

Olsson, Maria; Nordström Skans, Oskar

Working Paper

The Rules of the Game: Local Wage Bargaining and the Gender Pay Gap

IZA Discussion Papers, No. 17381

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Olsson, Maria; Nordström Skans, Oskar (2024) : The Rules of the Game: Local Wage Bargaining and the Gender Pay Gap, IZA Discussion Papers, No. 17381, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/307205>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION PAPER SERIES

IZA DP No. 17381

**The Rules of the Game:
Local Wage Bargaining and the Gender
Pay Gap**

Maria Olsson
Oskar N. Skans

OCTOBER 2024

DISCUSSION PAPER SERIES

IZA DP No. 17381

The Rules of the Game: Local Wage Bargaining and the Gender Pay Gap

Maria Olsson

BI Norwegian Business School

Oskar Nordström Skans

Uppsala University and IZA

OCTOBER 2024

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

The Rules of the Game: Local Wage Bargaining and the Gender Pay Gap*

We study how local bargaining institutions affect the within-job gender wage gap among Swedish blue collar workers. Collective agreements with varying degrees of local flexibility tend to cover blue-collar workers across different occupations within the same firm. As a consequence, workers performing the same tasks but in different firms are covered by different agreements. We show that the gender pay gap is substantially reduced in jobs covered by collective agreements that guarantee each worker a minimum pay raise every year. Bargaining constraints have a greater impact on gender equality in settings where females are underrepresented. Effects are smaller in more productive firms as these firms can share rents above the contractual minimum with less constraints, even when formal contracts are rigid. Overall, the results suggest that the specifics of local bargaining institutions can play an important role in shaping gender wage disparities among low-paid workers.

JEL Classification: J52, J51, J31, J16

Keywords: collective bargaining, gender equality, unions

Corresponding author:

Oskar Nordström Skans
Department of Economics
Uppsala University
Box 513
751 20 Uppsala
Sweden

E-mail: oskar.nordstrom_skans@nek.uu.se

* We are grateful to Katarina Warg and seminar participants at the Nordic Initiative for Corporate Economics for useful comments. We also thank Lena Hensvik who administrated the data collection for the Uppsala Center for Labor Studies. Skans received financial support from Vetenskapsrådet (2018-04581).

1 Introduction

A substantial body of research in economics and other social sciences has documented the prevalence of gender pay disparities across various settings. Beyond potential causes on the supply and demand sides, the literature has highlighted that wage-setting institutions may influence the nature and magnitude of these gaps (see, e.g., Blau and Kahn, 2017). On the policy front, several stakeholders—including international organizations such as the OECD, labor unions, and gender equality advocates—have promoted collective bargaining as a remedy for excessive pay disparities. Influential studies have demonstrated that gender pay gaps tend to be smaller in settings where wages are determined by collective bargaining (see, e.g., Blau and Kahn, 2003). However, evidence on how the *design* of local bargaining institutions affects the gender pay gap within unionized economies remains limited.¹ Considering that evidence from various settings suggests that women and men behave differently in bargaining situations (see, e.g., Exley et al., 2019; Säve-Söderbergh, 2019; Cortes et al., 2024, and references therein), it is plausible that the contractual structure of local negotiations may influence the gender pay gap, even in contexts where collective agreements are prevalent.

In this paper, we provide evidence on how the degree of flexibility in local wage setting influences the within-job gender wage gap. Our analysis utilizes data from Sweden, a highly unionized and gender-equal economy. In this context, collective agreements vary across firms, but a single agreement typically covers workers across different blue-collar occupations within each firm. Consequently, workers performing the same tasks in different firms are subject to distinct firm-specific bargaining regimes.

We combine register data on workers and firms with a targeted survey in which firms were asked which collective agreements they are covered by. We map these agreements to a classification of contract types provided by the National Mediation Office. Using this information, we distinguish between *rigid* contracts, which guarantee each worker a specified (annual) minimum wage increase, and *flexible* contracts, which do not include any specified individual guarantee. To validate the empirical relevance of this dichotomy, we demonstrate that the association between firm-level productivity (value added per worker) and wages is stronger in firms with more flexible local contracts.

We use a cross-sectional identification strategy, which we argue gives justice to our setting as firms

¹Biasi and Sarsons (2021) is a notable recent exception set in the US context, showing that increased wage flexibility widened the gender pay gap among teachers.

rarely change contracts, and contracts rarely change type according to the used classification. We therefore treat the contract type as a fixed firm-level attribute and all our results will reflect long-run, steady state, patterns.

We estimate the empirical relationship between bargaining flexibility at the local level and the gender wage gap *within job cells* where each job cell is defined as a combination of an occupation and a firm. Our models include occupation fixed effects to ensure that we identify the effects from differences across job cells where the workers perform similar tasks, but in different firms. Identification exploits the fact that the same tasks can be performed at firms with different types of local wage bargaining contracts. Our data also allow us to control for a rich set of firm-level factors capturing potentially important aspects such as firm size, gender composition, industry, and firm productivity. As it turns out, none of these factors are crucial for our results. We further validate that the results are not caused by sorting from the firm or worker side (see below for details).

Our main results indicate that the within-job gender wage gap is larger when local wage bargaining is more flexible. The magnitudes are notably meaningful within our context; the gap is reduced by 2 percentage points—approximately equivalent to half of the average within-job gender wage gap in our sample—when contracts are more rigid. Rigid contracts reduce the gap by about 2.3 percentage points at the mean, 2.1 at the 25th, and 2.8 at 75th percentiles of the within-job wage distribution. Given that the average gap is much smaller at the bottom of the wage distribution, rigid contracts have a greater influence, in relative terms, for lower-paid workers.

To ensure that the results are not driven by *firm*-side selection into flexible rather than rigid contracts, we construct an instrumental variable based on the average rigidity of other occupations within the firm, as predicted by the rigidity of these occupations in other firms. This instrument leverages the notion that firms sort into contracts based on their overall composition of blue-collar occupations. Consequently, the contract covering a specific job often depends on the composition of *other* jobs within the firm. Results using this double leave-out instrument (rigidity of other occupations in other firms) are larger than our baseline specification.²

To assess if the baseline effects are influenced by *worker*-side selection, we use AKM estimates (Abowd et al., 1999) of person effects from *previous* jobs when constructing the wage gap. The findings suggest

²Although our preferred IV estimates are statistically significant, they are more imprecisely estimated than the OLS estimates, as expected.

that worker sorting, if anything, may lead us to underestimate the effects of the contracts. This outcome aligns with the incentives of firms (but not workers); it is relatively cheaper for firms with contracts that compress gender wage gaps to employ otherwise highly paid men.

We further examine the relationship between the effects and (i) the gender composition within firms or occupations, and (ii) firm productivity. The results indicate that the effects are larger in male-dominated occupations and firms, as well as in low-productivity firms. The first finding suggests that a strong female presence can help women protect their wage progression in flexible bargaining settings. The second finding implies that institutional constraints on how wage increases are allocated at the local level are more critical for gender inequality when there are fewer rents to be distributed. This is consistent with the fact that most of our rigid contracts include a *minimum* guaranteed wage increase for all workers, a provision that appears to benefit female workers on average. Since all contracts permit local partners to flexibly distribute wage increases above the mandated minimum, our results suggest that (productive) firms disproportionately allocate these additional flexible wage components toward male workers, even under rigid contracts.³

Our analysis contributes to the extensive and active literature on gender wage disparities within and across firms and jobs (see, e.g., Olivetti and Petrongolo (2016) for a discussion). In particular, Card et al. (2016) demonstrated that part of the gender gap within Portuguese firms arises due to gender differences in the ability to extract firm-level rents during local wage bargaining. For recent studies in a similar vein, see Bruns (2019), Gallen et al. (2019), and Li et al. (2023). However, the literature still lacks evidence on exactly how and why the gender bargaining gap emerges and whether the design of collective agreements can influence it.⁴ A recent exception is Biasi and Sarsons (2021), which shows that female teachers in the US receive lower wages when more flexible bargaining institutions are implemented. Although many aspects of their setting and empirical design differ from ours, the main results are well aligned. From this perspective, our findings suggest that the insights from Biasi and Sarsons (2021) can be generalized to a broader range of institutional contexts. Furthermore, our results—indicating that bargaining flexibility has a more pronounced impact on the gender wage gap in less productive firms—underscore that bargaining

³We also explore the intersection between contracts and part-time work (a highly gendered phenomenon), and estimate separate effects for parents with small children. Our estimates do not indicate that any of these aspects interact with contract types in a meaningful way.

⁴An interesting study using Swedish data is S  ve-S  derbergh (2019), which shows gender differences in bargaining tactics and outcomes among Swedish graduates in Social and Economic Sciences, Computer Science, and Law.

institutions can influence gender wage gaps in ways that are not directly captured by a shift in a proportional bargaining (rent extraction) parameter. Instead, the institutional features of rigid contracts, which guarantee a minimum wage increase for each worker, have a more significant effect on the gender wage gap when rents are limited because even the most rigid contracts typically allow firms the flexibility to allocate wage increases that exceed the required minimum.

At a general level, our paper relates to a very active literature on collective bargaining in Europe. For an overview of bargaining systems, see Bhuller et al. (2022). A subset of studies in this literature analyze how the presence of collective agreements and/or works councils influence the gender wage gap (see, e.g., Antonczyk et al. (2010) and Kiessling et al. (2024) for evidence from Germany). For recent studies of collective bargaining and wage dispersion more broadly, see Dahl et al. (2013), Willén (2021), Card and Cardoso (2022), and Jäger et al. (2022). For previous studies of the gender wage gap in Sweden, see in particular Edin and Richardsson (2002) and Milgrom et al. (2001).⁵

Overall, our results suggest that bargaining institutions can influence the within-job gender wage gap in a highly unionized and gender-equal country like Sweden. The findings indicate that individually guaranteed annual wage increases contribute to reducing gender wage disparities, particularly in male-dominated occupations, among women at the lower end of the wage distribution, and in low-productivity firms where the rents distributed among workers are limited.

The paper is structured as follows. Section 2 outlines the most relevant aspects of the Swedish bargaining framework. Section 3 describes the data and empirical approach, Section 4 presents descriptive evidence. Section 5 shows our results and Section 6 concludes the article.

2 Institutional setting

This section presents an overview of Swedish bargaining institutions with a focus on aspects that matter for our study. Our analysis is focused on blue collar workers in the private sector and the presentation will center on aspects that are relevant for this part of the economy. Unless otherwise noted, the presentation

⁵For other related studies on Swedish data, see Olsson (2024) on rigid contracts and business cycle adjustments, Coglianese et al. (2023) on how monetary shocks transmissions are affected by contract rigidity, Grönqvist et al. (2022) on how institutionally determined wage increases affect teacher exits and Eliasson and Skans (2014) on how collective agreements with targeted wage increases for establishments with low-paid women affect actual wages and retention rates.

draws on information presented in the annual reports of the Swedish National Mediation Offices.⁶

Sweden's industrial relations are characterized by a multi-tiered "pattern bargaining" system. The National Mediation Office counts around 650 different collective agreements. Some agreements are tiny, with a short institutional distance to the covered workers. But most workers are employed by firms that are covered by what is typically referred to as "industry-level" collective agreements (although their coverage only rarely follow standard industry classifications). The agreements are signed between a union and an employer organization. The contract is valid for all workers (i.e. also workers who are not union members) in firms that are members of the signing employer organization. Outside firms can choose to join the contract without being members of the employer organization through "add-on" contracts (*hängavtal*).

The negotiation process generating the content of industry-level collective agreements starts with an initial round of coordinated agreements for all subsets of the manufacturing sector. These coordinated agreements are jointly referred to as *Industriavtalet* (IA). The IA stipulates an average rate of wage increases across a chosen duration (between 1 to 3 years) and this rate must be implemented in all other collective agreements as well. The National Mediation Office act as mediators in all negotiations after the IA and they vigilantly guard the implementation of the stipulated rate of wage increases, as does the central union and employer confederations (from different perspectives, of course). Because of the strong IA norm, variations in exactly *how* (indirectly, for *whom*) the wage increases should be implemented at the local level is the core of the wage part of the collective agreements.

The National Mediation Office classifies the collective agreements based on the degree of local autonomy. The scale moves from procedural agreements which only state how local negotiations should be conducted (known as "no-numbers" agreements) to highly rigid agreements where the wage increases are fully specified in the central contracts. In the mid range, where most blue collar workers are employed, contracts primarily vary in the extent to which the local partners can differentiate the wage increases across individual workers. Below we list the seven classes with the four types of contracts we specify as *Rigid* in *italics*:

1. Procedural agreements (no numbers)
2. Procedural agreements with a fall-back average wage increase if local negotiations fail
3. *Procedural agreements with a fall-back individual wage increase if local negotiations fail*

⁶The reports available at www.mi.se are in Swedish only, unfortunately. See in particular <https://www.mi.se/app/uploads/Kollektivavtal-vilka-tecknar-avtalen-och-hur-ar-loneavtalen-konstruerade.pdf>

4. Agreements with specified average wage increases at the local level, but no individual guarantee
5. *Agreements with specified average wage increase, with an individual guaranteed minimum raise, or an individual guarantee if negotiations fail*
6. *Specified general wage increase (for all), and an additional average wage increase to be distributed locally*
7. *Only a specified general wage increase (for all)*

Because we focus on wage differences within jobs, we include all contracts where there is a clause specifying an *individual guarantee* into the set of contracts we define as *Rigid*.⁷ These individual guarantees can either be formulated as minimum percent increases, as minimum absolute increases (specified in the local currency, SEK), or a combination of both. The amounts stipulated in the clauses are specific to each year and contract and can vary with aspects such as the current wage, experience, and occupation.

Actual wage increases are determined during local negotiations, except in the most rigid of contracts (group 7). Details on how local negotiations should be conducted are outlined in each contract. The formal rules vary, and it is natural to assume that the implementation is dispersed as well. Although some steps may be irrelevant in some settings, the process can be schematically summarized as follows: First, each worker is offered a salary meeting with a supervisor focusing on on-the-job performance since the last local bargaining round (usually one year ago). After the round of individual talks, the employer discusses with the local union representatives regarding, among other things, gender wage disparities.⁸ In the end, the employer compiles a list of proposed wage increases for the workers. The list comes into effect if the union accept it as being in line with the collective agreement. The union can always enforce the fall-back clause in settings where such a clause exist. For our purposes, we consider the contracts with individual guarantees in the fall-back clauses as *Rigid* since firms need to convince the unions to accept any deviations from that clause (presumably by offering higher increases to other workers).

During local negotiations, the partners are bound by a “peace obligation” because they have a binding (central) collective agreement. This means that unions and employers are prevented from engaging in strikes or lock-outs. Disagreements can be resolved through arbitration by the central partners. The agreements do in almost all cases allow for larger local wage increases than what is stipulated, creating a scope for additional upwards flexibility, or “wage drift” as discussed in Hibbs and Locking (1996).

⁷The classification into *Rigid* follows the sector classifications used in Olsson (2024) and Coglianese et al. (2023).

⁸Gender discrimination is illegal and firms are required to monitor gender wage gaps among their employees. Successful lawsuits based on discriminatory wage gaps are rare.

As is evident from this discussion, the core of the agreements are the *wage increases*, not minimum wages or wage levels. This also means that wages of new hires tend to be less constrained by the collective agreements. Most collective agreements do include minimum wages, but they only bite for a trivial share of workers in most agreements, see Forslund et al. (2014). Exceptions where minimum wages do bite include restaurants, retail, and transportation. These contracts will all be considered as rigid according to the classification system described above.

Unions are organized in three confederations representing blue-collar workers (LO), non-academic white collar workers (TCO), and academic white collar workers (SACO). Each firm tends to be covered by a separate contract with (a union within) each of these confederations. Most larger firms employ workers in categories which are covered by all 3 confederations and they will therefore have multiple contracts. Our focus in this paper is on the blue collar agreements as their coverage is easiest to reconstruct in register data, whereas it is more complex to separate between the two groups of white collar workers in the registers. In most firms, all blue collar workers are covered by the same contract.

The process leading up to exactly which blue collar agreement each firm is covered by is not entirely transparent. Contracts are signed between an employer organisation and a union. The firm chooses which employer organisation to join. These organisations are structured according to type of production, e.g. “The Swedish Construction Federation” and the “Green Employers”. The boundaries are, however, not always sharp. There is also a set of parallel employer organisations related to the ownership (e.g. public sector or NGO). Unions on their side also organize themselves in relation to what is produced, but not always along parallel lines with the employers. Our impression is that their coverage is more responsive to production inputs, i.e. the occupations of the blue collar workers. A consequence of the imperfect overlap between employer organisations and unions is a web of contracts for each active combination. Firms that are on the margin can of course change employer organisation, but in most cases they do not. Almost all blue collar unions belong to the same confederation (LO) and therefore do not actively compete with each other for contracts. As a consequence, each firm tends to be covered by the same blue-collar contract across time.

An additional source of differences in flexibility across firms arise because some firms do not sign any collective agreements at all. This is fully possible unless the firm is a member of an employer association. Uncovered firms are, in principle, free to set wages as they wish since Sweden do not have a minimum

wage. In contrast to neighboring countries, collective agreements in Sweden are never made legally binding for firms that do not sign them. However, the vast majority of workers are employed by firms with collective agreements. Most uncovered firms have few employees. Once firms become larger, having a collective agreements tend to be expected by workers and by unions. If a firm refuses to sign a collective agreement after being approached by a union, it may be subjected to a blockade whereby unionized workers in *other* firms and public agencies are not allowed to perform tasks that are related to the relevant firm. Famous cases include refusal to remove garbage, handle financial transactions, or sending number plates to new cars (see the current ongoing conflict regarding TESLA). For completeness, we will include firms without contracts in the “flexible” group.

3 Data and empirical strategy

3.1 Data

Our dataset combines individual-level information on demographics, wages and occupations with firm-level information about collective agreements and firm-level annual accounts. Demographic data are drawn from population wide registers. Wage and occupation data are drawn from the Swedish Structure of Wage Statistics, a survey directed to employers. Firms in the sample frame are required by law to respond to the survey. The survey covers all large firms (500+ employees) and a sample of smaller employers. In total, the data cover around 50 percent of all private sector employees (SCB, 2013). Data are collected separately for white collar and blue collar workers, and we only rely on the file containing blue collar workers in our main analysis.

In 2013, the Uppsala Center for Labor Studies (UCLS) administrated a survey to 7,098 employers that were already covered by the Structure of Wages Statistics in that year. The response rate was 42 percent. A key part of the questionnaire was an open ended question about which collective agreements the firm had signed, if any. Many firms responded that they do not have any agreements, but these firms are typically very small, employing few workers. As expected, many covered firms reported to have signed multiple agreements, but less than 5 percent reported more than one agreement with a counterpart from the blue collar confederation. For firms that did report multiple blue collar contracts, we assign the firm

to the contract that cover most of their workers (also asked in the survey).⁹ For each firm’s blue collar agreement, we add information from the National Mediation Office regarding the type of agreement. We classify them as “rigid” if the agreement includes an individual guaranteed wage increase during local bargaining, see Section 2. The survey responses were matched to longitudinal register data.

Our used sample includes blue collar workers in private sector firms that responded to the survey in 2013 and where the stated agreement could be found on the list of agreements from the National Mediation Office, or where the answer was “no collective agreement”. See Section A.1 for additional details. For the main analysis, we use wage data from 2013. Since the agreements, and the agreement types, tend to be fixed over time, the data should be interpreted as reflecting steady state patterns in a mature setting.

We calculate jobs using 4-digit occupations as defined by the ISCO-88 classification system which has around 389 occupations in total. We use a broad wage measure that accounts for supplementary wage benefits related to the performed task, but not overtime pay. Wages are adjusted for differences in working time. For our baseline sample we include only full-time workers.

We use the firm accounts to calculate labor productivity as the log of value added per worker.¹⁰ We residualize productivity by first regressing it on 3-digit industry indicators to get a measure of productivity that is less confounded by obvious differences in capital intensity across industries (i.e. to avoid comparing a mine to a restaurant).

3.2 The empirical strategy

As described in Section 2, we are analyzing a setting where contracts rarely change in terms of the general structure (at the level we are analyzing them), and where firms tend to be covered by the same contract across time.¹¹ We will therefore exploit a cross-sectional identification strategy by comparing workers across firms and occupations. An advantage of this approach is that we directly estimate how contractual arrangements are related to the gender pay gap after adjustment processes—in terms of the interplay between entry wages and wage increases—have matured.

We focus on the gender pay gap within *jobs* defined as an occupation x firm combination. To make the analysis of the relationship between contracts on the gender wage gap explicit, we create an outcome

⁹See Section A.1 for additional details. Very few firms replied multiple blue-collar contracts that differ in the dimension that we use for our analysis.

¹⁰We use the firm accounts in 2012 which is the last observation of this variable in our data.

¹¹In Appendix A.1.3 and Table A1 we elaborate on the assignment of contracts to types. The results are similar.

variable measuring the gender wage gap within each such *job*. In practice, we define the gap as the mean log difference in wages between men and women within each job. Thus, for firm j and occupation occ , the gender wage gap is computed as:

$$G_{j,occ} = \frac{1}{M_{j,occ}} \sum_{i \in j,occ} \ln wage_i \times m_i - \frac{1}{F_{j,occ}} \sum_{i \in j,occ} \ln wage_i \times f_i \quad (1)$$

where m_i (f_i) is an indicator variable taking the value 1 if individual i is classified as male (female) in national registers and zero otherwise and where $M(F)$ denotes the number of men (women) in each job. In some of our exercises, we replace the raw wages with alternative wage-related measures.

Our empirical strategy relates the gender-wage gap within firms and occupations, defined as in equation (1), to collective agreement rigidity after accounting for occupation fixed effects and firm-level attributes:

$$G_{j,occ} = \alpha + \gamma Rigid_j + \beta X_j + \chi_{occ} + \epsilon_{j,occ}. \quad (2)$$

Rigid takes the value 1 for blue collar jobs covered by a collective agreement that includes a minimal individual guaranteed wage increase in local negotiations. It takes the value zero for other more flexible contracts, including the absence of contracts. Key for our identification are the occupation fixed effects in χ_{occ} . These ensure that we estimate the association between a rigid collective agreement and the gender wage gap between men and women within the same job *relative* to the gender wage gap of other jobs within the same occupation. Note also that the definition of the outcome variable nets out any wage-level differences across firms that are shared across both genders.¹²

The vector X_j contains firm-level covariates that capture other differences across firms that may correlate with both contractual rigidity and the gender wage gap. Because the willingness to sign contracts are correlated with firm size, which could affect career prospects of men and women differently, our baseline specification include 3 indicators for bins of firm-level employment. We also control for the overall share of female workers, and the share of immigrants, within the firm. In robustness exercises, we include similar variables at the job-level, i.e. we control for the number of workers and the share of female workers within the relevant job (occupation x firm). We also estimate alternatives that account for productivity and

¹²The corresponding model at the individual-level with individual log wages as the outcome would include a gender dummy, a job (firm x occupation) fixed effect, gender-specific occupation fixed effects, and gender-specific coefficients on the firm-level attributes we control for. Collapsing the individual data at the gender x job level and then taking the first (gender) difference within jobs to remove the job fixed effects would result in our used model.

the industry of the employing firm. In all specifications, we cluster standard errors at the firm level.

3.2.1 Threats to identification

Our empirical model relates the gender pay gap to the type of collective agreement the firm has signed. Two types of concerns may influence the exact interpretation of this relationship. The first is that the *Rigid* indicator may be correlated with other aspects of the firm that have an independent impact on the gender pay gap. To account for this, we control for first-order suspects such as the size and female-share of the firm and show robustness checks controlling for additional firm-level attributes as discussed above. To further ensure that the results are not due to endogenous self-selection based on idiosyncratic firm specific factors, we provide instrumental variables estimates where we rely on the (historical) occupational structure of the firm as an instrument, for details see Section 5.2.

The second possible concern relates to the sorting of workers. It is possible that the estimated effects in part reflect differences in skills of men and women across firms with more or less rigid contracts. We study the role of sorting by replacing the actual wages by person effects from an AKM-model (see Abowd et al., 1999) estimated using only data from *previous jobs* and then compute a gender wage gap in terms of these estimated fixed person effects as well as in terms of residual wages, for details see Section 5.3.

4 Data description

4.1 Descriptive statistics

Our used data are described in Table 1. Each of the 1142 jobs in our data is treated as one observation. These jobs are found across 519 firms, each of which has a reported collective agreement and at least one man and women within each job. Additionally, we use only jobs in occupations that also exist in at least one more firm with an identified collective agreement. Section A.1 and Table A2 further describes the sample restrictions.¹³ 85 percent of jobs have Rigid contracts. Sampled firms are, as expected due to the sampling frame, fairly large. The average firm has 4-500 workers. But the data do also cover a set of smaller firms, the median firm has around 100 workers. The share of women is fairly low (about 1/4), because we only include blue collar workers in the private sector.

¹³36 percent of all jobs in the wage statistics in 2013 have at least one man and one women, and 26 percent of those are identified with a contract name.

Overall, the differences between observations in the Rigid and Flexible contracts are modest. Firms and jobs are larger in the Rigid contracts whereas the flexible firms tend to be more productive. Turning to the wages, the table show that men (women) earn on average 5 (3) percent higher wages in flexible than in rigid firms. The gender wage gap is 5 (3) percent in flexible (rigid) firms. In Appendix Figure A1, we show that the gender pay gap has a similar distribution in our used data as in the overall data. We further show the distribution of Rigid contract intensities across occupations in Figure A2.

Table 1: Descriptive statistics

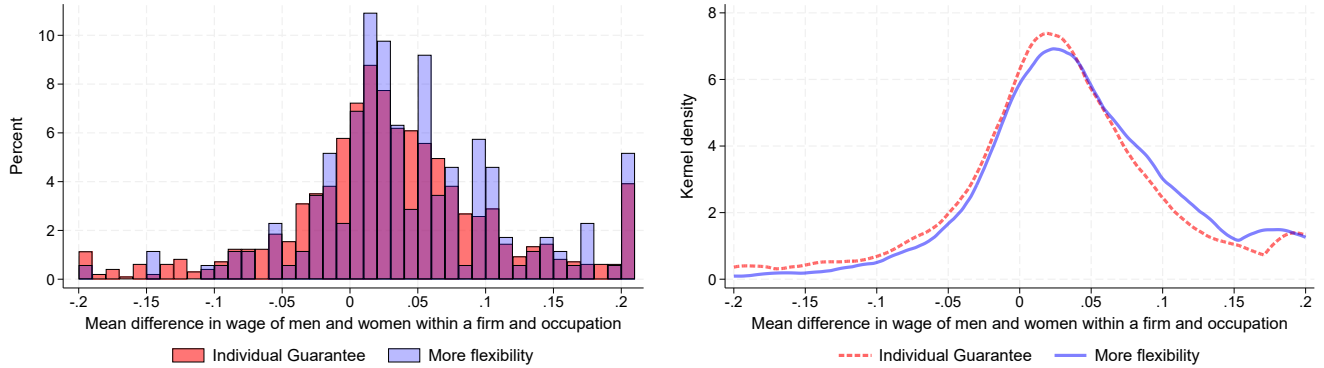
	Flexible		Rigid	
	Mean	Sd.	Mean	Sd.
Firm size	493	(713)	536	(1376)
Female share	0.27	(0.22)	0.25	(0.20)
Immigrant share	0.16	(0.17)	0.17	(0.14)
Productivity 5-quintiles	3.11	(1.73)	2.99	(1.43)
Part-time share	0.09	(0.17)	0.09	(0.17)
Jobb size	57.54	(113.72)	85.31	(251.82)
Jobb female share	0.31	(0.23)	0.30	(0.24)
Male occupation	0.79	(0.41)	0.82	(0.38)
<i>Av. wage men</i>	10.21	(0.19)	10.16	(0.13)
<i>P25 wage men</i>	10.16	(0.19)	10.10	(0.14)
<i>P75 wage men</i>	10.26	(0.20)	10.21	(0.14)
<i>Av. wage women</i>	10.16	(0.18)	10.13	(0.13)
<i>P25 wage women</i>	10.12	(0.18)	10.08	(0.13)
<i>P75 wage women</i>	10.21	(0.19)	10.17	(0.14)
Number of jobs	174		968	
Number of firms	89		430	
Number of occupations	74		106	

Notes: Mean and standard deviation within job-cells by contract type. Size, female share, immigrant share and part-time share is measured using the linked employer-employee data in the wage structure statistics. Productivity is 5-quintiles of log(value added per FTE worker) residualized within 3-digit industries. Male occupation is an indicator for at least 50 percent men in the occupation (across all firms).

As a prelude to our main regression results, we show the dispersion of gender pay gaps across jobs by contract type in Figure 1. The graphs show, in histogram and kernel form, that more of the jobs with really large male pay premia are found in firms with more flexible agreements. It is useful to study the tails, although they will be influenced by random factors, as the coverage of the two types of contracts are strikingly different at the far left and the far right. Jobs where men have more than 10 percent higher wages than women are more than twice as likely to be in flexible contract firms as (the fewer) jobs where

women earn 10 percent more than men (19 vs. 8 percent flexible).

Figure 1: Distribution of gender pay-gaps across contract types



Notes: Figures display the mean difference in wages for men and women within jobs by contract types. Jobs covered by an individually guaranteed wage growth are marked in orange, and jobs without in blue. Left histogram in percent, right kernel density. Year 2013. Winzorized at ± 0.2 .

4.2 Contracts and rent sharing

Our primary focus is on how rigid bargaining protocols at the local level affect wage differences across men and women within jobs. But rigid wage-setting practices may also affect the pass-through of economic shocks to workers' wages, see e.g. Olsson (2024) for a study of how the pass-through from the Great Recession varied across agreement types using data similar to ours.

To illustrate how the pass-through from firm productivity to worker wages differs across contract types within and across occupations, we run a set of cross-sectional regressions where we explain firm-level wages by our measure of firm productivity.¹⁴ As dependent variables we use residualized wages. The first two columns of Table 2 use wages that are residualized at the 3-digit industry level (i.e. similar to how we residualize productivity). The following columns instead use wages that are residualized by occupation dummies, and industry-occupation combinations respectively. The table shows separate results for flexible and rigid contract firms. As is evident, all estimated associations between firm productivity and wages are less than one-third as large in firms with rigid contracts. We return to the interaction between productivity, contracts, and gender wage gaps in Section 5.4.

¹⁴As in all our analyzes, we use productivity measure which is residualized from 3-digit industry dummies.

Table 2: Residual rent sharing by rigidity

	(1)	(2)	(3)	(4)	(5)	(6)
	Res. wage Flexible	Res. wage Rigid	Res. wage Flexible	Res. wage Rigid	Res. wage Flexible	Res. wage Rigid
Res. prod	0.121 (0.026)***	0.0356 (0.008)***	0.0948 (0.022)***	0.0332 (0.011)***	0.0960 (0.021)***	0.0295 (0.008)***
Constant	-0.0444 (0.010)***	-0.0232 (0.004)***	-0.0699 (0.010)***	-0.0360 (0.004)***	-0.0429 (0.009)***	-0.0240 (0.004)***
Observations	214	555	214	555	214	555
Residualized	Industry	Industry	Occ.	Occ.	Industry Occ.	Industry Occ.

Notes: Outcome variable is residual wage. Columns 1-2 residualizes the wage at the 3-digit industry level, in columns 3-4 at the 3-digit occupation level, and in 5-6 by industry and occupation. Productivity is the residual value added per worker of a firm relative to their 3-digit industry. Uneven columns are for flexible and even columns for rigid. Observation level firm.

5 Results

5.1 Contracts and gender pay gaps

Our main results based on Equation 2 are presented in Table 3. The table presents the estimated effect of a rigid collective agreement on the gender wage gap within jobs. For completeness, we first estimate the raw association without any controls, and as expected from the histogram presented above, this association is negative and statistically significant (p-value 0.054). We then move to our preferred, baseline, specification which compares the gender gap within occupations across firms after controlling for firm size, the overall gender composition, and the immigration share at the firm. The magnitudes suggest that a rigid wage contract is associated with 2 pp smaller gender gap from a grand mean of 3.6 pp. We then add controls for the size and gender composition in the job-cell, with similar estimates. In the last column, we add 1-digit industry controls to our baseline specification.

In the Appendix Table A1, we provide further robustness results with sample variations and variations in the estimated model. Removing *Verkstadsavtalet* an agreement that do change type across years leads to marginally larger estimates. This is consistent with the notion that rigid contracts matter more if they are in place for a longer period of time. Focusing only on workers with at least 2 years of tenure (stayers) does not change the results.¹⁵

¹⁵Further tests include controlling for productivity, using more detailed controls, changing our occupational controls to a less detailed level; with a very limited impact on the estimates. Controlling for the existence of any collective agreement leads to

Table 3: Union contracts and mean wage differences

	(1)	(2)	(3)	(4)
	Raw	Baseline	Cell controls	Industry controls
Rigid	-0.0137 (0.007)*	-0.0227 (0.010)**	-0.0218 (0.011)**	-0.0240 (0.012)**
Observations	1142	1142	1142	1142
Mean dep.	0.0354	0.0354	0.0354	0.0354
Controls:				
Occupation FE		✓	✓	✓
Firm Size		✓	✓	✓
Firm Female Share		✓	✓	✓
Firm Immigrant Share		✓	✓	✓
Cell Size			✓	
Cell Female Share			✓	
1-digit Industry				✓

Notes: Regressions follow Equation 2. The outcome variable is the difference in mean wages of men and women within the same firm and occupation. Observation level is the job. Controls are 3-bins and included as indicator variables. Occupation fixed effects are at the 4-digit level. Std. errors clustered at the firm in parentheses. * $p < .10$. ** $p < .05$. *** $p < .01$.

Due to our set-up where we use the gender wage gap as the outcome, it is straightforward to modify the model to analyse the impact on different moments related to the gender-specific wage distribution within jobs. In Table 4, we show results from models that use the median, the 25th and 75th percentile in the job-and-gender specific wage distribution as the outcomes (instead of the means as in the main table).¹⁶ The effects are fairly similar across the distribution. As a consequence, we do not find any impact on the gender differences in wage dispersion (4th column). On the other hand, the impact tend to be larger in *relative* terms for the lower part of the wage distribution as the average gap is much smaller at the bottom. As shown in the table, the mean gap is 2.1 percent at the 25th percentile and 4.6 percent at the 75th percentile. In the last column of Table 4 we further investigate how contract rigidity differentially impact the wages of men and women by comparing the wage difference between higher paid men relative to lower paid women within each job. We find that jobs with more rigid contracts have about 20 percent smaller pay difference (3 pp from a mean of 13 pp).

marginally lower estimates with a p-value of 0.052.

¹⁶For P25 and P75, we require at least 1 man and 2 women, or 2 men and 1 women, within each cell. For cell-by-gender observations with fewer than 5 observations, we replace P25 with the lowest paid worker and P75 with the highest paid worker.

Table 4: Union contracts and wage differences across the distribution

	(1) P50	(2) P25	(3) P75	(4) P75/P25	(5) P75 _M /P25 _F
Rigid	-0.0221 (0.011)**	-0.0206 (0.010)**	-0.0280 (0.013)**	-0.00735 (0.013)	-0.0302 (0.013)**
Observations	1142	1085	1085	1085	1085
Mean dep.	0.0342	0.0214	0.0456	0.0242	0.133

Notes: Regressions follow the baseline specification in Equation 2 with occupation fixed effects and controls for firm size, female share, and immigrant share. The outcome variable in columns 1-4 is the difference between men and women at a given point in the distribution, i.e. column 4 is $(\log wage_{men,P75} - \log wage_{men,P25}) - (\log wage_{female,P75} - \log wage_{female,P25})$. The outcome variable in column 5 is the difference in log wages between men at the 75th percentile and women at the 25th percentile. Std. errors clustered at the firm in parentheses. * $p < .10$. ** $p < .05$. *** $p < .01$.

5.2 A validation using instrumental variables

A potential limitation of our empirical approach is that firms may self-select into contracts due to some unobservable factors that independently affect the gender wage gap. To validate the robustness of our main results, we exploit the composition of occupations within each firm to construct an instrumental variable. We compute a “double leave-out” instrument as follows: First, we calculate a “Rigidity Score” for each occupation and firm by averaging our *Rigid*-indicator across all other firms within that occupation. Second, we calculate, separately for each job, the weighted average of the Rigidity Score for all *other occupations* in the firm, excluding the focal job. This approach yields a job-specific measure of the probability that the contract is rigid, based on the composition of other occupations within the same firm. Each step of the instrument construction is weighted by the total employment in each job.¹⁷

The main intuition behind the instrument is that unions and employer organizations sort across firms based on the aggregate composition of occupations within each firm. The identifying assumption is that the gender wage gap within a given job (occupation by firm) is not influenced by factors correlated with the composition of *other occupations* within the firm.

Results where we use the instrumental to estimate the baseline version of Equation 2 are presented in Table 5. The table reports results from two versions of the instrument: one using the *Concurrent* composition of occupations in the firm and the other using the *Historical* composition based on the occupational

¹⁷Thus other occupations in the firm, weighted by the employment within each job, and rigidity in other firms weighted by their employment.

structure in the first year (since 2005) that the firm appears in our data. We prefer the historical version due to its stronger argument for exogeneity, but we present estimates from both versions for completeness.

Overall, the IV results support the conclusions from our baseline model. As expected, the IV estimates are less precise than the OLS baseline. Our preferred estimate relying on the historical composition remains significant at the 10 percent level (p-value 0.058) whereas the version using the current composition is not. Reassuringly, however, both point estimates are consistent with, and even larger than, our baseline specification.

Table 5: Instrumental variable

	(1)	(2)	(3)	(4)
	IV-estimate		First stage	
	Historical	Concurrent	Historical	Concurrent
Rigid	-0.0349 (0.018)*	-0.0332 (0.022)	0.9353 (0.1522)***	0.9516 (0.1629)***
Observations	1142	1040		
	First stage F-statistic			
	Cragg-Donald		233.9	290.9
	Kleibergen-Papp		37.77	34.11

Notes: IV-estimate in columns 1–2, and first stage in columns 3–4. Column 1 uses the historical composition of jobs within the firm (2005-2013) to calculate the instrument. Column 2 uses the contemporaneous job-structure. For both columns, we weight jobs by the size of the cell before computing the instrument. The first stage in column 3 corresponds to column 1, and 4 is the first stage for column 2. All regressions have occupation fixed effects and 3 bins for firm size, female share, and immigrant share at the firm. Std. errors clustered at the firm in parentheses. * $p < .10$. ** $p < .05$. *** $p < .01$.

5.3 Sorting

Our data shows that gender wage differences are smaller when jobs are performed within firms that have more rigid contracts. But a key question for interpretation is if this pattern reflects differential sorting of men and women into these firms or differences in wage outcomes for workers with similar wage potentials. In essence, high wage women may have comparatively stronger incentives than high-wage men to work in firms with rigid contracts. Firms' incentives may point in the opposite direction, a compressed gender wage gap makes it cheaper for rigid firms to hire high-wage men instead of high-wage women.

To analyze this issue we estimate if firms with more rigid contracts attract women with a higher earnings capacity (relative to men) than firms with more flexible contractual arrangements. In practice we

estimate an AKM model, as in the vast literature arising from Abowd et al. (1999), using data from *previous jobs* only and then replace the actual wage by these pre-estimated person effects. A key advantage with this approach is that it allows us to study sorting using a metric that is comparable to the actual wage effects and we can therefore directly decompose the effects into worker sorting vs. effects conditional on sorting. Formally, we first estimate:

$$\ln wage_{it} = \theta_i + \psi_{j(i,t)} + \beta Z_{it} + u_{it} \quad (3)$$

using all employment spells (firms by worker) except those in our main used data. The fixed effects θ_i are the person effects of interest and $\psi_{j(i,t)}$ capture the shared firm-level component of wages. The vector Z includes time dummies and the quadratic and cube (the linear term is captured by time effects) of $Age - 45$, as is the convention. The used data span 1997-2012. In the next step, we recalculate the gap as in Equation 1 but instead of the actual wage, we use $\hat{\theta}_i$, as well as the residual $w_i - \hat{\theta}_i$.¹⁸

We then estimate our main model using these transformed outcome variables. It should be noted that the sample is smaller in this exercise as we do not observe all workers in other jobs, partly due to the sampled nature of our data, partly due to limited mobility. Furthermore, we can only include jobs where we could estimate person effects from previous jobs for both men and women, in the same current job. The results, presented in Table 6, suggest that sorting on earnings capacities is not driving our main results. As a prelude, the first column shows that the average person effect of all workers (i.e. irrespective of gender) are very similar between rigid and flexible jobs. The following columns focus on gender differences. The point estimates suggest that the gender gap in terms of pre-estimated earnings capacities are larger in more rigid contracts, which is in the firms' interest as discussed above. As a consequence, the gap in terms of residual wages are larger than the impact on the raw gap. Although it should be acknowledged that the statistical power is insufficient for precise statements about the nature of sorting (column 2), we find it reassuring that the main results remain robust (even becoming larger) when we account for sorting on portable earnings capacities.

As a complementary analysis, we have also residualized the wages using a gender-specific Mincer-

¹⁸We estimate the model jointly for men and women even though the firm-effects may be different for the two genders. This simplification has the advantage of producing estimated person effects that are normalized according to the same baseline and therefore directly comparable across men and women.

Table 6: AKM individual effects

<u>Panel A: AKM-model</u>				
	(1) $\hat{\theta}_i$	(2) Diff. $\hat{\theta}_i$	(3) Diff. $(wage - \hat{\theta}_i)$	(4) Diff. $wage$
Rigid	-0.00555 (0.011)	0.0324 (0.021)	-0.0470 (0.022)**	-0.0197 (0.011)*
Observations	805	805	805	805
Mean dep.	-0.168	0.0438	-0.0138	0.0317
<u>Panel B: Mincer-model</u>				
	(1) $\hat{\delta}_i$	(2) Diff. $\hat{\delta}_i$	(3) \hat{v}_i	(4) Diff. $wage$
Rigid	0.00228 (0.002)	-0.00369 (0.003)	-0.0190 (0.011)*	-0.0227 (0.010)**
Observations	3168	1142	1142	1142
Mean dep.	10.11	0.000971	0.0344	0.0354

Notes: In panel A, the outcome variables comes from the AKM model Equation 3. In column 1, the outcome variable is the average AKM individual fixed effect in the job-cell. In columns 2, the outcome variable is the mean difference in AKM individual fixed effects between men and women in the same job-cell. In column 3, the outcome variable is the mean gender difference in the wage net of individual fixed effects at the job-cell level. Column 4 reproduces the baseline model from Table 3 in this sample. The outcome variables in panel B comes from a Mincer-regression Equation 4. Column 1 uses the average predicted value, column 2 the gender difference in that predicted value, and column 3 the gender difference in the residual. All regressions have occupation fixed effects and 3 bins for firm size, female share, and immigrant share at the firm. Std. errors clustered at the firm in parentheses. * $p < .10$. ** $p < .05$. *** $p < .01$.

style regression g :

$$\ln wage_{itg} = \delta_g X_{ig} + v_{itg}, \quad (4)$$

where X includes demographic indicators (education, age, age square, number of children).¹⁹ We allow for differential demographic effects by gender by estimating the model separately for the two genders. From $\hat{\delta}_g$, we recover the predicted wage for each individual as well as the residual wage \hat{v}_{itg} . We then recalculate the gender-gap for the predicted wage and residualized wage using Equation 1 analogous to what we did with the AKM-predictions. The results from this exercise is presented in panel B of Table 6. In this case, we see a mild gender difference in terms of predicted wages (column 2), but the point estimate is small (less than one fifth of the main effect). As a consequence, the residualized wage difference (column

¹⁹We include one indicator for each 10-year age interval, 3 separate indicators for number of young children, teenagers, and total number of children, respectively. Education is recorded at the 3-digit level (SUNnivå) and we include indicators for all 49 categories.

3) is close to our baseline estimate.

We have also analysed three other aspects of sorting. We have explored if women are more (or less) likely to be sorted into rigid jobs. We find no evidence of such sorting. We have also studied if workers who want to work part-time, or workers with children, are sorted into jobs with more rigid wage contracts. We find no statistically significant differences. Detailed results are presented in Appendix Table A3.

5.4 Heterogeneity

We focus on two crucial dimensions of heterogeneity that we interpret as being of first-order importance. The first dimension relates to gender sorting. We categorize our data into male- and female-dominated occupations (based on national data) and estimate a modified version of Equation 2 that allows for differential effects of rigidity for each group.²⁰ The results, presented in the first column of Table 7, indicate that the main effect is driven by jobs in occupations that are predominantly performed by males. These occupations also exhibit a large average wage gap of 4 percent, compared to a 1.4 percent gap in female-dominated occupations. A similar pattern emerges when firms are divided into female- versus male-dominated firms. The effect of rigid contracts is somewhat larger in firms where the share of men is above the median (column 2), and more pronounced in jobs at firms where the share of men exceeds the occupation average (columns 3). Overall, these results suggest that the gender wage gap is particularly sensitive to contract type when female workers constitute a more marginal group.

The second dimension of interest is the relationship to productivity. On one hand, productive firms with flexible contracts have more rents to share, suggesting that contractual flexibility may primarily inflate the gender wage gap in more productive firms. On the other hand, contractual guarantees may be more binding in firms with limited rents to share, implying that contract type should have a greater impact in the least productive firms. The results presented in Table 8 suggest that the effects are driven by low-productivity firms that, as shown in the descriptive section, tend to pay lower wages on average. We also analyze how this pattern manifests across different moments of the within-job-and-gender wage distribution. The estimated average gender wage gap is 4 percent in firms with productivity below the median, but only 1 percent (and statistically insignificant) in firms with productivity above the median. These differences are primarily driven by the wage gap at the top of the wage distribution.

²⁰We estimate the model jointly to have the same coefficients on the controls and fixed effects across groups.

Table 7: Wage gaps by rigidity and gender dominated occupations

	(1)	(2)	(3)
Gender dominated:	Occupation	Firm	Firm to Occupation
Rigid \times Male dominated	-0.0249 (0.011)**	-0.0343 (0.014)**	-0.0393 (0.012)***
Rigid \times Female dominated	-0.0114 (0.023)	-0.0100 (0.014)	-0.00850 (0.014)
Observations	1142	1142	1142

Notes: Regressions follow a modified version of Equation 2 where we interact Rigid with groups that portion the data based on the gender composition. Column 1 is for occupations with above/below 50 percent women (full-time). Column 2 for firms with above/below median female share at the firm. Column 3 is for firms where the share of men is above/below the occupation average. In all columns, the outcome variable is the difference in mean wages of men and women within the same firm and occupation. All regressions have occupation fixed effects and 3 bins for firm size, female share, and immigrant share at the firm. Std. errors clustered at the firm in parentheses. * $p < .10$. ** $p < .05$. *** $p < .01$.

These patterns should be interpreted in relation to the institutional setting described in Section 2. Firms are less constrained when they distribute *top-up* wage increases as long as they provide the guaranteed minimum as stipulated in the contracts. In some contracts they can also avoid paying the minimum guarantees if the unions accept the full package, which should be more likely to occur if the firm pays above minimum wage increases. As a result, productive firms with large rents to share are less constrained by rigid contracts.

6 Conclusions

Our paper analyzes how guaranteed wage increases affect the gender wage gaps within jobs among Swedish blue collar workers. We use individual-level micro data to construct a measure of the wage gap between male and female workers who perform the same tasks within the same firm. We then relate these within-job wage gaps to the type of bargaining contract the firm is covered by. Our regression models control for occupation fixed effects and firm-level characteristics. The results show that the wage gap is smaller in firms with rigid contracts which stipulate a yearly minimum wage raise for each individual worker, as compared to similar jobs in firms with more flexible local bargaining procedures. We validate the results using instrumental variables techniques that rely on a (double) leave-out strategy where rigid contracts are predicted from the contract types (in other firms) of other occupations in the firm. We analyse

Table 8: Productivity and wage gaps by rigidity

	(1)	(2)	(3)
Gap measure:	Mean	P25	P75
Rigid×Low productivity	-0.0450 (0.013)***	-0.0304 (0.014)**	-0.0523 (0.016)***
Rigid×High productivity	-0.0104 (0.012)	-0.0120 (0.012)	-0.0106 (0.013)
Observations	1105	1044	1044

Notes: Regressions follow a modified version of Equation 2 where we interact Rigid with groups that portion the data based on firm productivity. High productivity firms have above average productivity, and low productivity firms have below average productivity (in the final sample). The outcome variable is the difference in wages of men and women within the same firm and occupation at different points in the distribution. In columns 1 this is measured as the difference at the mean, columns 2 difference at P25, columns 3 difference at P75. Quintiles are calculated by gender within each firm-occupation. Productivity is the residual value added per worker of a firm relative to their 3-digit industry. Firms where information about productivity is unobserved are excluded. All regressions have occupation fixed effects and 3 bins for firm size, female share, and immigrant share at the firm. Std. errors clustered at the firm in parentheses. * $p < .10$. ** $p < .05$. *** $p < .01$.

sorting using AKM-person effects estimated from previous jobs and show that gender-specific sorting on person effects, if anything, appears to dilute the impact of rigid contracts on the gender wage gap.

Further results indicate that rigid contracts have a greater impact on the gender wage gap in male-dominated occupations and firms, suggesting that contractual rigidities reduce the gender pay gap more in settings where female representation is lower. In addition, rigid contracts matter more in low-productive firms where rents are scarce, consistent with the fact that all contracts, even the rigid ones, allow (productive) firms to allocate wage increases above the mandated minimum with less constraints.

Overall, our paper shows that bargaining institutions can play an important role in shaping the anatomy of the gender wage gap in highly unionized settings. In a broader perspective, our results highlight that the group-specific impact of collective bargaining institutions may not always be accurately described by models with proportional bargaining parameters. Instead, our results show that contractual designs may have very different effects on marginalized groups' ability to extract rents in cases when rents are small vs. cases where firms choose to share rents above the contractual minimum.

References

Abowd, John, Francis Kramarz, and David Margolis (1999) "High Wage Workers and High Wage Firms," *Econometrica*, 67 (2), 251–333.

- Antonczyk, Dirk, Bernd Fitzenberger, and Katrin Sommerfeld (2010) "Rising wage inequality, the decline of collective bargaining, and the gender wage gap," *Labour Economics*, 17 (5), 835–847.
- Bhuller, Manudeep, Karl Ove Moene, Magne Mogstad, and Ola L. Vestad (2022) "Facts and Fantasies about Wage Setting and Collective Bargaining," *Journal of Economic Perspectives*, 36 (4), 29–52, [10.1257/jep.36.4.29](https://doi.org/10.1257/jep.36.4.29).
- Biasi, Barbara and Heather Sarsons (2021) "Flexible Wages, Bargaining, and the Gender Gap*," *The Quarterly Journal of Economics*, 137 (1), 215–266.
- Blau, Francine D. and Lawrence M. Kahn (2003) "Understanding Differences in the Gender Pay Gap," *Journal of Labor Economics*, 21 (1), 106–144.
- (2017) "The Gender Wage Gap: Extent, Trends, and Explanations," *Journal of Economic Literature*, 55 (3), 789–865.
- Bruns, Benjamin (2019) "Changes in Workplace Heterogeneity and How They Widen the Gender Wage Gap," *American Economic Journal: Applied Economics*, 11 (2), 74–113.
- Card, David E. and Ana Rute Cardoso (2022) "Wage Flexibility Under Sectoral Bargaining," *The Journal of the European Economic Association*, 20 (5), 2013–2061.
- Card, David E., Ana Rute Cardoso, and Patrick Kline (2016) "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *The Quarterly Journal of Economics*, 131 (2), 633–686.
- Coglianesi, John, Maria Olsson, and Christina Patterson (2023) "Monetary Policy and the Labor Market: A Quasi-Experiment in Sweden," *mimeo*.
- Cortes, Patricia, Jacob French, Jessica Pan, and Basit Zafar (2024) "Gender Differences in Negotiations and Labor Market Outcomes: Evidence from an Information Intervention with College Students," *mimeo*.
- Dahl, Christian M., Daniel le Maire, and Jakob R. Munch (2013) "Wage Dispersion and Decentralization of Wage Bargaining," *Journal of Labor Economics*, 31 (3), 501–533.
- Edin, Per-Anders and Katarina Richardsson (2002) "Swimming with the tide: Solidarity wage policy and the gender earnings gap," *Scandinavian Journal of Economics*, 104 (1), 49–67.

- Eliasson, Tove and Oskar Nordström Skans (2014) "Negotiated wage increases and the labor market outcomes of low-wage workers: Evidence from the Swedish public sector," IFAU Rapport 2014:10.
- Exley, Christine L., Muriel Niederle, and Lise Vesterlund (2019) "Knowing When to Ask: The Cost of Leaning In," *Journal of Political Economy*, 128, 816–854.
- Forslund, Anders, Lena Hensvik, Oskar Nordström Skans, Alexander Westerberg, and Tove Eliasson (2014) "Avtalslöner, löner och sysselsättning," IFAU Rapport 2014:18.
- Gallen, Yana, Rune V. Lesner, and Rune Vejlin (2019) "The labor market gender gap in Denmark: Sorting out the past 30 years," *Labour Economics*, 56, 58–67.
- Grönqvist, Erik, Lena Hensvik, and Anna Thoreson (2022) "Teacher career opportunities and school quality," *Labour Economics*, 77 (101997).
- Hibbs, Douglas A and Håkan Locking (1996) "Wage compression, wage drift and wage inflation in Sweden," *Labour Economics*, 3 (2), 109–141.
- Jäger, Simon, Shakked Noy, and Benjamin Schoefer (2022) "The German Model of Industrial Relations: Balancing Flexibility and Collective Action," *Journal of Economic Perspectives*, 36 (4), 53–80.
- Kiessling, Lukas, Pia Pinger, Philipp. Seegers, and Jan Bergerhoff (2024) "Gender differences in wage expectations and negotiation," *Labour Economics*, 87.
- Li, Jiang, Benoit Dostie, and Gaëlle Simard-Duplain (2023) "Firm Pay Policies and the Gender Earnings Gap: The Mediating Role of Marital and Family Status," *ILR Review*, 76 (1), 160–188.
- Milgrom, Eva M. Meyersson, Trond Petersen, and Vemund Snartland (2001) "Equal Pay for Equal Work? Evidence from Sweden and a Comparison with Norway and the U.S.," *The Scandinavian Journal of Economics*, 103 (4), 559–583.
- Olivetti, Claudia and Barbara Petrongolo (2016) "The evolution of gender gaps in industrialized countries," *Annual Review of Economics*, 8, 405–434.
- Olsson, Maria (2024) "Labor cost adjustments during the Great Recession: wages, separations, and labor market frictions," *Mimeo*, <https://sites.google.com/view/mariaolsson>.

- SCB (2013) "Lönestruktur hela ekonomin 2013," Technical report, Statistics Sweden, https://www.scb.se/contentassets/866428f1e4c346ccabe0c37ec2f33170/am0110_b_s2013_l140612.pdf.
- Säve-Söderbergh, Jenny (2019) "Gender gaps in salary negotiations: Salary requests and starting salaries in the field," *Journal of Economic Behavior & Organization*, 161, 35–51.
- Willén, Alexander (2021) "Decentralization of wage determination: Evidence from a national teacher reform," *Journal of Public Economics*, 198, 104388, <https://doi.org/10.1016/j.jpubeco.2021.104388>.

A Appendix

A.1

A.1.1 Data on collective agreements and the survey

We use data on collective agreements from the Swedish mediation office as recorded in the ALEX-database (IFAU). This data includes detailed information about collective agreements between 2000-2011. The sample in ALEX is based on the yearly reports from the Swedish mediation office where all larger contracts are reported. Smaller contracts are thus excluded and unmatched to the survey.

The survey was sent to 7,098 firms. 2,950 unique firms replied with an answer to the question if they have or do not have a collective agreement. 1,982 firms replied that they have a collective agreement, 808 that they do not have a collective agreement, and 160 replied an add-on agreement (“hängavtal”). After excluding firms that reported an add-on agreement or contracts referring to the Swedish Church, sports or interest organizations, or only white-collar agreements—752 firms provided the name of a blue-collar agreement which could be directly identified in the ALEX-database along with information on the contract type.

For firms that replied only an employee organization without more details on the contract, and if that employee organization had multiple different contracts, the firms were matched to an agreement only if the contract could be inferred from the answers of other firms in the same 5-digit industry. For example, a firm that replied only “If Metall” would not be directly matched as this association have multiple contracts. But if most firms in that same 5-digit industry reported additional information on the used If Metall contract, for example “If Metall mining” we would use that information to assign a contract to the firm.

For firms that reported multiple contracts that correspond to at least two blue collar agreements in ALEX, we use the contract that cover most of their blue collar workers. If only one of the contracts is recorded in ALEX, we use that contract (the largest aggregate coverage according to the Swedish mediation office). If the share of workers on different contracts is not reported, we use the answers from other firms in the same industry to assign the most used contract. We exclude firms where this mapping is unclear.

A.1.2 The sample

Table A2 shows descriptive statistics for the firms and jobs in our dataset. A firm-occupation cell (job) is included in the wage data (column 1) if the firm was sampled in the wage structure statistics in 2013, had at least one man and one women working fulltime in a blue collar occupation, which also exist in another firm. We further restrict our data to job-pairs (column 2) where we require at least one man and one women working fulltime at the same firm and occupation.²¹ With these restrictions, the average firm and job is larger (mean 435 workers vs. 299 workers) and slightly more productive than the average firm in the initial wage structure statistics, but looks similar across our wage measures. Our final sample further restrict attention to the job-pairs with an identified blue-collar agreement (column 3). These are jobs at firms who replied to the survey with the name of a blue-collar agreement that could be identified in the ALEX-database or to "no union contract", and where the occupation is observed in at least two firms given the sample requirements described above.

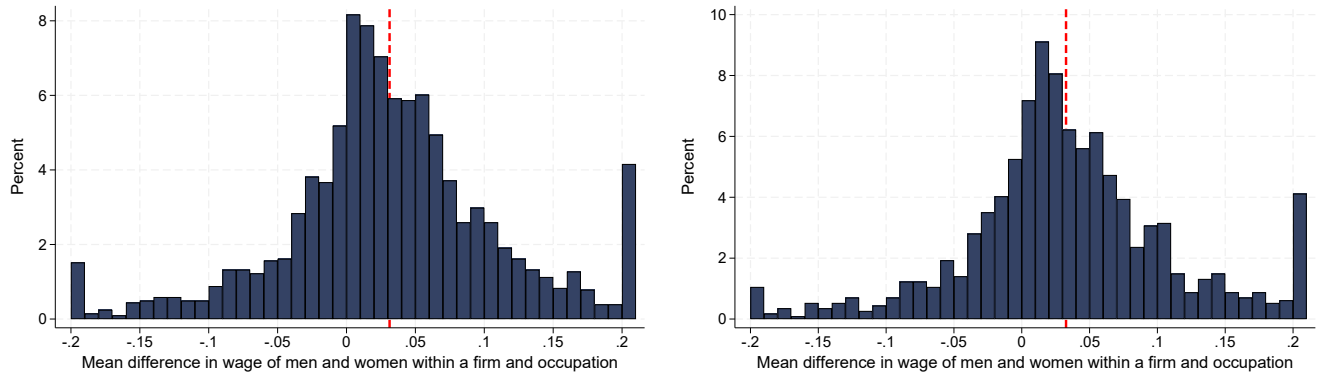
A.1.3 The assignment of contract types

The main specification uses the mode contract type over 2000-2011 from information in ALEX. Most contracts have stable contract types over time. One exception here is a larger contract called "verkstadsavtalet". In Table A1 the baseline results are confirmed while excluding this contract from the subsequent analysis.

²¹For outcome variables that refer to the distribution, Table 4 and Table 8, we require at least 1 man and 2 women, or 2 men and 1 women, in each firm-occupation.

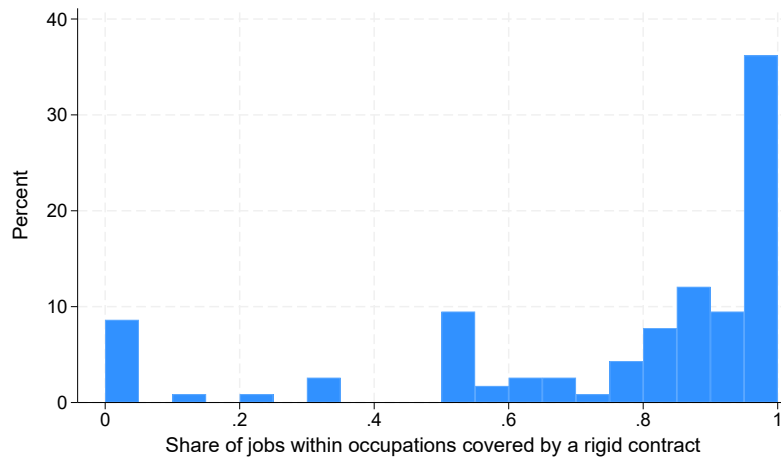
A.2 Figures and Tables

Figure A1: Histogram of gender-pay differences



Notes: These figures plot the mean difference in wage of men and women within a firm and occupation, as calculated in Equation 1. Year 2013. WinzORIZED at ± 0.2 . Left figure uses all jobs in the data. Right is for the baseline sample of occupations where we have information on contract rigidity. The red vertical line indicates the mean difference (0.033 in left and 0.035 in right).

Figure A2: Distribution of rigidity by occupations



Notes: Share of jobs within 4-digit occupations covered by a rigid union contract.

Table A1: Alternative specifications

	(1) Productivity	(2) Controls in 5-bins	(3) 3-dig occ	(4) Any contract	(5) Stayers only	(6) Excluding "verkstads"
Rigid	-0.0227 (0.011)**	-0.0219 (0.011)**	-0.0216 (0.010)**	-0.0206 (0.011)*	-0.0185 (0.009)**	-0.0256 (0.010)**
Observations	1142	1142	1142	1142	1081	880
Mean dep.	0.0354	0.0354	0.0354	0.0354	0.0358	0.0372

Notes: Regressions follow the baseline specification in Table 3 for alternate controls and samples. Column 1 adds control for productivity (3-bins of residualized value added per FTE worker). In column 2, we use 5-bins (instead of 3-bins) for the firm controls and in column 3 the 4-digit occupation fixed effects is replaced with 3-digit. Column 4 control for any collective bargaining agreement at the firm. In column 5, we exclude new hires and in column 6 "verkstadsavtalet" is excluded. All regressions have occupation fixed effects and controls for firm size, female share, and immigrant share at the firm. Std. errors clustered at the firm in parentheses. * $p < .10$. ** $p < .05$. *** $p < .01$.

Table A2: Sample statistics

	(1) Wage data		(2) + Job-pairs		(3) + Contract Info.	
	Mean	Sd.	Mean	Sd.	Mean	Sd.
Firm size	299	(775)	435	(970)	529	(1297)
Female share	0.23	(0.24)	0.29	(0.22)	0.25	(0.21)
Immigrant share	0.16	(0.17)	0.17	(0.16)	0.17	(0.14)
Productivity 5-quintiles	2.47	(1.72)	2.50	(1.73)	3.09	(1.47)
Part-time share	0.09	(0.16)	0.10	(0.17)	0.09	(0.17)
Jobb size	29.69	(182.99)	73.04	(301.02)	81.08	(236.23)
Jobb female share	0.22	(0.34)	0.33	(0.24)	0.30	(0.24)
Male occupation	0.83	(0.37)	0.82	(0.39)	0.82	(0.39)
Av. wage men	10.17	(0.15)	10.15	(0.14)	10.18	(0.14)
P25 wage men	10.12	(0.15)	10.09	(0.14)	10.11	(0.15)
P75 wage men	10.21	(0.16)	10.20	(0.15)	10.22	(0.16)
Av. wage women	10.10	(0.14)	10.11	(0.14)	10.13	(0.14)
P25 wage women	10.07	(0.15)	10.07	(0.14)	10.09	(0.14)
P75 wage women	10.14	(0.16)	10.15	(0.15)	10.17	(0.15)
Number of jobs	12461		4444		1142	
Number of firms	3910		2227		519	
Number of occupations	199		172		116	

Notes: This table shows the mean and standard deviation for firm and job characteristics across three samples. Column 1 includes blue-collar jobs sampled in the wage structure statistics in 2013. Column 2 further restrict attention to jobs (firm-occupation pairs) with at least one man and one women. Finally, column 3 condition on an identified collective agreement for blue-collar workers at the firm. For all samples, in addition to these requirements, we include only jobs if their (4-digit) occupations is observed in at least two firms.

Table A3: Sorting of women, part-time, and workers with young children

	(1) Female share	(2) Part-time share	(3) Children share	(4) Rigid
Rigid	0.000543 (0.014)	-0.00218 (0.007)	0.000105 (0.015)	
Part-time share				-0.0184 (0.061)
Female share				0.00234 (0.037)
Children share				0.0000653 (0.033)
Constant	0.323 (0.089)***	0.0667 (0.040)*	0.128 (0.063)**	0.819 (0.087)***
Observations	3168	3168	3168	3168
Mean dep.	0.186	0.0568	0.217	0.800
Fixed effects	Occ.	Occ.	Occ.	Occ.

Notes: Outcome variable in column 1 is the share of female workers in the firm-occupation cell, in column 2 the share of part-time workers, and in column 3 the share of workers with young children (age 10 or younger). Column 4 relates a rigid contract to the share of part-time, female workers, and workers with young children in the cell. In contrast to the main tables, this analysis does not require both genders to be present in each cell (job), hence the sample size is larger. All regressions have occupation fixed effects. Std. errors clustered at the firm in parentheses. * $p < .10$. ** $p < .05$. *** $p < .01$.