love for variety, while σ_k governs the degree of national-level market power in industry k .

Elasticity of Demand Facing National-Level Varieties.—Following equation (2), the demand for aggregate variety ji,k is a function of total income in market i , Y_i , and the entire vector of origin \times industry-specific consumer price indexes in that market: namely, $Q_{ji,k} = \mathcal{D}_{ji,k}(Y_i, \tilde{\mathbf{P}}_i)$. To keep track of changes in demand, we define the elasticity of demand for national-level variety ji,k with respect to the price of variety ni,g as follows:

$$\varepsilon_{ji,k}^{(ni,g)} \equiv \frac{\partial \ln \mathcal{D}_{ji,k}(Y_i, \tilde{\mathbf{P}}_i)}{\partial \ln \tilde{P}_{ni,g}} \sim \text{price elasticity of demand.}$$

Under Cobb-Douglas preferences (i.e., zero cross-substitutability between industries), the national-level demand elasticities are fully determined by the upper-tier

¹²Namely, $\tilde{P}_{ji,k} = \left[\sum_{\omega \in \Omega_{j,k}} \varphi_{ji,k}(\omega) \tilde{p}_{ji,k}(\omega)^{1-\gamma_k} \right]^{1/(1-\gamma_k)}$ and $\tilde{P}_{i,k} = \left(\sum_{j \in \mathbb{C}} \tilde{P}_{ji,k}^{1-\sigma_k} \right)^{1/(1-\sigma_k)}$.

CES parameter σ_k and national-level expenditure shares. Specifically, $\varepsilon_{ji,k}^{j,i,g} = 0$ if $g \neq k$, while

$$\varepsilon_{ji,k} \sim \varepsilon_{ji,k}^{(j,i,k)} = -1 - (\sigma_k - 1)(1 - \lambda_{ji,k}); \quad \varepsilon_{ji,k}^{(j,i,k)} = (\sigma_k - 1)\lambda_{ji,k} \quad (J \neq j),$$

where $\lambda_{ji,k} \equiv \tilde{P}_{ji,k} Q_{ji,k} / \sum_J \tilde{P}_{ji,k} Q_{ji,k}$ denotes the (within-industry) share of expenditure on ji, k . In the presence of cross-substitutability between industries, the demand elasticity will feature an additional term that accounts for cross-industry demand effects.

In our setup, optimal policy internalizes the entire matrix of own- and cross-demand elasticities. To present our optimal policy formulas concisely, we use the following matrix notation to track the elasticity of demand with respect to goods sourced from various origins and industries.

DEFINITION (D1): Let $K = |\mathbb{K}|$ denote the number of industries. The $K \times K$ matrix $\mathbf{E}_{ji}^{(ni)}$ describes the elasticity of demand for origin $j \in \mathbb{C}$ goods with respect to the price of origin $n \in \mathbb{C}$ goods in market i :

$$\mathbf{E}_{ji}^{(ni)} \equiv \begin{bmatrix} \varepsilon_{ji,1}^{(ni,1)} & \cdots & \varepsilon_{ji,1}^{(ni,K)} \\ \vdots & \ddots & \vdots \\ \varepsilon_{ji,K}^{(ni,1)} & \cdots & \varepsilon_{ji,K}^{(ni,K)} \end{bmatrix}.$$

To simplify the notation, we use $\mathbf{E}_{ji} \sim \mathbf{E}_{ji}^{(ji)}$ to denote the elasticity of origin j goods with respect to origin j prices and use the $K \times (N - 1)K$ matrix, $\mathbf{E}_{ji}^{(-ii)} = [\mathbf{E}_{ji}^{(ni)}]_{n \neq i}$, to summarize the elasticity of demand for origin j goods with respect to price of all import varieties in market i (i.e., all varieties source from any origin $n \neq i$). Important for our analysis, \mathbf{E}_{ji} is an invertible matrix—the proof of which is provided in online Appendix E using the primitive properties of Marshallian demand.

B. Production and Firms

Each economy $i \in \mathbb{C}$ is populated with a mass $M_{i,k} = |\Omega_{i,k}|$ of single-product firms in industry $k \in \mathbb{K}$ that compete under monopolistic competition. Labor is the only factor of production. Firm entry into industry k is either free or restricted. Under restricted entry, $M_{i,k} = \bar{M}_{i,k}$ is invariant to policy. Under free entry, a pool of ex ante identical firms can pay an entry cost $w_i f_k^e$ to serve industry k from origin i . After paying the entry cost, each firm $\omega \in \Omega_{i,k}$ draws a productivity $z(\omega) \geq 1$ from distribution $G_{i,k}(z)$ and faces a marginal cost $\tau_{ij,k} w_i / z(\omega)$ for producing and delivering goods to destination $j \in \mathbb{C}$, where $\tau_{ij,k}$ denotes a flat iceberg transport cost. Collecting these assumptions, the “producer” price index of composite good ij, k (which aggregates over firm-level varieties associated with origin i –destination j –industry k) is

$$(4) \quad P_{ij,k} = \frac{\gamma_k}{\gamma_k - 1} \tau_{ij,k} \bar{a}_{i,k} w_i M_{i,k}^{-\frac{1}{\gamma_k - 1}},$$

where $\bar{a}_{i,k} \equiv \left[\int_1^\infty z^{\gamma_k-1} dG_{i,k}(z) \right]^{1/(1-\gamma_k)}$ denotes the average unit labor cost in origin i .¹³ Following Kucheryavy, Lyn, and Rodríguez-Clare (2023a), we refer to $1/(\gamma_k - 1) = -(\partial \ln P_{ij,k} / \partial \ln M_{i,k})$ as the industry-level scale elasticity:

$$\mu_k \equiv \frac{1}{\gamma_k - 1} \sim \text{scale elasticity} \sim \text{markup.}$$

Considering equation (4), μ_k represents both the constant firm-level markup in industry k (i.e., $1 + \mu_k = \gamma_k / (\gamma_k - 1)$), and the elasticity by which (variety-adjusted) TFP increases with industry-level employment $L_{i,k}$ (noting that $L_{i,k} \propto M_{i,k}$).¹⁴ The equivalence between markup and scale elasticity is not a universal property but a specific feature of our baseline Krugman model. We take advantage of this equivalence to simplify notation, but it is not essential for the theoretical results that follow. As shown in Section IV, our analytical formulas for optimal policy extend to alternative models where the scale elasticity and markup levels diverge.

Expressing Producer Prices in Terms of Profit-Adjusted Wages.—Our optimal policy framework reveals a tight connection between the restricted and free entry cases—even though misallocation stems from markup distortions in the former and scale distortions in the latter. To illustrate this connection and integrate optimal policy results for the two cases, we specify producer prices as a function of profit-adjusted wage rates. The idea is that net profits (if any) are rebated back to workers. The profit-adjusted wage rate in country i is defined as

$$\hat{w}_i \equiv (1 + \bar{\mu}_i)w_i \sim \text{profit-adjusted wage,}$$

where $\bar{\mu}_i$ denotes economy i 's average profit margin across all industries. Namely,

$$(5) \quad \bar{\mu}_i = \begin{cases} 0, & \text{if entry is free;} \\ \frac{\sum_{k \in \mathbb{K}} \sum_{j \in \mathbb{C}} \frac{\mu_k}{1 + \mu_k} P_{ij,k} Q_{ij,k}}{\sum_{k \in \mathbb{K}} \sum_{j \in \mathbb{C}} \frac{1}{1 + \mu_k} P_{ij,k} Q_{ij,k}}, & \text{if entry is restricted.} \end{cases}$$

Under free entry, profits are drawn to zero, resulting in $\bar{\mu}_i = 0$. Under restricted entry, the average profit margin is positive and depends on the industrial composition of country i 's output—with a higher $\bar{\mu}_i$ reflecting more sales in high-markup (high- μ) industries. Appealing to our definitions for \hat{w}_i and μ_k , we can reformulate equation (4) to express producer prices as a function of *profit-adjusted* wages:

$$(6) \quad P_{ij,k} = \begin{cases} Q_{ij,k} \left[\sum_{j \in \mathbb{C}} \tau_{ij,k} Q_{ij,k} \right]^{-\frac{\mu_k}{1 + \mu_k}} \hat{w}_i, & \text{if entry is free;} \\ Q'_{ij,k} \frac{1 + \mu_k}{1 + \bar{\mu}_i} \hat{w}_i, & \text{if entry is restricted.} \end{cases}$$

¹³ Notice that $\bar{a}_{i,k}$ is constant in our baseline model. This is no longer true in the Melitz (2003) extension of our model explored in Section IV, in which firms incur a fixed cost to serve individual markets.

¹⁴ With free entry and constant markups, it follows immediately that $L_{i,k} = \bar{c}_{i,k} M_{i,k}$, where $\bar{c}_{i,k}$ is a constant.

In the above formulation, $Q_{ij,k} \equiv (1 + \mu_k)\tau_{ij,k}\bar{a}_{i,k}^{1/(1+\mu_k)}(\mu_k/f_k^e)^{-\mu_k/(1+\mu_k)}$ and $Q'_{ij,k} \equiv \tau_{ij,k}\bar{a}_{i,k}\bar{M}_{i,k}^{-\mu_k}$ are constant price shifters, and $\sum_{j \in \mathbb{C}}(\tau_{ij,k}Q_{ij,k})$ denotes *origin i -industry k 's gross output*.¹⁵ As we explain shortly, the above formulation of producer prices is useful for tracking the terms of trade gains from policy in open economies. These gains require a contraction of producer prices in the rest of the world, which can occur through alterations to production scale, $\sum_{j \in \mathbb{C}}(\tau_{ij,k}Q_{ij,k})$, under free entry or average profit margins, $\bar{\mu}_i$, under restricted entry.

C. The Instruments of Policy

The government in country i is afforded a complete set of revenue-raising trade and domestic policy instruments, namely,

- (i) **import tax**, $t_{ji,k}$, applied to all goods imported from origin $j \neq i$ in industry k ;
- (ii) **export subsidy**, $x_{ij,k}$, applied to all goods sold to market $j \neq i$ in industry k ;
- (iii) **industrial subsidy**, $s_{i,k}$, applied to industry k 's output irrespective of where it is sold.

Our specification of policy is quite flexible, as it accommodates import subsidies or export taxes ($-1 \leq t < 0$ or $-1 \leq x < 0$) as well as production taxes ($-1 \leq s < 0$). We disregard *consumption taxes*, as they are redundant given the availability of the other tax instruments (see online Appendix A). There is a simple intuition behind this redundancy: country $i \in \mathbb{C}$ has access to $2(N - 1) + 2$ different tax instruments in each industry (where $N \equiv |\mathbb{C}|$ denotes the number of countries). These $2(N - 1) + 2$ tax instruments can directly manipulate $2(N - 1) + 1$ consumer price indexes: $N - 1$ export prices, $N - 1$ import prices, and one price associated with the domestically produced and consumed variety (namely, $\tilde{P}_{ii,k}$). So, by construction, one of the $2(N - 1) + 2$ tax instruments in each industry is redundant. Here, we treat the industry-level consumption tax as a redundant instrument.¹⁶

The above tax instruments create a wedge between consumer price indexes, $\{\tilde{P}_{ji,k}\}$, and producer price indexes, $\{P_{ji,k}\}$, as follows:

$$(7) \quad \tilde{P}_{ji,k} = \frac{1 + t_{ji,k}}{(1 + x_{ji,k})(1 + s_{j,k})} P_{ji,k}, \quad \forall j, i \in \mathbb{C}, k \in \mathbb{K}.$$

¹⁵Under free entry, the total cost of entry must equal gross profits across all markets. In particular,

$$w_i f_k^e M_{i,k} = \sum_{j \in \mathbb{C}} \left(\frac{\mu_k}{1 + \mu_k} P_{ij,k} Q_{ij,k} \right) \quad (\text{Free Entry Condition}).$$

Replacing $P_{ij,k}$ in the above equation with 4 yields $M_{i,k} = \left[\frac{\mu_k}{f_k^e} \sum_j (\bar{a}_{ij,k} Q_{ij,k}) \right]^{1/(1+\mu_k)}$. Equation (6), then, follows from plugging the expression for $M_{i,k}$ back into equation (4).

¹⁶With more than two countries ($N > 2$), country i has access to $2(N - 1) + 2$ instruments per industry. These instruments can manipulate $2(N - 1) + 1$ price variables, which implies the same redundancy.

These tax instruments also generate/exhaust revenue for the tax-imposing country. The combination of all taxes imposed by country $i \in \mathbb{C}$ produces a tax revenue equal to

$$(8) \mathcal{R}_i = \overbrace{\sum_{k \in \mathbb{K}} \left[\left(\frac{1}{1 + s_{i,k}} - 1 \right) P_{ii,k} Q_{ii,k} \right]}^{\text{industrial subsidies}} + \underbrace{\sum_{k \in \mathbb{K}} \sum_{j \neq i} \left\{ \frac{t_{ji,k}}{(1 + x_{ji,k})(1 + s_{j,k})} P_{ji,k} Q_{ji,k} + \left[\frac{1}{(1 + x_{ij,k})(1 + s_{i,k})} - 1 \right] P_{ij,k} Q_{ij,k} \right\}}_{\text{import taxes + export subsidies}}$$

Tax revenues are rebated to the consumers in a lump-sum fashion. After we account for tax revenues, total income in country i equals the sum of profit-adjusted wage payments, $\dot{w}_i L_i = (1 + \bar{\mu}_i) w_i L_i$, and tax revenues. Namely, $Y_i = \dot{w}_i L_i + \mathcal{R}_i$, where \mathcal{R}_i can be positive or negative depending on whether country i 's policy consists of net taxes or subsidies.

D. General Equilibrium

For convenience, we refer to *profit-adjusted* wages as just wages going forward, using $\mathbf{w} \equiv \{\dot{w}_i\}$ to denote the global vector of wages. We also assume throughout the paper that the underlying parameters of the model are such that the necessary and sufficient conditions for the uniqueness of equilibrium are satisfied.¹⁷ To present our theory, we express all equilibrium outcomes (except wages) as an explicit function of global taxes (\mathbf{x} , \mathbf{t} , and \mathbf{s}) and wages \mathbf{w} , with the understanding that \mathbf{w} is itself an equilibrium outcome. As detailed in online Appendix E, this formulation derives from solving a system that imposes all equilibrium conditions aside from the labor market clearing conditions. For future reference, we outline this formulation of equilibrium variables below and provide a summary of key variables in Table 1.

NOTATION: For a given vector of taxes and wages $\mathbf{T} = (\mathbf{t}, \mathbf{x}, \mathbf{s}; \mathbf{w})$, equilibrium outcomes $Y_i(\mathbf{T})$, $P_{ji,k}(\mathbf{T})$, $\tilde{P}_{ji,k}(\mathbf{T})$, $Q_{ji,k}(\mathbf{T})$ are determined such that (i) producer prices are characterized by equation (6); (ii) consumer prices are given by equation (7); (iii) industry-level consumption choices are a solution to Problem (1) with demand for national-level varieties, $Q_{ji,k}$, given by equation (2); and (iv) total income (which dictates total expenditure by country i) equals profit-adjusted wage payments plus tax revenues:

$$Y_i(\mathbf{T}) = \dot{w}_i L_i + \mathcal{R}_i(\mathbf{T}),$$

where tax revenues $\mathcal{R}_i(\mathbf{T})$ are described by equation (8).

¹⁷ Following Kucheryavy, Lyn, and Rodríguez-Clare (2023a), this assumption holds in the two-country case if $\gamma_k \geq \sigma_k$ and holds otherwise if trade costs are sufficiently small.

TABLE 1—SUMMARY OF KEY VARIABLES

Variable	Description
$\tilde{P}_{ji,k}$	Consumer price index (origin j –destination i –industry k)
$P_{ji,k}$	Producer price index (origin j –destination i –industry k)
Y_i	Total income in country i
\mathcal{R}_i	Total tax revenue in country i (equation (8))
w_i and \hat{w}_i	pure and profit-adjusted wage rates in country i : $\hat{w}_i = (1 + \bar{\mu}_i)w_i$
$x_{ji,k}$	Export subsidy applied to good ji,k (if $j \neq i$)
$t_{ji,k}$	Import tax applied on good ji,k (if $j \neq i$)
$s_{i,k}$	Industrial subsidy applied to all goods from origin i –industry k
$\lambda_{ji,k}$	Within-industry expenditure share (good ji,k): $\tilde{P}_{ji,k} Q_{ji,k} / \sum_j \tilde{P}_{ji,k} Q_{ji,k}$
$r_{ji,k}$	Within-industry sales share (good ji,k): $P_{ji,k} Q_{ji,k} / \sum_l P_{jl,k} Q_{jl,k}$
$e_{i,k}$	Industry-level expenditure share (destination i –industry k)
$\rho_{i,k}$	Industry-level sales share (origin i –industry k)
μ_k	industry-level markup \sim industry-level scale elasticity
$\bar{\mu}_i$	Average profit margin in origin i (equation (5))
σ_k	Cross-national CES parameter $\sim (1 + \text{trade elasticity})$
$\varepsilon_{ji,k}^{(ni,g)}$	Elasticity of demand for good ji,k w.r.t. the price of ni,g
$\omega_{ji,k}$	Inverse of good ji,k 's supply elasticity (online Appendix equation E.11)

Considering the above formulation of equilibrium variables, welfare, too, can be expressed as an explicit function of taxes and wages, $\mathbf{T} = (\mathbf{t}, \mathbf{x}, \mathbf{s}; \mathbf{w})$. Namely,

$$W_i(\mathbf{T}) \equiv V_i(Y_i(\mathbf{T}), \tilde{\mathbf{P}}_i(\mathbf{T})).$$

Since \mathbf{w} is itself an equilibrium outcome, vector $\mathbf{T} = (\mathbf{t}, \mathbf{x}, \mathbf{s}; \mathbf{w})$ is feasible if and only if \mathbf{w} is the equilibrium wage consistent with \mathbf{t} , \mathbf{x} , and \mathbf{s} . Accordingly, our objective in this paper is to study problems where the government in i chooses $\mathbf{T} = (\mathbf{t}, \mathbf{x}, \mathbf{s}; \mathbf{w})$ to maximize $W_i(\mathbf{T})$ subject to the noted feasibility constraint, which is formally defined below.

DEFINITION (D2): *The set of feasible policy–wage vectors, \mathbb{F} , consists of any vector $\mathbf{T} = (\mathbf{t}, \mathbf{x}, \mathbf{s}; \mathbf{w})$ where \mathbf{w} satisfies the labor market clearing condition in every country, given \mathbf{t} , \mathbf{x} , and \mathbf{s} :*

$$\mathbb{F} = \left\{ \mathbf{T} = (\mathbf{t}, \mathbf{x}, \mathbf{s}; \mathbf{w}) \mid \hat{w}_i L_i = \sum_{k \in \mathbb{K}} \sum_{j \in \mathbb{C}} [P_{ij,k}(\mathbf{T}) Q_{ij,k}(\mathbf{T})]; \forall i \in \mathbb{C} \right\}.$$

There is a basic reason for why we formulate equilibrium outcomes as a function of $\mathbf{T} = (\mathbf{t}, \mathbf{x}, \mathbf{s}; \mathbf{w})$ instead of just $(\mathbf{t}, \mathbf{x}, \mathbf{s})$. This choice of formulation allows us to articulate an important intermediate result regarding tax neutrality. This result, which is stated below, simplifies our theoretical derivation of optimal policy to a great degree.

LEMMA 1 (Tax Neutrality): For any a and $\tilde{a} \in \mathbb{R}_+$, (i) if $\mathbf{T} = (\mathbf{1} + \mathbf{t}_i, \mathbf{t}_{-i}, \mathbf{1} + \mathbf{x}_i, \mathbf{x}_{-i}, \mathbf{1} + \mathbf{s}_i, \mathbf{s}_{-i}; \dot{w}_i, \mathbf{w}_{-i}) \in \mathbb{F}$, then $\mathbf{T}' = (a(\mathbf{1} + \mathbf{t}_i), \mathbf{t}_{-i}, a(\mathbf{1} + \mathbf{x}_i), \mathbf{x}_{-i}, (1/\tilde{a})(\mathbf{1} + \mathbf{s}_i), \mathbf{s}_{-i}; (a/\tilde{a})\dot{w}_i, \mathbf{w}_{-i}) \in \mathbb{F}$. Moreover, (ii) welfare is preserved under \mathbf{T} and \mathbf{T}' : $W_n(\mathbf{T}) = W_n(\mathbf{T}')$ for all $n \in \mathbb{C}$.

The above lemma is proven in online Appendix B and connects two fundamental tax neutrality principles: the Lerner symmetry (Lerner 1936; Costinot and Werning 2019) and the welfare-neutrality of uniform subsidies or markups (Lerner 1934; Samuelson 1948). Importantly, Lemma 1 implies that there are multiple optimal tax combinations for each country i —a result that simplifies our forthcoming task of characterizing optimal policy. To give some detail, the contribution of general equilibrium wage and income effects to the optimal tax schedule is often summarized by aggregate tax shifters that are industry-blind. The neutrality established by Lemma 1 simplifies the task of handling these aggregate terms.

II. Sufficient Statistics Formulas for Optimal Policy

This section derives sufficient statistics formulas for optimal trade and industrial policies. These formulas are later employed to quantify the ex ante gains from policy among many countries. Before proceeding to the derivation, let us review the two rationales for policy intervention in our setup. A noncooperative, welfare-maximizing government seeks to (i) restrict trade and reap unexploited *terms of trade* (ToT) gains vis-à-vis the rest of the world and (ii) correct misallocation in the domestic economy. Misallocation, notice, stems from the cross-industry heterogeneity in markups or scale elasticities, leading to inefficiently low output in high-profit or high returns-to-scale (high- μ) industries. A crucial difference between these policy objectives is that ToT manipulation is inefficient from a global standpoint, as it distorts international prices to transfer surplus from the rest of the world to the tax-imposing country.

A. Efficient Policy from a Global Standpoint

As a useful benchmark, we first characterize the efficient policy from a global standpoint. Efficient policies, by definition, are the solution to a central planner's problem that maximizes global welfare via taxes and lump-sum international transfers. Let δ_i denote the Pareto weight assigned to country i in the planner's objective function. The globally efficient policy solves the following planning problem *subject to* the availability of lump-sum transfers:

$$\max_{(\mathbf{t}, \mathbf{x}, \mathbf{s}; \mathbf{w}) \in \mathbb{F}} \sum_{i \in \mathbb{C}} \delta_i \log W_i(\mathbf{t}, \mathbf{x}, \mathbf{s}; \mathbf{w}).$$

Keep in mind that the above problem affords the planner enough instruments to obtain their first-best. Good-specific taxes allow the planner to restore allocative efficiency, while lump-sum transfers allow her to redistribute internationally based on the Pareto weights, δ_i . This point is expanded on in online Appendix F, where it is

shown that the efficient tax policy involves zero trade taxes and Pigouvian subsidies that restore marginal-cost-pricing globally:¹⁸

$$(9) \quad t_{ji,k}^* = x_{ji,k}^* = 0, \quad \forall ji,k; \quad 1 + s_{i,k}^* = 1 + \mu_k, \quad \forall i,k.$$

The above characterization applies to both the free and restricted entry cases—with the understanding that μ_k assumes different interpretations in each case. Appealing to this result, online Appendix E establishes a basic point about international cooperation. Welfare-maximizing governments will settle on the efficient policy only if they are unable to influence consumer/producer prices in the rest of the world. Otherwise, they will defect to take advantage of terms of trade gains. This result indicates that the pursuit of ToT gains is the sole reason welfare-maximizing governments deviate from efficient policy choices—at least when they are afforded sufficient policy instruments.¹⁹ This result echoes the argument in Bagwell and Staiger (2001, 2004), generalizing it to settings with many countries and differentiated industries.

In the next section, we characterize the unilaterally optimal policy of noncooperative governments. This exercise elucidates two issues. First, it determines how governments deviate from the cooperative policy choice when ToT considerations are taken into account. Second, it clarifies how governments approach industrial policy when they view first-best Pigouvian subsidies as politically infeasible. Once we settle these two issues, we argue that the implementation of globally efficient policies requires both a “*shallow agreement*” to discipline trade policy choices and a “*deep agreement*” to coordinate industrial policy implementation.

B. Unilaterally Optimal Policy Choices

First-Best: Unilaterally Optimal Trade and Domestic Policies.—We now characterize a noncooperative country’s unilaterally optimal policy. We consider cases where a noncooperative country $i \in \mathbb{C}$ selects taxes, $\mathbf{t}_i \equiv \{t_{ji,k}\}$, $\mathbf{x}_i \equiv \{x_{ij,k}\}$, and $\mathbf{s}_i \equiv \{s_{i,k}\}$, taking policy choices elsewhere as given. Countries in the rest of the world are passive in their use of taxes but actively maintain internal cooperation.²⁰ We begin with the unilaterally *first-best* case where the government in i is afforded

¹⁸To be specific, the implementation of the efficient allocation involves the above taxes plus lump-sum international transfers based on Pareto weights. The logic is that the planner maximizes global output by restoring marginal-cost pricing and redistributes the corresponding income gains between countries via efficient transfers. Absent transfers, implementing $\{\mathbf{t}^*, \mathbf{x}^*, \mathbf{s}^*\}$ would deliver a Kaldor-Hicks improvement (Kaldor 1939; Hicks 1939) but not necessarily a Pareto improvement relative to laissez-faire—see online Appendix F for more details.

¹⁹This need not be true if governments are prohibited from using domestic taxes and afforded only import tax instruments. Following Venables (1987) and Ossa (2011), welfare-maximizing governments will erect tariffs in that case, even if they perceive world prices as invariant to their policy choice. By doing so, they improve allocative efficiency in the domestic economy but impose a negative firm-delocation externality on the rest of the world.

²⁰One aspect of internal cooperation requires further clarification: country i ’s policy could, in principle, disrupt the balance of market access concessions in the rest of the world, leading to a deterioration of one cooperative country’s ToT relative to another. Cooperation *within* the rest of the world entails that these extraterritorial ToT effects be neutralized via buffers that preserve w_n/w_j for all $n, j \neq i$ —see online Appendix G for further details.

all possible tax instruments. The first-best unilaterally optimal policy solves the following problem:²¹

$$\max_{(\mathbf{t}, \mathbf{x}_i, \mathbf{s}_i; \mathbf{w})} W_i(\mathbf{t}_i, \mathbf{x}_i, \mathbf{s}_i; \mathbf{w})$$

subject to

$$(P1) \quad (\mathbf{t}_i, \mathbf{x}_i, \mathbf{s}_i; \mathbf{w}) \in \mathbb{F}.$$

We analytically solve Problem (P1) under both the restricted and free entry cases. We perceive the restricted entry case to be a more appropriate benchmark if governments are concerned with *short-run* gains from policy. The free entry case, on the other hand, is more relevant if governments are concerned with *long-run* gains. These two cases exhibit an important difference: producer prices respond differently to contractions in export supply under restricted and free entry—as we elaborate next.

Conditional Export Supply Elasticity: The terms of trade gains from policy, in our framework, channel through changes in the price of imported and exported goods. The government in $i \in \mathbb{C}$ cannot directly dictate the *producer* price of say, good ji,k that is imported from origin $j \neq i$. Instead, it can deflate its producer price ($P_{ji,k}$) *indirectly* by contracting or expanding its export supply ($Q_{ji,k}$). The contraction in $Q_{ji,k}$ also affects the producer price of goods supplied by other locations through general equilibrium linkages. Our theory indicates that, for optimal policy analysis, the *conditional inverse export supply elasticity* is sufficient to track these effects. To present this elasticity, let $\tilde{\mathbb{P}}_i$ contain the consumer price of all goods either produced by or consumed in country i . These are prices that country i 's government can fully control via taxes. We define the *conditional inverse export supply elasticity* of good ji,k as

$$\omega_{ji,k} \equiv \frac{1}{r_{ji,k} \rho_{j,k}} \sum_{g \in \mathbb{K}} \left[\frac{\dot{w}_i L_i}{\dot{w}_j L_j} \rho_{i,g} \left(\frac{\partial \ln P_{ii,g}}{\partial \ln Q_{ji,k}} \right)_{\mathbf{w}, \mathbf{Y}, \tilde{\mathbb{P}}_i} + \sum_{n \neq i} \frac{\dot{w}_n L_n}{\dot{w}_j L_j} r_{ni,g} \rho_{n,g} \left(\frac{\partial \ln P_{ni,g}}{\partial \ln Q_{ji,k}} \right)_{\mathbf{w}, \mathbf{Y}, \tilde{\mathbb{P}}_i} \right],$$

where $r_{ni,g} \equiv (P_{ni,g} Q_{ni,g}) / (\sum_{\iota} P_{ni,g} Q_{ni,g})$ and $\rho_{n,g} \equiv (\sum_{\iota} P_{ni,g} Q_{\iota,g}) / (\sum_{\iota,s} P_{ni,s} Q_{\iota,s})$ respectively denote the *good-specific* and *industry-wide* sales shares associated with origin n . Notice, $\omega_{ji,k}$ is a conditional elasticity that describes how the *producer* prices linked to economy i respond to a change in $Q_{ji,k}$, holding $\tilde{\mathbb{P}}_i$ and the entire vector of wage and income levels constant. This elasticity encapsulates different economic forces under free and restricted entry, as we detail next.

Under restricted entry, producer prices from origin $j \in \mathbb{C}$ are fully determined by the (profit-adjusted) wage rate, \dot{w}_j , and the aggregate profit margin, $\bar{\mu}_j$

²¹ Given that the rest of the world is passive in their use of taxes (i.e., $\mathbf{t}_{-i} = \mathbf{x}_{-i} = \mathbf{s}_{-i} = \mathbf{0}$), we condense the notation by specifying equilibrium variables as a function of only $(\mathbf{t}_i, \mathbf{x}_i, \mathbf{s}_i; \mathbf{w})$.

(see equation (6)). Policy, thus, has two distinct effects on producer prices under restricted entry: one effect that channels through wages, \mathbf{w} , and another that channels through aggregate profit margins. To explain the latter, hold \mathbf{w} constant: contracting the export supply of good ji,k with taxes will alter all producer prices associated with origin j through a change in origin j 's aggregate profit margin, $\bar{\mu}_j$. The change in $\bar{\mu}_j$ derives from the fact that industries have differential markup margins and that taxing good ji,k alters the industrial composition of output in origin $j \in \mathbb{C}$.

Under free entry, producer prices from origin $j \in \mathbb{C}$ are determined by the wage rate, w_j , and the *origin j -industry k* -specific scale of production. So, aside from wage-related effects, policy has a second effect on producer prices that channels through industry-level scale economies. To elaborate, consider an import tax on good ji,k (origin j -destination i -industry k). Such a tax contracts the supply of ji,k and the scale of production in *origin j -industry k* . Given equation (6), this contraction in scale increases the entire vector of producer price indexes associated with *origin j -industry k* —all through additional firm entry.

In both cases, $\omega_{ji,k}$ describes how expanding or contracting good ji,k 's export supply impacts country i 's terms of trade via either profit-shifting or industry-level scale economies. Importantly, $\omega_{ji,k}$ can be characterized (to a first-order approximation) as a simple function of sales shares, scale elasticities, and Marshallian demand elasticities (see online Appendix E):²²

$$(10) \quad \omega_{ji,k} \approx \begin{cases} \frac{-\frac{\mu_k}{1+\mu_k} r_{ji,k}}{1 - \frac{\mu_k}{1+\mu_k} \sum_{i \neq j} r_{ji,k} \varepsilon_{ji,k}} \left[1 - \frac{\mu_k}{1+\mu_k} \frac{w_i L_i}{w_j L_j} \sum_{n \neq i} \rho_{i,k} r_{in,k} \varepsilon_{in,k}^{(jn,k)} \right], & \text{if entry is free;} \\ \frac{\left(1 - \frac{1+\bar{\mu}_j}{1+\mu_k}\right) \sum_g r_{ji,g} \rho_{j,g}}{1 + \sum_g \sum_{i \neq j} \left[1 + \left(1 - \frac{1+\bar{\mu}_j}{1+\mu_g}\right) r_{ji,g} \rho_{j,g} \varepsilon_{ji,g} \right]}, & \text{if entry is restricted.} \end{cases}$$

The above formulation for $\omega_{ji,k}$ is quite intuitive. Under restricted entry, $\omega_{ji,k}$ governs the relationship between export supply and the average markup paid on imports. Accordingly, $\omega_{ji,k}$ is nonzero only when industries exhibit differential markup levels. Otherwise, $\omega_{ji,k}$ collapses to zero, as the average markup (or profit margin) paid on imports is constant and invariant to changes in export supply, i.e., $\bar{\mu}_j = \mu_k = \mu \Rightarrow \omega_{ji,k} = 0$. Under free entry, $\omega_{ji,k}$ regulates the terms of trade gains from policy that channel through scale economies. Accordingly, in the limit where industries operate based on constant returns-to-scale, $\omega_{ji,k}$ once again collapses to zero—namely, $\lim_{\mu_k \rightarrow 0} \omega_{ji,k} = 0$.

Three-Step Dual Approach to Characterizing Optimal Policy: Our characterization of optimal policy employs the dual approach and is presented in online Appendix E. Below, we provide a verbal summary of our approach, which involves three main steps.

First, we simplify Problem (P1) by reformulating it into a problem where country i 's government chooses the vector of prices $\tilde{\mathbb{P}}_i = \{\tilde{\mathbf{P}}_{ii}, \tilde{\mathbf{P}}_{ji}, \tilde{\mathbf{P}}_{ij}\}$ associated with its

²²The above approximation derives from Wu et al.'s (2013) first-order approximated inverse of a diagonally dominant matrix. Figure E.1 in online Appendix E illustrates the precision of this approximation. Online Appendix E also presents an exact (approximation-free) formulation for $\omega_{ji,k}$.

own economy. Country i 's optimal tax/subsidy schedule $\mathbb{T}_i^* \equiv (\mathbf{t}_i^*, \mathbf{x}_i^*, \mathbf{s}_i^*)$ is then recovered as the wedge between the optimal price vector $\tilde{\mathbb{P}}_i^*$ and producer prices.

Second, we derive the first-order conditions (F.O.C.) associated with country i 's reformulated optimal policy problem. We use two technical tricks to overcome the complications related to general equilibrium analysis. First, we use the envelope conditions associated with optimal demand choices to net out redundant behavioral responses. Second, we identify additional welfare neutrality conditions specific to Problem (P1). Most importantly, we observe that terms in the F.O.C.s that account for general equilibrium wage and income effects are redundant in the neighborhood of the optimum. That is, we could specify the F.O.C.s associated with (P1) as if wages were constant and Marshallian demand functions were income-inelastic.²³

Third, we combine the F.O.C.s and solve them as part of one system. In this process, we appeal to the tax neutrality result specified by Lemma 1 to eliminate redundant tax shifters, which are difficult to characterize. We then appeal to well-known properties of Marshallian demand functions (e.g., *Cournot aggregation* and *homogeneity of degree zero*) to establish that our system of F.O.C.s admits a unique solution.²⁴ Together, these steps lead us to simple sufficient statistics formulas for unilaterally optimal policies, as summarized by the following theorem.²⁵

THEOREM 1: *Country i 's unilaterally optimal policy is unique up to two uniform tax shifters $1 + \bar{s}_i$ and $1 + \bar{t}_i \in \mathbb{R}_+$ and is implicitly given by*

$$\begin{aligned} \text{[domestic subsidy]} \quad & 1 + s_{i,k}^* = (1 + \mu_k)(1 + \bar{s}_i), \\ \text{[import tax]} \quad & 1 + t_{ji,k}^* = (1 + \omega_{ji,k})(1 + \bar{t}_i), \\ \text{[export subsidy]} \quad & \mathbf{1} + \mathbf{x}_{ij}^* = -\mathbf{E}_{ij}^{-1} \mathbf{E}_{ij}^{(-ij)} (\mathbf{1} + \mathbf{t}_i^*), \end{aligned}$$

where $\omega_{ji,k}$ denotes the good ji,k 's inverse export supply elasticity as given by equation (10), while $\mathbf{E}_{ij} \sim \mathbf{E}_{ij}^{(ij)}$ and $\mathbf{E}_{ij}^{(-ij)}$ denote matrixes of Marshallian demand elasticities as defined under (D1).²⁶

The uniform tax shifters, \bar{s}_i and \bar{t}_i account for the multiplicity of optimal policy equilibria (as indicated by Lemma 1). These shifters can be assigned any arbitrary value, provided that $1 + \bar{s}_i$ and $1 + \bar{t}_i \in \mathbb{R}_+$. For instance, if we assign a sufficiently high value to \bar{t}_i and \bar{s}_i , the optimal policy will involve import tariffs, export subsidies, and industrial subsidies. Conversely, if we assign a sufficiently low value to \bar{t}_i and \bar{s}_i , the

²³Farrokhi and Lashkaripour (2021) streamline the dual approach developed in this paper, extending our result about the welfare neutrality of wages and income effects to settings with arbitrary international externalities.

²⁴To be clear, it is possible that our model admits multiple optimal policy equilibria. Yet the optimal policy formulas are uniquely specified by Theorem 1 in each case.

²⁵We later combine the formulas specified by Theorem 1 with micro-estimated parameter values to quantify the *ex ante* gains from policy. In online Appendix H, we test the accuracy and speed of our formulas by performing 150 numerical simulations in which the underlying model parameters are repeatedly sampled from a uniform distribution. The theoretical policy predictions are then compared to those obtained from numerical optimization.

²⁶ $\mathbf{E}_{ij}^{(-ij)} = [\mathbf{E}_{ij}^{(ij)}]_{n \neq i}$ is a $K \times (N - 1)K$ matrix, and $\mathbf{1} \equiv \mathbf{1}_{(N-1)K \times 1}$ is a column vector of ones. Also, in the general case with asymmetric income elasticities of demand, \mathbf{E}_{ij} should be replaced with $\tilde{\mathbf{E}}_{ij} \equiv \left[\frac{e_{ij,g}}{e_{ij,k}} \varepsilon_{ij,g}^{(j,k)} \right]_{g,k}$. Otherwise, the symmetry of the Slutsky matrix implies that $\frac{e_{ij,g}}{e_{ij,k}} \varepsilon_{ij,g}^{(j,k)} = \varepsilon_{ij,k}^{(i,g)}$, which implies that $\mathbf{E}_{ij} = \tilde{\mathbf{E}}_{ij}$.

optimal policy will involve import subsidies, export taxes, and industrial production taxes.

Intuition behind Optimal Tax Formulas: Theorem 1 states that country i 's unilaterally optimal policy consists of (i) Pigouvian subsidies that restore marginal cost pricing in economy i ; (ii) import taxes/subsidies that exploit country i 's collective import market power, delivering an optimal markdown on the producer price of imported goods $P_{ji,k}$; and (iii) export taxes/subsidies that exploit country i 's collective export market power, charging the optimal national-level markup on the consumer price of exported goods $\tilde{P}_{ij,k}$.²⁷

Theorem 1 has two additional implications worth highlighting. The first is that first-best optimal tariffs and export subsidies are *misallocation-blind* but not necessarily blind to the overall magnitude of scale economies. In particular, \mathbf{t}_{ji}^* and \mathbf{x}_{ij}^* are misallocation-blind in that allocative efficiency is restored exclusively via industrial subsidies under the first-best. At the same time, \mathbf{t}_{ji}^* and \mathbf{x}_{ij}^* are sensitive to scale economies because raising the average scale elasticity modifies \mathbf{t}_{ji}^* and \mathbf{x}_{ij}^* regardless of the underlying degree of misallocation.²⁸

The optimal export tax-cum-subsidy, furthermore, depends on the entire matrix of own- and cross-price demand elasticities, echoing our previous assertion that $x_{ij,k}^*$ (in Theorem 1) corresponds to the optimal markup of a multiproduct monopolist. To better understand this point, assign $\bar{t}_i = 0$, in which case $x_{ij,k}^*$ represents a tax on good ij, k (rather than a subsidy). The optimal tax rate on ij, k is equal to the optimal markup on that good if country i 's government was pricing its exports as a multiproduct monopolist rather than an individual single-product firm. The government's optimal pricing decision, accordingly, internalizes the effect of raising $\tilde{P}_{ij,k}$ on its sales of other products in destination j .

Lastly, Theorem 1 can be useful for quantitative applications. It specifies optimal policy as a function of Marshallian demand elasticities and inverse export supply elasticities, both of which are fully determined by observable shares ($r_{ij,k}$ and $\lambda_{ij,k}$) and industry-level trade and scale elasticities (σ_k and μ_k). In other words, Theorem 1 characterizes optimal policy in terms of a set of observable or estimable *sufficient statistics*. We capitalize on this feature to simplify our quantitative analysis of optimal policy in Section VI.

Optimal Tariffs Are Uniform, Absent Scale Economies or Profits: A canonical special case of Theorem 1 is the multi-industry Armington case, in which $\mu_k = 0$ for all $k \in \mathbb{K}$. In that case, $\omega_{ji,k} = 0$ for all ji, k , implying that optimal import tariffs

²⁷Theorem 1 reveals that once scale/markup distortions are corrected via domestic subsidies, optimal trade taxes resemble those derived by Dixit and Norman (1980) in perfectly competitive, constant returns-to-scale environments (see also Matsuyama 2008). The main difference between our formula and Dixit and Norman (1980) is that our Theorem 1 equates the optimal tariff to a conditional elasticity, $\omega_{ji,k}$, free of general equilibrium wage-and-income effects. These difficult-to-characterize general equilibrium effects, we prove, are redundant in the neighborhood of the optimum policy and can be disregarded. This specific feature of our optimal policy formulas makes them useful for quantitative analysis—as we demonstrate later in Section VI.

²⁸By the degree of misallocation, we mean the log welfare distance to the efficient frontier, \mathcal{L}_i . In a closed economy with Cobb-Douglas-CES preferences, $\mathcal{L}_i = E_{\rho_i}[\mu \log \mu] - E_{\rho_i}[\mu] \log E_{\rho_i}[\mu]$, where $E_{\rho_i}[\mu]$ denotes the sales-weighted average scale elasticity. Note that scale economies do not necessarily lead to misallocation. If the scale elasticity is strictly positive but uniform across industries, then $\mathcal{L}_i = 0$.

are uniform; i.e., $t_{ji,k}^* = \bar{t}_i$ for all ji,k . Intuitively, in the absence of scale economies or profits, import tariffs cannot influence the producer price of imports on a good-by-good basis. At best, they can elicit a uniform reduction in import prices by deflating w_{-i} relative to w_i , which is best achieved via a uniform import tax rate. Section IV shows that in the absence of scale economies and markups, first-best optimal tariffs remain uniform even with input-output linkages.²⁹

Special Case with Cobb-Douglas Preferences across Industries: To gain deeper intuition about Theorem 1, consider a special case where preferences are Cobb-Douglas across industries. In that case, the formulas specified by Theorem 1 reduce to³⁰

$$\begin{aligned}
 \text{[domestic subsidy]} \quad & 1 + s_{i,k}^* = (1 + \mu_k)(1 + \bar{s}_i), \\
 \text{[import tax]} \quad & 1 + t_{ji,k}^* = (1 + \omega_{ji,k})(1 + \bar{t}_i), \\
 \text{(11) [export subsidy]} \quad & 1 + x_{ij,k}^* = \frac{(\sigma_k - 1) \sum_{n \neq i} [(1 + \omega_{ni,k}) \lambda_{nj,k}]}{1 + (\sigma_k - 1)(1 - \lambda_{ij,k})} (1 + \bar{t}_i).
 \end{aligned}$$

A well-known special case of the above formula is the *single-industry* \times *two-country* formula in Gros (1987). To demonstrate this, drop the industry subscript k and reduce the global economy into two countries, i.e., $\mathbb{C} = \{i, j\}$. Noting that $1 - \lambda_{ij} = \lambda_{ji}$ in the two-country case, we can deduce from the above formulas that

$$\frac{1 + t_{ji}^*}{1 + x_{ij}^*} = 1 + \frac{1}{(\sigma - 1)\lambda_{ji}}.$$

By the Lerner symmetry, export and import taxes are equivalent in the single-industry model.³¹ Hence, without loss of generality, we can set $x_{ij}^* = 0$ and arrive at the familiar-looking optimal tariff formula in Gros (1987), i.e., $t_{ji}^* = 1/(\sigma - 1)\lambda_{ji}$.

The Cobb-Douglas case of Theorem 1 is also a strict generalization of the formula derived concurrently by Bartelme et al. (2019) for a small open economy with multiple sectors. Specifically, enforcing the *small open economy* assumption—i.e., setting $\omega_{ji,k} \approx \lambda_{ij,k} \approx 0$; $\lambda_{jj,k} \approx 1$ —our optimal policy formulas in the Cobb-Douglas case reduce to

$$(12) \quad 1 + s_{i,k}^* = 1 + \mu_k; \quad t_{ji,k}^* = 0; \quad 1 + x_{ij,k}^* = \frac{\sigma_k - 1}{\sigma_k}.$$

²⁹To be clear, these results hinge on the restriction that country i 's policy does not influence aggregate relative wages in the rest of the world. This restriction holds trivially in the two-country case but requires internal cooperation in the rest of the world otherwise (see online Appendix G).

³⁰In the Cobb-Douglas case, $\varepsilon_{nj,k}^{(ij,k)} = -\sigma_k \mathbf{1}_{n=j} + (\sigma_k - 1)\lambda_{ij,k}$ and $\varepsilon_{nj,g}^{(ij,k)} = 0$ if $g \neq k$, which when plugged into the restricted entry case of equation (6), delivers

$$\omega_{ji,k} \approx \frac{\left(1 - \frac{\bar{\mu}_j}{\mu_k}\right) \sum_g r_{ji,g} \rho_{j,g}}{1 + \sum_g \sum_{i \neq j} \left\{1 - \left(1 - \frac{\bar{\mu}_i}{\mu_g}\right) r_{ji,g} \rho_{j,g} \left[1 + (\sigma_g - 1)(1 - \lambda_{j,g})\right]\right\}}.$$

The parameterization of $\omega_{ji,k}$ under free entry can be derived similarly.

³¹The Lerner symmetry is a special case of the equivalence result presented under Lemma 1. Also, note that the decentralized equilibrium is efficient in the single industry Krugman model studied by Gros (1987). As such, the optimal industrial subsidy can also be normalized to zero, i.e., $s_i^* = 0$.

Second-Best: Unilaterally Optimal Import Tariffs and Export Subsidies.— Suppose the government in $i \in \mathbb{C}$ cannot use domestic subsidies due to say institutional barriers or political pressures. It is optimal, in that case, to use trade taxes as a second-best policy to restore allocative efficiency in the domestic economy. In this section, we derive analytic formulas for *second-best* optimal trade taxes in such circumstances. Country i 's optimal policy problem, in this case, includes an added constraint that $\mathbf{s}_i = \mathbf{0}$:

$$(P2) \quad \max_{(\mathbf{t}_i, \mathbf{x}_i, \mathbf{s}_i; \mathbf{w})} W_i(\mathbf{t}_i, \mathbf{x}_i, \mathbf{s}_i; \mathbf{w}) \quad \text{subject to} \quad \begin{cases} (\mathbf{t}_i, \mathbf{x}_i, \mathbf{s}_i; \mathbf{w}) \in \mathbb{F}; \\ \mathbf{s}_i = \mathbf{0}. \end{cases}$$

Using the dual approach discussed earlier, we analytically solve Problem (P2) and derive sufficient statistics formulas for second-best optimal trade taxes. The following theorem presents these formulas, with a formal proof provided in online Appendix I.

THEOREM 2: *Suppose the government is unable (or unwilling) to apply domestic industrial subsidies. In that case, the second-best optimal import tariffs and export subsidies are unique up to a uniform tax shifter $\bar{t}_i \in \mathbb{R}_+$ and are implicitly given by*

$$\begin{aligned} \text{[import tariff]} \quad \mathbf{1} + \mathbf{t}_{ji}^{**} &= (1 + \bar{t}_i)(\mathbf{1} + \mathbf{\Omega}_{ji}) \oslash \left[\mathbf{1} + \mathbf{E}_{-ii}^{-1} \mathbf{E}_{-ii}^{(ii)} \left(1 - \frac{1 + \mu_k}{1 + \bar{\mu}_i} \right)_k \right], \\ \text{[export subsidy]} \quad \mathbf{1} + \mathbf{x}_{ij}^{**} &= -(1 + \bar{t}_i) \left[\mathbf{E}_{ij}^{-1} \mathbf{E}_{ij}^{(-ij)} (\mathbf{1} + \mathbf{\Omega}_{-ii}) \right] \odot \left(\frac{1 + \mu_k}{1 + \bar{\mu}_i} \right)_k, \end{aligned}$$

where $\mathbf{\Omega}_{ji} = [\omega_{ji,k}]_k$ is a vector of inverse export supply elasticities (equation (10)); $\bar{\mu}_i$ denotes the output-weighted average markup in economy i (equation (5)); and \mathbf{E}_{-ii} , $\mathbf{E}_{-ii}^{(ii)}$, \mathbf{E}_{ij} , and $\mathbf{E}_{ij}^{(-ij)}$ are matrixes of Marshallian demand elasticities as defined under Definition (D1).³²

Theorem 2 asserts that, when governments cannot use industrial subsidies, the optimal export subsidy is adjusted to promote exports in high returns-to-scale (high- μ) industries and the optimal import tax is adjusted to restrict import competition in high returns-to-scale (high- μ) industries. Intuitively, the government's objective when solving (P2) is to mimic Pigouvian industrial subsidies with trade taxes/subsidies. To reach this objective, import taxes and export subsidies should increase in high returns-to-scale industries relative to the *first-best* benchmark. While these adjustments elevate domestic production in high- μ industries, they are insufficient for obtaining the unilaterally first-best allocation.

³²Letting N and K denote the number of countries and industries, $\mathbf{e}_{-ii} \sim \mathbf{E}_{-ii}^{(-ii)} = [\mathbf{E}_{ni}^{(j)}]_{n \neq i, j \neq i}$ is a square $(N-1)K \times (N-1)K$ matrix, where $\mathbf{E}_{ni}^{(j)} \equiv [\varepsilon_{ni,k}^{(j,g)}]_{k,g}$ as defined under Definition (D1). Likewise, $\mathbf{E}_{ij}^{(-ij)} = [\mathbf{E}_{ij}^{(nj)}]_{n \neq i}$ and $\mathbf{E}_{-ii}^{(ii)} = [\mathbf{E}_{ni}^{(ii)}]_{n \neq i}$ are respectively, $K \times (N-1)K$ and $(N-1)K \times K$ matrixes. In all the equations, $\mathbf{1} \equiv \mathbf{1}_{(N-1)K \times 1}$ is a columns vector of ones. Meanwhile, $\mathbf{\Omega}_{-ii} = [\omega_{ni,k}]_{n \neq i, k}$ is a $(N-1)K \times 1$ vector, and the operators \odot and \oslash denote element-wise multiplication and division.

Special Case with Cobb-Douglas Preferences across Industries: We can invoke the Cobb-Douglas assumption to further elucidate the second-best tax formulas under Theorem 2. Under this assumption, there are zero cross-demand effects between industries, and the optimal policy formulas specified by Theorem 2 can be simplified as follows:

$$\begin{aligned}
 \text{[import tariff]} \quad 1 + t_{ji,k}^{**} &= \frac{1 + (\sigma_k - 1)\lambda_{ii,k}}{1 + \frac{1 + \bar{\mu}_i}{1 + \mu_k}(\sigma_k - 1)\lambda_{ii,k}}(1 + t_{ji,k}^*), \\
 \text{[export subsidy]} \quad 1 + x_{ij,k}^{**} &= \frac{1 + \mu_k}{1 + \bar{\mu}_i}(1 + x_{ij,k}^*),
 \end{aligned}$$

where $1 + t_{ji,k}^* = (1 + \omega_{ji,k})(1 + \bar{t}_i)$ and $1 + x_{ji,k}^* = \frac{(\sigma_k - 1)\sum_{n \neq i} [(1 + \omega_{ni,k})\lambda_{nj,k}]}{1 + (\sigma_k - 1)(-\lambda_{ij,k})}(1 + \bar{t}_i)$ denote the first-best optimal rate (equation (11)). For a *small open economy*, the formulas further reduce to

$$\begin{aligned}
 1 + t_{ji,k}^{**} &= \frac{1 + (\sigma_k - 1)\lambda_{ii,k}}{1 + \frac{1 + \bar{\mu}_i}{1 + \mu_k}(\sigma_k - 1)\lambda_{ii,k}}(1 + \bar{t}_i); \\
 1 + x_{ij,k}^{**} &= \frac{1 + \mu_k}{1 + \bar{\mu}_i} \left(\frac{\sigma_k - 1}{\sigma_k} \right) (1 + \bar{t}_i).
 \end{aligned}$$

In summary, the above formulas indicate that second-best *import* taxes are higher in industries with a greater-than-average markup and industries in which country *i* has a comparative advantage (i.e., high- $(\sigma_k - 1)\lambda_{ii,k}$ industries). These two properties allow second-best import taxes to mimic Pigouvian subsidies to the best extent possible. Likewise, second-best *export* subsidies feature a *misallocation-correcting* component that favors industries with a higher-than-average scale elasticity or markup.

Importantly, if the markup or scale elasticity is uniform across industries (i.e., $\mu_k = \mu = \bar{\mu}_i$), the above formulas yield the *first-best* or purely ToT-improving tax rate—i.e., $t_{ji,k}^{**} = t_{ji,k}^*$ and $x_{ij,k}^{**} = x_{ij,k}^*$. The intuition is that the Krugman model *without* cross-industry markup heterogeneity is efficient, leaving no room for policy interventions to restore allocative efficiency.

Third-Best: Unilaterally Optimal Import Tariffs.—Now suppose that, in addition to restrictions on industrial subsidies, the use of export subsidies is also restricted. The government’s optimal policy problem in this case features two additional constraints, $s_i = x_i = \mathbf{0}$:

$$\text{(P3)} \quad \max_{(\mathbf{t}_i, \mathbf{x}_i, \mathbf{s}_i; \mathbf{w})} W_i(\mathbf{t}_i, \mathbf{x}_i, \mathbf{s}_i; \mathbf{w}) \quad \text{subject to} \quad \begin{cases} (\mathbf{t}_i, \mathbf{x}_i, \mathbf{s}_i; \mathbf{w}) \in \mathbb{F}; \\ \mathbf{s}_i = \mathbf{x}_i = \mathbf{0}. \end{cases}$$

Some variation of the above problem has been studied by an expansive literature on optimal tariffs. However, nearly all existing studies are limited to partial equilibrium two-by-two models. Here, we use the same dual approach described earlier

to analytically solve Problem (P3) within our multicountry, multi-industry general equilibrium framework. Our derivation, as before, yields simple sufficient statistics formulas for optimal third-best import taxes. The following theorem presents these formulas, with a formal proof provided in online Appendix J.³³

THEOREM 3: *Suppose the government is unable to apply domestic industrial subsidies or export subsidies. In that case, third-best optimal import tariffs are uniquely given by*

$$\mathbf{1} + \mathbf{t}_{ji}^{***} = (1 + \bar{\tau}_i^*)(\mathbf{1} + \mathbf{\Omega}_{ji}) \odot \left(\mathbf{1} + \mathbf{E}_{-ii}^{-1} \mathbf{E}_{-ii}^{(ii)} \left[1 - \frac{1 + \mu_k}{1 + \bar{\mu}_i} \right]_k \right),$$

where $\bar{\tau}_i^* = \left[-\sum_{g,s} \sum_{j \neq i} \chi_{ij,g} (1 + \varepsilon_{ij,g}^{(ij,s)}) \right]^{-1}$ is a uniform tariff shifter that represents the elasticity of international demand for country i 's labor (with $\chi_{ij,g} \equiv P_{ij,g} Q_{ij,g} / \sum_{n \neq i} \mathbf{P}_{in} \cdot \mathbf{Q}_{in}$ denoting export shares). $\bar{\mu}_i$ denotes the output-weighted average markup in economy i as described by equation (5), and \mathbf{E}_{-ii} and $\mathbf{E}_{-ii}^{(ii)}$ are matrixes of Marshallian demand elasticities as defined under Definition (D1).

Unlike Theorems 1 and 2, the third-best optimal tariff schedule identified by Theorem 3 is unique. That is because the multiplicity implied by Lemma 1 no longer applies when both export and industrial subsidies are restricted to zero. Nevertheless, the third-best tariff specified by Theorem 3 differs from the second-best tariffs (in Theorem 2) by only a uniform tariff shifter, $1 + \bar{\tau}_i^*$. So, barring the uniform component, $1 + \bar{\tau}_i^*$, we can understand the above formula based on the same intuition provided under Theorem 2.

The uniform tariff component, $1 + \bar{\tau}_i^*$, compensates for the unavailability of export tax-cum-subsidies to the government. By the Lerner symmetry, which is implicit in Lemma 1, import taxes can perfectly mimic a uniform export tax. This ability was previously redundant (under Theorems 1 and 2) because export taxes/subsidies were directly applicable, and there was no point in using other instruments to mimic them. But since export taxes are restricted under Problem (P3), it is optimal to uniformly raise all tariffs by a factor $1 + \bar{\tau}_i^*$, using them as a second-best substitute for optimal export taxes/subsidies.

III. Discussion: The Efficacy of Trade and Industrial Policy

This section, guided by Theorems 1–3, discusses the efficacy of trade and industrial policy in distorted open economies. We conjecture that stand-alone trade policy measures can be ineffective even when chosen optimally. Unilateral scale correction via industrial policy can also backfire, underscoring the importance of international coordination. We later test these conjectures by estimating model parameters and utilizing our optimal policy formulas.

³³In the special case where entry is restricted and countries are sufficiently small, the optimal tariff formula presented under Theorem 4 reduces to the formula used by Lashkaripour (2021) to examine global tariff wars.

Tension between Allocative Efficiency and ToT: Theorems 2 and 3 reveal that second-best trade policies seek to strike a balance between (a) improving the terms of trade, which requires contracting exports in nationally differentiated (low- σ) industries, and (b) correcting misallocation, which requires expanding output in high returns-to-scale (high- μ) industries. Obtaining this balance becomes difficult if not impossible when $\text{cov}(\sigma_k, \mu_k) < 0$ —which is the empirically relevant case based on our forthcoming estimation. To navigate this tension, a welfare-maximizing government must tailor its (second-best) trade policy in a way that curtails the ToT gains without necessarily correcting misallocation. This balancing act erodes the gains from second-best trade policies and can even render them industry-blind—unable to beneficially correct interindustry misallocation *or* manipulate industry-specific export market power.³⁴

The tension between allocative efficiency and ToT can be theoretically demonstrated for a local change in policy (see online Appendix L). But how this tension precisely modifies the ex ante gains from trade policy is an empirical matter. Our conjecture is that if industry-level trade and scale elasticities exhibit a *strong* negative correlation, the welfare gains from trade policy are limited—a claim we evaluate quantitatively in Section VI.

CONJECTURE 1: *If industry-level trade and scale elasticities exhibit a strong negative correlation ($\text{cov}(\sigma_k, \mu_k) \ll 0$), stand-alone trade policy measures deliver limited welfare gains even when set optimally.*

A formal evaluation of this conjecture requires estimating model parameters, as performed in Sections V and VI. Nevertheless, we can numerically illustrate this point using artificial parameter values, as presented in Figure 1. The left panel demonstrates that the gains from second-best trade policies diminish rapidly as $\text{cov}(\sigma_k, \mu_k)$ is artificially lowered from positive to negative values. In each case, the trade elasticities are held constant, meaning that the scope for ToT gains remains the same. The only thing that changes is the tension between ToT and corrective gains from policy, which amplifies as $\text{cov}(\sigma_k, \mu_k)$ becomes more negative.³⁵

Unilateral Scale Correction Can Cause Immiserizing Growth: The flip side of the noted tension is that if $\text{cov}(\sigma_k, \mu_k) < 0$, unilateral implementation of scale-correcting Pigouvian subsidies worsens the ToT, resulting in possibly adverse welfare consequences. These arguments echo the *immiserizing growth* paradox in

³⁴In the canonical Krugman (1980) model where $\mu_k = 1/(\sigma_k - 1)$, optimal (second-best) import tariffs and export subsidies are industry-blind for a small open economy. In particular, applying Theorem 2 to this special case yields

$$1 + t_{j,k}^{**} = 1 + \bar{t}_i; \quad 1 + x_{j,k}^{**} = (1 + \bar{t}_i) \left(1 - \frac{1}{\bar{\sigma}_i} \right),$$

where $\bar{t}_i \in \mathbb{R}$ is an arbitrary tax shifter and $\bar{\sigma}_i = (\sum_k \rho_{i,k} / \sigma_k)^{-1}$ is the sales-weighted average trade elasticity facing country i . The optimal trade tax in each industry is evidently blind to misallocation (μ_k) or industry-specific export market power (σ_k)—reflecting the difficulty to reconcile these two policy considerations. All this policy choice can achieve is to improve country i 's aggregate ToT by inflating its wage relative to the rest of the world.

³⁵This tension is distinct from the *targeting principle* (Bhagwati and Ramaswami 1963), which applies irrespective of the sign of $\text{cov}(\sigma_k, \mu_k)$. Indeed, second-best trade taxes become more potent despite the targeting principle if $\text{cov}(\sigma_k, \mu_k) > 0$; but we focus on the case where $\text{cov}(\sigma_k, \mu_k) < 0$, as it aligns with our forthcoming estimation.

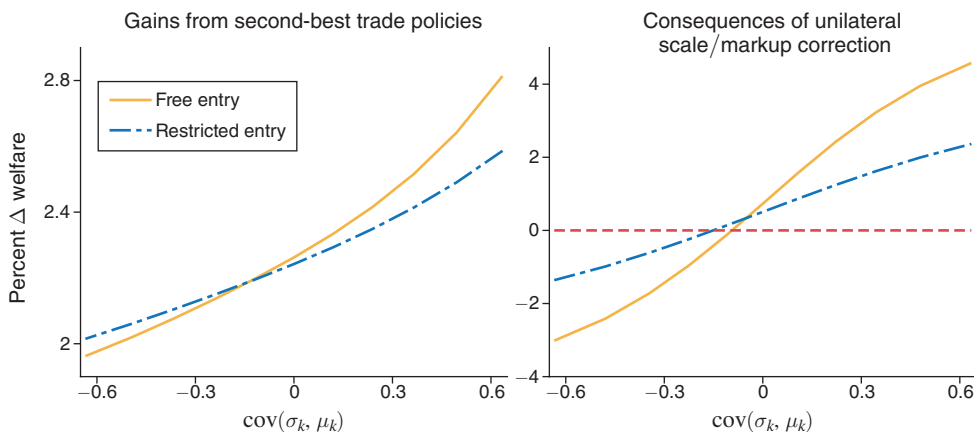


FIGURE 1. TENSION BETWEEN TOT AND ALLOCATIVE EFFICIENCY, WHEN $\text{cov}(\sigma_k, \mu_k) < 0$, CAN YIELD DIRE POLICY OUTCOMES

Notes: This figure corresponds to a two-country and two-industry model with symmetric countries and Cobb-Douglas preferences across industries. The left panel reports the gains from second-best trade taxes, specified by Theorem 2. The trade elasticities in industries 1 and 2 are assigned values $\sigma_1 = 1.5$ and $\sigma_2 = 3$, and scale elasticities are adjusted to vary $\text{cov}(\sigma_k, \mu_k)$. The right panel reports the welfare consequences of unilateral scale or markup correction. Industry-level scale elasticities are assigned values $\mu_1 = 0.5$ and $\mu_2 = 0.2$, and trade elasticities are adjusted to vary $\text{cov}(\sigma_k, \mu_k)$.

Bhagwati (1958) and follow a simple logic. If $\text{cov}(\sigma_k, \mu_k) < 0$, scale-correcting Pigouvian subsidies expand domestic output in high- σ industries, which are nationally differentiated industries in which countries hold significant export market power. Elevating output and, thus, exports in these industries could worsen the ToT to the point of triggering immiserizing welfare effects.

Online Appendix L demonstrates theoretically that if $\text{cov}(\sigma_k, \mu_k) < 0$, unilateral scale/markup correction worsens the ToT. But whether these adverse ToT effects outweigh the allocative efficiency gains from scale/markup correction is an empirical matter. We conjecture that if industry-level trade and scale elasticities exhibit a *strong* negative correlation, the adverse ToT effects from scale correction are large enough to cause immiserizing growth.

CONJECTURE 2: *If industry-level trade and scale elasticities exhibit a strong negative correlation ($\text{cov}(\sigma_k, \mu_k) \ll 0$), unilateral scale (or markup) correction via industrial policy will likely worsen national welfare, echoing the immiserizing growth paradox.*

We evaluate this conjecture with micro-estimated parameters in Section VI, but a numerical illustration with artificial parameter values is also provided in Figure 1 (right panel). This figure is produced by fixing the degree of interindustry misallocation and artificially adjusting trade elasticities to vary $\text{cov}(\sigma_k, \mu_k)$ from positive to negative values. In regions where $\text{cov}(\sigma_k, \mu_k) < 0$, a unilateral adoption of scale-correcting subsidies compromises welfare, causing immiserizing growth. These immiserizing consequences, as we discuss next, can be an obstacle for industrial policy implementation for countries committed to efficient trade policies under shallow agreements.

TABLE 2—THE INDUSTRIAL POLICY IMPLEMENTATION GAME FACING COOPERATIVE GOVERNMENTS WHEN $\text{cov}(\sigma_k, \mu_k) < 0$

		Country j ($\% \Delta W_j$)	
		$s_j = \mathbf{0}$	$s_j = \mu$
		Country i ($\% \Delta W_i$)	
$s_i = \mathbf{0}$	(0%, 0%)	(3.7%, -1.2%)	
$s_i = \mu$	(-1.2%, 3.7%)	(2.7%, 2.7%)	

Notes: This table corresponds to a two-country and two-industry model with symmetric countries and CES preferences across industries with substitution elasticity 1.2. Industry-level trade elasticities are $\sigma_1 = 1.5$ and $\sigma_2 = 3$, and scale elasticities are $\mu_1 = 0.2$ and $\mu_2 = 0.5$, implying $\text{cov}(\sigma_k, \mu_k) \approx -0.225$. The unique Nash equilibrium is a race to the bottom, wherein $s_i = s_j = \mathbf{0}$.

Industrial Policy Coordination via Deep Agreements: Immiserizing growth can be a major obstacle to industrial policy implementation in open economies. To convey this point, we adopt the common view that international negotiations involve two stages.³⁶ In the first-stage, governments negotiate over policy space to ensure each party restricts itself to the globally efficient policy choice (equation (9)). In the second stage, governments negotiate a deeper agreement to ensure the implementation of efficient (misallocation-correcting) policies, with each country having the choice to either implement efficient policies or withhold implementation.

Table 2 illustrates the second-stage implementation game facing cooperative countries. The game involves two cooperative countries (i and j) that can take two actions: (i) implement scale-correcting Pigouvian subsidies, i.e., $s_i = \mu$, or (ii) withhold implementation and stick to business as usual, i.e., $s_i = \mathbf{0}$. The efficient outcome is the implementation of Pigouvian subsidies in both countries, which will boost welfare by 2.7 percent across the board. But this outcome is not sustainable without formal coordination because each country has an incentive to free ride on the other country’s implementation. The outcome of this game is a race to the bottom, wherein no party is willing to correct misallocation in domestic industries without violating its commitments to shallow cooperation (i.e., zero trade taxes).³⁷ A deep agreement that ensures reciprocity in industrial policy implementation can remedy this problem.

In principle, the first-mover country can renegotiate their tariffs to nullify the adverse ToT effects of unilateral scale correction—for example, under Article XXVIII of the GATT. But these negotiations are costly, suggesting that governments might prefer to refrain from implementation and opt for free riding. Alternatively, countries may pair their industrial policy with *hidden* trade restric-

³⁶Modeling international cooperation as a two-stage game consisting of *enactment* and *implementation* stages is commonplace in the global governance literature (Shaffer and Pollack 2009). In the trade and environmental policy literature, many studies treat international negotiations as multistage games where initial stages restrict policy choices and latter stages ensure implementation (e.g., Murdoch, Sandler, and Vijverberg 2003; Drazen and Limão 2008; Kosfeld, Okada, and Riedl 2009).

³⁷We must emphasize that a race to the bottom will not occur if $\text{cov}(\sigma_k, \mu_k) \geq 0$. In that case, a unilateral adoption of corrective subsidies improves rather than deteriorates the terms of trade, making unilateral implementation a dominant strategy. We nevertheless maintain our focus on the case where $\text{cov}(\sigma_k, \mu_k) < 0$ because it aligns with our forthcoming estimation of scale and trade elasticities.

tions, enabling them to avoid adverse ToT effects without incurring the administrative cost of tariff renegotiation. The Chinese government, for instance, has coupled its proactive industrial policy measures with hidden export taxes, which are applied via partial value-added tax rebates and enhance China's terms of trade (Garred 2018).

IV. Extensions and Application to Other Canonical Models

In this section, we first show that our theoretical results readily apply to two other canonical trade models. We then extend our baseline theoretical results to richer environments featuring input-output linkages and political economy pressures.

A. Application to Other Canonical Trade Models

The optimal policy formulas specified by Theorems 1–3 apply to two other canonical trade models. However, parameters σ_k and μ_k in these formulas adopt different interpretations, which reflects the different micro-foundation underlying these frameworks.

The Eaton-Kortum Model with Marshallian Externalities.—Consider a multi-industry Eaton and Kortum (2002) model where industry-level production is subject to agglomeration economies. Let ψ_k denote the constant agglomeration elasticity in industry k and θ_k denote the Eaton-Kortum Fréchet shape parameter. Theorem 1 characterizes the optimal policy in this model under the following reinterpretation of parameters: $\mu_k^{EK} = \psi_k$ and $\sigma_k^{EK} = 1 + \theta_k$. The tension between the *ToT* and *allocative efficiency* manifests itself in this model if $\text{cov}(\psi_k, \theta_k) < 0$. The fact that our theory readily extends to the Eaton-Kortum model echoes the isomorphism established in Kucheryavyy, Lyn, and Rodríguez-Clare (2023a). Online Appendix C shows that this isomorphism is even more profound. The nested-CES import demand function implied by (A1), we demonstrate, may analogously arise from within-industry specialization a la Eaton and Kortum.

The Melitz-Pareto Model.—Consider a multi-industry Melitz (2003) model that features the same nested-CES demand function specified by (A1). Suppose the firm-level productivity distribution is Pareto in each industry with a shape parameter, ϑ_k . Online Appendix D establishes that the Melitz-Pareto model is isomorphic to our baseline Krugman model insofar as macro-level representation is concerned. Hence, Theorem 1 characterizes the optimal policy in the Melitz-Pareto model under the following reinterpretation of parameters: $\mu_k^{Melitz} = \frac{\gamma_k \vartheta_k}{(\gamma_k - 1)(\vartheta_k + 1) - \vartheta_k} - 1$ if entry is restricted and $\mu_k^{Melitz} = 1/\vartheta_k$ if entry is free; and $\sigma_k^{Melitz} = 1 + \vartheta_k \left[1 + \vartheta_k \left(\frac{1}{\sigma_k - 1} - \frac{1}{\gamma_k - 1} \right) \right]^{-1}$. This mapping indicates that we need to estimate parameter ϑ_k , in addition to σ_k and γ_k , to quantify the gains from policy under firm-selection effects—a procedure we formally undertake and elaborate on in Section VI.

B. *Extension 1: Accounting for Input-Output Networks*

Suppose production employs both labor and intermediate inputs, which are distinguished from final goods by superscript \mathcal{I} . Cost minimization entails that the producer price of good ij, k (origin i –destination j –industry k) depends on the wage rate in origin i and the price of all intermediate inputs, $\tilde{\mathbf{P}}_i^{\mathcal{I}} \equiv \{\tilde{P}_{nj,k}^{\mathcal{I}}\}$, available to firms in origin i . Namely,

$$(13) \quad P_{ij,k} = \bar{\rho}_{ij,k} C_{i,k}(w_i, \tilde{\mathbf{P}}_i^{\mathcal{I}}) \mathcal{Q}_{i,k}^{-\frac{\mu_k}{1+\mu_k}},$$

where $C_{i,k}(\cdot)$ is a homogeneous of degree one function with respect to w_i and $\tilde{\mathbf{P}}_i^{\mathcal{I}}$.³⁸ The dependence of $P_{ij,k}$ on gross industry-level output, $\mathcal{Q}_{i,k} \equiv \sum_{j \in \mathbb{C}} (\bar{a}_{ij,k} \mathcal{Q}_{ij,k})$, reflects scale economies under free entry. The formal definition of general equilibrium under input-output (IO) linkages is presented in online Appendix K. The same Appendix characterizes optimal policy using an augmented version of our dual approach that leverages additional supply-side envelope conditions. Our characterization reveals that the formulas for optimal industrial subsidies and import tariffs are IO-blind—i.e., they are similar to Theorem 1. There is a simple intuition for this result: although a tariff on imported inputs can influence prices in the rest of the world through reexportation, any potential ToT gains from tariff reexportation are already internalized by the government’s optimal choice vis-à-vis export tax-cum-subsidies. Consistent with this intuition, the optimal export subsidy formula features an explicit adjustment for export subsidy reimportation via the IO network. The following theorem formalizes these results.

THEOREM 4: *Under IO linkages, the unilaterally first-best import tariffs and domestic subsidies are IO-blind; but first-best export subsidies exhibit an upward-adjustment that account for reimportation via the IO network. More formally, the unilaterally first-best policy schedule is given by*

$$\begin{aligned} \text{[domestic subsidy]} \quad & 1 + s_{i,k}^* = (1 + \mu_k)(1 + \bar{s}_i^{\mathcal{C}})^{-1}, \\ \text{[import tax]} \quad & 1 + t_{ji,k}^* = (1 + \omega_{ji,k})(1 + \bar{t}_i)(1 + \bar{s}_i^{\mathcal{C}}), \\ \text{[export subsidy]} \quad & \mathbf{1} + \mathbf{x}_{ij}^* = -\mathbf{E}_{ij}^{-1} \left[\mathbf{E}_{ij}^{(-ij)} (\mathbf{1} + \mathbf{t}_i^*) + \Lambda_{ij} (1 + \bar{t}_i) (1 + \bar{s}_i^{\mathcal{C}}) \right], \end{aligned}$$

where elements of $\Lambda_{ij} \equiv [\Lambda_{ij,k}]_k$ correspond to the fraction of good ij, k that is reimported and $\bar{s}_i^{\mathcal{C}}$ is an arbitrary tax shifter that assumes a positive value if the taxed item is a final good and zero otherwise.³⁹

³⁸ Without loss of generality, we assume that good ji, k can be used as either an intermediate input or a final consumption good, with taxes being applied on a good irrespective of the intended final use, i.e., $\tilde{P}_{ij,k}^{\mathcal{I}} = P_{ij,k}$. This assumption is innocuous because we can fragment every industry k into a final good version k' and an intermediate good version k'' . Since we do not impose any restrictions on the number of industries, our theory extends to the case where differential taxes are imposed on fragments k' and k'' .

³⁹ If country i is a small open economy, $\Lambda_{ij,k} \approx 0$. Correspondingly, optimal policy formulas for a small open economy under IO linkages perfectly overlap with the baseline formulas specified under equation (12).

The above theorem indicates that the tax equivalent of export subsidies is relatively lower on intermediate inputs to mitigate export tax-reimportation via the IO network.⁴⁰ Moreover, there is a uniform wedge between final and intermediate input taxes as represented by the final good tax shifter, \bar{s}_i^C .⁴¹ This detail aside, the ToT-improving motive for policy still requires export contraction in low- σ industries, while the misallocation-correcting objective asks for output expansion in high- μ industries. So, unless intermediate inputs exhibit a systemically lower σ , our conjecture about the inefficacy of stand-alone trade policy measures withstands.⁴²

C. Extension 2: Accounting for Political Economy Pressures

Suppose optimal policy choices internalize political economy pressures a la Ossa's (2014) adaptation of Grossman and Helpman (1994). That is, the government maximizes a politically weighted welfare function, $W_i \equiv V_i(w_i L_i + \mathcal{R}_i + \sum_k \pi_{i,k} \Pi_{i,k}, \tilde{\mathbf{P}}_i)$, where $\pi_{i,k}$ is the political economy weight assigned to industry k 's profits (with $\sum_k \pi_{i,k}/K = 1$). It follows trivially from Theorem 1 that the first-best policy in the political setup consists of the same trade tax/subsidy formulas but a politically adjusted industrial subsidy rate. Namely,

$$1 + s_{i,k}^* = (1 + \mu_{i,k}^P)(1 + \bar{s}_i),$$

where $\mu_{i,k}^P = \mu_k / [\pi_{i,k} - (1 - \pi_{i,k})\mu_k]$ is the *political economy-adjusted* markup of industry k . Considering the above formulas: if $\text{cov}(\pi_{i,k}, \mu_k) < 0$, the optimal policy may tax high- μ industries to the detriment of social welfare. In that case, even if $\text{cov}(\sigma_k, \mu_k) > 0$, the misallocation-correcting and ToT motives for trade taxation will clash. However, if $\text{cov}(\pi_{i,k}, \mu_k) \geq 0$, our conjecture about the inefficacy of stand-alone trade policy measures withstands.

V. Estimating the Key Policy Parameters

Based on our theory, policy evaluation in open economies requires credible estimates for industry-level *trade elasticities*, σ_k , and industry-level *scale elasticities*, $\mu_k \sim 1/(\gamma_k - 1)$. The former governs the degree of national-level market power, while the latter regulates the extent of love for variety. The trade literature has paid considerable attention to estimating σ_k but less to estimating μ_k . The existing policy literature typically normalizes μ_k in one of two ways: (i) $\mu_k = 1/(\sigma_k - 1)$ in

⁴⁰In the absence of scale economies or profits, $\omega_{j,i,k} = 0$. So, Theorem 1 posits that optimal import tariffs are uniform by choice of $\bar{s}_i^C = 0$ in the multi-industry Armington and Eaton-Kortum models (see online Appendix M).

⁴¹Theorem 4 indicates that uniform markup pricing is not a necessary condition for efficiency. Consider, for instance, a vertical production economy where goods are used as either final goods or intermediate inputs but not both. The efficient policy, in this setup, must restore uniform markup pricing within input and final good segments but not across—see Antràs et al. (2022) for further exploration of this issue.

⁴²Theorem 4 offers insight into the structure of second-best import tariffs and export subsidies under IO linkages. Several papers, including Blanchard, Bown, and Johnson (2016); Beshkar and Lashkaripour (2020); Caliendo et al. (2021); and Antràs et al. (2022), examine in more detail how IO linkages impact third-best import tariff choices. Beshkar and Lashkaripour (2020), moreover, adopt a special case of Theorem 4, with Cobb-Douglas production and no scale economies, to examine the cost of trade wars.

imperfectly competitive models, and (ii) $\mu_k = 0$ in perfectly competitive models.⁴³ Both normalizations impose artificial restrictions on $\text{cov}(\mu_k, \sigma_k)$, which following the discussion in Section III, can bias the estimated gains from policy.

Against this backdrop, we seek to estimate σ_k and μ_k in a way that ascertains mutual consistency, giving us a credible assessment of $\text{cov}(\mu_k, \sigma_k)$. To this end, we propose a new methodology that simultaneously estimates σ_k and μ_k from the same data.⁴⁴ Our approach involves fitting the structural *firm-level* import demand function implied by (A1) to the universe of Colombian import transactions from 2007 to 2013. We outline this approach below, starting with a description of the data used in our estimation.

Data Description.—Our estimation uses data on import transactions from the Colombian Customs Office during 2007–2013.⁴⁵ The data include detailed information about each transaction, such as the Harmonized System 10-digit product category (HS10), country of origin, importing and exporting firm IDs, quantity, f.o.b. (free on board), and c.i.f. (customs, insurance, and freight) transaction values, freight, insurance, and value-added tax in US dollars. As a unique feature, our data report the identities of all foreign firms exporting to Colombia. This feature allows us to define import varieties as firm-product combinations rather than country-product combinations, which is the standard approach. Table N.1 (in the online Appendix) reports a summary of basic trade statistics in our data. Working with firm-level data presents two challenges. First, exporters are not identified by unique standardized IDs. Instead, they are identified by a name, a number, and an address. We handle this problem by standardizing the spelling and lengths of firms' names and using the information on firms' phone numbers (see online Appendix N). Second, Colombia changed the HS10 classification for some products between 2007 and 2013. Fortunately, the Colombian Statistical Agency, DANE, keeps track of these changes. We utilize this information to concord the Colombian HS10 codes over time, using the guidelines in Pierce and Schott (2012). Overall, changes in HS10 codes between 2007 and 2013 affect less than 0.1 percent of our data points.

A. Estimating Equation

Since we are focusing on one importer, we hereafter drop the importer's subscript i and add a year subscript t to account for the time dimension of our data. With this switch in notation, the demand facing firm ω located in country j and supplying product k in year t is given by

$$(14) \quad q_{j,kt}(\omega) = \varphi_{j,kt}(\omega) \left[\frac{\tilde{P}_{j,kt}(\omega)}{\tilde{P}_{j,kt}} \right]^{-\gamma_k} \left(\frac{\tilde{P}_{j,kt}}{\tilde{P}_{kt}} \right)^{-\sigma_k} Q_{kt}$$

⁴³ See Ossa (2016) and Costinot and Rodríguez-Clare (2014) for a synthesis of the previous literature. Under restricted entry, μ_k denotes the firm-level markup, which can be alternatively estimated with firm-level production data (e.g., De Loecker and Warzynski 2012; De Loecker et al. 2016).

⁴⁴ In the presence of firm-selection effects, our estimated parameters are necessary but not sufficient to pin down the trade and scale elasticities—see online Appendixes D and Y for details.

⁴⁵ The data are obtained from Datamyne, a company that specializes in documenting import and export transactions in the Americas. For more detail, please see www.datamyne.com. Our estimation also employs data on monthly average exchange rates, which are taken from the Bank of Canada (2017). A detailed description of all the datasets used in this paper is provided in online Appendix N.

Subscript k , in our theoretical model, designated industries. In our estimation, k denotes an HS10 product—the most disaggregated product classification in our data. The quadruplet “ ωjkt ” accordingly denotes a unique variety corresponding to *firm* ω –*country of origin* j –*HS10 product* k –*year* t . Let $\tilde{x}(\omega) \equiv \tilde{p}(\omega)q(\omega)$ denote gross sales. Rearranging equation (14) yields the following log-linear import demand function facing individual varieties:

$$(15) \quad \ln \tilde{x}_{j,kt}(\omega) = (1 - \sigma_k) \ln \tilde{p}_{j,kt}(\omega) + \left(1 - \frac{\sigma_k - 1}{\gamma_k - 1}\right) \ln \lambda_{j,kt}(\omega) \\ + D_{kt} + \ln \varphi_{j,kt}(\omega),$$

where $D_{kt} \equiv \ln P_{kt}^{\sigma_k} Q_{kt}$ can be treated as a product-year fixed effect and $\lambda_{j,kt}(\omega)$ denotes the share of expenditure on firm ω conditional on buying good k from country of origin j ,

$$\lambda_{j,kt}(\omega) \equiv \varphi_{j,kt}(\omega) \left[\frac{\tilde{p}_{j,kt}(\omega)}{\tilde{P}_{j,kt}} \right]^{1-\gamma_k} = \frac{\tilde{x}_{j,kt}(\omega)}{\sum_{\omega' \in \Omega_{j,kt}} \tilde{x}_{j,kt}(\omega')}.$$

We assume that $\varphi_{jkt}(\omega) = \bar{\varphi}_{j,k}(\omega) \times \varphi_{\omega jkt}$ can be decomposed into a time-invariant *firm-and-product-specific* quality component, $\bar{\varphi}_{j,k}(\omega)$, and a time-varying component $\varphi_{\omega jkt}$ that encompasses idiosyncratic variations in consumer taste, measurement errors, and/or omitted variables that account for dynamic demand optimization. To eliminate $\bar{\varphi}_{j,k}(\omega)$ from the estimating equation, we employ a first-difference estimator, which also drops observations pertaining to one-time exporters. We deem the first-difference estimator appropriate given the possibility that $\varphi_{\omega jkt}$'s are sequentially correlated. As a robustness check, we also report estimation results based on a two-way fixed effects estimator in online Appendix Q.⁴⁶ Stated in terms of first-differences, our estimating equation takes the following form:

$$(16) \quad \Delta \ln \tilde{x}_{j,kt}(\omega) = (1 - \sigma_k) \Delta \ln \tilde{p}_{j,kt}(\omega) + \left(1 - \frac{\sigma_k - 1}{\gamma_k - 1}\right) \Delta \ln \lambda_{j,kt}(\omega) \\ + \Delta D_{kt} + \Delta \ln \varphi_{\omega jkt},$$

where $\Delta \ln \varphi_{\omega jkt}$, roughly speaking, represents a variety-specific demand shock and ΔD_{kt} is a product-year fixed effect.⁴⁷ Of the remaining variables, $\Delta \ln \tilde{p}_{j,kt}(\omega)$ and $\Delta \ln \tilde{x}_{j,kt}(\omega)$ are directly observable for each import variety. The change in the *within-national* market share, $\Delta \ln \lambda_{j,kt}(\omega)$, can be calculated using the universe of firm-level sales to Colombia.

⁴⁶ Following Boehm, Levchenko, and Pandalai-Nayar (2020), the first-difference estimation offers a partial remedy for omitted variable bias and reverse causality. Both issues pose a serious challenge to traditional log-level estimations of import demand. Depending on the application, though, the first-difference estimator may not necessarily identify the desired long-run elasticity. As detailed in online Appendix Q, this limitation is less severe in our firm-level estimation—as we explicitly control for the extensive margin of trade. We illustrate this point formally in online Appendix Q by reestimating equation (15) in levels and comparing the estimation results to the baseline values.

⁴⁷ To handle outliers, we trim our sample to exclude observations that report a price change, $\Delta \ln \tilde{p}_{j,kt}(\omega)$, above the ninety-ninth percentile or below the first percentile of the HS10 product code k in year t .

Recovering Scale and Trade Elasticities from Demand Parameters.—Equation (16) allows us to estimate demand parameters, σ_k and γ_k , from which we can recover the scale and trade elasticities as follows (see Section I for the underlying theoretical foundation):

$$\mu_k = \frac{1}{\gamma_k - 1} \sim \text{scale elasticity}, \quad \sigma_k \sim \text{trade elasticity}.$$

The reason we can infer μ_k from demand parameters is that the scale elasticity in the generalized Krugman model reflects the extent of *love for variety*—the social benefits of which are not internalized by firms' entry decisions. In the Melitz-Pareto case, we also need estimates for the shape of the Pareto productivity distribution (in addition to γ_k and σ_k) to recover the scale elasticity—see online Appendix P. Our estimation is, of course, unable to detect scale externalities unrelated to love for variety. These externalities can be estimated using techniques that leverage scale-related moments but under stronger assumptions about the variability of production inputs. We discuss the relative advantages of each technique in online Appendix R.⁴⁸

Breaking the Sample into Broadly Defined Industries.—As noted earlier, k indexes an HS10 product category in equation (16). To conduct our forthcoming quantitative analysis, we must estimate demand parameters for 14 broadly defined industries based on the World Input-Output Database (WIOD) industry classification. Considering this, we pool all HS10 products belonging to the same WIOD industry \mathcal{K} together and estimate equation (16) on this pooled sample assuming that σ_k and γ_k are uniform across products within the same industry (i.e., $\gamma_k = \gamma_{\mathcal{K}}$ and $\sigma_k = \sigma_{\mathcal{K}}$ for all $k \in \mathbb{K}_{\mathcal{K}}$). In principle, we can also estimate the import demand function separately for each HS10 product category to attain HS10-level elasticities. However, such elasticities will be of little use for quantitative policy analysis, as multicountry data on trade, production, and expenditure shares are scarce at such levels of disaggregation.

B. Identification Strategy

The identification challenge we face is that $\Delta \ln \tilde{p}_{j,kt}(\omega)$ and $\Delta \ln \lambda_{j,kt}(\omega)$ are endogenous variables that can covary with the demand shock, $\Delta \ln \varphi_{\omega jkt}$.⁴⁹ Traditional *country-level* import demand estimations overcome a similar challenge by instrumenting for prices with plausibly exogenous tariff rates.⁵⁰ This strategy, however, does not suit our *firm-level* estimation because tariffs discriminate by country of origin but not across firms from the same country.

⁴⁸In the restricted entry case, our demand estimation identifies the firm-level markup in each industry. Following an old tradition in the industrial organization literature, we assume that market conduct is monopolistic competition and recover firm-level markups as $\mu_k = 1/(\gamma_k - 1)$. Online Appendix W recovers markups under alternative market conduct assumptions.

⁴⁹Another challenge is that unit price data may be contaminated with measurement errors, as they are averaged across many transactions. Following Berry (1994), this type of measurement error is fairly innocuous when dealing with log-linear demand functions. Furthermore, our instrumental variable approach will handle measurement errors, provided that lagged monthly sales patterns are uncorrelated with concurrent measurement errors.

⁵⁰A prominent example is Caliendo and Parro (2015), who use tariff data to identify the trade elasticity.

We employ a shift-share research design to overcome our *firm-level* identification challenge. Our approach is rooted in two data observations. First, a typical variety is imported under multiple invoices spread across different months in a given year. As a matter of accounting, the annual-level price of a variety is the quantity-weighted average of monthly prices. Namely, $\tilde{p}_{j,kt}(\omega) = \sum_{m \in \mathcal{M}} \tilde{s}_{j,kt}(\omega; m) \tilde{p}_{j,kt}(\omega; m)$, where m denotes month and $\tilde{s}(\omega; m)$ and $\tilde{p}(\omega; m)$ denote the quantity share and price associated with month m . Second, the month m price of an imported variety (in Colombian pesos) is equal to *markup-plus-taxes* \times *marginal input cost invoiced in local currency* \times *exchange rate in month m* . More formally, $\tilde{p}_{j,kt}(\omega; m) = \tau_{j,kt}(\omega) \times C_{j,kt}(\omega) \times \mathcal{E}_{jt}(m)$, where τ and C respectively denote markup-plus-tax and marginal cost, while $\mathcal{E}_{jt}(m)$ represents the exchange rate between origin j 's currency and the Colombian peso in month m of year t . To a first-order approximation, the annual change in variety-level prices in response to monthly exchange rate fluctuations can be specified as

$$\Delta \ln \tilde{p}_{j,kt}(\omega) \approx \sum_{m \in \mathcal{M}} s_{j,kt}(\omega; m) \Delta \ln \mathcal{E}_{jt}(m),$$

where $\Delta \ln \mathcal{E}_{jt}(m)$ denotes the year-over-year change in origin j 's exchange rate with the Colombian peso in month m and $s_{j,kt}(\omega; m)$ is the share of month m in variety ωjkt 's annual export sales to Colombia:

$$s_{j,kt}(\omega; m) = \frac{\tilde{x}_{j,kt}(\omega; m)}{\sum_{m' \in \mathcal{M}} \tilde{x}_{j,kt}(\omega; m')} \sim \text{share of month } m \text{ in annual export sales.}$$

Capitalizing on the above observation, we construct our shift-share instrument as the inner product of *lagged* monthly export shares and monthly exchange rate changes. Namely,

$$z_{j,kt}(\omega) = \sum_{m \in \mathcal{M}} s_{j,kt-1}(\omega, m) \Delta \ln \mathcal{E}_{jt}(m).$$

Stated verbally, $z_{j,kt}(\omega)$ measures exposure to exchange rate fluctuations at the *firm* \times *origin* \times *product* \times *year* level. The idea being that aggregate exchange rate movements have differential effects on individual firms depending on the monthly composition of their prior export activity to Colombia.⁵¹ Encouragingly, the validity of our instrument is corroborated by the strong and statistically significant correlation between z and $\Delta \ln \tilde{p}$. Online Appendix O illustrates this relationship visually using two US-based firms as an example.

Exclusion Restriction.—Our instrument utilizes lagged export shares, which depend on lagged prices and a set of market-level indexes—namely, $s_{jkt-1}(\omega, m) = s(\tilde{p}_{jkt-1}(\omega, m); \dots)$. Considering this, the exclusion restriction in our setup, $E[z \Delta \ln \varphi] = 0$, rests on two conditions:

- (C1) Prior price-setting decisions (and thus lagged export shares) are orthogonal to concurrent demand shocks: $E[\tilde{p}_{j,kt-1}(\omega) \Delta \ln \varphi_{\omega jkt}] = 0$.

⁵¹ Leveraging exchange rate data to construct firm-level input cost shifters has also been explored by Piveteau and Smagghue (2019) and Amiti, Itskhoki, and Konings (2019).

(C2) Monthly national-level exchange rate shocks are orthogonal to variety-level demand shocks: $E[\Delta \ln \mathcal{E}_{jt}(m) \Delta \ln \varphi_{\omega jkt}] = 0$.

Since our sample features many firms and a finite number of months, Condition (C1) is sufficient for the consistency of our estimates (see Proposition 2.1 in Goldsmith-Pinkham, Sorkin, and Swift 2020). In fact, our two-stage least squares (2SLS) estimator is numerically equivalent to a generalized method of moments (GMM) estimator with lagged monthly export shares as instruments and a weight matrix constructed from monthly aggregate exchange rate shocks. Condition (C2), meanwhile, is more crucial for the finite sample properties of our estimator. Both conditions can, in principle, be violated if there are cross-inventory linkages or if individual export varieties account for a significant fraction of national exports to Colombia. We discuss and address these issues in Section D.

Instruments for $\Delta \ln \lambda_{j,kt}(\omega)$.—Following Khandelwal (2010), we construct two standard instruments for the annual variation in the within-national market shares: (i) annual changes in the total number of origin j firms serving the Colombian market in product category k and (ii) changes in the total number of HS10 product categories actively served by firm ω in year t . These count measures will be correlated with $\Delta \ln \lambda_{j,kt}(\omega)$ but uncorrelated with $\Delta \ln \varphi_{\omega jkt}$ if variety-level entry and exit occurs prior to, or independent of, the demand shock realization of competing varieties. As noted by Khandelwal (2010), this assumption is widely invoked when estimating discrete choice demands curves—see also Berry, Levinsohn, and Pakes (1995).⁵²

C. Estimation Results

Table 3 reports our industry-level estimation results. We also report results corresponding to a pooled sample of all industries in Table S.1 of the online Appendix. The same table compares the 2SLS and OLS estimates to ensure that our IV strategy operates in the expected direction. Our estimates point to a median trade elasticity of $\sigma - 1 = 3.9$ and a median scale elasticity of $\mu \approx 0.20$. Our pooled estimation yields a heteroskedasticity-robust Kleibergen-Paap Wald rk F -statistic of 259, rejecting the null of weak instruments given the Stock-Yogo critical values. A similar albeit weaker outcome emerges from the industry-level estimation.

The industry-level elasticities reported in Table 3 display considerable variation across industries. The estimated scale elasticity or markup margin is highest in the “Electrical and Optical Equipment” ($\mu = 0.55$) and “Petroleum” ($\mu = 1.2$) industries; both of which are associated with high R&D or fixed costs. The estimated scale elasticity is lowest in “Agricultural and Mining” ($\mu = 0.14$) and “Machinery” ($\mu = 0.12$) industries. Furthermore, with the exception of “Agriculture and Mining,” we cannot reject the prevalence of scale economies due to *love for variety*.⁵³

⁵²Border taxes tend to be a weak instrument for firm-level prices in our sample, but we include to comply with the past literature. These include applied ad valorem tariffs and the Colombian value-added tax (VAT). We exclude the VAT component in the “Transportation” and “Petroleum” industries since the VAT in these industries discriminates by the delivery method and level of luxury—both of which may be correlated with $\Delta \ln \varphi_{\omega jkt}$.

⁵³The finding that returns-to-scale are negligible in the agricultural sector aligns with a large body of evidence on the inverse farm-size productivity (IFSP) relationship—see Sen (1962) and subsequent references to IFSP.

TABLE 3—INDUSTRY-LEVEL ESTIMATION RESULTS

Sector	ISIC4 codes	Estimated parameter			Obs.	Weak Ident. Test
		$\sigma_k - 1$	$\frac{\sigma_k - 1}{\gamma_k - 1}$	μ_k		
Agriculture and Mining	100-1499	6.227 (2.345)	0.891 (0.148)	0.143 (0.059)	11,568	2.40
Food	1500-1699	2.303 (0.765)	0.905 (0.046)	0.393 (0.132)	19,615	6.27
Textiles, Leather and Footwear	1700-1999	3.359 (0.353)	0.753 (0.022)	0.224 (0.024)	125,120	66.65
Wood	2000-2099	3.896 (1.855)	0.891 (0.195)	0.229 (0.120)	5,872	1.41
Paper	2100-2299	2.646 (1.106)	0.848 (0.061)	0.320 (0.136)	37,376	3.23
Petroleum	2300-2399	0.636 (0.464)	0.776 (0.119)	1.220 (0.909)	3,973	2.83
Chemicals	2400-2499	3.966 (0.403)	0.921 (0.025)	0.232 (0.024)	133,142	38.01
Rubber and Plastic	2500-2599	5.157 (1.176)	0.721 (0.062)	0.140 (0.034)	106,398	7.16
Minerals	2600-2699	5.283 (1.667)	0.881 (0.108)	0.167 (0.056)	27,952	3.53
Basic and Fabricated Metals	2700-2899	3.004 (0.484)	0.627 (0.030)	0.209 (0.035)	153,102	20.39
Machinery	2900-3099	7.750 (1.330)	0.927 (0.072)	0.120 (0.023)	263,797	12.01
Electrical and Optical Equipment	3100-3399	1.235 (0.323)	0.682 (0.017)	0.552 (0.145)	257,775	26.27
Transport Equipment	3400-3599	2.805 (0.834)	0.363 (0.036)	0.129 (0.041)	85,920	5.50
N.E.C. and Recycling	3600-3800	6.169 (1.012)	0.938 (0.090)	0.152 (0.029)	70,264	11.57

Notes: Estimation results of equation (16). Standard errors in parentheses. The estimation is conducted with HS10 product-year fixed effects. All standard errors are simultaneously clustered by product-year and by origin-product, which is akin to the correction proposed by Adao, Kolesár, and Morales (2019). The weak identification test statistics are the F -statistics from the Kleibergen-Paap Wald test for weak identification of all instrumented variables. The test for overidentification is not reported due to the pitfalls of the standard overidentification Sargan-Hansen J -test in the multidimensional large datasets pointed by Angrist, Imbens, and Rubin (1996).

Importantly, our estimates indicate that $(\sigma_k - 1)/(\gamma_k - 1) \neq 1$ in nearly all industries. This finding rejects the arbitrary link often assumed between the firm-level and national-level degrees of market power in the literature. Our estimates also indicate that

$$\text{cov}(\sigma_k, \mu_k) \approx -0.65,$$

which corroborates the innate tension between the *ToT-improving* and *misallocation-correcting* objectives at the core of Conjectures 1 and 2.⁵⁴

⁵⁴Our trade elasticity estimates are one of the few based on firm-level data. Traditional estimates of σ_k are typically based on country-level data (e.g., Simonovska and Waugh 2014; Caliendo and Parro 2015).

D. Challenges to Identification

Our two conditions for identification, (C1) and (C2), can be contested under certain circumstances. Below, we discuss these issues and present additional evidence to address them.

Within-Cluster Correlation in Error Terms.—Adao, Kolesár, and Morales (2019) show that identification based on shift-share instruments exhibits an overrejection problem if regression errors are cross-correlated. In the context of our estimation, this problem will arise if demand shocks are correlated across *firm-origin-product-year* varieties with a similar monthly export composition. We adopt a conservative two-way clustering of standard errors by *product-year* and *origin-product* to handle this issue. Clustering standard errors in this manner is akin to the correction proposed by Adao, Kolesár, and Morales (2019).

Dynamic Cross-Inventory Effects.—Lagged inventory clearances can challenge our identifying assumptions on two fronts. First, firms' optimal pricing decisions may be forward-looking, violating Condition (C1). To address this concern, we reconstruct our shift-share instrument using four lags instead of one. If inventories clear in at most four years, we can deduce that pricing decisions do not internalize expected demand shocks beyond the four-year mark. Hence, $E[\tilde{p}_{jkt-4}(\omega)\Delta\ln\varphi_{\omega jkt}] = 0$, and the new instrument will satisfy the exclusion restriction. The trade-off is that we lose observations, as the instrument is constructible for only firms that continuously export in the four different years. The *top* panel of Figure P.1 (in online Appendix P) compares estimation results under this alternative instrument to the baseline results. The new estimation preserves the ordering and magnitude of our estimated elasticities, retaining the sign of $\text{cov}(\sigma_k, \mu_k)$, which is pivotal for optimal policy outcomes.

Second, with cross-inventory linkages, $\Delta\ln\varphi_{\omega jkt}$ may encompass omitted variables that reflect firms' dynamic inventory management decisions. One of these omitted variables is presumably the exchange rate. If so, $E[\Delta\ln\mathcal{E}_{jt}(m)\Delta\ln\varphi_{\omega jkt}] \neq 0$, which violates Condition (C2). To address this concern, we reestimate equation (16) while directly controlling for the annual change in the exchange rate, $\Delta\ln\mathcal{E}_{jt}$. Even if changes in inventory-related demand depend on the changes in the exchange rate, we can still assert that $E[z_{j,kt}(\omega)\Delta\ln\varphi_{\omega jkt}|\Delta\ln\mathcal{E}_{jt}] = 0$ —i.e., the exclusions restriction is satisfied with the added control, $\Delta\ln\mathcal{E}_{jt}$. The *middle* panel of Figure P.1 (in online Appendix P) compares estimation results from this alternative specification to the baseline results. Reassuringly, the new estimation preserves the ordering and magnitude of our estimated elasticities and the negative correlation between σ_k and μ_k .

Export Varieties with Significant Market Share.—Our identification can come under threat if individual varieties account for a significant fraction of a country's sales to Colombia. In such a case, variety-specific demand shocks can influence the bilateral exchange between the Colombian peso and the origin country's currency, thereby violating Condition (C2). This concern, however, does not apply to our sample of exporters. The variety with the highest ninety-ninth percentile within-national

market share accounts for only 0.1 percent of the origin country's total exports to Colombia. The variety with the highest ninetieth percentile within-national market share accounts for only 0.0008 percent of the origin country's total exports to Colombia.

One may remain concerned about large multiproduct firms that export multiple product varieties to Colombia in a given year. Consider, for instance, a multiproduct firm ω that exports goods k and g to Colombia in year t . If demand shocks are correlated across varieties supplied by this firm (i.e., $E[\Delta \ln \varphi_{\omega jkt} \Delta \ln \varphi_{\omega jgt}] \neq 0$), Condition (C2) may be violated despite each variety's market share being infinitesimally small. To address this issue, we reestimate equation (16) on a restricted sample that drops excessively large firms with a total within-national market share that exceeds 0.1 percent. The *bottom* panel of Figure P.1 (in online Appendix P) compares estimation results from the trimmed sample to the baseline results. Encouragingly, the ordering and magnitude of the estimated elasticities are preserved across industries. The new estimation also retains the negative correlation between σ_k and μ_k .

E. Plausibility of Estimates and Limitations

We conclude this section by discussing the plausibility of our estimates. We do so by exploring their macro-level implications and comparing them to counterparts in the literature.

Plausibility from the Lens of Macro-Level Predictions.—Our scale elasticity estimates can be evaluated based on their prediction about the income-size elasticity. As pointed out by Ramondo, Rodríguez-Clare, and Saborío-Rodríguez (2016), the factual relationship between real per capita income and population size (i.e., the income-size elasticity) is negative and statistically insignificant. Quantitative trade models featuring the normalization $\mu(\sigma - 1) = (\sigma - 1)/(\gamma - 1) = 1$, however, predict a strong and positive income-size elasticity that remains significant even after the introduction of domestic trade frictions. Ramondo, Rodríguez-Clare, and Saborío-Rodríguez (2016) call this observation the income-size elasticity puzzle. Considering this puzzle, in online Appendix S we compute the income-size elasticity implied by our estimated value of $(\sigma - 1)/(\gamma - 1) \approx 0.67$. Encouragingly, our estimated value for $(\sigma - 1)/(\gamma - 1)$ completely resolves the aforementioned puzzle. In other words, our micro-estimated elasticities are consistent with the macro-level cross-national relationship between population size and real per capita income.

Comparison to Counterparts in the Literature.—Reassuringly, our estimates align closely with well-known industry-level case studies. Take, for example, our elasticity estimates for the “Petroleum” industry, which appear somewhat extreme. First, our estimate for σ_k aligns with the consensus in the Energy Economics literature that national-level demand for petroleum products is price-inelastic.⁵⁵ Second, our estimated μ_k for the “Petroleum” industry closely resembles existing estimates

⁵⁵See Pesaran, Smith, and Akiyama (1998) for specific estimates and Fattouh (2007) for a survey of this literature.

in the Industrial Organization literature. Considine (2001), for instance, estimates $\mu \approx 1.15$ using detailed data on the US petroleum industry. Moreover, our finding that the “Petroleum” industry is the most scale-intensive industry is consistent with the finding in Antweiler and Trebler (2002), which is based on more aggregated data. Likewise, consider the “Transportation” or auto industry, where our estimated $\mu_k = 0.13$ implies an optimal markup of 13 percent. This estimate aligns with existing estimates from various industry-level studies. Recently, Coşar et al. (2018) have estimated markups for the auto industry that range between roughly 6 percent to 13 percent. Previously, Berry, Levinsohn, and Pakes (1995) have estimated markups of around 20 percent in the US auto industry using data from 1971 to 1990.

Limitations of Our Estimation Technique.—A potential limitation of our estimation technique is that we do not directly leverage scale-related moments in our estimation. The scale elasticity can be alternatively estimated using *direct* estimation techniques that leverage such moments, as reviewed in online Appendix R. These techniques have the advantage of detecting scale externalities unrelated to *love for variety*—albeit under additional assumptions about the variability of production inputs. While we are mindful of this limitation, we think our approach may be appropriate for policy evaluation in certain settings, as it separately identifies the trade elasticity from the scale elasticity, delivering credible estimates for $\text{COV}(\sigma_k, \mu_k)$.

VI. Quantifying the Consequences of Trade and Industrial Policy

This section uses our estimated values for μ_k and σ_k to quantify the ex ante gains from trade and industrial policy among many countries. We begin by describing the macro-level data needed to calibrate our quantitative model.

Trade, Production, and Tariff Data.—We take macro-level data on domestic and international production and expenditure from the 2014 World Input-Output Database (Timmer et al. 2015). This database spans 56 industries and 43 countries plus an aggregate of the rest of the world. The list of countries in the sample includes all 27 members of the European Union plus 16 other major economies—all of which are listed in Table 4. Following Costinot and Rodríguez-Clare (2014), we aggregate the 56 WIOD industries into 15 traded industries (for which we have estimated μ_k and σ_k) plus a service sector. Details for our industry aggregation are reported in Table U.1 of the online Appendix. Our baseline analysis normalizes $\mu_k = 0$ and $\sigma_k = 11$ for all service-related industries. In online Appendix Y, however, we test the sensitivity of our results to alternative normalization choices. We also need to take a stance on applied tariffs and subsidies. We construct data on applied tariffs, $t_{j,k}$, following the guidelines in Kucheryavy, Lyn, and Rodríguez-Clare (2023b).⁵⁶ International data on domestic and export subsidies are not as widely available.

⁵⁶We follow Kucheryavy, Lyn, and Rodríguez-Clare (2023b) to download the data from the UNCTAD Trade Analysis and Information System (TRAINS) and match it with WIOD. To make the data consistent with our theoretical model, we also purge them from trade imbalances following the procedure described in Costinot and Rodríguez-Clare (2014).

TABLE 4—THE GAINS FROM NONCOOPERATIVE POLICIES AND THE CONSEQUENCES OF RETALIATION

Country	Restricted entry				Free entry			
	1st-best	2nd-best trade tax	3rd-best import tax	Post-retaliation	1st-best	2nd-best trade tax	3rd-best import tax	Post-retaliation
AUS	0.90%	0.21%	0.14%	-0.83%	2.30%	0.58%	0.35%	2.85%
AUT	1.31%	0.66%	0.45%	-1.67%	2.12%	1.11%	0.58%	-2.34%
BEL	1.31%	0.69%	0.50%	-3.07%	1.66%	0.90%	0.55%	-4.13%
BGR	2.01%	0.63%	0.52%	0.04%	5.48%	1.88%	0.83%	0.23%
BRA	1.87%	0.23%	0.17%	0.93%	4.31%	0.64%	0.36%	3.21%
CAN	1.72%	0.57%	0.45%	-1.10%	3.42%	1.19%	0.46%	0.34%
CHE	1.04%	0.56%	0.46%	-0.92%	1.28%	0.76%	0.53%	-1.26%
CHN	1.76%	0.27%	0.24%	1.43%	3.70%	0.39%	0.28%	2.90%
CYP	1.75%	0.68%	0.61%	-4.20%	5.83%	1.56%	1.43%	-10.88%
CZE	1.68%	0.88%	0.54%	-1.70%	2.86%	1.60%	0.73%	-2.13%
DEU	1.71%	0.78%	0.49%	-0.37%	2.60%	1.34%	0.65%	-0.10%
DNK	1.22%	0.58%	0.48%	-2.41%	1.94%	0.91%	0.49%	-4.78%
ESP	1.53%	0.52%	0.39%	-0.15%	2.51%	1.03%	0.49%	0.29%
EST	1.18%	0.62%	0.45%	-4.86%	2.69%	1.42%	0.56%	-7.90%
FIN	1.40%	0.54%	0.25%	-0.76%	2.00%	0.83%	0.47%	-1.08%
FRA	1.20%	0.45%	0.34%	-0.78%	2.52%	1.10%	0.50%	0.01%
GBR	1.09%	0.48%	0.42%	-0.65%	2.10%	1.03%	0.58%	0.17%
GRC	1.82%	0.59%	0.53%	-0.20%	2.75%	1.08%	0.68%	-0.04%
HRV	1.03%	0.55%	0.45%	-2.27%	1.79%	0.73%	0.52%	-2.99%
HUN	2.25%	1.06%	0.73%	-2.85%	4.08%	2.28%	0.96%	-3.33%
IDN	2.00%	0.36%	0.26%	-0.37%	4.81%	1.45%	0.50%	2.42%
IND	1.81%	0.35%	0.31%	1.60%	4.20%	1.10%	0.37%	2.82%
IRL	0.86%	0.68%	0.53%	-2.07%	1.49%	0.89%	0.40%	-3.32%
ITA	1.50%	0.46%	0.26%	0.36%	2.75%	0.95%	0.48%	0.80%
JPN	1.48%	0.32%	0.22%	0.31%	2.97%	0.75%	0.42%	1.55%
KOR	2.12%	0.65%	0.49%	0.14%	4.37%	1.62%	0.73%	1.25%
LTU	2.50%	0.93%	0.76%	-1.38%	3.54%	1.26%	0.85%	-2.47%
LUX	0.93%	0.81%	0.78%	-3.59%	1.52%	1.15%	1.01%	-4.90%
LVA	0.91%	0.54%	0.44%	-4.26%	1.57%	0.80%	0.46%	-4.85%
MEX	2.24%	0.60%	0.44%	-0.92%	4.96%	1.36%	0.76%	1.96%
MLT	1.35%	0.92%	0.82%	-3.75%	2.08%	1.39%	1.06%	-4.81%
NLD	1.35%	0.66%	0.54%	-3.44%	1.95%	0.98%	0.59%	-3.73%
NOR	1.19%	0.40%	0.27%	-0.74%	2.02%	0.81%	0.41%	-0.33%
POL	2.19%	0.76%	0.65%	-0.26%	4.03%	1.62%	0.81%	0.80%
PRT	2.04%	0.74%	0.65%	-0.36%	3.92%	1.72%	0.80%	0.73%
ROU	2.05%	0.77%	0.67%	-0.42%	3.92%	1.72%	0.98%	1.15%
RUS	2.39%	0.32%	0.27%	0.84%	5.25%	1.39%	0.37%	2.80%
SVK	2.01%	1.06%	0.79%	-2.17%	3.50%	2.09%	1.08%	-2.73%
SVN	1.42%	0.87%	0.67%	-3.09%	1.35%	1.20%	0.90%	-4.92%
SWE	1.21%	0.62%	0.45%	-0.95%	1.54%	0.77%	0.49%	-1.48%
TUR	1.43%	0.46%	0.32%	-1.52%	3.52%	1.34%	0.61%	-6.33%
TWN	2.18%	0.69%	0.56%	-0.94%	4.92%	1.85%	0.79%	0.81%
USA	1.53%	0.32%	0.27%	0.69%	3.04%	0.80%	0.30%	2.15%
Average	1.59%	0.60%	0.47%	-1.23%	3.05%	1.19%	0.63%	-1.20%

Notes: The data source is the 2014 World Input-Output Database (WIOD) (Timmer et al. 2015; WIOD 2021). The first-best policy is characterized by Theorem 1; second-best trade taxes are characterized by Theorem 2; and third-best import taxes are characterized by Theorem 3. Post-retaliation corresponds to a situation where the home country sets its first-best policies and the rest of the world retaliates.

Considering that these subsidy measures are generally prohibited by the WTO, we extrapolate that $x_{ij,k} \approx s_{i,k} \approx 0$.

A. Mapping Optimal Policy Formulas to Data

The formulas provided by Theorems 1–3 allow us to compute the *maximal* gains from policy without resorting to numerical optimization. This approach can be useful considering that numerical optimization routines (like Matlab’s `fmincon`) have well-known limitations when applied to nonlinear models with many free-moving variables. Our *optimization-free* procedure (a) specifies optimal tax/subsidy choices as a function of equilibrium variables (e.g., expenditure shares) using the formulas

provided by Theorems 1–3 and (b) specifies equilibrium variables as a function of optimal tax/subsidy choices based on equilibrium constraints. The system of equations implied by (a) and (b) is then jointly solved to compute welfare outcomes under optimal policy. Our method utilizes the exact hat-algebra technique whereby $\hat{z} = z^*/z$ denotes the counterfactual change in a generic variable after implementing optimal policies. Our presentation here is focused on first-best policies under free entry. Online Appendix T demonstrates how the same approach can be applied to other optimal policy scenarios.

To map our theory to data, we must take a stance on the cross-industry utility aggregator, which we assume is Cobb-Douglas—i.e., $U_i(\mathbf{Q}_i) = \prod_k Q_{i,k}^{\epsilon_{i,k}}$. Under this parameterization, we need data on observable shares, national accounts, and applied taxes, which we denote by $\mathbb{D} = \{\lambda_{ji,k}, r_{ji,k}, e_{i,k}, \rho_{i,k}, Y_i, W_i L_i, x_{ij,k}, t_{ji,k}, s_{i,k}\}_{j,i,k}$.⁵⁷ We also need estimates for the industry-level scale and trade elasticities, $\Theta = \{\sigma_k, \mu_k\}$, which were obtained in Section V.

As discussed under Theorem 1, country i 's first-best optimal policy schedule in the Cobb-Douglas case is described by the following set of formulas:

$$1 + s_{i,k}^* = 1 + \mu_k; \quad 1 + t_{ji,k}^* = 1 + \omega_{ji,k}^*;$$

$$1 + x_{ij,k}^* = \frac{(\sigma_k - 1) \sum_{n \neq i} [(1 + \omega_{ni,g}^*) \lambda_{nj,k}^*]}{1 + (\sigma_k - 1)(1 - \lambda_{ij,k}^*)},$$

where superscript “*” indicates that a variable is evaluated in the counterfactual optimal policy equilibrium. Using the hat-algebra notation and our expression for the good-specific supply elasticity, $\omega_{ji,k}$ (equation (10)), we can write the above formulas in changes as follows.⁵⁸

[optimal import tax]

$$1 + t_{ji,k}^* = \frac{-\frac{\mu_k}{1 + \mu_k} \hat{r}_{ji,k} r_{ji,k} \Phi_{ji,k}^*}{1 - \frac{\mu_k}{1 + \mu_k} \sum_{\nu \neq i} \left\{ \hat{r}_{\nu,k} r_{\nu,k} \left[1 + (\sigma_k - 1)(1 - \hat{\lambda}_{\nu,k} \lambda_{\nu,k}) \right] \right\}},$$

[optimal export subsidy]

$$1 + x_{ij,k}^* = \frac{(\sigma_k - 1) \sum_{n \neq i} [(1 + t_{ni,g}^*) \hat{\lambda}_{nj,k} \lambda_{nj,k}]}{1 + (\sigma_k - 1)(1 - \hat{\lambda}_{ij,k} \lambda_{ij,k})},$$

(17) [change in taxes]

$$\widehat{1 + s_{i,k}} = \frac{1 + \mu_k}{1 + s_{i,k}}; \quad \widehat{1 + t_{ji,k}} = \frac{1 + t_{ji,k}^*}{1 + t_{ji,k}}; \quad \widehat{1 + x_{ij,k}} = \frac{1 + x_{ij,k}^*}{1 + x_{ij,k}}.$$

⁵⁷ As explained in Section I, under free entry, the number of firms operating in origin n -industry k can be expressed as $M_{i,k} = \bar{m}_{i,k} \rho_{i,k}$, where $\bar{m}_{i,k}$ is composed of parameters and variables that are invariant to policy. We can, therefore, use $\rho_{i,k}$ to track scale economies that channel through entry—as detailed under equation (6).

⁵⁸ The multiplier $\Phi_{ji,k}^* = 1 - (1 - \frac{1}{\mu_k})(\sigma_k - 1) \sum_{\nu \neq i} \left(\frac{\hat{r}_{i,k} r_{i,k}}{\hat{r}_{\nu,k} r_{\nu,k}} \hat{\lambda}_{\nu,k} \lambda_{\nu,k} \right) \frac{\hat{\rho}_{i,k} \rho_{i,k} \hat{w}_i w_i L_i}{\hat{\rho}_{\nu,k} \rho_{\nu,k} \hat{w}_\nu w_\nu L_\nu}$ accounts for cross-demand effects in foreign markets—see equation (10).

Since the rest of the world is passive in their use of taxes, $\widehat{1 + s_{n,k}} = \widehat{1 + t_{jn,k}} = \widehat{1 + x_{nj,k}} = 1$ for all $n \neq i$. To determine the change in expenditure shares, $\widehat{\lambda}_{ji,k}$, we need to determine the change in consumer price indexes. Invoking the CES structure of within-industry demand, we can express the change in market i –industry k 's consumer price index as

$$(18) \text{ [price indexes]} \quad \widehat{P}_{i,k} = \sum_{n \in \mathbb{C}} \left\{ \lambda_{ni,k} \left[\frac{\widehat{1 + t_{ni,k}}}{(\widehat{1 + x_{ni,k}})(\widehat{1 + s_{n,k}})} \widehat{w}_n \widehat{\rho}_{n,k}^{-\mu_k} \right]^{1-\sigma_k} \right\}^{\frac{1}{1-\sigma_k}}.$$

Recall that $\rho_{n,k} = L_{n,k}/L_n$ denotes industry k 's sales share in origin n , which—under free entry—is equal to the share of origin n 's workers employed in that industry. The above formulation uses the fact that, by free entry, $\widehat{M}_{i,k} = \widehat{\rho}_{i,k}$. Given $\widehat{P}_{i,k}$, we can calculate the change in expenditure and revenue shares as follows:

$$(19) \text{ [expenditure shares]} \quad \widehat{\lambda}_{ji,k} = \left[\frac{\widehat{1 + t_{ji,k}}}{(\widehat{1 + x_{ji,k}})(\widehat{1 + s_{j,k}})} \widehat{w}_j \widehat{\rho}_{j,k}^{-\mu_k} \right]^{1-\sigma_k} \widehat{P}_{i,k}^{\sigma_k-1},$$

$$\text{ [revenue shares]} \quad \widehat{r}_{ji,k} = \left(\frac{\widehat{1 + x_{ji,k}}}{\widehat{1 + t_{ji,k}}} \widehat{\lambda}_{ji,k} \widehat{Y}_i \right) \left(\sum_{n \in \mathbb{C}} r_{jn,k} \frac{\widehat{1 + x_{jn,k}}}{\widehat{1 + t_{jn,k}}} \widehat{\lambda}_{jn,k} \widehat{Y}_n \right)^{-1}.$$

The change in the wage rate, \widehat{w}_i , and industry-level sales shares, $\widehat{\rho}_{i,k}$, are dictated by the labor market clearing (LMC) condition, which ensures that industry-level sales match wage payments, industry by industry:

$$(20) \text{ [LMC]} \quad \widehat{\rho}_{i,k} \rho_{i,k} \widehat{w}_i w_i L_i = \sum_{j \in \mathbb{C}} \left[\frac{(1 + x_{ij,k}^*)(1 + s_{i,k}^*)}{1 + t_{ij,k}^*} \widehat{\lambda}_{ij,k} \lambda_{ij,k} e_{j,k} \widehat{Y}_j Y_j \right];$$

$$\sum_{k \in \mathbb{K}} \widehat{\rho}_{i,k} \rho_{i,k} = 1.$$

The change in national expenditure, \widehat{Y}_i , is governed by the balanced budget (BB) condition, which ensures that total expenditure matches total income from wage payments and tax revenues:

$$(21) \text{ [BB]} \quad \widehat{Y}_i Y_i = \widehat{w}_i w_i L_i - \sum_{k \in \mathbb{K}} (s_{i,k}^* \widehat{\lambda}_{ii,k} \lambda_{ii,k} e_{i,k} \widehat{Y}_i Y_i)$$

$$+ \sum_{j \neq i} \sum_{k \in \mathbb{K}} \left[\frac{t_{ji,k}^*}{1 + t_{ji,k}^*} \lambda_{ji,k} \widehat{\lambda}_{ji,k} e_{i,k} \widehat{Y}_i Y_i \right.$$

$$\left. + \frac{1 - (1 + x_{ij,k}^*)(1 + s_{i,k}^*)}{1 + t_{ij,k}^*} \lambda_{ij,k} \widehat{\lambda}_{ij,k} e_{j,k} \widehat{Y}_j Y_j \right].$$

Equations (17)–(21) represent a system of $2N + NK + [2(N - 1) + 1]K$ independent equations and unknowns. The independent unknown variables

are \widehat{w}_i (N unknowns), \widehat{Y}_i (N unknowns), $\widehat{\rho}_{i,k}$ (NK unknowns), $\widehat{1 + s_{i,k}}$ (K unknowns), $\widehat{1 + t_{j,i,k}}$ ($(N - 1)K$ unknowns), and $\widehat{1 + x_{ij,k}}$ ($(N - 1)K$ unknowns). Solving the aforementioned system is possible with information on observable data points, \mathbb{D} , and estimated parameters, $\Theta \equiv \{\mu_k, \sigma_k\}$. Once we solve this system, the welfare consequences of country i 's optimal policy are automatically determined. The following proposition outlines this result.⁵⁹

PROPOSITION 1: *Suppose we have data on observable shares, national accounts, and applied taxes, $\mathbb{D} = \{\lambda_{j,i,k}, r_{j,i,k}, e_{i,k}, \rho_{i,k}, Y_i, w_i L_i, x_{ij,k}, t_{j,i,k}, s_{i,k}\}_{j,i,k}$, and information on structural parameters, $\Theta \equiv \{\mu_k, \sigma_k\}$. We can determine the economic consequences of country i 's optimal policy by calculating $\mathbb{X} = \{\widehat{Y}_i, \widehat{w}_i, \widehat{\rho}_{i,k}, \widehat{1 + s_{i,k}}, \widehat{1 + t_{j,i,k}}, \widehat{1 + x_{ij,k}}\}$ as the solution to the system of equations (17)–(21). After solving for \mathbb{X} , we can fully determine the welfare consequence of country i 's optimal policy as*

$$\widehat{W}_i = \widehat{Y}_i / \prod_{k \in \mathbb{K}} \widehat{P}_{i,k}^{e_{i,k}}, \quad (\forall n \in \mathbb{C}),$$

where $\widehat{P}_{i,k}$ is determined by equation (18) as a function of \mathbb{X} and data, \mathbb{D} .

To take stock, the optimization-free procedure described by Proposition 1 simplifies the task of computing the gains from *first-best* trade and industrial policies. We can use a similar procedure (based on Theorems 2 and 3) to compute the gains from second-best trade policies—see online Appendix T. Without Proposition 1, we would have to rely on numerical optimization to recover country i 's optimal policy.⁶⁰ As noted earlier, numerical optimization can become increasingly difficult to implement when dealing with many free-moving policy variables. Furthermore, in many instances, obtaining credible results from numerical optimization requires specialized commercial solvers like Snopt or Knitro. Propositions 1's optimization-free procedure allows us to bypass such complications, delivering notable gains in computational speed and accuracy.

B. The Consequences of Noncooperative Policies

Our first set of results elucidates policy consequences when governments act noncooperatively. These results convey two basic points. First, noncooperative trade taxes are ineffective at correcting misallocation in domestic industries—even without retaliation. Second, the cost of retaliation is sizable. Together these results point to little justification for noncooperative trade taxation even in second-best economies plagued with misallocation.

Table 4 reports the gains from optimal noncooperative policies under free and restricted entry. The first three columns in each case report welfare gains assuming the rest of the world does not retaliate. The fourth column reports net welfare effects after retaliation. The first-best noncooperative policy consists of Pigouvian

⁵⁹ Under Proposition 1, the optimal policy specification uses our approximation for $\omega_{j,i,k}$. In online Appendix V, we examine the accuracy of our approximation and outline how our optimization-free approach can be alternatively conducted with an exact formula for $\omega_{j,i,k}$.

⁶⁰ Such a problem is typically formulated as a Mathematical Programming with Equilibrium Constraints (MPEC) problem—see Ossa (2014) for further details.

subsidies, import tariffs, and export subsidies (Theorem 1); the second-best consists of only import tariffs and export subsidies (Theorem 2); the third-best consists of only import tariffs (Theorem 3).

We can draw two main conclusions from Table 4. First, trade taxes are a poor second-best substitute for Pigouvian subsidies. Trade taxes can replicate only 1/3 of the welfare gains attainable under the first-best policy that combines Pigouvian subsidies with trade taxes.⁶¹ Under free entry, the first-best noncooperative policy increases welfare by 3.05 percent on average, whereas second-best trade taxes/subsidies raise welfare by only 1.19 percent. Third-best import tariffs (not paired with export subsidies) are half as effective.

These results reflect the tension between the terms of trade and allocative efficiency emphasized in Section III. Since our estimated scale and trade elasticities satisfy $\text{cov}(\sigma_k, \mu_k) < 0$, correcting interindustry misallocation with trade policy worsens the terms of trade—making it difficult for second-best trade policies to strike a balance between these two policy targets. We elucidate this point further in online Appendix W by artificially raising $\text{cov}(\sigma_k, \mu_k)$ and recomputing the gains from policy. The results displayed in Figure W.1 of the online Appendix point to a sharp rise in the efficacy of second-best trade taxes as $\text{cov}(\sigma_k, \mu_k)$ is artificially inflated.

Second, we find that retaliation more than wipes out the gains from noncooperative taxation. The net effect of noncooperative policies *after retaliation* is a welfare loss of 1.20 percent under free entry. Retaliation, in our calculation, occurs through the reciprocal adoption of optimal trade taxes by trading partners. As noted in Section III, the cost of retaliation may not deter a short-sighted government from erecting trade taxes—at least when trade taxes are a less politically controversial instrument for correcting misallocation. In such cases, our finding that trade policy is incapable of improving misallocation should serve as a deterrent.

C. *The Gains from International Cooperation*

Suppose governments recognize the danger of noncooperation and limit themselves to the cooperative policy outlined in Section II. As we show next, cooperative countries risk *immiserizing growth* if they take the lead in policy implementation. This issue can cause a race to the bottom in industrial policy implementation but can be resolved via a deep agreement.

The Immiserizing Effects of Unilateral Industrial Policy Implementation.—Suppose countries limit themselves to the efficient policy consisting of zero trade taxes and scale-correcting industrial policies. Unilateral industrial policy implementation, under this arrangement, worsens the ToT and can lead to immiserizing growth (Conjecture 2). In this section, we show that these immiserizing welfare effects do, in fact, occur in most countries.

Table 5 reports the welfare consequences of unilateral and coordinated industrial policy implementation. The policy applied in each case is a set of scale- or

⁶¹This finding echoes the numerical result in Balistreri and Markusen (2009) that optimal tariffs yield smaller gains in the presence of positive firm-level markups.

TABLE 5—IMMISERIZING EFFECTS OF NONCOORDINATED INDUSTRIAL POLICIES

	Restricted entry		Free entry	
	Unilateral	Coordinated	Unilateral	Coordinated
Gains from <i>corrective</i> industrial policies	−0.32%	1.67%	−2.78%	3.42%

Notes: The data source is the 2014 World Input-Output Database (Timmer et al. 2015; WIOD 2021). The columns titled “Unilateral” report welfare gains when a country unilaterally adopts industrial subsidies that restore marginal cost pricing in the domestic economy. The columns titled “Coordinated” report welfare gains when all countries simultaneously adopt industrial subsidies that restore marginal cost pricing globally. The average gains are calculated as the simple average across all 43 countries in the WIOD sample. Country-level results are reported in online Appendix X.

markup-correcting industrial policies that restore marginal cost-pricing in the local economy. Unilateral implementation corresponds to a scenario where the home country implements its industrial policy but trading partners do not reciprocate and stick to business as usual. Coordinated implementation corresponds to a reciprocal implementation of industrial policies worldwide.

The results in Table 5 confirm the strong immiserizing growth effects of unilateral scale correction. Real income in the average country drops by more than 2.7 percent if corrective industrial policies are implemented unilaterally. By comparison, welfare increases by more than 3.4 percent under coordinated or reciprocal implementation. These results suggest that we may be witnessing a *race to the bottom* in industrial policy implementation—without a deep agreement to ensure reciprocity in implementation. As things stand, cooperative countries have two choices: (i) implement scale correction and risk immiserizing growth or (ii) violate their commitments to cooperation by pairing corrective subsidies with trade restrictions.

It is worth emphasizing that the immiserizing growth effects reported in Table 5 stem from the tension between misallocation and the terms of trade. Given that our estimated scale and trade elasticities satisfy $\text{cov}(\sigma_k, \mu_k) < 0$, restoring allocative efficiency with Pigouvian subsidies worsens one’s ToT, to the point of causing immiserizing growth. We confirm this point in online Appendix W by artificially raising $\text{cov}(\sigma_k, \mu_k)$ and recomputing the gains from unilateral scale (or markup) correction. As $\text{cov}(\sigma_k, \mu_k)$ is artificially inflated relative to its estimated value, immiserizing growth effects fade and are even reversed (see Figure W.2 of the online Appendix).

The Gains from Deep versus Shallow Cooperation.—Recall from Section III that we can model international cooperation as a two-stage process:

- (i) The first stage involves a *shallow agreement* that disciplines noncooperative trade taxes helping countries avert a full-fledged trade war.
- (ii) The second stage involves a *deep agreement* that ensures reciprocity in industrial policy implementation, helping countries avoid a race to the bottom.

Figure 2 reports the welfare gains associated with each stage. The blue bars correspond to the welfare gains brought by the existing nexus of shallow agreements. These gains are computed relative to a counterfactual equilibrium where all countries




FIGURE 2. THE WELFARE GAINS FROM DEEP AND SHALLOW COOPERATION

Notes: The data source is the 2014 World Input-Output Database (Timmer et al. 2015; WIOD 2021). Welfare effects are computed under restricted entry. The gains from deep cooperation correspond to welfare gains when moving from the status quo to the globally efficient equilibrium in which all countries coordinate their corrective Pigouvian subsidies. The gains from shallow cooperation correspond to the avoided welfare losses when moving from the status quo to a noncooperative equilibrium where all countries adopt Nash trade taxes and subsidies.

adopt their noncooperative trade taxes. The red bars correspond to the prospective but unrealized welfare gains from deep cooperation. These gains are computed as the welfare gains associated with a universal implementation of (markup-correcting) industrial policies.⁶² As discussed earlier, a *deep agreement* is necessary to uncover these welfare gains.

The welfare gains from shallow cooperation are, on average, 3.2 percent. That is, the average country is poised to lose 3.2 percent of its real income if trade taxes are counterfactually raised to their noncooperative level everywhere. The existing nexus of shallow agreements have already materialized these gains. The prospective gains

⁶²The gains from deep cooperation can be computed with the aid of the optimal policy formulas specified under equation (9) and the logic presented earlier in Section VIA.

from deep cooperation are 1.6 percent under restricted entry. That is, if countries can agree to a reciprocal implementation of markup-correcting industrial policies, they can boost their real income by an additional 1.6 percent. Similar but larger welfare gains will occur under free entry.

Our estimated gains from shallow cooperation relate to the theoretical arguments in Bagwell and Staiger (2001, 2004). As explained earlier, shallow cooperation is sufficient for global efficiency if (i) $\text{cov}(\mu_k, \sigma_k) \geq 0$ or (ii) governments play a one-shot game where they simultaneously choose and implement their best-policy response with the belief that others do the same. Our micro-level estimation rejected the former condition. Exhaustive literature on the history of international policy coordination disputes the latter condition. If we suspect that governments move sequentially in the implementation stage and are not convinced about reciprocity in policy implementation, then deep cooperation is necessary for global efficiency.

A Stronger Case for International Cooperation.—The standard argument for international cooperation recognizes that governments can reap short-term gains if they adopt noncooperative trade taxes. But these short-term gains will turn into losses if trading partners retaliate. The standard argument may, thus, fail to deter a short-sighted government from taking the noncooperative route. After all, from the lens of traditional theories, the only way to reap short-term welfare gains (relative to the status quo) is to adopt noncooperative trade restrictions and hope for delayed retaliation.

Our quantitative analysis unveils a stronger argument for international cooperation. Starting from the status quo, a government seeking welfare improvements has two options: (i) engage in a coordinated industrial policy effort or (ii) raise noncooperative trade taxes to reap short-lived ToT benefits. We find that the former option not only delivers sustainable welfare gains but strictly dominates the latter option even before we adjust for the cost of retaliation.

This point is illustrated in Figure 3. The x-axis corresponds to the unrealized gains from deep cooperation. The y-axis corresponds to the maximal short-term gains from (noncooperative) unilateral policies, which transpire before retaliation. For most economies, the unrealized gains from deep cooperation dominate even the short-term gains from unilateral policy interventions. This finding indicates that the spillover gains from corrective policies in the rest of the world exceed the short-term ToT gains from noncooperative taxes—echoing our earlier claim that the ToT gains from policy are limited in scope, even before retaliation.⁶³

D. Sensitivity Analysis

In online Appendix Y, we recalculate the gains from policy under several alternative specifications. First, we recompute the gains assuming that the data-generating process is a *Melitz-Pareto* model. Second, we recompute the gains based on alternative

⁶³ Interestingly, the gains from deep cooperation favor small countries that have a comparative disadvantage in high-returns-to-scale (or high-profit) industries, e.g., Estonia, Malta, and Slovenia. The intuition is that these countries depend relatively more on imported varieties in high-returns-to-scale industries and under deep cooperation, these industries are subsidized across the globe.

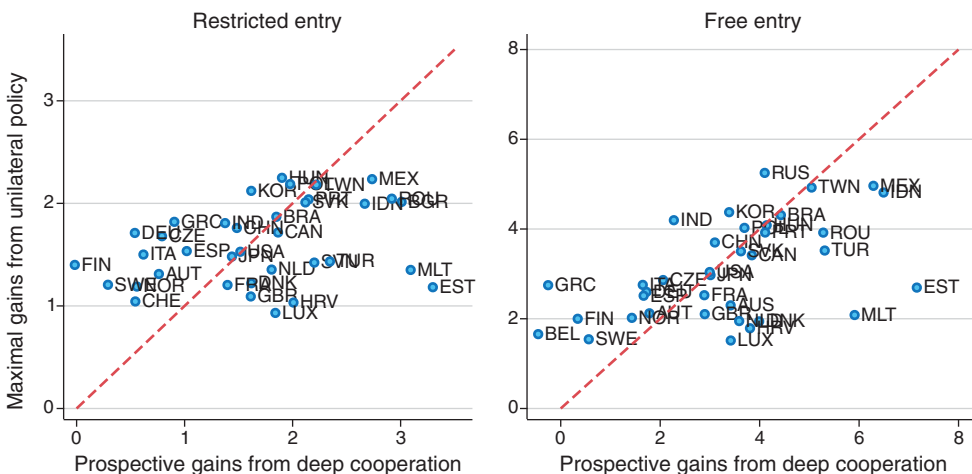


FIGURE 3. DEEP COOPERATION VERSUS UNILATERAL POLICY INTERVENTIONS

Notes: The data source is the 2014 World Input-Output Database (Timmer et al. 2015; WIOD 2021). The gains from first-best noncooperative policy are the gains when each country implements the policy characterized by Theorem 1 and the rest of the world is passive. The gains from global cooperation correspond to a scenario where all countries forgo trade taxation and apply industrial subsidies that restore marginal cost pricing.

values for σ_k and μ_k , which are estimated via a two-ways fixed effects estimation (as reported in online Appendix Q). Lastly, we recompute the gains from policy under a more conservative set of values assigned to μ_k and σ_k in services. In all cases, trade policy turns out to be a poor second-best instrument for resorting allocative efficiency. Another noteworthy observation is that accounting for firm-selection effects a la Melitz (2003) magnifies the gains from (first-best) optimal policies. However, these greater gains are primarily driven by the larger misallocation-correcting gains. If anything, second-best trade taxes/subsidies are even less effective at replicating the first-best policy gains in the presence of firm-selection into export markets.

What Parameter Values Would Imply Larger Gains from Policy?—We analyze this question in online Appendix Z, noting that the gains from optimal policy are increasing in two statistics: (i) the cross-industry variance of the scale elasticities, $\text{var}(\log \mu_k)$, and (ii) the average of the (inverse) trade elasticities, $E[1/(\sigma_k - 1)]$. In online Appendix Z, we adjust our estimated parameter values to artificially increase both of these statistics and recompute the gains from policy under these artificial parameter values. The results are reported in Figure Z.1 of the same Appendix. They reveal that the gains from optimal policy nearly double for all countries if we artificially increase $\text{var}(\log \mu_k)$ by a factor of about three. The policy gains for different countries, however, exhibit different degrees of sensitivity to an artificial increase in $E[1/(\sigma_k - 1)]$ —with the gains for larger countries like the United States or China being noticeably less sensitive. The intuition is that $\text{var}(\log \mu_k)$ governs the gains from correcting misallocation, whereas $E[1/(\sigma_k - 1)]$ regulates the extent to which countries can improve their ToT. For large countries, where trade accounts

for a smaller fraction of the GDP, there is less scope for raising real GDP via ToT improvements—hence, the lower sensitivity of policy gains to $E[1/(\sigma_k - 1)]$.

VII. Concluding Remarks

Correcting misallocation driven by economies of scale has often served as a justification for controversial trade and industrial policy practices. Yet we know surprisingly little about the efficacy of trade and industrial policy in increasing returns-to-scale industries. Against this backdrop, we employed micro-level trade data to examine the consequences of industry-level scale economies, offering insights into the effectiveness of various policy interventions. Our findings reveal that trade restrictions are a poor second-best policy choice for correcting misallocation in the domestic economy. Unilateral industrial policy measures can be equally ineffective, as they lead to immiserizing growth. Meanwhile, industrial policies coordinated internationally via deep agreements deliver welfare gains that are more transformative than any unilateral policy alternative.

The scale elasticity estimates obtained in this paper can be used for further exploration in several areas. First, our scale elasticity estimates can help disentangle the relative contribution of scale economies and Ricardian comparative advantage to intersectoral specialization. This is an old topic of interest, for which our empirical understanding is somewhat limited. Second, our estimates can perhaps shed fresh light on the puzzlingly large income gap between advanced and emerging economies. Economists have long hypothesized that a fraction of this income gap stems from specialization across low- and high returns-to-scale industries. Empirical assessments of this hypothesis have proven difficult due to a lack of comprehensive estimates for industry-level scale elasticities. Our micro-level estimates, however, can pave the way for such explorations. Finally, we document a negative cross-industry correlation between trade and scale elasticities. This relationship is crucial for policy evaluation in open economies, as it points to a tension between the terms of trade and allocative efficiency objectives. While our theoretical model is purposely agnostic about the origins of this negative relationship, further exploration into this matter is warranted.

REFERENCES

- Adao, Rodrigo, Michal Kolesár, and Eduardo Morales. 2019. “Shift-Share Designs: Theory and Inference.” *Quarterly Journal of Economics* 134 (4): 1949–2010.
- Aiginger, Karl, and Dani Rodrik. 2020. “Rebirth of Industrial Policy and an Agenda for the Twenty-First Century.” *Journal of Industry, Competition and Trade* 20 (2): 1–19.
- Amiti, Mary, Oleg Itskhoki, and Jozef Konings. 2019. “International Shocks, Variable Markups, and Domestic Prices.” *Review of Economic Studies* 86 (6): 2356–2402.
- Angrist, J., G. Imbens, and D. Rubin. 1996. “Identification of Causal Effects Using Instrumental Variables.” *Journal of the American Statistical Association* 91 (434): 444–55.
- Antràs, Pol, Agustín Gutiérrez, Teresa C. Fort, and Felix Tintelnot. 2022. “Trade Policy and Global Sourcing: A Rationale for Tariff Escalation.” Unpublished.
- Antweiler, Werner, and Daniel Trefler. 2002. “Increasing Returns and All That: A View from Trade.” *American Economic Review* 92 (1): 93–119.
- Bagwell, Kyle, and Robert W. Staiger. 2001. “Domestic Policies, National Sovereignty, and International Economic Institutions.” *Quarterly Journal of Economics* 116 (2): 519–62.
- Bagwell, Kyle, and Robert W. Staiger. 2004. *The Economics of the World Trading System*. Cambridge, MA: MIT Press.

- Bagwell, Kyle, and Seung Hoon Lee.** 2018. "Trade Policy under Monopolistic Competition with Firm Selection." Unpublished.
- Balistreri, Edward J., and James R. Markusen.** 2009. "Sub-National Differentiation and the Role of the Firm in Optimal International Pricing." *Economic Modelling* 26 (1): 47–62.
- Bank of Canada.** 2017. "Exchange Rates: Historical Noon and Closing Rates." <https://www.bankof-canada.ca/rates/exchange/legacy-noon-and-closing-rates/>
- Baqace, David, and Emmanuel Farhi.** 2019. "Networks, Barriers, and Trade." NBER Working Paper 26108.
- Bartelme, Dominick, Arnaud Costinot, Dave Donaldson, and Andres Rodriguez-Clare.** 2019. "The Textbook Case for Industrial Policy: Theory Meets Data." Unpublished.
- Berry, Steven, James Levinsohn, and Ariel Pakes.** 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841–90.
- Berry, Steven T.** 1994. "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics* 25 (2): 242–62.
- Beshkar, Mostafa, and Ahmad Lashkaripour.** 2019. "Interdependence of Trade Policies in General Equilibrium." Unpublished.
- Beshkar, Mostafa, and Ahmad Lashkaripour.** 2020. "The Cost of Dissolving the WTO: The Role of Global Value Chains." Unpublished.
- Bhagwati, Jagdish.** 1958. "Immiserizing Growth: A Geometrical Note." *Review of Economic Studies* 25 (3): 201–205.
- Bhagwati, Jagdish, and Vangal K. Ramaswami.** 1963. "Domestic Distortions, Tariffs and the Theory of Optimum Subsidy." *Journal of Political Economy* 71 (1): 44–50.
- Bhagwati, Jagdish N.** 1988. *Protectionism*. Cambridge, MA: MIT Press.
- Blanchard, Emily J., Chad P. Bown, and Robert C. Johnson.** 2016. "Global Supply Chains and Trade Policy." NBER Working Paper 21883.
- Boehm, Christoph E., Andrei A. Levchenko, and Nitya Pandalai-Nayar.** 2020. "The Long and Short (Run) of Trade Elasticities." NBER Working Paper 27064.
- Caliendo, Lorenzo, and Fernando Parro.** 2015. "Estimates of the Trade and Welfare Effects of NAFTA." *Review of Economic Studies* 82 (1): 1–44.
- Caliendo, Lorenzo, Robert C. Feenstra, John Romalis, and Alan M. Taylor.** 2015. "Tariff Reductions, Entry, and Welfare: Theory and Evidence for the Last Two Decades." NBER Working Paper 21768.
- Caliendo, Lorenzo, Robert C. Feenstra, John Romalis, and Alan M. Taylor.** 2021. "A Second-best Argument for Low Optimal Tariffs." NBER Working Paper 28380.
- Campolmi, Alessia, Harald Fadinger, and Chiara Forlati.** 2014. "Trade Policy: Home Market Effect versus Terms-of-Trade Externality." *Journal of International Economics* 93 (1): 92–107.
- Campolmi, Alessia, Harald Fadinger, and Chiara Forlati.** 2018. "Trade and Domestic Policies in Models with Monopolistic Competition." Unpublished.
- Chaney, Thomas.** 2008. "Distorted Gravity: The Intensive and Extensive Margins of International Trade." *American Economic Review* 98 (4): 1707–21.
- Considine, Timothy J.** 2001. "Markup Pricing in Petroleum Refining: A Multiproduct Framework." *International Journal of Industrial Organization* 19 (10): 1499–1526.
- Coşar, A. Kerem, Paul L. E. Grieco, Shengyu Li, and Felix Tintelnot.** 2018. "What Drives Home Market Advantage?" *Journal of International Economics* 110: 135–50.
- 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. Amsterdam: Elsevier.
- Costinot, Arnaud, and Iván Werning.** 2019. "Lerner Symmetry: A Modern Treatment." *American Economic Review: Insights* 1 (1): 13–26.
- Costinot, Arnaud, Andrés Rodríguez-Clare, and Iván Werning.** 2016. "Micro to Macro: Optimal Trade Policy with Firm Heterogeneity." NBER Working Paper 21989.
- Costinot, Arnaud, Dave Donaldson, Jonathan Vogel, and Iván Werning.** 2015. "Comparative Advantage and Optimal Trade Policy." *Quarterly Journal of Economics* 130 (2): 659–702.
- DATAMYNE Inc.** "International Database." <https://www.datamyne.com> (accessed between May 8th of 2014 and November 8th of 2016).
- De Blas, Beatriz, and Katheryn N. Russ.** 2015. "Understanding Markups in the Open Economy." *American Economic Journal: Macroeconomics* 7 (2): 157–80.
- De Loecker, Jan, and Frederic Warzynski.** 2012. "Markups and Firm-Level Export Status." *American Economic Review* 102 (6): 2437–71.
- De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik.** 2016. "Prices, Markups, and Trade Reform." *Econometrica* 84 (2): 445–510.

- Demidova, Svetlana.** 2017. "Trade Policies, Firm Heterogeneity, and Variable Markups." *Journal of International Economics* 108: 260–73.
- Demidova, Svetlana, and Andres Rodríguez-Clare.** 2009. "Trade Policy under Firm-Level Heterogeneity in a Small Economy." *Journal of International Economics* 78 (1): 100–112.
- Dhingra, Swati, and John Morrow.** 2019. "Monopolistic Competition and Optimum Product Diversity under Firm Heterogeneity." *Journal of Political Economy* 127 (1): 196–232.
- Dixit, Avinash, and Victor Norman.** 1980. *Theory of International Trade: A Dual, General Equilibrium Approach*. Cambridge, UK: Cambridge University Press.
- Drazen, Allan, and Nuno Limão.** 2008. "A Bargaining Theory of Inefficient Redistribution Policies." *International Economic Review* 49 (2): 621–57.
- Eaton, Jonathan, and Samuel Kortum.** 2002. "Technology, Geography, and Trade." *Econometrica* 70 (5): 1741–79.
- Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu.** 2015. "Competition, Markups, and the Gains from International Trade." *American Economic Review* 105 (10): 3183–3221.
- Farrokhi, Farid, and Ahmad Lashkaripour.** 2021. "Can Trade Policy Mitigate Climate Change?" Unpublished.
- Fattouh, Bassam.** 2007. *The Drivers of Oil Prices: The Usefulness and Limitations of Non-Structural Model, the Demand-Supply Framework and Informal Approaches*. Oxford, UK: Oxford Institute for Energy Studies.
- Felbermayr, Gabriel, Benjamin Jung, and Mario Larch.** 2013. "Optimal Tariffs, Retaliation, and the Welfare Loss from Tariff Wars in the Melitz Model." *Journal of International Economics* 89 (1): 13–25.
- Garred, Jason.** 2018. "The Persistence of Trade Policy in China after WTO Accession." *Journal of International Economics* 114: 130–42.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift.** 2020. "Bartik Instruments: What, When, Why, and How." *American Economic Review* 110 (8): 2586–2624.
- Gros, Daniel.** 1987. "A Note on the Optimal Tariff, Retaliation and the Welfare Loss from Tariff Wars in a Framework with Intra-Industry Trade." *Journal of International Economics* 23 (3–4): 357–67.
- Grossman, Gene, and Elhanan Helpman.** 1994. "Protection for Sale." *American Economic Review* 84 (4): 833–50.
- Haaland, Jan I., and Anthony J. Venables.** 2016. "Optimal Trade Policy with Monopolistic Competition and Heterogeneous Firms." *Journal of International Economics* 102: 85–95.
- Harrison, Ann, 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, 4039–4214. Amsterdam: Elsevier.
- Hicks, John R.** 1939. "The Foundations of Welfare Economics." *Economic Journal* 49 (196): 696–712.
- Irwin, Douglas A.** 2017. *Clashing over Commerce: A History of US Trade Policy*. Chicago: University of Chicago Press.
- Kaldor, Nicholas.** 1939. "Welfare Propositions of Economics and Interpersonal Comparisons of Utility." *Economic Journal* 49 (195): 549–52.
- Khandelwal, Amit.** 2010. "The Long and Short (of) Quality Ladders." *Review of Economic Studies* 77 (4): 1450–76.
- Kosfeld, Michael, Akira Okada, and Arno Riedl.** 2009. "Institution Formation in Public Goods Games." *American Economic Review* 99 (4): 1335–55.
- Krugman, Paul.** 1980. "Scale Economies, Product Differentiation, and the Pattern of Trade." *American Economic Review* 70 (5): 950–59.
- Kucheryavy, Konstantin, Gary Lyn, and Andrés Rodríguez-Clare.** 2023a. "Grounded by Gravity: A Well-Behaved Trade Model with Industry-Level Economies of Scale." *American Economic Journal: Macroeconomics* 15 (2): 372–412.
- Kucheryavy, Konstantin, Gary Lyn, and Andrés Rodríguez-Clare.** 2023b. "Replication data for: Grounded by Gravity: A Well-Behaved Trade Model with Industry-Level Economies of Scale." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.1257/mac.20190156>
- Lane, Nathaniel.** 2020. "The New Empirics of Industrial Policy." *Journal of Industry, Competition and Trade* 20 (2): 209–34.
- Lashkaripour, Ahmad.** 2021. "The Cost of a Global Tariff War: A Sufficient Statistics Approach." *Journal of International Economics* 131: 103419.
- Lerner, Abba P.** 1934. "The Concept of Monopoly and the Measurement of Monopoly Power." *Review of Economic Studies* 1 (3): 157–75.
- Lerner, Abba P.** 1936. "The Symmetry Between Import and Export Taxes." *Economica* 3 (11): 306–13.
- Matsuyama, Kiminori.** 2008. "A General Theory of Competitive Trade." Unpublished.

- Melitz, Marc.** 2003. "The Impact of Trade on Aggregate Industry Productivity and Intra-Industry Reallocations." *Econometrica* 71 (6): 1695–1725.
- Murdoch, James C., Todd Sandler, and Wim P. M. Vijverberg.** 2003. "The Participation Decision versus the Level of Participation in an Environmental Treaty: A Spatial Probit Analysis." *Journal of Public Economics* 87 (2): 337–62.
- Ossa, Ralph.** 2011. "A "New Trade" Theory of GATT/WTO Negotiations." *Journal of Political Economy* 119 (1): 122–52.
- Ossa, Ralph.** 2014. "Trade Wars and Trade Talks with Data." *American Economic Review* 104 (12): 4104–46.
- Ossa, Ralph.** 2016. "Quantitative Models of Commercial Policy." In *Handbook of Commercial Policy*, Vol. 1, edited by Kyle Bagwell and Robert W. Staiger, 207–59. Amsterdam: Elsevier.
- Pesaran, M. Hashem, Ron P. Smith, and Takamasa Akiyama.** 1998. *Energy Demand in Asian Developing Economies*. Oxford, UK: Oxford University Press.
- Pierce, Justin R., and Peter K. Schott.** 2012. "Concording U.S. Harmonized System Codes over Time." *Journal of Official Statistics* 28 (1): 53–68.
- Piveteau, Paul, and Gabriel Smagghue.** 2019. "Estimating Firm Product Quality using Trade Data." *Journal of International Economics* 118: 217–32.
- Ramondo, Natalia, Andrés Rodríguez-Clare, and Milagro Saborío-Rodríguez.** 2016. "Trade, Domestic Frictions, and Scale Effects." *American Economic Review* 106 (10): 3159–84.
- Samuelson, Paul Anthony.** 1948. "Foundations of Economic Analysis." *Science and Society* 13 (1).
- Sen, Amartya K.** 1962. "An Aspect of Indian Agriculture." *Economic Weekly* 14 (4–6): 243–46.
- Shaffer, Gregory C., and Mark A. Pollack.** 2009. "Hard vs. Soft Law: Alternatives, Complements, and Antagonists in International Governance." *Minnesota Law Review* 94: 706–99.
- Simonovska, Ina, and Michael E. Waugh.** 2014. "The Elasticity of Trade: Estimates and Evidence." *Journal of International Economics* 92 (1): 34–50.
- Spearot, Alan.** 2016. "Unpacking the Long-Run Effects of Tariff Shocks: New Structural Implications from Firm Heterogeneity Models." *American Economic Journal: Microeconomics* 8 (2): 128–67.
- Timmer, Marcel P., Erik Dietzenbacher, Bart Los, Robert Stehrer, and Gaaitzen J. De Vries.** 2015. "An Illustrated User Guide to the World Input–Output Database: The Case of Global Automotive Production." *Review of International Economics* 23 (3): 575–605.
- Venables, Anthony J.** 1987. "Trade and Trade Policy with Differentiated Products: A Chamberlinian-Ricardian Model." *Economic Journal* 97 (387): 700–717.
- WIOD**, 2021, "World Input-Output Database 2016 Release, 2000-2014," Groningen Growth and Development Centre [distributor]. <https://doi.org/10.34894/PJ2MIC>, DataverseNL, V2.
- Wu, Michael, Bei Yin, Aida Vosoughi, Christoph Studer, Joseph R. Cavallaro, and Chris Dick.** 2013. "Approximate Matrix Inversion for High-Throughput Data Detection in the Large-Scale MIMO Uplink." In *2013 IEEE International Symposium on Circuits and Systems*, 2155–58. New York, NY: IEEE.