Wage markups and buyer power in intermediate input markets*

Leonard Treuren†

February 2023

Abstract

A rapidly growing literature suggests that monopsony power is common in US labor markets. I examine whether this result generalizes to Europe, where collective bargaining agreements characterize labor markets. I use Dutch firm-level manufacturing data from 2007 to 2018, together with an efficient bargaining model and revenue function estimation. Wages are typically above the marginal revenue contribution of employees. This is not in line with monopsonistic labor markets but precisely what is expected when employees have bargaining power and can extract rents from their employers. In addition, I provide evidence of buyer power in intermediate input markets and show that firms that underpay their input suppliers on the margin set higher wage markups. This suggests that firms share rents generated in intermediate input markets with their employees. Firm-time-specific rent sharing elasticities indicate that firms increase wages on average by 0.22 percent following a 1 percent increase in quasi-rents per employee.

Keywords: Monopsony; Rent sharing; Buyer power; Revenue function estimation

JEL Codes: D43; J31; J42; J50; L10

*I thank Dan Ackerberg, Eric Bartelsman, Filippo Biondi, Jan De Loecker, Katharina Erhardt, Nadine Hahn, Paul Heidhues, Matthias Mertens, Michael Rubens, Tobias Salz, Jo Seldeslachts, Joel Stiebale, Jo Van Biesebroeck, Frank Verboven, and seminar participants at DICE, KU Leuven, and the 2021 EARIE (Bergen), 2021 CAED (Coimbra), 2022 IIOC (Boston), 2022 EEA (Milan), and 2022 JEI (Las Palmas) conferences for useful comments and discussions. Results based on calculations using non-public microdata from Statistics Netherlands (CBS). This research is funded by the Open Data Infrastructure for Social Science and Economic Innovations (ODISSEI) in the Netherlands (www.odissei-data.nl). Remaining errors are my own.

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

1 Introduction

A rapidly growing literature suggests that monopsonistic labor markets are the norm in the United States, allowing firms to mark down wages relative to the marginal revenue product of labor, which leads to misallocation and reduces welfare (e.g., Yeh et al. (2022); Berger et al. (2022); Lamadon et al. (2022)).[1] These findings have led to a surge of policy proposals.[2] Unlike in the US, European labor markets are characterized by collective bargaining agreements that could cause firms to share rents with their employees, eliminating monopsony concerns. Determining whether monopsony results translate to European labor markets is of first-order policy concern and is, therefore, the first aim of this paper.

Buyer power in intermediate input markets is also documented by an emerging literature (e.g., Morlacco (2020); Rubens (2021), Avignon and Guigue (2022)). When firms are not wage takers in the labor market, buyer power in other input markets can influence wages. However, studies of imperfectly competitive labor markets typically ignore imperfections in other input markets. The second aim of this paper is to fill this gap by studying intermediate input market imperfections and their relation to wages.

In this paper, I study imperfect competition in labor markets and buyer power in intermediate input markets in Dutch manufacturing from 2007 to 2018. Collective bargaining agreements are central to Dutch labor markets, while intermediate input markets are fairly concentrated and rely heavily on imports. I choose this setting because Dutch labor markets are representative of western European labor markets, and because buyer power has been documented in comparable intermediate input markets in France (Morlacco, 2020).

I develop a simple theoretical framework where firms bargain with collectively organized workers in the labor market and potentially possess buyer power in the market for intermediate inputs. In this setting, wages are marked up above the marginal revenue product of labor. In contrast, when firms have monopsony power, wage markdowns are expected. Likewise, if firms have buyer power in the market for intermediates, intermediate input prices are marked down relative to the marginal revenue product of intermediates. Finally, the extent of the wage markup depends positively on the rents a firm generates. As buyer power increases rents, firms with more buyer power should pay higher wages, all else equal.

I estimate the input wedges of labor and intermediate inputs to test these theoretical predictions. An input wedge is the ratio of an input's marginal revenue product to its price. I show that input wedges can be expressed as the revenue elasticity of an input divided by its expenditure share in revenue. Akin to Petrin and Sivadasan (2013), I estimate a revenue function using a control function approach to recover revenue elasticities of labor

[1] Manning (2021) and Card (2022) survey the recent monopsony literature.

[2] See the Department of Justice's 2016 Antitrust Guidance for Human Resource Professionals, the Federal Trade Commission's 2018 hearing on Multi-Sided Platforms, Labor Markets and Potential Competition (Hearing #3, October 15–17), and calls by academics to reform antitrust practice concerning labor markets (e.g., Naidu et al. (2018); Marinescu and Posner (2019)).

and intermediate inputs.[3] In contrast, the US-based monopsony results of Yeh et al. (2022) are based on labor wedges identified using the production approach.

The production approach is the prevalent approach to identifying labor wedges. It relies on production function estimation and requires the existence of a variable input that is frictionlessly adjustable and for which firms are price takers.[4] This literature tends to select intermediate inputs – often referred to as "materials" – as this variable input. I show that this approach underestimates the labor wedge if firms have buyer power in the market for intermediate inputs. As buyer power might be a concern in the Dutch setting I study, the production approach is not suitable. The main advantage of my identification strategy is that it allows for buyer power in all input markets.

I find that wage markups are prevalent in Dutch manufacturing, covering more than 75 percent of all observations. At the median, wages are marked up by 16 percent over the marginal revenue product of labor. This is not in line with monopsony power, which would result in wage markdowns as observed in US manufacturing by Yeh et al. (2022), but is in line with employees using their bargaining power to extract rents from firms. In contrast, intermediate input wedges suggest that buyer power for intermediates is common in Dutch manufacturing. The median intermediates price markdown is 15 percent. This finding is in line with Morlacco (2020), who finds evidence of buyer power for imported intermediate inputs in French manufacturing.

The labor wedge distribution and the intermediate input wedge distribution show substantial dispersion. This dispersion is not caused by variation over time, suggesting that adjustment frictions do not play an important role. Within-industry differences, rather than between-industry differences, are the primary driver of input wedge variation, implying that firm-specific factors shape variation in wage markups more than industry-specific ones. I show that within-industry wage markup variation is largely unrelated to the revenue-generating abilities of employees but can be explained by wage variation. In addition, firms that underpay their input suppliers on the margin set higher wage markups. The higher a firm's intermediate input price markdown, the higher its wage. The data, therefore, support the hypothesis that the distribution of wage markups is shaped not by worker ability but by firms sharing rents generated in the intermediate input market with their employees.

To quantify rent sharing, I estimate the firm-level responsiveness of wages to rents. In my efficient bargaining model, firms bargain about wages and employment, consistent with the Dutch setting where employee associations trade off wage increases with unemployment. In this setting, firms act as if they are wage takers facing the employees' outside option as the

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

[4] Applications are in Lu et al. (2019), Mertens (2021, 2022), Caselli et al. (2021), Brooks et al. (2021), Yeh et al. (2022), and in papers that estimate measures closely related to the labor wedge based on the marginal revenue product of labor and the wage (e.g., Mertens (2020); Dobbelaere et al. (2020); Dobbelaere and Wiersma (2020)).

going wage – they set the marginal revenue product of labor equal to the employees' outside option. Relative bargaining power between employers and employees then determines how much of the firm's rents are captured by employees in the form of wages exceeding their outside option. These insights allow me to use revenue function estimates to identify the employees' outside option. Rent sharing elasticities are obtained by comparing actual wages to implied outside options.

I show that the elasticity of wage with respect to a firm's quasi-rents can be written as a Lerner index wage markup: the ratio of wage minus the marginal revenue product of labor to the wage. The standard approach in the rent sharing literature is to regress wages on a measure of rents and several controls (see Card et al. (2018) for a survey). This method delivers an average of the underlying firm-specific rent sharing elasticities. In addition, as both the employees' and the firms' outside option are typically unobserved, instruments for a firm's rents that do not also shift the outside options are required. In contrast, my approach recovers firm-time-specific elasticities using an efficient bargaining model and revenue function estimation.

I find that firms pay their employees on average 0.22 percent more following a 1 percent increase in quasi-rents per employee. The distribution of elasticities is right skewed, with wages increasing by less than 0.5 percent following a 1 percent increase in quasi-rents per employee in 95 percent of all cases and increasing by less than 0.1 percent in roughly 25 percent of all cases. The mean elasticities I report are in line with, but slightly below, the rent sharing elasticities reported in Van Reenen (1996) and Kline et al. (2019). In addition, I show that the rent sharing elasticity I obtain using a standard regression approach is comparable to the mean of the firm-time specific elasticities.

The main implication of this paper is that concerns about widespread monopsony power are unlikely to be warranted in European labor markets characterized by collective bargaining agreements. The second implication is that researchers should be careful when identifying labor market imperfections by restricting imperfections in other input markets. Specifically, in my Dutch sample, using the prevalent production approach to identify the labor wedge would substantially bias estimates of the labor wedge and ignore the connection between the two input markets. The third implication of this paper is that there exists substantial variation in rent-sharing elasticities, which can not be entirely identified by regressing wages on a measure of rents. This paper shows how an efficient bargaining model and revenue function estimation allow consideration of these implications

Key results in this paper require estimates of revenue elasticities of labor and intermediate inputs. A firm's revenue function depends on primitives of demand and supply, each potentially containing unobserved determinants. Identification of revenue functions, therefore, potentially requires strong assumptions on which unobservables are relevant and how they evolve over time. I show that all key results are robust to a host of alternative approaches to obtaining revenue elasticities. In particular, all key results remain valid when ignoring revenue function estimation altogether and simply calibrating a single revenue elasticity for each input. The reason is that my results are driven by variation of revenue shares of input expenditure, which are observed, and not by variation of revenue elasticities.

This paper contributes to three different strands of literature. First, the literature on monopsony power.[5] 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 have wage markdowns. In Germany, a considerable number of firms are characterized by wage markdowns, although substantially less than in the US. These wage markdowns are concentrated in firms setting high average wages (Mertens, 2021) and are markedly less likely to occur where collective bargaining agreements cover employees (Dobbelaere et al., 2020). In line with the current paper, these findings suggest that wage markdowns are unlikely to occur when collective bargaining agreements are central to labor markets. Finally, wage markdowns are found to substantially decrease labor shares in China and India (Brooks et al., 2021), and Germany (Mertens, 2022).[6]

The second strand of literature related to this paper studies rent sharing. An extensive literature regresses wages on a measure of rents and controls to obtain an average rent sharing elasticity (see Card et al. (2018) for a recent survey). I contribute by instead relying on revenue function estimation and an efficient bargaining model to identify firm-time-specific rent sharing elasticities. Several papers provide evidence of rent sharing based on production function estimation. Dobbelaere and Mairesse (2013) classify the majority of French manufacturing firms as operating in labor markets characterized by efficient bargaining.[7] Caselli et al. (2021) report sizeable wage markups in French manufacturing and show that these wage markups decline in response to Chinese import competition. Finally, 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.

The third strand of related work is an emerging literature on buyer power in intermediate input markets. Morlacco (2020) finds evidence of buyer power for imported intermediate inputs in French manufacturing. Rubens (2021) finds that ownership consolidation in Chinese cigarette manufacturing has increased intermediate input price markdowns by 30 percent. A fundamental difficulty in this literature is that prices of intermediate inputs are typically unobserved. This prevents the current paper from making detailed inferences on the origins of buyer power. Input prices are observed by Avignon and Guigue (2022), who show that

[5] I focus on papers estimating labor wedges using the production approach. Other common approaches include estimating labor supply elasticities and relating labor market concentration to wage measures. Manning (2021) and Card (2022) survey the monopsony literature.

[6] Several papers study the relation between decreasing protectionism and monopsony power (e.g., Lu et al. (2019); Mertens (2020); Dobbelaere and Wiersma (2020)). Results depend heavily on the country, and the particular deregulation studied.

[7] See also Crépon et al. (2005) and Dobbelaere and Mairesse (2018). These papers jointly estimate markups in output markets and imperfections in labor markets.

French dairy manufacturers have buyer power in the market for raw milk. Atalay (2014) does not directly study buyer power, but does find dispersion in materials' prices within narrowly defined industries producing relatively homogeneous goods in the US, and shows that within-supplier markup differences can explain part of this dispersion.

This paper contributes to the aforementioned articles by considering both labor and intermediate input market imperfections. The production approach identifies labor wedges by ruling out buyer power in another input market. I show that using this approach to study Dutch manufacturing would overestimate the extent of wage markups and rule out buyer power for intermediates. My primary contribution is permitting buyer power for both labor and intermediates. This allows me to study imperfections in both input markets as well as their relation.

The remainder of this paper proceeds as follows. In Section 2, the theoretical framework and the Dutch setting on which it is based are introduced. Section 3 outlines the empirical approach to identifying input wedges and the data. Section 4 provides the results and robustness checks, followed by concluding remarks in Section 5.

2 Theoretical framework

This section presents the theoretical framework and the Dutch institutional setting. I introduce a simple model that links a firm's optimality conditions for labor and intermediate inputs to input wedges. An input wedge is the ratio of an input's marginal revenue product to its price. The model provides an interpretation of these wedges and several empirical hypotheses which are explored in this paper.

2.1 Setting

In the Netherlands, employees' compensation and labor market legislation are primarily determined by collective bargaining and dialogue between federations of employee associations, firms, and, occasionally, the government. Collective agreements are central to Dutch labor markets. Collective bargaining coverage in the Netherlands in 2016 was 78.6 percent, compared to 11.5 percent in the US (OECD, 2019a).[8] High collective bargaining coverage characterizes all western European labor markets and is the main difference between labor markets in the United States and Europe.[9]

Bargaining is an institutional feature of Dutch labor markets that is supported by the

[8]The collective bargaining coverage is the share of all employees whose terms of employment are governed by at least one collective agreement. The 2016 OECD average is 32.4 percent.

[9]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 (OECD, 2019a).

central government. Each year, firms and associations of employees meet to set guidelines for the subsequent wage adjustments and other outcomes such as mandatory social security contributions (Visser, 2016). Collective agreements apply automatically to all workers employed by firms involved in a particular agreement, regardless of whether those employees are members of an employee association (Hijzen et al., 2019). Unionization rates and union membership status are, therefore, no longer directly tied to collective bargaining coverage. This has allowed coverage to remain stable and high while unionization rates decreased rapidly the past 30 years.[10]

The Netherlands has a "two-level" bargaining system. Bargaining at the sector level sets wage floors, while bargaining at the firm level determines wages and employment. Collective agreements at the sector level specify various pay scales determined by factors such as seniority and job description. Firms can and do form their own agreements with employee associations, creating firm-level variation in collective agreements within industries. In addition, firms usually have some discretion when allocating employees to pay scales and awarding premiums, ensuring that collective agreements act primarily as wage floors. Bargaining in most of western Europe is organized similarly, and a robust result from the literature is that actual wages tend to exceed wage floors substantially (e.g., Cardoso and Portugal (2005); Card et al. (2014); Bhuller et al. (2022)).[11]

Low unemployment and high minimum wages characterize the Dutch labor market. In 2016, the real annual minimum wage in the Netherlands was 55.24 percent higher than the real annual minimum wage in the US. Compared to Germany, the concomitant percentage differences was only 6.77 (OECD, 2019b).[12] Unemployment as a percentage of the labor force increased following the financial crisis of 2007-2008, but remained relatively low at on average 5.3 percent throughout the sample period (OECD, 2019c).

Collective agreements are absent in the Dutch intermediate input market, so interactions between firms and their input suppliers are bilateral. Intermediate input markets are fairly concentrated and Dutch firms frequently rely on imports. Dutch imports totaled 75.16 percent of GDP in 2015, compared to 15.4 percent in the US.[13] In 2018, more than 80

[10]Decreasing unionization rates have forced unions to rely less on members' contributions to finance their operations. For instance, from 2019 to 2020, the largest Dutch employee association maintained a constant budget while seeing contributions from private firms and the government increase at the expense of members' contributions. See `https://www.fnv-magazine. nl/fnv-jaarverslag-2020`.

[11]The 1982 Wassenaar agreement between the leading employee association and employer representatives set the stage for decentralization of collective bargaining over wages and other labor market outcomes. This decentralization allowed firms to deviate from national or industry-level collective agreements by increasing the scope for firm-level bargaining. Under "semi-binding law", deviations from higher-level agreements can be specified in written agreements between firms and their employees (Hijzen et al., 2019).

[12]Real annual minimum wages are computed by converting statutory minimum wages into a common hourly and annual pay period and then converting the sums into a common currency unit (USD) using Purchasing Power Parities (PPPs) for private consumption expenditure.

[13]Obtained from `https://data.worldbank.org/indicator/NE.IMP.GNFS.ZS?most_`

percent of Dutch imports accrued to either manufacturing or wholesale and retail trade, with manufacturing imports exceeding 100 billion (CBS, 2019b, p.69). According to the OECD, low business dynamism characterizes Dutch manufacturing and a relatively large share of business is concentrated in larger firms.[14]

Two main differences emerge between input markets in the US and the Netherlands. First, while in the US, relations between firms and their input suppliers are primarily bilateral. In the Netherlands, firms interact mainly with employee associations in the labor market. Second, sellers and buyers in input markets are mainly domestic in the US, while intermediate input suppliers are often foreign firms in the Netherlands. In French manufacturing, Morlacco (2020) finds evidence of buyer power for imported intermediate inputs. Therefore, a suitable theoretical framework should incorporate the collective nature of Dutch labor markets while simultaneously allowing – but not imposing – firms to have buyer power in intermediate input markets.

2.2 Model

To model the key characteristics of the Dutch labor market, I rely on the efficient bargaining setting of McDonald and Solow (1981).[15] In this setting, employee associations have bargaining power and capture part of the surplus generated by firms. I allow firms to have finite input supply elasticities in the market for intermediate inputs. This allows for – but does not impose – buyer power in the intermediate input market.

Firm i's production function at time t is given by

$$Q_{it} = F_{it}^{Q}(K_{it}, L_{it}, M_{it})\Omega_{it}^{Q}, \tag{1}$$

where K_{it}, L_{it}, and M_{it} are, respectively, a firm's capital, labor, and intermediate inputs, and Ω_{it}^{Q} is Hicks-neutral total factor productivity. The production function is assumed to be twice differentiable with respect to its arguments, with $\frac{\partial Q_{it}}{\partial x} > 0$ and $\frac{\partial^2 Q_{it}}{\partial x^2} < 0$ for $x \in \{K_{it}, L_{it}, M_{it}\}$. Inverse demand is given by

$$P_{it}(Q_{it}) = F_{it}^{P}(Q_{it})\Omega_{it}^{P}, \tag{2}$$

where Ω_{it}^{P} is a demand shock. $P_{it}(Q_{it})$ is assumed to be twice differentiable with respect to quantity, with $\frac{\partial P_{it}}{\partial Q_{it}} < 0$. The profit of firm i at time t is given by

$$\Pi_{it} = R_{it}(Q_{it}) - P_{it}^{K}K_{it} - W_{it}L_{it} - P_{it}^{M}(M_{it})M_{it}, \tag{3}$$

`recent_year_desc=false&locations=OE-NL.`

[14] "Larger firms" are defined as firms with 50+ employees, and Dutch manufacturing is compared to manufacturing in a group of EU countries. See OECD (2021) for details.

[15] Efficient bargaining models are standard in the rent sharing literature and have been used before to model labor markets in western European manufacturing (e.g., Dobbelaere and Mairesse (2013)).

where $R_{it}(Q_{it}) = P_{it}(Q_{it})Q_{it}$ is revenue, P^K_{it} is the user cost of capital, W_{it} the wage, and $P^M_{it}(M_{it})$ the price of intermediate inputs. I assume that $\frac{\partial R_{it}}{\partial x} > 0$ and $\frac{\partial^2 R_{it}}{\partial x^2} < 0$ for $x \in \{K_{it}, L_{it}, M_{it}\}$.[16]

In the labor market, a firm's wage and employment result from bargaining with an association of employees. The association maximizes

$$U(L_{it}, W_{it}) = L_{it}(W_{it} - \bar{W}_{it}), \tag{4}$$

where $\bar{W}_{it} > 0$ is the outside option of the association's workers.[17] Equation (4) is standard in the literature on rent sharing and captures the notion that the employee association puts weight on both employment and wages (Card et al., 2018). If bargaining fails, employees are assumed to obtain their outside option. If the employee association has any bargaining power at all, therefore, $W_{it} > \bar{W}_{it}$. Profit if bargaining fails is denoted by $\bar{\Pi}_{it} > 0$. Therefore, firm profit should exceed $\bar{\Pi}_{it}$ if firm i has bargaining power. Following the literature, $\bar{W}_{it}$ and $\bar{\Pi}_{it}$ are assumed to be independent of W_{it} and L_{it}.

By modeling bargaining as a single interaction between firms and employee associations, I abstract from the two-level bargaining that occurs in practice. Unfortunately, data limitations ensure that I can not distinguish the wage floors resulting from sectoral bargaining from wages and employment resulting from bargaining at the firm level. I, therefore, follow Card et al.'s (2014) approach to modeling two-level bargaining in Italy and represent the negotiation process as a single efficient bargaining game between a firm and an employee association. I interpret results as measuring the sum of wage premiums received by workers, including those resulting from firm-level deviations from sectoral wage floors.[18]

I assume that labor, wages, and intermediate inputs can be adjusted each period. Not all sectoral collective agreements are adjusted yearly so this assumption can be questioned for labor and wages. In western Europe, the majority of firms adjust wages each year, and wages tend to substantially exceed sectoral wage floors (e.g., Bhuller et al. (2022)).[19] These findings suggest that the firm-level tier of bargaining is sufficiently flexible to overcome sector-level rigidities. In line with this, Caloia et al. (2021) show that the staggered setting of collective

[16] Positive first derivatives of the revenue function require that marginal revenue is positive, as for all inputs x we have $\frac{\partial R_{it}}{\partial x} = \frac{\partial R_{it}}{\partial Q_{it}}\frac{\partial Q_{it}}{\partial x}$ and $\frac{\partial Q_{it}}{\partial x} > 0$ by assumption. $\frac{\partial^2 R_{it}}{\partial x^2} < 0$ requires $\frac{\partial^2 P_{it}}{\partial Q_{it}^2}$ to either be negative, or positive but not too large, as $\frac{\partial^2 R_{it}}{\partial x^2} = \frac{\partial R_{it}}{\partial Q_{it}}\frac{\partial^2 Q_{it}}{\partial x^2} + \frac{\partial^2 R_{it}}{\partial Q_{it}\partial x}\frac{\partial Q_{it}}{\partial x}$ and $\frac{\partial^2 R_{it}}{\partial Q_{it}\partial x} = \frac{\partial Q_{it}}{\partial x}(2\frac{\partial P_{it}}{\partial Q_{it}} + \frac{\partial^2 P_{it}}{\partial Q_{it}^2}Q_{it})$. A simple example where all assumed properties hold is an iso-elastic demand curve together with a Cobb-Douglas production function.

[17] This is typically thought of as a firm-level average of employees' "unemployment compensation benefits, but should really include all the other contributions to the standard of living that would not be received if workers were employed by the bargaining firm"(McDonald and Solow, 1981, p.899).

[18] Alternatively, one could think of the employees' outside option being shifted by the sectoral wage floor and my theoretical framework as firm-level bargaining where sectoral outcomes are taken as given.

[19] Based on a survey covering 15 countries, Fabiani et al. (2010) show that about 75 percent of all firms adjust wages at least once a year.

agreements does not materially affect firm-level employment in the Netherlands. Therefore, treating labor and wages as variable appears reasonable. I focus on the firm's two variable inputs and assume that capital is predetermined.[20]

The generalized Nash-bargaining solution to the bargaining process between the firm and the employee association solves

$$\max_{L_{it},M_{it},W_{it}} (L_{it}(W_{it} - \bar{W}_{it}))^{\phi_{it}}(R_{it}(Q_{it}) - P^K_{it}K_{it} - W_{it}L_{it} - P^M_{it}(M_{it})M_{it} - \bar{\Pi}_{it})^{1-\phi_{it}}, \tag{5}$$

where ϕ_{it} denotes the bargaining power of the employee association ($0 < \phi_{it} < 1$), which is taken as given by the firm.[21] I allow – but do not impose – bargaining power to be firm-specific.

The first-order condition with respect to W_{it} is

$$W_{it} = \bar{W}_{it} + \frac{\phi_{it}}{1-\phi_{it}} \left(\frac{R_{it}(Q_{it}) - P^K_{it}K_{it} - W_{it}L_{it} - P^M_{it}(M_{it})M_{it} - \bar{\Pi}_{it}}{L_{it}} \right), \tag{6}$$

and the first-order condition with respect to L_{it} is

$$W_{it} = MRPL_{it} + \phi_{it} \left(\frac{R_{it}(Q_{it}) - P^K_{it}K_{it} - MRPL_{it}L_{it} - P^M_{it}(M_{it})M_{it} - \bar{\Pi}_{it}}{L_{it}} \right), \tag{7}$$

where $MRPL_{it} = \frac{\partial R_{it}}{\partial Q_{it}} \frac{\partial Q_{it}}{\partial L_{it}}$, the marginal revenue product of labor. Combining first-order conditions (6) and (7) shows that $MRPL_{it} = \bar{W}_{it}$, which implies that the firm's optimal input selection results in the same revenue product as the input choices of a firm operating in a perfectly competitive labor market facing wage $\bar{W}_{it}$. This insight will allow me to identify bargaining power and rent sharing elasticities later on.

The optimality conditions given in equations (6) and (7) can be related to the labor wedge. Denote the labor wedge by

$$\gamma^L_{it} = \frac{MRPL_{it}}{W_{it}}. \tag{8}$$

Define a firm's quasi-rents as $QR_{it} = R_{it}(Q_{it}) - P^K_{it}K_{it} - \bar{W}_{it}L_{it} - P^M_{it}(M_{it})M_{it} - \bar{\Pi}_{it}$, which are positive whenever a firm engages in bargaining. A firm's quasi-rents represent the surplus to be divided between the firm (profit in excess of the firm's outside option) and its employees (wages in excess of the employees' outside option). From first-order conditions (6) and (7) it is then clear that

$$W_{it} - MRPL_{it} = \phi_{it} \frac{QR_{it}}{L_{it}}, \tag{9}$$

[20] Within a period, $P^K_{it}K_{it}$ represents the yearly payment flow of predetermined capital – shaped, for instance, by depreciation. Card et al. (2014) study a two-period efficient bargaining model with dynamic capital formation.

[21] I follow the rent sharing literature and leave endogenizing bargaining power for future work as this would greatly complicate the model and detract from the main contributions of this paper.

which shows that $W_{it} > MRPL_{it}$, or, equivalently, that $\gamma_{it}^L < 1$. Wages are marked up relative to the marginal revenue product of labor – wage markups occur. Equation (9) states that the extent of the wage markup depends on the bargaining power of the employee association and the quasi-rents per employee.

In the market for intermediate inputs, I allow for buyer power by assuming that $P_{it}^M(M_{it})$ is continuous and monotonically increasing in M_{it}, and that the supply elasticity of intermediate inputs, $(\varepsilon_{it}^M)^{-1} = \frac{\partial P_{it}^M}{\partial M_{it}} \frac{M_{it}}{P_{it}^M}$, is positive but finite ($0 < (\varepsilon_{it}^M)^{-1} < \infty$). This approach nests the competitive case in the limit as $(\varepsilon_{it}^M)^{-1}$ approaches 0. The first-order condition of (5) with respect to M_{it} is

$$MRPM_{it} = P_{it}^M(M_{it})(\frac{\varepsilon_{it}^M + 1}{\varepsilon_{it}^M}), \tag{10}$$

where $MRPM_{it} = \frac{\partial R_{it}}{\partial Q_{it}} \frac{\partial Q_{it}}{\partial M_{it}}$, the marginal revenue product of intermediate inputs. Equation (10) states that when the intermediates supply elasticity is finite, a wedge is driven between the marginal revenue product of intermediates and their price. Denote this intermediate input wedge by

$$\gamma_{it}^M = \frac{MRPM_{it}}{P_{it}^M}. \tag{11}$$

Equation (10) and $0 < (\varepsilon_{it}^M)^{-1} < \infty$ imply that $\gamma_{it}^M > 1$. When the intermediate input wedge is above unity, the price of intermediates is marked down relative to their marginal revenue product – intermediate input price markdowns occur. The extent of the markdown depends on the elasticity of intermediate input supply. The less elastic is supply, the larger the intermediates markdown.

My approach to modeling intermediate input markets does not impose buyer power and does not require taking a stand on the origins of buyer power. This prevents me from having to make strong assumptions on unobservables and generates testable predictions that hold in a wide variety of models.[22] However, the lack of data on supplier networks and input prices implies that I can not identify the sources of buyer power. By focusing on the measurement and description of buyer power and the relation between labor markets and intermediate input markets, this paper sets a first step that future work on buyer power for intermediate inputs can build on.

My theoretical framework makes three main predictions regarding input wedges. First, we should expect wage markups – labor wedges below unity – given that collective bargaining agreements are prevalent in Dutch manufacturing. Second, if buyer power exists in the intermediate input market, we should find intermediate input price markdowns – wedges above unity. Finally, labor wedges and intermediate input wedges should be negatively

[22] In Appendix C, I give several examples of models where the intermediate input wedge captures the extent of buyer power.

related. Firms will exercise buyer power only if it increases quasi-rents and the extent of a wage markup depends on these quasi-rents. The following section outlines the empirical approach I use to investigate these predictions. Below, I first discuss the generality of this simple model.

While standard in the rent sharing literature, my theoretical framework abstracts from a more general bargaining setting with many employee associations and firms. Using more sophisticated bargaining models and allowing for adjustment frictions would substantially increase the notational complexity without altering the main predictions or the empirical analysis. The critical insight is that when employees have bargaining power, they will capture a share of their employer's quasi-rents leading to wage markups.

Note that not all firms are expected to have wage markups. Firms will only participate in bargaining if this leads to a weakly higher profit than their outside option. Firms with access to a sufficiently large competitive or monopsonistic labor market need not rely on an employee association to purchase labor. Likewise, when a firm with monopsony power bargains with an employee association, both wage markups and wage markdowns are possible outcomes, depending on the relative bargaining power of the two players (Falch and Strøm, 2007). Labor wedges should, therefore, not be expected to be below unity for all firms. However, given that collective bargaining agreements cover roughly 80 percent of Dutch employment, wage markups are expected to characterize most firms.

When inputs are not homogeneous, the input wedges introduced in this section are still informative. I model both labor and intermediate inputs as homogeneous goods because I do not observe input heterogeneity in my data, as is the case for most datasets on balance sheets and income statements such as Worldscope and Orbis. In reality, labor is heterogeneous, with high-skilled employees generally receiving higher wages and generating more revenue than low-skilled employees. In addition, intermediate inputs are typically a bundle comprising several inputs such as raw materials and energy. The input wedges introduced in equations (8) and (11) can be expressed as weighted averages of the wedges of the different input types, with weights that depend on revenue elasticities.[23] Therefore, the labor wedge should be interpreted as the average labor wedge across all labor markets the firms is active in, and likewise for the intermediate input wedge. Mertens (2021) provides a detailed theoretical underpinning of the "average input-wedge" interpretation of homogeneous input wedges.

3 Empirical approach

This section presents the empirical approach used to obtain input wedges and test key predictions of the theoretical model discussed in Section 2. Identification of input wedges is based

[23]For example, assume that there are two types of labor, indexed by $L1$ and $L2$. Equation (12) and $\frac{W_{it}L_{it}}{R_{it}} = \frac{W_{it}^{L1}L_{it}^{L1}}{R_{it}} + \frac{W_{it}^{L2}L_{it}^{L2}}{R_{it}}$ together imply that $\frac{\theta_{it}^{L}}{\gamma_{it}^{L}} = \frac{\theta_{it}^{L1}}{\gamma_{it}^{L1}} + \frac{\theta_{it}^{L2}}{\gamma_{it}^{L2}}$.

on revenue elasticities and is introduced in Section 3.1. Section 3.2 outlines the estimation of revenue elasticities, while Section 3.3 presents the data and descriptive statistics.

3.1 Identifying input wedges

Multiplying equation (8) by $\frac{L_{it}R_{it}}{L_{it}R_{it}}$ and equation (11) by $\frac{M_{it}R_{it}}{M_{it}R_{it}}$ gives

$$\gamma_{it}^{L} = \frac{\theta_{it}^{L}}{LS_{it}}, \text{ and } \gamma_{it}^{M} = \frac{\theta_{it}^{M}}{MS_{it}} \tag{12}$$

where $\theta_{it}^{L} = \frac{\partial R_{it}}{\partial L_{it}}\frac{L_{it}}{R_{it}}$ and $\theta_{it}^{M} = \frac{\partial R_{it}}{\partial M_{it}}\frac{M_{it}}{R_{it}}$ are the revenue elasticities of labor and intermediate inputs, LS_{it} is the labor share of revenue $\frac{W_{it}L_{it}}{R_{it}}$, and MS_{it} is the intermediate input share of revenue $\frac{P_{it}^{M}M_{it}}{R_{it}}$. Equation (12) requires firms to be active in both input markets and inputs to be substitutable, in line with the theoretical framework.[24] Input shares of revenue are observed in the data while revenue elasticities need to be estimated as they are not observed. Before discussing how I estimate revenue elasticities, I first compare my approach to other approaches to estimating input wedges.

The identification approach used in this paper is closely related to Petrin and Sivadasan (2013), who estimate a revenue function to obtain marginal revenue products of materials and electricity and compare these to industry-specific input price indices.[25] In contrast, equation (12) relies on an input's revenue elasticity and revenue share. My approach sidesteps problems that arise when input expenditure is deflated with an industry-wide price index but input prices are firm-specific.[26] Also related are papers that assume industry-specific input prices and use dispersion of calibrated marginal revenue products to study misallocation (e.g., Hsieh and Klenow (2009)).

In contrast, a rapidly growing literature estimates the labor wedge (e.g. Lu et al. (2019); Mertens (2021, 2022); Caselli et al. (2021); Brooks et al. (2021): Yeh et al. (2022)), or closely related measures based on the marginal revenue product of labor and the wage (e.g., Mertens (2020); Dobbelaere et al. (2020); Dobbelaere and Wiersma (2020)) using an approach relying on output elasticities and first-order conditions from cost minimization. This production 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

[24] For measuring input wedges when inputs are perfect complements, I refer to Rubens (2021).

[25] Petrin and Sivadasan (2013) are primarily interested in estimating a production function to identify an alternative labor wedge – the value of the marginal product of labor divided by the wage. In the case of competitive output markets, their measure is identical to the labor wedge discussed in this paper.

[26] Recently, Hashemi et al. (2022) suggest that researchers use revenue elasticities to identify input wedges. I provide one approach to doing this in practice.

markup of price over marginal cost by $\mu_{it} = \frac{P_{it}}{\lambda_{it}}$. For any input V_{it}, let

$$\mu_{it}(V_{it}) = \frac{\tilde{\theta}_{it}^{V}}{P_{it}^{V} V_{it}/R_{it}}, \tag{13}$$

where $\tilde{\theta}_{it}^{V}$ is the output elasticity of input V_{it} and P_{it}^{V} is its price. Under the assumptions that 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 $\mu_{it}(V_{it}) = \mu_{it}$.[27] The choice of V_{it} is crucial in recovering the true markup μ_{it}. If the required assumptions on V_{it} do not hold, $\mu_{it}(V_{it})$ is a joint measure of μ_{it} and imperfections in the input market for V_{it}, so that $\mu_{it} \neq \mu_{it}(V_{it})$. This insight has led the production literature to estimate the labor wedge by comparing markup estimates obtained using different inputs as V_{it}. Most recent work attempts to identify the labor wedge by comparing $\mu_{it}(L_{it})$ to $\mu_{it}(M_{it})$. In particular,

$$\frac{\mu_{it}(L_{it})}{\mu_{it}(M_{it})} = \frac{\tilde{\theta}_{it}^{L} P_{it}^{M} M_{it}}{\tilde{\theta}_{it}^{M} W_{it} L_{it}} = \frac{\theta_{it}^{L}/LS_{it}}{\theta_{it}^{M}/MS_{it}} = \frac{\gamma_{it}^{L}}{\gamma_{it}^{M}}, \tag{14}$$

where the second equality follows from $\frac{\partial R_{it}}{\partial x} = \frac{\partial R_{it}}{\partial Q_{it}} \frac{\partial Q_{it}}{\partial x}$ for $x \in \{L_{it}, M_{it}\}$. 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 (14) identifies the labor wedge. In general, however, the production approach obtains the labor wedge relative to the wedge of a different input, in this case intermediate inputs. Applying the production approach in the Dutch context is problematic. Due to the institutional setting discussed in Section 2.1, I do not want to *a priori* rule out buyer power in the intermediate input market by assuming that $\gamma_{it}^{M} = 1$. A particularly attractive feature of my approach is that it does not require a frictionlessly adjustable input for which firms are price takers to identify input wedges. Therefore, I can investigate the relation between input market imperfections in labor markets and intermediate input markets instead of focusing only on labor markets.[28]

3.2 Estimating revenue elasticities

Identifying input wedges using equation (12) requires estimating revenue elasticities. Several approaches to estimating revenue functions exist, each with their own assumptions. This

[27] The conditional cost function refers to the static cost function conditional on other choice variables of the firm, which are potentially determined by a dynamic maximization problem.

[28] An alternative is to give up identifying the labor wedge and instead use equation (14) to identify the labor wedge relative to the intermediate input wedge. For instance, Morlacco (2020) identifies the wedge of imported intermediates relative to the wedge of domestic intermediates. Another approach, also taken in Morlacco (2020), is to make parametric assumptions on demand and competition and then use calibration to obtain a measure of μ_{it} so that the imported intermediates wedge can be backed out of the relative wedge equation.

section outlines the control function approach I use to generate revenue elasticities. In Section 4.4, I show that my results are robust to several alternative approaches to obtaining revenue elasticities. Appendix B contains a more detailed description of the estimation routine outlined below.

Like Petrin and Sivadasan (2013), I use insights from the literature on production function estimation to identify revenue elasticities. Note that I observe revenue, but not output, so that revenue function estimation is not plagued by the output price bias that would occur when deflated revenue is used in the place of output when estimating a production function (Klette and Griliches, 1996; De Loecker and Goldberg, 2014). Using revenue elasticities to identify input wedges, therefore, sidesteps strong identification concerns known for related ratio estimators based on output elasticities – particularly markups (Bond et al., 2021; De Ridder et al., 2022).[29]

Consider the revenue function of firm i at time t

$$R_{it} = F_{it}(K_{it}, L_{it}, M_{it})\Omega_{it}, \tag{15}$$

where Ω_{it} is Hicks-neutral revenue productivity which is potentially known to the firm at time t, but unobserved by the econometrician. As revenue results from multiplying quantity (1) and inverse demand (2), the key assumption made to ensure that revenue is given by equation (15) is that $F^P_{it}(F^Q_{it}(K_{it}, L_{it}, M_{it})\Omega^Q_{it})$ is multiplicatively separable in $F^Q_{it}(K_{it}, L_{it}, M_{it})$ and Ω^Q_{it}. This ensures that a single term that enters (15) multiplicatively contains all demand and supply shocks: $\Omega_{it} = F^P_{it}(\Omega^Q_{it})\Omega^P_{it}$.[30]

Taking logs and allowing for log-additive mean-zero deviations from planned revenue, ϵ_{it}, gives

$$r_{it} = f_{it}(k_{it}, l_{it}, m_{it}; \theta) + \omega_{it} + \epsilon_{it}, \tag{16}$$

where lowercase letters denote the natural logarithm of the concomitant uppercase letter and θ is a vector containing coefficients. The main challenge to identifying θ 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} – which is clearly the case in the theoretical framework introduced in Section 2.2 – estimates of the revenue elasticities can be biased. A similar problem has long been recognized in the literature on production function estimation (Marschak and Andrews, 1944).

I use the control function approach, due to Olley and Pakes (1996) and Levinsohn and Petrin (2003), to deal with the correlation of unobserved revenue productivity and input choices.[31] A static input demand equation, demand for intermediate inputs, is used to

[29] An alternative solution to the output price bias is explicitly modeling the demand side (e.g., De Loecker (2011)).

[30] This holds, for instance, if $P_{it}(Q_{it}) = Q_{it}^{-\frac{1}{\epsilon}}\Omega^P_{it}$.

[31] See also Ackerberg et al. (2015) and Gandhi et al. (2020).

control for ω_{it},

$$m_{it} = m_d(\omega_{it}, k_{it}, l_{it}, \mathbf{x}^m_{it}; \delta^m), \tag{17}$$

where d is a 2-digit NACE industry, $\mathbf{x}^m_{it}$ contains additional control variables such as the firm's wage and 4-digit NACE industry intermediate input share, and δ^m contains all coefficients. Inverting equation (17) result in a control for ω_{it} based on observables which can be substituted into equation (16) under the "scalar unobservable" assumption – the assumption that revenue productivity ω_{it} is the sole unobservable in equation (17). In addition, $m_{it}(\cdot)$ is assumed to be one-to-one in ω_{it}. Note that revenue productivity ω_{it} potentially contains price variation due to demand shocks in output markets. Therefore, controlling for such price variation is unnecessary to satisfy the scalar unobservable assumption.[32]

Intermediate input demand can be derived from first-order condition (10) if one is willing to make parametric assumptions on a firm's demand, production function, and inverse intermediates supply. For expository purposes, consider a time-invariant Cobb-Douglas production function and an iso elastic demand so that the revenue function in equation (1) is given by $R_{it} = K_{it}^{\theta^K} L_{it}^{\theta^L} M_{it}^{\theta^M} \Omega_{it}$. Assuming inverse intermediates supply is iso elastic and given by $P^M_{it} = M_{it}^{(\varepsilon^M_{it})^{-1}}$, the first-order condition of intermediate inputs in equation (10) can be rewritten to provide an expression for intermediates demand

$$M_{it} = \left(\frac{\varepsilon^M_{it}}{1+\varepsilon^M_{it}} \theta^M K_{it}^{\theta^K} L_{it}^{\theta^L} \Omega_{it} \right)^{\frac{1}{1+(\varepsilon^M_{it})^{-1}-\theta^M}}, \tag{18}$$

which shows that satisfying the scalar unobservable assumption depends crucially on controlling for the intermediate input supply elasticity, in addition to capital and labor. This is done by including the firm's intermediate input share in $\mathbf{x}^m_{it}$. The intermediate input share controls for buyer power in a wide variety of models, as explained in Appendix C.

If the assumptions on demand, production, and intermediates supply are relaxed, the main point stands: satisfying the scalar unobservable assumption requires intermediate input demand to depend on other inputs and the supply elasticity of intermediate inputs.[33] Crucially, it is not necessary to control for bargaining power ϕ_{it}. Bargaining power affects intermediate input demand through its effect on labor, but not directly.

[32] The one-to-one assumption rules out adjustment frictions in intermediate inputs that are so severe that demand for intermediates does not respond to changes in revenue productivity. The main idea behind the control function approach is that the control, here intermediate inputs, can pin down unobserved revenue productivity.

[33] For instance, consider using a translog revenue function instead of a Cobb-Douglas revenue function. The key difference is that revenue elasticities are now firm-time-specific. However, revenue elasticity variation across firms and time is entirely due to difference in input use, so that including capital and labor in equation (17), in addition to a control for the intermediate input supply elasticity, is sufficient to satisfy the scalar unobservable assumption.

I use a translog revenue function, which results in firm-time specific revenue elasticities and is given by

$$f_d(k_{it}, l_{it}, m_{it}) = \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + \beta_{kk} k_{it}^2 + \beta_{ll} l_{it}^2 + \beta_{mm} m_{it}^2 + \beta_{kl} k_{it} l_{it} + \beta_{km} k_{it} m_{it} + \beta_{lm} l_{it} m_{it}, \tag{19}$$

where d indicates the 2-digit NACE industry at which the revenue function is estimated and the revenue elasticities are given by $\theta_{it}^x = \frac{\partial f_d(\cdot)}{\partial x_{it}}$, where $x \in \{k, l, m\}$. Note that the translog revenue function does not restrict revenue returns to scale and allows for a flexible elasticity of substitution compared to a standard CES revenue function. A translog revenue function would result, for instance, if a firm's production function is translog and its demand is iso-elastic. Alternatively, equation (19) can be seen as a second-order approximation to an unspecified revenue function.

While I observe the number of full-time equivalent employees, I use deflated intermediate input expenditure and the deflated book value of capital in place of intermediate inputs and capital. If firm-level input prices differ from the input-specific 2-digit industry level deflators, and these differences correlate with input choices, estimates of the revenue elasticities will be biased. The input price bias is especially relevant for intermediate inputs, as the revenue elasticity of intermediates is needed to construct its input wedge, and firm-specific buyer power could create within-industry input price differences.[34]

To control for within-industry input price differences, I use as control function

$$b_{it} = b_d((1, k_{it}, l_{it}, m_{it}) \times \mathbf{x}_{it}^b; \delta^b), \tag{20}$$

where the notation indicates that inputs enter b_d only interacted with the control variables in $\mathbf{x}_{it}^b$ – due to the translog specification – and δ^b contains all coefficients. Following De Loecker et al. (2020), I control for price variation due to unobserved quality differences by using market shares and 4-digit NACE 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.

To control for within-industry intermediate input price differences stemming from buyer power, I include a firm's 4-digit NACE industry intermediate input share in $\mathbf{x}_{it}^b$. Intermediate input shares control for intermediate input price variation in a wide class of models of buyer power. Examples are given in Appendix C. Note that input price biases only occur if input price variation around the 2-digit industry-specific deflators exists. Therefore, models of buyer power with a single input price as equilibrium outcome do not produce an input price bias – for example, a Cournot-style input market for intermediates. Finally, note that $b_{it}(\cdot)$

[34] Input price biases are well known but mostly ignored in production function estimation. See De Loecker and Syverson (2022) for a review of the literature.

can 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 arguments of equations (19), (20), and the inverse of equation (17), $\omega_{it} = m_d^{-1}(k_{it}, l_{it}, m_{it}, \mathbf{x}_{it}^m; \delta^m)$.[35] This first step is not meant to identify any coefficients, but rather to separate observed revenue into planned revenue and a revenue shock: $r_{it} = \hat{r}_{it} + \hat{\epsilon}_{it}$. 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 (12) by $\hat{LS} = \frac{W_{it}L_{it}}{R_{it}/\exp(\epsilon_{it})}$ and $\hat{MS} = \frac{P_{it}^M M_{it}}{R_{it}/\exp(\epsilon_{it})}$.

In the second step, I use as first-order Markov law of motion of revenue productivity

$$\omega_{it} = g_d(\omega_{it-1}; \delta^g) + \xi_{it}, \tag{21}$$

where ξ_{it} is a mean zero revenue productivity shock, $g_d(\cdot)$ a stochastically increasing function, and δ^g contains all coefficients. 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, $\hat{\xi}_{it}$, conditional on the still to be estimated coefficients.[36] Estimates are based on the following moment conditions

$$\mathbb{E}(\xi_{it}\mathbf{Z}_{it}) = 0, \tag{22}$$

where $\mathbf{Z}_{it}$ includes terms in $g_d(\cdot)$, a second-order polynomial in k_{it} and l_{it}, and interactions of contemporaneous capital and labor with all lagged inputs. The full unbalanced panel is used, and separate revenue functions are estimated for each 2-digit NACE industry. Estimates of revenue elasticities are reported in Table B1 of Appendix B.

Estimating a revenue function comes with several challenges, particularly concerning the treatment of unobserved revenue productivity. Equation (15) posits a Hicks-neutral revenue productivity term, while equation (21) states that revenue productivity evolves according to a first-order Markov process. These requirements are significant and, in particular, place strong restrictions on the underlying demand and productivity shocks. One solution is to specify particular demand and production functions. Akin to Petrin and Sivadasan (2013), assuming that $F(\cdot)^Q$ in equation (1) is a translog function and that demand is iso elastic, $P_{it}(Q_{it}) = Q_{it}^{-\frac{1}{\epsilon}}$, would alleviate concerns regarding the treatment of revenue productivity. To ensure that my results hold more generally, robustness to alternative ways to obtain

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

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

revenue elasticities is crucial. Section 4.4 shows that my results are robust to, among other things, moments based on different timing assumptions, estimating the revenue function at the 4-digit industry instead of the 2-digit industry level, estimating year-by-2-digit industry revenue functions, and calibrating sample-wide revenue elasticities so all variation is driven by revenue shares.

3.3 Data and descriptive statistics

I construct a yearly firm-level dataset covering Dutch manufacturing firms with at least one employee over the period 2007 to 2018, using non-public data obtained from Statistics Netherlands (CBS).[37] I combine data from the "General Firm Registry" (ABR) and the "Financial Statistics of Non-financial Firm" (NFO). These two yearly firm-level datasets use as primary sources registries 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, the book value of capital, and earnings before interest and taxes (EBIT).

The CBS routinely controls quality by contacting firms when reporting errors are suspected. In addition, I remove outliers and observations that report internally inconsistent statistics. I further 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 21,293 firms for a total of 121,057 firm-year observations. Appendix A provides a detailed overview of all variables and the sample selection procedure. Table A2 presents a breakdown of observations by year and 2-digit NACE 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. In the absence of labor market imperfections other than taxes, firms can be expected to equalize the marginal revenue product of labor and the per-employee variable expenditure on labor, not the marginal revenue product of labor and take-home wage of the employee. With slight abuse of terminology, I refer to labor expenditure as "wages" in order to remain in line with the broader academic literature, which refers to total labor expenditure in this

[37]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	730	1,739	4,502	6,998	71,418
Capital	96	372	1,081	1,679	14,921
Labor (FTE)	4	10	23	23.33	56.24
Labor expenditure	233	522	1,229	1,435	5,742
Wage	43	53	66	58	23
Labor share	0.22	0.31	0.41	0.32	0.15
Intermediate input expenditure	377	970	2,731	4,939	65,464
Intermediate input share	0.48	0.58	0.69	0.58	0.15
EBIT	9	87	289	417	3,141

Notes: Summary statistics for key variables based on the full sample of 121,057 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.

fashion.[38] I construct the wage W_{it} as the ratio of a firm's total expenditure on labor to its FTE employees, which fits both the homogenous labor theoretical framework of Section 2.2 and the average labor wedge interpretation of the labor wedge when there is unobserved labor heterogeneity. 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.[39]

When interpreting W_{it}, note that net wages are roughly half of total labor expenditure. This gap is due to mandatory social security contributions, such as pension contributions by both employers and employees, and high income tax rates. In 2019, three percent of total labor expenditure consisted of recruiting and on-the-job training costs. Employers' mandatory social contributions towards employees constituted another 21 percent of total labor expenditure, leaving gross salary at 76 percent of total labor expenditure. Employees' mandatory social security contributions and income taxes accounted for another 24 percent, leaving net salary at 52 percent of a firm's total labor expenditure (CBS, 2020a, p.77).[40]

Eurostat's NACE Rev. 2 industry classification is used throughout this paper (Eurostat, 2008). The most aggregated industry classification is the NACE section, while NACE

[38]In the Netherlands, the term "wage" is typically used to refer to gross salary.

[39]This includes the revenue elasticity of labor. Unobserved employment enters unobserved revenue productivity, ω_{it}, which the estimation routine controls for.

[40]These percentages are fairly stable throughout the sample period. For instance, in 2018, net wages made up 51 percent of total labor expenditure (CBS, 2019a, p.87).

divisions (2-digit), NACE groups (3-digit), and NACE codes (4-digit) are increasingly disaggregated industry classifications.[41] All NACE divisions in manufacturing with sufficient observations such that the revenue function can be estimated are included in the final sample. Table A1 in Appendix A list the 18 2-digit NACE industries that are included in the final sample.

In the revenue function estimation, firm location is considered by indicating the NUTS1 region (nomenclature of territorial units for statistics) where a firm's headquarters is located. The NUTS classification is Eurostat's system for dividing the economic territories of the EU in order to enable socio-economic analyses of regions (Eurostat, 2020).

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, publicly-traded firms found in international datasets such as Compustat. Compared to those datasets, my sample includes many smaller firms. The median number of employees is 10, while the mean is 23.33. The presence of small firms does not drive my results, which are robust to the omission of firms with fewer than, for instance, 5 employees. Total labor cost has a (rounded to thousands) median of 522,000 euros and a mean of 1,435,000 euros. The median book value of capital is 372,000 euros, while the mean is 1,679,000 euros. For revenue, the median and mean are 1,739,000 euros and 6,998,000 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.

Table 2: Summary statistics of input wedges

	p(5)	p(25)	p(50)	p(75)	p(95)
Labor wedge	0.53	0.73	0.86	0.99	1.26
Intermediate input wedge	0.99	1.09	1.18	1.30	1.57
Relative input wedge	0.36	0.56	0.72	0.89	1.24

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 121,057 firm-year observations covering the years 2007 to 2018.

[41] As an example: NACE section C is "manufacturing", which consists of NACE divisions 10-33. NACE division 10 is "Manufacture of food products" and includes nine different NACE groups. NACE group 101 is "Processing and preserving of meat and production of meat products", which includes three different NACE codes. NACE code 1011 is "preserving and processing of meat".

Figure 1: Density function of firm-level labor wedges (panel a) and intermediate input wedges (panel b)

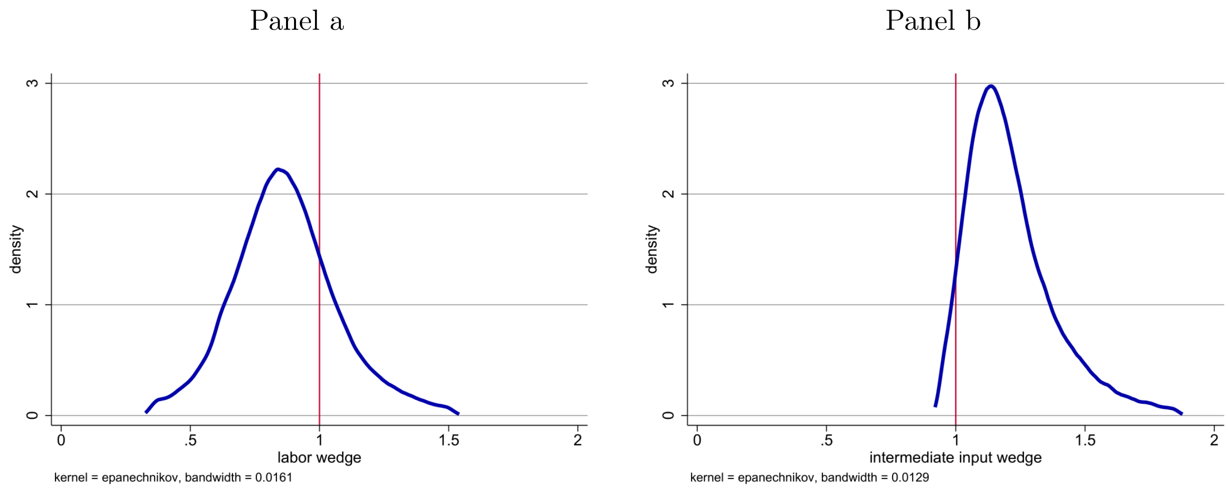

Notes: Kernel density functions of firm-level input wedges. Based on the full samples of 121,057 observations covering the years 2007 to 2018. Top and bottom percentiles trimmed.

4 Results

In Section 4.1, I analyze the distributions of firm-level input wedges in Dutch manufacturing, both in the cross-section and over time. Section 4.2 explores the relation between the intermediate input wedge and the labor wedge. Section 4.3 provides firm-time-specific rent sharing elasticities and estimates of bargaining power. Robustness checks are discussed in Section 4.4.

4.1 Evidence of wage markups and buyer power

Figure 1 displays the distributions of the firm-level labor and intermediate input wedges in the entire sample, and Table 2 gives several percentiles of these distributions. Wage markups are prevalent in Dutch manufacturing. The median firm-year has a labor wedge equal to 0.86 – that is, the wage is 16.28 percent higher than the marginal revenue product of labor. In total, 76.47 percent of all observations are characterized by wage markups.

Wage markups are not consistent with monopsony power but are expected when employees have bargaining power and can extract rents from their employers. The median wage markup of roughly 16 percent is in line with the 20 percent wage markup that employees in Portugal receive on average – a country also characterized by high collective bargaining coverage (Card and Cardoso, 2022).[42] In contrast, Yeh et al. (2022) report that nearly 90

[42] Card and Cardoso (2022) estimate the markup of wages over sectoral wage floors. Note that in the model discussed in Section 2.2, optimality conditions guarantee that the outside option $\bar{W}_{it}$ is equal to $MRPL_{it}$.

percent of all plant-year observations in US manufacturing are characterized by wage markdowns and interpret this as evidence of monopsony power. I conclude that the data support the hypothesis that employees generally have bargaining power in Dutch manufacturing, but not the hypothesis that firms have monopsony power.

Intermediate inputs price markdowns characterize Dutch manufacturing. The median intermediate input wedge is 1.18, corresponding to a markup of the marginal revenue product of intermediates over their price of 17.84 percent. This implies that, at the median firm, intermediate input suppliers receive only 85 cents for the marginal euro of revenue they generate. Price markdowns cover more than 95 percent of all 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 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 (14) 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 shifts the distribution of the labor wedge to the left compared to the labor wedge distribution identified based on equation (12). Wage markups are even more prevalent in my sample when using the production approach. This is unsurprising, given that intermediate input wedges are typically above unity in Dutch manufacturing.

Figure 1 shows that both input wedge distributions exhibit substantial dispersion. One possibility is that this dispersion is driven by variation over time. Adjustment frictions, for instance, could make input wedges sensitive to economic shocks. Alternatively, there might be significant between-industry differences, for instance, due to regulation. Finally, within-industry differences between firms might drive the observed dispersion.

Figure 2 plots selected percentiles of the two input wedge distributions over time. Both distributions are remarkably stable over the sample period. The cross-sectional dispersion of both input wedge distributions is much larger than the variation over time of a given percentile. 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 drivers of input wedge dispersion are primarily cross-sectional factors, not temporal factors such as economic shocks combined

Therefore, under the assumption that $\bar{W}_{it}$ is the sectoral wage floor, wage markups measures in this paper are directly comparable to those discussed in Card and Cardoso (2022).

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

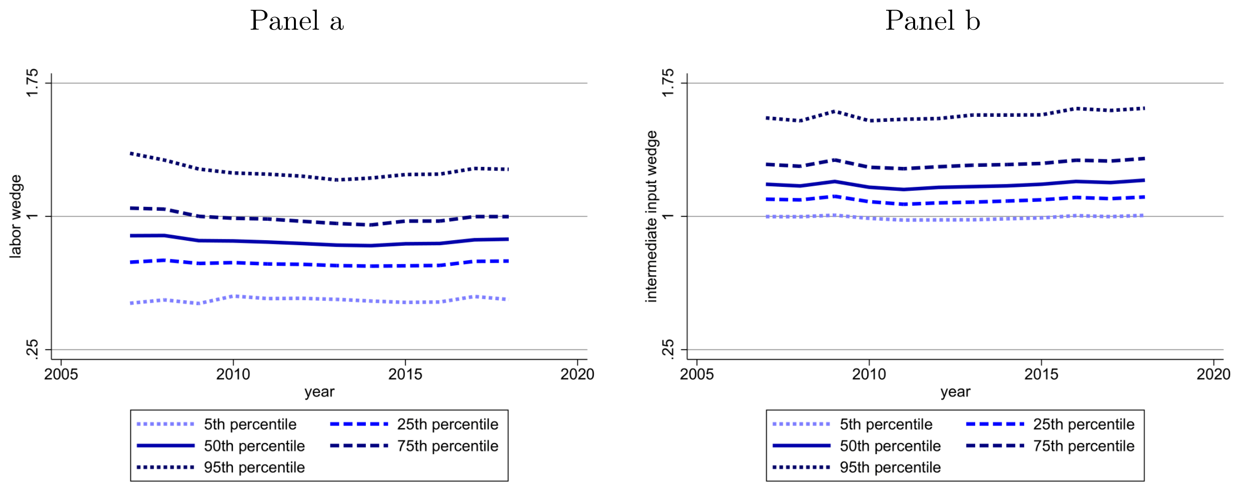

Notes: Selected percentiles of input wedge distributions over time. Based on the full samples of 121,057 observations covering the years 2007 to 2018.

with adjustment frictions.

Observed cross-sectional dispersion of input wedges is primarily due to within-industry variation, not between-industry variation. Table 3 reports statistics based on percentiles of the two input wedge distributions at different levels of industry aggregation. The average over all industries of the labor wedge distribution's inter-quartile range is nearly identical at the 2-digit (0.26), 3-digit (0.28), and 4-digit industry (0.27) levels. Moreover, the standard deviation of this average is also comparable at the 2-digit (0.06), 3-digit (0.06), and 4-digit (0.09) levels, implying that between-industry differences in the labor wedge distribution's inter-quartile range are not increasing as industry definitions become increasingly narrow. The same is true for the distribution of the intermediate input wedge, where the average of the inter-quartile range lies between 0.19 (4-digit level) and 0.21 (2-digit level), with standard deviations between 0.05 (2-digit level) and 0.07 (4-digit level). Performing the same analysis on the median and the difference between the 95th and 5th percentile yields similar results, showing that the input wedge distributions are quite similar at different levels of aggregation. The observed dispersion of input wedges is therefore likely related to idiosyncratic factors, such as rent sharing, instead of industry-level factors such as regulation or industry institutions.

Three main conclusions emerge from analyzing the input wedge distributions. First, labor wedges in Dutch manufacturing are not in line with monopsony power but rather indicate that employees extract rents from their employers. This difference with the US-based monopsony findings of Yeh et al. (2022) is not due to methodological differences.

Table 3: Input wedge statistics at different levels of industry aggregation

	Labor wedge			
Level	average p(50)	average (p(75) - p(25))	average (p(95) - p(5))	observations
Full sample	0.86 (-)	0.26 (-)	0.73 (-)	1
Division (2-digit)	0.84 (0.07)	0.26 (0.06)	0.74 (0.13)	18
Group (3-digit)	0.86 (0.08)	0.28 (0.09)	0.75 (0.18)	78
Code (4-digit)	0.86 (0.10)	0.27 (0.09)	0.73 (0.19)	167
	Intermediate input wedge			
Full sample	1.18 (-)	0.21 (-)	0.58 (-)	1
Division (2-digit)	1.18 (0.04)	0.21 (0.05)	0.58 (0.10)	18
Group (3-digit)	1.17 (0.07)	0.20 (0.06)	0.58 (0.13)	78
Code (4-digit)	1.17 (0.07)	0.19 (0.07)	0.54 (0.16)	167

Notes: Input wedge statistics at different levels of NACE industry classifications (more digits is a more narrowly defined industry). Averages are taken over all industries at the level, standard deviation in brackets. Number of units on which the average is based listed under "observations". 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. Based on the full sample of 121,057 observations covering the years 2007 to 201

Second, intermediate input wedges suggest that buyer power for intermediates is common in Dutch manufacturing. Third, within-industry differences between firms drive the substantial input wedge dispersion.

If employees have bargaining power, equation (9) shows that their wage depends on the quasi-rents a firm generates. Given the prevalence of collective bargaining agreements and wage markups, the evidence points to such bargaining power existing in the Netherlands. A firm's quasi-rents are, therefore, a likely candidate for explaining within-industry labor wedge variation. Buyer power in the intermediate input market leads to higher quasi-rents per employee, suggesting a negative relation between the two input wedges. The following section investigates the determinants of within-industry labor wedge dispersion – in particular, to what extent dispersion of the intermediate input wedge can explain the dispersion of wage markups.

4.2 Determinants of wage markup heterogeneity

Within-industry wage markup heterogeneity could be driven primarily by differences in the revenue-generating ability of employees across firms, or by differences in labor compensation.

To separate these explanations, the following non-parametric regression is used to correlate the labor wedge with a variable of interest, x_{it},

$$ln(\gamma_{it}^{L}) = \beta_0 + \sum_{dc=2}^{10} \beta_{dc}^{x} \mathbb{I}_{x_{it} \in X_{dc}} + NACE \times Year_{it} + \varepsilon_{it}, \tag{23}$$

where $\mathbb{I}_{x_{it} \in X_{dc}}$ is an indicator that equals 1 if x_{it} lies between the the dc^{th} and $dc - 1^{th}$ decile of the distribution of x_{it} in the full sample and $NACE \times Year_{it}$ is a set of year-by-4-digit industry 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 4-digit industry level to take into account any within-industry dependence between observations.

Figure 3 plots point estimates and confidence intervals from regressions relating the labor wedge to revenue per employee (panel a) and the wage (panel b). Except for average revenue's top two deciles, revenue per employee and the labor wedge are unrelated. For the top two deciles, firms with higher revenue per employee grant their employees higher wages on the margin. For the vast majority of firms, however, the average revenue their employees create is uncorrelated to the extent of wage markups. This suggests that cross-firm differences in labor compensation are the primary driver of within-industry labor wedge variation. Indeed, a very pronounced negative relationship exists between a firm's wage and its labor wedge. Moving from the bottom to the top decile of the wage distribution induces all variation in the labor wedge distribution. High wage markups go hand-in-hand with high wages, but not with highly productive employees.

Figure 4 displays coefficients from regressions associating the labor wedge (panel a) and the wage (panel b) to the intermediate input wedge. A pronounced negative association between the two input wedges exists. Firms with high price markdowns for intermediate inputs have large wage markups. In other words, firms that pay intermediate input suppliers less than their marginal revenue contribution pay their employees more than their marginal revenue contribution. As revenue per employee is mostly unrelated to the labor wedge, the relation between the two input wedges is likely driven by the compensation of employees. Indeed, panel b of Figure 4 shows that intermediate input wedges and wages are strongly positively related. Overall, the data suggest that firms with buyer power in the intermediate input market share those rents with their employees by setting high wages.[43]

Highlighting the role that intermediate input markets play in within-industry wage markup dispersion does not imply that other idiosyncratic factors are irrelevant. Figure D1 in Appendix D shows that the labor wedge is positively associated with both total employment and capital. Figure D2 in Appendix D reveals that firms with higher labor wedges make more revenue and have slightly higher accounting profit than firms with low labor wedges.

[43] A similar result is observed in Neven and Röller (1996), who provide evidence from the European airline industry that lax competition in output markets induces rent sharing by setting excessive wages.

Figure 3: Regressions of firm-level labor wedges on revenue per employee (panel a) and wage (panel b), over time

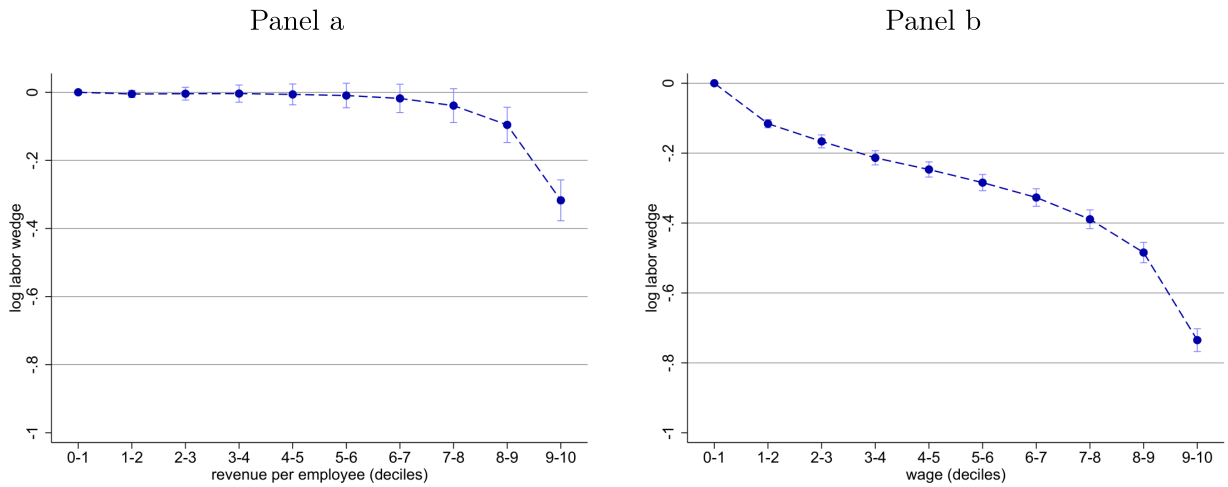

Notes: Regression coefficients and 95 percent confidence intervals from a non-parametric regression of log labor wedge on revenue per employee (panel a), and wage (panel b). Explanatory variables are divided into 10 categories based on the deciles of their distribution. The coefficient on the indicator of the first decile is normalized to 0. Year-by-4-digit NACE industry fixed effects included. Standard errors are clustered at the 4-digit industry level. Based on the full samples of 121,057 observations covering the years 2007 to 2018.

However, the variation in the labor wedge that these factors can induce is significantly lower than the variation induced by the wage and the intermediate input wedge. In addition, all results presented in this section are both qualitatively unchanged and quantitatively virtually identical when controlling for firm size in the regressions.

Summing up, a key predictor of the extent of wage markups is the extent to which firms can generate rents in the intermediate input market. This link operates mainly through wages rather than through the marginal revenue product of labor. As differences in revenue per employee can not explain the dispersion of wage markups, between-firm quality differences of employees are unlikely to drive my results. I conclude that the data support the hypothesis that the distribution of wage markups is shaped by the sharing of rents generated in one input market with suppliers in a different input market. In line with this, the remainder of the paper focuses on obtaining rent sharing elasticities and estimates of bargaining power.

Figure 4: Regression of firm-level labor wedges (panel a) and wages (panel b) on the intermediate input wedge, over time

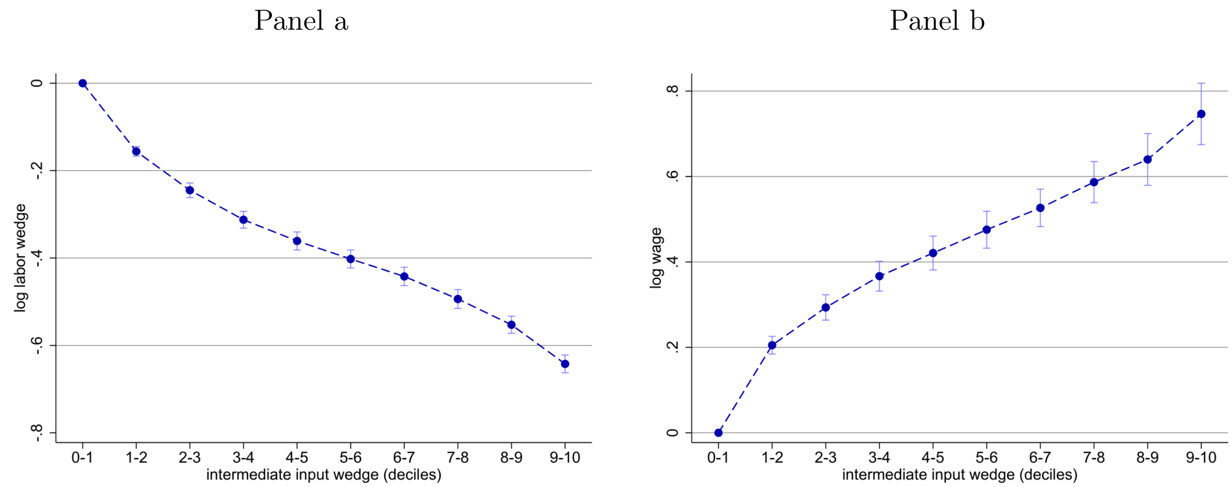

Notes: Regression coefficients and 95 percent confidence intervals from a non-parametric regression of log labor wedge (panel a), and log wage (panel b) on the intermediate input wedge. Explanatory variables are divided into 10 categories based on the deciles of their distribution. The coefficient on the indicator of the first decile is normalized to 0. Year-by-4-digit NACE industry fixed effects included. Standard errors are clustered at the 4-digit industry level. Based on the full samples of 121,057 observations covering the years 2007 to 2018.

4.3 Rent sharing elasticities

Rent sharing elasticities can be identified based on the bargaining framework introduced in Section 2.2. Combining first-order conditions (6) and (7) gives

$$W_{it} = \bar{W}_{it} + \phi_{it}\left(\frac{QR_{it}}{L_{it}}\right), \tag{24}$$

which shows that when employees have bargaining power ($\phi_{it} > 0$), they receive a fraction of the firm's quasi-rents, driving their wages up above the outside option $\bar{W}_{it}$. The empirical rent sharing literature is motivated by such relationships between quasi-rents per worker and wages (see Card et al. (2018) for a survey). The central object of interest in this literature is the elasticity of wage with respect to quasi-rents per employee, which can be expressed as

$$\varepsilon_{it}^{W,QR} = \frac{\phi_{it}(\frac{QR_{it}}{L_{it}})}{\bar{W}_{it} + \phi_{it}(\frac{QR_{it}}{L_{it}})}. \tag{25}$$

As a firm's quasi-rents are typically unobserved, most researchers have resorted to estimating the elasticity of wage with respect to value added per worker instead. In practice, researchers regress log wage on log value added per employee and several controls, such as capital per employee. As employees' outside option $\bar{W}_{it}$ is typically unobserved, instruments for value-added that do not shift $\bar{W}_{it}$ are required. This regression-based approach delivers an average rent sharing elasticity, rather than the firm-time specific elasticities given in equation (25).

I identify firm-time-specific rent sharing elasticities without needing to observe the employees' outside option or the firm's quasi-rents. Identification relies on the theoretical framework introduced in Section 2.2 and estimates of the revenue elasticity of labor. Note that first-order conditions (6) and (7) imply that $\bar{W}_{it}$ equals $MRPL_{it}$ – that is, firms facing employees with outside option $\bar{W}_{it}$ act like firms in a competitive labor market facing wage $\bar{W}_{it}$. Together with equation (24), this implies that a firm's quasi-rents per employee can be written as $\frac{W_{it}-MRPL_{it}}{\phi_{it}}$. Plugging this into equation (25) and again using $\bar{W}_{it} = MRPL_{it}$ allows the elasticity of wage with respect to quasi-rents per employee to be written as

$$\varepsilon_{it}^{W,QR} = \frac{W_{it}L_{it} - \theta_{it}^{L}R_{it}}{W_{it}L_{it}} = \frac{W_{it} - MRPL_{it}}{W_{it}}, \tag{26}$$

which shows two things. First, the rent sharing elasticity is a Lerner-index measure of wage markups. Second, the firm-time specific elasticity of wage with respect to quasi-rents per employee can be recovered from an estimate of the revenue elasticity of labor and data on revenue and total wage expenditure. Essentially, the theoretical framework shows that revenue function estimates can be used to obtain information on the outside option $\bar{W}_{it}$ and a firm's quasi-rents.

Figure 5 plots the density function of rent sharing elasticities and Table 4 gives several percentiles of their distribution. The median elasticity of wage with respect to quasi-rents per employee is 0.20. Firms pay their employees on average 0.22 percent more following a 1-percent increase in quasi-rents per employee. The distribution of elasticities is right-skewed, with wages increasing by more than 0.5 percent following a 1 percent increase in quasi-rents per employee for only 5.2 percent of all firm-year observations. The significant dispersion observed in Figure 5 shows the importance of allowing for firm-time-specific elasticities. However, the mean elasticity is comparable to estimates obtained in the regression-based rent sharing literature surveyed in Card et al. (2018). Van Reenen (1996), instrumenting for quasi-rents per employee using innovation, reports an elasticity of 0.29 for UK manufacturing. Using patent-induced shocks to labor productivity, Kline et al. (2019) estimates an elasticity of wage with respect to value added per employee of 0.47, which corresponds to a quasi-rents based elasticity of roughly 0.25.[44]

The rent sharing elasticity alone can not pin down the employees' bargaining power. To isolate the bargaining power of employees, use $MRPL_{it} = \bar{W}_{it}$ and rewrite equation (25) as

$$\phi_{it} = \frac{\varepsilon_{it}^{W,QR}}{1 - \varepsilon_{it}^{W,QR}} \frac{\theta_{it}^{L}R_{it}}{QR_{it}}, \tag{27}$$

[44]Multiplying an elasticity based on value added by $\frac{VA_{it}}{QR_{it}}$ gives the corresponding elasticity based on quasi-rents. The literature shows that a good rule of thumb is $\frac{VA_{it}}{QR_{it}} = 2$, see Card et al. (2018).

Figure 5: Density function of firm-level rent sharing elasticities

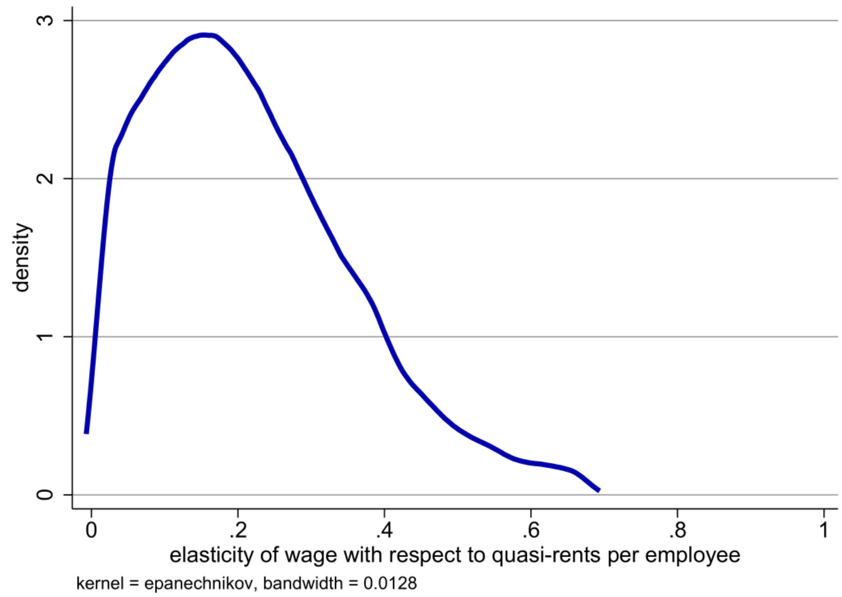

Notes: Kernel density function of the firm-level elasticity of wage with respect to quasi-rents per employee. Based on the sample of 92,568 firm-year observations for which $\gamma_{it}^{L} < 1$. Top and bottom percentiles trimmed.

which makes clear that the only missing element needed to identify bargaining power is quasi-rents QR_{it}. Two elements of a firm's quasi-rents are unobserved: the rental rate of capital P_{it}^{K} and the firm's outside option $\bar{\Pi}_{it}$. Table 4 lists selected percentiles of the distribution of employees' bargaining power assuming either $P_{it}^{K} = 0.10$ or $P_{it}^{K} = 0.12$. Following most of the literature, I assume that $\bar{\Pi}_{it} = 0$, which implies that the estimates in Table 4 should be seen as a lower bound on ϕ_{it}.

Table 4 shows that bargaining power is distributed reasonably symmetrically between firms and employees. The median bargaining power of employees is about 0.5, with roughly 50 percent of all firm-year instances falling between 0.30 and 0.65. The calibrated rental rate of capital appears to have little effect on the estimates. These results indicate that, conditional on $\bar{\Pi}_{it} = 0$, the median firm splits increases in quasi-rents per employee approximately 50-50 between profit and higher wage expenditure. Recall from the discussion in Section 3.3 that this entails an increase of roughly 25 cents in net wages following a 1 euro increase in quasi-rents per employee.

The unobservability of outside option $\bar{W}_{it}$ is one of the key problems in the regression-based rent sharing literature. To see the importance of using $MRPL_{it}$ to control for $\bar{W}_{it}$, and as a comparison to the rent sharing literature, Table 5 reports regression output that relates

Table 4: Summary statistics of rent sharing elasticities and employees' bargaining power

	p(5)	p(25)	p(50)	p(75)	p(95)
Rent sharing elasticity	0.03	0.11	0.20	0.31	0.51
Employees' bargaining power ($P^K = 0.10$)	0.07	0.29	0.47	0.65	0.90
Employees' bargaining power ($P^K = 0.12$)	0.08	0.30	0.49	0.67	0.95

Notes: Summary statistics of the elasticity of wage with respect to quasi-rents per employee and employees' bargaining power, ϕ_{it}. P^K refers to the price of capital that is assumed when constructing bargaining power. p(5), p(25), p(50), p(75), and p(95) refer to the 5th, 25th, 50th, 75th, and 95th percentile of the rent sharing elasticity distribution, respectively. Percentiles are rounded to two decimal points. Statistics based on the sample of 92,568 firm-year observations for which $\gamma_{it}^L < 1$.

wages to quasi-rents per employee.[45] Column (1) shows that regressing wage on quasi-rents per employee results in an implausibly large coefficient estimate of 1.35. Including year-by-4-digit NACE industry fixed effects effectively halves the point estimate, showing that the specification is extremely sensitive to including additional control variables. The regression in column (3) includes the marginal revenue product of labor as a control variable, so that first-order condition (7) is being estimated while restricting $\phi_{it} = \phi$. In line with the first-order condition, the coefficient on $MRPL_{it}$ is statistically indistinguishable from 1. Bargaining power is estimated to be about 0.5, in line with the median ϕ_{it} reported in Table 4. In addition, after controlling for the marginal revenue product of labor, estimated coefficients are stable when including additional controls, as shown by column (4). Columns (5) to (8) show that the same insights hold when estimating the elasticity of wage with respect to quasi-rents per employee. Erroneously omitting the marginal revenue product of labor – not controlling for $\bar{W}_{it}$ – leads to implausibly large estimated elasticities that are sensitive to the regression specification. Including $MRPL_{it}$ leads to stable estimates in line with the median $\varepsilon_{it}^{W,QR}$ obtained in Table 4.

Summing up, the mean rent sharing elasticity in Dutch manufacturing is 0.22. This is in line with, but slightly below, estimates of Van Reenen (1996) and Kline et al. (2019). At the median firm, bargaining power is distributed equally between the firm and its employees. Obtaining firm-time specific elasticities and bargaining power relies on estimates of the marginal revenue product of labor and first-order condition from the efficient bargaining model discussed in Section 2.2. Alternatively, when identifying average rent sharing elasticities using regressions, the marginal revenue product of labor can be used to control for the employees' outside option.

[45] I assume that $P_{it}^K = 0.1$ and $\bar{\Pi}_{it} = 0$. The results are robust to relaxing these assumptions.

Table 5: Rent sharing regressions

	Dependent variable: W_{it}			
	(1)	(2)	(3)	(4)
$\frac{QR_{it}}{L_{it}}$	1.35	0.69	0.53	0.56
	(0.032)***	(0.016)***	(0.012)***	(0.012)***
$MRPL_{it}$			0.97	1.02
			(0.010)***	(0.021)***
NACE × Year FE	no	yes	no	yes
R^2	0.79	0.97	0.98	0.98
observations	92,568			
	Dependent variable: $log(W_{it})$			
	(5)	(6)	(7)	(8)
$\log \frac{QR_{it}}{L_{it}}$	1.08	0.41	0.25	0.29
	(0.004)***	(0.043)***	(0.007)***	(0.005)***
$\log MRPL_{it}$			0.79	0.66
			(0.006)***	(0.039)***
NACE × Year FE	no	yes	no	yes
R^2	1	1	1	1
observations	92,230			

Notes: Table 5 reports output of regressions relating (log) wage to (log) quasi-rents per employee and the (log of the) marginal revenue product of labor, omitting a constant term and based on the sample of firm-year observations for which $\gamma_{it}^L < 1$. NACE×Year FE refers to whether indicator variables for 4-digit NACE industry-year combinations are included. Quasi-rents per employee constructed under the assumption that $P_{it}^K = 0.1$. Standard errors clustered at the 4-digit NACE industry level in brackets. *** indicates significance at the 1% level.

4.4 Robustness

This section discusses the robustness of the input wedge distributions, the negative correlation between the two input wedges, and the distribution of rent sharing elasticities. In particular, I show that all key results are robust to alternative identification strategies regarding the revenue elasticities. Indeed, simply calibrating a single revenue elasticity for each input preserves all key results. The reason is that my results are driven by variation of the revenue shares of input expenditure, not by variation of revenue elasticities. In what follows, I refer to the estimation approach introduced in Section 3.2 as the "baseline specification". Tables E1 to E4 in Appendix E provide results for all robustness checks. More details on revenue function estimation are in Appendix B.

Calibrating a single revenue elasticity for each input My first approach is to 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 its median value in the total sample of 121,057 observations. The revenue elasticity of intermediate inputs is calibrated similarly.[46] This calibration approach is not subject to the identification concerns surrounding revenue function estimation. Investigating the input wedges in equation (12) shows that key results on correlations between input wedges and their dispersion are now entirely driven by revenue shares of input expenditure.

Tables E1 and E2 show that calibrating revenue elasticities does not change the fact that, in most cases, firms have wage markups and intermediate input price markdowns. Unsurprisingly, not allowing revenue elasticities to differ across firms or time introduces spurious input wedge dispersion because firms tend to spend more on inputs that have higher revenue elasticities. Table E3 shows that the two input wedges are still strongly negatively correlated within industry-year. Finally, Table E4 shows that while the tails of the distribution of rent sharing elasticities are similar compared to the baseline results, mass has shifted to the right so that the median firm now increases wages by 0.3 percent following a 1 percent increase in quasi-rents. Summing up, key results are qualitatively unchanged.

Omitting the input price control function Next, I estimate revenue elasticities while omitting the control function for unobserved firm-specific input price variation around the 2-digit input price deflator. Instead, I employ a more parsimonious first stage. Specifically, the input price control function in equation (20) is omitted and all variables in $\mathbf{x}_{it}^{m}$ are dropped from the intermediate input demand function given in equation (17). This simplified specification – referred to as the "parsimonious specification" – allows for a direct comparison to papers that estimate labor wedges using similar first stages, in particular Yeh et al. (2022). In addition, this specification serves as a starting point for the remaining robustness checks. As the baseline specification is quite demanding of the data, further refinements are largely infeasible there.

Tables E1 and E2 show that the intermediate input wedge distribution is almost unchanged when using the parsimonious specification. The labor wedge distribution is similar except in the far tails, which are now more spread out. A clear negative relation between the input wedges still exists (Table E3). Due to the longer left tail of labor wedge distribution, the distribution of rent-sharing elasticities moves slightly to the right, with a 1 percent

[46] Alternative calibration choices, such as using a reasonable guess for the two elasticities or taking the median within a firm's 2- or 4-digit industry, yield similar results.

increase in quasi-rents now associated with an 0.27 percent increase in the wage (Table E4). These results are in line with research on production function estimation and markups, which shows that a simple control function for unobserved productivity is already sufficient to obtain accurate estimates of elasticities and markups (De Ridder et al., 2022), and that omitting the input price control tends to not matter much in practice (Collard-Wexler and De Loecker, 2015).

Estimating revenue functions at the year-by-2-digit level The baseline specification used in the main text rules out non-Hicksian revenue productivity. Considerable attention has recently been devoted to labor-augmenting technological change which is ruled out by the Hicks neutral framework used in this paper (e.g., Doraszelski and Jaumandreu (2018); Raval (2019); Oberfield and Raval (2021)). 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. To allow for even more flexibility and to address potential deviations from Hicks neutrality, I estimate the parsimonious specification separately for each year-by-2-digit industry combination.[47] Then, for a given firm, revenue elasticities are allowed to differ flexibly over time not only because input use changes but also because the derivatives of the revenue function with respect to inputs changes.

Estimating revenue functions at the 4-digit level Relatedly, the translog revenue function used in the baseline specification allows elasticities to vary in the cross section due to differences in input utilization but restricts the function of inputs that determines these elasticities to be constant within a 2-digit industry. However, underlying demand and production functions might differ within a 2-digit industry. I, therefore, estimate the parsimonious specification separately for each 4-digit industry, allowing the revenue function to vary within a 2-digit industry.[48]

Alternative timing assumption The moments given in equation (22) – used to identify the revenue function in the baseline specification – assume that contemporaneous labor is orthogonal to the contemporaneous mean-zero revenue productivity shock ξ_{it}. In the setting introduced in Section 2.2 this would be true, for instance, if a firm's information set contains

[47]Note that this implies dropping 2007 from the sample, as estimation requires one-year lags of a firm's inputs. Estimating the baseline specification used in the main text on 3-year-by-2-digit cells is feasible, and all main results are robust to such an approach.

[48]The most flexible approach given the data, estimating revenue functions at the year-by-4-digit level, is infeasible even when using the parsimonious specification. Even 3-digit industries typically do not have sufficient observations to estimate a revenue function with the baseline specification. For similar reasons, two 4-digit industries are omitted from the 4-digit level parsimonious specification (NACE codes 1624 and 2053).

$\mathbb{E}(\omega_{it}) = g(\omega_{it-1})$ but not ξ_{it}. If we are unwilling to make such a timing assumption, l_{it} could depend on ξ_{it}, invalidating some of the moments used to obtain the baseline estimates. I, therefore, estimate the parsimonious specification without assuming that l_{it} and ξ_{it} are orthogonal to each other. This is done by removing all moments relying on l_{it}.

Tables E1 to E4 show that all key results are remarkably robust to the alternative timing assumption, the 4-digit estimation approach, and the year-by-2-digit estimation approach. In fact, the results for these robustness checks are quantitatively nearly identical to the parsimonious specification results. Overall, all main results are unchanged for all robustness checks.

The robustness checks discussed above show that my results are not meaningfully affected by the identification of revenue elasticities. This does not imply that I claim that all maintained assumptions in the baseline specification hold in general. Instead, where assumptions might be violated, this section has shown that all key results are robust. The reason is that revenue shares of input expenditure are the primary drivers of my results, not estimates of revenue elasticities.

5 Concluding remarks

This paper studies imperfect competition in labor markets and buyer power in intermediate input markets in Dutch manufacturing from 2007 to 2018. Most firms are characterized by wage markups, in line with an efficient bargaining model where employees have bargaining power but not in line with firms having monopsony power. In addition, I provide evidence of buyer power for intermediate inputs. I show that idiosyncratic factors, primarily wages, drive variation of wage markups. Firms with high intermediate input price markdowns pay their employees more even though these employees do not generate more revenue, suggesting that firms share rents generated in the intermediate input market with their employees. I find that bargaining power is (on average) evenly distributed between firms and employees and that the average firm increases wages by 0.22 percent following a 1-percent increase in quasi-rents per employee.

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, my approach does not require the existence of a variable input that is frictionlessly adjustable and for which firms are price takers. The ability of this paper to estimate firm-time-specific rent sharing elasticities is due to the efficient bargaining framework and revenue function estimation. While the control function approach to revenue function estimation makes several strong assumptions on unobserved revenue productivity, I show that these assumptions do not drive my results.

The main implication of this paper is that concerns of widespread monopsony power are unlikely to be warranted in European labor markets characterized by collective bargaining agreements. The second implication is that researchers should be careful when identifying labor market imperfections by restricting imperfections in other input markets. The third implication is that there exists substantial variation in rent-sharing elasticities, which can not be entirely pinned down by regressing wages on a measure of firm rents. This paper provides researchers with possible tools to consider these implications in their own work.

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 or not. I leave these considerations for further work.

References

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

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

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

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

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

Bhuller, M., Moene, K. O., Mogstad, M., and Vestad, O. L. (2022). Facts and fantasies about wage setting and collective bargaining (no. w30437) National Bureau of Economic Research.

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.

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.

Caloia, F., Parlevliet, J., and Mastrogiacomo, M. (2021). Staggered wages, unanticipated shocks and firms' adjustments. DNB Working Paper no. 711.

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.

Cardoso, A. R. and Portugal, P. (2005). Contractual wages and the wage cushion under different bargaining settings. *Journal of Labor Economics*, 23(4):875–902.

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

CBS (2019a). De arbeidsmarkt in cijfers 2018. Statistical Offices Netherlands.

CBS (2019b). Dutch trade in facts and figures 2019. Statistical Offices Netherlands.

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.

Crépon, B., Desplatz, R., and Mairesse, J. (2005). Price-cost margins and rent sharing: Evidence from a panel of French manufacturing firms. *Annales d'Economie et de Statistique*, 79/80:583–610.

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. (2022). The hitchhiker's guide to markup estimation. *mimeo*.

Dobbelaere, S., Hirsch, B., Müller, S., and Neuschaeffer, G. (2020). Organised labour, labour market imperfections, and employer wage premia. Tinbergen Institute Discussion Paper 2020-081/V.

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.

Dobbelaere, S. and Mairesse, J. (2018). Comparing micro-evidence on rent sharing from two different econometric models. *Labour Economics*, 52:18–26.

Dobbelaere, S. and Wiersma, Q. (2020). The impact of trade liberalization on firms' product and labor market power. IZA Discussion Paper No. 12951.

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

Eurostat (2008). NACE Rev. 2. Statistical classification of economic activities in the European community.

Eurostat (2020). Statistical regions in the European Union and partner countries.

Fabiani, S., Kwapil, C., Rõõm, T., Galuscak, K., and Lamo, A. (2010). Wage rigidities and labor market adjustment in Europe. *Journal of the European Economic Association*, 8(2-3):497–505.

Falch, T. and Strøm, B. (2007). Wage bargaining and monopsony. *Economics Letters*, 94(2):202–207.

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.

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(1).

Horvát, P. and Webb, C. (2020). The OECD STAN database for industrial analysis: Sources and methods. OECD Science, Technology and Industry Working Paper 2020/10.

Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and manufacturing TFP in China and India. *Quarterly Journal of Economics*, 124(4):1403–1448.

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

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

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.

Lu, Y., Sugita, Y., and Zhu, L. (2019). Wage markdowns and FDI liberalization. HIAS-E-83.

Manning, A. (2021). Monopsony in labor markets: A review. *ILR Review*, 74(1):3–26.

Marinescu, I. and Posner, E. A. (2019). Why has antitrust law failed workers? *Cornell L. Rev.*, 105:1343–1394.

Marschak, J. and Andrews, W. H. (1944). Random simultaneous equations and the theory of production. *Econometrica*, 12(3,4):143–205.

McDonald, I. M. and Solow, R. M. (1981). Wage bargaining and employment. *American Economic Review*, 71(5):896–908.

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

Mertens, M. (2021). Labour market power and wage (in)equality. *mimeo*.

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

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

Naidu, S., Posner, E. A., and Weyl, G. (2018). Antitrust remedies for labor market power. *Harvard Law Review*, 132:536–601.

Neven, D. J. and Röller, L.-H. (1996). Rent sharing in the European airline industry. *European Economic Review*, 40:933–940.

Oberfield, E. and Raval, D. (2021). Micro data and macro technology. *Econometrica*, 89(2):703–732.

OECD (2019a). Collective bargaining coverage. Obtained from: `https://stats.oecd.org/Index.aspx?DataSetCode=EPL_T#`.

OECD (2019b). Real minimum wages. Obtained from: `https://stats.oecd.org/Index.aspx?DataSetCode=EPL_T#`.

OECD (2019c). ALFS summary tables : Employment. Obtained from: `https://stats.oecd.org/Index.aspx?DataSetCode=EPL_T#`.

OECD (2021). Insights on productivity and business dynamics - March. Obtained from: `https://www.oecd.org/sti/ind/oecd-business-dynamics-insights-netherlands.pdf`.

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.

Raval, D. R. (2019). The micro elasticity of substitution and non-neutral technology. *RAND Journal of Economics*, 50(1):147–167.

Rubens, M. (2021). Market structure, oligopsony power, and productivity. *mimeo.*

Van Reenen, J. (1996). The creation and capture of rents: Wages and innovation in a panel of UK companies. *Quarterly Journal of Economics*, 111(1):195–226.

Visser, J. (2016). What happened to collective bargaining during the great recession? *IZA Journal of Labor Policy*, 5(1):1–35.

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.[49]

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.[50] 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".[51] 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-

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

[50]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.

[51]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.[52] 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. Eurostat (2008) provides a complete description of all NACE classifications and the conversion to other international industry classification codes. 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 NACE divisions (2-digit NACE 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 NACE division that are included.

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 (Horvát and Webb, 2020). 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 (Eurostat, 2020).

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.

[52] 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.

- Intermediate inputs (M_{it}): Deflated total expenditure on intermediate goods, energy, and other intermediate expenses, in euros.
- Earnings before interest and taxes ($EBIT_{it}$): Revenue minus cost of goods sold and all operating expenses, plus non-operating income, 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 Winsorization

Within each 2-digit-industry-by-year cell, I winsorize the top and bottom percent of the distributions of revenue, capital, total labor expenditure, and total intermediate input expenditure. After these steps, all 2-digit NACE 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 121,057 observations in the final sample by year and 2-digit NACE industry.

Table A1: NACE divisions covered in the final sample

Division	Description
10	Manufacture of food products
13	Manufacture of textiles
16	Manufacture of wood and of products of wood and cork, except furniture; manufacture of articles of straw and plaiting materials
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 NACE divisions (2-digit NACE industries) covered by the final sample. Manufacturing covers NACE divisions 10 to 33.

Table A2: Observation count, by year and NACE division

	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	total
10	788	809	837	854	889	911	904	922	948	964	893	977	10,696
13	204	215	206	212	208	208	207	195	204	200	188	207	2,454
16	393	410	387	390	412	402	374	380	411	431	391	434	4,815
17	139	128	129	141	145	144	145	148	150	147	139	141	1,696
18	990	946	912	935	892	849	821	776	758	725	651	665	9,920
20	199	212	219	242	262	271	276	279	281	297	263	294	3,095
22	458	487	474	491	501	490	511	521	523	528	486	535	6,005
23	278	298	294	284	296	294	296	289	297	296	277	305	3,504
24	83	95	102	89	106	98	105	100	104	104	95	106	1,187
25	2,076	2,183	2,234	2,286	2,412	2,356	2,373	2,397	2,396	2,491	2,326	2,522	28,052
26	276	305	318	333	377	360	384	352	371	381	338	382	4,177
27	256	271	285	302	312	316	323	323	315	316	289	308	3,616
28	984	1,030	1,035	1,085	1,126	1,161	1,180	1,177	1,193	1,218	1,103	1,204	13,496
29	200	204	208	198	209	204	198	198	200	211	189	203	2,422
30	164	179	176	192	204	198	194	195	198	205	168	194	2,267
31	538	570	569	602	624	609	593	613	611	644	601	668	7,242
32	491	516	519	534	572	582	591	607	613	615	538	614	6,792
33	559	644	681	695	775	798	830	890	931	963	844	1,011	9,621
	9,076	9,502	9,585	9,865	10,322	10,251	10,305	10,362	10,504	10,736	9,779	10,770	121,057

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

B Revenue function estimation

This appendix provides additional information on the estimation of the revenue elasticities and lists means, medians, and standard deviations of revenue elasticities and revenue returns to scale by 2-digit NACE industry in Table B1.

Recall that labor L_{it} is a headcount, while intermediate inputs M_{it} are measured by deflating intermediate input expenditure using a 2-digit-year specific intermediates price deflator, $\bar{P}^M_{dt}$. Similarly, capital K_{it} is measured as the book value of capital deflated using a 2-digit-year specific gross fixed capital price deflator. While it is common to assume industry-specific rental rates of capital, assuming industry-specific prices for intermediate inputs is potentially more problematic. Indeed, the Dutch manufacturing context and the estimation results point to buyer power in the market for intermediate inputs, potentially creating within-industry input price dispersion. Within-industry intermediate input price deviations from $\bar{P}^M_{dt}$ will enter the structural error term of the revenue function, scaled by the revenue elasticities.[53]

Collecting all input price bias terms in $b_{it}(\cdot)$, equation (16) can be written as

$$r_{it} = f_d(k_{it}, l_{it}, m_{it}; \theta) + b_{it}(\cdot) + \omega_{it} + \epsilon_{it}. \tag{B1}$$

I control for the input price bias $b_{it}(\cdot)$ using equation (20), repeated here for clarity

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

where $\mathbf{x}^b_{it}$ contains firm i's intermediate inputs as a share of total intermediate inputs in the firm's 4-digit NACE 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 4-digit NACE industry. 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 (see De Loecker et al. (2016)). This closely follows the implementation of De Loecker et al. (2016) (see also De Loecker et al. (2020)), who provide the theoretical foundations showing that market shares, industry indicators, and location indicators can be used to control for unobserved input price variation around the input price deflator. I add the intermediate input share to soak up input price variation due to buyer power in the market for intermediate inputs.[54] Appendix C discusses conditions under which this is a sufficient control for buyer power induced input price variation. Note that all input price variation related to revenue productivity ω_{it} is already controlled for by the inclusion of the control

[53]For instance, it is easy to verify that for a Cobb-Douglas revenue function the term $\theta^M_{it} \log(\frac{P^M_{it}}{\bar{P}^M_{dt}})$ will enter the error term.

[54]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 market shares instead.

function in equation (B3), and that including the control function (B2) is recommended even if it does not perfectly control for the input price bias.

Inverting the demand for intermediate inputs given in equation (17) provides a control for revenue productivity

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

where $\mathbf{x}_{it}^m$ 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 NACE industry.

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

$$r_{it} = f_d(k_{it}, l_{it}, m_{it}; \theta) + b_d((1, k_{it}, l_{it}, m_{it}) \times \mathbf{x}_{it}^b; \delta^b) + m_d^{-1}(k_{it}, l_{it}, m_{it}, \mathbf{x}_{it}^m; \delta^m) + \epsilon_{it}, \tag{B4}$$

where, as discussed, d denotes firm i's NACE division (2-digit NACE industry). The revenue function $f_d(\cdot)$ is approximated by a translog function, so that the revenue elasticity of input $x \in \{k, l, m\}$ is given by $\theta_{it}^x = \frac{\partial f_d}{\partial x_{it}}$. The control functions for revenue productivity and the input price bias are treated non-parametrically as their functional forms depend on the exact specification of demand and intermediates supply, which are left unspecified.

The first step of the two-step approach is based on an approximation to equation (B4). It consists of regressing logged revenue on a third-degree polynomial in all right-hand side variables except for indicator variables and time trends, which are added linearly.[55] This first step is not meant to identify any of the coefficients in θ, δ^b, or δ^m, but rather to separate observed revenue into planned revenue and a revenue shock: $r_{it} = \hat{r}_{it} + \hat{\epsilon_{it}}$. 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 θ and δ^b from equation (B4)

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

From the law of motion of revenue productivity in equation (21), we can then obtain an estimate of the mean zero revenue productivity shock that depends on unknown coefficients θ, δ^b, and δ^g

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

Estimates of θ, δ^b, and δ^g are obtained by forming moments on ξ_{it}. I follow the literature and assume $\mathbb{E}(\xi_{it}, k_{it}) = 0$, as capital adjusts slowly due to adjustment frictions. Labor is assumed

[55] Higher-order polynomial approximations are not computationally feasible. Higher-order polynomials approximations are feasible for the parsimonious specification discussed below but do not alter any of the main results presented in this paper.

quasi-fixed, so that $\mathbb{E}(\xi_{it}, l_{it}) = 0$. While this is commonly done in the literature (e.g., Morlacco (2020); Mertens (2022)) and can be rationalized by adding timing assumptions to the model introduced in Section 2.2, I show in a robustness check that allowing for labor to respond to ξ_{it} does not affect the key results of this paper. I allow for the possibility that m_{it} can be adjusted in response to the realization of ξ_{it}, so that $\mathbb{E}(\xi_{it}, m_{it}) \neq 0$, Together, the following moments result

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

where $\mathbf{Z}_{it}$ includes a second-order polynomial in k_{it} and l_{it} and and interactions of contemporaneous capital and labor with lagged inputs. In addition, $\mathbf{Z}_{it}$ includes all terms in $g_d(\cdot)$, which is approximated by a third order polynomial, and lags of all terms in $b_d(\cdot)$, which is approximated by a second order polynomial. In both cases, the polynomial is in all the function's arguments except for indicators and time trends, which are added linearly. That is, $\mathbf{Z}_{it}$ contains k_{it}, k_{it}^2, $k_{it}k_{it-1}$, $k_{it}l_{it-1}$, $k_{it}m_{it-1}$, l_{it}, l_{it}^2, $l_{it}k_{it-1}$, $l_{it}l_{it-1}$, $l_{it}m_{it-1}$, $k_{it}l_{it}$, k_{it-1}, k_{it-1}^2, l_{it-1}, l_{it-1}^2, m_{it-1}, m_{it-1}^2, $k_{it-1}l_{it-1}$, $k_{it-1}m_{it-1}$, $l_{it-1}m_{it-1}$, mms_{it-1}, mms_{it-1}^2, $mms_{it-1}k_{it-1}$, $mms_{it-1}l_{it-1}$, $mms_{it-1}m_{it-1}$, t, t^2, indicators for all but one NUTS1-region, and indicators for all but one 4-digit NACE industry in NACE division d.

The remainder of this appendix provides additional information on the robustness checks discussed in Section 4.4. In the parsimonious specification, $b_d(\cdot;\delta^b)$ and $\mathbf{x}_{it}^m$ are removed from equation (B4). Other than that, the procedure to obtain the revenue productivity shocks is as discussed above. Identification of the coefficients is based on instruments $\mathbf{Z}_{it}^p$, which contains k_{it}, k_{it}^2, $k_{it}k_{it-1}$, $k_{it}l_{it-1}$, $k_{it}m_{it-1}$, l_{it}, l_{it}^2, $l_{it}k_{it-1}$, $l_{it}l_{it-1}$, $l_{it}m_{it-1}$, $k_{it}l_{it}$, k_{it-1}, k_{it-1}^2, l_{it-1}, l_{it-1}^2, m_{it-1}, m_{it-1}^2, $k_{it-1}l_{it-1}$, $k_{it-1}m_{it-1}$, and $l_{it-1}m_{it-1}$.

For the year-by-2-digit level estimates, the parsimonious specification is used. However, the revenue function and revenue productivity control are now indexed by dt so that each 2-digit industry has a different revenue function each year. Likewise, for the 4-digit level estimates, the parsimonious specification is used, but the revenue function and revenue productivity control are indexed by n, where n indicates a 4-digit NACE industry. Finally, when altering the timing assumptions the parsimonious specification is used, but coefficients are identified based on instruments $\mathbf{Z}_{it}^t$, which contains k_{it}, k_{it}^2, $k_{it}k_{it-1}$, $k_{it}l_{it-1}$, $k_{it}m_{it-1}$, k_{it-1}, k_{it-1}^2, l_{it-1}, l_{it-1}^2, m_{it-1}, m_{it-1}^2, $k_{it-1}l_{it-1}$, $k_{it-1}m_{it-1}$, and $l_{it-1}m_{it-1}$.

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	0.20 *0.19* (0.08)	0.74 *0.75* (0.09)	0.06 *0.06* (0.03)	1.00 *1.00* (0.01)	10,696
13	0.25 *0.26* (0.11)	0.70 *0.69* (0.13)	0.04 *0.04* (0.02)	0.99 *0.99* (0.01)	2,454
16	0.22 *0.21* (0.07)	0.74 *0.75* (0.06)	0.03 *0.03* (0.01)	0.99 *0.99* (0.02)	4,815
17	0.21 *0.21* (0.06)	0.76 *0.76* (0.05)	0.04 *0.04* (0.02)	1.00 *1.00* (0.02)	1,696
18	0.32 *0.32* (0.11)	0.64 *0.63* (0.11)	0.05 *0.05* (0.02)	1.00 *1.00* (0.03)	9,920
20	0.20 *0.19* (0.06)	0.76 *0.76* (0.05)	0.05 *0.05* (0.02)	1.00 *1.00* (0.02)	3,095
22	0.23 *0.23* (0.07)	0.72 *0.72* (0.08)	0.05 *0.04* (0.02)	0.99 *0.99* (0.01)	6,005
23	0.25 *0.24* (0.09)	0.69 *0.70* (0.07)	0.04 *0.04* (0.02)	0.99 *0.99* (0.02)	3,504
24	0.22 *0.22* (0.07)	0.73 *0.73* (0.07)	0.05 *0.05* (0.03)	1.00 *1.00* (0.01)	1,187
25	0.31 *0.30* (0.10)	0.64 *0.64* (0.11)	0.05 *0.05* (0.02)	0.99 *0.99* (0.02)	28,052
26	0.29 *0.28* (0.11)	0.68 *0.68* (0.13)	0.03 *0.03* (0.01)	0.99 *0.99* (0.04)	4,177
27	0.25 *0.25* (0.09)	0.72 *0.72* (0.13)	0.04 *0.04* (0.01)	1.00 *1.01* (0.06)	3,616
28	0.27 *0.27* (0.10)	0.71 *0.70* (0.13)	0.03 *0.03* (0.02)	1.00 *1.00* (0.04)	13,496
29	0.25 *0.24* (0.09)	0.73 *0.73* (0.08)	0.03 *0.03* (0.01)	1.00 *1.00* (0.02)	2,422
30	0.28 *0.28* (0.12)	0.66 *0.66* (0.19)	0.04 *0.04* (0.04)	0.98 *0.98* (0.08)	2,267
31	0.27 *0.26* (0.09)	0.70 *0.71* (0.11)	0.03 *0.03* (0.02)	1.00 *1.00* (0.02)	7,242
32	0.31 *0.30* (0.09)	0.58 *0.59* (0.11)	0.04 *0.04* (0.01)	0.93 *0.94* (0.03)	6,792
33	0.30 *0.30* (0.13)	0.66 *0.65* (0.15)	0.03 *0.03* (0.03)	0.98 *0.98* (0.04)	9,621
full sample	0.27 *0.27* (0.11)	0.69 *0.68* (0.12)	0.04 *0.04* (0.02)	0.99 *0.99* (0.03)	121,057

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 Models of buyer power

This appendix gives examples of models in which the intermediate input wedge captures the extent to which buyer power is exercised. In addition, I discuss for each model to what extent the input price bias control function given in equation (20) controls for unobserved buyer power. This appendix is by no means meant as an exhaustive list, as intermediate input wedges capture buyer power in a much wider class of models.[56]

In a textbook oligopsonistic intermediate input market, the intermediates wedge straightforwardly captures a firm's buyer power. Take $\pi(M_i) = R(M_i) - P^M(M)M_i$, where $M = \sum_i M_i$. The first-order condition can be rewritten as $\frac{MRPM_i}{P^M(M)} = 1 + \frac{M_i}{M}(\varepsilon^M)^{-1}$. As such models result in a single market price for intermediate inputs, there is no input price bias. Recall that intermediate input expenditure is deflated with a 2-digit-year specific intermediates deflator. In addition, unobserved revenue productivity and input use are held constant, so an input price bias requires price variation around the deflator conditional on these factors. I provide two such examples.

Berger et al. (2022) develop a tractable general equilibrium model of oligopsony in labor markets based on Atkeson and Burstein (2008). Their approach could also be applied to model buyer power in concentrated intermediate input markets where input suppliers see firms' orders as imperfect substitutes. While, in the baseline model, a representative household is assumed to make all input supply decisions, Berger et al. (2022) show that the supply side of their model can also be founded on heterogeneous input suppliers making independent decisions. Output markets are assumed to be perfectly competitive and productivity differs across firms.

Equation (7) of Berger et al. (2022) shows that an application of their model to buyer power in the market for intermediates would result in an intermediate input wedge in a particular input market equal to

$$\gamma_{it}^M = \frac{\varepsilon_{it}^M + 1}{\varepsilon_{it}^M}, \tag{C1}$$

which corresponds exactly to equation (10) of the current paper. Moreover, the intermediate input supply elasticity within each intermediate input market would then be given by

$$\varepsilon_{it}^M = (\frac{1}{\eta^{in}} + (\frac{1}{\eta^{out}} - \frac{1}{\eta^{in}})\frac{P_{it}^M M_{it}}{\sum_j^n P_{jt}^M M_{jt}})^{-1}, \tag{C2}$$

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. Equation (C2) shows that a firm's share of intermediate input expenditure in a particular intermediate input market, in the current paper, a 2-digit NACE industry, accounts for

[56] For instance, Rubens (2021) models input supply using a discrete choice model with differentiated firms.

all within-market variation of the input wedge. Therefore, in this setting, including the intermediates expenditure share in the price control function soaks up all variation in buyer power within a 2-digit NACE industry. See Berger et al. (2022) for details.

Morlacco (2020) studies input wedges of imported and domestic intermediate inputs in French manufacturing. Firms are assumed to buy from one input supplier, who can supply several firms simultaneously. Firms bargain with their input supplier to determine the price of intermediate inputs, and the solution concept is Nash-in-Nash.

Under the assumptions that there is a single supplier of inputs in a 2-digit NACE industry and a single variety of intermediate inputs, and as marginal revenue equals marginal cost, equations (10) and (55) in Morlacco (2020) can be written as

$$\gamma_{it}^{M} \simeq 1 + \left(\frac{1-\rho}{\rho}\right)\frac{M_{it}}{\sum_j M_{jt}}\left(1 - \frac{\eta_{it}}{2}\left(1 - \frac{M_{it}}{\sum_j M_{jt}}\right)\right), \tag{C3}$$

where ρ is the input supplier's returns to scale, j indexes all firms in firm i's 2-digit industry, and $\eta_{it} \in (0, 1)$ is a constant that depends on market conditions upstream and downstream that are exogenous to the firm. See Morlacco (2020) for details.

Equation (C3) clarifies that, conditional on η_{it}, including firm i intermediate input market share controls for all within-industry variation in the intermediate input wedge. As long as the factors in η_{it} are constant within a 2-digit NACE industry, or controlled for by factors included in the estimating equation (B4), including the intermediate input market share in the input price control function corrects for all remaining within-industry buyer power variation. Recall that including an input price control function is always preferred, even if some within-industry buyer power variation remains conditional on intermediate input market shares.

D Additional figures

Figure D1: Regressions of firm-level labor wedges on employment (panel a) and capital (panel b), over time

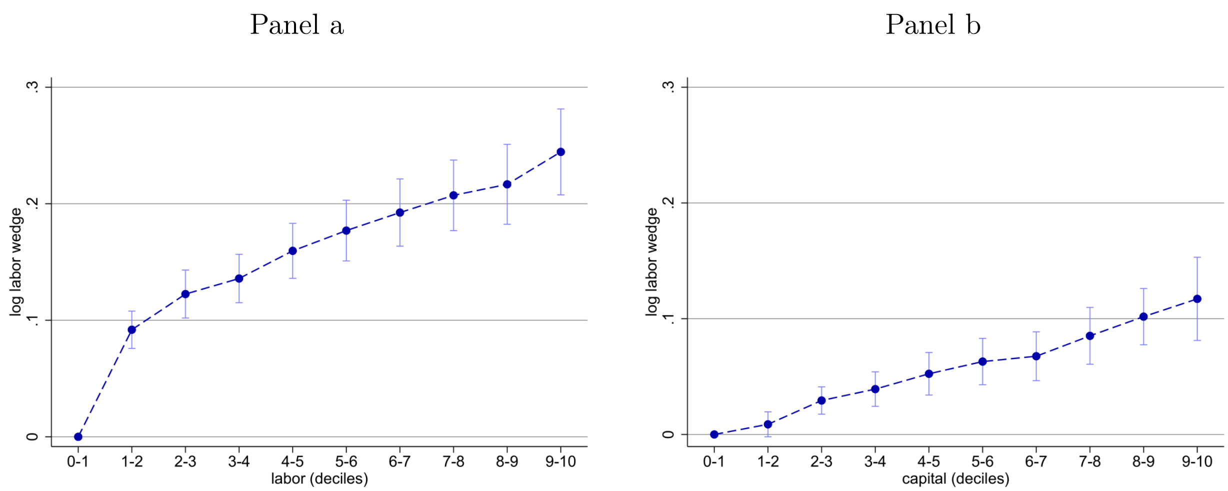

Notes: Regression coefficients and 95 percent confidence intervals from a non-parametric regression of log labor wedge on employment (panel a), and capital (panel b). Explanatory variables are divided into 10 categories based on the deciles of their distribution. The coefficient on the indicator of the first decile is normalized to 0. Year-by-4-digit NACE industry fixed effects included. Standard errors are clustered at the 4-digit industry level. Based on the full samples of 121,057 observations covering the years 2007 to 2018.

Figure D2: Regressions of firm-level labor wedges on revenue (panel a) and EBIT (panel b), over time

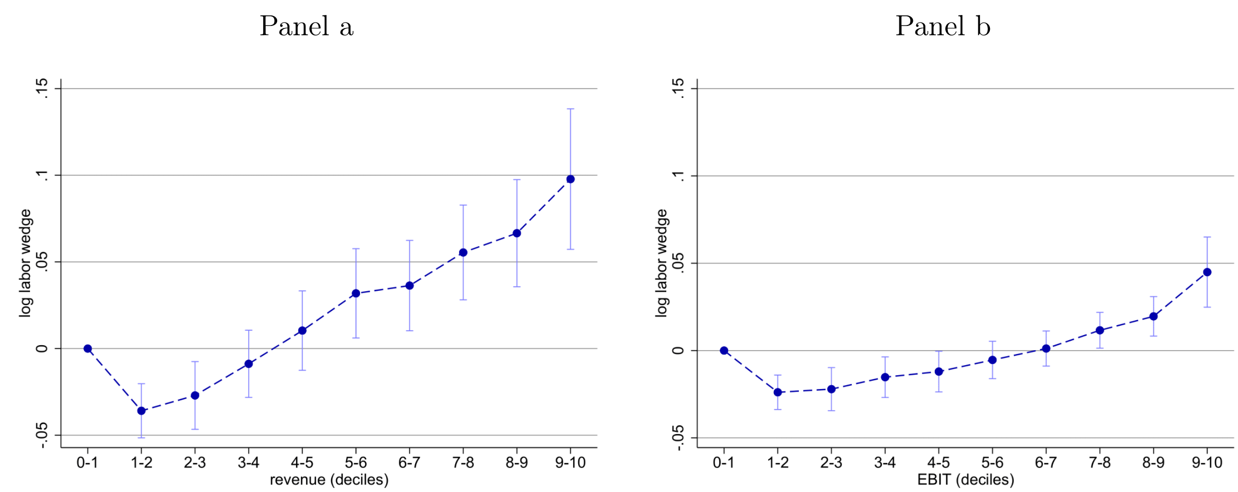

Notes: Regression coefficients and 95 percent confidence intervals from a non-parametric regression of log labor wedge on revenue (panel a), and earnings before interest and taxes (panel b). Explanatory variables are divided into 10 categories based on the deciles of their distribution. The coefficient on the indicator of the first decile is normalized to 0. Year-by-4-digit NACE industry fixed effects included. Standard errors are clustered at the 4-digit industry level. Based on the full samples of 121,057 observations covering the years 2007 to 2018.

E Additional tables

Table E1: Summary statistics of labor wedges, by robustness check

	p(5)	p(25)	p(50)	p(75)	p(95)	observations
Baseline estimates	0.53	0.73	0.86	0.99	1.26	121,057
Median elasticities imposed	0.48	0.65	0.85	1.20	2.28	121,057
No input price control	0.43	0.65	0.82	1.01	1.43	121,453
Year-by-2 digit estimates	0.40	0.65	0.83	1.04	1.47	110,465
4-digit estimates	0.37	0.61	0.78	0.99	1.45	119,430
Alternative timing assumption	0.42	0.65	0.81	1.00	1.39	120,460

Notes: Summary statistics of the labor wedge. p(5), p(25), p(50), p(75), and p(95) refer to the 5th, 25th, 50th, 75th, and 95th percentile of the labor wedge distribution, respectively. Labor wedges are rounded to two decimal points. Observations = number of observations in the full sample of the robustness check under consideration.

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

	p(5)	p(25)	p(50)	p(75)	p(95)	observations
Baseline estimates	0.99	1.09	1.18	1.30	1.57	121,057
Median elasticities imposed	0.86	1.02	1.19	1.44	2.04	121,057
No input price control	1.05	1.12	1.18	1.27	1.50	121,453
Year-by-2 digit estimates	1.04	1.12	1.18	1.27	1.49	110,465
4-digit estimates	1.02	1.12	1.20	1.31	1.54	119,430
Alternative timing assumption	1.05	1.13	1.19	1.29	1.51	120,460

Notes: Summary statistics of 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 intermediate input wedge distribution, respectively. Intermediate input wedges are rounded to two decimal points. Observations = number of observations in the full sample of the robustness check under consideration.

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

	$\hat{\beta}$	s.e.($\hat{\beta}$)	observations
Baseline estimates	-1.22	0.026***	121,057
Median elasticities imposed	-1.44	0.039***	121,057
No input price control	-0.42	0.059***	121,453
Year-by-2-digit estimates	-0.51	0.044***	110,465
4-digit estimates	-0.57	0.063***	119,430
Alternative timing assumption	-0.43	0.054***	120,460

Notes: Regression coefficient on log intermediate input wedge and concomitant standard error from a regression of log labor wedge on log intermediate input wedge. Year-by-4-digit NACE industry fixed effects included. Standard errors are clustered at the 4-digit industry level. Observations = number of observations in the full sample of the robustness check under consideration. *** indicates statistical significance at the 1% level.

Table E4: Summary statistics of rent sharing elasticities, by robustness check

	p(5)	p(25)	p(50)	p(75)	p(95)	observations	full sample
Baseline estimates	0.03	0.11	0.20	0.31	0.51	92,568	121,057
Median elasticities imposed	0.04	0.17	0.30	0.41	0.55	76,522	121,057
No input price control	0.03	0.14	0.27	0.40	0.60	89,389	121,453
Year-by-2-digit estimates	0.03	0.14	0.27	0.41	0.64	78,630	110,465
4-digit estimates	0.04	0.17	0.30	0.44	0.66	90,346	119,430
Alternative timing assumption	0.03	0.15	0.27	0.40	0.61	90,791	120,460

Notes: Summary statistics of the elasticity of wage with respect to quasi-rents per employee. p(5), p(25), p(50), p(75), and p(95) refer to the 5th, 25th, 50th, 75th, and 95th percentile of the rent sharing elasticity distribution, respectively. Percentiles are rounded to two decimal points. observations = number of observations for which $\gamma_{it}^{L} < 1$; full sample = number of observations in the full sample of the robustness check under consideration.