ETLA Working Papers | 141 15.6.2026

Beggar-Thy-Neighbor by Other Means

Jaakko Markkanen
ETLA Economic Research, Finland
jaakko.markkanen@etla.fi

Markku Siikanen
VATT Institute for Economic Research, Finland
markku.siikanen@vatt.fi

Nelli Valmari
Aalto University and Helsinki GSE, Finland
nelli.valmari@aalto.fi

Suggested citation:
Markkanen, Jaakko, Siikanen, Markku & Valmari, Nelli (15.6.2026). "Beggar-Thy-Neighbor by Other Means". ETLA Working Papers No 141. https://pub.etla.fi/ETLA-Working-Papers-141.pdf

Abstract

A country that cannot devalue its currency can still cut its exporters' costs through industrial policy and steal business from foreign rivals. We call this beggar-thy-neighbor by other means and measure it for Finland's 2017–2019 internal devaluation policy. We estimate the export demand for nine large manufacturing industries, together accounting for roughly 4 percent of Finnish GDP, using a BLP-style demand model and a sufficient-statistic identity for cost incidence. We document super-pass-through to export prices, averaging about 1.18, above the CES gravity ceiling. The realized policy cut labor costs by 3.6 percent and raised Finnish export revenue by €239.0 million over 2017–2020, 0.6 percent of baseline. A more ambitious original government proposal with 5 percent cost decrease would have shifted €567.8 million in revenue away from rival exporters in the same destinations. A hypothetical four-day work week would have cost €2.4 billion with wage costs rising 28 percent. Internal devaluation captures export-market share from foreign rivals. The cross-border revenue transfer is comparable in magnitude to the domestic gains.

Tiivistelmä

”Oma suu lähinnä” -politiikka toisin keinoin – Teollisuuspolitiikan vaikutukset vientiin ja kilpailijoiden markkinaosuuksiin

Maat, jotka eivät voi devalvoida valuuttaansa, voivat alentaa vientiyritystensä kustannuksia teollisuuspolitiikan keinoin ja viedä markkinaosuuksia ulkomaisilta kilpailijoiltaan. Kutsumme tällaista politiikkaa toisin keinoin harjoitetuksi ”oma suu lähinnä” -politiikaksi ja mittaamme sen vaikutuksen Suomen vuosien 2017–2019 kilpailukykysopimuksen eli sisäisen devalvaation osalta. Tarkastelemme yhdeksää suurta teollisuusmarkkinaa, jotka tuottavat yhdessä noin neljä prosenttia Suomen bruttokansantuotteesta, ja mallinnamme niiden vientikysynnän BLP-tyyppisellä kysyntämallilla. Ratkaisemme työvoiman kustannusten kohtaannon uudella tyhjentävään tunnuslukuun perustuvalla menetelmällä. Tulostemme perusteella muutokset työvoimakustannuksissa välittyvät vientihintoihin yli täysimääräisesti, keskimäärin noin kertoimella 1,18, mikä on enemmän kuin CES-painovoimamallin teoreettinen yläraja mahdollistaisi. Toteutunut kilpailukykysopimus alensi työvoimakustannuksia 3,6 prosenttia ja kasvatti Suomen vientituloja noin 239 miljoonalla eurolla vuosina 2017–2020, mikä vastaa noin 0,6 prosentin lisäystä. Hallituksen alkuperäinen, kunnianhimoisempi esitys noin viiden prosentin kustannusalennuksesta olisi siirtänyt 567,8 miljoonaa euroa vientituloja pois kilpailijoilta samoilla markkinoilla. Hypoteettinen nelipäiväinen työviikko nostaisi palkkakustannuksia 28 prosenttia ja vähentäisi vientituloja noin 2,4 miljardia euroa. Sisäinen devalvaatio on tehokas keino vallata markkinaosuuksia ulkomaisilta kilpailijoilta, ja rajojen yli siirtyvä vientitulo on samaa suuruusluokkaa kuin kotimaahan jäävä hyöty.

D.Sc. (Econ.) **Jaakko Markkanen** is a Senior Researcher at ETLA Economic Research.
D.Sc. (Econ.) **Markku Siikanen** is a Senior Researcher at VATT Institute for Economic Research.
D.Sc. (Econ.) **Nelli Valmari** is a Postdoctoral Researcher at Aalto University and Helsinki GSE and ETLA Senior Fellow.

KTT **Jaakko Markkanen** on Elinkeinoelämän tutkimuslaitoksen vanhempi tutkija.
KTT **Markku Siikanen** on Valtion taloudellisen tutkimuskeskuksen erikoistutkija.
KTT **Nelli Valmari** toimii tutkijatohtorina Aalto-yliopistossa sekä Helsinki GSE:ssä ja on Etlan Senior Fellow.

Acknowledgements: We thank Niku Määttänen and Salla Kalin for comments and Petteri Larjos for helping with our computational needs. We gratefully acknowledge financial support from OP Research Foundation. Jaakko Markkanen also gratefully acknowledges support from TT-Säätiö and the Alfred Kordelin Foundation. All errors are ours.

Artificial intelligence has been used to support human work in the preparation of this working paper in accordance with Etla's ethical guidelines (version 14.1.2026, see https://www.etla.fi/en/ai-ethics).

Kiitokset: Kiitämme Niku Määttästä ja Salla Kalinia kommenteista sekä Petteri Larjosta avusta laskentaresurssien kanssa. Olemme kiitollisia OP Ryhmän Tutkimussäätiölle hankkeen rahoituksesta. Jaakko Markkanen kiittää lisäksi TT-Säätiötä ja Alfred Kordelinin säätiötä. Kaikki virheet ovat omiamme.

Tämän työpaperin työstämisessä on hyödynnetty tekoälyä ihmistyön tukena Etlan eettisen ohjeiston mukaisesti (versio 14.1.2026, ks. https://www.etla.fi/ai-etiikka).

Keywords: International trade, Markups, Pass-through, Industrial policy, Internal devaluation

Asiasanat: Kansainvälinen kauppa, Markkinavoima, Teollisuuspolitiikka, Sisäinen devalvaatio

JEL: F12, F14, L11, L13, L52

1 Introduction

Industrial policy is back on the agenda of advanced economies, and with it an old concern: in the *Wealth of Nations*, Adam Smith observed that nations had been taught that their interest "consisted in beggaring all their neighbours" (Smith 1776). That concern is usually associated with tariffs and competitive currency devaluations. The policies we study can reach the same end by other means. A government can lower its own producers' costs through a business subsidy, a corporate-tax cut, or a labor-cost reform. Such cost-side policies can move the global market equilibrium: lower domestic costs let home firms steal business from rivals abroad, redistributing revenue from foreign producers to domestic exporters. Yet foreign buyers pay the lower export prices that take those sales, so the policy beggars thy neighbor's producers and enriches thy neighbor's consumers.

The size of this redistribution depends on the structure of competition. When firms have market power and demand curvature is high enough, pass-through can exceed one, so a cost cut lowers a firm's price by more than the saving itself and takes business from rivals. We study nine large manufacturing markets in forest products and stainless steel, which together turn over about €134 billion a year. Because these markets are global, even a cost policy adopted in a small economy can move hundreds of millions of euros across borders and redistribute welfare between producers and consumers.

We measure this cross-border redistribution for Finland's 2017–2019 internal devaluation policy, which cut employer social-insurance contributions and extended working time. We evaluate three scenarios in these nine markets: the realized policy, a more ambitious original government proposal, and a four-day-week benchmark that moves labor costs in the opposite direction. Relative to a no-policy counterfactual, the realized policy increased Finnish export revenue by €239.0 million over 2017–2020, or 0.6 percent of baseline. Under the more ambitious proposal, rival-exporter revenue in the same destinations would have fallen by €567.8 million. We interpret internal devaluation as a form of industrial policy (Juhász, Lane, and Rodrik 2024; Van Biesebroeck and Zhang 2025) with a measurable cross-border revenue-redistribution or business-stealing component.

Our application is an unusually clean setting for this measurement. Finland's membership in the euro area rules out an offsetting nominal devaluation, so the wage cut is a genuine internal devaluation and a cost-side substitute for the exchange-rate tool (Farhi, Gopinath, and Itskhoki 2014). Following the global financial crisis and the decline of key export sectors, Finnish policymakers became concerned that high labor costs were limiting the recovery of export-oriented industries. The realized policy and the more ambitious original government proposal share an implementation channel, making the proposal a direct benchmark. Forest and steel are exactly the sectors that internal devaluation policies aim to protect, with bilateral customs flows that make the effects observable.

We make three contributions to the industrial organization and international trade literatures. First, we quantify the equilibrium effects of three policy reforms on domestic exporters and their foreign rivals. Reversing the realized policy reduces Finnish firms' market shares by 0.19 to 1.80 percent across the nine four-digit Harmonized System (HS) markets relative to the no-policy counterfactual. The more ambitious original government proposal would have increased market shares by 0.40 to 3.52 percent. The four-day work week, modeled as one fewer workday per week at unchanged pay, implies a much larger market-share decline of 2.23 to 19.23 percent. The cross-scenario pattern shows that labor-cost interventions affect export competitiveness in proportion to the underlying cost shock. The mirror image falls on foreign rivals. Under the more ambitious proposal, competing exporters in the same destinations lose €567.8 million in revenue. The loss is a cross-border business-stealing policy externality (Mankiw and Whinston 1986), in which one government's cost policy redistributes sales toward its own firms and away from competitors abroad. Most closely related at the firm level, Benzarti and Harju (2021) use Finnish payroll-tax variation to study firm input substitution. We measure export-side incidence at the firm-destination-product level rather than through aggregate or domestic-employment outcomes.

Second, we develop and apply an empirical framework that combines Berry-Levinsohn-Pakes (BLP)-style discrete-choice demand estimation (Berry, Levinsohn, and Pakes 1995) on bilateral customs data with a sufficient-statistics approach to cost incidence. Identification comes from instruments that generate variation in market shares and firms' marginal costs. On the demand side, we use traditional BLP instruments and differentiation instruments. On the cost side, we use tariffs, transport margins, and exchange rates. The demand side accommodates multiple destination markets and endogenous export prices, separates the roles of price and distance, and allows heterogeneous price sensitivity through a lognormal random coefficient. Unlike the standard constant elasticity of substitution (CES)-plus-iceberg gravity workhorse, this specification does not impose unit pass-through.

On the supply side, we invert the standard Hall (1988) and De Loecker and Warzynski (2012) identity. The standard approach recovers the markup from an output elasticity estimated on a production function. We instead take the markup from the demand inversion and combine it with the observed accounting input cost revenue share. This combination identifies the input's cost-share object that governs incidence under Shephard's lemma. The approach requires only the markup and the accounting input cost revenue share at the firm-year level. Because it never estimates an output elasticity, it sidesteps the production-function non-identification concerns raised by Bond, Hashemi, Kaplan, and Zoch (2021) and Doraszelski and Jaumandreu (2021). Because the only firm-year observables are a markup and an accounting input cost revenue share, the framework extends to any single-input price shock: labor costs, tariffs, trade costs, materials, or energy. Moreover, the framework does not require data on other input costs.

Third, we provide new evidence on markup dynamics in nine manufacturing markets (forest products and stainless steel) over 1999–2021. Markups follow a common U-shaped pattern across all nine markets: an early-2000s peak, a post-financial-crisis trough, and a partial recovery toward initial levels. In three markets—sawn wood, tissue, and other paper grades—markups decline by two to eight percentage points across the sample, and in the remaining six they are flat within roughly one percentage point net. Our demand-based estimates therefore do not support a systematic increase in markups across the studied markets, contrasting with the broad rise documented for the United States by De Loecker, Eeckhout, and Unger (2020). The recent exchange between Benkard, Miller, and Yurukoglu (2025a), De Loecker, Eeckhout, and Unger (2025), and Benkard, Miller, and Yurukoglu (2025b) highlights how sensitive production-function-based markup estimates are to assumptions about input shares and technology. Because our markups come from demand-side substitution rather than from input cost shares, they are not subject to the same identification concerns.

The remainder of the article is organized as follows. Section 2 describes the relevant Finnish labor-market institutions and regulations. Section 3 presents the data. Section 4 introduces the demand model and discusses identification. Section 5 presents the supply model and the sufficient-statistics approach for the cost incidence. Section 6 describes the counterfactual simulation framework and the counterfactual policy scenarios. Section 7 presents the results. Section 8 concludes.

2 Institutional Background

2.1 Labor Market Regulation

Statutory labor law and collective bargaining agreements jointly govern Finnish labor markets. National legislation provides the general legal framework for employment relationships, whereas collective agreements negotiated between trade unions and employer associations determine many of the concrete terms and conditions of employment. In practice, sectoral collective agreements regulate wages, working hours, holiday entitlements, pay supplements, sick pay arrangements, and other employment conditions.

A distinctive feature of the Finnish labor market is the broad coverage of collective agreements and the role of generally binding agreements (Paukkeri, Ravaska, and Riihelä 2024). Collective agreements directly bind the employers and employees represented by the organizations that negotiated them. In addition, a large share of collective agreements are confirmed as generally binding within a sector. In such cases, even employers that are not members of the employer association party to the agreement must comply with the minimum terms of the agreement. This general-binding mechanism gives collective bargaining a regulatory function beyond the organized parties. It extends negotiated labor standards to non-unionized workers and non-organized employers in the sector. In 2017, collective agreements covered all public-sector workers, including state and municipal employees, and 83.1 percent of private-sector workers (Ahtiainen 2019). At the same time, around 60 to 65 percent of the workforce was unionized which is high by OECD standards (Böckerman, Kalin, Juuti, Kauhanen, and Suhonen 2025).

The importance of collective bargaining is reinforced by the absence of a statutory minimum wage in Finland. Sectoral collective agreements, rather than legislation, set minimum wages. Employers may not ignore the wage provisions of an applicable collective agreement. As a result, wage floors vary across sectors and occupations instead of being fixed by a single national legal minimum wage. This institutional arrangement makes collective bargaining the central mechanism for minimum-wage regulation in Finland.

The Finnish wage-bargaining system has nevertheless evolved considerably over the past two decades. Until 2006, wages were negotiated primarily through centralized tripartite agreements between employer confederations, employee confederations, and the state (Kiander, Sauramo, and Tanninen 2011).[1] Since then, wage bargaining has become more decentralized. Coordination has nonetheless remained important: settlements in export-oriented sectors often serve as benchmarks for wage agreements in other industries and sectors (Jonker-Hoffrén 2019). The forest industry provides a particularly clear example of this decentralization process. In 2020, it moved away from industry-level collective agreements and shifted toward firm-level wage bargaining.

Social insurance costs. Social insurance and health insurance contributions are split between employers and employees. Employers pay compulsory employment-related contributions, whereas employees contribute through deductions from their wages. For example, in 2015, the employer's health insurance contribution was 2.08 percent of wages. Employees paid 2.10 percent in total, split between a 0.78 percent daily allowance and a 1.32 percent medical care component (Confederation of Finnish Industries 2015).

These statutory contributions cover several areas, including earnings-related pension insurance, unemployment insurance, health insurance, accident insurance, and other social security schemes. The employer is usually responsible for withholding the employee's contribution from wages and remitting both the employer and employee contributions to the relevant institutions. As a result, Finnish labor costs consist not only of gross wages but also of compulsory employer contributions, whereas employees' take-home pay is reduced by their own statutory insurance contributions. Changes in contribution rates therefore affect both firms' labor costs and workers' disposable income. Our counterfactual analyses use social insurance costs to study how shocks to wage cost pass through to export prices and what happens to export revenue.

Working-time regulation. The Working Hours Act governs statutory working time, setting regular hours at a maximum of eight per day and 40 per week. As with wages, sectoral collective agreements in practice set the operative limits. These agreements fix regular annual working time—on the order of 1,720 hours—and regulate overtime, shift and Sunday work, and the corresponding pay premiums. Because working time is itself a core term of these agreements, the employee and employer confederations can lengthen or shorten it through the same bargaining channel that

1. The internal devaluation policy discussed in Subsection 2.2 was also negotiated through tripartite negotiations in 2016, but these negotiations were not part of the regular wage-bargaining process.

sets wages. Through this same channel, the internal devaluation policy, discussed in Section 2.2, extended annual working time by 24 hours at unchanged annual pay.

2.2 Internal Devaluation Policy

Juha Sipilä's government took office on 29 May 2015 with cost competitiveness as one of the flagship objectives of its programme (Finnish Government 2015). In July 2015, the government set a target of reducing unit labor costs by 5 percent. The initial aim was to achieve this through a tripartite "Society Agreement" negotiated between the central trade union and employer confederations. However, the negotiations collapsed in autumn 2015, after which the government introduced the "forced laws" (or *pakkolait* in Finnish), a package of draft bills unilaterally reducing sick-leave pay, Sunday and overtime premiums, public-sector annual leave, and holiday bonuses to achieve the unit labor cost reduction objective (Ministry of Finance in Finland 2015).

Faced with the prospect of unilateral legislation, the labor market partners returned to the negotiating table. They reached a preliminary tripartite agreement on 29 February 2016 and signed the Competitiveness Pact, abbreviated (*kiky*) on 14 June 2016 (EK, SAK, STTK and Akava 2016). The realized pact reduced unit labor costs by 3–4 percent, falling short of the government's original 5 percent target. It also narrowed the holiday-bonus cut to public-sector workers. Upon the signing of the pact, the government formally withdrew the unilateral legislation package.

The policy responded to the prevailing narrative about Finland's economic troubles. In that narrative, unit labor costs had drifted above the level productivity could support and eroded the export sector's cost competitiveness. With the euro ruling out a nominal devaluation, the adjustment had to come through labor costs directly. Policymakers saw internal devaluation as a tool to shrink the output gap by increasing export revenues and employment.

Realized policy and proposed policy For tradable-sector employers, the Competitiveness Pact had two cost-side components during 2017–2019. First, it extended standard annual working time under sectoral collective agreements by 24 hours without increasing annual compensation. Relative to the pre-pact baseline, this extension corresponds to an effective hourly-wage reduction of about 1.4 percent ($24/1720$). Second, the pact reduced employer-side social insurance contribution rates by approximately 1.6, 2.2, and 2.2 percentage points in 2017, 2018, and 2019, respectively.

Offsetting increases in employee-side contributions absorbed part of this reduction, implying that the policy lowered employer labor costs partly by shifting social insurance costs to workers. Although the industrial-policy literature recognizes tax and social-insurance relief as common policy instruments (Juhász, Lane, and Rodrik 2024), the distinctive feature of the Finnish policy was that it reduced firms' unit labor costs through a negotiated reallocation of labor costs within the wage-setting system. The pact then partially unwound in 2020: the 24-hour working-time extension lapsed, and the employer-side contribution reduction fell to roughly 0.6 percentage points. We translate these year-specific policy changes into the labor-price shock used in the sufficient-statistic identity in Section 6.1.

The realized pact was less ambitious than the cost-competitiveness package the Sipilä government proposed, which targeted a roughly 5 percent reduction in unit labor costs (Prime Minister's Office of Finland 2015). We treat that package as a separate counterfactual ("Sipilä" in Section 6) to bracket the policy's potential cost-side incidence. Because our sample is private-sector exporters, we use the cost decreases that fall on private labor cost—the employer-contribution cut, the reductions to Sunday, overtime, and sick-leave premiums, and two public holidays converted to unpaid workdays—and set aside the measures aimed mainly at the public sector, such as the cuts to public-sector annual leave and holiday bonuses. In this counterfactual unit labor costs decrease by 4 percent and annual hours increase by 16 hours.

Previous internal devaluation policies. The kiky pact joins a set of European internal devaluation episodes that pair payroll-tax cuts with wage moderation or working-time changes. Comparators include the Belgian wage-norm mechanism (Geis 2023), the French *Crédit d'impôt pour la compétitivité et l'emploi* of 2013 (Cahuc, Carcillo, and Le Barbanchon 2019), the German Hartz reforms of the mid-2000s (Dustmann, Fitzenberger, Schönberg, and Spitz-Oener 2014), and the Spanish 2012 labor market reform (Bentolila, Cahuc, Dolado, and Le Barbanchon 2012). The Latvian post-crisis adjustment provides the canonical fixed-exchange-rate analog (Blanchard, Griffiths, and Gruss 2013).

Commentary on these episodes has long invoked the language of beggar-thy-neighbor competitive devaluation (Caballero, Farhi, and Gourinchas 2021), conducted through wages once the exchange rate is fixed. Internal devaluation of this kind is cost-side and not inherently protectionist, unlike the tariffs and trade barriers that Fajgelbaum, Goldberg, Kennedy, and Khandelwal (2020) study, in the sense that consumers are better off through lower prices in the output market.

This distinction has a welfare counterpart. A tariff is a trade-volume wedge that destroys surplus, whereas an internal devaluation shifts the price level without restricting volume, and because the super-pass-through we estimate lets export prices fall by more than costs, it compresses markups rather than opening a wedge. Its bite is therefore distributional rather than a deadweight loss: across borders, it falls on rival producers, at home on the workers who fund the contribution cut, and the revenue rivals shed is in part competed away to the consumers who import the now-cheaper goods. However, industrial policy can itself be protectionist, as are many of China's industrial subsidies. Existing work has focused largely on aggregate labor-market and current-account outcomes, but we study how the policy shifts demand for Finnish exports in global markets, where the cross-border revenue redistribution occurs.

3 Data

We use population-wide register data on Finnish firms from Statistics Finland and the Finnish Customs. We complement the register data with open access trade, tariff, transportation, and insurance margin data. We list our data sources in Table 1 and Appendix Section A.2 defines key variables.

Our first primary register is the product-level customs data from Finnish Customs, which records annual Finnish firm-destination import and export values and quantities at the 8-digit Combined Nomenclature (CN) level.[2] We augment these customs records with bilateral trade flows from the BACI database (Gaulier and Zignago 2010) to estimate the demand for Finnish products in global export markets.[3]

We combine the Finnish customs data with global trade data (BACI). These data contain annual HS6-level bilateral trade flows between countries. We replace Finnish aggregate observations in BACI with the firm level observations from the customs data and we also aggregate the Finnish observations to the same HS6-level as the rest of BACI. Finally, we convert Finnish observations from euros to U.S. dollars (USD).

We also combine BACI with country-year level demographic and macroeconomic data from The Global Macro Database by Müller, Xu, Lehbib, and Chen (2025). We use this dataset to control for observed

2. CN is the European Union (EU) 8-digit extension of the international 6-digit HS classification; the first six digits coincide with HS and the final two encode EU-specific subcategories.

3. CEPII: Centre d'Études Prospectives et d'Informations Internationales; BACI: Base pour l'Analyse du Commerce International.

Table 1: Data Sources

Source	Data
Statistics Finland	Firm balance sheet data
Finnish Customs	Firm export and import data
Gaulier and Zignago (2010)	Bilateral trade flows (CEPII-BACI)
Müller, Xu, Lehbib, and Chen (2025)	Global macro indicators
Teti (2024)	Historical tariff data
Fiallos, Liberatore, and Cassimon (2024)	Transportation and insurance margins

Notes: Table lists the data sources used in the empirical analysis.

cross-country heterogeneity in the demand estimation. We augment trade values (and prices) by tariffs, transportation and insurance margins, because importers pay these components on top of Free-On-Board (FOB) prices. Historical tariff data comes from Teti (2024) and transportation and insurance margins are from Fiallos, Liberatore, and Cassimon (2024).

The production-side analysis relies on firm-level accounting data from Statistics Finland. The data cover the universe of Finnish firms and are indexed by business identifier. We use information on annual wage sums, compulsory employer contributions, and gross value of production.

3.1 Demand Estimation Sample Descriptive Statistics

We define product markets using the HS level 4 classification. This classification is a global standardized trade classification and we treat each four-digit HS code as a separate product market.[4] Table A.2 in Appendix A.4 presents the official definition of each HS4-category. The markets studied are forest-related manufacturing industries, with the exception of stainless steel.

Table 2 presents descriptive statistics for the nine HS4 categories included in the demand-estimation sample over the period 1999–2021. For non-Finnish exporters, the unit of observation is a country–HS6 product–destination–year combination. We observe Finnish exports at the firm level, so the corresponding unit of observation is a firm–HS6 product–destination–year combination. The full estimation sample con-

4. Each HS4 can nest multiple more specific HS6 categories. For example, HS4 7219 covers flat-rolled stainless-steel products at least 600 mm wide, whereas HS6 721911 covers the more specific category of hot-rolled stainless-steel coils more than 10 mm thick.

Table 2: Demand Estimation Descriptive Statistics by HS4 category

HS4	Obs.	HS6	Markets	Products	Value
4407	144,715	7	990	1,663	636
4412	99,605	7	1,100	1,156	288
4703	27,286	4	704	348	581
4801	20,996	1	1,100	142	184
4802	126,053	8	1,430	1,135	395
4803	17,789	1	698	143	56
4804	151,525	12	1,364	1,447	254
4805	121,879	11	1,336	1,196	209
7219	163,532	14	814	1,488	476
Total	873,380	65	9,536	8,718	3,078

Notes: This table lists the demand estimation sample descriptive statistics by HS4 categories. Value is reported in billions of USD. All monetary values are deflated to 2010 values.

tains 873,380 observations.

The number of observations varies substantially across product categories. Stainless steel (7219) has the largest sample, with 163,532 observations, followed by paper and paperboard (4804), with 151,525 observations, and sawn wood (4407), with 144,715 observations. Category 4803 has the smallest sample, with 17,789 observations. The categories also differ in their degree of product differentiation within the HS classification. Stainless steel contains 14 distinct HS6 categories, whereas categories 4801 and 4803 each contain only one category.

Within each HS4 category, we define a market as a destination–year pair and a product as a firm- or country–HS6 combination. Category 4802 spans the largest number of markets, with 1,430 destination–year pairs, followed by categories 4804 and 4805, with 1,364 and 1,336 markets, respectively. Category 4803 spans the fewest markets, with 698. Sawn wood (4407) contains the largest number of products, with 1,663, followed by stainless steel (7219), with 1,488, and category 4804, with 1,447. Categories 4801 and 4803 contain the fewest products, with 142 and 143, respectively.

The final column reports total trade value over the study period in billions of constant 2010 U.S. dollars. Sawn wood (4407) is the largest category by trade value, at 636 billion, followed by chemical pulp (4703), at 581 billion, and stainless steel (7219), at 476 billion. Category 4803 is the smallest, with a total trade value of 56 billion.[5]

5. Empirical analysis uses nominal prices, but descriptive statistics in Table 2 use values in constant 2010 U.S. dollars to make the Value column economically interpretable.

The appendix Table A.3 presents descriptive statistics for the sample of Finnish companies that export their products during the study period.

3.2 Sample Selection

For each HS4 market, we retain only destination countries in which Finland has been an active exporter for at least one year during our sample period, as we estimate the demand model on markets relevant to Finnish industrial policy. Within each market, we rank importers by total import quantity and keep those that cumulatively account for 95 percent of total trade volume, removing small and noisy destination markets while preserving the vast majority of trade. We exclude Venezuela as both an importer and an exporter due to the extreme inflation episodes and exchange rate instability during the study period. We also drop enclaves in our data (Lesotho, San Marino, Monaco), whose trade observations are likely to be re-exports or customs artifacts rather than domestic production and produce extreme unit values that could distort the demand inversion. We exclude observations with quantity below one metric tonne to remove negligible shipments that may represent samples, returns, or recording errors. Finally, we apply the outlier-replacement procedure of Piveteau and Smagghue (2023): we regress unit prices on year dummies with exporter–importer-pair and HS6 product fixed effects. We then replace observations whose absolute log-price residual exceeds two with the regression's predicted price.

4 Demand Model

4.1 A Discrete Choice Demand Model of International Trade

Our demand model is based on Piveteau and Smagghue (2023), who first used a BLP-like demand specification to study demand in an international-trade framework. In the canonical model, consumer i receives utility from consuming good $j \in \mathcal{J}$ in market $t \in \mathcal{T}$. In our application, a product is a combination of a 6-digit HS code $k \in \mathcal{K}$ and a firm $f \in \mathcal{F}$.[6] Furthermore, a market is a destination country $d \in \mathcal{D}$ in a given year. Our data combines aggregate country-level trade data with Finnish micro-level export data. The set of firms $\mathcal{F}$ therefore consists of both aggregate countries (such as Sweden or Germany) and Finnish firms. In practice, we replace aggregate

6. Table A.1 in Appendix A.1 lists the notation used throughout the paper.

Finland with our observations of the micro-level exporters. The utility model can be presented as

$$U_{ijt} = \delta_{jt} + \mu_{ijt} + \epsilon_{ijt}, \tag{1}$$

in which the idiosyncratic taste shock ϵ_{ijt} follows a Type-I Extreme Value distribution, μ_{ijt} is the consumer-specific price term defined below, and δ_{jt} is the mean utility from the exogenous product characteristics:

$$\delta_{jt} = \beta \mathbf{X}^{ex}_{jt} + \xi_{jt}. \tag{2}$$

The parameters β of the exogenous product characteristics $\mathbf{X}^{ex}_{jt}$ capture distance and a set of fixed effects: HS6 product, exporter, destination, and year.[7] The term ξ_{jt} is the unobserved demand shock for product j in market t, which can be correlated with prices. We discuss our identification strategy with respect to price endogeneity in Section 4.2.

Price Coefficient and Pass-through. The term μ_{ijt} in Equation (1) carries the entire price response, with a lognormal random coefficient on price that is heterogeneous across destinations and consumers:

$$\begin{aligned} \mu_{ijt} &= \alpha_i \, p_{jt}, \\ \alpha_i &= -\exp\big(\pi + \pi_{gdp}\,\mathrm{GDP}_{dt} + \sigma_\alpha\, \nu_i^p\big), \\ \nu_i^p &\sim \mathcal{N}(0,1). \end{aligned} \tag{3}$$

The price coefficient α_i is strictly negative by construction. The intercept π sets the mean log price-sensitivity, π_{gdp} scales it with destination GDP per capita, σ_α governs the unobserved consumer-level dispersion, and ν_i^p is the standardized consumer-level taste draw.

Specifying the price coefficient as a lognormal random coefficient ensures demand slopes downward for every consumer and allows for more flexible demand elasticities and cost pass-through, as shown in Miravete, Seim, and Thurk (2023). This specification choice plays a central role in determining the pass-through results we report below. Logit and CES gravity specifications cap pass-through at unity by construction. When consumers face only horizontal differentiation with a constant elasticity of substitution, the equilibrium markup is bounded, so incomplete pass-through follows mechanically. More generally, Mrázová and Neary (2017) show that any well-behaved demand function is summarized by its demand manifold, a smooth curve relating the elasticity and convexity

7. The HS6 product fixed effect is dropped in the single-HS6 markets 4801 and 4803, where it has no within-market variation.

Table 3: Instruments

Instrument	D/C	Variation level
Tariff rate	C	$d \times k \times t$
Transport margin (HS4)	C	$f \times t$
Exchange rate	C	$f \times t$
Rival sums	D	$d \times t \times j$
Rival sums by k	D	$d \times t \times j$
Diff. IV (local)	D	$d \times t \times j$
Diff. IV (quadratic)	D	$d \times t \times j$
Diff. IVs by k	D	$d \times t \times j$
GDP interactions	D	As above

Notes: D = demand shifter, C = cost shifter. Diff. IVs follow Gandhi and Houde (2019). Each diff. IV block includes instruments based on distance, tariff, margin, exchange rate, and price. "By k" variants compute rival sets within HS6 groups.

of demand. CES is the knife-edge case whose manifold collapses to a single point. Therefore, the choice of functional form is the substantive modeling choice. Mixed-logit demand with skewed random coefficients does not impose this cap. Miravete, Seim, and Thurk (2023) characterize the curvature conditions under which discrete-choice demand admits pass-through above unity, and a lognormal random coefficient on price falls within that admissible class. Whether pass-through is above, below, or at unity is therefore determined by the data within our specification rather than fixed in advance by functional form.

4.2 Demand Model Identification

Identification of Price Elasticity of Demand. We need instruments to estimate the mean price-sensitivity parameters π and π_{gdp} and the random coefficient on prices, σ_α. As shown by Berry and Haile (2014), a BLP demand system also requires instruments that shift quantities or market shares. We group our instruments into demand and cost shifters and present them in Table 3.

First, we use tariff and transportation cost data as instruments. Our tariff rates from Teti (2024) provide product × market-level variation in prices that we treat as orthogonal to consumer demand. This identification requires that countries do not set their tariff rates in response to unobserved quality shocks of foreign manufacturers. Our data on transportation margins, observed at the four-digit HS level, vary at the exporter × market level. Both of these instruments are cost shifters. Because our demand models include interactions of prices with income, we add the interactions of our cost shifters with income as additional instruments.

We also construct several Gandhi and Houde (2019) differentiation instruments.[8] The first two are based on the two cost shifters. These two differentiation instruments capture the changes in rivals' costs and therefore shift the quantities or market shares of the firm in question. We also construct a differentiation instrument from the one exogenous product characteristic in our model: the distance between the destination country and the exporting country. Our final differentiation instrument is computed from predicted prices, which identifies the price heterogeneity parameter. We obtain the predicted prices by regressing observed prices on only exogenous regressors—product characteristics, fixed effects, and the excluded cost-shifter instruments. The predicted-price instrument therefore inherits exogeneity from the regressors used in the prediction.

Identification of Distance Elasticity of Demand. Distance enters as the only exogenous product characteristic with a free coefficient in $\mathbf{X}_{jt}^{ex}$. Because the firm set $\mathcal{F}$ includes both Finnish micro-firms and aggregate-country exporters with materially different distances to each destination, distance varies at the exporter–destination–year level. The mean-utility fixed effects on exporter, destination, year, and HS6 absorb the corresponding main effects but not the exporter–destination interaction, which is the residual variation that identifies the distance coefficient. The recovered elasticities reflect a structural preference for proximity, net of multilateral resistance and aggregate trade-cost confounders. This identification argument is in the spirit of Anderson and van Wincoop (2003) on multilateral resistance and Novy (2013) on non-CES gravity.

4.3 Estimation

We estimate the demand model separately for each four-digit HS product market. This per-market design lets parameter estimates vary across broad product categories at the cost of estimating a large number of computationally demanding models. Most empirical industrial organization applications estimate a single demand system per product market (cars, cereals, beer). Döpper, MacKay, Miller, and Stiebale (2025) is one exception, estimating thousands of demand models using the covariance restrictions of MacKay and Miller (2025). We aim instead for good coverage across

8. Our differentiation instruments take the local and/or quadratic (sum-of-squared-differences) forms.

several industries at the four-digit HS level, while still using a standard instrumental-variable approach.

From our trade data, we observe only exports to each market. This export-only coverage leaves domestic production and the market share of the outside option unobserved. Demand estimation therefore requires a market-potential input, which determines market shares from observed quantities. Traditionally, researchers have set the market potential using simple population-based measures, or imposed a two-nest structure for the structural error term ϵ_{ijt} with inside goods in one nest and the (unobserved) outside good in another nest, which yields a random-coefficients nested logit model (Grigolon and Verboven 2014).

Zhang (2023) show that the market potential can be estimated as part of a slightly modified BLP procedure. The procedure requires only the quantities sold for the products in our data. The modification relative to the canonical BLP estimation concerns the fixed-point contraction. In Zhang (2023), the original contraction mapping that matches predicted and observed market shares,

$$\delta_{jt}^{h+1} = \delta_{jt}^{h} + \ln(s_{jt}) - \ln(s_{jt}(\boldsymbol{\delta}, \sigma)) \tag{4}$$

is replaced by

$$\delta_{jt}^{h+1} = \delta_{jt}^{h} + \ln(q_{jt}) - \ln(\phi\Psi_t) - \ln(s_{jt}(\boldsymbol{\delta}, \sigma)). \tag{5}$$

Here, q_{jt} denotes quantities and Ψ_t the market potential. This procedure adds an additional parameter, ϕ, to be estimated alongside the usual non-linear parameter σ_α.

In our application, we set Ψ_t equal to the observed sales of the inside goods and restrict ϕ to a single scalar within each per-HS4 generalized method of moments (GMM) estimation, common across the destination-year markets (d, t) pooled in that estimation. Because we estimate the model separately on each of our nine HS4 products, Tables 4 and 5 report one ϕ per HS4. The model's outside option then captures both domestic consumption and the true outside option. We therefore estimate the usual demand parameters $\theta = (\beta, \pi, \pi_{gdp}, \sigma_\alpha)$ from Equations (2) and (3) together with the per-HS4 market-potential parameter ϕ.

Otherwise, our estimation procedure follows the best practices laid out in Conlon and Gortmaker (2020). We absorb all of our fixed effects and use Gaussian quadrature with 10 nodes to integrate over consumer-level taste draws ν_i^p (see Train (2009) for further details). Following Illanes, Bhattacharya, Stillerman, Kreps, and Salas (2024), we replace all of our instruments, distance, and income data with their respective Z-scores before estimation, which improves the numerical stability of our weighting matrices. We solve the demand systems with the trust-constr algorithm using analytical gradients, with the Varadhan and Roland (2008) and Reynaerts, Varadhan, and Nash (2012) SQUAREM inner loop run to tolerance 10^{-14} and the trust-constr outer loop run to gradient tolerance 10^{-5}.

5 Supply Model

5.1 Marginal Costs

Firms set prices in Bertrand-Nash competition. Each firm maximizes:

$$\underset{p_{jt}}{\text{maximize}} \quad \sum_{j \in J^f} (p_{jt} - c_{jt}) \cdot \phi \Psi_t \cdot s_{jt}(\mathbf{p}, \mathbf{x}, \boldsymbol{\xi}, \theta), \tag{6}$$

in which ϕ is the estimated market-potential scaling factor and s_{jt} denotes market shares as a function of prices, product characteristics, demand shocks, and the demand parameters. This specification yields the first-order conditions (FOCs) relating prices to marginal costs c_{jt}:

$$\sum_{m \in J^f} (p_{mt} - c_{mt}) \cdot \frac{\partial s_{mt}}{\partial p_{jt}} + s_{jt}(\mathbf{p}_t, \mathbf{x}_t, \boldsymbol{\xi}_t, \theta) = 0. \tag{7}$$

This FOC yields the following expression for the markup:

$$\mathbf{p}_t - \mathbf{c}_t = - \left[\boldsymbol{\Omega} \odot \frac{\partial \mathbf{s}_t}{\partial \mathbf{p}_t} \right]^{-1} \mathbf{s}_t. \tag{8}$$

in which $\boldsymbol{\Omega}$ denotes the ownership matrix of the inside goods—both the Finnish firm-level exporters and the non-Finnish country-level aggregates—and $\odot$ denotes element-wise multiplication. The prices $\mathbf{p}_t$, the market shares $\mathbf{s}_t$, and the ownership matrix are data. We obtain our estimate for the demand jacobian $\partial \mathbf{s}_t / \partial \mathbf{p}_t$ from our demand model. We define the price-cost markup as $\mathcal{M}_{jt} \equiv p_{jt}/c_{jt}$. We assign ownership at the exporter level: a Finnish firm owns the HS6 varieties it exports, and each non-Finnish exporting country owns the varieties it ships, so a country aggregate enters the pricing game as a single multi-product exporter rather than a competitive fringe. The same $\boldsymbol{\Omega}$ governs both the markup inversion here and the counterfactual re-pricing of Section 6, in which non-Finnish exporters reprice their full HS6 portfolio at their unshocked baseline marginal costs. As the common-ownership literature stresses (Backus, Conlon, and Sinkinson 2020, 2021), aggregating each non-Finnish country into a single owner is itself an assumption: it imposes within-country price coordination that the true firm-level structure need not exhibit. A misspecified ownership matrix maps into biased markups and hence biased recovered marginal costs.

Because the demand model is estimated on Cost, Insurance, and Freight (CIF) prices, the marginal costs c_{jt} recovered from Equation (8) are also on a CIF basis. For the cost-side analysis of Section 5.2 the relevant object is the producer-side (FOB) marginal cost. We obtain it by dividing the recovered c_{jt} by the same exporter–destination–year tariff and transportation-margin factors used to construct CIF prices in Section 4: $c_{jt}^{\text{FOB}} = c_{jt}/[(1+\tau_{jt})(1+m_{jt})]$. Here τ_{jt} is the ad valorem tariff from Teti (2024) and m_{jt} the transportation-and-insurance margin from Fiallos, Liberatore, and Cassimon (2024). Appendix A.3 derives this transformation from the producer-side Bertrand–Nash FOC. All subsequent uses of c_{jt} refer to this stripped, producer-side marginal cost.

5.2 A Sufficient Statistic for Cost-Side Incidence

To evaluate Finland's 2017–2019 internal devaluation policy, we recover the cost-side incidence of the labor-cost change: how a firm's marginal cost responds to an input-price shock. It follows from two objects already in hand, without estimating a production or cost function: the firm-year markup $\bar{\mathcal{M}}_{ft}$ from Section 5.1 and the input's revenue share from the financial statements. The markup combines the BLP demand estimates with the Bertrand-Nash FOC. It is the cost-weighted within-firm-year mean of the product markups $\mathcal{M}_{jt} = p_{jt}/c_{jt}$, computed within each HS4 market. A firm spanning several markets therefore carries a distinct markup in each. The mechanism is the Hall (1988) and De Loecker and Warzynski (2012) identity: a variable input's output elasticity equals the markup times its revenue share, and under constant returns to scale (CRS) that elasticity is the input's cost share. Previous empirical studies use the identity to recover the markup from an estimated elasticity. We invert it: we recover the markup from demand and observe the revenue share, so the identity returns the cost share that governs incidence.

Setup. Index firms by f, years by t, and products by j. A firm may make several products. Let V collect the variable inputs whose prices may move and K_{ft} the predetermined capital stock. The variable-cost function delivers output Q_{ft} at minimum outlay over the variable inputs, holding

K_{ft} fixed:

$$C_{ft}(\mathbf{w}_{ft}, Q_{ft}) = \min_{\mathbf{x}\geq 0} \mathbf{w}_{ft}\cdot\mathbf{x} \quad \text{s.t.} \quad F_{ft}(\mathbf{x}; K_{ft}) \geq Q_{ft}. \tag{9}$$

Input v enters through two shares: its revenue share divides input spending by revenue, its cost share by variable cost,

$$s^R_{v,ft} \equiv \frac{w_{v,ft}\, x_{v,ft}}{R_{ft}}, \qquad s^C_{v,ft} \equiv \frac{w_{v,ft}\, x^{\star}_{v,ft}}{C_{ft}}. \tag{10}$$

The revenue share is observed in the firm financial accounts. We do not observe total variable costs, so the cost share, the Shephard's-lemma object that governs incidence, is not observed either. The markup connects the two.

The identity. The result rests on five conditions: an exogenous input-price shock, CRS in the variable inputs, an interior cost minimum, a markup recovered from demand analysis, and observed accounting shares.

Assumption 1 (Exogenous variable-input price). The perturbed input price $w_{v,ft}$ is independent of the input quantity $x_{v,ft}$ within the year and carries no adjustment costs.

Restriction 1 (Variable-input CRS). $C_{ft}(\mathbf{w}_{ft}, Q) = c_{ft}(\mathbf{w}_{ft})\cdot Q$, equivalently $\gamma_{ft} \equiv \sum_{v\in V}\beta_{v,ft} = 1$, in which γ_{ft} is the returns-to-scale parameter and $\beta_{v,ft}$ the output elasticity of input v (distinct from the demand-side parameter vector β in Equation (2)).

Regularity Condition 1 (Interior, single-valued conditional factor demand). Problem (9) admits an interior optimum with single-valued $\mathbf{x}^{\star}_{ft}(\mathbf{w}_{ft}, Q)$.

Assumption 2 (BLP markup recoverability). $\bar{\mathcal{M}}_{ft}$ is identified from the Bertrand-Nash FOC of Section 5.1, using only demand-side data.

Assumption 3 (Firm-level accounting). $w_{v,ft}\, x_{v,ft}$ and R_{ft} are observed at the firm-year level for any v whose price is perturbed.

Proposition 1 (Sufficient statistic for cost-side incidence). Under Restriction 1, Regularity 1, and Assumptions 1, 2, and 3, a log shock $d\ln w_{v,ft}$ to one variable-input price moves marginal cost by

$$d\ln c_{ft} = s^C_{v,ft}\cdot d\ln w_{v,ft}, \qquad s^C_{v,ft} = \bar{\mathcal{M}}_{ft}\cdot s^R_{v,ft}. \tag{11}$$

Shephard's lemma gives the cost-side incidence directly, $\partial \ln C_{ft}/\partial \ln w_{v,ft} = s^C_{v,ft}$. The first equality is an envelope result, so the firm's optimal input substitution is second-order and does not enter—and CRS makes marginal cost move with the cost function, $\partial \ln c_{ft}/\partial \ln w_{v,ft} = \partial \ln C_{ft}/\partial \ln w_{v,ft}$. That is the first equality. For the second, the cost-minimization FOC is $w_{v,ft} = c_{ft}\, \partial F_{ft}/\partial x_v$, so the marginal product is $\partial F_{ft}/\partial x_v = w_{v,ft}/c_{ft}$ and the output elasticity is $\beta_{v,ft} \equiv (\partial F_{ft}/\partial x_v)(x_v/Q_{ft}) = w_{v,ft}\, x_v/(c_{ft} Q_{ft})$. With the markup $\bar{\mathcal{M}}_{ft} = p_{ft}/c_{ft} = R_{ft}/C_{ft}$, this becomes $\beta_{v,ft} = (p_{ft}/c_{ft})\,(w_{v,ft} x_v/R_{ft}) = \bar{\mathcal{M}}_{ft}\, s^R_{v,ft}$. Under CRS the output elasticity equals the cost share, $\beta_{v,ft} = s^C_{v,ft}$, which is the second equality.

The ratio $\bar{\mathcal{M}}_{ft} = R_{ft}/C_{ft}$ is the firm's gross markup of revenue over variable cost, and cost-weighting the firm's products makes the firm-year markup and this ratio the same object: $\bar{\mathcal{M}}_{ft} = \sum_j p_j Q_j / \sum_j c_j Q_j = R_{ft}/C_{ft}$. A revenue share understates the cost share by exactly this factor (it divides input spending by revenue rather than by cost), so scaling $s^R_{v,ft}$ by R_{ft}/C_{ft} recovers $s^C_{v,ft}$.

Corollary 1 (Bertrand-Nash equilibrium response)**.** Under Restriction 1, Regularity 1, Assumptions 1, 2, and 3, and BLP demand parameters θ_D from Section 4.2,

$$d\ln c_{ft} \;=\; \sum_{v \in V'} \bar{\mathcal{M}}_{ft} \cdot s^R_{v,ft} \cdot d\ln w_{v,ft}, \qquad d\ln \mathbf{p} \;=\; \Phi\big(d\ln \mathbf{c};\, \theta_D\big), \tag{12}$$

in which $\Phi(\cdot;\theta_D)$ is the equilibrium price-response map induced by the Bertrand-Nash FOCs of Section 5.1 at demand parameters θ_D.

The identity covers any single variable input by relabeling v. The devaluation policy counterfactuals in Section 6 are the $v = L$ case, with the labor revenue share. A tariff or imported-input shock is $v = M$, an energy-price shock is $v = E$. In each case the BLP markup and the matching observed revenue share deliver the cost-side incidence.

Identification. The identity $\beta_{v,ft} = \bar{\mathcal{M}}_{ft} \cdot s^R_{v,ft}$ is usually read left-to-right: estimate $\beta_{v,ft}$ from a production function, divide by the observed revenue share, and recover the markup. Bond, Hashemi, Kaplan, and Zoch (2021) show that this direction is not identified when only firm-level revenue is observed, because separate price and quantity variation is unavailable to disentangle a revenue elasticity from a markup. Doraszelski and Jaumandreu (2021) reach a related conclusion. Our route reverses the identity. We recover $\bar{\mathcal{M}}_{ft}$ from demand inversion (Assumption 2), combine it with the directly observed $s^R_{v,ft}$ (Assumption 3), and read off $s^C_{v,ft}$ (equivalently $\beta_{v,ft}$). Because we do not estimate revenue elasticity from a production function, the Bond, Hashemi, Kaplan, and Zoch (2021) critique does not apply.

5.3 Extensions

Multi-product allocation. Our sufficient statistic for cost-side incidence allows multiproduct firms to have product-specific production technologies, even if input allocation across products is unobservable. Estimation of these firms' production functions would require estimating the product-specific input allocation as, for example, in Orr (2022) or Valmari (2023). However, the identity $\beta_{v,jt} = \mathcal{M}_{jt} \cdot s^R_{v,jt}$ holds product by product, provided variable inputs are split across products in proportion to product revenue. This split is exact under separability with a common variable-input technology, and a first-order approximation otherwise. We don't observe product-specific input quantities, so we define the cost share as $s^C_{v,jt} = \bar{\mathcal{M}}_{ft} \cdot s^R_{v,jt}$, with $s^R_{v,jt}$ being defined on product-j revenue and the firm-year markup broadcast across the firm's products.

Relaxing CRS. Restriction 1 can be dropped at a cost. With decreasing returns ($\gamma_{ft} < 1$), marginal cost is no longer scale-free: $c_{ft}(Q) \propto Q^{1/\gamma_{ft}-1}$. The incidence read from $\bar{\mathcal{M}}_{ft}\, s^R_{v,ft}$ then rescales by $1/\gamma_{ft}$, and the equilibrium response becomes a fixed point in Q that requires an external value of γ_{ft}.

6 Counterfactual Simulation

We estimate three industrial policy counterfactuals. The first counterfactual simulates the equilibrium effects of Finland's 2017–2019 internal devaluation policy ("kiky"). The policy extended employees' annual working time by 24 hours without changing annual compensation and phased in employer-side social-security contribution cuts. The second counterfactual studies the more ambitious alternative devaluation policy, a flat reduction in employer contributions paired with a smaller hours extension. The third counterfactual demonstrates how the demand for Finnish exports responds to a substantial increase in labor costs. We model this shock as the removal of one standard eight-hour workday per week, corresponding to an approximately 24 percent reduction in annual working hours. This scenario reverses the sign of the hours channel relative to the realized devaluation policy: working time contracts rather than expands.

In each counterfactual scenario the simulation translates the different labor-cost shocks through the sufficient-statistic cost-incidence identity of Section 5.2 into a marginal-cost change and re-solves the demand-side Bertrand-Nash pricing game under the new costs.

6.1 From Policy to Labor Price

The counterfactual policies bundled two cost-side changes acting on the effective price of labor: an extension (or reduction) of the standard annual working hours at unchanged nominal pay, and a reduction in employer-side social-insurance contribution rates. We translate both into a single wage-equivalent log change that varies by year.

Hours extension. Let H denote baseline annual working hours and e_t the additional hours mandated by the counterfactual policy in year t. At unchanged annual compensation, the effective hourly labor cost falls in proportion:

$$d\ln P_{ft}^{L,\text{hours}} = \ln\left(\frac{H}{H+e_t}\right) \approx -\frac{e_t}{H}, \tag{13}$$

with $H = 1720$, the baseline annual working hours under the prevailing Finnish industrial collective agreements (broadly 37.5 hours per week times 45.87 effective working weeks, net of statutory holidays and vacation). The realized devaluation policy ("kiky") set $e_t = 24$ in 2017–2019 and $e_t = 0$ in 2020, when the pact rolled back the hours extension. The hours channel therefore contributes about -0.014 per year in 2017–2019 and zero in 2020.

Employer social security contribution cut. Let $\Delta\tau_t$ denote the cumulative percentage-point reduction in total employer labor cost in year t relative to the pre-pact 2016 baseline, driven by the cuts to employer-side social-insurance contribution rates negotiated in the pact. The log change in employer labor cost W_t is then $d\ln W_t = \ln(1 + \Delta\tau_t/100)$.[9] The realized cuts were $\Delta\tau_t = -1.59, -2.19, -2.23, -0.58$ pp in 2017–2020 respectively, reflecting a phased schedule that peaked in 2018–2019 and largely unwound in 2020.

9. We treat the statutory employer-side contribution cut as passing fully into the employer's effective labor price, leaving employee earnings unchanged.

Composite. The total log change in the effective hourly labor cost is the year-specific sum:

$$d\ln P_t^L \;=\; d\ln W_t + \ln\!\left(\frac{H}{H+e_t}\right), \tag{14}$$

which evaluates to approximately -0.030 in 2017, -0.036 in 2018–2019 (the full bite), and only -0.006 in 2020 (residual contribution cut, no hours extension). The first internal devaluation policy counterfactual reverses these year-specific shocks one for one. The more ambitious devaluation policy and four-day-week counterfactuals instead apply a flat shock across 2017–2020, because they are alternative permanent policies rather than the phased realized pact. The more ambitious devaluation policy sets $d\ln P^L \approx -0.0499$ (Prime Minister's Office of Finland 2015).

The four-day-week shock removes the equivalent of one standard eight-hour workday per week across all 52 weeks, $e = -52 \times 8 = -416$ hours, applied uniformly across 2017–2020 with $\Delta\tau = 0$.[10] Equation (13) then yields $d\ln P_{ft}^{L,\text{hours}} = \ln(1720/1304) \approx +0.277$ in every treated year.

6.2 From Labor Price to Counterfactual Equilibrium

The sufficient-statistic identity translates the applied labor-price shock from Section 6.1 into a firm-year marginal-cost change. Setting $v = L$ in Proposition 1 with $w_{L,ft} = P_{ft}^L$, and using the year-specific $d\ln P_t^L$ from the composite of Equation (14), we obtain:

$$d\ln c_{ft}^{\text{cf}} \;=\; \bar{\mathcal{M}}_{ft} \cdot \frac{P_{ft}^L L_{ft}}{R_{ft}} \cdot d\ln P_t^L. \tag{15}$$

The markup-corrected labor cost share $\bar{\mathcal{M}}_{ft}\cdot(P_{ft}^L L_{ft}/R_{ft})$ rescales the firm-year revenue share $P_{ft}^L L_{ft}/R_{ft}$ from accounting data (Assumption 3) by the BLP-recovered markup $\bar{\mathcal{M}}_{ft}$ of Section 5.1. Shephard's lemma delivers this rescaling. Here $\bar{\mathcal{M}}_{ft}$ is the firm-year markup $\sum_j p_{jt}Q_{jt}/\sum_j c_{jt}Q_{jt}$, the cost-weighted within-firm-year average of p_j/c_j across the firm's HS6 products j. We apply Equation (15) at this firm-year level with the year-specific shock $d\ln P_t^L$ of Equation (14). The HS4-level means reported in Table A.4 are summary statistics rather than inputs to the simulation.

10. Relative to the 1720-hour Finnish baseline this is a 24.2 percent contraction in annual hours. We adopt the "drop one standard eight-hour workday" convention because it matches the public framing of the four-day-week proposal more directly than a proportional 20 percent trim of 1720 hours (which would give $e = -344$ hours).

The counterfactual marginal cost for treated (Finnish) firms in years $t \in \{2017, 2018, 2019, 2020\}$ is

$$c_{jt}^{\text{cf}} \;=\; c_{jt}^{\text{bl}} \cdot \left(1 + d\ln c_{ft}^{\text{cf}}\right), \tag{16}$$

with c_{jt}^{bl} recovered via the Bertrand-Nash markup inversion of Equation (8). Untreated (non-Finnish) firms retain their baseline costs. Equation (16) is the first-order Taylor approximation of the marginal-cost response. Because $\partial \ln c_{ft} / \partial \ln w_{v,ft} = s_{v,ft}^{C}$ from Equation (11), expanding c_{ft} to first order in the input-price log-change gives $c_{jt}^{\text{cf}} = c_{jt}^{\text{bl}}\,(1 + d\ln c_{ft}^{\text{cf}})$. The sufficient statistic identifies the local, marginal response, so we apply it to first order rather than imposing the constant-elasticity extrapolation $\exp(d\ln c_{ft}^{\text{cf}})$, which would impose curvature the identity does not deliver. The two formulations coincide for small shocks and diverge only for the large four-day-week shock. We apply Equation (16) uniformly across the firm's HS6 product lines under the labor-proportional-to-revenue allocation discussed in Section 5.2.

Given the counterfactual costs $\mathbf{c}^{\text{cf}}$ delivered by Equation (16), we resolve for the Bertrand-Nash equilibrium prices using the Morrow and Skerlos (2011) fixed-point algorithm. We simulate two equilibria for each studied scenario, one with baseline costs $\mathbf{c}^{\text{bl}}$ and one with counterfactual costs $\mathbf{c}^{\text{cf}}$, and report changes between them. A simulated baseline ensures both sides use identical numerical procedures, and avoids discrepancies between the contraction mapping that estimation uses and the equilibrium solver that the simulation uses.

Four-day-week counterfactual. The four-day-week scenario is a separate hypothetical policy rather than an internal devaluation policy variant. It reuses the full pipeline of the devaluation scenarios and differs only in the policy parameters. Using the flat hours and contribution parameters introduced in the composite paragraph above, it feeds Equations (13)–(14) into the same firm-year sufficient-statistic identity of Equation (15), the same counterfactual marginal cost of Equation (16), and the same simulation as the realized devaluation policy.

6.3 Outcome Measures in Counterfactual Analysis

We measure the counterfactual at the product-destination-year level as the change in prices, quantities, and revenue between the simulated baseline and the post-shock equilibrium. For each observation in an affected market, we compute:

$$
\begin{aligned}
\Delta p_{jt} &= p_{jt}^{\mathrm{cf}} - p_{jt}^{\mathrm{bl}}, && \text{(price change)} \\
\text{pass-through}_{jt} &= \frac{\Delta p_{jt}}{\Delta c_{jt}}, && \text{(treated firms only)} \\
\Delta s_{jt} &= s_{jt}^{\mathrm{cf}} - s_{jt}^{\mathrm{bl}}, && \text{(market share change)} \\
\Delta R_{jt} &= p_{jt}^{\mathrm{cf}} q_{jt}^{\mathrm{cf}} - p_{jt}^{\mathrm{bl}} q_{jt}^{\mathrm{bl}}, && \text{(revenue change)}
\end{aligned}
$$

in which quantities are $q_{jt} = s_{jt} \cdot \hat{\phi} \cdot \Psi_t$, with $\hat{\phi}$ the estimated market-potential parameter from Equation (5) and Ψ_t the observed inside-good sales (in tonnes). Let $i(j)$ denote the exporter country of variety j. We aggregate the per-observation revenue change separately for Finnish (treated) and non-Finnish (rival) exporters:

$$
\begin{aligned}
\Delta R_t^{\mathrm{FI}} &= \sum_{j:\, i(j)=\mathrm{FI}} \Delta R_{jt}, && \text{(Finnish exporter revenue change)} \\
\Delta R_t^{\mathrm{rivals}} &= \sum_{j:\, i(j)\neq\mathrm{FI}} \Delta R_{jt}, && \text{(rival exporter revenue change)}
\end{aligned}
$$

with the net market response given by $\Delta R_t^{\mathrm{net}} = \Delta R_t^{\mathrm{FI}} + \Delta R_t^{\mathrm{rivals}}$. The same Finland/rivals split applies to Δp and Δs where reported, but not to pass-through, which is defined for treated (Finnish) firms only because non-Finnish rivals face no cost shock. We aggregate over the treated years $t \in \{2017, \ldots, 2020\}$ and across the nine HS4 markets in the headline tables.

7 Results

7.1 Demand side

Tables 4 and 5 report the demand-side GMM parameter estimates for each HS4 market: the distance coefficient β, the lognormal random coefficient on prices σ, the demographic interaction π and its GDP gradient π_{gdp}, and the market-potential scaling parameter ϕ from Equation (5). We report implied price sensitivities $E[|\alpha_i|]$ at the GDP mean and at ± 1 standard deviation around it, showing how income shifts price sensitivity across destinations.

Three patterns hold across the nine markets. First, the GDP gradient π_{gdp} is negative in every market, so higher-GDP destinations are systemat-

Table 4: Two-step GMM estimates of the BLP demand system (Equations (2)–(3)) by HS4 market (Part 1 of 2).

Market	4407	4412	4703	4801
β, Dist.	-0.077	-0.105	-0.086	-0.066
	(0.005)	(0.004)	(0.008)	(0.015)
σ, Price	1.092	0.758	0.696	0.400
	(0.032)	(0.071)	(0.082)	(0.083)
π	2.687	2.106	2.380	2.425
	(0.099)	(0.139)	(0.095)	(0.236)
π_{gdp}	-0.066	-0.216	-0.089	-0.166
	(0.016)	(0.055)	(0.025)	(0.056)
ϕ	1.233	2.030	1.072	1.164
	(0.176)	(0.212)	(0.067)	(0.178)
$E[\|\alpha_i\|]$, GDP mean	26.663	10.955	13.769	12.242
$E[\|\alpha_i\|]$, GDP +1sd	24.953	8.829	12.598	10.367
$E[\|\alpha_i\|]$, GDP -1sd	28.489	13.594	15.050	14.456

Notes: Robust standard errors in parentheses. Each column corresponds to a separate HS4 product market. All models estimated by two-step GMM with analytical gradients.

ically less price-sensitive—consistent with the standard income-elasticity prior in trade demand. Second, the lognormal random coefficient on price, σ, is identified separately in every market and spans 0.28 (kraft paper) to 1.09 (sawn wood), recovering within-market heterogeneity in price sensitivity throughout. Third, the mean price-sensitivity level $E[|\alpha_i|]$ at the GDP mean ranges from ≈ 2.4 (stainless steel) to ≈ 26.7 (sawn wood), a wide dispersion that reflects cross-market differences in price elasticity.

Tables 6 and 7 summarize the BLP-recovered demand-side outcomes for each HS4 market: mean own-price elasticities, marginal costs, markups, and Lerner indices.

Quantity-weighted mean own-price elasticities range from -2.9 (stainless steel) to -6.8 (newsprint) and distance elasticities from -0.2 (kraft paper) to -0.6 (plywood): demand responds more sharply to prices than to distance in every market.[11]

11. These distance elasticities are $\partial \ln q / \partial \ln \text{dist}$ evaluated at observed distances—the object directly comparable to a gravity log-distance coefficient. The linear distance coefficient β in Tables 4–5 enters mean utility in rescaled levels (km/1000) and is not the object compared to the gravity benchmark.

Table 5: Two-step GMM estimates of the BLP demand system (Equations (2)–(3)) by HS4 market (Part 2 of 2).

Market	4802	4803	4804	4805	7219
β, Dist.	-0.093	-0.100	-0.046	-0.074	-0.119
	(0.003)	(0.013)	(0.003)	(0.005)	(0.002)
σ, Price	1.032	0.634	0.283	0.899	0.706
	(0.032)	(0.126)	(0.196)	(0.032)	(0.064)
π	2.036	2.429	1.327	2.367	0.614
	(0.059)	(0.665)	(0.960)	(0.051)	(0.247)
π_{gdp}	-0.072	-0.050	-0.496	-0.024	-0.357
	(0.009)	(0.049)	(0.142)	(0.010)	(0.026)
ϕ	1.209	11.699	1.986	1.178	3.115
	(0.028)	(23.637)	(12.423)	(0.009)	(1.473)
$E[\lvert\alpha_i\rvert]$, GDP mean	13.043	13.876	3.924	15.988	2.371
$E[\lvert\alpha_i\rvert]$, GDP +1sd	12.138	13.200	2.389	15.615	1.659
$E[\lvert\alpha_i\rvert]$, GDP −1sd	14.014	14.587	6.445	16.370	3.387

Notes: Robust standard errors in parentheses. Each column corresponds to a separate HS4 product market. All models estimated by two-step GMM with analytical gradients.

Comparisons to previous markup literature. Our Lerner indices of 23 to 61 percent overlap the range that demand- and production-side studies report for manufacturing. De Loecker, Eeckhout, and Unger (2020) document U.S. markup ratios rising from 1.21 in 1980 to 1.61 by 2016 (Lerner indices of 17 to 38 percent), driven by the upper tail. Our results do not support the notion of rising markup trends across the studied product markets. Figure A.6 shows that estimated markups follow a common U-shaped pattern across all nine markets: an early-2000s peak, a post-financial-crisis trough, and a partial recovery toward initial levels.

Differentiated paper grades and stainless steel have the highest Lerner indices, with kraft paper at 61 percent and stainless steel at 53 percent, whereas the more commodity pulp has the lowest Lerner index, with chemical pulp at 23 percent. The pattern matches the right-tail emphasis in the markup-heterogeneity literature (De Loecker, Eeckhout, and Unger 2020; Albrecht and Decker 2026): pricing power concentrates in firms producing differentiated, harder-to-substitute varieties, whereas commodity grades operate closer to competitive benchmarks. Published macro or productivity literature estimates are typically markup ratios $\mu = p/c$ derived from production-function inversion; ours are Lerner

Table 6: Quantity-weighted means of BLP-recovered own-price elasticities, distance elasticities, marginal costs, and markups for Finnish exporters by HS4 market (Part 1 of 2).

Market	4407	4412	4703	4801
Price elasticity	-4.444	-4.910	-6.524	-6.827
	(1.699)	(1.062)	(1.758)	(2.969)
Distance elasticity	-0.456	-0.622	-0.522	-0.314
	(0.370)	(0.457)	(0.407)	(0.266)
Marginal Cost	0.326	0.737	0.491	0.326
	(0.283)	(0.429)	(0.163)	(0.386)
Markup	0.179	0.328	0.138	0.324
	(0.156)	(0.190)	(0.081)	(0.361)
Lerner index (%)	42.533	33.716	22.657	49.772
	(57.366)	(21.016)	(15.780)	(52.574)

Notes: Each column is a quantity-weighted mean across treated observations of the corresponding per-observation statistic: Marginal Cost = $E_q[c]$, Markup = $E_q[p-c]$, Lerner index = $100 \times E_q[(p-c)/p]$. The Lerner index column is a mean of ratios, not a ratio of column means, so by Jensen's inequality it need not equal Markup/(Markup + Marginal Cost). Standard deviations of the per-observation statistic in parentheses. Elasticities, marginal costs, and markups are computed from the demand model estimates.

indices $(p-c)/p$ from a Bertrand–Nash FOC under BLP demand.[12]

Comparisons to previous trade literature. Distance enters trade flows through two channels: transport costs, which raise the delivered price of imported goods, and consumer preferences for proximity or country-of-origin attributes (Italian leather, French wines, German engineering). Standard CES gravity bundles both channels into a single distance coefficient typically estimated near -1. Our specification separates them. Tariffs and transportation margins from Teti (2024) and Fiallos, Liberatore, and Cassimon (2024) convert reported FOB values to CIF, so the price coefficient α_i already absorbs the transport-cost channel. The residual distance coefficient is then the preference object alone, and for the differentiated intermediate goods in our sample—pulp, paper, sawn wood, and stainless steel, purchased on technical specification rather than perceived geographic origin—that object is small. Price elasticities of -2.9 to -6.8 dwarf the distance elasticities of -0.2 to -0.6, putting the response on delivered price, not on origin.

12. $\mu = \frac{1}{1-\text{Lerner}}$.

Table 7: Quantity-weighted means of BLP-recovered own-price elasticities, distance elasticities, marginal costs, and markups for Finnish exporters by HS4 market (Part 2 of 2).

Market	4802	4803	4804	4805	7219
Price elasticity	-4.555	-6.266	-3.645	-4.806	-2.876
	(1.454)	(1.913)	(2.438)	(1.473)	(0.994)
Distance elasticity	-0.476	-0.495	-0.236	-0.288	-0.592
	(0.404)	(0.440)	(0.205)	(0.295)	(0.510)
Marginal Cost	0.633	1.123	0.336	0.306	1.255
	(0.457)	(0.689)	(0.403)	(0.996)	(1.035)
Markup	0.307	0.300	0.429	0.279	1.120
	(0.270)	(0.089)	(0.280)	(0.942)	(0.378)
Lerner index (%)	33.492	23.012	60.745	47.878	52.909
	(28.451)	(7.774)	(39.171)	(136.003)	(21.241)

Notes: Each column is a quantity-weighted mean across treated observations of the corresponding per-observation statistic: Marginal Cost = $E_q[c]$, Markup = $E_q[p - c]$, Lerner index = $100 \times E_q[(p - c)/p]$. The Lerner index column is a mean of ratios, not a ratio of column means, so by Jensen's inequality it need not equal Markup/(Markup + Marginal Cost). Standard deviations of the per-observation statistic in parentheses. Elasticities, marginal costs, and markups are computed from the demand model estimates.

Gravity-literature distance elasticities cluster near -0.9, larger in magnitude than the -0.2 to -0.6 we recover. The Disdier and Head (2008) meta-analysis of 1,467 estimates from 103 studies reports a mean of -0.9 and a median of -0.89, with the central 50 percent between -1.1 and -0.7. The Head and Mayer (2014) update across 2,508 estimates reports a mean of -0.93 and a typical range of -0.7 to -1.2, re-confirmed in Larch, Shikher, and Yotov (2025). These estimates identify a composite distance coefficient. In the canonical CES gravity model of Anderson and van Wincoop (2003), $\beta_{\text{dist}} = (1 - \sigma)\,\rho$: the trade-cost elasticity ρ and the demand elasticity σ load on the same coefficient. Subtracting the transport-cost channel from a typical -0.9 leaves a preference residual of the order we recover in our demand analyses.

7.2 Counterfactual Analysis: The Effect of Internal Devaluation Policy and Reduced Work Hours on Finnish Exporters

The counterfactuals tell a consistent story: each policy moves export revenue between Finnish and foreign producers by far more than it changes the total revenue. The cross-border reallocation is the first-order effect and the net change in producer revenue second-order, demonstrating the beggar-thy-neighbor mechanism in action.

Table 8 reports the aggregate revenue effects of the three counterfactual scenarios on Finnish exporters and their non-Finnish competitors in the same destination markets. The effects are summed over the nine HS4 markets and the four treated years (2017–2020).[13] We use the markup-corrected variant presented in Subsection 5.2 as the preferred specification. Under this specification, the more ambitious devaluation policy (*Sipilä*) would have raised Finnish export revenue by €447.0 million and lowered rival revenue by €567.8 million. Eliminating the realized devaluation policy, reported as the (*No Kiky*) scenario, would have reduced Finnish export revenue by €239.0 million. A four-day work week would have reduced Finnish export revenue by €2,394.4 million. Relative to the 2017–2020 Finnish baseline of €37.8 billion in the nine HS4 markets, these effects correspond to changes of +1.2, −0.6, and −6.3 percent respectively.

The per-HS4 breakdown in Table A.5 shows substantial heterogeneity across markets. Chemical pulp exhibits the largest proportional response in every scenario: reversing the realized policy reduces its revenue by 1.07 percent, the more ambitious proposal raises it by 2.01 percent, and the four-day work week reduces it by 11.18 percent. Other paper grades and sawn wood also respond strongly, whereas kraft paper and especially stainless steel are considerably less sensitive. Despite having the largest baseline export value, stainless steel changes by only −0.06 percent under the reversal, +0.13 percent under the more ambitious proposal, and −0.69 percent under the four-day work week. This heterogeneity reflects differences in labor cost shares and demand elasticities across the markets studied.

Role of super-pass-through. The *Net* Δ column in Table 8 measures the super-pass-through wedge. In the cost-rise scenarios, total inside-good revenue increases because rival price increases generate more additional revenue for non-Finnish firms than the Finnish firms lose. Mean pass-through, defined as the simple mean across the the nine HS4 markets of the per-HS4 mean $\partial p/\partial c$ on treated observations, is approximately 1.18 across devaluation scenarios and 1.19 for the four-day week. Pass-through above one arises from the curvature of the lognormal random coefficient

13. The treated window is held fixed at 2017–2020 across all three scenarios for comparability. Because the realized kiky pact unwound in 2020 whereas the more ambitious devaluation policy and four-day-week shocks are permanent, the 2020 contributions differ across scenarios. Table A.6 makes this asymmetry visible.

on price, consistent with the mechanism documented by Miravete, Seim, and Thurk (2023) in discrete-choice demand modeling. When rivals raise prices alongside the treated Finnish firms, residual demand becomes locally less elastic, and the markup component of prices increases by more than the cost shock. The stability of pass-through across scenarios reflects local linearity around a common baseline. The first-order cost mapping of Equation (16) adds no curvature of its own, and our three shocks span implied $d\ln c$ from only ≈ 0.003 to ≈ 0.086. The above-unity pass-through is therefore a feature of the estimated demand system rather than independent evidence of curvature in the cost mapping.

The net column makes the cross-border reallocation explicit. Under the two cost-rise scenarios, foreign rivals gain more revenue than Finnish exporters lose, so total inside-good revenue rises (a net €143.8 million under No-kiky and €452.4 million under the four-day week), with the entire increase accruing to foreign competitors. The Finnish policy direction is the mirror image: under the second counterfactual *Sipilä*, Finnish firms gain €447.0 million, whereas rivals lose €567.8 million, a net €120.8 million loss of producer surplus competed away to importing consumers. The policy redistributes export revenue across borders rather than raising it for producers, reallocating sales toward or away from foreign competitors (Mattoo, Mishra, and Subramanian 2017). This result depends on the flexibility of the demand system. A less flexible specification, such as a plain logit model, would mechanically restrict cost pass-through to $1-s$ for a single-product seller and therefore rule out super-pass-through. However, the cross-scenario stability would remain, because it reflects local linearity around the same baseline rather than the particular level of pass-through.

The Naive variant in Table 8 replaces the markup-corrected labor-cost term, $(R_{ft}/C_{ft}) \cdot s_L^R$, with the raw labor revenue share, s_L^R. Across all three scenarios, the corresponding Finnish ΔR values are approximately 81 percent of the markup-corrected estimates. This implies a revenue-weighted firm-year R/C ratio of about 1.24 in the simulation. The implied ratio is below the simple per-HS4 mean $R/C \approx 1.83$ reported in Table A.4, because high-R/C markets such as kraft paper ($R/C \approx 2.75$) receive relatively small revenue weight in the aggregate.

Revenue reallocation. Table 9 reports the revenue effects of the three counterfactuals for non-Finnish exporters in the markup-corrected specification. In aggregate, reversing the realized internal devaluation policy *No Kiky* increases rival revenue by €382.9 million, whereas implementing the more ambitious internal devaluation policy *Sipilä* reduces rival revenue

Table 8: Aggregate counterfactual revenue effects of three labor-cost scenarios on Finnish exporters and non-Finnish exporters in the same destination markets, 2017–2020.

Scenario	Variant	Finland ΔR (M€)	Rivals ΔR (M€)	Net Δ (M€)	Pass-thru (mean)
No kiky	R/C-corr.	-239.0 (-0.63%)	+382.9	+143.8	1.18
	Naive	-194.4 (-0.51%)	+329.6	+135.2	1.18
Sipilä	R/C-corr.	+447.0 (+1.18%)	-567.8	-120.8	1.18
	Naive	+359.7 (+0.95%)	-463.2	-103.5	1.18
Four-day week	R/C-corr.	-2,394.4 (-6.33%)	+2,846.8	+452.4	1.19
	Naive	-1,938.2 (-5.13%)	+2,309.1	+370.9	1.19

Notes: Counterfactual revenue changes in millions of euros, summed across the nine HS4 markets and the four treated years (2017–2020). Finland ΔR aggregates the Finnish exporter revenue change; the percentage in parentheses below each Finland value normalizes by the 2017–2020 Finnish baseline of 37,799 M€in these nine markets from BACI. Rivals ΔR aggregates the non-Finnish exporter revenue change in the same destination markets, reflecting Bertrand-Nash price adjustments. Net Δ is the sum. Pass-thru is the simple mean across the nine HS4 markets of the per-HS4 mean pass-through on treated observations. *Variants*: Naive applies $d\ln c = s_L^R \cdot d\ln_{\text{total}}$; R/C-corrected applies $d\ln c = (R_{ft}/C_{ft}) \cdot s_L^R \cdot d\ln_{\text{total}}$, where R_{ft}/C_{ft} is the BLP-recovered firm-year revenue-to-cost ratio; this variant is the paper's preferred specification. Revenue figures are on FOB basis (deflated from the BLP solver's landed CIF+tariff prices by a per-observation transport-and-tariff wedge), matching BACI's reporting convention. USD→EUR at year-averaged rate of 0.875.

by €567.8 million. The four-day-work-week counterfactual produces a substantially larger reallocation toward foreign exporters, increasing their combined revenue by €2.8 billion. The countries with the largest absolute responses include major Finnish competitors in pulp and forest-products markets, such as Sweden, Brazil, Russia, Canada, and Indonesia, as well as trade and re-export hubs such as Hong Kong and the Netherlands. Under the *No Kiky* scenario, the Netherlands, Hong Kong, and Japan record the largest revenue gains, suggesting that exporters participating broadly across affected destination markets can benefit disproportionately even from a relatively small Finnish cost increase.

Among the major producer countries, the revenue changes are generally small relative to their baseline revenues. Under the more ambitious devaluation policy *Sipilä* scenario, for example, the percentage changes for the principal forest-products competitors remain below 0.2 percent in absolute value, whereas the corresponding effects under the four-day-work-week scenario are typically below 1 percent. Consequently, the ranking of countries by absolute revenue change largely reflects their initial trade volumes and exposure to the affected markets. However, the responses are not fully proportional. Hong Kong is an outlier, with revenue changes of 65.1 million under *No Kiky* and -69.9 million under more ambitious policy, despite its small baseline revenue. Moreover, some exporters, including Sweden and Germany, respond differently from the aggregate rival pattern in the smaller counterfactuals. These differences illustrate the heterogeneous substitution patterns permitted by the estimated demand system.

Heterogeneous impact of counterfactual policies. The counterfactual shocks are applied annually and Finnish trade volumes vary over 2017–2020. The year-on-year decomposition in Table A.6 shows that under the No-kiky scenario the estimated effect shrinks sharply in 2020 (about −€8.6 million, compared with effects between −€67.0 million and −€81.8 million in 2017–2019). This decline reflects the partial unwinding of the realized devaluation policy in 2020. By contrast, the more ambitious devaluation policy and four-day-week shocks are time-invariant in logarithmic terms, so their year-to-year variation reflects changes in Finnish export volumes rather than changes in policy intensity.

We report the per-market revenue effects in Table A.5. Chemical pulp accounts for approximately 39 percent of the Finnish counterfactual revenue change in every scenario, with sawn wood as the second largest contributor. The remaining revenue effects are distributed across the different paper grades, plywood, and stainless steel. Table A.4 shows the corresponding per-HS4 mean pass-through estimates, which range from approximately 1.05 to 1.32 and are above one in all markets. The paper grades cluster between 1.05–1.11, whereas wood, chemical pulp, and stainless steel fall between 1.17–1.32.

7.3 Scope of the Counterfactuals

This subsection states the assumptions behind the counterfactuals and the scope of the estimated ΔR. Its main message is that two of the channels we

Table 9: Non-Finnish exporter revenue changes under each counterfactual (R/C-corrected variant).

Country	Baseline rev. (M€)	No kiky (M€)	Sipilä (M€)	Four-day week (M€)
Sweden	38,493	-5.1 (-0.01%)	-13.6 (-0.04%)	+335.3 (+0.87%)
Brazil	36,378	+24.8 (+0.07%)	-48.8 (-0.13%)	+272.7 (+0.75%)
Russia	30,722	+16.2 (+0.05%)	-36.9 (-0.12%)	+231.7 (+0.75%)
Canada	51,778	+17.7 (+0.03%)	-37.0 (-0.07%)	+226.9 (+0.44%)
Germany	31,475	-0.4 (-0.00%)	+10.0 (+0.03%)	+180.8 (+0.57%)
Indonesia	31,696	+13.4 (+0.04%)	-26.9 (-0.08%)	+149.1 (+0.47%)
United States	45,587	+25.7 (+0.06%)	-34.6 (-0.08%)	+139.8 (+0.31%)
Chile	13,989	+11.0 (+0.08%)	-21.2 (-0.15%)	+119.0 (+0.85%)
Portugal	8,189	+7.5 (+0.09%)	-15.4 (-0.19%)	+108.1 (+1.32%)
Austria	14,304	+7.3 (+0.05%)	-13.9 (-0.10%)	+97.2 (+0.68%)
Hong Kong	565	+65.1 (+11.52%)	-69.9 (-12.37%)	+1.1 (+0.19%)
Netherlands	10,787	+67.2 (+0.62%)	-57.9 (-0.54%)	+64.7 (+0.60%)
Japan	7,983	+43.6 (+0.55%)	-46.4 (-0.58%)	+18.9 (+0.24%)
China	27,690	+25.6 (+0.09%)	-25.5 (-0.09%)	+37.8 (+0.14%)
Malaysia	7,944	+16.9 (+0.21%)	-17.9 (-0.23%)	+11.9 (+0.15%)
Other rivals (122)	154,463	+46.4 (+0.03%)	-112.0 (-0.07%)	+851.6 (+0.55%)
Total rivals	512,044	+382.9 (+0.07%)	-567.8 (-0.11%)	+2,846.8 (+0.56%)

Notes: Revenue changes in millions of euros for non-Finnish exporters into the destination markets. Each counterfactual (R/C-corrected variant), aggregated across 9 HS4 product markets and four treated years. The percentage in parentheses below each euro value is that scenario's revenue change as a percent of the country's own baseline revenue. Listed rivals are the union of the ten largest by absolute revenue change under each of the three counterfactuals, ordered by their largest absolute change in any scenario. *Baseline rev.*: total non-Finnish exporter revenue in the affected destination markets before the shock. Revenue figures are on FOB basis. USD→EUR at year-averaged rate of 0.875.

hold fixed, input prices and entry and investment, bias the estimated $|\Delta R|$ toward zero, whereas the omitted destination-market domestic-supply response probably has the opposite sign. The remaining points clarify what ΔR does and does not measure and state one identification assumption.

Our counterfactual simulations are all partial equilibrium exercises. We re-solve the product-market pricing game under the labor-cost shock and do not model wage-setting, employment, or aggregate feedback. Two features of the setting make this appropriate. The pact was a negotiated statutory change rather than a market wage, so firms take it as given and there is no endogenous wage response to recover. Furthermore, it targeted unit labor costs that were high relative to productivity in a slack labor market in which the cut is not bid back through higher wages by construction.[14] Therefore, the policy passes to input prices without a first-order labor-market offset.

First, the Finnish ΔR is the relevant proxy for how the size of the export market ('cake') changes. Rival revenue movements in the same destinations are largely internal distribution among foreign exporters competing for the same buyers. The exception is stainless steel, where material imports cross borders and the redistribution is not purely within destinations. Forestry has analogous cross-border input flows in Russian roundwood and wood chips, although the channel was already in retreat during the kiky window and collapsed after 2022.[15]

Second, the non-Finnish exporters in our inside-good set reprice strategically (Section 5.1). What we omit is the response of destination-market domestic producers—local supply in each destination—which the demand estimation captures only through the outside option. In the counterfactuals, the local suppliers do not re-price, because we lack data on these producers. Foreign-rival prices in our simulation move in the same direction as Finland's in every scenario, rising alongside Finnish prices under cost-rise shocks and falling alongside them under cost-cut shocks due to rival strategic responses. To the extent that destination-market domestic supply responds with the same sign, the omitted responses bias $|\Delta R|$ away from zero, not toward it. A substitute that moves with Finland's price partly offsets Finland's relative price change and dampens its share response, so holding it fixed overstates the reallocation. With mixed-logit demand this co-movement is not guaranteed by primitives. Quint (2014)

14. The *No kiky* and *Sipilä* scenarios should be interpreted as labor market agreements, whereas for the four-day work week wages can be assumed to be sticky.

15. Russian roundwood and wood chip (HS 4401 and 4403) flows to Finland in BACI peaked at \$891M in 2008 and fell after Russia's 2007 export tariffs, ran at roughly \$280–400M per year over 2017–2019, then collapsed to \$65M in 2022 after Russia's invasion of Ukraine and Finland's import boycott.

shows that discrete-choice pricing is a game of strategic complements in the Bulow, Geanakoplos, and Klemperer (1985) sense under log-concave product-level preferences, but that this does not extend to BLP-style random coefficients like ours. Our demand system can therefore deliver either strategic complements or strategic substitutes. Foreign rivals co-move with Finland in every scenario, so the same-sign extrapolation to unmodeled domestic supply is empirical, not a theorem. The bias is in any case bounded by Finland's small share in each destination.

Third, we treat firms as price-takers in their input markets, including labor: each takes the pact's statutory cost changes as given, and the prices of non-labor variable inputs—materials and energy—do not respond to the counterfactual. This assumption is probably partly offset, because the policies reach the input suppliers themselves: where a firm sources from other Finnish firms also subject to the pact—not necessarily among our nine export industries—those suppliers' costs move and their prices with them, so the buyer's input prices would not in fact hold fixed.

Finnish chemical pulp is the one such linkage inside our sample: it is itself an input to the paper and packaging markets we study, so when the counterfactual cost shock moves pulp costs those downstream markets see a correlated, same-signed input-cost movement that our market-by-market setup does not pass through.

Fourth, entry, exit, and investments are exogenous in the counterfactuals. We expect this to understate ΔR: A decrease in labor costs increases the return on capital investments, which increases the optimal level of investment and thereby also ΔR. The opposite is likely to hold for an increase in labor costs.

Fifth, the BLP-recovered $\bar{\mathcal{M}}_{ft}$ is identified on export markets, but $s^{R}_{L,ft}$ is a firm-wide accounting object: we assume our export-market weighted-average markup matches domestic-market markups for affected firms.

Finally, ΔR is a cross-border producer-revenue measure and therefore does not identify who ultimately gains or loses within each economy (Ignatenko, Lashkaripour, Macedoni, and Simonovska 2025). In particular, the analysis cannot determine how changes in export revenue are distributed between firm owners and workers, nor can the current framework estimate the employment effects of the studied counterfactual policies. In addition, we do not characterize efficiency effects of the devaluation policy or the other counterfactuals.[16] Subsection 7.4 provides back-of-the-envelope employment calculations, but these calculations rely primarily on publicly available industry-level aggregates rather than on the struc-

16. Firms with higher labor cost shares are likely to benefit more from the devaluation policy.

tural estimates of this article.

7.4 Back of the Envelope Employment Calculations

Our empirical framework does not allow us to estimate the effect of increased revenue on firm employment. Translating the estimated ΔR into employment terms using public figures nonetheless puts the magnitudes in perspective. Forest industries (forestry, wood products, pulp and paper) employ about 38,400 people directly in Finland and about 82,800 across the broader value chain (Finnish Forest Industries Federation 2024).[17] A forest-industry production value on the order of €25 billion implies revenue per direct employee of about €653,000 (Statistics Finland 2025). Dividing the cumulative export-revenue change by this figure gives direct full-time equivalent (FTE) years, which we scale by the value-chain-to-direct employment ratio of about 2.16 for the broader total. The €239.0 million four-year cumulative export-revenue loss from reversing the realized kiky thus corresponds to roughly $239.0/0.653 \approx 370$ direct FTE-years, or about 800 value-chain FTE-years, in the nine markets we cover. The €2.4 billion cumulative loss under the four-day work week scales these to roughly $2{,}394.4/0.653 \approx 3{,}700$ and about 8,000 FTE-years respectively.

8 Conclusion

This study develops an empirical framework that combines demand estimation tools from industrial organization with a sufficient-statistics approach to cost incidence. The framework inverts the Hall–De Loecker–Warzynski identity, $\beta_v = \bar{\mathcal{M}} \cdot s_v^R$: rather than estimating β_v from a production function, we recover $\bar{\mathcal{M}}$ from a BLP demand inversion and combine it with observed revenue shares. The demand model provides a non-CES, non-gravity, firm-level structural estimate using bilateral customs data. This flexibility allows us to jointly identify heterogeneous price elasticities, distance elasticities, and pass-through to export prices—margins that are restricted by construction in the standard CES iceberg-cost gravity framework. We apply the framework to Finland's 2017–2020 internal devaluation policy across nine HS4 export markets to study how the policy affected Finnish exporters and their foreign competitors.

The counterfactual results show that labor-cost policies can have economically meaningful effects on export-market outcomes, but that the

17. The broader forest-sector definition, including forestry, places direct employment near 64,000 in 2017 (Finnish Forest Association 2018).

magnitude and distribution of these effects depend on equilibrium pricing responses. Abolishing the realized devaluation policy would have reduced Finnish export revenue by €239.0 million. By contrast, the more ambitious devaluation policy counterfactual would have increased Finnish export revenue by €447.0 million, whereas a four-day-week scenario would have reduced export revenue by €2.4 billion. These effects are shaped by super-pass-through: when Finnish firms face higher costs, foreign rivals expand revenues by more than Finnish firms lose, generating net revenue gains of €143.8 million and €452.4 million in the relevant cost-increase scenarios. These results imply that policies affecting domestic input costs can generate downstream price and market-share effects that differ substantially from those implied by logit Lerner-rule intuition.

More broadly, the results quantify the international spillovers associated with industrial policy designed to improve the competitiveness of domestic firms. Internal devaluation can increase the export performance of domestic firms, but foreign competitors' strategic pricing responses mediate its effects. As a result, policies that improve the relative cost position of domestic exporters may redistribute revenue across firms and countries without necessarily increasing the overall size of the market. This mechanism is beggar-thy-neighbor in a precise, redistributive sense: by cutting costs a country steals business from rival producers abroad, yet the lower export prices it generates are passed through to the consumers who buy those goods. It beggars thy neighbor's producers and enriches thy neighbor's consumers.

Because the rents rivals lose are competed away to buyers rather than destroyed by a trade-volume wedge, the cross-border effect is distributional, not a deadweight loss: its incidence falls on rival producers and on the domestic workers who fund the cut. The welfare consequences of internal devaluation thus turn not only on domestic cost savings, but on how these equilibrium price responses redistribute surplus across firms, workers, and consumers worldwide. Seen this way, internal devaluation is privately rational for the adopting country, whose firms gain, yet negative-sum for producers worldwide: rivals lose more than domestic firms gain, and the difference is competed away to consumers.

References

Albrecht, Brian C., and Ryan A. Decker. 2026. "Markups and Business Dynamism across Industries." *International Journal of Industrial Organization* (January 18, 2026): 103260.

Anderson, James E., and Eric van Wincoop. 2003. "Gravity with Gravitas: A Solution to the Border Puzzle." *American Economic Review* 93 (1): 170–192.

Backus, Matthew, Christopher Conlon, and Michael Sinkinson. 2020. "Theory and Measurement of Common Ownership." *AEA Papers and Proceedings* 110 (May): 557–560.

———. 2021. *Common Ownership and Competition in the Ready-to-Eat Cereal Industry.* Working Paper 28350. National Bureau of Economic Research, January.

Benkard, C. Lanier, Nathan H. Miller, and Ali Yurukoglu. 2025a. *The Rise of Market Power and the Macroeconomic Implications: Comment.* Working Paper 34513. National Bureau of Economic Research, November.

———. 2025b. "Rejoinder to 'The Rise of Market Power and the Macroeconomic Implications: Reply to Benkard, Miller, and Yurukoglu (2025)'." October.

Bentolila, Samuel, Pierre Cahuc, Juan J. Dolado, and Thomas Le Barbanchon. 2012. "Two-Tier Labour Markets in the Great Recession: France versus Spain." *Economic Journal* 122 (562): F155–F187.

Benzarti, Youssef, and Jarkko Harju. 2021. "Using Payroll Tax Variation to Unpack the Black Box of Firm-Level Production." *Journal of the European Economic Association* (jvab010 2021).

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841–890.

Berry, Steven T., and Philip A. Haile. 2014. "Identification in Differentiated Products Markets Using Market Level Data." *Econometrica* 82 (5): 1749–1797.

Blanchard, Olivier, Mark Griffiths, and Bertrand Gruss. 2013. "Boom, Bust, Recovery: Forensics of the Latvia Crisis." *Brookings Papers on Economic Activity* 2013 (2): 325–388.

Böckerman, Petri, Salla Kalin, Toni Juuti, Merja Kauhanen, and Timo Suhonen. 2025. "Miksi järjestäytymisaste laskee ja mitä siitä seuraa?" *Talous & Yhteiskunta*, no. 4.

Bond, Steve, Arshia Hashemi, Greg Kaplan, and Piotr Zoch. 2021. "Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data." *Journal of Monetary Economics* 121 (July 1, 2021): 1–14.

Bulow, Jeremy I., John D. Geanakoplos, and Paul D. Klemperer. 1985. "Multimarket Oligopoly: Strategic Substitutes and Complements." *Journal of Political Economy* 93 (3): 488–511.

Caballero, Ricardo J, Emmanuel Farhi, and Pierre-Olivier Gourinchas. 2021. "Global Imbalances and Policy Wars at the Zero Lower Bound." *The Review of Economic Studies* 88, no. 6 (November 1, 2021): 2570–2621.

Cahuc, Pierre, Stéphane Carcillo, and Thomas Le Barbanchon. 2019. "The Effectiveness of Hiring Credits." *The Review of Economic Studies* 86, no. 2 (March 1, 2019): 593–626.

Conlon, Christopher, and Jeff Gortmaker. 2020. "Best practices for differentiated products demand estimation with PyBLP." *The RAND Journal of Economics* 51 (4): 1108–1161.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *The Quarterly Journal of Economics* 135, no. 2 (May 1, 2020): 561–644.

———. 2025. "The Rise of Market Power and the Macroeconomic Implications: Reply to Benkard, Miller, and Yurukoglu (2025)." November.

De Loecker, Jan, and Frederic Warzynski. 2012. "Markups and Firm-Level Export Status." *American Economic Review* 102, no. 6 (May): 2437–2471.

Disdier, Anne-Célia, and Keith Head. 2008. "The Puzzling Persistence of the Distance Effect on Bilateral Trade." *Review of Economics and Statistics* 90 (1): 37–48.

Döpper, Hendrik, Alexander MacKay, Nathan H. Miller, and Joel Stiebale. 2025. "Rising Markups and the Role of Consumer Preferences." *Journal of Political Economy* 133, no. 8 (August 1, 2025): 2462–2505.

Doraszelski, Ulrich, and Jordi Jaumandreu. 2021. *Reexamining the De Loecker & Warzynski (2012) method for estimating markups.* CEPR Discussion Paper 16027. Paris & London: CEPR Press, April 1, 2021.

Dustmann, Christian, Bernd Fitzenberger, Uta Schönberg, and Alexandra Spitz-Oener. 2014. “From Sick Man of Europe to Economic Superstar: Germany’s Resurgent Economy.” *Journal of Economic Perspectives* 28 (1): 167–188.

Fajgelbaum, Pablo D, Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal. 2020. “The Return to Protectionism.” *The Quarterly Journal of Economics* 135, no. 1 (February 1, 2020): 1–55.

Farhi, Emmanuel, Gita Gopinath, and Oleg Itskhoki. 2014. “Fiscal Devaluations.” *The Review of Economic Studies* 81, no. 2 (April 1, 2014): 725–760.

Fiallos, A., A. Liberatore, and S. Cassimon. 2024. *CIF/FOB Margins: Insights on Global Transport and Insurance Costs of Merchandise Trade.* Working Paper 2024/05. OECD Publishing.

Gandhi, Amit, and Jean-François Houde. 2019. *Measuring Substitution Patterns in Differentiated-Products Industries.* Working Paper 26375. National Bureau of Economic Research.

Gaulier, Guillaume, and Soledad Zignago. 2010. *BACI: International Trade Database at the Product-Level. The 1994–2007 Version.* Working Paper 2010–23. CEPII.

Geis, André. 2023. *Wage Indexation and International Competitiveness in Belgium: An Uneasy Coexistence: Belgium.* International Monetary Fund.

Grigolon, Laura, and Frank Verboven. 2014. “Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation.” *The Review of Economics and Statistics* 96 (5): 916–935.

Hall, Robert E. 1988. “The Relation between Price and Marginal Cost in U.S. Industry.” *Journal of Political Economy* 96 (5): 921–947.

Head, Keith, and Thierry Mayer. 2014. “Gravity Equations: Workhorse, Toolkit, and Cookbook.” In *Handbook of International Economics,* 4:131–195. Elsevier.

Ignatenko, Anna, Ahmad Lashkaripour, Luca Macedoni, and Ina Simonovska. 2025. “Making America great again? The economic impacts of Liberation Day tariffs.” *Journal of International Economics* 157 (September 1, 2025): 104138.

Illanes, Gaston, Vivek Bhattacharya, David Stillerman, Avner Kreps, and Jose Salas. 2024. *A Large-Scale Evaluation of Merger Simulations,* 5177604, Rochester, NY, December 1, 2024.

Jonker-Hoffrén, Paul. 2019. "Finland: Goodbye centralised bargaining? The emergence of a new industrial bargaining regime." *Collective bargaining in Europe: towards an endgame* 1:197–215.

Juhász, Réka, Nathan Lane, and Dani Rodrik. 2024. "The New Economics of Industrial Policy." *Annual Review of Economics* 16:213–242.

Kiander, Jaakko, Pekka Sauramo, and Hannu Tanninen. 2011. "Finnish incomes policy as corporatist political exchange: development of social capital and the social wage." *Transfer: European Review of Labour and Research* 17 (4): 515–531.

Larch, Mario, Serge Shikher, and Yoto V. Yotov. 2025. "Estimating Gravity Equations: Theory Implications, Econometric Developments, and Practical Recommendations." *Review of International Economics* 33 (5): 1066–1092.

MacKay, Alexander, and Nathan H. Miller. 2025. "Estimating Models of Supply and Demand: Instruments and Covariance Restrictions." *American Economic Journal: Microeconomics* 17, no. 1 (February): 238–281.

Mankiw, N. Gregory, and Michael D. Whinston. 1986. "Free Entry and Social Inefficiency." *The RAND Journal of Economics* 17 (1): 48–58.

Mattoo, Aaditya, Prachi Mishra, and Arvind Subramanian. 2017. "Beggar-Thy-Neighbor Effects of Exchange Rates: A Study of the Renminbi." *American Economic Journal: Economic Policy* 9, no. 4 (November): 344–366.

Miravete, Eugenio J., Katja Seim, and Jeff Thurk. 2023. *Elasticity and Curvature of Discrete Choice Demand Models.* CEPR Discussion Paper 18310. CEPR Press.

Morrow, W. Ross, and Steven J. Skerlos. 2011. "Fixed-Point Approaches to Computing Bertrand-Nash Equilibrium Prices Under Mixed-Logit Demand." *Operations Research* 59 (2): 328–345.

Mrázová, Monika, and J. Peter Neary. 2017. "Not So Demanding: Demand Structure and Firm Behavior." *American Economic Review* 107, no. 12 (December 1, 2017): 3835–3874.

Müller, Karsten, Chenzi Xu, Mohamed Lehbib, and Ziliang Chen. 2025. *The Global Macro Database: A New International Macroeconomic Dataset.* Working Paper 33714. National Bureau of Economic Research, April.

Novy, Dennis. 2013. "International Trade without CES: Estimating Translog Gravity." *Journal of International Economics* 89 (2): 271–282.

Orr, Scott. 2022. "Within-Firm Productivity Dispersion: Estimates and Implications." *Journal of Political Economy* 130, no. 11 (November): 2771–2828.

Paukkeri, Tuuli, Terhi Ravaska, and Marja Riihelä. 2024. "Examining inequalities: from labour markets to social outcomes in Finland." *Fiscal Studies* 45 (2): 225–241.

Piveteau, Paul, and Gabriel Smagghue. 2023. "Foreign Competition along the Quality Ladder." *The Economic Journal* (December 7, 2023): uead107.

Quint, Daniel. 2014. "Imperfect competition with complements and substitutes." *Journal of Economic Theory* 152 (July 1, 2014): 266–290.

Reynaerts, Jo, Ravi Varadhan, and John C. Nash. 2012. *Enhancing the Convergence Properties of the BLP (1995) Contraction Mapping.*

Smith, Adam. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations.* London: W. Strahan / T. Cadell.

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

Train, Kenneth E. 2009. *Discrete Choice Methods with Simulation.* Cambridge University Press, June 30, 2009.

Valmari, Nelli. 2023. "Estimating Production Functions of Multiproduct Firms." *The Review of Economic Studies* 90, no. 6 (November 8, 2023): 3315–3342.

Van Biesebroeck, Johannes, and Hongsong Zhang. 2025. "Lessons on industrial policy from industrial organization: Introduction to the special issue." *International Journal of Industrial Organization,* 103243.

Varadhan, Ravi, and Christophe Roland. 2008. "Simple and Globally Convergent Methods for Accelerating the Convergence of Any EM Algorithm." *Scandinavian Journal of Statistics* 35 (2): 335–353.

Zhang, Linqi. 2023. "Identification and Estimation of Market Size in Discrete Choice Demand Models."

A Appendix

A.1 Notation

Table A.1: Notation

Symbol	Description
Indices	
i	Consumer index (set $\mathcal{I}_{dt}$)
j	Product index, firm–HS6 combination (f, k) (set $\mathcal{J}_{dt}$)
d, t	Destination country and year (sets $\mathcal{D}, \mathcal{T}$)
k	HS6 product code within a 4-digit HS market (set $\mathcal{K}$)
f	Firm (exporter) index (set $\mathcal{F}$)
(d, t)	Market definition: destination–year pair within an HS4 market
Demand model	
p_{jt}, q_{jt}, s_{jt}	Unit price, quantity, market share of product j in market t
δ_{jt}	Mean utility of product j in market t
α_i	Consumer-specific price coefficient (depends on market via per-capita GDP_{dt})
β	Mean taste parameters for exogenous characteristics
$\pi, \pi_{gdp}, \sigma_\alpha$	Intercept, GDP gradient, and standard deviation of log price-sensitivity
ϕ, Ψ_t	Market-potential scaling parameter; observed inside-good sales
ξ_{jt}	Demand shock (unobserved product quality)
ϵ_{ijt}	Type-I Extreme Value taste shock
Supply model and tariff/margin	
c_{jt}	Producer-side (FOB) marginal cost of product j in market t
τ_{jt}, m_{jt}	Ad valorem tariff; transportation-and-insurance margin
$\mathcal{M}_{jt}$	Product-level price-cost markup, p_{jt}/c_{jt}
$\bar{\mathcal{M}}_{ft}$	Firm-year markup, cost-weighted within-firm-year mean of $\mathcal{M}_{jt}$
η_{jt}	Absolute price-cost markup, $p_{jt} - c_{jt}$
$\boldsymbol{\Omega}$	Ownership matrix of inside goods
Cost-side incidence and production primitives	
R_{ft}	Firm revenue
$C_{ft}(\mathbf{w}, Q)$	Total cost function
$c_{ft}(\mathbf{w})$	Unit cost function (under CRS)

(continued on next page)

Symbol	Description
$F_{ft}(\cdot)$	Production function
V	Set of variable inputs
$w_{v,ft}, x_{v,ft}$	Price and optimal quantity of variable input v
$s^R_{v,ft}, s^C_{v,ft}$	Revenue share and cost share of input v
$\beta_{v,ft}$	Output elasticity of variable input v
γ_{ft}	Returns-to-scale parameter, $\gamma_{ft} = \sum_{v\in V} \beta_{v,ft}$
Q_{ft}, K_{ft}	Output and predetermined capital stock
L_{ft}, M_{ft}	Labor and materials inputs
P^L_{ft}, P^M_{ft}	Price of labor; price of materials

Notes: In the demand model, market-level variables are indexed by destination–year pairs (d, t); we often suppress d for notational simplicity. Each product $j \in \mathcal{J}_{dt}$ is a firm–HS6 combination (f, k). The cost-side identity in Section 5.2 indexes objects at the firm–year level (f, t).

A.2 Key Variable Definitions

Firm identity. We define Finnish firms through the ownership structure. This ownership definition aggregates all parent-company subsidiaries into a single entity and leaves independent firms as they are in the data. Parent company–subsidiary links are missing mainly before 2003, but we fill these missing links using data from later years.

Labor Cost. We measure the firm labor cost as the wage bill plus other personnel expenses.[18] Other personnel expenses include pension expenses, social security contributions, statutory and voluntary personal insurance contributions. We take the wage sum and personnel expenses from Statistics Finland balance sheet data.

Gross Value of Production. Gross value of production is based on Statistics Finland balance sheet data and it measures the actual production of an establishment. In addition to turnover, production activity includes all production-related income, including production for own use and production delivered to the enterprise's other establishments. Purchases of merchandise are deducted from income so that production activity includes only the margin generated by the sale of merchandise.

18. However, the incidence of employer social insurance costs is based on the wage sum or gross wages to avoid double accounting.

Quantity. The relevant quantity measure in all demand analyses is the weight of exported product (HS6 category) in metric tonnes.

Unit Price. We define unit price for each product as the ratio of sales value in US dollars and quantity in metric tons for each exporter-product-destination-year group. All prices are measured in dollars and we use nominal prices in demand estimation. Because our BACI data from Gaulier and Zignago (2010) and Finnish Customs are reported at FOB values, we convert our prices to CIF prices using the tariff and transportation cost data.

Distance. The BACI database reports the annual population-weighted distance in kilometers for each importer-exporter. In the demand model, we use distance to calculate distance elasticities.

A.3 From CIF to FOB marginal cost

The producer-side (FOB) marginal cost we use in the cost-side analysis follows from the producer-side Bertrand–Nash first-order condition. Writing $p_{jt}^{\text{CIF}} = p_{jt}^{\text{FOB}}(1+\tau_{jt})(1+m_{jt})$, the markup inversion of Equation (8) on CIF prices delivers $c_{jt}^{\text{CIF}} = p_{jt}^{\text{CIF}} - \eta_{jt}^{\text{CIF}}$, in which $\eta_{jt} \equiv p_{jt} - c_{jt}$ denotes the per-unit absolute markup. Under the multiplicative wedge, the producer's true first-order condition on FOB prices and FOB costs has $\eta_{jt}^{\text{FOB}} = \eta_{jt}^{\text{CIF}}/[(1+\tau_{jt})(1+m_{jt})]$. Hence $c_{jt}^{\text{FOB}} = p_{jt}^{\text{FOB}} - \eta_{jt}^{\text{FOB}} = c_{jt}^{\text{CIF}}/[(1+\tau_{jt})(1+m_{jt})]$.

A.4 HS4 Categories

Table A.2 lists the nine 4-digit HS product markets that define our analysis sample, together with short product descriptions.

A.5 Additional Tables

Table A.3 reports descriptive statistics for the sample of Finnish exporting firms used in the production-side and counterfactual analysis.

A.6 Demand Model Results

Figures A.1–A.7 report the demand-side distributional outputs for Finnish exporters at the HS4 level: own-price, aggregate, and distance elasticities, marginal costs, Lerner indices, and their year-on-year evolution.

Figure A.1: Own-Price Elasticities by HS4 Market

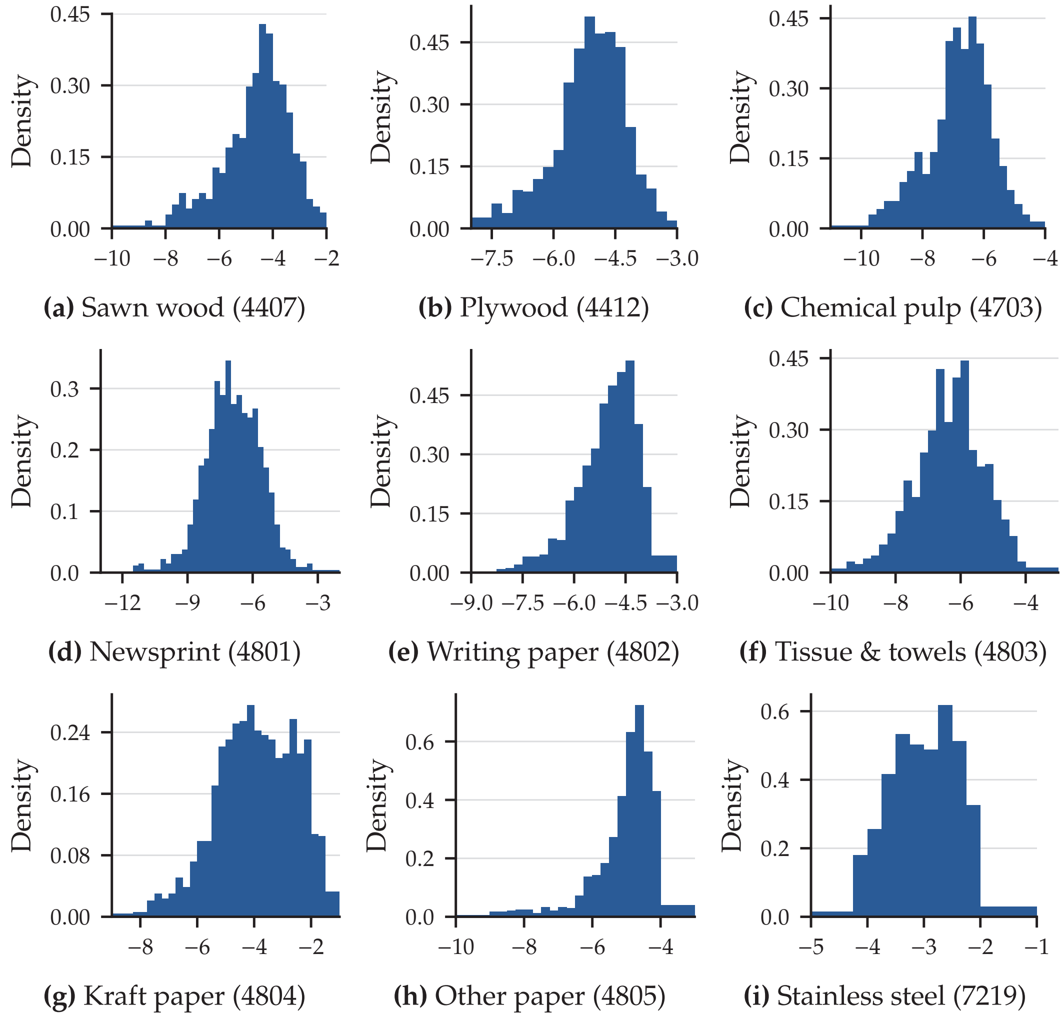

(a) Sawn wood (4407) **(b)** Plywood (4412) **(c)** Chemical pulp (4703)

(d) Newsprint (4801) **(e)** Writing paper (4802) **(f)** Tissue & towels (4803)

(g) Kraft paper (4804) **(h)** Other paper (4805) **(i)** Stainless steel (7219)

Notes: Per-observation own-price elasticities for Finnish exporters, recovered from the BLP demand estimates of Section 4 over the full sample (1999–2021). Each panel is a histogram for one HS4 product market; horizontal axis is the own-price elasticity, vertical axis is density.

Figure A.2: Aggregate Elasticities by HS4 Market

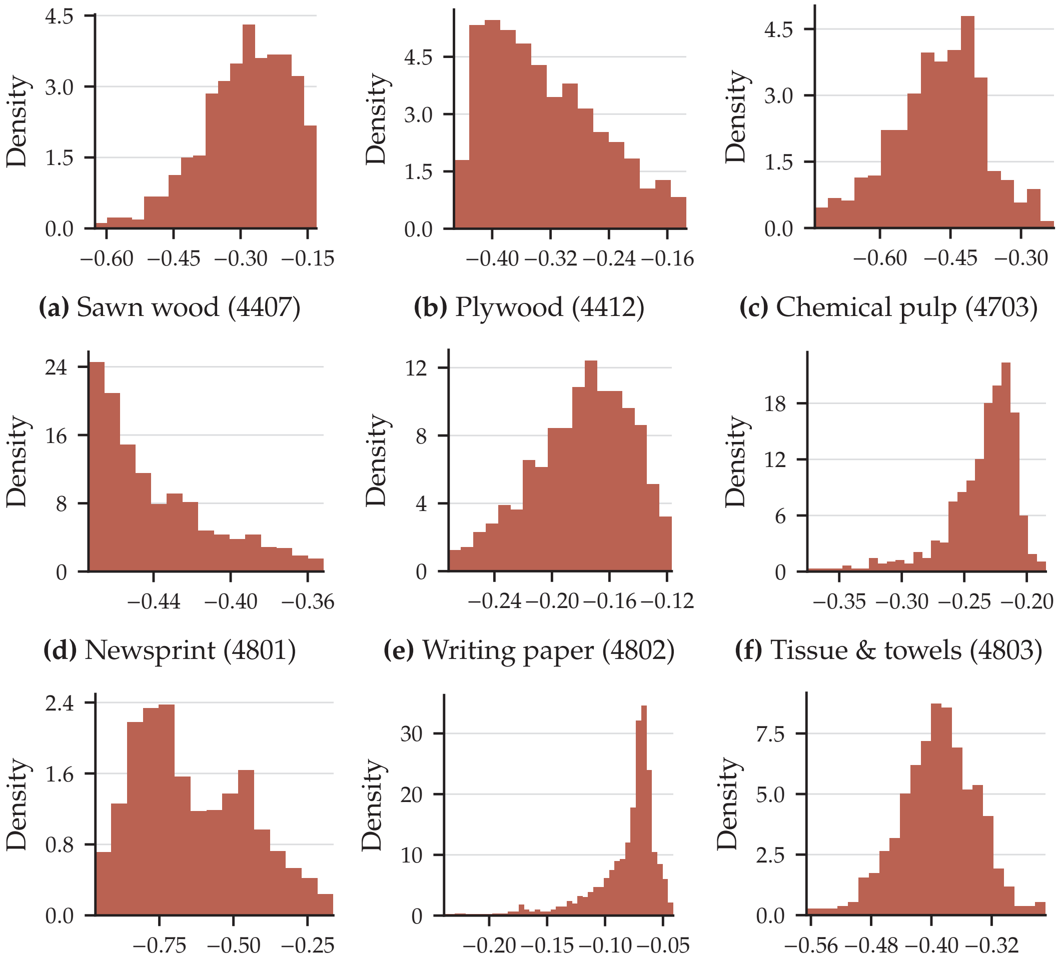

(a) Sawn wood (4407) **(b)** Plywood (4412) **(c)** Chemical pulp (4703)

(d) Newsprint (4801) **(e)** Writing paper (4802) **(f)** Tissue & towels (4803)

(g) Kraft paper (4804) **(h)** Other paper (4805) **(i)** Stainless steel (7219)

Notes: Per-observation Finland-aggregate own-price elasticities (sum of the per-firm elasticities for Finnish exporters in each destination-year), recovered from the BLP demand estimates over 1999–2021. Each panel is a histogram for one HS4 product market.

Figure A.3: Distance Elasticities by HS4 Market

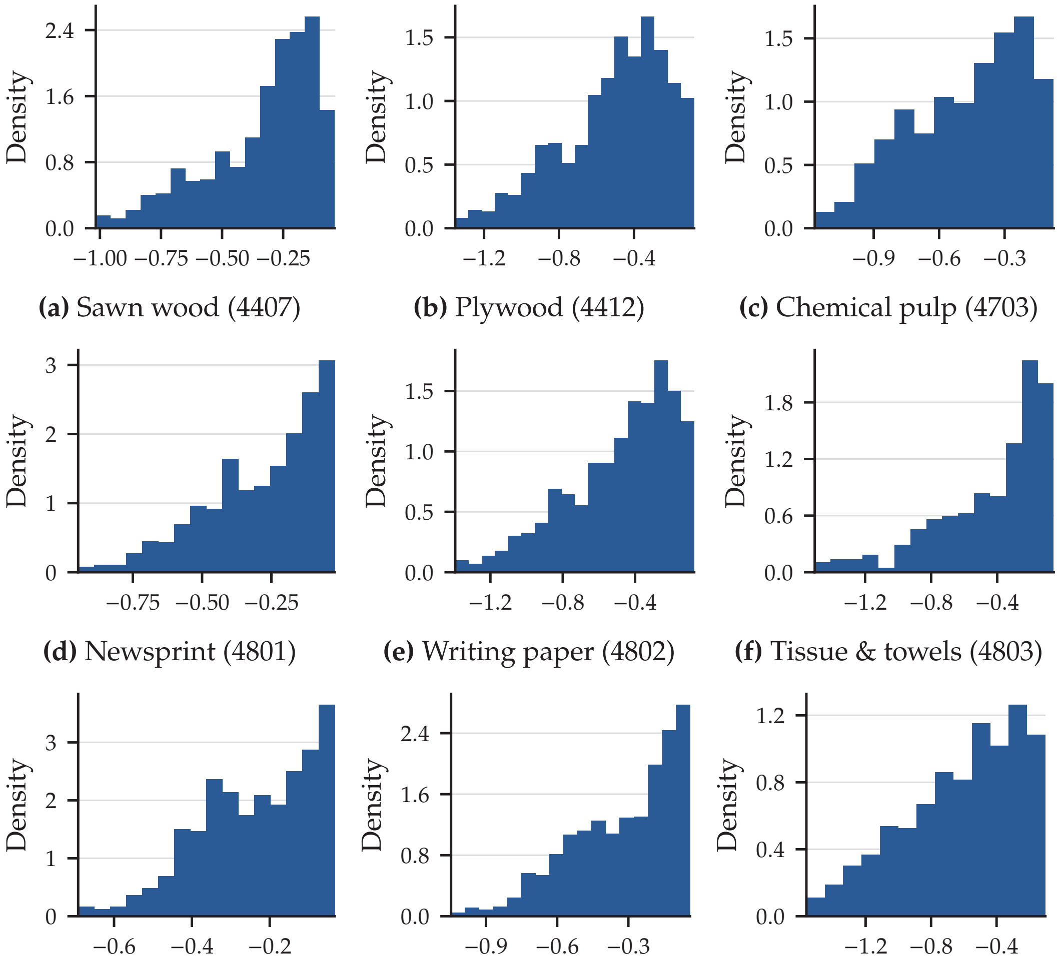

(a) Sawn wood (4407) **(b)** Plywood (4412) **(c)** Chemical pulp (4703)

(d) Newsprint (4801) **(e)** Writing paper (4802) **(f)** Tissue & towels (4803)

(g) Kraft paper (4804) **(h)** Other paper (4805) **(i)** Stainless steel (7219)

Notes: Per-observation demand elasticities with respect to log bilateral distance, recovered from the BLP estimates over 1999–2021. Each panel is a histogram for one HS4 product market.

Figure A.4: Marginal Cost Distributions by HS4 Market

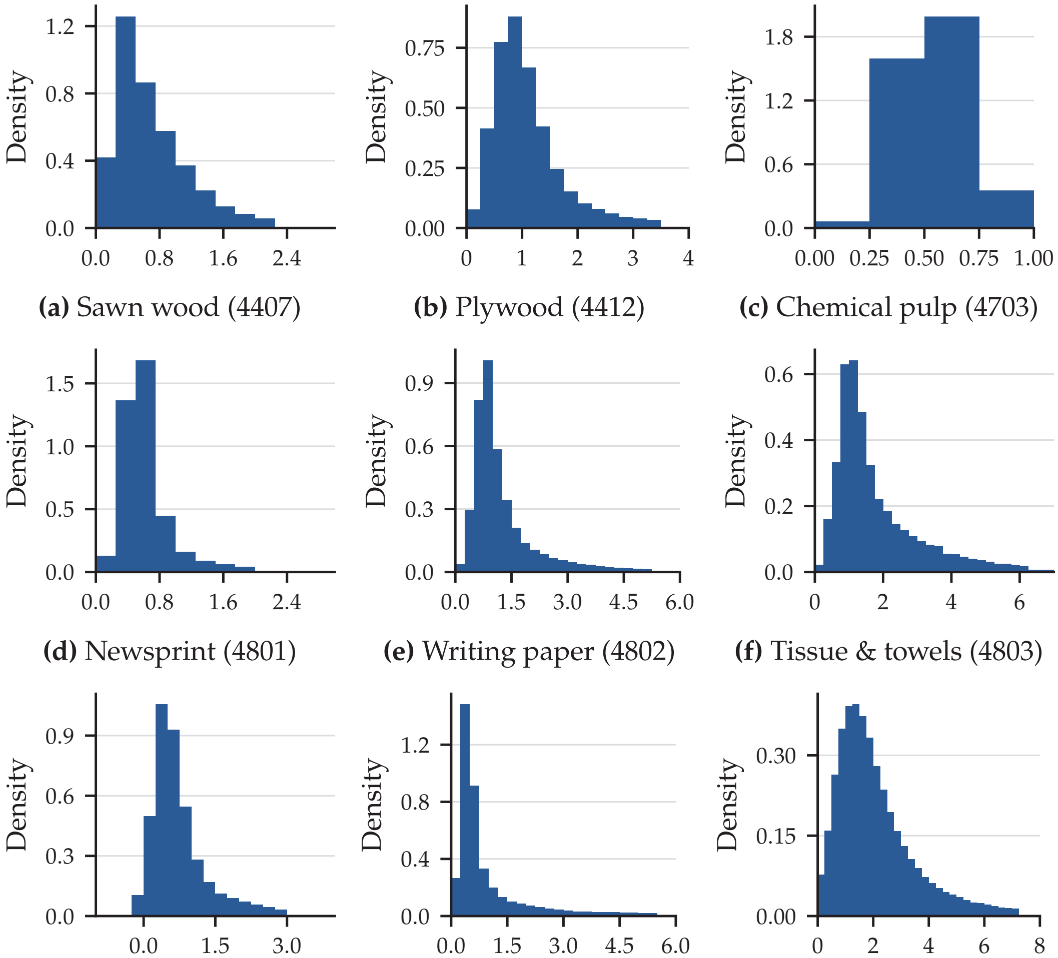

(a) Sawn wood (4407) **(b)** Plywood (4412) **(c)** Chemical pulp (4703)

(d) Newsprint (4801) **(e)** Writing paper (4802) **(f)** Tissue & towels (4803)

(g) Kraft paper (4804) **(h)** Other paper (4805) **(i)** Stainless steel (7219)

Notes: Producer-side (FOB) marginal costs c_{jt} for Finnish exporters, recovered from the Bertrand-Nash markup inversion (Equation (8)) over 1999–2021, in thousands of USD per tonne. Each panel is a histogram for one HS4 product market.

Figure A.5: Lerner Index Distributions by HS4 Market

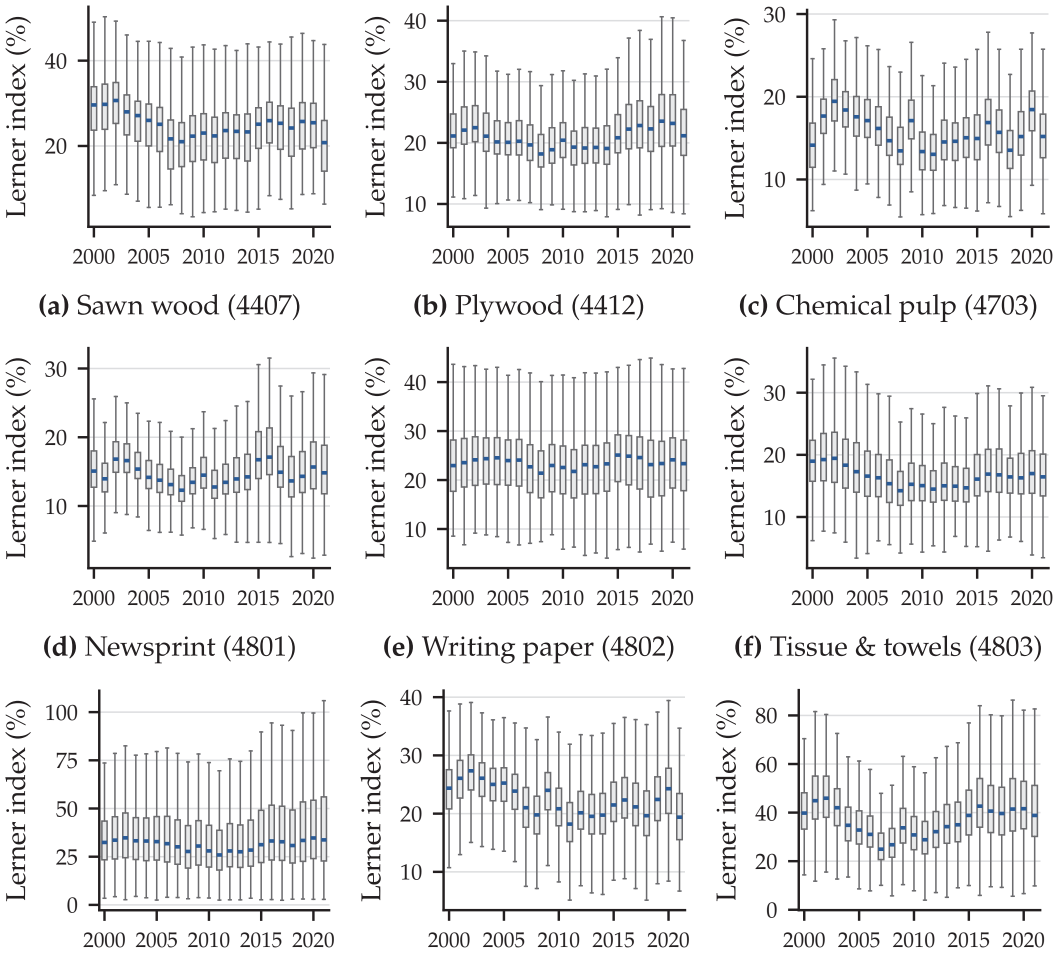

(a) Sawn wood (4407) **(b)** Plywood (4412) **(c)** Chemical pulp (4703)

(d) Newsprint (4801) **(e)** Writing paper (4802) **(f)** Tissue & towels (4803)

(g) Kraft paper (4804) **(h)** Other paper (4805) **(i)** Stainless steel (7219)

Notes: Lerner indices $(p_{jt} - c_{jt})/p_{jt}$ (in percent) for Finnish exporters, recovered from the demand-side inversion over 1999–2021. Each panel shows the per-year box-and-whisker distribution for one HS4 product market.

Figure A.6: Lerner Index Trends by HS4 Market

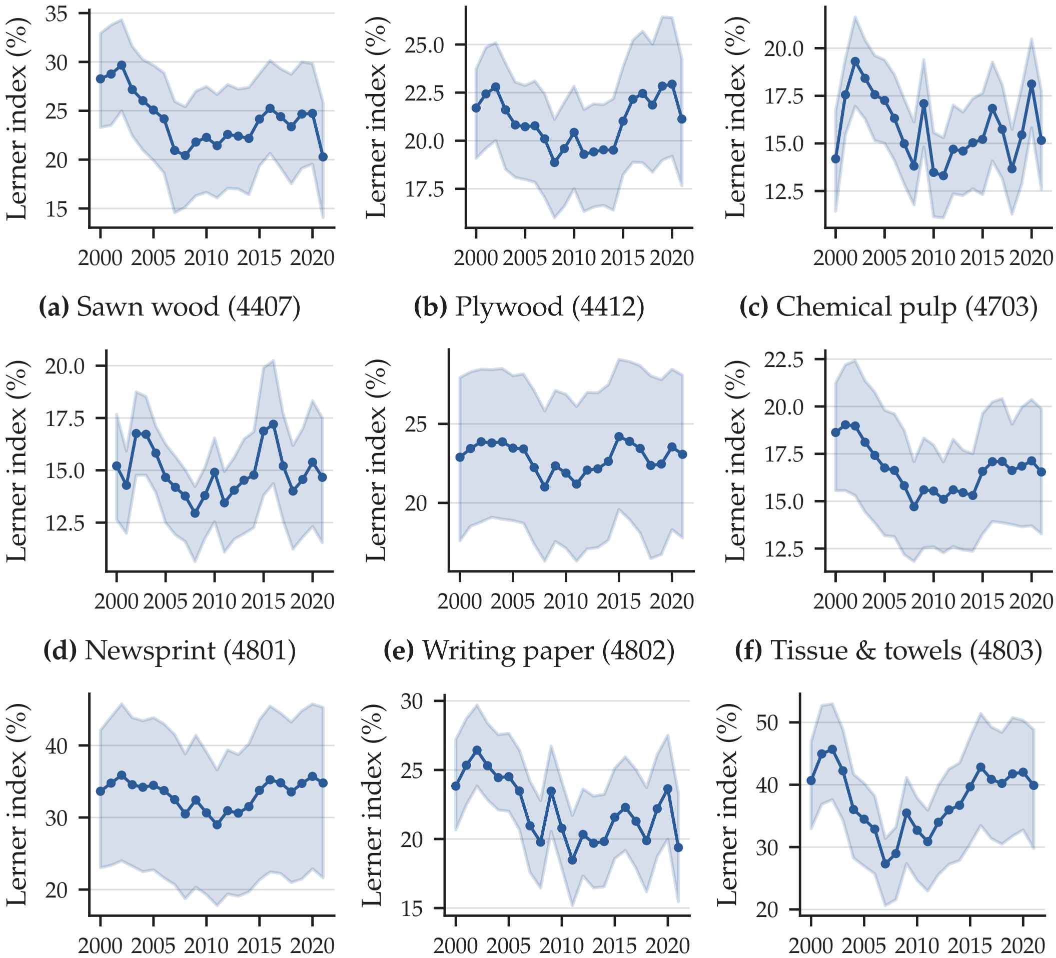

(a) Sawn wood (4407) **(b)** Plywood (4412) **(c)** Chemical pulp (4703)

(d) Newsprint (4801) **(e)** Writing paper (4802) **(f)** Tissue & towels (4803)

(g) Kraft paper (4804) **(h)** Other paper (4805) **(i)** Stainless steel (7219)

Notes: Year-by-year mean Lerner index $(p_{jt} - c_{jt})/p_{jt}$ (in percent) for Finnish exporters, 1999–2021, recovered from the demand-side inversion, with the interquartile range shaded. Each panel shows one HS4 product market; the horizontal axis is the year.

Figure A.7: Marginal Cost Trends by HS4 Market

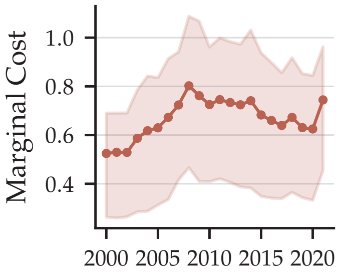

(a) Sawn wood (4407)

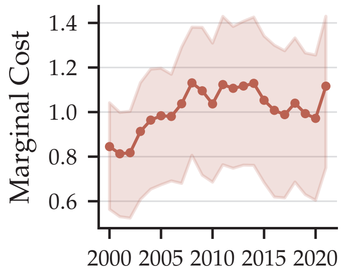

(b) Plywood (4412)

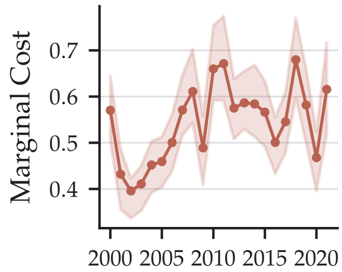

(c) Chemical pulp (4703)

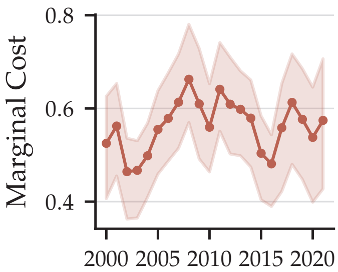

(d) Newsprint (4801)

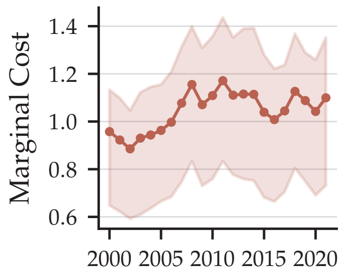

(e) Writing paper (4802)

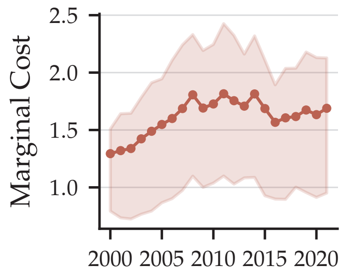

(f) Tissue & towels (4803)

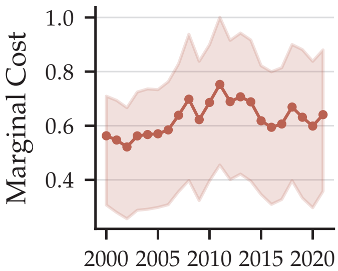

(g) Kraft paper (4804)

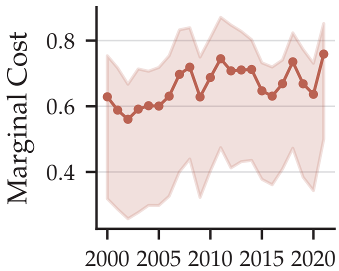

(h) Other paper (4805)

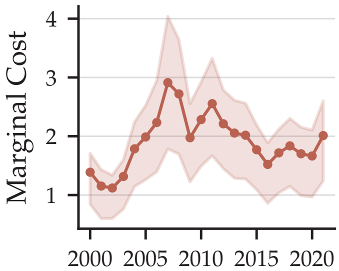

(i) Stainless steel (7219)

Notes: Year-by-year median producer-side marginal costs for Finnish exporters, 1999–2021, in thousands of USD per tonne. Each panel shows one HS4 product market; horizontal axis is the year, vertical axis is the median c_{jt}.

Table A.2: HS4 Product Markets

HS4	Description	Market
HS 44: Wood and articles of wood		
4407	Wood sawn or chipped lengthwise, sliced or peeled, whether or not planed, sanded or end-jointed, of a thickness exceeding 6mm	Sawn wood
4412	Plywood, veneered panels and similar laminated wood	Plywood
HS 47: Pulp of wood		
4703	Chemical wood pulp, soda or sulphate, other than dissolving grades	Chemical pulp
HS 48: Paper and paperboard		
4801	Newsprint, in rolls or sheets	Newsprint
4802	Uncoated paper and paperboard, used for writing, printing or other graphics, non perforated punch-cards and punch tape paper, in rolls or rectangular sheets, of any size, other than paper of heading 4801 or 4803; hand-made paper and paperboard	Writing paper
4803	Tissue, towel, napkin stock or similar; for household or sanitary uses, cellulose wadding, webs of cellulose fibres, in rolls over 36cm in width or rectangular sheets with one side exceeding 36cm when unfolded	Tissue & towels
4804	Uncoated kraft paper and paperboard, in rolls or sheets, other than that of heading no. 4802 or 4803	Kraft paper
4805	Uncoated paper and paperboard n.e.c., in rolls or sheets	Other paper
HS 72: Iron and steel		
7219	Stainless steel; flat-rolled products of width of 600mm or more	Stainless steel

Notes: The Market column gives the short name used to refer to each HS4 product market in the text and figures. Each HS4 market is treated as a separate product market and estimated independently throughout the analysis.

Table A.3: Descriptive Statistics for Finnish Firms

HS4	Obs.	HS6	Markets	Products	Firms	Value	Share
4407	20,084	7	781	597	453	26.10	4.10
4412	5,043	7	804	247	118	10.40	3.60
4703	2,165	4	547	28	15	23.00	4.00
4801	1,401	1	750	17	17	5.30	2.90
4802	5,830	6	1,045	166	89	27.10	6.90
4803	488	1	265	21	21	0.60	1.10
4804	7,378	12	982	149	58	8.10	3.20
4805	3,442	11	937	114	56	5.50	2.60
7219	6,825	14	689	176	46	31.00	6.50
Total	52,656	63	6,800	1,515	873	137.20	4.50

Notes: This table lists the descriptive statistics for the Finnish firms by HS4 categories. Value is reported in billions of USD. All monetary values are deflated to 2010 values. Share represents the revenue share of Finnish firms from the full sample value.

Table A.4: Per-HS4 sufficient-statistic components and Finnish counterfactual response under Sipilä's proposal (R/C-corrected variant), aggregated over 2017–2020.

		Sufficient-statistic components			Finnish response			
HS4	Market	s_L^R	R/C	s_L^C	Δp	Δq	ΔR	PT
4407	Sawn wood	0.09	1.21	0.11	-0.69	+2.59	+125.1	1.32
4412	Plywood	0.11	1.37	0.15	-0.93	+1.98	+26.5	1.19
4703	Chem. pulp	0.10	1.16	0.11	-0.57	+3.52	+176.5	1.17
4801	Newsprint	0.10	1.29	0.13	-0.57	+2.00	+22.6	1.05
4802	Writing paper	0.11	1.29	0.14	-0.73	+2.69	+41.1	1.11
4803	Tissue stock	0.11	1.37	0.15	-1.02	+2.07	+1.4	1.10
4804	Kraft paper	0.10	2.75	0.31	-0.64	+1.03	+11.4	1.06
4805	Other paper	0.12	1.25	0.15	-0.91	+3.44	+28.9	1.31
7219	Stainless steel	0.04	4.73	0.21	-0.77	+0.40	+13.7	1.32
Mean		0.10	1.83	0.16	-0.76	+2.19	+49.7	1.18

Notes: Quantity-weighted means across Finnish firm-years (arithmetic for shares, harmonic for R/C). $s^R_{L,ft} = wL_{ft}/R_{ft}$, labor revenue share (FSS panel); $R_{ft}/C_{ft} = \sum_j p_{jft}Q_{jft} / \sum_j c_{jft}Q_{jft}$, firm-year revenue-to-cost ratio over HS6 products j (BLP cache); $s^C_{L,ft} = (R_{ft}/C_{ft})\, s^R_{L,ft}$, Shephard-consistent labor cost share. R/C is a ratio, not a BLP markup. Because the columns aggregate firm-year objects by different statistics, the displayed $s^C_L \neq (R/C) \times s^R_L$; the simulation applies the per-firm-year R_{ft}/C_{ft} at observation level (Eq. (15)). Δp (simple mean) and Δq (quantity-weighted) are percentage changes on treated Finnish observations; ΔR is the total revenue change (M€); PT is the mean pass-through $\partial p/\partial c$. Labor-cost shock $d\ln_{\text{total}} = -0.0499$ (Prime Minister's Office of Finland 2015); revenue FOB; USD→EUR at 0.875.

A.7 Counterfactual: additional tables

Table A.4 reports the per-HS4 calibration shares (s^R_L, R/C, s^C_L) and the corresponding Finnish counterfactual response under Sipilä's proposal (R/C-corrected variant). Table A.5 disaggregates the headline Finnish revenue change across markets, reporting both euro amounts and percentages of the 2017–2020 BACI baseline, which exposes the cross-market heterogeneity in cost-shock incidence. Table A.6 decomposes the same revenue change across years.

Table A.5: Finnish revenue change by HS4 market and counterfactual scenario (R/C-corrected variant), summed across 2017–2020.

HS4	Market	Baseline (M€)	No kiky (M€)	Sipilä (M€)	Four-day week (M€)
4407	Sawn wood	7,332	-65.7 (-0.90%)	+125.1 (+1.71%)	-658.2 (-8.98%)
4412	Plywood	1,982	-13.7 (-0.69%)	+26.5 (+1.34%)	-146.7 (-7.40%)
4703	Chem. pulp	8,774	-94.3 (-1.07%)	+176.5 (+2.01%)	-980.8 (-11.18%)
4801	Newsprint	1,880	-12.5 (-0.67%)	+22.6 (+1.20%)	-126.8 (-6.74%)
4802	Writing paper	2,706	-23.7 (-0.88%)	+41.1 (+1.52%)	-186.5 (-6.89%)
4803	Tissue stock	126	-0.7 (-0.59%)	+1.4 (+1.08%)	-7.7 (-6.12%)
4804	Kraft paper	2,642	-5.9 (-0.22%)	+11.4 (+0.43%)	-64.4 (-2.44%)
4805	Other paper	1,578	-15.7 (-1.00%)	+28.9 (+1.83%)	-149.4 (-9.47%)
7219	Stainless steel	10,779	-6.8 (-0.06%)	+13.7 (+0.13%)	-74.0 (-0.69%)
Total		37,799	-239.0 (-0.63%)	+447.0 (+1.18%)	-2,394.4 (-6.33%)

Notes: Total Finnish exporter revenue change in millions of euros by HS4 market and counterfactual scenario, summed across 2017–2020 treated years (R/C-corrected variant). Baseline is the 2017–2020 sum of Finnish exports in each HS4 market from BACI HS92; the percentage in parentheses below each € value normalizes the scenario ΔR by that baseline. *No kiky*: kiky-pact effects reversed (employer social-insurance contributions restored to pre-2017 levels, 24 extra annual hours removed). *Sipilä's proposal*: the government's proposed pakkolait package per (Prime Minister's Office of Finland 2015). *Four-day week*: 32-hour work week (416 fewer annual hours, no contribution change). Revenue figures are on FOB basis, matching BACI's reporting convention. USD→EUR at year-averaged rate of 0.875.

Table A.6: Annual Finnish exporter revenue change by counterfactual scenario (R/C-corrected variant), in M€, 2017–2020.

Year	No kiky	Sipilä	Four-day week
2017	-67.0	+115.8	-606.6
2018	-81.3	+122.4	-661.4
2019	-81.8	+121.3	-638.9
2020	-8.6	+87.1	-484.8
Total	-238.7	+446.6	-2,391.7

Notes: Annual Finnish exporter revenue change in millions of euros by counterfactual scenario (R/C-corrected variant). Each year converted at its own USD→EUR average rate (2017: 0.886, 2018: 0.847, 2019: 0.893, 2020: 0.876); the cumulative total here therefore differs from the cross-year aggregate in Tables 8 and A.5 (which use the flat year-averaged rate) by ≈ €0.3–2.7M, a mechanical FX-convention difference. The *No kiky* effect contracts in 2020 because the realized kiky pact was already largely unwound that year. *Sipilä's proposal* and *Four-day week* shocks are constant across treated years, so their year-to-year variation reflects Finnish trade volume only. Revenue figures are on FOB basis (deflated from the BLP solver's landed CIF+tariff prices by a per-observation transport-and-tariff wedge), matching BACI's reporting convention.

Data and policy sources

Ahtiainen, Lasse. 2019. *Työehtosopimusten kattavuus vuosina 2017/2018.* TEM Reports. Ministry of Economic Affairs and Employment of Finland.

Confederation of Finnish Industries. 2015. *Sosiaalivakuutusmaksut vuonna 2015.* https://ek.fi/wp-content/uploads/sosiaalivakuutusmaksut_2015.pdf.

EK, SAK, STTK and Akava. 2016. *Kilpailukykysopimus.* Tripartite agreement signed 14 June 2016, June.

Finnish Forest Association. 2018. *Forest Sector in Finland.* https://forest.fi/forest-facts/forest-sector-in-finland/.

Finnish Forest Industries Federation. 2024. *Economy and International Trade.* https://metsateollisuus.fi/en/economy-and-international-trade/.

Finnish Government. 2015. *Ratkaisujen Suomi: Pääministeri Juha Sipilän hallituksen strateginen ohjelma 29.5.2015.* Hallituksen julkaisusarja 10/2015.

Ministry of Finance in Finland. 2015. *Rakenneyksikön arvioita yksikkötyökustannuksia alentavista toimenpiteistä.* Memo, VM/BO/RY, 7.9.2015, September.

Prime Minister's Office of Finland. 2015. *Hallituksen esitykset kustannuskilpailukyvyn parantamiseksi.*

Statistics Finland. 2025. *Structural business and financial statement statistics.* https://stat.fi/en/statistics/yrti.

ETLA

Elinkeinoelämän tutkimuslaitos

ETLA Economic Research

ISSN-L 2323-2420
ISSN 2323-2420 (print)
ISSN 2323-2439 (pdf)

Publisher: Etlatieto Oy

Tel. +358-9-609 900
www.etla.fi
firstname.lastname@etla.fi

Arkadiankatu 23 B
FIN-00100 Helsinki