Home-Foreign Elasticity Schedules and the Gains from Trade*

Jonathan Vogel
University of California, Los Angeles and NBER

September 2026

Abstract

Gains from trade are typically calculated using a constant elasticity estimated from substitution across foreign countries. For general homothetic preferences, however, gains from trade depend only on the domestic expenditure share and the elasticity schedule between expenditures on the home good and a Hicks composite of foreign goods. We further show that every finite point on this price path can be implemented in general equilibrium, so suitable trade-policy variation can, in principle, identify this one-dimensional schedule directly from equilibrium responses. To take this welfare result to the data, we directly estimate sector-level home-foreign elasticity schedules in a multi-sector framework using changes in home relative to foreign expenditure and import-tariff changes as shifters of foreign relative prices. Estimated home-foreign elasticities generally rise with domestic expenditure shares. Compared with holding each country-sector elasticity fixed at its initial value, allowing those elasticities to rise along the path to autarky lowers measured gains from trade for every country, regardless of how expenditure is aggregated across sectors.

1 Introduction

How large are the welfare gains from international trade? In a broad class of quantitative trade models, the answer depends on two statistics: the domestic expenditure share and a trade elasticity that governs both substitution between home and foreign goods and substitution across foreign sources (Arkolakis et al., 2012). In practice, this elasticity is typically estimated from expenditure reallocation across foreign origins in response to changes in their relative prices. Thus, while the underlying theory restricts two distinct margins of substitution to be equally elastic, the empirical implementation identifies this common elasticity using variation only across foreign sources. This distinction matters: the welfare consequences of moving from trade to autarky generally depend on the home-foreign elasticity schedule, but not at all on the across-foreign-source elasticity that is used to identify it.

*I thank Arnaud Costinot and Andres Santos for very helpful comments.

The basic idea of this paper is simple. Costinot and Rodríguez-Clare (2018) observe that foreign goods can be treated as a Hicks composite along a path on which their prices move proportionally relative to home prices. We first formalize this observation: for general homothetic preferences, gains from trade can be calculated from the initial domestic expenditure share and the schedule of elasticities between home and foreign expenditures along this path. Substitution across foreign sources does not enter this welfare calculation. The proportional ray can be used to calculate gains from trade even if its intermediate points are not equilibrium outcomes. However, if this ray could not be implemented in general equilibrium, the low-dimensional welfare representation would not yield a low-dimensional identification strategy. Under the general homothetic preferences considered here, recovering the schedule from off-ray price movements would generally require estimating a much richer, high-dimensional demand system of the sort studied by Adão et al. (2017). Because wages adjust endogenously, it is not obvious that equilibrium price variation can trace the ray. We show that it can. Every finite point on the ray can be implemented in general equilibrium using bilateral tariffs or subsidies imposed by the home country, with revenue rebated through a lump-sum transfer. Our implementation result has a strong implication for identification: trade-policy variation can, in principle, identify the one-dimensional welfare-relevant elasticity schedule directly from equilibrium responses of the domestic-to-foreign expenditure ratio to proportional foreign-price shifters.

In practice, however, the price variation available in the data does not trace this ray: source-specific prices move by different amounts, even when underlying trade-cost changes are proportional across sources. To use this variation for identification, we introduce multiple sectors and impose weak separability between home and foreign goods within each sector and weak separability across sectors. Our general welfare results continue to apply under this additional structure but do not require it. The additional structure serves two roles in our empirical implementation. First, and most importantly, within-sector weak separability provides a sector-level foreign price aggregate for these source-specific price changes. Changes in this aggregate relative to the home price play the same role as the proportional foreign-price shifter in our general theory and provide the variation used to identify the sector-level home-foreign elasticity schedule. Second, the inclusion of many sectors allows us to pool variation in expenditure shares and tariff changes across sectors within broader sectoral aggregates when estimating how the home-foreign elasticity varies with the domestic expenditure share.

We first treat each sector-level home-foreign elasticity as an arbitrary function of the domestic expenditure share. Our baseline identification approach assumes that this elasticity schedule is common across countries within each sector; countries with different domestic expenditure shares can, therefore, have different elasticities within sectors. We also assume that tariff changes satisfy a conditional exclusion restriction. Under these assumptions, tariff variation identifies each sectoral elasticity schedule nonparametrically over the supported range of domestic expenditure shares. In sensitivity, we partially relax the cross-country restriction by allowing the level of each sector's schedule to vary with observable destination-country characteristics, while retaining a common slope. For estimation, we approximate each schedule with a polynomial. Within this polynomial family, a constant home-foreign elasticity as in CES models is the zeroth-order approximation, while our baseline linear schedule is the first-order approximation. In a sensitivity exercise, we use a quadratic, second-order approximation.

Our empirical strategy estimates these sector-level home-foreign elasticity schedules directly, rather than inferring them from substitution across foreign sources. The dependent variable is the change in expenditure on the home sector relative to aggregate foreign expenditure in the same sector. Import-tariff changes shift the price of the foreign aggregate relative to the home price. Because producer costs are endogenous and unobserved, we estimate their changes jointly with the elasticity parameters. Under the maintained exclusion restriction, variation in the full foreign-relative-price shifter, which combines tariff changes with estimated country-level productivity-adjusted wage changes, identifies the level of the home-foreign elasticity, while its interaction with the initial domestic expenditure share identifies how that elasticity varies with the share.

The baseline estimator requires the observed tariff and estimated country-cost components to satisfy exclusion restrictions. We assess these restrictions using two distinct instrumental-variable strategies. The first uses lagged tariff shifters and domestic expenditure shares to address correlated measurement error between the dependent variable and the regressor that identifies the elasticity-schedule slope. The second combines tariff changes implied by predetermined trade agreements with leave-sector-out estimates of country cost changes to address potential endogeneity in the tariff-and-country-cost shifter. This exercise implements a parametric analogue of the tariff-based nonparametric identification argument outlined above.

Because sectoral domestic expenditure is constructed under many assumptions and measured less directly than international trade flows, we implement this strategy in two independently constructed global data sets. The first combines the Inter-Country Input-Output (ICIO) tables of the Organisation for Economic Co-operation and Development (OECD) with tariff changes from the Global Tariff Database. It covers 80 destination countries and 19 manufacturing sectors. In our baseline, we use changes between 1995 (the first year in the dataset) and 2007 (the year before the financial crisis). The second combines expenditure data from the Global Trade Analysis Project (GTAP) Data Base with the same tariff source and covers 120 destination countries and 27 manufacturing sectors between 2004 and 2014 (the first and last years in the dataset). In a sensitivity exercise, we also estimate the OECD ICIO model from 2004 to 2013, closely aligning its sample period with the GTAP analysis. Both data sets measure total absorption, including intermediate and final uses. We therefore impose the assumption that final and intermediate users purchase the same sectoral composite and consequently have the same origin shares and home-foreign elasticity, as in, e.g., Caliendo and Parro (2015).

We find that home-foreign elasticities generally increase with domestic expenditure shares across data sets, sample periods, and sectoral classifications. For instance, with elasticity-schedule intercepts that are sector-specific and slopes that are common across all manufacturing sectors, the estimated slopes are 5.45 and 2.39 in the OECD ICIO and GTAP data, respectively; both 95 percent confidence intervals based on standard errors clustered by destination country exclude zero. These estimates imply that the average fitted home-foreign elasticity rises from 4.01 at the sector-specific 25th percentiles of the domestic expenditure share to 5.88 at the 75th percentiles in the OECD ICIO, and from 2.55 to 3.48 in GTAP. When slopes are allowed to vary across consumer- and producer-serving manufacturing, both slopes are positive in both data sets. When slopes are allowed to vary across seven manufacturing aggregates, all seven estimates are positive in GTAP and five of seven are positive in the OECD ICIO.

We subject these findings to an extensive sensitivity analysis that varies samples, pooling restrictions, tariff construction, functional form, influential observations, and identifying variation. These exercises include a specification that allows the level of each sectoral schedule to vary with observable destination-country characteristics and the two IV strategies described above. Across these exercises, the fitted elasticity levels evaluated at the 25th, 50th, and 75th percentiles of the domestic expenditure share distribution remain broadly comparable in magnitude, and the conclusion that home-foreign elasticities generally increase with domestic expenditure shares is stable. A Monte Carlo exercise based on the quantitative multi-sector extended Armington model developed in Appendix C shows that the estimator produces slopes centered around and near zero when the true slopes are zero and accurately recovers positive slopes generated by linear elasticity schedules.

Finally, we combine the estimated elasticity schedules (both intercepts and slopes) with alternative assumptions about substitution across sectors and a parsimonious structure for intermediate inputs to calculate gains from trade. These calculations reveal a novel relationship between the initial domestic expenditure share and the elasticity relevant for gains from trade. When the home-foreign elasticity increases with the domestic expenditure share, a sector with a lower initial domestic expenditure share is not only more exposed to trade; it also begins at a lower elasticity and travels along a longer portion of the elasticity schedule on the path to autarky. The initial domestic expenditure share therefore determines both trade exposure and the effective home-foreign elasticity along that path, even when the same elasticity schedule applies across sectors. Using the OECD ICIO estimates, the median gain from trade across countries, defined as the log change in real income from autarky to the trade equilibrium, is 10.59 log points. Holding each country-sector elasticity fixed at its initial value raises this median to 12.64 log points. In general (regardless of the particular assumptions imposed on sectoral aggregation), allowing estimated elasticities to vary along the path to autarky rather than fixing them at their initial levels lowers measured gains from trade for every country.

This paper contributes to two closely related literatures. First, it contributes to the literature on sufficient statistics for the gains from trade. Arkolakis et al. (2012) show that, within a broad class of quantitative trade models, changes in welfare can be recovered from changes in domestic expenditure shares and a constant trade elasticity. Their CES restriction on import demand implies that the welfare-relevant home-foreign elasticity equals the elasticity governing substitution across source countries. They recommend estimating this common elasticity using the gravity relationship between bilateral imports and variable trade costs; that is, they suggest leveraging variation across foreign source countries. Costinot and Rodríguez-Clare (2018) characterize the welfare-relevant object more generally. They interpret gains from trade as the area under the compensated demand curve for foreign factor services and emphasize that, outside the constant-elasticity case, the relevant object is the elasticity of that demand along the path to autarky. It is their insight that we extend and bring to the data. Ansari et al. (2026) provide a nice bridge between the two previous papers. They show that gains-from-trade predictions can be highly sensitive to small departures from workhorse models, even when taking an approach motivated by Arkolakis et al. (2012) and matching the gravity moment used in standard estimation to identify substitution across foreign sources. Across the broader class of models they consider, satisfying this moment does not pin down the welfare-relevant home-foreign elasticity. Motivated by the sensitivity documented by Ansari et al. (2026) and building on the insights of Costinot and Rodríguez-Clare (2018), we allow the home-foreign elasticity to differ from the elasticity across foreign sources, estimate the home-foreign elasticity schedule directly rather than infer it from the bilateral elasticity conventionally identified by the across-foreign-source gravity moment, and show that the proportional ray can also be implemented in general equilibrium.[1] Circling back to the quantitative applications that motivated Arkolakis et al. (2012), Sections C, D, and E of the theoretical appendix develop fully specified Armington, Krugman, and Melitz-Pareto structural gravity models that generate the baseline linear home-foreign elasticity schedule exactly (thereby nesting the case of CES demand) while also allowing the home-foreign elasticity to differ from the elasticity across foreign sources. These models provide tractable parametric implementations of the general results. The Krugman model also shows that, within this parametric setting, the gains-from-trade formula does not rely on the Armington supply structure imposed in the main theoretical analysis.

The second is the literature on nonparametric approaches to trade counterfactuals. Adão et al. (2017) show that counterfactual trade and welfare changes can be recovered from a reduced factor-demand system without imposing the functional-form restrictions of standard gravity models. We focus on the one-dimensional demand system sufficient for measuring gains from trade. Our implementation result shows that suitable trade-policy variation can, in principle, trace this one-dimensional object directly along the proportional ray. We then provide conditions under which the non-proportional tariff variation observed in the data identifies the same object. Relatedly, Adão et al. (2026) estimate functions that summarize firm heterogeneity on the extensive and intensive export margins and use a monopolistic-competition model to quantify their implications for gains from trade. Whereas their approach estimates bilateral firm-export responses, we directly estimate the home-foreign elasticity schedule that enters the gains-from-trade calculation.

The paper is also related to the broader literature on trade elasticities. A large gravity literature estimates constant trade elasticities from bilateral variation across source countries (e.g., Eaton and Kortum, 2002; Simonovska and Waugh, 2014; Caliendo and Parro, 2015). We estimate home-foreign elasticity schedules, which govern substitution between the home good and all foreign sources jointly, which is the margin that enters the gains-from-trade calculation. The home-foreign elasticity is also central to international macroeconomics, where it shapes trade-balance dynamics and the international transmission of shocks (e.g., Backus et al., 1994; Heathcote and Perri, 2002; Corsetti et al., 2008).[2] Closest to our empirical analysis, Feenstra et al. (2018) distinguish the "macro" elasticity between home and imported goods from the "micro" elasticity across foreign sources and estimate both using detailed U.S. data. Like them, we distinguish between these two margins. Unlike them, we allow the home-foreign elasticity to vary with the domestic expenditure share and esti-

[1]Relatedly, Baqaee et al. (2025) address the sensitivity of surplus measurement to demand-curve extrapolation by measuring the area under firms' input demand curves directly. We instead estimate the welfare-relevant elasticity schedule and use it to integrate along the path to autarky.

[2]Our estimated home-foreign elasticities tend to be higher than those estimated in the international macro literature, although we use changes over substantially longer horizons: 1995-2007 in the OECD ICIO and 2004-2014 in GTAP.

mate its sector-specific schedule using variation across countries. Appendix E embeds this schedule in their Melitz-Pareto environment, with their constant home-foreign elasticity as the constant-schedule case. Our data are suitable only for long differences, so we provide no evidence on whether the home-foreign elasticity varies with the adjustment horizon, a dimension that Boehm et al. (2023) show is important for bilateral trade-flow elasticities.

The remainder of the paper proceeds as follows. Section 2 derives the welfare results and establishes the feasibility of the proportional foreign-price path. Section 3 introduces the multi-sector structure used for identification. Section 4 describes the empirical implementation and presents the elasticity estimates. Section 5 uses these estimates to quantify the gains from trade and Section 6 concludes.

2 Gains from Trade

Theoretical Framework. On the production side, the model is a standard Armington model of international trade (Armington, 1969). Let $\mathcal{N}$ denote a finite set of countries. Each country $i \in \mathcal{N}$ produces one country-specific differentiated good using an exogenously supplied and internationally immobile factor.[3] Country i has an exogenous factor endowment $L_i > 0$ and linear production technology $y_i = a_i L_i$, with $a_i > 0$. Markets are competitive and each country has a representative consumer. Endogenous wages are w_i, so nominal income is $Y_i = w_i L_i$ and the unit production cost is $c_i = w_i / a_i$. Delivering one unit of country i's good to country k requires shipping $\mathcal{D}_{ki} \geq 1$ units with $\mathcal{D}_{ii} = 1$, where $\mathcal{D}_{ki}$ is the iceberg cost of shipping. The consumer price in destination k is therefore $p_{ki} = \mathcal{D}_{ki} c_i$. Goods markets clear and trade is balanced.

Preferences, however, are more general. Country k has preferences over the vector $\boldsymbol{x}_k = (x_{ki})_{i \in \mathcal{N}}$, where $x_{ki} \geq 0$ denotes the quantity of country i's good delivered to country k. These preferences are represented by a utility function u_k that is continuous, increasing, concave, and homogeneous of degree one. For strictly positive consumption bundles, u_k is twice continuously differentiable and strictly quasi-concave, and Hicksian demand is interior and continuously differentiable in prices at every strictly positive price vector.[4] Country n's preferences may contain arbitrary interactions between the domestic variety and the foreign varieties.

Here, we focus on country n. Because u_n is homogeneous of degree one, define country n's unit expenditure function as

$$e_n(\boldsymbol{p}_n) \equiv \min_{\boldsymbol{x}_n \geq 0} \left\{ \sum_{i \in \mathcal{N}} p_{ni} x_{ni} : u_n(\boldsymbol{x}_n) \geq 1 \right\}$$

and let

$$h_{ni}(\boldsymbol{p}_n) \equiv \frac{\partial e_n(\boldsymbol{p}_n)}{\partial p_{ni}}$$

[3]The one-good-per-country formulation is for expositional convenience. Section 3 shows how the gains-from-trade result applies in a multi-sector environment. Appendices D and E replace the Armington supply side with, respectively, a Krugman model and a Melitz-Pareto model.

[4]Interiority is imposed only for tractability: if aggregate imports reach zero at a finite choke price, the welfare result can be re-stated over the range in which imports are positive, with autarky defined as the limiting endpoint.

denote the corresponding Hicksian demand corresponding to one unit of utility. The assumptions above imply that e_n is positive, increasing, concave, and homogeneous of degree one. By the assumed continuous differentiability of Hicksian demand and Shephard's lemma, e_n is twice continuously differentiable at every strictly positive price vector.

Let $X_{ni} \equiv p_{ni}x_{ni}$ denote expenditure by country n on goods from origin i, let $\lambda_{ni} \equiv X_{ni}/\sum_j X_{nj}$ denote its share of country n's total expenditure, and define $X_{nF} \equiv \sum_{i \neq n} X_{ni}$ as n's spending across all foreign sources. Country n's expenditure shares sum to $1 - \lambda_{nn}$ across all foreign sources. Denote by $W_n \equiv Y_n/e_n(\boldsymbol{p}_n)$ country n's real income.

Welfare Along a Proportional Foreign-Price Ray. Costinot and Rodríguez-Clare (2018) observe that foreign goods can be treated as a Hicks composite along a path on which their prices move proportionately relative to the domestic price. We now formalize this observation.

Consider a path in country n's consumer-price space indexed by $m \geq 1$, beginning at the consumer-price vector in the initial equilibrium, where $m = 1$. The scalar m is the common multiplier applied to every foreign price relative to the home price: as m rises, all foreign relative prices rise in the same proportion. Let $p_{ni}(m)$ denote the consumer price in country n of the good produced by country i at point m on this path, and let $\boldsymbol{p}_n(m) = (p_{ni}(m))_{i \in \mathcal{N}}$ denote the corresponding consumer-price vector. Define the normalized price vector by $\pi_{nn}(m) = 1$ and, for every $i \neq n$,

$$\pi_{ni}(m) = m\frac{p_{ni}(1)}{p_{nn}(1)}. \tag{1}$$

The vector $\boldsymbol{\pi}_n(m)$ is obtained from the normalized consumer-price vector in the initial equilibrium by holding its domestic component fixed and multiplying every foreign component by m. In what follows, we restrict attention to a proportional foreign-price ray, defined by the requirement that, for every $i \neq n$,

$$\frac{p_{ni}(m)}{p_{nn}(m)} = \pi_{ni}(m). \tag{2}$$

Define unit expenditure along the normalized ray by

$$H_n(m) \equiv e_n(\boldsymbol{\pi}_n(m)),$$

where $H_n(m)$ is expressed in units of the domestic good because $\pi_{nn}(m) = 1$. By construction, $\boldsymbol{p}_n(m) = p_{nn}(m)\boldsymbol{\pi}_n(m)$, so homogeneity implies

$$e_n(\boldsymbol{p}_n(m)) = p_{nn}(m)H_n(m). \tag{3}$$

Along the ray, homotheticity and Shephard's lemma imply that each origin's expenditure share satisfies

$$\lambda_{ni}(m) = \frac{\pi_{ni}(m)h_{ni}(\boldsymbol{\pi}_n(m))}{H_n(m)}. \tag{4}$$

Euler's theorem implies that these shares sum to one.

Because every foreign component of $\boldsymbol{\pi}_n(m)$ is proportional to m, differentiating $H_n(m)$ and applying Shephard's lemma gives

$$\frac{\mathrm{d}\ln H_n(m)}{\mathrm{d}\ln m} = \sum_{i\neq n}\lambda_{ni}(m) = 1-\lambda_{nn}(m).$$

Moreover, since $Y_n/p_{nn} = a_n L_n$, equation (3) implies

$$W_n(m) = \frac{a_n L_n}{H_n(m)}$$

and hence

$$\mathrm{d}\ln W_n = -\left(1-\lambda_{nn}\right)\mathrm{d}\ln m.$$

Thus, along the ray, the welfare effect of a proportional increase in all foreign prices relative to the home price depends only on the aggregate foreign expenditure share. It does not depend at all on how that expenditure is allocated across foreign sources.

The assumption that Hicksian demand is continuously differentiable in prices at every strictly positive price vector makes expenditure shares continuously differentiable, so we can define the home-foreign elasticity along the ray as

$$\varepsilon_n(m) \equiv \frac{\mathrm{d}\ln[X_{nn}(m)/X_{nF}(m)]}{\mathrm{d}\ln m}. \tag{5}$$

This is both the Hicksian and Marshallian elasticity of the home-to-foreign expenditure ratio.[5] The home-foreign elasticity determines how the domestic expenditure share changes along the ray. For any two values $m \geq m_0 \geq 1$, we have

$$\ln\frac{\lambda_{nn}(m)}{1-\lambda_{nn}(m)} = \ln\frac{\lambda_{nn}(m_0)}{1-\lambda_{nn}(m_0)} + \int_{m_0}^{m}\varepsilon_n(z)\mathrm{d}\ln z. \tag{6}$$

Hence, the initial domestic expenditure share, $\lambda_{nn}(m_0)$, and the elasticity schedule indexed by m, $\varepsilon_n(m)$, determine the path of domestic expenditure shares. If, in addition, the elasticity is strictly positive throughout the interval between m_0 and m, the domestic expenditure share is strictly increasing in m and can itself index the path. In this case, we write the elasticity schedule as $\varepsilon_n(\lambda) \equiv \varepsilon_n(m(\lambda))$. The next proposition gives welfare changes in the general case and an alternative expression indexed by the domestic expenditure share when the elasticity is positive. Proofs of all propositions and corollaries are provided in Appendix B.

Proposition 1. *Consider two points, indexed by* 0 *and* 1*, on the proportional foreign-price ray in equation* (2)*, corresponding to* $m = m_0$ *and* $m = m_1$*. Hold country* n*'s preferences, technology, and factor endowment fixed. Then*

$$\ln\frac{W_n^1}{W_n^0} = -\int_{m_0}^{m_1}\left[1-\lambda_{nn}(m)\right]\mathrm{d}\ln m. \tag{7}$$

[5] For any price vector $\boldsymbol{p}_n$, Hicksian demand at utility U is $U\boldsymbol{h}_n(\boldsymbol{p}_n)$ and Marshallian demand at expenditure Y_n is $[Y_n/e_n(\boldsymbol{p}_n)]\boldsymbol{h}_n(\boldsymbol{p}_n)$. The common scalar cancels from the ratio X_{nn}/X_{nF}, so its response to m is the same under Hicksian and Marshallian demand.

If, in addition, $\varepsilon_n(m) > 0$ along the portion of the ray connecting the two points, then

$$\ln \frac{W_n^1}{W_n^0} = \int_{\lambda_{nn}^1}^{\lambda_{nn}^0} \frac{\mathrm{d} \ln \lambda}{\varepsilon_n(\lambda)}. \tag{8}$$

Together, equation (6) and Proposition 1 show that, for given endpoints m_0 and m_1 on the proportional ray, welfare changes depend only on the initial domestic expenditure share and the one-dimensional home-foreign elasticity schedule. This schedule can be indexed by m in general and by the domestic expenditure share when $\varepsilon_n(m) > 0$. Elasticities governing substitution across foreign sources do not enter equations (7) and (8).

For a finite counterfactual along the ray, the terminal point is itself an outcome to be determined. The autarky counterfactual is special because it corresponds to $m \to \infty$. To focus on the case of finite gains from trade and avoid separating the autarky result into two cases (finite and infinite), assume that foreign goods are not jointly essential: one unit of utility can be attained using a finite quantity of the home good alone.[6] This assumption implies that the domestic expenditure share converges to one as m tends to infinity along the proportional foreign-price ray, without requiring a sign restriction on $\varepsilon_n(m)$.[7]

Corollary 1. *Let T denote an initial trade equilibrium at $m = 1$, and let A denote autarky for country n. Hold country n's preferences, technology, and factor endowment fixed. Suppose one unit of utility can be attained using a finite quantity of the home good alone. Then $\lambda_{nn}(m) \to 1$ as $m \to \infty$, $W_n^A > 0$, and*

$$\ln \frac{W_n^T}{W_n^A} = \int_1^{\infty} [1 - \lambda_{nn}(m)] \, \mathrm{d} \ln m. \tag{9}$$

If, in addition, $\varepsilon_n(m) > 0$ for every finite $m \geq 1$, then

$$\ln \frac{W_n^T}{W_n^A} = \int_{\lambda_{nn}^T}^{1} \frac{\mathrm{d} \ln \lambda}{\varepsilon_n(\lambda)}. \tag{10}$$

We define gains from trade as the log change in real income from autarky to the trade equilibrium. Equation (6) and Corollary 1 imply that calculating gains from trade requires only the initial domestic expenditure share and the home-foreign elasticity schedule along the proportional ray. When $\varepsilon_n(m) > 0$ for every finite $m \geq 1$, this schedule can be indexed by the domestic expenditure share from λ_{nn}^T to one.

Implementing the Proportional Foreign-Price Ray. The preceding welfare results use the proportional ray as a price path along which to calculate gains from trade. Can

[6] This is a condition on the level of utility at the boundary of the consumption set. It is distinct from standard Inada conditions, which restrict marginal utilities as consumption approaches the boundary.

[7] If this assumption is violated, equation (7) continues to apply on every finite portion of the ray, but real income converges to zero as $m \to \infty$, so gains from trade are infinite. Cobb-Douglas preferences over the domestic good and a foreign aggregate illustrate this case: the domestic expenditure share is fixed at the Cobb-Douglas parameter and therefore does not converge to one as the foreign relative price tends to infinity.

suitable trade-policy variation trace the ray in equilibrium and identify its elasticity schedule directly? If equilibrium variation could not trace the ray, recovering the one-dimensional schedule from price movements off the ray would generally require estimating a richer, high-dimensional demand system. Because wages respond endogenously to the policies used to change consumer prices, it is not obvious that equilibrium variation can trace the ray. We show that every finite $m \geq 1$ can be implemented using origin-specific tariffs or subsidies imposed by country n along with lump-sum rebates. These lump-sum transfers affect disposable income (factor income plus the net lump-sum transfer) at intermediate points on the ray, but the welfare formula for gains from trade compares the initial and autarky endpoints, at both of which tariff revenue and, therefore, the net transfer are zero.[8]

Taking country n's good as the numeraire, we represent the policy on imports from country i by a specific trade tax, measured in units of country n's good per unit of good i delivered to country n. Because the tax is specific, it enters the consumer price additively, with a negative value representing a subsidy. Net revenue is returned to the representative country-n household through a lump-sum transfer; the transfer is negative when subsidy payments exceed tariff revenue. The initial equilibrium has no tariffs. For each finite $m \geq 1$, we seek an equilibrium satisfying equation (2) for every foreign origin $i \neq n$. Preferences, productivities, and endowments remain fixed in every country, while wages and other equilibrium prices adjust. By Corollary 1, country n approaches autarky as m tends to infinity, and its imports and tariff revenue are zero at that endpoint.

Proposition 2. *Fix an initial equilibrium without tariffs at $m = 1$. Hold all primitive parameters fixed. For every finite $m \geq 1$, there exist origin-specific tariffs or subsidies imposed by country n (with revenue rebated lump sum) such that there is a competitive equilibrium satisfying equation (2) for every foreign origin $i \neq n$. At $m = 1$, the implementing trade taxes can be set to zero. As $m \to \infty$, country n approaches autarky, where imports and tariff revenue are zero.*

The implementation of each finite point does not require a sign restriction on $\varepsilon_n(m)$.

Fix a finite $m \geq 1$. Because wages (and, therefore, producer costs) are endogenous in general equilibrium and respond to changes in the economic environment, including country n's trade policy, the trade policy required to generate particular consumer-price ratios depends on the equilibrium wages induced by that same policy. This creates a fixed-point problem: the policies needed to implement the target consumer-price ratios depend on equilibrium wages, while those wages themselves depend on the policies. The argument can be understood as follows.

We construct an auxiliary economy with the same endowments, technologies, iceberg trade costs, and resource constraints as the original economy. In the auxiliary economy, every country $k \neq n$ has preferences u_k (as in the original economy), while country n's preferences differ from those in the original economy. Specifically, n's preferences are Leontief and chosen so that its consumption is proportional to $\boldsymbol{h}_n(\boldsymbol{\pi}_n(m))$, its Hicksian demand

[8] A similar approach can be used to establish feasibility using exogenous iceberg trade-cost changes rather than trade taxes. These trade-cost changes generate no fiscal revenue for country n, so disposable income equals factor income at every finite point along the ray. The construction guarantees $\mathcal{D}_{ni}(m) > 0$ but may yield $\mathcal{D}_{ni}(m) < 1$, which is why we prefer the trade-tax formulation.

under its original preferences at the target prices along the ray. An equilibrium of this auxiliary economy determines a vector of unit costs and a feasible allocation. Two features of the construction avoid fixed-point problems. First, we use the unit costs in the auxiliary economy without trade-policy wedges to define country n's trade policies that implement the target consumer-price ratios $\boldsymbol{\pi}_n(m)$ in the original economy. We then verify that the auxiliary-equilibrium wages and allocation are an equilibrium of the original economy under those policies. This approach avoids the loop, described in the previous paragraph, between trade policies and endogenous wages in the original economy. Second, Leontief preferences fix country n's consumption shares across goods independent of unit costs, which are themselves endogenous. This allows auxiliary preferences to be specified before those costs are determined and, therefore, avoids an additional fixed-point problem between consumption shares and endogenous unit costs within the auxiliary economy.

The auxiliary economy plays a role analogous to that of a planner's problem in an implementation argument. One could instead formulate a planner's problem subject to the relevant foreign-equilibrium and trade-balance constraints. However, a trade-balance constraint is restriction on values that must be evaluated at endogenous (supporting) prices. That alternative approach would therefore have to either solve for prices jointly with quantities or substitute equilibrium conditions that express prices as functions of quantities into the planner's problem before establishing a decentralization.[9] For our feasibility question, our auxiliary-equilibrium formulation is relatively direct. Country n's Leontief preferences impose the target Hicksian consumption shares across goods. And retaining the original endowments and the original preferences of every other country allows the equilibrium to deliver balanced trade, a feasible allocation, and the supporting unit costs needed to construct the implementing policies.

We prove the proposition in four steps. First, we invoke a standard equilibrium existence result for the auxiliary economy. Second, we use this equilibrium to construct the country-n policies that implement the target price ratios in the original economy. Third, we show that the resulting allocation is an equilibrium of the original economy under these policies: consumers and firms optimize, markets clear, and trade is balanced. Finally, we establish the claims concerning the initial equilibrium and the autarky limit. The proof is contained in Appendix B.

3 Weak Separability and Identification

The implementation result in Section 2 shows that suitable trade-policy variation can trace the proportional foreign-price ray under the general homothetic preferences considered there. Price changes available in the data do not generally follow this ray. In Section 3.1, we therefore introduce a multi-sector model that imposes weak separability between home and

[9] Kemp and Wan (1976) hold a customs union's aggregate net trade with non-members fixed, obtain supporting prices for an efficient allocation, and use those prices to construct a common external tariff vector. We fix the relative consumer prices faced by one country, use an auxiliary equilibrium to obtain the supporting unit costs, and construct origin-specific trade policies that implement those prices. See also Costinot et al. (2015) and Costinot et al. (2020) for related primal approaches that characterize allocations subject to foreign-equilibrium and trade-balance constraints and then construct taxes that implement them.

foreign goods within each sector and weak separability across sectors. Within-sector weak separability provides a foreign price aggregate for arbitrary, potentially non-proportional source-specific price changes. The ratio of that price index to the domestic price then plays the role of the index m of the price ray in Section 2. Hence, in spite of not being on the ray, we avoid estimating the full matrix of cross-price responses. These restrictions are not required for the welfare and autarky results in Section 2. Indeed, along the proportional price path described below, the gains-from-trade result in Section 2 applies to this model using the aggregate domestic expenditure share and the corresponding aggregate home-foreign elasticity.

We derive the formula for evaluating gains from trade in the multi-sector environment in Section 3.1. Section 3.2 then establishes conditions under which cross-country variation in tariffs and domestic expenditure shares identifies the sector-specific home-foreign elasticity schedules that enter this formula.

3.1 A Multi-Sector Model

Now let $s \in \mathcal{S}$ index sectors. Suppose preferences in country n are weakly separable across sectors, each sectoral sub-preference is homothetic, and the cross-sector aggregator is homothetic. At each nest, we maintain the utility and demand assumptions of Section 2. Within each sector, suppose preferences are weakly separable between a home good and a foreign aggregate, with homothetic sub-preferences within the foreign aggregate and homothetic preferences across the two aggregates. Let $e_{ns}(p_{nns}, P_{nFs})$ denote the sector-s unit expenditure function and let $\mathcal{E}_n((e_{ns})_{s\in\mathcal{S}})$ denote the cross-sector unit expenditure function, whose arguments are the sectoral unit expenditure functions. Weak separability implies that the allocation of sectoral spending between the home good and foreign aggregate depends only on p_{nns} and P_{nFs}. Hence, given weak separability we can study demand within sector s without specifying the allocation of total expenditure across sectors (or demand in other sectors).

The demand system above is stated for a representative consumer. However, the data we use in estimation measures total absorption, which aggregates across household consumption, government purchases, investment, and intermediate-input demand. We aim to estimate the elasticity schedules without separately modeling each user. To achieve this, we impose a common sourcing assumption within each destination-sector cell. Specifically, every user of sector s in country n buys the same sector-s composite at the same prices. If a indexes users and X_{ans} denotes user a's total expenditure in sector s, then expenditure on origin i is $X_{anis} = \lambda_{nis} X_{ans}$, where the origin share λ_{nis} is common across users. Summing across users therefore leaves the origin shares unchanged. This assumption concerns sourcing within sectors and does not require users to allocate expenditure across sectors in the same way. This assumption parallels Caliendo and Parro (2015), where the same country-sector composite is used for final consumption and as an intermediate input.

Let X_{nns} and X_{nFs} denote expenditure on the home good and foreign aggregate in sector s, let λ_{nns} denote the home expenditure share within sector s, and define the sectoral home-foreign elasticity in sector s as

$$\varepsilon_{ns}(\lambda_{nns}) \equiv \frac{\partial \ln(X_{nns}/X_{nFs})}{\partial \ln(P_{nFs}/p_{nns})},$$

which we assume is strictly positive throughout the relevant interval of sectoral domestic expenditure shares. This restriction allows the domestic expenditure share to index the relative-price path within each sector, as in the expenditure-share representation of Proposition 1. Weak separability between home and foreign goods allows us to treat demand within each sector as a choice between the home good and the foreign aggregate, whose price is P_{nFs}. The demand-side argument used to establish Proposition 1 then applies separately in each sector. In particular, expressing the sectoral unit expenditure function relative to the domestic sectoral price [10]

$$\mathrm{d}\ln\left(\frac{e_{ns}}{p_{nns}}\right) = \frac{\mathrm{d}\ln\lambda_{nns}}{\varepsilon_{ns}(\lambda_{nns})}. \tag{11}$$

Equation (11) determines the change in each sectoral price index relative to the corresponding domestic price. Combining these sectoral changes into aggregate welfare requires the cross-sector aggregator $\mathcal{E}_n$.

For two equilibria, indexed by 0 and 1, across which the sectoral and cross-sector preference shifters are fixed, define

$$I_{ns} \equiv \ln\left(\frac{e^1_{ns}}{p^1_{nns}}\right) - \ln\left(\frac{e^0_{ns}}{p^0_{nns}}\right),$$

as the change in the log sectoral unit expenditure function, denominated in units of country n's good in sector s, relative to the domestic sectoral price. Integrating equation (11) between the two equilibria gives

$$I_{ns} = \int_{\lambda^0_{nns}}^{\lambda^1_{nns}} \frac{\mathrm{d}\ln\lambda}{\varepsilon_{ns}(\lambda)}.$$

Equivalently, exponentiating the definition of I_{ns} yields

$$e^1_{ns} = e^0_{ns}\frac{p^1_{nns}}{p^0_{nns}}\exp(I_{ns}).$$

Therefore, without imposing a particular cross-sector aggregator,

$$\ln\frac{W^0_n}{W^1_n} = \ln\frac{Y^0_n}{Y^1_n} + \ln\frac{\mathcal{E}_n\left((e^0_{ns}(p^1_{nns}/p^0_{nns})\exp(I_{ns}))_{s\in\mathcal{S}}\right)}{\mathcal{E}_n\left((e^0_{ns})_{s\in\mathcal{S}}\right)}. \tag{12}$$

Suppose nominal income and all domestic sectoral prices change proportionally, so that $Y^1_n/Y^0_n = p^1_{nns}/p^0_{nns} = q_n$ for every sector s. This condition holds, for example, in a one-factor environment with fixed sectoral productivities. Homogeneity permits dividing equilibrium 1 nominal income and all prices by q_n without changing real income. Because all domestic-sector prices change by q_n, this normalization returns them to their equilibrium 0 levels, while sector s's unit expenditure becomes

$$\frac{e^1_{ns}}{q_n} = e^0_{ns}\exp(I_{ns}).$$

[10]The result follows from the same demand-side argument used in Proposition 1. For a two-good problem with home price p, foreign-aggregate price P, domestic expenditure share λ, and home-foreign elasticity $\varepsilon(\lambda)$, we have $\mathrm{d}\ln(e/p) = (1-\lambda)\mathrm{d}\ln(P/p)$ and $\mathrm{d}\ln[\lambda/(1-\lambda)] = \varepsilon(\lambda)\mathrm{d}\ln(P/p)$. Combining these expressions gives $\mathrm{d}\ln(e/p) = \mathrm{d}\ln\lambda/\varepsilon(\lambda)$.

Hence, using homogeneity of $\mathcal{E}_n$, equation (12) becomes

$$\ln \frac{W_n^0}{W_n^1} = \ln \frac{\mathcal{E}_n\left((e_{ns}^0 \exp(I_{ns}))_{s\in\mathcal{S}}\right)}{\mathcal{E}_n\left((e_{ns}^0)_{s\in\mathcal{S}}\right)}. \tag{13}$$

Equation (13) generally aggregates the sectoral unit-expenditure changes through $\mathcal{E}_n$. If we additionally assume that every foreign price changes by a common factor m relative to the corresponding domestic-sector price, then after the preceding normalization (i) the entire vector of domestic prices is fixed at its equilibrium 0 value, while (ii) the entire vector of foreign prices is multiplied by m. Relative prices are therefore fixed within the domestic price vector and, separately, the foreign price vector. Hence, the domestic goods and foreign goods can each be treated as a Hicks composite along this path. The full multi-sector problem then reduces to the proportional foreign-price-ray problem in Section 2, with the domestic composite playing the role of the home good. The representation indexed by m in Proposition 1 therefore applies using the aggregate domestic expenditure share and the corresponding aggregate home-foreign elasticity, without requiring the aggregate domestic expenditure share to be monotone along the ray.

The additional common-m restriction is not required to calculate gains from trade from equation (13). If equilibrium 0 is the trade equilibrium and equilibrium 1 is autarky, then $\lambda_{nns}^1 = 1$ in every sector. The gains-from-trade calculation generally depends on the cross-sector aggregator $\mathcal{E}_n$. As a familiar benchmark, suppose $\mathcal{E}_n$ is Cobb-Douglas with expenditure weights α_{ns}, a specification commonly used in quantitative trade models (e.g., Caliendo and Parro, 2015). Then

$$\ln \frac{W_n^T}{W_n^A} = \sum_{s\in\mathcal{S}} \alpha_{ns} \int_{\lambda_{nns}^T}^{1} \frac{\mathrm{d}\ln\lambda}{\varepsilon_{ns}(\lambda)}. \tag{14}$$

The baseline quantification in Section 5 instead uses the more flexible nested CES cross-sector aggregator described in Section 5.2.

3.2 Identification

Nonparametric Identification Under Cross-Country Invariance. Weak separability permits identification of the sectoral elasticity function from within-sector demand variation without specifying the cross-sector aggregator $\mathcal{E}_n$. One might imagine using alternative sources of variation to achieve this identification. For instance, one approach involves assuming that the elasticity schedule is stable over time within a given country-sector pair. Given this assumption, one could then leverage repeated observations over time within that pair to trace the elasticity schedule. We instead assume that the elasticity schedule is common across destinations within each sector. Hence, use cross-sectional variation in a single time change to trace the elasticity schedule within each sector. This approach requires only one time change and exploits variation across destinations in domestic expenditure shares and relative-price changes. The key cross-country restriction in our baseline is therefore $\varepsilon_{ns}(\lambda) = \varepsilon_s(\lambda)$ for every destination n within sector s. Because we estimate in changes, this restriction permits time-invariant differences in demand levels across destinations.

To connect this restriction to the variation used in estimation, let $\Delta\kappa_{ns}$ denote the change in the foreign-versus-home demand shifter in destination n and sector s. By the definition of the home-foreign elasticity, a first-order expansion around the initial domestic expenditure share gives

$$\Delta \ln \frac{X_{nns}}{X_{nFs}} = \varepsilon_s(\lambda_{nns,0}) \left(\Delta \ln \frac{P_{nFs}}{p_{nns}} - \Delta\kappa_{ns} \right) + \eta^H_{ns}, \tag{15}$$

where η^H_{ns} is the error from applying the local relationship over a finite change; Section 4.4 assesses this approximation by re-estimating the model using the exact relationship between the initial and terminal domestic expenditure shares.

To decompose the relative-price change in equation (15), we use a first-order approximation to the foreign price-index change,

$$\Delta \log P_{nFs} = \sum_{i \neq n} \mu_{nis,0} \Delta \log p_{nis} + \eta^P_{ns},$$

where $\mu_{nis,0}$ is origin i's initial-period share of country n's foreign expenditure within sector s, with the shares summing to one across foreign sources. The term η^P_{ns} contains the foreign-source demand-shifter component of the first-order approximation and the approximation error. One benefit of this approach is that it does not require an estimate of the across-foreign-source elasticity. Let $c_{is} = w_i/a_{is}$ denote the unit production cost of country i in sector s. Decompose productivity growth into a country average and a sector-specific deviation:

$$\Delta \log a_{is} = \overline{a}_i + \widetilde{a}_{is},$$

where $\overline{a}_i$ is average productivity growth across sectors in country i and $\widetilde{a}_{is}$ is sector s's deviation from that average. Define $C_i \equiv \Delta \log w_i - \overline{a}_i$. Unit-cost growth then satisfies

$$\Delta \log c_{is} = C_i - \widetilde{a}_{is},$$

where C_i is average unit-cost growth across sectors in country i and $-\widetilde{a}_{is}$ is sector s's deviation from that average. For each source country i, consumer prices satisfy

$$\Delta \log p_{nis} = \Delta \log(1 + \mathcal{T}_{nis}) + C_i - \widetilde{a}_{is} + \Delta \log \mathcal{D}_{nis},$$

where $\mathcal{T}_{nis}$ is the tariff and $\mathcal{D}_{nis}$ is the non-tariff trade cost and where $\mathcal{T}_{nns} = 0$ and $\mathcal{D}_{nns} = 1$. The tariff-induced component of the change in foreign prices relative to home prices is

$$Q^{\mathcal{T}}_{ns} \equiv \sum_{i \neq n} \mu_{nis,0} \Delta \log(1 + \mathcal{T}_{nis}).$$

Let $\boldsymbol{C} = (C_i)_i$ collect the country-level unit-cost shifters. The corresponding country-level unit-cost component of this relative-price change is

$$Q^C_{ns}(\boldsymbol{C}) \equiv \sum_{i \neq n} \mu_{nis,0}(C_i - C_n) = \sum_{i \neq n} \mu_{nis,0} C_i - C_n.$$

Define their sum as the tariff-and-country-cost component of the change in foreign prices relative to home prices:

$$Q_{ns}(\boldsymbol{C}) \equiv Q^{\mathcal{T}}_{ns} + Q^C_{ns}(\boldsymbol{C}) = \sum_{i \neq n} \mu_{nis,0} \left[\Delta \log(1 + \mathcal{T}_{nis}) + C_i - C_n \right].$$

Combining the previous expressions gives

$$\Delta \log P_{nFs} - \Delta \log p_{nns} = Q_{ns}(\boldsymbol{C}) + R_{ns},$$

where R_{ns} collects relative sector-specific unit-cost changes, changes in non-tariff trade costs, and the foreign-price-index residual:

$$R_{ns} \equiv -\sum_{i \neq n} \mu_{nis,0} \widetilde{a}_{is} + \widetilde{a}_{ns} + \sum_{i \neq n} \mu_{nis,0} \Delta \log \mathcal{D}_{nis} + \eta^P_{ns}.$$

Finally, substituting into equation (15) gives the estimating relationship underlying both the nonparametric identification argument and the parametric implementation:

$$\Delta \ln \frac{X_{nns}}{X_{nFs}} = \varepsilon_s(\lambda_{nns,0}) Q_{ns}(\boldsymbol{C}) + u_{ns}, \tag{16}$$

where

$$u_{ns} \equiv \varepsilon_s(\lambda_{nns,0}) \left(R_{ns} - \Delta\kappa_{ns}\right) + \eta^H_{ns}.$$

The error term u_{ns} combines sector-specific deviations in productivity growth from country averages, changes in non-tariff trade costs, and the foreign-price-index residual (collected in R_{ns}); changes in the home-foreign demand shifter (collected in $\Delta\kappa_{ns}$); and the finite-change approximation error η^H_{ns}.

The following proposition gives conditions under which variation in $Q^{\mathcal{T}}_{ns}$ identifies the elasticity schedule $\varepsilon_s(\lambda)$.

Proposition 3. *Fix a sector s and let $\mathcal{L}_s \subset (0,1)$ be an interval contained in the support of $\lambda_{nns,0}$ across destinations in that sector. Let V_{ns} collect predetermined shares and any additional controls. Suppose equation (16) holds for a cross-section of destination-sector cells after one time change. For each $\lambda \in \mathcal{L}_s$, suppose that*

$$\operatorname{Var}(Q^{\mathcal{T}}_{ns} \mid \lambda_{nns,0} = \lambda, Q^C_{ns}(\boldsymbol{C}), V_{ns}, s) > 0$$

and

$$\operatorname{Cov}(Q^{\mathcal{T}}_{ns}, u_{ns} \mid \lambda_{nns,0} = \lambda, Q^C_{ns}(\boldsymbol{C}), V_{ns}, s) = 0.$$

Then the elasticity function is identified on $\mathcal{L}_s$ as

$$\varepsilon_s(\lambda) = \frac{\operatorname{Cov}\left(\Delta \ln \frac{X_{nns}}{X_{nFs}}, Q^{\mathcal{T}}_{ns} \mid \lambda_{nns,0} = \lambda, Q^C_{ns}(\boldsymbol{C}), V_{ns}, s\right)}{\operatorname{Var}(Q^{\mathcal{T}}_{ns} \mid \lambda_{nns,0} = \lambda, Q^C_{ns}(\boldsymbol{C}), V_{ns}, s)}. \tag{17}$$

Proof. Conditional on $(\lambda_{nns,0}, Q^C_{ns}(\boldsymbol{C}), V_{ns}, s)$, take the covariance of equation (16) with $Q^{\mathcal{T}}_{ns}$. Because $Q^C_{ns}(\boldsymbol{C})$ is held fixed, $\operatorname{Cov}(Q^{\mathcal{T}}_{ns} + Q^C_{ns}(\boldsymbol{C}), Q^{\mathcal{T}}_{ns} \mid \cdot) = \operatorname{Var}(Q^{\mathcal{T}}_{ns} \mid \cdot)$. Orthogonality of $Q^{\mathcal{T}}_{ns}$ and u_{ns} then gives equation (17). □

The tariff component $Q^{\mathcal{T}}_{ns}$ is an observed component of the relative-price change that enters the regressor in equation (16). Conditional on the country-level unit-cost component, its variation identifies the coefficient on that regressor. A sufficient orthogonality condition is that $Q^{\mathcal{T}}_{ns}$ be conditionally orthogonal to R_{ns}, the demand-shifter change $\Delta\kappa_{ns}$, and the finite-change error η^H_{ns}. This condition permits $\Delta\kappa_{ns} \neq 0$; it requires only that this change be conditionally orthogonal to tariff-induced price variation.

The country-level unit-cost shifters entering $Q^C_{ns}(\boldsymbol{C})$ may be conditioned on directly or recovered jointly under a normalization and a rank condition. Each combines the country's wage change with its country-average productivity change, as defined above. Our baseline estimator uses the latter approach, while our main IV specification uses separately estimated country-cost components and conditions on them.

This nonparametric argument identifies $\varepsilon_s(\lambda)$ only on the range of domestic shares for which there is both support in $\lambda_{nns,0}$ and residual tariff variation within sector s. Calculating gains from trade requires the elasticity function between the observed trade share and one. If that interval extends beyond the support of the identifying variation, the gains calculation necessarily uses functional-form or shape restrictions to extrapolate the elasticity function. As shown in Figure 1 after introducing the data, both empirical samples contain substantial support in initial domestic expenditure shares near one, reducing the extrapolation required along this dimension.

Parametric Identification via a Polynomial Elasticity Schedule. In what follows, we introduce a polynomial approximation, which has the standard approximation rationale: because domestic expenditure shares lie in $[0, 1]$, any continuous elasticity schedule can be uniformly approximated arbitrarily well on this domain by polynomials. If the elasticity is bounded away from zero, a sufficiently accurate approximation also remains positive.

Proposition 3 treats $\varepsilon_s(\lambda)$ as an unrestricted function and identifies it locally at each domestic expenditure share using tariff variation across destinations within sector s. A finite-order polynomial replaces this collection of local identification problems with a finite-dimensional parameter vector, allowing observations at different initial shares to identify the elasticity schedule. Consider the following sector-specific approximation of order J,

$$\varepsilon_{s,J}(\lambda) = \sum_{j=0}^{J} \theta_{sj}\lambda^j. \tag{18}$$

The corresponding finite-change estimating equation is

$$\Delta \ln \frac{X_{nns}}{X_{nFs}} = \left[\sum_{j=0}^{J} \theta_{sj}\lambda^j_{nns,0}\right] [\Delta \ln P_{nFs} - \Delta \ln p_{nns}] + r^J_{ns},$$

where r^J_{ns} contains the demand-shifter change and finite-change error in equation (15), the error from approximating the true elasticity function by an order-J polynomial, and any measurement error. The coefficients are identified by the interactions of the tariff-induced foreign relative price change with powers of the initial domestic expenditure share, provided these interactions satisfy the corresponding exogeneity and rank conditions. Setting $J = 0$ recovers the constant-elasticity case, setting $J = 1$ allows the home-foreign elasticity to vary linearly with the domestic expenditure share, and higher values of J allow non-linear variation. For a fixed J, a polynomial is a flexible approximation to a smooth elasticity function on the supported range of shares. For better or for worse, a polynomial also allows us to extrapolate outside of this range of shares when the gains-from-trade calculation requires it.

Appendix C provides a quantitative Armington gravity model that delivers a linear $J = 1$ home-foreign elasticity schedule, strictly nesting the standard $J = 0$ constant-elasticity framework. Appendices D and E provide corresponding Krugman and Melitz-Pareto gravity models.

4 Estimating Home-Foreign Elasticities

Building on the identification results in Section 3.2, this section estimates sector-specific home-foreign trade elasticity schedules. We first describe the parametric implementation, then the data and measurement, and finally the estimation results and a range of sensitivity.

4.1 Parametric Implementation

Parametric Specification and Estimating Equation. As in Section 3.2, the empirical sample consists of one T-year change for each destination-sector cell. This time change removes the time-invariant component of destination-sector demand heterogeneity, while changes in demand shifters remain in the residual. We use the estimating relationship in equation (16), where Δ denotes this single time change and $\lambda_{nns,0}$ denotes the initial-period domestic expenditure share. Section 4.4 replaces the local approximation underlying that equation with the exact relationship as a sensitivity exercise.

Our baseline starts from the sector-specific elasticity schedule in equation (18), which is common across destination countries up to the initial domestic expenditure share. We make two additional parametric choices. First, we use the linear $J = 1$ specialization. Second, we partially pool slopes across sectors by imposing $\theta_{s1} = \theta_{g(s)1}$, where $g(s)$ denotes the aggregate containing sector s. The resulting home-foreign elasticity schedule, evaluated for country n and sector s at the domestic expenditure share λ_{nns}, is

$$\varepsilon_s(\lambda_{nns}) = \theta_{s0} + \theta_{g(s)1}\lambda_{nns}.$$

As in equation (18), the intercept and slope are common across destination countries. We allow the intercept to vary freely across sectors and impose that the slope is common across sectors within group $g(s)$. Across the baseline specifications, we vary the extent to which sectors are pooled into groups. We consider three sensitivity exercises related to these modeling choices. First, we relax linearity by moving to the quadratic $J = 2$ specialization. Second, we relax the cross-country invariance embedded in equation (18) by allowing the intercept to vary with observable country characteristics. Third, we relax the pooling of slopes across sectors by setting $g(s) = s$.

Specializing equation (16) to the linear elasticity schedule yields the estimating equation

$$\Delta \log \frac{X_{nns}}{X_{nFs}} = \left(\theta_{s0} + \theta_{g(s)1}\lambda_{nns,0}\right) Q_{ns}(\boldsymbol{C}) + u_{ns}, \tag{19}$$

where u_{ns}, defined in Section 3.2, collects the remaining relative-price component, the demand-shifter change, and the finite-change error.

Variation Used in Estimation. The initial foreign expenditure shares $\mu_{nis,0}$ weight the destination-origin-specific tariff changes $\Delta\log(1+\mathcal{T}_{nis})$ in $Q^{\mathcal{T}}_{ns}$ and country-cost change differences $C_i - C_n$ in $Q^{C}_{ns}(\boldsymbol{C})$. Hence, the tariff component, $Q^{\mathcal{T}}_{ns}$, predicts a larger relative decline in foreign-to-home prices (within destination-sector ns) when initial foreign expenditure is concentrated in sources whose tariffs decline more. Similarly, the country-cost component, $Q^{C}_{ns}(\boldsymbol{C})$, predicts a larger relative decline in foreign-to-home prices when initial foreign expenditure is concentrated in sources whose country-common unit costs decline more than the destination's, as measured by low C_i relative to C_n.

Conditional on the elasticity parameters, variation across destination-sector cells identifies the country-cost shifters. The country-cost component in cell ns equals the average of source-country shifters weighted by foreign expenditure less the destination's own shifter:

$$Q^{C}_{ns}(\boldsymbol{C}) = \sum_{i\neq n} \mu_{nis,0} C_i - C_n.$$

This component varies across cells because destinations and their initial foreign expenditure shares, $\mu_{nis,0}$, differ. For each candidate elasticity vector, we first choose the country-cost shifters that best fit the estimating equation and then evaluate the fit of that candidate using those chosen shifters.

At each candidate elasticity vector, the observed tariff component and the country-cost shifters chosen in this way determine $Q_{ns}(\boldsymbol{C})$, the tariff-and-country-cost component of the change in foreign prices relative to home prices. Variation across cells in $Q_{ns}(\boldsymbol{C})$ identifies the levels of the sectoral elasticity schedules, while variation in $\lambda_{nns,0}Q_{ns}(\boldsymbol{C})$ identifies how the home-foreign elasticities vary with the domestic expenditure share.

Moment Conditions and Identification. Here, we describe the moment conditions and identification restrictions underlying our baseline estimator. Section 4.4 relaxes many of these restrictions using two distinct instrumental-variable strategies.

The identifying assumptions can be written as moment conditions corresponding to the NLLS first-order conditions. We have

$$Q_{ns}(\boldsymbol{C}) = Q^{\mathcal{T}}_{ns} + \sum_{i\neq n} \mu_{nis,0}(C_i - C_n).$$

Because the country cost shifters enter $Q_{ns}(\boldsymbol{C})$ only through differences $C_i - C_n$, adding the same constant to every C_j leaves $Q_{ns}(\boldsymbol{C})$ unchanged, so the level of the country cost shifters is not identified. Hence, we normalize $C_{\text{USA}} = 0$. Collect the elasticity parameters $(\theta_{s0})_s$ and $(\theta_{g1})_g$ and the remaining country cost shifters in $\boldsymbol{\gamma}$, and let $Y_{ns} \equiv \Delta\log(X_{nns}/X_{nFs})$. The nonlinear regression function is

$$m_{ns}(\boldsymbol{\gamma}) = \left[\theta_{s0} + \theta_{g(s)1}\lambda_{nns,0}\right]\left[Q^{\mathcal{T}}_{ns} + \sum_{i\neq n} \mu_{nis,0}(C_i - C_n)\right],$$

where $C_{\text{USA}} = 0$. At the true parameter vector $\boldsymbol{\gamma}_0$, let $u_{ns} = Y_{ns} - m_{ns}(\boldsymbol{\gamma}_0)$ and let $D_{ns}(\boldsymbol{\gamma}_0) \equiv \partial m_{ns}(\boldsymbol{\gamma}_0)/\partial\boldsymbol{\gamma}$. The implemented nonlinear least-squares estimator uses the population moment conditions

$$\mathbb{E}\left[D_{ns}(\boldsymbol{\gamma}_0)u_{ns}\right] = 0.$$

Writing $\varepsilon_{ns,0} \equiv \theta_{s0} + \theta_{g(s)1}\lambda_{nns,0}$, the moment conditions can be expressed as

$$
\begin{aligned}
&\mathbb{E}\left[\mathbf{1}\{s=\ell\}Q_{ns}(\boldsymbol{C})u_{ns}\right]=0 && \text{for every sector } \ell,\\
&\mathbb{E}\left[\mathbf{1}\{g(s)=g\}\lambda_{nns,0}Q_{ns}(\boldsymbol{C})u_{ns}\right]=0 && \text{for every } g,\\
&\mathbb{E}\left[\varepsilon_{ns,0}\left(\mathbf{1}\{n\neq j\}\mu_{njs,0}-\mathbf{1}\{n=j\}\sum_{i\neq n}\mu_{nis,0}\right)u_{ns}\right]=0 && \text{for every } j\neq \text{USA}.
\end{aligned}
$$

The three lines are, respectively, the first-order conditions for the sector-specific intercepts $(\theta_{\ell 0})_\ell$, the group-specific slopes $(\theta_{g1})_g$, and the normalized country-cost shifters $(C_j)_{j\neq\text{USA}}$. The first condition requires the residual to be orthogonal to the tariff-and-country-cost shifter that identifies the elasticity schedule intercept. The second condition requires the residual to be orthogonal to interaction between the tariff-and-country-cost shifter and the initial domestic expenditure share that identifies the elasticity schedule slope. The third condition corresponds to each country-cost shifter C_j. Country j enters $Q_{ns}(\boldsymbol{C})$ with coefficient $\mu_{njs,0}$ when j is a foreign source and with coefficient $-\sum_{i\neq n}\mu_{nis,0}=-1$ when j is the destination. The derivative of the regression function with respect to C_j multiplies this coefficient by $\varepsilon_{ns,0}$, so the third condition requires the residual to be orthogonal to this derivative. Since C_j includes country-common productivity change, the restriction concerns the remaining sector-specific unit-cost component and the other terms in the residual.

These conditions can also be presented as a just-identified GMM system, with one moment for each parameter in $\boldsymbol{\gamma}$. The first two moments use the fitted tariff-and-country-cost shifter $Q_{ns}(\boldsymbol{C})$, including both its tariff and recovered country-cost components, as its own instrument. Thus, the NLLS criterion uses the fitted tariff-and-country-cost shifter, $Q_{ns}(\boldsymbol{C})$, as its own instrument; using $Q^{\mathcal{T}}_{ns}$ as the instrument would instead define a different nonlinear IV/GMM estimator.[11] The NLLS first-order conditions set all of these moments equal to zero. If they have a unique zero, NLLS and exactly identified GMM based on these moments therefore have the same solution. We take NLLS as the estimator because its objective determines these moment conditions.

A convenient sufficient condition for all of these moments is

$$\mathbb{E}\left[u_{ns}\mid Q^{\mathcal{T}}_{ns},\lambda_{nns,0},(\mu_{njs,0})_j,s\right]=0. \tag{20}$$

Equation (20) is stronger than the orthogonality condition used in Proposition 3. Equation (20) requires orthogonality with every NLLS derivative, including the country-exposure terms used to estimate C_j; but Proposition 3 requires orthogonality only with the tariff component conditional on $Q^{C}_{ns}(\boldsymbol{C})$. We require the stronger condition here because we jointly estimate the country cost shifters and the elasticity parameters by nonlinear least squares rather than identifying the elasticity parameters only from residualized tariff variation. In Section 4.4, we relax different parts of these exclusion restrictions using two alternative instrumentation strategies. The first uses a lagged tariff shifter and its interaction with the lagged domestic expenditure share to address correlated measurement error. The second combines tariff changes implied by pre-determined trade agreements with leave-sector-out

[11] Section 4.4 considers a related approach that constructs the country-cost component separately and uses agreement-implied tariff changes as excluded instruments for the resulting tariff-and-country-cost shifter.

estimates of country cost changes, implementing the core identification logic of Proposition 3 by instrumenting the tariff-and-country-cost shifter with excluded tariff variation.

For local identification and regular inference, the expected outer product

$$\mathbb{E}[D_{ns}(\boldsymbol{\gamma}_0)D_{ns}(\boldsymbol{\gamma}_0)']$$

must be nonsingular. We assess the sample analogue of this condition using numerical rank and conditioning diagnostics for the sample derivative matrix.

Estimator and Inference. We estimate equation (19) by unrestricted nonlinear least squares. Let $u_{ns}(\boldsymbol{\theta}, \boldsymbol{C})$ denote its residual for elasticity parameters $\boldsymbol{\theta}$ and normalized country cost shifters $\boldsymbol{C}$, where $C_{\text{USA}} = 0$. For every candidate $\boldsymbol{\theta}$, define

$$\widehat{\boldsymbol{C}}(\boldsymbol{\theta}) \equiv \arg\min_{\boldsymbol{C}} \sum_{n,s} u_{ns}(\boldsymbol{\theta}, \boldsymbol{C})^2.$$

Conditional on $\boldsymbol{\theta}$, the country cost shifters enter linearly, so $\widehat{\boldsymbol{C}}(\boldsymbol{\theta})$ is obtained by least squares. The unrestricted NLLS estimator of $\boldsymbol{\theta}$ is

$$\widehat{\boldsymbol{\theta}} \equiv \arg\min_{\boldsymbol{\theta}} \sum_{n,s} u_{ns}\Big(\boldsymbol{\theta}, \widehat{\boldsymbol{C}}(\boldsymbol{\theta})\Big)^2.$$

Thus, the numerical minimization is over only the sector-specific intercepts $(\theta_{s0})_s$ and group-specific slopes $(\theta_{g1})_g$. Appendix Section A.2 provides details of the numerical implementation.

Here, "unrestricted" means only that we do not impose the inequalities $\theta_{s0} \geq 0$ and $\theta_{s0} + \theta_{g(s)1} \geq 0$, which make each linear elasticity schedule nonnegative on $[0, 1]$. We use the unrestricted estimator as our baseline to show what the data call for under the maintained specification before imposing positivity. Section 4.4 reports how often the resulting fitted elasticities are negative at the observed initial domestic expenditure shares. For the specification used in the baseline quantification, positivity at those shares is sufficient for positivity along the path to autarky because the estimated slopes tend to be positive. Nevertheless, Section 4.4 also re-estimates the model imposing these restrictions in sensitivity; and the quantification uses the resulting estimates.

We report destination-country-clustered sandwich standard errors and apply the conventional finite-sample degrees-of-freedom adjustment, multiplying the sandwich covariance matrix by $G/(G-1) \times (N-1)/(N-K)$, where G is the number of destination-country clusters, N is the number of observations, and K is the total number of estimated elasticity and country-cost parameters (C_{USA} is excluded because it is normalized to zero). Standard errors for functions of the estimated parameters use the delta method, and confidence intervals use the corresponding normal approximation. Let $\widehat{\boldsymbol{\gamma}}$ denote the full estimated parameter vector, consisting of $\widehat{\boldsymbol{\theta}}$ and $\widehat{\boldsymbol{C}}(\widehat{\boldsymbol{\theta}})$. Because the country-cost shifters are estimated jointly with the elasticity parameters, the covariance matrix is computed using the full derivative vector $D_{ns}(\widehat{\boldsymbol{\gamma}})$, thereby accounting for estimation of the country-cost shifters.

Reported Functions of Estimates. We report each group-specific slope $\widehat{\theta}_{g1}$. To summarize the levels of the implied elasticity schedules, let $\mathcal{S}_g \equiv \{s : g(s) = g\}$ denote the set of sectors in group g, and let $\lambda_{p,s}$ denote the pth percentile of the initial domestic expenditure share across destination countries within sector s. For each group g, we report

$$\widehat{\varepsilon}_{p,g} = \frac{1}{|\mathcal{S}_g|} \sum_{s \in \mathcal{S}_g} \left(\widehat{\theta}_{s0} + \widehat{\theta}_{g1} \lambda_{p,s} \right), \qquad p \in \{25, 50, 75\}.$$

Thus, $\widehat{\varepsilon}_{p,g}$ is the unweighted average across sectors in group g of the fitted sectoral elasticity evaluated at each sector's own pth percentile of the domestic expenditure share.

4.2 Data and Measurement

This section describes the empirical objects required to estimate equation (19) and how we construct them using two separate data sets. One implementation uses the OECD ICIO tables for 1995 and 2007; the other uses GTAP 10 for 2004 and 2014. Both use tariff changes from the Global Tariff Database (Teti, 2024).

Unlike bilateral trade flows, sectoral domestic expenditure is not observed in a single harmonized international source. The OECD ICIO and GTAP construct such measures separately from national supply-use and input-output tables, supplemented and harmonized using additional sources and assumptions (Organisation for Economic Co-operation and Development, 2025b; Aguiar et al., 2019). Because these independently constructed data sets provide distinct measures of the dependent variable, we put them on equal footing throughout the empirical analysis.

In both data sets, $X_{nis,t}$ measures total absorption from origin i within sector s at time t by destination n summed across intermediate and final users within n. Under the common-sourcing assumption in Section 3.1, every user in a destination-sector cell has the same origin shares and home-foreign elasticity. Aggregating their expenditures therefore preserves the common domestic expenditure share and home-foreign elasticity that enter the estimating equation.

Objects to Be Measured. For each destination country, origin country, sector, and endpoint year, estimation requires bilateral expenditures $X_{nis,t}$ and bilateral tariffs $\mathcal{T}_{nis,t}$. Bilateral expenditures are used to construct the change in relative domestic expenditure, the initial domestic expenditure share $\lambda_{nns,0}$, and the initial foreign-origin shares $\mu_{nis,0}$. Bilateral tariffs are used to construct the tariff-induced foreign price shifter $Q^{\mathcal{T}}_{ns}$. The country-level unit-cost components C_i are estimated jointly with the elasticity parameters under the normalization $C_{\text{USA}} = 0$. In robustness, we allow the elasticity schedule to vary across countries as a function of destination country characteristics, $\boldsymbol{Z}_n$, which contains log PPP GDP per capita, log population, remoteness, and landlocked status.

OECD ICIO Measurement. Our first expenditure data set comes from the 2025 regular release of the OECD ICIO tables (Organisation for Economic Co-operation and Development, 2025b). We use 1995 and 2007, and the estimation sample includes 80 destination countries and the 19 OECD ICIO manufacturing sectors. For each origin-sector row, we construct destination absorption by summing over every intermediate-use and final-demand column belonging to the destination. The diagonal origin-destination cell gives $X_{nns,t}$, and the sum across all foreign origins gives $X_{nFs,t}$.

Tariff changes come from the ISIC two-digit file in the Global Tariff Database. We concord these industries to OECD ICIO sectors, averaging with equal weight when multiple ISIC industries map to one ICIO sector; details are provided in the Empirical Appendix. Following Teti's recommendation to use the winsorized tariff measures, we use `tariff95`, the unweighted average bilateral applied tariff winsorized at the 95th percentile (Teti, 2024). For each importer, exporter, and sector, we compute

$$\Delta\tau_{nis} = \log\left(1 + \frac{\mathtt{tariff95}_{nis,2007}}{100}\right) - \log\left(1 + \frac{\mathtt{tariff95}_{nis,1995}}{100}\right). \tag{21}$$

Let t_0 denote the initial year, which is 1995 for the OECD ICIO estimates. We construct origin weights using initial-year expenditure shares among foreign origins with tariff observations in both endpoint years:

$$\mu_{nis,0} = \frac{X_{nis,t_0}}{\sum_{j\in\mathcal{O}_{ns}} X_{njs,t_0}} \quad \text{for } i \in \mathcal{O}_{ns}, \tag{22}$$

where $\mathcal{O}_{ns}$ is the set of foreign origins with observed tariff values in both endpoint years for destination n and sector s, including origins whose tariff change is zero. Thus the origin shares are normalized to sum to one within each destination-sector cell over origins with tariff observations in both endpoint years.

GTAP Measurement. Our second expenditure data set comes from GTAP 10 (Aguiar et al., 2019). We use the benchmark years 2004 and 2014, the first and last years available, and the estimation sample includes 120 destination countries and 27 manufacturing sectors. For foreign expenditure, we use `VIMS`, which is defined as bilateral imports at importer market prices. For domestic expenditure on domestic output, we use the domestic-use analogue of the same absorption concept:

$$X_{nns,t} = \sum_u \mathtt{VDFM}_{sunt} + \mathtt{VDPM}_{snt} + \mathtt{VDGM}_{snt},$$

where `VDFM` is domestic purchases by firms and the capital-goods sector and we sum across using sectors u, `VDPM` is private household purchases of domestic goods, and `VDGM` is government purchases of domestic goods. We then set

$$X_{nFs,t} = \sum_{i\neq n} \mathtt{VIMS}_{sint}.$$

Tariff changes come from the GTAP-sector file in the Global Tariff Database. We construct $\Delta\tau_{nis}$ as in equation (21), replacing 1995 and 2007 with 2004 and 2014. We construct the initial-year bilateral-import shares as in equation (22), with $t_0 = 2004$.

In both data sets, $X_{nFs,t}$ includes imports from all foreign origins, including the rest of the world (ROW) region in the OECD ICIO and the residual regional aggregates in GTAP;

we refer to these regions collectively as ROW. Hence, imports from ROW enter both the dependent variable and $\lambda_{nns,0}$. Because ROW lacks matched bilateral tariff observations, we exclude it as an estimation destination and from the origin weights entering the tariff-and-country-cost shifter $Q_{ns}(\boldsymbol{C})$, which are normalized over origins with tariffs observed at both endpoints. Across destination-sector observations in each baseline sample, the median share of initial foreign expenditure accounted for by origins with tariffs observed at both endpoints is 98.94 percent in the OECD ICIO and 99.57 percent in GTAP.

The Distribution of $\lambda_{nns,0}$ and Extrapolation. The gains-from-trade calculation evaluates each sector's elasticity schedule from every destination's initial domestic expenditure share to one. Although the estimating equation identifies elasticities locally at observed initial shares, cross-country variation in those shares provides information about the elasticity schedule over the range relevant for the path to autarky, under the assumption that each sector's schedule is common across countries. The extent of the required extrapolation therefore depends on the support of the initial domestic expenditure shares.

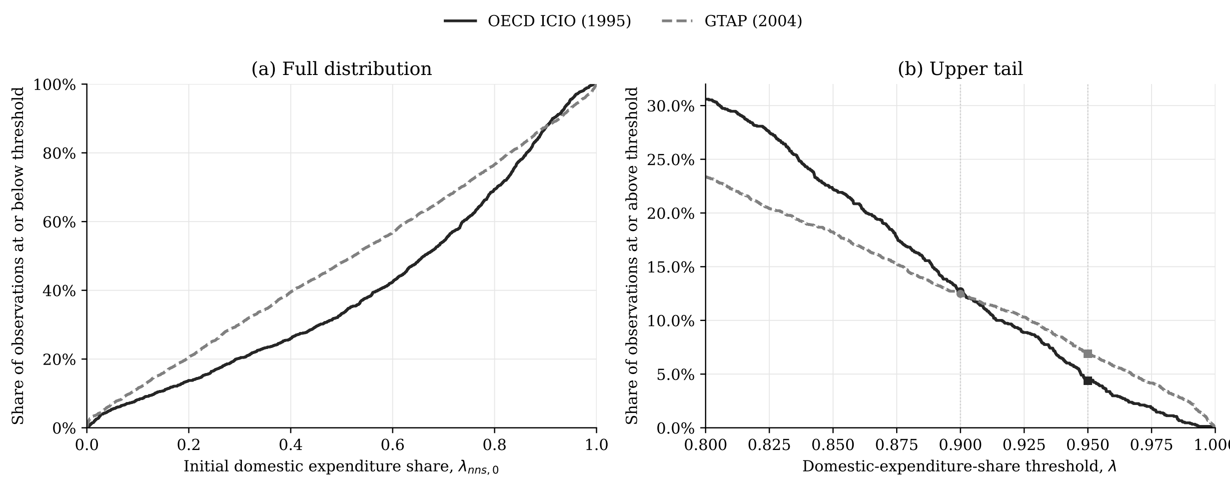

Figure 1: Distribution of Initial Domestic Expenditure Shares in Manufacturing

Notes: The figure plots the empirical distributions of the initial domestic expenditure share, $\lambda_{nns,0}$, across the manufacturing destination-sector observations used in the baseline estimates. The OECD ICIO distribution uses 1995 observations for 80 destinations and 19 manufacturing sectors; the GTAP distribution uses 2004 observations for 120 destinations and 27 manufacturing sectors. Each destination-sector observation receives equal weight. The left panel plots the empirical CDF: it reports the share of observations at or below each threshold. The right panel plots the corresponding upper-tail share: it reports the share at or above each threshold for thresholds of at least 0.8.

Figure 1 displays the distribution of initial domestic expenditure shares across the manufacturing destination-sector observations in the two baseline samples: the 1995 OECD ICIO data and the 2004 GTAP data. The data contain substantial support at high domestic expenditure shares. In the OECD ICIO and GTAP samples, respectively, 12.7 and 12.5 percent of observations have domestic expenditure shares of at least 0.90, while 4.4 and 6.9 percent have shares of at least 0.95. We conclude that the samples used to estimate the elasticity schedules include a substantial share of observations near one. Under the maintained restriction that $\varepsilon_{ns}(\lambda) = \varepsilon_s(\lambda)$, observations with higher initial domestic expenditure shares in a sector provide identifying variation for the elasticity schedule along the path to autarky for destinations that begin with lower shares. Within the range of shares observed in a sector, the functional form interpolates across observations. Functional-form extrapolation is required only for the portion of the path above the highest domestic expenditure share observed in that sector.

Additional Measurement. To allow the home-foreign elasticity to differ across broad types of manufacturing sectors, we split manufacturing into consumer and producer sectors separately in each data set. For each manufacturing sector, we calculate the global final-use share of absorption in the initial year and classify sectors weakly above the within-data-set median as consumer manufacturing. We also estimate seven-group specifications using analogous manufacturing aggregates in the two classifications. Appendix Tables 5 and 6 report the OECD ICIO and GTAP allocations, respectively.

In our baseline and in sensitivity exercises focusing on early periods ending in 2007, we use tariffs from the Global Tariff Database. In a sensitivity exercise, we use OECD ICIO data for 2004-2013 and the more product-disaggregated MAcMap-HS6 tariff data (Guimbard et al., 2012), which begin later. The OECD ICIO contains eight non-manufacturing goods sectors; seven have matched tariff products. GTAP contains 17 corresponding non-manufacturing goods sectors after omitting raw milk, which is not separately reported in the tariff data. We add these sectors in robustness and in the quantification.

For the country-characteristics sensitivity, PPP GDP per capita and population come from Penn World Table 10.01 (Feenstra et al., 2015), using 1995 values for the OECD ICIO exercise and 2004 values for the GTAP exercise. PPP GDP per capita is expenditure-side real GDP at chained PPPs divided by population. Remoteness is the corresponding-year partner-PPP-GDP-weighted mean of log population-weighted bilateral distance, using the full set of partner countries available in Penn World Table 10.01. Population-weighted bilateral distance and landlocked status come from CEPII GeoDist (Mayer and Zignago, 2011).

4.3 Estimation Results

Tables 1 and 2 report estimates of the group-specific slope $\theta_{g(s)1}$ and the unweighted average across sectors within each g of the fitted sectoral elasticities evaluated at each sector s's own 25th, 50th, and 75th percentiles of the initial domestic expenditure share. Table 1 uses data from the OECD ICIO and focuses on the period 1995-2007, whereas Table 2 uses data from GTAP and focuses on the period 2004-2014. The OECD ICIO estimates use 19 manufacturing sectors and 80 destination countries from 1995 to 2007; the GTAP estimates use 27 manufacturing sectors and 120 destination countries from 2004 to 2014.

Specification (A) estimates a common slope across all manufacturing sectors while allowing each sector to have its own elasticity intercept. The estimated slope is 5.45 in the OECD ICIO and 2.39 in GTAP, and both 95 percent confidence intervals exclude zero. In the OECD ICIO, the average fitted home-foreign elasticity rises from 4.01 at the sector-specific 25th percentiles of the domestic expenditure share to 5.88 at the 75th percentiles. The corresponding increase in GTAP is from 2.55 to 3.48.

Specification (B) allows the slope to differ between consumer and producer manufacturing. In the OECD ICIO, the consumer-manufacturing slope is 8.06 and its confidence

Table 1: OECD ICIO Estimates of the Home-Foreign Trade Elasticity, 1995-2007

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	5.45	4.01	5.12	5.88	0.26
		(1.88)	(1.12)	(1.32)	(1.49)	
(B)	Consumer manuf.	8.06	2.98	4.67	5.76	0.27
		(2.44)	(0.81)	(1.12)	(1.39)	
	Producer manuf.	2.05	5.19	5.60	5.89	
		(2.49)	(1.56)	(1.56)	(1.65)	
(C)	Food, bev., tob.	6.90	3.04	3.73	4.03	0.28
		(3.91)	(0.74)	(0.80)	(0.87)	
	Textiles/apparel/leather	12.35	0.84	4.19	6.11	
		(4.09)	(0.92)	(1.00)	(1.48)	
	Wood/paper/printing	-1.32	3.34	3.13	3.01	
		(4.15)	(1.04)	(0.87)	(0.98)	
	Petrol./chem./pharm./rubber	3.70	4.53	5.18	5.60	
		(2.80)	(1.30)	(1.45)	(1.61)	
	Minerals/metals	-1.48	7.18	6.90	6.71	
		(3.14)	(2.18)	(2.03)	(2.01)	
	Machinery/electrical	5.50	3.00	4.06	4.93	
		(2.99)	(1.07)	(1.22)	(1.50)	
	Transport/other manuf.	9.25	0.68	3.35	5.25	
		(2.92)	(0.72)	(0.97)	(1.46)	

Notes: This table reports the OECD ICIO 1995–2007 baseline, imposing $\varepsilon_{ns}(\lambda) = \theta_{s0} + \theta_{g(s)1}\lambda$ in the 19 manufacturing sectors and using the ISIC two-digit `tariff95` measure from the Global Tariff Database. Specification (A) estimates one manufacturing slope; (B) estimates consumer- and producer-manufacturing slopes; and (C) estimates seven manufacturing-aggregate slopes. Each sector has its own intercept. The column $\widehat{\theta}_1$ reports the slope for the aggregate shown. For each aggregate, $\widehat{\varepsilon}_p$ is the unweighted mean across sectors of the fitted sectoral elasticities evaluated at each sector's own pth percentile of the initial domestic expenditure share. Each specification uses 1,510 destination–sector observations. Standard errors are clustered by destination country and reported in parentheses.

interval excludes zero, while the producer-manufacturing slope is 2.05 and is imprecisely estimated. In GTAP, both estimates are positive and their confidence intervals exclude zero: 2.12 for consumer manufacturing and 2.71 for producer manufacturing. Specification (C) allows the slope to vary across seven manufacturing aggregates. Five estimates are positive and two are negative in the OECD ICIO; two of the positive estimates have confidence intervals that exclude zero, while neither negative estimate does. All seven GTAP estimates are positive, ranging from 1.76 for textiles, apparel, and leather to 3.64 for transport and other manufacturing; four have confidence intervals that exclude zero.

4.4 Empirical Sensitivity

We consider three broad types of sensitivity exercises. Our baseline specifications provide our preferred quantitative estimates. The exercises below show how the estimated magnitudes respond to alternative samples, specifications, variable constructions, and sources of

Table 2: GTAP Estimates of the Home-Foreign Trade Elasticity, 2004-2014

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	2.39	2.55	3.04	3.48	0.30
		(0.82)	(0.51)	(0.57)	(0.66)	
(B)	Consumer manuf.	2.12	3.01	3.47	3.88	0.30
		(0.90)	(0.59)	(0.63)	(0.71)	
	Producer manuf.	2.71	2.02	2.55	3.00	
		(0.93)	(0.48)	(0.55)	(0.65)	
(C)	Food, bev., tob.	2.04	3.95	4.40	4.79	0.30
		(1.43)	(0.81)	(0.82)	(0.91)	
	Textiles/apparel/leather	1.76	1.70	2.10	2.55	
		(0.96)	(0.40)	(0.42)	(0.56)	
	Wood/paper/printing	1.93	0.82	1.18	1.47	
		(1.30)	(0.34)	(0.38)	(0.50)	
	Petrol./chem./pharm./rubber	2.34	1.68	2.30	2.66	
		(1.05)	(0.41)	(0.51)	(0.62)	
	Minerals/metals	2.79	2.63	3.14	3.64	
		(1.34)	(0.80)	(0.81)	(0.89)	
	Machinery/electrical	3.21	2.30	2.69	3.17	
		(1.21)	(0.61)	(0.62)	(0.67)	
	Transport/other manuf.	3.64	1.92	2.61	3.26	
		(1.15)	(0.51)	(0.55)	(0.66)	

Notes: This table reports results using GTAP data and the GTAP-sector `tariff95` measure from the Global Tariff Database over the single 2004–2014 time change, imposing $\varepsilon_{ns}(\lambda) = \theta_{s0} + \theta_{g(s)1}\lambda$. Specification (A) estimates one manufacturing slope; (B) estimates consumer- and producer-manufacturing slopes; and (C) estimates seven manufacturing-aggregate slopes. Each sector has its own intercept. The column $\widehat{\theta}_1$ reports the slope for the aggregate shown. For each aggregate, $\widehat{\varepsilon}_p$ is the unweighted mean across sectors of the fitted sectoral elasticities evaluated at each sector's own pth percentile of the initial domestic expenditure share. Each specification uses 3,240 destination–sector observations. Standard errors are clustered by destination country and reported in parentheses.

identifying variation. We use both the magnitudes of the resulting estimates and the persistence of the central qualitative pattern—that home-foreign elasticities generally increase with domestic expenditure shares—to assess robustness. First, we examine alternative samples, specifications, and variable constructions. Second, we consider potential endogeneity. We use lagged variables to address correlated measurement error in the dependent and independent variables. Separately, we combine tariff changes implied by predetermined trade agreements with leave-sector-out estimates of country cost changes to address potential endogeneity in the tariff-and-country-cost shifter, thereby returning to the tariff-based identification logic of Proposition 3. Third, Monte Carlo exercises ask whether the estimator mechanically generates positive elasticity-schedule slopes when the true schedules are flat and whether it correctly identifies positive slopes when the true schedules are linear. We generate the model data using the quantitative multi-sector version of the extended Armington model in Appendix C, in which the home-foreign elasticity can differ from the elasticity across foreign sources and is linear in the domestic absorption share.

Alternative Samples and Specifications. The empirical appendix reports parallel robustness exercises for both data sets, generally replicating Specifications (A)-(C). The first sensitivity exercise expands the estimation sample to incorporate non-manufacturing goods-producing sectors. This asks whether the positive relationship between the home-foreign elasticity and the domestic expenditure share applies more broadly than manufacturing. Appendix Tables 7 and 22 report the results.

The second set of sensitivity exercises examines how the baseline restrictions on variation in the elasticity schedule across sectors and countries affect the results. In one exercise, we allow the elasticity schedule's intercept to vary with destination-country characteristics, checking whether the restriction that the intercept is common across destinations confounds variation in the domestic expenditure share with systematic differences across countries in income, size, or geography. This exercise replaces the baseline schedule, $\varepsilon_s(\lambda)$, with

$$\varepsilon_{ns}(\lambda) = \underbrace{\theta_{s0} + \boldsymbol{Z}_n'\boldsymbol{\delta}}_{\theta_{ns0}} + \theta_{g(s)1}\lambda,$$

where $\boldsymbol{Z}_n$ contains PPP GDP per capita, population, remoteness, and landlocked status. Appendix Tables 8 and 23 report the results. We next restrict variation across sectors by pooling the intercepts within sector aggregates. This imposes more structure than the baseline and asks whether the slope estimates rely on allowing every sector to have its own elasticity level (Appendix Tables 9 and 24). Moving in the opposite direction, we allow the elasticity-schedule slope to vary freely by sector, which asks whether the conclusion that the home-foreign elasticity increases with the domestic expenditure share survives without pooling slopes within manufacturing-sector aggregates (Appendix Tables 10 and 25).

A third set examines the construction of the tariff shifter. The baseline tariff shifter, $Q^{\mathcal{T}}_{ns}$, uses only foreign origins with observed tariff levels in both end years and defines $\mu_{nis,0}$ by renormalizing initial foreign expenditure shares to sum to one across those origins within each destination-sector cell. To assess sensitivity to incomplete tariff coverage, we re-estimate the model retaining only destination-sector cells in which these origins account for at least 50 percent of total initial foreign expenditure. This exercise asks whether the results rely on cells in which the tariff shifter represents only a small portion of foreign expenditure; Appendix Tables 11 and 26 report the results. We next use trade-weighted rather than unweighted tariffs. This allows us to determine whether the results depend on how product-level tariffs are aggregated (Appendix Tables 12 and 27). Finally, we repeat the analysis using MFN tariffs. This allows us to deal with any endogeneity induced from applied tariffs including bilateral preferential rates that may respond endogenously to productivity or other sectoral shocks (Appendix Tables 13 and 28).

A fourth set examines the parametric specification and the influence of individual observations. We replace the linear elasticity schedule with a quadratic one, thereby relaxing the assumption that its slope is constant over the observed range of domestic expenditure shares (Appendix Tables 14 and 29). Because a quadratic schedule has no single slope, these tables replace $\widehat{\theta}_{g1}$ with the derivative of the fitted schedule evaluated at the average of the sector-specific median domestic expenditure shares within each group, giving every sector equal weight. Finally, we re-estimate Specifications (A)-(C) after omitting each destination country and each manufacturing sector in turn to determine whether an influential country or sector drives the results (Appendix Figures 4 and 5).

The magnitudes of individual slope estimates and fitted elasticities typically change modestly across these exercises, with some appreciable changes. Across the exercises, the home-foreign elasticity generally increases with the domestic expenditure share. The common slope in Specification (A) remains positive in every sensitivity that retains the linear schedule in both data sets. When slopes vary freely by sector, 14 of the 19 OECD ICIO estimates and 23 of the 27 GTAP estimates are positive. The quadratic estimates generally imply increasing elasticities over the supported range, and the leave-one-out estimates show that no individual country or sector drives the positive Specification (A) slope.

A fifth set of exercises restricts estimates to impose positivity of the elasticity schedule and addresses the finite-change approximation. The baseline empirical analysis does not impose positivity on the elasticity schedules. In Specification (B), for example, the fitted elasticity at the initial domestic expenditure share is negative for only 24 of 1,510 OECD ICIO manufacturing destination-sector observations and 89 of 3,240 GTAP observations. Because both estimated slopes are positive in each data set, the fitted elasticities in all remaining observations stay positive as their domestic expenditure shares rise toward one. For quantification, we use the expenditure-share representation with a positive elasticity along the full path to autarky for every observation. We therefore re-estimate the model imposing $\theta_{s0} \geq 0$ and $\theta_{s0} + \theta_{g(s)1} \geq 0$ for every sector. The estimates are broadly robust to imposing these positivity restrictions. Appendix A.4 provides details of the constrained estimator, and Appendix Tables 15 and 30 report the estimates used in Section 5.

We next retain these constraints, but replace the elasticity evaluated at the initial domestic expenditure share with the exact finite-change relationship. The exact estimator solves demand forward from the observed initial share for each potential elasticity schedule and vector of country cost changes. It minimizes the difference between the predicted and the observed terminal log home-foreign expenditure ratios. Appendix A.4 provides details of the exact finite-change estimator, and Appendix Tables 16 and 31 report the results. Comparing the exact estimates with the corresponding constrained initial-share estimates, the common-manufacturing slope changes from 3.90 to 4.04 in the OECD ICIO and from 1.78 to 1.91 in GTAP, while the signs of every group-specific slope in Specifications (B) and (C) remain unchanged. Thus, the initial-share approximation does not drive the estimated positive relationship.

Inference under the positivity restrictions would require methods that account for parameters on or near the boundary. Andrews (1999) derives asymptotic distributions for parameters on the boundary, while Chernozhukov et al. (2023) develops uniformly valid inference for constrained conditional moment models. Because we use the constrained estimates only as quantitative inputs or sensitivity point estimates and do not report standard errors or confidence intervals for them, we do not conduct constrained inference.

Finally, the longer OECD ICIO time series allows us to check the extent to which the OECD baseline results depend on using the earlier period than used in the GTAP sample. Appendix Table 17 repeats the OECD ICIO analysis using expenditure changes from 2004 to 2013 and the more disaggregated MAcMap-HS6 tariff changes, as described in Section 4.2; we end this analysis in 2013 because it is the closest available MAcMap-HS6 benchmark year to 2014. The common-manufacturing slope is 8.77 and its confidence interval excludes zero. The consumer- and producer-manufacturing slopes are 5.37 and 10.66, and all seven slopes in Specification (C) are positive. This exercise shows that the positive relationship in the OECD ICIO data is equally evident over a sample period closely aligned with the GTAP analysis.

Endogeneity I: Correlated Measurement Error. We first address a concern about measurement error. The initial domestic expenditure share, $\lambda_{nns,0}$, and the initial home-foreign expenditure ratio are constructed, in part, from the same measurements, since $\log(X_{nns,0}/X_{nFs,0})$ can be re-expressed as $\log[\lambda_{nns,0}/(1-\lambda_{nns,0})]$. This enters the dependent variable with a negative sign because $\Delta \log(X_{nns}/X_{nFs})$ subtracts the initial expenditure ratio. Consequently, measurement error that (e.g.) overstates initial domestic expenditure relative to foreign expenditure raises the measured initial domestic expenditure share but lowers the measured change in the home-foreign expenditure ratio. Measurement errors in the dependent variable and in $\lambda_{nns,0}$ are therefore mechanically negatively correlated. Because the regressor that identifies the elasticity-schedule slope is $\lambda_{nns,0}Q_{ns}(\boldsymbol{C})$, this correlation can bias the estimated slope even if the price shifter itself is measured without error.

Appendix Table 18 uses 1995 data to construct lagged instruments for an estimating equation based on changes from 1996 to 2007, thereby separating the initial-year measurements entering the dependent variable from those entering the instruments. Both panels use expenditure and tariff changes from 1996 to 2007, rather than our baseline 1995 start date. In the instrumented panel, the 1995-2007 tariff shifter instruments the 1996-2007 tariff shifter, while the interaction of the 1995-2007 tariff shifter with the 1995 domestic expenditure share instruments the interaction of the 1996-2007 tariff shifter with the 1996 share.[12] The instrument for the slope regressor therefore does not use the 1996 expenditure measurements that enter the dependent variable. The common slope changes from 4.87 without instruments to 5.81 with instruments, while the average fitted elasticity at the sector-specific median domestic expenditure shares, $\widehat{\varepsilon}_{50}$, changes from 5.66 to 7.03. To the extent that measurement errors in the 1995 and 1996 domestic expenditure shares are uncorrelated, and the lagged tariff instruments otherwise satisfy the exclusion restriction, the similar magnitudes of both estimates suggest that correlated measurement error is not driving our results.

Endogeneity II: The Price Shifter. We next address the possibility that destination-sector shocks entering the residual are correlated with either applied tariff changes or the country-cost changes estimated jointly with the elasticity parameters. We construct a first tariff instrument from changes required by the May 2004 accession of ten countries to the European Union. The entrants eliminated tariffs on imports from the other members of the enlarged European Union and adopted the common external tariff on imports from other origins, while the incumbent EU15 eliminated tariffs on imports from the entrants. A second, broader instrument adds changes associated with the 2007 EU enlargement—including the elimination of EU15 tariffs on imports from Bulgaria and Romania—the EU-Türkiye customs union, the Gulf Cooperation Council, and the West African Economic and Monetary Union. In constructing each instrument, we set bilateral tariff changes not covered by these agreements to zero. We aggregate using exactly the same initial foreign expenditure shares, $\mu_{nis,0}$, as in $Q^{\mathcal{T}}_{ns}$. We identify agreement membership and implementation dates from the WTO Regional Trade Agreements Database and official institutional sources

[12] The first year with coverage in the OECD ICIO data is 1995.

and construct the implied sectoral tariff changes from the Global Tariff Database (World Trade Organization, 2025; Teti, 2024); Appendix A.1 provides details.

For each specification and sector s, we re-estimate the unrestricted NLLS model after excluding all observations in sector s. Within this auxiliary estimation, we choose the country cost shifters by least squares at every candidate elasticity vector for the retained sectors and then choose the elasticity vector that minimizes the profiled sum of squared residuals. Let $\widehat{\boldsymbol{C}}^{(-s)}$ denote the country cost shifters evaluated at this leave-sector-out NLLS optimum. We define

$$\widehat{Q}_{ns}^{(-s)} \equiv Q_{ns}\left(\widehat{\boldsymbol{C}}^{(-s)}\right)$$

Excluding sector s constructs the country-cost component for observation (n, s) using outcomes from the other sectors. This prevents the outcome in (n, s) from feeding mechanically into its own constructed price shifter. In spite of this, it is possible that the leave-sector-out country-cost estimate may remain correlated with the residual. For this reason, we treat the constructed price shifter as endogenous in the instrumental-variable estimation. At the same time, retaining the other sectors for destination n allows its own cost change, $\widehat{C}_n^{(-s)}$, to be identified from the role of the domestic price in those sectors.

The baseline estimator is nonlinear because the elasticity parameters and country-cost changes are estimated jointly: changing an elasticity parameter changes the country-cost estimates, which in turn changes the tariff-and-country-cost shifter. However, once $\widehat{\boldsymbol{C}}^{(-s)}$ has been estimated from the other sectors, we hold $\widehat{Q}_{ns}^{(-s)}$ fixed; hence, here estimation becomes linear. Specifically, let $\boldsymbol{C}_0$ denote the true vector of country cost changes, so that $Q_{ns}(\boldsymbol{C}_0)$ is the true tariff-and-country-cost shifter. Proxying $Q_{ns}(\boldsymbol{C}_0)$ with $\widehat{Q}_{ns}^{(-s)}$, the resulting estimating equation is linear in the elasticity parameters:

$$\Delta \ln \frac{X_{nns}}{X_{nFs}} = \theta_{s0}\widehat{Q}_{ns}^{(-s)} + \theta_{g(s)1}\lambda_{nns,0}\widehat{Q}_{ns}^{(-s)} + \widetilde{u}_{ns},$$

where

$$\widetilde{u}_{ns} = u_{ns} - \varepsilon_{ns,0}\left[\widehat{Q}_{ns}^{(-s)} - Q_{ns}(\boldsymbol{C}_0)\right].$$

We estimate it by two-stage least squares, instrumenting $\widehat{Q}_{ns}^{(-s)}$ and $\lambda_{nns,0}\widehat{Q}_{ns}^{(-s)}$ with the agreement-implied tariff shifter and its interaction with $\lambda_{nns,0}$. This construction implements the core identification logic of Proposition 3: the agreement-implied tariff shifter supplies excluded tariff variation after the country-cost component has been constructed separately. The leave-sector-out country-cost component appears in the endogenous regressors but not in the excluded instruments.

The exclusion restriction requires the agreement-implied tariff instrument and its interaction with $\lambda_{nns,0}$ to be orthogonal to the composite error $\widetilde{u}_{ns}$. A sufficient condition is that the agreement-implied tariff changes are unrelated to non-price destination-sector shocks, so that their effect on relative home and foreign expenditure operates through relative prices. They must also be unrelated to the component of the country-cost proxy error inherited from the retained sectors. Because the instruments are constructed using initial foreign expenditure shares and one is interacted with the initial domestic expenditure share, these shares must also be predetermined with respect to the composite error. Cross-sector correlation in residuals does not by itself violate the exclusion restriction; it is a concern only if the agreement instruments are correlated with shocks in other sectors that enter the country-cost proxy.

To account for uncertainty from estimating the leave-sector-out country-cost components, we use a destination-cluster bootstrap that re-estimates the leave-sector-out NLLS models, country-cost changes, and linear IV model in each draw. Appendix A.3 describes the bootstrap procedure. Appendix Tables 19 and 20 report the full-sample estimates, medians across 999 bootstrap draws, and 95 percent confidence intervals.

Each table additionally reports the smallest destination-clustered Sanderson-Windmeijer conditional first-stage F statistic across all endogenous regressors (Sanderson and Windmeijer, 2016; Windmeijer, 2024). For the EU-enlargement instrument, these minima are 33.18, 14.83, and 16.44 in Specifications (A), (B), and (C); for the broader instrument, they are 20.01, 22.41, and 7.36. Table 21 reports the full vector for every specification and instrument. As a simpler benchmark, in Specification (A), conventional scalar first-stage regressions of $\widehat{Q}_{ns}^{(-s)}$ on the agreement-implied tariff shifter, controlling for destination and sector effects, yield destination-clustered F statistics of 104.2 for the EU-enlargement instrument and 111.9 for the broader instrument. Thus, the scalar first-stage statistics indicate a strong first-stage relationship.[13]

The common-manufacturing slope is 8.83 using the EU-enlargement instrument and 10.75 using the broader customs-union instrument; the corresponding bootstrap medians are 5.28 and 6.91. The 95 percent bootstrap confidence interval is $[-1.77, 18.08]$ using the EU-enlargement instrument and $[1.20, 14.79]$ using the broader instrument; the latter excludes zero. The full-sample fitted-elasticity profiles in Specification (A) are quantitatively similar to the baseline. The triplet $(\widehat{\varepsilon}_{25}, \widehat{\varepsilon}_{50}, \widehat{\varepsilon}_{75})$ equals $(4.01, 5.12, 5.88)$ in the baseline, $(3.62, 5.42, 6.65)$ using the EU-enlargement instrument, and $(2.76, 4.96, 6.45)$ using the broader instrument. Each baseline value lies within the corresponding 95 percent bootstrap confidence interval under both instruments. The agreement-IV estimates therefore place the fitted elasticity near the center of the observed domestic-share distribution at a level similar to the baseline and imply a steeper increase across that distribution.

Both consumer- and producer-manufacturing slope estimates and their bootstrap medians are positive with either instrument. With the broader instrument, the confidence interval excludes zero for the consumer-manufacturing slope, $[0.62, 12.97]$, and has a lower endpoint of -0.16 for the producer-manufacturing slope, $[-0.16, 19.41]$; both intervals include zero with the EU-enlargement instrument. In Specification (C), six of seven point estimates are positive under each instrument; under the broader instrument, the remaining estimate is -0.18, or essentially zero. The average fitted elasticities at the sector-specific 50th and 75th percentiles are positive for every group under both instruments. At the 25th percentiles, they are negative for machinery/electrical and transport/other manufacturing in Specification (C) under both instruments.

We conduct this exercise only with the OECD ICIO data because the 2004 GTAP tariffs already incorporate the May 2004 accession changes.

Monte Carlo Exercises. Appendix H reports Monte Carlo exercises designed to answer two questions. First, does our estimation procedure mechanically generate positive elasticity-schedule slopes when the true slopes are zero? Second, does it accurately recover

[13]The EU-enlargement instrument is nonzero in 25 destination clusters, and the broader instrument is nonzero in 32 destination clusters.

positive slopes when the data are generated by a model featuring linear home-foreign elasticity schedules with positive slopes? We answer these questions using the quantitative multi-sector version of the extended Armington model developed in Appendix C. We consider four calibrations. Each calibration matches the same initial equilibrium and therefore the same domestic share in every country-sector pair, but differs in its home-foreign elasticity schedules. For each calibration, we use the same sequence of 200 independently generated tariff-shock matrices. To construct each matrix, we select every destination-origin-sector cell with a foreign origin independently with probability 0.45. In each selected cell, we reduce the tariff by a random amount. The selected cells and the magnitudes of their tariff reductions vary across the 200 draws, while the tariff-shock matrix for a given draw is held fixed across the four calibrations. Conditional on the calibrated economy, the variation across draws comes only from the incidence and magnitude of the tariff reductions. For each of the 200 matrices and each calibration, we then solve for the new equilibrium, and estimate Specifications (A)-(C) in the model-generated data using our baseline procedure. We report the distribution of estimates.

The first calibration determines whether our estimation procedure mechanically generates positive elasticity-schedule slopes when the true slopes are zero. It sets $\theta_{s1} = 0$ in every sector, while allowing the resulting constant home-foreign elasticity, θ_{s0}, to vary across sectors. We also set θ_{s0} equal to the elasticity of substitution across foreign sources. Under these restrictions, demand is CES across all origins within each sector, yielding the standard constant-elasticity expenditure-share system of many-sector Armington and Eaton-Kortum models. Table 34 reports estimates of Specifications (A)-(C) using the model-generated data. The mean estimated slope in Specification (A) is 0.001; the estimates in Specifications (B) and (C) are likewise close to and centered around zero. We conclude that the estimation procedure does not mechanically produce the pattern found in the data within canonical constant-elasticity environments.

The remaining three calibrations determine whether our estimation procedure accurately recovers positive slopes when the data are generated by linear home-foreign elasticity schedules with positive slopes. The second calibration sets the slope equal to 4 in every sector. The third assigns slopes of $2.5, 3.0, \ldots, 5.5$ to the seven manufacturing groups used in Specification (C). The fourth allows every sector to have a distinct positive slope. Tables 35-37 report the results. When the pooling imposed by the estimator matches the calibration, the estimates closely recover the true group-specific slopes. For every pooled group within which true sector slopes vary, all 200 draw-specific slope estimates underlying the averages reported in the tables remain positive and lie within the range of true slopes in that group. We conclude that our estimation procedure accurately recovers elasticity-schedule slopes in the generalized Armington environment.

5 Quantification

In this section and the Quantitative Appendix, we quantify gains from trade relative to autarky, building on equation (13), but now allowing traded goods to be used as intermediate inputs. We take equilibrium 0 to be the 2007 trade equilibrium when using the OECD ICIO data and estimates and the 2014 trade equilibrium when using GTAP. Equilibrium 1 is autarky, so that $\lambda^1_{nns} = 1$ in every traded sector. We therefore raise each sector's domestic expenditure share from its observed value in the corresponding trade equilibrium to one and calculate the resulting increase in aggregate unit expenditure. This increase determines the ratio of real income in the initial trade equilibrium to real income under autarky. We focus on the OECD ICIO quantification in the body and report the corresponding GTAP exercise in Appendix G.

According to the theory in Section 3.1, the increase in each traded sector's unit expenditure function as its domestic expenditure share rises to one is determined by its initial domestic expenditure share and its home-foreign elasticity schedule. As also shown there, translating these sectoral changes into aggregate welfare requires the demand system that aggregates consumption across sectors. Section 3.1 allows us to estimate the elasticity schedules from total absorption by imposing common sourcing across intermediate and final users within each sector, but it leaves their roles in production and welfare unspecified. For quantification, intermediate-input use determines how sectoral cost changes feed into production costs, while the treatment of final uses connects total absorption to household welfare. We next specify this additional structure.

We retain 26 goods-producing sectors from the OECD ICIO: 19 manufacturing sectors and seven non-manufacturing goods sectors with matched tariff data. Let $\mathcal{S}_M$ and $\mathcal{S}_G$ denote these two sets. We combine the remaining sectors, including the non-manufacturing goods sector without matched tariffs, into a single composite, indexed by 0, which we treat as non-traded. For manufacturing, we use Specifications (A), (B), and (C) of Appendix Table 15. To obtain the complete set of schedules required for quantification (which includes the non-manufacturing goods sectors not included in the baseline), we estimate the corresponding constrained specifications on all goods, retain the manufacturing schedules from the manufacturing-only estimates, and use the all-goods estimates for the non-manufacturing schedules. Our primary exercise uses Specification (B): the 19 manufacturing-sector intercepts and consumer- and producer-manufacturing slopes come from Appendix Table 15, while the seven non-manufacturing-sector intercepts and common non-manufacturing slope come from the corresponding constrained all-goods specification.

5.1 Production and Final Uses

The assumptions in this subsection are used only for quantification; neither the elasticity estimates nor the welfare results in Sections 2 and 3.1 require them. In country n, every production sector has the same value-added share $\beta_n \in (0, 1)$ and uses as its intermediate input the same aggregate composite used for household consumption, government purchases, and investment. Let $P_n \equiv \mathcal{E}_n((e_{ns})_{s\in\mathcal{S}})$ denote the price index for this composite. The unit cost in sector s is

$$c_{ns} = a_{ns}^{-1} w_n^{\beta_n} P_n^{1-\beta_n}.$$

Thus, intermediate producers and final users have the same sourcing system within sectors and the same aggregation across sectors. This deliberately parsimonious structure yields the familiar scalar intermediate-input amplification of gains from trade (Arkolakis et al., 2012). The production structure can be generalized to allow sector-specific value-added shares and input-output linkages, as in Caliendo and Parro (2015), but our focus is on the role of home-foreign elasticity schedules rather than heterogeneity in production networks.

Let γ_{nC}, γ_{nG}, and γ_{nI} denote the fixed shares of nominal income allocated to household consumption, government purchases, and investment, with $\gamma_{nC} + \gamma_{nG} + \gamma_{nI} = 1$. We treat government purchases and investment expenditures as current final demands that absorb resources. We hold capital stocks fixed, so investment does not alter productive capacity through capital accumulation. Household real consumption is therefore $\gamma_{nC} Y_n / P_n$. Because γ_{nC} is fixed across the trade and autarky equilibria, it cancels from their proportional welfare comparison, which is governed by Y_n / P_n.

Because every final and intermediate user has the same cross-sector expenditure shares and the same origin shares within each sector, summing across users preserves both sets of shares. We can therefore calibrate the cross-sector weights from total absorption and estimate the sectoral sourcing elasticities from total absorption. We calibrate β_n as economy-wide value added divided by gross output in the relevant trade equilibrium. For the OECD ICIO exercise, this is the sum of the `VA` row divided by the sum of the `OUT` row across country n's production sectors in 2007. For GTAP, it is primary-factor cost divided by primary-factor plus domestic and imported intermediate-input costs at agents' prices in 2014. The median β_n is 0.48 in the OECD ICIO quantification sample and 0.42 in the GTAP sample.

It remains to specify preferences across sectors, which we treat as fully general in Section 3.1. We turn to this next.

5.2 Preferences

We use nested CES preferences across sectors. At the lower tier, manufacturing sectors are aggregated into a manufacturing composite and non-manufacturing goods sectors are aggregated into a non-manufacturing-goods composite. For $b \in \{M, G\}$, let E_{nb} denote the exact price index (unit expenditure function) for composite (b),

$$E_{nb}\left((e_{ns})_{s\in\mathcal{S}_b}\right) = \left[\sum_{s\in\mathcal{S}_b} \omega_{nbs} e_{ns}^{1-\sigma_L}\right]^{1/(1-\sigma_L)},$$

where σ_L is the elasticity of substitution across sectors within each goods-producing composite. We impose the same σ_L within the manufacturing and non-manufacturing-goods composites. At the upper tier, the unit expenditure function aggregates the manufacturing composite, the non-manufacturing-goods composite, and the non-goods composite:

$$\mathcal{E}_n(E_{nM}, E_{nG}, e_{n0}) = \left[\omega_{nM} E_{nM}^{1-\sigma_U} + \omega_{nG} E_{nG}^{1-\sigma_U} + \omega_{n0} e_{n0}^{1-\sigma_U}\right]^{1/(1-\sigma_U)},$$

where σ_U is the elasticity of substitution across the three upper-tier composites. Our baseline imposes $\sigma_L > \sigma_U$. The preference weights are country specific and calibrated to reproduce each country's 2007 expenditure shares.

For each goods-producing sector, define its unit-expenditure increase moving to autarky, relative to its domestic price, by

$$\widehat{e}_{ns} \equiv \exp\left(\int_{\lambda_{nns}^T}^{1} \frac{\mathrm{d}\ln\lambda}{\varepsilon_{ns}(\lambda)}\right).$$

By construction, $\widehat{e}_{n0} = 1$ for the non-traded non-goods composite. Let $\alpha_{ns|b}$ denote sector s's initial expenditure share within composite b and let α_{nM}, α_{nG}, and α_{n0} denote the initial upper-tier expenditure shares. The exact CES changes in the two lower-tier price indexes are

$$\widehat{E}_{nb} = \left[\sum_{s \in \mathcal{S}_b} \alpha_{ns|b} \widehat{e}_{ns}^{\,1-\sigma_L} \right]^{1/(1-\sigma_L)}, \qquad b \in \{M, G\}.$$

Under the production structure in Section 5.1, the resulting gains from trade are

$$\ln \frac{W_n^T}{W_n^A} = \frac{1}{\beta_n (1-\sigma_U)} \ln \left[\alpha_{nM} \widehat{E}_{nM}^{1-\sigma_U} + \alpha_{nG} \widehat{E}_{nG}^{1-\sigma_U} + \alpha_{n0} \right]. \tag{23}$$

Thus, initial domestic expenditure shares and the home-foreign elasticity schedules determine the sectoral costs of autarky, $\widehat{e}_{ns}$. The lower-tier elasticity σ_L aggregates these sectoral costs into the composite costs $\widehat{E}_{nb}$, while σ_U aggregates the composite costs and $1/\beta_n$ captures their amplification through intermediate inputs. Appendix F derives equation (23).

We choose σ_U using estimates from the structural-change literature, in which substitution across broad sectors is a central concern. Our baseline sets $\sigma_U = 0.6$. Comin et al. (2021) estimate an elasticity of 0.57 using a country-level panel of 39 countries with sectoral employment, value added, and prices together with consumption per capita. Cravino and Sotelo (2019) obtain a final-consumption estimate of 0.59 from U.S. input-output use tables for 1977-2012, using final-consumption shares at producer prices and sectoral gross-output price indexes. Our baseline sets $\sigma_L = 2$. Berlingieri et al. (2026, Appendix Table D.1, column 3) estimate a between-sector elasticity of 2.03 when French sectors are defined at the two-digit level, an aggregation roughly comparable to the OECD ICIO manufacturing sectors used here. In robustness we set $\sigma_L = 1$ and $\sigma_L = 3$. The former reproduces the Cobb-Douglas aggregation across sectors commonly imposed in quantitative trade models (e.g., Caliendo and Parro, 2015), while the latter provides a higher-substitution comparison.

5.3 Quantitative Results

Table 3: Gains from Trade Across Elasticity-Schedule Specifications

	p10	p25	Median	Mean	p75	p90
Specification (A)	4.45	6.01	10.34	12.53	15.64	23.15
Specification (B)	4.44	6.05	10.59	12.82	15.91	23.74
Specification (C)	4.34	6.14	10.56	12.80	15.92	23.47

Notes: Entries report $100 \ln(W_n^T/W_n^A)$ calculated using equation (23), where β_n is 2007 economy-wide value added divided by gross output. The percentiles and mean are calculated across the 80 countries in the OECD quantification sample. All three specifications set $\sigma_L = 2$ and $\sigma_U = 0.6$.

Table 3 reports the distribution of gains from trade across the 80 OECD ICIO countries using elasticity schedules estimated in Specifications (A), (B), and (C). Under Specification (B), the median and mean gains from trade are 10.59 and 12.82 log points, while the 10th and 90th percentiles are 4.44 and 23.74 log points. Thus, the estimated gains are economically substantial for the typical country and heterogeneous across countries. The results are also very similar across the three specifications. Given this similarity, we use Specification (B) in what follows.

Cross-country variation in gains from trade is closely related to the initial aggregate import share. Figure 2 plots this relationship under Specification (B). The import share is measured as expenditure on imports in the 26 goods-producing sectors divided by total expenditure, including expenditure on the outside composite. Its correlation with gains from trade is 0.67. This relationship is intuitive: the aggregate import share measures how far the initial trade equilibrium is from autarky. The relationship is not one-to-one, however. The same aggregate import share can arise from different combinations of sectoral expenditure shares and within-sector domestic expenditure shares. Because these objects interact with sector-specific elasticity schedules, countries with similar aggregate import shares can have different gains from trade. The two largest deviations in Figure 2 illustrate this composition effect: the OECD ICIO records zero domestic expenditure for Singapore (SGP) in several non-manufacturing-goods sectors, while Cambodia (KHM) devotes a large share of expenditure to highly import-dependent manufacturing sectors whose fitted elasticities are low. In addition, countries differ in their value-added shares. This is a separate source of variation: holding the other objects fixed, a lower β_n raises measured gains through the intermediate-input multiplier $1/\beta_n$.

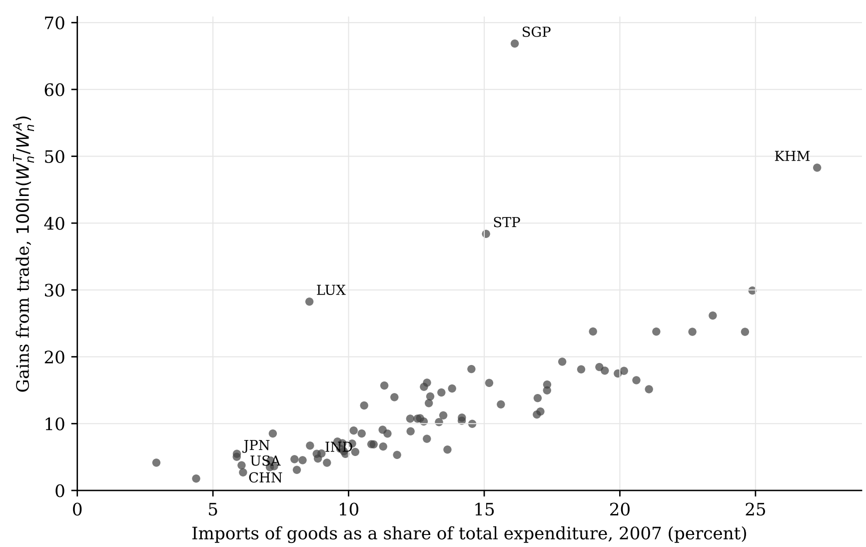

Figure 2: Gains from Trade and Initial Import Shares

Notes: Each observation is one of the 80 countries in the OECD ICIO quantification sample. The vertical axis reports gains from trade, $100\ln(W_n^T/W_n^A)$, under constrained Specification (B), with $\sigma_L = 2$, $\sigma_U = 0.6$, and the country-specific intermediate-input multiplier $1/\beta_n$. The horizontal axis reports 2007 expenditure on imports in the 26 goods-producing sectors divided by total expenditure, including the non-traded composite. Selected countries are labeled.

To assess the quantitative implications of allowing the home-foreign elasticity to vary along the path to autarky, we consider fixed-elasticity comparisons. There is no unique way to construct such a comparison. We consider two examples here. First, in the left panel of Figure 3, we hold each country-sector elasticity fixed at its fitted 2007 value,

$\varepsilon_{ns,2007} \equiv \widehat{\theta}_{s0} + \widehat{\theta}_{g(s)1}\lambda_{nns}^{2007}$. This comparison isolates the implications of allowing the elasticity to vary after the initial equilibrium. Second, in the right panel, let $\lambda_{s,50}$ denote the median 2007 domestic expenditure share in sector s across the 80 countries and define $\varepsilon_{s,50} \equiv \widehat{\theta}_{s0} + \widehat{\theta}_{g(s)1}\lambda_{s,50}$. We apply this sector-specific elasticity to every country and hold it fixed along the path to autarky. This comparison produces one fixed elasticity per sector, as in standard multi-sector constant-elasticity specifications, with each elasticity evaluated at a central value of the domestic expenditure share. Both comparisons retain Specification (B) and the baseline values of the cross-sector elasticities. Moreover, both use home-foreign elasticities estimated directly from variation in home relative to foreign expenditure. Thus, even the fixed-elasticity calculations differ from the standard approach of using elasticities estimated from substitution across foreign sources. Figure 3 plots the gains from trade under each fixed-elasticity calculation against the baseline estimates using the variable elasticity schedules. The dashed lines are 45-degree lines.

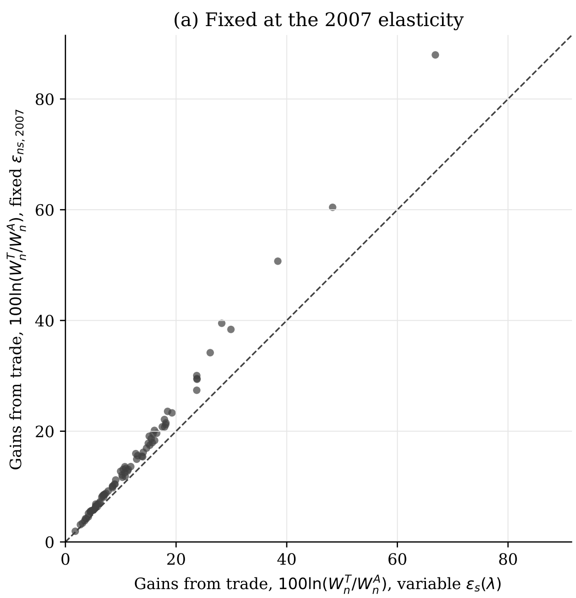

Figure 3: Gains from Trade with Variable and Fixed Home-Foreign Elasticities

Notes: Each observation is one of the 80 countries in the OECD ICIO quantification sample. The horizontal axis reports $100\ln(W_n^T/W_n^A)$ under the variable-elasticity schedules from constrained Specification (B), with $\sigma_L = 2$, $\sigma_U = 0.6$, and the country-specific intermediate-input multiplier $1/\beta_n$. In the left panel, the vertical axis holds each country-sector elasticity fixed at its fitted 2007 value. In the right panel, it holds each sector's elasticity fixed at the value implied by that sector's median 2007 domestic expenditure share and applies that value to every country. Dashed lines are 45-degree lines.

The two measures in the left panel are closely related. Holding each elasticity fixed at its 2007 value increases gains from trade for every country because all slopes used in the quantification are positive. The intuition for this result is simple. Because the baseline elasticity rises along the path to autarky, the fixed elasticity lies below the baseline elasticity at every point along the path except for the original point, where they are equal. This implies larger gains from trade under the fixed-elasticity calculation. These effects are not small: the median rises from 10.59 to 12.64 log points and the mean rises from 12.82 to 15.60 log points.

The two measures in the right panel are also closely related. For a country whose initial domestic expenditure share lies above the corresponding sectoral median in every sector, $\varepsilon_{s,50}$ lies below the baseline elasticity throughout the path to autarky. The fixed-elasticity calculation therefore produces larger gains from trade. This intuition helps explain why countries with lower gains from trade, which tend to have higher domestic expenditure shares as reflected in Figure 2, generally lie above the 45-degree line in the right panel of Figure 3.

Conversely, consider a country whose initial domestic expenditure shares lie below the corresponding sectoral medians. Until a sector's domestic expenditure share reaches its median, the baseline elasticity lies below the fixed elasticity, increasing the gains from trade under the baseline specification. Once the domestic expenditure share exceeds its median, this ordering reverses. The comparison in the right panel therefore depends on the relative importance of these two portions of the path. For countries with sufficiently low initial domestic expenditure shares in important traded sectors, the first portion dominates, placing them below the 45-degree line. This intuition helps explain why countries with higher gains from trade, which tend to have lower domestic expenditure shares as reflected in Figure 2, generally lie below the 45-degree line in the right panel of Figure 3.

These calculations are related to the between-sector mechanism emphasized by Ossa (2015): cross-sector heterogeneity in constant trade elasticities can magnify gains from trade. Under Cobb-Douglas aggregation, a sector's contribution to gains is the product of its trade exposure, $\alpha_{ns}[-\ln \lambda^T_{nns}]$, and its inverse trade elasticity. Trade exposure is approximately $\alpha_{ns}(1-\lambda^T_{nns})$, which is the share of country n's aggregate expenditure that is spent on imports in sector s. Holding the cross-sector distributions of these two objects fixed, increasing the covariance between trade exposure and inverse trade elasticities further magnifies gains from trade. Our framework nests this calculation when cross-sector aggregation is Cobb-Douglas and the home-foreign elasticity in each sector is constant, and therefore incorporates the between-sector mechanism. Figure 3, however, does not isolate its quantitative importance: both axes in both panels retain differences in elasticity levels across sectors. Instead, the figure compares the variable-elasticity schedules with two distinct ways of holding home-foreign elasticities fixed along the path to autarky.

Our framework also introduces a distinct within-sector mechanism, which is most directly isolated in the left panel of Figure 3. When the home-foreign elasticity rises with the domestic expenditure share, a sector with a lower initial domestic expenditure share, for a given elasticity schedule, begins at a lower elasticity and and travels along a longer portion of the elasticity schedule on the path to autarky; consequently, a larger portion of its gains-from-trade integral is evaluated at low elasticities. For a given sectoral expenditure share, its initial domestic expenditure share therefore determines both its trade exposure and the portion of the elasticity schedule relevant for its contribution to gains from trade. The connection to the between-sector mechanism is that both give greater welfare weight to low elasticities in sectors with greater trade exposure. The within-sector mechanism, however, operates even if every sector has the same home-foreign elasticity schedule.

Sensitivity. Finally, Table 4 reports robustness of the Specification (B) results to alternative values of the two cross-sector elasticities. The first row replicates the second row of

Table 4: Gains from Trade Sensitivity to Cross-Sector Substitution

σ_L	σ_U	p10	p25	Median	Mean	p75	p90
2	0.6	4.44	6.05	10.59	12.82	15.91	23.74
1	0.6	4.70	7.57	18.65	∞	∞	∞
3	0.6	4.05	5.58	9.37	10.68	13.70	20.72
2	1	4.38	6.00	10.26	11.99	15.33	22.96

Notes: Entries report $100\ln(W_n^T/W_n^A)$ calculated using equation (23), where β_n is 2007 economy-wide value added divided by gross output. The percentiles and mean are calculated across the 80 countries in the OECD quantification sample. The baseline sets $\sigma_L = 2$ and $\sigma_U = 0.6$; each subsequent row changes the elasticities shown.

Table 3. The next two rows lower and raise σ_L from 2 to 1 and 3. Raising σ_L mitigates the impact of a move to autarky because countries can more easily substitute consumption away from sectors in which price indices rise the most. When $\sigma_L = 1$, 20 countries have infinite gains because the OECD ICIO records zero domestic expenditure in at least one traded sector and Cobb-Douglas preferences make every sector essential. The median among all countries remains finite at 18.65 log points, but the mean and upper quartile are infinite. Raising σ_L to 3 lowers the median to 9.37 log points. The final row raises σ_U from 0.6 to one while holding $\sigma_L = 2$. The higher upper-tier elasticity allows consumption to shift more easily toward the outside composite as the prices of the two goods-producing composites rise; the median falls more modestly, from 10.59 to 10.26 log points.

Appendix G then repeats the quantification using the GTAP data and estimates, taking the 2014 trade equilibrium as the initial equilibrium. It reports and discusses the corresponding distribution of gains from trade, relationship with initial import shares, fixed-elasticity comparisons, and sensitivity to cross-sector substitution.

6 Conclusion

Two broad messages emerge from our analysis. The first is negative. The elasticity relevant for gains from trade describes how expenditure on all foreign goods changes relative to expenditure on home output as foreign prices rise relative to the home price. In practice, this elasticity is typically estimated from the relationship between expenditure across foreign countries and their relative prices. Using such an estimate to calculate gains from trade either incorrectly measures gains or requires strong assumptions equating very different margins of substitution. Moreover, typical approaches additionally impose a constant elasticity along the path to autarky. These assumptions may or may not be appropriate in particular contexts, but they restrict the demand system beyond homotheticity and are not required for quantifying the gains from trade.

The second message is more positive. For measuring gains from trade, a more direct empirical approach exists: estimate how home expenditure changes relative to aggregate foreign expenditure as foreign prices change relative to the home price. The logic parallels Arkolakis et al. (2012). They argue that leveraging the relationship between bilateral imports and bilateral trade costs is attractive for welfare analysis because it measures the welfare-relevant response of trade to trade costs rather than recovering that response indirectly from a micro-level parameter. We apply the same argument, but take it a step further. Leveraging the relationship between home-relative-to-foreign expenditure and home-relative-to-foreign price shifters directly targets the elasticity that enters the gains-from-trade calculation, rather than inferring it from substitution across foreign sources. This direct approach remains tractable because gains from trade remain a low-dimensional object. Conditional on the initial domestic expenditure share, measuring gains requires only the home-foreign elasticity schedule along a path on which all foreign prices rise proportionally relative to the home price. Our implementation result has an implication for measuring that schedule: suitable trade-policy variation can, in principle, identify the elasticity schedule directly from proportional changes in foreign prices. In our empirical implementation, weak separability allows non-proportional price changes to identify the same home-foreign elasticity schedule. Thus, the welfare-relevant object remains low-dimensional in theory and can be targeted directly in the data.

For future structural applications, we develop Armington, Krugman, and Melitz-Pareto gravity models in which the home-foreign elasticity is exactly linear in the domestic expenditure share rather than constant, while substitution across foreign sources is governed by a separate elasticity. These models retain the tractability of familiar gravity frameworks while permitting variable home-foreign elasticity schedules in quantitative counterfactual analysis.

In the empirical analysis, we estimate the welfare-relevant elasticity schedule directly from changes in home relative to foreign expenditure. Across two independently constructed global data sets, the OECD ICIO and GTAP, the estimated home-foreign elasticities generally rise with domestic expenditure shares. The Specification A estimate is positive and statistically significant in both data sets. The broad pattern is stable across alternative samples, sectoral classifications, variable constructions, and instrumental-variable strategies.

The paper introduces a novel within-sector connection between trade exposure and the elasticity relevant for welfare. A destination-sector pair with a lower initial domestic expenditure share begins the path to autarky at a lower home-foreign elasticity. Hence, a larger portion of its gains-from-trade integral is evaluated at low elasticities. Quantitatively, using the OECD ICIO estimates, the median gain from trade is 10.59 log points when the elasticity is allowed to rise along the path to autarky, compared with 12.64 log points when each country-sector elasticity is held fixed at its initial value. Allowing the elasticity to vary lowers measured gains from trade for every country, regardless of how expenditure is aggregated across sectors. This ordering is a consequence of the increasing schedules found in the data, rather than a general implication of allowing elasticities to vary.

The analysis also indicates where further evidence would be most valuable. Sectoral domestic expenditure is constructed less directly than international trade flows, although the similarity of the findings across the OECD ICIO and GTAP data is reassuring. In addition, the calculation of gains from trade requires evaluating elasticity schedules all the way to autarky. Both datasets include a substantial share of observations with high domestic expenditure shares; nevertheless, the exercise relies on the functional form of the elasticity schedule and on restrictions that allow observations from different countries and, in some specifications, sectors to inform one another. Better measurement of domestic expenditure and additional price variation over a wider range of domestic expenditure shares would permit these restrictions to be relaxed.

References

Adão, Rodrigo, Arnaud Costinot, and Dave Donaldson, "Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade," *American Economic Review*, 2017, *107* (3), 633–689.

_ , Costas Arkolakis, and Sharat Ganapati, "From Heterogeneous Firms to Heterogeneous Trade Elasticities: The Aggregate Implications of Firm Export Decisions," February 2026. Conditionally accepted at the Journal of Political Economy; NBER Working Paper No. 28081.

Aguiar, Angel, Maksym Chepeliev, Erwin L. Corong, Robert McDougall, and Dominique van der Mensbrugghe, "The GTAP Data Base: Version 10," *Journal of Global Economic Analysis*, 2019, *4* (1), 1–27.

Andrews, Donald W. K., "Estimation When a Parameter Is on a Boundary," *Econometrica*, 1999, *67* (6), 1341–1383.

Ansari, Habib Chenguiti, Dave Donaldson, and Edward Wiles, "Quantifying the Sensitivity of Quantitative Trade Models," June 2026. Working paper.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare, "New Trade Models, Same Old Gains?," *American Economic Review*, 2012, *102* (1), 94–130.

Armington, Paul S., "A Theory of Demand for Products Distinguished by Place of Production," *IMF Staff Papers*, 1969, *16* (1), 159–178.

Arrow, Kenneth J. and Gérard Debreu, "Existence of an Equilibrium for a Competitive Economy," *Econometrica*, 1954, *22* (3), 265–290.

Backus, David K., Patrick J. Kehoe, and Finn E. Kydland, "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?," *American Economic Review*, March 1994, *84* (1), 84–103.

Baqaee, David, Ariel Burstein, Cédric Duprez, and Emmanuel Farhi, "Consumer Surplus From Suppliers: How Big Is It and Does It Matter for Growth?," *Econometrica*, November 2025, *93* (6), 2043–2081.

Berlingieri, Giuseppe, Filippo Boeri, Danial Lashkari, and Jonathan Vogel, "Capital-Skill Complementarity in Firms and in the Aggregate Economy," June 2026. Manuscript.

Boehm, Christoph E., Andrei A. Levchenko, and Nitya Pandalai-Nayar, "The Long and Short (Run) of Trade Elasticities," *American Economic Review*, April 2023, *113* (4), 861–905.

Caliendo, Lorenzo and Fernando Parro, "Estimates of the Trade and Welfare Effects of NAFTA," *The Review of Economic Studies*, 2015, *82* (1), 1–44.

Chernozhukov, Victor, Whitney K. Newey, and Andres Santos, "Constrained Conditional Moment Restriction Models," *Econometrica*, 2023, *91* (2), 709–736.

Comin, Diego, Danial Lashkari, and Martí Mestieri, "Structural Change With Long-Run Income and Price Effects," *Econometrica*, January 2021, *89* (1), 311–374.

Corsetti, Giancarlo, Luca Dedola, and Sylvain Leduc, "International Risk Sharing and the Transmission of Productivity Shocks," *The Review of Economic Studies*, April 2008, *75* (2), 443–473.

Costinot, Arnaud and Andrés Rodríguez-Clare, "The US Gains from Trade: Valuation Using the Demand for Foreign Factor Services," *Journal of Economic Perspectives*, 2018, *32* (2), 3–24.

_ , Andrés Rodríguez Clare, and Iván Werning, "Micro to Macro: Optimal Trade Policy With Firm Heterogeneity," *Econometrica*, November 2020, *88* (6), 2739–2776.

_ , Dave Donaldson, Jonathan Vogel, and Iván Werning, "Comparative Advantage and Optimal Trade Policy," *The Quarterly Journal of Economics*, None 2015, *130* (2), 659–702.

Cravino, Javier and Sebastian Sotelo, "Trade-Induced Structural Change and the Skill Premium," *American Economic Journal: Macroeconomics*, July 2019, *11* (3), 289–326.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Feenstra, Robert C., Philip Luck, Maurice Obstfeld, and Katheryn N. Russ, "In Search of the Armington Elasticity," *The Review of Economics and Statistics*, March 2018, *100* (1), 135–150.

_ , Robert Inklaar, and Marcel P. Timmer, "The Next Generation of the Penn World Table," *American Economic Review*, 2015, *105* (10), 3150–3182.

Guimbard, Houssein, Sébastien Jean, Mondher Mimouni, and Xavier Pichot, "MAcMap-HS6 2007, an Exhaustive and Consistent Measure of Applied Protection in 2007," *International Economics*, 2012, *130*, 99–121.

Hanoch, Giora, "Production and Demand Models with Direct or Indirect Implicit Additivity," *Econometrica*, 1975, *43* (3), 395–420.

Heathcote, Jonathan and Fabrizio Perri, "Financial Autarky and International Business Cycles," *Journal of Monetary Economics*, April 2002, *49* (3), 601–627.

Kemp, Murray C. and Henry Y. Wan Jr., "An Elementary Proposition Concerning the Formation of Customs Unions," *Journal of International Economics*, February 1976, *6* (1), 95–97.

Krugman, Paul R., "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 1980, *70* (5), 950–959.

Mayer, Thierry and Soledad Zignago, "Notes on CEPII's Distances Measures: The GeoDist Database," CEPII Working Paper 2011-25, CEPII 2011.

Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 2003, *71* (6), 1695–1725.

Organisation for Economic Co-operation and Development, "Bilateral Trade in Goods by End-Use: HS to End-Use Conversion Key," OECD Data Explorer 2025.

_ , "Inter-Country Input-Output Tables: 2025 Edition," OECD Data Explorer 2025.

Ossa, Ralph, "Why Trade Matters After All," *Journal of International Economics*, November 2015, *97* (2), 266–277.

Præstgaard, Jens and Jon A. Wellner, "Exchangeably Weighted Bootstraps of the General Empirical Process," *The Annals of Probability*, October 1993, *21* (4), 2053–2086.

Rubin, Donald B., "The Bayesian Bootstrap," *The Annals of Statistics*, January 1981, *9* (1), 130–134.

Sanderson, Eleanor and Frank Windmeijer, "A Weak Instrument F-Test in Linear IV Models with Multiple Endogenous Variables," *Journal of Econometrics*, February 2016, *190* (2), 212–221.

Simonovska, Ina and Michael E. Waugh, "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 2014, *92* (1), 34–50.

Teti, Feodora A., "Missing Tariffs," CESifo Working Paper 11590, CESifo 2024.

Windmeijer, Frank, "Testing Underidentification in Linear Models, with Applications to Dynamic Panel and Asset Pricing Models," *Journal of Econometrics*, 2024, *240* (2), 105104.

World Trade Organization, "Regional Trade Agreements Database," WTO Data Portal 2025. Accessed July 30, 2026.

Empirical Appendix

A Additional Empirical Details

A.1 Additional Data Details

Country Characteristics. For the country-characteristics sensitivity, log PPP GDP per capita, log population, and remoteness use 1995 values for the OECD ICIO exercise and 2004 values for the GTAP exercise. Within each exercise, we subtract the mean across destination countries and divide by the sample standard deviation. We demean the landlocked indicator across the corresponding destinations.

OECD ICIO Measurement. For each origin-sector row in the ICIO table, we construct $X_{nis,t}$ by summing absorption in destination n across intermediate-use columns and final-demand columns. The final-demand columns are `HFCE`, `NPISH`, `GGFC`, `GFCF`, `INVNT`, and `DPABR`. Negative changes in inventories are allowed to offset other uses when constructing absorption, after which any negative aggregate origin-destination-sector expenditure is set to zero. The diagonal origin-destination cell gives $X_{nns,t}$, while $X_{nFs,t} = \sum_{i \neq n} X_{nis,t}$. The log-change second step keeps cells with positive domestic and foreign expenditure in both years.

The OECD ICIO baseline uses the ISIC two-digit `tariff95` measure from the Global Tariff Database in 1995 and 2007. The exercise using a lagged instrument uses the same source in 1995, 1996, and 2007. To use connect the two datasets, we concord ISIC industries to OECD ICIO goods sectors. When several ISIC industries concord into one ICIO sector, we require each industry to be observed at both endpoints and average their tariff levels with equal weight before taking the log change. Where the ICIO classification is more disaggregated than the public ISIC file, we allocate the same ISIC tariff change to each of the corresponding ICIO sectors. In particular, chemicals and pharmaceuticals share ISIC sector 24, ferrous and non-ferrous basic metals share sector 27, and ships and other non-motor-vehicle transport equipment share sector 35. The resulting crosswalk covers all 19 manufacturing sectors and seven of the eight non-manufacturing goods sectors used in the paper.

Trade-Agreement Tariff Instruments. We identify customs-union membership and implementation dates using the WTO Regional Trade Agreements Database and official sources from the European Union, the Gulf Cooperation Council, and the West African Economic and Monetary Union. Both instruments use the 1995 and 2007 ISIC two-digit `tariff95` measures from the Global Tariff Database, which are concorded to OECD ICIO sectors as described above.

For the May 2004 EU entrants, we set terminal tariffs as follows: tariffs are set to zero on imports from the future EU25 and tariffs are set to the average tariff applied by EU15 destinations in 2007 for each external origin and sector. For incumbent EU15 destinations, we set terminal tariffs as follows: tariffs are set to zero on imports from the ten entrants.

The instrument is the change in the log gross tariff from each destination's observed 1995 tariff to the terminal target.

The broader instrument adds the 2007 EU accession, the EU-Türkiye customs union, the Gulf Cooperation Council, and the West African Economic and Monetary Union. Consider first the 2007 EU expansion and the EU-Türkiye customs union. When considering as destinations the set of Bulgaria, Romania, and Türkiye: (i) we set tariffs on imports from EU members to zero and (ii) we set tariffs on imports from non-EU members equal to the 2007 EU15 tariff for each origin-sector pair. When considering as destinations the incumbent EU15 countries, we set the tariff to zero on imports from Bulgaria and Romania. The EU had already eliminated tariffs on Turkish industrial goods before our 1995 baseline, so the EU-Türkiye component assigns tariff changes only to Türkiye as a destination. Consider next the Gulf Cooperation Council and the West African Economic and Monetary Union. We set internal tariffs to zero and use, for each destination-origin-sector triplet, the average 2007 tariff among the other union members observed in the OECD ICIO sample. The leave-one-out construction does not use a destination's own 2007 tariff in its predicted change.

In both instruments, bilateral tariff changes not covered by the agreements are set to zero. We aggregate the resulting bilateral changes using exactly the same initial foreign expenditure shares, $\mu_{nis,0}$, as in the baseline tariff shifter.

MAcMap-HS6 Tariffs. For the 2004-2013 period sensitivity, we use CEPII's MAcMap-HS6 data and construct a sector-level measure designed to match the `tariff95` measure from the Global Tariff Database as closely as possible. The MAcMap variable `ADV` is an ad-valorem equivalent in decimal units. Let h index HS6 products, n the importer, i the exporter, s the OECD ICIO sector, and $t \in \{2004, 2013\}$ the year. For each year, we pool the valid bilateral `ADV` observations among the concorded HS6 products and countries in the ICIO sample and calculate their 95th percentile, denoted by q_t^{95}. The resulting caps are $q_{2004}^{95} = 0.30$ and $q_{2013}^{95} = 0.20$. We then cap each product-level tariff from above:

$$\mathtt{ADV}_{nih,t}^{95} \equiv \min\{\mathtt{ADV}_{nih,t}, q_t^{95}\}.$$

We concord HS2007 products to OECD ICIO sectors using the OECD's Bilateral Trade in Goods by End-Use conversion key (Organisation for Economic Co-operation and Development, 2025a). The key assigns each HS commodity to a unique detailed CPA activity, which we aggregate to the OECD ICIO activity classification. Let m_{hs} denote the concordance weight from HS6 product h to OECD ICIO sector s, which equals one for the sector to which h is assigned and equals zero for any other sector.

Set $t_0 = 2004$. For each importer-exporter-sector cell, let $\mathcal{H}_{nis}^{t_0}$ denote the set of concorded HS6 products for which valid tariffs are observed in both t_0 and 2013. We average the capped tariff levels across these products without using trade weights:

$$\overline{\mathcal{T}}_{nis,t}^{95,t_0} \equiv \frac{\sum_{h \in \mathcal{H}_{nis}^{t_0}} m_{hs} \mathtt{ADV}_{nih,t}^{95}}{\sum_{h \in \mathcal{H}_{nis}^{t_0}} m_{hs}}.$$

The concordance weights assign products to ICIO sectors; no HS6 import values enter this average. The importer-exporter-sector tariff change used in estimation is the change in the log gross tariff constructed from these unweighted averages:

$$\Delta\tau_{nis}^{T95,t_0} \equiv \log\left(1+\overline{T}_{nis,2013}^{95,t_0}\right) - \log\left(1+\overline{T}_{nis,t_0}^{95,t_0}\right).$$

In words: (i) We first winsorize product-level tariff levels separately by year. (ii) We then average these winsorized tariff levels, without trade weights, within importer-exporter-sector cells. (iii) Finally, we take the log change. The output of this approach is a log change of unweighted average tariff levels.

We keep importer-exporter-sector cells for which $\Delta\tau_{nis}^{T95,t_0}$ is observed. Origin weights are constructed from initial-year OECD ICIO bilateral expenditures. We use

$$\mu_{nis,0} = \frac{X_{nis,t_0}}{\sum_{j\in\mathcal{O}_{ns}} X_{njs,t_0}} \quad \text{for } i \neq n,$$

where $\mathcal{O}_{ns}$ is the set of foreign origins with an observed tariff change (including observed zero tariff changes) for destination n and sector s. Thus the origin shares are normalized to sum to one within each destination-sector cell over origins with observed tariff changes.

Consumer and Producer Manufacturing Allocation. This appendix reports the sector allocations used in specifications that allow separate elasticities for consumer and producer manufacturing. The allocations are constructed separately using initial-year absorption data from the OECD ICIO and GTAP. For each manufacturing sector s, we compute the final-use share at the global level as follows:

$$\chi_s = \frac{\text{world final consumption of sector } s}{\text{world absorption of sector } s}.$$

The numerator combines domestic and imported final consumption. The denominator is total sectoral absorption, including intermediate and final uses, again summing over domestic and imported sources. Sector s is classified as consumer manufacturing if χ_s is weakly above the median value of χ_s across manufacturing sectors in the relevant data set. It is classified as producer manufacturing otherwise.

HS-to-ICIO Crosswalk. To aggregate the product-level MAcMap-HS6 tariffs to OECD ICIO sectors, we use the OECD's Bilateral Trade in Goods by End-Use conversion key (Organisation for Economic Co-operation and Development, 2025a). This key assigns every HS2007 commodity to a unique detailed CPA activity. We aggregate the detailed CPA activities to the goods activities in the OECD ICIO 2025 classification. The resulting crosswalk therefore assigns each HS2007 product to one OECD ICIO sector rather than splitting products across sectors.

Table 5: OECD ICIO Manufacturing Sector Allocation

Code	Description	Final-use share	Allocation
C10T12	Food, beverages, tobacco	70.3	Consumer
C13T15	Textiles, apparel, leather	59.0	Consumer
C19	Coke and refined petroleum	25.9	Consumer
C20	Chemicals	20.2	Consumer
C21	Pharmaceuticals	51.5	Consumer
C26	Computer, electronic, optical	20.2	Consumer
C27	Electrical equipment	21.6	Consumer
C29	Motor vehicles	37.6	Consumer
C302T309	Rail, aircraft, other transport	16.6	Consumer
C31T33	Furniture, other manufacturing, repair	32.8	Consumer
C16	Wood products	6.7	Producer
C17_18	Paper and printing	12.4	Producer
C22	Rubber and plastics	15.8	Producer
C23	Other non-metallic minerals	10.7	Producer
C24A	Iron and steel	0.7	Producer
C24B	Non-ferrous metals	1.7	Producer
C25	Fabricated metals	8.7	Producer
C28	Machinery and equipment	7.0	Producer
C301	Ships and boats	11.9	Producer

Notes: The final-use share is the 1995 share of worldwide absorption accounted for by household, nonprofit-institution, and government final consumption in the OECD ICIO tables. Sectors weakly above the median manufacturing-sector final-use share are assigned to consumer manufacturing; the remaining sectors are assigned to producer manufacturing.

Table 6: GTAP Manufacturing Sector Allocation

Sector	Code	Description	Final-use share	Allocation
19	cmt	Cattle meat	63.5	Consumer
20	omt	Other meat	68.2	Consumer
21	vol	Vegetable oils	39.1	Consumer
22	mil	Dairy products	59.3	Consumer
23	pcr	Processed rice	75.8	Consumer
24	sgr	Sugar and molasses	45.6	Consumer
25	ofd	Other food	58.2	Consumer
26	b_t	Beverages and tobacco	67.4	Consumer
27	tex	Textiles	25.9	Consumer
28	wap	Wearing apparel	78.1	Consumer
29	lea	Leather products	62.0	Consumer
34	bph	Pharmaceuticals	28.0	Consumer
43	mvh	Motor vehicles	29.5	Consumer
45	omf	Other manufacturing	27.5	Consumer
30	lum	Wood products	10.7	Producer
31	ppp	Paper and printing	6.4	Producer
32	p_c	Petroleum and coke	20.3	Producer
33	chm	Chemicals	18.5	Producer
35	rpp	Rubber and plastics	11.5	Producer
36	nmm	Non-metallic minerals	7.2	Producer
37	i_s	Iron and steel	0.6	Producer
38	nfm	Non-ferrous metals	1.1	Producer
39	fmp	Fabricated metals	4.4	Producer
40	ele	Computer, electronic, and optical products	13.8	Producer
41	eeq	Electrical equipment	7.1	Producer
42	ome	Machinery and equipment n.e.c.	8.4	Producer
44	otn	Other transport equipment	15.3	Producer

Notes: The table lists GTAP manufacturing sectors and their allocation to the consumer and producer manufacturing aggregates. The final-use share is the 2004 share of worldwide absorption accounted for by private household and government purchases, using domestic and imported GTAP flows. Sectors with final-use shares weakly above the median manufacturing-sector final-use share are assigned to consumer manufacturing; the remaining sectors are assigned to producer manufacturing.

A.2 Numerical Implementation of the Baseline Estimator

The baseline estimator described in Section 4 profiles out the normalized country cost shifters: for each candidate elasticity vector $\boldsymbol{\theta}$, we solve by least squares for $\widehat{\boldsymbol{C}}(\boldsymbol{\theta})$, with $C_{\text{USA}} = 0$, and substitute this solution into the sum of squared residuals. We minimize the resulting profiled NLLS objective over $\boldsymbol{\theta}$ using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) quasi-Newton algorithm and supply its analytic gradient with respect to $\boldsymbol{\theta}$. BFGS approximates the curvature of the objective function using changes in this gradient. When we construct the gradient, we hold $\widehat{\boldsymbol{C}}(\boldsymbol{\theta})$ fixed rather than differentiating its least-squares solution with respect to $\boldsymbol{\theta}$. Nevertheless, this gives the exact gradient of the profiled objective by an envelope condition. This makes BFGS a convenient optimizer: it uses the exact gradient of the profiled scalar objective without requiring the Jacobian of the profiled residual vector, which would have to account for the dependence of $\widehat{\boldsymbol{C}}(\boldsymbol{\theta})$ on $\boldsymbol{\theta}$.

For each baseline linear specification, we use four starting vectors: we set all sector-specific intercepts, $(\theta_{s0})_s$, and all group-specific slopes, $(\theta_{g1})_g$, respectively, to $(2.2, 0)$, $(1, 2)$, $(4, -1)$, and $(0, 0)$. We require no starting values for the country cost shifters, since they are profiled out. We run BFGS from each starting elasticity vector. We keep the run with the smallest sum of squared residuals. Across the baseline OECD ICIO and GTAP specifications, all four starting vectors lead to the same objective up to numerical precision. At the final estimate $\widehat{\boldsymbol{\theta}}$, we solve again for $\widehat{\boldsymbol{C}}(\widehat{\boldsymbol{\theta}})$.

The normal approximation used for inference in Section 4 requires an asymptotic sequence with an increasing number of independent destination clusters, increasing information for each country- and sector-specific nuisance parameter, and nuisance-estimation bias that is negligible relative to the standard error of each reported statistic.

A.3 Bootstrap Inference for the IV Estimates

For the OECD ICIO IV estimates, we construct the leave-sector-out country-cost components before estimating the elasticity-schedule parameters using 2SLS. Inference must account for uncertainty from estimating these components. A conventional destination-cluster bootstrap can omit some destinations from a draw, even though an omitted country's cost change is still needed (since that country appears as a foreign origin in the price shifters of retained destinations). We therefore use a positive-weight destination-cluster bootstrap. These weights retain every destination in every draw while varying the influence of its observations in every estimation stage.

This is an exchangeably weighted bootstrap in the sense of Præstgaard and Wellner (1993). In bootstrap draw b, we draw independent $E_n^b \sim \text{Exponential}(1)$ and assign destination n the weight

$$\omega_n^b = \frac{80 E_n^b}{\sum_k E_k^b}.$$

The normalized exponential weights are equivalent to 80 times a Dirichlet$(1, \ldots, 1)$ draw (Rubin, 1981). They are strictly positive and sum to 80, so every destination remains in every draw. We apply the same destination weight in every sector and estimation stage, re-estimating the leave-sector-out NLLS models, country-cost changes, and linear IV model.

Appendix Tables 19 and 20 report the full-sample estimates, medians across the 999 bootstrap draws, and 95 percent confidence intervals. The intervals are the empirical 2.5th and 97.5th percentiles of the bootstrap estimates.[14]

A.4 Additional Empirical Results

Table 7: OECD ICIO Estimates Including Non-Manufacturing Sectors

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All goods	4.98	3.85	4.93	5.58	0.18
		(1.55)	(0.87)	(1.04)	(1.17)	
(B)	Consumer manuf.	6.92	2.59	4.05	4.99	0.19
		(2.00)	(0.65)	(0.86)	(1.06)	
	Producer manuf.	0.54	4.67	4.77	4.85	
		(2.50)	(1.25)	(1.10)	(1.13)	
	Non-manufacturing	7.69	4.08	5.99	6.83	
		(2.92)	(1.08)	(1.39)	(1.61)	
(C)	Food, bev., tob.	1.25	3.27	3.39	3.45	0.19
		(3.31)	(0.73)	(0.70)	(0.74)	
	Textiles/apparel/leather	10.79	0.94	3.87	5.55	
		(3.88)	(1.00)	(0.91)	(1.31)	
	Wood/paper/printing	0.64	3.01	3.11	3.17	
		(4.00)	(0.94)	(0.78)	(0.90)	
	Petrol./chem./pharm./rubber	3.81	4.39	5.07	5.50	
		(2.41)	(1.07)	(1.12)	(1.24)	
	Minerals/metals	-4.29	7.22	6.41	5.87	
		(3.95)	(1.93)	(1.54)	(1.44)	
	Machinery/electrical	4.95	2.47	3.41	4.20	
		(2.66)	(0.93)	(0.95)	(1.15)	
	Transport/other manuf.	9.10	0.08	2.71	4.58	
		(2.54)	(0.70)	(0.69)	(1.06)	
	Non-manufacturing	6.93	4.30	6.02	6.77	
		(2.75)	(1.10)	(1.35)	(1.55)	

Notes: This table replicates Table 1, but adds the seven non-manufacturing goods sectors with matched tariff data. Standard errors are clustered by destination country and reported in parentheses.

[14] We retain all 999 requested draws. No draw is rank deficient, and no leave-sector-out estimation fails.

Table 8: OECD ICIO Estimates with Country Characteristics in the Elasticity Intercept

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	5.92	3.26	4.47	5.29	0.28
		(1.85)	(0.80)	(1.02)	(1.22)	
(B)	Consumer manuf.	8.71	2.37	4.19	5.38	0.28
		(2.33)	(0.60)	(0.90)	(1.16)	
	Producer manuf.	3.59	4.39	5.10	5.61	
		(2.46)	(1.16)	(1.31)	(1.50)	
(C)	Food, bev., tob.	16.60	0.86	2.52	3.23	0.29
		(6.14)	(0.50)	(0.70)	(0.91)	
	Textiles/apparel/leather	16.15	-1.13	3.25	5.76	
		(5.09)	(0.96)	(0.80)	(1.48)	
	Wood/paper/printing	0.79	2.55	2.67	2.74	
		(2.39)	(0.81)	(0.76)	(0.81)	
	Petrol./chem./pharm./rubber	8.40	2.73	4.21	5.16	
		(3.11)	(0.85)	(1.19)	(1.48)	
	Minerals/metals	0.56	4.16	4.26	4.33	
		(2.46)	(1.34)	(1.37)	(1.47)	
	Machinery/electrical	7.53	0.77	2.22	3.42	
		(2.58)	(0.59)	(0.79)	(1.10)	
	Transport/other manuf.	10.32	-1.02	1.96	4.08	
		(3.63)	(0.88)	(0.64)	(1.22)	

Notes: This table replicates Table 1, but allows sector elasticity intercepts to vary with 1995 PPP GDP per capita, population, remoteness, and landlocked status. Standard errors are clustered by destination country and reported in parentheses.

Table 9: OECD ICIO Estimates with Intercepts Pooled within Sector Aggregates

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	1.57	3.20	3.64	3.91	0.17
		(1.34)	(1.00)	(1.13)	(1.26)	
(B)	Consumer manuf.	3.04	2.37	3.25	3.86	0.18
		(1.26)	(0.76)	(0.97)	(1.17)	
	Producer manuf.	-1.40	3.90	3.54	3.33	
		(2.03)	(1.31)	(1.24)	(1.30)	
(C)	Food, bev., tob.	3.90	3.07	3.46	3.62	0.21
		(3.21)	(0.86)	(0.88)	(0.92)	
	Textiles/apparel/leather	11.75	0.75	3.94	5.76	
		(4.35)	(0.91)	(1.10)	(1.64)	
	Wood/paper/printing	-1.68	3.14	2.90	2.74	
		(3.55)	(1.12)	(0.95)	(0.97)	
	Petrol./chem./pharm./rubber	2.09	4.12	4.43	4.80	
		(2.16)	(1.42)	(1.50)	(1.67)	
	Minerals/metals	-5.89	6.44	5.07	4.34	
		(4.73)	(2.32)	(1.95)	(1.99)	
	Machinery/electrical	3.42	3.15	3.84	4.54	
		(2.11)	(1.25)	(1.36)	(1.59)	
	Transport/other manuf.	4.76	0.95	2.61	3.74	
		(2.20)	(0.76)	(0.85)	(1.23)	

Notes: This table replicates Table 1, but pools intercepts within manufacturing aggregates. Standard errors are clustered by destination country and reported in parentheses.

Table 10: OECD ICIO Estimates with Sector-Specific Elasticity Schedules

Industry	$\widehat{\theta}_{s0}$	$\widehat{\theta}_{s1}$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$
C10T12 Food, beverages, tobacco	-2.53 (3.18)	7.11 (3.89)	3.00 (0.70)	3.71 (0.77)	4.01 (0.85)
C13T15 Textiles, apparel, leather	-4.06 (2.27)	12.21 (4.06)	0.88 (0.92)	4.20 (0.98)	6.09 (1.46)
C16 Wood products	5.58 (3.78)	-2.33 (4.70)	4.14 (1.35)	3.65 (1.20)	3.46 (1.35)
C17_18 Paper and printing	3.17 (3.17)	-0.98 (4.12)	2.53 (0.78)	2.42 (0.65)	2.32 (0.78)
C19 Coke and refined petroleum	2.81 (3.19)	5.05 (4.46)	5.74 (1.92)	6.93 (2.20)	7.41 (2.44)
C20 Chemicals	2.34 (2.28)	7.21 (4.12)	5.34 (1.84)	6.62 (2.09)	7.36 (2.33)
C21 Pharmaceuticals	2.74 (3.68)	2.19 (5.04)	3.88 (1.43)	4.27 (1.14)	4.53 (1.29)
C22 Rubber and plastics	3.26 (2.91)	-0.74 (4.35)	2.88 (1.04)	2.79 (0.91)	2.69 (1.09)
C23 Other non-metallic minerals	3.41 (2.72)	-1.61 (3.14)	2.16 (0.70)	2.03 (0.67)	1.95 (0.70)
C24A Iron and steel	7.44 (3.13)	1.83 (3.38)	8.21 (2.45)	8.61 (2.37)	8.90 (2.46)
C24B Non-ferrous metals	14.27 (6.43)	0.40 (7.46)	14.39 (5.06)	14.51 (4.53)	14.59 (4.73)
C25 Fabricated metals	6.00 (2.66)	-4.69 (3.21)	3.35 (1.07)	2.60 (0.81)	2.09 (0.79)
C26 Computer, electronic, optical	0.63 (1.52)	5.53 (4.56)	1.47 (1.12)	2.34 (1.07)	3.31 (1.50)
C27 Electrical equipment	2.14 (1.92)	6.16 (3.90)	4.05 (1.43)	5.50 (1.66)	6.23 (1.94)
C28 Machinery and equipment	2.74 (1.77)	4.88 (4.59)	3.40 (1.48)	4.29 (1.46)	5.18 (1.87)
C29 Motor vehicles	-1.20 (1.32)	9.45 (3.30)	0.08 (0.97)	2.75 (0.78)	5.32 (1.38)
C301 Ships and boats	-1.99 (2.96)	12.35 (5.60)	-0.25 (2.42)	4.84 (2.09)	8.17 (3.03)
C302T309 Rail, aircraft, other transport	-0.25 (1.26)	8.90 (3.62)	0.96 (1.08)	3.90 (1.46)	5.71 (2.04)
C31T33 Furniture, other manufacturing, repair	-0.99 (1.78)	4.26 (2.34)	1.80 (0.74)	2.34 (0.77)	2.67 (0.84)

Notes: This table replicates Table 1, but allows every manufacturing sector to have its own intercept and slope. Standard errors are clustered by destination country and reported in parentheses. The specification uses 1,510 observations and has $R^2 = 0.28$.

Table 11: OECD ICIO Estimates Requiring at Least 50 Percent Tariff Coverage

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	5.46	4.01	5.12	5.88	0.26
		(1.87)	(1.12)	(1.32)	(1.49)	
(B)	Consumer manuf.	8.02	2.99	4.66	5.74	0.27
		(2.43)	(0.82)	(1.13)	(1.39)	
	Producer manuf.	2.11	5.15	5.57	5.87	
		(2.48)	(1.56)	(1.56)	(1.66)	
(C)	Food, bev., tob.	6.97	3.03	3.73	4.03	0.28
		(3.93)	(0.74)	(0.80)	(0.88)	
	Textiles/apparel/leather	12.31	0.84	4.18	6.10	
		(4.08)	(0.91)	(1.00)	(1.48)	
	Wood/paper/printing	-1.20	3.33	3.14	3.03	
		(4.14)	(1.03)	(0.87)	(0.98)	
	Petrol./chem./pharm./rubber	3.60	4.57	5.19	5.59	
		(2.79)	(1.31)	(1.45)	(1.61)	
	Minerals/metals	-1.45	7.16	6.89	6.70	
		(3.15)	(2.18)	(2.03)	(2.01)	
	Machinery/electrical	5.54	2.99	4.05	4.93	
		(2.99)	(1.07)	(1.22)	(1.51)	
	Transport/other manuf.	9.25	0.69	3.35	5.25	
		(2.92)	(0.71)	(0.97)	(1.46)	

Notes: This table replicates Table 1, but requires origins with observed tariff changes to account for at least 50 percent of initial foreign expenditure. Standard errors are clustered by destination country and reported in parentheses.

Table 12: OECD ICIO Estimates Using Trade-Weighted Winsorized Tariffs

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	4.52	3.43	4.35	4.99	0.27
		(1.53)	(0.86)	(1.02)	(1.17)	
(B)	Consumer manuf.	6.73	2.48	3.87	4.81	0.27
		(1.99)	(0.63)	(0.86)	(1.08)	
	Producer manuf.	1.58	4.38	4.69	4.91	
		(2.05)	(1.19)	(1.22)	(1.33)	
(C)	Food, bev., tob.	10.03	2.26	3.27	3.67	0.28
		(4.75)	(0.62)	(0.71)	(0.83)	
	Textiles/apparel/leather	10.47	0.47	3.37	5.06	
		(3.53)	(0.84)	(0.86)	(1.28)	
	Wood/paper/printing	-0.36	2.90	2.85	2.82	
		(3.62)	(0.86)	(0.78)	(0.90)	
	Petrol./chem./pharm./rubber	3.13	3.73	4.28	4.63	
		(2.45)	(0.98)	(1.07)	(1.21)	
	Minerals/metals	-1.20	5.79	5.57	5.41	
		(2.40)	(1.56)	(1.47)	(1.48)	
	Machinery/electrical	4.49	2.59	3.44	4.20	
		(2.42)	(0.79)	(0.89)	(1.15)	
	Transport/other manuf.	7.34	0.72	2.80	4.35	
		(2.16)	(0.57)	(0.77)	(1.13)	

Notes: This table replicates Table 1, but uses the trade-weighted winsorized applied-tariff measure `tariff95_w` from the Global Tariff Database. Standard errors are clustered by destination country and reported in parentheses.

Table 13: OECD ICIO Estimates Using Winsorized MFN Tariffs

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	4.33	3.21	4.10	4.70	0.26
		(1.65)	(1.03)	(1.23)	(1.40)	
(B)	Consumer manuf.	6.61	2.45	3.83	4.73	0.26
		(2.22)	(0.78)	(1.10)	(1.35)	
	Producer manuf.	1.74	4.17	4.52	4.76	
		(2.07)	(1.39)	(1.41)	(1.49)	
(C)	Food, bev., tob.	5.81	2.73	3.31	3.56	0.27
		(3.39)	(0.84)	(0.95)	(1.03)	
	Textiles/apparel/leather	11.35	0.73	3.81	5.57	
		(4.00)	(0.78)	(1.12)	(1.64)	
	Wood/paper/printing	-0.65	2.78	2.68	2.62	
		(3.46)	(1.04)	(0.86)	(0.91)	
	Petrol./chem./pharm./rubber	2.99	3.64	4.17	4.51	
		(2.29)	(1.23)	(1.36)	(1.49)	
	Minerals/metals	-0.92	5.66	5.49	5.37	
		(2.59)	(1.97)	(1.85)	(1.84)	
	Machinery/electrical	4.01	2.32	3.09	3.73	
		(2.37)	(0.94)	(1.09)	(1.32)	
	Transport/other manuf.	7.52	0.46	2.63	4.18	
		(2.78)	(0.57)	(0.91)	(1.41)	

Notes: This table replicates Table 1, but uses the winsorized MFN measure `mfn95` from the Global Tariff Database. Standard errors are clustered by destination country and reported in parentheses.

Table 14: OECD ICIO Estimates with a Quadratic Elasticity Schedule

Spec.	Aggregate	$\widehat{\varepsilon}'_{50}$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	7.17	3.32	4.28	5.34	0.27
		(2.45)	(0.99)	(1.16)	(1.41)	
(B)	Consumer manuf.	8.98	2.46	3.90	5.15	0.27
		(2.78)	(0.77)	(1.02)	(1.31)	
	Producer manuf.	5.24	4.21	4.51	5.38	
		(3.01)	(1.29)	(1.31)	(1.57)	
(C)	Food, bev., tob.	18.96	2.01	3.40	4.31	0.29
		(6.75)	(0.62)	(0.73)	(0.94)	
	Textiles/apparel/leather	11.49	0.75	4.00	5.74	
		(4.54)	(0.87)	(1.13)	(1.39)	
	Wood/paper/printing	12.51	1.62	2.71	4.12	
		(6.33)	(0.66)	(0.78)	(1.32)	
	Petrol./chem./pharm./rubber	6.16	3.67	4.50	5.30	
		(4.10)	(1.25)	(1.28)	(1.46)	
	Minerals/metals	9.55	5.06	5.15	6.64	
		(3.45)	(1.57)	(1.52)	(1.83)	
	Machinery/electrical	4.84	2.68	3.32	4.27	
		(2.74)	(0.93)	(1.11)	(1.39)	
	Transport/other manuf.	8.61	0.65	2.93	4.75	
		(2.89)	(0.67)	(0.91)	(1.33)	

Notes: This table replicates Table 1, but estimates $\varepsilon_{ns}(\lambda) = \theta_{s0} + \theta_{g(s)1}\lambda + \theta_{g(s)2}\lambda^2$. The derivative column is evaluated at the unweighted average across sectors of each sector's median initial domestic expenditure share. Standard errors are clustered by destination country and reported in parentheses. Each specification uses 1,510 destination–sector observations.

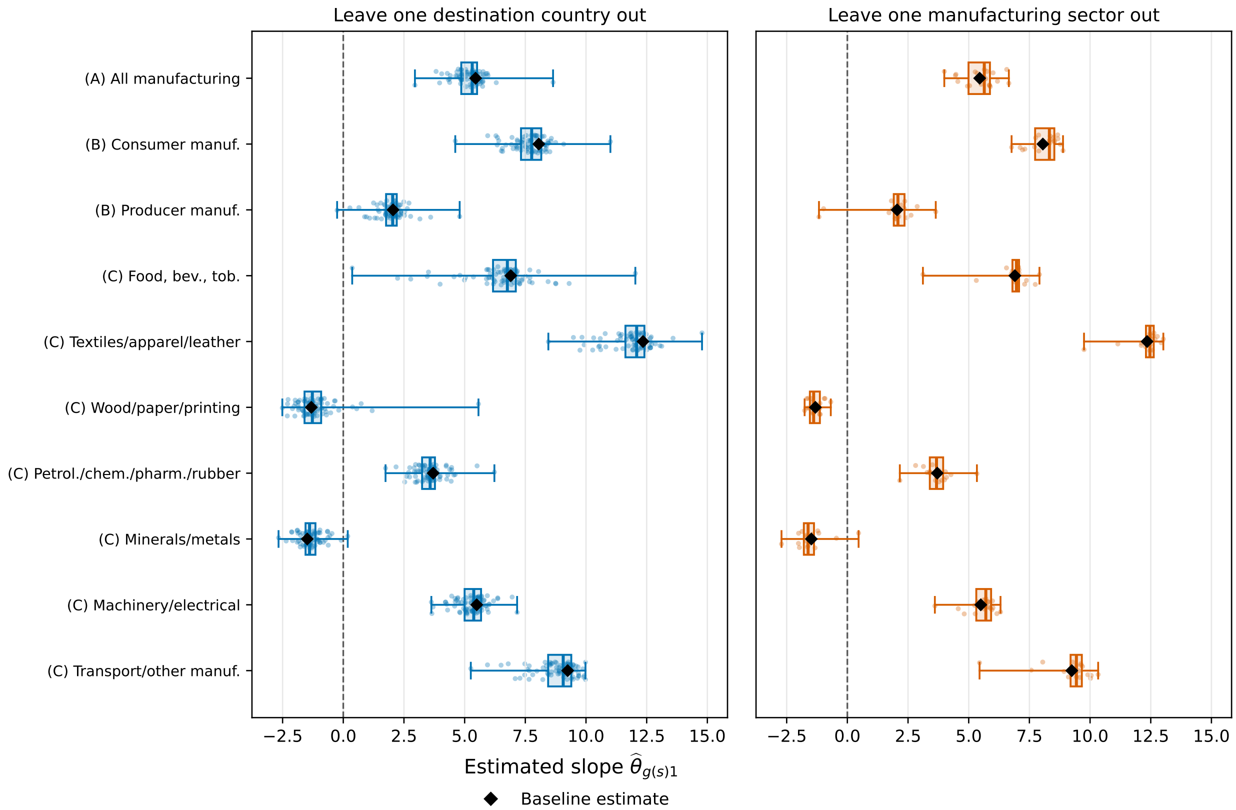

Figure 4: OECD ICIO Leave-One-Out Estimates of the Home-Foreign Elasticity Slopes
Notes: The left panel re-estimates Specifications (A)-(C) after omitting each of the 80 destination countries in turn. The right panel re-estimates the same specifications after omitting each of the 19 manufacturing sectors in turn. Boxes show the interquartile range, vertical lines within the boxes show the median, and whiskers show the full range of leave-one-out estimates. Each point is an individual leave-one-out estimate, and black diamonds are the baseline estimates.

Constrained and Exact Finite-Change Estimates. The constrained estimator retains the estimating equation and NLLS objective from Section 4, but imposes

$$\theta_{s0} \geq 0 \qquad \text{and} \qquad \theta_{s0} + \theta_{g(s)1} \geq 0 \quad \text{for every } s.$$

Because the elasticity schedule is linear in the domestic expenditure share, these endpoint restrictions ensure that it is non-negative on $[0, 1]$. Conditional on each elasticity vector, we obtain the normalized country cost shifters by least squares, exactly as in the unrestricted estimator. We then minimize the resulting profiled sum of squared residuals over the elasticity parameters subject to $\theta_{s0} \geq 0$ and $\theta_{s0} + \theta_{g(s)1} \geq 0$ for every sector, using a sequential quadratic programming algorithm that incorporates these linear constraints into each optimization step. Thus, the objective function and profiling step are unchanged from the baseline estimator, while the elasticity parameters are optimized over the restricted parameter space. We initialize the constrained optimizer from three feasible elasticity vectors, setting all sector-specific intercepts and group-specific slopes, respectively, to $(2.2, 0)$, $(1, 2)$, and $(4, -1)$, and retain the feasible solution with the smallest sum of squared residuals. Appendix Tables 15 and 30 report the resulting point estimates. These are the estimates used in the quantification.

The exact finite-change estimator retains the same samples, specifications, and positivity constraints but does not evaluate the elasticity schedule only at its initial value. For each candidate elasticity schedule and vector of country cost changes, let $\widehat{\lambda}_{nns,1}(\boldsymbol{\theta}, \boldsymbol{C})$ solve

$$Q_{ns}(\boldsymbol{C}) = \int_{\lambda_{nns,0}}^{\widehat{\lambda}_{nns,1}(\boldsymbol{\theta},\boldsymbol{C})} \frac{\mathrm{d}\lambda}{\lambda(1-\lambda)[\theta_{s0} + \theta_{g(s)1}\lambda]}.$$

We estimate $\boldsymbol{\theta}$ and $\boldsymbol{C}$ jointly by minimizing

$$\sum_{n,s} \left\{ \ln \frac{\lambda_{nns,1}}{1-\lambda_{nns,1}} - \ln \frac{\widehat{\lambda}_{nns,1}(\boldsymbol{\theta},\boldsymbol{C})}{1-\widehat{\lambda}_{nns,1}(\boldsymbol{\theta},\boldsymbol{C})} \right\}^2 .$$

In the baseline estimator, the country cost changes enter linearly (conditional on the elasticity schedule) and are obtained by least squares for each candidate elasticity vector. In the exact finite-change objective, they enter nonlinearly, so we minimize jointly over the elasticity parameters and normalized country cost changes. We use a sequential quadratic programming algorithm subject to the same linear endpoint restrictions imposed by the constrained initial-share estimator. At each candidate parameter vector, we use a fourth-order Runge-Kutta algorithm to solve the finite-change demand equation and calculate its parameter derivatives from analytically differentiated equations. We then use these derivatives to compute the gradient of the objective supplied to the optimizer. When $\theta_{g(s)1} = 0$, the exact equation reduces to $\Delta \ln(X_{nns}/X_{nFs}) = \theta_{s0} Q_{ns}(\boldsymbol{C})$, so the exact and initial-share equations coincide in the constant-elasticity case. Appendix Tables 16 and 31 report the exact point estimates.

We use three starting vectors for the joint optimization. The first combines the constrained initial-share elasticity estimates with their corresponding least-squares country cost changes. The other two multiply the initial elasticity estimates by 0.75 and 1.25, with adjustments when needed to satisfy the endpoint restrictions. We retain the solution with the smallest sum of squared residuals. All three starting vectors yield the same objective up to numerical precision in every OECD ICIO and GTAP specification. For each observation, we divide the finite relative-price change into 64 equal steps when solving the demand equation by Runge-Kutta Doubling this number changes the predicted terminal log home-foreign expenditure ratio by less than 2×10^{-10} in every observation.

Table 15: Constrained OECD ICIO Estimates Used in Quantification

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	3.90	4.20	5.00	5.54	0.26
(B)	Consumer manuf.	4.74	3.45	4.45	5.09	0.26
	Producer manuf.	1.42	5.13	5.41	5.61	
(C)	Food, bev., tob.	4.37	3.40	3.83	4.02	0.27
	Textiles/apparel/leather	6.63	2.68	4.48	5.51	
	Wood/paper/printing	-2.13	3.64	3.30	3.11	
	Petrol./chem./pharm./rubber	3.91	4.79	5.48	5.92	
	Minerals/metals	-1.75	7.59	7.26	7.04	
	Machinery/electrical	5.64	3.35	4.44	5.33	
	Transport/other manuf.	5.05	1.97	3.42	4.46	

Notes: This table reports constrained NLLS estimates using the OECD ICIO over 1995–2007. The estimates impose $\theta_{s0} \geq 0$ and $\theta_{s0} + \theta_{g(s)1} \geq 0$ for every sector and are used in the quantification. The table reports point estimates only. Each specification uses 1,510 destination–sector observations.

Table 16: Exact Finite-Change Constrained OECD ICIO Estimates

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	4.04	4.57	5.40	5.96	0.26
(B)	Consumer manuf.	5.17	3.77	4.85	5.55	0.26
	Producer manuf.	0.84	5.45	5.62	5.74	
(C)	Food, bev., tob.	4.18	3.74	4.15	4.33	0.27
	Textiles/apparel/leather	7.41	3.00	5.01	6.16	
	Wood/paper/printing	-2.96	3.77	3.31	3.04	
	Petrol./chem./pharm./rubber	3.15	5.21	5.77	6.12	
	Minerals/metals	-3.17	7.97	7.37	6.97	
	Machinery/electrical	8.08	3.66	5.21	6.50	
	Transport/other manuf.	5.83	2.20	3.88	5.08	

Notes: This table replicates Table 15, replacing the initial-share approximation by solving the exact demand equation forward from the observed initial domestic expenditure share and comparing its predicted terminal log home-foreign expenditure ratio with the observed ratio. The estimates use the OECD ICIO over 1995–2007. The table uses the same sample, Specifications (A)–(C), sector classifications, and constraints $\theta_{s0} \geq 0$ and $\theta_{s0} + \theta_{g(s)1} \geq 0$ as the comparison table. The table reports point estimates only. For each aggregate, $\widehat{\varepsilon}_p$ is the unweighted mean across sectors of the fitted sectoral elasticities evaluated at each sector's own pth percentile of the initial domestic expenditure share. The country-cost changes and elasticity parameters are estimated jointly. The residual and R^2 remain in log home-foreign expenditure-ratio units.

Table 17: OECD ICIO Estimates, 2004-2013

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	8.77	-0.22	1.42	2.74	0.23
		(3.81)	(0.26)	(0.80)	(1.35)	
(B)	Consumer manuf.	5.37	0.39	1.33	2.17	0.25
		(2.33)	(0.33)	(0.66)	(1.00)	
	Producer manuf.	10.66	-0.12	2.00	3.55	
		(4.37)	(0.33)	(1.01)	(1.63)	
(C)	Food, bev., tob.	4.02	0.67	0.94	1.22	0.26
		(3.35)	(0.41)	(0.56)	(0.76)	
	Textiles/apparel/leather	2.68	0.43	0.93	1.57	
		(1.98)	(0.48)	(0.61)	(0.98)	
	Wood/paper/printing	4.37	0.20	0.80	1.18	
		(3.18)	(0.26)	(0.53)	(0.79)	
	Petrol./chem./pharm./rubber	1.21	0.84	1.07	1.23	
		(1.43)	(0.54)	(0.65)	(0.77)	
	Minerals/metals	11.16	0.15	2.44	3.97	
		(5.57)	(0.40)	(1.32)	(2.06)	
	Machinery/electrical	8.18	-0.63	1.07	2.37	
		(4.23)	(0.39)	(0.61)	(1.26)	
	Transport/other manuf.	8.82	-0.69	1.14	2.97	
		(4.51)	(0.40)	(0.67)	(1.58)	

Notes: This table replicates Table 1, but uses OECD ICIO outcomes over 2004–2013 and MAcMap-HS6 tariff changes rather than the baseline period and tariff source. Standard errors are clustered by destination country and reported in parentheses.

Table 18: OECD ICIO Estimates with Lagged Instruments

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
Panel A. Ordinary 1996–2007 estimates						
(A)	All manufacturing	4.87	4.69	5.66	6.32	0.26
		(2.18)	(1.22)	(1.45)	(1.65)	
(B)	Consumer manuf.	7.01	3.77	5.23	6.16	0.26
		(2.36)	(0.94)	(1.22)	(1.46)	
	Producer manuf.	2.28	5.62	6.05	6.37	
		(2.81)	(1.61)	(1.80)	(2.02)	
(C)	Food, bev., tob.	0.62	3.44	3.50	3.53	0.27
		(4.14)	(0.77)	(0.65)	(0.68)	
	Textiles/apparel/leather	11.23	2.09	4.78	6.48	
		(4.37)	(1.02)	(1.15)	(1.63)	
	Wood/paper/printing	-3.16	4.19	3.68	3.41	
		(4.53)	(1.17)	(1.03)	(1.15)	
	Petrol./chem./pharm./rubber	5.13	5.25	6.17	6.73	
		(2.98)	(1.53)	(1.78)	(1.98)	
	Minerals/metals	1.56	7.78	8.06	8.26	
		(3.33)	(2.27)	(2.54)	(2.78)	
	Machinery/electrical	2.83	4.48	5.03	5.47	
		(2.75)	(1.40)	(1.42)	(1.57)	
	Transport/other manuf.	8.38	0.88	3.13	4.85	
		(2.77)	(0.63)	(0.85)	(1.31)	
Panel B. Instrumented estimates						
(A)	All manufacturing	5.81	5.87	7.03	7.83	0.25
		(2.61)	(1.60)	(1.75)	(1.93)	
(B)	Consumer manuf.	8.52	4.65	6.42	7.55	0.26
		(2.86)	(1.22)	(1.43)	(1.67)	
	Producer manuf.	2.27	7.02	7.45	7.77	
		(3.43)	(2.08)	(2.20)	(2.41)	
(C)	Food, bev., tob.	-1.20	4.03	3.91	3.86	0.26
		(6.59)	(0.95)	(0.71)	(0.77)	
	Textiles/apparel/leather	9.42	2.64	4.90	6.32	
		(6.80)	(1.76)	(1.17)	(1.73)	
	Wood/paper/printing	-5.72	5.10	4.18	3.69	
		(6.55)	(1.40)	(1.25)	(1.51)	
	Petrol./chem./pharm./rubber	8.05	7.00	8.46	9.33	
		(3.86)	(1.81)	(2.09)	(2.34)	
	Minerals/metals	0.71	10.58	10.71	10.80	
		(4.18)	(2.93)	(3.18)	(3.44)	
	Machinery/electrical	5.47	5.78	6.85	7.70	
		(4.76)	(2.02)	(1.84)	(2.03)	
	Transport/other manuf.	10.20	1.29	4.03	6.13	
		(3.61)	(0.94)	(1.08)	(1.63)	

Notes: This table replicates Table 1, but both panels use 1996–2007 outcomes and tariff changes. Panel B uses the 1995-based lagged instruments described in the text. Standard errors are clustered by destination country and reported in parentheses. Both panels use the same 1,510 destination–sector observations.

Table 19: OECD ICIO IV Estimates Using the 2004 EU Enlargement

		Full-sample estimate (bootstrap median)				
Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	Min. cond. F
(A)	All manufacturing	8.83 (5.28)	3.62 (2.33)	5.42 (3.67)	6.65 (4.47)	33.18
		[-1.77,18.08]	[0.12,7.71]	[0.91,8.15]	[1.21,9.98]	
(B)	Consumer manuf.	5.20 (3.36)	4.24 (2.57)	5.33 (3.54)	6.04 (4.01)	14.83
		[-8.85,14.06]	[0.45,9.45]	[1.08,8.78]	[1.26,9.24]	
	Producer manuf.	10.51 (6.55)	3.91 (2.57)	5.99 (4.13)	7.48 (5.25)	
		[-2.57,22.46]	[-0.85,11.34]	[0.97,11.79]	[1.33,12.05]	
(C)	Food, bev., tob.	15.57 (11.53)	2.70 (2.00)	4.25 (3.18)	4.92 (3.70)	16.44
		[2.71,22.59]	[0.64,3.33]	[1.14,4.98]	[1.32,5.75]	
	Textiles/apparel/leather	26.27 (14.94)	0.92 (0.86)	8.05 (5.72)	12.13 (7.98)	
		[-53.04,59.64]	[-11.19,27.25]	[0.76,12.81]	[0.68,16.99]	
	Wood/paper/printing	16.10 (11.70)	1.06 (0.52)	3.59 (2.40)	5.05 (3.52)	
		[0.48,34.70]	[-2.10,2.61]	[0.68,5.50]	[1.06,7.98]	
	Petrol./chem./pharm./rubber	-8.48 (-2.90)	9.08 (5.22)	7.58 (4.68)	6.62 (4.21)	
		[-25.34,10.35]	[1.12,14.28]	[1.37,11.72]	[1.19,10.84]	
	Minerals/metals	6.27 (4.48)	5.77 (3.11)	6.95 (4.43)	7.75 (5.15)	
		[-33.93,46.91]	[-6.24,16.67]	[0.80,12.34]	[1.07,13.52]	
	Machinery/electrical	15.37 (7.79)	-1.30 (-0.34)	1.65 (1.11)	4.10 (2.44)	
		[-2.55,35.09]	[-5.85,2.59]	[-0.43,3.59]	[0.50,7.38]	
	Transport/other manuf.	9.77 (6.09)	-0.35 (-0.22)	2.47 (1.53)	4.48 (2.82)	
		[-0.34,20.67]	[-3.37,2.45]	[0.21,4.37]	[0.58,8.10]	

Notes: This table reports linear IV estimates using the 1995–2007 OECD ICIO sample. For each specification and sector s, the unrestricted NLLS model is re-estimated after excluding sector s at every destination, jointly re-estimating the retained-sector elasticity parameters and country-cost changes. The resulting country-cost changes are held fixed in the linear IV estimation, and their leave-sector-out component is combined with the observed tariff component in the structural price shifter. The excluded instruments are the tariff changes required by the May 2004 EU enlargement, including changes applied by both the ten entrants and the incumbent EU15. Bilateral tariff changes not covered by these agreements are set to zero, and the instrument uses the same initial foreign-expenditure weights as the baseline tariff shifter. The instrument is also interacted with the initial domestic expenditure share. For each aggregate, $\widehat{\varepsilon}_p$ is the unweighted mean across sectors of the fitted sectoral elasticities evaluated at each sector's own pth percentile of the initial domestic expenditure share. In each coefficient or elasticity cell, the first number is the full-sample point estimate, the number in parentheses is the median across 999 bootstrap estimates, and brackets contain percentile 95 percent confidence intervals, given by the 2.5th and 97.5th percentiles of those estimates. In each bootstrap draw, independent exponential weights are assigned to destination countries and normalized to sum to 80. The same destination weight is applied to all sectors and estimation stages, and the leave-sector-out NLLS models, country-cost changes, and linear IV model are re-estimated. The bootstrap uses the predetermined draw IDs 0 through 998; no draw is discarded based on its coefficient magnitude or numerical condition number. Because the positive weights retain every destination, the intervals condition on the observed country and sector sets, trade network, agreement support, and initial expenditure shares. Complete draw-level coefficient ranges, condition-number distributions, effective cluster counts, rank diagnostics, and first-stage estimation failures are retained in the accompanying numerical-summary files. The final column reports, once for each specification, the smallest destination-clustered Sanderson-Windmeijer conditional first-stage F statistic across all endogenous regressors for the instrument used in this table; Table 21 reports the full vector.

Table 20: OECD ICIO IV Estimates Using the Broader Customs-Union Instrument

Spec.	Aggregate	Full-sample estimate (bootstrap median) $\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	Min. cond. F
(A)	All manufacturing	10.75 (6.91)	2.76 (1.81)	4.96 (3.34)	6.45 (4.37)	20.01
		[1.20,14.79]	[0.20,5.77]	[0.86,7.09]	[1.06,8.34]	
(B)	Consumer manuf.	8.60 (5.93)	2.75 (1.76)	4.55 (3.18)	5.72 (4.07)	22.41
		[0.62,12.97]	[0.15,4.97]	[1.01,5.88]	[1.39,7.10]	
	Producer manuf.	11.64 (7.24)	3.47 (2.25)	5.77 (3.90)	7.43 (5.07)	
		[-0.16,19.41]	[-0.72,8.46]	[0.93,10.66]	[1.34,12.07]	
(C)	Food, bev., tob.	9.70 (7.01)	2.54 (1.82)	3.51 (2.56)	3.92 (2.86)	7.36
		[0.10,13.92]	[0.63,3.17]	[0.91,4.08]	[1.03,4.55]	
	Textiles/apparel/leather	11.08 (8.04)	1.78 (0.76)	4.78 (3.61)	6.50 (4.77)	
		[-14.38,34.98]	[-4.60,11.21]	[0.80,8.26]	[1.29,11.83]	
	Wood/paper/printing	19.03 (13.75)	0.68 (0.36)	3.68 (2.55)	5.40 (3.82)	
		[3.48,32.03]	[-1.34,1.54]	[0.78,4.97]	[1.19,7.82]	
	Petrol./chem./pharm./rubber	-0.18 (0.76)	6.31 (3.90)	6.28 (4.07)	6.26 (4.13)	
		[-14.30,10.02]	[1.14,9.48]	[1.33,8.10]	[1.26,8.52]	
	Minerals/metals	6.85 (2.83)	5.77 (3.61)	7.06 (4.38)	7.93 (4.89)	
		[-34.82,19.82]	[-0.05,18.60]	[0.92,13.77]	[1.04,12.38]	
	Machinery/electrical	15.34 (8.36)	-0.81 (-0.21)	2.14 (1.36)	4.58 (2.87)	
		[-1.43,35.16]	[-6.05,2.74]	[-0.45,3.64]	[0.66,7.29]	
	Transport/other manuf.	11.69 (8.08)	-0.88 (-0.54)	2.49 (1.69)	4.90 (3.43)	
		[2.17,18.41]	[-3.09,1.24]	[0.42,3.82]	[1.02,7.04]	

Notes: This table reports linear IV estimates using the 1995–2007 OECD ICIO sample. For each specification and sector s, the unrestricted NLLS model is re-estimated after excluding sector s at every destination, jointly re-estimating the retained-sector elasticity parameters and country-cost changes. The resulting country-cost changes are held fixed in the linear IV estimation, and their leave-sector-out component is combined with the observed tariff component in the structural price shifter. The excluded instruments are the tariff changes associated with the 2004 and 2007 EU enlargements, the EU–Türkiye customs union, the Gulf Cooperation Council, and the West African Economic and Monetary Union. Bilateral tariff changes not covered by these agreements are set to zero, and the instrument uses the same initial foreign-expenditure weights as the baseline tariff shifter. The instrument is also interacted with the initial domestic expenditure share. For each aggregate, $\widehat{\varepsilon}_p$ is the unweighted mean across sectors of the fitted sectoral elasticities evaluated at each sector's own pth percentile of the initial domestic expenditure share. In each coefficient or elasticity cell, the first number is the full-sample point estimate, the number in parentheses is the median across 999 bootstrap estimates, and brackets contain percentile 95 percent confidence intervals, given by the 2.5th and 97.5th percentiles of those estimates. In each bootstrap draw, independent exponential weights are assigned to destination countries and normalized to sum to 80. The same destination weight is applied to all sectors and estimation stages, and the leave-sector-out NLLS models, country-cost changes, and linear IV model are re-estimated. The bootstrap uses the predetermined draw IDs 0 through 998; no draw is discarded based on its coefficient magnitude or numerical condition number. Because the positive weights retain every destination, the intervals condition on the observed country and sector sets, trade network, agreement support, and initial expenditure shares. Complete draw-level coefficient ranges, condition-number distributions, effective cluster counts, rank diagnostics, and first-stage estimation failures are retained in the accompanying numerical-summary files. The final column reports, once for each specification, the smallest destination-clustered Sanderson-Windmeijer conditional first-stage F statistic across all endogenous regressors for the instrument used in this table; Table 21 reports the full vector.

Table 21: First-Stage SW F Statistics for the Agreement Instruments

	EU enlargement			Broader customs unions		
Regressor associated with	(A)	(B)	(C)	(A)	(B)	(C)
Sector-specific intercepts, θ_{s0}						
C10T12	154.89	72.52	188.53	55.35	37.84	71.70
C13T15	111.37	72.30	17.84	52.63	45.04	17.23
C16	145.63	144.57	38.06	48.53	65.35	21.92
C17_18	182.83	146.97	42.67	51.41	61.80	21.39
C19	70.73	68.52	38.75	52.99	47.42	21.71
C20	223.08	105.24	48.42	43.93	39.17	18.25
C21	199.07	91.63	38.67	37.21	34.99	18.66
C22	162.35	189.73	29.25	60.45	71.97	27.44
C23	209.62	168.01	24.40	50.17	56.33	43.87
C24A	76.74	50.40	23.40	25.01	33.72	34.91
C24B	33.18	14.83	16.57	20.01	22.41	30.86
C25	176.67	130.94	24.57	42.33	50.07	41.54
C26	125.98	76.29	75.50	36.03	30.08	9.52
C27	165.28	83.03	105.12	35.29	24.96	7.36
C28	189.39	152.27	210.15	53.43	73.62	13.50
C29	148.81	81.47	142.43	79.57	52.38	142.03
C301	155.12	161.90	181.11	40.00	55.33	82.23
C302T309	238.25	121.49	211.60	31.59	29.75	31.29
C31T33	95.16	58.81	118.13	53.70	41.86	69.21
Group-specific slopes, θ_{g1}						
All manufacturing	175.83	—	—	40.02	—	—
Consumer manufacturing	—	76.52	—	—	31.74	—
Producer manufacturing	—	133.55	—	—	45.52	—
Food, beverages, tobacco	—	—	214.14	—	—	82.37
Textiles/apparel/leather	—	—	16.44	—	—	14.33
Wood/paper/printing	—	—	45.72	—	—	17.25
Petrol./chem./pharm./rubber	—	—	38.15	—	—	17.41
Minerals/metals	—	—	21.78	—	—	40.09
Machinery/electrical	—	—	150.17	—	—	8.68
Transport/other manufacturing	—	—	187.13	—	—	53.11
Minimum	33.18	14.83	16.44	20.01	22.41	7.36

Notes: Destination-clustered Sanderson-Windmeijer conditional first-stage F statistics for every endogenous regressor in Tables 19 and 20 (Sanderson and Windmeijer, 2016; Windmeijer, 2024). Intercept regressors are sector indicators times $\widehat{Q}_{ns}^{(-s)}$; slope regressors are group indicators times $\lambda_{nns,0}\widehat{Q}_{ns}^{(-s)}$. Dashes denote regressors absent from a specification. Every statistic has one numerator degree of freedom.

Table 22: GTAP Estimates Including Non-Manufacturing Sectors

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All goods	3.73	-0.56	0.32	1.09	0.20
		(1.56)	(0.31)	(0.25)	(0.51)	
(B)	Consumer manuf.	3.30	0.24	0.94	1.59	0.20
		(1.43)	(0.27)	(0.43)	(0.68)	
	Producer manuf.	4.04	-0.03	0.75	1.42	
		(1.88)	(0.24)	(0.40)	(0.68)	
	Non-manufacturing	3.94	-1.64	-0.53	0.45	
		(1.72)	(0.76)	(0.40)	(0.43)	
(C)	Food, bev., tob.	4.04	0.20	1.09	1.85	0.20
		(1.75)	(0.33)	(0.51)	(0.80)	
	Textiles/apparel/leather	2.57	0.17	0.76	1.41	
		(1.31)	(0.31)	(0.39)	(0.65)	
	Wood/paper/printing	3.95	-1.27	-0.54	0.06	
		(1.84)	(0.51)	(0.31)	(0.36)	
	Petrol./chem./pharm./rubber	1.97	0.49	1.00	1.30	
		(1.37)	(0.32)	(0.47)	(0.64)	
	Minerals/metals	4.92	0.16	1.05	1.93	
		(2.35)	(0.37)	(0.56)	(0.93)	
	Machinery/electrical	4.11	0.09	0.60	1.22	
		(2.05)	(0.32)	(0.36)	(0.57)	
	Transport/other manuf.	4.11	0.12	0.90	1.64	
		(1.85)	(0.33)	(0.46)	(0.73)	
	Non-manufacturing	3.89	-1.61	-0.51	0.45	
		(1.69)	(0.74)	(0.39)	(0.42)	

Notes: This table replicates Table 2, but adds the 17 non-manufacturing goods sectors. Standard errors are clustered by destination country and reported in parentheses.

Table 23: GTAP Estimates with Country Characteristics in the Elasticity Intercept

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	3.62	1.70	2.44	3.10	0.31
		(1.06)	(0.39)	(0.48)	(0.62)	
(B)	Consumer manuf.	3.45	2.28	3.01	3.69	0.31
		(1.11)	(0.48)	(0.58)	(0.72)	
	Producer manuf.	3.87	1.07	1.83	2.47	
		(1.24)	(0.38)	(0.43)	(0.56)	
(C)	Food, bev., tob.	3.86	3.19	4.04	4.77	0.31
		(1.41)	(0.68)	(0.78)	(0.95)	
	Textiles/apparel/leather	2.33	1.19	1.72	2.31	
		(1.11)	(0.45)	(0.45)	(0.59)	
	Wood/paper/printing	2.17	0.15	0.54	0.88	
		(1.63)	(0.45)	(0.38)	(0.48)	
	Petrol./chem./pharm./rubber	3.14	1.03	1.85	2.33	
		(1.15)	(0.33)	(0.42)	(0.55)	
	Minerals/metals	3.77	1.61	2.29	2.97	
		(1.42)	(0.59)	(0.62)	(0.74)	
	Machinery/electrical	4.89	1.09	1.70	2.43	
		(1.58)	(0.50)	(0.47)	(0.53)	
	Transport/other manuf.	5.15	0.95	1.92	2.84	
		(1.53)	(0.48)	(0.46)	(0.59)	

Notes: This table replicates Table 2, but allows sector elasticity intercepts to vary with 2004 PPP GDP per capita, population, remoteness, and landlocked status. Standard errors are clustered by destination country and reported in parentheses.

Table 24: GTAP Estimates with Intercepts Pooled within Sector Aggregates

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	2.45	2.55	3.23	3.87	0.26
		(0.70)	(0.53)	(0.64)	(0.77)	
(B)	Consumer manuf.	2.59	2.67	3.56	4.20	0.26
		(0.78)	(0.54)	(0.70)	(0.85)	
	Producer manuf.	2.09	2.17	2.66	3.14	
		(0.96)	(0.54)	(0.61)	(0.74)	
(C)	Food, bev., tob.	1.51	3.91	4.38	4.62	0.27
		(0.89)	(0.76)	(0.86)	(0.94)	
	Textiles/apparel/leather	2.60	1.73	2.24	3.03	
		(1.13)	(0.44)	(0.48)	(0.69)	
	Wood/paper/printing	1.22	1.23	1.46	1.64	
		(1.44)	(0.32)	(0.39)	(0.55)	
	Petrol./chem./pharm./rubber	2.07	1.75	2.18	2.68	
		(1.11)	(0.42)	(0.49)	(0.67)	
	Minerals/metals	2.34	2.54	3.13	3.60	
		(1.59)	(0.87)	(0.80)	(0.89)	
	Machinery/electrical	4.66	2.26	2.81	3.72	
		(1.49)	(0.57)	(0.65)	(0.84)	
	Transport/other manuf.	2.96	1.94	2.61	3.28	
		(0.93)	(0.52)	(0.57)	(0.68)	

Notes: This table replicates Table 2, but pools intercepts within manufacturing aggregates. Standard errors are clustered by destination country and reported in parentheses.

Table 25: GTAP Estimates with Sector-Specific Elasticity Schedules

Industry	$\widehat{\theta}_{s0}$	$\widehat{\theta}_{s1}$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$
19 Cattle meat	-8.38 (4.65)	15.37 (6.15)	3.10 (1.12)	5.29 (1.47)	6.66 (1.87)
20 Other meat	2.62 (1.25)	3.02 (2.16)	4.63 (1.07)	5.30 (1.39)	5.55 (1.52)
21 Vegetable oils	3.52 (1.10)	1.34 (1.38)	3.85 (0.97)	4.16 (0.95)	4.48 (1.04)
22 Dairy products	2.10 (1.08)	1.48 (1.40)	3.06 (0.67)	3.36 (0.75)	3.50 (0.82)
23 Processed rice	2.52 (1.39)	3.93 (2.51)	3.00 (1.22)	4.02 (1.07)	6.15 (1.79)
24 Sugar/molasses	4.49 (1.76)	-0.90 (3.15)	4.28 (1.32)	3.92 (1.38)	3.67 (1.95)
25 Other food	2.94 (1.65)	0.64 (2.00)	3.35 (0.77)	3.43 (0.74)	3.50 (0.79)
26 Beverages/tobacco	1.26 (1.68)	1.53 (2.07)	2.27 (0.64)	2.54 (0.62)	2.67 (0.69)
27 Textiles	-0.13 (0.50)	1.39 (1.34)	0.14 (0.34)	0.36 (0.31)	0.62 (0.44)
28 Wearing apparel	1.04 (0.59)	3.68 (1.27)	1.92 (0.51)	3.04 (0.64)	4.09 (0.91)
29 Leather products	2.72 (0.73)	-0.04 (1.10)	2.71 (0.64)	2.71 (0.58)	2.69 (0.64)
30 Wood products	0.10 (0.94)	-0.10 (1.31)	0.05 (0.42)	0.03 (0.35)	0.01 (0.43)
31 Paper/printing	0.16 (0.87)	3.49 (1.65)	1.54 (0.51)	2.16 (0.56)	2.63 (0.69)
32 Petroleum/coke	1.32 (0.99)	-0.00 (1.84)	1.32 (0.86)	1.32 (0.98)	1.32 (1.27)
33 Chemicals	1.10 (0.63)	3.45 (1.36)	1.85 (0.57)	2.47 (0.62)	2.87 (0.70)
34 Pharmaceuticals	0.87 (0.46)	2.79 (1.20)	1.22 (0.40)	1.69 (0.40)	2.06 (0.47)
35 Rubber/plastics	1.24 (0.69)	1.96 (1.31)	1.93 (0.51)	2.29 (0.56)	2.57 (0.66)
36 Non-metallic minerals	0.11 (1.35)	2.96 (1.92)	1.85 (0.59)	2.25 (0.61)	2.56 (0.69)
37 Iron/steel	2.22 (1.12)	0.99 (1.79)	2.37 (0.93)	2.61 (0.75)	2.82 (0.77)
38 Non-ferrous metals	3.37 (2.62)	1.99 (3.82)	3.81 (1.96)	4.14 (1.60)	4.60 (1.44)
39 Fabricated metals	-0.06 (0.56)	4.61 (1.27)	1.95 (0.52)	2.77 (0.66)	3.57 (0.83)
40 Computers/electronics/optical	1.00 (0.60)	3.08 (1.52)	1.48 (0.52)	1.74 (0.52)	2.17 (0.58)
41 Electrical equipment	3.71 (1.38)	4.05 (2.04)	4.47 (1.27)	5.19 (1.27)	5.96 (1.38)
42 Machinery/equipment n.e.c.	0.27 (0.74)	2.13 (2.52)	0.44 (0.61)	0.67 (0.49)	0.93 (0.53)
43 Motor vehicles	1.40 (0.64)	3.04 (1.22)	1.65 (0.61)	2.21 (0.60)	2.82 (0.68)
44 Other transport equip.	0.37 (0.96)	5.34 (2.63)	1.29 (0.71)	2.31 (0.73)	3.31 (1.02)
45 Other manufacturing	0.73 (0.60)	3.26 (0.97)	2.15 (0.53)	2.78 (0.60)	3.27 (0.69)

Notes: This table replicates Table 2, but allows every manufacturing sector to have its own intercept and slope. Standard errors are clustered by destination country and reported in parentheses. The specification uses 3,240 observations and has $R^2 = 0.31$.

Table 26: GTAP Estimates Requiring at Least 50 Percent Tariff Coverage

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	2.45	2.56	3.06	3.51	0.30
		(0.83)	(0.51)	(0.57)	(0.67)	
(B)	Consumer manuf.	2.18	3.02	3.49	3.91	0.30
		(0.91)	(0.59)	(0.63)	(0.72)	
	Producer manuf.	2.80	2.02	2.57	3.03	
		(0.93)	(0.48)	(0.55)	(0.65)	
(C)	Food, bev., tob.	2.17	3.96	4.44	4.85	0.30
		(1.44)	(0.81)	(0.82)	(0.92)	
	Textiles/apparel/leather	1.74	1.74	2.13	2.57	
		(0.97)	(0.40)	(0.43)	(0.56)	
	Wood/paper/printing	1.88	0.84	1.19	1.48	
		(1.29)	(0.34)	(0.38)	(0.50)	
	Petrol./chem./pharm./rubber	2.55	1.70	2.37	2.76	
		(1.05)	(0.42)	(0.52)	(0.63)	
	Minerals/metals	2.89	2.62	3.14	3.66	
		(1.35)	(0.80)	(0.81)	(0.89)	
	Machinery/electrical	3.27	2.29	2.69	3.18	
		(1.22)	(0.61)	(0.62)	(0.67)	
	Transport/other manuf.	3.70	1.90	2.61	3.27	
		(1.15)	(0.51)	(0.55)	(0.66)	

Notes: This table replicates Table 2, but requires origins with observed tariff changes to account for at least 50 percent of initial foreign expenditure. Standard errors are clustered by destination country and reported in parentheses.

Table 27: GTAP Estimates Using Trade-Weighted Winsorized Tariffs

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	1.90	1.82	2.21	2.55	0.30
		(0.68)	(0.45)	(0.53)	(0.62)	
(B)	Consumer manuf.	1.75	2.18	2.55	2.89	0.30
		(0.72)	(0.54)	(0.60)	(0.68)	
	Producer manuf.	2.07	1.43	1.83	2.17	
		(0.76)	(0.40)	(0.49)	(0.58)	
(C)	Food, bev., tob.	1.72	2.92	3.29	3.62	0.30
		(1.06)	(0.74)	(0.78)	(0.87)	
	Textiles/apparel/leather	1.37	1.32	1.63	1.97	
		(0.71)	(0.37)	(0.38)	(0.46)	
	Wood/paper/printing	1.39	0.64	0.89	1.10	
		(1.05)	(0.26)	(0.32)	(0.44)	
	Petrol./chem./pharm./rubber	1.97	1.17	1.68	1.98	
		(0.92)	(0.31)	(0.45)	(0.56)	
	Minerals/metals	1.93	1.91	2.25	2.60	
		(1.01)	(0.64)	(0.68)	(0.76)	
	Machinery/electrical	2.40	1.62	1.91	2.27	
		(0.87)	(0.49)	(0.51)	(0.57)	
	Transport/other manuf.	2.96	1.27	1.84	2.37	
		(0.93)	(0.40)	(0.47)	(0.59)	

Notes: This table replicates Table 2, but uses the trade-weighted winsorized applied-tariff measure `tariff95_w` from the Global Tariff Database. Standard errors are clustered by destination country and reported in parentheses.

Table 28: GTAP Estimates Using Winsorized MFN Tariffs

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	1.81	1.81	2.18	2.51	0.30
		(0.65)	(0.49)	(0.55)	(0.63)	
(B)	Consumer manuf.	1.52	2.21	2.53	2.83	0.30
		(0.64)	(0.60)	(0.64)	(0.71)	
	Producer manuf.	2.11	1.37	1.78	2.13	
		(0.78)	(0.39)	(0.48)	(0.57)	
(C)	Food, bev., tob.	1.32	2.94	3.23	3.48	0.30
		(0.99)	(0.81)	(0.82)	(0.88)	
	Textiles/apparel/leather	1.32	1.27	1.57	1.91	
		(0.73)	(0.41)	(0.46)	(0.56)	
	Wood/paper/printing	1.60	0.52	0.81	1.06	
		(1.01)	(0.28)	(0.32)	(0.41)	
	Petrol./chem./pharm./rubber	1.76	1.20	1.66	1.93	
		(0.86)	(0.34)	(0.45)	(0.55)	
	Minerals/metals	2.20	1.82	2.22	2.61	
		(1.06)	(0.62)	(0.66)	(0.75)	
	Machinery/electrical	2.35	1.54	1.83	2.18	
		(0.98)	(0.47)	(0.51)	(0.59)	
	Transport/other manuf.	2.72	1.28	1.80	2.29	
		(0.92)	(0.40)	(0.46)	(0.57)	

Notes: This table replicates Table 2, but uses the winsorized MFN measure `mfn95` from the Global Tariff Database. Standard errors are clustered by destination country and reported in parentheses.

Table 29: GTAP Estimates with a Quadratic Elasticity Schedule

Spec.	Aggregate	$\widehat{\varepsilon}'_{50}$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	3.07	2.29	2.67	3.39	0.31
		(0.88)	(0.46)	(0.52)	(0.66)	
(B)	Consumer manuf.	3.60	2.69	3.08	3.96	0.31
		(1.06)	(0.53)	(0.57)	(0.74)	
	Producer manuf.	2.85	1.81	2.21	2.81	
		(0.91)	(0.44)	(0.51)	(0.63)	
(C)	Food, bev., tob.	5.83	3.34	3.92	5.11	0.31
		(1.65)	(0.70)	(0.76)	(0.97)	
	Textiles/apparel/leather	1.50	1.48	1.38	2.46	
		(0.87)	(0.37)	(0.39)	(0.56)	
	Wood/paper/printing	3.99	-0.06	0.27	1.19	
		(1.57)	(0.49)	(0.46)	(0.47)	
	Petrol./chem./pharm./rubber	1.96	2.01	2.79	2.93	
		(1.28)	(0.54)	(0.75)	(0.75)	
	Minerals/metals	4.11	1.92	2.19	3.17	
		(1.35)	(0.65)	(0.63)	(0.75)	
	Machinery/electrical	3.07	2.27	2.63	3.12	
		(1.96)	(0.63)	(0.70)	(0.79)	
	Transport/other manuf.	3.75	1.76	2.35	3.11	
		(1.17)	(0.47)	(0.52)	(0.65)	

Notes: This table replicates Table 2, but estimates $\varepsilon_{ns}(\lambda) = \theta_{s0} + \theta_{g(s)1}\lambda + \theta_{g(s)2}\lambda^2$. The derivative column is evaluated at the unweighted average across sectors of each sector's median initial domestic expenditure share. Standard errors are clustered by destination country and reported in parentheses. Each specification uses 3,240 destination–sector observations.

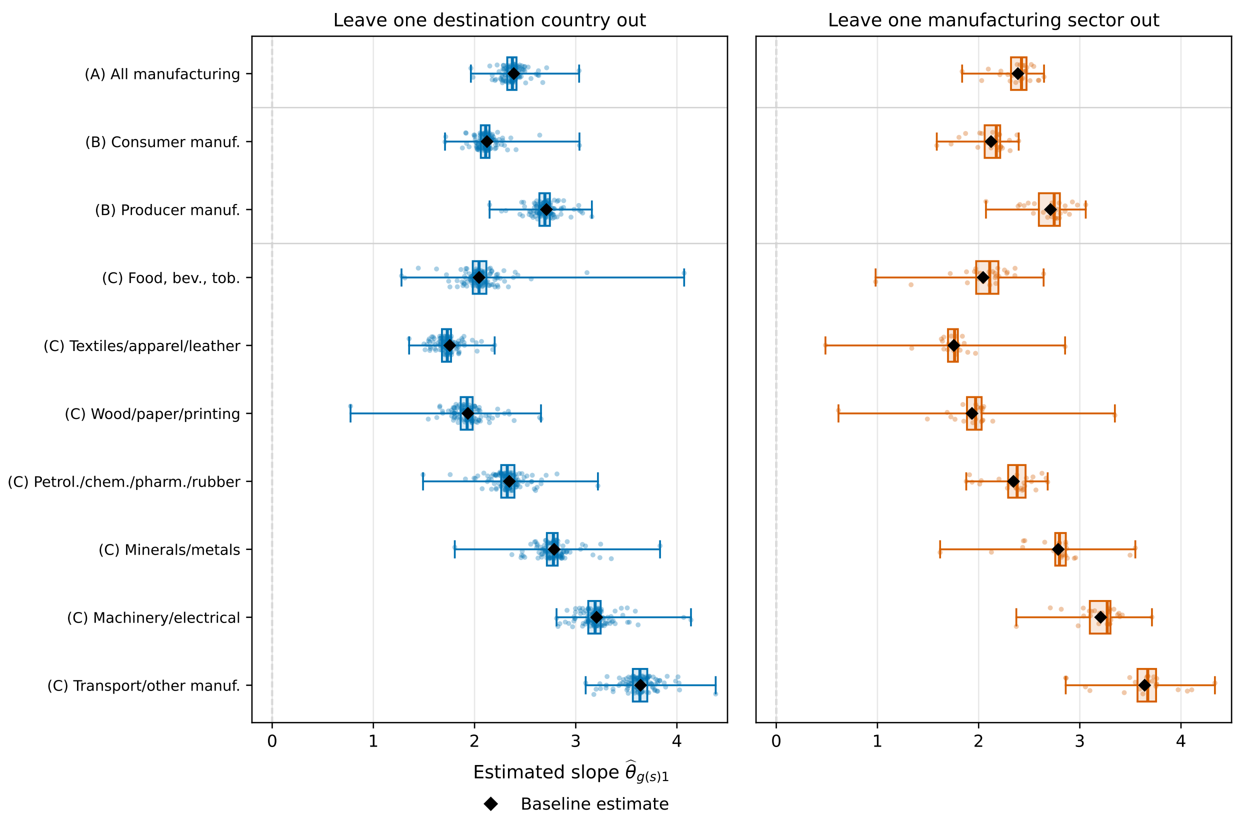

Figure 5: GTAP Leave-One-Out Estimates of the Home-Foreign Elasticity Slopes

Notes: The left panel re-estimates Specifications (A)-(C) after omitting each of the 120 destination countries in turn. The right panel re-estimates the same specifications after omitting each of the 27 manufacturing sectors in turn. Boxes show the interquartile range, vertical lines within the boxes show the median, and whiskers show the full range of leave-one-out estimates. Each point is an individual leave-one-out estimate, and black diamonds are the baseline estimates.

Table 30: Constrained GTAP Estimates Used in Quantification

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	1.78	2.72	3.09	3.41	0.30
(B)	Consumer manuf.	1.76	3.10	3.48	3.82	0.30
	Producer manuf.	1.81	2.32	2.67	2.96	
(C)	Food, bev., tob.	1.93	3.96	4.38	4.75	0.30
	Textiles/apparel/leather	1.47	1.85	2.19	2.56	
	Wood/paper/printing	0.46	1.29	1.37	1.44	
	Petrol./chem./pharm./rubber	2.29	1.69	2.29	2.64	
	Minerals/metals	2.73	2.66	3.15	3.65	
	Machinery/electrical	3.14	2.33	2.72	3.19	
	Transport/other manuf.	3.56	1.95	2.63	3.26	

Notes: This table reports constrained NLLS estimates using GTAP over 2004–2014. The estimates impose $\theta_{s0} \geq 0$ and $\theta_{s0} + \theta_{g(s)1} \geq 0$ for every sector and are used in the quantification. The table reports point estimates only. Each specification uses 3,240 destination–sector observations.

Table 31: Exact Finite-Change Constrained GTAP Estimates

Spec.	Aggregate	$\widehat{\theta}_1$	$\widehat{\varepsilon}_{25}$	$\widehat{\varepsilon}_{50}$	$\widehat{\varepsilon}_{75}$	R^2
(A)	All manufacturing	1.91	2.80	3.19	3.53	0.30
(B)	Consumer manuf.	1.97	3.17	3.59	3.97	0.30
	Producer manuf.	1.86	2.40	2.76	3.07	
(C)	Food, bev., tob.	1.91	4.01	4.43	4.80	0.30
	Textiles/apparel/leather	1.95	1.96	2.40	2.90	
	Wood/paper/printing	0.52	1.30	1.40	1.48	
	Petrol./chem./pharm./rubber	2.36	1.77	2.39	2.75	
	Minerals/metals	2.77	2.84	3.34	3.84	
	Machinery/electrical	3.69	2.49	2.94	3.50	
	Transport/other manuf.	3.81	2.14	2.87	3.55	

Notes: This table replicates Table 30, replacing the initial-share approximation by solving the exact demand equation forward from the observed initial domestic expenditure share and comparing its predicted terminal log home-foreign expenditure ratio with the observed ratio. The estimates use GTAP over 2004–2014. The table uses the same sample, Specifications (A)–(C), sector classifications, and constraints $\theta_{s0} \geq 0$ and $\theta_{s0} + \theta_{g(s)1} \geq 0$ as the comparison table. The table reports point estimates only. For each aggregate, $\widehat{\varepsilon}_p$ is the unweighted mean across sectors of the fitted sectoral elasticities evaluated at each sector's own pth percentile of the initial domestic expenditure share. The country-cost changes and elasticity parameters are estimated jointly. The residual and R^2 remain in log home-foreign expenditure-ratio units.

Theoretical Appendix

B Proofs

B.1 Proof of Proposition 1

Because utility is homogeneous of degree one, a consumer attaining utility W_n has Hicksian demand $W_n h_{ni}(\boldsymbol{p}_n)$ and total expenditure $W_n e_n(\boldsymbol{p}_n)$. The standard expenditure-share expression is therefore

$$\lambda_{ni} = \frac{p_{ni} h_{ni}(\boldsymbol{p}_n)}{e_n(\boldsymbol{p}_n)}.$$

Along the ray, equation (2), together with $\pi_{nn}(m) = 1$, gives

$$p_{ni}(m) = p_{nn}(m)\pi_{ni}(m)$$

for every $i \in \mathcal{N}$. Hicksian demand is homogeneous of degree zero in prices, so

$$h_{ni}(\boldsymbol{p}_n(m)) = h_{ni}(\boldsymbol{\pi}_n(m)),$$

and equation (3) gives

$$e_n(\boldsymbol{p}_n(m)) = p_{nn}(m) H_n(m).$$

Substituting these three relationships into the standard expenditure-share expression yields

$$\lambda_{ni}(m) = \frac{p_{nn}(m)\pi_{ni}(m) h_{ni}(\boldsymbol{\pi}_n(m))}{p_{nn}(m) H_n(m)} = \frac{\pi_{ni}(m) h_{ni}(\boldsymbol{\pi}_n(m))}{H_n(m)},$$

which is equation (4). Euler's theorem applied to the homogeneous of degree one expenditure function gives

$$H_n(m) = \sum_{i \in \mathcal{N}} \pi_{ni}(m) h_{ni}(\boldsymbol{\pi}_n(m)),$$

so the expenditure shares in equation (4) sum to one.

Differentiating $H_n(m) = e_n(\boldsymbol{\pi}_n(m))$ yields

$$\begin{aligned}
H_n'(m) &= \sum_{i \in \mathcal{N}} \frac{\partial e_n(\boldsymbol{\pi}_n(m))}{\partial \pi_{ni}} \frac{\mathrm{d}\pi_{ni}(m)}{\mathrm{d}m} \\
&= \frac{1}{m} \sum_{i \neq n} \frac{\partial e_n(\boldsymbol{\pi}_n(m))}{\partial \pi_{ni}} \pi_{ni}(m) \\
&= \frac{1}{m} \sum_{i \neq n} \pi_{ni}(m) h_{ni}(\boldsymbol{\pi}_n(m)),
\end{aligned}$$

where we used $\pi_{nn}(m) = 1$ and $\mathrm{d}\pi_{ni}(m)/\mathrm{d}m = \pi_{ni}(m)/m$ for every $i \neq n$ in the second equality and we applied Shephard's lemma in the third. Using equation (4), it follows that

$$\frac{\mathrm{d}\ln H_n(m)}{\mathrm{d}\ln m} = \sum_{i \neq n} \lambda_{ni}(m) = 1 - \lambda_{nn}(m),$$

which is exactly the equation displayed in the body.

Since $p_{nn} = w_n/a_n$, we have $Y_n = w_n L_n = p_{nn} a_n L_n$. Thus, equation (3) implies

$$W_n(m) \equiv \frac{Y_n(m)}{e_n(\boldsymbol{p}_n(m))} = \frac{a_n L_n}{H_n(m)}.$$

Hence

$$\mathrm{d}\ln W_n = -(1-\lambda_{nn})\mathrm{d}\ln m.$$

Integrating between m_0 and m_1 gives equation (7). Moreover, since $X_{nn}/X_{nF} = \lambda_{nn}/(1-\lambda_{nn})$, the definition of $\varepsilon_n(m)$ gives

$$\mathrm{d}\ln \frac{\lambda_{nn}}{1-\lambda_{nn}} = \varepsilon_n(m)\mathrm{d}\ln m.$$

Integrating this identity gives equation (6).

If, in addition, $\varepsilon_n(m) > 0$ throughout the relevant interval, $\lambda_{nn}(m)$ is strictly increasing and can index the path. Since

$$\mathrm{d}\ln \lambda_{nn} = (1-\lambda_{nn})\varepsilon_n(m)\mathrm{d}\ln m,$$

we obtain

$$\mathrm{d}\ln W_n = -\frac{\mathrm{d}\ln \lambda_{nn}}{\varepsilon_n(\lambda_{nn})}.$$

Changing variables from m to λ_{nn} and integrating between the two points gives equation (8).

B.2 Proof of Corollary 1

It remains to establish that the proportional foreign-price ray converges to autarky. Denote by H_n^A the finite amount of the domestic good needed to reach one util when consuming only the domestic good. It is finite given the assumption preceding the corollary. This bundle is feasible in the expenditure minimization problem at every m, so $H_n(m) \leq H_n^A$. The function $H_n(m)$ is increasing and concave, and therefore converges to a finite limit. Moreover, H_n' is nonnegative, nonincreasing, and integrable on $[1, \infty)$, which implies $mH_n'(m) \to 0$. Since the foreign expenditure share is $mH_n'(m)/H_n(m)$, it follows that $\lambda_{nn}(m) \to 1$.

To identify the limiting expenditure, let $\boldsymbol{h}_n(\boldsymbol{\pi}_n(m))$ denote Hicksian demand. The bound $H_n(m) \leq H_n^A$ implies $h_{ni}(\boldsymbol{\pi}_n(m)) \leq H_n^A/\pi_{ni}(m) \to 0$ for every foreign i. Any limit of the domestic component must therefore attain one unit of utility using the domestic good alone, so $\lim_{m\to\infty} H_n(m) = H_n^A$. Consequently, factor-income welfare converges to positive autarky welfare, $W_n^A = a_n L_n/H_n^A > 0$. Applying equation (7) between $m = 1$ and a finite m and taking $m \to \infty$ yields equation (9). If, in addition, $\varepsilon_n(m) > 0$ for every finite $m \geq 1$, applying equation (8) and using $\lambda_{nn}(m) \to 1$ yields equation (10).

B.3 Proof of Proposition 2

1. Existence of an auxiliary equilibrium. Fix a finite $m \geq 1$ and consider the auxiliary economy described in the text. All primitives are the same as in the original economy except that country n has Leontief preferences

$$\widetilde{u}_n^m(\boldsymbol{x}_n) \equiv \min_{i \in \mathcal{N}} \left\{ \frac{x_{ni}}{h_{ni}(\boldsymbol{\pi}_n(m))} \right\}.$$

The auxiliary economy has finitely many commodities, consumers, and production activities. Linear production and iceberg transportation generate closed convex constant-returns-to-scale production sets. Producing a positive amount of country i's good requires country i's factor, and each country has a positive endowment of its factor. Every factor can be transformed directly into a differentiated good desired by consumers. Preferences are continuous, convex, and locally nonsatiated. Country n's fixed-proportions auxiliary preference has the same properties because $h_{ni}(\boldsymbol{\pi}_n(m)) > 0$ for every i at every finite m.

These primitives satisfy the conditions of Theorem II of Arrow and Debreu (1954).[15] Therefore, the auxiliary economy has a competitive equilibrium.

2. Construction of the country-n policies that implement the target price ratios in the original economy. The auxiliary economy has no trade-policy wedges. Let $\boldsymbol{c}(m) \equiv (c_i(m))_{i \in \mathcal{N}}$ denote the vector of unit production costs in an auxiliary equilibrium, normalized by $c_n(m) = 1$.

We use these unit costs to construct the source-specific policies $\{t_{ni}(m)\}_{i \in \mathcal{N}}$ that implement the target price ratios $\boldsymbol{\pi}_n(m)$ in the original economy. Recall that equation (2) requires $p_{ni}(m) = p_{nn}(m)\pi_{ni}(m)$. Since $p_{nn}(m) = c_n(m)$, implementing each target relative price requires $p_{ni}(m) = c_n(m)\pi_{ni}(m)$. The untaxed delivered cost is $\mathcal{D}_{ni}c_i(m)$, so set $t_{nn}(m) = 0$ and, for every $i \neq n$, define the source-specific per-unit tariff or subsidy by

$$t_{ni}(m) = c_n(m)\pi_{ni}(m) - \mathcal{D}_{ni}c_i(m). \tag{24}$$

The wedge may be positive, representing a tariff, or negative, representing a subsidy. Whenever $\mathcal{D}_{ni}c_i(m) > 0$, the equivalent ad valorem rate is $t_{ni}(m)/[\mathcal{D}_{ni}c_i(m)]$. The consumer prices in country n are

$$p_{nn}(m) = c_n(m)$$

and, for every $i \neq n$,

$$p_{ni}(m) = \mathcal{D}_{ni}c_i(m) + t_{ni}(m) = c_n(m)\pi_{ni}(m).$$

Consequently,

$$\frac{p_{ni}(m)}{p_{nn}(m)} = \pi_{ni}(m),$$

which is the proportional ray in equation (2).

[15] Their Theorem II allows a finite collection of commodities differentiated by location, closed convex production sets, continuous convex non-satiated preferences, and consumers endowed with productive labor. Linear one-factor production and iceberg transportation satisfy their production assumptions, and each country's factor directly produces a commodity desired by consumers.

3. Verification of equilibrium in the original economy. Because the auxiliary and original economies have the same endowments and production and transportation technologies, we take the auxiliary-equilibrium unit costs $\boldsymbol{c}(m)$ and consumption allocation $\boldsymbol{x}(m) \equiv (x_{ki}(m))_{k,i\in\mathcal{N}}$ as the corresponding objects in a candidate for an equilibrium of the original economy. For this candidate, set $w_i(m) = a_i c_i(m)$ and impose the country-n trade policies constructed in Step 2. Because production technologies, wages, and producer prices are unchanged from the auxiliary economy, firms' profit-maximization conditions continue to hold in the candidate original equilibrium. We now verify consumer optimization, market clearing, and balanced trade in the original economy.

3a. Country n. Since physical output is $y_n = a_n L_n$, country n's factor income in the auxiliary economy is $c_n(m)y_n$. Its fixed-proportions demand therefore satisfies

$$x_{ni}(m) = s_n(m)h_{ni}(\boldsymbol{\pi}_n(m)),$$

where

$$s_n(m) = \frac{c_n(m)y_n}{\sum_{j\in\mathcal{N}} \mathcal{D}_{nj}c_j(m)h_{nj}(\boldsymbol{\pi}_n(m))} \tag{25}$$

is the scale of country n's fixed-proportions consumption bundle in the auxiliary equilibrium. Because the candidate original-economy allocation uses these same consumption quantities and Hicksian demand is homogeneous of degree zero in prices, this bundle is proportional to country n's Hicksian demand for one unit of utility at $\boldsymbol{p}_n(m) = c_n(m)\boldsymbol{\pi}_n(m)$.

Tariff revenue, with a negative value representing the fiscal cost of subsidies, is

$$\mathcal{R}_n(m) \equiv \sum_{i\neq n} t_{ni}(m)x_{ni}(m).$$

Country n's consumer expenditure in the candidate equilibrium satisfies

$$\begin{aligned}\sum_{i\in\mathcal{N}} p_{ni}(m)x_{ni}(m) &= \sum_{i\in\mathcal{N}} \mathcal{D}_{ni}c_i(m)x_{ni}(m) + \mathcal{R}_n(m)\\ &= c_n(m)y_n + \mathcal{R}_n(m),\end{aligned}$$

where the first equality follows from the construction of country n's trade policy in Step 2—with $p_{ni}(m) = \mathcal{D}_{ni}c_i(m) + t_{ni}(m)$ for all $i \neq n$—and the definition of $\mathcal{R}_n(m)$, while the second equality follows because country n exhausts its budget in the auxiliary equilibrium. The final expression is labor income plus the lump-sum rebate of tariff revenue. Country n therefore affords the bundle constructed above. Because this bundle is proportional to country n's Hicksian demand at the implemented consumer prices and exhausts its disposable income, homotheticity implies that it is optimal under country n's original preferences. It remains to show that every other country is also optimizing under its original preferences and that markets clear and trade is balanced.

3b. Other countries. For every country $k \neq n$, its consumption bundle in the auxiliary equilibrium maximizes $u_k(\boldsymbol{x}_k)$ subject to

$$\sum_i \mathcal{D}_{ki}c_i x_{ki} \leq c_k y_k.$$

This is also its optimization problem in the candidate original equilibrium. Thus every other country is also optimizing under its original preferences and trade costs.

Goods-market clearing in the auxiliary economy is

$$y_i = \sum_{k \in \mathcal{N}} \mathcal{D}_{ki} x_{ki},$$

which is also goods-market clearing in the candidate original equilibrium. Since $y_i = a_i L_i$, this also gives factor-market clearing, with

$$w_i(m) = a_i c_i(m).$$

To establish balanced trade, country k's budget in the auxiliary equilibrium is

$$\sum_{i \in \mathcal{N}} \mathcal{D}_{ki} c_i x_{ki} = c_k y_k.$$

Define its import expenditure at source values, inclusive of the resources used in transportation, by

$$M_k \equiv \sum_{i \neq k} \mathcal{D}_{ki} c_i x_{ki}.$$

The budget implies $M_k = c_k y_k - c_k x_{kk}$. Goods-market clearing for country k's variety gives

$$y_k - x_{kk} = \sum_{\ell \neq k} \mathcal{D}_{\ell k} x_{\ell k}.$$

Therefore,

$$M_k = c_k \sum_{\ell \neq k} \mathcal{D}_{\ell k} x_{\ell k},$$

which is country k's export revenue. Imports and exports are equal country by country.

4. Initial and autarky endpoints. At $m = 1$, the absence of tariffs implies $\sum_i \mathcal{D}_{ni} c_i(1) x_{ni}(1) = c_n(1) y_n$. Homotheticity makes $\boldsymbol{x}_n(1)$ proportional to $\boldsymbol{h}_n(\boldsymbol{\pi}_n(1))$. Every other country's initial consumption remains optimal in the auxiliary economy, and goods-market clearing is unchanged. The initial equilibrium allocation and unit costs therefore constitute an equilibrium of the auxiliary economy. Selecting this auxiliary equilibrium gives

$$\pi_{ni}(1) = \mathcal{D}_{ni} \frac{c_i(1)}{c_n(1)},$$

so equation (24) yields $t_{ni}(1) = 0$.

The proof of Corollary 1 shows that $h_{ni}(\boldsymbol{\pi}_n(m)) \to 0$ for every $i \neq n$ and that domestic Hicksian demand converges to the finite amount required to attain one unit of utility using the domestic good alone. Equation (25) therefore implies that $s_n(m)$ remains bounded, so delivered import quantities converge to zero. At the autarky endpoint there is no trade involving country n, and tariff revenue is zero.

C An Armington Gravity Model

Here, we restrict preferences to a parametric expenditure system. Our goal is to provide a simple general-equilibrium gravity model that strictly generalizes the canonical CES framework. The substantive generalization is in the upper nest between the home good and the foreign aggregate. We retain CES aggregation across foreign sources but replace the CES upper nest with a constant-difference-of-elasticities (CDE) expenditure system (Hanoch, 1975). Whereas the CES restriction generates the constant, zeroth-order elasticity schedule corresponding to $J = 0$ in equation (18), the CDE upper nest generates exactly the linear, first-order schedule corresponding to $J = 1$.

A non-nested CDE system over national varieties could also be used to construct a gravity model with variable substitution elasticities. We use the nested specification because one of our focuses is the distinction between substitution across foreign sources and substitution between the home good and the foreign aggregate. The nested system allows the home-foreign elasticity schedule to vary separately from the elasticity across foreign sources.

Sections C.1 and C.2 develop the one-sector version, corresponding to a simplification of Section 2. Section C.3 applies the same demand system within each sector of the multi-sector environment in Section 3. Section D shows that the gains-from-trade formula is unchanged when the Armington goods are replaced by Krugman varieties with free entry. Section E introduces heterogeneous firms and Pareto productivity and derives the corresponding welfare-relevant elasticity schedule.

C.1 Environment

Except for the demand system, the environment and notation are the same as in Section 2.[16]

Notation. For later use, let $E_n \equiv \sum_{i\in\mathcal{N}} X_{ni} = Y_n$ denote country n's total expenditure and, for each foreign source $i \neq n$, let $\mu_{ni} \equiv X_{ni}/X_{nF}$ denote its share of country n's foreign expenditure. Thus, $X_{ni}/E_n = (1-\lambda_{nn})\mu_{ni}$.

Preferences. We use a nested homothetic CDE-CES expenditure system: a CDE upper nest over home goods and an aggregation of foreign goods, and a constant-elasticity-of-substitution (CES) foreign-source aggregator. Because preferences are homothetic, they can be represented by an implicit unit expenditure function.[17]

The unit expenditure function $e_n(\boldsymbol{p}_n)$ is defined implicitly by

$$\omega_{nn}\left(\frac{e_n}{p_{nn}}\right)^{\theta_0} + \left[\sum_{i\neq n}\omega_{ni}\left(\frac{e_n}{p_{ni}}\right)^{\rho}\right]^{(\theta_0+\theta_1)/\rho} = 1, \tag{26}$$

where $\theta_0 > 0$, $\theta_0 + \theta_1 > 0$, $\rho > 0$, $\omega_{nn} > 0$, and $\omega_{ni} > 0$. The parameters θ_0, θ_1, and ρ shape substitution elasticities, as shown below. The parameters ω_{ni} shape relative expenditure

[16] The Armington formulation is not essential; see Sections D and E.

[17] We use a CES lower nest for simplicity only. A single-nested CDE model would already be strictly more general than the canonical quantitative trade models.

shares across foreign sources, while ω_{nn} shifts demand for the home good relative to the foreign aggregation. When $\theta_1 = 0$, the elasticity between the home aggregate and the foreign aggregate is constant and equal to θ_0. When, in addition, $\rho = \theta_0$, equation (26) is equivalent to the standard CES unit expenditure function/price index, where the elasticity of substitution is $1 + \theta_0$.[18]

C.2 Characterization

Define $A_{nn} \equiv \omega_{nn} \left(e_n/p_{nn}\right)^{\theta_0}$ and $A_{nF} \equiv \left[\sum_{i \neq n} \omega_{ni} \left(e_n/p_{ni}\right)^{\rho}\right]^{(\theta_0+\theta_1)/\rho}$. According to equation (26),

$$A_{nn} + A_{nF} = 1. \tag{27}$$

Let

$$B_n \equiv \theta_0 A_{nn} + (\theta_0 + \theta_1) A_{nF}.$$

Holding preference weights fixed and differentiating equation (27) with respect to prices gives

$$\theta_0 A_{nn}(d \ln e_n - d \ln p_{nn}) + (\theta_0 + \theta_1) A_{nF} \left(d \ln e_n - \sum_{i \neq n} \mu_{ni} d \ln p_{ni} \right) = 0,$$

where

$$\mu_{ni} \equiv \frac{\omega_{ni} p_{ni}^{-\rho}}{\sum_{k \neq n} \omega_{nk} p_{nk}^{-\rho}}.$$

Solving the previous expression for $d \ln e_n$ and combining with Shephard's lemma yields the domestic expenditure share

$$\lambda_{nn} = \frac{X_{nn}}{E_n} = \frac{\partial \ln e_n}{\partial \ln p_{nn}} = \frac{\theta_0 A_{nn}}{B_n}, \tag{28}$$

and, for each foreign source $i \neq n$,

$$\lambda_{ni} = \frac{X_{ni}}{E_n} = \frac{\partial \ln e_n}{\partial \ln p_{ni}} = \frac{(\theta_0 + \theta_1) A_{nF}}{B_n} \mu_{ni}. \tag{29}$$

The restricted CDE system in equation (26) can be written as a CDE system across two aggregates,

$$\omega_{nn} \left(\frac{e_n}{p_{nn}} \right)^{\theta_0} + \left(\frac{e_n}{P_{nF}} \right)^{\theta_0 + \theta_1} = 1, \tag{30}$$

given the following definition

$$P_{nF} \equiv \left(\sum_{i \neq n} \omega_{ni} p_{ni}^{-\rho} \right)^{-1/\rho}.$$

[18] The word "restricted" refers to the fact that we restrict the exponents to be common across foreign sources. A non-nested homothetic CDE system would be $\sum_i \omega_{ni} \left(e_n/p_{ni}\right)^{\gamma_{ni}} = 1$. Our specification instead imposes a common within-foreign exponent ρ.

The restriction in equation (26) is that the many foreign sources enter the foreign aggregate with the common exponent ρ. At the upper level, however, equation (30) is an unrestricted two-good CDE system, up to the normalization of the foreign-aggregator's weight.

The upper-level exponents are θ_0 and $\theta_0 + \theta_1$, while ρ governs substitution across foreign sources inside the foreign aggregate. The share of foreign expenditure on source country i is

$$\mu_{ni} \equiv \frac{X_{ni}}{X_{nF}} = \frac{\omega_{ni} p_{ni}^{-\rho}}{\sum_{k \neq n} \omega_{nk} p_{nk}^{-\rho}}.$$

Therefore, for any two foreign source countries $i, j \neq n$,

$$\frac{X_{ni}}{X_{nj}} = \frac{\omega_{ni}}{\omega_{nj}} \left(\frac{p_{ni}}{p_{nj}} \right)^{-\rho}.$$

Relative expenditure across foreign source countries is CES, and the across-foreign-source trade elasticity is ρ.

According to equations (28) and (29),

$$\frac{X_{nn}}{X_{nF}} = \frac{\lambda_{nn}}{1 - \lambda_{nn}} = \frac{\theta_0 A_{nn}}{(\theta_0 + \theta_1) A_{nF}}.$$

Log differentiating the previous expression holding all ω terms fixed,

$$d \ln \frac{X_{nn}}{X_{nF}} = d \ln A_{nn} - d \ln A_{nF},$$

where, again holding all ω terms fixed, we have

$$d \ln A_{nn} = \theta_0 (d \ln e_n - d \ln p_{nn}),$$

and

$$d \ln A_{nF} = (\theta_0 + \theta_1)(d \ln e_n - d \ln P_{nF}).$$

Using Shephard's lemma, we have

$$d \ln e_n = \lambda_{nn} d \ln p_{nn} + (1 - \lambda_{nn}) d \ln P_{nF}.$$

The previous expressions therefore yield

$$d \ln \frac{X_{nn}}{X_{nF}} = (\theta_0 + \theta_1 \lambda_{nn})\, d \ln \left(\frac{P_{nF}}{p_{nn}} \right).$$

The previous expression implies

$$\varepsilon_n(\lambda_{nn}) \equiv \frac{\partial \ln(X_{nn}/X_{nF})}{\partial \ln(P_{nF}/p_{nn})} = \theta_0 + \theta_1 \lambda_{nn}.$$

Discussion. The three parameters in the unit expenditure function in equation (26) have distinct roles. First, the across-foreign-source trade elasticity is ρ. Second, the value of the home-foreign trade elasticity at $\lambda_{nn} = 0$ is θ_0. Third, the slope of the home-foreign trade elasticity in the domestic expenditure share λ_{nn} is θ_1. Thus the model allows the across-foreign elasticity, the slope of the home-foreign elasticity, and the gap between the two elasticities at a given domestic expenditure share to vary separately.

Gains from Trade. Substituting the linear elasticity schedule into Proposition 1 gives the following closed form for two equilibria, indexed by 0 and 1, across which country n's productivity, labor, and preference weights are fixed:

$$\ln \frac{W_n^0}{W_n^1} = \frac{1}{\theta_0} \ln \left[\frac{\lambda_{nn}^1 \left(\theta_0 + \theta_1 \lambda_{nn}^0\right)}{\lambda_{nn}^0 \left(\theta_0 + \theta_1 \lambda_{nn}^1\right)} \right].$$

If equilibrium 0 is the trade equilibrium, indicated with T, and equilibrium 1 is autarky, then $\lambda_{nn}^1 = 1$ and gains from trade are

$$\ln \frac{W_n^T}{W_n^A} = \frac{1}{\theta_0} \ln \left[\frac{\theta_0 + \theta_1 \lambda_{nn}^T}{(\theta_0 + \theta_1)\lambda_{nn}^T} \right]. \tag{31}$$

As $\theta_1 \to 0$, equation (31) converges to the CES/ACR formula (Arkolakis et al., 2012)

$$\lim_{\theta_1 \to 0} \ln \frac{W_n^T}{W_n^A} = \frac{-1}{\theta_0} \ln \left(\lambda_{nn}^T\right). \tag{32}$$

The elasticity θ_0 in equation (32) is the elasticity between home expenditure and aggregate foreign expenditure. The across-foreign elasticity ρ does not enter the gains formula. When $\theta_1 = 0$ and $\rho = \theta_0$, the home-foreign elasticity equals the elasticity across foreign sources and using ρ for welfare is therefore correct.

C.3 The Multi-Sector Version

The environment and notation are those of Section 3.1. We impose Cobb-Douglas aggregation across sectors,

$$\mathcal{E}_n \left((e_{ns})_{s \in \mathcal{S}}\right) = \prod_{s \in \mathcal{S}} \left(\frac{e_{ns}}{\alpha_{ns}}\right)^{\alpha_{ns}},$$

where $\alpha_{ns} > 0$ and $\sum_{s \in \mathcal{S}} \alpha_{ns} = 1$, and use the CDE-CES expenditure system from Section C within each sector:

$$\omega_{nns} \left(\frac{e_{ns}}{p_{nns}}\right)^{\theta_{s0}} + \left[\sum_{i \neq n} \omega_{nis} \left(\frac{e_{ns}}{p_{nis}}\right)^{\rho_{ns}}\right]^{(\theta_{s0} + \theta_{g(s)1})/\rho_{ns}} = 1,$$

where $\theta_{s0} > 0$, $\theta_{s0} + \theta_{g(s)1} > 0$, $\rho_{ns} > 0$, and all preference weights are positive. As in the empirical specification, $g(s)$ denotes the sectoral aggregate containing sector s.

The derivations above apply separately within each sector. In particular, relative expenditure across foreign sources is CES:

$$\frac{X_{nis}}{X_{njs}} = \frac{\omega_{nis}}{\omega_{njs}} \left(\frac{p_{nis}}{p_{njs}}\right)^{-\rho_{ns}} \qquad \text{for all } i, j \neq n,$$

while the sectoral home-foreign elasticity is

$$\varepsilon_{ns}(\lambda_{nns}) = \theta_{s0} + \theta_{g(s)1} \lambda_{nns}.$$

Thus ρ_{ns} is the elasticity across foreign sources, θ_{s0} is the intercept of the home-foreign elasticity schedule, and $\theta_{g(s)1}$ is its slope in the domestic expenditure share. Under Cobb-Douglas aggregation, substituting this linear elasticity schedule into equation (14) gives

$$\ln \frac{W_n^T}{W_n^A} = \sum_{s \in \mathcal{S}} \frac{\alpha_{ns}}{\theta_{s0}} \ln \left[\frac{\theta_{s0} + \theta_{g(s)1}\lambda_{nns}^T}{(\theta_{s0} + \theta_{g(s)1})\lambda_{nns}^T} \right].$$

D A Krugman Gravity Model

We now replace the Armington good produced by each country with an endogenous mass of differentiated firm varieties, as in Krugman (1980). We retain the CDE upper tier between the domestic and foreign aggregates and CES demand with exponent ρ across firm varieties. We show that the resulting gains-from-trade formula is exactly the same as in the Armington model.

Origin i has a mass M_i of symmetric firms. Each firm has marginal cost $c_i = w_i/a_i$, and entry requires f_i^e units of country i's factor. Delivering a variety from country i to country n requires the iceberg trade cost $\mathcal{D}_{ni}$, with $\mathcal{D}_{nn} = 1$. Because the elasticity of substitution across firm varieties is $1 + \rho$, the consumer price of each variety is

$$q_{ni} = \frac{1+\rho}{\rho} c_i \mathcal{D}_{ni}.$$

Let $\zeta_{ni} > 0$ denote a fixed destination-origin demand shifter. The price index for the composite of country i's varieties sold in country n is

$$P_{ni} = \left[M_i \zeta_{ni}^{\rho} q_{ni}^{-\rho} \right]^{-1/\rho}.$$

The domestic and foreign price indexes are therefore

$$P_{nn} = \left[M_n \zeta_{nn}^{\rho} q_{nn}^{-\rho} \right]^{-1/\rho}$$

and

$$P_{nF} = \left[\sum_{i \neq n} M_i \zeta_{ni}^{\rho} q_{ni}^{-\rho} \right]^{-1/\rho}.$$

Country n's unit expenditure function is defined by the same CDE upper tier as in equation (30):

$$\omega_{nn} \left(\frac{e_n}{P_{nn}} \right)^{\theta_0} + \left(\frac{e_n}{P_{nF}} \right)^{\theta_0 + \theta_1} = 1.$$

This upper tier implies that the elasticity between domestic and foreign expenditure varies with the domestic expenditure share. However, within the CES variety aggregate, each atomistic firm faces the constant demand elasticity $1 + \rho$ and therefore charges the same constant markup $(1 + \rho)/\rho$ as in the standard CES Krugman model.

A constant markup, the free-entry condition, and factor-market clearing together imply that the mass of firms is independent of trade costs and expenditure patterns, as in the standard (CES) Krugman model. Let $\overline{R}_i$ denote the total revenue of a representative country-i firm across all destinations. Operating profits are $\overline{R}_i/(1+\rho)$, so free entry implies

$$\frac{\overline{R}_i}{1+\rho} = w_i f_i^e.$$

The variable factor bill is the fraction $\rho/(1+\rho)$ of revenue and therefore equals $\rho w_i f_i^e$. Including the entry cost, each firm uses $(1+\rho)f_i^e$ units of the factor. Factor-market clearing then gives

$$M_i = \frac{L_i}{(1+\rho)f_i^e}. \tag{33}$$

Thus M_i depends on the factor endowment and entry cost but not on trade costs or expenditure patterns.

For country n, the domestic price index can be written as

$$P_{nn} = \frac{1+\rho}{\rho}\frac{w_n}{a_n}\frac{1}{\zeta_{nn}M_n^{1/\rho}}.$$

Since $Y_n = w_n L_n$,

$$\frac{Y_n}{P_{nn}} = \frac{\rho}{1+\rho} a_n L_n \zeta_{nn} M_n^{1/\rho}.$$

The previous expression implies that Y_n/P_{nn} is unchanged across equilibria for fixed values of technology, a_n, factor endowments, L_n, demand shifters, ζ_{nn}, and entry costs. The same cancellation between nominal income and the domestic price is used to derive the Armington welfare formula.

Let λ_{nn} now denote the share of country n's expenditure allocated to the domestic firm composite. The derivation in Section C.2 applies directly to the CDE upper tier:

$$\mathrm{d}\ln\left(\frac{e_n}{P_{nn}}\right) = \frac{\mathrm{d}\ln\lambda_{nn}}{\theta_0+\theta_1\lambda_{nn}}.$$

More generally, before imposing equation (33), the welfare comparison between equilibria 0 and 1 would be

$$\ln\frac{W_n^0}{W_n^1} = \int_{\lambda_{nn}^0}^{\lambda_{nn}^1}\frac{\mathrm{d}\ln\lambda}{\theta_0+\theta_1\lambda} + \frac{1}{\rho}\ln\frac{M_n^0}{M_n^1}.$$

Equation (33) makes the second term zero. Therefore, if equilibrium 0 is the trade equilibrium and equilibrium 1 is autarky,

$$\ln\frac{W_n^T}{W_n^A} = \int_{\lambda_{nn}^T}^{1}\frac{\mathrm{d}\ln\lambda}{\theta_0+\theta_1\lambda} = \frac{1}{\theta_0}\ln\left[\frac{\theta_0+\theta_1\lambda_{nn}^T}{(\theta_0+\theta_1)\lambda_{nn}^T}\right].$$

This is exactly the Armington gains-from-trade formula in equation (31). When $\theta_1 = 0$, it reduces to the canonical Krugman result in Arkolakis et al. (2012).

E A Melitz-Pareto Gravity Model

We next replace the symmetric firms in Section D with heterogeneous firms and endogenous selection into each destination, as in Melitz (2003), imposing Pareto distributed productivities. We retain the CDE upper tier between the domestic and foreign aggregates and CES demand with exponent ρ across firm varieties. This extension nests the constant-elasticity welfare result in Feenstra et al. (2018) and shows how that result changes when the home-foreign elasticity varies with the domestic expenditure share.

Each origin has a fixed mass of potential firms, normalized to one. A firm in country i draws productivity φ from a Pareto distribution with $\Pr(\Phi \geq \varphi) = \varphi^{-\kappa}$ for $\varphi \geq 1$, where $\kappa > \rho$. Serving country n requires f_{ni} units of country i's factor, in addition to the iceberg trade cost $\mathcal{D}_{ni}$. Let φ^*_{ni} denote the productivity cutoff for serving country n and let $N_{ni} = (\varphi^*_{ni})^{-\kappa}$ denote the mass of active country-i firms in that market. We assume active selection as is standard in the literature, and impose this throughout the path to autarky: every cutoff for a market served remains strictly above the Pareto lower support, $\varphi^*_{ni} > 1$, including domestic cutoffs in the autarky limit. Since the elasticity of substitution across firm varieties is $1 + \rho$, firms charge the constant markup $(1 + \rho)/\rho$. Integrating over the active firms gives the price index

$$P_{ni} = A_{ni} c_i \mathcal{D}_{ni} N_{ni}^{1/\kappa - 1/\rho},$$

where $A_{ni} > 0$ collects the markup, Pareto, and fixed demand-shifter terms. The zero-profit condition for the cutoff firm and the Pareto distribution also imply

$$X_{ni} = \frac{(1 + \rho)\kappa}{\kappa - \rho} w_i f_{ni} N_{ni}.$$

Thus, changes in the mass of firms serving a market can be recovered from changes in expenditure and the origin wage.

Country n's unit expenditure function is defined by the same CDE upper tier as in equation (30), replacing the Armington prices with the domestic and foreign firm-variety price indexes. Let

$$\varepsilon_n^{HF}(\lambda) \equiv \theta_0 + \theta_1 \lambda$$

denote the home-foreign expenditure elasticity generated by this upper tier. Consumers own the domestic firms, so Y_n includes factor income and firm profits. In the one-sector model, combining factor-market clearing and the balanced-budget condition yields

$$Y_n = \frac{1 + \rho}{1 + \rho - \rho/\kappa} w_n L_n.$$

The coefficient multiplying factor income is constant across equilibria. Hence, $d \ln Y_n = d \ln w_n$. Since $X_{nn} = \lambda_{nn} Y_n$, the preceding expenditure equation implies

$$d \ln N_{nn} = d \ln X_{nn} - d \ln w_n = d \ln \lambda_{nn}.$$

The domestic firm-variety price index therefore satisfies

$$d \ln P_{nn} = d \ln w_n + \left(\frac{1}{\kappa} - \frac{1}{\rho}\right) d \ln \lambda_{nn},$$

and thus

$$d \ln \left(\frac{Y_n}{P_{nn}}\right) = \left(\frac{1}{\rho} - \frac{1}{\kappa}\right) d \ln \lambda_{nn}.$$

The CDE upper tier implies

$$d \ln \left(\frac{e_n}{P_{nn}}\right) = \frac{d \ln \lambda_{nn}}{\varepsilon_n^{HF}(\lambda_{nn})}.$$

Using $W_n = Y_n/e_n$ and combining the previous two expressions gives

$$d \ln W_n = -\left[\frac{1}{\varepsilon_n^{HF}(\lambda_{nn})} + \frac{1}{\kappa} - \frac{1}{\rho}\right] d \ln \lambda_{nn}.$$

It is therefore useful to define the welfare-relevant Melitz-Pareto elasticity schedule by

$$\frac{1}{\varepsilon_n^{MP}(\lambda)} \equiv \frac{1}{\theta_0 + \theta_1 \lambda} + \frac{1}{\kappa} - \frac{1}{\rho}. \tag{34}$$

Provided this expression is positive and the interior-selection restriction above holds along the path to autarky, gains from trade are

$$\begin{aligned}
\ln \frac{W_n^T}{W_n^A} &= \int_{\lambda_{nn}^T}^{1} \frac{d \ln \lambda}{\varepsilon_n^{MP}(\lambda)} \\
&= \frac{1}{\theta_0} \ln \left[\frac{\theta_0 + \theta_1 \lambda_{nn}^T}{(\theta_0 + \theta_1)\lambda_{nn}^T}\right] + \left(\frac{1}{\kappa} - \frac{1}{\rho}\right) \ln \left(\frac{1}{\lambda_{nn}^T}\right).
\end{aligned}$$

The first term is the Armington and Krugman expression in equation (31); the second term captures firm selection.

To relate equation (34) to Feenstra et al. (2018), set $\theta_1 = 0$, let $\omega = 1 + \theta_0$ denote their macro elasticity of substitution, and let $\sigma = 1 + \rho$ denote their micro elasticity of substitution. The welfare-relevant elasticity is then constant and can be written

$$\varepsilon_n^{MP} = \frac{\kappa(\omega - 1)}{(\sigma - \omega)\left(\kappa/(\sigma - 1) - 1\right) + \sigma - 1},$$

which is their trade-elasticity formula. When $\sigma = \omega$, or equivalently $\rho = \theta_0$, this elasticity equals the Pareto shape parameter κ, as in the standard Melitz-Pareto model. Equation (34) therefore replaces their constant macro elasticity with the linear home-foreign elasticity schedule while retaining their firm-selection mechanism.

Quantitative Appendix

F Intermediate-Input Multiplier

This section derives equation (23). Let R_n denote the increase in country n's aggregate unit expenditure moving from trade to autarky while holding all domestic sectoral unit costs fixed. Under the nested CES system in Section 5.2,

$$\ln R_n = \frac{1}{1-\sigma_U} \ln \left[\alpha_{nM} \widehat{E}_{nM}^{1-\sigma_U} + \alpha_{nG} \widehat{E}_{nG}^{1-\sigma_U} + \alpha_{n0} \right].$$

Let $\widehat{P}_n \equiv P_n^A / P_n^T$, $\widehat{w}_n \equiv w_n^A / w_n^T$, and $\widehat{c}_n$ denote the common proportional change in domestic sectoral unit costs. Homogeneity of the expenditure system implies

$$\widehat{P}_n = \widehat{c}_n R_n.$$

Because every sector has the same value-added share and uses the same intermediate composite, its unit-cost change is

$$\widehat{c}_n = \widehat{w}_n^{\beta_n} \widehat{P}_n^{1-\beta_n}.$$

Combining these equations gives

$$\left(\frac{\widehat{P}_n}{\widehat{w}_n} \right)^{\beta_n} = R_n.$$

Fixed factor endowments imply that nominal income changes in proportion to the wage. Hence,

$$\frac{W_n^T}{W_n^A} = \frac{\widehat{P}_n}{\widehat{w}_n} = R_n^{1/\beta_n},$$

which yields equation (23). When $\beta_n = 1$, the formula reduces to the no-intermediate-input expression in equation (13).

G GTAP Quantification

This appendix repeats the OECD ICIO quantification but now using GTAP. We use the 2014 trade equilibrium, 27 manufacturing sectors and 17 non-manufacturing goods sectors, combine raw milk with dairy products, and treat the remaining 20 non-goods sectors as a single non-traded composite. We construct the manufacturing home-foreign elasticity schedules from the constrained estimates in Table 30 and the non-manufacturing home-foreign elasticity schedules from the corresponding constrained all-goods specification. We calculate each country's value-added share from GTAP primary-factor and intermediate-input costs; see Section 5.1. As in the baseline exercise, the subsequent results use Specification (B), $\sigma_L = 2$, and $\sigma_U = 0.6$ unless otherwise stated.

Table 32 reports the distribution across the 120 GTAP countries. Under Specification (B), the median and mean gains from trade are 30.61 and 41.61 log points. The results remain similar across Specifications (A)-(C).

The larger GTAP estimates are not driven solely by its broader country sample. Among the 75 countries common to the two quantifications, the median gain from trade is 28.83 log points using GTAP and 10.43 log points using the OECD ICIO. Three differences push the results in this direction. First, trade exposure is larger in GTAP in 2014 than in the OECD ICIO in 2007: among these countries, the median share of total expenditure devoted to goods imports is 16.9 percent in GTAP and 12.5 percent in the OECD ICIO, while the

Table 32: GTAP Gains from Trade Across Elasticity-Schedule Specifications

	p10	p25	Median	Mean	p75	p90
Specification (A)	12.20	20.44	30.77	41.65	51.79	78.75
Specification (B)	12.16	20.42	30.61	41.61	51.74	78.77
Specification (C)	11.42	20.08	30.63	42.09	52.20	80.55

Notes: Entries report $100 \ln(W_n^T / W_n^A)$. The percentiles and mean are calculated across the 120 countries in the quantification sample. The calculation includes the country-specific intermediate-input multiplier $1/\beta_n$, where β_n is the 2014 economy-wide value-added share. All three specifications set $\sigma_L = 2$ and $\sigma_U = 0.6$.

median share devoted to goods-producing sectors is 41.7 and 37.1 percent. Second, the GTAP estimates imply lower home-foreign elasticities. Under Specification (B), the average fitted elasticities at the sector-specific median domestic expenditure shares are 3.48 and 2.67 in GTAP for consumer and producer manufacturing, compared with 4.45 and 5.41 in the OECD ICIO; the corresponding non-manufacturing-goods values are 3.87 and 5.91. Third, the median value-added share among these countries is 0.41 in GTAP and 0.48 in the OECD ICIO, so the intermediate-input multiplier is larger in GTAP. Greater trade exposure places the initial equilibrium farther from autarky, while lower elasticities and value-added shares increase the gains from trade.

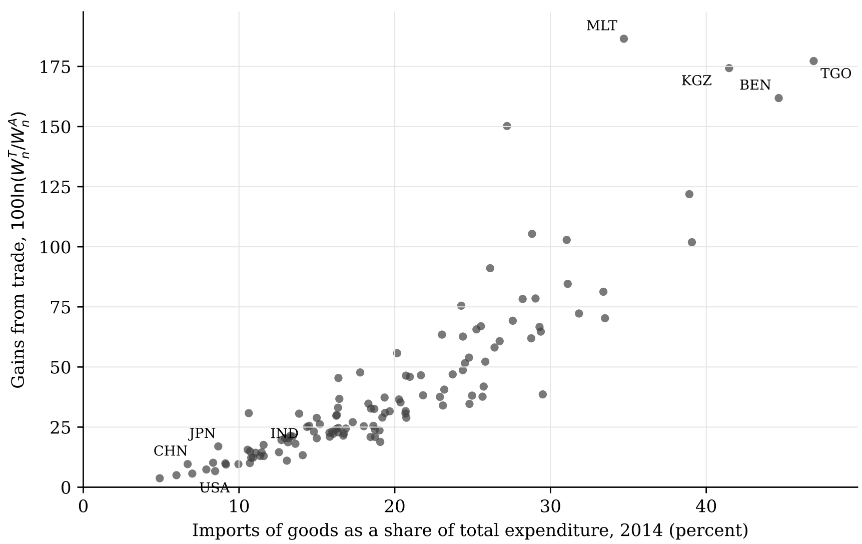

Figure 6: GTAP Gains from Trade and Initial Import Shares

Notes: Each observation is one of the 120 countries in the GTAP quantification sample. The vertical axis reports gains from trade, $100 \ln(W_n^T / W_n^A)$, under constrained Specification (B), with $\sigma_L = 2$, $\sigma_U = 0.6$, and the country-specific intermediate-input multiplier $1/\beta_n$. The horizontal axis reports 2014 expenditure on imports in the 44 goods-producing sectors divided by total expenditure, including the non-traded composite. Selected countries are labeled.

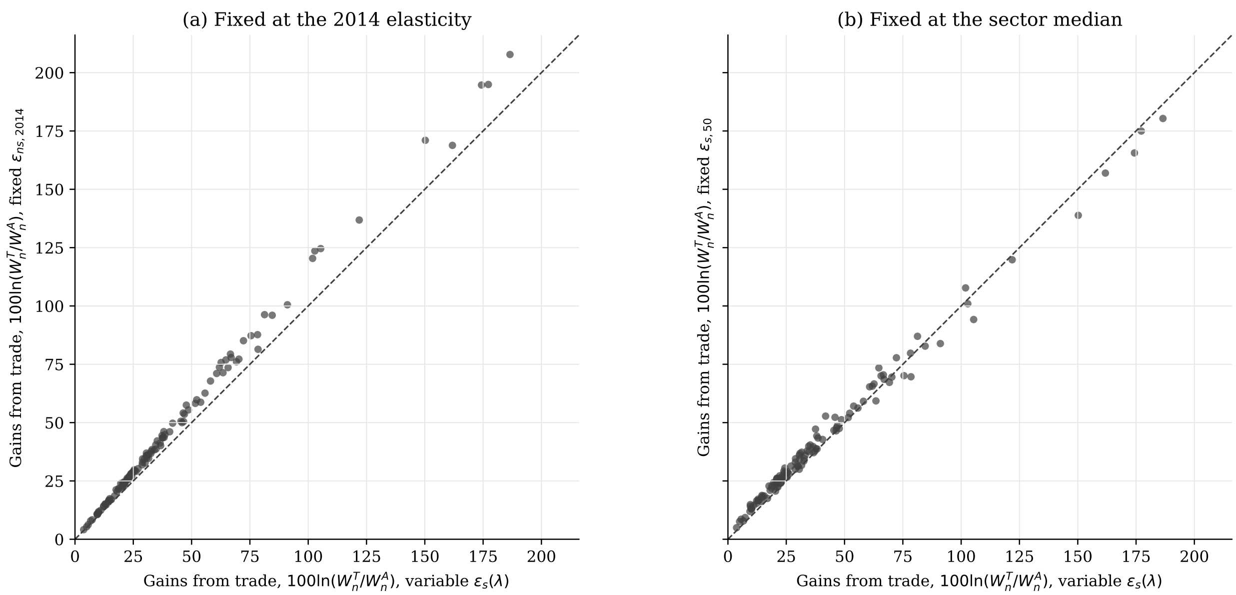

Figure 7: GTAP Gains from Trade with Variable and Fixed Home-Foreign Elasticities

Notes: Each observation is one of the 120 countries in the GTAP quantification sample. The horizontal axis reports $100\ln(W_n^T/W_n^A)$ under the variable-elasticity schedules from constrained Specification (B), with $\sigma_L = 2$, $\sigma_U = 0.6$, and the country-specific intermediate-input multiplier $1/\beta_n$. In the left panel, the vertical axis holds each country-sector elasticity fixed at its fitted 2014 value. In the right panel, it holds each sector's elasticity fixed at the value implied by that sector's median 2014 domestic expenditure share and applies that value to every country. Dashed lines are 45-degree lines.

Figures 6 and 7 repeat the two main figures. Holding each country-sector elasticity fixed at its 2014 value raises the median gain from 30.61 to 34.79 log points and raises gains for every country. Holding each sector's elasticity fixed at the value implied by its median domestic expenditure share produces a median of 33.25 log points.

Table 33: GTAP Gains from Trade Sensitivity to Cross-Sector Substitution

σ_L	σ_U	p10	p25	Median	Mean	p75	p90
2	0.6	12.16	20.42	30.61	41.61	51.74	78.77
1	0.6	13.09	29.59	57.66	909.21	112.38	190.48
3	0.6	10.79	16.66	23.54	32.05	38.70	57.35
2	1	11.81	19.49	29.36	38.63	48.84	72.43

Notes: Entries report $100\ln(W_n^T/W_n^A)$ under Specification (B). The percentiles and mean are calculated across the 120 countries in the GTAP quantification sample. The calculation includes the country-specific intermediate-input multiplier $1/\beta_n$, where β_n is the 2014 economy-wide value-added share. The baseline sets $\sigma_L = 2$ and $\sigma_U = 0.6$; each subsequent row changes the elasticities to the values shown.

Table 33 varies the cross-sector elasticities. Lowering σ_L to one raises the median to 57.66 log points, while raising it to three lowers the median to 23.54 log points. Setting $\sigma_U = 1$ lowers the median to 29.36 log points.

H Details of Monte Carlo Exercises

This appendix uses model-generated data to answer two questions. First, does our estimation procedure mechanically generate positive elasticity-schedule slopes when the true slopes are zero? Second, does our estimation procedure accurately recover positive slopes when the data are generated by the linear home-foreign elasticity schedules that we estimate? We answer these questions using the quantitative multi-sector version of the extended Armington model developed in Appendix C.

We simulate an economy with 40 countries and 27 sectors, the same number of sectors as in the GTAP manufacturing sample. Preferences are Cobb-Douglas across sectors. Within each sector, demand follows the CDE-CES system

$$\omega_{nns}\left(\frac{e_{ns}}{p_{nns}}\right)^{\theta_{s0}} + \left[\sum_{i\neq n}\omega_{nis}\left(\frac{e_{ns}}{p_{nis}}\right)^{\rho_s}\right]^{(\theta_{s0}+\theta_{s1})/\rho_s} = 1.$$

As shown in Appendix C, the elasticity across foreign sources is ρ_s, while the home-foreign elasticity is

$$\varepsilon_s(\lambda_{nns}) = \theta_{s0} + \theta_{s1}\lambda_{nns}.$$

Thus, the model allows the home-foreign elasticity to differ from the elasticity across foreign sources and generates the linear schedule imposed in the empirical analysis.

We hold labor endowments, destination-specific Cobb-Douglas sectoral expenditure shares, iceberg trade costs, initial tariffs, and sector-specific across-foreign-source elasticities fixed throughout the exercise. We draw each of these objects once. Initial tariffs range from zero to 20 percent. To parameterize the across-foreign-source elasticities, we draw $\log\widetilde{\rho}_s \sim \mathcal{N}(\log 2, 0.35^2)$ independently across sectors and set $\rho_s = \min\{4.5, \max\{0.9, \widetilde{\rho}_s\}\}$. We impose $\theta_{s0} = \rho_s$ throughout.

We then consider four calibrations of the elasticity-schedule slopes. Each calibration contains 27 sectors, matching the GTAP analysis.

1. **Calibration 1** sets $\theta_{s1} = 0$ in every sector. Because $\theta_{s0} = \rho_s$, demand is CES across all origins within each sector. This is the traditional multi-sector CES Armington environment and generates the same expenditure-share system as a multi-sector Eaton-Kortum model.

2. **Calibration 2** sets $\theta_{s1} = 4$ in every sector, so the pooling imposed by Specification (A) is correct.

3. **Calibration 3** assigns slopes of 2.5, 3.0, 3.5, 4.0, 4.5, 5.0, and 5.5 to the seven manufacturing groups used in Specification (C), with the same sectors assigned to each group as in the GTAP analysis.

4. **Calibration 4** draws a positive slope separately for each sector from a uniform distribution on $[2, 6]$; the realized slopes range from 2.29 to 5.95.

We first solve Calibration 1 for its initial equilibrium wages and bilateral expenditure shares. We then choose the CDE preference weights in Calibrations 2-4 so that, at the same wages and initial tariffs, they reproduce the Calibration 1 bilateral shares. The four calibrations therefore have the same initial equilibrium and domestic-share distribution. They differ in the home-foreign elasticity schedules that govern their responses to the tariff shocks.

We generate 200 tariff-shock matrices. For each matrix, we select every foreign destination-origin-sector cell independently with probability 0.45. In each selected cell, we reduce the log gross tariff by a fraction equal to 0.25 times an independent Beta$(2, 2)$ draw; tariffs in the remaining cells are unchanged. We apply the same 200 shock matrices to all four calibrations and solve for the new equilibrium wages and expenditure shares after each shock. Within each calibration, the model parameters and initial equilibrium remain fixed across draws. The only variation across draws comes from which tariffs are reduced and by how much.

For each calibration and each of the 200 equilibrium changes, we construct the variables used in the empirical analysis: the dependent variable, initial domestic absorption share, tariff shifter, and country-exposure terms. Because the initial equilibrium is common across calibrations and fixed across tariff shocks, the initial domestic absorption shares and country-exposure terms are the same in every model-generated data set.

We apply the paper's unrestricted NLLS estimator to each data set (there are 200 datasets given the 200 tariff-shock matrices), jointly estimating the country cost shifters and imposing no positivity constraints. For every calibration, we estimate Specifications (A), (B), and (C) using the same code as in the GTAP analysis. The tables report the mean and the 5th and 95th percentiles of each slope estimate across the 200 tariff-shock matrices. When every sector represented in a row has the same true slope, the first column, with the "true sector slope" reports that common value. When a row pools sectors with different true slopes, the column reports their range.

Calibration 1: Constant Home-Foreign Elasticities. Table 34 reports the results when every true slope is zero. The estimates are centered close to zero throughout Specifications (A), (B), and (C). The mean Specification-(A) slope is 0.001, with a 5th-95th percentile range of $[-0.004, 0.006]$. Across the seven Specification-(C) groups, mean slopes range from -0.001 to 0.003, and every percentile range contains zero.

Table 34: Armington Monte Carlo with Constant Home-Foreign Elasticities

Spec.	Aggregate	True sector slope(s)	Mean $\widehat{\theta}_1$	5th pct.	95th pct.
(A)	All manufacturing	0.00	0.001	-0.004	0.006
(B)	Consumer manuf.	0.00	0.002	-0.004	0.007
	Producer manuf.	0.00	0.001	-0.008	0.009
(C)	Food, bev., tob.	0.00	0.003	-0.003	0.008
	Textiles/apparel/leather	0.00	-0.001	-0.022	0.016
	Wood/paper/printing	0.00	0.002	-0.007	0.011
	Petrol./chem./pharm./rubber	0.00	-0.001	-0.018	0.015
	Minerals/metals	0.00	0.000	-0.012	0.013
	Machinery/electrical	0.00	0.001	-0.010	0.011
	Transport/other manuf.	0.00	0.002	-0.007	0.011

Notes: Every sector has $\theta_{s1} = 0$ and $\rho_s = \theta_{s0}$. The table reports results from Specifications (A), (B), and (C) across 200 independently drawn tariff shocks, holding the realized 40-country, 27-sector Armington economy fixed. The mean, 5th percentile, and 95th percentile refer to the distribution of the estimated pooled slope across tariff shocks. The true-slope column reports the common true slope among sectors in the row when one exists and their range otherwise; a row spanning different true slopes has no single structural pooled coefficient. Every regression uses the same unrestricted NLLS estimation routine as the real-data tables, including the jointly estimated country cost shifters.

Calibration 2: A Common Positive Slope. Table 35 reports the results when the true slope is 4 in every sector. Specification (A), which imposes the correct pooling structure, has a mean estimate of 4.040 and a 5th-95th percentile range of $[4.030, 4.050]$. Specifications (B) and (C) estimate additional slopes, each of which also closely recovers the common true value.

Table 35: Armington Monte Carlo with a Common Positive Slope

Spec.	Aggregate	True sector slope(s)	Mean $\widehat{\theta}_1$	5th pct.	95th pct.
(A)	All manufacturing	4.00	4.040	4.030	4.050
(B)	Consumer manuf.	4.00	4.035	4.024	4.045
	Producer manuf.	4.00	4.046	4.030	4.063
(C)	Food, bev., tob.	4.00	4.033	4.021	4.045
	Textiles/apparel/leather	4.00	4.043	4.006	4.079
	Wood/paper/printing	4.00	4.030	4.013	4.051
	Petrol./chem./pharm./rubber	4.00	4.049	4.015	4.080
	Minerals/metals	4.00	4.061	4.034	4.089
	Machinery/electrical	4.00	4.026	4.003	4.048
	Transport/other manuf.	4.00	4.045	4.028	4.064

Notes: Every sector has $\theta_{s1} = 4$, so Specification (A) matches the data-generating pooling structure. The table reports results from Specifications (A), (B), and (C) across 200 independently drawn tariff shocks, holding the realized 40-country, 27-sector Armington economy fixed. The mean, 5th percentile, and 95th percentile refer to the distribution of the estimated pooled slope across tariff shocks. The true-slope column reports the common true slope among sectors in the row when one exists and their range otherwise; a row spanning different true slopes has no single structural pooled coefficient. Every regression uses the same unrestricted NLLS estimation routine as the real-data tables, including the jointly estimated country cost shifters.

Calibration 3: Seven Positive Group Slopes. Table 36 reports results when the slopes are common across the sectors within each of the seven groups in Specification (C). When estimating Specification (C), estimates are 2.523, 3.031, 3.528, 4.048, 4.568, 5.029, and 5.559, compared with true values of 2.5, 3.0, 3.5, 4.0, 4.5, 5.0, and 5.5. Although Specifications (A) and (B) pool sectors with different true slopes, their estimates remain positive and summarize the slopes within the corresponding broader groups.

Table 36: Armington Monte Carlo with Seven Positive Group Slopes

Spec.	Aggregate	True sector slope(s)	Mean $\widehat{\theta}_1$	5th pct.	95th pct.
(A)	All manufacturing	$[2.50, 5.50]$	3.745	3.673	3.815
(B)	Consumer manuf.	$[2.50, 5.50]$	3.256	3.173	3.360
	Producer manuf.	$[3.50, 5.50]$	4.380	4.322	4.438
(C)	Food, bev., tob.	2.50	2.523	2.514	2.533
	Textiles/apparel/leather	3.00	3.031	2.998	3.064
	Wood/paper/printing	3.50	3.528	3.512	3.546
	Petrol./chem./pharm./rubber	4.00	4.048	4.017	4.080
	Minerals/metals	4.50	4.568	4.540	4.597
	Machinery/electrical	5.00	5.029	5.003	5.054
	Transport/other manuf.	5.50	5.559	5.538	5.581

Notes: Slopes equal 2.5, 3.0, 3.5, 4.0, 4.5, 5.0, and 5.5 across the seven manufacturing aggregates, so Specification (C) matches the data-generating pooling structure. The table reports results from Specifications (A), (B), and (C) across 200 independently drawn tariff shocks, holding the realized 40-country, 27-sector Armington economy fixed. The mean, 5th percentile, and 95th percentile refer to the distribution of the estimated pooled slope across tariff shocks. The true-slope column reports the common true slope among sectors in the row when one exists and their range otherwise; a row spanning different true slopes has no single structural pooled coefficient. Every regression uses the same unrestricted NLLS estimation routine as the real-data tables, including the jointly estimated country cost shifters.

Calibration 4: Sector-Specific Positive Slopes. Table 37 reports the results when all 27 sector slopes are positive and distinct. Every slope estimate in Specifications (A), (B), and (C) is positive. The mean Specification-(A) estimate is 4.139, the two Specification-(B) means are 4.107 and 4.178, and the seven Specification-(C) means range from 3.373 to 4.828.

Table 37: Armington Monte Carlo with Sector-Specific Positive Slopes

Spec.	Aggregate	True sector slope(s)	Mean $\widehat{\theta}_1$	5th pct.	95th pct.
(A)	All manufacturing	$[2.29, 5.95]$	4.139	4.063	4.215
(B)	Consumer manuf.	$[2.46, 5.93]$	4.107	4.009	4.210
	Producer manuf.	$[2.29, 5.95]$	4.178	4.068	4.313
(C)	Food, bev., tob.	$[2.65, 5.93]$	4.702	4.624	4.778
	Textiles/apparel/leather	$[3.38, 4.38]$	3.531	3.481	3.573
	Wood/paper/printing	$[3.26, 5.95]$	4.828	4.601	5.109
	Petrol./chem./pharm./rubber	$[2.29, 5.76]$	4.071	3.787	4.384
	Minerals/metals	$[2.58, 4.88]$	3.471	3.342	3.604
	Machinery/electrical	$[3.49, 5.46]$	4.208	4.120	4.292
	Transport/other manuf.	$[2.46, 4.61]$	3.373	3.240	3.505

Notes: Every sector has a distinct positive slope. The table reports results from Specifications (A), (B), and (C) across 200 independently drawn tariff shocks, holding the realized 40-country, 27-sector Armington economy fixed. The mean, 5th percentile, and 95th percentile refer to the distribution of the estimated pooled slope across tariff shocks. The true-slope column reports the common true slope among sectors in the row when one exists and their range otherwise; a row spanning different true slopes has no single structural pooled coefficient. Every regression uses the same unrestricted NLLS estimation routine as the real-data tables, including the jointly estimated country cost shifters.

In summary, the Monte Carlo exercises show two things. First, our estimation procedure does not mechanically produce the pattern found in the data when data is generated by canonical constant-elasticity trade models. Second, our estimation procedure successfully recovers positive elasticity-schedule slopes when the data-generating model is one with positive linear schedules.