Monopsony and Wage Inequality in Sweden*

The Effect of Markdowns on Immigrant and Gender Wage Gaps

Theodor Friederiszick† Axel Malmcrona‡ Nelson Sunuwar§

Click here for the latest version

November 3, 2026

Abstract

How do differences in wage markdowns contribute to gender and origin earnings gaps? We combine Swedish firm accounts with administrative worker records and estimate production functions with separate labor inputs by gender and origin. We measure markdowns as the marginal revenue product of labor (MRPL) relative to earnings and examine their differences across industries, employers and worker groups over 1998–2023. In 2023, women face a markdown 9% higher than men at the employment-weighted median within-firm ratio. Their allocation across industries and employers partly offsets this within-employer disadvantage. Holding estimated MRPL and employment allocation fixed, assigning women the male markdown within their employer closes nearly half of the gender earnings gap. Turning to origin, foreign-born workers face a markdown 20% higher than native-born workers at the employment-weighted median within-firm ratio in 2023. Their relative markdowns decline substantially over the period, primarily reflecting falling relative estimated MRPL while relative earnings change little. Assigning foreign-born workers the native-born markdown within their employer, under the same fixed-MRPL and allocation benchmark, closes most of the origin earnings gap in 2023. These findings show why employment allocation and changes in productive contributions matter for interpreting markdown differences, and why a narrowing markdown gap need not imply a narrowing earnings gap.

*Theodor Friederiszick and Nelson Sunuwar gratefully acknowledge funding from the University Research Priority Program (URPP) Equality of Opportunity.

†University of Zurich. Email: theodor.friederiszick@econ.uzh.ch

‡University of Gothenburg. Email: axel.malmcrona@economics.gu.se

§University of Zurich. Email: nelson.sunuwar@econ.uzh.ch

1 Introduction

Earnings differences between groups of workers require careful interpretation. They may reflect differences in workers' characteristics, the tasks they perform, their contributions to production, or discrimination (Acemoglu and Autor, 2011). A growing literature emphasizes firms' labor market power and its implications for pay (Manning, 2021; Yeh, Macaluso, and Hershbein, 2022; Lamadon, Mogstad, and Setzler, 2022). Wage markdowns, defined as the ratio of the marginal revenue product of labor (MRPL) to wages, describe how much workers receive relative to their contribution to production. Recent studies document higher markdowns for immigrants than native-born workers in Canada (Tino, 2026) and for women than men in Brazil (Sharma, 2023). Related production-based evidence from Denmark shows that productivity differences do not account for the pay disadvantage of women without children (Gallen, 2024). These findings raise the question of how much differences in markdowns contribute to earnings gaps between worker groups.

This paper examines how differences in markdowns contribute to gender and origin earnings gaps in Sweden. We combine firm accounts with administrative records covering every worker employed by each firm over 1998–2023. Our data cover a broad range of industries, and the worker records complement the firm accounts with information on each employer's workforce, individual earnings, and characteristics such as age, education and occupation. This allows us to estimate separate markdowns for worker groups and examine whether the differences persist after accounting for workforce characteristics. We study workers of different gender and origin within a common production framework and estimation sample to examine how differences within employers and allocation across employers and industries jointly shape markdown gaps. This distinction matters because where workers work can reinforce or offset the markdown differences they face at the same employer. Following these relationships over more than two decades also allows us to distinguish narrowing markdown gaps driven by rising relative earnings from those driven by falling relative MRPL. We compare these group-specific estimates with conventional employer-level markdowns to assess which differences and changes are concealed when labor is treated as a single input.

We estimate production functions following Hall, Blanchard, and Hubbard (1986), Levinsohn and Petrin (2003) and Ackerberg, Caves, and Frazer (2015). In our specifications, we treat worker groups as separate labor inputs. In practice, we use the linked worker records to divide each firm's personnel costs between groups according to their payroll shares. Following the production-based approach to measuring labor markdowns in Dobbelaere and Mairesse (2013) and Yeh, Macaluso, and Hershbein (2022), we combine estimated output elasticities with expenditure shares to recover group-specific markdowns, measuring each

group's estimated MRPL relative to its earnings. Our model allows relative MRPL and markdowns across worker groups to vary with each employer's combination of labor, capital and materials. We estimate separate production functions for 61 industry groups, allowing the relationship between inputs and output to differ across industries. Within each industry, we account for unobserved firm productivity and allow it to evolve over time.

Identification draws on variation in output and the inputs firms use, across firms within the same industry, and over time. Variation in each worker group's labor expenditure, holding other inputs fixed, helps us estimate how output responds to that group's labor input. A challenge for identification is that more productive firms may also choose to use more inputs, making it difficult to distinguish the contribution of those inputs from differences in underlying productivity. We account for this by using firms' materials demand as a proxy for unobserved productivity, together with assumptions about how productivity evolves and when firms choose their inputs (Levinsohn and Petrin, 2003; Ackerberg, Caves, and Frazer, 2015). Under these assumptions, capital chosen in advance and earlier labor and materials choices help isolate variation in current inputs that is unrelated to unexpected changes in productivity. Changes in materials prices over time then provide additional variation by changing firms' input choices under the assumption that input prices do not directly change firms' production technology (Gandhi, Navarro, and Rivers, 2020; Yeh, Macaluso, and Hershbein, 2022). Under these assumptions we estimate worker-group output elasticities, and combine them with observed expenditure shares to recover markdowns.

We find that markdowns differ substantially across industries, employers and worker groups. We begin by estimating a conventional production function with a single labor input, combining expenditure on all workers. In 2023, the mean markdown across firms is 1.10 while the median is 0.97. The median implies that the typical firm pays its workers close to their estimated MRPL, while the higher mean reflects a distribution skewed to the right. There is nevertheless substantial dispersion: the middle half of firms have markdowns between 0.70 and 1.31. Most of the variation in markdowns occurs between firms in the same industry, accounting for 62% of the total; differences between industries account for the remaining 38%. We use these estimates to examine whether women and foreign-born workers are more likely than their comparison groups to work in industries or firms with higher markdowns. However, this specification assigns the same markdown to all workers at a firm, concealing any differences between groups working for the same employer.

Distinguishing female and male labor inputs allows us to explore *within-firm markdown differences*. In 2023, at the employment-weighted median, women face a markdown approximately 9% higher than men at the same firm. For workers with the same MRPL, this translates into approximately 8% lower earnings for women. The difference in within-firm

markdowns remains after accounting for workforce composition: adding age, education and occupation controls reduces the within-firm gap by around one third. At the same time, women are more likely to work in both industries and firms with lower female markdowns, offsetting some of the within-firm difference. Women also have lower estimated MRPL: in 2023, their mean estimated MRPL is 80% of men's, while their mean earnings are 75% of men's. Finally, we quantify what within-firm markdown differences imply for the gender earnings gap. Assigning women the same markdown as men within firms, while leaving estimated MRPL and the allocation of workers across industries and firms fixed, raises female relative earnings from 75% of men's to 87%, closing almost half of the gender earnings gap in our sample. The corresponding exercise closes approximately half of the earnings gap in 1998 as well.

Turning to foreign-born and native-born workers, we find that their markdown gap narrows substantially over the period. The employment-weighted mean markdown for foreign-born workers is 54% higher than for native-born workers in 1998, falling to 18% higher in 2023. This narrowing comes primarily from differences within firms. Foreign-born workers' estimated MRPL falls relative to that of native-born workers, while their relative earnings change little. The decline in the within-firm markdown gap remains after accounting for age, education, occupation and time since recorded immigration. Over the same period, foreign-born workers account for a growing share of employment. Within the estimated production model, growth in foreign-born employment relative to native-born employment is associated with a decline in its relative marginal contribution to revenue. One interpretation is that the roles filled by foreign-born workers early in the sample contributed more to revenue at the margin, relative to those filled by native-born workers, than the roles filled as their employment expanded. Finally, assigning foreign-born workers the native-born markdown at their employer, while holding estimated MRPL and employment allocation fixed, raises their relative earnings from 87% to 131% of native-born earnings in 1998, compared with 88% to 98% in 2023. Thus, within-employer markdown differences more than account for the initial earnings shortfall and account for most of the remaining shortfall in 2023 under the baseline estimates.

Previous literature. Our paper is closely related to studies comparing earnings and productive contributions across worker groups. Gallen (2024) examines gender productivity and pay differences in Denmark, while Tino (2026) studies productivity, markdowns and allocation in the immigrant–native pay gap in Canada. Sharma (2023) finds that women's lower labor-supply elasticities and the concentration of attractive employment alternatives in fewer industries contribute to higher markdowns and the gender pay gap in Brazil. Related studies examine productivity differences by origin in Sweden (Ek, 2024) and migrant productivity

and pay in Britain (Hall and Manning, 2024). We build on this work by asking whether allocation across employers and industries reinforces or offsets within-employer markdown differences. We also examine what conventional employer-level markdowns conceal about these differences and their evolution. We study gender and origin within a common production framework and the same firm-year sample over 1998–2023, allowing comparisons across groups and over time.

Our approach builds on production-based comparisons of worker productivity and pay (Hellerstein, Neumark, and Troske, 1999; Hellerstein and Neumark, 2007) and the use of output elasticities and expenditure shares to recover markdowns (Dobbelaere and Mairesse, 2013; Yeh, Macaluso, and Hershbein, 2022). This framework already allows for heterogeneous labor, with studies estimating separate markdowns for production and nonproduction workers or for workers with different skill levels (Yeh, Macaluso, and Hershbein, 2022; Hong, Liu, and Pang, 2025; Caselli, Locatelli, and Schiavo, 2025; Wong, 2026). We apply it to gender and origin groups and emphasize the interpretation of within-employer comparisons. As Treuren (2025) shows, buyer power in materials markets complicates the interpretation of production-based markdown levels. A common materials-market wedge nevertheless cancels from the ratio of worker-group markdowns within an employer, conditional on identifying the output elasticities and maintaining the labor-input assumptions. This observation, related to the group-wedge relationships in Deb et al. (2024) and Mertens (2023), allows us to interpret relative markdowns without assuming competitive materials markets.

We also relate to studies examining group earnings differences through labor-supply responses and employer wage premiums. Hirsch, Schank, and Schnabel (2010) study gender differences in labor supply to employers, while Hirsch and Jahn (2015) examine corresponding differences between immigrants and natives. Webber (2016) emphasizes employer allocation in gender differences in labor-supply elasticities. A related literature examines how employer pay setting and allocation contribute to earnings gaps. Card, Cardoso, and Kline (2016) distinguish between women working at firms that pay lower premiums and women receiving smaller premiums than men at the same firm. For origin, Åslund et al. (2025) examine how firm productivity, allocation and pay setting relate to immigrant–native earnings differences in Sweden. Our estimates complement this evidence by comparing each group's earnings with its estimated MRPL at the same employer. Working at a low-paying firm need not imply facing a high markdown if estimated MRPL is also low. We can therefore distinguish where workers earn less from where they receive a smaller share of their estimated productive contribution.

More broadly, our paper contributes to the literature on employer wage setting and earnings inequality (Manning, 2021; Kline, 2025; Card et al., 2018; Lamadon, Mogstad, and

Setzler, 2022; Berger, Herkenhoff, and Mongey, 2022). We examine how the relationship between pay and estimated productive contributions differs across worker groups, connecting employer markdowns to gender and origin earnings gaps.

Our paper proceeds as follows. Section 2 describes the data and estimation sample. Section 3 presents the production-function framework and worker-type markdown measures. Section 4 examines the identifying assumptions and the interpretation of markdowns in the Swedish setting. Section 5 presents the empirical findings, and Section 6 concludes.

2 Data

This section introduces the data, worker classifications, and common estimation sample. Using production function estimation to compare worker types requires two kinds of data: the firm financial accounts typically used for production-function estimation and detailed information identifying worker characteristics of the employees within each production unit.

2.1 Swedish Firm Data with Detailed Employee Records

We construct an annual panel of Swedish non-financial private-sector firms over 1997–2023. The unit of observation is the firm-year. Firm accounts from Statistics Sweden's Business Statistics for Non-Financial Firms provide common production function inputs, such as production value, materials expenditure, capital and personnel costs. We combine these firm records with linked administrative employee data, which identify workers' gender and whether they were born in Sweden for each firm's workforce and the earnings paid to each worker type. A full list of datasets can be found in Appendix A.1.

A benefit of our setting is that our data cover firms beyond manufacturing, a traditional focus of production function studies (Yeh, Macaluso, and Hershbein, 2022; Mertens, 2022; Treuren, 2025), including a broad range of non-financial private-sector industries. This coverage allows us to examine worker-type markdowns across industries with different production technologies, workforce compositions, and wage-setting institutions.

We use the administrative worker records to identify each firm's workers, their characteristics and the earnings paid to them. Our main results use Primary-employer assignment: each worker is assigned exclusively to the employer paying their highest annual earnings. We classify assigned workers by gender and birthplace, and sum their positive earnings paid by the firm to construct the worker-type payroll shares described in Section 3.1. Assigned-worker counts measure neither hours nor full-time equivalents.

Our worker-type specifications divide the workforce by gender and by foreign-born or native-born status. These dimensions may capture differences in outside options, mobility, and bargaining positions relative to the firm. Gender and origin are separate specifications,

Table 1: Firm-Year Summary Statistics: Primary-Employer Estimation Sample

	Median	SD	P25	P75
Panel A: Firm Scale and Inputs Monetary values in million SEK, 2020 prices				
Output	12.48	314.96	5.36	34.30
Materials cost	3.14	143.32	1.15	10.31
Labor cost	4.71	51.87	2.03	11.91
Capital	1.29	817.12	0.30	6.01
Assigned workers	14.0	82.6	7.0	29.0
Panel B: Workforce Composition (%)				
Gender				
Women	33.3	24.6	16.7	57.1
Men	66.7	24.6	42.9	83.3
Origin				
Native-born	81.0	21.9	63.9	90.0
Foreign-born	19.0	21.9	10.0	36.1
Gender × origin				
Native-born women	25.0	22.0	11.1	44.4
Native-born men	45.5	26.5	23.1	66.7
Foreign-born women	4.5	13.7	0.0	14.3
Foreign-born men	10.9	17.6	3.7	24.2
Panel C: Wage-Bill Composition by Worker Type (%)				
Women	28.2	25.3	12.4	51.8
Men	71.8	25.3	48.2	87.6
Native-born	82.5	25.0	61.8	92.0
Foreign-born	17.5	25.0	8.0	38.2
Panel D: Annual Earnings per Assigned Worker Thousand SEK, 2020 output prices				
All workers	244.27	128.94	162.53	328.69
Women	202.38	131.06	121.30	293.59
Men	268.32	159.40	177.49	355.91
Native-born	248.07	142.59	154.42	336.42
Foreign-born	230.79	159.45	137.05	330.74
Panel E: Full-Period Sample Coverage				
Firm-year observations (total)		729,358		
Distinct firms		147,239		
Production-function families		61		
Years observed		26 (1998–2023)		
FEK employment (annual average)		667,848		
Assigned workers (annual average)		871,660		

Notes: Unweighted firm-year statistics for the common Primary-employer estimation sample, before markdown reporting screens. Monetary values use 2020 prices; materials use industry P2pp deflators, and labor costs and earnings use the output deflator. Amounts are expressed in SEK. Workers are assigned to their primary employer; counts do not measure hours or full-time equivalents. Workforce and wage-bill shares include unclassified records in the denominators. Joint gender–origin shares use the same workforce denominator. Earnings are annual payroll per assigned worker at the primary employer.

but to follow the differences in results across worker types we construct one common estimation sample.

2.2 Estimation Sample

To obtain our estimation sample, we require each firm-year to have valid output, capital, materials, labor costs, and the lagged observations required by our instruments. This is the regular requirement for production function estimation. To add our worker groups we also require each firm-year to have positive linked employment and positive labor inputs for both gender groups and both origin groups. We intersect the retained estimation samples across the aggregate, gender, and origin specifications, so that all three use the same firm-year sample. This allows us to hold the sample fixed and facilitates comparisons between markdown differences in the three models.

We estimate production functions separately by industry, allowing production technologies to differ across activities. The Primary-employer estimation sample covers 61 industry groups. Appendix Table A1 reports the industry classification and sample coverage, while Appendix Table A2 documents the sample restrictions.

3 Method

We use the production approach to recover worker-type-specific markdowns within firms. De Loecker and Warzynski (2012) recover firm-level *markups*, the ratio of price to marginal cost, from output elasticities and variable-input expenditure shares under cost minimization. The production approach has been extended to measuring labor market power, instead recovering *markdowns*, the ratio of the marginal revenue product of labor to its wage (Dobbelaere and Mairesse, 2013; Brooks et al., 2021; Yeh, Macaluso, and Hershbein, 2022).

We estimate markdowns separately for worker types by replacing the conventional single labor input with separate worker-type inputs, allocating the firm's labor costs across groups using their payroll shares. This allows us to estimate worker-type-specific output elasticities, which we combine with the corresponding expenditure shares to recover a markdown for each type within the same firm. Under our identifying assumptions, each worker-type elasticity is identified from variation in that worker type's input conditional on the firm's other inputs rather than, as in the conventional estimation, varying any worker that counts towards the total labor cost. We can therefore compare markdowns across worker types within firms and examine how these differences evolve over time.

Our approach relies on the production-function and input-choice assumptions described below. We highlight selected conditions and an interpretation of the recovered wedges that we examine further in Section 4. We first construct the worker-type labor inputs and describe

production-function estimation, then explain how we recover markdowns and relate them to the conventional single-labor-input measure.

3.1 Labor Type Inputs

We treat different worker types as separate inputs in the production function, building on previous studies that have allowed for heterogeneous labor inputs (Hellerstein, Neumark, and Troske, 1999; Hellerstein and Neumark, 2007; Gallen, 2024; Yeh, Macaluso, and Hershbein, 2022). We observe each employer's total personnel costs in the firm accounts and the earnings paid to individual workers in the employee records. We use the employee records to divide personnel costs between worker groups. For example, if women receive 40% of the employer's recorded payroll, we assign 40% of its personnel costs to female labor. This assumes that costs beyond recorded earnings are proportional to payroll across groups. We then deflate these group-specific costs to construct the labor inputs used in estimation.

Let i index firms, t years, and $j = 1, \ldots, J$ worker types within a given classification. The conventional single-labor-input model is nested as the special case $J = 1$. Let B_{itj} denote the positive earnings paid by firm i to workers of type j in year t, constructed from the employer-employee links described in Section 2. We calculate the type's share of the firm's classified linked-worker payroll as

$$s^W_{itj} = \frac{B_{itj}}{\sum_{h=1}^{J} B_{ith}}. \tag{1}$$

Each classification margin permits at most 5 percent unclassified headcount and payroll.[1] As in the conventional single-labor-input specification, we measure labor using the firm's total personnel costs recorded in the accounting data, C^L_{it}. We allocate these costs across worker types using their payroll shares:

$$C^L_{itj} = s^W_{itj} C^L_{it}. \tag{2}$$

This construction preserves the firm's total labor cost while allowing its allocation across worker types to vary across firms and over time.

For estimation, we deflate the worker-type labor costs constructed in Section 3.1 and take logarithms:

$$L_{itj} = \frac{C^L_{itj}}{P^L_{it}}, \qquad \ell_{itj} = \log L_{itj}, \tag{3}$$

where P^L_{it} denotes the industry output-price deflator, expressed as a price factor equal to

[1] Unclassified payroll is allocated proportionally to the classified groups, leaving the shares in equation (1) unchanged.

one in 2020. Output, materials and capital are also deflated before taking logarithms; the corresponding price indices and input transformations are described in Appendix A.1.

We use the industry output-price index to express personnel costs in constant output-price units. This produces a measure of real labor expenditure, not hours worked or a physical quantity of labor. Differences in compensation and labor quality within worker groups therefore remain relevant to the interpretation of the estimates.

3.2 Production-Function Specification and Estimation

We estimate an industry-specific gross-output translog production function. We treat capital as predetermined and allow current labor and materials choices to respond to productivity observed by the firm but unobserved by the econometrician. A direct regression of output on inputs would therefore mix the contribution of those inputs with differences in productivity. We use materials demand to help account for this unobserved productivity, then identify the production-function parameters using restrictions on the timing of input choices.

Let y_{it} denote log real gross output, k_{it} log real capital, and m_{it} log real materials. Production is given by

$$y_{it} = f_{g(i)}\left(k_{it}, m_{it}, \ell_{it1}, \ldots, \ell_{itJ}; \boldsymbol{\beta}_{g(i)}\right) + \omega_{it} + \varepsilon_{it}, \tag{4}$$

where $g(i)$ denotes the firm's industry, ω_{it} is firm productivity observed by the firm before the variable inputs are chosen, and ε_{it} is an output measurement error. The baseline holds production-function coefficients fixed within each industry while allowing firm productivity to change annually. Thus, productivity can rise even when the coefficients are constant. We also estimate separate technologies for three periods to assess whether allowing the relationship between inputs and output to change affects our results.

Suppressing the industry subscript, the translog production function is

$$\begin{aligned} f(\cdot) = \beta_0 &+ \beta_K k_{it} + \beta_M m_{it} + \sum_{j=1}^{J} \beta_{L_j} \ell_{itj} \\ &+ \beta_{KK} k_{it}^2 + \beta_{MM} m_{it}^2 + \beta_{KM} k_{it} m_{it} \\ &+ \sum_{j=1}^{J} \left(\beta_{KL_j} k_{it} \ell_{itj} + \beta_{ML_j} m_{it} \ell_{itj} + \beta_{L_j L_j} \ell_{itj}^2\right) \\ &+ \sum_{j<h} \beta_{L_j L_h} \ell_{itj} \ell_{ith}. \end{aligned} \tag{5}$$

This specification allows the output elasticity of each input to vary with the firm's entire

input mix. In particular, the output elasticity of worker type j is

$$\theta^L_{itj} = \beta_{L_j} + 2\beta_{L_jL_j}\ell_{itj} + \beta_{KL_j}k_{it} + \beta_{ML_j}m_{it} + \sum_{h\neq j}\beta_{L_jL_h}\ell_{ith}, \tag{6}$$

The corresponding materials output elasticity is

$$\theta^M_{it} = \beta_M + 2\beta_{MM}m_{it} + \beta_{KM}k_{it} + \sum_{j=1}^{J}\beta_{ML_j}\ell_{itj}. \tag{7}$$

The estimation follows the proxy-variable approach of Levinsohn and Petrin (2003) and Ackerberg, Caves, and Frazer (2015), using materials demand to control for unobserved productivity. Materials demand is assumed to be strictly increasing in productivity, conditional on capital and the worker-type labor inputs. Conditional on these inputs, productivity is assumed to be the sole unobservable entering the materials-demand proxy.[2]

The first stage estimates the predictable component of output, which combines the production function and productivity observed by the firm. It does not separately identify their contributions. In the second stage, each candidate set of production-function coefficients implies a productivity series. We select the coefficients using the assumption that unexpected productivity changes are unrelated to the predetermined and lagged inputs used as instruments.

Let $\widehat{\phi}_{it}$ denote the fitted first stage component of the log output. In the second stage, productivity implied by a candidate parameter vector is

$$\omega_{it}(\boldsymbol{\beta}) = \widehat{\phi}_{it} - f\left(k_{it}, m_{it}, \ell_{it1}, \ldots, \ell_{itJ}; \boldsymbol{\beta}\right). \tag{8}$$

We assume that productivity follows a first-order Markov process,

$$\omega_{it} = G(\omega_{i,t-1}) + \xi_{it}, \tag{9}$$

where ξ_{it} is the productivity innovation. The parameters are then identified from moment conditions of the form

$$\mathbb{E}\left[\mathbf{z}_{it}\xi_{it}(\boldsymbol{\beta})\right] = \mathbf{0}, \tag{10}$$

Here, $\mathbf{z}_{it}$ contains current capital, lagged materials and worker-type labor inputs, and

[2] Our inputs are measured using deflated expenditures rather than physical quantities. Unobserved differences in input prices and quality may therefore introduce additional heterogeneity beyond the scalar productivity term (Yeh, Macaluso, and Hershbein, 2022). Distinguishing worker types accounts for some of the observable workforce heterogeneity, but does not eliminate differences in compensation or labor quality within each type.

their corresponding squared and interaction terms. Each worker-type elasticity is identified using these moments and variation in its input conditional on the firm's other inputs.

Beyond the input-timing restrictions, identification requires relevant variation in input choices. For gross-output production functions, materials use must vary beyond firms' responses to productivity (Gandhi, Navarro, and Rivers, 2020). Yeh, Macaluso, and Hershbein (2022), whose panel covers 1976–2014 (39 years), argue that changes in materials prices over a long panel provide this variation. Our panel covers 27 years, 1997–2023.

3.3 Identification and Interpretation of Markdowns

We assume that firms minimize costs and that materials are chosen statically, without adjustment costs, and used only to produce output (Yeh, Macaluso, and Hershbein, 2022). The production function must be twice continuously differentiable with respect to materials and each worker-type labor input, a requirement satisfied by our translog specification. Materials must also satisfy the price-taking requirement:

Condition 1. *Firms take material input prices as given.*

Together with the other material-input assumptions stated above, Condition 1 allows the materials first-order condition to identify the product markup. Section 4.1 examines this condition and its implications for our estimates. Observation 1 below shows that within-firm worker-type markdown ratios do not require this condition, conditional on identified elasticities and the maintained labor-input assumptions.

Condition 2. *Each worker-type labor input is chosen statically within the annual period and entails no adjustment costs.*

Hiring, training and separation costs may make labor-input choices forward-looking, so that contemporaneous wedges also reflect these costs and the anticipated benefits of current employment for future production (Chan et al., 2026). Section 4.2 assesses this condition at our annual observation frequency and examines the implications of departures from it.

Under these input-choice conditions, we recover worker-type markdowns by combining the estimated output elasticities with observed input expenditure shares. Let $\widetilde{R}_{it}$ denote firm revenue after correcting for output measurement error, C_{it}^{M} materials expenditure, and C_{itj}^{L} the labor cost of worker type j. Their expenditure shares are

$$\alpha_{it}^{M} = \frac{C_{it}^{M}}{\widetilde{R}_{it}}, \qquad \alpha_{itj}^{L} = \frac{C_{itj}^{L}}{\widetilde{R}_{it}}. \tag{11}$$

The firm's product markup is recovered from the materials first-order condition:

$$\mu_{it} = \frac{\theta_{it}^{M}}{\alpha_{it}^{M}}. \tag{12}$$

Combining this markup with the first-order condition for worker type j gives the worker-type markdown:

$$\nu_{itj} = \frac{\theta_{itj}^{L}/\alpha_{itj}^{L}}{\theta_{it}^{M}/\alpha_{it}^{M}} = \frac{\theta_{itj}^{L}}{\alpha_{itj}^{L}}\mu_{it}^{-1}. \tag{13}$$

When $J = 1$, equation (13) produces the conventional firm-level markdown. When $J > 1$, it produces a separate markdown for every firm-worker-type pair. A markdown above one implies that the worker type is paid less than its marginal revenue product; $1/\nu_{itj}$ is the share of that marginal revenue product paid to the worker.

Comparing two worker groups at the same employer removes a common component of their estimated wedges: the firm's product markup. Relative markdowns can therefore be recovered without using materials to identify that markup. This does not remove the assumptions needed to estimate the worker-type elasticities or interpret firms' labor choices.

Observation 1 (Identification of relative worker-type markdowns)**.** *Suppose the worker-type output elasticities are identified and the maintained labor-input conditions hold, including Condition 2. For two worker types j and h employed by firm i in year t, with positive labor costs and output elasticities, the relative markdown is identified as*

$$\frac{\nu_{itj}}{\nu_{ith}} = \frac{\theta_{itj}^{L}/\alpha_{itj}^{L}}{\theta_{ith}^{L}/\alpha_{ith}^{L}} = \frac{\theta_{itj}^{L}}{\theta_{ith}^{L}}\frac{C_{ith}^{L}}{C_{itj}^{L}}. \tag{14}$$

This result does not require Condition 1 and remains valid when firms exercise buyer power in material input markets.

Proof. Under the maintained labor-input conditions, the ratio of the output elasticity to the expenditure share satisfies

$$\frac{\theta_{itj}^{L}}{\alpha_{itj}^{L}} = \mu_{it}\nu_{itj} \tag{15}$$

for each worker type, where μ_{it} denotes the firm's product markup. Dividing this relationship for types j and h eliminates the common markup. Substituting the expenditure-share definitions gives the final expression. Neither step requires identifying the markup from the materials first-order condition. □

Recovering relative worker-type markdowns therefore requires fewer restrictions than recovering their individual levels. If firms exercise buyer power in material markets, recovered

labor markdowns are measured relative to the materials wedge. Any common materials wedge cancels in the worker-type markdown ratio within the firm-year, even when that wedge varies across firms and over time.

To report earnings and estimated marginal revenue products on a common scale, let U_{itj} denote the assigned-worker counts under Primary-employer assignment described in Section 2. Using the positive earnings B_{itj} paid to the group, we define

$$e_{itj} = \frac{B_{itj}}{U_{itj}}, \qquad \widehat{\mathrm{MRPL}}_{itj} = \widehat{\nu}_{itj} e_{itj}. \tag{16}$$

The reported MRPL is thus an employer–worker-type benchmark expressed in annual earnings units. This normalization uses worker earnings rather than total personnel costs and leaves the estimated markdown unchanged. It does not measure individual or hourly productivity; the production-function labor inputs remain deflated personnel costs.

The mapping from recovered labor wedges to employer monopsony power depends on the wage-setting mechanism. Bargaining over wages alone and joint bargaining over wages and employment imply different relationships between wages and marginal revenue products (Dobbelaere and Mairesse, 2013). Worker bargaining power can offset employer monopsony power through rent sharing, so that the observed wedge need not identify either side's power separately (Wong, 2026). Binding wage constraints can also separate observed wedges from employers' potential wage-setting power (Bhuller et al., 2025). We interpret our recovered markdowns as follows:

Interpretation 1. *Recovered worker-type markdowns measure wedges under the prevailing wage-setting mechanism, rather than employer monopsony power independently of worker bargaining.*

A higher markdown means that a group receives a smaller share of its estimated marginal revenue product in pay. It does not establish whether this reflects weaker outside options, less worker bargaining power or other features of wage setting. Similarly, a markdown below one describes pay above estimated MRPL under the maintained assumptions; it does not independently identify stronger worker bargaining power. Section 4.3 discusses the scope for firm-level wage setting under Swedish collective agreements and how negotiated requirements may bind differently across worker types.

3.4 Relationship to the Single Labor Input Markdown

Treuren (2025) provides an aggregation benchmark relating worker-type markdowns to the single-labor-input markdown (Lemma 1).

This benchmark requires worker types to enter a continuously differentiable revenue function through a single, continuously differentiable labor aggregate with constant returns to scale (Treuren, 2025). Holding labor inputs fixed, other inputs therefore cannot change the types' relative marginal contributions. Scaling all worker-type inputs proportionally scales the labor aggregate by the same factor; worker types may nevertheless differ in productivity and complement one another.

Under these conditions the single labor input markdown is the payroll-weighted average of worker type markdowns:

$$\nu_{it}^{L} = \sum_{j=1}^{J} s_{itj}^{W} \nu_{itj}, \qquad s_{itj}^{W} = \frac{C_{itj}^{L}}{\sum_{h=1}^{J} C_{ith}^{L}} = \frac{B_{itj}}{\sum_{h=1}^{J} B_{ith}} \tag{17}$$

The aggregate markdown compares the marginal revenue contributions of all worker types with their combined labor cost. Under a constant-returns labor aggregator, these contributions sum to the contribution of the aggregate input. Each worker type therefore enters the aggregate markdown with a weight equal to its share of total labor costs.

The weights depend on compensation. Holding a worker type's total marginal-revenue-product contribution fixed, a higher markdown implies a smaller labor-cost share and therefore a lower weight in the aggregate markdown.[3]

An aggregate markdown close to one can conceal substantial differences across worker types, including markdowns above and below one. Changes in the aggregate can also reflect changes in labor-cost weights, rather than changes in worker-type markdowns alone.

Our unrestricted translog specification does not impose these aggregation restrictions, so the equality need not hold between the separately estimated single-input and worker-type models. Section 5.4 compares markdowns from the single-labor-input model with the labor-cost-weighted aggregates of worker-type markdowns defined above.

4 Recovering Markdowns in the Swedish Context

This section examines selected conditions highlighted in Section 3 and their implications for interpreting our estimated markdowns. We first assess the assumption that firms take material input prices as given (Condition 1). We then consider firms' ability to adjust

[3] Under the stated aggregation restrictions, the same aggregate can be expressed as an elasticity-weighted harmonic mean:

$$\nu_{it}^{L} = \frac{\sum_j \theta_{itj}^{L}}{\sum_j \theta_{itj}^{L} / \nu_{itj}}. \tag{18}$$

For positive markdowns and elasticities, this is weakly below the arithmetic mean using the same elasticity weights.

worker-type labor inputs within a year (Condition 2). Finally, we discuss how Sweden's collective bargaining institutions shape the interpretation of recovered worker-type wedges (Interpretation 1).

4.1 Buyer Power in Material Input Markets

Condition 1 provides the competitive materials benchmark used to recover labor markdowns. This benchmark follows established applications of the production approach (Dobbelaere and Mairesse, 2013; Yeh, Macaluso, and Hershbein, 2022). As Yeh, Macaluso, and Hershbein (2022) emphasize, materials are often traded in broad, potentially international markets, providing a rationale for treating individual firms as price takers.

If firms exercise buyer power in material markets, recovered labor markdowns are measured relative to the materials wedge, understating their levels when that wedge exceeds one. Distinguishing worker types nevertheless provides an identification advantage: as Observation 1 establishes, the common materials wedge cancels from within-firm markdown ratios, conditional on identifying the worker-type output elasticities and maintaining the labor-input assumptions.

Imperfections in material and labor markets may also be linked through rent sharing. For example, in Dutch manufacturing, Treuren (2025) finds that exchange-rate-induced increases in material-market rents among firms importing intermediates from China raise wages without corresponding increases in labor's marginal revenue product. Such wage responses are an economic consequence of materials buyer power, distinct from the measurement issue addressed by the observation.

We next examine whether within-firm worker-type markdown ratios differ across sourcing environments. Although access to foreign suppliers may broaden sourcing opportunities, international transactions can also involve substantial buyer bargaining power (Morlacco, 2020; Alviarez et al., 2023). We compare firms with and without recorded goods imports, without interpreting import status as a direct measure of materials buyer power.

Appendix Figure B1 reports this comparison using the primary-employer specification used in the main results. After controlling for industry, firm size, ownership, export status and workforce composition, the estimated differences are small: the female/male markdown ratio is 0.85 percent higher among recorded importers, while the foreign/native ratio is 0.61 percent higher. Neither difference is statistically distinguishable from zero at the 5 percent level. Recorded goods imports are therefore associated with limited differences in relative worker-type markdowns among observably similar firms; this comparison does not directly test the competitive-materials assumption.

4.2 Labor Adjustment and Forward-looking Choices

Condition 2 requires each worker-type labor input to be chosen statically within the annual period and without adjustment costs. Hiring, training and separation costs and anticipated future staffing needs can make firms' labor-input decisions forward-looking, so observed wage wedges may reflect dynamic employment decisions as well as employer market power (Chan et al., 2026). For example, a firm may retain workers during a temporary downturn to avoid future recruitment and training costs. Their current marginal revenue product may then fall below their wage, even though retaining them is profitable over time. Conversely, hiring costs can leave current marginal revenue products above wages without the entire gap reflecting employer market power. In Sweden, statutory employment protection and collective agreements govern the conditions and procedures for separations, potentially making employment adjustment costly.[4]

An advantage of our within-firm comparisons is that adjustment costs and other forward-looking considerations cancel if they affect both worker types' markdowns proportionally. Relative markdowns may nevertheless be affected if hiring, training or separation costs or anticipated future staffing needs differ across worker types, for example because these groups work in different occupations. Using Danish data, Chan et al. (2026) find that static monopsony power accounts for most of the average markdown level, although dynamic employment costs explain nearly all its cross-sectional dispersion. In French manufacturing, Wong (2026) finds that accounting for quadratic labor adjustment costs leaves his estimated worker bargaining power broadly similar to the baseline. Long-term wage contracts also affect this interpretation: Agostinelli et al. (2025) show that low firm-specific labor supply elasticities need not imply large instantaneous markdowns when workers value future returns to firm-specific human capital.

We therefore interpret our estimates as contemporaneous wedges between MRPL and earnings, which may also reflect dynamic employment decisions. Within-firm comparisons remove these influences when they affect worker types proportionally, but cannot separate differences in adjustment costs from differences in employer market power.

4.3 Collective Bargaining and the Interpretation of Markdowns

Interpretation 1 treats recovered markdowns as wedges under the prevailing wage-setting institutions. Accounting for worker bargaining is a general consideration in the production-based literature. Yeh, Macaluso, and Hershbein (2022, Online Appendix O.1) report an imprecise negative association between union coverage and markdowns in U.S. manufactur-

[4] The Employment Protection Act (LAS) regulates notice periods and priority in redundancy dismissals; collective agreements may derogate from specified provisions. See Sections 2b, 11 and 22.

ing and discuss how bargaining affects their interpretation. Wong (2026) explicitly models employer market power and worker bargaining jointly in French manufacturing.

Collective bargaining is a central feature in the Swedish labor market. Collective agreements are negotiated between unions and employer organizations, typically at the sector level. Collective-agreement coverage is high in an international context (Bhuller et al., 2022), but varies across sectors. Appendix Figures C1 and C2 report sectoral coverage and its descriptive relationship with median firm markdowns in 2021. Agreements cover employees of covered employers regardless of union membership. Collective bargaining coverage does not imply that individual wages are fully determined at the sector level.

Sectoral agreements generally cover periods of two to three years and establish a framework for wage revisions, with the degree of local discretion varying across agreements. They may specify the *wage-increase pool* (löneutrymme) to be allocated locally, *individual guarantees* of a minimum increase for each employee, *minimum wage levels*, and procedures for local negotiations. Some agreements leave both the size and allocation of wage increases to local wage setting, without a centrally specified numerical increase (Medlingsinstitutet, 2025, pp. 32–34).

The second tier of wage bargaining is local. The employer and local union representatives negotiate over the size of the wage-increase pool, its allocation across employees, or both. The scope for local and individual wage setting varies across agreements, but centrally agreed increases generally do not impose a ceiling (Medlingsinstitutet, 2025, pp. 32–34). Collective agreements also accommodate market considerations in wage setting. Medlingsinstitutet (2026, p. 10) notes that starting wages are generally negotiated individually and documents agreement provisions that explicitly recognize the influence of market forces. Collective agreements therefore leave room for wages to respond to firm conditions and for employees within the same firm to receive different adjustments.

Collective bargaining does not eliminate employer wage-setting power. Consistent with this, Dobbelaere et al. (2024) find that collective bargaining and works councils in Germany are associated with a smaller prevalence of wage markdowns and higher wages relative to marginal revenue products. The relevant distinction is instead between collective-agreement coverage and binding wage-setting constraints. As discussed above, Swedish firms bound by collective agreements may still pay above contractual minimum wages and differentiate wages across employees, with even larger discretion over new-hires' wages. Evidence from Norway shows that even when negotiated requirements bind, the observed wage still reflects the firm's incentives to attract and retain workers (Bhuller et al., 2025). More generally, Dobbelaere and Mairesse (2013) show that the interpretation of production-based labor wedges depends on whether wages and employment are determined through monopsony or bargaining. Con-

ditional on the production-function and input-choice assumptions, our approach recovers wage–marginal-revenue-product wedges without separately identifying employer monopsony power and worker bargaining power. We consequently interpret recovered markdowns as wedges under the prevailing bargaining institutions, rather than as measures of employer power in the absence of collective bargaining.

Collective bargaining may also affect worker types differently within the same firm. These differences need not reflect agreements explicitly targeting gender or immigrant status; they may arise because worker types are unevenly distributed across occupations and pay levels. For example, collective agreements in private care have allocated larger wage-increase pools to assistant nurses, a female dominated occupation, than to other employees (Medlingsinstitutet, 2017, p. 55). Contractual wage floors and negotiated wage revisions can bind more strongly for some employees than for others. Differences in recovered markdowns may therefore reflect differences in the constraints imposed by collective agreements, alongside differences in workers' outside options and employers' wage-setting discretion. The coexistence of employer market power and worker bargaining power also features in Wong (2026), who models their joint role in wage determination. Our within-firm comparisons capture the resulting differences in wage–marginal-revenue-product wedges, but do not separately identify these mechanisms.

5 Results

We examine how markdowns differ across employers and worker groups, how these differences evolve over time, and what they imply for earnings inequality. We begin with the conventional specification that assigns one markdown to all workers at a firm, establishing differences across industries and employers. We then distinguish female and male labor inputs to examine gender differences within firms and their contribution to the earnings gap. Finally, we turn to foreign-born and native-born workers, focusing on the decline in their markdown gap and the roles of relative earnings and estimated marginal revenue products.

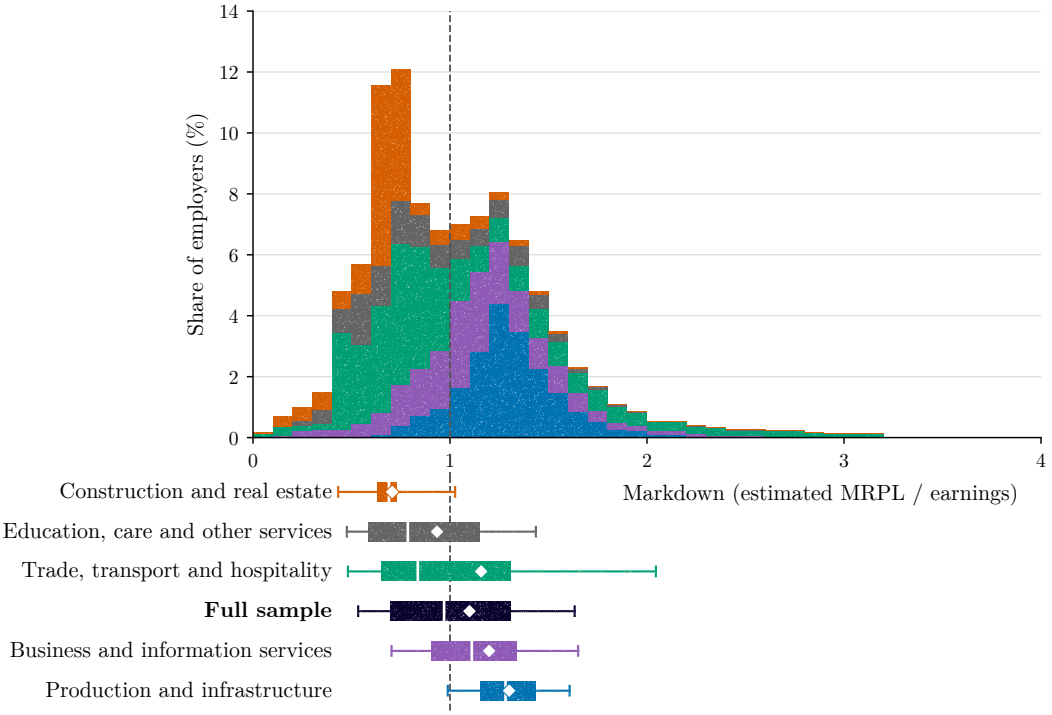
5.1 Markdowns Across Employers and Sectors

Our starting point is the single-input specification, where all worker expenditures enter the production function through the same labor input. This gives one markdown for all workers at a firm each year and allows us to compare our estimates with the literature (see Appendix D). We examine how much markdowns vary across and within industries, and whether worker groups face different markdowns because of where they work.

Figure 1 shows the distribution of employer markdowns in 2023, giving each firm equal weight. The mean is 1.10 and the median is 0.97, indicating a right-skewed distribution.

The median close to one indicates that the typical employer pays earnings close to estimated MRPL. The interquartile range is however 0.70–1.31 so there is substantial dispersion, both based on the industry of the firm but also between firms inside the same industry. For workers, that means that their markdown depends on which industry they are working in, where a worker at the typical firm in *Production and Infrastructure* sectors receives around 78% of their MRPL, while a worker in *Construction and real estate* earns 145% of theirs. However, even comparing these two industries at each end of the distribution does not determine a worker's markdown: in *Construction and real estate*, some employers have higher markdowns than the median employer in *Production and infrastructure*.

Figure 1: Single-input markdowns by industry group, 2023: primary-employer assignment

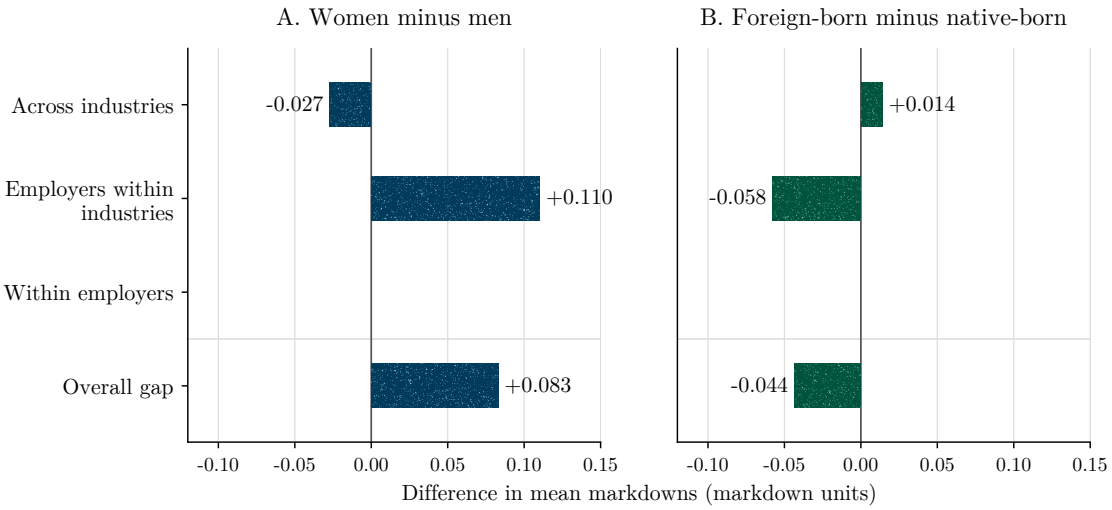


Notes: Primary employment; 28,397 equally weighted employers on common three-specification support. Stacks show employer shares per 0.1 markdown unit. The tail above 3.2 is omitted without rescaling. Boxplots retain the full distributions: P25–P75 boxes, P10–P90 whiskers, median lines and mean diamonds. Groups are ordered by median. The dashed line marks MRPL equal to earnings.

Most of the variation in markdown levels, 61.5%, occurs between firms within the same industry, while differences between industries account for the remaining 38.5%. This calculation gives equal weight to each of the 28,397 firms in our 2023 sample and uses the 57 production-function industries represented that year, rather than the five broad sector groups shown in Figure 1. These industries distinguish, for example, food, beverages and tobacco manufacturing from wood products manufacturing, land transport from water transport,

and legal and accounting activities from architectural and engineering services. Industry affiliation therefore captures meaningful differences in markdown levels, but variation among employers within the same industry accounts for a larger share.[5]

Figure 2: Group differences in exposure to single-input markdowns, 2023



Notes: Primary employment; 28,397 employers on common three-specification support. Panels show differences in arithmetic mean single-input markdowns using each group's employment weights. Industry allocation replaces actual industry shares with pooled all-worker employment shares; the remainder is employer allocation within industries. The within-employer row is blank because the specification assigns everyone at an employer the same markdown.

The single-labor input estimation assigns the same markdown to all workers at a firm and therefore we cannot reveal differences between workers at the same firm. However, it can reveal differences in groups' exposure to employers with different markdowns. Figure 2 shows how differences in markdowns between women and men (Panel A) and foreign-born and native-born workers (Panel B) can be decomposed across and within industries. Women are more likely to work in industries that have a lower markdown, but within those industries they are found at firms that impose higher markdowns. Taken together women have a higher markdown compared to men in this exercise, implying earnings losses for a woman with the same MRPL as a man from working at firms with higher markdowns of 6% compared to men. The conclusion is that women's distribution across industries lowers their average markdown relative to men's, but their distribution across employers within industries more than offsets this.

Foreign-born workers' employment allocation instead results in lower markdowns than

[5] This decomposition concerns the variance of markdown levels. For log markdowns, the within-industry share is 36.1% and the between-industry share is 63.9%.

those of native-born workers. Compared to Swedish-born workers they are more likely to work in industries with higher markdowns, but within industries they are found at firms with relatively lower markdowns. Taken together their distribution to firms with lower markdowns leaves them with earnings that are approximately 3.5% higher than Swedish-born workers' with the same MRPL. This exercise, as the previous one for gender, only accounts for markdown differences, comparing MRPL to pay, which does not imply the actual earnings differences.

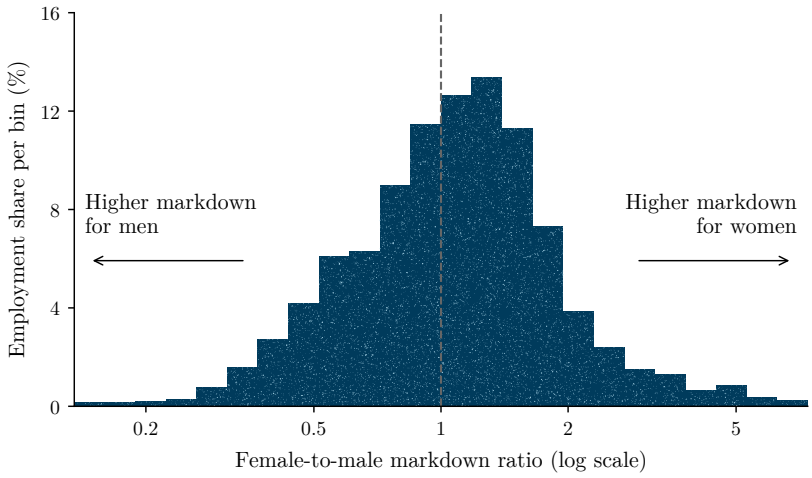
We will now move to estimating two separate worker-type specifications, which allows us to calculate one markdown for each worker group and firm, each year. The markdowns differ based on the industry common productivity of each worker type, as well as each firm's choice of input mix between capital, materials and each worker group. This allows us to reveal the bar that is missing from Figure 2, what the difference is in markdowns between workers at the same firm. We begin with an exposition of what this means for female and male markdown gaps and after that move to the gap between foreign-born and native workers.

5.2 Women Earn Less and Face Higher Markdowns

Do women face higher markdowns than men even when they work *at the same employer*? We answer this by estimating the same production function as before, only now we divide the labor input into expenditure on male labor and female labor. This allows us to calculate separate MRPL for each worker group, and allows for separate markdowns for women and men.

Figure 3 plots the distribution of the ratios of female markdowns to male markdowns for the same 28,397 firms in 2023 as the previous exercise. Employers are now weighted to give more importance to firms employing more workers. The weighting gives Figure 3 a labor focus, answering what markdown ratio a worker is likely to face at their firm. A ratio of 1 means that male and female workers face the same markdown. As Observation 1 shows, the common materials-market wedge cancels from these within-employer ratios, so recovering them does not require competitive materials markets. This result remains conditional on identifying the worker-type output elasticities and maintaining the labor-input assumptions; it does not extend to markdown levels or differences. The median ratio is 1.09: for at least half of workers their markdown is 9% higher or more if they are female. The mean is around 1.3, which means that for many workers the difference is much larger. We now examine whether these within-employer differences persist after accounting for observed workforce composition, how they combine with employer allocation and evolve over time, and their accounting implications for the gender earnings gap.

Figure 3: Within-employer female-to-male markdown ratios, 2023



Notes: Primary employment; 28,397 employers weighted by combined female and male employment. Ratios above one indicate higher markdowns for women. Bars show employment shares in equal-width log-ratio bins, labelled in ratios. Tails outside the displayed range are omitted without rescaling.

We next ask whether the gender markdown difference persists after accounting for differences in age, education and occupational composition between women and men within employers. Table 2 reports the *difference* between female and male markdowns, measured in markdown units, whereas Figure 3 reports their *ratio*. The table also reports standard errors and confidence intervals to assess whether these differences are statistically distinguishable from zero. In the raw comparison of markdown differences in the new specification with separate markdowns women have a higher markdown of around 0.08 when each worker is given equal weight, implying that for men and women with the same MRPL a woman would receive around 6.5% less pay under these group-average markdowns. Column 2 adds employer fixed effects, meaning that we now compare women and men within the same firm, similar to Figure 3. The within-firm markdown difference is on average 0.157 units, translating to women receiving around 12% lower pay than men due only to within-firm markdown differences. Accounting for differences in age and educational composition between women and men within employers reduces the estimated gap. Adding occupational composition reduces it further, to 0.106 markdown units. At equal MRPL, these adjusted markdowns imply earnings around 8.3% lower for women than for men.

Table 2: Gender markdown differences with employer and workforce controls

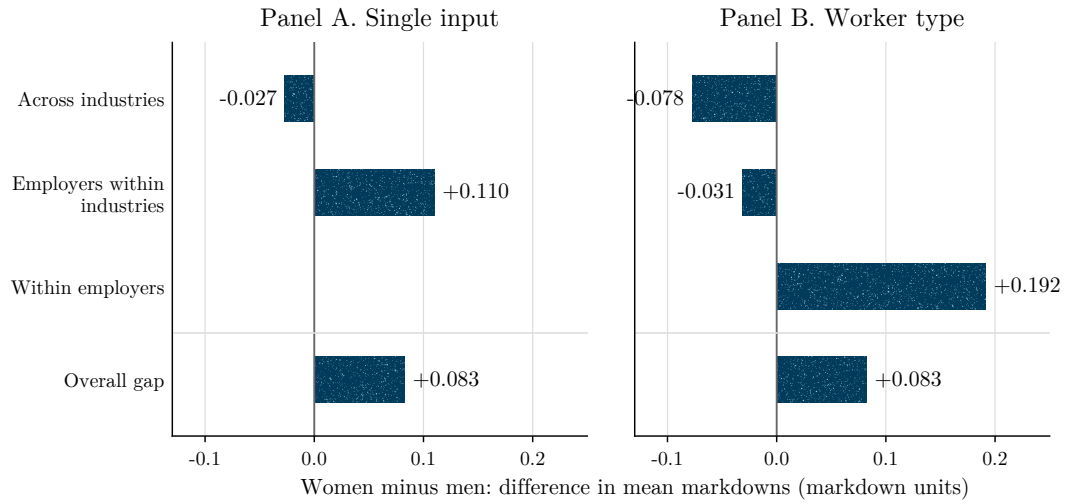
	(1) Pooled	(2) Employer FE	(3) Age and education	(4) Add occupation
Women minus men	0.082	0.157	0.136	0.106
Standard error	(0.017)	(0.025)	(0.027)	(0.041)
95% confidence interval	[0.049, 0.115]	[0.107, 0.207]	[0.082, 0.189]	[0.025, 0.187]
Employer fixed effects	No	Yes	Yes	Yes
Age and education	No	No	Yes	Yes
Occupation	No	No	No	Yes
Employers	28,147	28,147	28,147	28,147
Eligible primary employees	794,343	794,343	794,343	794,343

Notes: Primary employment, 2023; female-minus-male differences in markdown units. All columns use the same eligible workers at employers with both genders represented and recorded age and education. Each worker has unit weight; groups are not reweighted to equal employment within employers. Column (1) reports the difference in group means. Controls are employer-group means of age and squared age and education shares; column (4) adds broad occupation shares, including unknown occupation. Reported standard errors cluster by employer and condition on the estimated production functions. These are descriptive adjustments to employer-group benchmarks, not causal effects. The reporting sample differs from Figure 3 because of covariate availability.

We measure the control characteristics at the employer–gender level to match the level at which the production function identifies markdowns. The composition of women and men within each employer may contribute to their estimated markdown differences. Individual controls would introduce variation among workers assigned the same estimated markdown, rather than directly accounting for differences in composition between the groups being compared. The difference in markdowns between the groups remains and is statistically distinguishable from zero in all comparisons. The controls together account for about one third of the markdown gap, while two thirds of the gap remains.

Having established that women face higher markdowns than men at the same firm, we now examine how this difference combines with their allocation across industries and employers. Figure 4 decomposes the overall gender markdown gap into industry allocation, employer allocation within industries, and within-employer differences. Comparing the single-input and worker-type specifications shows how allowing gender-specific markdowns changes the contribution of each component.

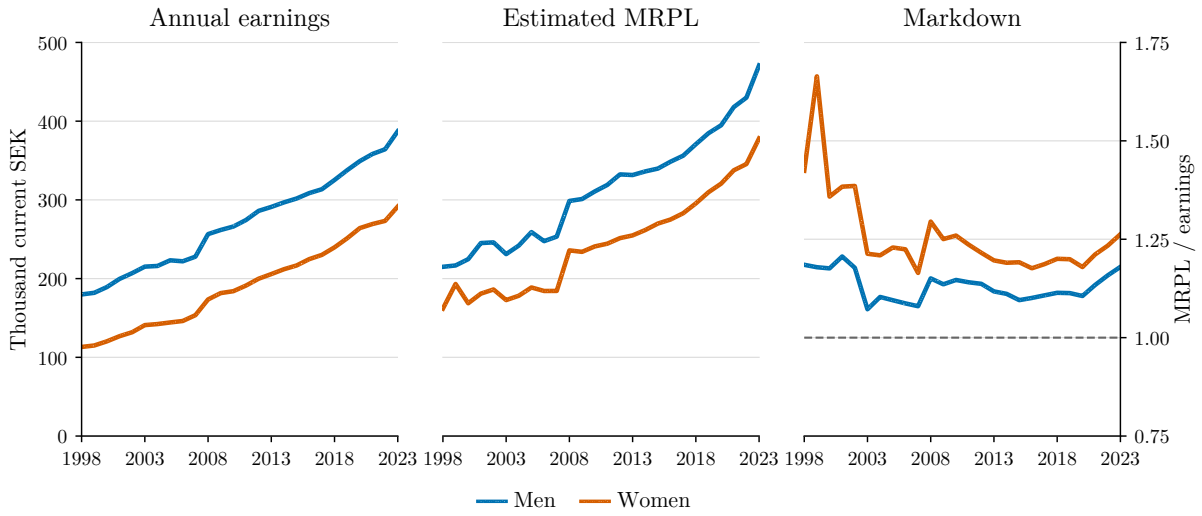
Figure 4: Gender markdown gaps: single-input and worker-type decompositions, 2023



Notes: Primary employment; the same 28,397 employers in both panels. Gaps are female-minus-male differences in arithmetic mean markdowns. Overall gaps use own-group employment weights; within-employer gaps use pooled employment weights. Industry allocation replaces actual industry shares with pooled shares; employer allocation within industries is the residual. Components sum to the total. The single-input within-employer row is blank by construction.

Women and men work at the same employers in both panels; what changes is whether we assign them a common or gender-specific markdown. The industry-allocation component becomes more negative, moving from -0.027 to -0.078 markdown units. This implies that women are concentrated in industries where their own markdown is lower than the common firm markdown suggests. The single-input model therefore understates how strongly industry allocation reduces the female–male markdown gap. The effect of employer allocation of women and men *within* industries changes sign between the two specifications. However, a high common employer markdown does not necessarily imply a high markdown for women at that employer. Instead, women's concentration at firms with high single-input markdowns does not translate into an equally unfavorable allocation when we use gender-specific markdowns. The largest difference between the panels is that we now estimate the within-firm difference in markdowns. The worker-type specification reveals a positive within-employer difference of 0.192 markdown units. The large positive within-employer component offsets the negative industry and employer allocation components, leaving the overall gender markdown gap virtually unchanged at 0.083 markdown units across the two specifications. We now move to placing the 2023 results into our full estimation sample, starting in 1998.

Figure 5: Earnings, estimated MRPL and markdowns for women and men, 1998–2023



Notes: Primary employment. Arithmetic means on annual common three-specification support, weighted by each group's own employment. Employers and weights vary annually. Earnings and MRPL are in thousands of current SEK per assigned worker. Mean markdowns average employer–group ratios and need not equal mean MRPL divided by mean earnings. The dashed line marks one.

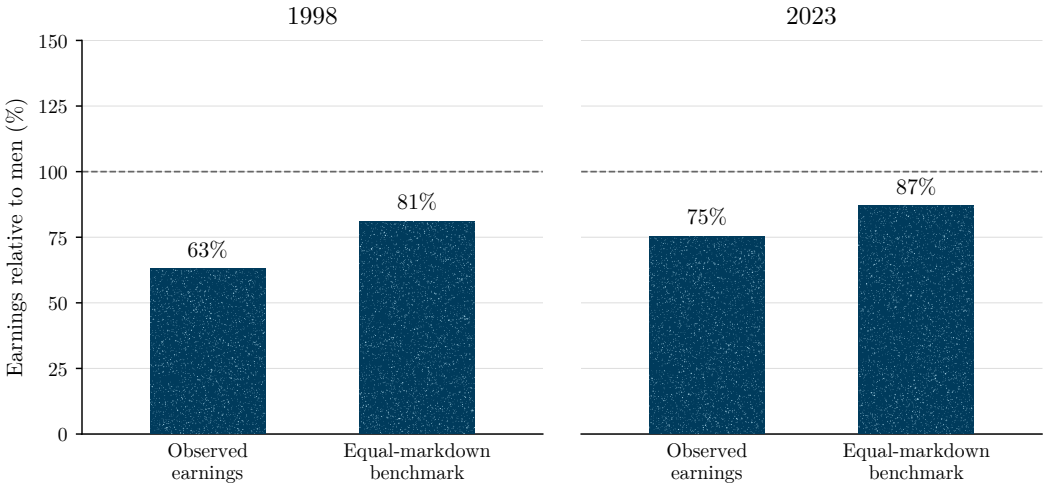
Figure 5 shows how women's markdowns fall over the period we study, moving from a mean markdown of 1.42 to 1.26, while men's markdowns are approximately 1.18 at both endpoints. Much of the closing of the gap occurs in the first half of the period: women's mean markdown gap is close to the 2023 gap already in 2013. The three panels allow us to understand where the markdown evolution comes from. Women's average earnings have risen from around 113,000 SEK to 292,000 SEK in 2023. This is an increase relative to men's earnings. Women's average earnings go from 63% to 75% of men's. At the same time, women's estimated MRPL rises from 75% to 80% of men's. Women's earnings thus increase relative to men's by more than their estimated MRPL does, alongside a decline in their relative markdowns. These series use each group's own employment weights and reflect both changes in employer-level outcomes and changes in where workers are employed.

A concern is that holding production technology fixed over the full period may misrepresent how women's and men's labor contributes to output at different points in time, thereby affecting their estimated markdown differences. Appendix Figure F1 shows that allowing technology to vary across three periods brings women's and men's estimated markdowns closer together, particularly early in the sample. The size of the estimated gender markdown gap therefore depends in part on how we allow production technology to change over time.

We have established that women have both lower average earnings and higher average markdowns than men. To conclude this section we ask how much of the gender pay gap

would be closed if we removed the difference between women's and men's markdowns. Having documented the contribution of industry and employer allocation, we now isolate the earnings implications of markdown differences within employers. We hold employment allocation fixed, preserving where women and men actually work, and assign women the male markdown at their existing employer. This benchmark therefore asks how much of the overall earnings shortfall is associated with within-employer markdown differences, without requiring assumptions about workers moving between employers or industries. We restrict the exercise to within-firm differences for two reasons: holding employment allocation fixed isolates the potential importance of pay-setting practices within workplaces; it also avoids assumptions about how workers' marginal products, employment, and earnings would change if they moved between firms or industries.

Figure 6: Gender earnings gaps and the equal-markdown benchmark



Notes: Primary employment. Mean earnings are expressed as percentages of the reference group's mean, using each group's own employment weights. The benchmark gives women the markdown of men at the same employer, holding estimated MRPL, reference-group earnings and employment allocation fixed. Both series use the same annual sample. The dashed line marks equal mean earnings. This is an accounting exercise, not a causal prediction.

Figure 6 shows that this adjustment closes nearly half of the observed gender earnings gap. In 2023, women earn 75% of men's average earnings; assigning them the male markdown at their existing employer raises this share to 87%. The earnings shortfall falls by 12 percentage points, equivalent to 47% of the observed gap. The corresponding share in 1998 is similar, at 49%, although the larger initial earnings gap implies a greater percentage-point contribution. Thus, despite the narrowing of the gender earnings gap over time, within-employer markdown differences account for approximately half of the remaining shortfall at both endpoints under the baseline specification. This is an accounting contribution rather than a predicted policy

effect, but it shows why where women work is only part of the picture: the share of their estimated contribution they receive at those workplaces also matters.

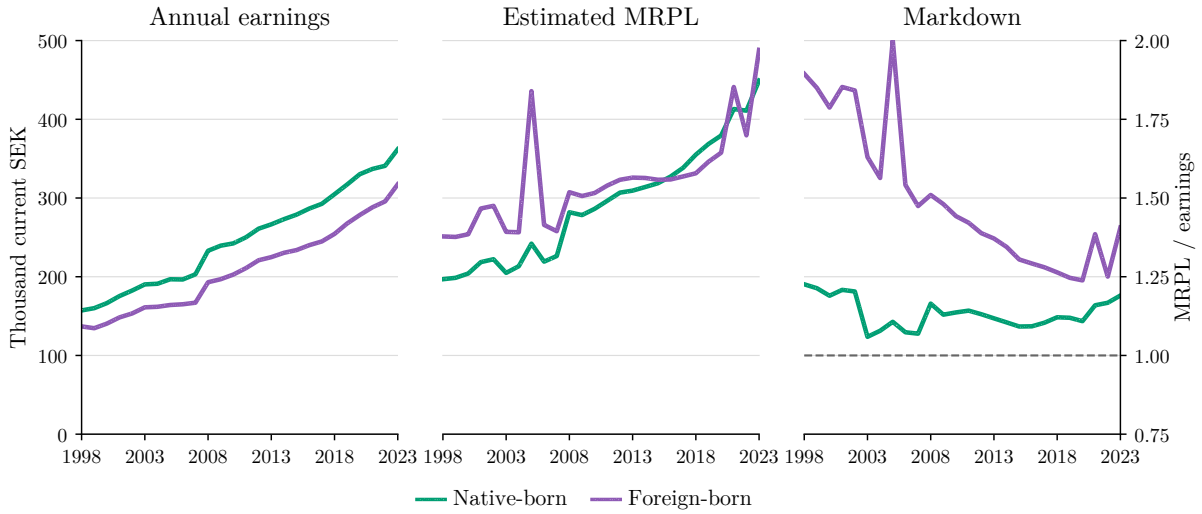
We repeat this exercise allowing production technology to differ across three periods, keeping employers and employment weights identical across specifications. On this matched sample, the benchmark closes 45.8% of the gender earnings shortfall in 2023 under the baseline technology, compared with 26.6% under period-specific technologies. In 1998, the corresponding shares are 38.3% and 13.2%. Within-employer markdown differences therefore retain a substantial accounting contribution to the earnings gap, although its magnitude is smaller when technology can change over time.

5.3 Foreign-born Workers' Markdowns Fall

Foreign-born workers face higher markdowns than native-born workers within employers. In 2023, the employment-weighted median within-employer ratio is 1.20, implying a 20% higher markdown for foreign-born workers. At equal MRPL, this corresponds to earnings approximately 17% lower than for native-born workers at the same employer. Appendix Figures G1–G2 report these comparisons and show how employment allocation partly offsets this within-employer disadvantage.

Foreign-born workers' average markdown falls substantially between 1998 and 2023. Figure 7 shows both the fall in average markdowns and the evolution of its components. Unlike the narrowing gender earnings gap, foreign-born workers' relative earnings change little: they earn 87% of native-born workers' earnings at the beginning and 88% at the end of the period. At the same time, foreign-born workers' mean markdown falls from 1.89 to 1.41, while native-born workers' changes relatively little, from 1.23 to 1.19. The gap narrows from 0.669 to 0.216 markdown units, approximately two-thirds. The narrowing markdown gap instead comes from foreign-born workers' mean estimated MRPL falling from 128% to 109% of the native-born mean. As panel b shows, neither MRPL falls, but the relative gain in MRPL in the period favors native-born workers. In this section we ask what accounts for the decline in foreign-born workers' relative estimated MRPL, and how much the remaining markdown difference contributes to their earnings shortfall. We first examine whether the narrowing markdown gap reflects changes in industry and employer allocation or differences within employers. We then assess its sensitivity to changes in production technology and investigate how changes in workforce composition contribute to the evolution of relative estimated MRPL within the model. Finally, we quantify how much of the remaining earnings shortfall is associated with within-employer markdown differences using the equal-markdown benchmark.

Figure 7: Earnings, estimated MRPL and markdowns for foreign-born and native-born workers, 1998–2023



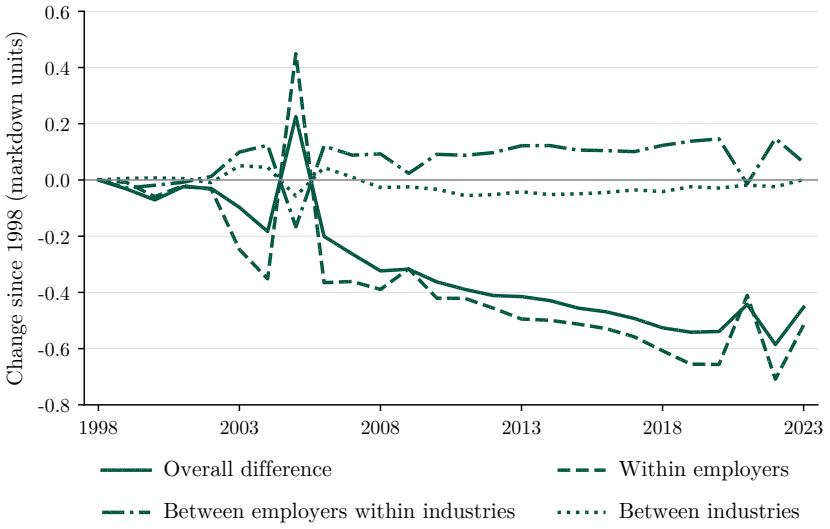
Notes: Primary employment. Arithmetic means on annual common three-specification support, weighted by each group's own employment. Employers and weights vary annually. Earnings and MRPL are in thousands of current SEK per assigned worker. Mean markdowns average employer–group ratios and need not equal mean MRPL divided by mean earnings. The dashed line marks one.

To narrow the possible explanations for the evolution of the markdown gap, we conduct a similar exercise to the gender gap decomposition, separating differences between industries, between firms within industries, and within firms, but following these components over the sample period. This allows us to distinguish changes in allocation from changes in the within-firm component. The overall gap between foreign-born and native-born workers falls by around 0.45 markdown units, from 0.67 to 0.22, a reduction of approximately 68%. To illustrate the magnitude, the group-average markdown levels imply an earnings difference at equal MRPL of around 35% in 1998 and 15% in 2023. These are illustrative comparisons based on average markdowns, rather than estimates of the earnings effect of equalizing markdowns within firms. Figure 8 decomposes the change, normalizing each component to zero in 1998. The contribution of industry allocation remains almost unchanged over the period. Allocation across firms within industries contributes modestly to widening the gap, as its contribution becomes less negative, moving from -0.47 to -0.41 markdown units. The decline instead comes from the within-firm component, which falls from 1.19 to 0.68 markdown units. Its decline therefore more than accounts for the overall narrowing, with changes in allocation partly offsetting it. Together with the evolution of earnings and estimated MRPL, this points towards changes in the relative marginal contribution of foreign-born labor, rather than substantial improvements in relative earnings. Such changes could reflect a changing workforce mix under a given technology or changes in production technology itself.

We next examine how the pattern depends on allowing technology to change over time.

Appendix Table H1 examines whether the decline in the within-employer origin markdown gap remains after accounting for workforce characteristics. On the same eligible primary-employer sample within each year, the unadjusted gap falls from 1.207 markdown units in 1998 to 0.685 in 2023. Adding age, education, occupation and time since latest recorded immigration changes these estimates to 1.525 and 0.694. The controls therefore affect the level of the gap, particularly early in the period, but do not remove its decline. These are descriptive comparisons on annual samples; occupation coverage varies over time.

Figure 8: Changes in the origin markdown gap, 1998–2023

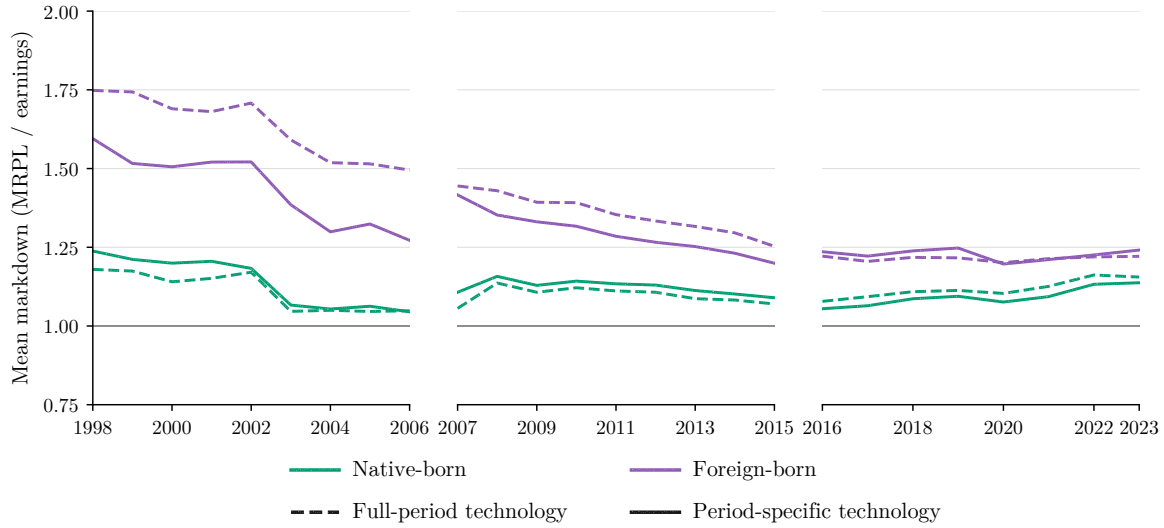


Notes: Primary employment, worker-type model. Changes since 1998 in the foreign-born-minus-native-born arithmetic mean markdown gap. Overall gaps use own-group employment weights; within-employer differences use pooled employment weights. Industry allocation standardizes industry shares to pooled shares; employer allocation within industries is the residual. The three components sum to the overall change. Employers and weights vary annually.

To examine whether the decline depends on holding production technology fixed over the full sample, we estimate separate technologies for 1998–2006, 2007–2015 and 2016–2023. Figure 9 compares these estimates to the baseline, using the same firms and employment weights under both specifications in each year. Allowing technology to change reduces the early markdown difference: in 1998, the foreign-born markdown falls from 1.75 to 1.60, while the native-born markdown rises from 1.18 to 1.24. The gap is therefore smaller, but it still narrows within each period. Under the period-specific technologies, the gap falls from 0.36 to 0.23 in the first period, from 0.31 to 0.11 in the second, and from 0.18 to 0.10 in the last. The last period differs from the earlier ones because foreign-born markdowns remain close to 1.24, while native-born markdowns rise. Taken together, allowing technology to change

affects the size of the estimated differences, particularly early in the sample, but does not remove the narrowing within periods.

Figure 9: Native-born and foreign-born markdowns under alternative technologies



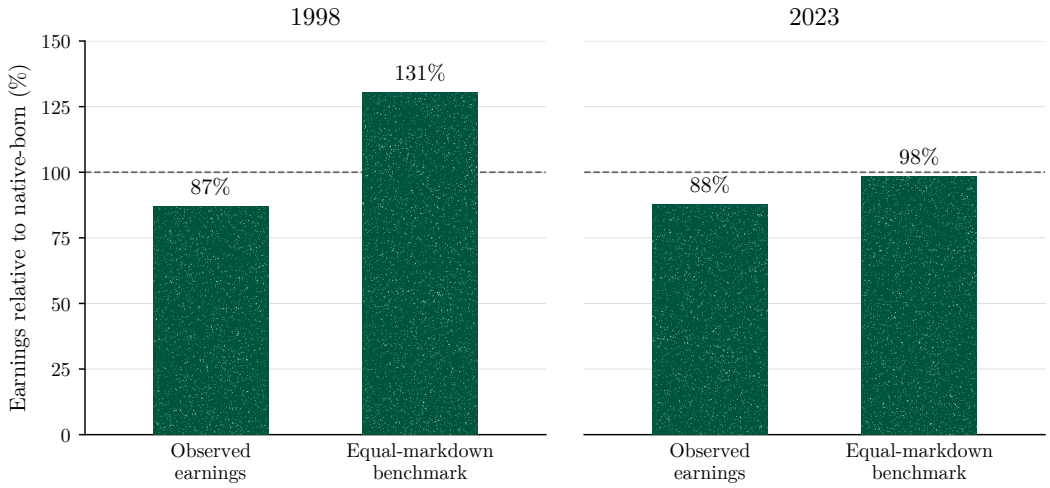
Notes: Primary employment. Arithmetic mean markdowns use each group's employment weights. Dashed lines use full-period technology; solid lines use technology estimated separately for each displayed period. Both fits use identical employers and weights within each year, on narrower support than Figure 8 (26,103 employers in 2023). Employers and weights vary annually. These group means include employer allocation, unlike Figure 9's within-employer component.

We next ask what changes in firms' inputs can tell us about the decline in relative estimated MRPL. We follow continuing firms within each technology period, holding employer weights fixed, and separate the contributions of changes in employment mix, employment scale, labor costs per worker, capital and materials. Within the fitted model, changes in the foreign-born to native-born employment mix account for most of the decline in relative MRPL, while capital contributes comparatively little. The model implies that as foreign-born employment grows relative to native-born employment, its contribution to output does not increase proportionately, lowering its relative marginal product. Relative earnings change more modestly, so the markdown difference narrows. This pattern holds under both full-period and period-specific technologies. The exercise therefore provides an accounting explanation for the decline within the production model, rather than evidence that individual foreign-born workers became less productive. It does not establish why firms changed their workforce mix.

To conclude this section, we ask how much of the remaining earnings shortfall is associated with foreign-born and native-born workers facing different markdowns at the same employer. We repeat the gender exercise, assigning foreign-born workers the native-born markdown at their existing employer while holding estimated MRPL, native-born earnings

and employment allocation fixed. Figure 10 shows that in 2023, this raises foreign-born workers' average earnings from 88% to 98% of native-born earnings. The adjustment closes approximately 87% of the observed earnings shortfall, leaving the two groups close to earnings parity. In 1998, the same exercise raises foreign-born earnings from 87% to 131% of native-born earnings, more than eliminating the shortfall. Equal markdowns therefore do not imply equal earnings: the benchmark preserves differences in estimated MRPL and where the groups work. Although the observed relative earnings shortfall is similar at the two endpoints, the benchmark increase falls from 44 to 11 percentage points. Taken together, the results suggest that within-employer markdown differences have become less important for foreign-born workers' relative earnings, but still account for most of their remaining earnings shortfall under the baseline estimates. This is an accounting contribution rather than a predicted policy effect, and its magnitude depends on the estimated production technology, particularly early in the sample.

Figure 10: Origin earnings gaps and the equal-markdown benchmark



Notes: Primary employment. Mean earnings are expressed as percentages of the reference group's mean, using each group's own employment weights. The benchmark gives foreign-born workers the markdown of native-born workers at the same employer, holding estimated MRPL, reference-group earnings and employment allocation fixed. Both series use the same annual sample. The dashed line marks equal mean earnings. This is an accounting exercise, not a causal prediction.

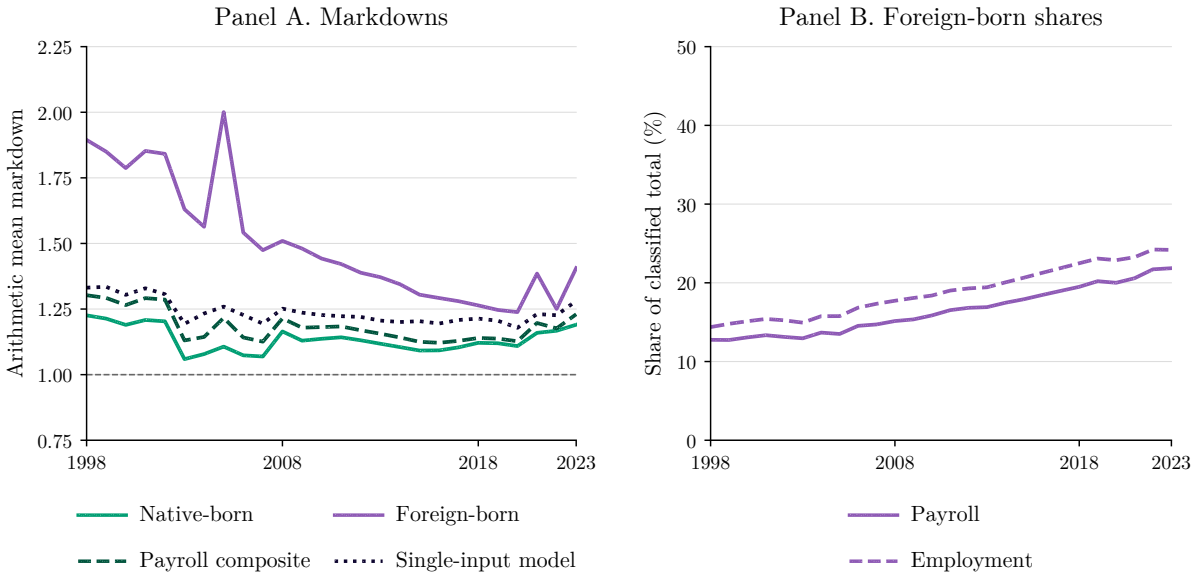
The narrowing origin markdown gap thus reflects a decline in relative estimated MRPL alongside little change in relative earnings. Over the same period, foreign-born workers account for a growing share of employment. Within the estimated production model, this expansion is associated with a lower marginal contribution of foreign-born labor relative to native-born labor. One interpretation is that the roles filled by foreign-born workers early in the sample contributed more to revenue at the margin, relative to those filled by native-born

workers, than the roles filled as their employment expanded. This interpretation concerns the roles workers fill within firms, rather than a decline in the productivity of individual foreign-born workers.

5.4 What a Single-input Markdown Conceals

The decline in foreign-born workers' markdowns raises a final question: how much of this change is visible in a single markdown for each employer? We compare the separately estimated single-input markdown with a composite of native-born and foreign-born markdowns, weighted by each group's share of payroll within the employer. Section 3.4 provides the theoretical aggregation benchmark, although our unrestricted translog models do not impose the conditions under which the composite and the single-input estimate coincide. Figure 11 shows that the foreign-born mean markdown falls from 1.89 in 1998 to 1.41 in 2023, while the native-born mean changes little, from 1.23 to 1.19. The payroll-weighted composite falls more modestly, from 1.30 to 1.23, and the single-input mean from 1.33 to 1.28.

Figure 11: Origin heterogeneity and the single-input markdown



Notes: Primary employment, 1998–2023. Group curves use own-group employment-weighted arithmetic means. The payroll composite weights the two groups' markdowns by their payroll shares within each employer. Its arithmetic mean and the single-input mean use the same combined native-born and foreign-born employment weights. Panel B reports pooled payroll and employment shares, not weights applied to national group means. Firms and weights vary annually. Differences between separately estimated models do not measure bias or causal effects.

Foreign-born workers account for 12.8% of classified payroll in 1998 and 21.9% in 2023, compared with employment shares of 14.4% and 24.2%. These pooled shares provide context; the composite uses payroll shares within each employer before averaging across employers.

A higher foreign-born markdown receives less weight where foreign-born workers account for a smaller share of payroll. Consequently, modest changes in the composite and single-input markdown can coexist with a substantial decline in foreign-born workers' markdowns. This comparison illustrates what employer-level summaries conceal, rather than measuring bias in the single-input estimates. The corresponding gender comparison in Appendix Figure I1 shows the same distinction: women's mean markdown falls from 1.42 to 1.26, while men's remains near 1.18. The payroll-weighted composite falls only from 1.25 to 1.20.

6 Conclusion

We examine how wage markdowns differ between women and men and between foreign-born and native-born workers using Swedish firm accounts linked to worker records for 1998–2023. Markdowns vary across industries and employers, but also between groups working at the same firm. Women and foreign-born workers face higher within-employer markdowns, while allocation across industries and employers partly offsets these differences. A single markdown for each employer therefore conceals important differences in how workers' pay compares with their estimated marginal revenue products.

Changes over time show why this distinction matters for interpreting earnings inequality. Women's relative earnings rise as their markdown disadvantage narrows. For foreign-born workers, the substantial decline in relative markdowns instead reflects a narrowing of their estimated MRPL advantage, while relative earnings change little. Allowing production technology to differ across periods changes the magnitude of the markdown gaps but preserves their narrowing within periods. A smaller markdown gap therefore need not imply that workers have closed their earnings shortfall.

The remaining markdown differences nevertheless have substantial implications for earnings. Holding estimated MRPL and employment allocation fixed, assigning women the male markdown at their employer closes nearly half of the gender earnings gap in our 2023 sample. Assigning foreign-born workers the native-born markdown closes most of their earnings shortfall. These are accounting benchmarks whose magnitudes depend on the estimated production technology, rather than predictions of a policy intervention. They show that understanding earnings differences requires examining both where workers work and the share of their estimated marginal contributions they receive in pay.

An important question for future research is which differences in workers' outside options and firms' wage-setting practices account for the within-employer markdown gaps we document.

References

Acemoglu, Daron and David Autor, "Skills, Tasks and Technologies: Implications for Employment and Earnings," in "Handbook of Labor Economics," Vol. 4B, Elsevier, 2011, chapter 12, pp. 1043–1171.

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer, "Identification Properties of Recent Production Function Estimators," *Econometrica*, 2015, *83* (6), 2411–2451.

Agostinelli, Francesco, Domenico Ferraro, Giuseppe Sorrenti, and Leonard Treuren, "Employment Relationships, Wage Setting, and Labor Market Power," Working Paper 34439, National Bureau of Economic Research October 2025.

Alviarez, Vanessa I., Michele Fioretti, Ken Kikkawa, and Monica Morlacco, "Two-Sided Market Power in Firm-to-Firm Trade," Working Paper 31253, National Bureau of Economic Research May 2023. Revised May 2026.

Åslund, Olof, Cristina Bratu, Stefano Lombardi, and Anna Thoresson, "Firm productivity, manager origin, and immigrant-native earnings disparities," 2025. Working paper, July 7.

Berger, David, Kyle Herkenhoff, and Simon Mongey, "Labor Market Power," *American Economic Review*, April 2022, *112* (4), 1147–1193.

Bhuller, Manudeep, Karl Ove Moene, Magne Mogstad, and Ola L Vestad, "Facts and Fantasies about Wage Setting and Collective Bargaining," *Journal of Economic Perspectives*, 2022, *36* (4), 29–52.

_ , Lukas Delgado-Prieto, Santiago Hermo, and Linnea Lorentzen, "Wage-Setting Constraints and Firm Responses to Demand Shocks," December 2025.

Brooks, Wyatt J., Joseph P. Kaboski, Yao Amber Li, and Wei Qian, "Exploitation of Labor? Classical Monopsony Power and Labor's Share," *Journal of Development Economics*, May 2021, *150*, 102627.

Card, David, Ana Rute Cardoso, and Patrick Kline, "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women *," *The Quarterly Journal of Economics*, May 2016, *131* (2), 633–686.

_ , _ , Joerg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2018, *36* (S1), 13–70.

Caselli, Mauro, Andrea Locatelli, and Stefano Schiavo, "Labour market frictions across worker types," *Economia Politica*, 2025, *42* (3), 799–827.

Chan, Mons, Elena Mattana, Sergio Salgado, and Ming Xu, "Dynamic Wage Setting: The Role of Monopsony Power and Employment Costs," August 2026. Working paper, August 25, 2026 version.

De Loecker, Jan and Frederic Warzynski, "Markups and Firm-Level Export Status," *American Economic Review*, October 2012, *102* (6), 2437–2471.

Deb, Shubhdeep, Jan Eeckhout, Aseem Patel, and Lawrence Warren, "Walras–Bowley Lecture: Market Power and Wage Inequality," *Econometrica*, 2024, *92* (3), 603–636.

Díez, Federico J., Davide Malacrino, and Ippei Shibata, "The Divergent Dynamics of Labor Market Power in Europe," IMF Working Paper 22/247, International Monetary Fund December 2022.

Dobbelaere, Sabien and Jacques Mairesse, "Panel Data Estimates of the Production Function and Product and Labor Market Imperfections," *Journal of Applied Econometrics*, 2013, *28* (1), 1–46.

_ , Boris Hirsch, Steffen Mueller, and Georg Neuschaeffer, "Organized Labor, Labor Market Imperfections, and Employer Wage Premia," *ILR Review*, 2024, *77* (3), 396–427.

Ek, Andreas, "Cultural Values and Productivity," *Journal of Political Economy*, 2024, *132* (1), 295–335.

Gallen, Yana, "Motherhood and the Gender Productivity Gap," *Journal of the European Economic Association*, June 2024, *22* (3), 1055–1096.

Gandhi, Amit, Salvador Navarro, and David A. Rivers, "On the Identification of Gross Output Production Functions," *Journal of Political Economy*, 2020, *128* (8), 2973–3016.

Hall, Robert E., Olivier Jean Blanchard, and R. Glenn Hubbard, "Market Structure and Macroeconomic Fluctuations," *Brookings Papers on Economic Activity*, 1986, *1986* (2), 285–338.

Hall, Tessa and Alan Manning, "Only Human? Immigration and Firm Productivity in Britain," CEP Discussion Paper 2060, Centre for Economic Performance, London School of Economics and Political Science 2024.

Hellerstein, Judith K. and David Neumark, "Production Function and Wage Equation Estimation with Heterogeneous Labor: Evidence from a New Matched Employer-Employee Data Set," in Ernst R. Berndt and Charles R. Hulten, eds., *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, University of Chicago Press, 2007, chapter 2.

_ , _ , and Kenneth R. Troske, "Wages, Productivity, and Worker Characteristics: Evidence from Plant-Level Production Functions and Wage Equations," *Journal of Labor Economics*, July 1999, *17* (3), 409–446.

Hirsch, Boris and Elke J. Jahn, "Is There Monopsonistic Discrimination against Immigrants?," *ILR Review*, 2015, *68* (3), 501–528.

_ , **Thorsten Schank, and Claus Schnabel**, "Differences in Labor Supply to Monopsonistic Firms and the Gender Pay Gap: An Empirical Analysis Using Linked Employer-Employee Data from Germany," *Journal of Labor Economics*, 2010, *28* (2), 291–330.

Hong, Joonkyo, Jin-Tan Liu, and Xinle Pang, "Trade Liberalization in Oligopsony," 2025. Working paper, April 17.

Kirov, Ivan and James Traina, "Labor Market Power and Technological Change in US Manufacturing," October 2023. Working paper.

Kline, Patrick, "Labor Market Monopsony: Fundamentals and Frontiers," in "Handbook of Labor Economics," Vol. 6, Elsevier, 2025, pp. 655–728.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler, "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market," *American Economic Review*, January 2022, *112* (1), 169–212.

Levinsohn, James and Amil Petrin, "Estimating Production Functions Using Inputs to Control for Unobservables," *The Review of Economic Studies*, April 2003, *70* (2), 317–341.

Manning, Alan, "Monopsony in Labor Markets: A Review," *ILR Review*, January 2021, *74* (1), 3–26.

Medlingsinstitutet, "Avtalsrörelsen och lönebildningen 2016," Annual report, Medlingsinstitutet, Stockholm February 2017.

_ , "Avtalsrörelsen och lönebildningen 2024," Annual report, Medlingsinstitutet, Stockholm February 2025.

_ , "Lönesättning och relativlöner ur ett jämställdhetsperspektiv," Report, Medlingsinstitutet, Stockholm March 2026.

Mertens, Matthias, "Micro-Mechanisms behind Declining Labor Shares: Rising Market Power and Changing Modes of Production," *International Journal of Industrial Organization*, March 2022, *81*, 102808.

_ , "Labor Market Power and Between-Firm Wage (In)Equality," *International Journal of Industrial Organization*, December 2023, *91*, 103005.

Morlacco, Monica, "Market Power in Input Markets: Theory and Evidence from French Manufacturing," October 2020. Working paper, University of Southern California.

Sharma, Garima, "Monopsony and Gender," 2023. Working paper, April 28.

Statens offentliga utredningar, "Förutsättningarna för en ny kollektivavtalad arbetslöshetsförsäkring," SOU 2023:41, Regeringskansliet 2023.

Tino, Stephen, "Labor Market Power, Firm Productivity, and the Immigrant-Native Pay Gap," September 2026. Working paper, September 19, 2026 version.

Treuren, Leonard, "Wage Markups and Buyer Power in Intermediate Input Markets," Discussion Paper Series DPS22.06, KU Leuven, Faculty of Economics and Business May 2025.

Webber, Douglas A., "Firm-Level Monopsony and the Gender Pay Gap," *Industrial Relations: A Journal of Economy and Society*, 2016, *55* (2), 323–345.

Wong, Horng Chern, "Understanding High-Wage Firms: Monopoly, Monopsony, and Bargaining Power," *American Economic Review*, 2026, *116* (7), 2504–2541.

Yeh, Chen, Claudia Macaluso, and Brad Hershbein, "Monopsony in the US Labor Market," *American Economic Review*, July 2022, *112* (7), 2099–2138.

A Data and Sample Construction

A.1 Data sources

We combine administrative records from Statistics Sweden (SCB) with firm accounts to construct an annual panel of Swedish non-financial firms and their workers. The main production-function estimation uses observations from 1997–2023. The unit of observation is the firm-year, identified by the organisation identifier in the delivered data. Worker records are linked to this identifier before constructing firm-level employment and payroll measures. The data are accessed through SCB's secure MONA environment.

Firm accounts and industry. The financial variables come from *Företagens ekonomi* (FEK, Structural Business Statistics), delivered in the annual FE tables. These records provide production value, value added, expenditure on raw materials and merchandise, personnel costs, the book value of fixed assets, and average employment. Industry information comes from the annual firm-register files. Historical SNI92 and SNI2002 codes are converted to SNI2007 using SCB correspondence tables. Production functions are estimated separately by industry family: non-manufacturing families follow SNI2007 divisions, while manufacturing is assigned to 14 production-function families. Financial activities are excluded.

Employment relationships and worker characteristics. JOBB records workers and their annual earnings from each employer through 2021; the 2022–2023 extension uses BAS worker records and newer population and origin snapshots. Primary-employer assignment assigns each worker exclusively to the employer paying their highest annual earnings. Assigned-worker counts measure neither hours nor full-time-equivalent employment.

Population records provide registered sex and year of birth. The delivered LISA-derived origin information identifies native-born and foreign-born workers. These sources determine the two worker classifications used in the main estimation: male and female workers, and native-born and foreign-born workers. The classifications are used separately. Unclassified workers remain visible in the underlying headcount and payroll totals and are handled by the coverage rules below.

Price indices. We use public SCB price series, normalised to 2020 equals 100. The output deflator is an industry gross-output implicit price index, calculated from current-price and chain-volume National Accounts series. Materials are deflated by the industry-specific National Accounts intermediate-consumption price index (P2pp), which covers goods and services. Capital is deflated by an industry investment-price index based on gross fixed capital formation. The industry mapping uses broader groups where required for consistent coverage; selected manufacturing combinations use fixed 2020 output weights. All labor-cost

inputs use the output deflator of the firm's industry family. The materials deflator covers a broader expenditure basket than the accounting input, and deflating booked assets does not turn them into a replacement-cost measure of productive capital.

A.2 Definition of variables

Let i denote a firm, t a year, and g a worker group within one classification. Monetary amounts are harmonized to Swedish kronor before constructing expenditure shares and deflated inputs. The following definitions apply to the main estimation.

- **Output.** Nominal output, R_{it}, is FEK production value (*Produktionsvärde*). It is not replaced by turnover or by value added. The dependent variable in the production function is $y_{it} = \log[R_{it}/(P^Y_{st}/100)]$, where P^Y_{st} is the output index assigned to industry family s.

- **Capital.** The nominal capital measure, K^N_{it}, is total booked fixed assets (*SummaAnlTillg*). This delivered aggregate includes financial as well as tangible and intangible fixed assets. The production-function input is $k_{it} = \log[K^N_{it}/(P^K_{st}/100)]$. It is a deflated book-value proxy, rather than a perpetual-inventory capital stock.

- **Materials.** Nominal materials expenditure, C^M_{it}, is expenditure on raw materials and merchandise (*SummaKostnaderFörRåvarorOchHandelsvaror*), converted from the negative expense convention in the source accounts to a positive amount. The input is $m_{it} = \log[C^M_{it}/(P^M_{st}/100)]$. This is the narrow accounting measure; production value minus value added is not substituted for it.

- **Labor costs.** Total labor cost, C^L_{it}, is FEK personnel cost (*SummaPersonalkostnader*), likewise converted to a positive expenditure. The aggregate labor input is $\ell_{it} = \log[C^L_{it}/(P^Y_{st}/100)]$. FEK average employment and linked JOBB headcount are separate measures. Neither replaces deflated personnel costs as the labor input in the main production function, and linked headcount is not full-time-equivalent employment.

- **Worker-type labor inputs.** Let B_{itg} be classified JOBB/BAS payroll for group g, and let B_{itU} denote payroll that cannot be classified within the relevant margin. For an admitted firm-year, residual payroll is allocated proportionally to classified payroll:

$$B^*_{itg} = B_{itg} + B_{itU}\frac{B_{itg}}{\sum_h B_{ith}}, \qquad C^L_{itg} = C^L_{it}\frac{B^*_{itg}}{B_{it}}, \tag{19}$$

  where h runs over the two classified groups and $B_{it} = \sum_h B_{ith} + B_{itU}$. Thus group

costs sum to total personnel cost within each classification. The corresponding input is $\ell_{itg} = \log[C^L_{itg}/(P^Y_{st}/100)]$. JOBB/BAS earnings determine the allocation shares; they are not assumed to equal the firm's total personnel expenditure.

- **Expenditure shares and wedges.** The first-stage residual $\widehat{\epsilon}_{it}$ is used to construct corrected nominal output, $\widetilde{R}_{it} = R_{it}\exp(-\widehat{\epsilon}_{it})$. Materials and labor shares are $\alpha^M_{it} = C^M_{it}/\widetilde{R}_{it}$ and $\alpha^L_{itg} = C^L_{itg}/\widetilde{R}_{it}$. Given the estimated elasticities, the materials wedge is $\mu_{it} = \theta^M_{it}/\alpha^M_{it}$ and the labor wedge is $\lambda_{itg} = \theta^L_{itg}/\alpha^L_{itg}$. The reported markdown is $\nu_{itg} = \lambda_{itg}/\mu_{it}$. These shares use nominal expenditures and corrected nominal output, although the production function uses deflated inputs.

- **Reported earnings and estimated MRPL.** Let U_{itg} denote assigned-worker counts under Primary-employer assignment. Annual earnings per unit are $e_{itg} = B_{itg}/U_{itg}$, and the reported MRPL benchmark is $\widehat{\text{MRPL}}_{itg} = \widehat{\nu}_{itg}e_{itg}$, as in equation (16). The numerator B_{itg} is the group's observed positive earnings, rather than its allocated share of FEK personnel costs. Both outcomes are employer–group benchmarks; they do not measure individual or hourly productivity.

A.3 Sample construction and cleaning

A.3.1 Required accounts and worker coverage

1. We require a valid organisation identifier and industry assignment, positive production value, positive booked capital, positive average employment, and correctly signed materials and personnel expenses. Missing inputs and zero expenditures are excluded before taking logarithms.

2. Firm-years must have positive linked employment and payroll, with positive headcount and payroll in each of the two gender groups and each of the two origin groups.

3. Within each classification, unclassified workers may account for at most 5 percent of linked headcount and at most 5 percent of linked payroll. Both conditions must hold for gender and for origin.

4. We verify that classified and unclassified headcounts exhaust total linked employment and that their payrolls exhaust total linked payroll. Admitted unclassified payroll is allocated as in equation (19); the raw missing-classification measures are retained for checking these identities.

Gender and origin are separate classifications, not four mutually exclusive worker types.

A.3.2 Input screens and common support

The input screens follow the implemented procedure based on Yeh, Macaluso, and Hershbein (2022):

1. Materials expenditure divided by production value must lie strictly between zero and one.

2. Within each industry-family-by-year cell, we exclude observations outside the first and ninety-ninth percentiles of each log input: capital, materials, and the labor input or inputs of the relevant specification. All input cutoffs are calculated on the same pre-trimming population, so the order in which inputs are screened does not affect retention. Observations equal to a cutoff are retained.

3. After this trimming, we calculate materials expenditure as a share of materials plus personnel costs. We exclude observations for which this share is below 15 percent of its mean within the remaining industry-family-by-year cell.

4. We apply these screens separately for the aggregate, gender, and origin specifications and intersect their retained firm-year keys.

The screens use nominal inputs; deflation is applied after selection and does not redefine the common sample.

A.3.3 Panel support and estimation families

Estimation pools 1997–2023 within each industry family, subject to the following requirements:

1. A current observation contributes to the second-stage moments only when the immediately preceding year belongs to the same organisation and industry family and both observations satisfy the common input screens. Gaps in the panel and changes in industry family interrupt the lag relation. The first year supplies lag information, so current estimation moments begin in 1998.

2. Following the minimum-observation criterion in Wong (2026, Appendix A.3), we require at least 300 pooled current firm-year observations contributing to the moments in each industry family. This threshold is not a minimum number of firms or a separate annual requirement.

3. The aggregate, gender, and origin models use common support, with numerical fit checks applied to each model. First-stage and elasticity-evaluation observations are distinguished from observations with valid lags for the moments.

The estimation population is distinct from the population retained for reporting markdowns.

A.3.4 Post-estimation reporting screens

Following the post-estimation screening approach in Treuren (2025, Appendix A.3.4), we apply the following rules separately for each model and industry family:

1. We require finite, strictly positive materials and labor wedges for all labor groups in the model.

2. We calculate first and ninety-ninth percentile cutoffs for each wedge on the same jointly positive evaluation population, pooling years within the family.

3. We retain an observation only if all wedges lie within their respective cutoffs, including the endpoints. The trimming variables are the materials and labor wedges, before dividing the latter by the former to obtain markdowns.

A positive markdown below one is retained if it passes these rules. A negative estimated capital elasticity alone does not remove an observation, and the procedure does not discard an entire family because some estimated elasticities are negative. Reporting support is available both within a specification and on the intersection of valid results across all three specifications. Comparisons on the latter support use the same firm-years after the reporting screens as well as before estimation.

A.3.5 Sample attrition and worker-group composition

Table A3 describes the worker-requirement step using a 2023 snapshot based on primary-employer assignment. Unlike the annual averages in the Primary-employer funnel in Table A2, these counts refer to a single year. The requirements reduce the starting population from 171,408 to 51,121 firms. The table distinguishes group absence from incomplete classification and from groups containing workers with no positive recorded payroll. Classification coverage is measured at the firm level, both for complete classification and allowing up to 5 percent unclassified headcount and payroll.

Table A4 compares firms that pass and fail the worker requirements under Primary-employer assignment, before subsequent input screens and family admission. Over 1997–2023, the retained firms account for 25.7 percent of firm-years, 77.3 percent of FEK employment and 78.6 percent of linked workers in the starting population. Their mean annual shares of nominal production value, FEK personnel costs and linked payroll are 82.6, 80.5 and 80.3 percent, respectively. Average FEK employment is 36.4 per retained firm-year and 3.7 per excluded firm-year. Industry composition also changes: manufacturing accounts

for 14.0 percent of retained firm-years and 7.2 percent of excluded firm-years, while construction accounts for 9.3 and 19.2 percent, respectively. The worker-qualified population consequently gives greater representation to larger employers with both gender and birthplace groups present. These descriptive differences identify the population covered by the worker requirements; they do not establish the effect of selection on estimated markdowns.

Table A1: Industry families and coverage: Primary-employer estimation sample, 1998–2023

Industry family	SNI 2007 codes	Firm-years	Firms
Agriculture and hunting	01	19,127	3,879
Forestry and logging	02	2,794	738
Fishing and aquaculture	03	487	104
Other mining and quarrying	08	461	111
Food, beverages and tobacco	10.1–12.0	20,052	2,944
Textiles, apparel and leather	13.1–15.2	4,030	657
Wood products	16.1–16.2	8,689	1,259
Paper and printing	17.1–18.2	8,904	1,316
Chemicals and pharmaceuticals	20.1–21.2	4,817	579
Rubber and plastics	22.1–22.2	6,467	813
Nonmetallic mineral products	23.1–23.9	4,206	592
Basic metals	24.1–24.5	3,265	379
Fabricated metal products	25.1–25.9; 33.11	27,538	3,829
Machinery and equipment	28.1–28.9; 33.12; 33.2	16,421	2,333
Electronics and electrical equipment	26.1–27.9; 32.5; 33.13–33.14	12,312	1,799
Transport equipment	29.1–30.9; 33.15–33.17	6,339	929
Furniture and other manufacturing	31.0–32.9 (except 32.5); other 33.1	5,986	968
Electricity, gas and steam	35	2,812	351
Waste collection and recovery	38	3,048	486
Construction of buildings	41	20,175	4,437
Civil engineering	42	2,733	622
Specialised construction	43	64,720	13,495
Motor vehicle trade and repair	45	10,264	2,607
Wholesale trade	46	25,238	6,714
Retail trade	47	23,949	6,783
Land transport and pipelines	49	36,183	7,636
Water transport	50	979	225
Air transport	51	557	135
Warehousing and transport support	52	6,713	1,423
Accommodation	55	25,743	3,875
Food and beverage services	56	116,041	26,217
Publishing	58	7,114	1,513
Film, television and sound recording	59	5,140	1,323
Telecommunications	61	1,872	421
Computer programming and consultancy	62	20,019	5,234

Continued on next page

Table A1 continued

Industry family	SNI 2007 codes	Firm-years	Firms
Information services	63	1,535	470
Real estate	68	21,001	4,817
Legal and accounting services	69	7,532	2,453
Head offices and management consultancy	70	9,117	3,162
Architecture, engineering and testing	71	16,419	3,975
Research and development	72	3,981	841
Advertising and market research	73	12,941	3,284
Other professional and technical services	74	6,191	1,652
Veterinary services	75	2,152	334
Rental and leasing	77	5,530	1,371
Employment activities	78	5,366	1,948
Travel agencies and reservation services	79	3,396	827
Security and investigation	80	1,163	281
Building services and landscaping	81	17,494	4,006
Office and business support	82	4,575	1,222
Education	85	22,552	4,109
Human health	86	20,263	3,884
Residential care	87	8,868	1,461
Social work without accommodation	88	4,457	1,355
Creative, arts and entertainment	90	3,304	859
Libraries, archives and museums	91	806	132
Gambling and betting	92	1,125	262
Sports, amusement and recreation	93	9,901	2,117
Membership organisations	94	1,864	361
Computer and household-goods repair	95	1,560	394
Other personal services	96	11,070	2,774
Total: 61 industry families		729,358	147,239

Notes: Common Primary-employer estimation sample before markdown reporting screens, matching Table 1. Non-manufacturing families follow SNI 2007 divisions; manufacturing uses the custom groups listed above. Repair activities in division 33 are assigned by detailed industry code. Only admitted families are shown; the petroleum family has no observations in this sample. Firm-years are summed over 1998–2023; 1997 supplies lagged inputs. Firms are distinct within each family. The total counts distinct firms across the full sample and therefore need not equal the sum of family counts, since a firm can change industry. Industry labels are abbreviated.

Table A2: Primary-employer sample construction and coverage, 1997–2023

	Firms		FEK employment		Linked workers
Sample	Annual mean	%	Annual mean	%	Annual mean
A. Accounting and worker-data requirements					
Positive employment in included non-financial sectors; valid identifiers	204,280	100.0	2,149,710	100.0	2,690,018
Positive output, capital and employment; positive materials and personnel expenditures	153,180	75.0	1,819,065	84.6	2,280,688
All three specifications					
firms meeting worker requirements	38,496	18.8	1,401,152	65.2	1,792,237
Single-labor-input specification					
linked workers with positive total payroll	150,012	73.4	1,813,567	84.4	2,280,671
Gender specification					
male and female workers with positive payroll; at most 5% unclassified[a]	76,827	37.6	1,622,150	75.5	2,066,624
Birthplace specification					
native- and foreign-born workers with positive payroll; at most 5% unclassified[a]	47,732	23.4	1,449,030	67.4	1,852,973
B. Input screening and estimation requirements					
Firms passing input trimming and materials-share restrictions	28,114	13.6	680,148	31.1	888,289
Firms in admitted industry groups receiving estimated elasticities	28,052	13.5	667,848	30.6	871,660
C. Post-estimation screening[b]					
All three specifications					
firms with reportable markdowns	20,489	9.9	445,876	20.4	592,199
Single-labor-input specification					
firms with reportable markdowns	26,002	12.5	601,126	27.5	784,689
Gender specification					
firms with reportable markdowns for both male and female workers	23,506	11.3	537,582	24.6	706,327
Birthplace specification					
firms with reportable markdowns for both native- and foreign-born workers	23,425	11.3	526,647	24.1	697,190

Notes: Primary-employer assignment. Counts are annual averages, rounded to integers. Percentages are ratios of period sums relative to FEK firms with positive recorded employment in the included non-financial sectors and valid identifiers. Each panel uses the baseline for its own period. FEK employment sums recorded average employment. Linked workers count primary-employer assignments, not full-time equivalents. Input restrictions retain the common sample of the single-labor-input, gender and birthplace specifications.

[a] The indented rows in panel A show specification-specific worker requirements before input screening, not final estimation samples; the row above is their intersection. The 5 percent limit applies separately to unclassified headcount and payroll within each classification.

[b] Panels B and C cover 1998–2023; 1997 contributes to estimation but has no reported markdowns. The first row in panel C is the intersection of the three specification-specific reporting samples shown below it. Positive-wedge and percentile restrictions apply to materials and all labor-type wedges.

Table A3: Primary-employer worker-group support and classification coverage, 2023

	Firms	%
Linked workers with positive total payroll	171,408	
A. Classification coverage in the linked-payroll population		
All linked workers have classified gender	171,408	100.00
Gender: at most 5% unclassified headcount and payroll	171,408	100.00
All linked workers have classified birthplace	171,356	99.97
Birthplace: at most 5% unclassified headcount and payroll	171,390	99.99
B. First applicable reason for exclusion		
A required group is absent; classification is complete in that dimension	120,275	99.99
Otherwise, no classified workers in a required group; unclassified workers make absence uncertain	5	< 0.01
Otherwise, all groups are present, but a group has no positive recorded payroll	0	0.00
Otherwise, all groups have positive payroll, but unclassified headcount or payroll exceeds 5 percent	7	0.01
Total excluded by worker requirements	120,287	100.00
Retained after worker requirements	51,121	

Notes: Counts refer to firms in 2023 and use the linked primary-employer roster. The starting population uses the same definition as the linked-payroll stage in Panel A of Table A2, which reports annual averages over 1997–2023. Panel A percentages use all starting firms as the denominator; Panel B percentages use excluded firms. Classification coverage is measured at the firm level. The 5 percent allowance applies separately to unclassified headcount and payroll; both worker groups must still have positive payroll. Panel B assigns each excluded firm to its first applicable reason in the displayed order. No positive recorded payroll combines zero, negative and missing earnings. Gender and birthplace are separate classifications.

Table A4: Primary-employer selection at the worker-group requirements, 1997–2023

	Retained	Excluded	Retained (%)
A. Coverage and firm size			
Firms	38,496	111,516	25.7
FEK employment	1,401,152	412,415	77.3
Linked workers	1,792,237	488,433	78.6
Production value			82.6
FEK personnel costs			80.5
Linked payroll			80.3
FEK employment per firm	36.4	3.7	

B. Industry composition and employment retention

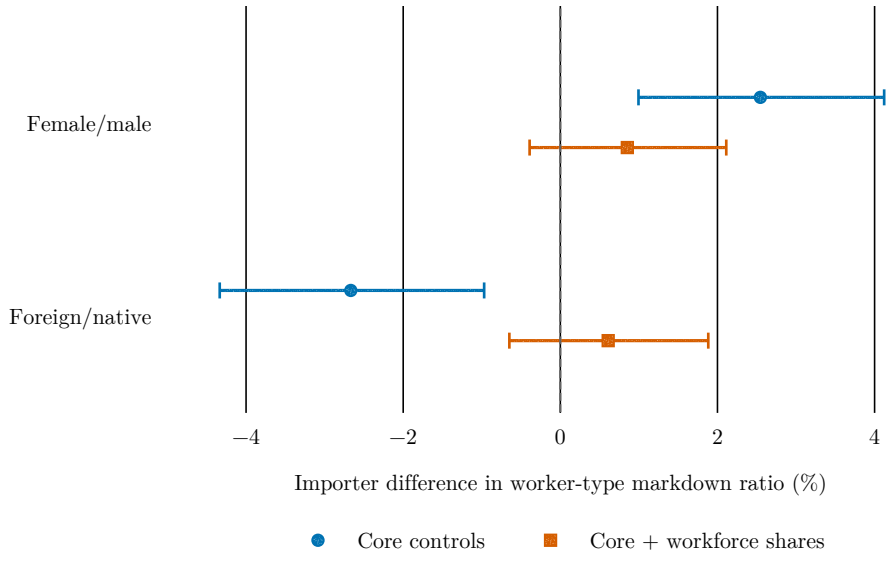
	Share of firms (%)		FEK employment retained after worker requirements (%)
Industry	Retained	Excluded	
Agriculture, forestry and fishing	2.5	5.5	41.2
Mining and quarrying	—	—	—
Manufacturing	14.0	7.2	90.8
Utilities and waste	0.8	—	—
Construction	9.3	19.2	58.6
Wholesale and retail trade	25.1	24.1	74.7
Transportation and storage	5.0	8.9	66.6
Accommodation and food	15.1	5.3	81.3
Information and communication	4.3	4.0	86.4
Real estate	2.4	3.3	69.6
Professional, scientific and technical	7.1	9.8	68.5
Administrative and support	4.7	3.1	85.7
Education	2.5	1.6	77.8
Health and social work	3.9	2.7	89.2
Arts, entertainment and recreation	1.7	1.8	72.0
Other services	—	—	—

Notes: Retained firms meet both positive-group-payroll requirements and both classification limits; excluded firms belong to the same linked-worker, positive-payroll population but fail at least one worker requirement. Retention is measured relative to the linked-worker, positive-payroll population in Table A2, before subsequent input, family-admission, and reporting restrictions. This is not a comparison with the final reporting sample. Counts are annual averages over all 27 years. Employment per firm divides period employment by period firm-years. Firm, employment and worker retention percentages are ratios of period sums within the starting population. Monetary measures report only retained shares: these are arithmetic means of the 27 annual nominal shares of production value, personnel costs and linked payroll. Industry shares divide each industry's retained or excluded firm-years by the corresponding national group; industry employment retention instead uses that industry's linked-payroll population. Linked workers count primary-employer assignments and are not full-time equivalents. A dash denotes a protected statistic or a statistic depending on a protected annual cell; no incomplete-period average or residual reconstruction is reported.

B Goods Imports and Worker-type Markdowns

Firm ownership. Serrano supplements the administrative records with firm ownership information collected from the Swedish Companies Registration Office. We use ownership information as a control in the importer sensitivity exercise; its availability is not an additional restriction on the main production-function estimation sample.

Figure B1: Worker-type markdown ratios and recorded goods imports, 2023



Notes: Conditional differences between firms with and without recorded goods imports in log female/male and foreign/native within-firm markdown ratios, reported as $100[\exp(\hat{\beta}) - 1]$. Bars: approximate 95% intervals (employer-clustered standard errors; 1.96). Core controls: industry fixed effects, quadratic log annual average employment, ownership and export status; the second specification adds female and foreign workforce shares. Both use 27,352 equally weighted firms (6,792 importers) and the primary-employer specification used in the main results, conditional on estimated elasticities. Goods imports do not directly measure materials buyer power.

C Collective-agreement Coverage and Markdowns

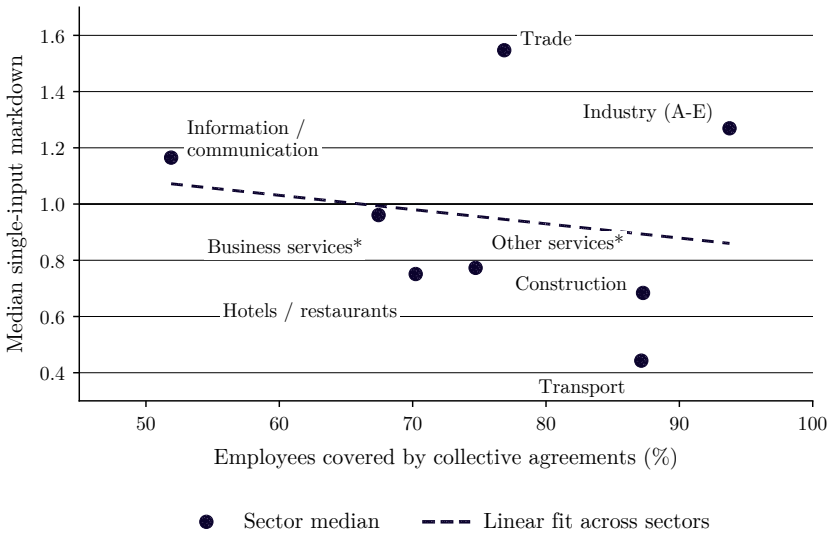
Collective-agreement coverage. We use sector-level counts of private-sector employees in 2021 from SCB and Fora, reported in Statens offentliga utredningar (2023, Table 10.5). Coverage is the share of employees working for employers bound by a collective agreement or a hängavtal, using the source's RAMS definition of employment. The published categories follow SNI2007 and combine some sectors: agriculture and industry (A–E), finance and business services (K–N), and education, health and other services (P–T). These measures characterize sector-level bargaining environments rather than individual firms' agreement status. They describe the source population, not coverage calculated within our estimation sample, and provide a 2021 cross-section rather than an annual coverage series.

Figure C1: Collective bargaining coverage among private-sector employees, 2021



Notes: Coverage is the employee share under regular collective agreements or hängavtal (RAMS). Sector pools are A–E, K–N and P–T (SNI 85–88 and 90–98); they describe the public source population, not the estimation sample.
Source: SCB and Fora, SOU 2023:41, Table 10.5.

Figure C2: Collective-agreement coverage and median firm markdowns, 2021



Notes: Points pair public employee coverage (including hängavtal) with sector medians of single-input markdowns: primary-employer assignment, common three-specification support, 27,414 equally weighted firms. Dashed line: equal-sector fit. Coverage describes the public source population, not sampled firms' agreement status. * Coverage includes finance and SNI 97–98, absent from the corresponding estimation sectors. The association is descriptive.
Source: Authors' estimates; SCB and Fora, Statens offentliga utredningar (2023, Table 10.5).

D Single-input Markdowns in the Literature

The single-input specification provides a benchmark for comparison with existing production-based markdown studies. Such comparisons require attention to sectoral coverage, reporting samples and aggregation. Much of the literature focuses on manufacturing, whereas our main results cover a broader set of nonfinancial industries. Figure 1 reports the distribution across equally weighted employers under primary-employer assignment in 2023, and Figure E1 provides the detailed industry comparison.

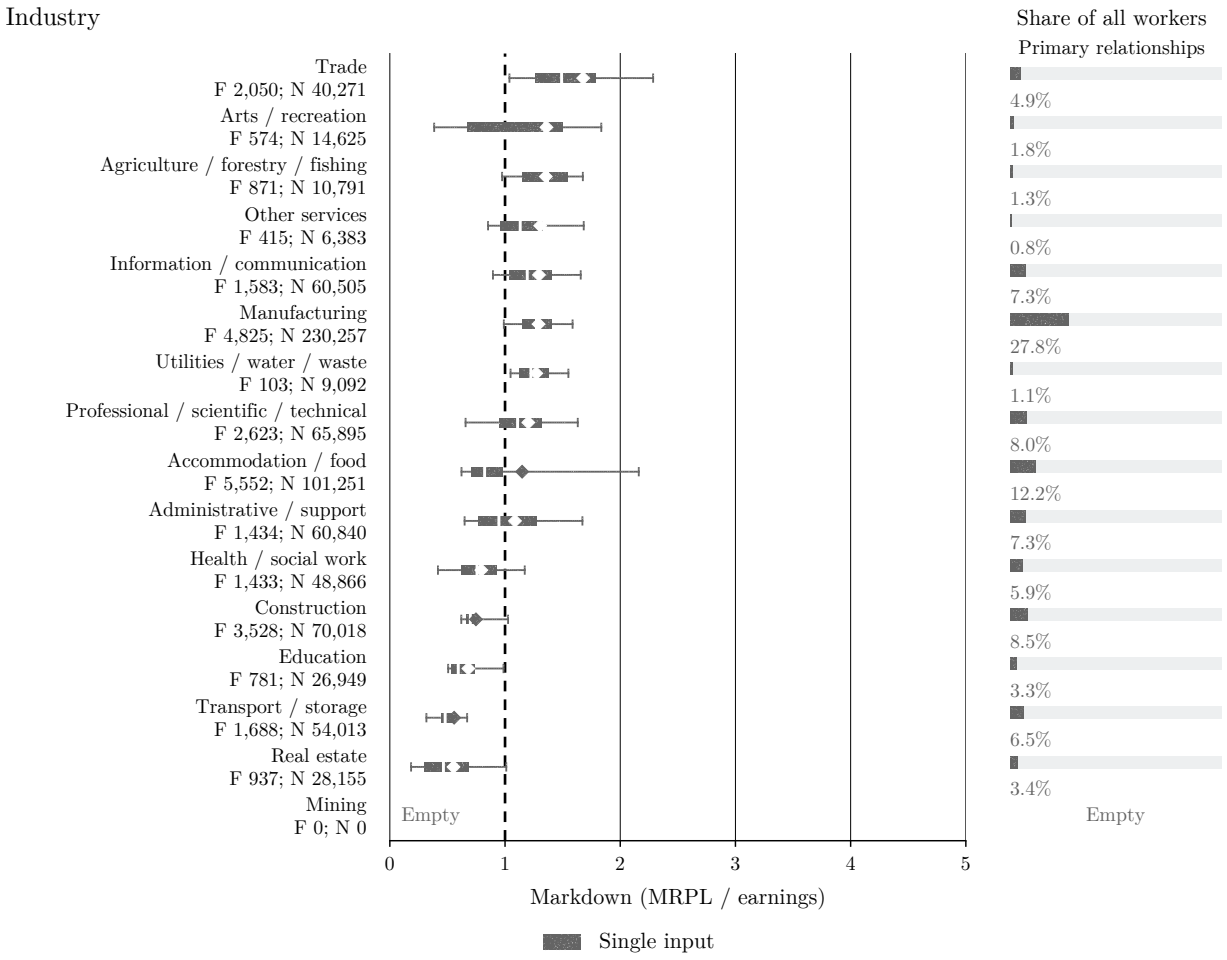
Markdown trends differ across studies. In U.S. manufacturing, Yeh, Macaluso, and Hershbein (2022) find that markdowns fell from the late 1970s to the early 2000s and subsequently rose sharply. Kirov and Traina (2023) instead report approximately stable employment-weighted markdowns before 1990, followed by increases, particularly after 2000. Their estimates concern production workers rather than a single aggregate labor input. In Europe, Mertens (2022) finds that the revenue-weighted aggregate markdown in German manufacturing rose modestly between 1995 and 2014 while the unweighted mean fell, whereas Treuren (2025) finds that the distribution of markdowns in Dutch manufacturing is stable between 2007 and 2018. In their December 2022 working paper, Díez, Malacrino, and Shibata (2022) report an increase in the wage-bill-weighted average markdown in Swedish manufacturing over 2006–2017. The contrasting German trends illustrate why changes in a weighted aggregate need not match changes in the distribution across equally weighted employers.

European evidence also shows substantial dispersion. Mertens (2022) reports a median markdown below one in Germany, while Treuren (2025) finds both wage markdowns and wage markups in Dutch manufacturing, with cross-sectional dispersion far exceeding variation over time. By contrast, Díez, Malacrino, and Shibata (2022) report a manufacturing median of 1.65 when pooling observations across ten European countries over 2000–2017. These estimates provide context for the dispersion on both sides of one in Figure 1. In our 2023 primary-employment sample, the equal-employer mean is 1.10, the median is 0.97, and the interquartile range is 0.70–1.31.

The industry differences in Figures 1 and E1 caution against treating manufacturing estimates as representative of the broader sample. Equally weighted employer distributions also answer a different question from revenue-weighted or wage-bill-weighted aggregates. Input-market assumptions also matter: Treuren (2025) allows for buyer power in materials, whereas our reported levels use a competitive-materials benchmark. A recovered markdown below one does not by itself identify worker bargaining power. The literature therefore provides context for the scale and dispersion of our estimates, rather than a directly comparable numerical benchmark.

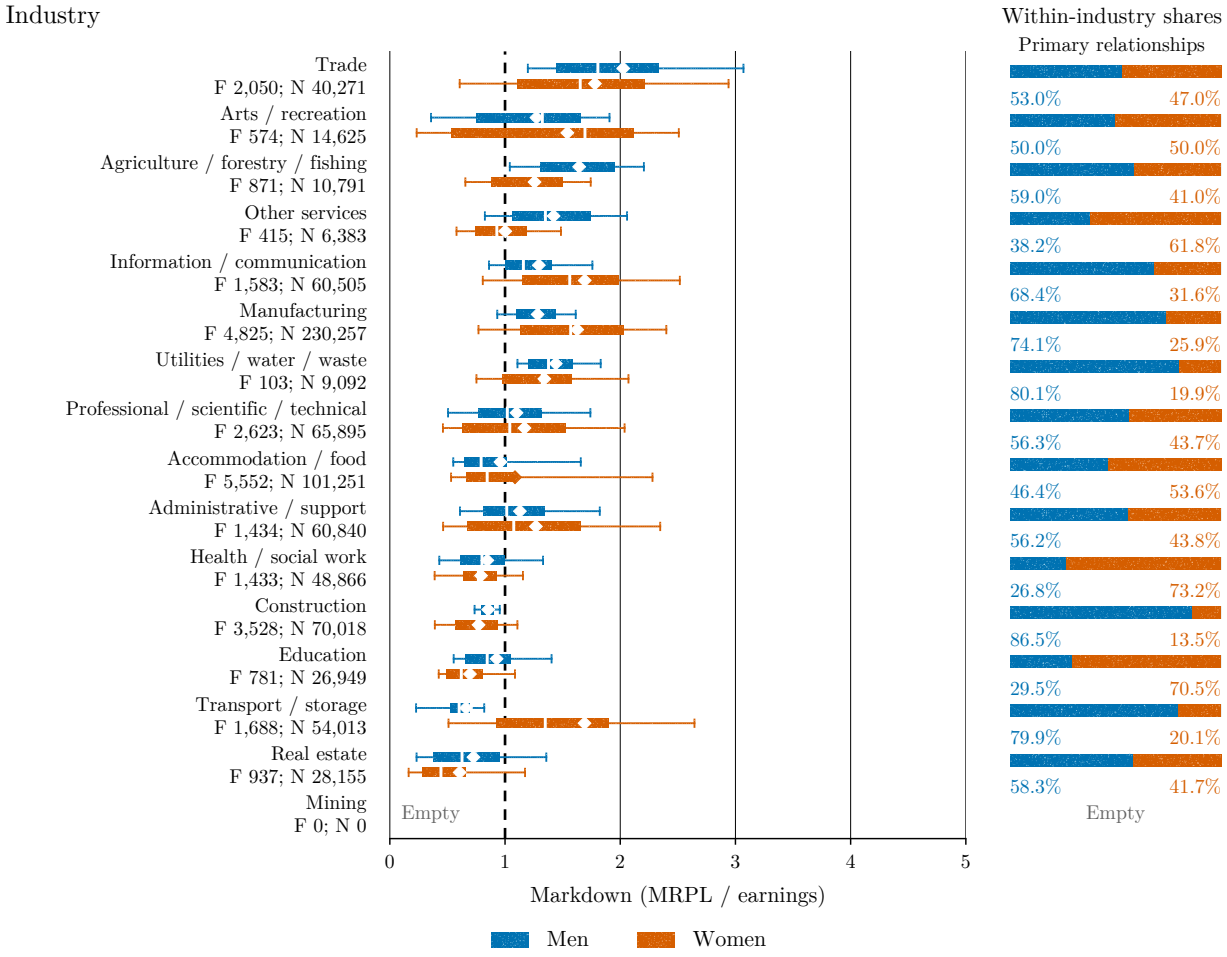
E Industry markdowns

Figure E1: Markdowns across industries: single input specification



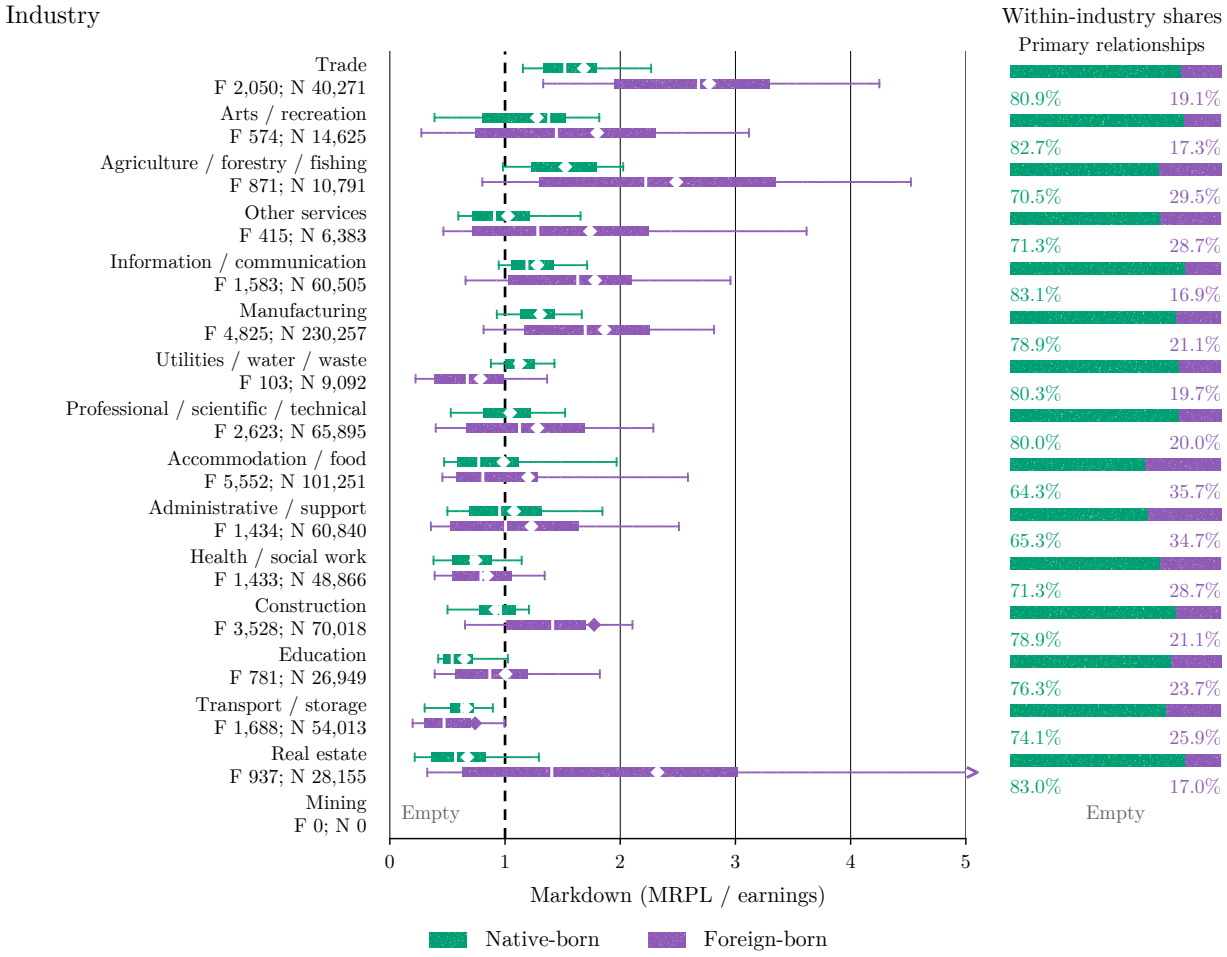
Notes: Primary-employer assignment, 2023; common reporting support across specifications; equal employer weights. Boxes: P25–P75; whiskers: P10–P90; diamonds: arithmetic means; ticks: medians. Common industry order follows single-input means. Worker shares use primary-employer relationships and refer to the national reporting workforce. F: firms; N: worker–employer relationships; mining is empty. Continuation marks indicate values above 5.

Figure E2: Markdowns across industries: gender specification



Notes: Primary-employer assignment, 2023; common reporting support across specifications; equal employer weights. Boxes: P25–P75; whiskers: P10–P90; diamonds: arithmetic means; ticks: medians. Common industry order follows single-input means. Worker shares use primary-employer relationships and refer to all workers in each industry. F: firms; N: worker–employer relationships; mining is empty. Continuation marks indicate values above 5.

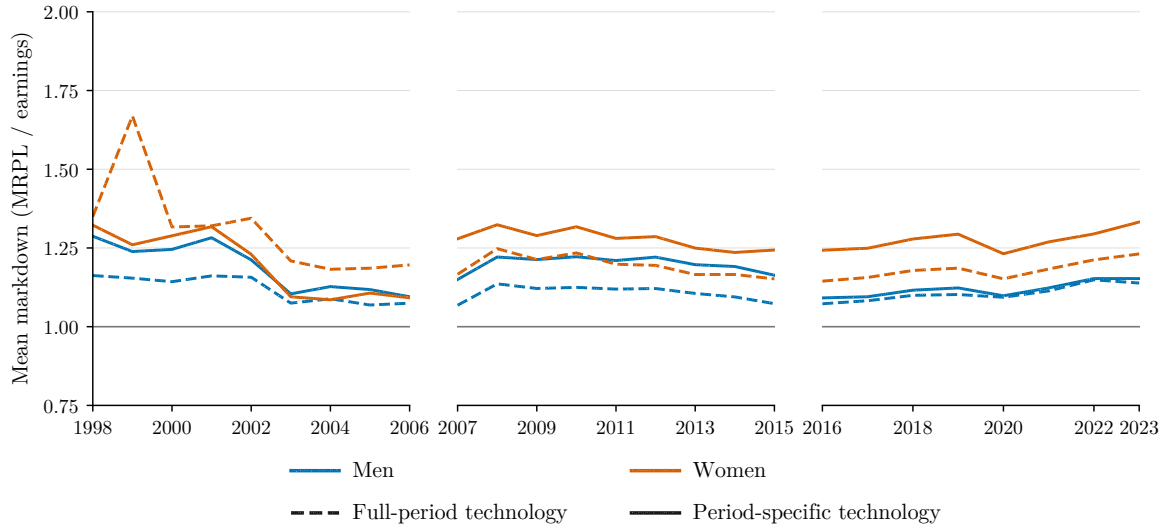
Figure E3: Markdowns across industries: origin specification



Notes: Primary-employer assignment, 2023; common reporting support across specifications; equal employer weights. Boxes: P25–P75; whiskers: P10–P90; diamonds: arithmetic means; ticks: medians. Common industry order follows single-input means. Worker shares use primary-employer relationships and refer to all workers in each industry. Unclassified origin remains in the denominator. F: firms; N: worker–employer relationships; mining is empty. Continuation marks indicate values above 5.

F Gender Markdowns Under Alternative Technologies

Figure F1: Women's and men's markdowns under alternative technologies

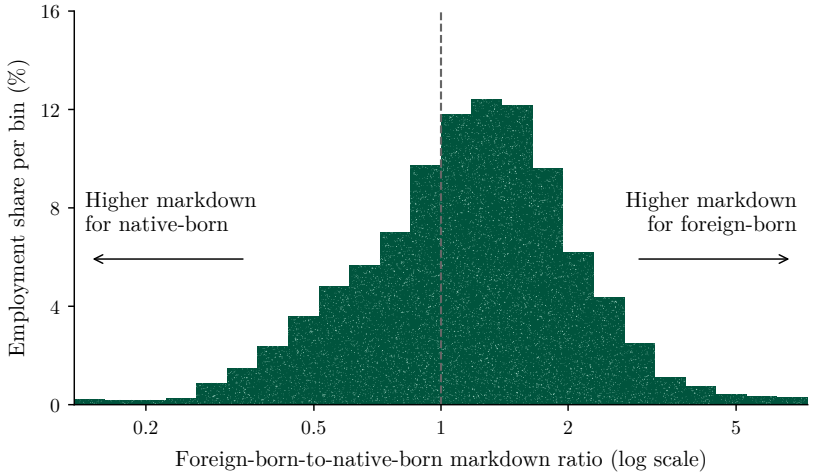


Notes: Primary employment. Arithmetic mean markdowns use each group's employment weights. Dashed lines use full-period technology; solid lines use technology estimated separately for each displayed period. Both fits use identical employers and weights within each year, on narrower support than Figure 5 (26,103 employers in 2023). Employers and weights vary annually. These are group means, including employer allocation, rather than within-employer gaps.

G Origin Markdown Differences in 2023

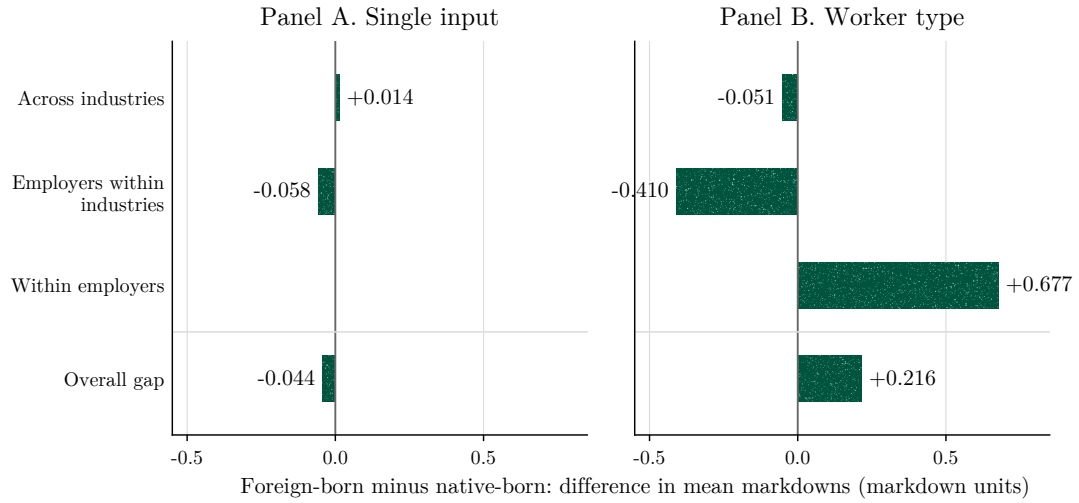
These figures provide the 2023 Origin counterparts to Figures 3 and 4, using the same employer sample, weighting conventions and presentation.

Figure G1: Within-employer foreign-born-to-native-born markdown ratios, 2023



Notes: Primary employment; 28,397 employers weighted by combined foreign-born and native-born employment. Ratios above one indicate higher markdowns for foreign-born workers. Bars show employment shares in equal-width log-ratio bins, labelled in ratios. Tails outside the displayed range are omitted without rescaling.

Figure G2: Origin markdown gaps: single-input and worker-type decompositions, 2023



Notes: Primary employment; the same 28,397 employers in both panels. Gaps are foreign-born-minus-native-born differences in arithmetic mean markdowns. Overall gaps use own-group employment weights; within-employer gaps use pooled employment weights. Industry allocation replaces actual industry shares with pooled shares; employer allocation within industries is the residual. Components sum to the total. The single-input within-employer row is blank by construction.

H Origin Trends and Workforce Characteristics

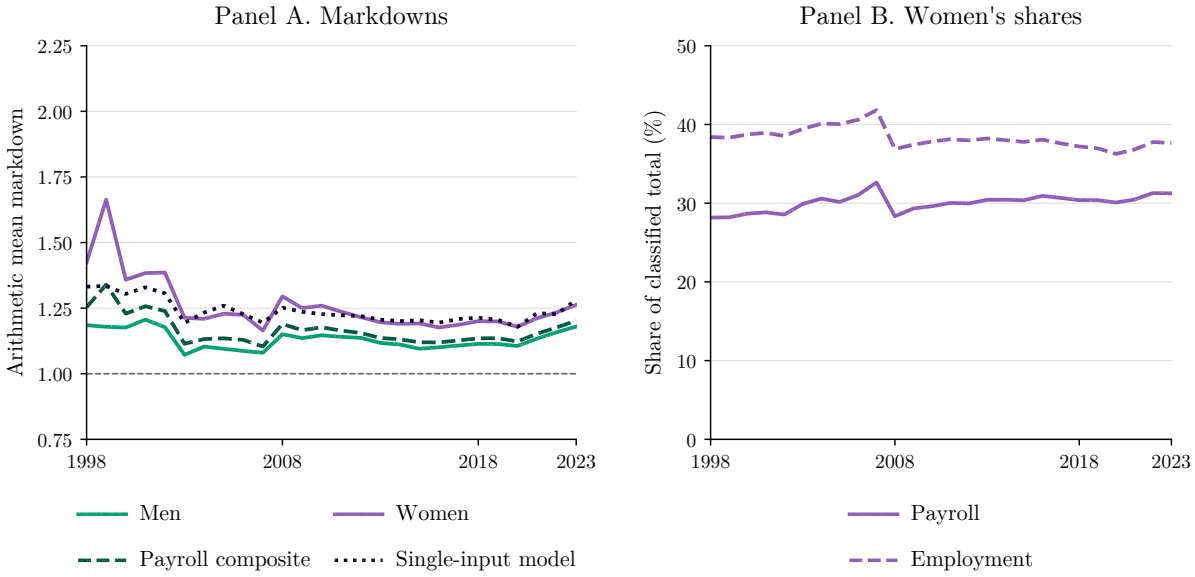
Table H1: Within-employer origin markdown gaps and workforce controls

Controls	1998	2023	Change
Unadjusted	1.207	0.685	-0.522
	(0.100)	(0.209)	
Age and education	1.525	0.732	-0.792
	(0.293)	(0.196)	
Also occupation	1.559	0.773	-0.786
	(0.313)	(0.230)	
Also immigration duration	1.525	0.694	-0.832
	(0.373)	(0.196)	
Employers	9,138	27,676	
Eligible primary employees	299,203	790,359	

Notes: Foreign-born minus native-born differences in markdown units under primary-employer assignment. Employer fixed effects and balanced group weights reproduce pooled eligible-employment weights across employers. All specifications use the same eligible sample within each year; annual samples differ. Age and education must be recorded. Controls are employer–origin means of age and squared age and education, broad occupation and immigration-duration shares. Unknown occupation is retained; occupation coverage varies over time. The final specification standardizes foreign-born workers' years since latest recorded immigration to the eligible 2023 distribution. Employer-clustered standard errors are in parentheses, conditional on the estimated production functions. Change is the 2023 estimate minus the 1998 estimate; no standard error for the change is reported. These descriptive adjustments concern employer-group benchmarks, not individual productivity.

I Gender Heterogeneity and the Single-input Markdown

Figure I1: Gender heterogeneity and the single-input markdown



Notes: Primary employment, 1998–2023. Group curves use own-group employment-weighted arithmetic means. The payroll composite weights the two groups' markdowns by their payroll shares within each employer. Its arithmetic mean and the single-input mean use the same combined male and female employment weights. Panel B reports pooled payroll and employment shares, not weights applied to national group means. Firms and weights vary annually. Differences between separately estimated models do not measure bias or causal effects.