Wage markups and buyer power in intermediate input markets*

Leonard Treuren†

May 2025

Abstract

How are imperfections in different input markets related? I show how to estimate the wedge between an input's price and its marginal revenue product using financial statements without requiring any input market to be perfectly competitive, and apply this approach to Dutch manufacturing. 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, but not a model that only allows for imperfections in a single input market. In line with this rent-sharing framework, I show that as the Euro appreciates, markdowns and rents in the intermediate input market of firms that import intermediates from China increase, and workers see their wages increase, even though labor does not become 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 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 such imperfections to be linked between input markets – for instance, if inputs are substitutable and firms face upward-sloping supply curves in several markets. Whether this is the case empirically remains an open question, as existing work focuses on one input market at a time.

This paper contributes by studying how market imperfections are related across different input markets. I do this in the context of labor and intermediate input markets in Dutch manufacturing from 2007 to 2018. I first discuss how one might estimate input wedges – the ratio of an input's marginal revenue product to its price – using widely available data on financial statements without requiring any input markets to be perfectly competitive, and then estimate labor and intermediate input wedges in Dutch manufacturing. Prices in the intermediate market are routinely marked down with respect to the marginal revenue product, consistent with buyer power. Both wage markups and wage markdowns with respect to labor's marginal revenue product are frequently observed. Moreover, the two wedges are negatively correlated. I show that these findings can be rationalized by a simple model that combines ingredients of the three strands of literature mentioned above – monopsony, rent sharing, and buyer power – but not by a model that considers imperfections in a single input market. I then show that importers whose intermediates markdowns increase due to exchange rate variations start paying their workers more, in absolute terms and relative to their marginal revenue product, consistent with the rent-sharing mechanisms outlined in the model. My findings imply that, once we depart from perfect competition in input markets, considering imperfections in several input markets can be required to understand price formation in a single input market.

I start by showing that input wedges can be expressed as the revenue elasticity of an input divided by its expenditure share in revenue, and that this measure remains informative in the presence of unobserved input heterogeneity. This approach differs from the predominant 'production approach' to estimating input wedges, which utilizes production function estimation, builds on the insights of Hall (1988) and De Loecker and Warzynski (2012),

[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.

and requires the existence of a variable input for which firms are price takers – typically, intermediate inputs are chosen. The approach that I take does not require the existence of such a variable input. While restrictions on demand and conduct in output markets are required, the same restrictions are (implictly) made by papers that use the production approach with revenue (but not output) data to identify input wedges in a single input market. The revenue approach is, therefore, attractive if one only observes financial statements and is interested in input market imperfections – a standard setting in much of the cited papers that estimate labor wedges. Expenditure shares are observed in commonly available data on financial statements, while revenue elasticities can be estimated using approaches that are well-established in industrial organization.

For my baseline results, following De Loecker (2011) and Petrin and Sivadasan (2013), I recover revenue elasticities by estimating a revenue function using a control function approach.[2] As identifying input wedges using revenue elasticities does not necessitate a control function approach, I verify that my results carry through for a variety of alternative estimation strategies. In particular, I consider an approach in the spirit of Blundell and Bond (2000) that does not rely on the inversion of a control function, and I do away with revenue function estimation altogether by calibrating elasticities of labor and intermediate inputs.

I find that both wage markups and wage markdowns are prevalent in Dutch manufacturing, covering 50 percent of all observations each. This is not in line with monopsony power, 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 14 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. This dispersion is due mainly to idiosyncratic differences rather than between-industry differences or variation over time. A key predictor of labor wedge dispersion is the degree of buyer power a firm has in the intermediate input market – as measured by the intermediate input wedge. The more a firm underpays its input suppliers on the margin, the higher its wages and wage markups. These forces are present both within firms over time and within industry-years.

To rationalize these findings, I combine elements from three theories of individual input market imperfections in a simple model. Like papers that study either monopsony power or buyer power, I assume that firms face upward-sloping supply curves in the labor market and intermediates market, respectively. From work on rent sharing, I borrow the assumption that wages are determined by bargaining between firms and their collectively organized workers – consistent with the Dutch setting and Western Europe more generally. If employees have no bargaining power, the model predicts that wages are marked down relative to the marginal revenue product of labor. If employees do have bargaining power, wages are above

[2] See Olley and Pakes (1996), Levinsohn and Petrin (2003), and Ackerberg et al. (2015).

monopsonistic levels as firms share part of their rents net of payments to suppliers of capital and intermediates. If those rents are sufficiently high, for instance, because firms have considerable buyer power for intermediates, wages exceed the marginal revenue product of labor – wage markups occur. These theoretical findings are consistent with the empirical results on input wedges and require input market imperfections in both input markets. In such a setting, the labor supply elasticity is no longer sufficient for measuring wage markdowns, as markdowns also depend on the firm's rents generated in other input markets.

The stylized theoretical framework suggests that one mechanism underlying the joint distribution of the two input wedges is that firms share rents generated in the intermediate input market with their employees. To test that prediction, I set out to find variation in such rents and quantify how they translate to wages and labor wedges. I focus on firms that import intermediate input from China and use variation in the degree to which firms are exposed to exchange rate changes due to firms' previous reliance on Chinese imports. Utilizing within-firm variation, I show that as the Euro appreciates, the intermediate input wedge increases, and the labor wedge decreases. Instrumenting the intermediates wedge by exposure to exchange rate fluctuations, I estimate that a 1-percent increase in the intermediate input wedge results in a 0.66 percent decrease in the labor wedge. This can be explained by increased value added per employee and higher wages, in line with rent-sharing, while employment and average or marginal revenue productivity are largely unchanged.

This paper contributes to several strands of literature. First, the literature on monopsony power.[3] Yeh et al. (2022) estimate plant-level labor wedges for US manufacturing 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.[4] In Germany, both wage markups and wage markdowns occur. However, wage markdowns are concentrated in large and highly productive firms (Mertens, 2023), and are markedly less likely to occur where collective bargaining agreements cover employees (Dobbelaere et al., 2024). As China's export demand grows, German firms are found to exercise more labor market power (Mertens, 2020).[5] I contribute by highlighting that imperfections in the intermediates market can influence wages and measures of labor market power, and by showing that large-scale monopsony concerns are likely unwarranted in Western European labor markets characterized by collective bargaining.

The second strand of literature related to this paper studies rent sharing (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

[3] 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.

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

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

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. (2023) find that product market power mitigates the incentive to mark down wages in the US construction industry. I contribute by linking rent sharing to buyer power for intermediate inputs. That is, existing work focuses on product market rents that are shared with employees, while I find evidence of redistribution of rents across different input markets consistent with buyer power for intermediates mitigating the incentive to mark down wages.

The third strand of related work is an 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. All these papers allow for imperfect competition in output markets, but are not interested in labor market imperfections. I contribute by studying the relationship between buyer power and labor market imperfections and by providing additional evidence that buyer power is much more common than previously thought.[6]

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). Angenhofer et al. (2025) 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 different type of countervailing power – one operating through other input markets. As buyer power for intermediate input increases, workers get paid more, which could negate monopsonistic markdowns. The broader implication is that antitrust policy aimed at preventing the build-up of purchasing power in one input market might come at the cost of input suppliers in 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 – the results of such an exercise are usually interpreted as output elasticities (De Loecker and Syverson, 2022). Authors typically observe revenue and possibly capital, labor, and intermediates inputs, but not output. A long line of work considers the conditions under which output elasticities are

[6]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.

indeed identified in this setting (e.g., Klette and Griliches (1996); Bond et al. (2021)). I argue that the restrictions on demand and conduct that are (implicitly or explicitly) made to estimate production functions in this context are sufficient for a revenue function to exist and be multiplicatively separable in a part that depends on output and a part that depends on unobserved shocks to demand and output. This allows one to apply existing techniques from production function estimation to a revenue function, which has the advantage that input wedges can be estimated without assuming that firms are price takers in any of their input markets – in contrast to production-function-based approaches to input wedge estimation.

The remainder of this paper proceeds as follows. Section 2 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. In Section 4, I construct a parsimonous model of input market inperfections that can rationalize my results, which suggests that one factor underlying the joint distibution is that firms share rents generated in the intermediate input market with their workers. This hypothesis is tested in Section 5. Concluding remarks follow in Section 6.

2 Empirical approach

This section discusses the identification of the firm-time-specific wedge between an input's marginal revenue product and its price – the 'input wedge' – using firm-level data on financial statements that are widely available. In particular, I consider a setting where a researcher does not observe output, but does observe revenue and expenditure on (and possibly quantity of) several input bundles, such as labor, intermediate inputs, and capital. Such data is routinely used in industrial organization and macroeconomics, among other fields.[7] In anticipation of my application, I focus on labor and intermediate inputs, but the approach can be applied to other inputs as desired.

2.1 Identifying input wedges

That typical datasets on financial statements only contain information on revenue and bundled input use or expenditure creates several challenges. To measure input wedges in this context, I make two assumptions that are used to prove Lemma 1.

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

[7]See De Loecker and Syverson (2022) for a discussion of such data and research on productivity more broadly. Related work using such data studies 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., Levinsohn and Petrin (2003); Petrin and Sivadasan (2013)), and the US (e.g., Yeh et al. (2022)), among many other countries.

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 with as entries the amount of labor for each variety in firm i's labor variety set $\mathcal{L}_{it}$ and $A^L(\cdot)$ is a CRS aggregator which 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 firms select $\mathbf{l}_{it}$ and $\mathbf{m}_{it}$.

Assumption 1 rules out revenue functions where inputs are perfect complements (i.e., Leontief). Assumption 2 helps confront the fact 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 $\theta^L_{it} = \frac{\partial R_{it}}{\partial L_{it}}\frac{L_{it}}{R_{it}}$ and $\theta^M_{it} = \frac{\partial R_{it}}{\partial M_{it}}\frac{M_{it}}{R_{it}}$ 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, $\frac{P^M_{it}M_{it}}{R_{it}}$.

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

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

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,

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

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

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

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. ■

While straightforward, Lemma 1 provides guidance on estimating and interpreting input market imperfections. In terms of interpretation, Lemma 1 states that the ratio of the observed input bundle's revenue elasticity to its revenue share is informative in the presence of unobserved input heterogeneity, as it equals the weighted average of the underlying varieties' input wedges. Note that the ratio of an input's revenue elasticity to its share of revenue equals that input's wedge in the absence of unobserved heterogeneity. Henceforth, I will refer to γ_{it}^L as the labor wedge, to γ_{it}^M as the intermediate input wedge, and to both as input wedges.[8]

Lemma 1 states that two things are required to estimate an input wedge: the input's expenditure share in revenue, which is routinely observed in financial statements, and the input's revenue elasticity, which needs to be estimated. It is important to note that any combination of demand and a production function that satisfies Assumption 1 in principle permits the identification of input wedges along the lines of Lemma 1. Any further restrictions placed on demand, conduct, or technology, therefore, are due to the method used to estimate the revenue function, not the approach described in this subsection.

The vast majority of existing work relies on production function estimation to identify input wedges so that a brief comparison of both approaches is in order. This 'production approach' to input wedge estimation builds on first-order conditions from cost minimization, requires information on output elasticities, and has primarily been used to identify the labor wedge.[9] This approach uses insights from the markup estimation methodology of De Loecker and Warzynski (2012) to identify γ_{it}^L. To understand the production approach to labor wedge estimation, first denote a firm's markup of price over marginal cost by $\mu_{it} = \frac{P_{it}}{\lambda_{it}}$. For any input V_{it}, let

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

where $\tilde{\theta}_{it}^V$ is the output elasticity of input V_{it} and P_{it}^V is its price. Under the assumptions that inputs are substitutable, V_{it} is frictionlessly adjustable, firms are price takers in the market for V_{it}, and firms select V_{it} to minimize their conditional cost function, De Loecker and Warzynski (2012) show that $\mu_{it}(V_{it}) = \mu_{it}$.[10] The choice of V_{it} is crucial in recovering the true markup μ_{it}. In particular, if firms are not price takers for V_{it}, $\mu_{it}(V_{it})$ is a joint measure of the markup and the input wedge of V_{it}: $\mu_{it}(V_{it}) = \mu_{it}\gamma_{it}^V$. This insight has led researchers to estimate the labor wedge and interpret it as a measure of labor market power by comparing markup estimates obtained using different inputs as V_{it}. Most recent work seeks to identify the labor wedge by comparing $\mu_{it}(L_{it})$ to $\mu_{it}(M_{it})$. In particular,

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

[8]Similar results to Lemma 1 have been derived for input wedges based on output elasticities instead of revenue elasticities (Morlacco, 2020).

[9]For example, see, 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.

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 intermediate inputs are frictionlessly adjustable, and firms are price takers in the intermediates market, profit maximization implies that $MRPM_{it} = P_{it}^M$, so that $\gamma_{it}^M = 1$ and equation (3) identifies the labor wedge. In general, however, the production approach obtains the labor wedge *relative* to the wedge of another input, in this case, intermediates.

The main advantage of the approach outlined in Lemma 1 over the production approach is that it does not require a frictionlessly adjustable input for which firms are price takers to identify input wedges. For my setting, this is essential, as I aim to investigate imperfections in the labor and intermediate input markets, so I am not willing to assume away imperfections in those markets. The one remaining input on which I have data is capital, which seems a particularly unsuitable candidate for a frictionlessly adjustable input. Of course, if one has data on an input for which the absence of any market imperfections is plausible, the production approach would suffice to estimate input wedges of other inputs and deliver an estimate of the markup.[11]

2.2 Estimating revenue elasticities

The revenue elasticities of labor and intermediate inputs are required to identify input wedges using Lemma 1. Obtaining such elasticities can be done in various ways, and the appropriate approach will differ based on the context and available data. One could simply calibrate the revenue elasticities, an approach I discuss at this subsection's end. For my baseline estimates and several robustness checks, however, I estimate a revenue function.

My starting point is a very large body of work that uses similar data to mine to estimate production functions (in addition to papers that I cited earlier, see De Loecker and Syverson (2022) for a review). As output data is missing, such research uses deflated revenue in its place, which is well-known to result in biased estimates if there is price variation around the industry-wide deflator (e.g., Klette and Griliches (1996)). As elaborately discussed in De Loecker and Goldberg (2014), regressing log (deflated) revenue on log inputs only delivers output elasticities if one (either implicitly or explicitly) restricts demand and conduct. In particular, to move beyond a setting without any price variation whatsoever, it is necessary to (implicitly) assume that firms are in monopolistic competition and face identical constant elasticity of substitution (CES) demand. This is commonly done by authors who rely on the production approach to estimate the labor wedge (e.g., Yeh et al. (2022)). I now show that specifying a revenue function is straightforward under those same restrictions.

[11] Another alternative approach, particularly suited to industry studies with more elaborate data, is to explicitly model and estimate input supply. See Rubens (2023) for a discussion.

Consider a translog production function

$$q_{it} = \tilde{\beta}_k k_{it} + \tilde{\beta}_{kk} k_{it}^2 + \tilde{\beta}_l l_{it} + \tilde{\beta}_{ll} l_{it}^2 + \tilde{\beta}_m m_{it} + \tilde{\beta}_{mm} m_{it}^2 + \tilde{\beta}_{kl} k_{it} l_{it} + \tilde{\beta}_{km} k_{it} m_{it} + \tilde{\beta}_{lm} l_{it} m_{it} + \tilde{\omega}_{it} \tag{4}$$

where lowercase letters denote the natural logarithm of the concomitant uppercase letter. Following De Loecker (2011), I consider a CES demand function of the form

$$P_{it} = P_{st} \left(\frac{Q_{it}}{Q_{st}} \right)^{\frac{1}{\eta_s}} \Xi_{it}, \tag{5}$$

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 an unobserved (to me) demand shock. The demand shock consists of industry-year, 3-digit industry g, and firm-year specific components: $\Xi_{it} = \Xi_{st} \Xi_g \tilde{\Xi}_{it}$. The resulting revenue function is translog

$$r_{it} = \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + \beta_{kk} k_{it}^2 + \beta_{ll} l_{it}^2 + \beta_{mm} m_{it}^2 + \beta_{kl} k_{it} l_{it} + \beta_{km} k_{it} m_{it} + \beta_{lm} l_{it} m_{it} + \gamma_{st} + \omega_{it} + \epsilon_{it}, \tag{6}$$

where $\omega_{it} = \frac{1}{\eta_s} \tilde{\omega}_{it} + \tilde{\xi}_{it}$ is unobserved revenue productivity, γ_{st} subsumes all industry-specific terms, and ϵ_{it} is introduced to capture mean-zero unanticipated deviations from planned revenue. Revenue elasticities are given by $\theta_{it}^x = \frac{\partial r_{it}(\cdot)}{\partial x_{it}}$, where $x \in \{k, l, m\}$ and, like output elasticities, are not restricted to sum to one. Although unimportant for the empirical implementation, note that the revenue elasticity of an input depends on that input's output elasticity and the elasticity of demand, $\theta_{it}^x = \frac{\eta_s + 1}{\eta_s} \tilde{\theta}_{it}^x$. Finally, note that while I estimate (6) at the 2-digit industry level in my baseline estimates following convention, revenue elasticities are firm-time specific due to the translog specification.[12]

Note that the revenue function that this demand system generates can also be derived from alternative demand systems, including a (nested) logit specification that is common in empirical demand estimation (see Section 6.2 of De Loecker (2011)). Also note that, compared to authors estimating production functions with deflated revenue, demand function (5) relaxes the assumption that firms face identical demand at the st level. While an extension to multiproduct firms is conceptually straightforward, this requires identifying which firms are multiproduct and the number of products they produce, neither of which I observe.

The primary identification challenge is controlling for unobserved (by the econometrician) revenue productivity captured in ω_{it}. If a firm's inputs at time t are at least partially determined by decisions made after the firm observes ω_{it}, estimates of the revenue elasticities can be biased. A similar concern arises in the estimation of production functions, where ω_{it}

[12] Utilizing 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 while keeping the estimation manageable.

would contain solely unobserved physical productivity. Indeed, if $\tilde{\xi}_{it}$ is i.i.d. across firms and time, then revenue productivity is physical productivity scaled by a constant. The two main approaches used in that literature to address the issue are i) methods that can allow for ω_{it} to evolve according to a general first-order Markov process but restrict competition or demand (e.g., Olley and Pakes (1996); Levinsohn and Petrin (2003); Ackerberg et al. (2015); Gandhi et al. (2020)), and ii) methods that assume that ω_{it} evolves according to an AR(1) process (e.g., work applying the insights of Arellano and Bond (1991) and Blundell and Bond (2000)).

In line with the few existing papers that estimate revenue functions, I rely on the first approach in my baseline estimates. In particular, both De Loecker (2011) and Petrin and Sivadasan (2013) apply a control function approach, where the latter uses the estimates to identify input wedges. De Loecker (2011) explicitly assumes that $\tilde{\xi}_{it}$ is i.i.d. across firms and time, but I will discuss the conditions under which the same approach permits serially correlated firm-specific demand shocks. Pham (2023), instead, relies on insights of Gandhi et al. (2020) to estimate a revenue function and quantify labor wedges. As this approach requires assuming away imperfections in the intermediate input market, $\gamma_{it}^{M} = 1$, it defeats the purpose of my paper, so I do not employ it.[13] All three papers assume that inputs are substitutable, as do I.[14]

I now briefly outline my baseline estimation routine, a more detailed description of which is in Appendix B. I follow De Loecker (2011) and Petrin and Sivadasan (2013) by using demand for intermediate inputs to proxy for ω_{it},

$$m_{it} = m_s(\omega_{it}, k_{it}, l_{it}, \mathbf{x}_{it}^{m}; \delta^{m}), \tag{7}$$

where, as before, s is the industry, $\mathbf{x}_{it}^{m}$ contains control variables, and δ^{m} contains all coefficients. If ω_{it} is the sole unobservable in equation (7) and $m_s(\cdot)$ is one-to-one in ω_{it}, inverting (7) results in a control for ω_{it} based on observables which can be substituted into equation (6). As ω_{it} contains the markup, physical productivity, and the demand shock, one does not need to control for these factors in (7), in contrast to proxy variable approaches to production function estimation.

The invertibility of input demand (7) is discussed in detail in Appendix B with parametric examples of intermediate input demand. For now, note that because the one-to-one requirement is on the joint revenue productivity term, and because $\tilde{\Xi}_{it}$ and $\tilde{\Omega}_{it}$ enter the revenue function multiplicatively, all that is required is that intermediate input demand is one-to-one in the product of the current realization of the demand and productivity shocks. This is

[13] Gandhi et al. (2020) raises identification issues for gross output production functions if output and input markets are perfectly competitive, and provides a solution under the same set of assumptions. Estimating a revenue function in the presence of imperfect competition allows one to sidesteps these issues.

[14] Hashemi et al. (2022) discuss how markup estimates based on revenue data might actually reflect input wedges, and suggest the use of revenue function estimation to identify input wedges.

satisfied for revenue function (6). Therefore, demand shock $\tilde{\Xi}_{it}$ can be serially correlated.[15]

The main concern related to ω_{it} being the sole unobservable in equation (7) is that intermediate input prices, which I do not observe, might vary between firms. As I deflate intermediates' expenditure with 2-digit-by-year deflators, the concern lies with input price differences at that level. Following De Loecker et al. (2020), I control for input price variation due to unobserved quality differences by using market shares and industry indicators. Location-based input price variation is controlled for by including indicator variables for the NUTS1 region where the firm is located. See De Loecker et al. (2016) for the theoretical underpinnings of this input price control function. Additionally, to control for within-industry intermediate input price differences stemming from buyer power, I include a firm's intermediate input share in its industry in $\mathbf{x}_{it}^{x}$, which controls for buyer power in a variety of models (see Appendix B). If, nonetheless, input price variation would remain, I find it unlikely that this would significantly impact my results, as my main results are robust to estimation approaches that do not require the inversion of intermediates demand – see the discussion at the end of this section.[16]

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

$$b_{it} = b_s((1, k_{it}, l_{it}, m_{it}) \times \mathbf{x}_{it}^{b}; \delta^{b}), \tag{8}$$

where the notation indicates that inputs enter $b_s(\cdot)$ interacted with the control variables in $\mathbf{x}_{it}^{b}$ – due to the translog specification – and δ^{b} contains all coefficients. The variables used to control for input price differences are as described in the previous paragraph. Input price biases are well-known but mostly ignored in production function estimation and papers that estimate input wedges. Note that a control function for input prices can at least partially control for input price variation even if it is misspecified so that its inclusion is always preferred (De Loecker et al., 2016).

To identify coefficients θ, I use a two-step estimation approach based on Ackerberg et al. (2015). In the first step, log revenue is regressed on a polynomial in the three inputs, the arguments of (8), and the arguments of the inverse of equation (7), $\omega_{it} = m_s^{-1}(k_{it}, l_{it}, m_{it}, \mathbf{x}_{it}^{m}; \delta^{m})$.[17] This first step is not meant to identify any coefficients, but rather

[15] If the revenue function is not log additive in the two shocks, satisfying invertibility is less straightforward. Dhyne et al. (2022) discuss how invertibility can be gauranteed more generally if one observes (at least) one proxy variable for each unobservable.

[16] That papers estimating input wedges or markups often ignore input price bias or use a control function that does not take input price variation into account has some basis in the literature, as De Ridder et al. (2025) show that a simple control function already allows for informative estimation of markup trends over time, and Collard-Wexler and De Loecker (2015) find that omitting the input price control tends to not matter much in their application.

[17] I use a third-degree polynomial in all variables except for indicator variables and time trends, which are added linearly.

to separate observed revenue into planned revenue and random deviations from planned revenue: $r_{it} = \hat{r}_{it} + \hat{\epsilon}_{it}$. This is important, as I need $\hat{\epsilon}_{it}$ to correct revenue shares when constructing the input wedges. Revenue shares need to be corrected as $R_{it}(\cdot)\exp(\epsilon_{it})$ is observed in the data, but firms base their decisions on planned revenue $R_{it}(\cdot)$. This is done by replacing the revenue shares in equations (1) by $\hat{LS}_{it} = \frac{W_{it}L_{it}}{R_{it}/\exp(\hat{\epsilon_{it}})}$ and $\hat{MS}_{it} = \frac{P^M_{it}M_{it}}{R_{it}/\exp(\hat{\epsilon}_{it})}$.

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

$$\omega_{it} = g_s(\omega_{it-1}; \delta^g) + \psi_{it}, \tag{9}$$

where ψ_{it} is a mean zero revenue productivity shock and δ^g contains all coefficients. De Loecker (2011) explicitly assumes that $\tilde{\xi}_{it}$ is i.i.d. over firms and time so that the Markov assumption relates only to physical productivity $\tilde{\omega}_{it}$, but note that the demand shock can be serially correlated, as long as the joint term ω_{it} has the first-order Markov property. Using $\omega_{it} = \hat{r}_{it} - f_{it}(\cdot;\theta) - b_{it}(\cdot;\delta^b)$ together with the law of motion of revenue productivity allows me to obtain an estimate of the revenue productivity shock, $\hat{\psi}_{it}$, conditional on the still to be estimated coefficients.[18] Estimates are based on the following moment conditions

$$\mathbb{E}(\psi_{it}\mathbf{Z}_{it}) = 0, \tag{10}$$

where I focus on the exactly identified case and $\mathbf{Z}_{it}$ contains a constant, and all contemporaneous values of variables, with the exception of labor, intermediate inputs, and $\hat{r}$, which are instrumented for using their first lag. These moments are based on the idea that labor and intermediate inputs are flexible and might, therefore, adjust to revenue productivity shocks within a period, while capital is a state variable that is given within a period (e.g., $K_{it} = (1-d)K_{it-1} + I_{it-1}$ with d denoting depreciation and I_{it-1} investments). Estimates of revenue elasticities are reported in Table B1 of Appendix B.

I refer to results based on the estimation approach described above as the “baseline specification”. I consider three alternative approaches to obtaining the revenue elasticities – as before, discussed in more detail in Appendix B.

First, I estimate the revenue function at the 2-digit-industry-by-2-year level. While the approach outlined in Lemma 1 does not require Hicks neutral revenue productivity, this restriction is imposed when estimating a revenue function. Considerable attention has recently been devoted to labor-augmenting technological change (e.g., Doraszelski and Jaumandreu (2018); Raval (2023); Demirer (2025)). Papers in the production approach that allow for labor-augmenting technological change rule out monopsony power by assumption, while papers that focus on monopsony power rule out factor-biased technology shocks by assumption (Rubens et al., 2024). Separating the two phenomena requires explicitly modeling labor market imperfections – an approach typically not taken by papers in either strand of literature.

[18] I approximate $g_s(\cdot)$ by a third-degree polynomial and $b_s(\cdot)$ by a second-degree polynomial in all their arguments except indicator variables and time trends, which are added linearly.

While conceptually, I follow the Hicks-neutral tradition, the translog revenue function used in the baseline specification allows for firm-time-specific revenue elasticities, as elasticities are a function of a firm's inputs. However, it restricts the function of those inputs to be fixed over time. The 2-digit-industry-by-2-year level estimates allow the firm-specific revenue elasticities to vary flexibly over time not only because input use changes but also because the derivatives of the revenue function with respect to input changes.

Second, in the spirit of Blundell and Bond (2000), I estimate the revenue function under the assumption that the law of motion of revenue productivity is an AR(1) process: $\omega_{it} = \rho\omega_{it-1} + \psi_{it}$. One can substitute for ω using the revenue function in equation (6), to arrive at the following moment conditions

$$\mathbb{E}((\psi_{it} + \epsilon_{it} - \rho\epsilon_{it-1})\,\mathbf{Z}_{it}^{AR1}) = 0, \tag{11}$$

where the difference between $\mathbf{Z}_{it}^{AR1}$ and $\mathbf{Z}_{it}$ is that variables have been lagged one (additional) time due to the moments including ϵ_{it-1}. The main advantage of this approach over the baseline specification is that ω_{it} is not proxied for using observables, so the invertibility of equation (7) is not required. The main drawback is that the AR(1) law of motion is more restrictive than equation (9).

Finally, I do away 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 value of the baseline estimates in the entire sample. The revenue elasticity of intermediate inputs is calibrated similarly.[19] By Lemma 1, variation in input wedges is then entirely driven 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. For several reasons, the Netherlands is a suitable setting to study market imperfections in labor and intermediates markets. First, the Netherlands is representative of Western Europe in that collective bargaining characterizes labor markets, which might overturn findings of wage markdown found in the US, where such bargaining is largely absent (Yeh et al., 2022).[20] Indeed, the Dutch bargaining system is similar to that of countries such as Italy and Portugal, which have been studied through the lens of

[19] 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.

[20] 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.

collective bargaining before (e.g., Card and Cardoso (2022); Card et al. (2014)).[21] In the Dutch intermediates market, collective agreements are absent, so interactions between firms and their suppliers are bilateral. Intermediate input markets are fairly concentrated, and Dutch firms heavily rely on imports.[22] Recent work identifies buyer power for imported intermediates in the French manufacturing sector, raising concerns that such power might also exist in the comparable Dutch context (Morlacco, 2020).

I construct a yearly firm-level dataset covering Dutch manufacturing firms over the period 2007 to 2018, using non-public data obtained from Statistics Netherlands (CBS).[23] I combine data from the "General Firm Registry" (ABR) and the "Financial Statistics of Non-financial Firm" (NFO). These two yearly firm-level datasets are based on registry data from the Dutch Chamber of Commerce, the Dutch tax authority, and the Dutch Ministry of Finance. The ABR and NFO aim to document the universe of all non-financial firms located in the Netherlands. The ABR contains yearly data on each firm's full-time equivalent (FTE) employment, the 4-digit NACE industry in which the firm is active, and the location of the firm's headquarters. The NFO contains yearly balance sheets and income statements from which I obtain revenue, the total expenditure on labor and intermediate inputs, and the book value of capital.

The CBS routinely checks data quality and contacts firms when reporting errors are suspected. In addition, I restrict the sample to firm-year observations with sufficient information to construct labor and intermediate input wedges. In particular, only observations with positive revenue, capital, intermediate input expenditure, and at least one FTE employee on the payroll are included. The final sample consists of 23,638 firms for a total of 132,722 firm-year observations. Appendix A provides a detailed overview of all variables and the sample selection procedure. Table A1 in Appendix A lists the 18 2-digit industries that are included in the final sample, and Table A2 presents a breakdown of observations by year and 2-digit industry.

I use total labor expenditure to construct my labor compensation variable, as total labor expenditure captures a firm's labor cost more accurately than employees' net or gross salary. I construct the wage W_{it} as the ratio of a firm's total expenditure on labor to its FTE employment. All results are valid only for employees on a firm's payroll, as 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 total labor expenditure in the Netherlands. This gap is due to mandatory social security contributions, such as pension

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

[22] 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.

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

Table 1: Summary statistics

variable	p(25)	p(50)	p(75)	mean	s.d.
Revenue	666	1,687	4,675	10,823	112,000
Capital	97	383	1,168	3,475	43,688
Labor (FTE)	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. EBIT is earnings before interest and taxes. p(25), p(50), and p(75) refer to the 25th, 50th, and 75th percentile 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.

contributions by both employers and employees, and high income tax rates.[24]

Table 1 reports summary statistics for key firm-level variables. Revenue and inputs are positively skewed, with means often substantially exceeding medians. The outliers in the right tales ensure that standard deviations are much larger than the interquartile range. 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. The presence of small firms does not drive my results, which are robust to the omission of smaller firms. However, dropping firms with only a handful of employees substantially reduces the sample, leading to convergence issues in several 2-digit industries when estimating the revenue function. This issue is compounded in robustness checks where I estimate the revenue function on smaller groups of observations. Therefore, I use the full sample for estimation and control for size throughout the paper. Total labor cost has a (rounded to thousands) median of 505,000 euros and a mean of 1,942,000 euros. The median book 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 inputs share at 0.58.

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

Figure 1: Kernel density function (panel a) and heatplot (panel b) of firm-level labor and intermediates wedges

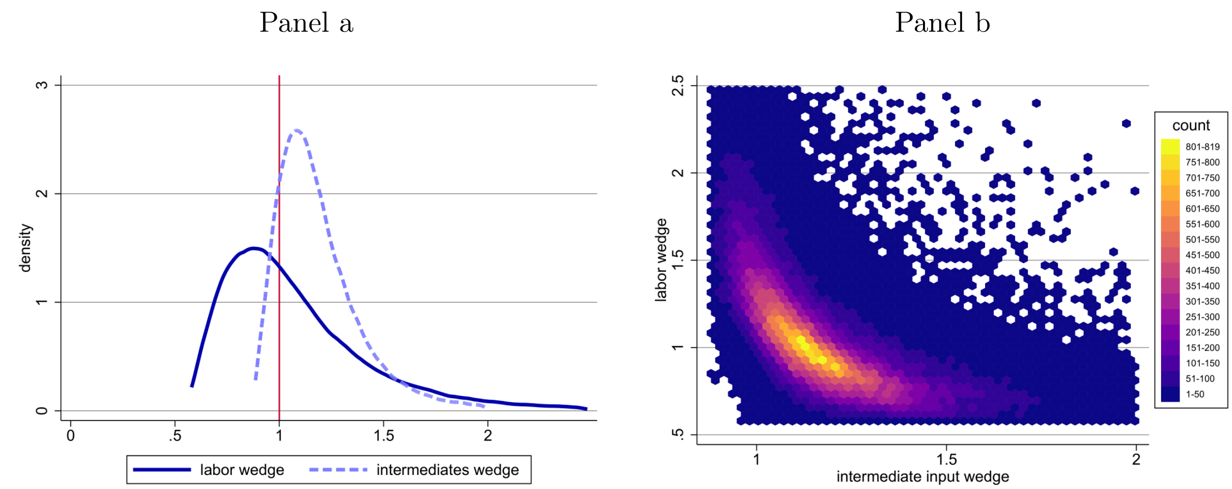

Notes: Based on the full samples of 132,722 observations covering the years 2007 to 2018 (top and bottom 2 percentiles of the wedges trimmed).

3 Results

In Section 3.1, I analyze the distributions of firm-level input wedges in Dutch manufacturing, as well as their correlation, both in the cross-section and over time. Section 3.2 explores idiosyncratic factors underlying the relationship between the intermediate input wedge and the labor wedge.

3.1 Evidence of wage markups and buyer power

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. In 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.

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

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.

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.

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.[25]

Panel b of Figure 1 contains a heat 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. 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 markdown in intermediate input markets and pay their employees somewhat more compared to their

[25]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.

marginal revenue product than smaller firms.

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

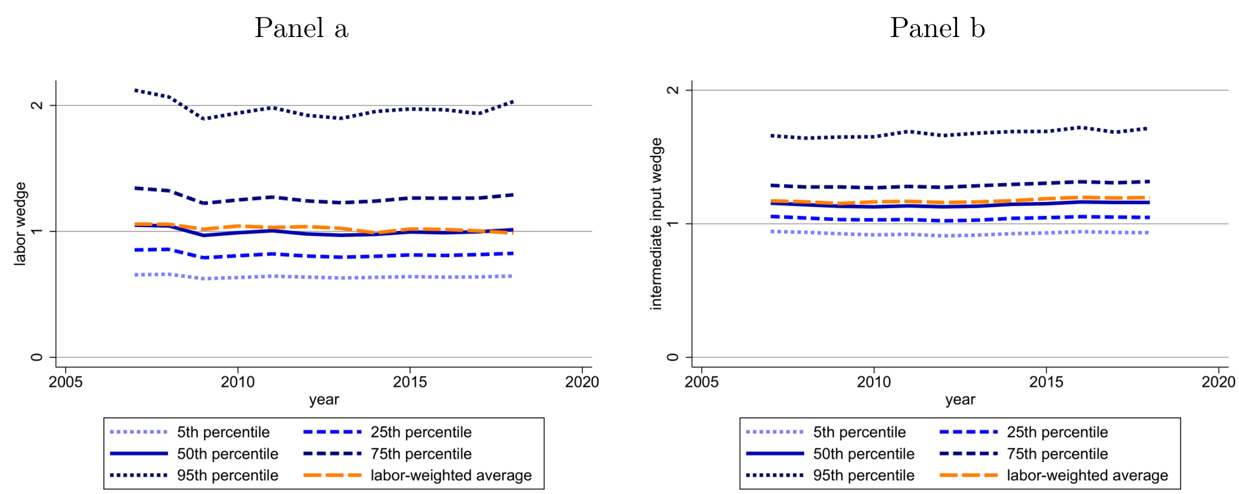

Notes: Selected percentiles of input wedge distributions (blue lines) and labor-weighted average of input wedges (orange lines), over time. Based on the full samples of 132,722 observations covering the years 2007 to 2018.

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 shocks. Figure 2 plots selected percentiles of the two input wedge distributions over time. Both distributions are remarkably stable over the sample period. Figure 2 does not imply that there is no time-series variation at all. The median labor wedge, for instance, drops slightly after the financial crisis and increases back to pre-crisis levels by 2018. 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 D1 in Appendix D 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. In the remainder of this paper, I focus on the joint distribution of the two input wedges rather than between-industry differences for a single wedge, as documenting the relationship between input market imperfections is the main contribution of this paper.

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 is 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 132,772 observations covering the years 2007 to 2018.

Summing up, 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. 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. 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, the following non-parametric regression is used to correlate the labor wedge with a variable of interest, x_{it},

$$ln(\gamma_{it}^L) = \beta_0 + \sum_{dc=2}^{10} \beta_{dc}^x \mathbb{I}_{x_{it} \in X_{dc}} + \mathbf{Controls}_{it} + \varepsilon_{it}, \tag{12}$$

where $\mathbb{I}_{x_{it} \in X_{dc}}$ is an indicator that equals 1 if x_{it} lies between the dc^{th} and $dc - 1^{th}$ 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 flexibly. 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 might contribute to the aggregate distribution of the labor wedge.

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 produce of labor (MRPL) (panel b). Both within industry and within firm, wages are strongly, positively, and monotonely correlated with the labor wedge. In addition, firms that have higher labor wedges have higher marginal revenue products of labor. Within firm, 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 C1 in Appendix C 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 and within year. 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 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

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

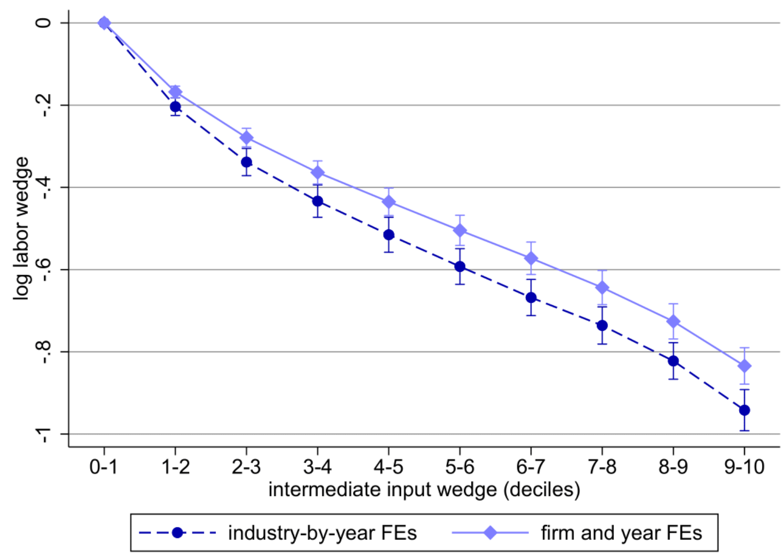

Notes: Point estimates and 95 percent confidence intervals from a non-parametric regression of the log labor wedge on the intermediate input wedge, controlling for employment; Explanatory variables are divided into 10 categories based on the deciles of their distribution, and the coefficient on the indicator of the first decile is normalized to 0; Either 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.

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 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 firm, 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.

Panel a of Figure C2 in Appendix C 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.

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

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

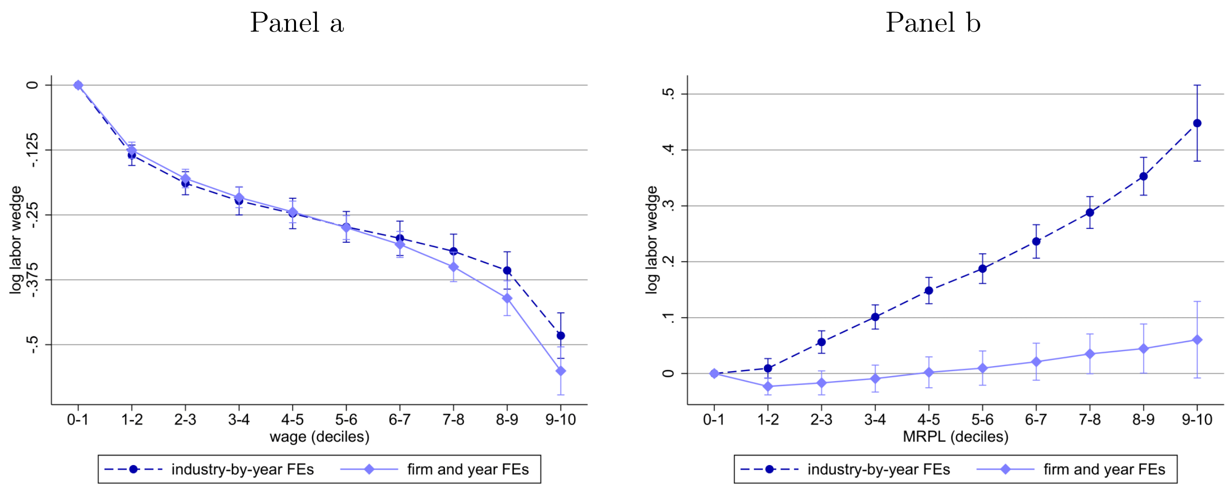

Notes: Point estimates and 95 percent confidence intervals from a non-parametric regression of the log labor wedge on the wage (panel a) and the marginal revenue product of labor (panel b), controlling for employment; Explanatory variables are divided into 10 categories based on the deciles of their distribution, and the coefficient on the indicator of the first decile is normalized to 0; Either 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.

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. What can explain these patterns?

Before moving on to an explanation for the observed patterns, I first stress that they are robust to alternative approaches to obtaining revenue elasticities. Table D2 in Appendix D 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 1.19 with a time-specific revenue function. Both wage markups and wage markdowns are common in all specifications. However, with a time-specific revenue function or an AR(1) revenue productivity process, the location of the distribution moves to the left – at the median wages are 12 to 18 percent above the marginal revenue product of labor. The reason is that the labor elasticities derived with these approaches are more similar to the OLS estimates – lower – than the control function estimates. Table D3 in Appendix D verifies that the two input wedges are negatively correlated both within firm and within industry-year for all estimation approach of revenue elasticities. I conclude that where the assumptions of a par-

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

Panel a   Panel b

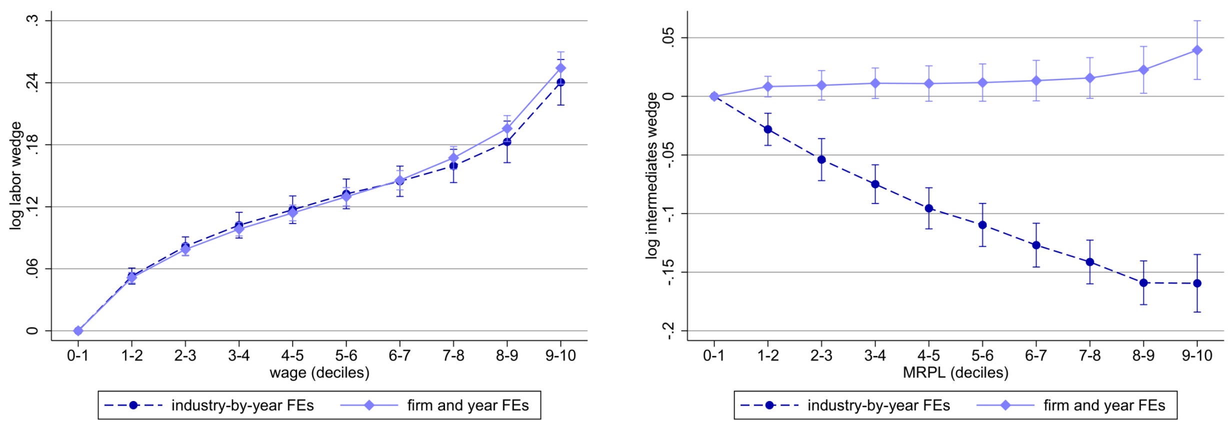

Notes: Point estimates and 95 percent confidence intervals from a non-parametric regression of the log intermediate input wedge on the wage (panel a) and the marginal revenue product of labor (panel b), controlling for employment; Explanatory variables are divided into 10 categories based on the deciles of their distribution, and the coefficient on the indicator of the first decile is normalized to 0; Either 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.

ticular estimation approach might be violated, this does not appear to affect my qualitative conclusions.

4 A simple model of input market imperfections

In this section, I present a stylized theoretical framework that combines key ingredients from existing work studying imperfections in labor and intermediate input markets. As I will not estimate or calibrate the theoretical model and aim for simplicity in exposition, I remain agnostic about many underlying features – such as the sources of upward sloping supply curves. The goal is not to suggest that this stylized model generates the only mechanism underlying the joint distribution of input wedges that I described in Section 3. Rather, I aim to show that combining standard models from separate strands of the literature studying input market imperfections generates predictions that can rationalize my findings, whereas a model from a single strand would fail to do so.

I combine assumptions from the three strands of literature mentioned in the introduction. First, from an extensive literature that studies rent sharing in (European) labor markets that abstracts from upward-sloping supply curves in input markets, I borrow the assumption that firms bargain with a union to determine wages (e.g., Card et al. (2018)). Second, from the literature on monopsony power in labor markets, I borrow the assumption that the labor supply curve slopes up (e.g., Yeh et al. (2022)). Finally, from papers that study buyer power in intermediate input markets, I borrow the assumption that firms also face an upward-sloping supply curve in the intermediates market (e.g., Morlacco (2020)). Combining all three generates a setting where wages can be marked up or down relative to labor's marginal revenue product, depending on the size of the firm's rents, which are, in turn, increasing a firm's buyer power for intermediate inputs.[26]

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. I assume that capital is predetermined, so the price of capital P_{it}^K can be interpreted as the flow cost of capital.

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 $x \in \{K_i, L_i, M_i\}$.[27] As a consequence, profit is $\Pi_i(K_i, L_i, M_i) = R_i(K_i, L_i, M_i) - P_i^K K_i - W_i L_i(W_i) - P_i^M M_i(P_i^M)$.

Intermediate input price determination Each firm plays a two-stage game. In Stage 1, wages and employment are determined by collective bargaining. In Stage 2, given the outcomes of Stage 1, firms select their intermediate inputs.[28]

Firm i selects M_i to maximize profit in Stage 2, which results in first-order condition

$$P_i^M = \frac{\varepsilon_i^M}{\varepsilon_i^M + 1} MRPM_i. \tag{13}$$

Equation (13) states that 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. First-order condition (13) implicitly defines the solution to Stage 2, conditional intermediate input price $P_i^M(L_i)$ and the corresponding quantity $M_i(L_i)$.

[26] For notational simplicity and following related work that I borrow from, I abstract from input heterogeneity and suppress the time subscript.

[27] A simple example where all assumed properties hold is the combination of demand and production specified in Section 2.2.

[28] As bargaining-induced adjustment frictions might make labor less easily adjustable than intermediate inputs, I assume wage and employment have been determined before intermediates are purchased, although all results in this section are unaffected by this timing assumption.

Wage determination In line with the European context, wages are determined by collective bargaining between workers and their employer. As is customary in the literature, I assume that workers collectively aim to maximize their employer's wage bill (W_iL_i) and that bargaining is efficient in the sense that wages are selected to maximize total rents conditional on product demand, labor supply, and the conditional intermediate input price and quantity that solve Stage 2.[29] 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

$$\max_{W_i} \Big(W_iL_i(W_i)\Big)^{\phi}\Big(\Pi_i\big(K_i, L_i(W_i), M_i(L_i(W_i))\big)\Big)^{1-\phi}, \tag{14}$$

where ϕ refers to workers' bargaining power.

The first-order condition can be written as

$$W_i = \phi\frac{R_i}{L_i}\left(1 - \frac{\tilde{\theta}_i^M}{\mu_i}\frac{P_i^M}{MRPM_i} - \frac{\tilde{\theta}_i^K}{\mu_i}\frac{P_i^K}{MRPK_i}\right) + (1-\phi)\frac{\varepsilon_i^L}{\varepsilon_i^L+1}MRPL_i, \tag{15}$$

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. Equation (15) shows that the wage is a weighted average of two familiar cases. If $\phi = 0$, workers have no bargaining power, and the wage equals the monopsonistic markdown with respect to the marginal revenue product of labor. If firms have no bargaining power ($\phi = 1$), all the residual rents that are bargained over accrue to the workers – residual rents being the rents per worker after suppliers of capital and intermediates have been paid. The size of these residual rents depends on several factors: the wedge between the price of intermediates and capital and their marginal revenue product, the output elasticities of those inputs, and the price-cost markup in the output market.[30] In particular, the more firms can mark down the price of intermediate inputs relative to their marginal revenue product, the larger the share of revenue per employee that workers receive, all else equal. In other words, buyer power for intermediates puts upward pressure on wages and redistributes rents from intermediate input suppliers to workers if $\phi > 0$ (and to firms if $0 < \phi < 1$).

To investigate the relation between wedges in different input markets, I rewrite equation

[29] Card et al. (2018) surveys papers that study rent sharing in labor markets using a similar collective bargaining framework. Wong (2023) combines this framework with an upward-sloping labor supply curve as I do but assumes that firms are price takers in other input markets.

[30] If firms have no bargaining power and there is perfect competition in output and input markets other than labor, suppliers of intermediates and capital receive the revenue shares of those inputs, and each worker receives a fraction $(1 - \frac{P_i^M M_i}{R_i} - \frac{P_i^K K_i}{R_i})$ of revenue per employee.

(15) as

$$\frac{W_i}{MRPL_i} = \phi \frac{R_i}{\tilde{\theta}_i^L} \left(1 - \frac{\tilde{\theta}_i^M}{\mu_i} \frac{P_i^M}{MRPM_i} - \frac{\tilde{\theta}_i^K}{\mu_i} \frac{P_i^K}{MRPK_i} \right) + (1-\phi) \frac{\varepsilon_i^L}{\varepsilon_i^L + 1}. \tag{16}$$

Equation (16) shows that the wedge between wage and the marginal revenue product of labor is also a weighted average of two extreme cases. If workers have no bargaining power – i.e., if workers are wage takers due to the absence of collective agreements – wages are marked down with respect to $MRPL_i$. However, if workers have at least some bargaining power, firms share part of the rents that they generate with their workers by way of higher wages. If the residual rents are large enough, for instance, because firms substantially mark down the price of intermediates relative to their marginal revenue product, wages could even exceed the marginal revenue product, so that wage markups, rather than wage markdowns, result. Note that allowing for only a single input market imperfection – a single upward-sloping supply curve or bargaining between workers and the firm – results in a model that can no longer rationalize the patterns presented in Section 3.

Summing up, the stylized theoretical framework predicts that intermediates prices will be marked down relative to their marginal revenue product if firms have buyer power. If workers collectively bargain with their employers to determine wages and an upward-sloping labor supply curve pins down labor, whether wages are marked up or down relative to the marginal revenue product of labor depends on the extent of the firm's rents. Increases in firms' rents, for instance, due to increased buyer power for intermediates, put upward pressure on wages in absolute terms, and relative to the marginal revenue product of labor. Of course, indirect effects could create additional links between buyer power and wages, as in equilibrium, changes in buyer power typically coincide with changes in firms' input use. Hence, the exact magnitude of the relation between imperfections in labor and intermediate input markets is an empirical question that depends on the particulars of underlying features of the model.

Combining assumptions from three strands of literature on input market imperfections has provided a potential mechanism underlying the joint distribution of the two wedges described in Section 3: firms share rents generated in the intermediate input market with their employees. To test this further, I next look for factors that affect rents in the intermediate input market and measure how wages and the labor wedge respond. Before moving on, observe from equation (15) that rents generated in output and capital markets also create upward pressure on wages, as is documented by existing work on rent sharing. I focus on the link between imperfections in the intermediate input market and the labor market as this aligns with the empirical evidence presented in Section 3 and as this link is the main contribution of this paper.

5 How exchange rate fluctuations affect employees

To test the predictions of the stylized theoretical framework introduced in the previous section, I approximate the relationship between the two input wedges by the following log-linear regression model

$$\ln(\gamma_{it}^{L}) = \beta_1 \ln(\gamma_{it}^{M}) + \beta_c' \mathbf{X}_{it} + \varepsilon_{it}, \tag{17}$$

where $\mathbf{X}_{it}$ is a vector of controls including 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. Including firm fixed effects not only helps condition on time-invariant differences between firms, but is also fitting as the rent sharing mechanism described in the previous section is a within-firm phenomenon. Even so, estimating equation (17) with OLS might yield biased estimates of β_1, as both input wedges are jointly determined in equilibrium. This joint determination makes it difficult to sign the bias of these coefficients, as any factor that affects input use in one input market might alter expenditure and wedges in both input markets. By identifying a shifter of intermediate input market rents, however, I aim to isolate one instance of how rents in the intermediate input market affect labor markets.

To isolate variation in exposure to shocks to the intermediate input market, I measure 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 likely being most exposed. For such rotations to generate a worthwhile sample, two conditions need to be met by a candidate country of origin: there needs to be meaningful exchange rate variation, and a non-negligible number of firms need to import from said country. China is the leading candidate that satisfies these requirements, as most Dutch imports originate in countries whose currency is the Euro or something (essentially) pegged to the Euro. To conduct this analysis, I obtain data from the "International Trade in Goods" (IHG) dataset from Statistics Netherlands. This data is available from 2009 onward and contains product-level value of imports by country of origin, but not output or unit prices.

I define the following instrument that interacts a firm's exposure to exchange rate fluctuations with the exchange rate realization in a given year

$$Z_{it}^{E/R} = \ln\left(IS_{io}^{china}\right) \ln\left(\frac{EUR}{RMB_t}\right), \tag{18}$$

where $\frac{EUR}{RMB_t}$ is the exchange rate defined as Euro per Chinese Yuan Renminbi, and IS_{io}^{china} is the share of imports from China in overall intermediate input expenditure in the first year that a firm shows up in the IHG data. The $Z_{it}^{E/R}$ notation is meant to alert readers to the fact that the value of the instrument decreases if the Euro appreciates – i.e. if the intermediates supply curve rotates out. Following the terminology of Borusyak and Hull (2023), my instrument measures non-random exposure (the shares) to an exogenous shock (the exchange rate realization). As said, the intuition is that firms that previously imported more from China are more likely to be affected by a given exchange rate fluctuation. The year a firm first imports from China in the IHG data is not used in estimation, other than to calculate the initial shares, for concerns that the contemporaneous import share might be endogenous to the input wedges. Note that my fixed effects specification implies that I control for the initial share, following Borusyak and Hull (2023), as it is time-invariant.

If conditional on all controls, the instrument's value is assigned as good as randomly, a regression of an outcome on the instrument has a causal interpretation. In my setting, this amounts to exchange rate fluctuations being random, as the initial share is controlled for. Indeed, it is well known that a random walk tends to forecast exchange rates better than economic models (Rossi (2013) surveys the field). Combining regressions of the labor wedge on the instrument – the 'reduced form' – and the intermediate input wedge on the instrument – the 'first stage' – provides an instrumental variable (IV) regression of γ_{it}^L on γ_{it}^M, where the latter is instrumented for. For the IV regression to have a causal interpretation, addition assumptions need to hold, in particular, the exclusion restriction: the instrument should influence the labor wedge only through the intermediates wedge, conditional on the controls, or $\mathbb{E}\left[\sum_i \sum_t Z_{it}^{E/R} \varepsilon_{it} | \mathbf{X}_{it}\right] = 0 \ \forall (i,t)$.

The main threats to the validity of the exclusion restriction are any direct effects of exchange rate fluctuations on labor markets and indirect effects operating through product markets. Note that a large body of work studies the effects of China opening up to trade on manufacturing in Western countries and finds that manufacturing employment and wages declined (Autor et al. (2016) surveys the literature). Likewise, if the Euro appreciates relative to the Chinese Yuan, Chinese demand – and firms' rents – should decrease, if change at all. However, the rent sharing theory developed in the previous section would predict that wages (in absolute terms and relative to marginal revenue product) *increase* as importing from China becomes cheaper. Reassuringly, therefore, violations of the exclusion restriction that appear most likely work against the hypotheses that I am attempting to test.

As not all firms import from China and the trade data is only available from 2009 onward, the sample I use for the IV estimation is considerably smaller than the full sample described in Section 3. In addition, for an observation to be in the sample, the firm needs to import from China, import from China the year before, and import from China in one earlier year to construct the initial share. This implies that I use within-firm variation to ask how firms that import from China in year $t-1$ and year t react to the exchange rate realization in year t, conditional on lagged input use, lagged revenue, and year (and firm) fixed effects. In total, this leaves 1,631 firms covering 8,469 observations. Table D4 in Appendix D reveals that the distribution of the two wedges in this reduced sample is very similar to their full-sample counterparts in Table 2. Moreover, the table 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 1.5 percent, while at the 75th and 95th percentiles, this share is 7.8 percent and 36.2 percent, respectively. Figure C3 in Appendix C shows that the exchange rate between the Netherlands and China varies considerably and non-monotonically over the years that are used in the estimation. At its highest point, in 2015, the exchange rate was 30 percent above its lowest value, which occurred in 2011. I conclude that both ingredients of the instrument display substantial variation.

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			-0.660	-1.574
			(0.263)**	(0.058)***
$Z^{E/R}$	0.098	-0.148		
	(0.041)**	(0.005)***		
Firm and year fixed effects	Yes	Yes	Yes	Yes
R^2	0.06	0.13	0.36	0.52
F-stat of instrument		764.14	764.14	
Observations	8,469	8,469	8,469	8,469

Notes: Table 4 reports 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 (18); All regressions include lagged inputs and revenue as controls; All variables in logs except the instrument; Standard error clustered at the firm level in parentheses; ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Table 4 presents IV estimates related to equation (17). 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 0.688. Such a firm would increase its intermediates markdowns by 10.18 percent and reduce its labor wedge by 6.74 percent. Note also that the F-stat of the instrument in the first stage regression is 764.14, implying that the instrument is highly relevant.

Column 3 gives IV estimates of equation (17). A one percent increase in the intermediate input wedge leads to firms reducing the labor wedge by 0.66 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 8.11 percent – i.e., its workers would be paid less relative to their marginal revenue product. Column 4 shows OLS estimates of (17) for reference. As anything that affects one input market can influence both wedges, signing the bias of the OLS estimates a priori is difficult. However, it turns out that the OLS estimates overstate the relationship between the two wedges that is induced by my instrument.

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		0.627	1.851		-0.033	0.277
		(0.181)***	(0.102)***		(0.176)	(0.075)***
$Z^{E/R}$	-0.093			0.005		
	(0.028)***			(0.026)		
Firm and year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.07	0.24	0.37	0.08	0.09	0.10
F-stat of instrument		764.14			764.14	
Observations	8,469	8,469	8,469	8,469	8,469	8,469

Notes: Table 5 reports 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 (18); All regressions include lagged inputs and revenue as controls; All variables in logs except the instrument; Standard error clustered at the firm level in parentheses; ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

The effect on the labor wedge of the intermediate input wedge changing is entirely due to wage changes. 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 7.7 percent. In contrast, the OLS estimates suggest that both wages and marginal revenue products increase, but wages by more (Columns 3 and 6).

To investigate the mechanisms underlying the estimates in Tables 4 and 5, Table 6 runs a series of reduced-form regressions relating several outcomes to the instrument. Columns 1 to 4 display that, as importing from China becomes cheaper, intermediates expenditure drops, while revenue, the number of employees, and, therefore, revenue per employee are unchanged. This corresponds to a setting where firms do not adjust their inputs in response to exchange rate changes that alter the intermediates price fluctuate. As the year-on-year exchange rate fluctuations seem too moderate to justify adjusting the production process, this is unsurprising. Note that, if the intermediates supply curve rotates out but input use is unaffected, the intermediates wedge should increase, which is indeed the case by the first stage regression.

Table 6: Reduced form estimates to shed light on adjustment mechanisms

	$P^M_{it}M_{it}$	L_{it}	R_{it}	$\frac{R_{it}}{L_{it}}$	$VADD_{it}$	$\frac{VADD_{it}}{L_{it}}$
$Z^{E/R}$	0.230	-0.003	0.049	0.052	-0.099	-0.097
	(0.080)***	(0.022)	(0.066)	(0.052)	(0.037)***	(0.024)***
Firm and year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.29	0.37	0.37	0.10	0.43	0.12
Observations	8,469	8,469	8,469	8,469	8,466	8,466

Notes: Table 6 reports results from regressions of several outcomes on the exchange rate instrument; $Z^{E/R}$ is defined in equation (18); All regressions include lagged inputs and revenue as controls; All variables in logs except the instrument; Standard error clustered at the firm level in parentheses; ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

These results suggest that exclusion restriction violations caused by exchange rate fluctuations impacting labor markets directly or through output markets are quantitatively unimportant. Factors impacting the output market likely affect revenue, which does not change. For labor markets, the primary reported direct effects of access to cheap imports from China are outsourcing or offshoring parts of the production process, particularly the low-skill and low-wage parts (e.g., Gu et al. (2022)). Such forces would impact the number of employees that I observe, as this measure does not include employees employed abroad, and would likely increase expenditure on intermediates. Finally, recall that the labor wedge that I measure can be interpreted as an average of underlying labor wedges by Lemma 1. Changes in the estimated wedge could be driven by (some combination of) changes in the sourcing strategy, the distribution of expenditure across different types of workers, or the underlying wedges. As changing the composition of employees is likely to affect employment, or if it does not affect employment, at least to influence the revenue per employee if any skill differences exist between different groups of employees, I conclude that the changes in the labor wedge that I observe most likely originate in changing wages of the different types of workers.

Columns 5 and 6 of Table 6 show that total and per-employee value added increases as importing from China becomes cheaper and that the magnitude of these effects is near-identical to the effect of the instrument on the labor wedge. Note that value added per worker is routinely used as a measure of firm rents by researchers studying rent sharing – see Card et al. (2018). Gathering results, we have that expenditure on intermediates falls, and markdowns of intermediate prices increase, as importing intermediates becomes cheaper. As labor and revenue do not change, value added per employee increases. In turn, workers get paid more even though their marginal and average revenue productivity does not change.

I conclude that the data offer support for firms sharing intermediate input market rents with their workers being one of the forces underpinning the joint distribution of the two input wedges. Hence, input market imperfections are related across input markets. As employment does not change while wages increase following an appreciation of the Euro, workers benefit from their employer's market power in the intermediates market in the context that I consider. Therefore, the data suggest that buyer power for intermediates partially shields workers from monopsony power. One implication of these findings is that reduced form estimates of labor supply elasticities are not sufficient statistics for wage markdowns, as a model that produces sufficiency rules out other input market imperfections. A second implication is that a policy that aims to reduce market power in one input market might hurt suppliers in other input markets.

Table D5 in Appendix D verifies that these conclusions are not driven by how revenue elasticities are estimated, how the shares are constructed, or output market effects, by repeating the IV estimates for the labor wedge, wage, and marginal revenue product of labor. Column 1 restricts the sample firms to firms that do not export to China, so they are not directly exposed to output market effects of exchange rate fluctuations. These firms increase wages by more than the sample of all firms that import from China, while marginal revenue products are not significantly affected by the intermediates wedge. The negative effect of the intermediate input wedge on the labor wedge is also slightly larger in this specification, suggesting that violations of the exclusion restriction that operate through output markets are quantitatively small and likely imply that the estimates in this section are a lower bound of the true effect. Column 2 verifies that my results are qualitatively robust to replacing the initial shares with lagged shares when constructing the instrument – although the magnitude of these wage and labor wedge effects is substantially larger, potentially signaling that lagged shares better proxy exposure than initial shares. Finally, Columns 3 to 5 verify that my results are qualitatively robust to using wedges based on the calibrated revenue elasticities, elasticities estimated based on an AR(1) revenue process, and time-specific revenue functions, respectively.

6 Concluding remarks

This paper studies how market imperfections in labor and intermediate input markets are related in Dutch manufacturing from 2007 to 2018. I provide evidence of buyer power for intermediate inputs. In the labor market, wages are commonly marked up or marked down with respect to the marginal revenue product of labor. Moreover, the two wedges are negatively correlated in the cross-section and within a firm over time. I show that idiosyncratic factors, primarily wages, drive variation in wage markups. These results can be rationalized by a model where firms face upward-sloping supply curves in both input markets and bargain over wages with their collectively organized employees. I find further empirical support for this model's predictions by studying how intermediate input market rents generated by exchange rate fluctuations affect the workers of firms that import intermediates from China.

The ability of this paper to simultaneously study market imperfections in labor markets and intermediate input markets is a result of my reliance on revenue function estimation. Compared to the production approach to identifying labor wedges, this approach does not require a variable input that is frictionlessly adjustable and for which firms are price takers. While the existence of a revenue function that can be estimated with widely used techniques in industrial organization requires one to place restrictions on demand and conduct, the same (and often stronger) restrictions are already made by authors who estimate production functions using revenue data.

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. In the setting studied in the paper, for instance, the labor supply elasticity is no longer a sufficient statistic for the extent to which wages are marked down. 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. In addition, studies using matched employer-employee data could unpack the firm-average labor wedges estimated in this paper into separate labor wedges of different types of employees and help pinpoint the employee-level determinants of whether a firm has labor market power over an employee. I leave these 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.

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.

Angenhofer, T., Collard-Wexler, A., and Weinberg, M. (2025). Oligopsony and collective bargaining. *mimeo.*

Arellano, M. and Bond, S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies*, 58(2):277—-297.

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

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.

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.

Collard-Wexler, A. and De Loecker, J. (2015). Reallocation and technology: Evidence from the US steel industry. *American Economic Review*, 105(1):131–171.

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., Eeckhout, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics*, 135(2):561–644.

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.

De Ridder, M., Grassi, B., and Morzenti, G. (2025). The hitchhiker's guide to markup estimation. *mimeo*.

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.

Dhyne, E., Petrin, A., Smeets, V., and Warzynski, F. (2022). Theory for extending single-product production function estimation to multi-product settings. (no. w30784) National Bureau of Economic Research.

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.

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.

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.

Kroft, K., Luo, Y., Mogstad, M., and Setzler, B. (2023). Imperfect competition and rents in labor and product markets: The case of the construction industry. (no. w27325). National Bureau of Economic Research.

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.

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. (2024). Exploiting or augmenting labor? *mimeo*.

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

Wong, H. C. (2023). Understanding high-wage firms. *mimeo*.

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

Appendices

A Data

A.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. These data are available to academic researchers, subject to conditions.[31]

The "General Firm Registry" (ABR) is a yearly firm-level dataset aiming to document all firms in the Netherlands. The ABR uses registry data of the Dutch Chamber of Commerce and the Dutch tax authority as primary sources.[32] The CBS uses daily updates from the Chamber of Commerce, and monthly updates from the tax authority, to construct anonymized firm identifiers. From the ABR, I take data on the total 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 firm-level dataset containing anonymized balance sheets and income statements of all identifiable firms active outside the financial sector. The NFO uses two primary data sources, depending on whether the firm in question is classified as "small" or "large".[33] For small firms, data is obtained each year from the Dutch Ministry of Finance, which documents data on balance sheets and income statements from tax returns. For large firms, data is obtained every year using surveys. Each firm receives a survey that is extensively checked for consistency with the data previously obtained by CBS. The firm is contacted again to verify the information if a reporting error is suspected. The cleaning procedure documented in this appendix is intended to eliminate any reporting errors that might have survived this process, particularly for small firms. From the NFO, I take data on revenue, expenditure on labor and intermediate inputs, the book value of capital, and earnings before interest and taxes.

In the process of anonymization, the CBS creates new firm identifiers at different levels of aggregation. The balance sheet and income statement data of the NFO come at the "organization group" (OG) level, which is considered to be the "actual agent in financial processes" (CBS, 2020b). The ABR is at the "firm unit" (BE) level, which is an "autonomous actor in the production process" (CBS, 2020c). For the vast majority of firms, there is a one-

[31] See https://www.cbs.nl/en-gb/our-services/customised-services-microdata for details.

[32] In particular, the "Nieuwe Handelsregister" (NHR) of the Chamber of Commerce and the "Beheer van Relaties" (BvR) of the Dutch tax authority are used. In addition, prior to April 1st, 2014, the "Basis Bedrijvenregister" (BBR) was employed, a partnership between the Chamber of Commerce, the tax authority, and the CBS.

[33] Before 2011, all firms with a balance sheet total of less than 23 million euros were classified as small. As of 2011, all firms with a balance sheet total of less than 40 million euros are classified as small.

to-one mapping from BEs to OGs, but for the largest firms – the TOPX, containing roughly 2000 firms each year – one OG can hold more than one BE. For these OGs, I aggregate ABR data to the OG level at which the balance sheets and income statements are available.[34] Firm identifiers at the OG level allow me to merge the NFO and the ABR.

Eurostat's NACE Rev. 2 industry classification is used throughout this paper. The most aggregated industry classification is the NACE section (one or more 2-digit NACE codes), while NACE divisions (2-digit), groups (3-digit), and codes (4-digit) are increasingly disaggregated industry classifications.[35] The ABR provides data on a firm's SBI08 code, the first four digits of which correspond to the firm's NACE Rev. 2 code. The NFO contains data on the first two digits of a firm's SBI08 code and is used as a consistency check for industry classification. Before 2008, the CBS only provides the SBI93 industry classification, which is first converted to SBI08 codes. All 2-digit industries in the NACE section "Manufacturing" with sufficient observations to estimate the revenue function are included in the final sample. Table A1 in Appendix A list the 18 industries that are covered.

When estimating the revenue function, all monetary variables are deflated to make them comparable across time using the appropriate deflators at the 2-digit NACE industry level obtained from the OECD STAN database.[36] Firm location is accounted for in the input price control function by including indicator variables for the NUTS1-region that a firm's headquarters is located in. The NUTS classification is Eurostat's system for dividing the economic territories of the EU in order to enable socio-economic analyses of regions.[37]

A.2 Definition of variables

The following list gives the definition of key 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.

[34]I aggregate categorical variables by selecting the mode of the separate BE-level observations as the value for the OG. My results are not sensitive to this procedure.

[35]See https://ec.europa.eu/eurostat/documents/3859598/5902521/KS-RA-07-015-EN.PDF for a complete description of all NACE classifications and the conversion to other international industry classification codes.

[36]Deflators are available at https://www.oecd.org/en/data/datasets/structural-analysis-database.html.

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

- 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 worker $\frac{W_{it}L_{it}}{L_{it}}$.
- 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}.

A.3 Data cleaning

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

A.3.1 Necessary variables

1. I drop observations for which revenue, total assets, (in)tangible assets, employment, wages, intermediate inputs, or depreciation are missing.
2. I drop observations for which revenue, capital, employment, wages, or intermediate inputs are incorrectly signed or zero.
3. I drop observations with missing NACE Rev.2 codes, and for which the NACE Rev.2 division from the ABR is inconsistent with the first two digits of the NACE Rev.2 code from the NFO.

A.3.2 Internal consistency of balance sheets and income statements

I check the internal consistency of balance sheets and income statements by comparing the sum of variables in some aggregate to the variable holding the aggregate. I construct the following ratios

1. The sum of tangible and intangible assets, total shareholdings, long and short receivables, inventories, debtors, and liquid assets, as a ratio of total assets.
2. The sum of domestic and foreign shareholdings, as a ratio of total shareholdings.
3. Revenue minus the sum of wages, intermediate inputs, and depreciation, as a ratio of earnings before interest and taxes (EBIT).
4. EBIT net of total shareholdings, interest income and charges, extraordinary income and charges, and other financial results, as a ratio of pre-tax income.

5. Pre-tax income net of corporate taxes and third party equity as a ratio of after-tax income.

Due to minor rounding errors, these ratios are not always equal to one, even if the individual components are otherwise correct. Therefore, I drop all observations for which the above ratios are smaller than 0.95 or larger than 1.05. This reduces the sample by less than 30 observations, confirming that the CBS' internal consistency checks do an excellent job of eliminating inconsistent reports.

A.3.3 Further quality checks

1. I drop firms which at some point report negative capital, employment, or tangible assets.

2. I drop observations with incorrectly signed, or zero, total liabilities.

3. I drop 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.

4. I drop observations with negative value-added, where value-added is constructed as revenue net of intermediate input expenditure.

5. I drop the top and bottom percent of the capital to wage ratio, the total assets to total funds ratio, and the wages to value-added ratio.

A.3.4 Trimming

Within each 2-digit-industry-by-year cell, I trim the top and bottom 0.1 percent of the distributions of revenue, capital, total labor expenditure, and total intermediate input expenditure. After these steps, all 2-digit industries with sufficient data to estimate the revenue function are kept. Finally, observations with a negative labor wedge or intermediate input wedge are dropped upon estimating the revenue function, and the top and bottom percent of both input wedge distributions are dropped. All results reported in this paper are qualitatively robust to alternative trimming procedures. Table A2 reports the distribution of the 132,722 observations in the final sample by year and 2-digit industry.

Table A1: 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: This table lists all 2-digit industries covered by the final sample. Manufacturing covers industries 10 to 33.

Table A2: 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.

B Revenue function estimation

This appendix discusses the invertibility of the control function and the estimation of the revenue elasticities, and lists 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 the demand for intermediate inputs. This requires revenue productivity to be the sole unobservable in intermediates demand and intermediates demand to monotonically increase in revenue productivity. Intuitively, the monotonicity condition rules out adjustment frictions that are so severe that demand for intermediates does not respond at all to changes in revenue productivity. Such severe frictions would imply that intermediates are 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 straightforwardly to the revenue function from equation (6). Let inverse intermediates demand be given by $P_{it}^M(M_{it})$ with intermediate supply elasticity $0 < \varepsilon_{it}^M < \infty$. The first-order condition of intermediate inputs in equation (13) can be rewritten to provide as expression for intermediates demand

$$M_{it}^* = \left(\frac{1}{P_{it}^M} \frac{\varepsilon_{it}^M}{\varepsilon_{it}^M + 1} \theta^M K_{it}^{\theta^K} L_{it}^{\theta^L} \Omega_{it} \right)^{\frac{1}{1-\theta^M}} . \tag{19}$$

Two assumptions need to be met for invertibility to hold. First, M_{it}^* needs to be one-to-one in Ω_{it}, which is evidently the case from equation (19). Recall from the revenue function in equation (6) that $\Omega_{it} = \tilde{\Omega}_{it}^{1/\eta} \tilde{\Xi}_{it}$. As both shocks enter the revenue function multiplicatively all that is required is that intermediates demand is one-to-one in their product, even though clearly here intermediates demand is one-to-one in the two individual shocks as well.

Second, the scalar unobservable assumption requires that the only variable that varies within the level of estimation is Ω_{it}. The two unobserved variables in equation (19) that might lead to violations of the scalar unobservable assumptions are the intermediates 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 between observations. This could occur, for instance, if firms are monopolists or in monopolistic competition and face inverse intermediates demand curve $P_{it}^M = M_{it}^{1/\varepsilon^M}$. To further allow for price and elasticity differences across time, at a more narrowly defined industry level, or by location, I include indicators for the year, a firm's 3-digit industry, and it's 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 intermediates 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 (19) 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).[38] 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 by equation (15) the theoretical framework of rent sharing implies that wages are related to intermediate input prices and supply elasticities through rent sharing. Finally, to allow for intermediates price differences stemming from input quality differences across firms, I include a firm's revenue share 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 utilizes a translog production function with a Hicks-neutral productivity term, which results in (log) revenue given by equation (6). Compared to the intermediates demand in equation (19), no new unobservables are introduced when switching from Cobb-Douglas to translog, so that the intermediates price and supply elasticity are still the main challenge to the scalar unobservable assumption. Compared to (19), revenue elasticities are firm-time-specific if the production function is translog. However, revenue elasticity variation across firms and time within the same industry is entirely driven by (observable) differences in input use. Moreover, by equation (6) the added demand-side variation from demand (5) is absorbed by fixed effects. Finally, as in the Cobb-Douglas case, revenue productivity enters revenue function (6) log additively, so M_{it}^* is still strictly increasing in Ω_{it}.

B.2 Estimation of revenue elasticities

Recall that labor is measured as a headcount, while intermediate inputs are measured by deflating intermediate input expenditure using a year-by-2-digit-industry intermediates price deflator, $\bar{P}_{st}^M$. Any within-industry intermediate input price deviations from $\bar{P}_{st}^M$ will enter the structural error term of the revenue function, scaled by the revenue elasticities.[39] The motivation for these controls is given in the previous subsection. 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. Collecting all input price bias terms in

[38] 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.

[39] For instance, for a translog revenue function the input price bias resulting from within-industry-year variation in the price of intermediates is $\log\left(\frac{P_{it}^M}{\bar{P}_{st}^M}\right)\left(\beta_m + \beta_{km} + \beta_{lm} + \beta_{mm}\left(2m_{it} + \log\left(\frac{P_{it}^M}{\bar{P}_{st}^M}\right)\right)\right)$.

$b_{it}(\cdot)$, equation (6) can be written as

$$r_{it} = f_s(k_{it}, l_{it}, m_{it}; \beta) + b_{it}(\cdot) + \gamma_{st} + \omega_{it} + \epsilon_{it}, \tag{B1}$$

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 – either a 2-digit industry or a year-by-2-digit industry combination. I control for the input price bias $b_{it}(\cdot)$ using equation (8), repeated here for clarity

$$b_{it} = b_s((1, k_{it}, l_{it}, m_{it}) \times \mathbf{x}^b_{it}; \delta^b), \tag{B2}$$

where $\mathbf{x}^b_{it}$ contains firm i's intermediate inputs as a share of total intermediate inputs in the firm's 4-digit industry at time t, mms_{it}, an indicator variable for the NUTS1 region in which firm i's headquarter is located, a linear and a quadratic time trend, and an indicator variables for the firm i's 3-digit industry.[40] The notation of b_{it} indicates that the inputs do not enter linearly, but only interacted with $\mathbf{x}^b_{it}$, which is a consequence of the translog specification. Note that all input price variation related to revenue productivity ω_{it} is already controlled for by the inclusion of the control function in equation (B3), and that including the control function (B2) can soak up part of the input price dispersion even if it does not perfectly control for its sources.

Inverting the demand for intermediate inputs given in equation (7) provides a control for revenue productivity

$$\omega_{it} = m_s^{-1}(k_{it}, l_{it}, m_{it}, \mathbf{x}^m_{it}; \delta^m), \tag{B3}$$

where $\mathbf{x}^m_{it}$ contains mms_{it}, logged wage at firm i in year t, a linear and a quadratic time trend, and indicator variables for the firm i's 4-digit industry. Invertibility is discussed at length in the previous subsection.

Next, substitute equations (B2) and (B3) into the revenue function given in equation (B1) to obtain

$$r_{it} = f_s(k_{it}, l_{it}, m_{it}; \beta) + b_s((1, k_{it}, l_{it}, m_{it}) \times \mathbf{x}^b_{it}; \delta^b) + m_d^{-1}(k_{it}, l_{it}, m_{it}, \mathbf{x}^m_{it}; \delta^m) + \gamma_{st} + \epsilon_{it}. \tag{B4}$$

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

The first step of the two-step approach is based on an approximation to equation (B4). It consists of regressing log revenue on a third-degree polynomial in all right-hand side variables

[40] 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.

except for indicator variables and time trends, which are added 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: $r_{it} = \hat{r}_{it} + \hat{\epsilon_{it}}$. As discussed in the main text, correcting for the unanticipated revenue shock is necessary to identify the input wedges as firms base their decisions on R_{it} while $R_{it}\exp(\epsilon_{it})$ is observed.

In the second step, using planned revenue $\hat{r}_{it}$, we can obtain an estimate of ω_{it} that is a function of unknown coefficients β and δ^b from equation (B4)

$$\hat{\omega}_{it} = \hat{r}_{it} - f_s(k_{it}, l_{it}, m_{it}; \beta) - b_s((1, k_{it}, l_{it}, m_{it}) \times \mathbf{x}^b_{it}; \delta^b). \tag{B5}$$

From the law of motion of revenue productivity in equation (9), we can then obtain an estimate of the revenue productivity shock that depends on unknown coefficients β, δ^b, and δ^g

$$\hat{\psi}_{it} = \hat{\omega}_{it} - g_s(\hat{\omega}_{it}; \delta^{\mathbf{g}}). \tag{B6}$$

Estimates of θ, δ^b, and δ^g are obtained by forming moments on ψ_{it}. I approximate $g_s(\cdot)$ using a third-degree polynomial in all its arguments except for indicator variables and time trends, which are added 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 for the possibility that both l_{it} and m_{it} can be adjusted in response to the realization of ψ_{it}, so that $\mathbb{E}(\psi_{it}, l_{it}) \neq 0$ $\mathbb{E}(\psi_{it}, m_{it}) \neq 0$.

Together, the following moments result

$$\mathbb{E}(\psi_{it}\mathbf{Z}_{it}) = 0, \tag{B7}$$

where I focus on the exactly identified case and $\mathbf{Z}_{it}$ contains a constant, and all contemporaneous values of variables with the exception of variables based on labor, intermediates, and $\hat{r}_{it}$ that are instead constructed using their first lag. That is, $\mathbf{Z}_{it}$ contains 1, k_{it}, k^2_{it}, l_{it-1}, l^2_{it-1}, m_{it-1}, m^2_{it-1}, $k_{it}l_{it-1}$, $k_{it}m_{it-1}$, $l_{it-1}m_{it-1}$, mms_{it-1}, mms^2_{it-1}, $mms_{it-1}k_{it}$, $mms_{it-1}l_{it-1}$, $mms_{it-1}m_{it-1}$, t, t^2, $\hat{r}_{it-1}$, indicators for all but one NUTS1-region, and indicators for all but one 3-digit industry in 2-digit industry s.

For the year-by-2-digit level estimates, estimation proceeds as described above, but the revenue function and revenue productivity control are now indexed by st so that each 2-digit industry has a different revenue function every 2 years.

Finally, in the spirit of Blundell and Bond (2000), I estimate the revenue function under the assumption that the law of motion of revenue productivity is an AR(1) process: $\omega_{it} = \rho\omega_{it-1} + \psi_{it}$. One can substitute for ω using the revenue function in equation (B1), to arrive at the following moment conditions

$$\mathbb{E}((\psi_{it} + \epsilon_{it} - \rho\epsilon_{it-1})\,\mathbf{Z}^{AR1}_{it}) = 0, \tag{B8}$$

where the difference between $\mathbf{Z}^{AR1}_{it}$ and $\mathbf{Z}_{it}$ is that variables have been lagged one (additional) time due to the moments being formed on a term that includes ϵ_{it-1}.

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 Additional figures

Figure C1: Regressions of the firm-level labor wedge on revenue per employee (panel a) and value-added per employee (panel b)

Panel a

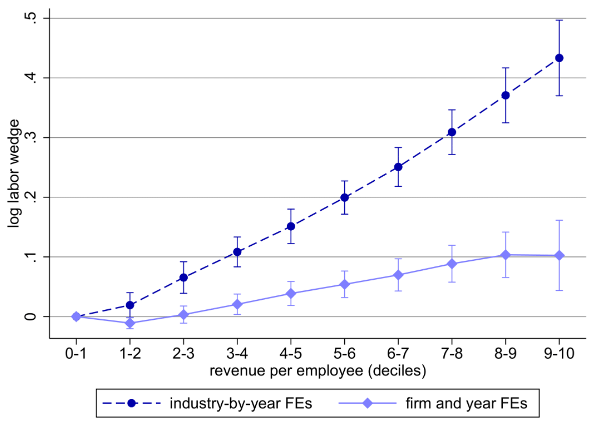

Panel b

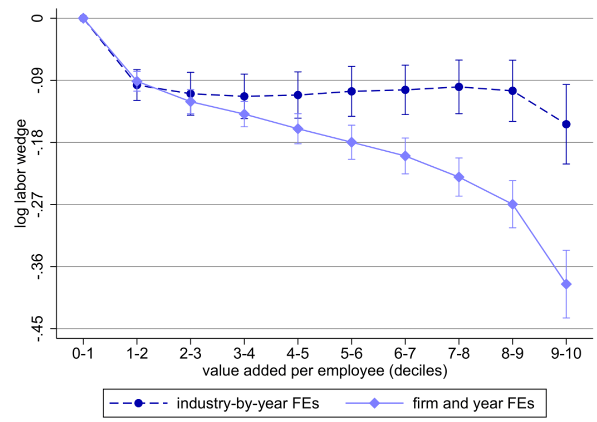

Notes: Point estimates and 95 percent confidence intervals from a non-parametric regression of the log labor wedge on revenue per employee (panel a) and value-added per employee (panel b), controlling for employment; Explanatory variables are divided into 10 categories based on the deciles of their distribution, and the coefficient on the indicator of the first decile is normalized to 0; Either 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 C2: Regressions of the firm-level intermediate input wedge on revenue per employee (panel a) and value-added per employee (panel b)

Panel a

Panel b

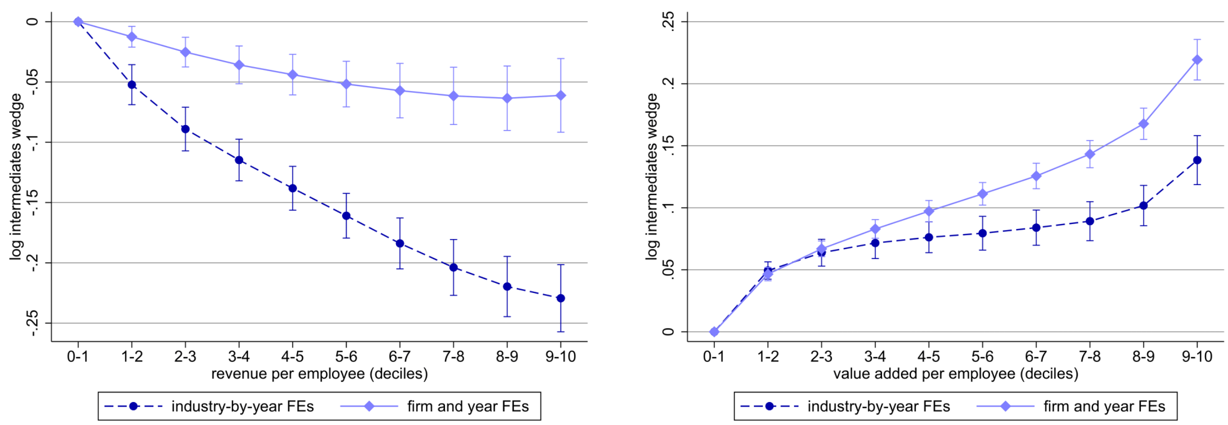

Notes: Point estimates and 95 percent confidence intervals from a non-parametric regression of the log intermediate input wedge on revenue per employee (panel a) and value-added per employee (panel b), controlling for employment; Explanatory variables are divided into 10 categories based on the deciles of their distribution, and the coefficient on the indicator of the first decile is normalized to 0; Either 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 C3: Yearly exchange rate EUR/RMB, over time

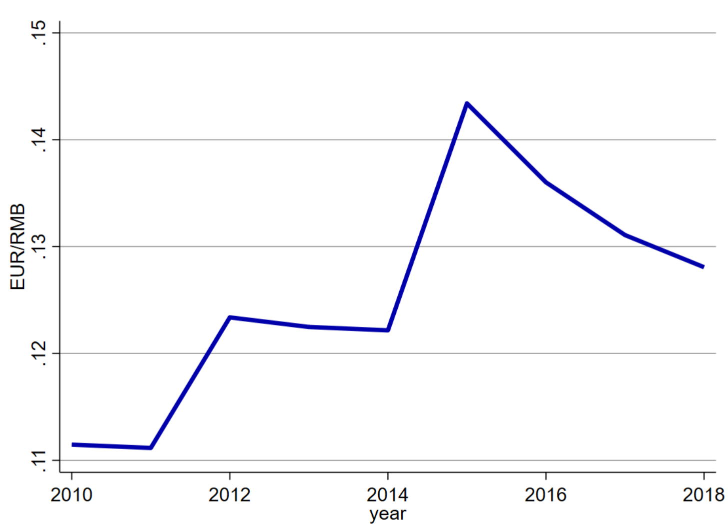

Notes: Yearly exchange rate – defined as Euro per Chinese Yuan Renminbi – over time from 2010 to 2018, the years used in regressions relying on firms' exposure to exchange rate fluctuations.

D Additional tables

Table D1: 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 A1.

Table D2: 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)	Observations
Baseline estimates	0.64	0.81	1.00	1.27	1.98	0.93	1.04	1.14	1.29	1.49	132,722
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
AR(1) revenue process	0.51	0.72	0.85	1.03	1.42	0.91	1.05	1.16	1.33	1.77	128,378
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

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 percentile of the input wedge distributions, respectively. Input wedges are rounded to two decimal points; Baseline estimates = Baseline estimates from the main text; Calibrated revenue elasticities = Input wedges based on calibrated revenue elasticities; AR(1) revenue process = Input wedges based on an AR(1) law of motion of unobserved revenue productivity; Time-specific revenue function = Input wedges based on 2-year-by-2-digit-industry-specific revenue functions.

Table D3: 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.($\hat{\beta}$)	Observations
Baseline estimates	-1.321	0.031***	-1.199	0.034***	132,722
Calibrated revenue elasticities	-0.900	0.036***	-0.449	0.030***	129,844
AR(1) revenue process	-0.279	0.022***	-0.079	0.025***	128,378
Time-specific revenue function	-0.944	0.062***	-2.230	0.198***	101,352

Notes: Table D3 reports results from regressions of the labor wedge on the intermediate input wedge; Baseline estimates = Baseline estimates from the main text; Calibrated revenue elasticities = Input wedges based on calibrated revenue elasticities; AR(1) revenue process = Input wedges based on an AR(1) law of motion of unobserved revenue productivity; Time-specific revenue function = Input wedges based on 2-year-by-2-digit-industry-specific revenue functions; All variables in logs; Standard errors are clustered at the 3-digit-industry level. *** indicates statistical significance at the 1% level.

Table D4: Summary statistics for IV sample

	p(5)	p(25)	p(50)	p(75)	p(95)	observations
Labor wedge	0.63	0.85	1.05	1.34	1.99	8,469
Intermediate input wedge	0.93	1.03	1.12	1.21	1.42	8,469
Chinese import share	0.000	0.003	0.015	0.078	0.362	8,469

Notes: Summary statistics for the sample used in the IV estimates of Section 5; p(5), p(25), p(50), p(75) and p(95) refer to the 5th, 25th, 50th, 75th and 95th percentile 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 D5: IV estimates of the relationship between the labor wedge and the intermediate input wedge, by robustness check

	Robustness check				
	Excluding exporters	Instrument using lagged shares	Calibrated revenue elasticities	AR(1) revenue process	Time-specific revenue function
	Panel a: Labor wedge				
Intermediate input wedge	-0.744	-1.256	-1.227	-0.542	-0.958
	(0.442)*	(0.339)***	(0.488)**	(0.255)**	(0.499)*
	Panel b: Wage				
Intermediate input wedge	1.041	1.390	0.627	0.885	0.565
	(0.267)***	(0.616)***	(0.214)***	(0.300)***	(0.191)***
	Panel c: Marginal revenue product of labor				
Intermediate input wedge	0.297	0.134	-0.600	0.344	-0.393
	(0.367)	(0.621)	(0.365)	(0.178)*	(0.424)
	Panel d: Intermediate input wedge (first stage)				
$Z^{E/R}$	-0.093	-0.185	-0.147	-0.110	-0.137
	(0.015)***	(0.037)***	(0.016)***	(0.007)***	(0.012)***
F-stat of instrument	36.83	24.47	82.73	264.46	138.47
Firm and year fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	4,698	8,469	8,351	8,246	7,551

Notes: Table D5 reports 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); Excluding exporters = Omits all firms that export to China in the year or year prior to the observation; Instrument using lagged shares = Exchange rate instrument based on the share of Chinese imports in all intermediate input expenditure in the year prior to the observation; Calibrated revenue elasticities = Input wedges based on calibrated revenue elasticities; AR(1) revenue process = Input wedges based on an AR(1) law of motion of unobserved revenue productivity; Time-specific revenue function = Input wedges based on 2-year-by-2-digit-industry-specific revenue functions; $Z^{E/R}$ is defined in equation (18); All regressions include lagged inputs and revenue as controls; All variables in logs except the instrument; Standard error clustered at the firm level in parentheses; ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.