Monopsony Power Dynamics during Trade Liberalization and Rising Employment Protection

Ignacio Horta Alejandro Micco*

September 14, 2026

ABSTRACT. We study how two simultaneous but opposing forces, trade liberalization and rising employment protection, jointly shaped employer wage-setting power in Chilean manufacturing over 1990–2018. Using plant-level data, we estimate firm-level wage markdowns via a production-function approach and document that the employment-weighted worker take-home share rose from 0.33 to 0.56 of each dollar of marginal revenue product, a broad-based gain coinciding with institutional reform common to all sectors. Trade worked in the opposite direction: a shift-share instrument identifies a causal effect of import penetration on markdowns with an elasticity that peaks near 0.13 at six years and becomes statistically significant from year four onward under exposure-robust inference. This effect is driven primarily by sustained markup compression, increased labor-market concentration, and firm exit. We develop and calibrate a structural model with oligopsonistic labor markets, firm-level unions, and endogenous exporter status that matches these empirical patterns. Welfare counterfactuals show that rapid liberalization yields 59.6% consumption-equivalent gains, compared to 47.5% under a gradual path. In the model, trade openness and institutional reform are welfare complements: their joint effect exceeds the sum of individual effects by 7.7 percentage points, as higher wages amplify the gains from lower prices.

Keywords: monopsony, wage markdowns, trade liberalization, Nash bargaining, local projections, shift-share, Chile.
JEL Codes: J42, F16, L13, J31, F13.

*Horta: University of Chile (ihorta@fen.uchile.cl). Micco: University of Chile. We thank colleagues for helpful comments and discussions.

How do trade liberalization and rising employment protection jointly shape employer wage-setting power, and what are the welfare stakes of their interaction? This question is particularly acute in emerging economies (where labor market institutions are weaker, outside options more constrained, and trade exposure more volatile) yet evidence from such settings remains scarce. This paper studies these questions in Chile, a small open economy that underwent both transformations simultaneously over three decades: from one of the world's most protected economies to one of its most open, and from a regime of limited collective bargaining to one of strengthened worker protections. That simultaneity allows us to identify both forces: a trade channel that raised employer wage-setting power through reallocation, and an institutional channel that reduced it through bargaining.

Using plant-level manufacturing data for 1990–2018, we estimate firm-level wage markdowns following the production-function approach developed in the industrial organization literature (De Loecker, 2011; Hall, 1986; Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Wooldridge, 2009; Ackerberg et al., 2015; Rovigatti and Mollisi, 2018; Gandhi et al., 2020). This method recovers input–output elasticities and allows us to infer wage markdowns under explicit assumptions about firms' cost minimization and their demand for flexible labor inputs as in De Loecker (2011) and Yeh et al. (2022). The estimates reveal substantial monopsony power: on average, the marginal revenue product of labor is approximately 2.2 times the wage, and workers capture roughly 45 percent of the value they generate, in the range documented for the United States, Europe, and emerging-economy peers (Yeh et al., 2022; Díez et al., 2022; Felix, 2026; Amodio et al., 2025; Huang et al., 2026).

Monopsony power also moved substantially over time, and in two directions at once. Markdowns declined broadly across firms, a secular improvement coinciding with the post-1990 institutional reforms common to all industries: sustained real minimum wage increases, expanded labor inspection capacity, and rapid growth in access to higher education. Yet the decline in simple average markdowns is considerably larger than in labor-weighted measures: even as the typical firm's wage-setting power fell, monopsony power became increasingly concentrated among larger employers. Understanding this divergence requires distinguishing forces that operate within firms from forces that operate through the reallocation of labor across firms.

Conditional on this secular institutional trend, trade openness acted as a countervailing force. Against the backdrop of deepening global trade integration, Chile's aggressive strategy of bilateral and multilateral trade agreements, initiated in the early 1990s, tripled the imports-to-GDP ratio between 1990 and 2018, and import penetration increased in more than two-thirds of manufacturing industries. Exploiting exogenous variation from supply shocks to pre-treatment origin countries' exports to non-Latin American destinations, we provide causal evidence that increased import penetration raised wage markdowns in more affected industries over the medium run, operating through a gradual sequence of markup compression, firm exit, and employment reallocation toward larger firms with greater wage-setting power.[1]

The interaction of these two forces produces three key empirical facts. First, the unweighted mean firm-level markdown declined from approximately 3.3 to 1.7 over three decades, while the employment-weighted measure fell from 3 to 1.9 (equivalently, the employment-weighted worker

[1] See Gutiérrez (2023) and Felix (2026) for related evidence on trade and labor market power.

take-home share rose from about 0.33 to 0.51). Second, this decline is common to all 2-digit manufacturing sectors, consistent with a shared institutional driver rather than industry-specific forces. Third, an Olley and Pakes (1996) decomposition reveals that the covariance between firm size and markdowns rose steadily: modestly *negative* in the early 1990s, when larger firms exhibited lower markdowns. It turned positive over the sample, so that by 2018 the covariance term accounted for about 12 percent of the labor-weighted aggregate, having peaked near 15 percent in the mid-2010s. This growing reallocation wedge (the increasing concentration of monopsony power in larger firms) is consistent with the trade channel, and its timing coincides with the acceleration of import competition following the FTA expansion of the 1990s and 2000s.

Because import penetration is simultaneously determined with labor market outcomes, OLS estimates are biased. We implement a Bartik instrumental-variables strategy (Bartik, 1991) within a panel local-projection framework at the 4-digit industry level. The instrument combines pre-determined origin-country import shares (computed over the 1985–1989 base period, before Chile's wave of bilateral trade agreements) with time-varying shifts in those countries' export capacity to destinations outside Latin America. This shift-share design isolates supply-side variation in import competition that is orthogonal to Chilean labor market conditions.

The LP-IV estimates trace a coherent, lagged adjustment consistent with a Melitz-type selection mechanism rather than any single contemporaneous channel. Import penetration has no significant effect on markdowns on impact but generates a persistent positive response that builds over the medium run, significant under exposure-robust inference from four years onward and peaking at an elasticity of approximately 0.13 at six years. The underlying margins line up in the sequence a selection mechanism predicts.

Markup compression is the clearest and most robustly estimated channel, significant at every horizon and deepening roughly monotonically from about −0.07 on impact to −0.24 by eight years. Labor's share of output is compressed over the middle horizons, and firm exit rises significantly from six to eight years out, precisely the timing a selection channel predicts once sustained margin compression erodes the viability of the least efficient producers, while entry is a precise null, exactly as expected when competitive pressure operates through destruction rather than creation of firms. Less precisely estimated payroll-weighted concentration follows a growing positive path consistent with employment reallocating toward larger higher-markdown employers, and average wages decline significantly over that same five- to eight-year window, falling by about nine percent by year seven to eight, a reduction that is not accompanied by a corresponding reduction in average employment. Taken together, the evidence describes a cumulative process: margins compress first, then the labor market concentrates and marginal producers exit while wages fall and labor reallocates toward larger employers.

To rationalize both forces and evaluate their welfare implications, we develop a structural model of oligopsonistic labor markets in a small open economy. Firms face upward-sloping labor supply curves and wages are determined through Nash bargaining. The bargaining power parameter β governs the division of the monopsony rent: when $\beta = 0$, the firm captures the entire rent (pure monopsony); as β increases, workers capture a larger share and wages rise toward the marginal revenue product. This framework nests the standard monopsony model of Berger et al. (2022a) as a special case and extends it by allowing institutional forces (captured

through β) to interact with trade-induced changes in market structure.

Calibrating the model to match the observed worker take-home share and import penetration at two points in time yields $\beta_1 = 0.04$ in 1990–95 and $\beta_2 = 0.29$ in 2013–18. This increase quantifies, within the model, the cumulative effect of three decades of institutional reform and developments on workers' bargaining position. In the early 1990s, Chilean manufacturing workers captured only 4 percent of monopsony rents through bargaining; by the 2010s, they captured 29 percent.[2]

We evaluate counterfactual trade scenarios using the synthetic control method (Abadie and Gardeazabal, 2003; Abadie et al., 2010; Abadie, 2021) to construct the import trajectory Chile would have followed under more gradual liberalization consistent with pre-1990 trends. The realized trade liberalization generated consumption-equivalent welfare gains of approximately 59.6%, compared with 47.5% under the gradual counterfactual. The welfare advantage of aggressive liberalization is substantial despite higher monopsony power: the realized path raises wage markdowns relative to the counterfactual, but this is more than offset by the consumption gains driven by lower domestic prices.

This paper establishes three findings. First, we identify the dynamic channel through which trade reshapes labor market power. The existing literature emphasizes the effects of import competition on wages and employment (Autor et al., 2013; Mendez, 2015) or studies how monopsony power interacts with firm-level trade decisions (Jha and Rodriguez-Lopez, 2021; Gutiérrez, 2023), but does not trace the full dynamic adjustment path. We show that the causal effect of import penetration on markdowns is not contemporaneous; instead, it builds gradually over three to eight years through a Melitz (2003)-type selection mechanism in which competitive pressure first compresses markups and, with a longer lag, induces firm exit. Markup compression, operating within firms, is the dominant, best-identified channel we document, and the delayed emergence of the full markdown response underscores the importance of dynamic estimation methods for detecting it.

Second, we develop a structural model that extends Berger et al. (2022a) in two directions and is parsimoniously suited to describe the observed phenomena. We embed the oligopsony framework in a small open economy with endogenous exporter status, so that changes in trade costs propagate through the firm-size distribution and directly shape aggregate labor market power. We then introduce firm-level Nash bargaining, in which each firm and its union split the monopsony rent with bargaining weight $\beta \in [0, 1]$, counteracting the firm's wage-setting power. Despite its parsimony, this framework jointly rationalizes the three empirical facts of Section IV: the secular, sector-wide decline in markdowns driven by rising bargaining power, common across industries, with a growing concentration of monopsony power in larger firms as trade reallocates employment. [3]

Third, we show that in this model trade liberalization and institutional reform are welfare complements. Implementing both reforms jointly yields welfare gains that exceed the sum

[2] Chile returned to democracy in 1990 after 17 years of military rule. A center-left coalition won the presidential election that year and remained in office until 2010. See French-Davis and Stallings (2001).

[3] A growing literature documents substantial monopsony power across countries: Yeh et al. (2022) for the United States, Díez et al. (2022) for Europe, Felix (2026) for Brazil, and Amodio et al. (2025) for Colombia. These studies measure the level of monopsony power but do not identify the institutional forces that move it.

of their effects in isolation, by approximately 7.7 consumption-equivalent percentage points. The mechanism is straightforward: trade liberalization expands total surplus by lowering trade costs, while stronger bargaining power increases the worker share of that surplus. As a result, improvements in bargaining power amplify the consumption gains from openness, and trade-induced income gains increase the returns to bargaining. Therefore, Chile's joint implementation of both reforms was potentially welfare-enhancing relative to either policy in isolation.

To our knowledge, no existing work decomposes the joint welfare effects of trade liberalization and labor-market institutional reform into additive channels within an oligopsonistic labor market. The two closest contributions bracket this result. Jha et al. (2026b) show that monopsony power amplifies the dynamic welfare gains from trade liberalization alone, capturing one side of the complementarity, but without an institutional-reform margin. In contrast, Amodio et al. (2025) find that a pro-competitive formalization policy can interact negatively with employer market power. This difference is not contradictory: their policy affects wages indirectly through workers' outside options, so its interaction with market power is theoretically ambiguous. By contrast, a reform that directly shifts bargaining power increases the worker share of trade-generated surplus, generating the positive complementarity documented here. More broadly, this paper contributes to the literature on trade shocks and labor market frictions (Gutiérrez, 2023; Felix, 2026) by showing that these interactions operate not only through market structure, but also through the institutional environment governing wage determination.

The remainder of the paper is organized as follows. Section I describes the institutional and trade policy context. Section II presents the production-function-based methodology used to estimate markups and wage markdowns. Section III describes the data. Section IV documents the level, evolution, and decomposition of monopsony power in Chilean manufacturing. Section V estimates the causal effect of trade exposure on labor market outcomes. Section VI presents the structural model and welfare results. Section VII concludes.

I. Institutional Background and Trade Liberalization in Chile

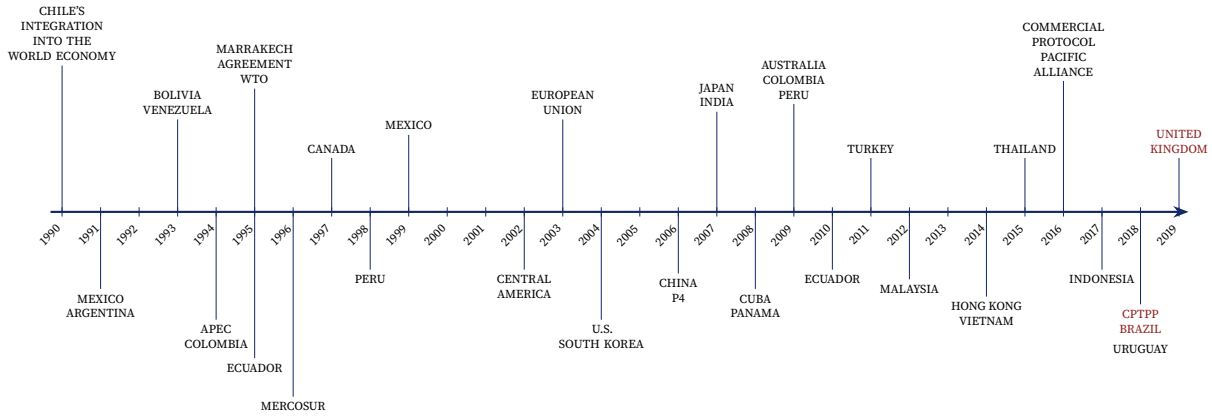
Chile's institutional trajectory over the past five decades provides a natural laboratory for studying the interaction between trade liberalization and labor market power. The country experienced profound changes in trade policy and labor market institutions, changes that moved in opposite directions with respect to firms' wage-setting power. This section describes the institutional context that motivates and identifies the forces examined in this paper.

Trade. Chile's trade liberalization process unfolded in distinct stages. The first stage began in the mid-1970s, during the period 1974–1979, when the country implemented a rapid and substantial unilateral reduction in tariff rates, lowering average tariffs from levels close to 100 percent to a uniform rate of 10 percent on most goods (Pavcnik, 2002). During this phase Chile also eliminated most non-tariff barriers to trade, abandoning the protectionist trade regime that had characterized earlier decades of import-substitution industrialization.

Despite this early commitment to trade openness, import volumes remained relatively subdued until the late 1980s. This delay was partly driven by the severe financial crisis of 1982–1983, during which real output contracted by approximately 8 percent. The recovery from this reces-

sion was gradual, and by the end of the decade import volumes had not yet fully returned to their pre-crisis levels (Pavcnik, 2002).

FIGURE 1—Chilean Free Trade Agreements timeline, 1990–2019



Notes: This figure shows the main Free Trade Agreements signed by Chile between 1990 and 2019. Red agreements corresponds to those that were under parliamentary discussion at that year.
Source: Subsecretaría de Relaciones Económicas Internacionales (SUBREI), Ministry of Foreign Affairs of Chile. Own elaboration.

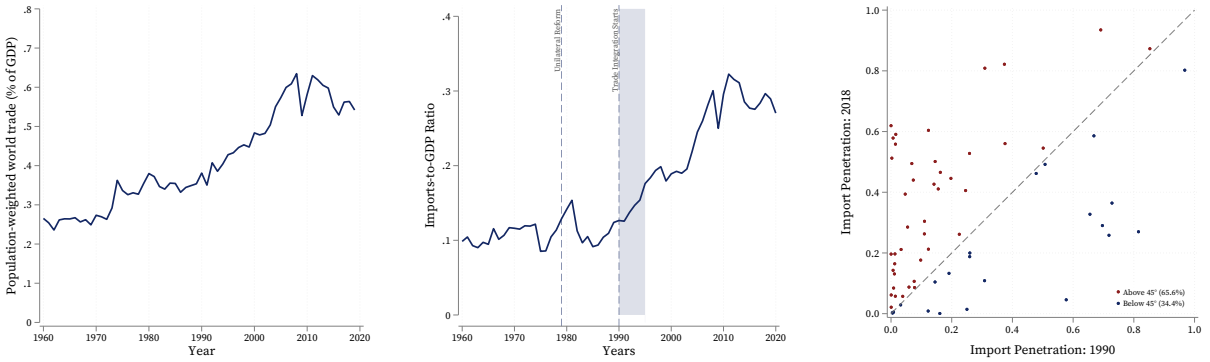
The 1990s marked a period of major political and economic transformation. Following the return to democracy in 1990, Chile entered a decade of sustained economic growth and institutional consolidation. During this period, trade policy shifted away from unilateral liberalization toward a strategy centered on bilateral and multilateral agreements aimed at deeper integration into the global economy. This approach materialized through the negotiation and implementation of numerous Free Trade Agreements (FTAs; see Figure 1). As a result, Chile developed one of the most extensive networks of trade agreements worldwide, ultimately securing preferential access to approximately 88 percent of global GDP (SUBREI, 2019). Accession to the World Trade Organization in 1995 marked an important milestone in Chile's integration into the multilateral trading system.

This transition toward an outward-oriented trade regime coincided with one of the strongest periods of economic performance in Chile's modern history. Economic growth averaged 6.2 percent annually during the 1990s and 4.4 percent during the 2000s (Banco Central de Chile, 2024), while social indicators improved substantially alongside the expansion of social protection programs and labor market reforms. With the onset of this new phase of trade liberalization, imports expanded rapidly and continued to grow throughout the following decades, exposing domestic manufacturing firms to increasing competitive pressure from foreign producers.

Figure 2 places Chile's experience in global context. World trade as a share of GDP has risen steadily since 1960, with the expansion accelerating through the 1990s and 2000s as transport costs fell and production fragmented across borders (Panel a). This sustained growth in origin countries' export capacity is what generates the supply-side variation that will identify the causal effect of import competition in Section V. Chile's imports-to-GDP ratio remained below 15 percent through the 1980s, rose steadily during the 1990s as bilateral agreements took effect, and accelerated sharply in the early 2000s, roughly tripling relative to its 1990 level by 2018 (Panel b).

This economy-wide pattern is mirrored within manufacturing: approximately two-thirds of 4-digit industries experienced higher import penetration in 2018 than in 1990, yet the dispersion is considerable: some industries saw increases of 40 percentage points or more, while others faced stable or declining import competition (Panel c). This cross-industry variation in the magnitude of the trade shock provides the basis for the identification strategy developed below.

FIGURE 2—Increased Trade Exposure: World, National, and Cross-Industry Variation



(a) World Trade-to-GDP Ratio, 1960–2020

(b) Chilean Imports-to-GDP Ratio, 1960–2020

(c) Import Penetration by Industry, 1990–2018

Notes: Panel (a): population-weighted average of total trade (imports plus exports) as a share of GDP across countries. Panel (b): Chilean imports as a share of GDP, 1960–2020. Panel (c): each point is a 4-digit ISIC Rev.2 manufacturing industry; import penetration is imports divided by apparent domestic consumption (production + imports − exports). *Sources:* Penn World Tables (Panels a–b); ENIA and the Atlas of Economic Complexity (Panel c); authors' calculations.

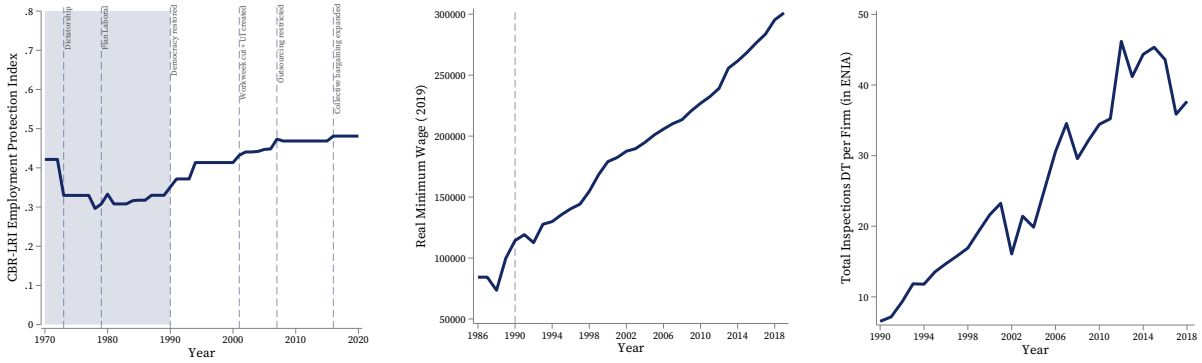
LABOR. Labor market institutions were profoundly reshaped during the authoritarian military regime (1973–1990). Union activity was suppressed during 1974–1979, and the *Plan Laboral* of 1979–1980 configured a labor market in which firms possessed wide latitude to set wages.[4] Collective bargaining was confined to the firm level, legal strikes were capped at sixty days (after which strikers could be dismissed without severance), and employers could freely replace striking workers. Dismissal without expression of cause was permitted from 1981, with mandatory severance capped at five months' wages.[5] The resulting combination of fragmented unions, minimal enforcement, and low separation costs left workers with few credible mechanisms to extract rents from employers.

The return to democracy partially reversed this trajectory. The first reform package (1990–1992) raised the cost of labor adjustment and expanded collective voice. Dismissal henceforth required an explicit justification, and the severance ceiling was raised from five to eleven months' wages. The sixty-day strike limit was eliminated, conditions for replacing striking workers were tightened, and union formation was facilitated through lower quorum requirements and stronger legal protections for union leaders. The *Dirección del Trabajo* (Chile's labor inspectorate) received a 50 percent increase in staffing and an expanded enforcement mandate, raising the expected cost of non-compliance for employers.[6]

[4] See Ffrench-Davis (2010), Edwards et al. (2000), and French-Davis and Stallings (2001), Chapter 6.
[5] See Edwards et al. (2000), Campero (2004), and Heckman and Pages (2004).
[6] See Campero (2004) and Frank (2002).

Subsequent reforms progressively closed the margins through which firms had adapted to this new regulatory environment. The 2001 reform (Ley 19.759) shortened the legal workweek, broadened the scope of collective bargaining, and created Chile's first unemployment insurance system, improving workers' outside option in wage negotiations. The 2007 Subcontracting Act (Ley 20.123) imposed joint liability on principal firms for their subcontractors' labor obligations, curbing outsourcing as a route around employment protections. A 2008 judicial reform more than tripled the number of specialized labor courts and expedited procedures, strengthening enforcement. The 2016 reform (Ley 20.940) banned the hiring of replacement workers during legal strikes and broadened the scope of bargainable subjects.[7] Figure 3 documents this trajectory quantitatively: the CBR-LRI employment protection index tracks the legislative sequence closely, rising sharply after each reform package, while the real minimum wage and the inspective capacity of the *Dirección del Trabajo* both rose almost continuously from 1990 onward.

FIGURE 3—Labor protection: Eroded, then restored



(a) CBR-LRI Employment Protection Index

(b) Chilean Real Minimum Wage

(c) DT Inspective Capacity

Notes: Panel (a): CBR-LRI employment protection index for Chile, 1970–2020. Source: Centre for Business Research, University of Cambridge. Panel (b): Real minimum wage in Chile, 1990–2018 (constant prices). Source: Chilean National Statistics Institute (INE). Panel (c): Labor inspections by the *Dirección del Trabajo* (DT) per ENIA plant, 1990–2018. Source: *Dirección del Trabajo*; authors' calculations.

Alongside these legislative changes, minimum wage policy served as a distinct and quantitatively important instrument: the real minimum wage rose by approximately 34.4 percent between 1990 and 2000 and by roughly 120 percent over the following decades,[8] progressively compressing the lower tail of the wage distribution and directly constraining firms' ability to mark down wages for low-skilled workers; by 2018 it stood at 46 percent of the average wage and 70 percent of the median, above the ratio in most OECD economies (OECD, 2009).

Taken together, these reforms shifted bargaining power toward workers through higher separation costs, broader collective rights, stronger enforcement, a binding wage floor, and improved outside options, the last reinforced by a more than threefold expansion of tertiary enrollment over the sample period, which broadened the pool of skilled workers and improved

[7] See Sánchez (2013) on the workweek reduction, OECD (2009) on unemployment insurance, Micco and Muñoz (2024) on subcontracting, and Pérez Ahumada (2021) on the 2016 reform.

[8] By contrast, during the previous two decades the real minimum wage declined by approximately 6.7 percent; see French-Davis and Stallings (2001).

fallback opportunities for a growing share of the labor force. This institutional trajectory provides the source of over-time variation in labor market conditions that disciplines our empirical analysis.

TRADE AND LABOR. Chile's two trajectories thus deliver, from a starting point combining a deregulated labor market with limited effective import exposure, two large and simultaneous changes that move in opposite directions with respect to firms' wage-setting power. The institutional strengthening of workers' bargaining position (captured in the structural framework of Section VI by the bargaining parameter β) predicts a secular, broad-based decline in wage markdowns. Trade-induced reallocation toward larger, more productive firms (Melitz, 2003) predicts the opposite for the labor-weighted aggregate: if larger firms face more inelastic labor supply, concentration raises the labor-weighted markdown even as within-firm markdowns decline. Neither force alone explains the evolution of Chilean labor markets; their interaction structures the empirical and structural analysis that follows.

II. Measuring Labor Market Power

MARKUPS. We estimate market power using the production-function approach, following De Loecker and Warzynski (2012) and Yeh et al. (2022). Suppose firms minimize costs subject to their technology:

$$\begin{aligned} \min_{X_{it} \in \mathcal{K}} \quad & \sum_{k=1}^{K} P_{it}^{k}(X_{it}^{k}) X_{it}^{k} + \Phi_{t}^{k}(X_{it}^{k}, X_{it-1}^{k}) \\ \text{s.t.} \quad & F(X_{it}; \omega_{it}) = Q_{it} \end{aligned} \tag{1}$$

Here P_{it}^{k} and X_{it}^{k} denote the price and quantity of input k, $F(X_{it}; \omega_{it})$ is the production technology, Q_{it} is output, and $\Phi_{t}^{k}(\cdot)$ captures adjustment costs. $\mathcal{K}$ is the set of $K > 1$ production inputs.

If there exists a flexible input $v \in \mathcal{K}$ not subject to adjustment costs, then following Hall (1986) the markup can be expressed as a function of input v's output elasticity (θ_{it}^{v}) and its revenue share (α_{it}^{v})[9].

$$\mu_{it} = \frac{\theta_{it}^{v}}{\alpha_{it}^{v}} \tag{2}$$

WAGE MARKDOWNS. Consider the conditional optimization problem for labor demand:

$$\begin{aligned} \min_{l_{it}} \quad & w_{it}(l_{it}) l_{it} \\ \text{s.t.} \quad & F(l_{it}, X_{-l,it}^{*}) = Q_{it} \end{aligned} \tag{3}$$

Here $X_{-l,it}^{*}$ is the vector of previously optimized inputs. Following Yeh et al. (2022), if the

[9] We rely on materials as flexible inputs, as is standard in the literature.

flexible input used to identify markups is not subject to monopsony forces, then plant-level wage markdowns (ν_{it}) are identified from the first-order conditions:

$$\nu_{it} \equiv 1 + \varepsilon_s^{-1}(l_{it}) = \mu_{it}^{-1}\left(\frac{\theta_{it}^{l}}{\alpha_{it}^{l}}\right) \tag{4}$$

Here $\varepsilon_s^{-1}(l_{it})$ is the firm's perceived inverse elasticity of labor supply, θ_{it}^{l} the labor output elasticity, and α_{it}^{l} the revenue labor share. Computing markups and markdowns thus requires only both inputs' revenue shares (directly observable) and output elasticities, which we estimate using the production-function method.[10]

Input-Output Elasticities. To estimate the relevant elasticities we follow the Ackerberg et al. (2015) estimation procedure under the assumption of a general but unknown production function $f(x_{it})$ that depends on a vector x_{it} of production inputs and is approximated by a Translog production function:

$$y_{it} = \beta_0 + f(l_{it}, k_{it}, m_{it}, e_{it}) + \omega_{it} + \varepsilon_{it} \tag{5}$$

Where y_{it} is the (log) gross output of plant i at year t, and $x_{it} = \{l_{it}, k_{it}, m_{it}, e_{it}\}$ represents the production inputs (capital, labor, materials, and energy)[11], ω_{it} corresponds to the systematic component of plant productivity, and ε_{it} is the idiosyncratic part. In this context, ω_{it} is the plant's total factor productivity (a state variable observed by the firm but not by the econometrician), and ε_{it} is an idiosyncratic shock. Because ω_{it} is unobserved, we follow the proxy-variable approach to control for it.

Estimation Algorithm. We rely on Ackerberg et al. (2015) (ACF) insights and estimate production function as follows: First, run a OLS regression of y_{it} on the inputs $\tilde{\mathbf{x}}_{it}$. Then, obtain nonparametric estimates of log output ϕ_{it} free of measurement error. Afterwards, construct an estimate of productivity as: $\omega_{it}(\tilde{\beta}) = \phi_{it} - f(\mathbf{x}_{it}; \tilde{\beta})$ and run a third-order polynomial regression of $\omega_{it}(\tilde{\beta})$ on $\omega_{it-1}(\tilde{\beta})$ to obtain estimates of productivity shocks $\xi_{it}(\tilde{\beta})$. We then obtain estimates $\hat{\beta}$ of the production function parameters through the GMM system induced by the moment conditions $\mathbb{E}[\xi_{it}(\tilde{\beta}) \cdot z_{it}] = \mathbf{0}_{Z\times 1}$ [12].

[10] The estimated input–output elasticities by sector are reported in Table D.2, Appendix D.

[11] Capital series are constructed by applying the perpetual inventories method. Labor is measured in number of workers. Materials is the gross expenditure in materials, and raw materials; and energy is a composite of fuels and electricity, measured in tons of oil equivalent units (TOE). All of them are written in logarithms.

[12] This methodology rests on several standard assumptions of the production-function literature, detailed exhaustively in Ackerberg et al. (2015) and Yeh et al. (2022). Formally, the procedure imposes: (i) the idiosyncratic productivity shock is not observed by the firm; (ii) the capital stock accumulates as a function of lagged capital and prior-period investment; (iii) production-function parameters are time-invariant and common within a sector; (iv) productivity evolves according to a first-order Markov process; (v) for some flexible input, the only unobservable to the econometrician in the input demand equation is productivity; and (vi) for the flexible input satisfying assumption (v), its demand is strictly monotone in productivity. Appendix 3 discusses how these assumptions apply in this paper's setting, taking Yeh et al. (2022) as the closest precedent.

III. Data

The main dataset used in this paper is the Annual National Industrial Survey (*Encuesta Nacional Industrial Anual*, ENIA), a publicly available census of Chilean manufacturing plants conducted annually by the Chilean National Statistics Institute (INE). ENIA provides plant-level information on revenues and input expenditures, including labor, intermediate inputs, electricity, and fuels, as well as measures of each plant's capital stock, and covers all manufacturing establishments with at least 10 employees, approximately 4,000 plants per year on average. The survey is available from 1979 onward. We use the full historical record to construct capital stocks via the perpetual inventory method and lagged industry-level controls, but restrict the estimation sample itself to 1990–2018. This restriction is deliberate: we already assume that the production-function parameters are constant within each sector over the estimation window, and extending the window back into the 1980s would additionally require them to be invariant across Chile's 1990 transition from military dictatorship to democracy, a structural break that reshaped trade policy, labor-market institutions, and the macroeconomic environment. Confining estimation to the post-1990 democratic period keeps this constancy assumption credible and aligns the sample with the reforms of interest, both of which are post-1990.

Additionally, we exclude plant-year observations with missing, zero, or negative values for any core production variable: labor, intermediate inputs, energy, capital, wages, or gross output. To limit the influence of extreme values and measurement error, we trim the top and bottom 1 percent of output, capital, materials, and energy per worker within year–sector cells. The resulting dataset contains 114,825 plant-year observations.[13] Table D.1 reports descriptive statistics. The average plant employs 85 workers per year, 47% of whom are skilled and 22.6% female. Incorporated firms (*sociedades anónimas*, open and closed, and limited-liability companies) account for 76% of plants. Cooperatives and other legal forms account for the remainder. Approximately one in five plants reports export income, and roughly 23% import materials.

To measure exposure to international trade at the industry level, we complement ENIA with data from the Atlas of Economic Complexity, which provides bilateral trade flows at the country–partner–product–year level for the period 1962–2021. We aggregate these flows to the manufacturing sector (4-digit ISIC Rev. 2)-year level, consistent with the sectoral classification used in ENIA, to construct the industry-level import penetration measure and the shift-share instrument used in Section V.

Because ENIA reports most variables in nominal terms, we construct industry-level deflators for both outputs and inputs spanning the full 1979–2018 period. For 1979–1986 and 1995–2007, we rely on previously constructed deflators from Liu (1993) and the replication materials of Cusolito et al. (2023), respectively. For the remaining years, we construct deflators from multiple official sources, including the Central Bank of Chile, INE, and the Chilean National Customs Service. This is standard practice in the literature, adopted to strengthen the assumption that production function estimation captures real-quantity variation rather than changes in nominal prices.[14]

[13] The *effective* sample size is 85,477 observations after excluding plants with economically meaningless estimated elasticities and trimming the top and bottom 1 percent of the markup and markdown distributions.

[14] This is the second-best solution the literature has adopted to address the broadly discussed problem of input price bias (Bond et al., 2021), which arises when estimation relies on deflated revenues rather than physical output. Despite this, Yeh et al. (2022) show that the concern is relevant for markup estimation but not for markdown

In addition to these core datasets, we use two auxiliary data sources to support specific components of the empirical analysis. First, the Penn World Table (Feenstra et al., 2015) provides internationally comparable macroeconomic data on output, factor accumulation, and productivity, constructed using harmonized national accounts and purchasing power parity adjustments. Its broad country coverage and long time span make it particularly suitable for the synthetic control-based definition of our counterfactual exercises employed in the paper, where Chile is benchmarked against a weighted combination of comparable economies.

Second, we use data from the Chilean National Socioeconomic Characterization Survey (CASEN), a large and nationally representative household survey conducted periodically by the Ministry of Social Development. CASEN contains detailed information on individuals' labor market status, occupation, industry affiliation, and location. Its primary role in our analysis is to construct within-region employment shares at the sector–*comuna* level, which pin down local labor market sizes in the structural model of Section VI. We additionally use CASEN to construct industry-level measures of occupational composition and self-employment intensity, which proxy for outside options in the heterogeneity analysis of Appendix 5.

A measurement caveat follows from ENIA's coverage threshold. Because the survey samples only plants with at least ten employees, the smallest firms, which face the most elastic labor supply and therefore the smallest markdowns, are unobserved. Markdowns computed on the ENIA sample alone are consequently overstated relative to the economy-wide object. Using administrative records (SII) to discipline the size and sales weight of the sub-threshold "competitive fringe," we quantify a model-based measure of this bias in Section VI and Appendix F: splicing the fringe into the aggregate reveals an 8.5% overstatement in 1990-1995 and a 4.4% overstatement in 2013-2018, as reported in Table E.1. The bias is therefore material early in the sample and small by the end; for level comparisons, the correction is applied explicitly in the structural calibration of Section VI, and it does not alter the direction of any result.

IV. The Level and Evolution of Wage Markdowns

A. Cross-Section of Monopsony Power

Table 1 establishes the central cross-sectional descriptive fact: Chilean manufacturing deviates substantially from the competitive benchmark. On average, the value of the marginal product of labor is more than twice the wage. Equivalently, employers capture more than half of the surplus each additional worker generates, with considerable heterogeneity across industries. Most sectors exhibit substantial wage-setting power. Food and beverages, metals, textiles, and apparel display the highest monopsony power, with average markdowns above 2 and workers receiving only 33–43% of the value of their marginal product. By contrast, paper and machinery face more competitive labor markets, with markdowns closer to 1, so pay tracks productivity far more tightly. The average markdown is 2.21 and the median firm's is 1.85, implying take-home shares of 45% and 54% of marginal product, respectively, the take-home share being simply the reciprocal of the markdown. The final column translates these markdowns into the average plant-level labor-supply elasticity they imply, which moves inversely with wage-setting power:

estimation: because markdowns are constructed as a *ratio*, the bias documented by Bond et al. (2021) cancels out.

the more inelastic the labor a firm faces, the more it can hold pay below marginal product without losing its workforce. At the economy-wide average of 0.79, a plant that reduced its wage offer by ten percent would contract its employment by under eight percent, so little labor is forgone by underpaying that markdowns of this magnitude are sustainable. The spread is economically large: metals, the highest-markdown industry, faces the most inelastic supply (0.43), whereas plastics and paper (near unity) are the least monopsonistic in the sample, with labor supply responding roughly proportionally to wages. These figures place the average Chilean manufacturing worker during 1990–2018 within the range of take-home shares documented for China (40%, Huang et al., 2026), Brazil (50%, Felix, 2026), the United States (65%, Yeh et al., 2022) and Colombia (70%, Amodio and de Roux, 2024), and are consistent with the Chilean cross-section reported in a recent 16-country Latin American and Caribbean sample (markdown ≈ 1.723, using 2010 data, Eslava et al., 2025).[15]

TABLE 1—Markdowns and Implied Labor Supply Elasticity by Economic Activity

	Labor Markdowns					Supply Elasticity
	Mean	Median	P25	P75	SD	Median
Food and Beverages	2.40	2.14	1.57	2.91	1.25	0.79
Textiles	2.33	1.87	1.29	2.90	1.59	0.81
Apparel	2.52	2.25	1.57	3.15	1.37	0.67
Paper and Woods	1.59	1.34	0.97	1.94	0.93	0.98
Chemicals	2.06	1.75	1.12	2.56	1.42	0.73
Plastics	1.91	1.71	1.28	2.28	0.98	1.09
Non-Metallic Manufactures	2.34	2.06	1.49	2.81	1.30	0.82
Metals	3.02	2.24	1.24	4.03	2.47	0.43
Machinery and Equipment	1.81	1.52	1.06	2.21	1.17	0.92
Average	2.21	1.85	1.26	2.70	1.48	0.79

Notes: Sectors are defined at the 2-digit ISIC Rev. 3 level.

These cross-sectional averages mask systematic heterogeneity in wage-setting power. Three firm-level patterns are diagnostic of monopsony (Appendix 5). First, the female share of employment is strongly and positively associated with markdowns (0.438): holding sector and year fixed, plants employing relatively more women hold pay further below marginal product, consistent with the pattern a monopsony account predicts if women face thinner outside options than comparable men, and with evidence that monopsony exposure contributes to the gender wage gap (Sharma et al., 2023; Blau and Kahn, 2017). Second, the unskilled-labor share carries a related positive coefficient (0.231): less-skilled workers typically have fewer alternative employers and search less intensively, leaving employers more room to underpay. Third, larger and more capital-intensive plants exhibit higher markdowns: the coefficient on log capital is 0.069, so a doubling of the capital stock is associated with about 0.05 additional markdown points. Appendix 5 documents the remaining margins in the same direction: exporters mark down modestly more than otherwise-similar non-exporters within sector and year, as they tend to be the larger, more productive employers with correspondingly greater wage-setting power, and

[15] In that year our average markdown estimate is 1.88.

organizational form matters, with worker cooperatives, partnerships, and limited companies compressing markdowns relative to sole proprietorships, signs consistent with these forms letting workers share in firm surplus rather than confront it as a pure counterparty.

Industry-level evidence from the CASEN household survey complements the plant-level picture along the outside-options margin: informality and self-employment are negatively associated with markdowns, consistent with a thicker informal margin widening workers' outside options and constraining employers' wage-setting power (Amodio et al., 2025), while industries intensive in portable-skill occupations (scientific professionals and administrative staff) display markedly lower markdowns and mechanical and craft workers are statistically indistinguishable from the unskilled baseline.[16]

B. Three Facts about the Evolution of Monopsony Power

The cross-sectional analysis does not capture the dynamics of market power over time. Aggregating wage markdowns annually reveals three facts that structure the remainder of the paper.

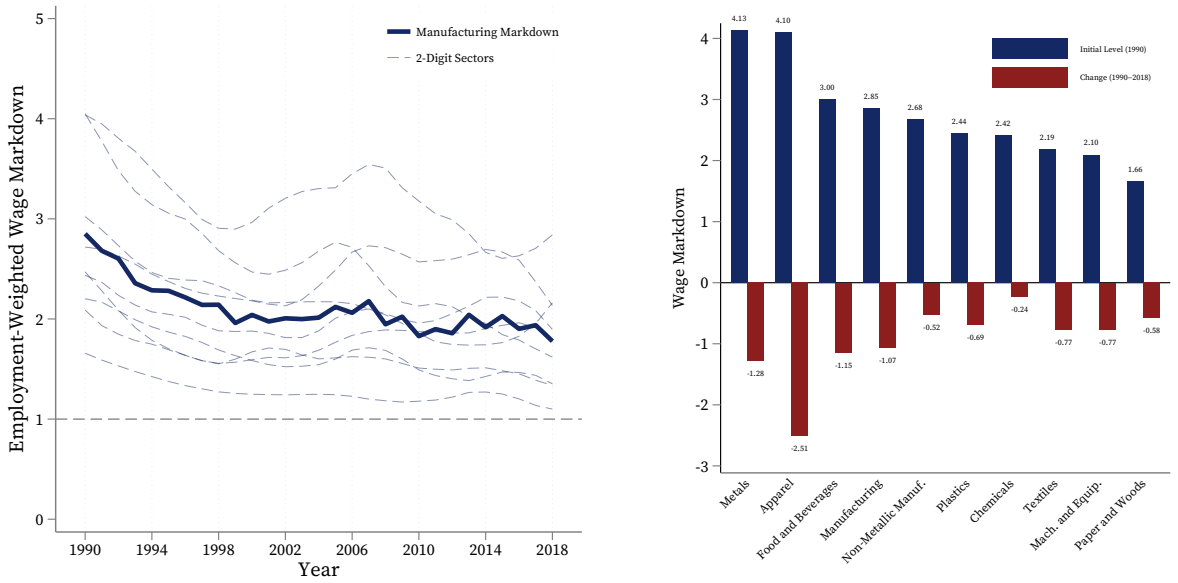
Fact 1: A secular, manufacturing-wide decline in monopsony power. Between 1990 and 2018, the employment-weighted mean markdown fell from approximately 3 to 1.9. The unweighted mean fell more steeply, from about 3.3 to 1.7. At the beginning of the 1990s, workers received roughly 33 cents for every dollar of marginal revenue product. By 2018, this figure had risen to 56 cents. These patterns are consistent with the institutional transformation documented in Section I: sustained real minimum wage increases, expanded labor inspection capacity, rapid growth in higher education access and overall expanded labor protection. This monotonic, three-decade decline is not the general pattern documented elsewhere: in Spain, markdowns fell after the 2008 financial crisis but partially reversed thereafter, with the resurgence concentrated among larger firms (Azar and Sanz-Espín, 2025), underscoring that Chile's sustained decline reflects a specific institutional trajectory rather than a universal or self-correcting tendency. Appendix C offers suggestive industry-level evidence broadly consistent with this channel: the pace of the markdown decline varies with predetermined industry skill intensity and initial concentration in a pattern one would expect if institutional reform bears differentially on the industries where monopsony power was initially strongest.

Fact 2: The decline is common to all manufacturing sectors. The secular decline is not concentrated in particular industries: it is shared across all nine 2-digit manufacturing sectors, which experienced broadly parallel reductions in monopsony power over the period. Remarkably, all nine sectors record a decline of at least 0.24 markdown points. The uniformity of this decline across industries with different trade exposures, technologies, and workforce compositions points to common, economy-wide forces rather than idiosyncratic industry-level shocks.[17]

[16] Full results on heterogeneity by firm characteristics, legal organization, and occupational composition are reported in Tables D.3 and D.4 and Figure D.4 in Appendix 5.

[17] In Appendix C we provide suggestive evidence that, even conditioning on the trend, the industry-level markdown reduction can be explained by the labor protection index, minimum wages, and DT inspection capacity.

FIGURE 4—Wage Markdown Trends Across Manufacturing



(a) National markdown trends by sector

(b) Initial level and 1990–2018 change

Notes: Left panel: employment-weighted wage markdown for aggregate manufacturing (solid line) and each 2-digit ISIC sector (dashed lines), 1990–2018. Right panel: initial markdown level (1990, blue) and total change (1990–2018, red) by 2-digit ISIC sector and manufacturing aggregate.

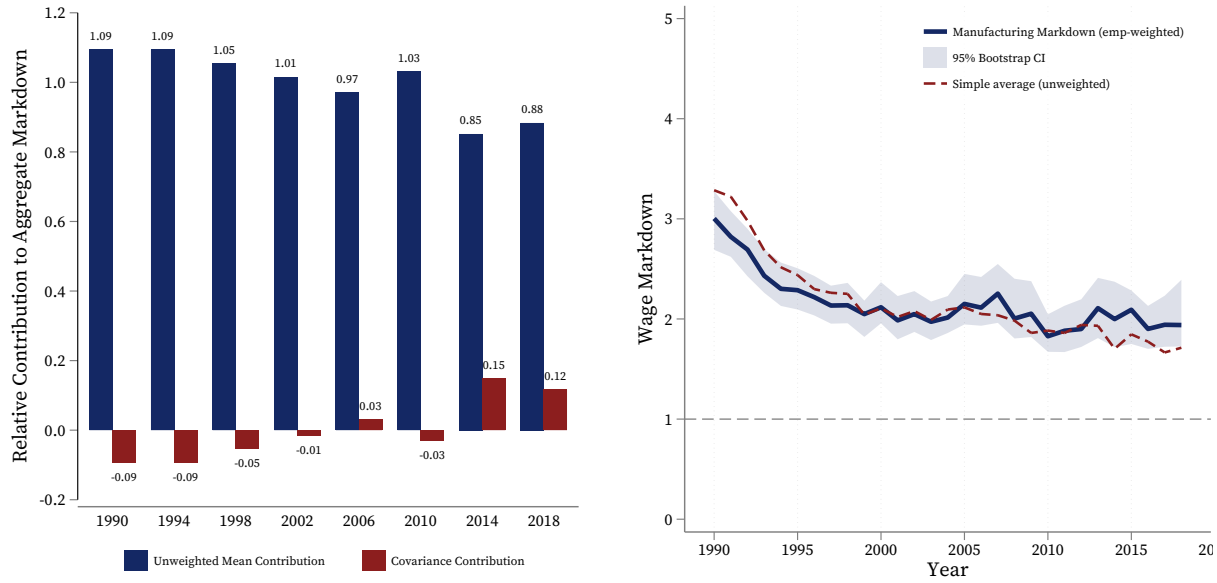
FACT 3: A GROWING CONCENTRATION OF MONOPSONY POWER IN LARGER FIRMS. While the typical firm's markdown fell, employment simultaneously shifted toward firms with higher wage-setting power. To formalize this distinction, we implement the Olley and Pakes (1996) decomposition, which separates the labor-weighted aggregate markdown Φ_t into the unweighted mean and the covariance between firm size and markdowns:

$$\Phi_t = \bar{\phi}_t + \text{cov}\left(s_{it}, \phi_{it}\right), \tag{6}$$

where $\bar{\phi}_t$ is the unweighted mean of firm-level markdowns, s_{it} is firm i's share of total employment, and $\text{cov}(s_{it}, \phi_{it})$ captures the extent to which larger firms exhibit systematically higher or lower markdowns than the average firm.

Figure 5 confirms all three facts and reveals their interaction. The right panel shows that the unweighted mean $\bar{\phi}_t$ drives the aggregate decline, falling from approximately 3.3 to 1.7 over three decades. This component captures the broad-based reduction in the typical firm's wage-setting power attributable to institutional reform. The covariance term, by contrast, was modestly *negative* in the early 1990s, when larger firms exhibited slightly *lower* markdowns than the typical firm, and rose steadily over the sample to about 12 percent of the labor-weighted aggregate by 2018, peaking near 15 percent around 2014. The left panel illustrates the consequence: the labor-weighted markdown declines less steeply than the unweighted mean, and the gap between the two widens progressively. By the end of the sample, about 12 percent of aggregate monopsony power, as measured by the employment-weighted average, arises not from the typical firm's behavior but from the disproportionate concentration of employment in high-markdown firms.

FIGURE 5—Olley–Pakes Decomposition of Aggregate Markdowns



(a) Mean and size–markdown covariance

(b) National markdown with bootstrap CI

Notes: Left panel: unweighted mean and covariance term of the Olley and Pakes (1996) decomposition, 1990–2018. Right panel: national wage markdown: block bootstrap 95% confidence-interval band around the employment-weighted aggregate, together with the unweighted simple average across firms.

The growing covariance admits two mechanically distinct interpretations. Employment may be reallocating toward firms that already possess high markdowns, the reallocation channel emphasized by Melitz-type trade models, in which import competition induces exit among smaller producers and expansion among larger ones. Alternatively, the markdowns of large firms may be rising relative to those of smaller firms, possibly because institutional reforms differentially benefit workers in smaller establishments whose initial bargaining power was weakest. The timing is suggestive: the covariance was modestly negative through the early 1990s, began rising noticeably in the mid-to-late 1990s as Chile's FTA strategy accelerated, turned positive, and continued increasing through the 2000s, coinciding with the implementation of major agreements with the United States, the European Union, China, and South Korea. While this temporal correlation is not causal, it motivates the formal test in Section V.

Facts 1 and 2 are consistent with institutional forces that improved workers' outside options economy-wide, while Fact 3 reflects a reallocation channel that partially offset this decline by concentrating employment in firms with greater wage-setting power. Section V provides causal evidence that trade liberalization drives this reallocation, and Section VI embeds both forces in a structural model to evaluate their welfare implications.

V. Trade and Monopsony Power

In this section, we estimate the effect of import penetration on wage markdowns and other labor market outcomes using a linear model. To achieve this, we compare variation across 4-digit manufacturing industries in the extent to which they were exposed to import competition

over the period 1990–2018. Specifically, we estimate the following equation for each horizon $h = 0, 1, 2, \ldots, 8$ in a panel local projection setting:

$$\Delta_h y_{j,t} = \alpha_j^h + \delta_t^h + \beta^h \ln(IP_{jt}) + \mathbf{X}'_{jt}\boldsymbol{\gamma}^h + \varepsilon_{jt}^h, \tag{7}$$

where $\Delta_h y_{j,t} \equiv y_{j,t+h} - y_{j,t-1}$ denotes the cumulative change (i.e., the long-difference) in the outcome of interest in industry j between the year immediately preceding the shock ($t-1$) and horizon $t+h$ (in the main specification, the log employment-weighted aggregate markdown). The variable IP_{jt} is import penetration, defined as the ratio of imports to apparent domestic consumption:

$$IP_{jt} = \frac{M_{jt}}{Q_{jt} + M_{jt} - X_{jt}}, \tag{8}$$

where M_{jt} denotes Chilean imports, X_{jt} Chilean exports, and Q_{jt} domestic production in industry j at time t.[18] The vector $\mathbf{X}_{jt}$ collects time-varying controls described in Section B. The terms α_j^h and δ_t^h are industry and year fixed effects. The coefficient β^h is the parameter of interest.

Given that the outcome is a long-difference, each β^h recovers the cumulative impulse response of the outcome to an import-penetration shock (Òscar Jordà, 2005, 2023). We prefer this formulation for two reasons. First, it targets the object of economic interest directly: the total change in the aggregate markdown accumulated by horizon h, rather than its level at $t+h$, so the sequence $\{\beta^h\}$ reads as a cumulative response that is additive in intuition. Second, re-basing every horizon to the common pre-treatment period $t-1$ nets out the level of the outcome prevailing at the moment of the shock and any pre-existing industry-specific trend embedded in it, which makes coefficients directly comparable across horizons. This is a deliberate specification choice: the local-projection framework of Òscar Jordà (2005, 2023); Jordà and Taylor (2025) validates estimating cumulative effects horizon by horizon in exactly this way, and we adopt the long-difference form as the baseline throughout. When $h = 0$, equation (7) reduces to a first-difference specification, $y_{j,t} - y_{j,t-1}$, relating the contemporaneous change in the outcome to import penetration. For $h > 0$, it implements the local projection approach, which traces the cumulative impulse response of each outcome to an import penetration shock over the medium run.

We begin by outlining the assumptions required for a causal interpretation of β^h when estimating equation (7) using OLS. Next, we introduce a shift-share instrumental variable approach to address potential identification threats. Finally, we present the instrumental-variable results.

A. Ordinary Least Squares

The coefficient β^h, if estimated by OLS, admits a causal interpretation under the conditional independence assumption $\mathbb{E}[\varepsilon_{jt}^h \mid IP_{jt}, \mathbf{X}_{jt}, \alpha_j, \delta_t] = 0$: conditional on industry fixed effects,

[18] Domestic production is inferred from export data and the export share of output computed from ENIA: $Q_{jt} = X_{jt}/ExportShare_{jt}$.

year fixed effects, and the control vector, import penetration must be uncorrelated with all unobserved determinants of labor market outcomes.

This assumption is unlikely to hold for three reasons, and the first two threaten identification through channels whose sign cannot be established *a priori*. First, an unobserved negative industry-specific demand shock mechanically raises measured import penetration, as apparent consumption in the denominator contracts faster than imports adjust; its effect on the employment-weighted markdown is itself unsigned, since a domestic contraction can *raise* the aggregate markdown by forcing marginal, low-markdown firms to exit and concentrating employment in high-markdown survivors or *compress* it if weaker labor demand erodes incumbents' wage-setting power. Second, import penetration is itself endogenous to labor market power through political economy (Trefler, 1993; Grossman and Helpman, 1994; Goldberg and Maggi, 1999). Politically organized industries obtain protection and hence lower import penetration, but whether protection accrues to sectors with *high* labor market power (concentrated, rent-rich sectors that lobby most effectively) or with *low* labor market power (declining, comparative-disadvantage sectors seeking shelter) is unclear. Additionally, monopsonistic firms may sort into less trade-exposed industries (Lindenlaub et al., 2022), though this largely static channel is absorbed by the industry fixed effects and the long-difference specification and bites only through differential pre-trends. Because neither channel admits a firm prior on its sign, we do not attempt to sign the net OLS bias.

Third, markdowns are estimated rather than directly observed: $\hat{\nu}_{jt}$ is a function of production-function estimates $(\hat{\theta}^{l}_{jt}, \hat{\theta}^{v}_{jt})$ and observed revenue shares. The idiosyncratic production-function shock ε_{it} contaminates the estimated output elasticities, so the resulting measurement error in $\hat{\nu}_{jt}$ is not trivial, as it inherits a component that is correlated with the structural error ε^{h}_{jt}. Under classical (uncorrelated) measurement error, OLS estimates remain consistent but the variance of $\hat{\nu}_{jt}$ is inflated, reducing precision. Under the more plausible correlated case, OLS is inconsistent and we require an instrument to filter out this source of bias. To address all three concerns jointly, we implement the shift-share instrumental variable approach described next.

B. Shift-Share Instrumental Variables

We address endogeneity and measurement error concerns by exploiting quasi-random variation in industries' exposure to foreign supply-side shocks. The idealized experiment behind the design randomizes each origin's third-market export surge across the industries that source from it: were those shifts literally assigned at random, the industry-specific demand shock that mechanically inflates measured import penetration (the leading OLS confounder identified in Subsection A) would be orthogonal to the instrument by construction. Following Borusyak et al. (2022), we implement a shift-share instrumental variables (SSIV) design that approximates this experiment: its identifying power comes from the quasi-random assignment of many origin-country export *shifts* rather than from the exogeneity of the import shares, a distinction we develop in Subsection C. The *shares* are constructed from the pre-determined composition of Chilean imports by origin country. For the *shift*, we use the expansion of origin countries' export capacity to destinations other than Chile and Latin America, reflecting supply-side developments (productivity growth, factor accumulation, trade liberalization in origin countries)

that are plausibly exogenous to Chilean industry-specific conditions. Formally, the instrument for industry j at time t is:

$$Z_{jt} = \sum_{o \in \mathcal{I}_j} s_{oj} \cdot \ln\left(\frac{E^{-LAC}_{o,jt}}{E^{-LAC}_{o,j,1989}}\right), \tag{9}$$

where $\mathcal{I}_j$ is the set of origin countries that supplied industry j in the base period,[19] s_{oj} is the (renormalized) share of origin country o in total Chilean imports of industry j, averaged over the base period 1985–1989, and $E^{-LAC}_{o,jt}$ is total exports from country o in industry j to all destinations excluding Chile and Latin American and Caribbean countries.[20]

1. Import shares (s_{oj})

The share component captures the pre-existing structure of Chilean import sourcing. We compute s_{oj} as the share of origin country o in total Chilean imports of industry j, averaged over the base period 1985–1989 and renormalized over the support $\mathcal{I}_j$:

$$s_{oj} = \frac{\sum_{t=1985}^{1989} M_{ojt}}{\sum_{t=1985}^{1989} M_{jt}}, \qquad M_{jt} = \sum_{o \in \mathcal{I}_j} M_{ojt}, \tag{10}$$

where M_{ojt} is the value of Chilean imports from country o in industry j at time t. Within the support $\mathcal{I}_j$ the renormalized shares sum to one by construction. By fixing shares in the pre-liberalization base period (after the unilateral tariff reforms of the late 1970s but before the wave of bilateral trade agreements beginning in the 1990s), the instrument avoids contamination from endogenous reorientation of import sourcing in response to subsequent trade shocks.

2. Export shifts ($E^{-LAC}_{o,jt}$)

The shift component captures supply-side developments in origin countries that expand their export capacity globally. For each origin country o in the support $\mathcal{I}_j$ and each industry j, we aggregate total exports to all destinations excluding LAC countries. The exclusion of Chilean and Latin American destinations purges the shift of demand-side influences originating in Chile's own economic conditions or those of geographically proximate and structurally similar economies. This is the same class of variation the China-shock literature exploits: Autor et al. (2013) isolate the labor-market impact of import competition using the growth of Chinese exports to *other* high-income markets, supply-driven expansion in a single origin's export capacity, plausibly orthogonal to destination-country demand. We generalize that logic beyond any one origin, letting the third-market export growth of *all* of Chile's pre-liberalization suppliers drive the shift; Chinese and broader East Asian supply shocks specific to the period are instead absorbed by a dedicated control (Subsection B), so the identifying variation is this broad, origin-specific

[19] Formally $\mathcal{I}_j = \{ o : E^{-LAC}_{o,j,1989} > 0 \}$, the origins with strictly positive 1989 exports in industry j. Restricting the sum to this support, and renormalizing the shares over it.

[20] The LAC exclusion set comprises Argentina, Bolivia, Brazil, Chile, Colombia, Costa Rica, Cuba, the Dominican Republic, Ecuador, El Salvador, Guatemala, Honduras, Mexico, Nicaragua, Panama, Paraguay, Peru, Uruguay, and Venezuela. Countries that ceased to exist during the sample period (Czechoslovakia, Yugoslavia, and the Netherlands Antilles) are excluded from the share computation.

component rather than the China shock alone.

3. Shift-share controls

The control vector $\mathbf{X}_{jt}$ is held fixed across the trade leg and is designed to absorb the confounders most likely to move with exposure to foreign supply shocks. It groups into three blocks. First, industry *pre-trends*. We include the five-year change (measured over the window five to ten years before t, so that each control is fixed well ahead of the shock) in the industry's (log) payroll-weighted labor-market Herfindahl index, exports, (average) capital share, output, female labor share, and real wage per worker. Conditioning on these prior trajectories, rather than on contemporaneous levels, absorbs the slow-moving industry dynamics in labor-market power, factor intensity, scale, workforce composition, and wages that were already under way before the shock and could otherwise drive markdown dynamics independently of import penetration.

Second, we control for two contemporaneous *global trends* that vary at the industry-year level: advanced-economy China-linked import exposure, which absorbs the common East-Asian supply shock of the period, and a world-export-market-share-weighted average of origin-country TFP growth, which nets out worldwide technical change in each product line. Together these isolate the origin-country variation specific to Chile's own base-period suppliers.

Third, and constructed analogously to the instrument, we include *shift-share-weighted origin-country fundamentals* in change form (the exposure-weighted changes in origin countries' log real GDP, labor share, investment share, log human capital, log terms of trade, WTO membership, and log exchange rate), built on the same base-period sourcing shares s_{oj} as the instrument, so that variation in Z_{jt} is not proxying for common macroeconomic movements transmitted through the identical sourcing composition.

Inference must confront a feature specific to shift-share designs: the instrument's shocks are correlated across industries that share a similar origin composition. Two industries that both source heavily from the same countries are exposed to the same origin-level export shocks, so their instrument realizations, and hence their regression residuals, are not independent, even after conditioning on the controls. Borusyak et al. (2022) show that conventional standard errors clustered on the outcome (industry) dimension ignore this cross-industry dependence and understate sampling uncertainty, and that valid inference must instead treat the origin-by-time shocks as the effective units of observation. We cluster at the origin level because the export-supply shifts are common to every industry sourcing from a given origin: the origin is the level at which shock realizations, and therefore residuals, are correlated. Because our identifying power derives from a finite set of such shocks (concentrated, as Appendix B documents, in a handful of large origins), this correction is particularly relevant. We therefore report Borusyak et al. (2022) exposure-robust standard errors, computed at the shock level and clustered on the shock origin, throughout.[21]

C. Identification

Two assumptions must be satisfied for a valid causal interpretation of β^h. The instrument must shift import penetration (relevance), and the origin-country export shocks it is built from must

[21] Adão et al. (2019) provides an alternative exposure-robust standard error estimation.

be exogenous to the markdown innovations they are used to explain (exclusion). The last one, in the case of local projection, is also named lead-lag exogeneity (Stock and Watson, 2018) and turns on a stronger assumption than conventional exclusion. They are discussed in this subsection.

RELEVANCE. Because the identifying variation is a shift-share of origin-country export shocks, we assess strength at the level of those shocks, following Borusyak et al. (2022), and also at the industry-year level. Table 2 reports the exposure-robust shock-level first stage, adding the control blocks cumulatively. In each column import penetration is residualized on that column's fixed effects and control block, aggregated to origin-by-industry-by-year cells (68,557 cells), and regressed on the raw origin-by-industry export shock, with standard errors clustered on the 83 shock origins[22]. The Kleibergen–Paap F ranges from 19.7 down to 8.8 in the preferred specification, comfortably rejecting the non-relevance null throughout.

TABLE 2—First Stage: Import Penetration on the Shift-Share Instrument, Shock-Level and Industry-Level

	(1)	(2)	(3)	(4)
Panel A: Shock-level first stage (exposure-robust)				
Export shock g_{ojt}	0.036***	0.036***	0.035***	0.026***
	(0.009)	(0.009)	(0.008)	(0.009)
Shock cells	68,557	68,557	68,557	68,557
Panel B: Industry-level first stage (exposure-robust)				
Shift-share instrument Z_{jt}	0.354***	0.358***	0.358***	0.299***
	(0.084)	(0.088)	(0.081)	(0.101)
Observations	1,815	1,815	1,815	1,815
Year FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
Lagged-trend controls	No	Yes	Yes	Yes
Global-trend controls	No	No	Yes	Yes
Shift-share controls	No	No	No	Yes
Kleibergen–Paap *F* (shock level)	17.8	16.5	19.7	8.8
p-value	0.000	0.000	0.000	0.004

Notes: First stage of import penetration on the shift-share instrument, at two aggregation levels, both under exposure-robust inference (Borusyak et al., 2022). *Panel A* is the shock-level regression on the raw origin × industry export-supply shock g_{ojt}, weighted by exposure and clustered on the 83 shock origins ("Shock cells"). *Panel B* is the location-level regression on the assembled instrument Z_{jt} at CIIU×año cells ("Observations"=1,815). The second-stage IV coefficient is numerically identical across panels; the first-stage slope is not, so Panel B's standard error is Panel A's origin-clustered error rescaled to Panel B's coefficient scale, no single origin cluster being defined at the CIIU×año grain. The Kleibergen–Paap F is computed once at the shock level. Columns add year and industry fixed effects, lagged-trend, global-trend, and shift-share origin-fundamental controls cumulatively. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[22] Because these shocks are indexed at the origin×industry×year level and load on relatively few origins, we cluster on the 83 shock origins and report exposure-robust standard errors (Borusyak et al., 2022). Weights imply about fourteen effective origins (Herfindahl ≈ 0.072), with the United States the largest single share by a wide margin (Appendix 2); this concentration bears on inference, not on consistency, which instead rests on some 2,984 origin×industry×year shock cells. Fourteen effective origins are below the Herfindahl of 1/20 at which Borusyak et al. (2022) validate the finite-sample coverage of exposure-robust inference in their own Monte-Carlo simulations.

EXCLUSION. This is the substantive requirement, and in a shift-share design the identifying burden can be placed on either of the two objects that constructs the instrument. Under share-exogeneity (Goldsmith-Pinkham et al., 2020), each initial import share s_{oj} acts as a separate instrument and must be uncorrelated with the latent determinants of subsequent markdown dynamics. That requirement is demanding here, because the base-period sourcing composition is not randomly assigned: an industry sources more heavily from a given origin because of its own technology, product mix, position on the quality ladder, input requirements, and incumbent supplier relationships, structural traits that plausibly also shape how its markdowns evolve, so that $\mathbb{E}[s_{oj}\,\varepsilon_{jt+h}] \neq 0$. Fixing the shares in the pre-liberalization base period (1985–1989) predates the bilateral agreements whose effects we estimate, but it does not make the shares exogenous.

We therefore rely instead on shift-exogeneity (Borusyak et al., 2022, 2025), which requires that the origin-country export shocks be as-good-as-randomly assigned conditional on controls and, crucially, permits the shares to be endogenous. This is plausible here because the identifying variation comes from differential growth in origin countries' export capacity to non-Latin American destinations, driven by those countries' evolving comparative advantage rather than by developments in Chilean manufacturing. Industry fixed effects absorb all time-invariant markdown determinants; year fixed effects remove shocks common to all industries, including economy-wide trade policy changes and, crucially, the nationwide institutional reforms that drive the secular trend documented in Subsection B. The LP-IV estimates therefore identify the differential effect of import penetration within industries, holding the aggregate institutional path fixed, and absorbing time-invariant industry characteristics.

The local-projection form sharpens what exogeneity must deliver. Valid identification requires the lead-lag condition of Stock and Watson (2018), $\mathbb{E}\left[Z_{jt}\,\varepsilon_{jt+h}\right] = 0$ for all $h \in \{0, \ldots, 8\}$: the instrument must be orthogonal not only to the contemporaneous error but to the entire future path of markdown innovations, of which the standard contemporaneous exclusion (that Z_{jt} affects y_{jt} only through IP_{jt}) is the special case $h = 0$. The threats to this condition are slow-moving confounders set off by the same global forces that generate the shifts and transmitted to Chilean labor markets with a lag through non-trade channels. The control vector is built to close these channels: the *global trends* block of $\mathbf{X}_{jt}$ absorbs threats operating through technical change and competition in the foreign markets of Chilean exporters. The shift-share-weighted *origin fundamentals* net out partner-country macroeconomic conditions transmitted through the same sourcing composition. The *own-industry lagged trends* capture internal confounders particular to each industry. In essence, identification requires the shocks to be conditionally assigned as if at random, and the testable shadow of that requirement is pre-trend balance. We test shift-exogeneity directly, in the spirit of Borusyak et al. (2022). Table 3 reports the balance exhibit the design's inference rests on.

Panel A regresses ten origin-country fundamentals on the normalized export shock; Panel B regresses eleven predetermined industry characteristics, each at its most recent fully-predetermined lag, on the normalized instrument. Both panels are estimated at the origin-by-industry-by-year grain at which exposure-robust inference is defined, 67,483–68,672 shock cells in Panel A and 69,049 in Panel B, weighted by base-period exposure shares. Each panel is estimated first with the baseline absorption (industry and year fixed effects plus the two global-trend confounds carried

in the headline specification) and then adding the industry-level control blocks the headline specification itself estimates in. The verdict is broad balance with two localized, surviving exceptions, and we state both plainly. Thirty-six of the forty-two coefficient cells are indistinguishable from zero at the ten-percent level, only two characteristics are significant at five percent and, most important for the outcome at hand, the two characteristics closest to the markdown itself, payroll-weighted labor-market concentration and the log real wage per worker, are balanced in every cell, providing robustness to the LP estimates against those concerns.

TABLE 3—Shift-Share Instrument Balance

	Baseline		+ Industry Controls	
Variable	Coef.	SE	Coef.	SE
Panel A: Origin-Level Fundamentals on the Export Shock				
Log Real Gross Domestic Product, Origin Country (PWT 10.01)	-0.032	(0.135)	-0.033	(0.135)
Log Human Capital Index, Origin Country (PWT 10.01)	-0.103	(0.070)	-0.103	(0.070)
Log Physical Capital per Worker, Origin Country (PWT 10.01)	-0.148	(0.102)	-0.148	(0.103)
Log Terms of Trade, Export Price over Import Price (Origin Country)	-0.084	(0.091)	-0.083	(0.091)
Log Real Exchange Rate, Origin Country (PWT 10.01)	-0.120	(0.081)	-0.120	(0.081)
Log Labor Productivity, Output per Worker (Origin Country)	-0.187	(0.115)	-0.187	(0.115)
Log Total Population, Origin Country (PWT 10.01)	0.062	(0.124)	0.062	(0.125)
Labor Share of Output, Origin Country (PWT 10.01)	0.096	(0.074)	0.096	(0.072)
Investment Share of Output, Origin Country (PWT 10.01)	0.120	(0.078)	0.121	(0.079)
Trade Openness, Absolute Imports Plus Exports over GDP †	-0.177**	(0.081)	-0.176**	(0.080)
Observations	67,483–68,672		67,483–68,672	
Panel B: Industry-Level Characteristics on the Instrument				
Mean Industry Output, Employment-Weighted (One-Year Lag)	0.169	(0.151)	0.150	(0.140)
Mean Industry Employment, Plant-Level Sum (One-Year Lag)	0.323*	(0.196)	0.293	(0.185)
Mean Female Employment Share, Industry Average (One-Year Lag)	0.022	(0.016)	0.018	(0.016)
Mean Unskilled-Worker Employment Share (One-Year Lag)	-0.006	(0.103)	0.005	(0.101)
Mean Capital Share of Value Added, Industry Average (One-Year Lag)	-0.103***	(0.037)	-0.126***	(0.036)
Mean Labor Share of Value Added, Industry Average (One-Year Lag)	0.007	(0.034)	-0.006	(0.045)
Mean Share of Plants Importing Intermediate Inputs †	0.004	(0.069)	0.005	(0.062)
Mean Share of Plants Exporting Final Output †	0.107*	(0.064)	0.085	(0.067)
Mean Skilled-Worker Wage-Bill Share, Industry Average	0.060	(0.124)	0.049	(0.123)
Labor-Market Concentration Index, Payroll-Weighted Herfindahl	-0.106	(0.100)	-0.131	(0.107)
Log Real Wage per Worker, Industry Average	0.054	(0.046)	0.047	(0.050)
Observations	69,049		69,049	

Notes: Panel A regresses (standardized) origin-country fundamental on the (standardized) export shock; Panel B regresses (standardized) predetermined CIIU characteristic (L_1) on the (standardized) shift-share instrument. *Baseline*: CIIU and year fixed effects. + *Industry Controls*: Panel A excludes the shift-share recentering controls as circular; Panel B includes them. Both panels weighted by base-period exposure shares. Standard errors clustered two-way on origin and the 3-digit CIIU sector. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D. Results

We present results from the LP-IV specification described in equation (7), estimated over an eight-year horizon. All specifications are employment-weighted and include industry and year fixed effects and the full set of controls.

Each figure reports two representations of the impulse response. The faint dots show the raw horizon-by-horizon IV coefficients, estimated independently at each h. The solid line and shaded confidence band show the smooth local projection (SLP), which imposes that the impulse response is a smooth function of the horizon by approximating the coefficient sequence with

penalized B-splines. The SLP methodology, developed by Barnichon and Brownlees (2019) for time-series applications and extended to panel settings by Hartley and Mejia (2025), addresses a well-known limitation of standard local projections: because each horizon is estimated as a separate regression, raw LP coefficients can exhibit substantial variability from one horizon to the next, producing "lumpy" impulse responses that are difficult to interpret economically. The B-spline penalty shrinks the coefficient path toward a low-order polynomial[23], trading a small bias for substantial efficiency gains. Hartley and Mejia (2025) show via Monte Carlo simulations that smooth panel local projections dominate standard panel LP in mean-squared error as long as the researcher has some prior confidence in the smoothness of the true impulse response, a condition that is natural in our setting, where economic theory predicts gradual adjustment to trade shocks rather than erratic horizon-to-horizon jumps. We report both representations to allow the reader to assess the sensitivity of our conclusions to the smoothing procedure: the raw coefficients confirm that the smoothed paths do not impose qualitative patterns absent from the underlying estimates.

1. Import Penetration and Wage Markdowns

Figure 6 plots the impulse response of employment-weighted markdowns to a one-percent increase in import penetration. The contemporaneous effect is small and statistically indistinguishable from zero: the smoothed point estimate at $h = 0$ is approximately 0.02, with the exposure-robust 90% confidence interval spanning both sides of zero. The effect remains near zero through the first two years, then turns positive and becomes statistically significant starting at $h = 4$, building steadily to a peak near 0.13 at $h = 6$ before easing to roughly 0.07 by $h = 8$. From $h = 4$ through $h = 8$ the exposure-robust confidence band lies entirely above zero, indicating a statistically significant positive effect that persists through the end of the horizon.

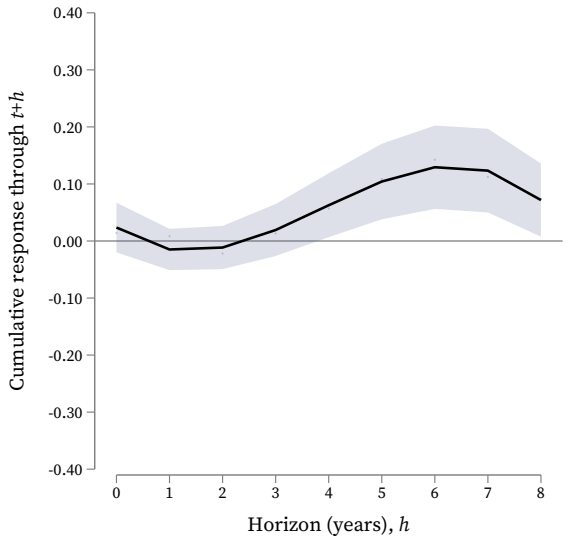
The delayed response is the central empirical finding of this paper. Import penetration does not raise monopsony power on impact. Instead, it sets in motion a sequence of adjustments (in product markets, firm selection, and labor reallocation) that over three to eight years gradually concentrates employment in firms with greater wage-setting power. The mechanisms underlying this buildup are examined next.

2. Mechanisms

Figure 7 reports impulse responses for the outcomes most directly connected to the channels through which import penetration affects markdowns; each panel isolates a distinct margin of the adjustment. The clearest and most robust response is the compression of markups. Panel (a) shows an elasticity of about −0.07 on impact that deepens roughly monotonically to about −0.24 by $h = 8$. In level terms, a 10 percent increase in import penetration compresses price-cost margins by about 0.7 percent on impact, widening to roughly 2.4 percent at the end of the horizon. This is a sustained erosion of product-market rents that continues to deepen rather than stabilizing. The exposure-robust confidence band excludes zero at every horizon $h = 0$ through $h = 8$, so markup compression is the one mechanism outcome that is precisely estimated throughout. The pattern rules out a transitory competitive shock: import penetration operates

[23] In this case, third-degree polynomials.

FIGURE 6—Import Penetration and Wage Markdown: LP-IV Impulse Response



Notes: The figure plots smoothed LP-IV estimates of the effect of import penetration on log employment-weighted markdowns at horizons $h = 0, \ldots, 8$. Shaded area: smoothed 90% confidence interval under exposure-robust (Borusyak-Hull-Jaravel, 2022) shock-level standard errors. Faint dots: raw horizon-by-horizon IV coefficients. All specifications include industry and year fixed effects, the full control set, and employment weights.

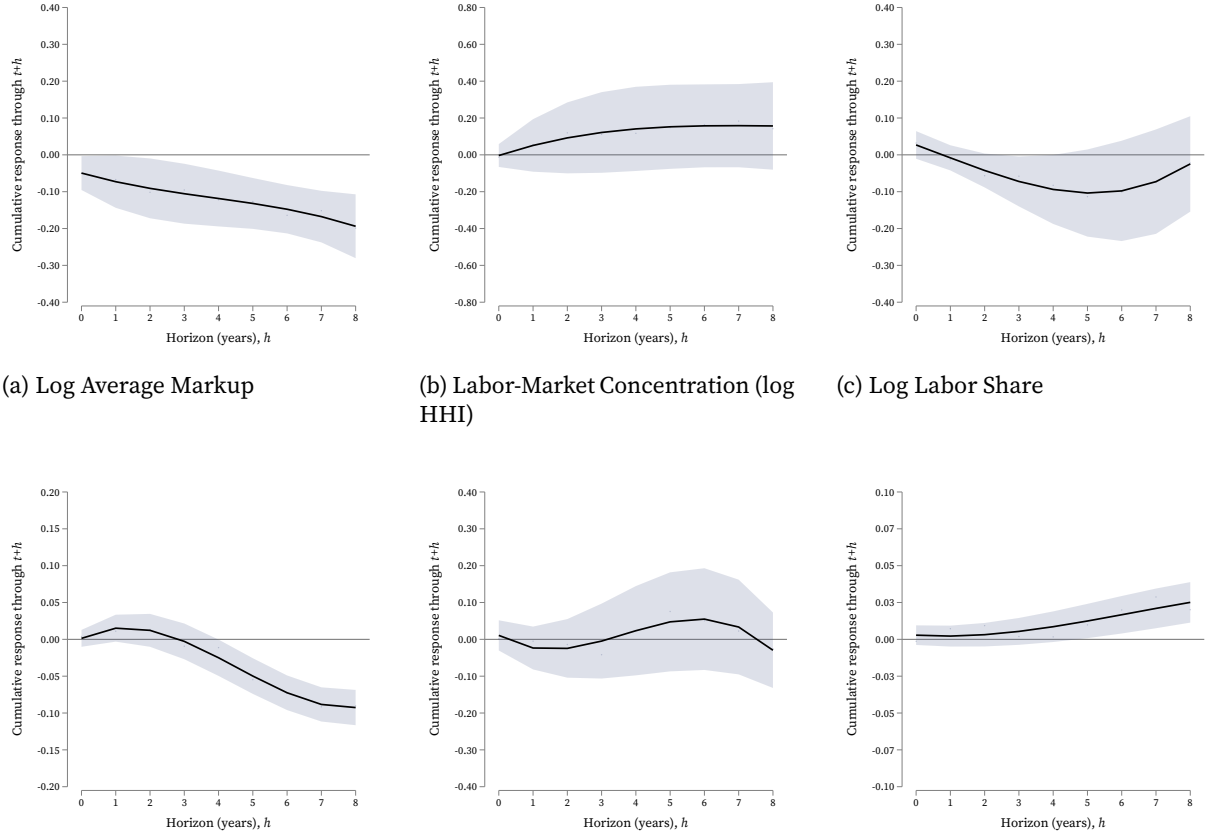
as a sustained pro-competitive force in the final-good market, progressively squeezing margins as foreign varieties establish market presence.

Labor-market concentration rises, but imprecisely. Panel (b) shows the log employment Herfindahl on a positive point path of roughly +0.10 to +0.15 from the first year onward, consistent with import competition concentrating employment among fewer employers within the local labor market, the reallocation margin through which surviving firms would acquire wage-setting power. Given the clear pattern of the IRF and marginal statistical insignificance, we interpret this pattern as a consequence of lack of statistical power rather than a meaningless estimate. Economic theory points in the same direction.

Labor's share of output is compressed over the middle horizons, consistent with monopsony rents expanding at labor's expense. Panel (c) shows a response near zero on impact that turns negative and is statistically significant from roughly $h = 1$ through $h = 5$, reaching about −0.13 around $h = 5$ before the band widens and the estimate loses significance at the longest horizons. The timing (compression concurrent with the markup decline, fading as the horizon lengthens) corroborates the reallocation reading, though we do not lean on the long-run magnitude, which is imprecisely estimated.

Mean wages decline significantly, with a lag that matches the rest of the sequence. Panel (d) is indistinguishable from zero through $h = 4$. A small, insignificant increase of about +0.01 at $h = 1$–2 gives way to a small, insignificant decline by $h = 3$–4. The estimate then turns sharply and significantly negative from $h = 5$ onward, at about −0.05 at $h = 5$ and deepening to −0.09 by $h = 7$–8. The exposure-robust band excludes zero at every horizon from $h = 5$ through $h = 8$. This is not merely the wage decline implied by the rising markdown; it is direct evidence that the

FIGURE 7—Local Projections: Mechanism Outcomes



(a) Log Average Markup

(b) Labor-Market Concentration (log HHI)

(c) Log Labor Share

(d) Log Average Wage

(e) Log Average Employment

(f) Probability of Exit

Notes: Each panel plots smoothed LP-IV estimates of the effect of import penetration on the indicated outcome at horizons $h = 0, \ldots, 8$. Shaded areas: smoothed 90% exposure-robust (Borusyak-Hull-Jaravel) confidence intervals, computed at the level of the origin-country export shocks. Faint dots: raw horizon-by-horizon IV coefficients. All specifications include industry and year fixed effects, the full control set, and employment weights.

same reallocation compressing margins and lifting concentration is, with a lag, pulling down the labor-weighted average wage.

Average plant employment is imprecisely estimated. Panel (e) shows a brief early dip (about −0.03 at $h = 2$) followed by a hump peaking near +0.09 at $h = 7$ before easing, with exposure-robust bands that include zero throughout. The point path carries the signature of selection (survivors expanding as marginal plants contract) but is not statistically significant, so we do not lean on the average-size margin.

The exit rate rises significantly, but only late. Panel (f) is essentially zero through $h = 4$ and touches significance only marginally at $h = 2$ and $h = 5$. It then clears zero cleanly from $h = 6$ through $h = 8$, rising to about +0.02 at $h = 6$, peaking near +0.03 at $h = 7$, and easing only slightly to +0.02 by $h = 8$. The exposure-robust band excludes zero throughout that window. Competitive selection, where present, arrives as a delayed attrition rather than a sharp early shakeout, and by $h = 6$–8 this margin is robustly identified.

Taken together, the mechanism evidence points to three robustly identified channels and

two corroborating but noisier ones. Markup compression (panel a) is precisely estimated and significant at every horizon under exposure-robust inference: import competition erodes product-market rents throughout. The exit rate (panel f) rises cleanly and significantly from $h = 6$ through $h = 8$, and mean wages (panel d) turn significantly negative over the same late window, $h = 5$ through $h = 8$, both channels operating on the same delayed clock as the selection mechanism, and both, unlike markup compression, silent until reallocation has had several years to work. The compression of labor's share (panel c) is significant over the middle horizons and corroborates the same reading. Rising labor-market concentration (panel b) carries the expected sign, but its exposure-robust band never excludes zero; average plant size (panel e) is the one true null over our window. The markdown response of Figure 6 (near zero on impact, significantly positive from $h = 4$ onward) emerges from the cumulative operation of these forces, in the tradition of Melitz (2003), as surviving firms acquire greater monopsony power.

This delayed, positive relationship between import penetration and monopsony power apparently contrasts with a recent literature that finds trade liberalization *reduces* labor-market power in China, where input-tariff cuts following WTO accession lowered firm-level wage markdowns (Kondo et al., 2024), narrowed markdown dispersion across firms (Xie et al., 2024), and lowered wage-setting power even as they raised product-market markups (Dobbelaere and Wiersma, 2025). We read this contrast as reflecting a difference in the economic force each shock captures rather than a contradiction: the underlying mechanism under input-tariff liberalization assessed by those papers is through reductions in the firms' costs and is a supply-side relief that need not displace incumbents, because it does not reflect competition in the final good. The import-penetration shock we study is output-market competition that displaces less productive domestic producers and reallocates employment toward larger, higher-markdown survivors, the reallocation channel that the mechanism panels trace through robust markup compression and, more tentatively, the rising concentration of panel (b). This reallocation channel parallels the exit-driven mechanism documented for Brazil's own early-1990s trade liberalization by Felix (2026), where concentration likewise rose through firm exit and worker reallocation toward surviving importers rather than through a contemporaneous wage response.

That markdowns respond only with a multi-year lag identifies monopsony power as an indirect consequence of the firm selection and reallocation that import competition induces, rather than of import competition per se. This is precisely the channel that the structural model of Section VI formalizes: a fall in trade costs reallocates employment toward high-productivity, high-payroll-share firms, lowering their firm-level labor-supply elasticity (equation (18)) and raising their markdowns. We do *not* use the long-run LP-IV elasticity as a calibration target. We instead reserve it as an out-of-sample check: the model is calibrated to match average import penetration and worker take-home share, and we then ask whether the model-implied response of the aggregate markdown to a marginal change in trade costs reproduces the reduced-form elasticity it never saw. The calibrated model implies an elasticity of 0.147: the same sign as, and comfortably within the exposure-robust confidence interval of, the cumulative $h = 8$ estimate (band $[0.015, 0.133]$), though in its lower quartile and about a third of the reduced-form peak. Because the model is static it delivers a single elasticity, whereas the reduced-form path starts near zero and rises above the static value once reallocation and selection have had time to operate,

peaking near 0.13 at $h = 6$ before easing back toward the static value by $h = 8$. The two are not commensurable objects (a comparative static against a cumulative long-difference impulse response), so this is a sign and order-of-magnitude check rather than a moment match. The gap over the middle horizons is itself informative: it is the empirical signature of the entry/exit margin that the static model abstracts from. This comparison, wholly outside the calibration moment set, connects the reduced-form evidence to the welfare analysis.

Finally, the trade-induced markdown response is concentrated where the selection mechanism predicts the most reallocation, but only over the medium run. Splitting the shift-share LP-IV on pre-determined (1990–92) industry characteristics, the implied markdown response for skill-intensive industries is indistinguishable from zero on impact, and the effect emerges only at longer horizons (0.199, standard error 0.057, at $h = 8$). The split on base-period labor-market concentration follows a very similar shape. Those results seem meaningful under the logic that skill-intensive industries could be paying higher wages and, therefore, can be far from employment protection caps like the minimum wage that indirectly sets a *maximum* wage markdown. By the concentration interpretation, one could interpret those results as consistent with the idea that for import penetration to have an impact on labor market power, an industry has to be relatively far from being competitive at the beginning. We report the full interacted term-structures, split-sample impulse responses, and diagnostics in Appendix 3. These same predetermined characteristics also carry the institutional heterogeneities documented in Appendix C.

VI. Model and Welfare Analysis

Up to this point we have documented the level and evolution of labor market power, its aggregate dynamics, and its causal link to trade exposure. Three questions remain open: how do trade liberalization and institutional reform interact in general equilibrium, what would the labor market have looked like under a less aggressive liberalization path, and what were the welfare stakes of the realized trajectory? Answering them requires a structural framework that connects the reduced-form estimates to welfare. This section develops the model, constructs a counterfactual trade path using the synthetic control method, and evaluates welfare under alternative scenarios.

The key model mechanism runs through the firm-level labor-supply elasticity. As trade liberalizes, employment reallocates toward high-productivity exporting firms, whose payroll shares within local labor markets increase. Through the nested-CES labor supply structure, a higher payroll share shifts weight toward the less elastic across-market margin, lowering the firm-level supply elasticity and raising monopsony power. The role of bargaining is to counteract this effect: a higher bargaining weight β raises workers' take-home share of the marginal revenue product, compressing the markdown wedge regardless of the market structure the firm faces. The interplay between these two forces (trade concentrating wage-setting power, institutions dispersing it) is what the model is designed to characterize.

A. A Model of Oligopsony, Bargaining, and Trade

We develop a static model of a small open economy in which wages are determined by Nash bargaining between oligopsonistic firms and firm-level unions. The framework nests Berger

et al. (2022a) as the special case $\beta = 0$ and extends it by embedding the oligopsony structure in an open economy with endogenous exporter status. A growing literature combines oligopsony and trade without a bargaining channel: Pham (2023) studies how China's WTO-era tariff liberalization affected labor-market distortion in a heterogeneous-firm oligopsony model, and Jha et al. (2026a) embed monopsonistic labor markets in a dynamic small-open-economy model with an endogenous exporter margin. Neither includes firm-level wage bargaining or an institutional channel; the joint bargaining-trade-institutions structure developed here is, to our knowledge, novel. We present the environment, the wage-setting protocol, the trade structure, and the aggregation that maps firm-level allocations into welfare; step-by-step derivations, the model's graphical schematics, and the numerical solution algorithm are collected in Appendix E.

Environment. The economy is static and consists of a representative household and a finite set of firms producing differentiated varieties. Labor is the sole factor of production. Varieties are indexed by i within a local labor market j; market j contains M_j firms, all of which produce (there is no entry, exit, or fixed-operating-cost margin[24]), and each firm bargains over its wage with a firm-level union of bargaining power $\beta \in [0, 1]$. Each firm's only discrete decision is whether to serve the domestic market only or also export, determined in equilibrium below.

Households. The representative household maximizes quasi-linear utility

$$U = C - \bar{\varphi}^{-\frac{1}{\varphi}} \frac{N^{1+1/\varphi}}{1+1/\varphi}, \tag{11}$$

where φ is the Frisch elasticity of labor supply, $\bar{\varphi}$ a labor-disutility shifter, and aggregate consumption is Cobb–Douglas over all varieties,

$$C = \prod_j \prod_{i \in \mathcal{F}_j} \left(\frac{c_{ij}}{\phi_{ij}} \right)^{\phi_{ij}}, \qquad \sum_j \sum_{i \in \mathcal{F}_j} \phi_{ij} = 1, \tag{12}$$

subject to the budget constraint

$$\sum_{ij} p_{ij} c_{ij} = \text{Income} = \sum_{ij} \pi_{ij} + \sum_{ij} w_{ij} n_{ij}, \tag{13}$$

where p_{ij} is the consumer price of variety ij, π_{ij} is firm profit, and $w_{ij} n_{ij}$ is labor income paid by the firm. Every variety enters the consumption basket regardless of whether it is produced domestically or imported[25]. Only its price and source of supply differ. Quasi-linearity and the Cobb–Douglas aggregator jointly imply that demand for each variety is

$$c_{ij} = \frac{\phi_{ij}}{p_{ij}} \cdot \text{Income}. \tag{14}$$

[24] This is the $f = 0$ limit of a model with a fixed operating cost and a minimum-wage floor, adopted because the calibration data (Subsection C) already condition on operating plants and the object of interest is reallocation across *incumbent* firms rather than net firm dynamics; the abstraction is discussed in Appendix F.

[25] Each variety is assumed to be locally produced by only one national firm that produces it, or imported from abroad. All varieties can be imported from a foreign supplier.

Labor supply. The household supplies labor only to domestic producers. Aggregate labor is a nested CES composite over firms,

$$N = \left(\sum_j N_j^{\frac{1+\theta}{\theta}} \right)^{\frac{\theta}{1+\theta}}, \qquad N_j = \left(\sum_{i \in \mathcal{F}_j} n_{ij}^{\frac{1+\eta}{\eta}} \right)^{\frac{\eta}{1+\eta}}, \tag{15}$$

with $\eta > \theta > 0$, so that substitution across firms within a market is easier than substitution across markets. The corresponding wage indices are

$$W = \left(\sum_j W_j^{1+\theta} \right)^{\frac{1}{1+\theta}}, \qquad W_j = \left(\sum_{i \in \mathcal{F}_j} w_{ij}^{1+\eta} \right)^{\frac{1}{1+\eta}}. \tag{16}$$

Optimal aggregate labor supply equates the marginal disutility of work to the wage index, $W = \bar{\varphi}^{-1/\varphi} N^{1/\varphi}$, equivalently $N = \bar{\varphi} W^{\varphi}$. Firm-level labor supply then follows from the nested-CES allocation,

$$n_{ij} = \bar{\varphi} \left(\frac{w_{ij}}{W_j} \right)^{\eta} \left(\frac{W_j}{W} \right)^{\theta} W^{\varphi}. \tag{17}$$

The firm-level labor supply elasticity is endogenous and determined by

$$\frac{1}{\epsilon_{ij}} = \frac{1 - s_{ij}}{\eta} + \frac{s_{ij}}{\theta}, \tag{18}$$

where

$$s_{ij} = \frac{w_{ij} n_{ij}}{\sum_{k \in \mathcal{F}_j} w_{kj} n_{kj}} \tag{19}$$

is firm ij's share of market j's wage bill. A larger payroll share shifts weight toward the less elastic across-market margin θ, lowering ϵ_{ij} and raising monopsony power: when trade-induced reallocation shifts employment toward high-productivity, high-share firms, their labor-supply elasticities fall and average monopsony power rises, even though the set of firms is fixed. This is the channel through which the reduced-form reallocation evidence of Section V operates in the model.

Firms. Each firm $i \in \mathcal{F}_j$ has productivity $a_{ij} > 0$ and produces with the linear technology $y_{ij} = a_{ij} n_{ij}$. Productivity decomposes as $a_{ij} = z \tilde{a}_{ij}$, where $\tilde{a}_{ij} = \exp(\hat{\omega}_{ij})$ is the firm-specific, unit-scale productivity draw recovered from the data and z is an economy-wide common scale (the same scale factor recovered in calibration) that converts model magnitudes into Chilean-peso levels. Because z is common across firms, it multiplies every nominal level (the value of output $p_{ij} a_{ij}$, wages, profit, and income) but cancels in every ratio, including the firm-level labor-supply elasticity ϵ_{ij}, the take-home share $\mu_{ij} = (\epsilon_{ij} + \beta)/(1 + \epsilon_{ij})$, and the normalized productivity shares $a_{ij}/\tilde{a}_j$; z therefore affects levels only and plays no role in the moments targeted below. Let p_{ij} denote the price the firm receives per unit of output, so the value of the

marginal product of labor is $p_{ij}\, a_{ij}$. Firms differ only through (a_{ij}, ϕ_{ij}) and their equilibrium export status. There is no fixed cost, and profit equals variable profit, $\pi_{ij} = \pi^{v}_{ij}$, defined in (26) below.

Wage determination. Each firm ij and its union bargain over the wage w_{ij} through the generalized Nash problem:

$$w^{*}_{ij} = \arg\max_{w_{ij}} \left[w_{ij}\, n_{ij}(w_{ij})\right]^{\beta} \cdot \left[(p_{ij}\, a_{ij} - w_{ij})\, n_{ij}(w_{ij})\right]^{1-\beta}, \tag{20}$$

where the union's payoff is labor income, the firm's payoff is variable profit, and disagreement payoffs are zero. The solution is the bargained wage:

$$w^{*}_{ij} = \mu_{ij} \cdot p_{ij}\, a_{ij}, \qquad \mu_{ij} = \frac{\epsilon_{ij} + \beta}{1 + \epsilon_{ij}}, \tag{21}$$

where μ_{ij} is the worker's take-home share of the marginal revenue product (i.e. the inverse of the wage-markdown)[26] Appendix 3 gives the full step-by-step first-order derivation.

PROPOSITION 1 (Additive Decomposition). *The equilibrium markdown admits an additive decomposition:*

$$\mu_{ij} = \underbrace{\frac{\epsilon_{ij}}{1 + \epsilon_{ij}}}_{\textit{Oligopsony}} + \underbrace{\frac{\beta}{1 + \epsilon_{ij}}}_{\textit{Bargaining premium}}. \tag{22}$$

Proof. $(\epsilon_{ij} + \beta)/(1 + \epsilon_{ij}) = \epsilon_{ij}/(1 + \epsilon_{ij}) + \beta/(1 + \epsilon_{ij})$. ∎

COROLLARY 1 (Nesting).

(i) $\beta = 0$ *(pure oligopsony):* $\mu_{ij} = \epsilon_{ij}/(1 + \epsilon_{ij})$*, recovering the standard oligopsony markdown.*
(ii) $\beta = 1$ *(full union power):* $\mu_{ij} = 1$*, so* $w_{ij} = p_{ij}\, a_{ij}$ *(competitive).*
(iii) $\epsilon_{ij} \to \infty$ *(flat supply):* $\mu_{ij} \to 1$ *for any* β.
(iv) $\epsilon_{ij} \to 0$ *(vertical supply):* $\mu_{ij} = \beta$ *(pure rent splitting).*

The additive structure of Proposition 1 delivers the two-channel interpretation that organizes the welfare analysis. The first term, $\epsilon_{ij}/(1 + \epsilon_{ij})$, is the oligopsony component: it rises as labor supply becomes less elastic and falls as competition intensifies, capturing the mechanism through which trade-induced reallocation raises monopsony power. The second term, $\beta/(1+\epsilon_{ij})$, is the bargaining premium: it rises with β and has a bigger marginal effect where monopsony power is strongest (small ϵ_{ij}), so institutional reform has its greatest wage effects precisely in the markets where firms' wage-setting power is most concentrated. Proposition 2 formalizes three implications of this structure that are central to the welfare analysis that follows.

[26] A note on notation. In the empirical sections II–IV, μ_{it} denotes the product *markup* and ν_{it} the wage *markdown*. In the model, μ_{ij} denotes the worker take-home share (the inverse of the markdown, $\mu_{ij} = 1/\nu_{ij}$), a distinct object from the empirical markup. We retain the symbol μ for the take-home share to align with Berger et al. (2022a); the two uses do not overlap within any single section.

PROPOSITION 2 (Inverse Markdown Properties).

(i) Monotonicity in bargaining power: $\dfrac{\partial \mu_{ij}}{\partial \beta} = \dfrac{1}{1+\epsilon_{ij}} > 0.$

(ii) Monotonicity in the supply elasticity: $\dfrac{\partial \mu_{ij}}{\partial \epsilon_{ij}} = \dfrac{1-\beta}{(1+\epsilon_{ij})^2} > 0$ *for* $\beta < 1.$

(iii) Complementarity between market structure and institutions: $\dfrac{\partial^2 \mu_{ij}}{\partial \beta \, \partial \epsilon_{ij}} = -\dfrac{1}{(1+\epsilon_{ij})^2} < 0.$

Proof. Direct differentiation of $\mu_{ij} = (\epsilon_{ij} + \beta)/(1+\epsilon_{ij})$:

$$\frac{\partial \mu_{ij}}{\partial \beta} = \frac{1}{1+\epsilon_{ij}} > 0,$$

$$\frac{\partial \mu_{ij}}{\partial \epsilon_{ij}} = \frac{(1+\epsilon_{ij}) - (\epsilon_{ij}+\beta)}{(1+\epsilon_{ij})^2} = \frac{1-\beta}{(1+\epsilon_{ij})^2} > 0 \quad \text{for } \beta < 1,$$

$$\frac{\partial^2 \mu_{ij}}{\partial \beta \, \partial \epsilon_{ij}} = \frac{\partial}{\partial \epsilon_{ij}}\left[\frac{1}{1+\epsilon_{ij}}\right] = -\frac{1}{(1+\epsilon_{ij})^2} < 0.$$

■

Proposition 1 implies that a firm-level union shifts the equilibrium wage by moving it along the labor supply curve between two polar cases: the pure oligopsony outcome ($\beta = 0$) and the competitive equilibrium ($\beta = 1$); Figure E.1 (Appendix E) traces this movement directly, with both wages and employment rising monotonically in β. Corollary 1 makes the nesting explicit: our model recovers Berger et al. (2022a) as the special case $\beta = 0$, delivers the competitive wage when either $\beta = 1$ or supply is perfectly elastic ($\epsilon_{ij} \to \infty$), and reduces to pure rent splitting when supply is perfectly inelastic ($\epsilon_{ij} \to 0$). Figure E.2 (Appendix E) illustrates these four cases geometrically.

Trade. The world price is normalized to one, $p^* = 1$, and is common across varieties. Iceberg trade costs satisfy $\delta \in (0,1)$: shipping one unit delivers δ units, so a smaller δ means a more costly trade regime.[27] Because foreign goods arrive melted, domestic consumers face the price $p_{ij} = p^*/\delta = 1/\delta$ for every variety. An active domestic producer therefore receives $1/\delta > p^*$ on domestic sales but only p^* (net of iceberg losses on its own shipments) on exports. Substituting the relevant price into (21), the bargained wage is

$$w_{ij} = \frac{\epsilon_{ij} + \beta}{1+\epsilon_{ij}} \cdot p_{ij}\, a_{ij}, \qquad p_{ij} \in \left\{\tfrac{1}{\delta},\, p^*\right\}, \tag{23}$$

where the relevant marginal price is $1/\delta$ for a domestic producer and p^* for an exporter. The domestic labor required to meet home demand for variety ij is

$$n_{ij}^{D} = \frac{\delta\, \phi_{ij}}{a_{ij}} \cdot \text{Income}, \tag{24}$$

[27] Under this convention, *trade liberalization* corresponds to a rise in δ toward one.

obtained by combining the variety demand (14) at price $1/\delta$ with the technology $y_{ij} = a_{ij} n_{ij}$. An active firm exports if and only if its equilibrium employment, evaluated at the exporter marginal price p^*, exceeds the labor needed to serve home demand:

$$n_{ij}(p^*) > n_{ij}^D. \tag{25}$$

Let $\mathcal{X}_j \subseteq \mathcal{F}_j$ denote the set of active exporters in market j. Variable profit is

$$\pi_{ij}^v = \begin{cases} \left(\frac{a_{ij}}{\delta} - w_{ij}\right) n_{ij} & \text{if } i \notin \mathcal{X}_j \quad \text{(non-exporter)}, \\ \left(\frac{a_{ij}}{\delta} - w_{ij}\right) n_{ij}^D + \left(a_{ij} - w_{ij}\right)\left(n_{ij} - n_{ij}^D\right) & \text{if } i \in \mathcal{X}_j \quad \text{(exporter)}. \end{cases} \tag{26}$$

The price discontinuity at n_{ij}^D generates three possible equilibrium regimes for each firm, illustrated in Figure E.3 (Appendix E). A *non-exporter* is sufficiently small that it is content to serve only the domestic market at the price $1/\delta$, so $n_{ij}(p) \le n_{ij}^D$. An *exporter* is sufficiently productive that, even at the lower world price p^*, it wishes to hire more than n_{ij}^D workers and ship the surplus abroad. The intermediate case, a *constrained non-exporter*, is a firm large enough to over-supply its variety at the domestic price but not profitable enough to export at the world price: formally, $n_{ij}(p^*) < n_{ij}^D < n_{ij}(p)$. Such a firm is pinned at $n_{ij} = n_{ij}^D$, producing exactly to meet home demand while earning the domestic price on every unit. Because exporters earn revenue from both markets, at the local-labor-market level a more productive firm tends to generate higher variable profit.[28]

Importantly, the trade channel links back to the labor supply elasticity through the payroll-share mechanism of equation (18). When trade costs fall (δ rises), the domestic price advantage of non-exporters narrows, employment shifts toward high-productivity exporters, and their payroll shares s_{ij} rise within local markets. Through equation (18), higher payroll shares lower the firm-level supply elasticity ϵ_{ij}, raising monopsony power even without any change in the set of firms. This general equilibrium effect is the structural counterpart to the reallocation channel identified in the LP-IV evidence of Section V.

Income and balanced trade. Aggregating the budget constraint (13) over firms, and valuing each active firm's revenue at the price relevant to its export status, nominal income satisfies:

$$\text{Income} = \frac{\delta \sum_j \sum_{i \in \mathcal{X}_j} a_{ij}\, n_{ij} \;+\; \delta^{-1} \sum_j \sum_{i \in \mathcal{F}_j \setminus \mathcal{X}_j} a_{ij}\, n_{ij}}{1 - \left(1 - \delta^2\right) \sum_j \sum_{i \in \mathcal{X}_j} \phi_{ij}}, \tag{27}$$

where the numerator collects exporters' revenue valued at the world price net of iceberg losses and non-exporters' revenue valued at the domestic price $1/\delta$, and the denominator is the general-equilibrium multiplier that corrects for the share of expenditure falling on exported varieties.[29]

[28] This is not uniformly true in the cross-section: differences in oligopsony power across markets break the monotonicity between productivity and profit; the statement holds within a local labor market.

[29] A step-by-step derivation from (13) and the balanced-trade condition (29), in the equivalent notation of a two-type (exporter/domestic) partition, is given in Appendix 2.

Because the Cobb–Douglas price index equals the common consumer price $1/\delta$, real consumption is given by:

$$C = \frac{\text{Income}}{1/\delta} = \delta \cdot \text{Income}. \tag{28}$$

The balanced-trade condition, which emerges directly from the household's budget constraint, requires that the value of exports equal the value of imports:

$$\underbrace{\sum_j \sum_{i \in \mathcal{X}_j} a_{ij} \left(n_{ij} - n_{ij}^D \right)}_{\text{exports}} = \underbrace{\sum_j \sum_{i \in \mathcal{F}_j \setminus \mathcal{X}_j} \left(\phi_{ij} \, \text{Income} - \tfrac{1}{\delta} \, a_{ij} \, n_{ij} \right)}_{\text{imports}}, \tag{29}$$

where the right-hand side is the residual domestic demand for each non-exporter's variety that the firm does not supply at its bargained employment and that is therefore met by imports of the same variety. Because every variety is produced by an active firm, there is no separate fully-imported-variety term.

Aggregation. Following Berger et al. (2022a), firm-level allocations aggregate into three welfare-relevant indices. Aggregate productivity is the nested-CES productivity aggregator

$$\tilde{A} = \left(\sum_j \tilde{a}_j^{1+\theta} \right)^{\frac{1}{1+\theta}}, \qquad \tilde{a}_j = \left(\sum_{i \in \mathcal{F}_j} a_{ij}^{1+\eta} \right)^{\frac{1}{1+\eta}}. \tag{30}$$

The aggregate worker take-home share is the productivity-weighted CES aggregator of firm inverse markdowns,

$$\mathcal{M} = \left[\sum_j \left(\frac{\tilde{a}_j}{\tilde{A}} \right)^{1+\theta} \mathcal{M}_j^{1+\theta} \right]^{\frac{1}{1+\theta}}, \qquad \mathcal{M}_j = \left[\sum_{i \in \mathcal{F}_j} \left(\frac{a_{ij}}{\tilde{a}_j} \right)^{1+\eta} \mu_{ij}^{1+\eta} \right]^{\frac{1}{1+\eta}}, \tag{31}$$

and the misallocation index is

$$\Omega = \sum_j \left(\frac{\tilde{a}_j}{\tilde{A}} \right)^{1+\theta} \left(\frac{\mathcal{M}_j}{\mathcal{M}} \right)^{\theta} \omega_j, \qquad \omega_j = \sum_{i \in \mathcal{F}_j} \left(\frac{a_{ij}}{\tilde{a}_j} \right)^{1+\eta} \left(\frac{\mu_{ij}}{\mathcal{M}_j} \right)^{\eta}. \tag{32}$$

The index equals one when markdowns are equalized across firms; dispersion in μ_{ij} moves it away from unity and measures the allocative-efficiency cost of that dispersion relative to the competitive benchmark. Following Berger et al. (2022a), these three indices are the channels through which firm-level monopsony aggregates: $\tilde{A}$ is aggregate productivity, $\mathcal{M}$ the economy-wide labor share (the aggregate counterpart of the firm take-home share), and Ω the allocative efficiency of the labor allocation across firms. These objects, together with aggregate consumption C and employment N, are the inputs to the welfare comparison in Subsection E.

B. Equilibrium

DEFINITION 1 (Equilibrium). *An equilibrium is a collection* $\{w_{ij}, n_{ij}, s_{ij}, \mathcal{X}_j\}$ *and an income level* Income *satisfying:*

(E1) Household optimization*: the household maximizes* (11) *subject to* (12), (13)*, and* (15)*, yielding variety demand* (14) *and firm-level labor supply* (17).

(E2) Wage bargaining*: for each firm* $i \in \mathcal{F}_j$*, the wage solves the Nash problem* (20)*, delivering* $w_{ij} = \mu_{ij}\, p_{ij}\, a_{ij}$ *with* $\mu_{ij} = (\epsilon_{ij} + \beta)/(1 + \epsilon_{ij})$.

(E3) Labor market clearing*: employment at each firm equals firm-level labor supply* (17)*, with wage indices computed via* (16)*, and the payroll shares* (19) *are mutually consistent within each market.*

(E4) Goods market clearing*: demand for each variety* (14) *is met by domestic production, by imports, or by a combination of the two when a domestic producer supplies less than home demand* ($n_{ij} < n_{ij}^D$).

(E5) Export status*:* $\mathcal{X}_j = \{i \in \mathcal{F}_j : n_{ij}(p^*) > n_{ij}^D\}$.

(E6) Income consistency and balanced trade*: income satisfies* (27) *and the balanced-trade condition* (29) *holds.*[30]

C. Calibration

We calibrate the model separately for two cross-sections, a baseline averaging years 1990–1995 ($\tau = 1$) and a final period averaging 2013–2018 ($\tau = 2$), taking the firm population and market sizes (a_{ij}, ϕ_{ij}, M_j) directly from the data. Both labor-supply elasticities are estimated in sample from the equilibrium bounding relation (18), $1/\epsilon_{ij} = (1 - s_{ij})/\eta + s_{ij}/\theta$, rather than borrowed from external calibrations. The relation ties a firm's supply elasticity to its labor-market share, and its two boundary cases identify the two elasticities separately. As a firm's payroll share approaches one, $s_{ij} \to 1$, its supply elasticity collapses to $\epsilon_{ij} \to \theta$ and its take-home share reduces to $\mu_{ij} \to (\theta + \beta)/(1 + \theta)$, which at $\beta = 0$ equals $\theta/(1 + \theta)$; we therefore set θ so that $\theta/(1 + \theta)$ matches the P95 of the employment-weighted markdown distribution (the tail populated by high-payroll-share firms), yielding $\theta = 0.24$.[31] This boundary is credible because it requires only a single near-monopsonistic firm at the top of the markdown distribution, which concentrated local labor markets readily supply. The opposite boundary identifies η: atomistic, low-payroll-share firms face the most elastic supply, $\epsilon_{ij} \to \eta$ as $s_{ij} \to 0$, and set the smallest markdowns, so we match the 5th percentile of the distribution, $\eta = 1/(P_5 - 1) = 9.34$. This limit is the more demanding of the two: because no firm has a payroll share of exactly zero, the P5 firm retains a small positive share, so $1/(P_5 - 1)$ is a slight *lower bound* on the true within-market elasticity. Reassuringly, the resulting value sits just below the external estimate of Berger et al. (2022b) ($\eta \approx 10.85$, for which Chile has no quasi-experimental analogue), consistent with this small downward approximation and lending it independent support.[32] The Frisch elasticity is $\varphi = 0.5$ (Chetty et al., 2011; Berger et al., 2022a), and the world price is normalized to $p^* = 1$.

[30] The latter is implied by the household budget constraint.

[31] We assume $\beta_{ij} = 0$ for the near-monopsonistic firm at P95 that identifies θ.

[32] As a robustness check, re-solving the model with η fixed at the Berger et al. (2022b) value of 10.85 (holding the calibrated (δ, β) and scale parameters at their baseline values) leaves the welfare results essentially unchanged: the realized welfare level rises marginally from $\lambda = 1.596$ to 1.606, and the realized-minus-gradual welfare gap is 0.122 versus 0.120.

Firm population and markets. The firm population is drawn non-parametrically. We populate each market with *empirical pairs* (a_{ij}, ϕ_{ij}) drawn jointly from the firm-level data, preserving the observed cross-sectional dependence between productivity a_{ij} and variety weight ϕ_{ij}. Market sizes M_j are obtained independently from ENIA market definitions (3-digit ISIC codes × administrative regions) and then tightened from the region to the *comuna* level using CASEN within-region employment shares, so that M_j captures local rather than regional concentration. Because ENIA surveys only plants with at least ten employees (precisely the firms that face the most inelastic labor supply and exhibit the highest markdowns), we augment the surveyed population with a *fringe* of sub-threshold firms, disciplined by administrative (SII) records on revenues and the ENIA-implied correlation between revenues and productivity: 1.60 fringe plants per surveyed plant (a count share of 0.62), accounting for only 3% of sales, domestic and non-exporting by construction. To avoid degenerate fringe-only markets, each market is required to contain at least one surveyed (ENIA) plant[33]. This approach corrects for ENIA's truncation and ensures the model represents the universe of manufacturing firms rather than only the upper tail of the size distribution; it also generates a model-based selection correction on aggregate markdowns that is most consequential in the early years of the sample, when sub-threshold firms represent a larger share of economic activity.

For each period we calibrate the iceberg trade cost δ_τ and the bargaining weight β_τ by minimizing the sum of squared relative residuals on two targeted moments over (δ, β): average import penetration $\overline{IP}_\tau$ and the employment-weighted worker take-home share $\bar{\mu}_\tau \equiv \sum_{ij}(n_{ij}/N)\,\mu_{ij}$. The productivity scale z and the labor-disutility shifter $\bar{\varphi}$ are recovered in a separate step that matches average firm size and average earnings per worker in the baseline[34]. Because the targeted moments are dimensionless and homogeneous of degree zero in scale, $(z, \bar{\varphi})$ affect levels only and do not feed back into the (δ, β) calibration.

The employment-weighted take-home share $\bar{\mu}_\tau$ pins β_τ: holding (η, θ) fixed, β enters every markdown additively through $\mu_{ij} = (\epsilon_{ij} + \beta)/(1 + \epsilon_{ij})$, so $\bar{\mu}$ maps monotonically into β. The import penetration moment $\overline{IP}_\tau$ pins δ_τ: holding β fixed, a lower trade cost (higher δ) raises import penetration monotonically by compressing the domestic price advantage and drawing in more foreign varieties. For the final period we impose $\delta_2 > \delta_1$, consistent with the liberalization documented in Section I.

Table 4 reports the calibrated values and the model fit. The calibrated bargaining weight rises from $\hat{\beta}_1 = 0.04$ to $\hat{\beta}_2 = 0.29$, matching the increase in the worker take-home share from $\bar{\mu}_1 = 0.33$ to $\bar{\mu}_2 = 0.56$ in the data. Within the model this rise in β is the institutional channel: trade liberalization alone would have raised monopsony power, by reallocating employment toward high-payroll-share firms, so only a higher bargaining weight can reconcile the observed compression of wage markdowns over the period. The iceberg cost falls (i.e. δ rises) from $\hat{\delta}_1 = 0.50$ to $\hat{\delta}_2 = 0.70$, matching the rise in import penetration from $\overline{IP}_1 = 0.17$ to $\overline{IP}_2 = 0.23$.

[33] The full fringe construction is detailed in Appendix F

[34] As defined above, z is a common scale factor on a_{ij} whose only role is to rescale model magnitudes to Chilean-peso levels.

TABLE 4—Model Calibration

Symbol	Description	Value	Target	Source
Panel A: Structural Parameters				
η	Within-market labor supply elasticity	9.34	—	Internal: markdown P_5
θ	Cross-market labor supply elasticity	0.24	—	Internal: markdown P_{95}
φ	Frisch elasticity of labor supply	0.50	—	Berger et al. (2022b)
P^*	World price	1.00	—	Normalisation
Panel B: Firm Population and Market Structure				
J	Number of synthetic markets	5000	—	Construction
$\bar{M}$	Mean firms per market	16.06	—	ENIA × CASEN
$\tilde{M}$	Median firms per market	2.00	—	ENIA × CASEN
$a_{ij} = \exp(\hat{\omega}_{ij})$	Productivity (PF Solow residual)	—	—	Ackerberg et al. (2015)
ϕ_{ij}	Cobb–Douglas weights ($\sum_{ij} \phi_{ij} = 1$)	—	—	SII (2005)
	Surveyed (ENIA) count share	0.41	0.38	ENIA × SII
	Fringe count share	0.59	0.62	ENIA × SII
	Fringe sales (ϕ) share	0.03	0.03	ENIA × SII
Panel C: Scale Parameters (Stage 2 recovery)				
z	Mean firm productivity	203959.11	—	Matches avg. firm size and earnings
$\tilde{\varphi}$	Labor disutility shifter	0.26	—	
	Average firm size	31.48	31.48	ENIA × SII, 1990–1995
	Average earnings per worker	214009.94	214009.94	ENIA × SII, 1990–1995
Panel D: Calibrated Parameters — Baseline (1990–1995 average)				
δ_1	Iceberg trade cost	0.50	—	Jointly calibrated
β_1	Nash bargaining power	0.04	—	Jointly calibrated
	Import penetration ($\overline{IP}$)	0.17	0.17	ENIA, 1990–1995
	Worker take-home share ($\tilde{\mu}$)	0.38	0.38	ENIA, 1990–1995
	Exporter share (%)	20.11	7.42	Untargeted
Panel E: Calibrated Parameters — Final (2013–2018 average)				
δ_2	Iceberg trade cost	0.70	—	Jointly calibrated
β_2	Nash bargaining power	0.29	—	Jointly calibrated
	Import penetration ($\overline{IP}$)	0.23	0.23	ENIA, 2013–2018
	Worker take-home share ($\tilde{\mu}$)	0.51	0.51	ENIA, 2013–2018
	Exporter share (%)	17.06	8.23	Untargeted

Notes: Structural parameters (Panel A) are set externally, except the labor-supply elasticities, pinned to the estimated markdown distribution: $\eta = 1/(P_5 - 1)$ and $\theta = 1/(P_{95} - 1)$. The period parameters $(\delta_\tau, \beta_\tau)$ (Panels D–E) are jointly calibrated on the two targeted moments: import penetration $\overline{IP}$ identifies δ, and the worker take-home share $\tilde{\mu}$ identifies β. The exporter share is untargeted.

D. Empirical Counterfactual

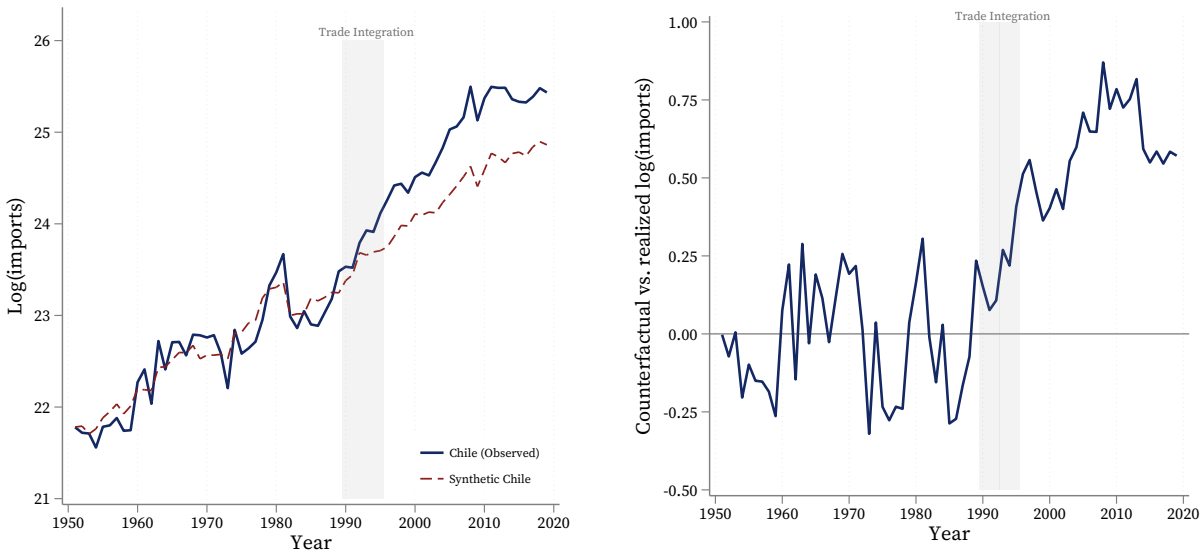
To evaluate counterfactual trade scenarios, we first construct the import trajectory that Chile would have followed under a more gradual liberalization path. We adopt the Synthetic Control Method (Abadie and Gardeazabal, 2003; Abadie et al., 2010; Abadie, 2021), which builds a counterfactual for a single treated unit from a weighted combination of untreated units that best reproduces the treated unit's pre-treatment outcomes and characteristics.

Chile is the treated unit and the treatment is the second stage of trade liberalization initiated in the early 1990s; we designate 1991 as the treatment year. The donor pool consists of countries that did not undergo comparable trade reforms over the same period and for which consistent manufacturing and trade data are available from the Penn World Table. The synthetic control is constructed to match Chile's pre-1990 import trajectory and a set of predictors including lagged imports, the labor share, the human-capital index, and real GDP.

Formally, let Y_{1t} denote the outcome for Chile at time t, and $Y_{0t} = \sum_{k=2}^{K+1} w_k Y_{kt}$ the outcome for the synthetic control, with non-negative weights w_k summing to one chosen to minimize

the pre-treatment distance in outcomes and predictors.

FIGURE 8—Synthetic Import Time Series



(a) Levels: Chile vs. synthetic

(b) Realized – synthetic

Notes: Left panel: import penetration for Chile (solid) and its synthetic control (dashed), 1960–2018; the vertical line marks the treatment year (1991). Right panel: difference between realized and synthetic imports. The synthetic control follows Abadie et al. (2010) with donor countries that did not undergo comparable trade reforms. *Source:* Penn World Table and Atlas of Economic Complexity. Authors' calculations.

Figure 8 presents the results. The synthetic control closely tracks Chile's import trajectory before the early 1990s, indicating a good pre-treatment fit. After liberalization, realized imports diverge sharply from the synthetic counterfactual, consistent with a substantial and persistent trade shock. Table 4 confirms that Chilean covariates are closely matched.[35] The zero weight assigned to Argentina, Brazil, and Colombia is informative rather than problematic: these economies also liberalized during the 1990s and are effectively treated units, so their exclusion strengthens the design.

We use the synthetic control to discipline the counterfactual trade cost. The synthetic (comparable-country) path reproduces a share 0.745 of Chile's realized change in import penetration (the component common to global integration), so only the residual $0.255 = 1 - 0.745$ reflects the post-1991 divergence of Chile's own, faster liberalization. The counterfactual iceberg cost drops this residual of the calibrated change in δ,

$$\delta_{\text{synth}} = \delta_1 + (1 - 0.255)\,(\delta_2 - \delta_1), \tag{33}$$

which feeds the more gradual import trajectory into the structural model.

[35] The donors receiving positive weight are Bolivia, Cyprus, Guatemala, the Netherlands, and Sri Lanka. Table B.6 in the Appendix reports the full weights and donor pool.

E. Welfare Analysis

The counterfactual markdown exercise isolates the trade channel's contribution to monopsony power, but it does not address the central welfare question: were workers and consumers better off under aggressive liberalization, even though it raised monopsony power? Answering this requires the structural model, which delivers consumption, employment, wages, and welfare under alternative scenarios.

We adopt the consumption-equivalent welfare measure of Berger et al. (2022a). Let C^o and N^o denote aggregate consumption and (CES) employment in the baseline equilibrium and C^c, N^c the corresponding outcomes in a counterfactual. The household requires λ times its baseline consumption to be indifferent between the two allocations:

$$\lambda C^o - \bar{\varphi}^{-\frac{1}{\varphi}} \frac{(N^o)^{1+\frac{1}{\varphi}}}{1+\frac{1}{\varphi}} = C^c - \bar{\varphi}^{-\frac{1}{\varphi}} \frac{(N^c)^{1+\frac{1}{\varphi}}}{1+\frac{1}{\varphi}}. \tag{34}$$

We evaluate six equilibria, indexed by a trade cost and a bargaining weight: the *baseline* (δ_1, β_1); *trade only* (δ_2, β_1), which applies the realized liberalization at the baseline institutional regime; *bargaining only* (δ_1, β_2); *both channels* (δ_2, β_2), the realized trajectory; *synthetic trade only* $(\tilde{\delta}, \beta_1)$; and *synthetic + bargaining* $(\tilde{\delta}, \beta_2)$, the gradual-liberalization counterfactual. For any outcome X the interaction term between the two channels is obtained from $X_{\text{joint}} - X_{\text{trade}} - X_{\text{barg}} + X_{\text{base}}$. A positive (negative) interaction indicates complementarity (substitutability) between trade and institutional reform. Table 5 reports the results: welfare λ, employment and consumption (normalized to baseline), nominal and real wages, aggregate markdown, the misallocation index Ω and exporter share of firms.

TABLE 5—Welfare Decomposition: Trade Liberalization and Bargaining

	λ	Cons.	W_n	W_r	Markdown	Misalloc.	Emp.	Exp.(%)
Baseline	1.000	1.000	1.000	1.000	1.000	1.000	1.000	20.09
Trade only ($\uparrow \delta$, β fixed)	1.372	1.434	0.920	1.297	1.053	0.993	1.141	21.62
Bargaining only (δ fixed, $\uparrow \beta$)	1.147	1.216	1.318	1.318	0.755	1.053	1.164	15.98
+ Interaction	+0.077	+0.116	-0.003	+0.126	-0.028	+0.005	+0.027	-0.45
Both channels (Reduced TC)	1.596	1.766	1.235	1.740	0.780	1.050	1.332	17.06
Synthetic trade only	1.273	1.318	0.936	1.221	1.042	0.994	1.107	21.32
+ Interaction	+0.056	+0.084	-0.003	+0.093	-0.022	+0.003	+0.020	-0.35
Synthetic + Bargaining (Reduced TC-C)	1.475	1.619	1.251	1.632	0.775	1.050	1.292	16.85
Realized − Counterfactual	+0.120	+0.148	-0.016	+0.108	+0.005	-0.000	+0.040	+0.21

Notes: λ is consumption-equivalent welfare relative to baseline ($\lambda = 1$). Consumption, nominal wage W_n, real wage W_r, markdown, misallocation, and employment normalised to baseline; exporter share in levels. Markdown is baseline-normalised worker take-home share $1/\tilde{\mu}$; higher values indicate lower monopsony power. Misalloc. is misallocation index Ω, normalised to baseline. Interaction = $X_{\text{joint}} - X_{\text{trade}} - X_{\text{barg}} + X_{\text{base}}$; positive values indicate complementarity. Realized − Counterfactual compares Both Channels against Synthetic + Bargaining.

Under the joint scenario (combining the realized reduction in trade costs with the calibrated increase in bargaining power), welfare rises by 59.6% relative to the baseline ($\lambda = 1.596$). The gain is driven overwhelmingly by consumption, which rises by roughly 77%, as lower trade costs reduce domestic prices and higher bargaining power raises wages; real wages increase by approximately 74%, and the model generates labor reallocation toward more productive exporting firms. Under the gradual counterfactual (the synthetic, more gradual import trajectory combined with the same increase in β), welfare rises by 47.5% ($\lambda = 1.475$). The realized path

therefore delivers 12.0 consumption-equivalent percentage points of additional welfare relative to the gradual alternative.[36]

The table makes explicit the central trade-off in open-economy labor markets. Trade liberalization simultaneously raises monopsony power (through the reallocation channel of Section V, by which employment shifts toward high-share, low-elasticity firms) and consumption, through lower domestic prices. Under the trade-only scenario the aggregate markdown rises by about 5.3% relative to baseline while consumption rises by roughly 43%: workers face a larger wedge between marginal product and wage, but cheaper goods, and the welfare comparison is governed by which effect dominates. Interventions that reduce trade exposure to protect workers' wage-setting position would forgo the consumption gains that this channel delivers.

Equally striking is the welfare contribution of the institutional channel in isolation. Under the bargaining-only scenario (holding trade costs at δ_1 while allowing β to rise to 0.29), welfare rises by 14.7% ($\lambda = 1.147$), as higher wages and employment lift household income and demand substantially. This result underscores that institutional reform is not merely a complement to trade liberalization but is independently powerful in raising welfare by compressing the monopsony rent that workers forgo.

The decomposition isolates the interaction between the two reforms. The welfare gain from reduced trade costs alone (holding β at β_1) is 37.2% ($\lambda = 1.372$); the gain from higher bargaining power alone (holding δ at δ_1) is 14.7% ($\lambda = 1.147$). If the two forces operated independently their combined effect would equal the sum of the components (+0.519 in λ), but the realized joint effect is +0.596, implying a positive interaction of +0.077, approximately 7.7 percentage points, or about 13% of the total realized welfare gain. The two reforms are therefore complements rather than independent forces.

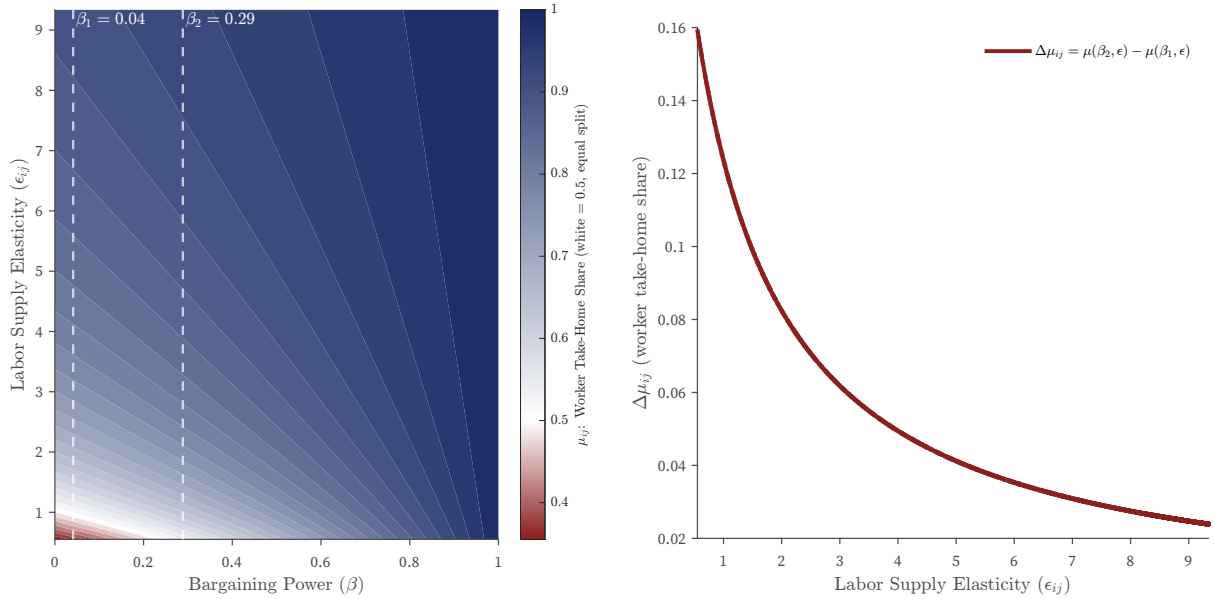
The mechanism behind this complementarity is rooted in the additive structure of Proposition 1. Stronger bargaining power raises wages and employment, lifting household income and demand for domestic varieties; when trade costs fall simultaneously, lower prices amplify the consumption gains from higher income. Conversely, lower trade prices amplify the real-income gains from higher wages. The cross-partial $\partial^2 \mu_{ij} / \partial \beta \, \partial \epsilon_{ij} < 0$ means that institutional reform has its largest wage effects precisely in the high-monopsony markets where trade-induced reallocation is most concentrated, so the two reforms are mutually reinforcing, and a country that liberalizes trade without strengthening workers' bargaining position would capture smaller welfare gains because the demand-side amplification is weaker.

Figure 9 makes this cross-partial visible rather than purely algebraic. Panel A plots the additive decomposition $\mu_{ij} = (\epsilon_{ij} + \beta)/(1 + \epsilon_{ij})$ over the (β, ϵ_{ij}) plane; white marks $\mu_{ij} = 0.5$, so blue regions are markets where workers keep the majority of marginal revenue product and red regions are markets where the firm keeps the majority. Panel B isolates the comparative static directly: it plots $\Delta\mu_{ij}(\epsilon_{ij}) = \mu_{ij}(\beta_2, \epsilon_{ij}) - \mu_{ij}(\beta_1, \epsilon_{ij}) = (\beta_2 - \beta_1)/(1 + \epsilon_{ij})$, the take-home-share gain from the $\beta_1 \to \beta_2$ reform as a function of ϵ_{ij} alone. The curve is strictly declining, which is exactly Proposition 2(iii): institutional reform raises wages by more in low-ϵ_{ij} (high-monopsony)

[36] This ranking is robust to the Frisch elasticity. Re-recovering the scale parameters for $\varphi \in \{0.10, 0.25, 0.50, 0.75, 0.90, 1.00\}$ and holding (δ, β) fixed, the realized-minus-counterfactual welfare gap remains positive and of similar magnitude throughout (between +0.08 and +0.19 in λ; Table F.1), so the conclusion that aggressive liberalization dominates the gradual path does not hinge on $\varphi = 0.5$.

markets. The overlaid points are the calibrated firm-level ϵ_{ij} distribution under the baseline equilibrium, plotted on the curve itself: their concentration toward the steep, low-ϵ_{ij} part of the curve shows that trade-induced reallocation pushes employment mass toward precisely the markets where institutional reform has the largest wage effect, which is why the two reforms reinforce one another in the welfare accounting above.

FIGURE 9—Complementarity between Bargaining Power and Labor Supply Elasticity



(a) Take-home share over (β, ϵ_{ij}) (b) Marginal gain from $\beta_1 \to \beta_2$

Notes: Panel (a): worker take-home share over the (β, ϵ_{ij}) plane, ϵ_{ij} spanning the 5th–95th percentile of the calibrated elasticity distribution; dashed white lines mark β_1 (baseline) and β_2 (final). Panel (b): the marginal take-home-share gain from $\beta_1 \to \beta_2$ as a function of ϵ_{ij}; points show the empirical ϵ_{ij} distribution under the baseline calibration.

Table F.1 (Appendix F) decomposes welfare into a reallocation component λ^{ω} and a markdown-level component λ^{μ}. The latter is computed in a collapsed economy with $\eta = \theta = e^*$, whose *uniform* markdown matches the employment-weighted aggregate take-home share, so the misallocation index $\Omega \to 1$ by construction and $\lambda^{\omega} = \lambda - \lambda^{\mu}$ is the residual reallocation term. The reallocation component λ^{ω} (the aggregate efficiency gain from the trade-induced resorting of labor across firms of differing productivity, and the welfare counterpart of the misallocation index Ω) accounts for approximately 79% of realized welfare at the baseline Frisch elasticity. This share is stable across the full range $\varphi \in \{0.10, \ldots, 1.00\}$, staying between 73.5% and 83.2%. The markdown-level component λ^{μ}, the welfare counterpart of the aggregate markdown $\mathcal{M}$ (the average wage wedge in a no-dispersion counterfactual), contributes the remaining 21%. This decomposition establishes that the welfare gains from Chile's liberalization derive primarily from the reallocation of labor toward more productive firms, not from changes in the average level of monopsony power. The monopsony dimension matters, however, through the complementarity channel: institutional reform ensures that workers capture a growing share of the productivity gains from reallocation, amplifying the consumption gains through higher household income and demand.

The welfare analysis establishes three results. First, the realized liberalization generates consumption-equivalent welfare gains relative to the gradual counterfactual despite raising monopsony power, because the consumption gains from lower prices are quantitatively large. Second, the joint effect of trade openness and institutional reform departs from the sum of their individual effects, so the two reforms interact rather than operate independently. Third, the appropriate policy response to trade-induced monopsony power is not trade protection but complementary institutional reform that strengthens workers' bargaining position.

VII. Conclusion

This paper studies how trade liberalization and rising employment protection jointly shape employer wage-setting power in Chilean manufacturing, and quantifies the welfare implications of their interaction. The analysis combines plant-level evidence, a shift-share local-projection IV design with exposure-robust inference, and a calibrated general-equilibrium model of oligopsonistic labor markets with firm-level Nash bargaining and endogenous exporter status.

Three results emerge. First, the effect of trade on labor market power is inherently dynamic. Import penetration has no contemporaneous effect on wage markdowns but generates a delayed and persistent increase: the response becomes statistically significant after four years, peaks at an elasticity of approximately 0.13 at six years, and remains positive thereafter. Static specifications therefore understate the impact of trade on monopsony power. Second, monopsony power in Chilean manufacturing is substantial but has declined over time. The median plant-level markdown is 1.85, while the employment-weighted markdown falls from roughly 3 to 1.9, implying a sizable increase in the worker share of marginal revenue product. This broad-based decline coincides with three decades of labor market reform that strengthened workers' bargaining position. Third, trade liberalization and institutional reform are welfare complements. In the calibrated model, the realized path of rapid liberalization delivers a consumption-equivalent gain of 59.6%, exceeding the 47.5% gain under gradual liberalization by 12.0 percentage points. Moreover, the joint effect of trade and institutional reform exceeds the sum of their individual contributions by 7.7 percentage points. Stronger bargaining power raises wages and amplifies the consumption gains from lower prices, while trade openness raises real income and increases the returns to bargaining.

The dynamics are consistent with a selection and reallocation mechanism. Import competition compresses markups at all horizons, induces medium-run labor-share adjustments, and shifts employment toward larger, higher-markdown firms, particularly in skill-intensive and initially concentrated industries. As a result, average wages decline while markdowns rise, even as aggregate welfare improves. These findings highlight a central tension: trade liberalization increases monopsony power, yet generates large welfare gains through lower prices. The welfare calculus is nonetheless unambiguous. The consumption gains from trade dominate the adverse effects on labor market power, implying that restricting trade to mitigate monopsony would be welfare-reducing. Instead, the appropriate policy response is complementary institutional reform. Strengthening workers' bargaining position offsets the rise in monopsony power and magnifies the gains from trade, making the two reforms jointly optimal.

While the empirical setting is Chile, the underlying mechanisms are general. Trade-induced

reallocation toward larger, more productive firms is a core prediction of heterogeneous-firm trade models (Melitz, 2003), and comparable levels of monopsony power have been documented in other emerging economies (Felix, 2026; Huang et al., 2026). The framework developed here—where institutional reform disciplines labor market power while trade reshapes firm composition—applies broadly to economies undergoing simultaneous product and labor market liberalization.

Several limitations suggest directions for future work. The welfare analysis abstracts from intertemporal dynamics, endogenous firm entry and exit, and richer product market interactions. Endogenizing the institutional channel—allowing bargaining power to respond to specific policies such as minimum wages, education expansion, or enforcement—would enable a structural decomposition of the sources of declining monopsony power. In addition, incorporating worker-level heterogeneity and geographic mobility would allow the analysis to capture distributional effects across gender, skill, and regions, which are likely central to understanding the full consequences of trade in imperfect labor markets.

Taken together, the evidence suggests a broader lesson: in economies with monopsonistic labor markets, the gains from trade liberalization depend critically on the institutional environment. In our framework, trade openness and pro-worker labor market reforms are not substitutes but complements, and their joint implementation yields the largest welfare gains.

References

Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature 59*(2), 391–425.

Abadie, A., A. Diamond, and J. Hainmueller (2010, 6). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American Statistical Association 105*, 493–505.

Abadie, A. and J. Gardeazabal (2003, 2). The economic costs of conflict: A case study of the basque country. *American Economic Review 93*, 113–132.

Ackerberg, D. A., K. Caves, and G. Frazer (2015). Identification properties of recent production function estimators. *Econometrica 83*, 2411–2451.

Adão, R., M. Kolešár, and E. Morales (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics 134*, 1949–2010.

Aghelmaleki, H., R. Bachmann, and J. Stiebale (2022). The China shock, employment protection, and european jobs. *ILR Review 75*(5), 1269–1293.

Amodio, F. and N. de Roux (2024, 10). Measuring labor market power in developing countries: Evidence from colombian plants. *Journal of Labor Economics 42*, 949–977.

Amodio, F., P. Medina, and M. Morlacco (2025, 9). Labor market power, self-employment, and development. *American Economic Review 115*, 3014–3057.

Autor, D. H., D. Dorn, and G. H. Hanson (2013, 10). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review 103*, 2121–2168.

Azar, J. and J. Sanz-Espín (2025). Labor market power in spain. *SERIEs 16*(2), 233–256.

Banco Central de Chile (2024). Cuentas nacionales: Base de datos estadísticos. Statistical database, Santiago, Chile.

Barnichon, R. and C. Brownlees (2019, 7). Impulse response estimation by smooth local projections. *The Review of Economics and Statistics 101*, 522–530.

Bartik, T. J. (1991, 9). *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute.

Berger, D., K. Herkenhoff, and S. Mongey (2022a, 4). Labor market power. *American Economic Review 112*, 1147–1193.

Berger, D., K. Herkenhoff, and S. Mongey (2022b, 4). Labor market power. *American Economic Review 112*, 1147–1193.

Billa, B., L. Bishop, S. Deakin, and K. Pourkermani (2025). The impact of labour laws on the labour share of national income, productivity, unemployment and employment: First results from the 2023 update of the CBR labour regulation index. Technical report, International Labour Organization.

Blau, F. D. and L. M. Kahn (2017, 9). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature 55*, 789–865.

Bond, S. (2005, 2). Adjustment costs and the identification of cobb douglas production functions.

Bond, S., A. Hashemi, G. Kaplan, and P. Zoch (2021). Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data. *Journal of Monetary Economics 121*, 1–14.

Borusyak, K., P. Hull, and X. Jaravel (2022, 1). Quasi-experimental shift-share research designs. *The Review of Economic Studies 89*, 181–213.

Borusyak, K., P. Hull, and X. Jaravel (2025). Design-based identification with formula instruments: A review. *The Econometrics Journal 28*(1), 83–108.

Botero, J. C., S. Djankov, R. La Porta, F. Lopez-de Silanes, and A. Shleifer (2004). The regulation of labor. *Quarterly Journal of Economics 119*(4), 1339–1382.

Campero, G. (2004). *Macroeconomic reforms, labour markets and labour policies : Chile, 1973-2000*. Employment Analysis Unit, Employment Strategy Dept., International Labour Office.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011). Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review 101*(3), 471–475.

Cusolito, A. P., A. Garcia-Marin, and W. F. Maloney (2023, 3). Proximity to the frontier, markups, and the response of innovation to foreign competition: Evidence from matched production-innovation surveys in chile. *American Economic Review: Insights 5*, 35–53.

De Loecker, J. (2011, 5). Recovering markups from production data. In *International Journal of Industrial Organization*, Volume 29, pp. 350–355.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2016). Prices, markups, and trade reform. *Econometrica 84*, 445–510.

De Loecker, J. and F. Warzynski (2012, 10). Markups and firm-level export status.

Dobbelaere, S. and Q. Wiersma (2025). The impact of trade liberalization on firms' product and labor market power. *Industrial and Corporate Change 34*(1), 210–233.

Doraszelski, U. and J. Jaumandreu (2021). Reexamining the De Loecker & Warzynski (2012) method for estimating markups. Discussion Paper 16027, CEPR. Revised 2023-10-30; journal status unconfirmed as of 2026-07-17.

Díez, F., D. Malacrino, and I. Shibata (2022). The divergent dynamics of labor market power in europe.

Edwards, S., A. C. Edwards, S. Bentolila, and J. Galí (2000). Economic reforms and labour markets: Policy issues and lessons from chile. *Economic Policy 15*(30), 183–229.

Eslava, M., A. García-Marín, and J. Messina (2025). Inequality and market power in latin america and the caribbean. *Oxford Open Economics 4*(Supplement_1).

Feenstra, R. C., R. Inklaar, and M. P. Timmer (2015). The next generation of the penn world table. *American Economic Review 105*(10), 3150–3182. Available for download at https://www.ggdc.net/pwt.

Felix, M. (2026). Trade, labor market concentration, and wages. NBER Working Paper 35018, National Bureau of Economic Research.

Ffrench-Davis, R. (2010). *Economic Reforms in Chile*. Palgrave Macmillan UK.

Frank, V. (2002, 1). The elusive goal in democratic chile: Reforming the pinochet labor legislation. *Latin American Politics and Society 44*, 35–68.

French-Davis, R. and B. Stallings (2001). *Reformas crecimiento y politicas sociales en Chile desde 1973*. LOM Ediciones and CEPAL.

Gandhi, A., S. Navarro, and D. A. Rivers (2020, 8). On the identification of gross output production functions. *Journal of Political Economy 128*, 2973–3016.

Garibaldi, P. and G. L. Violante (2005). The employment effects of severance payments with wage rigidities. *Economic Journal 115*(506), 799–832.

Goldberg, P. K. and G. Maggi (1999). Protection for sale: An empirical investigation. *American Economic Review 89*(5), 1135–1155.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020, 8). Bartik instruments: What, when, why, and how. *American Economic Review 110*, 2586–2624.

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

Gutiérrez, A. (2023). Labor market power and the pro-competitive gains from trade *.

Hall, R. E. (1986). Market structure and macroeconomic fluctuations. *Brookings Papers on Economic Activity 1986*, 285.

Hartley, J. and J. Mejia (2025). Smooth panel local projections.

Heckman, J. J. and C. Pages (2004). *Law and Employment*. University of Chicago Press.

Huang, L., C. Liu, and X. Zhang (2026). Capital-biased tax incentives and labor market power: Evidence from China. *Economics Letters 266*, 113057.

Jha, P. and A. Rodriguez-Lopez (2021, 11). Monopsonistic labor markets and international trade. *European Economic Review 140*, 103939.

Jha, P., A. Rodriguez-Lopez, and A. H. Spencer (2026a). Labor market monopsony power and the dynamic gains to openness reforms. International Economic Review, forthcoming.

Jha, P., A. Rodriguez-Lopez, and C. Spencer (2026b). Labor market monopsony power and the dynamic gains to openness reforms. *International Economic Review*.

Jordà, O. and A. M. Taylor (2025). Local projections. *Journal of Economic Literature 63*(1), 59–110.

Kondo, I. O., Y. A. Li, and W. Qian (2024). Trade liberalization and labor monopsony: Evidence from Chinese firms. *Journal of International Economics 152*, 104006.

Levinsohn, J. and A. Petrin (2003, 4). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies 70*, 317–341.

Lindenlaub, I., R. Oh, and M. Peters (2022, Nov). Firm sorting and spatial inequality. NBER Working Papers 30637, National Bureau of Economic Research, Inc.

Liu, L. (1993, 12). Entry-exit, learning, and productivity change evidence from chile. *Journal of Development Economics 42*, 217–242.

Ljungqvist, L. (2002). How do lay-off costs affect employment? *Economic Journal 112*(432), 829–853.

Melitz, M. J. (2003, 11). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica 71*, 1695–1725.

Mendez, O. (2015, 11). The effect of chinese import competition on mexican local labor markets. *The North American Journal of Economics and Finance 34*, 364–380.

Micco, A. and P. Muñoz (2024, 1). The impact of extending employment protection to agency workers on firms. *American Economic Journal: Macroeconomics 16*, 66–101.

OECD (2009, 4). *OECD Reviews of Labour Market and Social Policies: Chile 2009*. OECD.

Olley, G. S. and A. Pakes (1996, 11). The dynamics of productivity in the telecommunications equipment industry. *Econometrica 64*, 1263.

Pavcnik, N. (2002). Trade liberalization, exit, and productivity improvements: Evidence from chilean plants.

Pérez Ahumada, P. (2021). Why is it so difficult to reform collective labour law? associational power and policy continuity in chile in comparative perspective. *Journal of Latin American Studies 53*, 81–105.

Pham, H. (2023). Trade reform, oligopsony, and labor market distortion: Theory and evidence. *Journal of International Economics 144*.

Raval, D. (2023a). A flexible cost share approach to markup estimation. *Economics Letters 230*.

Raval, D. (2023b). Testing the production approach to markup estimation. *The Review of Economic Studies 90*(5), 2592–2611.

Rovigatti, G. and V. Mollisi (2018, 9). Theory and practice of total-factor productivity estimation: The control function approach using stata. *The Stata Journal: Promoting communications on statistics and Stata 18*, 618–662.

Sharma, G., D. Acemoglu, S. Caldwell, V. Corradini, J. Cohen, A. Dube, M. Felix, L. Ho, S. Jäger, N. Kala, T. Suri, D. Sverdlin-Lisker, M. Uccioli, S. Wang, and L. Weiwu (2023). Monopsony and gender.

Stock, J. H. and M. W. Watson (2018, 5). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal 128*, 917–948.

SUBREI (2019, January). Impacto de los tratados de libre comercio ı política comercial de chile. Hacia una Política Comercial Inclusiva.

Sánchez, R. (2013, 1). Do reductions of standard hours affect employment transitions?: Evidence from chile. *Labour Economics 20*, 24–37.

Trefler, D. (1993). Trade liberalization and the theory of endogenous protection: An econometric study of U.S. import policy. *Journal of Political Economy 101*(1), 138–160.

van der Wiel, K. (2010). Better protected, better paid: Evidence on how employment protection affects wages. *Labour Economics 17*(1), 16–26.

Wooldridge, J. M. (2009, 9). On estimating firm-level production functions using proxy variables to control for unobservables. *Economics Letters 104*, 112–114.

Xie, E., M. Xu, and M. Yu (2024). Trade liberalization, labor market power, and misallocation across firms: Evidence from China's WTO accession. *Journal of Development Economics 171*, 103353.

Yeh, C., C. Macaluso, and B. Hershbein (2022, 7). Monopsony in the us labor market. *American Economic Review 112*, 2099–2138.

Òscar Jordà (2005, 3). Estimation and inference of impulse responses by local projections. *American Economic Review 95*, 161–182.

Òscar Jordà (2023, 9). Local projections for applied economics. *Annual Review of Economics 15*, 607–631.

Appendix

This appendix collects the material that supports and stress-tests the results in the main text: the construction of the plant-level panel and its variables, the robustness of the production-function and markdown estimates, the identification and balance diagnostics behind the trade and institutional legs, and the derivation, calibration, and estimates of the structural model. Readers seeking to assess the empirical strategy or reproduce the quantitative results will find the relevant detail here.

A. Data Construction

This appendix describes the construction of the dataset used in the empirical analysis. The final estimation sample is a balanced-comparable panel of Chilean manufacturing plants spanning 1990–2018, built from multiple raw data sources that required harmonization, linkage, and cleaning. We organize the description into six parts: (i) raw data sources, (ii) panel construction and plant linkage, (iii) variable harmonization, (iv) deflation, (v) capital measurement, and (vi) sample restrictions and outlier treatment.

1. Raw Data Sources

The backbone of the dataset is the Annual National Industrial Survey (*Encuesta Nacional Industrial Anual*, ENIA), a census of manufacturing establishments with ten or more employees conducted every year by the Chilean National Statistics Institute (INE). ENIA provides plant-level information on revenues, input expenditures (labor, intermediate materials, energy), capital stocks, investment, and workforce composition disaggregated by skill, gender, and occupation type.

The raw ENIA microdata are publicly available, but they arrive in three historically distinct formats that differ in variable naming conventions, survey structure, and plant identifiers:

- *1979–1985 panel:* An early version of the survey with a unique plant identifier (`PADRON`) and raw variable codes (e.g., `XX13`–`XX26` for employment categories). Each annual cross-section is stored as a single file.
- *1986–1994 panel:* A second format with descriptive variable names (e.g., `WC_PROD_WORKES_MALE`) and a different identifier (`NUI`). Capital stock and depreciation variables are reported separately.
- *1995–2022 panel:* The modern ENIA structure, itself divided into two sub-periods: 1995–2007 (identifier `NUI`) and 2008–2022 (identifier `ID`). Each annual survey consists of two parts (pre-2000) or three parts (post-2000) that must be merged within year before appending across years.

We complement ENIA with four auxiliary datasets:

1. *Atlas of Economic Complexity* (Harvard Growth Lab): Bilateral trade flows at the country–partner–product–year level (SITC Rev. 2, 4-digit), covering 1962–2021. We aggregate these flows to 4-digit ISIC Rev. 2 manufacturing industries using a product concordance table supplemented by manual classification of unmatched goods.

2. *CASEN* (Ministry of Social Development): A nationally representative household survey providing individual-level data on labor market status, occupation (ISCO), industry affiliation, self-employment, and informality. We aggregate these data to construct industry-level measures of occupational composition and informal employment shares.
3. *Penn World Table* (Feenstra et al., 2015): Country-level macroeconomic indicators (GDP, factor shares, productivity) used in the synthetic control exercises.
4. *Central Bank of Chile and INE price data*: Industry-level price indices and input–output tables used to construct deflators (described below).

2. Panel Construction and Plant Linkage

Building a consistent plant-level panel across the full 1979–2019 period requires linking plants across the three ENIA formats, each of which uses a different identifier system. The linkage proceeds in three stages.

Stage 1: Within-format assembly. For each ENIA format, we merge the multiple parts of each annual survey using the within-year plant identifier (`PADRON`, `NUI`, or `ID`), dropping duplicate observations that arise from data entry errors. For the 1995–2022 period, each year's survey parts (Part 1, Part 2, and where available Part 3) are merged sequentially using exact 1:1 matches. The resulting annual cross-sections are then appended to form two preliminary panels: 1995–2007 and 2008–2022.

Stage 2: Cross-format linkage via the Seguimiento database. The INE maintains a plant-tracking database (*Base de Seguimiento*) that records plant identifiers across survey waves, allowing us to link the 1995–2007 panel (`NUI`-based) to the 2008–2022 panel (`ID`-based). We perform this linkage using a probabilistic record-linkage algorithm that matches on a large set of variables observed in both datasets, including employment counts by occupation and gender, wage bills, output, value added, energy consumption, industry classification, and legal organization. We require exact matches on industry code (CIIU Rev. 3), electricity consumption, and ownership shares (domestic vs. foreign) to ensure high match quality. The algorithm returns a match score; we retain only matches with zero discrepancy in critical employment variables (total male, total female, contracted, subcontracted, and total wage bill), and we break ties by retaining the most recent consistent identifier. This procedure yields a unique crosswalk between old and new identifiers, which we use to assign a single plant identifier across the full 1979–2019 period.

Stage 3: Backward linkage to early panels. For the 1979–1985 and 1986–1994 formats, we exploit a previously constructed longitudinal identifier available in the literature (the *Base Final* dataset), which provides a consistent plant code (`id`) that spans 1979–2007. We merge each early-format panel to this longitudinal file, retaining only matched plants.

After completing the three stages, we assign each plant a unique identifier (`new_id`) and fix time-invariant characteristics, such as industry classification and geographic region, to their within-plant modal values, ensuring consistency for plants that report minor coding changes over time.

3. Variable Harmonization

Because variable names, definitions, and units change across ENIA formats, we harmonize all key variables to a common set of definitions.

EMPLOYMENT. Total employment is constructed as the sum of all worker categories reported in the survey: white-collar workers (production, executive, and administrative, by gender), blue-collar workers (production and nonproduction, by gender), home workers, and salespersons on commission. We classify workers into *skilled* (white-collar plus home workers and salespersons) and *unskilled* (blue-collar) categories. From 2001 onward, ENIA reports total employment directly; for earlier years, we reconstruct it as the sum of contracted and subcontracted workers. Gender-specific and skill-specific employment counts are constructed analogously. We impose internal consistency: if the sum of male and female workers differs from total employment by one or two workers (due to rounding), we adjust the larger group downward; the same procedure is applied to the skilled–unskilled and gender–skill decompositions.

WAGES. Total wages are defined as the sum of all labor compensation items: base wages for white-collar and blue-collar workers, bonuses, and payments to subcontracted and home workers. Skilled wages aggregate white-collar compensation and subcontracted-worker payments; unskilled wages correspond to blue-collar compensation. From 2002 onward, the survey introduces finer categories for non-standard workers (e.g., subcontracted by qualification level), which we map into the same two-category classification used in earlier years.

MATERIALS. Intermediate material inputs are defined as the sum of raw material purchases and material expenditures, including materials transferred from other establishments within the same firm. This definition remains stable across ENIA formats, requiring only variable renaming.

CURRENCY OF REPORTING (2011–2019). From 2011 the ENIA questionnaire allows plants that keep their accounts in foreign currency to report monetary variables in that currency, flagged by a currency-of-report field. Such plants are a small share of establishments (roughly 4–6%) but, being systematically large, account for 15–19% of manufacturing employment each year, so mishandling them would distort aggregate wage-bill and output measures. The convention is not uniform: the questionnaire moved from unit-denominated foreign-currency reporting (2011–2013) to thousand-denominated reporting (2014 onward), and some flagged plants in fact continued to report in thousands of pesos. We recover the reporting scale of each foreign-currency plant-year with a two-stage classifier. The first stage compares currency-free unit values (the implied price of electricity in pesos per kWh and of individual fuels, and value added and wages per worker) against the same-year medians of peso-reporting plants in the same industry, and exploits within-plant jumps in these series across the recoding boundaries; the second stage assigns the residual plants using the verified questionnaire regime and the plausibility of their converted per-worker values. Classified plants' monetary variables are converted to Chilean pesos at annual average exchange rates, and conversions implying economically implausible cost shares or per-worker values are discarded. This procedure recovers 92.3% of foreign-currency plant-years (1,610 of 1,744); each recovered plant carries a confidence tier, so the sample can be

rebuilt on the high-confidence subset as a robustness check. After conversion, the wage- and output-per-worker distributions of foreign-currency plants coincide with those of peso-reporting plants up to the expected large-plant premium, with no level break in the aggregate labor share.

ENERGY. We construct a firm-level measure of total energy consumption expressed in tons of oil equivalent (toe) by aggregating all reported energy inputs, namely electricity (bought plus self-generated, net of sales), solid fuels (coal, coke, charcoal), liquid fuels (petroleum, diesel, gasoline, paraffin), liquefied petroleum gas, piped gas, natural gas, and firewood, after converting each source into a common physical unit using standard calorific equivalence factors. Because the units in which some fuels are reported change over survey waves (e.g., kilograms vs. tons for coal, liters vs. thousands of liters for gasoline), we inspect unit-of-measure variables and apply the corresponding rescaling; the calorific factors used to map each source into toe are reported in Table A.6. Observations with implausible energy values (e.g., LPG consumption exceeding 10 million kg) are dropped as outliers.

TABLE A.6—Conversion Factors to Tons of Oil Equivalent (toe)

Energy Source	**Original Unit**	**Conversion Factor**
Electricity	MWh	0.086 toe per MWh
Coal, coke, carboncillo	ton	0.7 toe per ton
Petroleum & diesel	1,000 liters	0.84 toe per 1,000 L
Benzine (gasoline)	1,000 liters	0.75 toe per 1,000 L
Liquefied gas (LPG)	1,000 kg	1.13 toe per 1,000 kg
Gas de cañería (piped gas)	m^3	0.00082 toe per m^3
Firewood	ton	0.3 toe per ton
Natural gas	m^3	0.00082 toe per m^3

Notes: The energy *share* used in the production-function robustness exercises (Appendix 4) is the monetary mirror of the toe basket above: purchased electricity plus the itemized fuel expenditures entering the toe aggregate, net of value-added tax. We build it from itemized fuel values rather than the survey's aggregated fuel total, unreliable for 2001–2007.

TRADE STATUS. We define a plant as an *exporter* if it reports positive export income in a given year, and as an *importer* if it reports positive expenditure on imported raw materials. These are indicator variables constructed from ENIA data.

INDUSTRY CLASSIFICATION. Throughout the paper, manufacturing industries are classified at the 4-digit ISIC Rev. 2 level. Because the modern ENIA uses ISIC Rev. 3, we construct a concordance by computing the modal ISIC Rev. 2 code for each ISIC Rev. 3 industry within the overlapping years where both classifications are available. We further group 4-digit industries into nine broad sectors (food and beverages, textiles, apparel and leather, wood/paper/furniture, chemicals, plastics, non-metallic minerals, metals, and machinery and equipment) for estimation of sector-specific production functions. Industries with very few observations (ISIC 314, 353, 354) and residual categories (ISIC 3903, 3909) are excluded.

4. Deflation

All nominal variables are deflated to express them in constant prices, so that their variation reflects changes in physical quantities rather than price movements. We construct separate deflator series for output, intermediate inputs, and capital.

OUTPUT PRICES. We assemble a continuous industry-level output price deflator by chaining together four sources: (i) for 1979–1986, industry-level deflators from Liu (1993), (ii) for 1986–1995, implicit GDP deflators by sector from the Central Bank of Chile, (iii) for 1995–2007, INE deflators used in Cusolito et al. (2023), and (iv) for 2007–2018, median product-level prices from INE microdata, aggregated to 4-digit industries and winsorized at the 5th and 95th percentiles of annual price variation. All series are rebased to 1980 = 1 using their percentage variations.

INPUT PRICES. Input price deflators are constructed following a similar logic, with an additional step that accounts for the multi-industry composition of intermediate demand. For each industry, we use demand shares from the Chilean input–output tables to compute a weighted average of supplying-industry price indices. This produces an input price index that reflects the actual composition of intermediate inputs purchased by each industry. For 1979–1986, we use Liu (1993)'s input deflators; for later periods, we apply the weighted-average procedure using the same underlying price sources described above.

CAPITAL PRICES. Separate deflators are constructed for buildings, machinery and equipment, and vehicles, drawing on series from both the literature and official sources. We normalize all capital deflators to a 1980 base year. Wages are deflated using the national Consumer Price Index (IPC).

5. Capital Measurement

ENIA reports gross capital stocks by asset type (buildings, machinery, vehicles, and other assets) only intermittently: the variable is available through 2014 but absent in later waves. To obtain a consistent capital series over the full sample, we implement a hierarchical imputation procedure:

1. *Observed stock:* If the gross capital stock for asset type z is reported in year t, we use it directly (after deflation).
2. *Recovery from depreciation:* If the stock is missing but depreciation of asset z is observed, we recover capital as $k_{zt} = \text{dep}_{zt}/\hat{\delta}_z$, where $\hat{\delta}_z$ is the median depreciation rate for asset type z computed from all firm-year observations with non-missing and positive stocks and depreciation. The estimated median rates are 3.1% for buildings, 9.1% for machinery, 14.2% for vehicles, and 22.4% for other assets.
3. *Backward imputation:* For firms whose first observed capital is still missing after step 2, we impute backward from the first available future observation using the law of motion $k_{zt} = k_{z,t+1}/(1-\hat{\delta}_z)^{(t'-t)}$, where t' is the first year with a positive stock.
4. *Steady-state imputation:* If no future observation exists, we impute initial capital using the long-run relationship $k_{z0} = \bar{I}_{zi}/(\hat{\delta}_z + \hat{g}_I)$, where $\bar{I}_{zi}$ is the firm's average investment in asset z and $\hat{g}_I$ is the aggregate investment growth rate in manufacturing.
5. *Prediction from observables:* For the remaining firms with zero or missing capital after all

previous steps, we predict capital from a cross-sectional regression of capital on labor, energy, materials, industry, and year fixed effects.

After obtaining initial capital stocks, investment in each asset type is computed as the implied residual from the law of motion: $I_{zt} = k_{z,t+1} - (1 - \hat{\delta}_z) \cdot k_{zt}$. Total capital is the sum across all asset types, deflated to constant prices.

6. Sample Restrictions and Outlier Treatment

Starting from the linked and harmonized panel, we apply the following restrictions to arrive at the estimation sample:

1. *Non-manufacturing exclusion:* We drop plants whose industry code falls outside the manufacturing range (ISIC 3000–3999) or that cannot be consistently mapped to one of the nine broad sectors.
2. *Missing or non-positive inputs:* We exclude plant-year observations with missing, zero, or negative values for any of the core production variables: total employment, total wages, material inputs, energy consumption, capital, or gross output.
3. *Legal organization:* We drop public enterprises and "other" legal forms (categories 7 and 8 in the ENIA classification), retaining only private firms (natural persons, limited liability companies, public corporations, and cooperatives).
4. *Intensity trimming:* We compute output, capital, materials, and energy per worker, and trim the top and bottom 1% of each ratio's distribution within each year–sector cell. This removes extreme productivity outliers that could distort production function estimates.
5. *Revenue share winsorization:* The labor share (wages/output) and material share (materials/output) are winsorized at the 1st and 99th percentiles within each year–sector cell to bound cost shares in a plausible range.

The resulting estimation sample contains 111,926 plant-year observations, covering approximately 4,000 plants per year. After production function estimation, we further exclude observations with economically implausible output elasticities and trim the 1st and 99th percentiles of the markup and markdown distributions, yielding an effective sample of 85,477 observations used in the main analysis.

7. Trade Data and Instrumental Variables

Industry-level trade exposure is measured using bilateral trade flows from the Atlas of Economic Complexity. We map SITC Rev. 2 product codes (4-digit) to ISIC Rev. 2 industries (4-digit) using existing concordance tables, supplemented by manual classification for approximately 200 unmatched products. All trade values are deflated to 2018 US dollars using the US GDP deflator.

For each industry–year, we compute total Chilean imports from the rest of the world, as well as imports from specific partners (notably China). These measures serve as the endogenous variable in the panel IV local projection framework.

8. Household Survey Data (CASEN)

The CASEN survey provides individual-level data on labor market participation, occupation, industry, and employment status. We use these data to construct two sets of industry-level covariates:

- *Informality and self-employment:* For each 2-digit manufacturing industry and survey year, we compute the share of workers classified as informal or self-employed. These measures capture the availability of outside options for workers in the formal manufacturing sector.
- *Occupational composition:* We compute the share of each ISCO occupation group within each industry, which allows us to characterize the skill and task composition of sectoral labor demand beyond the skilled–unskilled dichotomy available in ENIA.

Because CASEN is conducted at irregular intervals (roughly every two to three years), we linearly interpolate industry-level variables for non-survey years to obtain an annual panel.

9. Employment Protection Index (CBR-LRI)

The employment protection series plotted in Figure 3 and used as the institutional level series in the heterogeneity analysis of Appendix C is the Labour Regulation Index constructed by the Centre for Business Research (CBR), University of Cambridge (Billa et al., 2025), in the leximetric tradition of quantitative labor-law coding whose scope overlaps with the categories analyzed by Botero et al. (2004). The CBR-LRI codes de jure statutory and case law for Chile continuously from 1970 onward (through 2022 in the underlying release; we plot the series through 2020) across 40 individual indicators, each scored on a $[0, 1]$ scale where 0 denotes no protective regulation and 1 denotes the maximum protection observed in the coding scheme. The 40 indicators span five areas of labor law: (i) the definition of the employment relationship and the regulation of part-time, fixed-term, and temporary-agency work, (ii) working time, (iii) individual and collective dismissal, (iv) employee representation, and (v) industrial action. Following the dataset's own recommended aggregation (the compilers report the indicators unweighted, on the grounds that there is no principled basis for a differential weighting scheme), we construct the index used in the paper as the simple average of the 40 indicators in each year. The resulting series is a measure of the intensity of de jure statutory and judicial labor protection; it does not capture enforcement in practice (which we proxy separately with Dirección del Trabajo inspection data, Figure 3c) or the economic cost of compliance with a given rule.

B. Trade Leg: Identification and Robustness

This part collects the identification diagnostics and robustness material for the trade leg of Section V: the derivation of the exposure-robust standard errors from the equivalent shock-level regression, the origin composition of the shift-share instrument, the control-by-control first-stage ladder under exposure-robust inference, the shock-level robustness of the headline IV, pre-trend and shift-share balance checks, and the full heterogeneity term-structures.

1. *Exposure-Robust Inference: The Equivalent Shock-Level Regression*

Subsection B reports Borusyak et al. (2022) exposure-robust standard errors for the trade-leg local projections. This part derives the equivalent shock-level regression that produces them, showing how the industry-by-year specification in equation (7), with its fixed effects, covariates, and employment weights, maps *exactly* onto a regression at the origin-by-industry shock level, and how the exposure-robust standard errors are then nothing more than the ordinary cluster-robust errors of that regression.

THE INDUSTRY-LEVEL ESTIMATOR. Fix a horizon h and estimate equation (7) by two-stage least squares, instrumenting $\ln(IP_{jt})$ with the shift-share Z_{jt} of equation (9), weighting each industry-year (j,t) by its base-period employment weight ω_j, and controlling for $\mathbf{w}_{jt} = \{\alpha_j^h, \delta_t^h, \mathbf{X}_{jt}\}$. For any industry-year variable v_{jt}, let $\widetilde{v}_{jt}$ denote its residual from the ω-weighted linear projection on $\mathbf{w}_{jt}$. By the Frisch–Waugh–Lovell theorem the coefficient of interest is the ratio of ω-weighted sample covariances,

$$\widehat{\beta}^h = \frac{\sum_{j,t}\ \omega_j\, Z_{jt}\, \widetilde{\Delta_h y}_{jt}}{\sum_{j,t}\ \omega_j\, Z_{jt}\, \widetilde{\ln IP}_{jt}}, \tag{B1}$$

where only the outcome and the endogenous regressor need be residualized, not the instrument.

FROM INDUSTRIES TO SHOCKS. Substitute the shift-share $Z_{jt} = \sum_{o \in \mathcal{I}_j} s_{oj}\, g_{ojt}$ into equation (B1) and exchange the order of summation between industry-years (j,t) and origins o. Collecting the shock index $n \equiv (o,j,t)$, write the raw export shock as $g_n \equiv g_{ojt}$ and define the shock-level weight

$$s_n = \frac{\omega_j\, s_{oj}}{\sum_{n'} \omega_{j'}\, s_{o'j'}}, \qquad \text{so that} \quad \sum_n s_n = 1. \tag{B2}$$

Because the unit of observation is the industry itself, each origin-by-industry shock cell $n = (o,j,t)$ draws on the single industry-year (j,t) that sources from origin o; the exposure-weighted average that the aggregation assigns to cell n therefore coincides with the residualized industry-year value, $\overline{v}_n = \widetilde{v}_{jt}$. Equation (B1) then becomes

$$\widehat{\beta}^h = \frac{\sum_n\ s_n\, g_n\, \overline{\Delta_h y}_n}{\sum_n\ s_n\, g_n\, \overline{\ln IP}_n}. \tag{B3}$$

NUMERICAL EQUIVALENCE. Expression (B3) is exactly the second-stage coefficient of the just-identified instrumental-variables regression, run at the origin-by-industry-by-year shock level,

$$\overline{\Delta_h y}_n = a^h + \beta^h \overline{\ln IP}_n + \bar{\varepsilon}^h_n, \tag{B4}$$

in which the aggregated treatment $\overline{\ln IP}_n$ is instrumented by the raw export shock g_n and observations are weighted by s_n. The industry-level estimate from equation (7) and the shock-level estimate from equation (B4) coincide term for term: this is the equivalence of Borusyak et al. (2022), specialized to the present panel through the relabeling $(j,t) \mapsto \ell$ and $(o,j,t) \mapsto n$. The mapping is an algebraic identity (it holds whether or not the identifying assumptions hold), and the estimation confirms it numerically, reproducing $\widehat{\beta}^h$ to a relative difference below 10^{-6} at every horizon before any standard error is formed.

The three ingredients the specification carries (fixed effects, covariates, and weights) each pass through this equivalence cleanly.

FIXED EFFECTS. The industry effects α^h_j and year effects δ^h_t are partialled out of $\Delta_h y_{jt}$ and $\ln(IP_{jt})$ at the industry-year level, inside the residualization $\widetilde{(\cdot)}$, so the shock-level coefficient in equation (B4) is the fixed-effects-adjusted coefficient. Two features make this mapping cleaner than in the canonical incomplete-share application of Borusyak et al. (2022), both consequences of the shares s_{oj} being fixed in the 1985–89 base period and summing to one within each industry's covered support $\mathcal{I}_j$. First, the predetermined covered-import fraction $S_j = \sum_{o \in \mathcal{I}_j} \tilde{s}_{oj}$ is time-invariant and is absorbed by the industry effect α^h_j, which plays the role of the sum-of-shares control that an incomplete-share design would otherwise require. Second, with shares that sum to one, the year effects δ^h_t isolate within-year variation in the shocks directly, without the interaction with S_j that incomplete shares would demand.

COVARIATES. The control vector $\mathbf{X}_{jt}$ contains the two admissible types of control identified by Borusyak et al. (2022). The first are the origin-fundamental exposure-weighted shift-shares, built with the same shares s_{oj} as the instrument, applied to changes in origin characteristics such as real GDP, the terms of trade, and factor shares, which have shift-share structure and correspond to shock-level controls that are functions of the shock index. The second are predetermined industry characteristics (base-period labor-market concentration, capital share, real wages, and their pre-trends), ordinary industry-level controls that are asymptotically uncorrelated with the instrument and serve to sharpen precision. Both blocks are partialled out at the industry-year level inside the residualization, so $\widehat{\beta}^h$ in equation (B4) is the fully-controlled coefficient.

WEIGHTS. The industry-year estimation weights ω_j, each industry's base-period share of manufacturing employment, carry over to the shock level as $s_n \propto \omega_j s_{oj}$ in equation (B2), the product of the industry's regression weight and the origin's base-period exposure share within that industry, renormalized to sum to one across shock cells. Industries with larger base-period employment and origins with larger base-period import shares receive proportionally more weight, so equation (B4) reproduces the employment weighting of equation (7) rather than imposing an unweighted shock-level design. These are precisely the average-exposure weights $s_n = \sum_\ell e_\ell s_{\ell n}$ of Borusyak et al. (2022).

EXPOSURE-ROBUST STANDARD ERRORS. Writing $\widehat{\beta}^{h}$ as the coefficient of a regression whose unit of observation is the shock cell makes the appropriate variance estimator transparent. The export-supply shocks g_{ojt} recur across every industry that sources from a given origin o, so the instrument Z_{jt} and the residual ε^{h}_{jt} are correlated across all industries exposed to o; standard errors that cluster on the industry (outcome) dimension treat those industries as independent and understate sampling uncertainty. The exposure-robust standard error is instead the conventional cluster-robust standard error of the shock-level regression equation (B4), clustered on the shock origin o, the level at which the identifying variation is assigned. This is the Borusyak et al. (2022) prescription: once the estimator is expressed at the level of the shocks, ordinary heteroskedasticity- or cluster-robust inference applied there is valid, with no bespoke variance formula. Clustering on origin is material rather than cosmetic, because the identifying variation loads on few effective origins (Appendix 2): the effective number of independent shocks (the inverse Herfindahl index $1/\sum_{n} s_{n}^{2}$ of the shock weights) is far below the raw count of shock cells. We report these origin-clustered exposure-robust standard errors throughout; Adão et al. (2019) provide an asymptotically equivalent alternative. Their validity rests on the shift-exogeneity assumptions set out in Subsection C, not on the algebraic equivalence in equations (B1)–(B4), which holds unconditionally.

APPLICATION TO THE LOCAL PROJECTIONS. The equivalence is applied horizon by horizon. At each h the reported impulse response is the industry-level point estimate $\widehat{\beta}^{h}$ from equation (7), and its exposure-robust band is $\widehat{\beta}^{h} \pm 1.645\,\mathrm{SE}^{h}$, where SE^{h} is the origin-clustered standard error from the shock-level regression equation (B4) formed from the horizon-h long difference $\Delta_{h} y_{jt}$. Joint inference across horizons and the average response are obtained from a single stacked shock-level regression with horizon-interacted shocks under one origin-clustered covariance matrix. First-stage strength is assessed at the same level: the shift-share-appropriate effective F-statistic is the squared t-statistic on g_{ojt} in the s_{n}-weighted regression of $\overline{\ln IP}_{n}$ on the shock (Table 2), following Borusyak et al. (2022, §5.2).

This equivalence has a further consequence for a level (rather than cumulative long-difference) specification of the markdown–import-penetration relationship: re-estimating the location-level instrumental-variables regression on the restricted sample of CIIUs with a positive base-period import share, using the shares-and-shocks-reconstructed instrument and weighted by the normalized exposure weights s_{n}, reproduces the shock-level coefficient exactly, confirming equation (B3)–(B4) numerically. Collapsing the shock to its origin-aggregate component only (discarding the origin-by-industry interaction) makes the first stage collapse entirely (Kleibergen–Paap $F = 0.2$ and 0.1 across two instrument constructions), confirming that it is the interaction of origin export shifts with predetermined industry sourcing shares, not raw origin variation, that identifies import penetration.

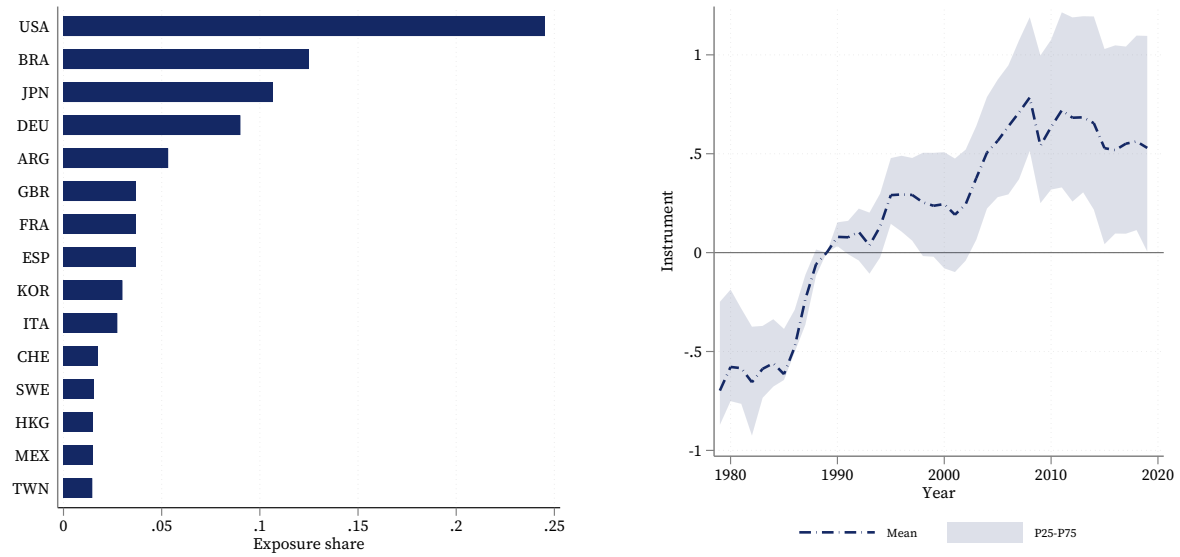
2. *Origin Composition of the Instrument*

The instrument's identifying variation is not spread evenly across Chile's trading partners; it is concentrated in a handful of large, diversified exporters. Figure B.1 ranks origins by their base-period exposure weight, and Table B.1 summarizes the concentration. The United States alone carries 24.5% of the exposure weight, followed by Brazil (12.5%), Japan (10.7%), Germany (9.0%),

and Argentina (5.3%); the five largest origins account for 62% of the total. The corresponding Herfindahl index is 0.106, an effective number of origins of only about nine.

This concentration is the empirical reason the exposure-robust inference of Subsection B is necessary rather than cosmetic. When the identifying shocks load on so few effective origins, industries that share those origins in their base-period import mix inherit correlated instrument variation, and standard errors that treat industries as independent overstate the design's precision. The shock-level standard errors we report as the headline are built precisely to accommodate this few-effective-origins structure.

FIGURE B.1—Top Origin Countries by Instrument Exposure, and Its Evolution over Time



(a) Top origin countries by exposure weight

(b) Instrument evolution over time

Notes: (a) Top 15 origin countries by base-period exposure weight in the headline instrument; the five largest account for 62% of exposure (Herfindahl 0.106). (b) Evolution of the exposure-weighted instrument over time, at the CIIU × year level: the line is the cross-industry mean, and the band spans the 25th–75th percentile across industries. *Source:* Atlas of Economic Complexity; authors' calculations.

TABLE B.1—Concentration of instrument exposure (intensive-margin, family D)

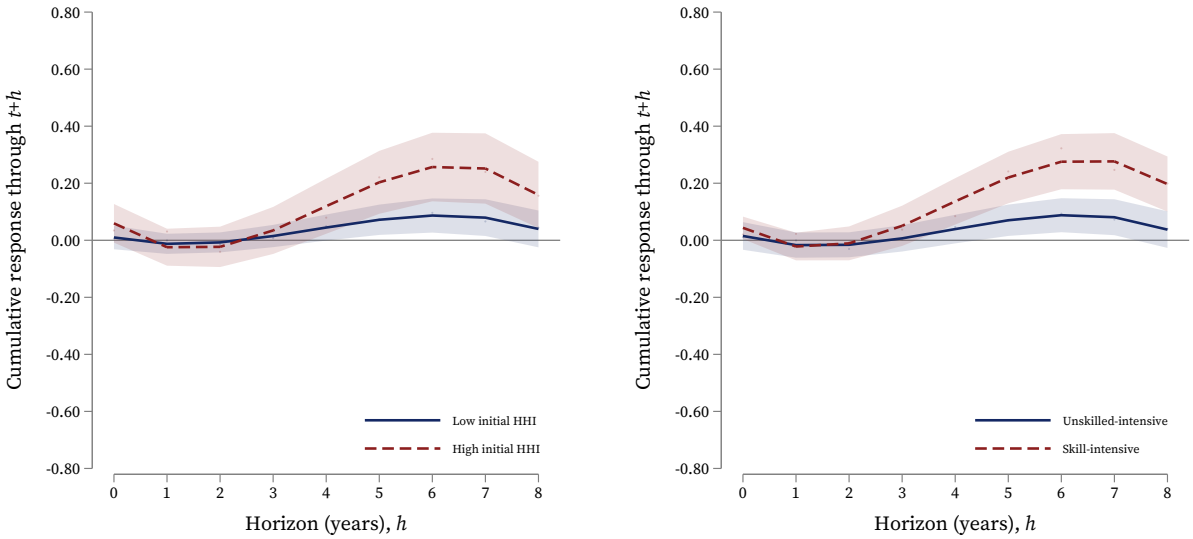
Statistic	Value
Origins with base-period support	83
CIIU manufacturing industries	87
Herfindahl index (import-wt)	0.1060
Effective number of origins (import-wt)	9.44
Effective number of origins (equal-wt)	12.53
Largest single origin weight (import-wt)	0.2451

Notes: HHI = $\sum_o w_o^2$ over origin exposure weights w_o (summing to one); effective number of origins = 1/HHI. A low effective number means a handful of countries carry the identifying variation.

3. *Trade Heterogeneity: Full Results*

Figure B.2 plots the split-sample markdown impulse responses. Tables B.2–B.4 report the full interacted LP-IV heterogeneity at horizons $h = 0$, 4, and 8, with the seven local-projection outcomes as columns and two panels for the pre-determined dimensions that the selection mechanism predicts: base-period labor-market concentration (Panel A) and skill intensity (Panel B). Each panel reports the base-group slope on import penetration, the group interaction (difference), and the implied slope for the high group, under exposure-robust (Borusyak–Hull–Jaravel) standard errors, together with the shock-level Kleibergen–Paap F.

FIGURE B.2—Import Penetration and Wage Markdowns, by Predetermined Industry Group



(a) High vs. low initial concentration

(b) High vs. low skill intensity

Notes: Split-sample smoothed LP-IV impulse responses of log employment-weighted markdowns to import penetration, by pre-determined (1990–92) industry group; each panel overlays the two group-specific responses. Estimation as in Figure 6; the shaded 90% bands are exposure-robust (Borusyak et al., 2022), computed at the level of the underlying origin shocks. Interaction coefficients and shock-level Kleibergen–Paap F statistics appear in Tables B.2–B.4.

4. *Synthetic Control*

Table B.5 summarizes the exercise: Chile's realized import growth splits into the component reproduced by the synthetic (comparable-country) path, the tailwind common to global integration, and the residual divergence attributable to Chile's own faster liberalization, which disciplines the counterfactual trade cost $\tilde{\delta}$ in Section E.

Table B.6 reports the weights assigned to each donor unit in constructing the synthetic import series, along with the balance check of the resulting estimates.

Table 4 provides the balance checks of synthetic counterfactual and realized imports of Chile with respect to the main covariates:

TABLE B.2—Heterogeneous Responses to Import Penetration at Horizon $h = 0$

	Markdown	Markup	Concentr.	Labor sh.	Wages	Employ.	Exit
Panel A: Base-period labor-market concentration							
ln Import penetration (low)	0.006	-0.027	0.014	0.038**	0.000	0.013	0.000
	(0.018)	(0.035)	(0.030)	(0.019)	(0.005)	(0.016)	(0.003)
× High concentration	0.028	-0.084	-0.091*	-0.020	0.013	-0.040	-0.006
	(0.027)	(0.055)	(0.048)	(0.022)	(0.009)	(0.042)	(0.004)
Implied (high group)	0.034	-0.110*	-0.077	0.018	0.013	-0.027	-0.006
(BHJ SE)	(0.034)	(0.060)	(0.058)	(0.027)	(0.012)	(0.049)	(0.005)
KP F (shock)	6.1	6.1	6.1	6.1	6.1	6.1	6.1
Observations	66,176	66,176	66,176	66,176	66,176	66,176	66,176
Panel B: Skill intensity							
ln Import penetration (low)	0.010	-0.047	0.064***	0.047**	0.004	0.025	0.001
	(0.022)	(0.032)	(0.020)	(0.019)	(0.007)	(0.022)	(0.004)
× Skill-intensive	0.010	-0.010	-0.194***	-0.038**	-0.001	-0.060**	-0.005
	(0.016)	(0.032)	(0.037)	(0.016)	(0.007)	(0.030)	(0.003)
Implied (high group)	0.020	-0.057*	-0.130***	0.009	0.003	-0.035	-0.004
(BHJ SE)	(0.020)	(0.030)	(0.043)	(0.019)	(0.009)	(0.035)	(0.003)
KP F (shock)	4.8	4.8	4.8	4.8	4.8	4.8	4.8
Observations	66,176	66,176	66,176	66,176	66,176	66,176	66,176

Notes: Each column is a separate BHJ (Borusyak et al., 2022) shock-level interacted local-projection IV of the cumulative response through $t + h$ to import penetration. Panel A splits at the employment-weighted median of base-period (1990–92) concentration; Panel B at the median skilled wage-bill share. Reports base-group slope, interaction, and implied high-group slope; exposure-robust (origin-clustered) SEs. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE B.3—Heterogeneous Responses to Import Penetration at Horizon $h = 4$

	Markdown	Markup	Concentr.	Labor sh.	Wages	Employ.	Exit
Panel A: Base-period labor-market concentration							
ln Import penetration (low)	0.049**	-0.060	0.169	-0.024	-0.029***	0.080**	0.005
	(0.020)	(0.065)	(0.129)	(0.031)	(0.009)	(0.035)	(0.005)
× High concentration	0.030	-0.248**	-0.208***	-0.235**	0.070***	-0.224	-0.015
	(0.041)	(0.114)	(0.076)	(0.097)	(0.025)	(0.218)	(0.015)
Implied (high group)	0.079	-0.307**	-0.039	-0.259**	0.042	-0.144	-0.009
(BHJ SE)	(0.049)	(0.129)	(0.118)	(0.111)	(0.032)	(0.233)	(0.017)
KP F (shock)	9.1	9.1	9.1	9.1	9.1	9.1	9.1
Observations	56,816	56,816	56,816	56,816	56,816	56,816	56,816
Panel B: Skill intensity							
ln Import penetration (low)	0.047**	-0.082	0.234*	-0.018	-0.007	0.080*	0.008
	(0.021)	(0.055)	(0.123)	(0.028)	(0.016)	(0.044)	(0.006)
× Skill-intensive	0.038	-0.145*	-0.428***	-0.238***	-0.014	-0.206	-0.023
	(0.046)	(0.080)	(0.120)	(0.091)	(0.029)	(0.158)	(0.015)
Implied (high group)	0.085*	-0.228**	-0.194**	-0.256***	-0.021	-0.126	-0.015
(BHJ SE)	(0.050)	(0.089)	(0.094)	(0.099)	(0.033)	(0.177)	(0.016)
KP F (shock)	13.4	13.4	13.4	13.4	13.4	13.4	13.4
Observations	56,816	56,816	56,816	56,816	56,816	56,816	56,816

Notes: Each column is a separate BHJ (Borusyak et al., 2022) shock-level interacted local-projection IV of the cumulative response through $t + h$ to import penetration. Panel A splits at the employment-weighted median of base-period (1990–92) concentration; Panel B at the median skilled wage-bill share. Reports base-group slope, interaction, and implied high-group slope; exposure-robust (origin-clustered) SEs. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE B.4—Heterogeneous Responses to Import Penetration at Horizon $h = 8$

	Markdown	Markup	Concentr.	Labor sh.	Wages	Employ.	Exit
Panel A: Base-period labor-market concentration							
ln Import penetration (low)	0.044	-0.122	0.217	0.084**	-0.085***	-0.031	0.018***
	(0.037)	(0.090)	(0.136)	(0.036)	(0.016)	(0.044)	(0.006)
× High concentration	0.112*	-0.293***	-0.286**	-0.406***	-0.020	0.019	0.009
	(0.059)	(0.102)	(0.121)	(0.104)	(0.026)	(0.135)	(0.013)
Implied (high group)	0.156**	-0.415***	-0.069	-0.321***	-0.105***	-0.011	0.026*
(BHJ SE)	(0.067)	(0.105)	(0.124)	(0.117)	(0.026)	(0.150)	(0.016)
KP F (shock)	13.3	13.3	13.3	13.3	13.3	13.3	13.3
Observations	47,456	47,456	47,456	47,456	47,456	47,456	47,456
Panel B: Skill intensity							
ln Import penetration (low)	0.040	-0.143*	0.175	0.065**	-0.060**	-0.028	0.017***
	(0.036)	(0.076)	(0.149)	(0.032)	(0.026)	(0.041)	(0.006)
× Skill-intensive	0.159***	-0.265***	-0.156	-0.414***	-0.146***	0.010	0.015
	(0.045)	(0.084)	(0.126)	(0.155)	(0.047)	(0.159)	(0.015)
Implied (high group)	0.199***	-0.408***	0.019	-0.349**	-0.205***	-0.018	0.032*
(BHJ SE)	(0.057)	(0.094)	(0.125)	(0.175)	(0.062)	(0.176)	(0.018)
KP F (shock)	9.3	9.3	9.3	9.3	9.3	9.3	9.3
Observations	47,456	47,456	47,456	47,456	47,456	47,456	47,456

Notes: Each column is a separate BHJ (Borusyak et al., 2022) shock-level interacted local-projection IV of the cumulative response through $t + h$ to import penetration. Panel A splits at the employment-weighted median of base-period (1990–92) concentration; Panel B at the median skilled wage-bill share. Reports base-group slope, interaction, and implied high-group slope; exposure-robust (origin-clustered) SEs. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE B.5—Synthetic Control: Decomposition of Chilean Import Growth

	$\Delta \ln$ Imports 1990–95 to 2013–18	Cumulative factor $e^{\Delta \ln}$	Share of realized
Chile (observed)	1.60	4.9×	1.000
Synthetic Chile	1.19	3.3×	0.745
Difference	0.41	1.5×	0.255

Notes: Synthetic control estimated on log imports (Penn World Table 10.01), matched over the pre-treatment period 1951–1990. Columns report the change in log imports between the 1990–95 and 2013–18 averages ($\Delta \ln$), the implied cumulative import factor $e^{\Delta \ln}$, and each path's $\Delta \ln$ as a share of Chile's realized growth. The Difference row nets out the small 1990–95 gap, isolating the post-1990 divergence. The Synthetic path is the component common to global integration; the Difference is attributable to Chile's own faster liberalization and disciplines the counterfactual trade cost δ_{synth} in the structural model (Section E).

TABLE B.6—Synthetic Control: Donor Weights

Country	Weight	Country	Weight
Argentina	0.000	Kenya	0.000
Australia	0.000	Luxembourg	0.000
Austria	0.000	Mauritius	0.000
Belgium	0.000	Mexico	0.000
Bolivia	**0.284**	Morocco	0.000
Brazil	0.000	Netherlands	**0.264**
Canada	0.000	New Zealand	0.000
Colombia	0.000	Norway	0.000
Costa Rica	0.000	Paraguay	0.000
Cyprus	**0.116**	Peru	0.000
Denmark	0.000	Philippines	0.000
Dominican Republic	0.000	Portugal	0.000
Ecuador	0.000	South Africa	0.000
Egypt	0.000	Spain	0.000
Finland	0.000	Sri Lanka	**0.245**
France	0.000	Sweden	0.000
Germany	0.000	Switzerland	0.000
Greece	0.000	Taiwan	0.000
Guatemala	**0.092**	Thailand	0.000
Iceland	0.000	Trinidad and Tobago	0.000
India	0.000	Turkey	0.000
Ireland	0.000	United Kingdom	0.000
Israel	0.000	United States	0.000
Italy	0.000	Uruguay	0.000
Japan	0.000	Venezuela	0.000

Notes: Donor weights in the synthetic control for Chile, which minimizes the pre-treatment (1951–1990) distance between Chile and the weighted donor combination in (log) imports and covariates. Positive weights (in bold) fall on 5 countries; all other donors receive zero weight. Pre-treatment RMSPE = 0.18. *Source*: Penn World Table 10.01.

C. Institutions: Employment Protection and the Markdown Gradient

The trade design of Section V deliberately absorbs the national institutional path through year fixed effects, so that the identified object is the differential response of markdowns to import penetration. Here we ask the complementary question that the year fixed effects leave open: as employment protection, the minimum wage, and enforcement rose over 1990–2018, did the markdown decline fall *differentially* across industries whose pre-determined characteristics make monopsony power stronger? Employment protection is a natural candidate channel: in search-and-matching environments, layoff costs and severance obligations raise workers' effective bargaining strength (Ljungqvist, 2002; Garibaldi and Violante, 2005), which in an oligopsonistic labor market operates directly on the wedge between marginal product and the wage. Figure 3 documents the trajectory of employment protection alongside the two other main institutional levers discussed in the main text.

TABLE B.7—Synthetic Control Balance Check

	Chile		
	Chile	Synthetic	Difference
ln(Imports)	22.590	22.613	-0.024
ln(GDP)	24.855	24.104	0.751
ln(Human Capital Index)	1.157	1.085	0.072
ln(Population)	2.362	2.037	0.325
ln(Exports)	22.263	22.613	-0.350
ln(Labor Share)	0.390	0.439	-0.049

Notes: The table reports average values of predictors for Chile, its synthetic control, and their difference (Chile minus Synthetic). Variables are averaged through the pre-treatment period 1951-1990. Imports covariates for each pre-treatment year are also included, but omitted for presentation reasons.

A positive answer would show the institutional force operating along the same margins as the trade force, though, because both gradients load on the same predetermined characteristics, this design does not separate the two forces from each other, and would anticipate the complementarity the structural model formalizes in Section VI. Aghelmaleki et al. (2022) document a related pattern across European countries, where the employment effects of the China trade shock are moderated by cross-country differences in EPL stringency; the exercise here differs in studying within-country industry heterogeneity in a single country's markdown response rather than cross-country heterogeneity in employment.

We relate the log employment-weighted markdown to three national institutional series, each in logs: the CBR-LRI employment-protection index (Botero et al., 2004; Billa et al., 2025), the real minimum wage, and Dirección del Trabajo labor inspections per firm. Table C.1 reports the level ladder, each series entered cumulatively, alone and then jointly, with baseline controls and industry (CIIU) fixed effects added in turn. *We do not read any of these level coefficients as the effect of a given institution.* All three series are national, one value per year, and correlated with each other and with the linear trend at 0.94–0.99; net of the trend, a level coefficient is a suggestive time-series association whose standard error is understated in the Moulton sense and which cannot be separated from other national trends of the period, trade opening or the schooling expansion among them. Causal identification of the underlying channel is available elsewhere: exploiting a 1999 notice-period reform in the Netherlands as a natural experiment, van der Wiel (2010) finds that an additional month of statutory notice raises wages by about three percent, direct micro evidence for the bargaining-power channel that these national time series can only suggest. We therefore report the full level ladder, and state these limitations, only in Appendix C; in the body we read the *gradient*, the cross-industry differential in the markdown–institution relationship, which the industry fixed effects and the common trend leave identified.

Because the three institutional levels move together so closely, jointly interacting all three with pre-determined (1990–92) industry indicators for high skill intensity and high labor-market concentration is a demanding specification. Table C.2 builds this joint interaction up from no fixed effects to industry fixed effects, then toggles each interaction block off in turn to check

TABLE C.1—Employment Protection, Inspections, and the Minimum Wage—OLS

	(1)	(2)	(3)	(4)	(5)
ln(CBR-LRI EPL index)	-1.202***	-1.132***	-1.108***	-1.002***	-0.893***
	(0.161)	(0.365)	(0.303)	(0.292)	(0.277)
ln(real minimum wage)				-0.528***	-0.455***
				(0.177)	(0.170)
ln(inspections per firm)					-0.085**
					(0.037)
Linear year trend	-0.001	-0.003	-0.005	0.011	0.013
	(0.003)	(0.006)	(0.004)	(0.008)	(0.008)
Industry fixed effects	None	None	Industry	Industry	Industry
Controls	No	Yes	Yes	Yes	Yes
Observations	1,885	1,885	1,885	1,885	1,885
R^2	0.080	0.170	0.800	0.803	0.803

Notes: Dependent variable is the log employment-weighted wage markdown, CIIU×year, 1990–2018. Columns (1)–(5) cumulatively add the EPL index, controls, industry FE, minimum wage, and inspections per firm, as indicated in the table. Each institutional series is national, so its level coefficient is a suggestive time-series association net of a linear trend. Employment-weighted; SEs clustered by industry (CIIU). * p < 0.10, ** p < 0.05, *** p < 0.01.

whether one characteristic is standing in for the other. None of the six level×characteristic interactions is estimated with much precision once all three institutional levels compete for the same identifying variation; the minimum-wage×skill interaction alone reaches marginal (ten-percent) significance once the concentration interactions are removed. Table C.3 instead examines each institutional series on its own, without the other two as competing regressors, against a fuller set of four pre-determined characteristics: skill intensity, female employment share, labor-market concentration, and export exposure. Freed from the collinearity that limits Table C.2, the skill and concentration gradients are precisely estimated and consistent across all three institutional measures.

The institution-era markdown decline is *flatter*, less steep, in high-skill and high-concentration industries. In Table C.3, the interaction of high skill intensity with the institutional level is positive and significant for every one of the three measures (0.585 for the EPL index, 0.208 for the minimum wage, and 0.099 for inspections, all significant at the five-percent level or better), and so is the interaction with high initial concentration (0.424, 0.161, and 0.075, respectively, significant at the ten-percent level or better). Because the base slope on each institutional level is negative, the implied slope for an industry that is both high-skill and high-concentration is markedly flatter than for a low-skill, low-concentration industry, consistently across all three institutions. Neither the female-share nor the export-exposure interaction is significant for any of the three. The joint specification of Table C.2 is directionally consistent with this pattern in most columns but cannot, on its own, distinguish the individual interactions from zero, the cost of asking three near-collinear national trends to identify six interaction terms at once.

TABLE C.2—Employment Protection, Inspections, and the Minimum Wage—Skill and Concentration Interactions

	(1)	(2)	(3)	(4)	(5)
ln(CBR-LRI EPL index)	-0.595	-0.408	-0.619*	-0.701**	-0.731**
	(0.582)	(0.563)	(0.316)	(0.293)	(0.304)
ln(real minimum wage)	-0.528	-0.476	-0.527***	-0.473***	-0.525***
	(0.369)	(0.346)	(0.156)	(0.170)	(0.154)
ln(inspections per firm)	-0.114	-0.089	-0.092*	-0.089**	-0.090**
	(0.118)	(0.111)	(0.046)	(0.043)	(0.043)
High skill intensity (>= median)	-2.971	-3.116			
	(3.183)	(2.947)			
High initial concentration (>= median)	-2.687	-3.251			
	(3.587)	(3.284)			
ln(CBR) × high skill	-0.456	-0.479	-0.310		-0.433
	(0.874)	(0.821)	(0.341)		(0.348)
ln(min. wage) × high skill	0.259	0.268	0.238**		0.280**
	(0.350)	(0.321)	(0.117)		(0.131)
ln(inspections) × high skill	0.040	0.050	0.031		0.047
	(0.166)	(0.152)	(0.064)		(0.063)
ln(CBR) × high concentration	-0.656	-0.735	-0.511	-0.600	
	(1.003)	(0.928)	(0.389)	(0.402)	
ln(min. wage) × high concentration	0.207	0.260	0.182	0.255	
	(0.394)	(0.360)	(0.146)	(0.163)	
ln(inspections) × high concentration	0.093	0.100	0.069	0.078	
	(0.187)	(0.174)	(0.056)	(0.057)	
Linear year trend	0.016	0.006	0.011	0.012	0.012
	(0.011)	(0.011)	(0.008)	(0.008)	(0.008)
Industry fixed effects	None	None	Industry	Industry	Industry
Controls	No	Yes	Yes	Yes	Yes
Observations	1,885	1,885	1,885	1,885	1,885
R^2	0.136	0.211	0.812	0.808	0.810

Notes: Dependent variable is the log employment-weighted wage markdown, CIIU×year, 1990–2018, regressed on ln(institution)×characteristic interactions; columns (1)–(2) have no industry FE, columns (3)–(5) add them. High skill/concentration = p50 splits of base-period (1990–92) levels (median assigned to the high group). Employment-weighted; columns (1)–(2) report robust SEs (cluster matrix singular), columns (3)–(5) CIIU-clustered SEs. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE C.3—Employment Protection, Inspections, and the Minimum Wage—Heterogeneity by Industry Characteristic

	(1) Index	(2) Min. Wage	(3) Inspections
ln(institution)	-1.450***	-0.839***	-0.234***
	(0.290)	(0.322)	(0.049)
High skill intensity (>= median)	0.349	-1.969**	-0.461**
	(0.221)	(0.830)	(0.197)
High female employment share (>= median)	-0.220	0.401	0.079
	(0.263)	(0.871)	(0.203)
High initial concentration (>= median)	0.222	-1.547**	-0.375**
	(0.229)	(0.659)	(0.156)
High base-period gross exports (>= median)	0.211	-0.671	-0.124
	(0.329)	(0.723)	(0.238)
× high skill	0.585**	0.208**	0.099***
	(0.238)	(0.094)	(0.037)
× high female share	-0.225	-0.049	-0.034
	(0.259)	(0.099)	(0.040)
× high concentration	0.424*	0.161**	0.075*
	(0.245)	(0.075)	(0.039)
× high exports	0.268	0.075	0.034
	(0.392)	(0.082)	(0.060)
Linear year trend	-0.001	0.012	0.000
	(0.003)	(0.011)	(0.003)
Observations	1,856	1,856	1,856
R^2	0.148	0.146	0.147

Notes: Dependent variable is the log employment-weighted wage markdown, CIIU×year, 1990–2018. Each column uses one national institutional series (CBR-LRI index, minimum wage, or inspections) with a linear trend and four characteristic interactions (p50 splits of 1990–92 levels). No controls, no fixed effects. Employment-weighted; columns (1),(3) CIIU-clustered SEs, column (2) robust (cluster matrix singular). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

We read the gradient as a differential association consistent with institutional reform compressing employer wage-setting power most where monopsony power is strongest (the same complementarity the structural model delivers) rather than as an isolated causal effect of any single institution. It is the empirical counterpart, on the institutional side, to the trade heterogeneity of Section 3: both forces bear most on the industries where pre-determined skill and concentration make labor-market power greatest.

D. Measurement: Descriptives and Robustness

This subsection reports supplementary tables referenced in the main text.

1. *Descriptive Statistics*

This part compiles additional information about descriptive statistics of main data sources. Table D.1 shows the summary of the main variables at the plant-level by decades:

TABLE D.1—Descriptive Statistics

	(1) All	(2) 1990s	(3) 2000s	(4) 2010s
Log Gross Output Value	13.53	13.07	13.60	14.24
	(1.70)	(1.55)	(1.72)	(1.66)
Labor	85.45	79.88	79.95	105.04
	(191.75)	(153.67)	(184.52)	(254.78)
Skilled Labor	39.89	26.78	42.05	60.09
	(119.26)	(73.10)	(118.68)	(172.94)
Female Labor	19.35	17.15	18.08	25.55
	(62.10)	(46.73)	(63.27)	(81.25)
Log Materials	12.62	12.25	12.64	13.25
	(1.79)	(1.61)	(1.82)	(1.85)
Log Capital Stock	9.98	9.90	9.60	10.77
	(2.02)	(1.80)	(2.06)	(2.08)
Energy	762.56	315.99	499.89	2026.67
	(22818.25)	(2725.22)	(5233.01)	(47955.71)
Legal Person Firm	0.20	0.21	0.21	0.17
	(0.40)	(0.40)	(0.41)	(0.38)
Public Corporation	0.76	0.71	0.78	0.82
	(0.42)	(0.45)	(0.41)	(0.38)
Cooperatives	0.01	0.01	0.01	0.00
	(0.09)	(0.11)	(0.09)	(0.05)
State and Others	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)
Exporter	0.21	0.20	0.21	0.23
	(0.41)	(0.40)	(0.41)	(0.42)
Importer	0.23	0.23	0.21	0.24
	(0.42)	(0.42)	(0.41)	(0.43)
Observations	114,825	46,113	43,403	25,309

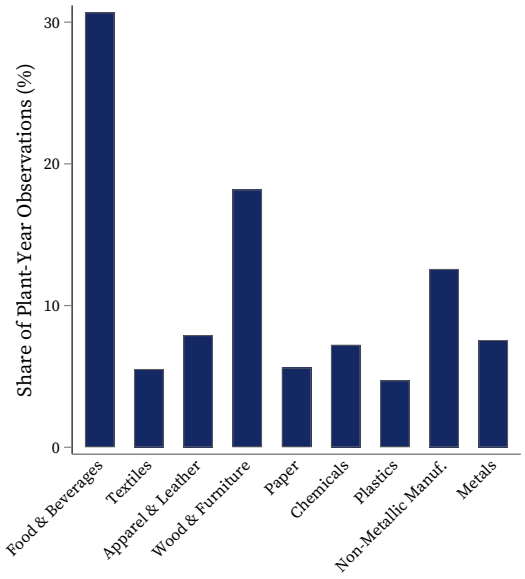
Notes: Mean values reported with standard deviations below.

The next figure illustrates the distribution of plants by industry:

2. *Input-Output Elasticities*

Table D.2 reports additional elasticity estimates used in the construction of firm-level markups and markdowns. These estimates complement the baseline results discussed in Section 3 and are used in robustness exercises.

FIGURE D.1—Overall Distribution of Firms by Sector



Notes: Distribution of manufacturing plants across 2-digit ISIC Rev. 2 sectors over the full sample period 1990–2018.
Source: ENIA.

TABLE D.2—Output Elasticities and Returns to Scale by Sector

	Elasticities				Returns
	Materials	Capital	Labor	Energy	Scale
Food and Beverages	0.56	0.05	0.40	0.06	1.07
	(0.14)	(0.03)	(0.15)	(0.03)	(0.05)
Textiles	0.53	0.04	0.43	0.04	1.05
	(0.13)	(0.02)	(0.16)	(0.02)	(0.06)
Apparel	0.52	0.04	0.56	0.03	1.15
	(0.13)	(0.02)	(0.15)	(0.02)	(0.04)
Paper and Woods	0.62	0.04	0.39	0.03	1.10
	(0.13)	(0.01)	(0.15)	(0.02)	(0.03)
Chemicals	0.52	0.17	0.31	0.06	1.06
	(0.09)	(0.05)	(0.11)	(0.04)	(0.04)
Plastics	0.55	0.07	0.39	0.04	1.06
	(0.14)	(0.01)	(0.16)	(0.01)	(0.02)
Non-Metallic Manufactures	0.58	0.07	0.51	0.06	1.22
	(0.19)	(0.04)	(0.22)	(0.03)	(0.07)
Metals	0.44	0.10	0.51	0.05	1.09
	(0.05)	(0.05)	(0.04)	(0.03)	(0.05)
Machinery and Equipment	0.50	0.06	0.49	0.05	1.09
	(0.15)	(0.03)	(0.17)	(0.01)	(0.04)
Average	0.55	0.06	0.43	0.05	1.09
	(0.14)	(0.04)	(0.16)	(0.03)	(0.06)

Notes: Economic Sectors are defined at the 3-digit ISIC revision 3 level.

3. *Validity of the Identifying Assumptions*

The wage markdown is recovered in Section II as the gap between labor's output elasticity and its revenue share. It therefore requires neither wage data nor an assumed model of labor-market conduct, and does not presuppose the monopsony power it is meant to measure; errors that move the estimated elasticity and the implied markup in the same direction partly cancel in this ratio. That structure makes the markdown more robust to violations of the underlying assumptions than each assumption taken in isolation might suggest, but the assumptions still matter for the level of the elasticity itself, and this appendix discusses them in the context of this paper rather than only in the abstract.

Conditions (ii)–(iv) of the footnote to Section II (on the timing of capital accumulation, the time-invariance of production-function parameters within a sector, and the first-order Markov evolution of productivity) are generic to any application of the Ackerberg et al. (2015) (ACF) proxy-variable estimator and are no more or less plausible in Chilean plant-level data than elsewhere. Conditions (v)–(vi), that materials (the flexible input used to identify elasticities) has productivity as the only unobservable in its demand equation and that this demand is strictly monotone in productivity, are the standard proxy-variable requirements; following Ackerberg et al. (2015), placing labor in the second stage of the estimation routine, rather than treating it as the flexible proxy as in Levinsohn and Petrin (2003), avoids the functional-dependence critique that arises when the proxy input and the input of ultimate interest are the same choice variable. Yeh et al. (2022) make the same argument for the same estimator on data of the same kind (an annual, plant-level, gross-output manufacturing panel with materials as the flexible input) and conclude the conditions are defensible; we rely on that precedent while being explicit that their setting is the U.S. Longitudinal Business Database rather than Chilean ENIA, so the argument transfers by analogy rather than by direct evidence for this sample.

The analogy is least comfortable on condition (v), for a reason specific to this paper's design. The identifying variation used throughout the paper is a trade-liberalization episode that changes the relative price of imported inputs, including materials, unevenly across plants. If part of that price variation is unobserved and correlated with a plant's import exposure, productivity is no longer the only unobservable in the materials-demand equation, and condition (v) fails in a way correlated with the paper's own regressor of interest, the concern De Loecker et al. (2016) raise for markup estimation under trade reform generally, and it would operate on the estimated elasticity in the numerator of the markdown rather than on the revenue share in the denominator, so it is not obviously absorbed by the ratio structure described above. They address the analogous concern with plant-level output and input prices, which separate physical productivity from markups directly; ENIA reports no plant-level prices for either outputs or inputs, only the industry-level deflators described in Section III, so that approach is not available here. Relatedly, Bond et al. (2021) show that revenue-based elasticities recover the product of the quantity elasticity and the inverse markup rather than the quantity elasticity itself when plant-level output prices are unobserved, which applies to the revenue-based ENIA data used throughout, a limitation this paper shares with most of the applied markup and markdown literature working from national industrial surveys, including Yeh et al. (2022). We cannot rule out either concern directly in this sample and report them as limitations rather than as

resolved; as a partial, indirect check, one of the seven alternative estimators in Appendix 4 replaces materials with energy as the flexible input and yields a very similar markdown, at least consistent with identification not resting solely on the input most directly exposed to the trade shock.

A separate concern is specific to the two-input construction used in this paper: Raval (2023b) shows, in manufacturing data that includes Chile, that markups recovered from different flexible inputs need not agree with one another. This matters here because materials identify the markup while labor's own first-order condition identifies the markdown, so the two objects are not built from the same flexible input and their mutual consistency is not guaranteed by the estimation procedure. We do not have a flexible input independent of both materials and energy against which to test this cross-input consistency directly in this sample; Raval (2023a) and Doraszelski and Jaumandreu (2021) propose estimators that relax the common-technology assumption underlying this critique, and adapting either to this setting is left to future work.

4. Production Function and Markdown Robustness

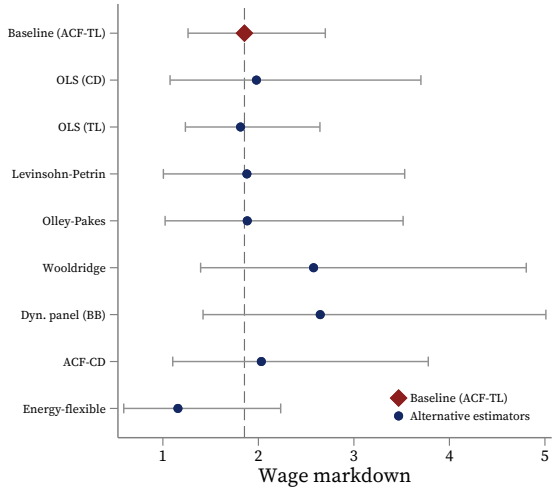
Markups and wage markdowns are functions of estimated output elasticities (Section II), so a natural concern is whether the results are an artifact of the baseline production-function estimator, a sector-specific translog gross-output specification estimated by the Ackerberg et al. (2015) procedure with lagged-investment instruments. We assess this by re-estimating the production function sector by sector with seven alternative estimators and rebuilding the wage markdown for each with the identical Yeh et al. (2022) formula, trimming rules, and estimation window as the baseline.

The alternatives span the main families of the production-function literature: ordinary least squares in Cobb–Douglas and translog form, the control-function proxy estimators of Olley and Pakes (1996) and Levinsohn and Petrin (2003), the GMM proxy estimator of Wooldridge (2009), a Blundell–Bond-type dynamic-panel estimator (Bond, 2005; Rovigatti and Mollisi, 2018), and a De Loecker and Warzynski (2012)-style GMM estimator in pure Cobb–Douglas form, estimated by the same in-house routine as the Ackerberg et al. (2015) translog baseline. A final specification instead treats energy as the flexible input, reusing the baseline elasticities. Every specification conditions on year fixed effects only, so that differences across rows reflect the estimator rather than the conditioning set.

The estimated labor elasticity and the resulting markdown are stable across estimators, and the implied labor-supply elasticities remain in the same range as the baseline. Figure D.2 summarizes the level and dispersion of firm-level markdowns by estimator; the baseline lies near the center of the distribution of point estimates. Figure D.3 reports the employment-weighted wage markdown by sector in levels, with 95% plant-cluster bootstrap confidence bands, in the same fashion as the whole-manufacturing series.

As is standard for gross-output proxy estimators, the flexible-input elasticity is identified partly off functional form (Gandhi et al., 2020); we therefore report point estimates without bootstrap standard errors, consistent with the baseline.

FIGURE D.2—Distribution of Firm-Level Wage Markdowns by Estimator



Notes: Markers are the median firm-level wage markdown for each estimator; horizontal whiskers span the 25th–75th percentiles of the pooled (cross-sector, cross-year) firm-level distribution. The baseline estimator (diamond) and its median (dashed vertical line) are highlighted. The bottom row treats energy rather than materials as the flexible input, using the baseline elasticities. Each estimator's statistics are computed over its own cleaned plant-year sample.

5. *Determinants of Monopsony Power*

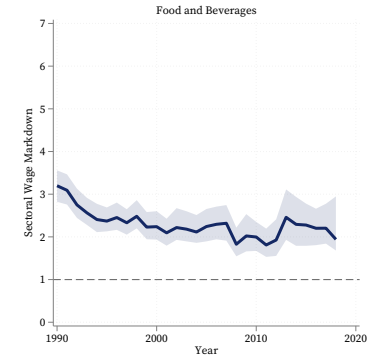
This subsection presents the full set of results on firm-level and industry-level determinants of wage markdowns summarized in Section IV. The analysis explores which characteristics of firms and industries are systematically associated with the degree of monopsony power, complementing the aggregate patterns documented in the main text.

FIRM-LEVEL DETERMINANTS. Table D.3 reports regressions of firm-level wage markdowns on measures of labor composition and firm attributes. The dependent variable is the estimated plant-level wage markdown. Explanatory variables include the female share of employment, the share of unskilled workers, the share of subcontracted workers, legal organization dummies, trade status (importer and exporter), and firm size proxied by the capital stock. All specifications include year fixed effects; columns progressively add sector and geographic fixed effects.

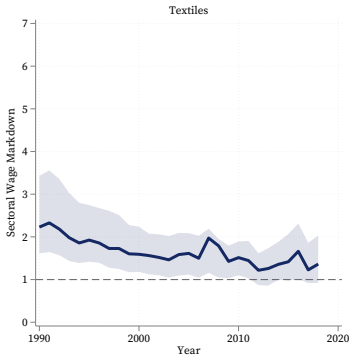
Several patterns are noteworthy. First, the female share of labor is positively and significantly correlated with monopsony power across all specifications. However, its coefficient declines substantially once sector fixed effects are included, suggesting that roughly three-quarters of the initial correlation reflects the disproportionate concentration of female employment in high-markdown sectors rather than a within-sector relationship. This finding is consistent with evidence in the literature that positions the relatively high monopsony exposure of women as a significant contributor to the observed gender wage gap (Sharma et al., 2023; Blau and Kahn, 2017).

Second, the relative presence of unskilled workers exhibits a similar pattern. Sectoral selection accounts for approximately half of the raw correlation, indicating that unskilled workers

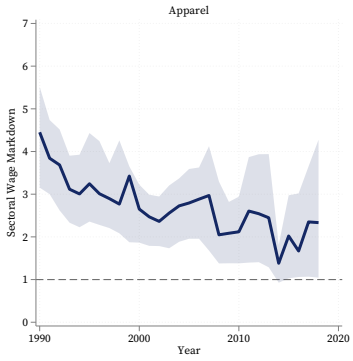
FIGURE D.3—Wage Markdown by Sector with Bootstrap Confidence Bands



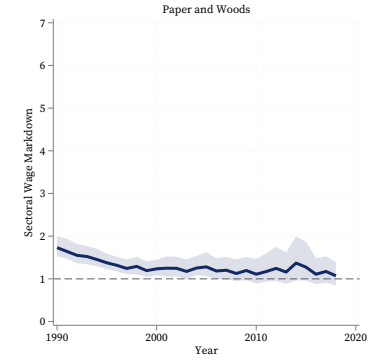
(a) Food and Beverages



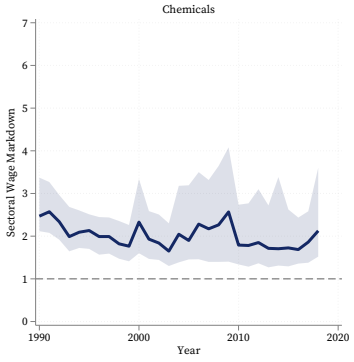
(b) Textiles



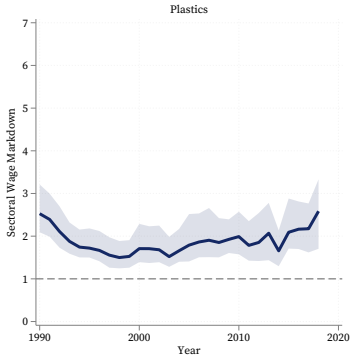
(c) Apparel



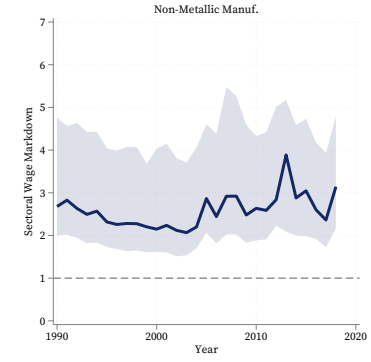
(d) Paper and Wood Products



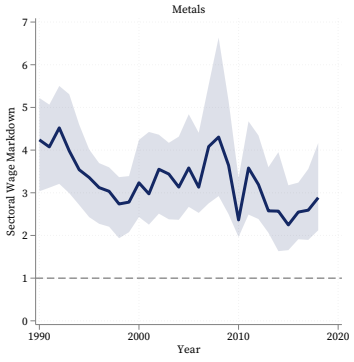
(e) Chemicals



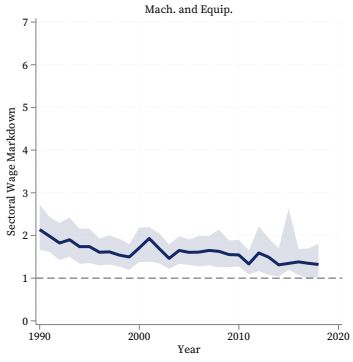
(f) Plastics



(g) Non-Metallic Mineral Manufacturing



(h) Metals



(i) Machinery and Equipment

Notes: Employment-weighted wage markdown by 2-digit manufacturing sector, 1990–2018, lowess-smoothed (navy line) with 95% plant-cluster bootstrap confidence bands (shaded, same smoothing) following Yeh et al. (2022), on a common vertical scale, in the same fashion as the whole-manufacturing series in Figure 5(b). The dashed horizontal line marks a markdown of one.

TABLE D.3—Labor Markdown and Firm-Level Covariates

	Years: 1990–2018				Years: 2000–2018			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Size(K)	0.063***	0.074***	0.069***	0.186***	0.089***	0.097***	0.080***	0.167***
	(0.009)	(0.009)	(0.008)	(0.011)	(0.010)	(0.011)	(0.008)	(0.014)
Women/Labor	0.759***	0.840***	0.438***	0.595***	0.585***	0.607***	0.486***	0.401***
	(0.057)	(0.056)	(0.059)	(0.114)	(0.067)	(0.066)	(0.081)	(0.149)
Unskilled/Labor	0.575***	0.310***	0.231***	0.079	0.195***	0.180***	0.138***	0.037
	(0.047)	(0.047)	(0.044)	(0.052)	(0.058)	(0.057)	(0.053)	(0.070)
Exports/Production	0.054	0.120**	0.126**	0.158	0.157**	0.131**	0.100	0.040
	(0.053)	(0.052)	(0.051)	(0.102)	(0.066)	(0.065)	(0.065)	(0.132)
Collectives	-0.736***	-0.881***	-1.157***		-0.399***	-0.440***	-0.736***	
	(0.088)	(0.096)	(0.105)		(0.151)	(0.152)	(0.173)	
Limited Companies	-0.191***	-0.198***	-0.197***		-0.096**	-0.104***	-0.109***	
	(0.027)	(0.028)	(0.025)		(0.038)	(0.038)	(0.034)	
Subcontracted/Labor					-0.055	-0.087	0.324***	0.873***
					(0.126)	(0.123)	(0.101)	(0.173)
Year FE	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Sector FE	No	No	Yes	No	No	No	Yes	No
Firm FE	No	No	No	Yes	No	No	No	Yes
Adjusted R^2	0.030	0.055	0.258	0.623	0.024	0.029	0.237	0.606
Observations	96542	96542	96542	95051	58575	58575	58575	56929

Notes: Dependent variable is the labor markdown. Columns (1)–(4) use the full sample 1990–2018; columns (5)–(8) restrict to 2000–2018 and add the subcontracted-labor share. Labor-weighted; robust standard errors. Size is the log of capital stock. Organizational form: baseline is sole proprietorships (Persona Natural), vs. Limited Companies (incorporated limited-liability entities) and Collectives (partnerships/cooperatives). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

also sort into high-markdown industries, though to a lesser extent than female workers. On average, after controlling for industry composition, an additional percentage point of unskilled workers in the firm's labor force is associated with a 0.96% increase in wage markdowns.

Third, subcontracted workers, despite positive coefficients, do not show a statistically significant correlation with wage markdowns. We cannot reject the null hypothesis of zero effect, suggesting that subcontracted labor constitutes a second-order determinant.

Regarding firm characteristics, importer status appears uncorrelated with markdowns, whereas exporter firms exhibit approximately 0.08 points lower monopsony power. This result may seem counterintuitive, as exporters tend to be larger, more productive, and pay higher wages. Conditioning on size, which captures the primary source of monopsony power, the remaining negative coefficient reflects the higher wages paid by exporters rather than a direct effect on their wage-setting power.

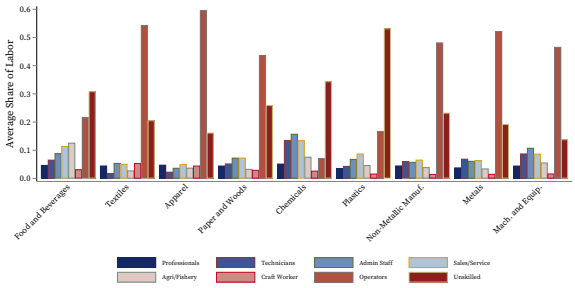
Finally, legal organization matters. Public corporations and cooperatives exhibit lower markdowns than legal person firms, which are not statistically different from other (less frequent) categories. Public corporations often provide additional benefits beyond wages, reducing the *effective* markdown. Cooperatives allow employees to share in firm earnings, providing an additional mechanism for mitigating monopsony power. These effects are quantitatively mean-

ingful, offsetting the markdown premia associated with nearly 40 percentage points of unskilled workers or 20 percentage points of the female labor share.

INDUSTRY-LEVEL DETERMINANTS. Firm-level evidence indicates that sectoral composition plays a substantial role in determining monopsony exposure. To explore this further, we complement ENIA with industry-level measures of informality, self-employment, and occupational composition derived from the CASEN household survey. Because ENIA does not provide detailed occupational information, CASEN, which reports individual occupations classified by ISCO codes and industry affiliation by ISIC, allows us to link industry-level wage markdowns with labor supply characteristics.

Figure D.4 illustrates the occupational profile of each 2-digit manufacturing sector. On average, manufacturing relies heavily on operators and unskilled workers, who account for roughly 50–70% of total employment.

FIGURE D.4—Sectoral Occupational Composition



Notes: Average share of sectoral employment accounted for by each occupational category, by 2-digit ISIC Rev. 2 manufacturing sector: professionals, technicians and mid-level professionals, administrative staff, service and sales workers, skilled agricultural and fishery workers, mechanical and craft workers, machine operators and assemblers, and unskilled laborers (ISCO-based; managerial and military categories excluded, as in the analysis below). *Source:* CASEN household survey, 1990–2019; authors' calculations.

To explore the relationship between industry-level monopsony power and these covariates, we estimate:

$$y_{jt} = \alpha_j + \alpha_t + \beta_1 \text{Informality}_{jt} + \beta_2 \text{Female}_{jt} + \beta_3 \text{Schooling}_{jt} + \sum_k \beta_{4k} \text{Occupation}_{k,jt} + \varepsilon_{jt},$$

where y_{jt} is the industry-level log wage markdown and occupation shares are normalized to sum to one, using unskilled workers as the baseline.[37] Table D.4 summarizes the results.

Consistent with the firm-level evidence, the female share and schooling are positively and negatively correlated with monopsony power, respectively. Informality and self-employment exhibit robust negative coefficients, suggesting that more attractive outside options for workers reduce monopsony power in the formal sector (Amodio et al., 2025).

Columns (4)–(6) incorporate the occupational profile. Relative to unskilled workers, nearly all occupations are associated with lower wage markdowns, with the exception of mechanical and

[37] Military and managerial categories are excluded: the former is irrelevant for manufacturing, and the latter is rarely observed in CASEN for these industries.

TABLE D.4—Plant-level Markdowns Heterogeneity by Skill Composition

	Log Aggregated Markdowns					
	(1)	(2)	(3)	(4)	(5)	(6)
Informality and Self Employment	-0.43***	-0.32***	-0.18***	-0.19***	-0.19***	-0.18***
	(0.06)	(0.06)	(0.06)	(0.07)	(0.07)	(0.07)
Female Share	0.31***	0.37***	0.17	0.12	0.12	0.13
	(0.06)	(0.05)	(0.16)	(0.16)	(0.16)	(0.16)
Schooling	-0.08***	-0.05***	0.01			0.01
	(0.01)	(0.01)	(0.01)			(0.01)
Scientific Professionals				-0.27*	-0.27*	-0.30**
				(0.14)	(0.14)	(0.15)
Technicians and Mid-level Professionals				0.28**	0.28**	0.22*
				(0.11)	(0.11)	(0.12)
Administrative Staff				-0.25**	-0.25**	-0.28**
				(0.12)	(0.12)	(0.12)
Service and Sales Workers				-0.22*	-0.22*	-0.24*
				(0.12)	(0.12)	(0.12)
Skilled Agricultural and Fishery Workers				-0.10	-0.10	-0.12
				(0.12)	(0.12)	(0.13)
Mechanical and Craft Workers				0.15	0.15	0.15
				(0.14)	(0.14)	(0.14)
Machine Operators and Assemblers				-0.08	-0.08	-0.09
				(0.08)	(0.08)	(0.08)
Year FE	No	Yes	Yes	Yes	Yes	Yes
Sector FE	No	No	Yes	Yes	Yes	Yes
Adjusted R^2	0.058	0.094	0.638	0.639	0.639	0.639
N	2098	2098	2098	2070	2070	2070

Notes: Markdowns are aggregated at the 4-digit CIIU rev. 2 level. Robust Standard errors are included. Share of occupations is normalized to sum 1. Omitted Category is *Unskilled Workers*. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

craft workers, for whom we find a positive correlation, possibly reflecting that these workers are more productive than unskilled workers but do not receive commensurate wage compensation. Occupations requiring higher education, such as scientific professionals, skilled agricultural and fishery workers, and administrative staff, are strongly associated with lower markdowns, consistent with greater bargaining power. Somewhat surprisingly, technical professionals do not differ markedly from unskilled workers in their exposure to monopsony power.

6. *Chilean Governments*

Table D.5 summarizes the political, economic and social orientation of Chilean governments during years 1973–2018:

TABLE D.5—Political, Economic, and Social Orientation of Chilean Governments Since 1973

Period	Government / Coalition	Political Orientation	Economic Orientation	Social Orientation
1973–1990	Military Dictatorship (Augusto Pinochet)	Authoritarian, military regime; suppression of political parties and civil liberties	Radical market liberalization; trade openness; privatization of state-owned enterprises; pension, labor, and education reforms inspired by neoliberal doctrine	Severe repression of unions and civil society; sharp reduction in social protection; targeted poverty alleviation only in late 1980s
1990–1994	Patricio Aylwin (Concertación)	Center-left democratic transition; institutional continuity with gradual political reform	Macroeconomic orthodoxy; preservation of market institutions; cautious regulatory strengthening	Expansion of social spending; poverty reduction programs; emphasis on reconciliation and human rights
1994–2000	Eduardo Frei Ruiz-Tagle (Concertación)	Center-left, technocratic governance	Export-led growth; infrastructure concessions; fiscal discipline	Incremental expansion of education and health coverage; continued poverty reduction
2000–2006	Ricardo Lagos (Concertación)	Center-left, social democratic orientation	Market-based economy with stronger regulation; competition policy; public investment	Expansion of social insurance; labor market reforms; strengthening of public health guarantees (AUGE)
2006–2010	Michelle Bachelet I (Concertación)	Center-left; emphasis on social inclusion	Countercyclical fiscal policy; sovereign wealth funds; continuity of open-economy model	Pension reform; expansion of childcare and social protection; gender-focused policies
2010–2014	Sebastián Piñera I (Coalición por el Cambio)	Center-right, pro-business	Market-oriented policies; reconstruction spending after 2010 earthquake; tax adjustments	Targeted social programs; limited redistribution; emphasis on efficiency and growth
2014–2018	Michelle Bachelet II (Nueva Mayoría)	Center-left to left; reform-oriented coalition	Structural reforms within market framework; tax reform; education finance overhaul	Expansion of social rights; free tertiary education for lower-income students; labor reform (Ley 20.940)

Notes: Author's periodization and characterization of Chilean national governments and their policy orientation, 1973–2018, compiled from the standard political-history periodization to contextualize the trade-liberalization and labor-market reforms discussed in the text.

E. The Model: Derivations and Equilibrium

This subsection collects the cost-minimization derivation of markups and markdowns used in the measurement framework, the national-income and Nash-bargaining step-by-step derivations, the model's graphical schematics, and the firm-level and general-equilibrium solution algorithm.

1. *Markups and Markdowns*

Consider the firm cost minimization problem given by:

$$\begin{aligned} \min_{X^v_{it}, X^f_{it}} \quad & P^v_{it} X^v_{it} + P^f_{it} X^f_{it} \\ \text{s.t.} \quad & F(X^v_{it}, X^f_{it}) = Q_{it} \end{aligned}$$

The associated Lagrangian is:

$$\mathcal{L} : P^v_{it} X^v_{it} + P^f_{it} X^f_{it} + \lambda_{it} \left[Q_{it} - F(X^v_{it}, X^f_{it}) \right]$$

Taking the First Order Condition with respect to the input X^v_{it}:

$$P^v_{it} - \lambda_{it} \frac{\partial F(X^v_{it}, X^f_{it})}{\partial X_{it^v}} = 0$$

Rearranging and multiplying both sides for $\frac{X^v_{it}}{Q_{it}}$:

$$\frac{P^v_{it}}{\lambda_{it}} \frac{X^v_{it}}{Q_{it}} = \frac{\partial F(X^v_{it}, X^f_{it})}{\partial X^v_{it}} \frac{X^v_{it}}{Q_{it}}$$

Now, we define the product markup as: $\mu_{it} \equiv \frac{P_{it}}{\lambda_{it}}$, where λ_{it} corresponds to the marginal cost since $\frac{\partial \mathcal{L}}{\partial Q_{it}} = \lambda_{it}$. In addition, we multiply the left side of the equation by a convenient one $(\frac{P_{it}}{P_{it}})$ to get:

$$\mu_{it} \left(\frac{P^v_{it}}{P_{it}} \frac{X^v_{it}}{Q_{it}} \right) = \underbrace{\frac{\partial F(X^v_{it}, X^f_{it})}{\partial X^v_{it}} \frac{X^v_{it}}{Q_{it}}}_{\theta^v_{it}}$$

Finally, we solve for μ_{it} to get the resulting expression for the markup:

$$\mu_{it} = \theta^v_{it} \frac{P_{it} Q_{it}}{P^v_{it} X^v_{it}} = \frac{\theta^v_{it}}{\alpha^v_{it}}$$

Where α^v_{it} denotes the revenue share of input X^v_{it}.

Wage Markdowns. Consider the following conditional optimization problem:

$$\begin{aligned} \min_{l_{it}} \quad & w_{it}(l_{it}) l_{it} \\ \text{s.t.} \quad & F(l_{it}, X^*_{-l,it}) = Q_{it} \end{aligned}$$

Where $X^*_{-l,it}$ is the vector of already optimized inputs, except for l_{it}. The Lagrangian of this problem is:

$$\mathcal{L} : w_{it}(l_{it})l_{it} + \lambda_{it}\left[Q_{it} - F(l_{it}, X^*_{-l,it})\right]$$

Taking the first order condition with respect to l_{it} gives:

$$\frac{\partial w_{it}(l_{it})}{\partial l_{it}} l_{it} + w_{it}(l_{it}) - \lambda_{it}\frac{\partial F(l_{it}, X^*_{-l,it})}{\partial l_{it}} = 0$$

Rearranging and dividing both sides by $w_{it}(l_{it})$:

$$\frac{\partial w_{it}(l_{it})}{\partial l_{it}} \frac{l_{it}}{w_{it}(l_{it})} + 1 = \lambda_{it}\frac{\partial F(l_{it}, X^*_{-l,it})}{\partial l_{it}} \frac{1}{w_{it}(l_{it})}$$

Multiplying by some convenient ones we can obtain:

$$\underbrace{\underbrace{\frac{\partial w_{it}(l_{it})}{\partial l_{it}} \frac{l_{it}}{w_{it}(l_{it})}}_{\varepsilon_s^{-1}(l_{it})} + 1}_{\nu_{it}} = \underbrace{\frac{\lambda_{it}}{P_{it}}}_{\mu_{it}^{-1}} \underbrace{\frac{\partial F(l_{it}, X^*_{-l,it})}{\partial l_{it}} \frac{l_{it}}{Q_{it}}}_{\theta^l_{it}} \underbrace{\frac{P_{it}Q_{it}}{w_{it}(l_{it})l_{it}}}_{\alpha^l_{it}}$$

$$\nu_{it} = \mu_{it}^{-1}\frac{\theta^l_{it}}{\alpha^l_{it}}$$

Where $\varepsilon_s^{-1}(l_{it})$ is the firm's perceived inverse elasticity of labor supply, ν_{it} is the wage markdown, μ_{it} the final good markup, θ^l_{it} the production elasticity to labor and α^l_{it} the revenue labor share.

Finally, if we replace the markup expression above, we get:

$$\nu_{it} = \left(\frac{\theta^v_{it}}{\alpha^v_{it}}\right)^{-1} \frac{\theta^l_{it}}{\alpha^l_{it}}$$

This is the wage markdown expression we were looking for.

2. Static Model: National Income

Let E and D denote the subsets of exporting and domestic producers, respectively. From the national budget constraint we know that consumption (C) equals income (Inc):

$$P_c C = Inc = \underbrace{\sum_{ij\in E} \phi Inc}_{\text{Local revenues of exporters}} + \underbrace{\sum_{ij\in E} a_{ij}\left[n_{ij} - n^D_{ij}\right]\delta P^*}_{\text{International revenues of exporters}} + \underbrace{\sum_{ij\in D} a_{ij} N_{ij}\frac{P^*}{\delta}}_{\text{Domestic revenues local producers}}$$

$$Income(1 - \phi N_{exp} + \delta^2 \phi N_{exp}) = \sum_{ij\in E} a_{ij} n_{ij}\delta P^* + \sum_{ij\in D} a_{ij} N_{ij}\frac{P^*}{\delta}$$

$$\text{Income} = \frac{\sum_{ij\in E} a_{ij} n_{ij} \delta P^* + \sum_{ij\in D} a_{ij} N_{ij} \frac{P^*}{\delta}}{(1 - \phi N_{exp} + \delta^2 \phi N_{exp})}$$

Where $N_{exp} \equiv \sum_{ij\in E} \mathbb{1}$ represents the total number of exporter firms.

3. Static Model: Nash-Bargaining

Given W_{ij}, employment is $N^S_{ij}(W_{ij})$ from the supply curve (17). The surpluses are:

UNION'S SURPLUS.

$$\Omega_{ij}(W_{ij}) = W_{ij} \cdot N^S_{ij}(W_{ij}). \tag{E1}$$

FIRM'S SURPLUS.

$$\pi_{ij}(W_{ij}) = (P_{ij}\, a_{ij} - W_{ij}) \cdot N^S_{ij}(W_{ij}). \tag{E2}$$

DISAGREEMENT PAYOFFS. $\Omega^0_{ij} = \pi^0_{ij} = 0$.

DEFINITION 2 (Nash Bargaining). *The bargained wage solves:*

$$W^*_{ij} = \arg \max_{W_{ij}\in(0,\, P_{ij} a_{ij})} \left[\Omega_{ij}(W_{ij})\right]^{\beta} \cdot \left[\pi_{ij}(W_{ij})\right]^{1-\beta}. \tag{E3}$$

Substituting (E1) and (E2):

$$\mathcal{N}(W_{ij}) = \left[W_{ij} \cdot N^S_{ij}(W_{ij})\right]^{\beta} \cdot \left[(P_{ij}\, a_{ij} - W_{ij}) \cdot N^S_{ij}(W_{ij})\right]^{1-\beta}. \tag{E4}$$

FIRST-ORDER CONDITION. Taking logarithms of (E4):

$$\ln \mathcal{N} = \beta \ln \Omega_{ij} + (1-\beta) \ln \pi_{ij}. \tag{E5}$$

Differentiating each term with respect to W_{ij}:

UNION TERM. Applying the product rule to $\Omega_{ij} = W_{ij} \cdot N^S_{ij}$:

$$\frac{d\Omega_{ij}}{dW_{ij}} = N^S_{ij} + W_{ij} \frac{dN^S_{ij}}{dW_{ij}}. \tag{E6}$$

Dividing by $\Omega_{ij} = W_{ij}\, N^S_{ij}$:

$$\frac{d \ln \Omega_{ij}}{dW_{ij}} = \frac{1}{W_{ij}} + \frac{1}{N^S_{ij}} \frac{dN^S_{ij}}{dW_{ij}}. \tag{E7}$$

FIRM TERM. Applying the product rule to $\pi_{ij} = (P_{ij} a_{ij} - W_{ij}) \cdot N_{ij}^S$:

$$\frac{d\pi_{ij}}{dW_{ij}} = -N_{ij}^S + (P_{ij} a_{ij} - W_{ij}) \frac{dN_{ij}^S}{dW_{ij}}. \tag{E8}$$

Dividing by $\pi_{ij} = (P_{ij} a_{ij} - W_{ij}) N_{ij}^S$:

$$\frac{d \ln \pi_{ij}}{dW_{ij}} = \frac{-1}{P_{ij} a_{ij} - W_{ij}} + \frac{1}{N_{ij}^S} \frac{dN_{ij}^S}{dW_{ij}}. \tag{E9}$$

ASSEMBLING THE FOC. Substituting (E7) and (E9) into the derivative of (E5):

$$\beta \left[\frac{1}{W_{ij}} + \frac{1}{N_{ij}^S} \frac{dN_{ij}^S}{dW_{ij}} \right] + (1 - \beta) \left[\frac{-1}{P_{ij} a_{ij} - W_{ij}} + \frac{1}{N_{ij}^S} \frac{dN_{ij}^S}{dW_{ij}} \right] = 0. \tag{E10}$$

The labor supply response $\frac{1}{N_{ij}^S} \frac{dN_{ij}^S}{dW_{ij}}$ appears in both terms with weights β and $(1 - \beta)$:

$$\frac{\beta}{W_{ij}} - \frac{1 - \beta}{P_{ij} a_{ij} - W_{ij}} + \underbrace{\left[\beta + (1 - \beta) \right]}_{=1} \cdot \frac{1}{N_{ij}^S} \frac{dN_{ij}^S}{dW_{ij}} = 0. \tag{E11}$$

Applying the Elasticity Definition:

$$\frac{1}{N_{ij}^S} \frac{dN_{ij}^S}{dW_{ij}} = \frac{\epsilon_{ij}}{W_{ij}}. \tag{E12}$$

Substituting (E12) into (E11):

$$\frac{\beta}{W_{ij}} + \frac{\epsilon_{ij}}{W_{ij}} = \frac{1 - \beta}{P_{ij} a_{ij} - W_{ij}} \tag{E13}$$

which simplifies to:

$$\frac{\beta + \epsilon_{ij}}{W_{ij}} = \frac{1 - \beta}{P_{ij} a_{ij} - W_{ij}}. \tag{E14}$$

SOLVING FOR THE WAGE MARKDOWN. Cross-multiplying (E14):

$$(\beta + \epsilon_{ij})(P_{ij} a_{ij} - W_{ij}) = (1 - \beta) W_{ij}. \tag{E15}$$

Expanding:

$$(\beta + \epsilon_{ij}) P_{ij} a_{ij} = W_{ij} \left[\beta + \epsilon_{ij} + 1 - \beta \right] = W_{ij} (1 + \epsilon_{ij}). \tag{E16}$$

Dividing both sides by $(1 + \epsilon_{ij}) P_{ij} a_{ij}$ and defining the markdown $\mu_{ij} \equiv W_{ij} / (P_{ij} a_{ij})$:

(E17) $$\mu_{ij}^{*} = \frac{\epsilon_{ij} + \beta}{1 + \epsilon_{ij}}$$

The equilibrium wage is:

(E18) $$W_{ij}^{*} = \frac{\epsilon_{ij} + \beta}{1 + \epsilon_{ij}} \cdot P_{ij}\, a_{ij}.$$

Survey truncation and the selection correction. Because the calibration targets the ENIA take-home share, the ten-employee threshold biases it: the smallest, most competitive firms are unobserved, so the ENIA-only markdown overstates the economy-wide object. The correction is *model-based*. Sub-threshold firms are assigned the take-home share that the calibrated oligopsony map implies for a near-zero-payroll-share firm, $(\eta+\beta)/(1+\eta)$, and the economy-wide aggregate splices this imputed fringe onto the surveyed sample, employment-weighted. Table E.1 reports the magnitude: the truncation overstates the markdown by 8.5% in 1990 (the aggregate worker share rises from 0.385 to 0.417 once the fringe is spliced in) but by only 4.4% by 2018 (0.511 to 0.534). The bias is therefore material early in the sample, when small low-markdown firms account for a larger share of activity, and modest by its end; crucially, it does not reverse the direction of any result.

TABLE E.1—Wage Markdowns: ENIA, Competitive Fringe, and the Economy-Wide Aggregate

	1990	2018
Worker take-home share $\bar{\mu}$ *(employment-weighted)*		
ENIA ($\geq$ 10 workers, targeted)	0.385	0.511
Competitive fringe ($<$ 10, model-implied)	0.577	0.649
Aggregate (economy-wide)	0.417	0.534
Implied markdown $1/\bar{\mu}$		
ENIA-only	2.60	1.96
Economy-wide (ENIA + fringe)	2.40	1.87
Selection bias from the 10-worker truncation		
Worker-share understatement ($\bar{\mu}_{agg} - \bar{\mu}_{ENIA}$)	0.033	0.022
Markdown overstatement (%)	8.50	4.38

Notes: $\bar{\mu}$ is the worker take-home share (inverse markdown), employment-weighted within each group. ENIA covers surveyed plants with $\geq$ 10 workers and is the sample to which (δ, β) are calibrated. The fringe (sub-10-worker plants, ~62% of firms, ~3% of sales) is assigned the model-implied worker share from the calibrated oligopsony map evaluated at its (near-zero) market share, so $\bar{\mu}_{fringe} \to (\eta+\beta)/(\eta+1)$. The aggregate splices ENIA and fringe. The selection bias is the gap between the ENIA-only and economy-wide objects.

4. Firm-Level Equilibrium

As shown in Berger et al. (2022a), the labor market block (given exporter status in our case) can be solved independently from aggregates, because firm-level equilibrium is determined by the following expressions:

$$\text{(E19)} \qquad \mu(s_{ij}) = \frac{\epsilon(s_{ij})}{1+\epsilon(s_{ij})}, \qquad W_{ij} = \frac{P_{ij}a_{ij}}{1+1/\epsilon_{ij}}, \qquad \frac{1}{\epsilon_{ij}} = \frac{1-s_{ij}}{\eta} + \frac{s_{ij}}{\theta},$$

Here s_{ij} is also a function of μ_{ij} and parameters. Given labor market shares, one can compute the labor supply elasticity and firm-level wage, then aggregate wages and employment at the industry and national levels.

Our model differs because capacity constraints arise from the fixed national demand for each variety and the resulting variable marginal productivity around the export cutoff. Solving the model also requires knowledge of each firm's exporting status. We therefore proceed in two layers:

1. Inner Layer (Labor): Given exporting status, solve for labor market equilibrium by iterating on labor market shares until convergence. Then compute all aggregate labor market variables.
    - Flag firms whose labor demand is inconsistent with their current exporting status.
2. Outer Layer (Trade): Starting with an initial guess of exporting status, iterate by updating the status of firms whose labor demand is inconsistent with their current classification.
    - Some firms are "constrained": large enough to export without transport costs, but too small to export profitably with trade costs. These firms oscillate between exporter and non-exporter status across iterations. We identify them and impose that they demand exactly the labor required to satisfy local demand, paying wages at the non-exporter marginal productivity.
3. Upon convergence of the Outer Layer, we obtain the final exporting status and compute the labor market equilibrium one final time to obtain equilibrium prices and quantities.

Figures E.1, E.2, and E.3 collect the model's graphical schematics (the bargained equilibrium on the labor supply curve, the four limiting cases of Corollary 1, and the three export regimes of Section VI), referenced from the main text.

5. General equilibrium

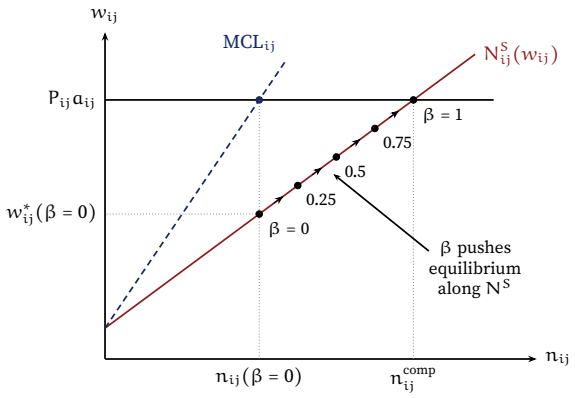
Given firm-level equilibrium, we aggregate to obtain the following steady-state quantities Y, N, C and wages W:

Output and aggregate budget constraint:

$$P_c C = Inc = \underbrace{\sum_{ij\in E} \phi Inc}_{\text{Local revenues of exporters}} + \underbrace{\sum_{ij\in E} a_{ij}\left[n_{ij} - n_{ij}^D\right]\delta P^*}_{\text{International revenues of exporters}} + \underbrace{\sum_{ij\in D} a_{ij} N_{ij} \frac{P^*}{\delta}}_{\text{Domestic revenues local producers}}$$

Wages:

FIGURE E.1—Equilibrium on the Supply Curve



Notes: As β increases from 0 to 1, the equilibrium moves along the supply curve from the monopsony point (where MCL = $P_{ij}a_{ij}$) to the competitive point (where $W^*_{ij} = P_{ij}a_{ij}$). At each β, $w^*_{ij} = [(\epsilon_{ij} + \beta)/(1 + \epsilon_{ij})]\, P_{ij}a_{ij}$. Both wages and employment increase monotonically in β.

$$W_j = \left(\sum_i^{M_{ji}} W_{ij}^{1+\eta} \right)^{1/(1+\eta)}, W = \left(\sum_j^{M_j} W_j^{1+\theta} \right)^{1/(1+\theta)}$$

Labor:

$$N_j = \left(\sum_i^{M_{ji}} N_{ij}^{\frac{1+\eta}{\eta}} \right)^{\eta/(1+\eta)}, N = \left(\sum_j^{M_j} N_j^{\frac{1+\theta}{\theta}} \right)^{\theta/(1+\theta)}$$

We also define the following indices:

Productivity:

$$A = \left[\sum_{j=1}^{M_j} a_j^{1+\theta} \right]^{\frac{1}{1+\theta}}, \qquad a_j = \left[\sum_{i=1}^{M_i} a_{ij}^{1+\eta} \right]^{\frac{1}{1+\eta}}$$

Markdown:

$$\mu = \left[\sum_{j=1}^{M_j} \left(\frac{a_j}{A} \right)^{1+\theta} \mu_j^{1+\theta} \right]^{\frac{1}{1+\theta}}, \qquad \mu_j = \left[\sum_{i=1}^{M_i} \left(\frac{a_{ij}}{a_j} \right)^{1+\eta} \mu_{ij}^{1+\eta} \right]^{\frac{1}{1+\eta}}$$

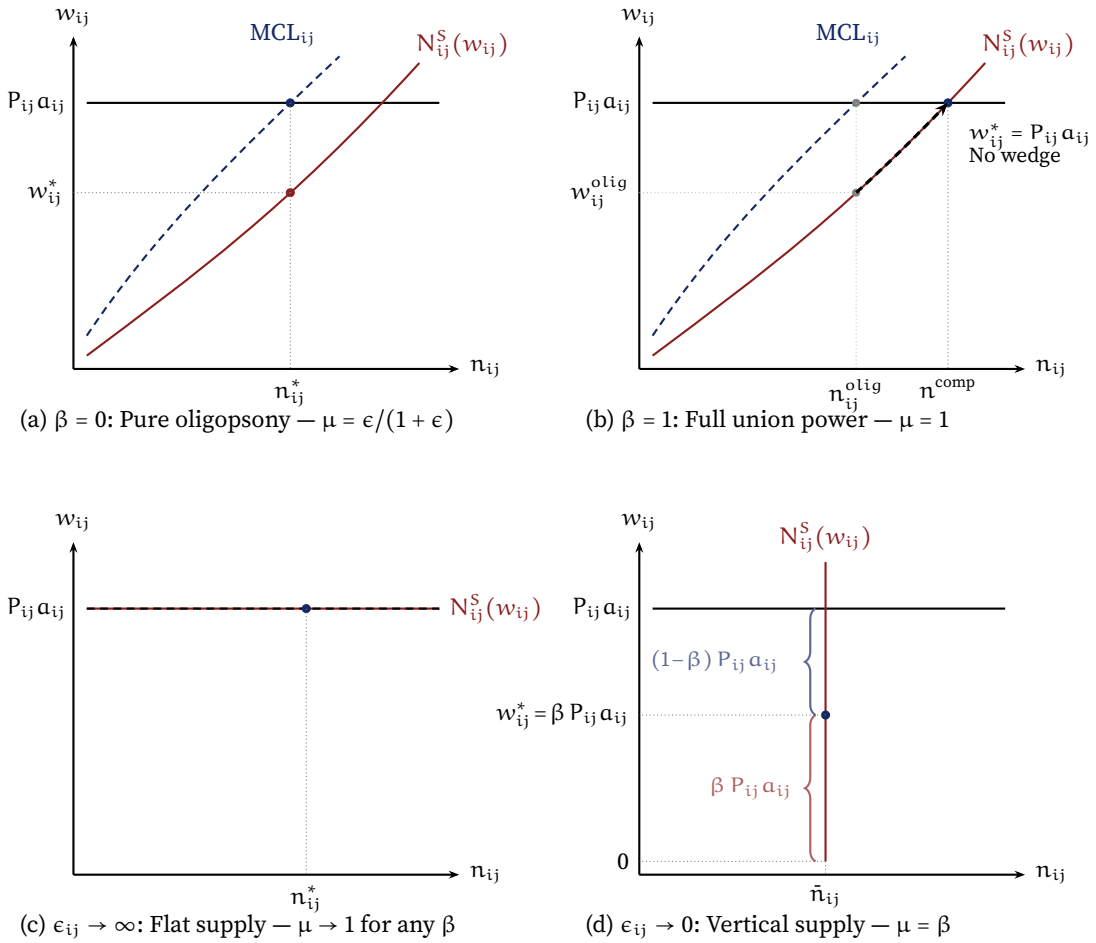
Misallocation:

$$\Omega = \sum_{j=1}^{M_j} \left(\frac{a_j}{A} \right)^{1+\theta} \left(\frac{\mu_j}{\mu} \right)^{\theta} \omega_j, \qquad \omega_j = \sum_{i=1}^{M_i} \left(\frac{a_{ij}}{a_j} \right)^{1+\eta} \left(\frac{\mu_{ij}}{\mu_j} \right)^{\eta}$$

F. Calibration

We calibrate the model separately for a baseline period averaging 1990–1995 ($\tau = 1$) and for a final period averaging 2013–2018 ($\tau = 2$). The procedure has three blocks. Block A fixes the

FIGURE E.2—The four limiting cases of $\mu_{ij} = (\epsilon_{ij} + \beta)/(1 + \epsilon_{ij})$.



(a) $\beta = 0$: Pure oligopsony — $\mu = \epsilon/(1+\epsilon)$

(b) $\beta = 1$: Full union power — $\mu = 1$

(c) $\epsilon_{ij} \to \infty$: Flat supply — $\mu \to 1$ for any β

(d) $\epsilon_{ij} \to 0$: Vertical supply — $\mu = \beta$

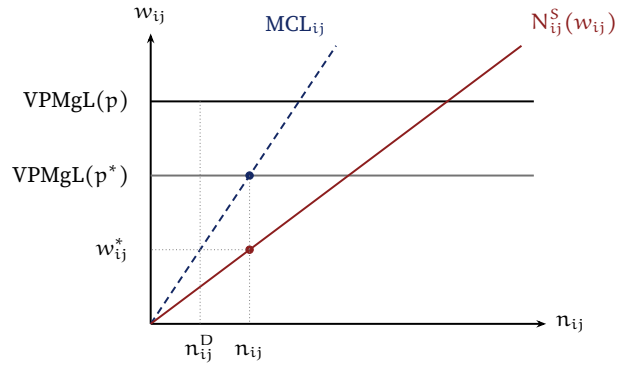
Notes: The four limiting cases of $\mu_{ij} = (\epsilon_{ij} + \beta)/(1 + \epsilon_{ij})$. *Panel (a):* pure monopsony ($\beta = 0$); the firm hires where MCL = $P_{ij}a_{ij}$ and pays the wage on the supply curve. *Panel (b):* full union power ($\beta = 1$); the wage equals $P_{ij}a_{ij}$ and the wedge vanishes. *Panel (c):* perfectly elastic supply; $W^* = P_{ij}a_{ij}$ regardless of β. *Panel (d):* perfectly inelastic supply; employment is fixed at $\bar{N}$ and the entire $P_{ij}a_{ij}$ is split by bargaining power.

elasticities and normalizations. Block B constructs the firm population non-parametrically from the firm-level estimates and augments it with a competitive fringe. Block C calibrates the period-specific $(\delta_\tau, \beta_\tau)$ to two targeted moments and recovers the scale parameters $(z, \bar{\varphi})$.

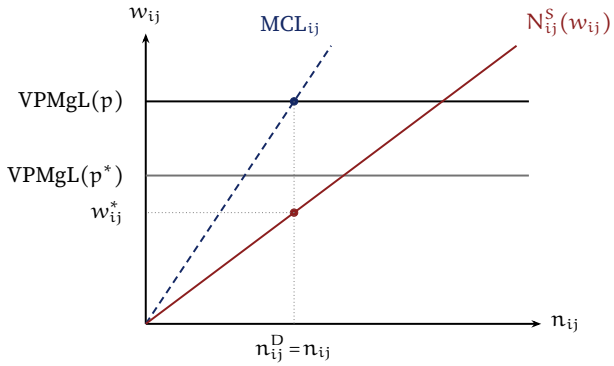
1. Block A: within-sample elasticities and normalizations

The within-market elasticity η is a within-sample estimate, identified from the *low*-markdown tail through the bounding relation (18) rather than imported externally. It governs substitution among atomistic, low-payroll-share firms, whose firm-level supply elasticity ϵ_{ij} tends to η as $s_{ij} \to 0$; these least-monopsonized firms set the smallest markdowns, so the 5th percentile of the employment-weighted markdown distribution pins $\eta = 1/(P_5 - 1) = 9.34$. Because no firm has a share of exactly zero, this is a slight lower bound on η; it lands just below the external estimate of Berger et al. (2022b) ($\eta \approx 10.85$, for which Chile has no quasi-experimental analogue), which supports the in-sample value and motivates the η-robustness check reported with the

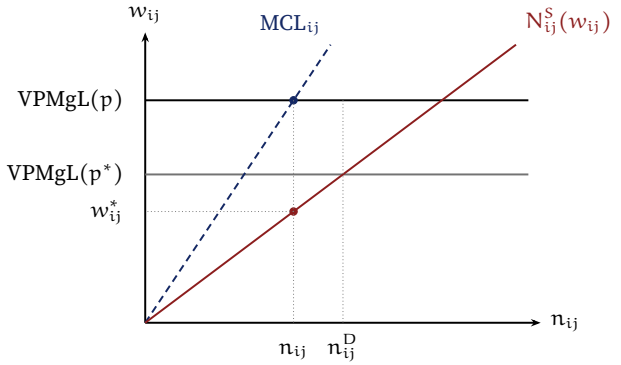
FIGURE E.3—Limiting Cases of Export Status



(a) Exporter: $n_{ij} > n^D_{ij}$



(b) Constrained Non-Exporter: $n_{ij} = n^D_{ij}$



(c) Non-Exporter: $n_{ij} < n^D_{ij}$

Notes: Types of firm-level equilibrium. The firm hires until the marginal labor cost (mc, navy) equals the value of the marginal product of the *marginal* worker, and the equilibrium wage w^*_{ij} is read off the labor supply curve n^S (brick dot). Output up to the domestic labor requirement n^D_{ij} is sold at home and valued at VPMgL(P); output beyond n^D_{ij} is exported and valued at VPMgL(P^*). The export decision compares n_{ij} with n^D_{ij}. *Panel (a):* an exporter ($n_{ij} > n^D_{ij}$); the marginal worker produces for export, so n_{ij} solves mc = VPMgL(P^*) and the surplus $n_{ij} - n^D_{ij}$ is shipped abroad. *Panel (b):* a constrained non-exporter ($n_{ij} = n^D_{ij}$); domestic demand exactly absorbs the firm's domestic optimum, so n_{ij} solves mc = VPMgL(P) and coincides with n^D_{ij}. *Panel (c):* a non-exporter ($n_{ij} < n^D_{ij}$); the firm reaches its domestic optimum mc = VPMgL(P) before domestic demand is exhausted, so it serves only the home market.

welfare results.

The across-market elasticity $\theta = 0.24$ is likewise a *within-sample estimate*. High-payroll-share firms have $\epsilon_{ij} \to \theta$ as $s_{ij} \to 1$ and therefore populate the low-take-home-share (high-markdown) tail, where $\mu_{ij} \to \theta/(1+\theta)$ at $\beta = 0$. Inverting this boundary condition at the 95th percentile of the employment-weighted markdown distribution (equivalently, the 5th percentile of the take-home-share distribution) gives the closed form $\theta = 1/(P_{95} - 1) = 0.24$, computed directly from the estimated markdowns rather than by an inner calibration loop.[38] The Frisch elasticity is $\varphi = 0.5$ (Chetty et al., 2011; Berger et al., 2022a), and the world price is the numeraire, $p^* = 1$,

[38] In the notation of the code, θ is stored as $x/(1-x)$ with $x = 0.1935$, the fifth percentile of the take-home-share distribution (equivalently, the P_{95} of the markdown distribution), giving $\theta = 0.1935/(1-0.1935) \approx 0.24$. Because θ is pinned by the low *tail* of the distribution while the bargaining weight β is pinned by the *aggregate* take-home share (Block C), the two are identified from distinct features of the data.

in both periods.

2. Block B: Firm population from empirical triplets and a fringe

We do not impose a parametric productivity distribution. Each firm is characterized by a triplet (a_{ij}, ϕ_{ij}, M_j): a productivity level $a_{ij} = \exp(\hat{\omega}_{ij})$ from the Ackerberg et al. (2015) estimates of Section IV, a Cobb–Douglas variety weight ϕ_{ij}, and a market size M_j. Triplets are drawn *jointly* (with replacement) from the empirical distribution, preserving the observed cross-sectional dependence among the three objects, and the same draw is fixed across calibration steps (common random numbers). Market sizes are tightened from the region to the *comuna* level using CASEN commuting-zone shares $p_{c|r}$, so that M_j reflects local rather than regional concentration.

Because the firm-level survey covers only plants with at least ten employees, we augment the surveyed population with a competitive *fringe* of smaller firms (source indicator = 1). The fringe is disciplined by administrative records: fringe plants number 1.60 per surveyed plant (a fringe count share of 0.616) but account for only 3.02% of total sales, and we assume no fringe plant exports. Fringe productivities are drawn so that their mean lies below the surveyed mean by the sales-share gap implied by the administrative data, with dispersion equal to the surveyed cross-firm standard deviation. To prevent degenerate markets we require at least one surveyed plant per market; the rare all-fringe draws are repaired by swapping one slot, which leaves the aggregate fringe count and sales shares essentially unchanged. The export threshold is set off the surveyed (non-fringe) productivity median, so that the large domestic-only fringe does not mechanically lower the export cutoff.

We hold the firm population fixed across periods and attribute the change in outcomes between 1990–1995 and 2018 to the change in trade costs and bargaining power, not to compositional change in the population. The model has no entry or exit margin; we therefore do not target net firm dynamics, and we discuss this abstraction below.

FRINGE CONSTRUCTION: FULL CALCULATIONS. The competitive fringe is disciplined by two administrative ratios from the 2005 manufacturing census (SII): the fringe-to-surveyed count ratio $r_f = N_{\text{fringe}}/N_{\text{ENIA}} = 1.6029$ (a fringe count share of $r_f/(1+r_f) = 0.616$) and the fringe share of total domestic sales $\Phi_f = 0.0302$. The market-size multiplier is $1 + r_f = 2.6029$. Fringe firms are constructed in four steps.

Step 1: productivity gap in ω units. We transport the *sales* gap between surveyed and subthreshold plants into the log-productivity (ω) scale using the within-ENIA relationship between ω and log sales. Within the 1990–1995 ENIA sample we regress log gross output on inputs with year×industry fixed effects, take the residual as $\hat{\omega}_{it}$, and regress $\hat{\omega}_{it}$ on log sales y_{it} to obtain the transport slope $b = \widehat{\partial\omega/\partial y}$. The expected log sales of surveyed plants, $\mathbb{E}[y \mid \text{ENIA}]$, is computed as a count-weighted average over the SII size bins, and $\mathbb{E}[y \mid \text{fringe}]$ is the corresponding subthreshold value. The absolute productivity gap is

$$\Delta\omega \;=\; b\left(\mathbb{E}[y \mid \text{ENIA}] - \mathbb{E}[y \mid \text{fringe}]\right) \;\approx\; 0.443. \tag{F1}$$

Step 2: surveyed triplets. Each surveyed plant (source = 0) contributes $\hat{\omega}_{ij}$ (its time-mean residual productivity), a variety weight $\phi_{ij} = \text{domestic sales}_{ij} / \sum_k \text{domestic sales}_k$, and a market

j defined by 2-digit industry × region, with firm count M_j.

Step 3: fringe draws. For each market j we generate $N_j^{\text{fringe}} = \text{round}(r_f\, N_j^{\text{ENIA}})$ fringe firms, each drawn independently as

$$\omega_{ij}^{\text{fringe}} \sim \mathcal{N}\big(\bar{\omega}_{\text{ENIA}} - \Delta\omega,\ \sigma^2_{\text{ENIA}}\big), \tag{F2}$$

where $\bar{\omega}_{\text{ENIA}}$ and σ_{ENIA} are the mean and cross-firm standard deviation of $\hat{\omega}$ among surveyed plants. The fringe distribution is thus the surveyed distribution shifted down by the sales-implied gap $\Delta\omega$ and given the same dispersion.

Step 4: market size and variety weights. Market sizes are scaled to $M_j \leftarrow \text{round}\big((1+r_f)\,M_j\big)$, and variety weights are renormalized so that surveyed plants retain $1-\Phi_f$ of total sales and the fringe receives Φ_f, split equally:

$$\phi_{ij} \leftarrow \phi_{ij}\,\frac{1-\Phi_f}{\sum_{k:\,\text{ENIA}}\phi_k}\ \text{(surveyed)}, \qquad \phi_{ij} = \frac{\Phi_f}{N_{\text{fringe}}}\ \text{(fringe)}, \tag{F3}$$

so that $\sum_{ij}\phi_{ij} = 1$. The per-firm fringe/surveyed expenditure ratio implied by (F3) is approximately 0.019, matching the SII per-firm benchmark. Because fringe firms are domestic by construction, the export cutoff is set off the *surveyed* (non-fringe) productivity median, so the large domestic-only fringe does not mechanically lower it; markets that would otherwise contain no surveyed plant have one slot swapped for a randomly drawn surveyed firm, leaving the aggregate fringe count and sales shares unchanged.

FRINGE MARKDOWNS AND THE ECONOMY-WIDE AGGREGATE. Fringe firms hold near-zero payroll shares ($s_{ij} \to 0$), so their firm-level supply elasticity tends to η and their model-implied worker take-home share tends to the atomistic limit

$$\mu^{\text{fringe}} \longrightarrow \frac{\eta+\beta}{\eta+1}. \tag{F4}$$

The economy-wide worker share splices the targeted ENIA object with this imputed fringe object, employment-weighted. As reported in Table E.1, the 10-employee truncation overstates the ENIA-only markdown relative to the economy-wide object by 8.5% in 1990 (the aggregate worker share rises from 0.385 to 0.417) and by 4.4% in 2018 (0.511 to 0.534): the selection bias is material early in the sample, when small low-markdown firms are a larger share of activity, and modest by the end.

3. Block C: Moment matching and scale recovery

TARGETED MOMENTS. For each period we calibrate $(\delta_\tau, \beta_\tau)$ to two moments,

$$\mathbf{m}_\tau = \big(\overline{IP}_\tau,\ \bar{\mu}_\tau\big)',$$

where $\overline{IP}_\tau$ is the employment-weighted mean of market-level import penetration $IP_j = IMP_j/(IMP_j + DOM_j - EXP_j)$ and $\bar{\mu}_\tau \equiv \sum_{ij}(n_{ij}/N)\,\mu_{ij}$ is the employment-weighted worker take-home share. The empirical targets are $(\overline{IP}_1, \bar{\mu}_1) = (0.200,\ 0.379)$ and $(\overline{IP}_2, \bar{\mu}_2) = (0.270,\ 0.506)$. The exporter

share ES_τ is reported in Table 4 as an untargeted model prediction, and it is where the static model shows its limits: the model over-predicts the level (20.11% against 7.42% in the baseline period, 17.06% against 8.23% in the final period) and misses the direction of change, as the data exporter share rises with liberalization while the model's falls. We read this as a consequence of the fixed firm population: the model abstracts from the entry, exit, and exporter-switching margins that the reduced-form evidence of Section V shows operate over multi-year horizons, so the exporter share marks the boundary of the static framework rather than serving as corroboration.

CALIBRATION OBJECTIVE. For each period we evaluate the model on a 40×40 grid in (δ, β), with $\delta \in [0.01, 0.99]$ and $\beta \in [0, 0.99]$ in the baseline and $\delta \in (\hat{\delta}_1, 0.99]$ in the final period (imposing $\delta_2 > \delta_1$). At each node we solve the equilibrium and compute the squared relative deviation of each targeted moment,

$$\hat{\psi}_\tau = \arg\min_{(\delta,\beta)} \left(\frac{\bar{\mu} - \bar{\mu}_\tau}{\bar{\mu}_\tau} \right)^2 + \left(\frac{\overline{IP} - \overline{IP}_\tau}{\overline{IP}_\tau} \right)^2, \tag{F5}$$

refining the argmin by linear interpolation onto a 3000×3000 grid and re-solving the equilibrium at the interpolated point for verification. Identification is as described in the main text: $\bar{\mu}_\tau$ disciplines β_τ and $\overline{IP}_\tau$ disciplines δ_τ.

SCALE RECOVERY. The level parameters $(z, \bar{\varphi})$ are recovered after (δ_1, β_1). We solve the baseline equilibrium in hat-units ($z = \bar{\varphi} = 1$) and match two level moments, average firm size $\overline{S} = 31.08$ workers and average earnings per worker $\overline{E} = 212{,}739$ (1990–1995 ENIA × SII), using the homogeneity of the equilibrium map. Earnings scale linearly in z and firm size scales in $\bar{\varphi}$ (up to the Frisch correction through the wage index), giving the closed form

$$\hat{z} = \frac{\overline{E}}{\overline{E}^{\text{hat}}}, \qquad \hat{\bar{\varphi}} = \frac{\overline{S}/\overline{S}^{\text{hat}}}{\left(\overline{E}/\overline{E}^{\text{hat}}\right)^{\varphi}}, \tag{F6}$$

where hat-superscripts denote the hat-unit equilibrium values. Because the targeted moments in (F5) are scale-free, $(z, \bar{\varphi})$ enter levels only and do not feed back into (δ, β).

4. Algorithm

1. Fix Block A parameters $(\eta, \theta, \varphi, p^*)$ and draw the firm population (Block B); hold the draw fixed throughout.
2. *Baseline.* Minimize (F5) over the (δ, β) grid with all firms active to obtain $(\hat{\delta}_1, \hat{\beta}_1)$.
3. *Scale.* Solve the baseline in hat-units and recover $(\hat{z}, \hat{\bar{\varphi}})$ from (F6); rescale productivities to levels.
4. *Final.* Minimize (F5) over the constrained grid $(\delta_2 > \hat{\delta}_1)$ to obtain $(\hat{\delta}_2, \hat{\beta}_2)$, holding $(\hat{z}, \hat{\bar{\varphi}})$ fixed.

The calibrated values and the model-vs-target fit are reported in Table 4.

ON FIRM DYNAMICS. The data exhibit both gross entry and gross exit between 1990–1995 and 2018, whereas the model holds the firm population fixed and has no activity margin. We therefore do not target the change in the active firm count and we do not interpret any parameter as a

fixed operating cost.

TABLE F.1—Welfare Decomposition across Frisch Elasticities

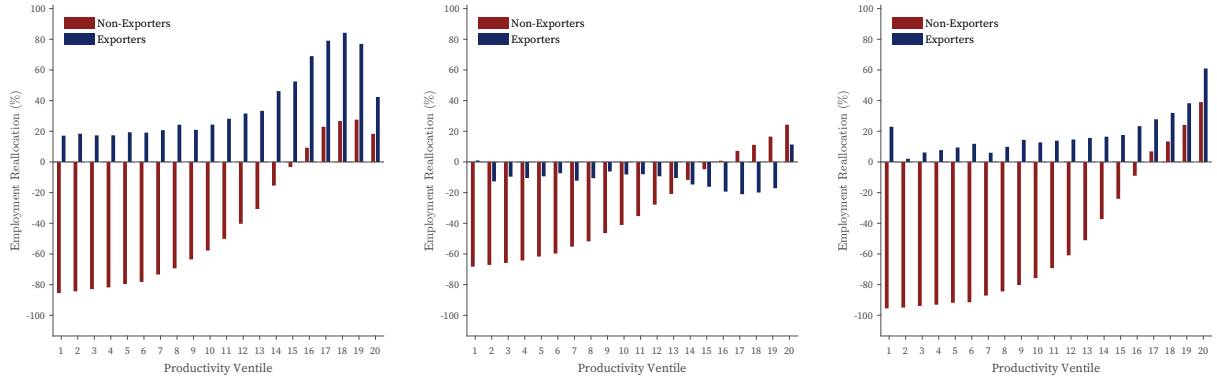
	Frisch elasticity φ					
Statistic	0.10	0.25	0.50	0.75	0.90	1.00
Panel A: Realized (Both channels, reduced TC)						
λ (welfare)	1.382	1.457	1.596	1.754	1.860	1.936
λ^{ω}/λ (%)	73.5	75.5	78.5	81.0	82.4	83.2
λ^{μ}/λ (%)	26.5	24.5	21.5	19.0	17.6	16.8
Panel B: Counterfactual (Synthetic + Bargaining)						
λ (welfare)	1.301	1.363	1.475	1.602	1.686	1.745
λ^{ω}/λ (%)	73.2	75.1	77.8	80.2	81.4	82.2
λ^{μ}/λ (%)	26.8	24.9	22.2	19.8	18.6	17.8
Panel C: Difference (Realized − Counterfactual)						
$\Delta\lambda$ (total)	+0.080	+0.094	+0.120	+0.152	+0.175	+0.191
$\Delta(\lambda^{\omega}/\lambda)$ (pp)	+0.3	+0.4	+0.6	+0.8	+0.9	+1.0

Notes: Each column re-recovers $(z, \bar{\varphi})$ for the stated Frisch elasticity φ, holding the calibrated (δ, β) fixed. λ is consumption-equivalent welfare relative to baseline. λ^{μ} is the welfare change in a no-dispersion economy; $\lambda^{\omega} = \lambda - \lambda^{\mu}$ is the residual reallocation component. Panel C reports Realized minus Counterfactual.

G. Model Estimates

This subsection illustrates key features of the model equilibrium and the effects of trade cost reductions. Figures G.1 and G.2 show employment reallocation, relative to the baseline equilibrium, under each of the three counterfactual channels (Trade Only, Bargaining Only, and Final), across bins of firm productivity:

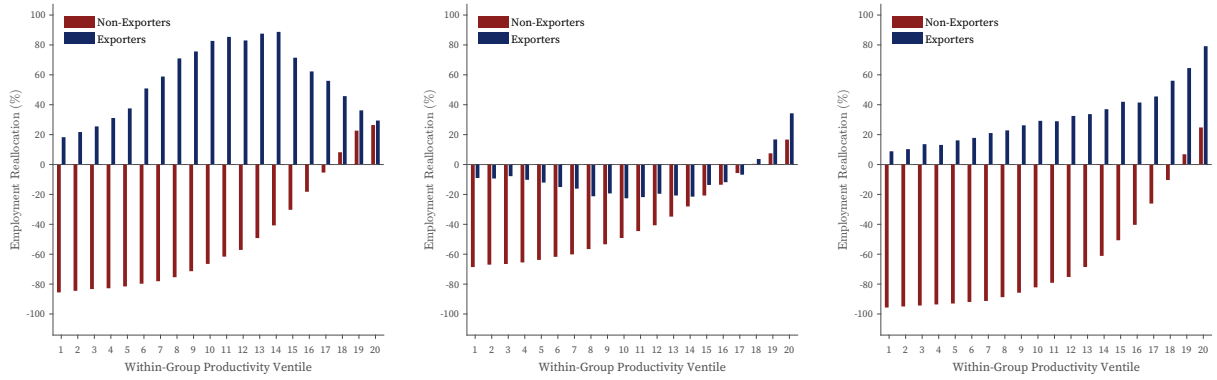
FIGURE G.1—Model Labor Reallocation by Channel — Pooled Productivity Bins, by Export Status



(a) Trade Only (b) Bargaining Only (c) Final (Trade + Bargaining)

Notes: Employment reallocation relative to the baseline equilibrium, by counterfactual channel (Trade Only, Bargaining Only, Final). Firms are grouped into productivity ventiles using the pooled (economy-wide) productivity distribution, then split by export status. Bars show the mean change in employment, by ventile and status, relative to baseline.

FIGURE G.2—Model Labor Reallocation by Channel — Within-Group Productivity Bins, by Export Status



(a) Trade Only (b) Bargaining Only (c) Final (Trade + Bargaining)

Notes: Same exercise as Figure G.1, except productivity ventiles are computed separately within each export-status group (within-group rank), isolating reallocation that occurs net of the productivity gap between exporters and non-exporters.