Wage markups, wage markdowns, and buyer power in intermediate input markets*

Leonard Treuren†

July 2026

Abstract

How are imperfections in different input markets related? I estimate wedges between marginal revenue products and prices in the labor and intermediate input markets, using a revenue-based approach that does not require either market to be competitive. I apply this approach to Dutch manufacturing and show that the wedges are negatively correlated. Firms mark down intermediate input prices relative to their marginal revenue product, while both wage markups and wage markdowns are common in labor markets. A model where firms face upward-sloping supply curves in both input markets, and collective bargaining determines wages, can rationalize these findings. Consistent with rent sharing linking market imperfections across input markets, I show that as the euro appreciates, importers of Chinese intermediates increase their intermediates markdowns and earn higher rents, and their workers' wages rise, even though labor becomes no more revenue productive on average or on the margin.

Keywords: Monopsony; Rent sharing; Buyer power; Revenue function estimation
JEL Codes: D43; J31; J42; J50; L10

*I thank the editor, John Asker, and three anonymous referees, Dan Ackerberg, Tommaso Alba, Marcus Asplund, Eric Bartelsman, Filippo Biondi, Giulia Brancaccio, Chris Conlon, Jan De Loecker, Katharina Erhardt, Harald Fadinger, Paul Grieco, Etienne Guigue, Nadine Hahn, Michael Lipsitz, Matthias Mertens, Monica Morlacco, Amil Petrin, Joris Pinkse, Michael Rubens, Jo Seldeslachts, Bradley Setzler, Joel Stiebale, Jo Van Biesebroeck, Frank Verboven, Stephen Yeaple, and seminar participants at CBS, CPB, DICE, DIW Berlin, KU Leuven, Norwich, NYU Stern, Oxford, Penn State, and the 2022 IIOC (Boston), 2022 EEA (Milan), 2022 JEI (Las Palmas), 2023 MaCCI (Mannheim), 2023 CRESSE (Rhodes), 2023 EARIE (Rome), 2024 Berlin IO Day (Berlin), 2024 RES (Belfast), 2025 M-Power (Leuven), and 2025 IIOC (Philadelphia) conferences for useful comments and discussions. Results based on calculations using non-public microdata from Statistics Netherlands (CBS). I gratefully acknowledge funding of ERC Consolidator grant 816638 and the Open Data Infrastructure for Social Science and Economic Innovations (ODISSEI) in the Netherlands. Remaining errors are my own.

†KU Leuven; leonard.treuren@kuleuven.be.

1 Introduction

Input market imperfections that drive a wedge between an input's price and its marginal revenue product can lead to resource misallocation and reduced welfare. Studying input price formation in the presence of market imperfections is, therefore, of considerable interest to researchers and policymakers alike. Recent examples include evidence of widespread monopsony power in US labor markets (e.g., Yeh et al. (2022); Lamadon et al. (2022)), studies documenting rent sharing in labor markets characterized by collective bargaining (e.g., Dobbelaere and Mairesse (2013); Card et al. (2014)), and an emerging strand of literature finding evidence of buyer power for intermediate inputs (e.g., Morlacco (2020); Rubens (2023)).[1] One might expect a firm's behavior and hence such imperfections to be linked across input markets. Whether this is the case empirically remains an open question, as existing work focuses on one input market at a time.

This paper investigates how market imperfections are related across input markets. I do this in the context of labor and intermediate input markets in Dutch manufacturing from 2007 to 2018. I estimate both labor and intermediate input wedges without normalizing either wedge to one, and document a strong negative relationship between them. Intermediate input prices are routinely marked down with respect to the marginal revenue product, consistent with buyer power. Both wage markups and markdowns with respect to labor's marginal revenue product are frequently observed. A simple model that combines rent sharing through wage bargaining, monopsony power, and buyer power can rationalize these findings. I use exchange rate variation among importers of Chinese intermediates to show that larger intermediates markdowns raise wages and lower labor wedges, while leaving employment and employees' revenue productivity largely unaffected, consistent with rent sharing creating a connection between imperfections in different input markets. My findings imply that, once perfect competition in input markets is relaxed, understanding price formation in any one input market may require accounting for imperfections in other input markets.

I identify input wedges by building on the revenue approach, which divides an input's revenue elasticity by its expenditure share in revenue (Hashemi et al., 2022). I extend this result to input bundles with unobserved heterogeneity, in which case the wedge is an expenditure-weighted average of the wedges of the underlying varieties. I then clarify how the revenue approach differs from the widely used production approach, which utilizes production function estimation and builds on Hall (1988) and De Loecker and Warzynski (2012). The production approach models a cost-minimization problem and identifies an input wedge relative to another input wedge. To recover an individual wedge, the production approach must normalize another input's wedge to unity – typically intermediate inputs are chosen. In contrast, the revenue approach builds on the revenue function's properties and does not

[1] In this paper, 'monopsony power' refers to firms' purchasing power in the labor market, and 'buyer power' to purchasing power in the intermediate input market.

require a unit wedge or the specification of an input choice problem.

To estimate a revenue function, restrictions on demand and conduct in output markets are required. I show that the restrictions imposed, often implicitly, when the production approach is applied to revenue data are sufficient to specify an estimable revenue function. This function can be estimated using established control-function or dynamic-panel methods originating in Olley and Pakes (1996) and Blundell and Bond (2000), respectively. More generally, such approaches can be applied with any conditional inverse demand function that is log-linear in own quantity with a slope common at the level of estimation and firm-specific demand heterogeneity entering additively.

Employing a variety of revenue function estimators, I find that both wage markups and wage markdowns are prevalent in Dutch manufacturing, each account for approximately half of the baseline observations. This is not in line with a pure monopsony model, which would result in wage markdowns, as observed in US manufacturing by Yeh et al. (2022). In contrast, intermediate input wedges suggest that buyer power for intermediates is common in Dutch manufacturing. The median intermediates price markdown is about 12 percent. The two wedges are negatively correlated: underpaying suppliers on the margin relative to their revenue product coincides with overpaying workers on the margin. Both input wedge distributions show substantial dispersion. A strong predictor of the labor wedge is the degree to which prices are marked down relative to the marginal revenue product in the intermediate input market. The more a firm underpays its input suppliers on the margin, the higher its wages and wage markups, although employees' revenue productivity is largely unchanged.

These patterns are suggestive of rent sharing creating a connection between market imperfections in different input markets, but are difficult to fully separate from factor-biased technical change. Raval (2023) documents a closely related pattern – output market markups based on labor and on intermediates are negatively correlated across a wide range of manufacturing datasets – and shows that it can be generated by unobserved labor-augmenting productivity even when input markets are competitive. Because both explanations operate through the same accounting objects, the negative correlation between the two wedges does not by itself distinguish rents from technology.

Variation in the intermediate input market that is not itself a change in technology allows me to make progress. I construct such variation by interacting firms' predetermined reliance on intermediates imported from China with fluctuations in the euro-renminbi exchange rate, which are rotations of the intermediates supply curve. To motivate the resulting regressions, I extend a standard bargaining framework of the rent-sharing literature with upward-sloping supply curves in both input markets: if workers have bargaining power, wages load on the firm's rents, so that buyer power for intermediates puts upward pressure on wages and can counteract monopsonistic wage markdowns. Consistent with this mechanism, I find that as the euro appreciates, exposed importers increase their intermediates markdowns and lower their labor wedges. Instrumenting the intermediates wedge with exchange rate exposure, a

one percent increase in the intermediate input wedge reduces the labor wedge by 0.68 percent. The response operates entirely through wages while employment, workforce composition, and workers' marginal revenue product are unchanged. Value added per worker rises even though revenue per worker does not, reflecting cheaper inputs rather than more productive workers. Overall, these results suggest that rent sharing creates a link between market imperfections in the labor and intermediate input markets.

This paper contributes to several strands of literature. First, the literature on monopsony power.[2] Yeh et al. (2022) estimate labor wedges for US manufacturing plants from 1976 to 2014 and, in line with monopsony power, find that most plants mark wages down with respect to the marginal revenue product of labor.[3] In Germany, however, both wage markups and markdowns occur, with markdowns concentrated among large, highly productive firms (Mertens, 2023) and substantially less common where workers are covered by collective agreements (Dobbelaere et al., 2024). As China's export demand grows, German firms are found to exercise more labor market power (Mertens, 2020).[4] I contribute by showing that imperfections in the intermediates market can influence wages and wage markdowns.

Second, the paper contributes to the rent-sharing literature (see Card et al. (2018) for a survey). Several papers jointly estimate markups in output markets and imperfections in labor markets, showing that wage markups are common in French manufacturing (Dobbelaere and Mairesse, 2013; Caselli et al., 2021). Card and Cardoso (2022) report a 20 percent wage markup over sectoral wage floors in Portugal, a country also characterized by high collective bargaining coverage. Allowing for imperfect competition in labor and product markets, Kroft et al. (2025) find that product market power mitigates the incentive to mark down wages in the US construction industry. Existing work focuses on product market rents that are shared with employees, while I contribute by providing evidence on redistribution of rents across different input markets.

Third, the paper contributes to the emerging literature on buyer power in intermediate input markets. Rubens (2023) finds that ownership consolidation in Chinese cigarette manufacturing has increased intermediate input price markdowns by 30 percent. Morlacco (2020) and Alviarez et al. (2023) estimate that buyer power for imported intermediates is common in, respectively, France and the United States. Avignon and Guigue (2022) show that French dairy manufacturers have buyer power in the market for raw milk. These papers allow for imperfect competition in output markets but do not study the interaction with labor market imperfections. I contribute by studying the relationship between buyer power

[2]I focus on papers estimating labor wedges using the production approach, surveyed in Syverson (2025). Other common approaches include estimating labor supply elasticities and relating labor market concentration to wage measures. Card (2022) survey the broader literature on firms' ability to set wages.

[3]The connection between labor and product market power in US manufacturing is studied by Kirov and Traina (2023) and Deb et al. (2024).

[4]In addition, wage markdowns are found to substantially decrease labor shares in China, India, and Germany (Brooks et al., 2021; Mertens, 2022).

and labor market imperfections and by providing additional evidence that buyer power is common.[5]

If both sides of a market have market power, an increase in buyer power could move prices closer to competitive levels and improve welfare. Such countervailing power has been studied since at least Galbraith (1952). The theoretical conditions under which buyer power is indeed countervailing are laid out by Avignon et al. (2024) and Demirer and Rubens (2025). Angerhofer et al. (2026) show that collective bargaining agreements mitigate wage markdowns for US public school teachers, and a broader literature studies countervailing power in relation to mergers (e.g., Barrette et al. (2022)). My findings point to a cross-market countervailing effect. As buyer power for intermediate inputs increases, workers get paid more, which could negate monopsonistic markdowns. The results suggest that interventions affecting purchasing power in one input market, such as trade restrictions, may have distributional spillovers to other input markets.

Finally, my work is related to a large body of work that regresses log revenue on logged measures of input use obtained from financial statements (De Loecker and Syverson, 2022). The results of such an exercise are usually interpreted as output elasticities. Authors typically observe revenue and possibly capital, labor, and intermediate inputs, but not output. A long line of work considers the conditions under which output elasticities or revenue elasticities are identified in this setting (e.g., Klette and Griliches (1996); Bond et al. (2021)). These issues motivate the observation of Hashemi et al. (2022) that previous work attempting to estimate product market markups might have inadvertently identified input wedges instead by using revenue elasticities in the place of output elasticities, and the applications of revenue function estimation to input wedges by Petrin and Sivadasan (2013) and Pham (2023).

Section 2 next outlines the empirical approach and introduces the setting and data. Section 3 provides results on the joint distribution of labor and intermediate input wedges in Dutch manufacturing, while Section 4 uses quasi-random variation to argue that rent sharing is one mechanism underpinning my results. Concluding remarks follow in Section 5.

2 Empirical approach

An input wedge is the ratio of an input's marginal revenue product to its price. This section expands on the 'revenue approach' to input wedge identification (Hashemi et al., 2022; Pham, 2023). Consider a setting where output is unobserved, but revenue and input bundles such as labor and intermediate inputs are observed.[6] In this setting, there is unobserved input

[5] A larger body of work studies input price dispersion while abstracting from buyer power (e.g., Atalay (2014)). Grieco et al. (2016) discuss production function estimation in the presence of unobserved input price dispersion in competitive input markets.

[6] Such data are routinely used in related work studying, for example, China and India (e.g., Brooks et al. (2021); Pham (2023)), France (e.g., Dobbelaere and Mairesse (2013); Caselli et al. (2021)), Chile (e.g.,

heterogeneity, it is not obvious that any observed input has a unit wedge, and estimating output elasticities is challenging. I show three things. First, input wedges can be interpreted as weighted averages of underlying input wedges in the presence of unobserved input heterogeneity. Second, the 'production approach' to input wedge estimation identifies one wedge relative to another wedge, while the revenue approach identifies individual wedges. Third, the assumptions on demand and conduct in output markets that are (often implicitly) made when the production approach is applied without output data are sufficient but not necessary to write down an estimable revenue function.

2.1 Identifying input wedges

To match the application, I focus on labor (L_{it}) and intermediate inputs (M_{it}), but the revenue approach can be applied to other inputs. Let Ω_{it}^R be a revenue productivity shifter and K_{it} denote capital. I make two assumptions.

Assumption 1 The revenue function $R_{it}(K_{it}, L_{it}, M_{it}, \Omega_{it}^R)$ is continuously differentiable in inputs so that the marginal revenue products are given by $MRPX_{it} = \frac{\partial R_{it}}{\partial X_{it}}$ for $X \in \{L, M\}$.

Assumption 2 Input bundles L_{it} and M_{it} are constant-returns-to-scale (CRS) aggregators of underlying input types. Formally, $L_{it} = A^L(\mathbf{l}_{it})$, where $\mathbf{l}_{it}$ is a vector whose entries are the amounts of labor of each variety in firm i's labor variety set $\mathcal{L}_{it}$ and $A^L(\cdot)$ is a CRS aggregator that is continuously differentiable in its elements. $M_{it} = A^M(\mathbf{m}_{it})$ and $\mathcal{M}_{it}$ are defined likewise. $\mathcal{L}_{it}$ and $\mathcal{M}_{it}$ are taken as given when firm i selects $\mathbf{l}_{it}$ and $\mathbf{m}_{it}$.

Assumption 1 rules out revenue functions in which inputs are perfect complements (i.e., Leontief). Assumption 2 addresses that I do not observe underlying input heterogeneity, which is the norm with financial statements. Assumption 2 is often explicitly made by papers that aggregate product-level import data to a bundle that is used as an input of a production function and is implicitly made by other papers that estimate input wedges in the presence of unobserved input heterogeneity. Let $\theta_{it}^L = \frac{\partial R_{it}}{\partial L_{it}} \frac{L_{it}}{R_{it}}$ and $\theta_{it}^M = \frac{\partial R_{it}}{\partial M_{it}} \frac{M_{it}}{R_{it}}$ denote the revenue elasticities of observed labor and intermediate inputs, respectively, LS_{it} the labor share of revenue, $\frac{W_{it}L_{it}}{R_{it}}$, and MS_{it} the intermediate input share of revenue, $\frac{P_{it}^M M_{it}}{R_{it}}$. Assumptions 1 and 2 imply the following lemma, which is proved in Appendix A.1.

Lemma 1 *The ratio of an input bundle's revenue elasticity to its revenue share equals the weighted average of the input wedges of the varieties contained in that bundle, with weights*

Levinsohn and Petrin (2003); Petrin and Sivadasan (2013)), and the US (e.g., Yeh et al. (2022)).

given by the share of a variety's expenditure in overall expenditure on that input bundle

$$\gamma_{it}^{L} \equiv \frac{\theta_{it}^{L}}{LS_{it}} = \sum_{X \in \mathcal{L}_{it}} \frac{P_{it}^{X} X_{it}}{W_{it} L_{it}} \frac{MRPX_{it}}{P_{it}^{X}} \quad \text{and} \quad \gamma_{it}^{M} \equiv \frac{\theta_{it}^{M}}{MS_{it}} = \sum_{X \in \mathcal{M}_{it}} \frac{P_{it}^{X} X_{it}}{P_{it}^{M} M_{it}} \frac{MRPX_{it}}{P_{it}^{X}}. \tag{1}$$

Lemma 1 shows that the ratio of a bundle's revenue elasticity to its expenditure share remains informative when input heterogeneity is unobserved: it equals the expenditure-weighted average of its varieties' wedges. Without such heterogeneity, the ratio equals the input wedge itself, as shown by Hashemi et al. (2022).[7] Estimating a bundle-level wedge therefore requires only its expenditure share in revenue, which is typically observed in financial statements, and its revenue elasticity, which must generally be estimated.

Lemma 1 is conditional on revenue elasticities and itself imposes no restrictions on input or output markets beyond Assumptions 1 and 2. The revenue approach can therefore in principle identify input wedges under a wide range of assumptions on conduct, demand, input supply, and technology. For instance, one does not need to specify whether input prices are bargained or posted, or whether there are adjustment frictions for labor. This flexibility is due to the fact that Lemma 1 is derived from properties of the revenue function rather than an optimization problem of a firm. In practice, estimating revenue elasticities requires restrictions, especially on output markets and technology, but these restrictions concern the estimation of the revenue function rather than the mapping from elasticities and shares to wedges in Lemma 1.

Note that Lemma 1 allows unobserved input quality to enter through both the varieties that firms use and the weights of specific varieties within the aggregators. This treatment of quality is common in the trade literature. For instance, the CES import aggregator of Halpern et al. (2015) is consistent with Assumption 2. Input varieties can have different wedges, which in turn influences the estimated bundle-level wedges. What Assumption 2 restricts is where complementarities can occur. Complementarity between the bundles L and M is allowed – it operates through $R(K, L, M, \Omega)$, on which only differentiability is imposed – but complementarity between individual elements of different bundles is not, as the aggregator then depends on inputs outside its bundle and no scalar bundle exists.[8]

Lemma 1 is also unaffected by factor-biased technical change, an issue that has received considerable recent attention (e.g., Doraszelski and Jaumandreu (2018); Raval (2023); Demirer (2025)). This literature considers labor-augmenting technology Ω_{it}^{L} that enters the

[7] Henceforth, I will refer to γ_{it}^{L} and γ_{it}^{M} as the labor and intermediate input wedges, respectively. Results similar to Lemma 1 have been derived for input wedges based on output elasticities (Morlacco, 2020).

[8] The same restriction applies broadly to research using financial statements: writing output as a function of observed bundles such as 'materials' assumes that underlying varieties matter only through those bundles. With across-bundle variety complementarity, output can differ between firms with identical productivity and bundles (but different composition), so no bundle-based function can represent the technology.

revenue function as $R_{it}(K_{it}, \Omega_{it}^L L_{it}, M_{it}, \Omega_{it})$. Because scaling observed labor L_{it} by 1 percent scales effective labor $\Omega_{it}^L L_{it}$ by 1 percent, the revenue elasticity of observed labor equals that of effective labor and no knowledge of Ω_{it}^L is required to recover wedges using Lemma 1. Changes in Ω_{it}^L affect the marginal revenue product of labor, while the wedge captures whether compensation adjusts proportionally. Indeed, in a textbook frictionless labor market, profit maximization equates the wage to the marginal revenue product of labor along any path of Ω_{it}^L. Note that factor-biased technical change can affect revenue elasticities beyond a Cobb-Douglas revenue function, but Lemma 1 is conditional on elasticities.

The production approach to input wedge estimation instead builds on first-order conditions from cost minimization, requires information on output elasticities, and has primarily been used to identify the labor wedge.[9] For any frictionlessly adjustable input V_{it} that is substitutable with other inputs and chosen to minimize the firm's conditional cost function, De Loecker and Warzynski (2012) show that[10]

$$\mu_{it}(V_{it}) \equiv \frac{\tilde{\theta}_{it}^V}{P_{it}^V V_{it}/R_{it}} = \mu_{it} \frac{MRPV_{it}}{P_{it}^V}, \tag{2}$$

where $\tilde{\theta}_{it}^V$ and P_{it}^V are the output elasticity and price of input V_{it}, and μ_{it} is the ratio of the firm's output price to its marginal cost. Using L_{it} and M_{it} in the place of V_{it} in equation (2) and dividing gives

$$\frac{\mu_{it}(L_{it})}{\mu_{it}(M_{it})} = \frac{\tilde{\theta}_{it}^L P_{it}^M M_{it}}{\tilde{\theta}_{it}^M W_{it} L_{it}} = \frac{\theta_{it}^L / LS_{it}}{\theta_{it}^M / MS_{it}} = \frac{\gamma_{it}^L}{\gamma_{it}^M}, \tag{3}$$

where the second equality follows from $\frac{\partial R_{it}}{\partial x_{it}} = \frac{\partial R_{it}}{\partial Q_{it}} \frac{\partial Q_{it}}{\partial x_{it}}$ for $x \in \{L, M\}$. If, in addition to the earlier assumptions, firms are price takers in the intermediates market, profit maximization implies that $MRPM_{it} = P_{it}^M$, so that equation (3) identifies the labor wedge. This has motivated a large recent literature to estimate labor wedges based on a static cost-minimization problem in which firms are price takers in intermediate input markets but might have monopsony power over frictionlessly adjustable labor (e.g., Caselli et al. (2021); Mertens (2022); Yeh et al. (2022)). In general, however, the production approach obtains the labor wedge *relative* to the wedge of another input.

The main advantage of the revenue approach is that it does not require a frictionlessly adjustable input for which firms are price takers to identify an individual input wedge. This is essential in my setting, where I aim to study imperfections in both labor and intermediate input markets and therefore do not want to assume either away. The only remaining input observed in my data is capital, which is unlikely to be frictionlessly adjustable. Moreover,

[9] For example, Caselli et al. (2021), Brooks et al. (2021), Yeh et al. (2022), and Mertens (2022, 2023), or, for closely related measures, Mertens (2020) and Dobbelaere et al. (2024).

[10] The conditional cost function refers to the static cost function conditional on other choice variables of the firm, such as capital, that are potentially determined by a dynamic maximization problem.

because Lemma 1 is based on properties of the revenue function instead of on a static cost-minimization problem, it allows more generally for input market imperfections such as adjustment frictions, which invalidate equation (2). A downside of not assuming that any input wedge equals unity is that equation (2) does not identify the firm's price-cost markup.

A second advantage is specific to the context that I consider. If output is unobserved, identifying output elasticities can be challenging. This motivates the Hashemi et al. (2022) insight that plugging an inadvertently estimated revenue elasticity into equation (2) identifies an input wedge rather than a markup.[11] The discussion above shows that even when output elasticities can be identified, the revenue approach has an advantage over the production approach: no frictionlessly adjustable input for which firms are price takers is required.

2.2 Estimating revenue elasticities

I obtain revenue elasticities using several approaches, most of which require specifying and estimating a revenue function. Key results are robust across approaches. My starting point is the large literature that estimates production functions using revenue data (see De Loecker and Syverson (2022) for a review). Using deflated revenue as a proxy for output can bias elasticity estimates when prices vary around the industry-wide deflator (e.g., Klette and Griliches (1996)). As discussed in De Loecker and Goldberg (2014), regressing log deflated revenue on log inputs delivers output elasticities only under restrictions on demand and conduct, imposed either implicitly or explicitly. In particular, moving beyond a setting with no price variation requires assuming monopolistic competition and identical constant elasticity of substitution (CES) demand, as is common in work that uses the production approach to estimate the labor wedge (e.g., Yeh et al. (2022)). Those assumptions also permit the specification of an estimable revenue function.

Following De Loecker (2011), I consider a CES demand function of the form

$$P_{it} = P_{st} \left(\frac{Q_{it}}{Q_{st}} \right)^{\frac{1}{\eta_s}} \Xi_{it}, \tag{4}$$

where P_{st} and Q_{st} denote the average price and the aggregate quantity shifter in industry s, respectively, η_s is the demand elasticity, and Ξ_{it} is a demand shock unobserved by the researcher. The demand shock consists of an industry-year, 3-digit industry g, and firm-year-specific component: $\Xi_{it} = \Xi_{st}\Xi_g\tilde{\Xi}_{it}$. Combining this demand system with a translog production function with Hicks-neutral productivity, $Q_{it} = Q_s(K_{it}, L_{it}, M_{it})\tilde{\Omega}_{it}$, results in a translog revenue function

$$r_{it} = \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + \beta_{kk} k_{it}^2 + \beta_{ll} l_{it}^2 + \beta_{mm} m_{it}^2 + \beta_{kl} k_{it} l_{it} + \beta_{km} k_{it} m_{it} + \beta_{lm} l_{it} m_{it} + \gamma_{st} + \omega_{it} + \epsilon_{it}, \tag{5}$$

[11] Hashemi et al. (2022) build on Bond et al. (2021), who show that markups are equal to unity if $\gamma_{it}^V = 1$ and the revenue elasticity of V is used instead of the output elasticity in equation (2).

where lowercase letters denote the natural logarithm of the corresponding uppercase letter, $\omega_{it} = \frac{\eta_s+1}{\eta_s}\tilde{\omega}_{it} + \tilde{\xi}_{it}$ is unobserved revenue productivity, γ_{st} subsumes all industry-year-specific terms, and ϵ_{it} is introduced to capture mean-zero unanticipated deviations from planned revenue. Revenue elasticities are given by $\theta^x_{it} = \frac{\partial r_{it}(\cdot)}{\partial x_{it}}$, where $x \in \{k, l, m\}$; like output elasticities, they are not restricted to sum to one. Note that the revenue elasticity of an input depends on that input's output elasticity and the elasticity of demand: $\theta^x_{it} = \frac{\eta_s+1}{\eta_s}\tilde{\theta}^x_{it}$. I estimate equation (5) at the 2-digit-industry level in the baseline specification, but revenue elasticities are firm-time specific due to the translog specification.[12]

The assumptions on CES demand and monopolistic competition are sufficient but not necessary. Equation (5) follows whenever conditional inverse demand is log-linear in own quantity, with a slope common at the level of estimation and firm-specific demand heterogeneity entering additively. This includes CES, nested CES, and standard or nested logit with log price, but excludes generic random-coefficients logit and other demand systems in which the slope of inverse demand varies with unobserved output prices so that estimation would require additional data. Conditional on the relevant market aggregates, the same estimating equation applies under standard price- or quantity-setting conduct, but the interpretation of wedges can change. Under oligopoly, measured wedges equal the wedges that govern firm behavior scaled by a common factor that increases in market share. This factor cancels from relative wedges, moves both measured wedges in the same direction and therefore cannot generate a negative correlation between the two, and vanishes as market shares approach zero. Details are in Appendix A.2.

I obtain revenue elasticities using six approaches, described in detail in Appendix B. Following De Loecker (2011) and Petrin and Sivadasan (2013), my baseline estimates use a control-function approach to identify the revenue function. This approach requires inverting an intermediate input demand function to proxy for unobserved revenue productivity. In contrast to production function estimation this inversion does not require controlling for markups, as they are contained in revenue productivity. I include a rich set of controls in intermediate input demand, as well as a control function in the spirit of De Loecker et al. (2016) to correct for unobserved intermediates prices.

Second, I estimate the revenue function at the two-digit-industry-by-two-year level. In Appendix B, I show that this allows for factor-biased technical change at the level of estimation. Moreover, it is a reduced-form way to partially capture such change at the firm level. In a translog function, factor-biased technical change can affect the intercept of the revenue function and the revenue elasticities of all inputs. These estimates allow the revenue intercept and the function of the revenue elasticities to vary at the 2-digit-industry-by-2-year level, and the elasticities to vary at the firm-year level through their dependence on inputs.

[12]Using a Cobb-Douglas production function would lead to qualitatively similar results, but distributions of input wedges are substantially more dispersed in that case. A translog specification strikes a balance between allowing for meaningful heterogeneity in elasticities and keeping estimation manageable.

Third, for the restricted sample of firms that import intermediates from China used in Section 4, I add information on unit prices of intermediate inputs, imported varieties, and trade intensity to the baseline estimates.

Fourth, sufficient input-price variation is required to identify the intermediate input elasticity (Ackerberg et al., 2015; Gandhi et al., 2020). Identification in the baseline estimates may be aided both by variables that shift intermediate input demand without entering the revenue function directly and by the use of more than a decade of data. To allow for the possibility that the available variation remains insufficient, I follow Flynn et al. (2019) and impose returns to scale, so that the intermediate input elasticity need not be estimated.[13]

Fifth, in the spirit of Blundell and Bond (2000), I estimate the revenue function assuming that revenue productivity follows an AR(1) process. The main advantage of this approach over the baseline specification is that ω_{it} is not proxied by observables, so intermediate input demand need not be inverted. For instance, control-function approaches struggle to incorporate strategic interaction as a firm's input demand may depend not only on its own productivity and predetermined inputs, but also on competitors' states and choices.[14] The AR(1) assumption sidesteps such issues at the cost of a more restrictive law of motion.

Finally, I dispense with revenue function estimation and calibrate the elasticities. For all firms, I set the revenue elasticity of labor equal to the median baseline estimate in the full sample. The revenue elasticity of intermediate inputs is calibrated similarly.[15] By Lemma 1, input wedge variation is then driven entirely by revenue shares.

2.3 Data and descriptive statistics

I apply the methodology outlined above to Dutch firm-level data covering the manufacturing sector from 2007 to 2018. The Netherlands is a suitable setting for several reasons. As in much of Western Europe, collective bargaining characterizes Dutch labor markets, which might overturn US-based wage markdown findings, where such bargaining is largely absent (Wong, 2026).[16] Indeed, the Dutch bargaining system is similar to those of countries such as Italy and Portugal, which have previously been studied through the lens of collective

[13] Alternatively, Pham (2023) estimates revenue elasticities using the approach of Gandhi et al. (2020). Like the production approach, this requires assuming that the intermediate input wedge equals one.

[14] Ackerberg and De Loecker (2024) show that, for several standard oligopoly models, the inversion can be restored by augmenting the control function with low-dimensional sufficient statistics for rivals' equilibrium behavior.

[15] Alternative calibration choices, such as using a reasonable guess for the two elasticities based on the literature or taking the median within a firm's industry, yield similar results.

[16] In 2016, employment terms of 78.6 percent of all employees were governed by at least one collective agreement, compared to only 11.5 percent in the United States. Collective bargaining coverage is stable throughout the sample period. OECD countries for which bargaining coverage was at least 70 percent in 2016 are Austria, Belgium, Denmark, Finland, France, Iceland, Italy, the Netherlands, Norway, Portugal, Slovenia, Spain, and Sweden.

bargaining (e.g., Card and Cardoso (2022); Card et al. (2014)).[17] In intermediate input markets, by contrast, collective agreements are absent, so interactions between firms and suppliers are bilateral. Dutch firms also rely heavily on imports.[18] Recent work identifies buyer power for imported intermediates in the French manufacturing sector, raising concerns that such power may also exist in the comparable Dutch context (Morlacco, 2020).

I construct an annual firm-level panel of Dutch manufacturers from 2007 to 2018, using non-public data from Statistics Netherlands (CBS).[19] I observe revenue, full-time-equivalent (FTE) employment, total expenditure on labor and intermediate inputs, the book value of capital, the firm's four-digit NACE industry, and the location of its headquarters. I do not observe R&D expenditure, product count, or overhead costs. CBS routinely checks data quality and contacts firms when reporting errors are suspected. I restrict the sample to firm-year observations with sufficient information to construct labor and intermediate input wedges. In particular, I include only observations with positive revenue, capital, labor, and intermediate input expenditure. The final sample consists of 23,638 firms and 132,722 firm-year observations. Appendix C provides a detailed overview of the data and the sample selection procedure. Table C1 in Appendix C lists the 17 2-digit industries included in the final sample, and Table C2 presents a breakdown of observations by year and 2-digit industry.

I use total labor expenditure as wage measure, as it captures a firm's labor cost more accurately than employees' net or gross salaries. I construct the wage W_{it} as the ratio of a firm's total labor expenditure to its FTE employment. All results apply only to employees on a firm's payroll, as the data on labor expenditure and employment cover only workers employed directly by firms. When interpreting W_{it}, note that net wages are roughly half of labor expenditure per worker in the Netherlands. This gap is mainly due to mandatory social security contributions and high income tax rates.[20]

Table 1 reports summary statistics for key firm-level variables. Revenue and inputs are positively skewed, with means often substantially exceeding medians. Right-tail outliers imply that standard deviations are much larger than interquartile ranges. These outliers are typically large firms found in international datasets such as Orbis or Compustat. The median number of employees is 10, while the mean is 30.13. To avoid convergence issues when estimating revenue elasticities, I use the full sample for estimation and control for size throughout the paper rather than dropping smaller firms. Total labor cost has a median of 505,000 euros and a mean of 1,942,000 euros, both rounded to thousands. The median book

[17] See Hijzen et al. (2019) for a detailed discussion of the collective bargaining institutions in the Netherlands and Portugal, for instance.

[18] Dutch imports totaled 75.16 percent of GDP in 2015, compared to 15.4 percent in the US. Statistics obtained from `https://data.worldbank.org/indicator/NE.IMP.GNFS.ZS?most_recent_year_desc=false&locations=OE-NL`.

[19] Under certain conditions, these data are accessible for research. See `https://www.cbs.nl/en-gb/our-services/customised-services-microdata` for details.

[20] For a detailed breakdown, see (CBS, 2020a, p.77).

Table 1: Summary statistics

variable	p(25)	p(50)	p(75)	mean	s.d.
Revenue	669	1,697	4,699	10,912	115,156
Capital	97	383	1,168	3,475	43,688
Labor	4	10	23	30.13	148.58
Labor expenditure	213	505	1,257	1,942	12,067
Wage	43	53	67	60	39
Labor share	0.22	0.31	0.41	0.32	0.15
Intermediate input expenditure	340	942	2,857	7,857	93,868
Intermediate input share	0.47	0.58	0.68	0.57	0.15

Notes: Summary statistics for key variables based on the full sample of 132,722 observations covering the years 2007 to 2018. p(25), p(50), and p(75) refer to the 25th, 50th, and 75th percentiles of the distribution, respectively. Mean and s.d. are the unweighted mean and standard deviation. Monetary values in thousands of nominal euros, rounded to whole numbers. Non-monetary variables rounded to two decimal points.

value of capital is 383,000 euros, while the mean is 3,475,000 euros. For revenue, the median and mean are 1,697,000 euros and 10,912,000 euros, respectively. Input shares of revenue are distributed symmetrically, with the median labor share at 0.31 and the median intermediate input share at 0.58.

3 Input wedges in Dutch manufacturing

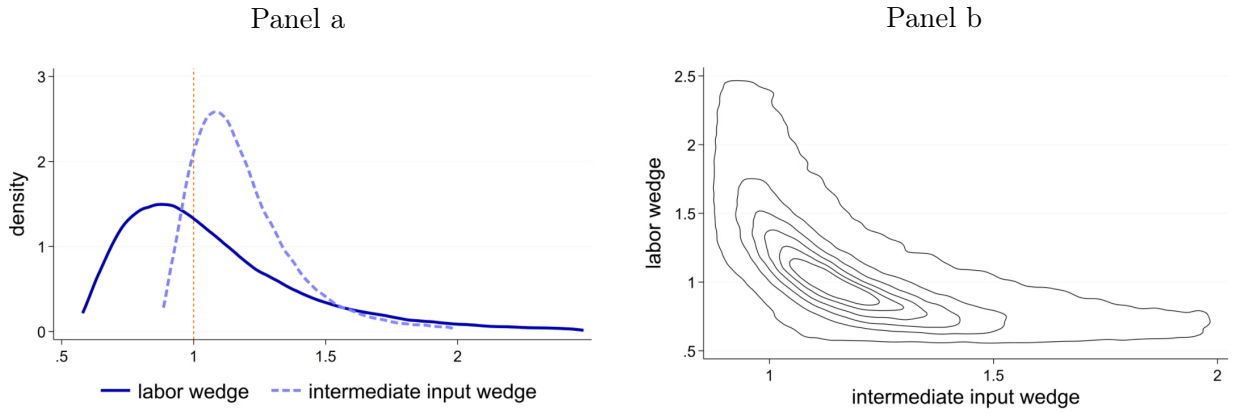
In Section 3.1, I analyze the joint distribution of firm-level input wedges in Dutch manufacturing. Section 3.2 explores idiosyncratic factors underlying this joint distribution. Section 3.3 discusses potential mechanisms.

3.1 The joint distribution of labor and intermediate input wedges

Panel a of Figure 1 displays the distributions of the firm-level labor and intermediate input wedges, and Table 2 gives several percentiles of these distributions. Both wage markups and wage markdowns are prevalent in Dutch manufacturing. The median firm-year has a labor wedge equal to unity – that is, the wage equals the marginal revenue product of labor. At the 25th percentile, wages are 23.46 percent above the marginal revenue product of labor, while at the 75th percentile, workers would receive only 78.74 cents for an additional euro of revenue generated by expanding labor. Wage markups are a common finding in Europe.

Dobbelaere et al. (2026) study Dutch firms and report wage markups in roughly 60 percent of all cases. For Germany, Mertens (2022) reports a median labor wedge of 0.94. Caselli et al. (2021) estimate the inverse of the labor wedge for French firms and report a value corresponding to a labor wedge of 0.64. In stark contrast, Yeh et al. (2022) report that nearly 90 percent of all plant-year observations in US manufacturing are characterized by wage markdowns, in line with widespread monopsony power.[21] I conclude that the data cannot be rationalized by a pure monopsony model, as wage markups are common.

Figure 1: Kernel density estimates of the marginal distributions (panel a) and joint distribution (panel b) of firm-level labor and intermediate input wedges



Notes: Panel a plots univariate kernel densities and a vertical line at 1. Panel b plots contour lines of the bivariate kernel density estimate of the joint distribution. All density estimates use a Gaussian kernel. The figure is based on the full sample of 132,722 observations covering the years 2007 to 2018, with the top and bottom 2 percentiles of each wedge distribution trimmed.

Intermediate input price markdowns characterize Dutch manufacturing. The median intermediate input wedge is 1.14, corresponding to a markup of the marginal revenue product of intermediates over their price of about 14 percent. This implies that, at the median firm, intermediate input suppliers receive only 88 cents for the marginal euro of revenue they generate. Price markdowns cover 111,024 firm-year observations, showing that intermediate input wedges above unity are the norm. In line with these findings, Morlacco (2020) reports average input wedges in French manufacturing for imported intermediate inputs ranging from 1.2 to 1.51 – depending on how markups are calibrated. I conclude that the data provide strong support for the existence of buyer power in Dutch intermediate input markets.

[21] Note that these authors use the production approach, so that the reported numbers underestimate the labor wedge in the presence of intermediate input markdowns. Note also that the ranking of how common wage markups are corresponds to the ranking of collective bargaining coverage for these countries: roughly 100 (France), 80 (Netherlands), 55 (Germany), and 10 percent (US) throughout the sample period. See `https://www.oecd.org/en/data/datasets/oecdaias-ictwss-database.html`.

Table 2: Summary statistics of input wedges

	p(5)	p(25)	p(50)	p(75)	p(95)
Labor wedge (γ_{it}^{L})	0.64	0.81	1.00	1.27	1.98
Intermediate input wedge (γ_{it}^{M})	0.93	1.04	1.14	1.29	1.68
Relative input wedge ($\gamma_{it}^{L}/\gamma_{it}^{M}$)	0.41	0.63	0.86	1.21	2.08

Notes: Summary statistics of the input wedges. The relative input wedge is the firm-time specific ratio of the labor wedge to the intermediate input wedge. p(5), p(25), p(50), p(75) and p(95) refer to the 5th, 25th, 50th, 75th and 95th percentile of the input wedge distribution, respectively. Input wedges are rounded to two decimal points. Statistics based on the full sample of 132,722 firm-year observations covering the years 2007 to 2018.

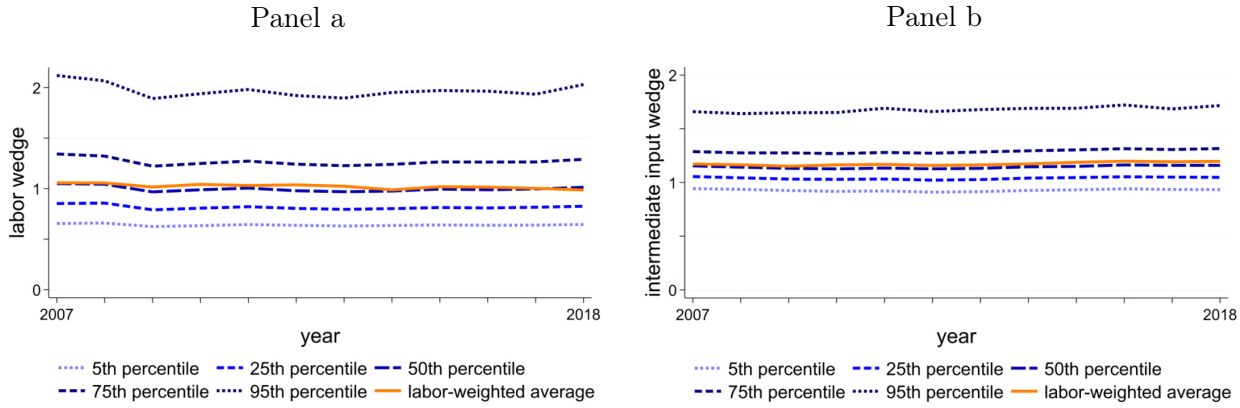
The stark contrast between the prevalence of wage markdowns in US manufacturing observed in Yeh et al. (2022) and the prevalence of wage markups in Dutch manufacturing is not due to differences in identification strategy. Recall from equation (3) that the production approach identifies the labor wedge relative to the intermediate input wedge and recovers the labor wedge by assuming that $\gamma_{it}^{M} = 1$. Table 2 lists selected percentiles of the relative wedge distribution in Dutch manufacturing. Incorrectly assuming that $\gamma_{it}^{M} = 1$ and identifying the labor wedge from the relative wedge increases the spread of the distribution and shifts most of the mass to the left. Therefore, wage markups are even more prevalent in my sample when using the production approach rather than the revenue approach, as intermediate input wedges are typically above unity in Dutch manufacturing.[22]

Panel b of Figure 1 contains a density plot showing that the two input wedges are strongly negatively correlated. Firms that tend to most strongly underpay their intermediate input suppliers relative to their revenue contribution on the margin tend to overpay their workers most on the margin. Size, therefore, cannot be correlated with purchasing power in both input markets. Indeed, if observations are weighted by firms' size – employment – this correlation is maintained, and the mass in the distribution of the labor wedge displayed in Figure 1 shifts slightly to the left, while the mass in the distribution of the intermediates input wedge shifts slightly to the right. That is, larger firms have somewhat higher markdowns in intermediate input markets, and pay their employees somewhat more relative to their marginal revenue product, than smaller firms.

Figure 1 shows that both input wedge distributions exhibit substantial dispersion. One potential explanation is that average differences between different years drive this dispersion. Adjustment frictions, for example, could make input wedges sensitive to aggregate economic

[22]Likewise, if buyer power for intermediates exists in the US, as suggested by Alviarez et al. (2023), the labor wedges in Yeh et al. (2022) would underestimate the degree of monopsony power in US labor markets.

Figure 2: Distribution of firm-level labor wedges (panel a) and intermediate input wedges (panel b), over time



Notes: Selected percentiles of input wedge distributions and labor-weighted average of input wedges, over time. Based on the full sample of 132,722 observations.

shocks. Figure 2 plots selected percentiles of the two input wedge distributions over time. Both distributions are remarkably stable over the sample period. Figures D1 and D2 in Appendix D verify that revenue elasticities, input shares, and hence the ratio of labor and intermediate input expenditure, are also stable over time. These results do not imply that there is no time-series variation whatsoever. Figure D3 in Appendix D plots year effects from regressions of the log wedges with industry or firm fixed effects: both wedges fall early in the sample, after which the intermediate input wedge recovers while the labor wedge remains below its 2007 level. Rather, Figure 2 shows that the cross-sectional dispersion of both input wedge distributions is much larger than the variation over time of a given percentile. In line with this, Table 3 shows that the standard deviation of the two input wedges and their correlation are virtually unchanged when we remove variation due to average differences between years from the data.

Alternatively, there might be significant differences in dispersion between industries. This is not the case. Table 3 shows that the standard deviation of the input wedges barely shrinks as we remove variation related to average industry-by-year differences from the data. The correlation between the wedges, on the other hand, becomes slightly larger. The observed dispersion of input wedges, and their correlation, therefore, is likely related to idiosyncratic factors, such as firm rents, rather than industry-level factors, such as regulation or industry institutions, or aggregate economic shocks that affect all firms in a similar fashion.

That the dispersion of the input wedge distributions – and the correlation between the wedges – is a within-industry phenomenon does not imply that there are no industry-specific differences in the location of those distributions. Table E1 in Appendix E displays industry-specific medians and means of the labor wedge and the intermediate input wedge, and

their correlation. Industries with the largest wage markdowns include the manufacture of chemicals, paper products, computer and electric products, and plastic and rubber. These industries are also among those with the largest wage markdowns in US manufacturing, although markdowns are substantially larger there (Yeh et al., 2022). In contrast, the manufacture of food, minerals, furniture, and motor vehicles exhibit the largest wage markups – a phenomenon largely absent in US manufacturing. While all industries are characterized by intermediate input price markdowns, on average, such markdowns stand out in, for instance, the manufacture of basic metals, fabricated metals, minerals, and furniture. These results are similar to those of Morlacco (2020) for French manufacturing, although the levels are higher in France than in the Netherlands.

Table 3: Standard deviation and correlation coefficient of input wedges

Input wedge variation used	Standard deviation γ_{it}^{L}	Standard deviation γ_{it}^{M}	Correlation γ_{it}^{L} and γ_{it}^{M}
Unconditional	0.448	0.263	-0.541
Conditional on year fixed effects	0.448	0.262	-0.544
Conditional on year-by-2-digit fixed effects	0.441	0.258	-0.556
Conditional on year-by-3-digit fixed effects	0.434	0.252	-0.553

Notes: Standard deviation of input wedges and their correlation coefficient, by variation in the data used. Input wedge data are the sum of the constant and the residual from a regression of the input wedge on a constant and the fixed effects mentioned in the row. Based on the full sample of 132,722 observations covering the years 2007 to 2018.

Summing up, the two input wedges are negatively correlated: firms with higher markdowns in the intermediate input market pay their employees more relative to their marginal revenue product than firms with lower markdowns. Both wage markups and wage markdowns are common in Dutch manufacturing. In the intermediate input market, on the other hand, price markdowns are the norm. There is substantial dispersion in the distribution of both wedges – particularly for the labor wedge. Both this dispersion and the negative correlation between the two wedges cannot be explained by average differences over time or between different industries and are, therefore, likely related to idiosyncratic factors. The next section investigates such factors.

3.2 Wage variation drives the correlation between the input wedges

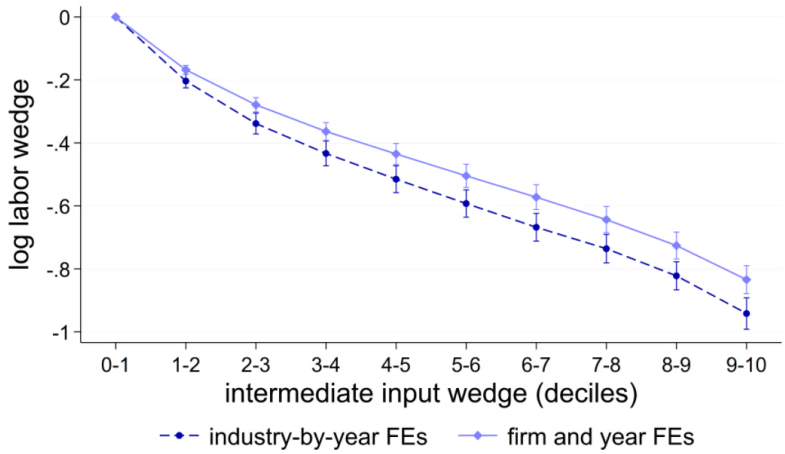
To decompose the relationship between the input wedges, I use the following non-parametric regression to correlate the labor wedge with a variable of interest, x_{it},

$$ln(\gamma_{it}^{L}) = \beta_0 + \sum_{dc=2}^{10} \beta_{dc}^{x} \mathbb{I}_{x_{it} \in X_{dc}} + \mathbf{Controls}_{it} + \varepsilon_{it}, \tag{6}$$

where $\mathbb{I}_{x_{it} \in X_{dc}}$ is an indicator that equals 1 if x_{it} lies between the $(dc-1)^{th}$ and dc^{th} decile of the distribution of x_{it} in the full sample and $\mathbf{Controls}_{it}$ is a vector of control variables that includes nine indicators for the top nine deciles of the size-distribution based on employment, and either firm and year fixed effects or 3-digit-industry-by-year fixed effects. This non-parametric regression, inspired by Haltiwanger et al. (2013), allows for non-monotone relations between the variables of interest. With firm and year fixed effects, the coefficients are identified from within-firm changes over time, net of common year shocks and size controls. With 3-digit-industry-by-year fixed effects, they are identified from cross-firm differences within the same 3-digit industry-year cell, conditional on firm size. In the remainder of this section, I refer to these as within-firm and within-industry-year variation, respectively. All regressions are run on the full sample, and standard errors are clustered at the 3-digit-industry level to account for any dependencies between observations within an industry and over time.

As a starting point, Figure 3 displays the relationship between the labor wedge and the intermediate input wedge. The correlation is negative, monotone, and strong – passing through all deciles of the distribution of the intermediates wedge induces all of the variation in the labor wedge. Moreover, the magnitude of the within-firm association is very similar to that of the within-industry-year association, suggesting that within-firm changes of the intermediate input wedge contribute to the aggregate distribution of the labor wedge.

Figure 3: Regression of the firm-level labor wedge on the firm-level intermediate input wedge



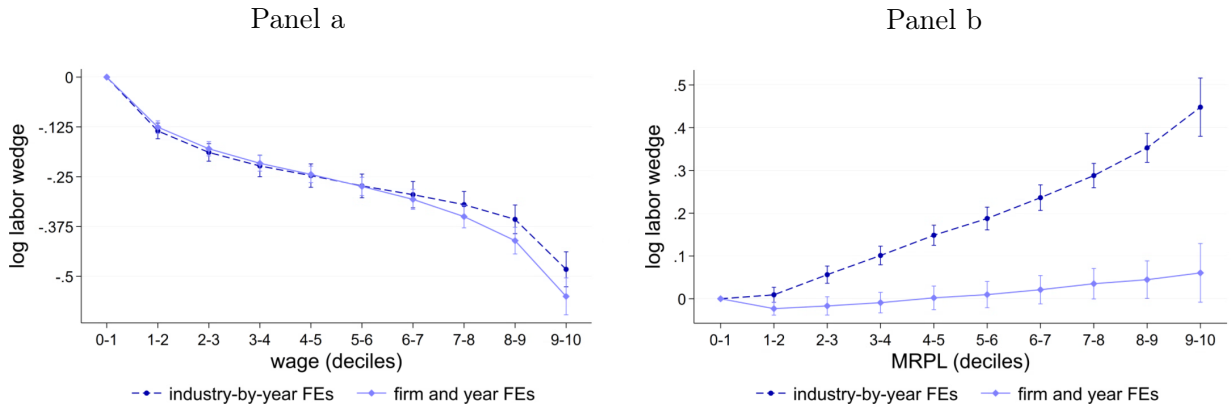
Notes: Point estimates and 95 percent confidence intervals from a non-parametric regression of the log labor wedge on the intermediate input wedge, controlling for employment. Explanatory variables are binned by decile, with the first decile normalized to 0. 3-digit-industry-by-year fixed effects (dashed line) or firm and year fixed effects (solid line) are included. Standard errors are clustered at the 3-digit-industry level. Based on the full sample of 132,722 observations covering the years 2007 to 2018.

To determine whether labor wedge variation is primarily due to wages or to marginal

revenue products, Figure 4 plots point estimates and confidence intervals from regressions relating the labor wedge to the wage (panel a) and the marginal revenue product of labor (MRPL) (panel b). Both within industries and within firms, wages are strongly, negatively, and monotonically correlated with the labor wedge. In addition, firms that have higher labor wedges have higher marginal revenue products of labor. Within firms, however, there is no relation between whether workers are over or underpaid on the margin and the revenue they generate on that margin – labor wedge variation is driven entirely by wage variation.

Panel a of Figure D4 in Appendix D verifies that similar patterns hold for the average revenue product of labor, although a very moderate positive association between revenue per worker and the labor wedge is observed. In contrast, panel b shows that value added per worker is strongly negatively correlated with the labor wedge, both within industry-year and within firms. As firms increase their value added per employee, therefore, they start paying their employees higher wages, even though their marginal and average revenue products have barely, if at all, changed. As the difference between revenue and value added per worker is intermediate input expenditure, these results suggest a role for the intermediates market in shaping labor market imperfections.

Figure 4: Regressions of the firm-level labor wedge on the wage (panel a) and the marginal revenue product of labor (panel b)

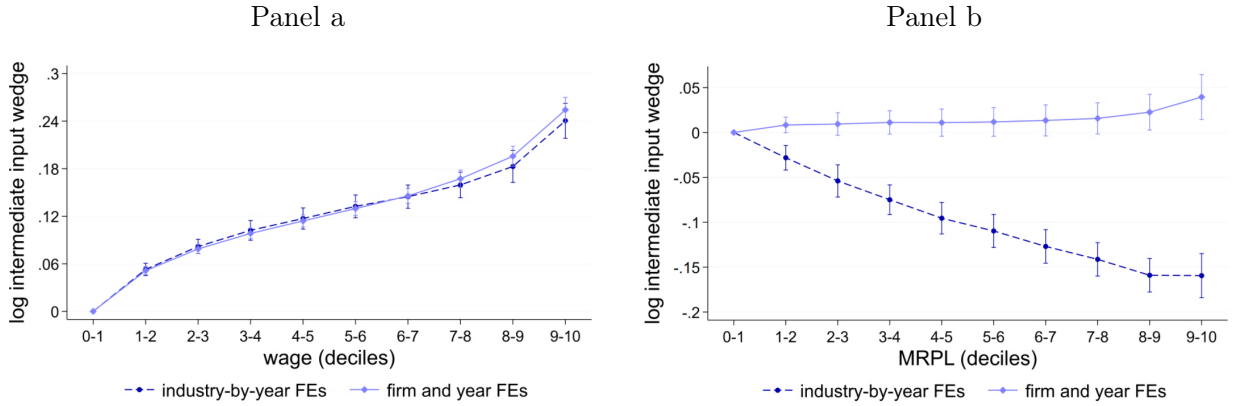


Notes: Point estimates and 95 percent confidence intervals from a non-parametric regression of the log labor wedge on the wage (panel a) and the marginal revenue product of labor (panel b), controlling for employment. Explanatory variables are binned by decile, with the first decile normalized to 0. 3-digit-industry-by-year fixed effects (dashed line) or firm and year fixed effects (solid line) are included. Standard errors are clustered at the 3-digit-industry level. Based on the full sample of 132,722 observations covering the years 2007 to 2018.

Figure 5 displays coefficients from regressions associating the intermediate input wedge to the wage (panel a) and the marginal revenue product of labor (panel b). High intermediate input wedges, consistent with substantial buyer power, are associated with high wages – both when comparing different firms in the same industry and when comparing different years of

the same firm. However, as with the labor wedge, a within-firm correlation between the intermediates wedge and the marginal revenue product of labor is almost entirely absent, while the concomitant within-industry-year relationship is negative and pronounced. Within firms, then, the degree to which a firm underpays its intermediates suppliers on the margin is largely unassociated with the marginal revenue productivity of its workers.

Figure 5: Regressions of the firm-level intermediate input wedge on the wage (panel a) and the marginal revenue product of labor (panel b)



Notes: Point estimates and 95 percent confidence intervals from a non-parametric regression of the log intermediate input wedge on the wage (panel a) and the marginal revenue product of labor (panel b), controlling for employment. Explanatory variables are binned by decile, with the first decile normalized to 0. 3-digit-industry-by-year fixed effects (dashed line) or firm and year fixed effects (solid line) are included. Standard errors are clustered at the 3-digit-industry level. Based on the full sample of 132,722 observations covering the years 2007 to 2018.

Panel a of Figure D5 in Appendix D shows that higher markdowns in the intermediates market are associated with lower revenue per employee, although the within-firm correlation is quantitatively small. Value added per worker, on the other hand, is positively associated with the intermediate input wedge in both specifications (panel b). Therefore, in years where my estimates suggest a firm has more buyer power in the intermediates market, wages and value added per worker are higher than when low buyer power is estimated, while the average and marginal revenue of workers is largely unchanged.

Finally, Figure D6 in Appendix D correlates input wedges to firm size, measured as employment or value added. Larger firms have lower labor wedges and slightly higher intermediate input wedges. The patterns are remarkably similar for both measures of firm size, but the point estimates are somewhat larger for value added. In short, larger firms generate more value added, pay workers more relative to their marginal revenue product, and pay their input suppliers less relative to their marginal revenue product than smaller firms. This contrasts with earlier results that found a positive relationship between firm size and labor wedges (e.g., Yeh et al. (2022) and Mertens and Mottironi (2025)).

Summing up, a key predictor of whether and how much workers are over or underpaid on the margin relative to their revenue product is the size of the markdown in the intermediates market. In years where firms mark down the intermediates price more, they have higher value added per worker and pay their employees more, both in absolute terms and relative to workers' marginal revenue product. This occurs even though workers are not more revenue productive on average, or on the margin, when the intermediates wedge is higher. Wage variation, in turn, explains most variation in the labor wedge over time, rather than differences in the marginal revenue product of labor.

Before moving on to an explanation for the observed patterns, I stress that they are robust to alternative approaches to obtaining revenue elasticities. Table E2 in Appendix E shows that intermediates price markdowns are the norm regardless of how revenue elasticities are obtained – the median ranges from 1.13 with calibrated elasticities to 1.25 with a CRS-constrained revenue function. Both wage markups and wage markdowns are common in all specifications. However, with a time-specific revenue function, an AR(1) revenue productivity process, or CRS-constrained estimate, the location of the distribution moves to the left – at the median wages are 12 to 32 percent above the marginal revenue product of labor. The reason is that the labor elasticities derived with these approaches are more similar to the lower OLS estimates than to the baseline estimates. Table E3 in Appendix E verifies that the two input wedges are negatively correlated both within firms and within industry-year for all but one exception, the within-firm AR(1) estimate, which is close to zero. I conclude that the negative correlation between the two input wedges, as well as the omnipresence of intermediate input markdowns and the prevalence of wage markups, are robust to how the revenue function is estimated.

3.3 Interpreting the negative correlation: rents or technology?

The results documented above have a natural interpretation in terms of rent sharing across input markets. Firms with larger intermediate input wedges pay higher wages and have lower labor wedges, both when comparing firms within the same industry and year and when comparing the same firm over time. Within firms, these relationships are primarily associated with wages rather than with workers' marginal or average revenue product. Whereas the existing literature has primarily considered rents generated in output markets, the evidence here suggests that rents generated through buyer power in the intermediate input market may be partly transmitted to workers. Such a mechanism would connect two literatures that have developed separately: work documenting that firms share product market rents with their employees (see Card et al. (2018) for a survey) and work documenting buyer power for intermediate inputs (e.g., Morlacco (2020); Rubens (2023)).

An alternative explanation is due to Raval (2023), who documents similar empirical results and shows that they can be rationalized as factor-biased technical change in the

absence of input market imperfections. Using seven manufacturing datasets, he computes the product of the output market markup and an input wedge using equation (2) separately for labor and for intermediate inputs and shows that the two objects are negatively correlated in the cross-section, have within-industry time trends in opposite directions, and that the labor estimate is more dispersed than the intermediate input estimate.[23] Raval (2023) assumes that both input bundles have wedges exactly equal to one, so that both estimates should deliver the same object. He attributes the disagreement to factor-biased technical change: when technology is not Hicks neutral, using Hicks-neutral production function estimators can bias the estimated output elasticities. Using a Monte Carlo exercise, in which plants minimize cost, face competitive input markets, and differ in labor-augmenting productivity and markups, he confirms that such differences alone can generate a negative correlation between the two estimates.[24]

Recall that Lemma 1 is unaffected by factor-biased technical change as captured by $R_{it}(K_{it}, \Omega_{it}^L L_{it}, M_{it}, \Omega_{it})$. The concern is not with the mapping from elasticities and revenue shares to wedges, but with the estimation of the elasticities themselves. Denoting the estimated revenue elasticity of input X by $\hat{\theta}_{it}^X$ and the true one by θ_{it}^X, equation (1) implies

$$\hat{\gamma}_{it}^X = \frac{\hat{\theta}_{it}^X}{\theta_{it}^X} \gamma_{it}^X . \tag{7}$$

Under the assumptions of Raval (2023), the second factor equals unity for every firm in every year and the estimated wedge captures only deviations of the estimated elasticity from the true elasticity caused by factor-biased technical change. Alternatively, if elasticities are identified without error, the first factor equals unity and equation (7) identifies the input wedge. Equation (2) identifies the product of the output market markup and an input wedge with accurate elasticity estimates, so that a negative correlation between the labor- and intermediate-input-based measures can instead reflect input market imperfections.

Distinguishing these interpretations with accounting data is difficult because they provide different explanations for variation in the same accounting objects. Under the revenue interpretation used in this paper,

$$\ln \hat{\gamma}_{it}^L - \ln \hat{\gamma}_{it}^M = \ln \left(\frac{\hat{\theta}_{it}^L}{\hat{\theta}_{it}^M} \right) - \ln \left(\frac{W_{it} L_{it}}{P_{it}^M M_{it}} \right) . \tag{8}$$

In Raval (2023), the same expression holds with input-specific markup estimates on the left and estimated output elasticities in the first term on the right. In both cases, the

[23] While the negative correlation and dispersion results hold in my data, the Dutch data show a common initial decline over time followed by a modest separation, which is more pronounced when identified from within-firm variation (Figure D3 in Appendix D).

[24] Most of the production function estimates in Raval (2023) use deflated revenue, so an alternative reading of his results put forward by Hashemi et al. (2022) following Bond et al. (2021) is that Raval (2023) divides a revenue elasticity by a revenue share and therefore identifies two input wedges in the sense of Lemma 1.

relative input-specific measure depends on the estimated elasticity ratio and on the labor-to-intermediate-input cost ratio. Factor-biased technology can move both objects by changing the relative productivity and equilibrium use of labor and intermediate inputs. Cross-market rent sharing can move the cost ratio because changes in firms' rents affect wages at a given marginal revenue product of labor, and change the ratio of elasticities through changed input use. The negative correlation between the two input-specific measures by itself therefore does not distinguish the two explanations. The same tension is present in the broader literature, where papers that allow for factor-biased technical change tend to rule out non-unit input wedges by assumption and vice versa (Rubens et al., 2026).

The broader empirical patterns do not distinguish the two interpretations perfectly either. Within industry-year, the labor wedge increases strongly with both the marginal revenue product of labor and revenue per employee, whereas the intermediate input wedge decreases with both. These patterns are consistent with the mechanism in Raval (2023). They are not unique to that mechanism, however, as more productive firms may have higher marginal revenue products and systematically different wedges, while wages need not adjust one-for-one with marginal revenue productivity. The documented within-industry-year relationships consequently show that technological heterogeneity is a plausible contributor, but do not separate it from heterogeneous rents and rent sharing.

The within-firm evidence that changes in input wedges coexist with wage changes while leaving average and marginal productivity of workers largely unchanged is consistent with a rent-sharing channel, but is not conclusive evidence either. Table E3 shows that the negative relationship between the two wedges is visible broadly in wedges relying on different estimation approaches. In particular, the wedges are negatively correlated with two-digit-industry-by-two-year-specific revenue functions, which absorb factor-biased technical change common to those cells. But this specification does not fully absorb factor-biased technical change if it is partly firm-time specific. However, by equation (7) it likely increases the signal-to-noise ratio if revenue productivity is indeed non-neutral.

Overall, I read the empirical evidence documented above and in Raval (2023) as indicating that factor-biased technical change and input market imperfections are complementary rather than competing explanations of the patterns documented above. Identifying the novel rent-sharing mechanism more sharply requires variation in the intermediate input market whose source is known and is not itself a change in technology. The next section constructs such variation by combining firm-level trade data with movements in the exchange rate between the euro and the Chinese renminbi, which change the cost of imported intermediate inputs for firms that source from China.

4 How exchange rate fluctuations affect employees

This section investigates rent sharing as a mechanism that links market imperfections in labor and intermediate input markets. Section 4.1 develops the empirical strategy, Section 4.2 presents the main estimates, Section 4.3 examines mechanisms, and Section 4.4 evaluates the exclusion restriction and alternative interpretations.

4.1 Identifying cross-market rent sharing

I connect to a large empirical literature on rent sharing that estimates how wages respond to rents per worker (Card et al., 2018). While that literature connects wages to measures of overall firm rents, I primarily focus on relating wages relative to employees' marginal revenue product to intermediate input prices relative to the marginal revenue product of intermediates. To make the connection to the rent-sharing literature precise, I extend a standard bargaining framework used to motivate rent-sharing regressions with upward-sloping supply curves in the labor and intermediate input markets – the central ingredient used to explain estimated input wedges by papers using the production approach (e.g., Yeh et al. (2022); Rubens (2023)). As is common in the rent-sharing literature, the model serves only to motivate the reduced-form empirical specifications.

In the model, wages are determined by Nash bargaining between a firm facing an upward-sloping labor supply curve and its collectively organized workers, in line with the high collective bargaining coverage in the Netherlands. In the intermediate input market, the firm unilaterally selects a point on its upward-sloping supply curve so that intermediates prices are marked down relative to the marginal revenue product of intermediates. Quasi-rents per worker – revenue net of payments to suppliers of intermediates and capital, per worker – can be written as a function of the intermediate input wedge, $QR_i(\gamma_i^M) \equiv \frac{R_i - P_i^M M_i - P_i^K K_i}{L_i} = \frac{R_i}{L_i}\left(1 - \frac{\theta_i^M}{\gamma_i^M} - \frac{P_i^K K_i}{R_i}\right)$. Quasi-rents per worker increase in the intermediate input wedge: the larger the markdown, the smaller the share of revenue paid to suppliers relative to their marginal contribution, and the more surplus remains per worker. Let $\widetilde{QR}_i(\gamma_i^M) \equiv QR_i(\gamma_i^M)/MRPL_i$ denote quasi-rents expressed in units of labor's marginal revenue product. Denoting workers' bargaining power by ϕ, the bargaining solution implies

$$\frac{W_i}{MRPL_i} = \phi\,\widetilde{QR}_i(\gamma_i^M) + (1-\phi)\frac{\varepsilon_i^L}{\varepsilon_i^L + 1}, \qquad W_i = \phi\,QR_i(\gamma_i^M) + (1-\phi)\frac{\varepsilon_i^L}{\varepsilon_i^L + 1}\,MRPL_i, \tag{9}$$

where ε_i^L is the labor supply elasticity. The expression on the right is the wage equation of the rent-sharing literature with two modifications. The upward-sloping labor supply curve replaces the exogenous outside option with the monopsonistic wage the firm could otherwise pay (Wong, 2026), and the upward-sloping intermediates supply curve makes the quasi-rents that are shared a function of the intermediate input wedge. The expression on the left is the

wage expression divided by the marginal revenue product of labor and connects the same rent-sharing logic to the objects of this paper. The full model is in Appendix A.3.

Equation (9) nests the two interpretations of the labor wedge that this paper engages. If workers have no bargaining power ($\phi = 0$), the labor wedge equals $\frac{\varepsilon_i^L+1}{\varepsilon_i^L}$: the pure monopsony benchmark of Yeh et al. (2022), in which the wedge measures labor market power, is pinned down by the labor supply elasticity. If workers do have bargaining power, the labor wedge also loads on quasi-rents per worker, so that buyer power for intermediates raises the surplus that bargaining divides, putting upward pressure on wages and downward pressure on labor wedges, all else equal. Rents originating in the output market put the same upward pressure on wages and are the focus of most of the rent-sharing literature. The aim of this section is to isolate rents originating in the intermediate input market. The empirical challenge is therefore to find plausibly exogenous variation where the intermediate input wedge is shifted.[25]

Motivated by the model, I follow the rent-sharing literature by specifying a linear regression that links a measure of worker compensation to a measure of firm rents

$$\ln(\gamma_{it}^L) = \beta_1 \ln(\gamma_{it}^M) + \beta_c' \mathbf{X}_{it} + \varepsilon_{it}, \tag{10}$$

where $\mathbf{X}_{it}$ collects a constant, lagged values of log input use and revenue, and firm and year fixed effects, and ε_{it} contains all other unobserved shocks affecting the firm's labor wedge. The coefficient of interest, β_1, measures the elasticity of the labor wedge with respect to the intermediate input wedge. Later, I decompose this elasticity into a wage elasticity and a marginal revenue product of labor elasticity. Firm fixed effects not only help condition on time-invariant differences between firms, but are also fitting as the rent-sharing mechanism described above is a within-firm phenomenon. Even so, estimating equation (10) with OLS might yield biased estimates of β_1, as the wedges are jointly determined in equilibrium.

Estimating β_1 therefore requires variation that shifts the intermediate input wedge and is unrelated to the other determinants of the labor wedge collected in ε_{it}. I obtain such variation from the exposure to exchange rate fluctuations of firms that import intermediates from China. The intuition is that exchange rate fluctuations correspond to rotations of the intermediates supply curve for importers, with firms that previously relied heavily on imports being most exposed. For such rotations to generate a worthwhile sample, a candidate country of origin needs to display meaningful exchange rate variation against the euro, and a non-negligible number of Dutch firms need to import from that country. China is the leading candidate that satisfies both criteria, as most Dutch imports originate in countries whose currency is the euro or a currency (essentially) pegged to the euro.

I define the following instrument, which interacts a firm's exposure to exchange rate

[25] Appendices A.4 and A.5 verify that the rent-sharing link between the two input wedges remains if intermediate input prices are bargained over, and if labor faces adjustment costs.

fluctuations with the exchange rate realization in a given year

$$Z_{it}^{E/R} = S_{io}^{import} \ln\left(\frac{EUR}{RMB}_t\right), \tag{11}$$

where $\frac{EUR}{RMB}_t$ is the exchange rate defined as euro per Chinese renminbi and S_{io}^{import} is the share of imports from China in overall intermediate input expenditure, measured in the first year in which the firm imports from China. I exclude that year from estimation and hold the share fixed thereafter. The $Z_{it}^{E/R}$ notation is meant to remind readers that the value of the instrument *decreases* if the euro appreciates (i.e., if the intermediates supply curve rotates out). To construct the instrument and the trade-related outcomes studied below, I combine the "International Trade in Goods" (IHG) dataset from Statistics Netherlands, available from 2009, with bilateral trade flows from BACI and annual-average exchange rates from the OECD. From these data, I construct each firm-year's import status, China import share, number of imported varieties, and an import price index. Appendix C contains more details.

The instrument in equation (11) combines exogenous shocks – exchange rate realizations – with predetermined but potentially non-random exposure – initial import shares. I therefore follow the non-random exposure to exogenous shocks framework of Borusyak and Hull (2023), in which reduced-form regressions of firm outcomes on $Z_{it}^{E/R}$ have a causal interpretation under two requirements. Let $e_t = \ln\left(\frac{EUR}{RMB}_t\right)$, let $\mathcal{F}_{i,t}$ denote the information available to firm i at the end of year t, and let o_i denote the year in which firm i's initial share is constructed.

Requirement 1 (Exogenous exchange rate innovations) Conditional on the controls, the exchange rate shocks underlying the instrument are as good as randomly assigned.

I formalize this by modeling the assignment process of the exchange rate as a random walk without drift,

$$e_t = e_{t-1} + u_t, \qquad \mathbb{E}[u_t \mid \mathcal{F}_{i,t-1}] = 0 \text{ for all firms}, \tag{12}$$

with innovations u_t orthogonal to contemporaneous firm-specific unobserved outcome determinants conditional on $\mathbf{X}_{it}$: $\mathbb{E}[u_t\, \varepsilon_{it} \mid \mathbf{X}_{it}] = 0$. The random walk assumption is motivated by the exchange rate forecasting literature and is consistent with my sample. A random walk without drift remains the most difficult forecasting benchmark to outperform (Rossi, 2013). Moreover, Table E4 in Appendix E reports autoregressive coefficients indistinguishable from unity using the underlying monthly exchange rate realizations of my sample, and finds neither the intercept nor the mean monthly change to be significantly different from zero.

Requirement 2 (Controls absorb the expected instrument) The instrument's expected value under the shock assignment process, $\mathbb{E}[Z_{it}^{E/R} \mid \mathcal{F}_{i,o_i}]$, lies in the span of the controls $\mathbf{X}_{it}$.

The random walk in equation (12) implies $\mathbb{E}[e_t \mid \mathcal{F}_{i,o_i}] = e_{o_i}$ for all $t > o_i$, so that the expected instrument,

$$\mathbb{E}[Z_{it}^{E/R} \mid \mathcal{F}_{i,o_i}] = S_{io}^{import} e_{o_i}, \tag{13}$$

is constant within firms and absorbed by the firm fixed effect. The fixed effects specification together with the random walk therefore satisfies both requirements. Importantly, neither requirement concerns firms' initial import shares, so S_{io}^{import} need not be randomly assigned. Initial import shares may be correlated with persistent differences in firms' sourcing strategies, productivity, and labor market outcomes, and any such level differences are absorbed by the firm fixed effects.

Under Requirements 1 and 2, variation in the instrument conditional on the controls reflects the particular exchange rate innovations that were realized rather than firms' non-random exposure to a typical realization of the exchange rate process. The reduced-form regressions of firm outcomes on $Z_{it}^{E/R}$ reported below therefore estimate the causal effect of exchange-rate-induced changes in the cost of importing, whatever the channels through which that effect operates. Interpreting IV estimates of β_1 in equation (10) as the effect of the intermediate input wedge on the labor wedge requires, in addition, an exclusion restriction.

Requirement 3 (Exclusion) Conditional on the controls, the instrument affects the labor wedge only through the intermediate input wedge.

Requirement 3 rules out, most prominently, direct effects of exchange rate fluctuations on labor markets and indirect effects operating through output markets. I defer a discussion of its plausibility to the end of this section, as the evidence on mechanisms presented there – in particular, the absence of responses of revenue, employment, and workforce composition – speaks directly to the channels through which violations would have to operate.

The identifying variation thus comes from comparing changes in outcomes, net of aggregate year effects and conditional on lagged input use and revenue, across firms with different predetermined exposure to common, unanticipated exchange rate movements. For an observation to enter the estimation sample, the firm must import from China in the current and the previous year, so the comparison is among continuing importers. When the euro appreciates, the instrument falls by more for a firm that heavily relied on Chinese imports at baseline than for a marginal importer, and the differential change in their outcomes identifies the effect of the induced supply curve rotation. Provided that an appreciation of the euro weakly raises the intermediate input wedge of every exposed importer, the IV estimand is

a weighted average of firm-year-specific elasticities of the labor wedge with respect to the intermediate input wedge, with nonnegative weights that are largest where the instrument shifts the intermediates wedge most (Angrist et al., 2000).

As not all firms import from China in consecutive years and the trade data are only available from 2009 onward, the estimation sample is considerably smaller than the full sample described in Section 3: it covers 2,203 firms and 10,630 observations. Table E5 in Appendix E reveals that the distributions of the two wedges in this reduced sample are similar to their full-sample counterparts in Table 2: the median labor and intermediate input wedges are 1.04 and 1.12, respectively. The table also indicates that the distribution of firms' share of imports from China in overall intermediate input expenditure is very skewed – as is customary in trade data. At the median, this share is 1.2 percent, while at the 75th and 95th percentiles it is 6.1 percent and 32.6 percent, respectively. Finally, Figure D7 in Appendix D shows substantial and non-monotone exchange rate variation between the Netherlands and China: at its highest point, in 2015, the exchange rate was 37 percent above its lowest value, which occurred in 2009.

4.2 Exchange-rate-induced changes in wedges and wages

Table 4 presents estimates of equation (10). Column 1 displays reduced-form estimates, which imply that as importing from China becomes cheaper, firms increase the compensation of their employees relative to their marginal revenue product – recall that the initial share is held fixed, and the value of the instrument declines as the euro appreciates. Column 2 shows that markdowns in the intermediate input market increase as the euro appreciates. To see that the instrument can induce meaningful variation, consider a firm that moves from the 95th to the 5th percentile of the distribution of the instrument – an absolute difference of 0.616. Such a firm would increase its intermediates markdowns by 9.61 percent and reduce its labor wedge by 6.09 percent. Note also that the F-stat of the instrument in the first stage regression is 485.01, implying that the instrument is highly relevant.

Column 3 gives IV estimates of equation (10). A one percent increase in the intermediate input wedge leads to firms reducing the labor wedge by 0.68 percent – the more a firm underpays its intermediates suppliers on the margin, the more it pays its workers relative to their marginal revenue contribution. Ultimately, this section aims to quantify if and to what extent intermediate input market imperfections affect labor market imperfections. Recall from Table 2 that the median firm has an intermediates wedge 14 percent above that of a firm operating in a competitive intermediates market. Table 4 implies that if this firm started operating in a competitive intermediates market, its labor wedge would increase by 9.4 percent – i.e., its workers would be paid less relative to their marginal revenue product. Column 4 shows OLS estimates of equation (10) for reference. As anything that affects one input market can influence both wedges, signing the bias of the OLS estimates a priori is

Table 4: Reduced form, first stage, IV, and OLS estimates of the relationship between the labor wedge and the intermediate input wedge

	Reduced form	First stage	IV	OLS
Intermediate input wedge			-0.683	-1.542
			(0.216)***	(0.054)***
$Z^{E/R}$	0.102	-0.149		
	(0.029)***	(0.007)***		
Firm and year fixed effects	Yes	Yes	Yes	Yes
R^2	0.04	0.10	0.36	0.51
F-stat of instrument		485.01	485.01	
Observations	10,630	10,630	10,630	10,630

Notes: Results from regressions of the labor wedge on the exchange rate instrument (Column 1), the intermediate input wedge on the exchange rate instrument (Column 2), the labor wedge on the intermediate input wedge instrumenting for the latter with the exchange rate instrument (Column 3), and the labor wedge on the intermediate input wedge (Column 4). $Z^{E/R}$ is defined in equation (11). All regressions include lagged inputs and revenue as controls; All variables in logs except the instrument. Standard errors clustered at the firm level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

difficult. It turns out that the OLS estimate is larger in magnitude than the IV.

The effect of the intermediate input wedge on the labor wedge operates entirely through the wage. Table 5 shows reduced form, IV, and OLS estimates relating either the wage or the marginal revenue product of labor to the intermediate input wedge. As importing from China becomes cheaper, Dutch manufacturing firms start paying higher wages, while the marginal revenue product of their employees does not change (Columns 1 and 4). Unsurprisingly, therefore, the IV estimates in Columns 2 and 5 indicate that as firms exhibit higher markdowns in the intermediate input market, they increase their workers' wages even though the MRPL does not change. These estimates imply that if the firm with the median intermediates wedge were placed in a competitive intermediate input market, its workers would see their wages reduced by 8.5 percent. In contrast, the OLS estimates suggest that both wages and marginal revenue products increase, but wages by more (Columns 3 and 6).

Table E6 in Appendix E verifies that these conclusions are not driven by how the estimation sample is selected and conditioned. Column 1 drops the requirement that firms import from China in the current and the previous year, and Column 2 removes the lagged input and revenue controls. These variations keep the point estimates similar and significance lev-

Table 5: Reduced form, OLS, and IV estimates of the relationship between either the wage or the marginal revenue product of labor and the intermediate input wedge

	Wage			Marginal revenue product of labor		
	Reduced form	IV	OLS	Reduced form	IV	OLS
Intermediate input wedge		0.677	1.712		-0.006	0.170
		(0.170)***	(0.111)***		(0.167)	(0.080)**
$Z^{E/R}$	-0.101			0.001		
	(0.023)***			(0.025)		
Firm and year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.06	0.22	0.31	0.08	0.08	0.10
F-stat of instrument		485.01			485.01	
Observations	10,630	10,630	10,630	10,630	10,630	10,630

Notes: Results from regressions of the wage on the exchange rate instrument (Column 1), the intermediate input wedge instrumented with the exchange rate instrument (Column 2), and the intermediate input wedge (Column 3), as well as regressions of the marginal revenue product of labor on the exchange rate instrument (Column 4), the intermediate input wedge instrumented with the exchange rate instrument (Column 5), and the intermediate input wedge (Column 6). $Z^{E/R}$ is defined in equation (11). All regressions include lagged inputs and revenue as controls. All variables in logs except the instrument. Standard errors clustered at the firm level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

els identical, although the first stage weakens somewhat, with F-stats of 219.91 and 74.50, respectively. In line with sample composition not responding to the identifying variation, Column 1 of Table E10 in Appendix E shows that import status itself is unaffected by the realization of the instrument.

Table E7 in Appendix E verifies that the negative effect of the intermediate input wedge on the labor wedge, and the positive effect on the wage, are robust to all alternative approaches to obtaining revenue elasticities described in Section 2.2. The labor wedge estimates range from -0.47 with the CRS-constrained revenue function to -1.23 with calibrated revenue elasticities, the wage estimates from 0.57 to 0.91, and the instrument is strong throughout, with first-stage F-stats between 73 and 634. As in the baseline, the marginal revenue product of labor does not respond to the intermediate input wedge in any specification, with the exception of a marginally significant negative estimate under calibrated elasticities. This specification ignores variation in revenue elasticities entirely, so that they cannot adjust to movements in revenue shares.

I next probe the identifying assumptions directly. While Requirements 1 and 2 do not re-

quire initial Chinese import shares to be randomly assigned, the design would be undermined if firms that eventually rely more heavily on Chinese imports were already on differential paths before their first observed trade year. In that case, the initial share would proxy for pre-existing changes in input market conditions or labor market outcomes rather than exposure to subsequent exchange rate innovations. Table E8 in Appendix E tests for differential pre-trends by regressing the two input wedges, the wage, and the marginal revenue product of labor on interactions between the initial import share and indicators for two and three years before the firm's first observed trade year, conditional on firm and first-trade-cohort-by-year fixed effects. Whether or not lagged controls are included, none of the interactions is statistically significant, and joint tests do not reject the absence of differential pre-exposure trends for any of the four outcomes.

The second exercise assesses timing. The random walk in equation (12) implies that the exchange rate change between years t and $t+1$ is unknown in year t, so that exposure to this future change should leave current outcomes unaffected – a response would indicate anticipated exchange rate movements or misspecified annual timing. Table E9 in Appendix E adds $S_{io}^{import}(e_{t+1} - e_t)$ to the reduced-form regressions while retaining the contemporaneous instrument, separating exposure to the subsequent innovation from exposure to the current exchange rate realization. Future exposure predicts neither the labor wedge, nor the wage, nor the marginal revenue product of labor, while the coefficients on the contemporaneous instrument closely mirror baseline estimates. Together, the two exercises support interpreting the estimates in this section as responses to unanticipated exchange rate innovations.

4.3 Mechanisms

The IV estimates establish that exchange-rate-induced increases in the intermediate input wedge raise wages relative to workers' marginal revenue product, but leave open how firms and workers arrive at this outcome. This subsection therefore examines the margins of adjustment: intermediates sourcing (Panel a of Table 6), employment (Panel b), and pay (Panel c). For the latter two, I match employee-level data to the firm panel, recording worker flows, contract and job characteristics, and a decomposition of labor costs into base and extra pay. Appendix C describes these data in more detail. The responses documented here also inform the discussion of the exclusion restriction in Section 4.4.

Panel a of Table 6 documents how firms adjust their intermediate input sourcing. Column 1 shows that, as importing from China becomes cheaper, total intermediate input expenditure falls – recall that the value of the instrument declines as the euro appreciates. This drop conceals a pronounced reallocation towards Chinese imports: the value of intermediates imported from China and the number of imported varieties both increase (Columns 2 and 3). To gauge magnitudes, consider again a firm that moves from the 95th to the 5th percentile of the distribution of the instrument, an absolute difference of 0.616. Such a firm reduces

Table 6: Reduced form regressions on mechanisms

	Panel a: Intermediate input adjustments					
	$P_{it}^{M}M_{it}$	$P_{it}^{import}M_{it}^{import}$	V_{it}^{import}	$VNew_{it}^{import}/V_{it}^{import}$	$PNew_{it}^{import}$	$PCont_{it}^{import}$
$Z^{E/R}$	0.205	-1.706	-0.325	0.379	0.204	0.045
	(0.071)***	(0.260)***	(0.056)***	(0.147)**	(0.102)**	(0.346)
R^2	0.24	0.03	0.02	0.03	0.01	0.02
Observations	10,630	10,630	10,630	8,461	8,459	8,476
	Panel b: Employment adjustments					
	L_{it}/HC_{it}	HC_{it}	HC_{it}^{exit}/HC_{it}	$HC_{it}^{exit+enter}/HC_{it}$	L_{it+1}	L_{it+2}
$Z^{E/R}$	-0.010	-0.006	0.064	0.078	-0.017	-0.002
	(0.014)	(0.026)	(0.017)***	(0.039)**	(0.017)	(0.023)
R^2	0.01	0.38	0.04	0.03	0.19	0.09
Observations	10,628	10,628	10,628	10,628	8,319	6,465
	Panel c: Wage and rent adjustments					
	$W_{it}L_{it}/HC_{it}$	W_{it}	W_{it}^{base}	W_{it}^{extra}	R_{it}/L_{it}	VA_{it}/L_{it}
$Z^{E/R}$	-0.111	-0.132	-0.008	-0.174	0.043	-0.100
	(0.030)***	(0.041)***	(0.015)	(0.079)**	(0.047)	(0.028)***
$Z^{E/R} * \mathbb{I}_{it}^{appreciate}$		0.005				
		(0.003)*				
R^2	0.29	0.06	0.06	0.03	0.08	0.10
Observations	10,630	10,630	10,628	10,335	10,630	10,627

Notes: Results from regressions of several outcomes on the exchange rate instrument. $Z^{E/R}$ is defined in equation (11). $\mathbb{I}_{it}^{appreciate}$ = indicator equal to one if the exchange rate appreciated between year $t-1$ and year t. All other variables are defined in Appendix C. All regressions include lagged inputs and revenue as controls and firm and year fixed effects. All continuous variables in logs except the instrument. Standard errors clustered at the firm level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

its total intermediates expenditure by 11.86 percent, while the value of its Chinese imports nearly triples and the number of varieties it imports from China rises by 22.17 percent. As the median firm in the sample sources only 1.2 percent of its intermediates from China (Table E5), a large proportional expansion of Chinese imports is consistent with a moderate decline in overall intermediates expenditure.[26]

[26]The absence of an MRPL response (Table 5) is informative about unobserved intermediates quantities.

Columns 4 to 6 clarify along which margins this import expansion occurs. As importing becomes cheaper, the share of newly imported varieties in all imported varieties falls by 20.82 percent for the same 95th-to-5th percentile firm. Combining this estimate with the increase in the total number of varieties implies that the number of newly added varieties changes little, so that the expansion in the variety count is driven by continued varieties. Firms, therefore, grow their import portfolio primarily by continuing to import varieties that they would have dropped had the exchange rate remained constant, rather than by adding varieties they did not previously import. In line with the exchange rate determining which varieties remain worthwhile, the price index of newly imported varieties drops by 11.81 percent (Column 5), whereas the price index of continued varieties falls by only 2.73 percent and not significantly so (Column 6). Taken together, Panel a shows that as the euro appreciates, firms substitute towards Chinese intermediates mainly by maintaining existing import varieties for longer, and that total intermediates expenditure falls even as reliance on imports deepens.

Panel b of Table 6 turns to employment adjustments. Recall that L_{it} measures employment in full-time equivalent units based on hours worked, so that its logarithm decomposes into full-time equivalents per worker and employee headcount: $\ln L_{it} = \ln(L_{it}/HC_{it}) + \ln(HC_{it})$. Columns 1 and 2 report both margins. As importing from China becomes cheaper, hours per worker and headcount both increase, but the effects are small and neither is statistically significant (nor is the increase in L_{it}, omitted for brevity). For the 95th-to-5th percentile firm, hours per worker rise by 0.62 percent and headcount by 0.37 percent, so that total employment increases by roughly 1 percent, with hours worked accounting for about two thirds of the increase.

Behind this modest expansion lies a decline in worker turnover. Column 4 shows that churn – the sum of newly hired workers and workers that leave the firm by the next year, relative to headcount – falls by 4.69 percent for the same firm. This decline operates through the exit margin: the share of workers that no longer work at the firm in the next year falls by 3.87 percent (Column 3). At the mean exit rate of 10.42 percent, the reduction in exits corresponds to 0.4 percentage points. Because exits are defined using employment in year $t+1$, their decline shows that greater retention contributes to the persistence of the modest employment response into the following year. Indeed, the limited contemporaneous employment response is not a sign of adjustment frictions delaying hiring. Columns 5 and 6 relate employment in years $t+1$ and $t+2$ to the instrument realization in year t. The coefficient on next year's employment is quantitatively very similar to the implied contemporaneous effect while the coefficient on employment two years ahead is close to zero, although the sample shrinks with every lead as later years drop out.

Under the translog revenue function, the elasticity of the marginal revenue product of labor with respect to intermediate inputs is $\theta^M_{it} + \beta_{lm}/\theta^L_{it}$. It equals 0.40 at the median of my sample, so that a substantial expansion of intermediates quantities would have raised the marginal revenue product of labor.

Relatedly, Table E10 in Appendix E suggests that firms do not adjust the composition of their workforce either. The shares of workers with full-time contracts, permanent contracts, managerial jobs, regular jobs, and other job types – such as interns and on-call workers – as well as average tenure are all unaffected by the instrument. None of the point estimates is statistically distinguishable from zero and they are small: the largest, on the permanent contract share, implies a change of 3.19 percent for the 95th-to-5th percentile firm. These results help interpret the response of the labor wedge. Recall from Lemma 1 that the measured labor wedge is an expenditure-share-weighted average of the wedges of the underlying worker types, so it can change either because the composition of the workforce shifts or because the underlying wedges themselves change. The results provide no evidence that the labor-wedge response is mechanically driven by reweighting toward workers with different observable contract or job characteristics. Taken together, these results point to a modest employment expansion, driven by hours worked and worker retention, with a stable workforce whose wedges respond to the instrument.

Panel c of Table 6 turns to compensation. The wage – labor expenditure per FTE – can be decomposed into labor expenditure per employee and FTE per employee: $\ln W_{it} = \ln(W_{it}L_{it}/HC_{it}) - \ln(L_{it}/HC_{it})$. As Panel b showed that neither hours per worker nor headcount responds meaningfully to the instrument, the wage response documented in Table 5 must be driven by pay per worker rather than by changes in the workforce. Column 1 confirms this: as importing from China becomes cheaper, labor expenditure per employee rises, by 7.08 percent for the 95th-to-5th percentile firm. Column 2 asks whether this response is symmetric by interacting the instrument with an indicator for years in which the euro appreciates against the renminbi. It largely is: for the 95th-to-5th percentile firm, the implied wage response is 8.14 percent in appreciation years and 8.47 percent in depreciation years.[27] The economically relevant conclusion is that wages move both up and down with the cost of importing, and do so near-symmetrically, suggesting that firms adjust a component of pay that is flexible in both directions.

Columns 3 and 4 decompose the average wage additively into the average base wage and the average extra wage. The base wage W_{it}^{base} measures the contracted part of compensation per FTE worker: regular gross salary and pay for additional hours at the regular rate. The extra wage W_{it}^{extra} is the difference between total labor cost per FTE and the base wage per FTE.[28] The base wage does not respond to the instrument – the point estimate implies an increase of 0.49 percent for the 95th-to-5th percentile firm and is statistically indistin-

[27] This is not in line with downward wage rigidity, or the loss-sharing that Mertens et al. (2026) emphasize for energy price shocks in German manufacturing. This small difference should not be overinterpreted. As Figure D7 shows, appreciations tend to be smaller in magnitude than depreciations.

[28] It comprises overtime pay at above-regular rates, special cash payments such as holiday allowances, thirteenth-month payments, bonuses, and profit-sharing payments, wage in kind, employer-side social security contributions, employer-side pension expense including additions to pension provisions, and other personnel costs such as reimbursements, training expenditure, and employee benefits.

guishable from zero – while the extra wage increases by 11.31 percent. The wage response to the instrument therefore operates almost entirely through the extra wage. Although not separately observable, several components of the extra wage are natural margins for rent sharing, as they can be adjusted annually in both directions and targeted at specific groups of employees: bonuses and profit-sharing payments, discretionary employee benefits, and employer-specific pension flows such as voluntary additional contributions. Statutory employer social security contributions and regular pension premiums, by contrast, are largely determined at the sector or national level by formulaic rates and are therefore unlikely to drive a firm-specific response.

The final two columns of Panel c of Table 6 provide further evidence that these payments constitute intermediate input rents being shared with workers. As importing from China becomes cheaper, revenue per worker does not change – the point estimate of 2.68 percent for the 95th-to-5th percentile firm is statistically insignificant – while value added per worker increases by 6.35 percent. As the difference between revenue and value added is intermediate input expenditure, and Panel a showed that intermediates expenditure falls, the rise in value added per worker reflects cheaper inputs rather than more productive workers – in line with the earlier finding that the marginal revenue product of labor does not respond to the instrument (Table 5). These results also mirror the full-sample patterns of Section 3, where value added per worker, but not revenue per worker, was strongly associated with both input wedges. Workers, then, generate the same revenue on the margin and on average, but capture part of the additional rents their employer earns in the intermediate input market through the flexible component of their pay.

The reduced-form estimates can be combined into a measure of how the rents generated by exchange rate fluctuations are divided between workers and firms. Let β_W and β_{VA} denote the reduced-form coefficients from regressions of logged W_{it} and VA_{it}/L_{it} on the instrument. The workers' share of the total instrument-induced euro change in value added per worker is

$$I_W = \frac{\beta_W}{\beta_{VA}} \sum_i \omega_i \frac{W_i L_i}{VA_i}, \qquad \omega_i = \frac{(VA_i/L_i)\, S_{io}^{import}}{\sum_j (VA_j/L_j)\, S_{jo}^{import}},$$

where firm-level variables are measured in the year that the initial shares are constructed. The ratio β_W/β_{VA} is the proportional compensation response per unit of proportional value added response, the labor-compensation share $W_i L_i/VA_i$ converts proportions into euros, and the weights ω_i reflect that firms with higher baseline value added per worker and greater initial exposure to Chinese imports contribute more euros to the total change. I estimate $I_W = 0.67$: of every euro of value added per worker generated by the exchange rate movement, 67 cents are used by the firm to compensate workers. This worker share is large but not incomparable to causal estimates of worker incidence for stable workforces: Kline et al. (2019) report that incumbent workers capture roughly 50 to 60 cents of every dollar of patent-induced surplus on average.

4.4 Exclusion restriction and alternative interpretations

For the IV regressions to have a causal interpretation, the instrument should affect the labor wedge only through the intermediate input wedge (Requirement 3). A first class of potential violations operates through the output market. An appreciation of the euro lowers the price of Chinese final goods outside China and makes Dutch products more expensive in China. This would likely dampen demand and reduce output market rents of Dutch firms, which is difficult to reconcile with wage increases. Moreover, neither employment nor revenue per worker responds significantly to the instrument (Table 6), so that total revenue is essentially unchanged. Nor does the marginal revenue product of labor respond (Table 5). In addition, Chinese import competition operates largely at the industry level and the results of Tables 4 and 5 are robust to adding industry-by-year fixed effects, which absorb any industry-wide demand, import-competition, or price shocks (Column 4 of Table E6 in Appendix E). Finally, excluding all firms that export to China in the current or previous year shuts down the direct export demand channel. The point estimates are, if anything, somewhat larger, -0.73 for the labor wedge and 1.02 for the wage, though less precisely estimated as the sample shrinks by more than 40 percent (Column 3 of Table E6).

A second class of violations operates directly through the labor market. Exchange rate movements could shift local labor market conditions if exposed importers cluster spatially, so that the wages of exposed firms change with regional labor market tightness. Adding NUTS1-region-by-year fixed effects, which absorb regional labor market shocks, leaves the estimates unaffected: -0.67 for the labor wedge and 0.70 for the wage (Table E6 in Appendix E). Moreover, exposure to Chinese import competition has typically been associated with worse outcomes for workers in Western manufacturing (Autor et al., 2016), and access to cheap imported inputs with outsourcing or offshoring, weakening workers' bargaining position (Gu et al., 2022). Such forces are difficult to reconcile with a largely unchanged workforce and increased wages. Hours per worker and headcount increase slightly and insignificantly (Table 6), and the shares of full-time, permanent, managerial, regular, and other jobs, as well as average tenure and the marginal revenue product of labor, are all unaffected (Tables E10 in Appendix E and 5). The margin of adjustment points in the same direction. Base wages, which are largely governed by sectoral collective agreements, do not respond, while the response operates through increased firm-level discretionary extra pay.

Finally, if variation in the instrument coincides with factor-biased changes in technology, the IV estimates could reflect technology rather than rents, as discussed in Section 3.3 (Raval, 2023). First, note that wedges based on 2-digit-industry-by-2-year revenue functions, which absorb factor-biased technical change at that level, deliver a labor wedge IV estimate of -0.98 with a wage response of similar magnitude to the baseline (Table E7 in Appendix E). The reduced-form results jointly discipline what any remaining firm-specific technology explanation must look like. It must be triggered by unanticipated exchange rate

innovations – firms are not on differential trends prior to importing and future exchange rate changes predict no outcomes (Tables E8 and E9 in Appendix E) – and it must reverse near-symmetrically when the exchange rate reverses. It must operate through varieties the firm keeps importing rather than through newly added ones, whereas the trade literature ties productivity gains from imported inputs precisely to the addition of new varieties and quality upgrading (Halpern et al., 2015). And it must raise wages while leaving revenue per worker, the marginal revenue product of labor, employment, and workforce composition largely unchanged (Table 6). A technology shock satisfying all of these restrictions would be observationally and economically close to a transitory change in firm rents.

Taken together, the evidence is most consistent with firms sharing rents generated in the intermediate input market with their workers: as the euro appreciates, intermediates markdowns deepen, value added per worker rises even though workers generate no more revenue on average or on the margin, and workers capture a substantial share of the additional surplus through the flexible component of their pay. This rent-sharing channel constitutes a novel explanation for the joint distribution of labor and intermediate input wedges documented in Section 3, and implies that input market imperfections, as measured by input wedges, are linked across input markets.

5 Concluding remarks

This paper studies how imperfections in labor and intermediate input markets interact, using Dutch manufacturing data from 2007 to 2018. Estimating both input wedges without normalizing either to unity, I find that the two wedges are strongly negatively correlated. Moreover, intermediates price markdowns are the norm, consistent with widespread buyer power, while wage markups and wage markdowns are equally common in the labor market. A model in which firms face upward-sloping supply curves in both input markets and bargain over wages with their collectively organized workers can rationalize these findings. Quasi-experimental evidence supports that the rent-sharing mechanism connects input market imperfections. Using exchange rate fluctuations interacted with firms' predetermined reliance on Chinese intermediates, I show that as the euro appreciates, exposed firms deepen their intermediates markdowns and raise wages, while employment, workforce composition, and workers' average and marginal revenue products are unchanged. Rents generated in one input market are thus partly passed on to suppliers in another.

The main implication of this paper is that once one departs from competitive input markets, market imperfections in multiple markets can be relevant to understand price formation in a single input market. Hence, approaches that estimate market imperfections in one input market by ruling them out in another are to be treated with care. A second implication is that widespread monopsony concerns might not be warranted in labor markets where collective bargaining plays an important role. Finally, reducing market power in one input market

might hurt input suppliers in other input markets, so that purchasing power in one market can act as countervailing power for input suppliers in a different input market.

While I make progress in studying input market imperfections and their interaction, a lot is still to be learned. In particular, as intermediate input prices are unobserved and input suppliers are often foreign firms, uncovering the origins of intermediate input price markdowns in Dutch manufacturing is challenging. Detailed industry- or firm-level studies using data on supplier networks and input prices would allow researchers to impose more structure, aiding inference on the origins of buyer power. I leave such considerations for future work.

References

Ackerberg, D. A., Caves, K., and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451.

Ackerberg, D. A. and De Loecker, J. (2024). Production function identification under imperfect competition. CEPR Discussion Paper DP19640. Technical report.

Alviarez, V. I., Fioretti, M., Kikkawa, K., and Morlacco, M. (2023). Two-sided market power in firm-to-firm trade (no. w31253) National Bureau of Economic Research.

Angerhofer, T., Collard-Wexler, A., and Weinberg, M. (2026). Oligopsony and collective bargaining. *mimeo*.

Angrist, J. D., Graddy, K., and Imbens, G. W. (2000). The interpretation of instrumental variables estimators in simultaneous equations models with an application to the demand for fish. *Review of Economic Studies*, 67(3):499–527.

Atalay, E. (2014). Materials prices and productivity. *Journal of the European Economic Association*, 12(3):575–611.

Atkeson, A. and Burstein, A. (2008). Pricing-to-market, trade costs, and international relative prices. *American Economic Review*, 98(5):1998–2031.

Autor, D. H., Dorn, D., and Hanson, G. H. (2016). The China shock: Learning from labor-market adjustment to large changes in trade. *Annual Review of Economics*, 8:205–240.

Avignon, R., Chambolle, C., Guigue, E., and Molina, H. (2024). Markups, markdowns, and bargaining in a vertical supply chain. *mimeo*.

Avignon, R. and Guigue, E. (2022). Markups and markdowns in the French dairy market. *mimeo*.

Barrette, E., Gowrisankaran, G., and Town, R. (2022). Countervailing market power and hospital competition. *Review of Economics and Statistics*, 104(6):1351–1360.

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

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25(2):242–262.

Björnerstedt, J. and Verboven, F. (2016). Does merger simulation work? Evidence from the Swedish analgesics market. *American Economic Journal: Applied Economics*, 8(3):125–164.

Blundell, R. and Bond, S. (2000). GMM estimation with persistent panel data: An application to production functions. *Econometric Reviews*, 19(3):321–340.

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

Borusyak, K. and Hull, P. (2023). Nonrandom exposure to exogenous shocks. *Econometrica*, 91(6):2155–2185.

Brooks, W. J., Kaboski, J. P., Li, Y. A., and Qian, W. (2021). Exploitation of labor? Classical monopsony power and labor's share. *Journal of Development Economics*, 150:102627.

Card, D. (2022). Who set your wage? *American Economic Review*, 112(4):1075–1090.

Card, D. and Cardoso, A. R. (2022). Wage flexibility under sectoral bargaining. *Journal of the European Economic Association*, 20(5):2013–2061.

Card, D., Cardoso, A. R., Heining, J., and Kline, P. (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(S1):S13–S70.

Card, D., Devicienti, F., and Maida, A. (2014). Rent-sharing, holdup, and wages: Evidence from matched panel data. *Review of Economic Studies*, 81(1):84–111.

Caselli, M., Nesta, L., and Schiavo, S. (2021). Imports and labour market imperfections: Firm-level evidence from France. *European Economic Review*, 131:103632.

CBS (2020a). De arbeidsmarkt in cijfers 2019. Statistical Offices Netherlands.

CBS (2020b). Documentatie statistiek financiën van niet-financiële ondernemingen. Statistical Offices Netherlands.

CBS (2020c). Documentatierapport algemeen bedrijven register 2019v1. Statistical Offices Netherlands.

De Loecker, J. (2011). Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity. *Econometrica*, 79(5):1407–1451.

De Loecker, J. and Goldberg, P. K. (2014). Firm performance in a global market. *Annual Review of Economics*, 6:201–227.

De Loecker, J., Goldberg, P. K., Khandelwal, A. K., and Pavcnik, N. (2016). Prices, markups, and trade reform. *Econometrica*, 84(2):445–510.

De Loecker, J. and Syverson, C. (2022). An industrial organization perspective on productivity. In Ho, K., Hortacsu, A., and Lizzeri, A., editors, *Handbook of Industrial Organization, Volume 4*. Elsevier.

De Loecker, J. and Warzynski, F. (2012). Markups and firm-level export status. *American Economic Review*, 102(6):2437–2471.

Deb, S., Eeckhout, J., Patel, A., and Warren, L. (2024). Walras–Bowley lecture: Market power and wage inequality. *Econometrica*, 92(3):603–636.

Demirer, M. (2025). Production function estimation with factor-augmenting technology: An application to markups. *mimeo*.

Demirer, M. and Rubens, M. (2025). Welfare effects of buyer and seller power. (no. w33371). National Bureau of Economic Research.

Dobbelaere, S., Fuss, C., and Vancauteren, M. (2026). Offshoring and labor market power: Comparing Belgian and Dutch firms. *Industrial Relations: A Journal of Economy and Society*, 65(1):80–99.

Dobbelaere, S., Hirsch, B., Müller, S., and Neuschaeffer, G. (2024). Organised labour, labour market imperfections, and employer wage premia. *ILR Review*, 77(3):396–427.

Dobbelaere, S. and Mairesse, J. (2013). Panel data estimates of the production function and product and labor market imperfections. *Journal of Applied Econometrics*, 28(1):1–46.

Doraszelski, U. and Jaumandreu, J. (2018). Measuring the bias of technological change. *Journal of Political Economy*, 126(3):1027–1084.

Edmond, C., Midrigan, V., and Xu, D. Y. (2023). How costly are markups? *Journal of Political Economy*, 131(7):1619–1675.

Flynn, Z., Traina, J., and Gandhi, A. (2019). Measuring markups with production data. *mimeo*.

Galbraith, J. K. (1952). *American Capitalism: The Concept of Counterveiling Power*. Houghton Mifflin.

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

Gopinath, G., Kalemli-Özcan, Ş., Karabarbounis, L., and Villegas-Sanchez, C. (2017). Capital allocation and productivity in South Europe. *Quarterly Journal of Economics*, 132(4):1915–1967.

Grieco, P. L., Li, S., and Zhang, H. (2016). Production function estimation with unobserved input price dispersion. *International Economic Review*, 57(2):665–690.

Gu, G., Malik, S., Pozzoli, D., and Rocha, V. (2022). Chinese import competition, offshoring and servitization. *Economic Inquiry*, 60(2):901–928.

Hall, R. E. (1988). The relation between price and marginal cost in US industry. *Journal of Political Economy*, 96(5):921–947.

Halpern, L., Koren, M., and Szeidl, A. (2015). Imported inputs and productivity. *American Economic Review*, 105(12):3660–3703.

Haltiwanger, J., Jarmin, R. S., and Miranda, J. (2013). Who creates jobs? Small versus large versus young. *Review of Economics and Statistics*, 95(2):347–361.

Hashemi, A., Kirov, I., and Traina, J. (2022). The production approach to markup estimation often measures input distortions. *Economics Letters*, 217:110673.

Hijzen, A., Martins, P. S., and Parlevliet, J. (2019). Frontal assault versus incremental change: A comparison of collective bargaining in Portugal and the Netherlands. *IZA Journal of Labor Policy*, 9(8):1–26.

Kirov, I. and Traina, J. (2023). Labor market power and technological change in US manufacturing. *mimeo*.

Klette, T. J. and Griliches, Z. (1996). The inconsistency of common scale estimators when output prices are unobserved and endogenous. *Journal of Applied Econometrics*, 11(4):343–361.

Kline, P., Petkova, N., Williams, H., and Zidar, O. (2019). Who profits from patents? Rent-sharing at innovative firms. *Quarterly Journal of Economics*, 134(3):1343–1404.

Kroft, K., Luo, Y., Mogstad, M., and Setzler, B. (2025). Imperfect competition and rents in labor and product markets: The case of the construction industry. *American Economic Review*, 115(9):2926–2969.

Lamadon, T., Mogstad, M., and Setzler, B. (2022). Imperfect competition, compensating differentials and rent sharing in the US labor market. *American Economic Review*, 112(1):169–212.

Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies*, 70(2):317–341.

Mertens, M. (2020). Labor market power and the distorting effects of international trade. *International Journal of Industrial Organization*, 68:102562.

Mertens, M. (2022). Micro-mechanisms behind declining labor shares: Rising market power and changing modes of production. *International Journal of Industrial Organization*, 81:102808.

Mertens, M. (2023). Labour market power and between-firm wage (in)equality. *International Journal of Industrial Organization*, 91:103005.

Mertens, M. and Mottironi, B. (2025). Do larger firms exert more market power? Markups and markdowns along the size distribution. *Economics Letters*, page 112500.

Mertens, M., Mueller, S., and Neuschaeffer, G. (2026). Identifying rent-sharing using firms' energy input mix. *Journal of the European Economic Association*, page jvag006.

Morlacco, M. (2020). Market power in input markets: Theory and evidence from French manufacturing. *mimeo*.

Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297.

Petrin, A. and Sivadasan, J. (2013). Estimating lost output from allocative inefficiency, with an application to Chile and firing costs. *Review of Economics and Statistics*, 95(1):286–301.

Pham, H. (2023). Trade reform, oligopsony, and labor market distortion: Theory and evidence. *Journal of International Economics*, 144:103787.

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

Rossi, B. (2013). Exchange rate predictability. *Journal of Economic Literature*, 51(4):1063–1119.

Rubens, M. (2023). Market structure, oligopsony power, and productivity. *American Economic Review*, 113(9):2382–2410.

Rubens, M., Wu, Y., and Xu, M. (2026). Exploiting or augmenting labor. *American Economic Review: Insights*, 8(1):72–89.

Syverson, C. (2025). Markups and markdowns. *Annual Review of Economics*, 17:57–76.

Wong, H. C. (2026). Understanding high-wage firms: Monopoly, monopsony, and bargaining power. *American Economic Review*, 116(7):2504–2541.

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

Appendices

A Theory appendix

A.1 Proof of Lemma 1

Proof Let the input wedge of intermediates variety X – suppressing the it subscripts – be denoted by $\gamma^X \equiv \frac{MRPX}{P^X}$ and its revenue elasticity by θ^X. Then,

$$\frac{P^X X}{R}\gamma^X = \theta^X \quad \iff \quad \frac{P^X X}{R}\gamma^X = \theta^M \frac{\partial M}{\partial X}\frac{X}{M},$$

as $\theta^X = \frac{\partial R}{\partial M}\frac{\partial M}{\partial X}\frac{X}{R}\frac{M}{M}$. Summing over all varieties in set $\mathcal{M}$ gives

$$\frac{P^M M}{R}\sum_{X\in\mathcal{M}}\frac{P^X X}{P^M M}\gamma^X = \theta^M \sum_{X\in\mathcal{M}}\frac{\partial M}{\partial X}\frac{X}{M} \iff \frac{\theta^M}{MS} = \sum_{X\in\mathcal{M}}\frac{P^X X}{P^M M}\gamma^X,$$

where the final equivalence holds as the elasticities of the varieties sum to one by the properties of the aggregator function $A^M(\cdot)$. γ^L can be derived similarly. ■

A.2 Demand, conduct, and the revenue function

This appendix characterizes the combinations of demand systems and conduct assumptions in output markets under which an estimating equation of the form of equation (5) exists. Throughout, I maintain the Hicks-neutral technology $q_{it} = \tilde{f}_s(k_{it}, l_{it}, m_{it}; \tilde{\beta}) + \tilde{\omega}_{it}$. As elsewhere, lowercase letters denote logs of concomitant uppercase letters.

Since $r_{it} = p_{it} + q_{it}$ and realized transactions lie on the inverse demand curve each firm faces, write realized log prices as $p_{it} = \mathcal{P}_s(q_{it}, \mathbf{A}_{c(i)t}) + \xi_{it}$, where ξ_{it} is a demand shock and $\mathbf{A}_{ct}$ is a vector of aggregates – price and quantity indices, inclusive values, outside shares – that take a common value for all firms in cell c at time t. The aggregates summarize rivals' actions and market conditions and may themselves be equilibrium objects that firm i influences.[29] The following assumption contains the required restrictions.

Assumption A1 Realized transactions lie on a conditional inverse demand curve of the form

$$p_{it} = a_{c(i)t} + \frac{1}{\eta_s}\, q_{it} + \xi_{it}, \tag{A1}$$

[29]Fixed effects in equation (5) absorb the realized value of any cell-common term regardless, and the conditions below can be read as holding at realized equilibria. Whether $\mathbf{A}_{ct}$ is exogenous to the firm matters for the moment conditions used in estimation, not for the existence of the estimating equation.

where a_{ct} is common to all firms in cell c at time t and the slope $\frac{1}{\eta_s}$ is constant at the level of estimation.

Proposition A1 *Suppose realized log prices can be written as $p_{it} = \mathcal{P}_s(q_{it}, \mathbf{A}_{c(i)t}) + \xi_{it}$ and that, holding $(\tilde{\omega}_{it}, \mathbf{A}_{ct}, \xi_{it})$ fixed, inputs trace out an interval of q_{it}. (i) Under Assumption A1 and a translog production function with a Hicks-neutral productivity shock, log revenue satisfies equation (5) with $\theta^x_{it} = \frac{\eta_s+1}{\eta_s}\tilde{\theta}^x_{it}$ for $x \in \{k, l, m\}$ and $\omega_{it} = \frac{\eta_s+1}{\eta_s}\tilde{\omega}_{it} + \tilde{\xi}_{it}$, where $\tilde{\xi}_{it}$ is the firm-specific component of the decomposition $\xi_{it} = \xi_{st} + \xi_{g(i)} + \tilde{\xi}_{it}$ – the log of the decomposition of Ξ_{it} in equation (4). (ii) Conversely, if log revenue can be written as $\frac{\eta_s+1}{\eta_s}\tilde{f}_s(\cdot)$ plus terms spanned by the included controls plus a scalar firm-time unobservable that does not vary with inputs conditional on $(\tilde{\omega}_{it}, \mathbf{A}_{ct}, \xi_{it})$, then Assumption A1 holds on the support.*

Proof (i) Substituting demand and production function into $r_{it} = p_{it} + q_{it}$ gives $r_{it} = \frac{\eta_s+1}{\eta_s}\tilde{f}_s(\cdot) + a_{c(i)t} + \frac{\eta_s+1}{\eta_s}\tilde{\omega}_{it} + \xi_{it}$. Decomposing ξ_{it} as in equation (4), the industry-year and 3-digit components are absorbed by γ_{st} and the included industry indicators, leaving $\tilde{\xi}_{it}$ in the composite ω_{it}. (ii) Fix $(\tilde{\omega}_{it}, \mathbf{A}_{ct}, \xi_{it})$. Then $r_{it} - q_{it} = \mathcal{P}_s(q_{it}, \mathbf{A}_{ct}) + \xi_{it}$ must equal $\frac{1}{\eta_s}q_{it}$ plus terms that do not vary with inputs, over an interval of q_{it}. Differentiating, $\frac{\partial \mathcal{P}_s}{\partial q_{it}} = \frac{1}{\eta_s}$ must be constant in q_{it}, in $\mathbf{A}_{ct}$, and across firms within the level of estimation. Integrating gives $\mathcal{P}_s(q_{it}, \mathbf{A}_{ct}) = \frac{1}{\eta_s}q_{it} + a_s(\mathbf{A}_{ct})$, and the requirement that the intercept consist of cell-common terms plus a log-additive scalar shifter is Assumption A1. ■

Note that the binding restriction concerns the slope: demand shocks may not enter the slope of conditional inverse demand, and firm-specific demand conditions must collapse to a scalar. Assumption A1 rules out variable elasticities along the conditional demand curve but it does not rule out variable markups, which can arise under oligopoly through the internalization of the aggregates. Demand system (4) satisfies Assumption A1 with $a_{st} = \ln P_{st} - \frac{1}{\eta_s}q_{st} + \ln \Xi_{st}$, which delivers equation (5). The remainder of this appendix discusses which other common demand specifications and conduct assumptions satisfy Assumption A1, and hence deliver an estimating equation of the same form as equation (5).

CES and nested CES demand Equation (4) is a property of preferences, not of equilibrium: it is the inverse demand curve itself, so it holds at realized transactions whether the firm is atomistic or an oligopolist – rivals matter only through the cell-common aggregates, whose realized values conduct determines but the fixed effects absorb. Note also that, under CES, taste and quality shifters are observationally equivalent for the existence of the revenue function: if quality Ξ^q_{it} scales quantity, $P_{it} = P_{st}\big(\frac{\Xi^q_{it}Q_{it}}{Q_{st}}\big)^{\frac{1}{\eta_s}}\Xi_{it}$, then $\frac{1}{\eta_s}\ln \Xi^q_{it} + \ln \Xi_{it}$ is again log-additive, so the two need not be distinguished.

Consider next a nested CES structure with within-nest elasticity of substitution σ_s and across-nest elasticity ϱ_s, $\sigma_s > \varrho_s > 1$: $Q_{it} = \big(\frac{P_{it}}{P_{gt}}\big)^{-\sigma_s}\big(\frac{P_{gt}}{P_{st}}\big)^{-\varrho_s}Q_{st}\,\Xi^{\sigma_s}_{it}$, where P_{gt} is the nest-g

price index. Inverting,

$$p_{it} = -\frac{1}{\sigma_s}\, q_{it} + \underbrace{\Big[\ln P_{gt} - \frac{\varrho_s}{\sigma_s}\big(\ln P_{gt} - \ln P_{st}\big) + \frac{1}{\sigma_s} q_{st}\Big]}_{\equiv\, a_{gt}} + \xi_{it}, \tag{A2}$$

so Assumption A1 holds with $\frac{1}{\eta_s} = -\frac{1}{\sigma_s}$ and a nest-by-time intercept or the nest aggregates included as observables. This is the market structure of Atkeson and Burstein (2008) on which much of the recent literature on oligopoly and market power builds (e.g., Edmond et al. (2023)).

Logit, nested logit, and random coefficients logit demand Let indirect utility contain $-\alpha_s \ln P_{it}$, $\alpha_s > 1$. This is the "constant expenditures" specification of empirical demand estimation, which has a (quasi-)constant price elasticity (Björnerstedt and Verboven, 2016). Let N_t denote market size and S_{0t} the outside good's share (capital S distinguishes choice shares from the industry index s). The inversion of Berry (1994), $\ln S_{it} - \ln S_{0t} = x_i'\beta - \alpha_s \ln P_{it} + \xi_{it}$ with $\ln S_{it} = q_{it} - \ln N_t$, with x_i denoting product characteristics, gives

$$r_{it} = \Big(1 - \frac{1}{\alpha_s}\Big) q_{it} + \frac{1}{\alpha_s}\big(\ln N_t + \ln S_{0t}\big) + \frac{x_i'\beta + \xi_{it}}{\alpha_s}, \tag{A3}$$

which satisfies Assumption A1 with $\frac{1}{\eta_s} = -\frac{1}{\alpha_s}$: market size and the outside share are market-time common, the mean-utility shock enters log-additively scaled by $\frac{1}{\alpha_s}$, and any persistent characteristics component folds into $\tilde{\Xi}_{it}$. With a nesting parameter $\sigma_N \in [0, 1)$ and within-nest share $S_{i|g,t}$, the inversion is $\ln S_{it} - \ln S_{0t} = x_i'\beta - \alpha_s \ln P_{it} + \xi_{it} + \sigma_N \ln S_{i|g,t}$ with $\ln S_{i|g,t} = q_{it} - q_{gt}$, so that

$$r_{it} = \Big[1 - \frac{1 - \sigma_N}{\alpha_s}\Big] q_{it} - \frac{\sigma_N}{\alpha_s}\, q_{gt} + \frac{1}{\alpha_s}\big(\ln N_t + \ln S_{0t}\big) + \frac{x_i'\beta + \xi_{it}}{\alpha_s}, \tag{A4}$$

which satisfies Assumption A1 with $\frac{1}{\eta_s} = -\frac{1-\sigma_N}{\alpha_s}$, provided the nest aggregate q_{gt} is spanned by nest-by-time controls or included as an observable regressor. Both equations are share inversions evaluated at realized transactions, conditional on the outside share and the nest aggregate. These are equilibrium objects that respond to a non-negligible firm's own actions. For the existence of the estimating equation this is immaterial, as these aggregates need only be cell-common and spanned by the controls.

The log-price specification is essential. With prices in levels, indirect utility contains $-\alpha_s P_{it}$ and the same inversion gives

$$P_{it} = \frac{1}{\alpha_s}\big[x_i'\beta + \xi_{it} + \ln N_t + \ln S_{0t} - q_{it}\big] \quad \Rightarrow \quad r_{it} = q_{it} + \ln P_{it}, \tag{A5}$$

so $\frac{\partial p_{it}}{\partial q_{it}} = -\frac{1}{\alpha_s P_{it}}$ varies with the price level. The demand elasticity depends on the price level and no fixed-effect structure results in Assumption A1 being satisfied. Note that Lemma

1 itself is unaffected as it requires only Assumptions 1 and 2, so the ratio of an input's revenue elasticity to its revenue share equals the (average) input wedge under any differentiable revenue function, including the one generated by (A5). What fails is the recovery of revenue elasticities from financial statements: under (A5), $\theta^x_{it} = \big(1 - \frac{1}{\alpha_s P_{it}}\big)\tilde{\theta}^x_{it}$ depends on the unobserved price level, so the elasticities cannot be estimated without price or quantity data. Random-coefficients logit, $S_{it} = \int S_{it}(\nu)\, dF(\nu)$ with $S_{it}(\nu)$ the choice share of consumer type ν, also fails to meet Assumption A1. The market-level inversion of Berry (1994) recovers mean utilities as a function of the entire vector of shares and characteristics. Rivals' identities matter, and the objects that summarize rivals are type-specific inclusive values so generic random-coefficients specifications do not admit a fixed, low-dimensional, cell-common aggregate conditional on which inverse demand is log-linear in own quantity. Even with a common coefficient α_s on log price, $\frac{\partial S_{it}(\nu)}{\partial \ln P_{it}} = -\alpha_s S_{it}(\nu)\big(1 - S_{it}(\nu)\big)$ implies

$$\frac{\partial \ln S_{it}}{\partial \ln P_{it}} = -\alpha_s \Big[1 - \frac{\int S_{it}(\nu)^2\, dF(\nu)}{\int S_{it}(\nu)\, dF(\nu)}\Big], \tag{A6}$$

where the bracketed term varies with the firm's own price, with rivals, and with the distribution of consumer types. Unlike in the homogeneous-consumer logit, conditioning on aggregate outside shares does not restore a constant slope, because the relevant conditioning object differs across consumer types. Again, with richer data, demand could be estimated and the revenue approach of Lemma 1 could nonetheless be applied.

Conduct Given a demand system satisfying Assumption A1, the estimating equation is invariant to conduct in the output market. Equation (5) is derived from the identity $r_{it} = p_{it} + q_{it}$ evaluated at realized transactions. Conduct determines which point on the conditional demand curve is realized, not the curve itself. Monopolistic competition (in which firms take $\mathbf{A}_{ct}$ as given), differentiated-products Cournot and Bertrand competition, conduct-parameter models, and collusion (in which $\mathbf{A}_{ct}$ is an equilibrium object that firms influence) all deliver the same estimating equation, because a fixed effect absorbs the realized value of a cell-common term regardless of how it is determined.

While the estimating equation is conduct-invariant, the interpretation of its coefficients is not in general. There are two notions of "holding all else fixed" when differentiating revenue with respect to an input. The *conditional* revenue elasticity, $\theta^x_{it} = \frac{\eta_s + 1}{\eta_s}\tilde{\theta}^x_{it}$, holds the aggregates $\mathbf{A}_{ct}$ fixed. It is the object identified by equation (5), whose fixed effects absorb the $\mathbf{A}_{ct}$, and hence the elasticity entering Lemma 1. The *decision-relevant* revenue elasticity holds rivals' strategies fixed but recognizes that the aggregates contain firm i's own decision. When a non-atomistic firm changes its behavior, the indices it is part of move mechanically, without any rival changing its behavior. A firm internalizes this, so its input choices are governed by the total derivative of revenue. Under monopolistic competition the two elasticities coincide, but under oligopoly they differ by a factor that market shares

summarize. Let ms_{it} denote firm i's revenue share within its nest. Under the nested CES demand of (A2), the elasticity of demand perceived by the firm is

$$\frac{1}{\varepsilon_{it}^{D}} = \frac{ms_{it}}{\varrho_s} + \frac{1 - ms_{it}}{\sigma_s} \;\; \text{(Cournot)} \qquad \text{and} \qquad \varepsilon_{it}^{D} = \sigma_s \left(1 - ms_{it}\right) + \varrho_s \, ms_{it} \;\; \text{(Bertrand)}, \tag{A7}$$

decreasing in the firm's share and converging to the conditional elasticity σ_s as $ms_{it} \to 0$ (Atkeson and Burstein, 2008). Both the conditional and the decision-relevant marginal revenue product of an input contain the same physical margin $\frac{\partial Q_{it}}{\partial X_{it}}$, which cancels from their ratio. This implies that the decision-relevant marginal revenue product of every input is κ_{it}^{-1} times its conditional counterpart, where

$$\kappa_{it} \equiv \frac{1 - \frac{1}{\sigma_s}}{1 - \frac{1}{\varepsilon_{it}^{D}}} = 1 + \frac{\frac{1}{\varrho_s} - \frac{1}{\sigma_s}}{1 - \frac{1}{\sigma_s}} \, ms_{it} + O\big(ms_{it}^2\big) \;\; \text{(Cournot)}, \tag{A8}$$

with $\kappa_{it} \geq 1$, increasing in ms_{it}, and $\kappa_{it} \to 1$ as $ms_{it} \to 0$. Under Bertrand the first-order coefficient is $\frac{\sigma_s - \varrho_s}{\sigma_s(\sigma_s - 1)}$, so κ_{it} rises more slowly in market share.[30] Dividing by the input's price, the wedge measured by Lemma 1 is $\gamma_{it}^{X} = \kappa_{it} \times \frac{MRPX_{it}^{d}}{P_{it}^{X}}$, where $MRPX_{it}^{d}$ denotes the decision-relevant marginal revenue product: the recovered wedge is κ_{it} times the wedge relative to the marginal revenue product that governs the firm's behavior.

Several implications follow. First, κ_{it} multiplies the measured wedge of every input identically, so it cancels exactly from the relative wedge $\frac{\gamma_{it}^{L}}{\gamma_{it}^{M}}$ in equation (3) and from cross-input comparisons more generally. Second, because κ_{it} enters both measured wedges multiplicatively and increases in market share, variation in κ_{it} moves them in the *same* direction: oligopoly internalization cannot by itself generate the negative correlation documented in Section 3. Third, the deviation of κ_{it} from one vanishes linearly in ms_{it} and vanishes under the maintained assumptions of Section 2.

A.3 Rent sharing with upward-sloping input supply curves

This appendix develops the model that motivates the empirical analysis of Section 4. The model combines a collective bargaining framework that is standard in the rent-sharing literature (Card et al., 2018) with upward-sloping supply curves in the labor and intermediate input markets, as in work on monopsony power and buyer power (e.g., Yeh et al. (2022); Rubens (2023)). I derive equation (9) of the main text and show that the combination of these ingredients, but none of them in isolation, can generate the joint distribution of input

[30] Under Cournot, $\frac{1}{\varepsilon_{it}^{D}}$ is linear in ms_{it}, so $\kappa_{it} = \left[1 - \frac{\sigma_s - \varrho_s}{\varrho_s(\sigma_s - 1)} \, ms_{it}\right]^{-1}$ exactly, and the expansion in equation (A8) is its geometric series, which is valid for all $ms_{it} \in [0, 1]$ as $\varrho_s > 1$. Under Bertrand, the analogous expansion of $\frac{1}{\varepsilon_{it}^{D}} = \frac{1}{\sigma_s}\big(1 - \frac{\sigma_s - \varrho_s}{\sigma_s} \, ms_{it}\big)^{-1}$ delivers the stated first-order coefficient, which is smaller than its Cournot counterpart because $\varrho_s < \sigma_s$.

wedges documented in Section 3. For notational simplicity and following the related work that the model borrows from, I abstract from input heterogeneity and suppress the time subscript.

Input supply Supply curves $M_i(P_i^M)$ and $L_i(W_i)$ are continuous and monotonically increasing in their arguments, and supply elasticities $\varepsilon_i^M = \frac{\partial M_i(P_i^M)}{\partial P_i^M}\frac{P_i^M}{M_i(P_i^M)}$ and $\varepsilon_i^L = \frac{\partial L_i(W_i)}{\partial W_i}\frac{W_i}{L_i(W_i)}$ are positive but finite ($0 < \varepsilon_i^M, \varepsilon_i^L < \infty$). This approach nests the competitive case in the limit as an input's supply elasticity approaches infinity. Similar to Yeh et al. (2022) and Rubens (2023), I do not specify the source of the finite supply elasticities. To focus on the link between labor and intermediates, I assume that capital is predetermined.

Revenue and profit Let firm i's revenue function be $R_i = F(K_i, L_i, M_i, \Omega_i)$, where Ω_i is revenue productivity, and assume it is twice differentiable with $\frac{\partial R_i}{\partial x} > 0$ and $\frac{\partial^2 R_i}{\partial x^2} < 0$ for $x \in \{K_i, L_i, M_i\}$. As a consequence, profit is $\Pi_i(K_i, L_i, M_i) = R_i(K_i, L_i, M_i, \Omega_i) - P_i^K K_i - W_i L_i(W_i) - P_i^M M_i(P_i^M)$. A simple example where all assumed properties hold is the combination of demand and production specified in Section 2.2.

Determination of wages and intermediate inputs In line with the European context, wages are determined by collective bargaining between workers and their employer, while firms choose their intermediate inputs independently. I assume that intermediate input suppliers are price takers and that labor can be adjusted costlessly, but relax both assumptions in Appendices A.4 and A.5 and show that the rent-sharing mechanism remains valid. As is customary in the literature, I assume that workers collectively aim to maximize their employer's wage bill ($W_i L_i$). I model wage determination using a generalized Nash bargaining solution between the firm and workers. Because the firm faces an upward-sloping labor supply curve, the negotiated wage affects both the division of rents and the amount of labor supplied (Wong, 2026). For simplicity, I assume that the outside option of both parties is 0. This assumption can be relaxed at the cost of additional notational complexity. The generalized Nash bargaining solution solves

$$\max_{W_i, M_i} \Big(W_i L_i(W_i)\Big)^{\phi} \Big(\Pi_i\big(K_i, L_i(W_i), M_i\big)\Big)^{1-\phi}, \tag{A9}$$

where ϕ denotes workers' bargaining power.

The first-order conditions can be written as

$$P_i^M = \frac{\varepsilon_i^M}{\varepsilon_i^M + 1} MRPM_i \;\text{ and }\; W_i = \phi\left(\frac{R_i - P_i^M M_i - P_i^K K_i}{L_i}\right) + (1-\phi)\frac{\varepsilon_i^L}{\varepsilon_i^L + 1} MRPL_i. \tag{A10}$$

When the intermediates supply elasticity is finite, the price of intermediates is marked down relative to their marginal revenue product. The less elastic intermediates supply is, the more

buyer power a firm has, and the larger the intermediates price markdown. The wage is a weighted average of two familiar cases. If $\phi = 0$, workers have no bargaining power, and the wage is marked down with respect to the marginal revenue product of labor. In the limit as $\phi \to 1$, workers capture all the residual rents that are bargained over.

To clarify the relation between wedges in different input markets, I rewrite (A10) as

$$\frac{W_i}{MRPL_i} = \phi \left(\frac{\mu_i}{\tilde{\theta}_i^L} - \frac{\tilde{\theta}_i^M}{\tilde{\theta}_i^L} \frac{P_i^M}{MRPM_i} - \frac{\tilde{\theta}_i^K}{\tilde{\theta}_i^L} \frac{P_i^K}{MRPK_i} \right) + (1 - \phi) \frac{\varepsilon_i^L}{\varepsilon_i^L + 1}, \tag{A11}$$

where, as before, $\tilde{\theta}^X$ is the output elasticity of input X and μ_i is the ratio of the output price to marginal cost – the markup. If workers have at least some bargaining power, firms share part of the rents that they generate by increasing their wages relative to their marginal revenue contribution. If the residual rents are large enough, wages could even exceed the marginal revenue product, resulting in wage markups. Equation (A11) makes it clear that such rents could originate in product market markups, the focus of the majority of the rent-sharing literature, or could originate in intermediate input markdowns, a novel prediction. In other words, buyer power for intermediates puts upward pressure on wages, downward pressure on wage markdowns, and redistributes rents from intermediate input suppliers to workers if $\phi > 0$. Of course, indirect effects could create additional links between buyer power and wages, as these objects are jointly determined in equilibrium together with, for instance, markups. Hence, the exact magnitude of the relation between imperfections in labor and intermediate input markets is an empirical question that depends on underlying features of the model.

A.4 Intermediate input price bargaining

This appendix shows that the rent-sharing mechanism in the theoretical model of Appendix A.3 is robust to allowing the intermediate input price to be determined through bilateral bargaining. I focus on two simple models, one where both price and quantity are bargained over, and one where firms face an upward-sloping supply curve but bargain over the price. I show that wage equation (A10) and, hence, the analysis of how increases in intermediate input market rents put upward pressure on wages is unaffected by the introduction of bargaining. What changes is the interpretation of the intermediates wedge, and the fact that it need not necessarily exceed unity when firms bargain. Similar results can be obtained with multiple input suppliers or in other bargaining models, for instance one where input suppliers, the union, and the firm jointly bargain over total surplus. That is not to say that theoretical results from the main text survive with any bargaining model. In particular, if the intermediates price is determined by bargaining and the firm can then purchase any quantity at that price, the firm will purchase intermediate inputs such that the intermediates wedge equals unity and an analysis of the relation between the two wedges becomes meaningless.

Assume, in contrast to the main text, that the firm does not purchase intermediates given supply curve $M_i(P_i^M)$, but bargains with a single input supplier. The supplier has total and marginal costs $C_i^M(M_i)$ and $c_i^M(M_i)$, with $C_i^M(0) = 0$, $C_i^M(M_i) \geq 0$ and $\frac{\partial c_i^M}{\partial M_i} \geq 0$. Denote the supplier's bargaining weight by $\phi^M \in [0,1]$ so that the firm's weight when bargaining with the intermediates supplier is $1 - \phi^M$.

Bargaining on P_i^M and M_i Assume that the firm conducts two simultaneous bilateral Nash bargains: one with the supplier over (P_i^M, M_i) and one with the union over W_i. In each bilateral bargain, the parties compute the gains from reaching that agreement relative to the passive-beliefs counterfactual in which the current bilateral agreement is absent and all other equilibrium agreements are held fixed. As both inputs are essential and input contracts are assumed production-contingent, failure to reach an agreement implies zero payoff for both parties.

Conditional on the labor agreement holding at (W_i^*, L_i^*), the firm's payoff from a supplier agreement at (P_i^M, M_i) is $R_i(K_i, L_i^*, M_i, \Omega_i) - P_i^K K_i - W_i^* L_i^* - P_i^M M_i$, and the supplier's payoff is $P_i^M M_i - C_i^M(M_i)$. The bilateral Nash bargain solves

$$\max_{P_i^M, M_i} \Big(R_i(K_i, L_i^*, M_i, \Omega_i) - P_i^K K_i - W_i^* L_i^* - P_i^M M_i\Big)^{1-\phi^M} \Big(P_i^M M_i - C_i^M(M_i)\Big)^{\phi^M}. \quad \text{(A12)}$$

Because the transfer $P_i^M M_i$ redistributes surplus without affecting total surplus, the first-order condition for M_i reduces to $MRPM_i = c_i^M(M_i)$: the bargained quantity equates the marginal revenue product of intermediates to the supplier's marginal cost and is privately efficient conditional on the labor agreement. The first-order condition for P_i^M yields the surplus-split condition

$$P_i^M M_i = C_i^M(M_i) + \phi^M \Big(R_i(K_i, L_i^*, M_i, \Omega_i) - P_i^K K_i - W_i^* L_i^* - C_i^M(M_i)\Big), \quad \text{(A13)}$$

so the supplier is paid its cost plus a share ϕ^M of the bilateral surplus generated by the intermediate input agreement, taking the labor agreement as fixed. The corresponding intermediate input wedge is

$$\gamma_i^M = \frac{MRPM_i}{P_i^M} = \frac{c_i^M(M_i) M_i}{C_i^M(M_i) + \phi^M \Big(R_i(K_i, L_i^*, M_i, \Omega_i) - P_i^K K_i - W_i^* L_i^* - C_i^M(M_i)\Big)}. \quad \text{(A14)}$$

Buyer power does not distort the bargained quantity here, it shows up in the division of surplus, by pushing the average price paid to the supplier below the marginal revenue product of intermediates. Because this marginal revenue product equals the supplier's marginal cost, the wedge exceeds one when the supplier receives an average payment below marginal cost. This is more likely when supplier costs are convex, so marginal cost exceeds average cost, and when the firm has enough bargaining power to prevent the supplier from capturing too much of the surplus.

The labor bargain holds the supplier agreement fixed at (P_i^{M*}, M_i^*), so it solves

$$\max_{W_i} \Big(W_i L_i(W_i)\Big)^{\phi} \Big(R_i(K_i, L_i(W_i), M_i^*, \Omega_i) - P_i^K K_i - W_i L_i(W_i) - P_i^{M*} M_i^*\Big)^{1-\phi}. \quad \text{(A15)}$$

The first-order condition for W_i yields exactly the same first-order condition as in the main text

$$W_i = \phi \left(\frac{R_i(K_i, L_i, M_i^*, \Omega_i) - P_i^K K_i - P_i^{M*} M_i^*}{L_i} \right) + (1-\phi) \frac{\varepsilon_i^L}{\varepsilon_i^L + 1} MRPL_i, \quad \text{(A16)}$$

only now $P_i^{M*} M_i^*$ reflects bilateral firm-supplier bargaining over the joint gain from trade rather than movement along a supply curve. The supplier transfer plays the same role here as intermediate input payments in the main text: it determines how much residual surplus is available to be shared with workers. A lower supplier transfer raises the intermediate input wedge and, when workers have positive bargaining power, raises the wage through the residual-surplus term.

Bargaining on P_i^M with an upward-sloping supply curve $M_i(P_i^M)$ Consider a case that mirrors the treatment of labor markets in the main text: the firm and the supplier bargain over the unit price only, while the supplied quantity is pinned down by the same upward-sloping supply curve $M_i(P_i^M)$ as in the main text. To define the supplier's payoff, let the total cost function $C_i^M(M_i)$ be such that it rationalizes the supply curve, so that $P_i^M(M_i) = c_i^M(M_i)$. As above, assume that beliefs are passive and production only occurs if both bargains end in agreement. Conditional on the labor agreement holding at (W_i^*, L_i^*), the supplier bargain solves

$$\max_{P_i^M} \Big(R_i(K_i, L_i^*, M_i(P_i^M), \Omega_i) - P_i^K K_i - W_i^* L_i^* - P_i^M M_i(P_i^M)\Big)^{1-\phi^M} \Big(P_i^M M_i(P_i^M) - C_i^M(M_i(P_i^M))\Big)^{\phi^M}. \quad \text{(A17)}$$

Because the supplier chooses its quantity optimally along the supply curve, the envelope theorem implies that a marginal increase in the price raises the supplier's surplus only through the inframarginal units it already sells. The firm's surplus, by contrast, responds both through the price it pays and through the induced movement along the supply curve. The first-order condition of P_i^M can be rewritten as

$$\gamma_i^M = 1 + \frac{1}{\varepsilon_i^M} \left(1 - \frac{\phi^M}{1-\phi^M} \frac{R_i(K_i, L_i^*, M_i(P_i^M), \Omega_i) - P_i^K K_i - W_i^* L_i^* - P_i^M M_i(P_i^M)}{P_i^M M_i(P_i^M) - C_i^M(M_i(P_i^M))} \right), \quad \text{(A18)}$$

which generalizes the buyer-power markdown of the main text. When the supplier has no bargaining power ($\phi^M = 0$), the firm sets the price unilaterally and the same price markdown as in the main text is obtained. As the supplier's bargaining power rises it commands

a higher price and the wedge falls. The wedge equals one, and the bargained quantity is privately efficient, only at the knife-edge weight equating $\phi^M\Big(R_i(K_i, L_i^*, M_i(P_i^M), \Omega_i) - P_i^K K_i - W_i^* L_i^* - P_i^M M_i(P_i^M)\Big)$ and $(1-\phi^M)\Big(P_i^M M_i(P_i^M) - C_i^M(M_i(P_i^M))\Big)$. For larger supplier weights the price exceeds the marginal revenue product and the wedge falls below one. Supplier bargaining power, therefore, can offset buyer power, and the intermediate input wedge reports the net outcome of these two forces.

Turning to the wage, the labor bargain holds the supplier agreement fixed at (P_i^{M*}, M_i^*) and solves $\max_{W_i}\Big(W_i L_i(W_i)\Big)^{\phi}\Big(R_i(K_i, L_i(W_i), M_i^*, \Omega_i) - P_i^K K_i - W_i L_i(W_i) - P_i^{M*} M_i^*\Big)^{1-\phi}$, which delivers exactly the wage equation of the main text and the previous paragraph,

$$W_i = \phi\left(\frac{R_i(K_i, L_i, M_i^*, \Omega_i) - P_i^K K_i - P_i^{M*} M_i^*}{L_i}\right) + (1-\phi)\frac{\varepsilon_i^L}{\varepsilon_i^L + 1} MRPL_i, \tag{A19}$$

the only difference being that $P_i^{M*} M_i^*$ now reflects the bargained intermediate input payment along the supply curve rather than the unilateral monopsony payment. The rent-sharing mechanism of the main text therefore operates unchanged.

A.5 Labor adjustment frictions

This appendix shows that the rent-sharing mechanism of Appendix A.3 is robust to allowing labor to be costly to adjust. When hiring and firing are costly, the firm's employment choice becomes dynamic and the static wage first-order condition (A10) no longer holds as stated. I show that Nash bargaining nevertheless delivers a similar wage equation, with two modifications: the monopsony term marks down the forward-looking marginal value of labor that nets out adjustment costs, rather than the current marginal revenue product, and the rents that are shared include the continuation value of the employment relationship. The intermediate input markdown is unaffected, and the cross-market channel through which buyer power for intermediates raises wages and reduces labor wedges survives. When adjustment costs vanish, the wage equation collapses to (A10).

Reinstating the time subscript suppressed in Appendix A.3, the firm enters period t with predetermined capital K_{it} and a workforce $L_{i,t-1}$ inherited from the previous period. As in the main text, it faces upward-sloping labor and intermediate input supply curves $L_{it}(W_{it})$ and $M_{it}(P_{it}^M)$ with finite positive elasticities ε_{it}^L and ε_{it}^M, and revenue $R_{it} = F(K_{it}, L_{it}, M_{it}, \Omega_{it})$ with the properties assumed in Appendix A.3. Adjusting the workforce away from its inherited level is costly: the firm pays $\Phi_i(L_{it}, L_{i,t-1}) \geq 0$, measured in revenue units, where Φ_i is continuously differentiable, convex in its first argument, and satisfies $\Phi_i(L, L) = 0$. I write $\Phi_{1,it}$ and $\Phi_{2,it}$ for the partial derivatives of Φ_i with respect to its first and second arguments.

Let $V_i(L_{i,t-1}, z_{it})$ denote the firm's value at the start of the period, where z_{it} collects the predetermined and exogenous state (K_{it}, Ω_{it}). Flow operating surplus before wage payments

is

$$\Pi^o_{it} \equiv R_{it} - P^K_{it} K_{it} - P^M_{it} M_{it} - \Phi_i(L_{it}, L_{i,t-1}), \tag{A20}$$

and the firm's value satisfies

$$V_i(L_{i,t-1}, z_{it}) = \Pi^o_{it} - W_{it} L_{it} + \beta \, \mathbb{E}_t V_i(L_{it}, z_{i,t+1}), \tag{A21}$$

where $\beta \in [0, 1)$ is the discount factor.

As in Appendix A.3, the wage is the outcome of Nash bargaining between the firm and its collectively organized workers, who maximize the wage bill $W_{it} L_{it}$, while the firm chooses intermediates. Bargaining takes place over the current wage, with the continuation value $V_i(\cdot, z_{i,t+1})$ taken as given.[31] The firm's surplus from agreement is

$$S^F_{it} = \Pi^o_{it} - W_{it} L_{it} + C_{it}, \qquad C_{it} \equiv \beta \, \mathbb{E}_t \big[V_i(L_{it}, z_{i,t+1}) - V_i(L_{i,t-1}, z_{i,t+1}) \big], \tag{A22}$$

where C_{it} is the continuation value of carrying the chosen workforce L_{it} into the next period rather than the inherited $L_{i,t-1}$; it is zero when employment is unchanged or adjustment is costless. The bargaining problem is

$$\max_{W_{it}, M_{it}} \big(W_{it} L_{it}(W_{it}) \big)^{\phi} \big(S^F_{it} \big)^{1-\phi}. \tag{A23}$$

Neither the union's objective, the adjustment cost, nor the continuation surplus depends on intermediates, so the first-order condition for M_{it} maximizes Π^o_{it} and reproduces the main-text markdown,

$$P^M_{it} = \frac{\varepsilon^M_{it}}{\varepsilon^M_{it} + 1} \, MRPM_{it}. \tag{A24}$$

Labor adjustment frictions therefore leave the buyer-power result for intermediates intact.

Define

$$J_{it} \equiv \frac{\partial \Pi^o_{it}}{\partial L_{it}} + \frac{\partial C_{it}}{\partial L_{it}} = MRPL_{it} - \Phi_{1,it} + \beta \, \mathbb{E}_t \Lambda_{i,t+1}, \qquad \Lambda_{i,t+1} \equiv \frac{\partial V_i(L_{it}, z_{i,t+1})}{\partial L_{it}}. \tag{A25}$$

J_{it} is the marginal contribution of a worker to firm surplus: the current marginal revenue product, net of the marginal adjustment cost incurred today, plus the discounted marginal value $\Lambda_{i,t+1}$ of carrying that worker into the next period. It replaces $MRPL_{it}$ as the relevant return to labor and reduces to it when adjustment is costless.

[31] This is the dynamic counterpart of the static bargain in (A9). As in the main text, outside options are set to zero, with the convention that disagreement yields no flow payoff and leaves the inherited workforce $L_{i,t-1}$ in place. The continuation value then enters the bargained surplus only through the effect of the chosen employment L_{it} on the value of carrying the workforce forward, which is governed by adjustment costs. This nests the static model when adjustment costs vanish and can be relaxed at the cost of additional notation.

The first-order condition for W_{it} in (A23) is $\frac{\phi}{W_{it}L_{it}}\frac{\partial(W_{it}L_{it})}{\partial W_{it}}+\frac{1-\phi}{S^F_{it}}\frac{\partial S^F_{it}}{\partial W_{it}}=0$. By the definition of the labor supply elasticity, $\frac{\partial L_{it}(W_{it})}{\partial W_{it}}=\varepsilon^L_{it}\frac{L_{it}}{W_{it}}$, so that $\frac{\partial(W_{it}L_{it})}{\partial W_{it}}=L_{it}(1+\varepsilon^L_{it})$ and $\frac{\partial S^F_{it}}{\partial W_{it}}=J_{it}\varepsilon^L_{it}\frac{L_{it}}{W_{it}}-L_{it}(1+\varepsilon^L_{it})$. Substituting these expressions and rearranging yields

$$W_{it}=\phi\left(\frac{\Pi^o_{it}+C_{it}}{L_{it}}\right)+(1-\phi)\,\frac{\varepsilon^L_{it}}{\varepsilon^L_{it}+1}\,J_{it}. \tag{A26}$$

Equation (A26) is the wage equation (A10) of the main text, with two modifications. First, the monopsony term marks down the dynamic marginal value J_{it} rather than the current marginal revenue product $MRPL_{it}$. Second, the per-worker rents that are shared, $(\Pi^o_{it}+C_{it})/L_{it}$, are forward-looking: they net out adjustment costs and add the continuation value of the relationship. Setting $\Phi_i\equiv 0$ yields $\Phi_{1,it}=0$ and an employment-independent continuation value, so that $C_{it}=0$, $\Lambda_{i,t+1}=0$, and $J_{it}=MRPL_{it}$. Equation (A26) then collapses to the main text wage equation.

Dividing (A26) by $MRPL_{it}$ and using $MRPL_{it}=(\tilde\theta^L_{it}/\mu_{it})R_{it}/L_{it}$ for the revenue term and the analogous identities for the intermediate input and capital terms gives the dynamic counterpart of (9),

$$\frac{W_{it}}{MRPL_{it}}=\phi\left(\frac{\mu_{it}}{\tilde\theta^L_{it}}-\frac{\tilde\theta^M_{it}}{\tilde\theta^L_{it}}\frac{P^M_{it}}{MRPM_{it}}-\frac{\tilde\theta^K_{it}}{\tilde\theta^L_{it}}\frac{P^K_{it}}{MRPK_{it}}-\frac{\Phi_{it}}{L_{it}\,MRPL_{it}}+\frac{C_{it}}{L_{it}\,MRPL_{it}}\right) \tag{A27}$$

$$+(1-\phi)\,\frac{\varepsilon^L_{it}}{\varepsilon^L_{it}+1}\left(1-\frac{\Delta_{it}}{MRPL_{it}}\right), \tag{14}$$

where $\Phi_{it}\equiv\Phi_i(L_{it},L_{i,t-1})$ and $\Delta_{it}\equiv MRPL_{it}-J_{it}=\Phi_{1,it}-\beta\,\mathbb{E}_t\Lambda_{i,t+1}$ is the wedge between the current and the dynamic marginal value of labor. Equation (A27) differs from equation (9) only through three dynamic terms, each of which is zero absent adjustment frictions: the flow adjustment cost Φ_{it} reduces the rents available for sharing, the continuation value C_{it} adds the value of carrying the chosen workforce forward, and Δ_{it} shifts the monopsony benchmark from the current to the dynamic marginal value of labor. The intermediate input wedge $\gamma^M_{it}=MRPM_{it}/P^M_{it}$ enters exactly as in the main text, through the term $-\frac{\tilde\theta^M_{it}}{\tilde\theta^L_{it}}\frac{1}{\gamma^M_{it}}$. Labor adjustment frictions thus leave the cross-input-market channel unchanged.

B Revenue function estimation

This appendix discusses the invertibility of intermediate input demand and the estimation of revenue elasticities. It also reports means, medians, and standard deviations of revenue elasticities and revenue returns to scale by 2-digit industry in Table B1.

B.1 Control function

I control for unobserved revenue productivity by inverting demand for intermediate inputs. This requires revenue productivity to be the sole unobservable in intermediate input demand and intermediate input demand to increase monotonically in revenue productivity. Intuitively, the monotonicity condition rules out adjustment frictions so severe that demand for intermediates does not respond to changes in revenue productivity. In that case, intermediates would be uninformative of revenue productivity.

To discuss invertibility, it is useful to derive an explicit expression for intermediate input demand, which requires taking a stand on demand and production. For simplicity, consider a time-invariant Cobb-Douglas production function $Q_{it} = K_{it}^{\tilde{\theta}^K} L_{it}^{\tilde{\theta}^L} M_{it}^{\tilde{\theta}^M} \tilde{\Omega}_{it}$ and an iso-elastic demand function $P_{it} = Q_{it}^{1/\eta} \Xi_{it}$. The revenue function is then given by $R_{it} = K_{it}^{\theta^K} L_{it}^{\theta^L} M_{it}^{\theta^M} \Omega_{it}$. Below, I discuss how the intuition developed here extends to the revenue function in equation (5). Let inverse intermediate input supply be given by $P_{it}^M(M_{it})$, with intermediate supply elasticity $0 < \varepsilon_{it}^M < \infty$. The first-order condition for intermediate inputs in equation (A10) can be rewritten to give

$$M_{it}^* = \left(\frac{1}{P_{it}^M} \frac{\varepsilon_{it}^M}{\varepsilon_{it}^M + 1} \theta^M K_{it}^{\theta^K} L_{it}^{\theta^L} \Omega_{it} \right)^{\frac{1}{1-\theta^M}}. \tag{B1}$$

Two assumptions need to be met for invertibility to hold. First, M_{it}^* needs to be one-to-one in Ω_{it}, which is immediate from equation (B1). Recall from the revenue function in equation (5) that $\Omega_{it} = \tilde{\Omega}_{it}^{1+\eta/\eta} \tilde{\Xi}_{it}$. As both shocks enter the revenue function multiplicatively, all that is required is for intermediate input demand to be one-to-one in their product, although here it is one-to-one in the two individual shocks as well.

Second, the scalar unobservable assumption requires that Ω_{it} is the only unobservable that varies within the level of estimation. The two unobserved variables in equation (B1) that might violate this assumption are the intermediate input price and supply elasticity. If both are constant at the level of estimation – either a 2-digit industry or a year-by-2-digit-industry – then Ω_{it} is the only unobservable that varies across observations. This could occur, for instance, if firms are monopsonists or in monopsonistic competition and face inverse intermediate input supply curve $P_{it}^M = M_{it}^{1/\varepsilon^M}$. To further allow for price and elasticity differences across time, more narrowly defined industries, or locations, I include indicators for the year, a firm's 3-digit industry, and its NUTS1 region in the control function. In addition, I allow for strategic interaction between firms in the intermediate input market

by adding a firm's share of total intermediate input expenditure in its industry to the control function. This follows from assuming that the industry-specific inverse intermediate input supply is a function of total industry purchases, $P_{it}^M = P_t^M(\sum_j M_{jt})$, so that $\frac{\varepsilon_{it}^M}{\varepsilon_{it}^M+1}$ in equation (B1) is replaced by $\frac{\varepsilon_{it}^M}{\varepsilon_{it}^M+\frac{P_{it}^M M_{it}}{\sum_j P_{jt}^M M_{jt}}}$ and the intermediates price is constant at the level of estimation, possibly conditional on fixed effects.[32] In case there is nonetheless remaining price or elasticity variation due to buyer power, I include a firm's wage in the control function, as equation (A10) implies that wages are related to intermediate input prices and supply elasticities through rent sharing. Finally, to allow for intermediate input price differences stemming from quality differences across firms, a firm's revenue share can be included in the control function – the formal basis for such an exercise is elaborately discussed in Appendix A of De Loecker et al. (2016).

My baseline empirical specification uses a translog production function with a Hicks-neutral productivity term, which yields log revenue as given by equation (5). Compared to the intermediate input demand in equation (B1), switching from Cobb-Douglas to translog introduces no new unobservables, so the intermediate input price and supply elasticity remain the main challenge to the scalar unobservable assumption. Compared to equation (B1), revenue elasticities are firm-time-specific under the translog production function. However, revenue elasticity variation across firms and time within the same industry is driven entirely by observable differences in input use. Moreover, by equation (5), the added demand-side variation from demand equation (4) is absorbed by fixed effects. Finally, as in the Cobb-Douglas case, revenue productivity enters revenue function (5) log additively, so M_{it}^* is still strictly increasing in Ω_{it}.

B.2 Estimation of revenue elasticities

This subsection explains how I obtain revenue elasticities. I first describe my baseline control-function approach, which follows De Loecker (2011) and Petrin and Sivadasan (2013). I then discuss five alternative approaches: a revenue function estimated separately by 2-digit industry and 2-year period, estimates that fix returns to scale, estimates using additional information on Chinese imports, estimates based on an AR(1) revenue productivity process, and calibrated revenue elasticities.

The baseline approach starts from equation (5) and faces two practical challenges. The first is controlling for unobserved revenue productivity, ω_{it}. If a firm's inputs at time t are at least partially determined after the firm observes ω_{it}, estimates of the revenue elasticities

[32] Intermediate input expenditure shares pin down variation due to supply elasticities in oligopsony models more broadly. See, for instance, equation (7) of Berger et al. (2022), $\varepsilon_{it}^M = (\frac{1}{\eta^{in}} + (\frac{1}{\eta^{out}} - \frac{1}{\eta^{in}})\frac{P_{it}^M M_{it}}{\sum_j^n P_{jt}^M M_{jt}})^{-1}$, where elasticities η^{in} and η^{out} are exogenous and measure, respectively, the substitutability of firms within the intermediate input market and between different intermediate input markets.

can be biased. A similar concern arises in production function estimation, where ω_{it} would contain only unobserved physical productivity. Indeed, if $\tilde{\xi}_{it}$ is i.i.d. across firms and time, revenue productivity is physical productivity scaled by a constant. The second challenge is that intermediate inputs are measured by deflating intermediate input expenditure using a year-by-2-digit-industry intermediate input price deflator, $\bar{P}^M_{st}$. Any within-industry intermediate input price deviations from $\bar{P}^M_{st}$ enter the structural error term of the revenue function, scaled by the revenue elasticities.[33] Collecting all input price bias terms in $b_{it}(\cdot)$, equation (5) can be written as

$$r_{it} = f_s(k_{it}, l_{it}, m_{it}; \beta) + b_{it}(\cdot) + \gamma_{st} + \omega_{it} + \epsilon_{it}, \tag{B2}$$

where $f_s(k_{it}, l_{it}, m_{it}; \beta)$ refers to the translog function, β contains all coefficients, and s refers to the level of estimation – a 2-digit industry in the baseline estimates.

Following De Loecker (2011) and Petrin and Sivadasan (2013), my baseline specification uses demand for intermediate inputs to proxy for ω_{it},

$$m_{it} = m_s(\omega_{it}, k_{it}, l_{it}, \mathbf{x}^m_{it}; \delta^m), \tag{B3}$$

where $\mathbf{x}^m_{it}$ contains control variables and δ^m contains all coefficients. If ω_{it} is the sole unobservable in equation (B3) and $m_s(\cdot)$ is one-to-one in ω_{it}, inverting (B3) yields a control for ω_{it} based on observables,

$$\omega_{it} = m_s^{-1}(k_{it}, l_{it}, m_{it}, \mathbf{x}^m_{it}; \delta^m). \tag{B4}$$

As ω_{it} contains the markup, physical productivity, and the demand shock, these factors need not be controlled for separately in (B3), in contrast to proxy variable approaches to production function estimation. Invertibility is discussed at length in the previous subsection. The controls in $\mathbf{x}^m_{it}$ are firm i's intermediate inputs as a share of total intermediate inputs in the firm's 4-digit industry at time t (mms_{it}), the logged wage at firm i in year t, and year, NUTS1, and 3-digit-industry fixed effects.[34]

To control for within-industry intermediate input price differences entering the revenue function directly, I use the input price control function

$$b_{it} = b_s((1, k_{it}, l_{it}, m_{it}) \times \mathbf{x}^b_{it}; \delta^b), \tag{B5}$$

[33]For instance, for a translog revenue function the input price bias resulting from within-industry-year variation in the price of intermediates is $\log\left(\frac{P^M_{it}}{\bar{P}^M_{st}}\right)\left(\beta_m + \beta_{km}k_{it} + \beta_{lm}l_{it} + \beta_{mm}\left(2m_{it} + \log\left(\frac{P^M_{it}}{\bar{P}^M_{st}}\right)\right)\right)$. Of course, buyer power does not necessitate input price variation. In a Cournot model of the intermediate input market, for instance, no input price variation exists.

[34]Intermediate input market shares are strongly correlated with market shares, and as a result including both variables in the input price control function leads to convergence problems. I therefore only include the intermediate input market share. All key results are robust to including revenue market shares instead.

where $\mathbf{x}_{it}^b$ contains mms_{it} as well as year, NUTS1, and 3-digit-industry fixed effects. The notation of $b_s(\cdot)$ indicates that the inputs do not enter linearly, but only interacted with $\mathbf{x}_{it}^b$, as a consequence of the translog specification. Input price variation related to revenue productivity ω_{it} is already controlled for by the control function in equation (B4), and including the control function (B5) can soak up part of the input price dispersion even if it does not perfectly control for its sources (De Loecker et al., 2016).

Substituting equations (B5) and (B4) into the revenue function in equation (B2) gives

$$r_{it} = f_s(k_{it}, l_{it}, m_{it}; \beta) + b_s((1, k_{it}, l_{it}, m_{it}) \times \mathbf{x}_{it}^b; \delta^b) + m_s^{-1}(k_{it}, l_{it}, m_{it}, \mathbf{x}_{it}^m; \delta^m) + \gamma_{st} + \epsilon_{it}. \tag{B6}$$

The revenue elasticity of input $x \in \{k, l, m\}$ is given by $\theta_{it}^x = \frac{\partial f_s}{\partial x_{it}}$, where the translog specification makes revenue elasticities firm-time specific. The control functions for revenue productivity and input price bias are treated non-parametrically, as their functional forms depend on the functional form of intermediate input supply, which I leave unspecified.

To identify the coefficients, I use a two-step estimation approach based on Ackerberg et al. (2015). The first step approximates equation (B6) by regressing log revenue on a third-degree polynomial in all right-hand-side variables except indicator variables, which enter linearly. This first step is not meant to identify any of the coefficients in β, δ^b, or δ^m, but rather to separate observed revenue into planned revenue and random deviations from planned revenue: $r_{it} = \hat{r}_{it} + \hat{\epsilon}_{it}$. This distinction is important because I need $\hat{\epsilon}_{it}$ to correct revenue shares when constructing input wedges. Revenue shares need to be corrected because $R_{it}(\cdot)\exp(\epsilon_{it})$ is observed in the data, while firms base their decisions on planned revenue $R_{it}(\cdot)$. I therefore replace the revenue shares in equations (1) by $\hat{LS}_{it} = \frac{W_{it}L_{it}}{R_{it}/\exp(\hat{\epsilon}_{it})}$ and $\hat{MS}_{it} = \frac{P_{it}^M M_{it}}{R_{it}/\exp(\hat{\epsilon}_{it})}$ and likewise correct revenue when investigating revenue or value added as outcomes.

In the second step, I assume that revenue productivity evolves according to the first-order Markov process

$$\omega_{it} = g_s(\omega_{it-1}) + \psi_{it}, \tag{B7}$$

where ψ_{it} is a mean-zero revenue productivity shock. De Loecker (2011) explicitly assumes that $\tilde{\xi}_{it}$ is i.i.d. across firms and time, so that the Markov assumption applies only to physical productivity $\tilde{\omega}_{it}$. Here, the demand shock may be serially correlated, as long as the joint term ω_{it} has the first-order Markov property.

Using planned revenue $\hat{r}_{it}$, I obtain an estimate of ω_{it} that is a function of unknown coefficients β and δ^b from equation (B6),

$$\hat{\omega}_{it} = \hat{r}_{it} - f_s(k_{it}, l_{it}, m_{it}; \beta) - b_s((1, k_{it}, l_{it}, m_{it}) \times \mathbf{x}_{it}^b; \delta^b). \tag{B8}$$

From the law of motion of revenue productivity in equation (B7), I then obtain an estimate of the revenue productivity shock that depends on unknown coefficients β and δ^b,

$$\hat{\psi}_{it} = \hat{\omega}_{it} - g_s(\hat{\omega}_{it-1}). \tag{B9}$$

I estimate the coefficients by forming moments on ψ_{it}. I approximate $g_s(\cdot)$ using a third-degree polynomial in all its arguments except indicator variables, which enter linearly. I assume $\mathbb{E}(\psi_{it}k_{it}) = 0$, in line with capital's significant adjustment frictions and its predetermined role in my theoretical framework. I allow both l_{it} and m_{it} to be adjusted in response to the realization of ψ_{it}, so that $\mathbb{E}(\psi_{it}l_{it}) \neq 0$ and $\mathbb{E}(\psi_{it}m_{it}) \neq 0$.

Together, this yields the following moments:

$$\mathbb{E}(\psi_{it}\mathbf{Z}_{it}) = 0, \tag{B10}$$

where I focus on the exactly identified case and $\mathbf{Z}_{it}$ contains a constant and all contemporaneous values of variables, except variables based on labor, intermediates, and $\hat{r}_{it}$, which are instead constructed using their first lags. That is, $\mathbf{Z}_{it}$ contains 1, k_{it}, k_{it}^2, l_{it-1}, l_{it-1}^2, m_{it-1}, m_{it-1}^2, $k_{it}l_{it-1}$, $k_{it}m_{it-1}$, $l_{it-1}m_{it-1}$, mms_{it-1}, mms_{it-1}^2, $mms_{it-1}k_{it}$, $mms_{it-1}l_{it-1}$, $mms_{it-1}m_{it-1}$, as well as year, NUTS1, and 3-digit-industry fixed effects. Estimates of revenue elasticities are reported in Table B1 of Appendix B.

I consider five alternative approaches to obtaining the elasticities of labor and intermediates. First, I estimate the revenue function at the 2-digit-industry-by-2-year level. This specification allows firm-specific revenue elasticities to vary over time not only because input use changes, but also because the derivatives of the revenue function with respect to inputs change. Estimation proceeds as in the baseline approach, but the revenue function and revenue productivity control are indexed by st, so that each 2-digit industry has a different revenue function every 2 years. To see what this specification accommodates, suppose log revenue is translog in effective labor $l_{it} + \omega_{it}^L$ rather than in l_{it}. Expanding equation (5) in observed inputs adds two terms: an additive component, $\beta_l\omega_{it}^L + \beta_{ll}(\omega_{it}^L)^2$, which is observationally equivalent to Hicks-neutral productivity, and an interaction component, $(2\beta_{ll}l_{it} + \beta_{kl}k_{it} + \beta_{lm}m_{it})\omega_{it}^L$, which changes revenue elasticities and exists only because the translog has second-order and cross terms. Decompose $\omega_{it}^L = \omega_{st}^L + v_{it}^L$, where st here indexes the estimation cell. If technical change is common within the cell ($v_{it}^L = 0$), the model is exactly a translog in observed inputs with cell-specific intercept $\gamma_{st} + \beta_l\omega_{st}^L + \beta_{ll}(\omega_{st}^L)^2$ and cell-specific first-order coefficients $\beta_{l,st} = \beta_l + 2\beta_{ll}\omega_{st}^L$, $\beta_{k,st} = \beta_k + \beta_{kl}\omega_{st}^L$, and $\beta_{m,st} = \beta_m + \beta_{lm}\omega_{st}^L$, with second-order coefficients unchanged. The remaining unobservable is the scalar ω_{it}, so the invertibility and Markov assumptions apply within each cell and, as the elasticity of observed labor equals that of effective labor, the estimates recover the true elasticities. If instead $v_{it}^L \neq 0$, the residual becomes $\omega_{it} + \beta_l v_{it}^L + \beta_{ll}(2\omega_{st}^L v_{it}^L + (v_{it}^L)^2) + (2\beta_{ll}l_{it} + \beta_{kl}k_{it} + \beta_{lm}m_{it})v_{it}^L$: the true elasticities exceed their cell-specific counterparts by $2\beta_{ll}v_{it}^L$, $\beta_{kl}v_{it}^L$, and $\beta_{lm}v_{it}^L$. The

cell-specific translog thus absorbs *st*-common factor-biased technical change exactly, but remains a reduced-form robustness exercise with respect to any firm-specific components.

Second, using the restricted sample of firms that import intermediates from China, I add information on intermediates unit prices, varieties, and trade intensity to my baseline estimates in two ways. First, I allow the law of motion of revenue productivity to depend on whether the firm imports intermediate inputs from China in year t:

$$\omega_{it} = g_s(\omega_{it-1}, \mathbb{I}_{it}^{import}) + \psi_{it}. \tag{B11}$$

Second, I include $\mathbb{I}_{it}^{import}$, and the logs of the intermediate input price index, the number of imported intermediate input varieties, and the share of intermediates imported in both $\mathbf{x}_{it}^m$ and $\mathbf{x}_{it}^b$. The resulting new moments are formed on the lags of these variables. This specification is estimated on the 2010 - 2018 sample, as Dutch trade data are only available from 2009 onward.

Third, to allow for the possibility that insufficient variation in intermediates demand conditional on inputs and revenue productivity exists (Gandhi et al., 2020), I apply the approach of Flynn et al. (2019) and fix returns to scale to unity, which closely mirrors the estimated revenue returns to scale from the baseline estimates. In particular, I impose the restrictions $\sum_{i\in k,l,m} \beta_i = 1$, $\beta_{km}+2\beta_{kk}+\beta_{kl} = 0$, $\beta_{lm}+2\beta_{ll}+\beta_{kl} = 0$, and $\beta_{km}+2\beta_{mm}+\beta_{lm} = 0$. Comparing the baseline estimates with these estimates is informative about whether the non-identification result of Gandhi et al. (2020) matters quantitatively for my estimates.

Fourth, in the spirit of Blundell and Bond (2000), I estimate the revenue function under the assumption that revenue productivity follows an AR(1) process: $\omega_{it} = \rho\omega_{it-1} + \psi_{it}$. Substituting for ω using the revenue function in equation (B2) yields the following moment conditions:

$$\mathbb{E}((\psi_{it} + \epsilon_{it} - \rho\epsilon_{it-1})\, \mathbf{Z}_{it}^{AR1}) = 0, \tag{B12}$$

where $\mathbf{Z}_{it}^{AR1}$ differs from $\mathbf{Z}_{it}$ only in that k_{it-1} is added to keep the model exactly identified in the presence of the new parameter ρ. The main advantage of this approach over the baseline specification is that ω_{it} is not proxied for using observables, so invertibility of equation (B3) is not required. The main drawback is that the AR(1) law of motion is more restrictive than equation (B7).

Finally, I dispense with revenue function estimation altogether and calibrate a single revenue elasticity for each input. For all firms, I set the revenue elasticity of labor equal to the median baseline estimate in the full sample. The revenue elasticity of intermediate inputs is calibrated similarly. By Lemma 1, variation in input wedges is then driven entirely by revenue shares. Alternative calibration choices, such as using a reasonable guess for the two elasticities based on the literature or taking the median within a firm's industry, yield similar results.

Table B1: Revenue elasticities and revenue returns to scale: median *mean* (standard deviation), by NACE division

NACE division	labor	intermediate inputs	capital	RRTS	observations
10	0.23 *0.23* (0.09)	0.71 *0.72* (0.10)	0.07 *0.07* (0.03)	1.01 *1.01* (0.02)	12,152
13	0.37 *0.38* (0.08)	0.61 *0.60* (0.10)	0.03 *0.04* (0.03)	1.01 *1.01* (0.05)	2,170
17	0.30 *0.30* (0.06)	0.67 *0.66* (0.09)	0.05 *0.06* (0.04)	1.02 *1.02* (0.04)	1,828
18	0.41 *0.41* (0.06)	0.59 *0.58* (0.08)	0.03 *0.03* (0.01)	1.02 *1.02* (0.01)	11,373
20	0.27 *0.27* (0.11)	0.70 *0.70* (0.10)	0.05 *0.05* (0.02)	1.02 *1.02* (0.02)	4,006
22	0.30 *0.30* (0.10)	0.67 *0.67* (0.13)	0.05 *0.05* (0.03)	1.02 *1.02* (0.02)	5,835
23	0.29 *0.28* (0.09)	0.68 *0.69* (0.09)	0.04 *0.04* (0.02)	1.01 *1.01* (0.03)	4,532
24	0.27 *0.26* (0.07)	0.68 *0.69* (0.09)	0.04 *0.04* (0.02)	1.00 *1.00* (0.03)	1,571
25	0.33 *0.33* (0.07)	0.64 *0.64* (0.09)	0.05 *0.04* (0.02)	1.01 *1.01* (0.02)	31,254
26	0.34 *0.34* (0.12)	0.66 *0.65* (0.13)	0.02 *0.02* (0.00)	1.02 *1.02* (0.01)	5,232
27	0.30 *0.30* (0.08)	0.70 *0.70* (0.09)	0.02 *0.02* (0.01)	1.01 *1.02* (0.03)	4,537
28	0.30 *0.30* (0.08)	0.68 *0.68* (0.10)	0.03 *0.03* (0.01)	1.01 *1.01* (0.02)	15,957
29	0.23 *0.24* (0.10)	0.73 *0.73* (0.11)	0.03 *0.03* (0.01)	1.00 *1.00* (0.02)	2,742
30	0.34 *0.34* (0.10)	0.64 *0.63* (0.13)	0.05 *0.05* (0.03)	1.02 *1.02* (0.04)	2,789
31	0.29 *0.30* (0.06)	0.69 *0.68* (0.10)	0.03 *0.04* (0.02)	1.01 *1.01* (0.02)	7,339
32	0.40 *0.39* (0.08)	0.59 *0.60* (0.08)	0.03 *0.03* (0.01)	1.01 *1.01* (0.04)	7,754
33	0.32 *0.34* (0.09)	0.67 *0.66* (0.10)	0.02 *0.02* (0.01)	1.00 *1.00* (0.02)	11,651
full sample	0.32 *0.32* (0.10)	0.65 *0.66* (0.10)	0.04 *0.04* (0.02)	1.01 *1.01* (0.02)	132,722

Notes: Summary statistics for firm-level revenue elasticities of labor, intermediate inputs, and capital, and the revenue returns to scale (RRTS). For each NACE division (2-digit NACE industry), the median, mean (in italics), and standard deviation (in brackets) are displayed. Medians, means, and standard deviations are rounded to 2 decimal points.

C Data

C.1 Data sources

This paper uses non-public yearly firm-level data from Statistics Netherlands (CBS), covering the universe of Dutch manufacturing firms from 2007 to 2018.[35]

The "General Firm Registry" (ABR) is a yearly firm-level dataset covering all firms in the Netherlands, built from Chamber of Commerce and Dutch tax authority registries,[36] with daily and monthly updates, respectively, used to construct anonymized firm identifiers. From the ABR, I take employment in full-time equivalent units (FTE), rounded to whole numbers, and each firm's 4-digit NACE industry and headquarters location.

The "Financial Statistics of Non-financial Firms" (NFO) is a yearly dataset of anonymized balance sheets and income statements of all identifiable non-financial firms. Data for "small" firms come from tax returns recorded by the Dutch Ministry of Finance; data for "large" firms come from surveys,[37] which the CBS checks against previously obtained data, recontacting firms when a reporting error is suspected. The cleaning procedure in this appendix eliminates reporting errors that survive this process, particularly for small firms. From the NFO, I take revenue, expenditure on labor and intermediate inputs, the book value of capital, and earnings before interest and taxes.

The CBS anonymizes firms at different levels of aggregation: the NFO comes at the "organization group" (OG) level, the "actual agent in financial processes" (CBS, 2020b), and the ABR at the "firm unit" (BE) level, an "autonomous actor in the production process" (CBS, 2020c). BEs map one-to-one to OGs except among the largest firms – the TOPX, roughly 2000 firms each year – where one OG can contain several BEs; for these, I aggregate ABR data to the OG level.[38] OG-level identifiers then allow me to merge the NFO and the ABR.

Industries follow Eurostat's NACE Rev. 2 classification: sections consist of one or more 2-digit divisions, with groups (3-digit) and codes (4-digit) increasingly disaggregated.[39] The ABR provides each firm's SBI08 code, whose first four digits are the NACE Rev. 2 code; the NFO's 2-digit SBI08 code serves as a consistency check. Pre-2008 SBI93 codes are converted

[35] These data are available to academic researchers, subject to conditions; see https://www.cbs.nl/en-gb/our-services/customised-services-microdata.

[36] In particular, the "Nieuwe Handelsregister" (NHR) of the Chamber of Commerce and the "Beheer van Relaties" (BvR) of the tax authority. Prior to April 1st, 2014, the "Basis Bedrijvenregister" (BBR), a partnership between both institutions and the CBS, was also employed.

[37] Before 2011, firms with a balance sheet total below 23 million euros were classified as small; from 2011 onward, the threshold is 40 million euros.

[38] Categorical variables are aggregated by taking the mode across BE-level observations. My results are not sensitive to this procedure.

[39] See https://ec.europa.eu/eurostat/documents/3859598/5902521/KS-RA-07-015-EN.PDF for a complete description and conversions to other international classifications.

to SBI08. The final sample includes all 2-digit industries in the section "Manufacturing" with sufficient observations to estimate the revenue function; Table C1 in Appendix C lists the 17 industries covered.

When estimating the revenue function, monetary variables are deflated using 2-digit NACE industry deflators from the OECD STAN database,[40] and firm location enters the input price control function as indicators for the headquarters' NUTS1 region, Eurostat's classification of EU economic territories for regional analysis.[41]

In Section 4, I add trade data from the "International Trade in Goods" (IHG), a yearly firm-level dataset available from 2009 onward that records the value of each firm's imports and exports by partner country, but no quantities and hence no unit values. I aggregate the IHG from the BE to the OG level as described above and retain only imports from China. For unit values, I turn to the BACI database, which provides bilateral trade values and quantities at the country-pair-by-6-digit-Harmonized-System-product level.[42] Since 6-digit HS codes are merged and split over time, I bundle all connected products in the sample period into time-consistent products. From the IHG and BACI, I construct each firm-year's import status, China import share, number of imported varieties and share of newly imported ones, and import price indices for all Chinese imports and separately for new and continued varieties. Annual-average exchange rates come from the OECD.

Finally, I match employee-level data from the "Spolisbus", a yearly CBS dataset available from 2009 onward recording wages, hours worked, and job and contract characteristics. Aggregated to the OG level, it yields each firm-year's headcount, worker exits, and churn; average base and extra wages; average tenure, measured within 2009 - 2018 and reset upon firm switches; the shares of full-time and permanent contracts; and workforce composition by job nature (managerial, regular, and other jobs, such as interns and on-call workers).

C.2 Definition of variables

- Revenue (R_{it}): Total revenue net of value-added taxes (VAT), in euros.
- Capital (K_{it}): Deflated total book value of fixed assets, in euros.
- Labor expenditure ($W_{it}L_{it}$): Total labor expenditure, consisting of gross wages and all other labor expenses, such as employers' mandatory social contributions, in euros.
- Labor (L_{it}): Full-time equivalent employment rounded to the nearest integer.
- Intermediate input expenditure ($P_{it}^{M}M_{it}$): Total expenditure on intermediate goods, energy, and other intermediate expenses, in euros.

[40] Available at https://www.oecd.org/en/data/datasets/structural-analysis-database.html.

[41] Documentation is available at https://ec.europa.eu/eurostat/web/products-manuals-and-guidelines/w/KS-GQ-23-010.

[42] See https://www.cepii.fr/DATA_DOWNLOAD/baci/doc/baci_webpage.html.

- Intermediate inputs (M_{it}): Deflated total expenditure on intermediate goods, energy, and other intermediate expenses, in euros.
- Value added (VA_{it}): Revenue minus intermediate input expenditure, in euros.
- Wage (W_{it}): Average labor expenditure per FTE, $\frac{W_{it}L_{it}}{L_{it}}$.
- Labor share (LS_{it}): Total labor expenditure $W_{it}L_{it}$ divided by total revenue R_{it}.
- Intermediate input share (MS_{it}): Total intermediate input expenditure $P_{it}^{M}M_{it}$ divided by total revenue R_{it}.
- Exchange rate ($\frac{EUR}{RMB\,t}$): Annual-average exchange rate, defined as euros per Chinese renminbi.
- Import indicator ($\mathbb{I}_{it}^{import}$): Indicator equal to one if firm i imports intermediate inputs from China in year t, and zero otherwise.
- Chinese import value ($P_{it}^{import}M_{it}^{import}$): Value of firm i's intermediate input imports from China in year t, in euros.
- China import share (S_{it}^{import}): Value of firm i's intermediate input imports from China in year t as a share of its total intermediate input expenditure $P_{it}^{M}M_{it}$.
- Number of imported varieties (V_{it}^{import}): Number of distinct 6-digit products that firm i imports from China in year t.
- New variety share ($VNew_{it}^{import}/V_{it}^{import}$): Share of distinct 6-digit products that firm i imports from China in year t that it did not import in year $t-1$.
- Import price index (P_{it}^{import}): Value-share-weighted average of the 6-digit-product import unit values of the products firm i imports from China in year t, with weights equal to each product's share in the firm's total China import value.
- Import price index, new varieties ($PNew_{it}^{import}$): Value-share-weighted average of the 6-digit-product import unit values of the products firm i imports from China in year t but did not import in year $t-1$, with weights equal to each product's share in the firm's China import value of those products.
- Import price index, continued varieties ($PCont_{it}^{import}$): Value-share-weighted average of the 6-digit-product import unit values of the products firm i imports from China in both year t and year $t-1$, with weights equal to each product's share in the firm's China import value of those products.
- Employment headcount (HC_{it}): Number of workers employed by firm i in year t.

- Exits (HC_{it}^{exit}): Number of workers employed by firm i in year t that no longer work at the firm in year $t+1$.

- Churn ($HC_{it}^{exit+enter}/HC_{it}$): Number of workers that are new at firm i in year t plus the number of workers employed in year t that leave the firm by year $t+1$, divided by the firm's headcount in year t.

- Base wage (W_{it}^{base}): The base wage consists of regular gross salary (which includes the employee's mandatory pension contribution) and pay for additional hours at the regular rate. Constructed by dividing firm-total base wage from the employee data by L_{it}.

- Extra wage (W_{it}^{extra}): Total labor expenditure per FTE net of the average base wage: $W_{it} - W_{it}^{base}$. It comprises overtime pay at above-regular rates, special cash payments – such as holiday allowances, thirteenth-month payments, bonuses, and profit-sharing payments – wage in kind, employer-side social security contributions, employer-side pension expense including additions to pension provisions, and other personnel costs such as reimbursements, training expenditure, and employee benefits.

- Full-time share ($S_{it}^{fulltime}$): Share of firm i's workers with a full-time contract in year t.

- Permanent share ($S_{it}^{permanent}$): Share of firm i's workers with a permanent contract in year t.

- Tenure ($Tenure_{it}$): Average number of years that firm i's workers in year t have been employed at the firm, measured over 2009 to 2018 and reset to one when a worker switches firms.

- Manager share ($S_{it}^{manager}$): Share of firm i's workers with a managerial job in year t.

- Regular share ($S_{it}^{regular}$): Share of firm i's workers with a regular job in year t.

- Other share (S_{it}^{other}): Share of firm i's workers with another job type in year t – interns, on-call workers, and workers with other special statuses, such as workers employed through a disability scheme.

C.3 Data cleaning

I clean the financial statements data in four steps, closely following the procedure that Gopinath et al. (2017) apply to Bureau van Dijk's AMADEUS balance sheet and income statement data.

First, I require that all necessary variables are present and valid. I drop observations with missing values for revenue, total assets, (in)tangible assets, employment, wages, intermediate

inputs, or depreciation, and observations in which revenue, capital, employment, wages, or intermediate inputs are zero or incorrectly signed. I further drop observations with missing NACE Rev. 2 codes, as well as those for which the NACE Rev. 2 division from the ABR conflicts with the first two digits of the NACE Rev. 2 code from the NFO.

Second, I verify the internal consistency of the balance sheets and income statements by comparing each aggregate to the sum of its components. Specifically, I construct five ratios: (i) the sum of tangible and intangible assets, total shareholdings, long and short receivables, inventories, debtors, and liquid assets, relative to total assets; (ii) the sum of domestic and foreign shareholdings, relative to total shareholdings; (iii) revenue net of wages, intermediate inputs, and depreciation, relative to earnings before interest and taxes (EBIT); (iv) EBIT net of total shareholdings income, interest income and charges, extraordinary income and charges, and other financial results, relative to pre-tax income; and (v) pre-tax income net of corporate taxes and third-party equity, relative to after-tax income. Because of minor rounding errors, these ratios can deviate from one even when the underlying components are correct. I therefore drop only those observations for which a ratio falls below 0.95 or exceeds 1.05. Doing so removes fewer than 30 observations, confirming that the CBS's own consistency checks eliminate virtually all inconsistent reports.

Third, I apply a set of further quality checks. I drop firms that at any point report negative capital, employment, or tangible assets; observations with zero or incorrectly signed total liabilities; and observations with incorrectly signed total shareholdings, long and short receivables, debtors, liquid assets, third-party equity, equalization reserves, provisions, long and short debt, or depreciation. I also drop observations with negative value added, defined as revenue net of intermediate input expenditure. Finally, I trim the top and bottom 0.1 percent of the capital-to-wage, total-assets-to-total-funds, and wages-to-value-added ratios.

Fourth, I restrict the sample for estimation. I keep all 2-digit industries with sufficient data to estimate the revenue function. After estimation, I drop observations with negative revenue elasticities or wedges, along with the top and bottom 1 percent of both input wedge distributions. Table C2 reports the distribution of the 132,722 observations in the final sample by year and 2-digit industry. When adding the worker-level data, I additionally drop workers whose change in log hourly base wages between two years exceeds one, drop observations with an hourly wage below 7 euros (roughly 20 percent below the legal minimum wage) or fewer than 40 hours worked in a year, and trim the top one percent of hourly base wages and hours worked.

Table C1: Industries covered in the final sample

Division	Description
10	Manufacture of food products
13	Manufacture of textiles
17	Manufacture of paper and paper products
18	Printing and reproduction of recorded media
20	Manufacture of chemicals and chemical products
22	Manufacture of rubber and plastic products
23	Manufacture of other non-metallic mineral products
24	Manufacture of basic metals
25	Manufacture of fabricated metal products, except machinery and equipment
26	Manufacture of computer, electronic and optical products
27	Manufacture of electrical equipment
28	Manufacture of machinery and equipment n.e.c.
29	Manufacture of motor vehicles, trailers and semi-trailers
30	Manufacture of other transport equipment
31	Manufacture of furniture
32	Other manufacturing
33	Repair and installation of machinery and equipment

Notes: Description of all 2-digit industries covered by the final sample. Manufacturing covers industries 10 to 33.

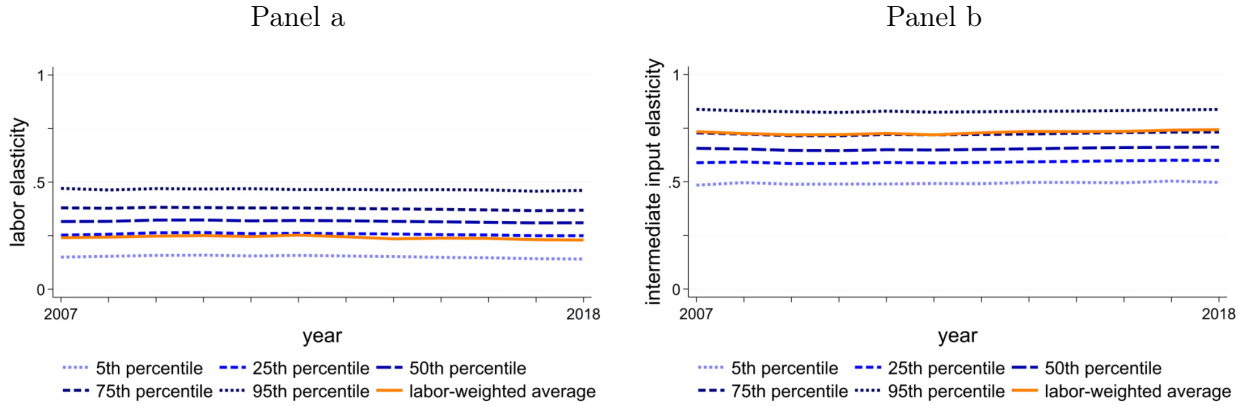
Table C2: Observation count, by year and industry

	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	total
10	907	897	979	991	1,023	1,028	1,048	1,035	1,047	1,060	1,036	1,101	12,152
13	170	176	194	195	193	189	182	178	177	171	170	175	2,170
17	159	147	150	147	156	152	155	153	159	151	148	151	1,828
18	1,141	1,074	1,084	1,087	1,048	979	935	878	838	828	735	746	11,373
20	289	282	291	316	339	346	348	355	341	364	361	374	4,006
22	470	492	492	486	471	471	480	488	500	491	472	522	5,835
23	376	379	395	386	387	367	380	369	368	372	363	390	4,532
24	113	125	128	127	134	129	136	134	141	140	129	135	1,571
25	2,357	2,451	2,564	2,568	2,666	2,634	2,621	2,619	2,637	2,725	2,595	2,817	31,254
26	356	377	419	434	469	455	467	426	444	472	438	475	5,232
27	352	346	383	386	394	385	383	380	379	392	366	391	4,537
28	1,193	1,209	1,302	1,312	1,333	1,344	1,369	1,380	1,383	1,411	1,310	1,411	15,957
29	228	237	249	253	234	230	226	218	217	224	211	215	2,742
30	210	212	235	253	251	255	238	224	223	234	215	239	2,789
31	493	504	545	653	661	655	631	631	617	650	620	679	7,339
32	582	585	614	630	657	657	676	680	676	681	620	696	7,754
33	642	706	783	840	932	992	1,022	1,067	1,109	1,181	1,095	1,282	11,651
	10,038	10,199	10,807	11,064	11,348	11,268	11,297	11,215	11,256	11,547	10,884	11,799	132,722

Notes: Observation count for the full sample, by year and 2-digit industry.

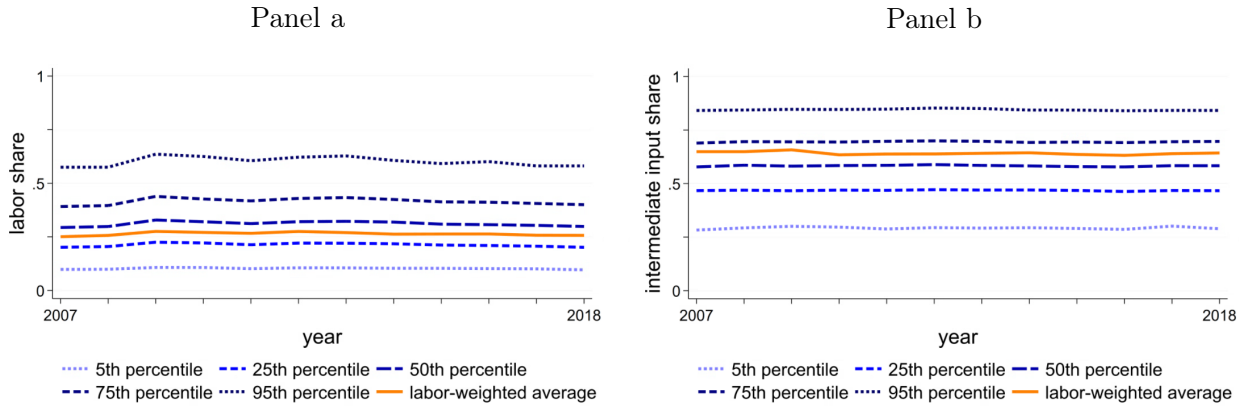
D Additional figures

Figure D1: Distribution of firm-level revenue elasticity of labor (panel a) and revenue elasticity of intermediate inputs (panel b), over time



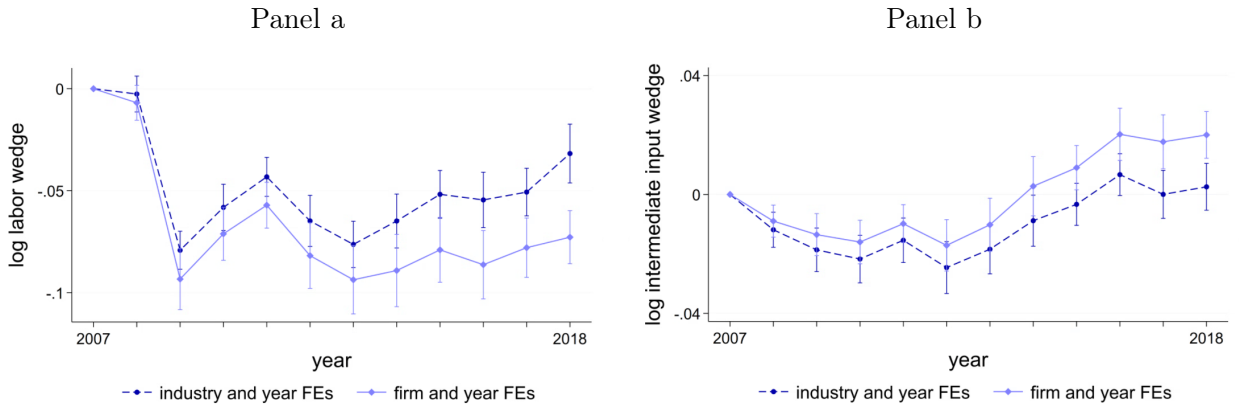
Notes: Selected percentiles of revenue elasticity distributions and labor-weighted average of revenue elasticities, over time. Based on the full sample of 132,722 observations covering the years 2007 to 2018.

Figure D2: Distribution of firm-level labor share of revenue (panel a) and intermediate input share of revenue (panel b), over time



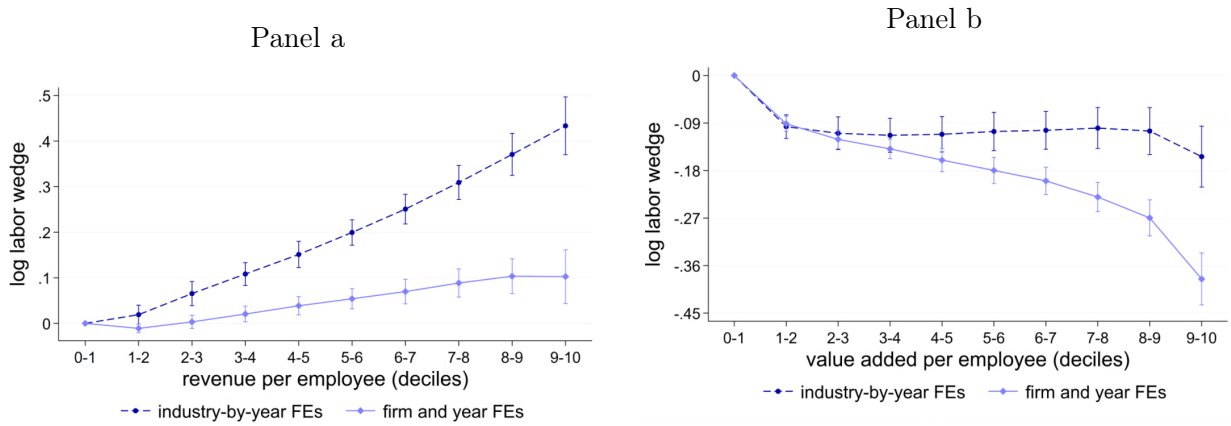
Notes: Selected percentiles of input share distributions and labor-weighted averages of input shares, over time. Based on the full sample of 132,722 observations covering the years 2007 to 2018.

Figure D3: Time-path of firm-level labor wedges (panel a) and intermediate input wedges (panel b)



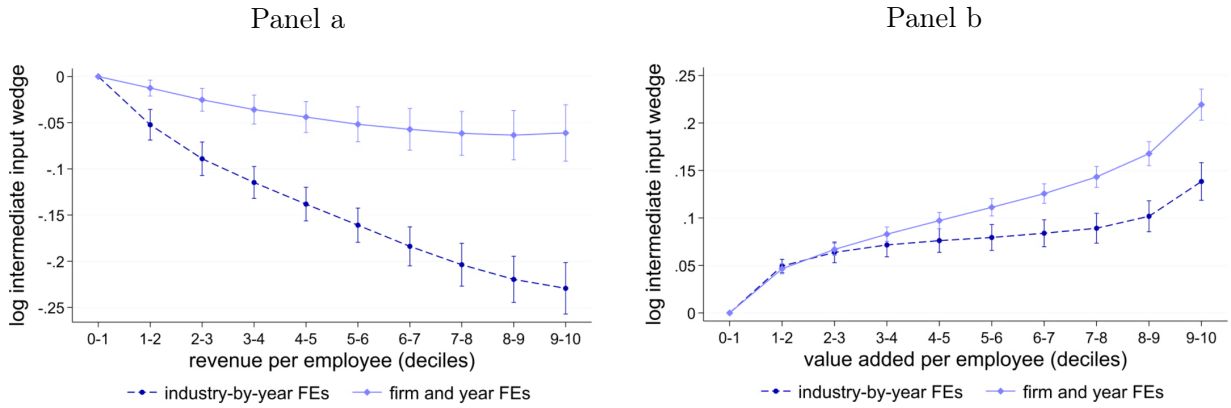
Notes: Point estimates and 95 percent confidence intervals from separate regressions of the log labor wedge (panel a) and log intermediate input wedge (panel b) on year fixed effects. 3-digit-industry (dashed line) or firm fixed effects (solid line) are included. 2007 normalized to 0. Standard errors are clustered at the 3-digit-industry level. Based on the full sample of 132,722 observations covering the years 2007 to 2018.

Figure D4: Regressions of the firm-level labor wedge on revenue per employee (panel a) and value added per employee (panel b)



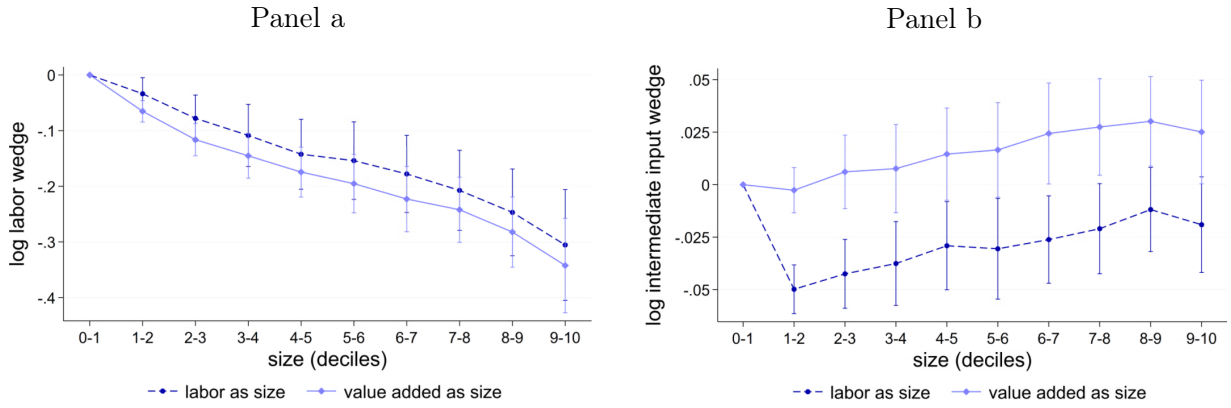
Notes: Point estimates and 95 percent confidence intervals from a non-parametric regression of the log labor wedge on revenue per employee (panel a) and value added per employee (panel b), controlling for employment. Explanatory variables are binned by decile, with the first decile normalized to 0. 3-digit-industry-by-year fixed effects (dashed line) or firm and year fixed effects (solid line) are included. Standard errors are clustered at the 3-digit-industry level. Based on the full sample of 132,722 observations covering the years 2007 to 2018.

Figure D5: Regressions of the firm-level intermediate input wedge on revenue per employee (panel a) and value added per employee (panel b)



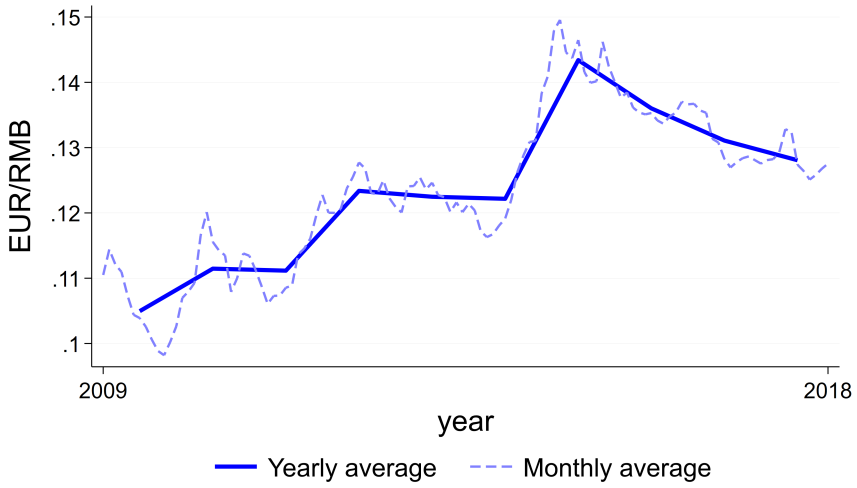
Notes: Point estimates and 95 percent confidence intervals from a non-parametric regression of the log intermediate input wedge on revenue per employee (panel a) and value added per employee (panel b), controlling for employment. Explanatory variables are binned by decile, with the first decile normalized to 0. 3-digit-industry-by-year fixed effects (dashed line) or firm and year fixed effects (solid line) are included. Standard errors are clustered at the 3-digit-industry level. Based on the full sample of 132,722 observations covering the years 2007 to 2018.

Figure D6: Regressions of the firm-level labor wedge (panel a) and intermediate input wedge (panel b) on firm size



Notes: Point estimates and 95 percent confidence intervals from non-parametric regressions of the log labor wedge (panel a) and the log intermediate input wedge (panel b) on firm size measured either by employment (dashed line) or value added (solid line). Explanatory variables are binned by decile, with the first decile normalized to 0. 3-digit-industry-by-year fixed effects are included. Standard errors are clustered at the 3-digit-industry level. Based on the full sample of 132,722 observations covering the years 2007 to 2018.

Figure D7: Exchange rate EUR per RMB, over time



Notes: Yearly and monthly average exchange rate – defined as euro per Chinese renminbi – over time. Yearly exchange rate plotted in the middle of the year.

E Additional tables

Table E1: Median and mean labor and intermediate input wedges, and their correlation, by industry

	Labor wedge		Intermediates wedge		Both wedges	
Industry	Median	Mean	Median	Mean	Correlation	Observations
10 - Food products	0.86	1.02	1.14	1.20	-0.56	12,152
13 - Textiles	0.99	1.10	1.14	1.19	-0.63	2,170
17 - Paper (products)	1.19	1.24	1.05	1.09	-0.56	1,828
18 - Printing and recording	1.08	1.16	1.09	1.14	-0.56	11,373
20 - Chemicals	1.22	1.28	1.08	1.15	-0.49	4,006
22 - Rubber and plastic products	1.08	1.16	1.11	1.13	-0.50	5,835
23 - Non-metallic mineral products	0.89	1.03	1.16	1.21	-0.51	4,532
24 - Basic metals	0.96	1.09	1.16	1.22	-0.61	1,571
25 - Fabricated metal products	0.96	1.10	1.17	1.23	-0.53	31,254
26 - Computer and electronic products	1.09	1.16	1.16	1.22	-0.59	5,232
27 - Electrical equipment	1.00	1.16	1.14	1.19	-0.54	4,537
28 - Machinery and equipment	1.01	1.10	1.15	1.19	-0.60	15,957
29 - Motor vehicles	0.82	0.91	1.13	1.16	-0.37	2,742
30 - Other transport equipment	1.07	1.24	1.12	1.20	-0.50	2,789
31 - Furniture	0.87	0.99	1.16	1.20	-0.56	7,339
32 - Other manufacturing	1.01	1.15	1.28	1.33	-0.63	7,754
33 - Repair and installation of machinery	1.05	1.16	1.13	1.21	-0.60	11,651
Full sample	1.00	1.11	1.14	1.20	-0.54	132,722

Notes: Median and mean labor and intermediate input wedges, as well as their correlation, rounded to two decimal points and by 2-digit industry. Full industry descriptions are in Table C1.

Table E2: Summary statistics of input wedges, by robustness check

	Labor wedge					Intermediate input wedge					
Robustness check	p(5)	p(25)	p(50)	p(75)	p(95)	p(5)	p(25)	p(50)	p(75)	p(95)	Obs
Time-specific revenue function	0.22	0.68	0.89	1.16	1.87	0.93	1.06	1.19	1.39	1.99	101,352
Trade-augmented revenue function	0.64	0.83	1.03	1.32	2.11	0.92	1.03	1.14	1.28	1.67	111,271
CRS-constrained revenue function	0.08	0.33	0.85	1.08	1.55	0.94	1.07	1.25	1.68	2.34	105,930
AR(1) revenue process	0.28	0.48	0.76	1.01	1.58	0.92	1.08	1.23	1.43	1.99	127,898
Calibrated revenue elasticities	0.54	0.76	1.01	1.46	2.91	0.79	0.95	1.13	1.40	2.19	129,844

Notes: Summary statistics of input wedges. p(5), p(25), p(50), p(75), and p(95) refer to the 5th, 25th, 50th, 75th, and 95th percentiles of the input wedge distributions, respectively. Input wedges are rounded to two decimal points. Time-specific revenue function = Input wedges based on 2-year-by-2-digit-industry-specific revenue functions. Trade-augmented revenue function = Includes trade data in the revenue process, input price control, and control function when estimating revenue functions. CRS-constrained revenue function = Imposes that the sum of revenue elasticities equals unity in estimation. AR(1) revenue process = Estimates based on an AR(1) law of motion of unobserved revenue productivity. Calibrated revenue elasticities = Estimates based on calibrated revenue elasticities.

Table E3: Correlation between labor wedge and intermediate input wedge, by robustness check

	Industry-by-year FE		Firm and year FE		
Robustness check	$\hat{\beta}$	s.e.($\hat{\beta}$)	$\hat{\beta}$	s.e.($\hat{\beta}$)	Obs
Time-specific revenue function	-0.944	0.062***	-2.230	0.198***	101,352
Trade-augmented revenue function	-1.329	0.040***	-1.207	0.036***	111,271
CRS-constrained revenue function	-0.363	0.206*	-0.552	0.189***	105,930
AR(1) revenue process	-0.230	0.063***	0.028	0.044	127,898
Calibrated revenue elasticities	-0.900	0.036***	-0.449	0.030***	129,844

Notes: Results from regressions of the labor wedge on the intermediate input wedge. Time-specific revenue function = Input wedges based on 2-year-by-2-digit-industry-specific revenue functions. Trade-augmented revenue function = Includes trade data in the revenue process, input price control, and control function when estimating revenue functions. CRS-constrained revenue function = Imposes that the sum of revenue elasticities equals unity in estimation. AR(1) revenue process = Estimates based on an AR(1) law of motion of unobserved revenue productivity. Calibrated revenue elasticities = Estimates based on calibrated revenue elasticities. All variables are in logs. Standard errors are clustered at the 3-digit-industry level. *** and * indicate statistical significance at the 1% and 10% level, respectively.

Table E4: Persistence and drift of the monthly EUR per RMB exchange rate

	Exchange rate		Exchange rate change
Lagged exchange rate	0.973	0.999	
	(0.019)***	(0.001)***	
Constant	-0.055		0.001
	(0.040)		(0.002)

Notes: Results from AR(1) regressions of the monthly exchange rate (defined as euro per Chinese renminbi) log (Columns 1 and 2), and change of logs of the monthly exchange rates (Column 3), on their first lags. Exchange rates are monthly averages. All regressions use 119 monthly observations from January 2009 through December 2018. Standard errors in parentheses. *** indicates statistical significance at the 1% level.

Table E5: Summary statistics for IV sample

	p(5)	p(25)	p(50)	p(75)	p(95)	observations
Labor wedge	0.63	0.84	1.04	1.32	1.96	10,630
Intermediate input wedge	0.94	1.03	1.12	1.22	1.45	10,630
Chinese import share	0.000	0.002	0.012	0.061	0.326	10,630

Notes: Summary statistics for the sample used in the IV estimates of Section 4; p(5), p(25), p(50), p(75), and p(95) refer to the 5th, 25th, 50th, 75th, and 95th percentiles of the variable's distribution, respectively. Chinese import share = Imports from China divided by total intermediate input expenditure. Input wedges are rounded to two decimal points. Chinese import share rounded to three decimal points.

Table E6: IV estimates of the relationship between the labor wedge and the intermediate input wedge, additional results

	Robustness check				
	Full importer sample	No lagged controls	Excluding exporters	Industry-year fixed effects	NUTS-year fixed effects
	Panel a: Labor wedge				
Intermediate input wedge	-0.608	-0.645	-0.732	-0.591	-0.673
	(0.154)***	(0.115)***	(0.410)*	(0.217)***	(0.208)***
	Panel b: Wage				
Intermediate input wedge	0.693	0.735	1.022	0.731	0.700
	(0.195)***	(0.198)***	(0.273)***	(0.164)***	(0.174)***
	Panel c: Marginal revenue product of labor				
Intermediate input wedge	0.085	0.090	0.290	0.140	0.027
	(0.163)	(0.167)	(0.369)	(0.182)	(0.158)
	Panel d: Intermediate input wedge (first stage)				
$Z^{E/R}$	-0.149	-0.134	-0.089	-0.149	-0.148
	(0.010)***	(0.016)***	(0.010)***	(0.007)***	(0.007)***
F-stat of instrument	219.91	74.50	87.71	467.98	436.40
Firm fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes		Yes
Industry-by-year fixed effects				Yes	
Observations	20,441	10,630	6,185	10,630	10,630

Notes: Results from IV regressions of the labor wedge (Panel a), the wage (Panel b), and the marginal revenue product of labor (Panel c) on the intermediate input wedge, as well as the first stage where the intermediate input wedge is regressed on the exchange rate instrument (Panel d). Full importer sample = Removes requirement that firms import from China in the year and the year prior to the observation. No lagged controls = Omits lagged inputs and revenue as controls. Excluding exporters = Omits all firms that export to China in the year or year prior to the observation. Industry-year fixed effects = Includes industry-by-year fixed effects. NUTS-year fixed effects = Includes NUTS-by-year fixed effects. $Z^{E/R}$ is defined in equation (11). All regressions except the no lagged controls specification include lagged inputs and revenue as controls. All variables in logs except the instrument. Standard errors clustered at the firm level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Table E7: IV estimates of the relationship between the labor wedge and the intermediate input wedge, by robustness check

	Robustness check				
	Time-specific revenue function	Trade-augmented revenue function	CRS-constrained revenue function	AR(1) revenue process	Calibrated elasticities
	Panel a: Labor wedge				
Intermediate input wedge	-0.979	-0.616	-0.473	-0.991	-1.226
	(0.400)**	(0.188)***	(0.136)***	(0.209)***	(0.404)***
	Panel b: Wage				
Intermediate input wedge	0.572	0.689	0.639	0.906	0.667
	(0.150)***	(0.166)***	(0.118)***	(0.284)***	(0.205)***
	Panel c: Marginal revenue product of labor				
Intermediate input wedge	-0.407	0.073	0.166	-0.084	-0.559
	(0.373)	(0.176)	(0.119)	(0.167)	(0.326)*
	Panel d: Intermediate input wedge (first stage)				
$Z^{E/R}$	-0.145	-0.147	-0.144	-0.114	-0.147
	(0.013)***	(0.006)***	(0.007)***	(0.011)***	(0.017)***
F-stat of instrument	119.38	634.09	488.33	109.11	73.24
Firm and year fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	9,452	10,361	7,752	10,238	10,483

Notes: Results from IV regressions of the labor wedge (Panel a), the wage (Panel b), and the marginal revenue product of labor (Panel c) on the intermediate input wedge, as well as the first stage where the intermediate input wedge is regressed on the exchange rate instrument (Panel d). Time-specific revenue function = Input wedges based on 2-year-by-2-digit-industry-specific revenue functions. Trade-augmented revenue function = Includes trade data in the revenue process, input price control, and control function when estimating revenue functions. CRS-constrained revenue function = Imposes that the sum of revenue elasticities equals unity in estimation. AR(1) revenue process = Estimates based on an AR(1) law of motion of unobserved revenue productivity.Calibrated revenue elasticities = Estimates based on calibrated revenue elasticities. $Z^{E/R}$ is defined in equation (11). All regressions include lagged inputs and revenue as controls. All variables in logs except the instrument. Standard errors clustered at the firm level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Table E8: Pre-exposure trends by initial Chinese import exposure

	Intermediate input wedge	Labor wedge	Wage	Marginal revenue product of labor
	Panel a: Without lagged controls			
Initial Chinese import share × 3 years before	0.021	-0.039	-0.103	-0.141
	(0.053)	(0.103)	(0.134)	(0.150)
Initial Chinese import share × 2 years before	0.012	-0.006	-0.089	-0.095
	(0.045)	(0.069)	(0.144)	(0.117)
Joint test p-value	0.924	0.890	0.739	0.620
Observations	5,524	5,524	5,524	5,524
	Panel b: With lagged inputs and revenue as controls			
Initial Chinese import share × 3 years before	0.112	0.018	-0.285	-0.268
	(0.101)	(0.261)	(0.429)	(0.346)
Initial Chinese import share × 2 years before	-0.028	0.007	-0.159	-0.152
	(0.033)	(0.068)	(0.162)	(0.128)
Joint test p-value	0.285	0.995	0.586	0.440
Observations	4,488	4,488	4,488	4,488

Notes: Pre-exposure regressions of the intermediate input wedge (Column 1), the labor wedge (Column 2), the wage (Column 3), and the marginal revenue product of labor (Column 4) on interactions between the initial Chinese import share and indicators for the number of years before the firm's first observed trade year (capped at three and post-2008). Panel a does not include lagged controls, while Panel b includes lagged capital, labor, intermediate inputs, and revenue as controls. One year before the first observed trade year is the omitted category. The joint test p-value is from a test that the coefficients on the interactions for three and two years before the first observed trade year jointly equal zero. All regressions include firm and first-trade-cohort-by-year fixed effects. All dependent variables and controls are in logs. Standard errors clustered at the firm level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Table E9: Reduced-form regressions including future exchange rate change exposure

	Labor wedge	Wage	Marginal revenue product of labor
$Z^{E/R}$	0.110	-0.121	-0.012
	(0.037)***	(0.033)***	(0.031)
$S_{io}^{import}(e_{t+1} - e_t)$	0.011	-0.030	-0.018
	(0.014)	(0.018)	(0.016)
R^2	0.04	0.06	0.08
Observations	10,630	10,630	10,630

Notes: Point estimates and standard errors from regressions of the labor wedge (Column 1), the wage (Column 2), and the marginal revenue product of labor (Column 3) on the exchange rate instrument and future exchange rate change exposure. $Z^{E/R}$ is defined in equation (11). $S_{io}^{import}(e_{t+1} - e_t)$ is the initial Chinese import share interacted with the difference between the log exchange rate realization in year $t+1$ and the log exchange rate realization in year t. All regressions include lagged inputs and revenue as controls and firm and year fixed effects. All dependent variables and controls are in logs. Standard errors clustered at the firm level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Table E10: Reduced form regression on workforce composition adjustments and import status

	$\mathbb{I}_{it}^{import}$	$S_{it}^{fulltime}$	$S_{it}^{permanent}$	$Tenure_{it}$	$S_{it}^{manager}$	$S_{it}^{regular}$	S_{it}^{other}
$Z^{E/R}$	0.008	-0.012	-0.051	-0.010	0.022	-0.019	-0.004
	(0.035)	(0.035)	(0.055)	(0.285)	(0.031)	(0.029)	(0.007)
R^2	0.02	0.01	0.02	0.85	0.02	0.01	0.00
Observations	20,441	10,628	10,628	10,628	10,628	10,628	10,628

Notes: Results from regressions of workforce composition outcomes and import status on the exchange rate instrument. $\mathbb{I}_{it}^{import}$ is an indicator equal to one if firm i imports intermediate inputs from China in year t; S_{it}^{x} denotes the share of firm i's workers belonging to category x, where the categories are full-time, permanent, managerial, regular, and other employment; other employment includes, among others, interns and on-call workers. $Tenure_{it}$ is the average tenure of firm i's workers, measured in years within the 2009-2018 sample period and reset when a worker switches firms. $Z^{E/R}$ is defined in equation (11). All regressions include lagged inputs and revenue as controls and firm and year fixed effects. All continuous variables are in logs except the instrument. Standard errors clustered at the firm level in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.