

Corradini, Viola; Lagos, Lorenzo; Sharma, Garima

Working Paper

Collective Bargaining for Women: How Unions Can Create Female-Friendly Jobs

IZA Discussion Papers, No. 15552

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Corradini, Viola; Lagos, Lorenzo; Sharma, Garima (2022) : Collective Bargaining for Women: How Unions Can Create Female-Friendly Jobs, IZA Discussion Papers, No. 15552, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/265773>

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. 15552

**Collective Bargaining for Women:
How Unions Can Create
Female-Friendly Jobs**

Viola Corradini
Lorenzo Lagos
Garima Sharma

SEPTEMBER 2022

DISCUSSION PAPER SERIES

IZA DP No. 15552

Collective Bargaining for Women: How Unions Can Create Female-Friendly Jobs

Viola Corradini

Massachusetts Institute of Technology

Lorenzo Lagos

Brown University and IZA

Garima Sharma

Massachusetts Institute of Technology

SEPTEMBER 2022

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

Collective Bargaining for Women: How Unions Can Create Female-Friendly Jobs*

Why aren't workplaces better designed for women? We show that changing the priorities of those who set workplace policies can create female-friendly jobs. Starting in 2015, Brazil's largest trade union federation made women central to its bargaining agenda. Neither establishments nor workers choose their union, permitting a difference-in-differences design to study causal effects. We find that "bargaining for women" increases female-centric amenities in collective bargaining agreements, which are then reflected in practice (e.g., more female managers, longer maternity leaves, longer job protection). These changes cause women to queue for jobs at treated establishments and separate from them less—both revealed preference measures of firm value. We find no evidence that these gains come at the expense of employment, workers' wages, or firm profits. Hence, prioritizing women's preferences in decision-making can lower within-firm gender inequality through more efficient bargaining.

JEL Classification: J31, J33, J51, J52

Keywords: gender disparities, amenities, collective bargaining, unions

Corresponding author:

Lorenzo Lagos
Brown University
Providence
Rhode Island 02912
USA
E-mail: lorenzo.lagos@brown.edu

* We thank Daron Acemoglu, David Atkin, David Autor, Abhijit Banerjee, Esther Duflo, Brandon Enriquez, Henry Farber, and Simon Jäger, as well as seminar participants at MIT, Brown, Princeton, NBER Summer Institute, SOLE, LERA, PUC-Rio, FGV-EESP, Insper, Central European University, and Norwegian Business School for helpful advice and comments. We are grateful to João Fernandes, Roberto Hsu Rocha, Stanley Gacek, Juvandia Moreira, Juliana da Penha Thomaz, and Beatriz Santos for their incredible help in better understanding the context, as well as Iacopo Morchio and Christian Moser for their generous help with the PageRank codes. Access to Brazil's RAIS database is governed by the Data Use Agreement between MIT and Brazil's former Ministry of Labor. We thank David Atkin and Mayara Felix for procuring MIT's access to the database, and Mayara Felix for de-identifying, harmonizing, and translating the RAIS datasets pursuant to MIT COUHES guidelines.

Despite significant labor market progress over the past decades, women continue to disproportionately suffer large earnings losses because they are in-charge at home (Kleven et al., 2019). Across 142 countries, over 30% of working women cite having to balance family and work as their main challenge (ILO and Gallup Inc., 2017). While governments and scholars have argued that making workplaces more female-friendly is key to lowering gender disparities—for example, Goldin (2014) argues that changing how jobs are structured may cause all remaining gender earnings gaps to vanish—little is known about whether labor market institutions can ameliorate the stark trade-offs working women face.

Per one view, making workplaces female-friendly—providing maternity leave, childcare, and flexible work schedules—is not worth the cost to employers since the marginal worker does not value these amenities. This paper tests an alternate view: that there is surplus to be shared, and the priorities of those deciding how to share that surplus determines workplace amenities. Because a few individuals typically decide workplace policies, their priorities take precedence and may not always feature women’s needs on top, even if these needs could be met. When these priorities change so too does the workplace. Unions provide a natural setting in which to test this hypothesis because, for about 20% of the world’s workers, a few union representatives negotiate pay and benefits (Visser, 2019). Since few union leaders are women, they may not represent women’s interests in collective bargaining.¹

In this paper we ask whether changing leader priorities in women’s favor changes workplaces and at what cost. The ideal experiment requires a top-down shift in priorities uncorrelated with changes to a firm’s labor demand or worker preferences. We exploit such a natural experiment in Brazil that pushed leaders at the largest trade union federation (or “union central”), the *Central Única dos Trabalhadores* (CUT), to prioritize women’s needs in collective bargaining.² Starting in 2015, CUT reserved half its leadership positions for women and prioritized female-centric policies in bargaining, such as 6 months of paid maternity leave, flexible work schedules, and childcare. Because unions seldom change affiliation to a union central and workers/establishments do not choose their union, this CUT reform serves as a top-down pro-women directive to union leaders unrelated to an establishment’s local labor demand and supply conditions. This motivates using a difference-in-differences design to compare amenities and related costs at establishments negotiating with CUT unions (treated) relative to non-CUT affiliates (comparison).

Unique to the Brazilian setting, our analysis relies on linking three rich sources of data: (i)

¹For example, nearly half of all workers but only 12% of union leaders are women in Brazil. In continental Europe, where collective bargaining covers a majority of workers, including Germany, Austria, and the Netherlands, fewer than 30% of union members are women (Skorge and Rasmussen, 2021).

²Union centrals are umbrella organizations that coordinate priorities among local unions. Half of Brazilian workers are covered by collective bargaining and 20% of unions affiliate with CUT.

establishment-level amenities from the text of all collective bargaining agreements (CBAs); (ii) worker outcomes from linked employer-employee data on the universe of formal sector workers; and (iii) union affiliation and leadership from the union registry. CBAs offer high quality information on 137 different types of amenities offered by establishments, including maternity leave, workplace safety, and rules governing absences or work hours. The administrative data tracks workers across employers over time and provides information on their gender, wages, and instances of maternity leave.

We start by using a revealed preference approach to identify which amenities are highly valued by women and which by men, relying on the idea that workers flock to employers with better work conditions. Employer-to-employer moves thus reveal valuable firms (Sorkin, 2018; Morchio and Moser, 2020) and correlating these values with CBA clauses reveals valuable amenities. We find that women value amenities that help strike work-life balance, including maternity protections, childcare payments, policies for dependents, absences, and workday reductions (“female-centric” amenities). In contrast, men value better pay and safer work conditions, including profit sharing, hazard pay, life-insurance, and safety equipment (“male-centric” amenities). In an out-of-sample sense check, we find that female amenities increase—and male amenities decrease—as women become a majority of an establishment’s workforce, hinting at the role that representation may play in how amenities are set at a workplace.

In the second part of our analysis we study the causal effect of making union leaders more pro-women on female and male-centric amenities, as well as downstream effects on workers and establishments, e.g., retention, employment, wages.

Our first main takeaway is that female-centric amenities increase on paper and in practice. On paper, female-centric amenities rise by 19% in treated relative to comparison contracts. This is a large effect, equivalent to moving from an average contract negotiated at a minority-female establishment to one where over 80% of the workforce are women. Interestingly, the biggest amenity gains are found where women have limited voice, i.e., where they are a minority of the workforce and where they are not well represented in union leadership. These amenity improvements on-paper translate into practice. After the reform, women are more likely to be managers, take longer maternity leaves, and enjoy job protection even after extending their maternity leaves. Although the composition of CBAs becomes more female-oriented (less male-oriented), we do not observe changes in the workplace environment that would suggest men are worse off, e.g., workplace injuries do not increase.

Our second main takeaway is that women value the new workplace environment being created at establishments affected by the CUT reform, ruling out a pure compensating differences story for the amenity gains. Specifically, we see women separating less and queuing

more at treated establishments, both of which are revealed preference measures of valuing one’s job (Krueger and Summers, 1988; Holzer et al., 1991) Among women, retention increases by 1.8% and is mostly driven by a reduction in voluntary separations. While we do not observe job queues directly, probationary contracts are commonly used by employers to screen applicants. We find that women’s share among probationary workers rises by 1.7%, which proxies for a higher share of women among job applicants. In sum, women flock to CUT affiliated establishments once female-friendly amenities go up.

If women are not losing out, perhaps men are. However, there is little evidence of this. Although men become a smaller share among probationary workers, incumbent male workers are more likely to stay at their baseline employer—albeit not as much as incumbent women. This suggests that, while treated establishments are perhaps not as attractive to potential male workers, the men already working there are happy with the changes to the workplace. Overall, our findings are consistent with models of the labor market wherein firms post utility offers for each gender. The reform causes firms affected by the new union priorities to increase posted utility for women without decreasing the posted utility for men.

Our third main takeaway is that improvements for women come without meaningful tradeoffs in employment or wages. Compensating differences would suggest that women’s wages fall to finance amenity improvements (Rosen, 1986). However, we find no effect on the mean log wages of both new and established workers for either gender group, ruling out even very small decreases at a 95% confidence level. We also find no wage decreases among incumbents or in the percentage wage adjustments negotiated in CBAs. Given no changes in wages, establishments might lower women’s employment because they are more costly to employ. There is no evidence of this either; in fact, employment remains unchanged while the share of women observed in the treated establishments rises by a small fraction (0.2pp relative to a baseline value of 36%).

If workers are not paying for the amenity improvements, perhaps firms are through lower profits. We provide empirical evidence and theoretical reasons against this explanation. Empirically, we find no increase in the wage bill or in establishment exit—which is a non-trivial margin of adjustment since 8.7% of establishments in the treatment group exit by the end of our sample period. For the subsample of establishments that report to Orbis, we also find null effects on profit margins. Theoretically, the CUT reform shifted union priorities rather than increase unions’ bargaining power, meaning that CUT unions were not positioned to capture a larger share of surplus, and thereby reduce profits. Indeed, while increasing union bargaining power generally predicts changes in employment—either moving right along the firm’s upward-sloping labor supply curve, or left along its labor demand curve, as in Lagos (2021)—we observe a precisely estimated zero.

How are these amenity gains paid for? Our results suggest that bargaining was inefficient ex-ante and the change in union priorities led to a Pareto improvement. There are at least two general models that explain these results. In one model, frictions in the bargaining process (e.g., information, contracting) or in how workers' interests are represented at the union level open up the possibility for win-win situations once unions are forced to focus on previously ignored issues. In another model, behavioral firms and unions do not think of providing female-centric amenities until changing union priorities put these issues front-and-center.

In the final part of our analysis we develop a method to quantify welfare changes from improving work environments using a revealed preference logic from consumer theory (Feenstra, 1994; Redding and Weinstein, 2016). Just as changes to consumer welfare from improving product varieties are measured through changes in the price index (i.e., how much more or less it costs to buy one additional util), changes to worker welfare from improving workplace conditions can be measured using changes to a wage index (i.e., how much more or less workers are paid to work one disutility-weighted hour.) This sufficient statistics approach allows us to remain agnostic on how different amenities enter workers' utility. Consistent with our reduced-form findings, we find the CUT reform improves women's welfare by about 6% while leaving men's welfare essentially unchanged.

Our paper contributes to four literatures. First, on unions and inequality. While firms care about the marginal worker, it's unclear who the union cares about (Farber, 1986). Unions have long struggled to organize workers with competing interests—see Hill (1996) for a historical review—and unionization has mixed effects for different worker groups, raising wages for low skill workers (Card, 1996; Farber et al., 2021) and black workers (Ashenfelter, 1972), but not necessarily women (DiNardo et al., 1996; Card et al., 2004, 2020; Bolotnyy and Emanuel, 2022). We provide quasi-experimental evidence that union leaders' bargaining priorities importantly determine whose interests they represent. When they prioritize women, unions can lower within-firm gender inequality.

Second, on the importance of leader priorities in how institutions function. Political leaders better represent their own group's preferences than the average constituent's (Chattopadhyay and Duflo, 2004; Pande and Ford, 2012). In the labor market, women negotiate less over pay than men (Dittrich et al., 2014; Leibbrandt and List, 2015; Biasi and Sarsons, 2022), suggesting that leaders could step in on their behalf. While women on company boards have only modest effects on gender gaps (Bertrand et al., 2018; Flabbi et al., 2019; Matsa and Miller, 2011; Maida and Weber, 2020), we find an important role for union leaders. Just as in politics, top-down changes in union leaders' priorities meaningfully change the workplace, in this case making it better for women.

Third, on the role of firms and worker voice in setting compensation. There is growing empirical evidence on employers’ wage-setting power in many labor markets (e.g., Manning, 2011; Card et al., 2018; Lamadon et al., 2022). Naturally, firms may also have some discretion in setting amenities, such as choosing health insurance networks (Tilipman, 2022). However, the impact of worker voice on wages and job quality remains rather mixed (e.g., Blandhol et al., 2020; Harju et al., 2021; Adhvaryu et al., 2022). This paper shows that worker voice—as expressed through unions—constrain employers’ discretion in setting compensation, and that amplifying the voices of working women can lead to Pareto improvements.³

Finally, our paper contributes to the revealed preference literature in three ways. We provide among the first quasi-experimental evidence that worker moves respond to amenity changes, consistent with a generation of papers that infer amenity values using such moves (Krueger and Summers, 1988; Sorkin, 2018; Taber and Vejlin, 2020; Morchio and Moser, 2020; Lagos, 2021; Lamadon et al., 2022). We also use worker moves and variation in amenities across establishments to identify what workers value, using a much richer set of amenities and higher stakes environment than possible through experiments (Mas and Pallais, 2017; Wiswall and Zafar, 2017; Maestas et al., 2018). Encouragingly, our results match this experimental work eliciting workers’ willingness-to-pay, e.g., women value flexibility. Finally, we bring an age-old method from consumer theory to the labor setting in order to obtain welfare estimates brought by changes in the firm environment.

The paper proceeds as follows. In Section 1, we discuss the institutional context and CUT reform. In Section 2, we describe our data and detail our approach for classifying amenities as female- or male-centric. In Section 3, we present our empirical strategy. In Section 4, we discuss our results on how changes in union priorities affect amenities, revealed preference measures of firm value, and labor market outcomes. In Section 5, we present our model quantifying the welfare impact of improving (female-centric) amenities on men and women. Section 6 concludes.

1 Institutional Context

We start by describing the collective bargaining structure in Brazil, emphasizing the distinction between unions (that represent workers in collective bargaining) and union centrals (that coordinate activities among affiliated unions). We then describe the 2015 reform enacted by a particular union central—namely, CUT—which provides the top-down pro-women shift in priorities at affiliated unions we use for identification.

³We find no evidence that providing female-centric amenities causes employers to lower wages (Gruber, 1994), or discriminate against women in hiring (Summers, 1989).

1.1 Collective Bargaining and Union Centrals

Types of CBAs There are two types of collective bargaining agreements in Brazil: sectoral and firm-level CBAs. In sectoral CBAs, unions negotiate with employer associations representing establishments in a specific industry and geography, e.g., car manufacturers in Curitiba. In firm-level CBAs, unions negotiate with individual employers, e.g., Volkswagen. Given their wider coverage, sectoral agreements typically set general floors for workers’ wage and non-wage benefits. Firm-level agreements build on these general floors to expand fringe benefits for workers at individual employers (Horn, 2009). We use firm-level CBAs in our main analysis (see Section 4.1), but also leverage information from sectoral CBAs to learn which clauses are highly valued by female and male workers (see Section 2.2).

Union determination The union that negotiates CBAs on behalf of workers at a given employer is not chosen by either the workers or the employer. This is because representation in collective bargaining is determined by two factors: 1) industry of employment (or category); and 2) the geographic location of work at the municipality level.⁴ Unions can join neighboring unions representing the same category to form a federation (state level) or confederation (national level). The left panel of Figure 1 depicts this within-category vertical structure. Examples of unions include the bank workers’ union of São Paulo and the teachers’ union of Florianopolis.

Neither workers (nor employers) can change the union that represents them (they negotiate with). Due to the legacy of Brazil’s corporatist past, the first union approved to represent a given category-geography cell enjoys a lifetime monopoly.⁵ As such, workers can only influence their union priorities from within, e.g., voting in union elections, running for leadership positions, voicing concerns to leadership, etc. At the same time, employers cannot avoid unions by virtue of this predetermined assignment to all category-geography cells. Naturally, union assignment by these cells produces an incredibly fragmented union landscape, with over 11 thousand labor unions in Brazil.⁶

⁴For a few professions, the worker’s occupation (rather than the industry of employment) becomes the relevant category for representation in collective bargaining, e.g., elevator operators, journalists, and musicians. These cross-industry, occupation-based unions comprise approximately only 5% of all unions in Brazil.

⁵President Getúlio Vargas instituted this “monopoly union” framework (known as *unicidade sindical*) in the late 1930s as a way to co-opt the labor movement by giving the federal government the power to determine which union would be given the rights over the representation of workers in collective bargaining.

⁶It’s worth noting that the assignment of representation rights (known as *enquadramento sindical*) is not always clear-cut, e.g., separate unions may claim the same set of workers and the employer may claim yet another union already holds the representation rights. All such matters are dealt by the labor courts.

CBA coverage Workers (and employers) cannot opt out of CBAs negotiated by their union. Coverage is universal, which means that workers need not be union members to enjoy negotiated benefits.⁷ Consequently, union membership is low (at around 20%) consisting of workers willing to pay membership dues in exchange for additional benefits that are not in CBAs, e.g., recreational facilities and private health insurance plans. Importantly, individual work contracts cannot take away benefits negotiated in CBAs, meaning that CBA provisions constitute a general floor for all represented workers. Similarly, CBAs cannot derogate provisions granted by the federal labor code, implying that CBA clauses generally build on top of these basic guarantees enjoyed by all workers.

Negotiation process The priorities of unions play a central role in determining what is discussed in CBA negotiations. Before an existing CBA expires, the union organizes a General Assembly where workers vote on the list of demands (or *pauta de reivindicações*) that they want to prioritize in the next negotiation. Union leaders typically select the topics that are discussed at these assemblies and are up for vote into the *pauta*. Negotiations officially start when the union sends these demands to employers and happen over several rounds. Most CBAs are signed for a duration of 12 months, giving rise to annual negotiations.⁸ Who sits at the bargaining table is decided by the union board and is not restricted to board members only.

Union centrals Unions can affiliate with union centrals (or *centrais sindicais*), which are somewhat analogous to trade union federations like the AFL-CIO in the United States. These are national level, umbrella organizations that help coordinate the activities of local unions and lobby for political favor (Liukkunen, 2019). While union centrals do not directly participate in collective bargaining, they are indirectly involved by coordinating union priorities *across* worker categories. For example, union centrals organize general strikes, plan conferences for union representatives, provide support to local unions, participate in public discussion forums on behalf of constituent unions, and steer union attention towards certain broad themes like gender or racial equality.

There are 9 union centrals in Brazil, depicted in the right panel of Figure 1. The *Central Única dos Trabalhadores* (CUT) is the largest of these, representing 30.4% percent of all organized workers as of 2016.⁹ CUT is Latin America’s largest union central, and among the

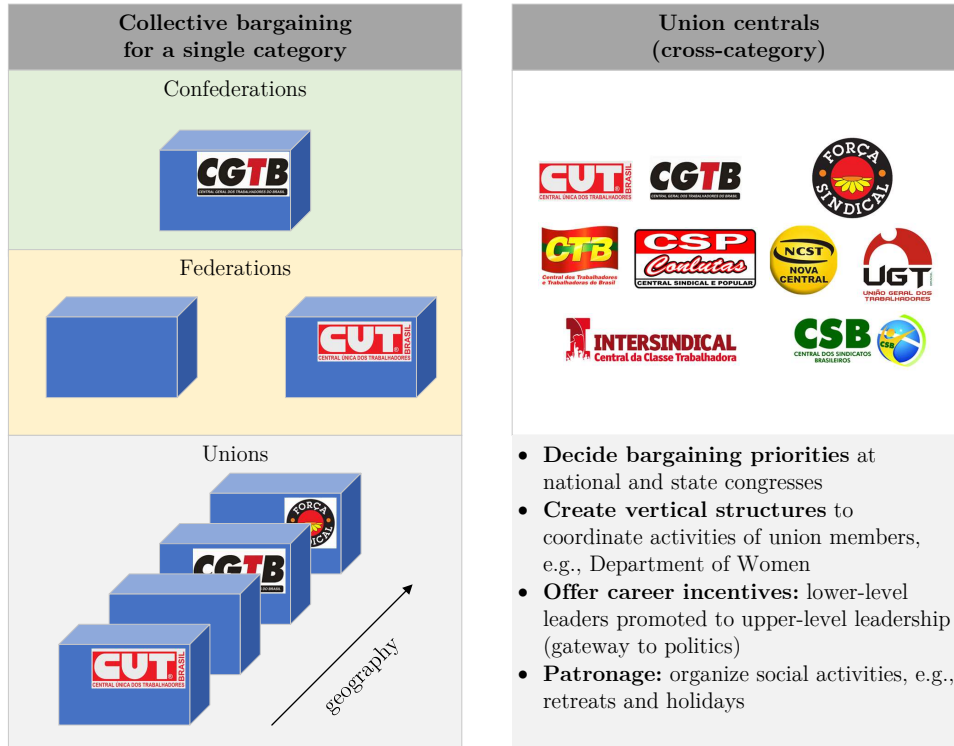
⁷Despite universal coverage, CBA coverage in Brazil is around 50% partly because not every union has a CBAs covering all (sometime any) of the municipalities they represent.

⁸In some cases negotiations take place every two years—the maximum possible duration for a CBA.

⁹The other union centrals are: *Força Sindical* (FS), *União Geral dos Trabalhadores* (UGT), *Central dos Trabalhadores e Trabalhadoras do Brasil* (CTB), *Nova Central Sindical de Trabalhadores* (NCST), *Central Geral dos Trabalhadores do Brasil* (CGTB), *Central dos Sindicatos Brasileiros* (CSB), *Intersindical—Central da*

largest in the world. It has close ties with the *Partido dos Trabalhadores* (PT), the Workers’ Party, which is Brazil’s most prominent left-leaning political party. President Luiz Inácio Lula da Silva (founding member of PT) was the leader of a metalworkers’ union within CUT before moving into politics—a common path for PT politicians (Lang and Gagnon, 2009).

Figure 1: Workers’ Representation Structure



Notes: Figure depicts the organizations representing workers in collective bargaining (as blue blocks on the left panel) and the union centrals they can affiliate with (as logos on the right panel). All workers in a category-geography cell (e.g., bank workers in São Paulo) are represented by a single union. Unions can integrate geographically within the same category, forming a federation (at the state level) or a confederation (at the national level). Local unions, federations and confederations can affiliate with union centrals (*centrais sindicais*), which are depicted in the figure as union central logos “stamped” on the blue blocks. Union centrals are associations of unions, representing cross-category interests and operating on a nationwide level, with political objectives and coordination functions. Union centrals cannot directly participate in collective bargaining.

CUT is vertically organized into congresses and executive boards at the regional, state, and national levels. Congresses are meetings of delegates who are elected by individual unions to develop a coherent agenda for unions within CUT.¹⁰ They meet once every three years to vote on CUT’s overarching priorities for the next 3 years, recorded in a book of resolutions or “fight plan”. Executive boards comprise a smaller group of leaders elected by

Classe Trabalhadora, and *Central Sindical e Popular Conlutas*.

¹⁰Elected delegates are typically local union leaders. The number of delegates that each union gets to elect to different levels depends on the number of workers it represents. Outlined in the CUT constitution [here](#).

congresses to oversee CUT’s day-to-day functioning. They manage CUT’s finances, oversee the fight plan’s execution, organize meetings and training of local union leaders, and have committees tackling specific topics like gender equality within CUT.¹¹

1.2 CUT Reform

The origins of the CUT reform we study arises from the tight link between this union central and the Workers’ Party (PT). In 2011, PT instituted a 50% quota for women in its leadership and Dilma Rousseff—the party’s candidate—was elected as Brazil’s first female president. This spawned a demand for greater gender parity within CUT. Prominent female leaders authored op-eds demanding greater say for women within CUT’s leadership and a similar quota (Godinho Delgado, 2017). They were successful, and CUT’s 2015 state and national congresses saw an unprecedented focus on women. The 2015 reform had two parts.

1) Gender quota First, CUT reserved 50% of seats in its state and national executive boards for women. This quota was voted in by the 2012 state and national congresses and came into effect in 2015. Figure 2a shows that the quota had bite: the share of women in CUT’s national board rose sharply from 35% to 50% right in 2015. To accommodate having more women in its national board, the board size was almost doubled from 30 to 50 members. Importantly, there is no indication that other union centrals directly reacted to CUT’s quota, maintaining a rather stable share of women on their national boards around 21-25% (averaged across union centrals).¹²

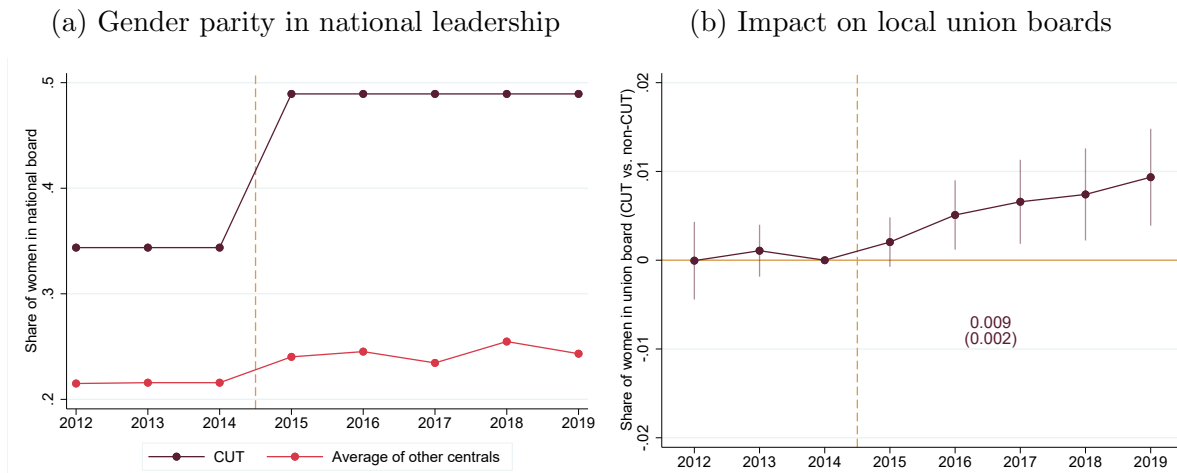
Along with this large increase at the union central level, the quota had spillover effects on the representation of women in CUT-affiliated unions. Figure 2b compares the share of women in local union boards between CUT and non-CUT affiliates using a difference-in-differences design. The effect is positive and statistically significant at about 3% relative to baseline. This estimates effect is not mechanical as the quota only applied at the union central level.¹³ Hence, the first part of the CUT reform should be interpreted as a leadership change favoring women mainly at the national level—where the involvement in collective bargaining is only indirect through, for example, the coordination of union activities and

¹¹For instance, CUT established the National Committee of Working Women in 1986 to campaign for universal childcare. In 2003, it gained a broader mandate to organize gender-related advocacy within CUT and became institutionalized as the Department of Working Women.

¹²The slight up-tick in 2015 for the share of women on the national boards for non-CUT centrals is driven by *Conlutas*—an even more combative left-leaning union central than CUT that has a very small number of affiliated unions. The main competitor to CUT is *Força Sindical* and it slightly decreased the share of women in its national board in 2017. See Figure B1 for details.

¹³If anything, the promotion of female leaders from the local union level to satisfy the quota could be dampening these positive spillovers.

Figure 2: The 2015 CUT Reform



Notes: The 2015 CUT reform consisted of two parts. The first is a 50% quota for women in CUT’s state and national executive bodies. The second is the adoption of a bargaining agenda more attentive to the needs of female workers. Figure 2a plots the annual share of women on CUT’s national executive committee and the average share in the other 7 union centrals (*Intersindical* is dropped due to missing information on its board). Refer to Figure B1 for the plots corresponding to each individual union central. Figure 2b shows how the reform had downstream effects on the gender composition of local union boards (for CUT affiliates relative to non-CUT affiliates as of 2012). The figure depicts the estimated coefficients for the interactions between a CUT affiliate dummy and year fixed effects, where the regression’s dependent variable is the share of women in the board for a given union-year observation. The event-study specification omits the baseline year 2014 and includes both union fixed effects and year fixed effects. Note that the average share of women across CUT affiliates unions in 2014 is around 33%. Confidence intervals at a 95% level are reported. Standard errors are clustered by union.

bargaining priorities—with limited positive spillover effects at the local union level.

2) Female-centric fight plan Second, the 2015 CUT national congress adopted a bargaining agenda more attentive to the needs of female workers. Its new fight plan featured a 14-page section on achieving gender equality in the workplace, which was the first time that such a section was authored in at least 10 years. Figure B2 shows the cover of the 2015 fight plan. Some of its demands included advocating for 6 months of paid maternity leave (up from the state-required 120 days), reduced work hours and flexible schedules to accommodate women’s household duties, and childcare as a universal right. The word *mulheres* (women) appeared 203 times in the 2015, compared with 46 in 2012 and 74 in 2009.

CUT’s 2015 fight plan also detailed a series of measures to promote gender parity at the local union level. These included giving women chairmanship of important committees (like finance and communications) and involving women in drafting the *pautas de reivindicações*, i.e., the union’s list of demands that are taken to employers for negotiation.¹⁴ Therefore, independently of the changes to women’s representation in local union boards, the aforementioned recommendations potentially translated into practices that improved gender parity at the local union level.

Summary In sum, starting in 2015, CUT had more female leaders and made its commitment to gender equality evident, especially to local union delegates who attended its congresses. CUT vowed to give women more power within the union central and affiliated unions, prioritizing their needs at the negotiating table. Importantly, the CUT reform did not change the bargaining power of unions relative to employers but simply shifted the priorities of CUT affiliates. Hence, any realized improvements for women brought through collective bargaining are unlikely to come from a lower share of surplus captured by employers.

2 Data and Amenity Classification

To study how the CUT reform affects the workplace for women and at what cost, we need establishment-level information on wages, amenities, and employment, as well as each negotiating union’s affiliation to a union central. After describing the data that satisfies these needs, this section details how we combine CBA clauses (our measure of amenities) and worker flows across establishments to classify clauses as being differentially valued by women relative to men—what we denote as female-centric amenities.

¹⁴These strategies were developed at the 2015 meeting of CUT Women, and voted in as official CUT policy by delegates at the 2015 national congress. The full text of the book of resolutions can be accessed [here](#).

2.1 Data Sources

Our analysis relies on linking three rich sources of data: (i) amenities at the establishment-level from the clauses of the CBAs providing coverage; (ii) worker outcomes from linked employer-employee data on the universe of formal sector workers; and (iii) union affiliation and leadership from the registry of unions.

For information on amenities, we use CBA clauses scraped from the Ministry of Labor’s *Sistema Mediador* registry, which tracks and stores every CBA signed in Brazil since 2009. To register an agreement, clauses need to be classified into 137 different clause types, e.g., overtime pay, childcare assistance, profit sharing, paid leave, etc.¹⁵ We extract the number of clauses in each type as a measure of amenities offered to workers.

For information on worker-level outcomes we use linked employer-employee data known as *Relação Anual de Informações Sociais* (RAIS). These are administrative data covering the universe of formal sector workers. Essentially, the federal government requires that employers send information about each worker at their establishments every year. For each work spell, RAIS reports average monthly earnings, leaves taken, and (6-digit) occupation; worker characteristics like gender, age, and education; and establishment characteristics like location (municipality) and industry (6-digit). We link RAIS to CBAs using a common establishment identifier—known as CNPJ—in both datasets.

For information on each union’s affiliation to a union central and its leadership composition over time we use the national registry of unions, or *Cadastro Nacional de Entidades Sindicais* (CNES). We infer the gender of leaders using the R package *genderBR*, which codes a name as female if most people with that name are women in the Brazilian census (and similarly for men).¹⁶ Among all union leaders between 2005 and 2019, 27.7% are women, 67% are men, and 5% are unclassified. CBAs record the same union identifier as CNES, which we use to link contracts to unions and thus union central affiliation and board composition.

2.2 Classifying Female-Centric Amenities

By matching CBAs to signing establishments in RAIS we can track workers across jobs, observing not only their wages but also some general amenities provided at each job. However, whether a CBA clause is differentially valued by women relative to men (i.e., a female-centric amenity) is not directly observed in these data. We adopt two approaches for classifying clauses as female-centric, which allow us to quantify how female-friendly the work environment is. Here we describe the main steps of each approach, with details in Appendix C.

¹⁵Figure B3 shows an example of a maternity leave clause.

¹⁶Developed by Fernando Meireles and posted on [GitHub](#).

1) Intuitive approach In the intuitive approach, we classify 20 of the 137 pre-specified clause types in *Sistema Mediador* as being disproportionately valued by female workers (see Table 1, Column 1). They fall into four broad types detailed in Table A1: (1) Leaves, e.g., following maternity, adoption, or a miscarriage; (2) Maternity and childcare, e.g., employment protection after maternity, childcare assistance, and policies for dependents; (3) Workplace harassment and discrimination, e.g., sexual harassment and equal opportunities in promotions; and (4) Flexibility and part-time work, e.g., workday controls, uninterrupted shifts, and part-time contracts. Themes (1)-(3) involve clauses that one could reasonably associate with women. The last theme was selected due to the existing literature documenting the value women place on flexibility in work hours (Goldin and Katz, 2011; Mas and Pallais, 2017; Maestas et al., 2018).

2) Data-driven approach In the data-driven approach, we aim to identify CBA clauses that correlate with women’s disproportionate desire to work at an establishment (relative to men). The underlying model motivating this approach is one where workers of gender $G \in \{F, M\}$ share some common ranking over establishments $j \in \mathcal{J}$. A worker’s utility from working at j is higher than at k because of the different wages and amenities offered to their group G in these workplaces. In particular, our model assumes that this gender-specific value (denoted as V_j^G) is a linear function of wages, amenities, and an unobserved component.

At a minimum, our approach requires measures for the value of employment, the wages paid, and the amenities provided at each establishment. For the value of employment, we estimate gender-specific PageRank values (Sorkin, 2018; Morchio and Moser, 2020). This is a revealed preference measure of the value of working at an establishment that leverages employer-to-employer worker flows to obtain information about the job ladder workers face.¹⁷ For wages, we estimate wage premiums (ψ_j^G) from gender-specific AKM models.¹⁸ For amenities, we use the over-year average count of clauses $a(z)_j$ for each of the 137 clause types $z \in Z$ coming from CBAs covering workers. Hence, we can obtain gender-specific estimates for the value of employment and wage premiums at each establishment, but we only observe a proxy for amenities without a sense of which clause types are valued by women.

¹⁷The intuition behind the fixed point problem that pins down the PageRank values is that good employers have many worker inflows—especially from other good employers—and few worker outflows (see Appendix D). For implementation details, refer to Appendix C.

¹⁸AKM is the acronym for Abowd et al. (1999), the original paper estimating firm-specific wage premiums with linked employer-employee data. The underlying model also assumes a common job ladder among workers and requires worker flows across establishments for identification (see Appendix D). For implementation details, refer to Appendix C.

In other words, the underlying model tells us that

$$V_j^G = \beta_w^G \psi_j^G + \sum_{z \in Z(G)} \beta_z^G a(z)_j + \epsilon_j^G. \quad (1)$$

The data allow us to get estimates of V_j^G and ψ_j^G . Our classification problem is finding the set of clauses for which $\beta_z^F - \beta_z^M$ is positive (which we denote as “female-centric”), as well as those for which this difference is negative (which we denote as “male-centric”).¹⁹ Our solution is to take the difference between the female and the male version of Equation (1) to identify the clause types that women value disproportionately more than men, and vice-versa.

Specifically, we use the following hedonic regression

$$V_j^F - V_j^M = \beta_w^F \psi_j^F - \beta_w^M \psi_j^M + \sum_{z \in Z} \beta_z a(z)_j + \epsilon_j \quad (2)$$

where β_z is capturing the value of the amenity for women relative to men, i.e., $\beta_z^F - \beta_z^M$. We estimate the regression in Equation (2) using lasso to select amenities that are the most predictive of utility differences between women and men (controlling for gender-specific wage premia). We deem the top 20 clauses with the highest β_z as “female-centric” and the bottom 20 with the lowest β_z as “male-centric”. To the best of our knowledge, this is the first time that such a rich description of the work environment can be combined with administrative data on worker flows to uncover which workplace features are valued by different groups of workers, i.e., men and women.²⁰

Omitted variable bias While the data-driven approach is a predictive exercise, mitigating omitted variable bias is still important. For example, establishments that want to hire women may redouble their recruitment efforts or provide other job features valued by women in addition to the observed clauses. Because we do not directly observe recruitment intensity or perfectly observe the work environment, we might erroneously identify a clause as valuable because it covaries with these unobserved features.²¹ Hence, we measure amenities $a(z)_j$ using sectoral rather than firm-level CBAs because the former are negotiated by employer associations and are therefore not subject to the demands of any individual employer. Using sectoral CBAs for classification is also important because we use firm-level CBAs to study the

¹⁹An advantage of the data-driven approach (relative to the intuitive case) is that it identifies male-centric clauses, allowing us to test for potential amenity tradeoffs by gender in response to the CUT reform.

²⁰Some papers elicit workers’ willingness-to-pay for workplace attributes like flexibility and wage growth (e.g. Mas and Pallais (2017) for workers on an online platform and Wiswall and Zafar (2017) for NYU college students). They find that women value flexibility in work schedules more than do men. In the same context as ours, Lagos (2021) quantifies the wage-equivalent value of CBA clauses irrespective of gender.

²¹Including ψ_j^G partly addresses this concern by accounting for recruitment efforts operating through wages.

effect of the CUT reform. As such, there is no mechanical relation between our data-driven approach and our estimates of the impact of the reform on these clause types.

Estimation sample The variation in Equation (2) comes from a cross-section of establishments, restricted to the subset for which we can estimate V_j^G , ψ_j^G , and $a(z)_j$. First, because we need to observe PageRank values for both genders, which can only be estimated for the largest super-connected set of employers (i.e., each establishment must hire from and lose a worker to some other establishment in the set), we are restricted to the 2009-2016 intersection of the super-connected set between genders. Second, AKM wage premia are estimated from the gender-specific largest connected sets of establishments meeting some criteria to reduce noise in ψ_j^G (e.g., mean establishment size ≥ 10 workers). Thus, we are also restricted to the 2006-2016 intersection of these largest connected sets between genders. Third, to reduce noise in $a(z)_j$ (i.e., the over-year average of the clause types), we restrict the sample to employers covered by at least four sectoral CBAs in 2009-2016.

Normalization The gender-specific PageRank values and AKM wage premiums require a normalization to make the female-to-male difference interpretable. In the case of the AKM premiums, we normalize ψ_j^F and ψ_j^M relative to the restaurant sector—a relatively competitive industry where it’s reasonable to assume that employers cannot pay wage premiums to either gender. For the PageRank values, V_j^F and V_j^M are unique up to an unknown multiplicative factor. Our results below are robust to three alternatives for $V_j^F - V_j^M$. The first chooses the establishment with the smallest wage premium gap as the normalizing establishment, and then adjusts female values relative to the male values. The second simply assumes the multiplicative factor is the same for both genders, i.e., no normalization is needed. The third is to use a normalized index from 0 to 100 of V_j^F and V_j^M . Our baseline data-driven approach uses a 50% random sample of our estimation sample and the normalized version of the gender gap in PageRank values as the dependent variable.

Results Table 1, Columns 2 and 3 list amenities identified as female and male-centric in the data-driven approach. Clauses are ranked in descending order based on the absolute value of $\hat{\beta}_z$. The clauses in red are those also intuitively classified as female-centric.

In line with our intuitive definition, the data-driven approach reveals that women disproportionately value clauses on leaves (e.g., following adoption and miscarriage), as well as childcare and maternity related clauses (e.g., childcare assistance, maternity protections, and policies for dependents). In addition to the 8 intuitively-defined female-centric clauses that are captured by our data-driven approach, Column 2 also highlights some female-centric

clauses that our intuitive definition was missing, e.g., absences, extension/reduction of workday, medical exams, and health education campaigns.

On the male side, we also get sensible results. Men highly value additional pay, e.g., on-call pay, profit sharing, hazard pay, workday compensation—especially stemming from tragic events like life insurance and death/funeral assistance. Men also place disproportionate value on workplace safety, e.g., protections for injured workers, machine and equipment maintenance, working environment conditions, safety equipment, etc.²²

The fact that “female workforce” clauses appear among the data-driven male clauses highlights the fact that our approach does not account for variation within the text of each clause type. For example, “female workforce” clauses vary a lot in content, including items that are clearly beneficial to women (e.g., free provision of sanitary pads), as well as others that are beneficial to men (e.g., forbidding women to cast concrete or work on scaffolding). It seems likely that our data-driven approach is capturing increases in the latter. While the availability of pre-specified clause types allows us to have a simple measure of CBA content than avoids issues with topic models—such as text preprocessing, choosing the number of topics, and noisy estimates—it is by no means a faultless measure.

Sense checks Out-of-sample sense checks indicate that both the “intuitive” and “data-driven” approaches are identifying clauses that women (or men) value disproportionately more than the other gender. Using firm-level CBAs signed in 2014—i.e., the year prior to the CUT reform—we see that female (male)-centric clauses are increasing with the share of women (men) at an establishment.²³ Figure 3a shows that intuitively classified female-centric clauses rise almost linearly with this share. Figure 3b shows a similar relationship for male and female-centric clauses as defined per the data-driven method. Specifically, the gap in male-to-female clauses starts at ≈ 1.5 in all-male workplaces and shrink to negligible levels as one moves toward all-female workplaces. Interestingly, the data-driven female clauses only begin to increase once women are a majority (around the 50% threshold). This suggests either that women successfully advocate for these amenities once they enter the majority or that establishments provide them to attract women—both implying that women value them.

²²Tables A2 and A3 show robustness of the data-driven approach to different normalizations for PageRank values. The tables also show that our classification is not driven by industry and geography since similar clauses are selected when adding state and industry fixed effects to the lasso. In fact, the rank correlation of the coefficients on the selected clauses with versus without these fixed effects is positive and statistically significant (0.56 with p-value below 0.01).

²³We also have evidence that the number of female clauses is strongly positively related to the difference between women and men’s PageRank valuation of an establishment (see Figure 3).


Table 1: Female- and Male-Centric Amenities

Intuitive definition Female clauses	Data-driven definition		Rank
	Top 20 female clauses	Top 20 male clauses	
Abortion leave	Childcare assistance	On-call pay	1
Abortion protections	Absences	Life insurance	2
Adoption leave	Adoption leave	Strike procedures	3
Childcare assistance	Other: holidays and leaves	Other: protections for injured workers	4
Equal opportunities	Seniority pay	Profit sharing	5
Female workforce	Maternity protections	Salary deductions	6
Maternity assistance	Abortion protections	Female workforce	7
Maternity leave	Paid leave	Transfers	8
Maternity protections	Night pay	Machine and equipment maintenance	9
On-call	Nonwork-related injury protections	Duration and schedule	10
Other: holidays and leaves	Abortion leave	Working environment conditions	11
Paid leave	Policy for dependents	Salary payment - means and timeframes	12
Part-time contracts	Extension/reduction of workday	Hazard pay (danger risk)	13
Paternity protections	Guarantees to union officers	Safety equipment	14
Policy for dependents	Renewal/termination of the CBA	CIPA: accident prevention committee	15
Sexual harassment	Medical exams	Other assistances	16
Special shifts	Unionization campaigns	Death/funeral assistance	17
Uninterrupted shifts	Health education campaigns	Workday compensation	18
Unpaid leave	Waiving union fees	Collective vacations	19
Workday controls	Salary adjustments/corrections	Tools and equipment	20

Notes: Table lists the clause types that were selected as “female-centric” based on intuition (column 1) and with our data-driven approach (column 2), which also allows us to define “male-centric” clauses (column 3)—refer to Section 2.2 for details on the data-driven approach. The clauses in column 1 are listed in alphabetical order while those selected with the data-driven approach are ranked on the basis of the coefficients β_z coming from the estimation of Equation (2). That is, the first female clause listed is the one with the highest estimate of β_z , the second is the one with the second highest value of β_z , etc. Similarly, the male clauses are ranked from the one with the lowest estimate of β_z to the one with the 20th lowest estimate. In columns 2 and 3, we highlight in red the clauses that also belong to the intuitive definition of female-centric clauses in column 1.

Figure 3: Sense Checks for Female- and Male-Centric Amenities

(a) Intuitive female clauses and share of women (b) Data-driven clauses and share of women



Notes: Figures depict binned scatterplots of the number of female-centric (and male-centric) clauses contained in firm-level CBAs signed at baseline (2014) by the share of women in the workforce of the establishment. The bins in the bottom figures are set to rounded values (in 0.05 increments) of the share of women at the establishment, with the size of the markers scaled to represent the number of pairs observed in a given bin. Figure 3a uses the intuitive definition of female-centric amenities, while Figure 3b uses the data-driven approach for both female- and male-centric amenities. The vertical line indicates 50% of women in the workforce. The sample consists of the establishments in our new contracts panel at baseline (2014). Regressing the y-axis variables in the bottom figures on the share of women at establishments reveals a positive (negative) and statistical significant relation between female (male) centric clauses and the share of women at the establishment. For the intuitive definition of female-centric clauses, the slope is 0.137 (SE 0.019). For the data-driven definition of female-centric clauses, the slope is 0.172 (SE 0.034). For the data-driven definition of male-centric clauses, the slope is -1.219 (SE 0.042).

3 Empirical Strategy

We employ a difference-in-differences strategy to study the CUT reform’s effect on amenities and labor market outcomes. In this section, we first describe the three analysis samples we use to study the effects of the reform on collective bargaining agreements, establishments, and workers. We then detail the empirical approach that exploits pre-reform variation in unions affiliation to CUT to define treatment and discuss identifying assumptions.

3.1 Analysis Samples

We construct three analysis samples to study effects on negotiated CBAs, establishments, and workers. For further details, refer to Appendix C.

1) Amenities sample To study the evolution of amenities, we construct a balanced panel of CBA clauses covering each establishment-union pair between 2012-2017. As noted, any category (usually industry) of workers in a given geography is represented by a single union—meaning that each pair can be thought of as constituting a unique worker group.²⁴ We focus on clauses in firm-level CBAs because most improvements in amenities and workers’ conditions are achieved through these agreements (Horn, 2009; Liukkunen, 2019).²⁵

While not every establishment-union pair renegotiates contracts every year, we obtain a balanced panel by exploiting the fact that CBAs are automatically extended until the next negotiation during our sample period (Lagos, 2021). Given that all CBAs had to be registered in *Sistema Mediador* starting 2009 and their maximum duration in 2 years, our panel paints an accurate picture of active CBAs 2012-2017, including zeros. Our results are robust to using an unbalanced panel that uses only years in which a new negotiation occurs.

2) Establishment sample To study effects on potential tradeoffs from changing amenities, we construct a sample of establishments signing a CBA in our *amenities sample*, following their outcomes in RAIS. Outcomes include employment, share of women among workers, and mean log wages. We make two additional sample restrictions. First, establishments must employ both men and women at baseline (2014) to ensure that any amenity change has possible downstream effects on the workplace. Second, we only considered signing establishments in the geographic coverage of the contracts they sign. A signing establishment

²⁴Most signing establishments (92%) negotiate with a single union over the entire study period, i.e., employers rarely negotiate with more than one worker category.

²⁵In conversation with the President of the bankers’ union of São Paulo, she also confirmed that most amenity improvements are achieved through firm-level CBAs. This is because sector-level negotiations typically involve several tens (or even hundreds) of employers, making it difficult to reach consensus on a rich set of amenities. Unions therefore typically reserve these topics for negotiation with individual employers.

might be outside of the geographic coverage of a contract is if it signs a CBA that only covers other establishments of the same firm, as in case of a headquarter signing contracts on behalf of subsidiaries.

3) Incumbent worker sample To study individual worker-level outcomes like wages and retention, we construct a sample of incumbent workers employed at establishments in the *establishment sample* at baseline (2014). We track these workers wherever they go, i.e., not conditional on staying at baseline employer.

Treatment definition Following the 2015 reform, CUT-affiliated unions made women central to their collective bargaining strategy. While the reform was enacted in 2015, gender quotas were approved in 2012 (see Section 1.2), suggesting that CUT’s pro-women pivot might have been anticipated and spurred unions to switch affiliation to avoid (or benefit from) the changes. While unions rarely switch union central affiliation, we define treatment using a union’s 2012 CUT affiliation to avoid bias from selection into or out of CUT affiliation. Figure B5 shows that neither treated nor comparison unions systematically switch affiliation away from or toward CUT after 2012. Thus, there is no evidence of endogenous selection into (or out of) CUT affiliation after the announcement of the gender quota in 2012.

Treatment is defined in the following way. In the *amenities sample*, a treated establishment-union pair is one where the negotiating union was affiliated with CUT in 2012. In the *establishment sample*, a treated establishment is one that belongs to a treated pair.²⁶ Finally, in the *incumbent worker sample*, a worker is treated if they worked at a treated establishment at baseline (2014).

Descriptive statistics Table 2 shows descriptive statistics for our starting sample, i.e., the *amenities sample*. Column 1 describes the full sample, while Columns 2 and 3 report information by treatment status.

Panel A describes sample sizes. Our sample contains more than 211 thousand firm-level CBAs signed by 89,920 establishment-union pairs. These pairs include 80,153 thousands signing establishments and 4,409 signing unions. On average, each pair signs new contracts in 2.4 out of the 6 years spanning our study (2012-2017). About 21% of pairs are treated and 79% are in the comparison group. This sample covers about 1.6% of all private-formal sector establishments in Brazil, highlighting that only a select set of employers negotiate firm-level CBAs. However, since these firms are typically large (about 143 workers on average, per Panel C), these establishments account for over 11% of total employment in 2014.

²⁶Over 93% of establishments negotiate with a single union and 98% with all unions with the same union central affiliation.

Panel B describes contract provisions at baseline (2014). CBA negotiations (at the pair-year observation level) feature 24.7 clauses on average, of which 3.2 are “female-centric” per our data-driven definition described in Section 2.2. These numbers are statistically indistinguishable across treated and control contracts. Although the share of female-specific clauses might seem small, this statistic need not be a good measure of the value and importance of a clause. On average, there are 1.7 more male clauses than female clauses.

Panels C and D document establishment- and union-level characteristics, respectively, at baseline (2014). Our sample consists of large employers (especially in the treated group) but both treated and control have slightly more than a third of workers being women. Treated and comparison establishments resemble each other along a number of observable dimensions, including their distribution of size, geographic region, and share of women working (see Figure B6). On the union side, treated unions have larger boards but with a similar share of women (i.e., around 23%), indicating no baseline differences between CUT and non-CUT affiliated unions in female representation. Only about 17% of unions have a female president.

3.2 Differences-in-Differences Design

To measure the causal effects of the CUT reform on negotiated amenities and labor market outcomes, we compare treated units of observation (i.e., pairs, establishments, or incumbent workers) with the comparison group using a dynamic difference-in-differences specification

$$Y_{it} = \sum_{j=2012}^{2017} \beta^{t=j} (D_i \times \delta_{t=j}) + \alpha_i + \gamma X_{it} + \varepsilon_{it} \quad (3)$$

where i indexes the unit of observation and t indexes a year. The treatment indicator D_i is interacted with year fixed effects δ_t . The specification also includes unit fixed effects α_i , as well as time-varying fixed effects X_{it} , i.e., industry-year and geography-year fixed effects.^{27,28} Idiosyncratic errors are captured by ε_{it} and standard errors are clustered by establishment.²⁹

Our coefficients of interest are β^t , capturing the effect of treatment in year t relative to the baseline year (β^{2014} is normalized to zero). The model allows for average differences between the treated and the comparison groups, absorbed by unit fixed effects α_i . The identifying variation occurs within the same unit, comparing outcomes in any year relative to 2014, and

²⁷For industry we use the first two digits of Brazil’s CNAE codes, which gives 87 possible values such as textile production, road transportation, etc. The largest imbalance between treated and control establishments at the industry level is in retail industry (16% among control versus 13% among treated)—see Figure B6.

²⁸For geography we use either states (27 in total) or microregions, which are neighboring municipalities grouped into 543 units that capture local labor markets.

²⁹Clustering by establishment assumes that establishments negotiate with unions that, as of 2012, were affiliated at random with a union central.

Table 2: Sample Descriptives

	All (1)	Treated (2)	Control (3)
<i>Panel A: Sample characteristics</i>			
Collective bargaining agreements	211,619	42,523	169,096
Establishment-union pairs	89,920	19,040	70,880
Signing establishments	80,153	17,190	62,963
Signing unions	4,409	886	3,523
Avg. years of CBA negotiation (per pair)	2.35	2.23	2.39
<i>Panel B: CBA negotiation characteristics</i>			
Avg. clause count	24.7	23.1	25.1
Avg. female clause count (intuitive)	1.66	1.81	1.63
Avg. female clause count (data-driven)	3.16	3.15	3.16
Avg. male clause count (data-driven)	4.87	4.59	4.94
<i>Panel C: Establishment-level characteristics (2014, baseline)</i>			
Avg. employment	143	197	128
Avg. share of women in workforce	0.38	0.36	0.38
Share employing both men and women	0.82	0.83	0.82
Share of single establishment firms	0.62	0.62	0.62
<i>Panel D: Union-level characteristics (2014, baseline)</i>			
Avg. size of union board	18.8	24.3	17.3
Avg. share of women in board	0.23	0.23	0.22
Share with female president or vice president	0.17	0.18	0.17

Notes: Table shows descriptive statistics for the sample of establishment-union pairs negotiating firm-level CBAs registered in *Sistema Mediador* between 2012 and 2017. All CBAs are valid, non-amendment, firm-level agreements that have a union counterpart with information on 2012 union central affiliation. We additionally drop contracts signed by more than one union if these unions have different CUT affiliation in 2012 (fewer than 0.33% of CBAs). On the signing establishment's side, we restrict to CBAs where the employer appears in RAIS and has active employees in 2014. Treated units are those where the union counterpart was affiliated to CUT in 2012. See Appendix C for more details. The starting sample described in Panel A has observations at the pair-year level for years when CBA negotiations occurred, i.e., the new contracts panel. Statistics in Panel B are averages across these pair-year observations. Panels C and D use unique establishment and union observations in the baseline year (2014), respectively.

within the same time period comparing treated and comparison establishments. The identifying assumption is that outcomes would have evolved in parallel at treated and comparison units absent the CUT reform, conditional on covariates. We assess the plausibility of this assumption by testing the null hypotheses for $\beta^{2012} = 0$ and $\beta^{2013} = 0$.

To summarize the average post-period impact of the CUT reform we run a “pooled” version of the above regression, which amounts to replacing the full interaction of D_i and δ_t with the simple interaction $D_i \times \delta_{t \geq 2015}$. Also, to make treatment effects in worker-level regressions interpretable as establishment-level averages, we weight each incumbent worker by the inverse of (own-gender) employment at their baseline employer (Jäger et al., 2021). Finally, it is worth noting that outcomes that may be considered downstream effects resulting from the changes in amenities (e.g., wages, retention, etc.) are unscaled by the amenity change since we do not directly observe the value workers assign to said amenities.

4 Impact of the CUT Reform

This section presents our main results. We start by analyzing how the CUT reform affects amenities, finding disproportionate gains for women both on paper and in practice. We then explore whether women value these changes by analyzing the impact of the reform on revealed preference measures of firm value, i.e., retention and job queues. We end the section by evaluating whether employment, wages, and firm profits are impacted by the reform in order to determine potential tradeoffs from the improvement in female-centric amenities.

4.1 Amenities: On Paper and In Practice

Negotiated amenities To formally explore the treatment effect of the CUT reform on female and male-centric clauses, Table 3 reports estimates from the pooled version of Equation (3). We summarize the effects below, comparing them to the average value of the dependent variable Y_{it} among the treated at baseline.

Columns 1 and 6 show that female clauses rise in number (intensive margin), in incidence (extensive margin), and as a share of all clauses. On the intensive margin (Panel A), the number of intuitive female clauses rises by 0.157 (SE 0.013)—a 17% increase relative to baseline. Using the data-driven classification, the number of female clauses rises by 0.301 (SE 0.021)—a 19% increase relative to baseline. These effects are not explained by a mere increase in the number of clauses for some clause type already being provided, e.g., going from 1 to 5 clauses on maternity leave. Panel B uses the sum of indicators (rather than counts) of the relevant clause types and we find a 12-18% increase relative to baseline. Hence, the

space of female-centric clause types being provided expands in response to the CUT reform.

We also find that the CUT reform increases the likelihood that a female clause exists in the CBA (Panel C). That is, on the extensive margin, treated contracts are 1.7pp (SE 0.003) more likely to have any intuitive female clause—a 5% increase relative to baseline. Using the data-driven approach, the extensive margin effect is 3.4pp (SE 0.003)—a 10% increase relative to baseline. In Panel D, we see that the share of intuitive female clauses also rises by 0.5pp (SE 0.001)—a 10% increase relative to baseline. The share effect with the data-driven classification is 2.1pp (SE 0.001)—a 30% increase relative to baseline.

Columns 2-5 show that all four themes of intuitive female-centric clauses rise, with 76% of the intensive margin effect driven by clauses on leaves and maternity/childcare. Hence, CUT-driven changes in amenities are likely to impact workers who value provisions on paid leaves following maternity, adoption, and miscarriage, as well as childcare assistance, maternity protections, and policies for dependents. In addition, the largest increase relative to baseline occurs among the harassment clauses (45%) due to the fact that these clauses are very uncommon at baseline.

Interestingly, we find some evidence that unions trade off men’s interests in favor of women’s, but only negligibly. Both the extensive margin and share of male amenities fall by small amounts: by 0.1pp (SE 0.003) compared to 46% at baseline for the former, and by 0.3pp point (SE 0.002) relative to 14% for the latter (Column 7). While there is a slight increase on the intensive margin, this is overshadowed by the rise in female-centric clauses.³⁰ Specifically, our preferred measure of the ratio between female-to-male clauses increases by 21% over baseline (Column 8). In summary, the composition of clauses is shifting to become more female-oriented, driven by an increase in female-centric clauses.³¹

Figure 4 shows parallel trends in female and male-centric amenities in the pre-period by plotting the time path of treatment-control differences, i.e., β^t in Equation (3). This implies that we are unable to reject parallel trends, suggesting that outcomes in the control group can be used to gauge what would have happened in the treatment group without the reform.³² Moreover, the rise in female amenities occurs right in 2015 (also present for all margins—see Figure B8). This is not the case for male-centric clauses, where the increase in the intensive margin occurs until 2017, indicating that this particular result is not a direct result of the CUT reform but perhaps a downstream response to its effects.

Turning to the question of *where* union priorities matter most for female amenities,

³⁰Only the increase in female clauses is robust to clustering standard errors at the union level (see Table A4).

³¹These results are robust to changes in the data-driven classification, the granularity of geography-year fixed effects, and conditioning on pairs that have coverage in 2014 (see Tables A5, A6 and A7).

³²Plotting the raw averages of female-centric clauses (see Figure B7) provides further evidence for the parallel trends assumption and shows that the effects are driven by increases among the treated group.


Table 3: Effect of CUT Reform on Negotiated Amenities

	Intuitive definition (female clauses)					Data-driven		
	All (1)	Leave (2)	Maternity (3)	Harassment (4)	Flexibility (5)	Female (6)	Male (7)	F/(F+M+1) (8)
<i>Panel A: Intensive margin (number)</i>								
$D_i \times \delta_{year \geq 2015}$	0.157*** (0.013)	0.078*** (0.006)	0.042*** (0.004)	0.009*** (0.001)	0.028*** (0.008)	0.301*** (0.021)	0.130*** (0.029)	0.032*** (0.002)
Mean outcome	0.95	0.25	0.24	0.02	0.44	1.58	2.55	0.15
<i>Panel B: Intensive margin (sum of indicators)</i>								
$D_i \times \delta_{year \geq 2015}$	0.123*** (0.010)	0.047*** (0.004)	0.042*** (0.004)	0.008*** (0.001)	0.027*** (0.004)	0.154*** (0.014)	0.067*** (0.017)	
Mean outcome	0.70	0.18	0.21	0.02	0.30	1.26	1.58	
<i>Panel C: Extensive margin</i>								
$D_i \times \delta_{year \geq 2015}$	0.017*** (0.003)	0.012*** (0.002)	0.020*** (0.002)	0.008*** (0.001)	0.022*** (0.003)	0.034*** (0.003)	-0.001 (0.003)	
Mean outcome	0.31	0.12	0.15	0.02	0.23	0.36	0.46	
<i>Panel D: As a share of all clauses</i>								
$D_i \times \delta_{year \geq 2015}$	0.005*** (0.001)	0.001*** (0.000)	0.001*** (0.000)	0.000*** (0.000)	0.003*** (0.001)	0.021*** (0.001)	-0.003*** (0.002)	
Mean outcome	0.05	0.01	0.01	0.00	0.03	0.07	0.14	
Observations	600,960	600,960	600,960	600,960	600,960	600,960	600,960	600,960


Notes: Table reports the coefficients for DID regressions—see Equation (3)—estimating the effect of the CUT reform on the female-centric and male-centric amenities included in CBAs. Panel A uses the total number of clauses per pair-year as an intensive margin measure. Panel B uses the sum of the corresponding clause type indicators, capturing how the space of female (male) clauses grows or shrinks. Panel C uses an indicator for pair-year observations with at least one corresponding clause as an extensive margin measure. Panel D uses the share of corresponding clauses with respect to the total contract clauses, capturing how the composition of CBAs change. Under each panel we report the mean of the dependent variable among the treated at baseline (2014). The sample is the filled panel of establishment-union pairs by year. All columns control for pair fixed effects, as well as time-varying state and industry fixed effects. Standard errors are clustered at the establishment level.

Figure 4: Effect of the CUT Reform on Female- and Male-Centric Amenities


(a) Female clauses: intensive margin




(b) Male clauses: intensive margin



(c) Female clauses: as a share of clauses



(d) Male clauses: as a share of clauses



Notes: Figures show estimates of the δ_t coefficients for $t \in [2012, 2017]$ (with 2014 omitted) from the DID specification in Equation (3) on the intensive margin (top figures) and shares (bottom figures) of female-centric (left side) and male-centric (right side) clauses, defined using the data-driven method. All figures use the filled panel. Confidence intervals at a 95% level are reported. Standard errors are clustered at the establishment level.

Table 4 shows that the largest effects are at establishments where women could not already advocate for themselves either as workers or in the union leadership. Column 1 shows the baseline results—which in light of Figure 3b can be interpreted as moving the mean treated establishment from having female centric clauses similar to those at minority female establishments to those observed at establishments with over 80% women. Columns 2-4 reveal that the effects of the CUT reform are significantly stronger where women lacked representation at baseline, whether that be a) few women in the establishment (Column 2); b) few women in the union board (Column 3); or c) no women president or vice-president (Column 4). Quite convincingly, the relation between the treatment effects and the share of women in the establishment is decreasing and monotonic (see Figure B9).

In terms of mechanisms, we cannot completely rule out that union leadership has an impact above and beyond the shift in priorities toward women among CUT-affiliated unions. Specifically, the CUT reform does have a small effect on the share of women at the local union boards. It does not, however, increase the likelihood that a woman takes the top leadership role of the union, i.e., union president or vice-president (see Figure B10). Regardless, women who were elected to the local union board—as well as the CUT board—are leaders who may have been instrumental in ensuring that the new priorities were implemented.

On a final note, it is worth highlighting that CBA clauses are equilibrium outcomes resulting from negotiations between unions and employers. As such, our results at the very least show that employers can agree to signing off on female-friendly amenities. Our upcoming analyses speak to whether this occurs because changes on paper are not reflected in practice, employers face adjustments on other margins (e.g., wages), or the proposed changes bring about Pareto improvements.

Actual amenities To show that changes on paper translate into practice, we draw on the text of female-centric clauses to identify three establishment-level outcomes that they can directly affect: (i) whether women are managers—corresponding to equal opportunity clauses; (ii) if women actually take longer maternity leaves—corresponding to leave extension clauses; and (iii) if women are less likely to leave their employer after maternity leave—corresponding to job protection clauses.

Per Figure 5, the reform positively affects actual outcomes along all three dimensions. The share of women among managers at treated establishments rises by 2% relative to baseline. Women also take longer maternity leaves, as reflected by the 14% increase in the share of women on maternity leave taking leaves longer than the state-mandated 120 days.³³

³³Matching data from *Empresa Cidadã*—a government program incentivizing extended maternity leave at firms—to our establishment sample, we find a 7.1% increase in take-up relative to baseline (p-value<0.01).


Table 4: Heterogeneity by Female Representation

	Full interaction: $D_i \times \delta_{year \geq 2015} \times H_i$			
	Baseline (1)	$H_i = \text{low } \%$ women in estab. (2)	$H_i = \text{low } \%$ women in union (3)	$H_i = \text{no}$ woman Pres/VP (4)
<i>Panel A: Intensive margin</i>				
$D_i \times \delta_{year \geq 2015}$	0.301*** (0.021)	0.139*** (0.028)	0.002 (0.038)	-0.058 (0.044)
$D_i \times \delta_{year \geq 2015} \times H_i$		0.307*** (0.040)	0.362*** (0.041)	0.396*** (0.049)
Sum of coefficients		0.446	0.364	0.338
p-value		[0.000]	[0.000]	[0.000]
Mean outcome	1.58	1.58	1.58	1.58
<i>Panel B: As a share of all clauses</i>				
$D_i \times \delta_{year \geq 2015}$	0.021*** (0.001)	0.009*** (0.001)	0.005*** (0.002)	-0.004** (0.002)
$D_i \times \delta_{year \geq 2015} \times H_i$		0.022*** (0.002)	0.020*** (0.002)	0.030*** (0.002)
Sum of coefficients		0.031	0.025	0.025
p-value		[0.000]	[0.000]	[0.000]
Mean outcome	0.07	0.07	0.07	0.07
Observations	600,960	600,960	592,344	592,344

Notes: Table tests for heterogeneity in the effect of the CUT reform on female-centric clauses (data-driven approach) according to the baseline representation of women among workers (column 2) and within union boards (columns 3-4). The dummy to test for heterogeneity in the effects (H_i) is fully interacted with the treatment dummy (D_i) and the post-period dummy ($\delta_{year \geq 2015}$). The table only reports the coefficients on the effects that determine the treatment effect for the baseline group ($H_i = 0$) and the differential effect relative to the baseline group—with the sum of both coefficients representing the treatment effect for the group of interest ($H_i = 1$). In column (2), H_i is an indicator for whether the share of women workers is below the median across our sample in 2014 (around 1/3). In column (3), H_i is an indicator for whether the share of women in union boards is below this 1/3 threshold in 2014. In column (4), H_i is an indicator for whether there is no women president or vice-president in the local union board as of 2014. All regressions use the filled panel sample and includes establishment-union pair fixed effects as well as time-varying state and industry fixed effects. Standard errors are clustered at the establishment level.

Moreover, although women at treated establishments are taking longer maternity leaves, they are no less likely to return to their employer after the leave, suggesting longer periods of post-maternity job security. Thus, changing union priorities not only improves negotiated clauses for women, it also leads to actual changes in the workplace environment.

Figure 5: Changes in Firm Environment



Notes: Figure reports results from four separate establishment-level DID regressions in Equation (3), with treatment effects reported relative to the mean among the treated at baseline (in percentage terms). The outcome variables are: 1) the share of women among managers; 2) the share of women on maternity leave who remain on leave longer than than the state-mandated 120 days (i.e., extended maternity leave); 3) the share of women taking maternity leave who remain employed at the employer where they took maternity leave (i.e., return from maternity leave); and 4) the share of workers taking leave due to a workplace injury. Each regression includes establishment fixed effects, industry-year fixed effects, and microregion-year fixed-effects. Two stars indicate significance at the 5% confidence level, while three stars indicate significance at the 1% level. Standard errors are clustered by establishment.

On a similar vein, we draw on the fact that male-centric amenities focus on workplace safety to explore whether less male-oriented CBAs led to more physically dangerous workplaces. The last bar in Figure 5 shows that this is not the case—as captured by the share of workers taking work-related injury leave. If anything, the point estimate implies a 3% decrease in this outcome. This suggests that there are no negative implications (at least on this dimension) from the small trade offs against men’s interests we observe on paper.

4.2 Revealed preference changes in firm value

The improvements in actual amenities we can document are limited to what is observable in RAIS. To get a more comprehensive answer as to whether workers value the changes in CUT workplaces, we analyze the reform’s impact on two revealed preference measures of job quality: retention and job queues.

Retention Retention serves as a revealed preference measure of the relative attractiveness of an employer (Krueger and Summers, 1988). Using the incumbent worker sample, we explore whether women and men are more likely to remain at their baseline employer. Figure 6a shows that women are 0.8pp (SE 0.003) more likely than men to stay at their baseline job relative to the gender difference for incumbents at comparison establishments. The overall retention effect for incumbent women is 1.8pp, which is the same magnitude as that estimated among women of childbearing age shown in Figure 6b. The baseline retention rate in treated establishments is 71% for women—thus, the treatment effect represents a 2.5% increase in women’s retention rate.

Since male-centric clauses fall as a share of all clauses, it is possible that men’s value from employment is also falling. However, retention among incumbent men also experiences a slight uptick amounting to 1.0pp (see Table A8). This effect represents a 1.5% increase in men’s retention rate, which is strong evidence against the hypothesis that men are worse off due to the prioritization of women by CUT. Thus, the reform disproportionately improves the work environment for incumbent women without pushing men to other jobs.

Higher retention need not imply that workers’ value of employment has increased since this may be driven by fewer involuntary separations (fires). However, we find that our retention effects are likely coming from fewer voluntary separations (quits). Specifically, since an incumbent is observed in RAIS either if they (i) stay at their baseline employer or (ii) make a *voluntary* employment-to-employment transition, the difference between effects on “employed in the formal sector” and “stay at baseline employer” tell us how voluntary transitions changed. Results in Table A8 show a 1.1pp (0.8pp) decrease in voluntary transition among incumbent women (men).³⁴ This further supports the hypothesis that workers value the changes in the work environment at treated establishments.³⁵


Job queues Another measure of revealed preference value is longer job queues (Holzer et al., 1991). Because we do not directly observe job applications, we use workers in the pro-

³⁴We simply use $0.7\text{pp} - 1.8\text{pp} = -1.1\text{pp}$ for women and $0.2\text{pp} - 1.0\text{pp} = -0.8\text{pp}$ for men.


³⁵Even in the cross-section of establishments in our sample for which we have PageRank values, there is a “CUT premium” for both men and women on this measure of firm value ($\approx 4\text{--}7$ log points).

Figure 6: Revealed Preference Measures of Firm Value


(a) Incumbent retention: women-men differential (b) Incumbent women’s retention: age 20-35



(c) Share of women among probationary workers



(d) Share of women in workforce



Notes: Figures test for revealed preference measures of whether women value the changes induced by the CUT reform in treated establishments. Top figures look at retention among incumbent workers, i.e., an indicator for whether the worker is observed at their baseline (2014) employer in year t . To make treatment effects in worker-level regressions interpretable as establishment-level averages, we weight each incumbent worker by the inverse of (own-gender) employment at their baseline employer. Figure 6a reports the differential in retention for women relative to men using a triple DID regression, which includes worker fixed effects, industry-year-gender fixed effects, microregion-year-gender fixed effects, and tenure-year-gender fixed effects. Figure 6b shows effects from the baseline DID specification in Equation (3) among women ages 20-35, which includes worker fixed effects, industry-year fixed effects, microregion-year fixed effects, and tenure-year fixed effects. Bottom figures look at the gender composition of spells observed at the establishment level using the DID specification in Equation (3). The outcome in Figure 6c is the share of women among probationary workers, i.e., those whose tenure at the establishment does not exceed 3 months. The outcome in Figure 6d is the share of women among all spells observed. Regressions include establishment fixed effects, industry-year fixed effects, and microregion-year fixed effects. Confidence intervals at a 95% level are reported. Standard errors are clustered at the establishment level.

bationary period (i.e., first 3 months of tenure) as a proxy measure. Per Brazilian labor law, employers can terminate probationary workers without any severance pay. Such contracts are thus used to screen applicants for jobs.

Using the establishment sample, we find that women disproportionately join the job queue at treated establishments. Figure 6c shows that women’s share among probationary workers rises by 0.6pp (SE 0.003), representing a 1.7% increase over baseline. This effectively translates into a 0.2pp increase in the share of workers who are women—see Figure 6d.

Although these precise estimates show more interest among women to work at treated establishments, the magnitudes are small. There are three factors that are likely dampening our estimates. The first (as discussed earlier) is the fact that we cannot directly observe changes in the value of amenities to scale our effects.³⁶ The second is information frictions. That is, even though CBAs are public information, most workers may not know which amenities their union has negotiated for them—let alone what has been negotiated at other establishments.³⁷ The third is potential screening during the hiring process, such that effects at the probationary stage are already muted. Putting these issues aside, in Section 5 we quantify by how much workers are better off after the CUT reform with a welfare model.

4.3 Possible tradeoffs

How are these improvements in female-centric amenities paid for? Firms might lower employment or change the composition of their workforce (Summers, 1989). They could also lower wages, as predicted by compensating differences (Rosen, 1986). Alternatively, these female-friendly jobs might come at the expense of lower profits. In Table 5 we explore these explanation using the establishment sample.

Employment Panel A shows that establishments do not lower employment in response to the reform. Column 1 reports a statistically insignificant 0.2 log point decline in employment. These estimates rule out negative effects larger than 1.6 log points at a 95% confidence level. While employment remains unchanged, Columns 2-3 show that treated establishments observe more female worker spells.³⁸ Therefore, there is no evidence that firms employ fewer workers, or fewer women in particular, to pay for the provision of female-centric amenities.

Churn also seems unaffected by the reform. Column 4 reports a statistically insignificant

³⁶Estimating PageRank values is very demanding in terms of the data, making pre- and post-period estimates of the measure infeasible given just 3 years of data in each period.

³⁷For example, we heard from an economics professor who thought she would have extended maternity because a co-worker at the same institution had already done so. However, this professor’s location was not covered by the same CBA as her colleague and was therefore ineligible for the extension.

³⁸These are the same results discussed in Section 4.2 concerning job queues.

0.9 log point decline in employment. These estimates rule out negative effects larger than 1.9 log points at a 95% confidence level. Hence, stable employment levels are not the result of mirroring increases (decreases) in both hiring and separations. While we see more women among new hires, there is an identical effect on the share of women among separations (Columns 5-6). It therefore seems that the women who are flocking to treated establishments are not treated any differently by employers in terms of firing or retention.

Since employers observe more female worker spells, perhaps they are able to select more productive workers to help pay for the new amenities. We do not find evidence for this argument (see Table A9). There is no indication that a higher proportion of the female workforce is poached from other employers. Moreover, the mean age, tenure, contracted hours, and schooling of the female workforce is unaffected by the reform. Hence, with the data available, we are unable to find evidence that the female workforce (on average) becomes more productive. However, we cannot rule out that the women who value these amenities may be more productive as a result of the changes in the workplace environment.

Wages If amenity improvements operate in a compensating differences world, we should observe a disproportionate decline in women’s wages.³⁹ Because employers cannot cut nominal wages for existing employees without the approval of the union, wage adjustments might only realize for new workers. We therefore use the mean of log wages by gender and by new versus established workers (i.e., tenure below/above 12 months) as dependent variables.

Panel B shows that establishments do not lower wages in any meaningful way in response to the reform. All point estimates are negative but very small and precise—the largest is -0.6 log points (SE 0.003) for new male workers. We can rule out negative effects larger than 1.2-1.3 log points (for new workers) and 0.7-0.8 log points (for established workers) at a 95% confidence level. Given the similar point estimates across men and women, gender wage gaps are also unaffected by the reform. Thus, we find no evidence that mean wages fall to cover for the costs of providing more female-centric amenities.

There are two important caveats to these results on mean log wages. First, the average worker may not capture the workers whose wages are affected by unions. We therefore extract the wage adjustment clauses negotiated in the CBAs covering these establishments and see whether the percentage adjustments are affected by the reform. We get a point estimate of 0.032pp (SE 0.021) which allows us to rule out decreases in the negotiated wage adjustments larger than 0.009pp. Second, changes in the composition of the workforce may dampen effects on wages. However, using the incumbent worker sample we also find

³⁹Our results showing an increase in women’s value of employment at treated establishments (using revealed preference measures) already rule out—more comprehensively than null effects on wages—a pure compensating differences story for the amenity gains.

Table 5: Impact of CUT Reform on Establishment-Level Outcomes

<i>Panel A: Employment</i>						
	Log employment (1a)	Share women [workforce] (2a)	Share women [probation] (3a)	Log hires (4a)	Share women [hires] (5a)	Share women [separations] (6a)
$D_i \times \delta_{year \geq 2015}$	-0.002 (0.007)	0.002** (0.001)	0.006** (0.003)	-0.009 (0.009)	0.004* (0.002)	0.004** (0.002)
Mean outcome	4.044	0.369	0.357	3.034	0.366	0.360
Observations	353,626	353,626	275,879	325,823	325,823	332,506
<i>Panel B: Wages</i>						
	Mean log(w) [women; $t > 12$] (1b)	Mean log(w) [men; $t > 12$] (2b)	Mean log(w) [women; $t \leq 12$] (3b)	Mean log(w) [men; $t \leq 12$] (4b)	Mean gender wage gap (5b)	CBA wage adjustments (6b)
$D_i \times \delta_{year \geq 2015}$	-0.004 (0.002)	-0.003 (0.002)	-0.005 (0.004)	-0.006* (0.003)	-0.001 (0.002)	0.032 (0.021)
Mean outcome	7.460	7.627	7.174	7.311	-0.150	0.781
Observations	323,271	329,960	260,956	289,334	334,562	123,432
<i>Panel C: Profits</i>						
	Log wage bill (1c)	Establishment exit (2c)	Profit margin (3c)			
$D_i \times \delta_{year \geq 2015}$	-0.010 (0.008)	-0.003 (0.003)	0.702 (1.167)			
Mean outcome	11.431	0.087	7.759			
Observations	351,593	61,716	2,874			

Notes: Table reports the coefficients for the establishment-level DID regression from Equation (3), comparing treated to comparison establishments on employment, wage, and profit outcomes. An establishment is treated if the union with which it negotiates is affiliated to CUT in 2012. Each regression includes establishment fixed effects, industry-year fixed effects, and microregion-year fixed effects. Panel A uses all spells observed at an establishment in a given year. The terms in brackets indicate the subsample among which the share of women is calculated, i.e., among all workers, among workers in probation, among hires, and among separated workers. Panel B uses workers' main spell in a given year. The terms in brackets indicate the subsample among which the mean of log wages is calculated, i.e., tenure > 12 months and tenure ≤ 12 months for either women or men. Panel C studies three imperfect measures of firm profits. Standard errors are clustered by establishment and reported in parentheses.

incredible precise null effects on wages (see Table A8). Hence, the bulk of the evidence does not support the hypothesis that lower wages compensate for female-friendly jobs.

Profits If workers are not paying the price for the improvements in amenities through either wages or employment, perhaps the burden falls on firms through lower profits. We provide some empirical evidence of null effects on profits, but this evidence is limited because profits are not directly observed in RAIS. Nonetheless, from a theoretical perspective, placing the CUT reform in the context of collective bargaining models also suggests that one would not expect profits to decrease.

Panel C shows no effects on firm profits in response to the reform, as measured by three different measures. First, the wage portion of labor costs (i.e., wage bill) is unchanged. We get a point estimate of -1.0 log points (SE 0.008) which allows us to rule out increases in the wage bill larger than 0.6 log points. Second, we observe zero treatment effects on establishment exit. Exit is a non-trivial margin of adjustment, e.g., 8.7% of treated establishments exit between 2014 and 2017. Third, using establishments with profit margin data available on Orbis in both the pre- and post-periods, we obtain a point estimate of 0.70pp (SE 1.17). As such, we rule out decreases in profit margins larger than 1.59pp in this restricted sample.

From a theoretical perspective, lower firm profits would require that unions capture a larger share of the surplus being negotiated. But there is no reason to believe this happens following the CUT reform, since it merely involved shifting priorities in favor of women rather than an increase in unions' bargaining power relative to employers. If anything, unions were in tougher bargaining positions given that Brazil fell into an economic recession from mid-2014 to 2016. Moreover, the power of CUT-affiliated unions was particularly vulnerable given the impeachment of President Dilma Rousseff from the left-wing Workers' Party—a process that started in December 2015 and ended with centrist Michel Temer taking power in May 2016. Therefore, the context we study in light of economic theory would predict that the CUT reform should not impact profits.

4.4 Discussion

The CUT reform that pushed union leaders to prioritize women's needs in collective bargaining improved the work environment for women relative to men, both on paper and in practice, as revealed through higher retention and longer job queues. This is consistent with a model of utility positing in the labor market, where the posted utility for women at treated establishments increases without generating a decrease for men. Moreover, these changes do not come at the expense of either women's or men's wages and employment. While men may be losing some rents through a smaller share of amenities in treated contracts, this does

not appear to be consequential in practice, suggesting that men are not marginal to these changes or that they may even reap positive spillovers from the new workplace environment.

These results imply that shifting union priorities can improve outcomes for the group being targeted and, more importantly, that these changes need not be redistributive. In our setting, female workers benefit neither at the expense of male workers nor at the expense of the employer. There are at least two models that could explain these results.

In one model, unions prevent employers from setting compensation to the preferences of the marginal worker (Dube et al., 2021; Lagos, 2021). But the compromise between union and firm objectives resulting from the collective bargaining process may not lead to a Pareto efficient compensation bundle.⁴⁰ As such, a shift in union priorities could lead to a win-win situation. In another model, firms (or unions) are behavioral and may consequently fail to implement all Pareto improvements (Heidhues and Kószegi, 2018). Once CUT pushes unions to focus on women’s issues, the improvements become salient and are implemented.

5 Quantifying the Welfare Effect of the CUT Reform

The CUT reform increases female-centric amenities and makes CUT establishments more valuable to women. By how much does women’s welfare increase? What about the reform’s impact on men’s welfare? We briefly describe our approach here with details in Appendix E.

Approach and Intuition To quantify how much the CUT reform improved workers’ welfare we adopt a sufficient statistics approach that (i) relies on a few parameters of interest that can be easily computed from the data; and (ii) does not take a stance on how different amenities enter workers’ utility. In particular, we adapt a framework used to evaluate changes in consumer welfare from introducing new or improved product varieties (Feenstra, 1994; Redding and Weinstein, 2016) to our labor market setting.

For tractability, we assume that workers have CES preferences over firms, as is common in consumer welfare calculations (Feenstra, 1994; Atkin et al., 2015). As shown by Anderson et al. (1992), a key advantage of CES is that it generates the same labor supply to firms as obtained by aggregating workers’ discrete choices over where to work. While it may seem strange to think of workers as consuming firms, workers choosing where to work based on where they are happiest is commonly used to model the labor market (Card et al., 2018; Sorkin, 2018; Berger et al., 2022; Lamadon et al., 2022). In Appendix E we microfound CES demand using such discrete choices.

⁴⁰This could be due to information frictions during negotiations, contracting frictions, mismatch between union and worker interests, among others.

As in the consumer setting, where changes to the price index—i.e., how much more (or less) expensive it is to buy one util—measure welfare changes for the representative consumer, in this model, welfare changes for the representative worker are measured by changes to the wage index—i.e., how much more (or less) she is paid to work one disutility-weighted hour.

Under CES, four sufficient statistics pin down the CUT reform’s effects on welfare. First, welfare increases with the share of labor income at CUT establishments, representing workers flocking to these firms in search of better amenities. Second, welfare is also higher as substitution across establishments becomes more inelastic because this implies that it takes a larger improvement in amenities to draw workers away. Third, the change in welfare is larger if the workers drawn to CUT firms are coming from non-CUT firms that were not initially valued highly by workers. Finally, welfare is higher if wages at non-CUT establishments rise, reflecting a pro-competitive effect of the reform.

Model In each period, a representative household with CES preferences over firms is willing to work a fixed amount of (dis)utility-weighted hours. They choose labor supply to each firm to maximize total income subject to the hours constraint

$$\max_{\{n_{jt}\}} \sum_{j \in \mathcal{J}_t} w_{jt} n_{jt} \quad s.t. \quad \left[\sum_j (b_{jt} n_{jt})^{\frac{1+\eta}{\eta}} \right]^{\frac{\eta}{\eta+1}} = N, \quad (4)$$

where \mathcal{J}_t is the set of firms operating at time t , n_j is the number of hours worked at firm j , w_j is the wage at firm j , b_j is the “taste” for working at firm j , and η is the elasticity of substitution across firms. Factors other than the wage that affects workers’ utility of working at firm j are captured by b_j , e.g., worsening amenities implies higher b_j . For simplicity we do not model the firm side and assume a utility-posting world. Jobs are not rationed and firms accept any worker that wants to work there.

The wage index measures how much the representative household is paid to work one more disutility-weighted hour, serving as a measure of welfare of the representative worker

$$\tilde{W} = \left[\sum_{j \in \mathcal{J}} \left(\frac{w_j}{b_j} \right)^{1+\eta} \right]^{\frac{1}{1+\eta}}$$

This is analogous to a consumer price index measuring how costly it is to purchase one util. An increase in the wage index from one period to the next represents a welfare improvement, captured by the ratio

$$\phi_{t-1,t} = \frac{\tilde{W}_t}{\tilde{W}_{t-1}}$$

The CUT reform changes amenities, or taste shifters b_{it} , at treated establishments. Since these b_{it} are unobserved, estimating welfare changes is not straightforward.

However, when consumer preferences are CES, welfare changes depend only on pre- and post-reform observed wages and employment at CUT and non-CUT establishments. Formally, the change in welfare is

$$\ln \phi_{t-1,t} = -\frac{1}{1+\eta} \ln \left(\frac{\lambda_t}{\lambda_{t-1}} \right) - \frac{1}{1+\eta} \ln \left(\frac{\bar{S}_t^*}{\bar{S}_{t-1}^*} \right) + \ln \left(\frac{\bar{w}_t^*}{\bar{w}_{t-1}^*} \right) \quad (5)$$

where λ_t is the overall share of labor income in t at non-CUT firms, \bar{S}_t^* is a geometric average of the share of labor income at each non-CUT firm in t , and \bar{w}_t^* is a geometric average of period t wages at non-CUT firms. The asterisk $*$ denotes that operations are taken over non-CUT firms.

Changes in welfare depend on three terms, as per Equation (5). The first, “variety-adjustment” term $\left(\frac{\lambda_t}{\lambda_{t-1}} \right)^{-\frac{1}{1+\eta}}$ is the ratio of the share of total labor income at non-CUT firms after versus before the reform. This ratio captures welfare changes through a revealed preference logic: workers substitute towards CUT firms if their amenities improve (lower b), lowering the share of the labor income at non-CUT firms and increasing welfare. The magnitude of this change depends on the elasticity of substitution across firms. If workers are inelastic (η is low), the same move towards amenity-improving CUT-firms implies a larger welfare increase because it takes a bigger improvement in amenities to draw workers away.

The term $\left(\frac{\bar{S}_t^*}{\bar{S}_{t-1}^*} \right)^{-\frac{1}{1+\eta}}$ captures the heterogeneity in labor income at non-CUT firms: welfare increases more as CUT firms draw away workers from less valued non-CUT firms, thereby increasing dispersion and lowering the geometric mean. As in the “variety-adjustment” term, the implied effects are larger as workers become more inelastic. The final term $\left(\frac{\bar{w}_t^*}{\bar{w}_{t-1}^*} \right)$ represents a change in wages at non-CUT firms, possibly as a pro-competitive response to the reform. As these “outside” wages increase so too does welfare.

Estimation To get at welfare changes by gender, we estimate Equation (5) separately for men and women. We use the establishment sample from Section 4.3 and calibrate an estimate of the cross-firm elasticity of substitution (η) from Felix (2022). Years 2012-2014 comprise the pre-reform period ($t - 1$) and 2015-2017 comprise the post-reform period (t).

We estimate the log change in \bar{w}^* and in \bar{S}^* with average changes across non-CUT establishments between $t - 1$ and t . That is, we run the following regression

$$y_{jt} = \alpha + \beta Post_t + \mu_j + \epsilon_{jt} \quad (6)$$

using as dependent variables the average of log earnings by establishment ($\log w_{jt}$) or the log of the establishment share of labor income among non-CUT firms ($\log s_{jt}$). The specification includes establishment fixed effects μ_j . Hence, the coefficient of interest is β which captures the average within-establishment change in the dependent variable between between $t - 1$ and t . Bootstrapped standard errors are clustered by establishment.

To estimate the change in λ we take a first order approximation of λ_t around λ_{t-1} . This allows to map the market-level change in the share of labor income at CUT establishments to changes in quantities that we can estimate through establishment-level regressions as in Equation (6). We refer the reader to Appendix E for more details.

Results Table 6 reports our results. Women’s welfare increases by 0.059 log points (or 6.1%) consistent with our reduced form results on higher retention and a higher share of women among new workers at CUT establishments.⁴¹ More than half of this increase is accounted for by workers’ movement across firms. That is, after the reform, women are more likely to work at CUT establishments, accounting for 15% of the welfare increase (i.e., a 1.8% rise in the share of CUT wage bill). In addition, the dispersion in the labor income across non-CUT firms is rising (i.e., S^* is falling), accounting for about 48% of the increase in welfare.

The remaining 37% of the welfare change is due to wage increases at non-CUT establishments after the CUT reform. Through the lens of this model, the increase in wages at non-CUT establishments should be seen as a pro-competitive effect of the CUT reform, also improving women’s welfare. We recognize, however, that the increase in real wages at non-CUT establishments after 2015 might be driven by factors other than the CUT reform. We therefore see the change in welfare due to worker movements across firms in our sample—which amounts to a 3.8% increase—as a more credible estimate of the welfare change for women that was driven by the CUT reform.⁴²

For men, welfare is slightly higher (1.3%) but remains essentially the same if one only considers the changes in welfare due to worker movements across firms in our sample (0.2%). Thus, the CUT reform improves welfare for women without reducing it for men, suggesting a possible Pareto improvement from directing more of the union’s attention towards women’s needs.⁴³ In more general terms, shifting the group-specific priorities of agents setting com-

⁴¹As intuited from the model, workers’ elasticity of substitution across firms amplifies (or dampens) the welfare effect of changes in employment across firms. For other reasonable values of η from the literature, ranging from 0.1 (Staiger et al., 2010) to 10.9 (Berger et al., 2022), the welfare increase for women ranges between 2.8% and 9.5%.

⁴²In Table A10 we also compute welfare separately for workers of child-bearing ages (i.e., between ages 20 and 35), noting that results are qualitatively very similar to those unrestricted by age.

⁴³As discussed before, the fact that CUT-affiliated union’s bargaining power is unlikely to be increasing

Table 6: Welfare Estimation

		Women (1)	Men (2)
$\ln\phi_{t-1,t}$		0.059 (0.007)	0.013 (0.005)
<i>Contribution by component:</i>			
Wage bill	$\ln(\lambda_{t,t-1}) - \ln(\lambda_{t-1,t})$	15%	22%
Dispersion	$\ln(\bar{S}_t^*) - \ln(\bar{S}_{t-1}^*)$	48%	-4%
Wages	$\ln(\bar{w}_t^*) - \ln(\bar{w}_{t-1}^*)$	37%	82%
η (calibrated)	1.015		
N establishments		60,651	60,651
N establishments in $\Omega_{t,t-1}$		47,195	47,195

Notes: Table reports the estimated welfare change for men and women. It also reports the contribution to the overall effect by each of the three components that make the welfare index, namely the Feenstra “new varieties” term $\ln(\lambda_{t,t-1}) - \ln(\lambda_{t-1,t})$, the change in the geometric average of the labor income shares of non-CUT firms $\ln(\bar{S}_t^*) - \ln(\bar{S}_{t-1}^*)$, and the change in the geometric average of the wages of non-CUT firms $\ln(\bar{w}_t^*) - \ln(\bar{w}_{t-1}^*)$. Standard errors in parenthesis come from the bootstrap procedure described in Appendix E.

pensation in the labor market need not result in a mere redistribution of surplus.

6 Conclusion

This paper shows that one reason that workplaces do not provide job features that are valuable to women is that decision-makers do not prioritize women’s preferences. Studying a top-down change in Brazil, where Latin America’s largest trade union federation, the *Central Única dos Trabalhadores*, adopted a bargaining plan more attentive towards women’s needs, we find that prioritizing women makes workplaces more female-friendly, both on paper and in practice, without leading to tradeoffs in wages, employment, or profits.

Just as political leaders’ priorities govern policy design (Chattopadhyay and Duflo, 2004; Pande and Ford, 2012), unions’ priorities govern what amenities firms provide, ultimately shaping within-firm inequality (Farber, 1986).⁴⁴ Prioritizing women caused a sharp increase in female-centric clauses in collective bargaining agreements, like those related to maternity leave and job protection, childcare allowances, and flexibility. As these contracts changed, so did the workplace, with women reaching managerial positions and taking longer maternity

during this period means that improvements for women are probably not financed through lower firm profits. Our reduced form results on exit, wages, and employment also support this claim.

⁴⁴A promising area of future study is to understand how the political economy of unions creates instances where they benefit or harm underrepresented groups of workers.

leaves. Women valued these changes, being less likely to leave and more likely to queue for jobs at CUT establishments. Although the reform may have made the firm environment less male-oriented, men were not marginal to these changes. Importantly, these improvements were not accompanied by lower wages, employment, or profits.

In sum, prioritizing women appears to usher in more efficient compensation for workers. The ex-ante inefficiency in the negotiated compensation bundle could be the result of frictions in collective bargaining or behavioral deviations from benchmark maximization (Bloom and Van Reenen, 2007; Massey and Thaler, 2013; Heidhues and Kőszegi, 2018). As such, the nature of the Pareto improvement remains an open question. For example, it may have involved a transfer of union rents to workers or perhaps these were sources of surplus that unions and employers were not considering. An alternative possibility is that the reform increased the total size of rents split between unions and employers. Turnover is typically costly to the firm, and our results suggest that separations among women decreased. Although we don't observe productivity, happier workers might also be more productive.

While gender differences in virtually all labor market outcomes have narrowed at a fast pace in the last century, more recently, closing gender gaps has proven harder (Goldin, 2014; Blau and Kahn, 2017, 2006). Policies and reforms aimed at changing women's representation in the workplace, such as those introducing gender quotas in firm boards of directors, have had mixed effects (Bertrand et al., 2018; Maida and Weber, 2020; Pande and Ford, 2012). Another possible lever to narrow gender gaps is collective bargaining. Past evidence has suggested that centralized pay setting, by compressing wage variation across industries and firms and by raising wages at the bottom of the earnings distribution, might in part explain why European countries have lower gender pay gaps than the United States (Blau and Kahn, 1996). We bring causal evidence to questions on the role of unions in shaping gender inequality and we argue that focusing only on wages does not give a complete picture of how collective bargaining might help reduce gender gaps. We consider a broader definition of worker compensation, that includes also amenities such as family allowances and leaves. We show that these might be an important lever for unions to change the workplace.

References

- Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High Wage Workers and High Wage Firms. *Econometrica*, 67(2):251–333.
- Adhvaryu, A., Molina, T., and Nyshadham, A. (2022). Expectations, Wage Hikes and Worker Voice. *The Economic Journal*, 132(645):1978–1993.
- Anderson, S. P., De Palma, A., and Thisse, J.-F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT press.
- Ashenfelter, O. (1972). Racial Discrimination and Trade Unionism. *Journal of Political Economy*, 80(3, Part 1):435–464.
- Atkin, D., Chaudhry, A., Chaudhry, S., Khandelwal, A. K., and Verhoogen, E. (2015). Markup and Cost Dispersion across Firms: Direct Evidence from Producer Surveys in Pakistan. *American Economic Review*, 105(5):537–44.
- Berger, D. W., Herkenhoff, K. F., and Mongey, S. (2022). Labor Market Power. *American Economic Review*.
- Bertrand, M., Black, S. E., Jensen, S., and Lleras-Muney, A. (2018). Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labour Market Outcomes in Norway. *The Review of Economic Studies*, 86(1):191–239.
- Biasi, B. and Sarsons, H. (2022). Flexible Wages, Bargaining, and the Gender Gap. *The Quarterly Journal of Economics*, 137(1):215–266.
- Blandhol, C., Mogstad, M., Nilsson, P., and Vestad, O. L. (2020). Do Employees Benefit from Worker Representation on Corporate Boards? *NBER Working Paper No. w28269*.
- Blau, F. D. and Kahn, L. M. (1996). Wage Structure and Gender Earnings Differentials: an International Comparison. *Economica*, 63:29–62.
- Blau, F. D. and Kahn, L. M. (2006). The U.S. Gender Pay Gap in the 1990S: Slowing Convergence. *ILR Review*, 60(1):45–66.
- Blau, F. D. and Kahn, L. M. (2017). The Gender Wage Gap: Extent, Trends, and Explanations. *Journal of Economic Literature*, 55(3):789–865.
- Bloom, N. and Van Reenen, J. (2007). Measuring and Explaining Management Practices Across Firms and Countries. *The Quarterly Journal of Economics*, 122(4):1351–1408.
- Bolotnyy, V. and Emanuel, N. (2022). Why Do Women Earn Less than Men? Evidence from Bus and Train Operators. *Journal of Labor Economics*.
- Card, D. (1996). The Effect of Unions on the Structure of Wages: A Longitudinal Analysis. *Econometrica*, pages 957–979.

- Card, D., Cardoso, A. R., Heining, J., and Kline, P. (2018). Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics*, 36(S1):S13–S70.
- Card, D., Lemieux, T., and Riddell, W. C. (2004). Unions and Wage Inequality. *Journal of Labor Research*, 25(4):519–559.
- Card, D., Lemieux, T., and Riddell, W. C. (2020). Unions and Wage Inequality: The Roles of Gender, Skill and Public Sector Employment. *Canadian Journal of Economics*, 53(1):140–173.
- Chattopadhyay, R. and Duflo, E. (2004). Women as Policy Makers: Evidence from a Randomized Policy Experiment in India. *Econometrica*, 72(5):1409–1443.
- DiNardo, J., Fortin, N., and Lemieux, T. (1996). Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach. *Econometrica*, 64(5):1001–44.
- Dittrich, M., Knabe, A., and Leipold, K. (2014). Gender differences in experimental wage negotiations. *Economic Inquiry*, 52(2):862–873.
- Dube, A., Naidu, S., and Reich, A. (2021). Power and Dignity in the Low-wage Labor Market: Theory and Evidence from Wal-Mart Workers.
- Farber, H. S. (1986). The Analysis of Union Behavior. *Handbook of Labor Economics*, 2:1039–1089.
- Farber, H. S., Herbst, D., Kuziemko, I., and Naidu, S. (2021). Unions and Inequality over the Twentieth Century: New Evidence from Survey Data*. *The Quarterly Journal of Economics*, 136(3):1325–1385.
- Feenstra, R. C. (1994). New Product Varieties and the Measurement of International Prices. *American Economic Review*, pages 157–177.
- Felix, M. (2022). Trade, labor market concentration, and wages. Technical report.
- Flabbi, L., Macis, M., Moro, A., and Schivardi, F. (2019). Do Female Executives Make a Difference? The Impact of Female Leadership on Gender Gaps and Firm Performance. *Economic Journal*, 129(622):2390–2423.
- Gerard, F., Lagos, L., Severnini, E., and Card, D. (2021). Assortative Matching or Exclusionary Hiring? The Impact of Employment and Pay Policies on Racial Wage Differences in Brazil. *American Economic Review*, 111(10):3418–57.
- Godinho Delgado, D. (2017). Trade Unions in Transformation. Building Trade Union Power with Gender Equality: The Case of the Unified Workers’ Central of Brazil. Technical report, Friedrich-Ebert-Stiftung.
- Goldin, C. (2014). A Grand Gender Convergence: Its Last Chapter. *American Economic Review*, 104(4):1091–1119.
- Goldin, C. and Katz, L. F. (2011). The Cost of Workplace Flexibility for High-Powered Professionals. *The Annals of the American Academy of Political and Social Science*, 638(1):45–67.

- Gruber, J. (1994). The Incidence of Mandated Maternity Benefits. *American Economic Review*, 84(3):622–641.
- Harju, J., Jäger, S., and Schoefer, B. (2021). Voice at work. Working Paper 28522.
- Heidhues, P. and Köszegi, B. (2018). Behavioral Industrial Organization. *Handbook of Behavioral Economics: Applications and Foundations 1*, 1:517–612.
- Hill, H. (1996). The Problem of Race in American Labor History. *Reviews in American History*, 24(2):189–208.
- Holzer, H. J., Katz, L. F., and Krueger, A. B. (1991). Job Queues and Wages. *The Quarterly Journal of Economics*, 106(3):739–768.
- Horn, C. H. (2009). *Collective Bargaining in Brazil: A Study on Joint Regulation of the Employment Relationship in Manufacturing*. VDM Verlag.
- ILO and Gallup Inc. (2017). Towards a Better Future for Women and Work: Voices of Women and Men. Technical report, ILO, Gallup.
- Jäger, S., Schoefer, B., and Heining, J. (2021). Labor in the Boardroom. *The Quarterly Journal of Economics*, 136(2):669–725.
- Kleven, H., Landais, C., and Sogaard, J. E. (2019). Children and Gender Inequality: Evidence from Denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.
- Krueger, A. and Summers, L. (1988). Efficiency Wages and the Inter-industry Wage Structure. *Econometrica*, 56(2):259–93.
- Lagos, L. (2021). Labor Market Institutions and the Composition of Firm Compensation: Evidence from Brazilian Collective Bargaining. *Working Paper*.
- Lamadon, T., Mogstad, M., and Setzler, B. (2022). Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market. *American Economic Review*, 112(1):169–212.
- Lang, K. and Gagnon, M.-J. (2009). Brazilian Trade Unions: In (In)Voluntary Confinement of the Corporatist Past. *Relations Industrielles / Industrial Relations*, 64(2):250–269.
- Leibbrandt, A. and List, J. A. (2015). Do Women Avoid Salary Negotiations? Evidence from a Large-Scale Natural Field Experiment. *Management Science*, 61(9):2016–2024.
- Liukkunen, U. (2019). *Collective Bargaining in Labour Law Regimes: A Global Perspective*, chapter The Role of Collective Bargaining in Labour Law Regimes: A Global Approach, pages 1–64. Springer Nature.
- Maestas, N., Mullen, K. J., Powell, D., von Wachter, T., and Wenger, J. B. (2018). The value of working conditions in the united states and implications for the structure of wages. Working Paper 25204.
- Maida, A. and Weber, A. (2020). Female Leadership and Gender Gap within Firms: Evidence from an Italian Board Reform. *ILR Review*, 0(0):0019793920961995.

- Manning, A. (2011). Imperfect Competition in the Labor Market. In *Handbook of Labor Economics*, volume 4, pages 973–1041. Elsevier.
- Mas, A. and Pallais, A. (2017). Valuing Alternative Work Arrangements. *American Economic Review*, 107(12):3722–59.
- Massey, C. and Thaler, R. H. (2013). The Loser’s Curse: Decision Making and Market Efficiency in the National Football League Draft. *Management Science*, 59(7):1479–1495.
- Matsa, D. A. and Miller, A. R. (2011). Chipping Away at the Glass Ceiling: Gender Spillovers in Corporate Leadership. *American Economic Review*, 101(3):635–39.
- Morchio, I. and Moser, C. (2020). The Gender Gap: Micro Sources and Macro Consequences. In *CEPR Discussion Paper No. 16383*.
- Pande, R. and Ford, D. (2012). Gender Quotas and Female Leadership: A Review. *Background Paper for the World Development Report on Gender. April, 2011*.
- Redding, S. J. and Weinstein, D. E. (2016). A Unified Approach to Estimating Demand and Welfare. *CEPR Discussion Paper No. DP11421*.
- Rosen, S. (1986). The Theory of Equalizing Differences. In *Handbook of Labor Economics*, volume 1, pages 641–692. Elsevier.
- Skorge, Ø. S. and Rasmussen, M. B. (2021). Volte-face on the welfare state: Social partners, knowledge economies, and the expansion of work-family policies. *Politics & Society*, page 00323292211014371.
- Sorkin, I. (2018). Ranking Firms Using Revealed Preference. *The Quarterly Journal of Economics*, 133(3):1331–1393.
- Staiger, D. O., Spetz, J., and Phibbs, C. S. (2010). Is There Monopsony in the Labor Market? Evidence from a Natural Experiment. *Journal of Labor Economics*, 28(2):211–236.
- Summers, L. H. (1989). Some Simple Economics of Mandated Benefits. *American Economic Review*, 79(2):177–183.
- Taber, C. and Vejlin, R. (2020). Estimation of a Roy/Search/Compensating Differential Model of the Labor Market. *Econometrica*, 88(3):1031–1069.
- Tilipman, N. (2022). Employer Incentives and Distortions in Health Insurance Design: Implications for Welfare and Costs. *American Economic Review*, 112(3):998–1037.
- Visser, J. (2019). Institutional Characteristics of Trade Unions, Wage Setting, State Intervention and Social Pacts (ICTWSS) Database, version 6.0. *Amsterdam Institute for Advanced Labour Studies (AIAS), University of Amsterdam*.
- Wiswall, M. and Zafar, B. (2017). Preference for the Workplace, Investment in Human Capital, and Gender. *The Quarterly Journal of Economics*, 133(1):457–507.

Online Appendix

A. Appendix Tables

B. Appendix Figures

C. Data Appendix

D. AKM and PageRank Model

E. Welfare Model

A Appendix Tables

Table A1: Clauses in the Intuitive Definition of Female-Centric Amenities

Group	Clause Type	Description
<i>Leaves</i>		
	Abortion leave	Leave in cases of miscarriage/abortion
	Adoption leave	Leave following the adoption of a child
	Maternity leave	Leave concerning the birth of a child
	Paid leave	Leave during which worker receives normal pay
	Unpaid leave	Leave during which worker does not receive normal pay
	Other: holidays and leaves	Provisions on holidays/leaves outside predefined clause types
	Female workforce	General provisions concerning female workers
<i>Maternity and childcare</i>		
	Childcare assistance	Payments to assist with childcare support
	Maternity assistance	Payments to assist with becoming a mother
	Abortion protections	Employment protections concerning miscarriage/abortion
	Maternity protections	Employment protections for mothers
	Paternity protections	Employment protections for fathers
	Policies for dependents	Workplace benefits that apply to dependents
<i>Workplace harassment and discrimination</i>		
	Sexual harassment	Rules/penalties pertaining to harassment in the workplace
	Equal opportunities	Initiatives/statements on equality of opportunity for workers
<i>Flexibility and part-time work</i>		
	Workday controls	Rules restricting the duration of the workday
	Special shifts	Work shifts for subgroups of workers, e.g., women, minors, students
	On-call	Rules on workers' availability outside of the normal workday
	Uninterrupted shifts	Rules concerning back-to-back shifts
	Part-time contracts	Directives on temporary/part-time employment contracts

Notes: Table lists the *Sistema Mediador* clause types used in our intuitive definition of female-centric amenities. The descriptions provided in this table are purposefully vague—clauses of a given type can vary to some degree. The clauses were chosen based on the content of CUT’s fight plan and the existing literature on workplace amenities valued by women, restricting ourselves to only 20 clause types.

Table A2: Robustness of Data-Driven Female-Centric Amenities

Clause type	Times selected: data-driven	Selected in baseline data-driven approach:	
	(out of 6 methods)	(no state and industry FEs)	(state and industry FEs)
Childcare assistance	6	1	1
Absences	6	1	1
Adoption leave	6	1	0
Other: holidays and leaves	6	1	1
Seniority pay	6	1	1
Maternity protections	6	1	1
Paid leave	6	1	1
Night pay	6	1	0
Abortion leave	6	1	0
Policy for dependents	6	1	0
Waiving union fees	6	1	1
Salary adjustments/corrections	6	1	0
Renewal/termination of the CBA	5	1	0
Nonwork-related injury protections	5	1	0
Extension/reduction of workday	5	1	1
Medical exams	5	1	0
Unionization campaigns	4	1	0
Abortion protections	4	1	0
Adoption protections	4	0	0
Guarantees to union officers	3	1	1
Health education campaigns	3	1	0
Military service protections	3	0	1
Separation/dismissal	2	0	1
Other employment protections	2	0	0
Awards	1	0	0
Moral harassment	1	0	1
Maternity leave	1	0	0

Notes: Table lists all of the clauses identified as female-centric in any of the 6 methods implemented based on the estimation of Equation (2). Methods vary in 1) the sample of establishments covered by sectoral CBAs used, i.e., a random sample or the full sample; and 2) the measure of PageRank values used to determine gender gaps, i.e., normalized, non-normalized, or rankings. The initial column simply shows the number of times the clause is picked as female-centric by one of these 6 methods (clauses in the table are sorted in descending order as per the values of this column). The next column is an indicator for whether the clauses is selected as a female-centric by the baseline method, i.e., using a random sample and normalized PageRanks. The final column is an indicator for whether the clause is selected as female-centric by the baseline method but where the lasso includes state and industry fixed effects. Note that the Spearman correlation of the coefficients on clauses using the data-driven lasso approach versus an OLS using these same clauses but adding state and industry fixed effects is 0.56 with p-value below 0.01.

Table A3: Robustness of Data-Driven Male-Centric Amenities

Clause type	Times selected: data-driven	Selected in baseline data-driven approach:	
	(out of 6 methods)	(no state and industry FEs)	(state and industry FEs)
On-call pay	6	1	1
Life insurance	6	1	1
Strike procedures	6	1	1
Other: protections for injured workers	6	1	1
Female workforce	6	1	1
Machine and equipment maintenance	6	1	1
Duration and schedule	6	1	1
Working environment conditions	6	1	0
Salary payment - means and timeframes	6	1	0
Hazard pay (danger risk)	6	1	0
Workday compensation	6	1	0
Tools and equipment	6	1	0
Profit sharing	5	1	1
Transfers	5	1	0
Safety equipment	5	1	0
Other assistances	5	1	0
Death/funeral assistance	5	1	0
Salary deductions	4	1	0
Equal opportunities	4	0	0
Collective vacations	3	1	0
Union fees	3	0	0
CIPA: accident prevention committee	2	1	1
Unpaid leave	2	0	0
Part-time contracts	2	0	0
Food assistance	1	0	0
Performance evaluation	1	0	0
Employment/hiring rules	1	0	0

Notes: Table lists all of the clauses identified as male-centric in any of the 6 methods implemented based on the estimation of Equation (2). Methods vary in 1) the sample of establishments covered by sectoral CBAs used, i.e., a random sample or the full sample; and 2) the measure of PageRank values used to determine gender gaps, i.e., normalized, non-normalized, or rankings. The initial column simply shows the number of times the clause is picked as male-centric by one of these 6 methods (clauses in the table are sorted in descending order as per the values of this column). The next column is an indicator for whether the clauses is selected as a male-centric by the baseline method, i.e., using a random sample and normalized PageRanks. The final column is an indicator for whether the clause is selected as male-centric by the baseline method but where the lasso includes state and industry fixed effects. Note that the Spearman correlation of the coefficients on clauses using the data-driven lasso approach versus an OLS using these same clauses but adding state and industry fixed effects is 0.56 with p-value below 0.01.

Table A4: Effect of CUT Reform on Negotiated Amenities (Cluster at Union-Level)

	Intuitive definition (female clauses)					Data-driven		
	All (1)	Leave (2)	Maternity (3)	Harassment (4)	Flexibility (5)	Female (6)	Male (7)	F/(F+M+1) (8)
<i>Panel A: Intensive margin (number)</i>								
$D_i \times \delta_{year \geq 2015}$	0.157* (0.083)	0.078** (0.040)	0.042* (0.023)	0.009** (0.004)	0.028 (0.031)	0.301** (0.144)	0.130 (0.159)	0.032* (0.018)
Mean outcome	0.95	0.25	0.24	0.02	0.44	1.58	2.55	0.15
<i>Panel B: Intensive margin (sum of indicators)</i>								
$D_i \times \delta_{year \geq 2015}$	0.123* (0.067)	0.047 (0.031)	0.042* (0.022)	0.008** (0.004)	0.027 (0.021)	0.154* (0.080)	0.067 (0.095)	
Mean outcome	0.70	0.18	0.21	0.02	0.30	1.26	1.58	
<i>Panel C: Extensive margin</i>								
$D_i \times \delta_{year \geq 2015}$	0.017 (0.015)	0.012 (0.011)	0.020* (0.012)	0.008** (0.004)	0.022 (0.015)	0.034* (0.020)	-0.001 (0.015)	
Mean outcome	0.31	0.12	0.15	0.02	0.23	0.36	0.46	
<i>Panel D: As a share of all clauses</i>								
$D_i \times \delta_{year \geq 2015}$	0.005 (0.004)	0.001 (0.001)	0.001 (0.001)	0.000 (0.003)	0.003 (0.015)	0.021 (0.006)	-0.003 (0.012)	
Mean outcome	0.05	0.01	0.01	0.00	0.03	0.07	0.14	
Observations	600,960	600,960	600,960	600,960	600,960	600,960	600,960	600,960

Notes: Table reports the coefficients for DID regressions—see Equation (3)—estimating the effect of the CUT reform on the female-centric and male-centric amenities included in CBAs. Panel A uses the total number of clauses per pair-year as an intensive margin measure. Panel B uses the sum of the corresponding clause type indicators, capturing how the space of female (male) clauses grows or shrinks. Panel C uses an indicator for pair-year observations with at least one corresponding clause as an extensive margin measure. Panel D uses the share of corresponding clauses with respect to the total contract clauses, capturing how the composition of CBAs change. Under each panel we report the mean of the dependent variable among the treated at baseline (2014). The sample is the filled panel of establishment-union pairs by year. All columns control for pair fixed effects, as well as time-varying state and industry fixed effects. Standard errors are clustered at the union level, instead of at the establishment level, which reduces the number of clusters from around 80 thousand to about 4.4 thousand.

Table A5: Effect of CUT Reform on Negotiated Amenities (CBA coverage in 2014)

	Intuitive definition (female clauses)					Data-driven		
	All (1)	Leave (2)	Maternity (3)	Harassment (4)	Flexibility (5)	Female (6)	Male (7)	F/(F+M+1) (8)
<i>Panel A: Intensive margin (number)</i>								
$D_i \times \delta_{year \geq 2015}$	0.96*** (0.015)	0.044*** (0.006)	0.020*** (0.004)	0.005*** (0.001)	0.028*** (0.010)	0.121*** (0.023)	0.111*** (0.031)	0.009*** (0.090)
Mean outcome	1.63	0.43	0.41	0.03	0.76	2.71	4.38	0.25
<i>Panel B: Intensive margin (sum of indicators)</i>								
$D_i \times \delta_{year \geq 2015}$	0.070*** (0.010)	0.023*** (0.004)	0.021*** (0.004)	0.003*** (0.001)	0.022*** (0.005)	0.076*** (0.014)	0.050*** (0.016)	
Mean outcome	1.21	0.31	0.36	0.03	0.51	2.17	2.71	
<i>Panel C: Extensive margin</i>								
$D_i \times \delta_{year \geq 2015}$	0.019*** (0.003)	0.012*** (0.002)	0.010*** (0.002)	0.004*** (0.001)	0.021*** (0.003)	0.005* (0.003)	0.009** (0.003)	
Mean outcome	0.53	0.21	0.25	0.03	0.40	0.62	0.79	
<i>Panel D: As a share of all clauses</i>								
$D_i \times \delta_{year \geq 2015}$	0.005*** (0.001)	0.001*** (0.000)	0.001*** (0.000)	0.000*** (0.000)	0.003** (0.001)	0.004*** (0.001)	0.001 (0.002)	
Mean outcome	0.08	0.01	0.01	0.00	0.06	0.11	0.25	
Observations	366,468	366,468	366,468	366,468	366,468	366,468	366,468	366,468

Notes: Table reports the coefficients for DID regressions—see Equation (3)—estimating the effect of the CUT reform on the female-centric and male-centric amenities included in CBAs. The sample is the filled panel of establishment-union pairs by year, restricted to establishment-union pairs with CBA coverage in 2014. Panel A uses the total number of clauses per pair-year as an intensive margin measure. Panel B uses the sum of the corresponding clause type indicators, capturing how the space of female (male) clauses grows or shrinks. Panel C uses an indicator for pair-year observations with at least one corresponding clause as an extensive margin measure. Panel D uses the share of corresponding clauses with respect to the total contract clauses, capturing how the composition of CBAs change. Under each panel we report the mean of the dependent variable among the treated at baseline (2014). All columns control for pair fixed effects, as well as time-varying state and industry fixed effects. Standard errors are clustered at the establishment level.

Table A6: Effect of CUT Reform on Female Amenities

	Female-Centric Clauses: Intensive Margin				
	(1)	(2)	(3)	(4)	(5)
<i>Panel A: Intuitive definition</i>					
$D_i \times \delta_{year \geq 2015}$	0.157*** (0.013)	0.157*** (0.013)	0.157*** (0.013)	0.194*** (0.014)	0.096*** (0.015)
Mean outcome	0.95	0.95	0.95	0.95	1.63
<i>Panel B: Data-driven definition</i>					
$D_i \times \delta_{year \geq 2015}$	0.301*** (0.021)	0.347*** (0.026)	0.262*** (0.017)	0.332*** (0.022)	0.121*** (0.023)
Mean outcome	1.58	2.05	1.17	1.58	2.71
Data-driven clauses	baseline	any	all	baseline	baseline
Geography-year FEs	state	state	state	microregion	state
CBA coverage in 2014	no	no	no	no	yes
Observations	600,960	600,960	600,960	600,960	366,468

Notes: Table reports the coefficients for DID regressions—see Equation (3)—estimating the effect of the CUT reform on female amenities included in CBAs. The dependent variable is the total number of clauses per pair-year as an intensive margin measure, with Panel A using the intuitive definition of female-centric clauses and Panel B using the data-driven approach. Columns (1)-(3) modify the dependent variable by changing the clauses that are chosen as female-centric in the data-driven approach: a) *baseline*: top 20 clauses using a random sample and normalized PageRank values for the gender gaps; b) *any*: counts any of the clauses selected across 6 approaches as female-centric; c) *all*: counts only those clauses that are selected in all 6 approaches as female-centric. Refer to Table A2 for a list of the clauses used in each of these scenarios. Column 4 adds more granular time-varying fixed effects at the geographic level, i.e., using micro-region instead of state. Column 5 requires that pairs are covered by a CBA at baseline to test whether effects are driven by changes in the amenities among units with active CBAs rather than by gains in coverage. Standard errors are clustered at the establishment level.

Table A7: Effect of CUT Reform on Female Amenities

	Female-Centric Clauses: As a Share of All Clauses				
	(1)	(2)	(3)	(4)	(5)
<i>Panel A: Intuitive definition</i>					
$D_i \times \delta_{year \geq 2015}$	0.005*** (0.001)	0.005*** (0.001)	0.005*** (0.001)	0.005*** (0.001)	0.005*** (0.001)
Mean outcome	0.05	0.05	0.05	0.05	0.08
<i>Panel B: Data-driven definition</i>					
$D_i \times \delta_{year \geq 2015}$	0.021*** (0.001)	0.021*** (0.001)	0.022*** (0.001)	0.017*** (0.001)	0.004*** (0.001)
Mean outcome	0.07	0.08	0.04	0.07	0.11
Data-driven clauses	baseline	any	all	baseline	baseline
Geography-year FEs	state	state	state	microregion	state
CBA coverage in 2014	no	no	no	no	yes
Observations	600,960	600,960	600,960	600,960	366,468

Notes: Table reports the coefficients for DID regressions—see Equation (3)—estimating the effect of the CUT reform on female amenities included in CBAs. The dependent variable is the share of female-centric clauses among all clauses per pair-year, with Panel A using the intuitive definition of female-centric clauses and Panel B using the data-driven approach. Columns (1)-(3) modify the dependent variable by changing the clauses that are chosen as female-centric in the data-driven approach: a) *baseline*: top 20 clauses using a random sample and normalized PageRank values for the gender gaps; b) *any*: counts any of the clauses selected across 6 approaches as female-centric; c) *all*: counts only those clauses that are selected in all 6 approaches as female-centric. Refer to Table A2 for a list of the clauses used in each of these scenarios. Column 4 adds more granular time-varying fixed effects at the geographic level, i.e., using micro-region instead of state. Column 5 requires that pairs are covered by a CBA at baseline to test whether effects are driven by changes in the amenities among units with active CBAs rather than by gains in coverage. Standard errors are clustered at the establishment level.

Table A8: Differential Effects by Gender for Incumbent Workers

	Stay at baseline employer (1)	Employed in formal sector (2)	Log wages (3)
$D_i \times \delta_{year \geq 2015}$	0.010*** (0.002)	0.002 (0.002)	-0.000 (0.001)
$D_i \times \delta_{year \geq 2015} \times Female_i$	0.008*** (0.003)	0.005** (0.002)	0.002 (0.002)
Observations	55,658,796	55,658,796	46,668,757
R^2	0.63	0.44	0.90

Notes: Table reports the coefficients for the gender-pooled DID regression estimating the effect of the CUT reform on retention, formal sector employment, and wages of incumbent workers. Treatment status of incumbent workers is based on the CUT-affiliation of the union negotiating with their baseline (2014) employer. These workers are tracked wherever they go. The regression interacts treatment status with dummy variables for the post period (after 2014) and gender. Regressions include worker fixed effects, industry-year-gender fixed effects, microregion-year-gender fixed effects, and tenure-year-gender fixed effects. To make treatment effects in worker-level regressions interpretable as establishment-level averages, we weight each incumbent worker by the inverse of employment at their baseline employer. Standard errors are clustered by establishment and reported in parentheses.

Table A9: Impact of CUT Reform on Female Workforce

	Share poached in (1)	Mean years of age (2)	Mean months of tenure (3)	Mean hours in contract (4)	Mean years of schooling (5)
$D_i \times \delta_{year \geq 2015}$	-0.001 (0.002)	-0.012 (0.041)	0.172 (0.215)	-0.033 (0.025)	-0.001 (0.010)
Mean outcome	0.209	33.5	43.1	42.0	11.3
Observations	342,207	342,207	342,207	342,207	342,207

Notes: Table reports the coefficients for the establishment-level DID regression from Equation (3), comparing treated to comparison establishments on characteristics of their female workforce. An establishment is treated if the union with which it negotiates is affiliated to CUT in 2012. Each regression includes establishment fixed effects, industry-year fixed effects, and microregion-year fixed effects. Standard errors are clustered by establishment and reported in parentheses.

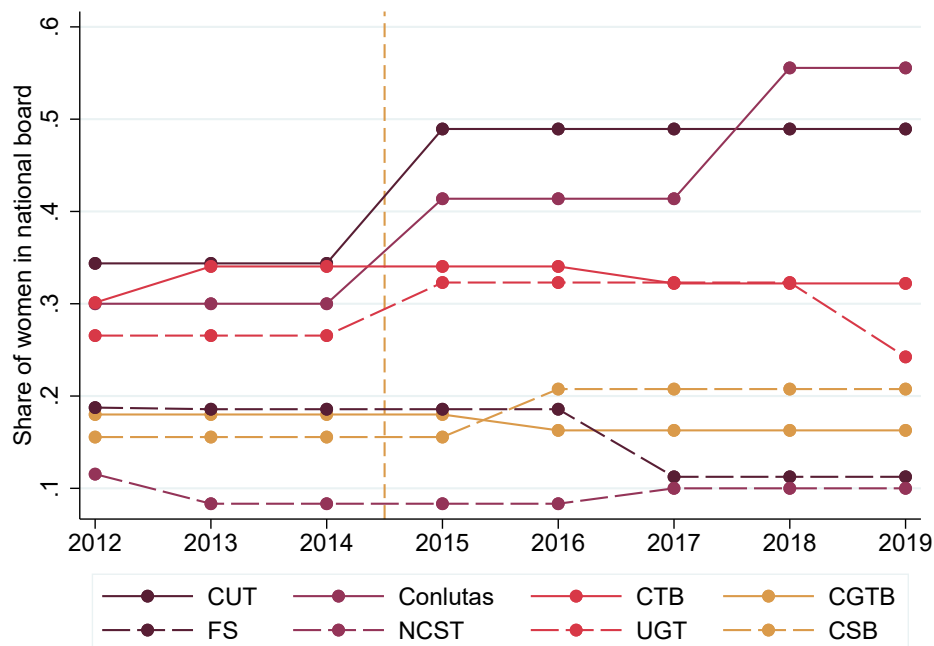
Table A10: Welfare Estimation

	Women 20-35	All women	Men 20-35	All men	
	(1)	(2)	(3)	(4)	
$\ln\phi_{t-1,t}$	0.044 (0.0062)	0.059 (0.0066)	-0.005 (0.0048)	0.013 (0.0045)	
<i>Components breakdown:</i>					
	$\ln(\lambda_{t,t-1}) - \ln(\lambda_{t-1,t})$	-0.012	-0.018	-0.005	-0.006
	$\ln(\bar{w}_t^*) - \ln(\bar{w}_{t-1}^*)$	0.015	0.022	-0.001	0.011
	$\ln(\bar{S}_t^*) - \ln(\bar{S}_{t-1}^*)$	-0.046	-0.058	0.013	0.001
η (calibrated)	1.015				
N establishments	58,417	60,651	59,438	60,651	
N establishments in $\Omega_{t,t-1}$	45,331	47,195	46,182	47,195	

Notes: Table reports the estimated welfare change for different groups of workers: women between 20 and 35 years old, all women, men between 20 and 35 years old, all men. It also reports estimates of the three components that make the welfare index, namely the Feenstra “new varieties” term $\ln(\lambda_{t,t-1}) - \ln(\lambda_{t-1,t})$, the change in the geometric average of the wages of non-CUT firms $\ln(\bar{w}_t^*) - \ln(\bar{w}_{t-1}^*)$, and the change in the geometric average of the labor income shares of non-CUT firms. Standard errors in parenthesis come from the bootstrap procedure described in Appendix E.

B Appendix Figures

Figure B1: Gender Parity in National Leadership by Union Central



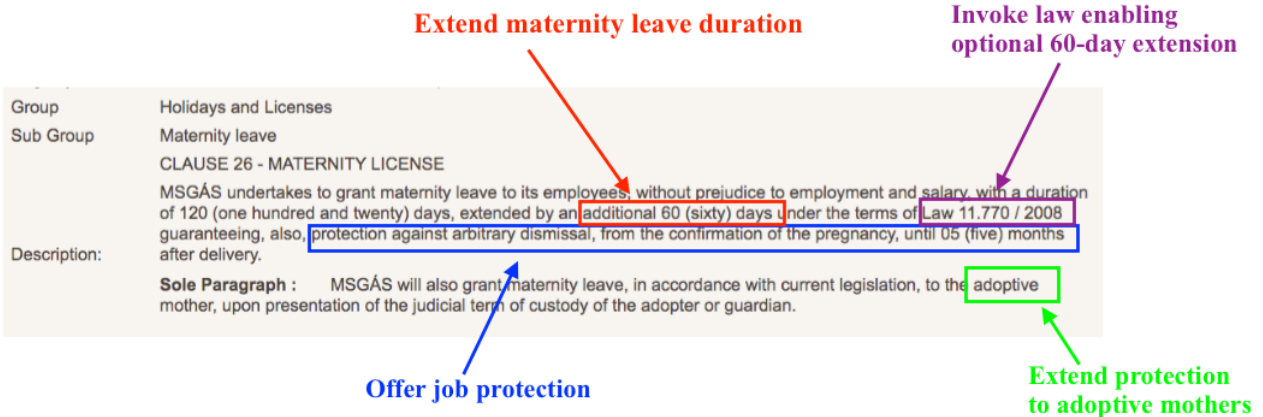
Notes: Figure plots the annual share of women on each union central’s national executive committee (*Inter-sindical* is dropped due to missing information on its board). The line for CUT is the same as in Figure 2a, while the unweighted average of all other union centrals make up the other line reported in Figure 2a. Solid lines refer to “combative” union centrals, while dashed lines represent “cooperative” union centrals. The second largest union central and main competitor to CUT is *Força Sindical* (FS).

Figure B2: Cover of CUT's Female-Centric "Fight Plan"



Notes: The 2015 CUT reform consisted of two parts. The first is a 50% quota for women in CUT's state and national executive bodies. The second is the adoption of a bargaining agenda more attentive to the needs of female workers. Figure B2 is the cover page of the book of resolutions (or "fight plan") developed at the 2015 meeting of CUT Women to detail concrete strategies for achieving parity in practice at all levels of unions within CUT. It recommends steps for giving women more actual voice in all levels of the union—like representation on committees and a say in union's list of demands (or *pautas*). It also specifies amenities like maternity leave extensions and subsidized childcare to highlight during collective bargaining. This book of resolutions was subsequently adopted by delegates at the 2015 CUT National Congress (full text [here](#)). The word count for *mulheres* (women) in the National Congress book of resolutions increased from 46 in 2012 to 203 in 2015.


Figure B3: Example of a Maternity Leave Clause



Notes: Figure shows an example of a maternity leave clause in a CBA. The clause is classified under the “Holidays and Licenses” broad group (10 in total) and the “Maternity Leave” subgroup (137 in total). This particular clause extends maternity leave duration from the state-mandated 120 days to 180 days—inclusive to adopting mothers. It also extends post-maternity job protection by 6 months. The paper relies on the subgroup classification of the different clauses, ignoring the variation in the text that may exist within each individual clause belonging to a specific subgroup.


Figure B4: Additional Sense Checks for Female- and Male-Centric Amenities

- (a) Value gaps and intuitive female clauses
- (b) Value gaps and data-driven female clauses



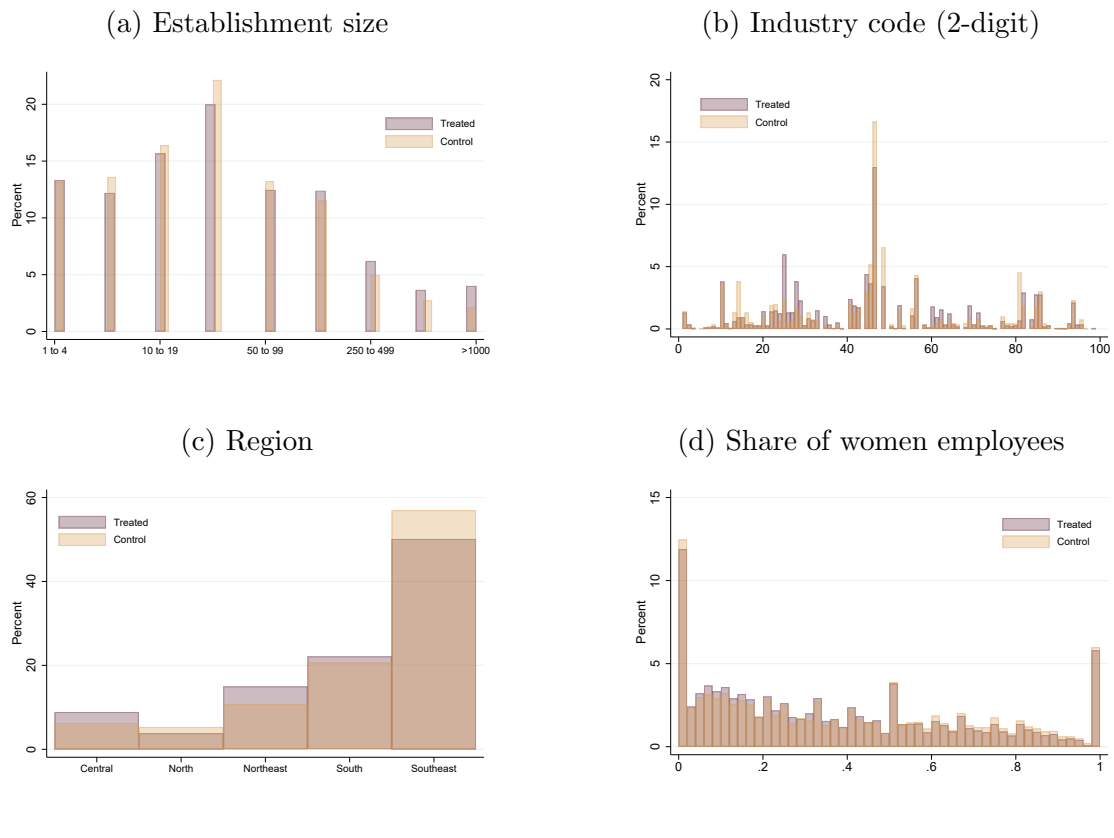
Notes: Figures depict binned scatterplots of the establishment-level gender gaps in PageRank values by the average female-centric clauses from sectoral CBAs applying to the establishment. Figure B4a uses the intuitive definition of female-centric amenities, while Figure B4b uses the data-driven approach. The sample used is the one used to estimate Equation (2), i.e., establishments in the intersection of the gender-specific super-connected sets covered by sectoral CBAs in at least 4 different years between 2009-2016.

Figure B5: Union Affiliation to CUT Over Time



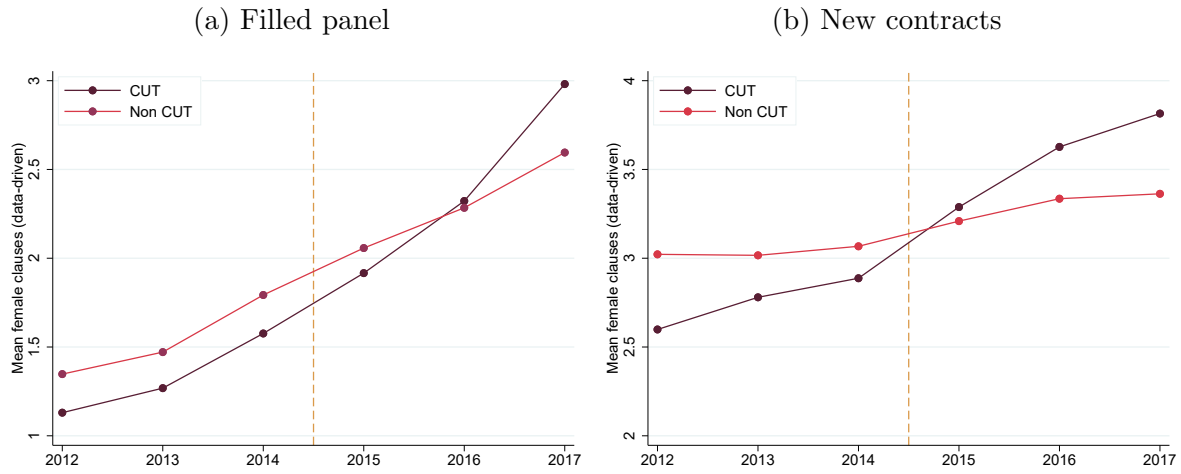
Notes: Figure plots changes in the probability of being affiliated to CUT between 2012 and 2016 separately for unions having either a high or a low share of women among the workers they represent (above or below the mean, i.e., 33% women). Coefficients represent the change with respect to 2012, in which the probability of being a CUT-affiliate is normalized to zero. Unions are weighted by the size of the workforce that they represent, computed by summing the 2012 worker count across establishments negotiating firm-level CBAs with the union. That is, if an establishment negotiates with n unions, we split the workforce count evenly to those n unions (results are robust to removing these weights). The sample is restricted to the unions in the filled panel, where only 3% of unions ever switch affiliation to or from CUT. Standard errors are clustered at the union level.

Figure B6: Baseline Characteristics of Treated and Control Establishments



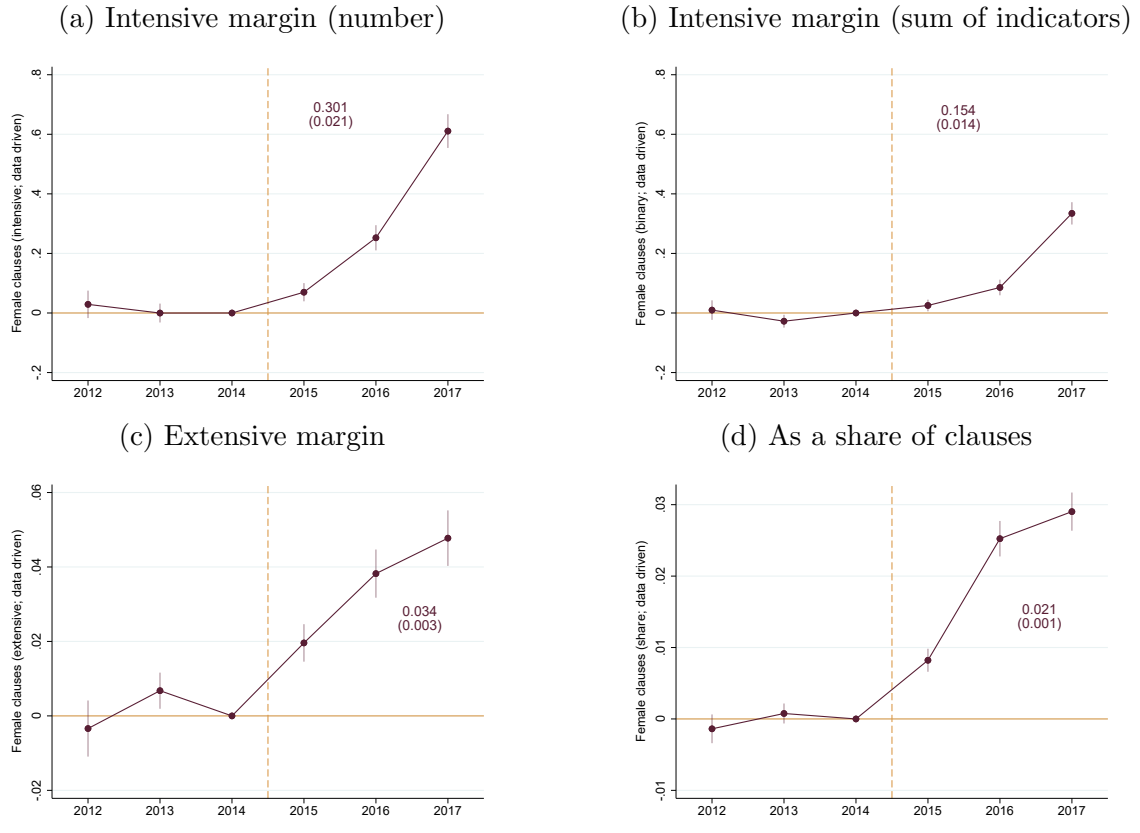
Notes: Figures shows that treated and control establishments distributions of size, 2-digit-industry, regional location, and female share of employment at baseline (using all unique worker-year observations during 2014). For some context regarding industry figure, the most represented industry in both treated and control groups are establishments in the retail industry. The establishments come from the starting sample detailed in Table 2.

Figure B7: Trends in Female-Centric Clauses (Data-Driven Approach)



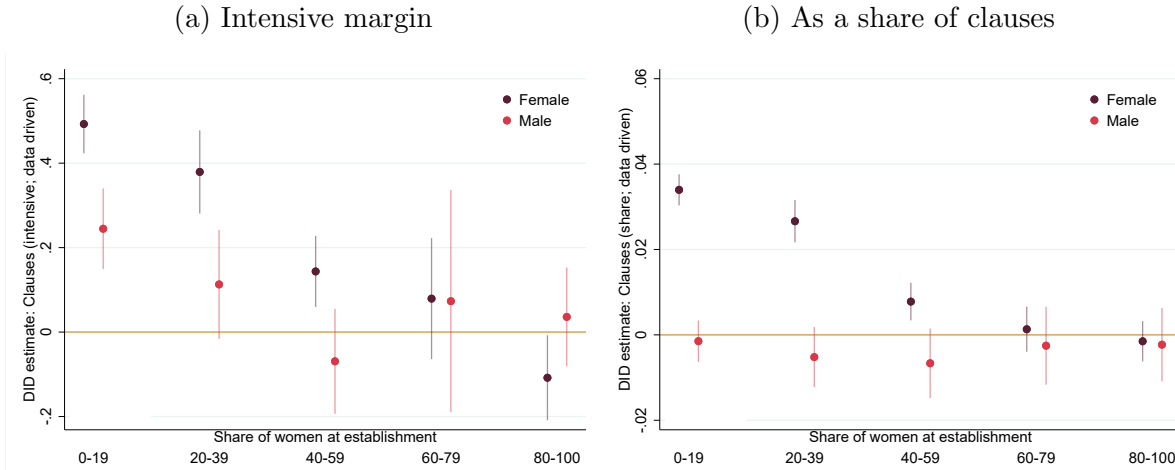
Notes: Figures plot the raw average number of female-centric clauses for treated (CUT) and control (Non CUT) establishment-union pairs over the years. Female-centric clauses are based on the data-driven classification. Figure B7a plots the average number of female-centric clauses for the filled panel, while Figure B7b plots the average number of female-centric clauses in newly signed contracts of the given year. Mean female clauses are lower in the filled panel and react slowly to changes in new contracts because of pairs that do not have CBA coverage in a given year.

Figure B8: Effect of the CUT Reform on Female-Centric Amenities



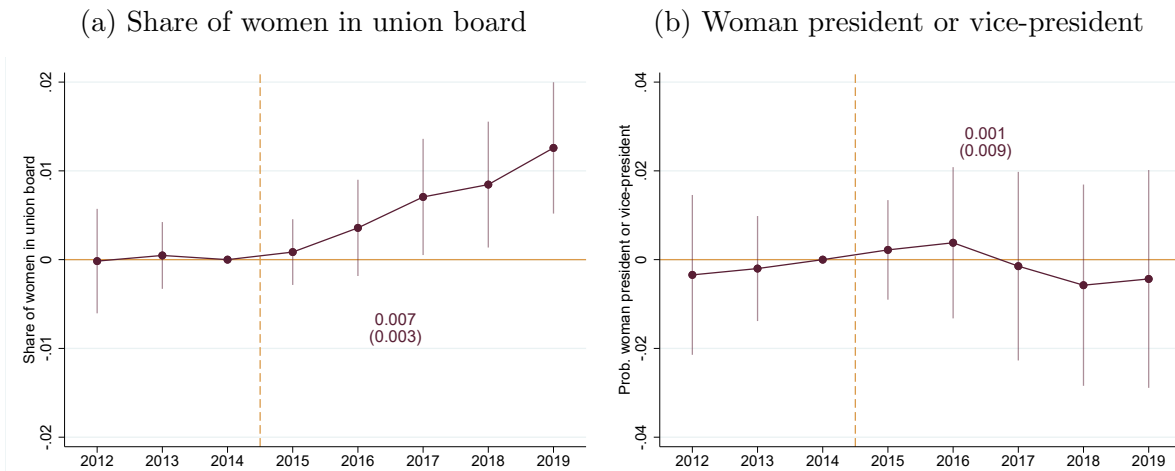
Notes: Figures show estimates of the δ_t coefficients for $t \in [2012, 2017]$ (with 2014 omitted) from the DID specification in Equation (3) on all margins considered for female-centric clauses, defined using the data-driven method. Confidence intervals at a 95% level are reported. Standard errors are clustered at the establishment level. All figures use the filled panel.

Figure B9: Effect on Amenities by Share of Female Workers at Establishment



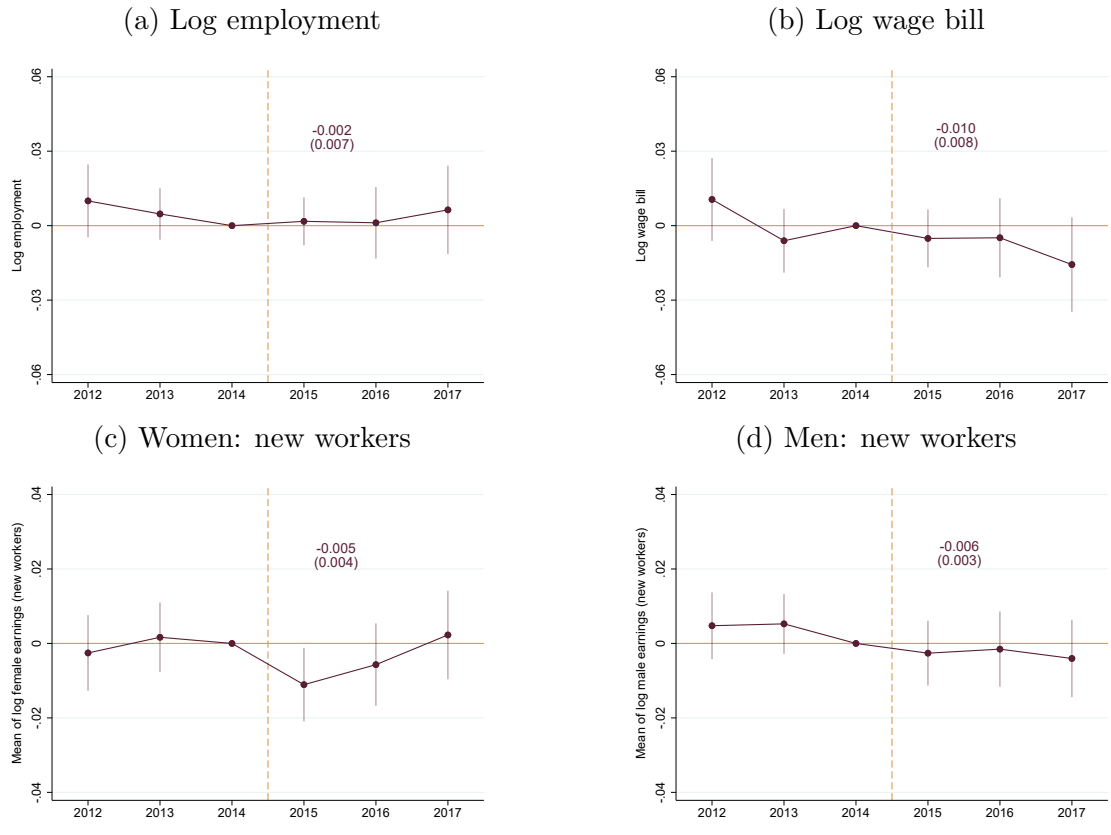
Notes: Figures show estimates of the treatment effect ($\delta_{year \geq 2015}$) from the DID specification in Equation (3) on the number of female- and male-centric clauses (data-driven approach) computed on subsamples of establishments divided according to the 2014 share of female workers. Figures use the filled panel. Confidence intervals at a 95% level are shown. Standard errors are clustered at the establishment level.

Figure B10: Impact on Gender Representation in Local Union Boards



Notes: Figures show estimates of the δ_t coefficients for $t \in [2012, 2019]$ (with 2014 omitted) from an event-study specification similar to the one in Equation (3) on measures of women representation within local union boards. The sample is restricted to unions in our analysis sample (unlike Figure 2b). The equation we estimate is slightly different from Equation (3) as the unit of observation here is the union-year so we include union fixed effects instead of establishment-union pair fixed effects. Figure B10a uses the share of women in the union board as a dependent variable, while Figure B10b uses a dummy indicating whether the union's president (or vicepresident) is a woman. Confidence intervals at a 95% level are reported. Standard errors are clustered by union.

Figure B11: Effects on Employment, Wages, and Profits



Notes: Figures report the results of the establishment-level DID regression in Equation (3) with outcome variables: log of total employment, log of the wage bill, mean log wages for new female hires, and mean log wages for new male hires. Each regression includes establishment fixed effects, industry-year fixed effects, and microregion-year fixed effects. The figure plots estimates of the δ_t coefficients for $t \in [2012, 2017]$ with 2014 omitted. Confidence intervals at a 95% level are reported. Standard errors are clustered by establishment.

C Data Appendix

C.1 Sample construction

To analyze the CUT reform’s impact on various outcomes, we construct three main analysis samples. The first is a sample to study changes in CBA clauses at the establishment-union pair level (henceforth, simply *pair* level). The second is a sample at the establishment level to study changes in the workplace. The third is a sample at the worker level used to track the labor market outcomes of incumbent workers. In addition to these three main samples, we also construct two panel datasets at the local union level and at the union central level to study the gender composition of their boards.

Amenities sample Amenities (on paper) are captured by CBA clauses signed by establishment-union pairs. We first construct a yearly panel of the new CBAs signed by a pair in a given year, i.e., new contracts. We then use this sample to construct a balanced panel containing the active clauses applying to a pair over time, i.e., filled panel.

1. New contracts: We construct this sample using the set of CBAs registered on *Sistema Mediador*. We restrict to valid, non-amendment, firm-level CBAs signed between 2012 and 2017 (inclusive). Each CBA contains information on who signs the agreement—the CNPJ identifiers of the employer(s) and union(s) signing it—and, importantly for our analysis, how many clauses it contains classified into clause types.⁴⁵

The union identifier allows us to merge these data with data on union affiliation to union centrals coming from CNES. The employer identifier allows us to merge these data with information in RAIS, e.g., industry, microregion, and employment. We drop CBAs signed by unions with missing information about their 2012 union central affiliation (around 1.5% of contacts).⁴⁶ We additionally drop contracts signed by multiple unions with different union central affiliations: this is fewer than 0.33% of CBAs.⁴⁷

Almost all pairs negotiate at most one contract per year: 96% of CBAs are the only agreement signed by a pair that year and 85% of pairs always negotiate at most one CBA per year during our study period. As for the remaining 15%, we take the maximum count

⁴⁵*Sistema Mediador* classifies clauses into 137 categories, e.g., maternity assistance, overtime pay, life insurance, procedures in relation to strikes and strikers, etc.

⁴⁶Unions that decide not to affiliate with any union central—which are registered in CNES as “Not-Affiliated”—are not dropped. The CBAs signed by these unions are part of the control group.

⁴⁷Of the remaining agreements, 89.8% are negotiated between a single establishment and a single union, 7.3% are negotiated by a single union with two or more establishments, 2.5% are signed by one establishment and two or more unions with the same CUT or non-CUT affiliation, and only 0.5% by multiple unions and multiple establishments.

of a given clause type across the CBAs negotiated by the pair in a given year.⁴⁸ In this way we obtain a sample of newly negotiated CBAs at the pair-year level, reporting the number of clauses for each clause type.

On the signing establishment’s side, we restrict to pairs that have non-missing industry and microregion information, and that employs workers at baseline (2014). These restriction drop 8.5% of observations. This comprises the starting sample with observations at the pair-year level reported in the descriptive statistics of Table 2.

2. Filled panel: This sample fills in the amenities information for pairs in the *new contracts* sample for years when a new firm-level CBA was not signed. In filling the panel, we consider the institutional context regarding the automatic extension of CBAs into the future. That is, for a given pair, contracts expiring after September 25 of 2012 are automatically extended into the future until a new CBA is signed (Lagos, 2021). Although CBAs expiring before that date were not extended, we observe contracts starting 3 years prior to our study period, i.e., starting in 2009. Since the maximum duration of a CBA is 24 months, by the start of our study period (i.e., 2012) we can already be certain whether any CBA applies to a given pair-year. As such, these institutional features allow us to generate a balanced panel at the pair-year level.

To aggregate amenities at the pair-year level, for each year we only consider the contract(s) covering at least 6 months of the year.⁴⁹ If more than one contract per pair-year remains, we take the maximum count of a given clause type across CBAs—similarly to what done for the *new contracts* sample. If a pair is not covered by a firm-level CBA in a given year (even after filling the panel), we set the clause count for each clause type to zero. As such, this procedure produces a yearly balanced panel at the establishment-union pair level.

Establishment sample To study changes in the workplace, we match the contracts in our *amenities sample* to the signing establishments in RAIS. Establishments covered by contracts negotiated by unions affiliated to CUT in 2012 form our treatment group, while establishments covered by CBAs signed by unions not affiliated to CUT in 2012 make up our comparison group.

We start with the list of establishments that are part of the pairs in our *new contracts* sample. We restrict to establishments employing both men and women at baseline, dropping 15,550 establishments. We further restrict this list to establishments in the geographic coverage of their “baseline CBA”, defined as the firm-level agreement closest to the 2015 CUT

⁴⁸We do this to avoid double-counting clauses as the multiple agreements per pair-year often result from misclassified CBA amendments or single-issue CBAs that are renegotiated more frequently than a year.

⁴⁹All other restrictions used in the *new contracts* sample apply.

reform among those signed by the establishment. The reason for this restriction is that, for multiple-establishment firms, the CNPJ listed as the employer counterpart in the CBA need not be covered.⁵⁰ Restricting to signing establishments in the geographic coverage of their baseline CBA further drops 8,684 establishments, leaving us with 61,752 establishments.

For each establishment in this list we compute outcomes at the establishment-year level, such as mean log wages or total female employment, either using all job spells registered at that establishment in the year or using workers’ “main job spell” in each year. We define the “main job spell” as the employment spell at which the worker worked the longest during the year. In case all job spells have the same duration, we break ties by keeping only one spell at random.

Because the same establishment can negotiate CBAs with more than one union, the final step to construct the *establishment sample* involves determining treatment status at the establishment level. We assign establishments to the treatment group as long as they are part of at least one treated pair. In practice, this decision is innocuous. Because the great majority (93.5%) of establishments always bargain with the same union, treatment assignment is trivially defined for most establishments. There are 4.4% of establishments that sign CBAs with more than one union over the time frame we consider, and all the unions they negotiate with have the same treatment status, e.g. they are all affiliated to CUT (or they are not) in 2012. The remaining 2.1% of establishments negotiate with more than one union over time and these unions have different treatment status. We conservatively assign this last group of establishment to the treatment group, which should run counter to finding effects if some of these establishments are not affected by the CUT reform.

Incumbent workers sample Incumbent workers are defined as those employed at a treated or comparison establishment as of 2014 (based on the *establishment sample*). Their treatment status depends on the treatment status of their baseline (2014) employer, as explained above in the description of the *establishment sample* construction. Leveraging the linked employer-employee feature of RAIS, incumbent workers are tracked across jobs from 2012 to 2017—that is, we are not restricting to job spells at employers in the *establishment sample*. In constructing this sample, we only consider the “main job spell” for each worker in each year.

Union and union central boards For each Brazilian union central, we construct a yearly panel with information on the gender composition of their national board between 2012 and

⁵⁰Firm-level CBAs apply to workers at all establishments of the signing firm that are in the geographic coverage specified in the contract. In case of multi-establishment firms, the establishment signing a CBA could be the firm headquarter but the contract might cover only subsidiaries located in other municipalities.

2019. The raw data contains the full name of all the board members, which allows us to infer their gender. We do so using the R package *genderBR*, which codes a name as female if most people with that name are women in the Brazilian census (and similarly for men).⁵¹ We use this data to check that the introduction of the CUT gender quota had bite.

We similarly construct a yearly panel with information on the gender composition of local union boards, the gender of their presidents and vice-presidents, and their affiliation to union centrals between 2012 and 2019. We use these data 1) to assign treatment status to unions; 2) to understand whether the reform had spillovers on local union boards; and 3) conduct heterogeneity analyses concerning women’s representation in unions.

C.2 Construction of variables

Amenities In the analysis we adopt two different ways of classifying clauses as female-centric amenities. The first is guided by intuition to select clause types that are of plausibly of greater value to women than men (intuitive definition). The second definition is data-driven, where we use lasso to pick clauses that are most predictive of women’s value of employment (relative to men) at an establishment in the cross-section. An important advantage of the data-driven approach—compared to the intuitive definition—is that it also identifies clauses that are valued relatively more by men, i.e., male-centric amenities.

We also generate four different outcome margins for clauses at the pair-year level. First, the *intensive margin (count)* measures the sum of the clause counts from the clause types categorized as either female- or male-centric in the corresponding contract. Second, the *intensive margin (sum of indicators)* measures the sum of clause type indicators for those categorized as either female- or male-centric in the corresponding contract. Third, the *extensive margin* simply indicates whether any female (or male) clause exists in the CBA of interest. Finally, we calculate the *share* of the intensive margin (count) relative to the total clause count in the CBA.

1. Intuitive definition: Guided by CUT’s “fight plan” and previous work documenting the value women place on flexibility (Goldin and Katz, 2011; Mas and Pallais, 2017; Maestas et al., 2018), we identified 4 themes as female-centric: 1) leaves; 2) maternity and childcare; 3) workplace harassment and discrimination; and 4) flexibility and part-time work. From these themes we restricted ourselves to select 20 clause types. These clauses are listed in Table A1—which includes clauses on maternity leave, childcare assistance, prevention of sexual harassment—all of which are conceivably of greater value to women than men.

⁵¹Developed by Fernando Meireles and posted on [GitHub](#).

2. Data-driven definition: The data-driven definition of amenities selects clauses that are most predictive of gender differences in the value of employment at an establishment, controlling for gender-specific wage premiums.⁵² In practice, we estimate the following cross-sectional specification using lasso:

$$V_j^F - V_j^M = \beta_w^F \psi_j^F - \beta_w^M \psi_j^M + \sum_{z \in Z} \beta_z a(z)_j + \epsilon_j$$

where V_j^G is the PageRank value of establishment j for workers of gender G , ψ_j^G is the establishment fixed-effect for workers of gender G at employer j from an AKM regression on wages, and $a(z)_j$ is the average clause count of amenity z (one among the 137 clause types) offered in the CBAs covering workers. We select the 20 clause types with the highest β_z and label them as “female-centric” amenities. Conversely, the 20 clause types with the lowest β_z comprise our “male-centric” amenities. Results are shown in Table 1.

PageRank values. To estimate PageRank values we take job spells of full-time workers, ages 18-54, on open-ended contracts, and earning monthly wages in private sector establishments from RAIS (2009-2016). For each gender, we find the largest strongly connected set of establishments based on worker flows, i.e., a link between two establishments is defined as having at least one inflow and one outflow. We restrict to establishments that have at least 10 hires overall, with at least one of these coming from non-employment. To solve for the vector of PageRank values (see Appendix D), we follow Morchio and Moser (2020) and only consider employment-to-employment flows to be month-to-month job transitions. In addition, we set the damping factor used in finding the fixed point in the linear system of normalized flows to 0.8—one of the standard values in computer science. That is, the “random surfer” moving through the labor market restarts his search at a new establishment with 80% probability. As shown in Sorokin (2018), PageRank values are unique up to an unknown multiplicative factor. Below we discuss robustness to assumptions about the multiplicative factor applying to women versus men to obtain $V_j^F - V_j^M$.

Wage premiums. To estimate the establishment fixed effect from AKM we take job spells of full-time workers, ages 18-54, on open-ended contracts, and earning monthly wages in private sector establishments from RAIS (2009-2016). For each gender, we find the largest strongly connected set of establishments based on worker flows, i.e., a link between two establishments is defined as having at least one inflow and one outflow. We restrict to establishments that have at least 10 workers (on average across years) and are observed at least 4 years in RAIS. Following Gerard et al. (2021), the model includes dummies for individual workers (α_i) and individual establishments (ψ_j), year dummies interacted with

⁵²Section 2.2 provides a detailed justification for this approach.

five education dummies, and quadratic and cubic terms in age interacted with the education dummies (X_{it})—see Appendix D. For the baseline year, the worker effects are measured as of age 40 to correspond to the approximate peaks of experience profiles. The establishment fixed effects for each gender—i.e., ψ_j^F and ψ_j^M —are normalized relative to the restaurant industry, where rents are assumed to be negligible.

Clause counts. To get a measure of $a(z)_j$ for each establishment, we take a yearly average of the number of clauses in each of the 137 clause groups found in sectoral CBAs negotiated between 2009 and 2016. To assign coverage from sectoral CBAs to establishments, we first need to map the signing employer association to the firms being represented. Using the equivalent of a FOIA request, we obtained the universe of establishments paying dues to employer associations. We then take sectoral CBAs and match them to all establishments paying dues to the signing employer association. The next step is to assign coverage only to establishments located in the geographic region specified in the CBA. Finally, to reduce overlap in CBA coverage, we exploit information on negotiated wage floors to assign a “main CBA” to each establishment-year.⁵³

Robustness. We check the robustness of our data-driven method on two dimensions: 1) two different ways of selecting the establishment sample used in the regressions: either a 50% random split-sample (used in our baseline approach) or the full estimation sample of establishments; and 2) three definitions of the gender gap in PageRank values, i.e., $V_j^F - V_j^M$. The first definition (used in our baseline approach) chooses the establishment with the smallest wage premium gap as the normalizing establishment, and then adjusts female values relative to the male values by multiplying the former by the ratio of the female-to-male PageRank values of the normalizing establishment. The second definition simply assumes the multiplicative factor is the same for both genders, i.e., no normalization is needed. The third definition uses a (within-gender) normalized index from 0 to 100 of V_j^F and of V_j^M .

Tables A2 and A3 show all the clause types selected by any of the combinations above. These tables also show how many of these 6 different combinations choose a given clause type as either female- or male-centric, as well as those selected under the baseline approach but adding state and industry fixed effects.

Labor market outcomes We briefly describe how we define the outcomes used for the establishment-level and incumbent worker-level analyses. While for all worker-level outcomes we use the main job spell, some establishment-level outcomes are constructed with all job spells. We first describe establishment-level outcomes derived with all job spells and then

⁵³Specifically, we first define an establishment’s “core union” to be the modal union involved in negotiating wage floors that have bite on the wage distribution. Among the CBAs negotiated by the “core union” in a given year, the “main CBA” is the one with the wage floor that has the largest mass of workers.

those derived using main job spells. Finally, we describe worker-level outcomes.

Establishment level outcomes - all job spells:

- Total employment. The total number of workers employed at an establishment in a given year.
- Share of women in the workforce. Share of women employed in a given establishment-year among all workers.
- Share of women in the probationary workforce. Share of women employed in a given establishment-year with less than 3 months of tenure among all workers with fewer than 3 months of tenure. Brazil's federal labor code allows for at most 3 months of probation, after which employment terminations imply severance payments.
- New hires. Number of workers recently hired by the establishment, defined as the number of workers employed in a given establishment-year with less than 12 months of tenure.
- Share of women among new hires. Share of women employed in a given establishment-year with less than 12 months of tenure among all workers with fewer than 12 months of tenure.
- Share of women among separating workers. Share of women among workers who separate from the establishment in that year. Separating workers are defined as those who are no longer employed at the establishment by the end of the year.
- Establishment exit. A dummy variable indicating whether the establishment does not appear in RAIS in 2017.

Establishment level outcomes - main job spell:

- Mean log wage. For any given worker subgroup, we take the mean of the wage outcome (defined below) in logs across all workers in the subgroup employed at the establishment in that year. This variable is defined for the following worker subgroups: women and men with more than 12 months of tenure, women and men with less than 12 months of tenure.
- Mean gender wage gap. The difference between the mean log wage for women and the mean log wage for men for a given establishment-year.

- Wage bill. The monthly wage bill for the establishment. That is, we sum the wage outcome (defined below) for all workers employed by the establishment in that year.
- Share of women poached in. Share of new female hires that are poached from another firm among all female workers. New hires are defined as workers with less than 12 months of tenure at that establishment in a given year. Poached hires are defined as workers who in the preceding year worked at another firm in RAIS, as opposed to being unemployed or out of the (formal) workforce.
- Age of female workforce. Mean age of female workers employed at an establishment in a given year.
- Tenure of female workforce. Average months of tenure of female workers employed at an establishment in a given year.
- Hours of female workforce. Average contracted hours of work per week of female workers employed at an establishment in a given year. Weekly contracted hours are those agreed upon hiring, and do not include overtime work.
- Education of female workforce. Average years of schooling of female workers employed at an establishment in a given year.
- Share of women among managers. The share of women among workers with an occupation code corresponding to a managerial role. Occupation codes corresponding to manager positions are those starting with 12, 13 or 14 (as per CBO: *Classificação Brasileira de Ocupações*).
- Maternity leave benefits. The share of women taking maternity leaves longer than 120 days among women employed at an establishment that start their maternity leave in a given year. We are able to identify women taking maternity leave thanks to detailed information on both the length and the reason of the three longest leave spells per job spell. We think that it is very unlikely that maternity leaves are not among the three longest leave spells in a year for a woman on maternity leave. For this reason we are confident that we are observing the near universe of maternity leave spells.
- Job protection after maternity. The share of women working at the same employer where they were working at the start of maternity leave by end-of-year for the year when their maternity leave ends, among women employed at said establishment who start their maternity leave in the same year.

- Injury leave. The share of workers taking leave due to a workplace injury among all workers employed at an establishment during a given year.

Establishment level outcomes - not in RAIS:

- CBA wage adjustments. The largest percentage wage adjustment negotiated among the firm-level CBAs covering an establishment. For years without a wage adjustment clause or without a negotiated CBA, the assigned wage adjustment is zero.
- Profit margin. The mean profit margin (in percentage terms) over 2012-2014 and 2015-2017. The sample is restricted to establishments reporting profit margin information to Orbis in both the pre- and post-reform periods.

Worker level outcomes - main job spell of incumbent workers

- Wages. The average monthly earnings that a worker makes during a job spell in a given year. We always use earnings in real terms by using the December CPI (i.e., the *Índice Nacional de Preços ao Consumidor* reported by IBGE) with 2015 as the base year.
- Retention. A dummy that indicates whether the worker is observed working at the baseline employer in any given year, where the baseline employer is defined as the (main) establishment of employment in 2014.
- Employed in formal sector. A dummy that indicates whether the worker is observed working in the formal sector in that year, i.e., they have a job spell registered in RAIS in that year.

D AKM and PageRank Model

Our data-driven approach to identify female- and male-centric amenities requires establishment-level estimates of gender-specific PageRank values and AKM wage premiums. This appendix presents the model underlying these estimates. For simplicity, we present the model without any reference to gender specificity. We also use establishment and firm interchangeably.

Denote \tilde{V}_j as the common value of employment for any worker i at firm j . Common value means that all workers agree on \tilde{V}_j such that a single job ladder exists ranking firms according to this value. All else equal, workers value higher compensation bundles so that one can write $\tilde{V}_j = h(w_j, a_j)$, where $h(\cdot)$ is strictly increasing in both the wage w_j and the amenity a_j arguments. The utility of workers from employment at the establishment, however, is heterogeneous and given by $u_{ij} = h(w_j, a_j) + \varepsilon_{ij}$, where ε_{ij} captures an individual's idiosyncratic preferences for working at j .

PageRank values

The starting point here is $u_{ij} = \tilde{V}_j + \varepsilon_{ij}$. In a market with only two firms and independently distributed type I Extreme Value ε_{ij} across workers, the probability that a worker prefers firm j over k is given by $\frac{\exp(\tilde{V}_j)}{\exp(\tilde{V}_j) + \exp(\tilde{V}_k)}$. With N workers and letting M_{jk} denote the number of workers choosing firm j over k , the following relation between employment decisions and valuations of firm-specific employment is simply $M_{kj}/M_{jk} = \exp(\tilde{V}_k)/\exp(\tilde{V}_j)$.

In a labor market with multiple firms $j \in \mathcal{J}$, the above condition imposes a restriction on each pair of firms, i.e.,

$$M_{kj} \exp(\tilde{V}_j) = M_{jk} \exp(\tilde{V}_k), \forall j \in \mathcal{J}. \quad (7)$$

Following Sorkin (2018), one can relax this condition by imposing a single restriction per firm that guarantees a consistent valuation of employers (e.g., no Condorcet cycles), as well as a unique set of firm-level values that best explains worker flows across establishments. Summing equation (7) across all employers and rearranging terms gives

$$\frac{\overbrace{\sum_{j \in \mathcal{J}} M_{kj} \exp(\tilde{V}_j)}^{\text{value-weighted entry}}}{\underbrace{\sum_{j \in \mathcal{J}} M_{jk}}_{\text{exits}}} = \underbrace{\exp(\tilde{V}_k)}_{\text{value}}, \quad (8)$$

which implies a single linear restriction per establishment.

The intuition behind equation (8) is that a valuable firm tends to be chosen over other valuable firms and has fewer workers leave it. This recursive definition of $\exp(\tilde{V}_j)$ is closely linked to Google's PageRank algorithm for ranking web-pages in a search. Along these lines, one can solve for $\exp(\tilde{V}_j)$ as a fixed point in a linear system. Moreover, a unique solution exists if the set of employers are strongly connected, i.e., an establishment has to both hire a worker from and have a worker hired by another establishment in the set.

AKM premiums

The starting point again is $u_{ij} = \tilde{V}_j + \varepsilon_{ij}$ but with the assumption that $\tilde{V}_j = \beta \log(w_j - b) + \eta \log(a_j - q)$. The parameters b and q are the workers' reference wage and amenity levels, and $\varepsilon_{i,j}$ refers to the idiosyncratic preferences from working at establishment j . Assuming that the $\{\varepsilon_{i,j}\}$ are independent draws from a Type I Extreme Value distribution and the number of establishments \mathcal{J} is very large, workers' choice probabilities are closely approximated by exponential probabilities.⁵⁴ Hence, the establishment-specific labor supply functions are approximated by:

$$\log(L_j) = \log(\lambda) + \beta \log(w_j - b) + \eta \log(a_j - q). \quad (9)$$

The employer's problem is to post the wages and amenities that minimize production costs given labor supply in (9). The posted wages and amenities are common to all workers since employers cannot discriminate on the basis of their idiosyncratic preferences $\{\varepsilon_{i,j}\}$.⁵⁵ The optimal choice is the solution to the following cost-minimization problem:

$$\min_{w,a} (w_j + \xi_j a_j) L(w_j, a_j) \quad \text{s.t.} \quad T_j f(L(w_j, a_j)) \geq \bar{Y}, \quad (10)$$

where ξ_j captures heterogeneity in the marginal cost of amenity provision across employers.

The first order conditions imply that the optimal compensation package is given by:

$$w_j = T_j f'(L_j) \mu_j \left(\frac{e_{wj}^L}{1 + e_{wj}^L + e_{aj}^L} \right) \quad (11)$$

$$a_j = T_j f'(L_j) \mu_j \left(\frac{e_{aj}^L}{\xi_j (1 + e_{wj}^L + e_{aj}^L)} \right). \quad (12)$$

Rearranging equations (11) and (12), one can write wages and amenities as weighted averages

⁵⁴The exponential probabilities are $p_j \approx \lambda \exp(\beta \log(w_j - b) + \eta \log(a_j - q))$, where λ is a constant common across all establishments in the market.

⁵⁵This asymmetry in information, rather than labor market concentration, is the source of monopsony power. Recall that \mathcal{J} is large so as to ignore strategic interactions in posting.

of the marginal revenue product of labor and their respective reference values, i.e.,

$$w_j = \left(\frac{\beta}{1 + \beta + e_{a_j}^L} \right) T_j f'(L_j) \mu_j + \left(\frac{1 + e_{a_j}^L}{1 + \beta + e_{a_j}^L} \right) b \quad (13)$$

$$a_j = \left(\frac{\eta}{\xi_j(1 + \eta + e_{w_j}^L)} \right) T_j f'(L_j) \mu_j + \left(\frac{1 + e_{w_j}^L}{1 + \eta + e_{w_j}^L} \right) q. \quad (14)$$

Assume a linear technology $f(L_j) = \theta L_j$ and price-taking employers in the output market to specify the marginal revenue product of labor: $T_j f'(L_j) \mu_j = T_j P_j \theta$. To simplify further, assume that reference wages and amenities are proportional to productivity ($b = \bar{b}\theta$ and $q = \bar{q}\theta$). Rearranging terms and taking logs results in

$$\log(w_j) = \log \left(\frac{\theta \bar{b}(1 + e_{a_j}^L)}{1 + \beta + e_{a_j}^L} \right) + \log(1 + \beta R_j^w) \quad (15)$$

$$\log(a_j) = \log \left(\frac{\theta \bar{q}(1 + e_{w_j}^L)}{1 + \eta + e_{w_j}^L} \right) + \log(1 + \eta R_j^a), \quad (16)$$

where $R_j^w = T_j P_j / [(1 + e_{a_j}^L) \bar{b}]$ and $R_j^a = T_j P_j / [\xi_j(1 + e_{w_j}^L) \bar{q}]$. With relatively small values of βR_j^w and ηR_j^a , log wages and log amenities are functions of a fixed worker component and a fixed establishment component as in Abowd et al. (1999)—henceforth AKM. Specifically,

$$\log(w_j) = \log \left(\frac{\bar{b}(1 + e_{a_j}^L)}{1 + \beta + e_{a_j}^L} \theta \right) + \beta R_j^w \quad (17)$$

$$\log(a_j) = \log \left(\frac{\bar{q}(1 + e_{w_j}^L)}{1 + \eta + e_{w_j}^L} \theta \right) + \eta R_j^a. \quad (18)$$

In short, equations (17) and (18) imply that the wages and amenities of workers can be written in the form $\log(w_j) = \alpha^w + \psi_j^w$ and $\log(a_j) = \alpha^a + \psi_j^a$, where $\psi^w = \beta R_j^w$ is an establishment-specific wage premium and $\psi^a = \eta R_j^a$ is an establishment-specific amenity premium. To separately identify these premiums from the worker fixed effects, one must focus on a set of firms that are connected through worker flows.

E Welfare Model

Following the CUT reform that increased female-centric amenities at CUT-affiliated establishments, we found that women separate from treated establishments less and queue for jobs there. These are revealed preference measures of firm value, suggesting that the reform causes treated establishments to become disproportionately valuable to women.

By how much did women’s welfare increase? To answer this question we adapt an approach measuring changes in welfare from the introduction of new or improved varieties in a consumer setting to our labor market setting. We model workers as choosing firms, just like consumers choose products. Because of the reform, the quality of CUT-affiliated firms is changing, differently by gender. From a modeling perspective, this is analogous to a situation in which the quality of certain goods is improving or when new, improved, good varieties are introduced in the market. This appendix describes the model used to estimate welfare change and the estimation strategy. It also discusses how the model underlies our data-driven classification of amenities.

Model

The model assumes that workers have CES preferences over firms. One advantage of the CES demand structure is that it can be microfounded using a continuum of workers making discrete choices over where to work—as shown in Anderson et al. (1992)—and derived below. This is a common way to model the labor market (Berger et al., 2022; Card et al., 2018; Lamadon et al., 2022; Sorkin, 2018).

Worker’s problem and solution A representative worker with CES preferences over \mathcal{J} firms chooses the number of hours to supply to each firm to maximize total income subject to a total hours constraint:

$$\max_{\{n_j\}} \sum_{j \in \mathcal{J}} w_j n_j \quad s.t. \quad \left[\sum_j (b_j n_j)^{\frac{1+\eta}{\eta}} \right]^{\frac{\eta}{\eta+1}} = N \quad (19)$$

where n_j is the number of hours worked at firm j , w_j is the wage at firm j , b_j is a taste-shifter governing the disutility of working at j , and η is the (constant) elasticity of substitution across firms. The parameter b_j captures workers’ valuation of firm attributes other than wages. The constraint is not hours but disutility-weighted hours. Because the representative worker solution is the same as aggregating discrete choices of a continuum of workers deciding where to work, n_j can also be seen as the measure of workers working at firm j .

Optimal labor supply to each firm is given by:

$$n_j^* = \left(\frac{w_j}{\tilde{W}} \right)^\eta \frac{1}{b_j^{1+\eta}} N \quad (20)$$

where \tilde{W} is a book-keeping term called the wage index, closely related to welfare (as seen below) and defined as follows:

$$\tilde{W} = \left[\sum_{j \in \mathcal{J}} \left(\frac{w_j}{b_j} \right)^{1+\eta} \right]^{\frac{1}{1+\eta}} \quad (21)$$

In addition, the share of “expenditure” (i.e., labor income) at any given firm is:

$$S_j = \frac{w_j n_j}{\sum_k w_k n_k} = \frac{\left(\frac{w_j}{b_j} \right)^{1+\eta}}{\sum_k \left(\frac{w_k}{b_k} \right)^{1+\eta}} \quad (22)$$

Wage index interpretation and welfare The wage index represents how much workers are paid to work one more disutility-adjusted hour and is thus a measure of worker welfare. This can be seen by taking the envelope condition around the optimal solution to the worker’s problem: $\sum_j w_j n_j = \tilde{W} N$. Formally

$$\frac{\partial}{\partial N} \sum_{j \in \mathcal{J}} w_j n_j^*(w_j, w_{-j}) = \tilde{W}$$

so that when \tilde{W} rises it means workers are now paid more for providing one additional unit of disutility-weighted labor supply, thereby increasing their welfare.⁵⁶

How welfare changes when firm attributes change When firms change attributes or amenities this changes the disutility of working there (b_{jt}). The change in welfare is measured by the ratio of the new and old wage indices:

$$\frac{\tilde{W}_t}{\tilde{W}_{t-1}} = \frac{\left[\sum_{j \in \mathcal{J}_t} \left(\frac{w_{jt}}{b_{jt}} \right)^{1+\eta} \right]^{\frac{1}{1+\eta}}}{\left[\sum_{j \in \mathcal{J}_{t-1}} \left(\frac{w_{j,t-1}}{b_{j,t-1}} \right)^{1+\eta} \right]^{\frac{1}{1+\eta}}} \quad (23)$$

where \mathcal{J}_t are the firms observed in period t .

⁵⁶In this way, the wage index is to welfare in the labor setting like the price index is to welfare in consumer theory. In consumer theory, the price index captures the cost of purchasing one util of utility. Welfare rises as it gets cheaper to purchase one more util.

The key challenge to estimating this change in welfare is that firm quality b_{jt} is unobserved or, in our case, is difficult to model because it would require specifying exactly how 137 clause types enter the worker’s utility function. However, as first shown in (Feenstra, 1994), assuming CES demand circumvents this problem because the welfare change depends on two terms that are observed in the data: 1) the wage index of firms whose quality (b_{jt}) remains unchanged and are “common” across periods; and 2) a variety-adjustment term to account for changes at firms that do change b_{jt} . That is, the welfare change is given by

$$\phi_{t-1,t} = \left[\frac{\lambda_t}{\lambda_{t-1}} \right]^{-\frac{1}{1+\eta}} \frac{\sum_{j \in \Omega_{t,t-1}} \left(\frac{w_{jt}}{b_{jt}} \right)^{1+\eta}}{\sum_{j \in \Omega_{t,t-1}} \left(\frac{w_{jt-1}}{b_{jt-1}} \right)^{1+\eta}} = \left[\frac{\lambda_t}{\lambda_{t-1}} \right]^{-\frac{1}{1+\eta}} \frac{\tilde{W}_t^*}{\tilde{W}_{t-1}^*} \quad (24)$$

Here $\Omega_{t,t-1} = \mathcal{J}_t \cap \mathcal{J}_{t-1}$ are firms common to both periods—in our case, non-CUT firms. The asterisk $*$ in W_t^* and W_{t-1}^* denotes that these are wage indices over the common set of firms. Finally, λ_t is the share of the wage bill at common firms in t (using wages at t).

To get an expression for $\tilde{W}_t^*/\tilde{W}_{t-1}^*$, we use Equations (21) and (22) to obtain

$$[\tilde{W}_t^*]^{1+\eta} = \frac{1}{S_{jt}^*} \left(\frac{w_{jt}}{b_{jt}} \right)^{1+\eta} \quad \forall j \in \Omega_{t,t-1} \quad (25)$$

Following Redding and Weinstein (2016), we take logs of both sides, difference over time, and sum over all $j \in \Omega_{t,t-1}$ to get

$$\log \left(\frac{\tilde{W}_t^*}{\tilde{W}_{t-1}^*} \right) = \log \left(\frac{\bar{w}_t^*}{\bar{w}_{t-1}^*} \right) - \frac{1}{1+\eta} \log \left(\frac{\bar{S}_t^*}{\bar{S}_{t-1}^*} \right) - \log \left(\frac{\bar{b}_t^*}{\bar{b}_{t-1}^*} \right) \quad (26)$$

where the bars indicate a geometric average and the last term is zero because we assume quality remains the same for these common firms. Thus, a change in welfare depends only on three terms that are observed in the data and η

$$\log \phi_{t-1,t} = -\frac{1}{1+\eta} \log \left(\frac{\lambda_t}{\lambda_{t-1}} \right) - \frac{1}{1+\eta} \log \left(\frac{\bar{S}_t^*}{\bar{S}_{t-1}^*} \right) + \log \left(\frac{\bar{w}_t^*}{\bar{w}_{t-1}^*} \right) \quad (27)$$

Microfoundation of CES demand using discrete choices

Following the CES demand in (Berger et al., 2022), workers’ utility for working at a firm has a component that is common across workers, encompassing wages and a common taste for the firm amenities, and an idiosyncratic shock that follows a logit distribution. Firms post utility offers—we don’t model the source of firm heterogeneity and assume that they exogenously differ. There is a unit measure of workers indexed by $i \in [0, 1]$. Each worker

has a disutility for working at firm j :

$$\nu_{ij} = \exp^{-\xi_{ij}} h_{ij} b_j$$

with ξ_{ij} iid across workers and drawn from a multivariate Gumbel distribution with parameter η . Each worker must earn $y \sim F(y)$, where earnings $y_i = w_j h_{ij}$. The worker chooses firm j to minimize disutility:

$$\min_j \{\log h_{ij} + \log b_j - \xi_{ij}\} = \max_j \{\log w_j - \log y_i - \log b_j + \xi_{ij}\}$$

Following McFadden (1973) on logit, the probability that worker i chooses to work at firm j is:

$$p_i(\tilde{w}) = \frac{\tilde{w}_j^{1+\eta}}{\sum_k \tilde{w}_k^{1+\eta}}$$

where $\tilde{w}_j := \frac{w_j}{b_j}$. The aggregate labor supply to firm j is then found by integrating the probability that a worker works at that firm times the hours supplied by that worker, over the mass of all workers:

$$\begin{aligned} n_j &= \int p_i(\tilde{w}) \cdot h_{ij} \cdot dF(y) \quad \text{where } h_{ij} = y_i/w_j \\ n_j &= \frac{\tilde{w}_j^{1+\eta}}{\sum_k \tilde{w}_k^{1+\eta}} \frac{1}{w_j} \int y_i dF(y) \\ &= \left(\frac{w_j}{\tilde{W}}\right)^\eta \frac{1}{b_j^{1+\eta}} N \end{aligned}$$

This is exactly the aggregate labor supply to firm j as in the representative worker's problem with CES demand. The last line follows from the fact that in equilibrium:

$$Y = \int y_i dF(y) = \sum_{j \in \mathcal{J}} w_j n_j^* = \tilde{W} N$$

Estimation

To get at welfare changes by gender, we estimate equation (27) separately for men and women. Starting from the same establishment-year panel that we use to study labor market outcomes, we compute the average earnings and total employment for each group of workers employed at an establishment during two periods: the pre-reform period (2012-2014), denoted by $t - 1$, and the post-reform period (2015-2017), denoted by t . To do that, we take averages of establishment level quantities across years.

We separately estimate each one of the terms in the right hand side of equation (27), that is, $\log\left(\frac{\lambda_t}{\lambda_{t-1}}\right)$, $\log\left(\frac{\bar{w}_t^*}{\bar{w}_{t-1}^*}\right)$ and $\log\left(\frac{\bar{S}_t^*}{\bar{S}_{t-1}^*}\right)$ and we combine them with an estimate of η that we calibrate from Felix (2022).

The ideal experiment to estimate the welfare change due to the CUT reform would be to randomly shock some labor markets with the reform while leaving other markets unaffected. Lacking this ideal setting, we estimate the welfare components from pre-post comparisons within establishments. As any pre-post strategy, we recognize that it might also pick up the effect of other things changing during the period under study that might affect wages or employment within establishments over time.

Changes in \bar{w}_t^* and in \bar{S}_t^* can be directly estimated through an establishment-level regression. Note that the difference in the log of the geometric mean of a variable x is equivalent to the average change in $\log(x)$ between t and $t-1$ across units. In our case

$$\log\left(\frac{\bar{w}_t^*}{\bar{w}_{t-1}^*}\right) = \log \bar{w}_t^* - \log \bar{w}_{t-1}^* = \frac{1}{N_\Omega} \left(\sum_{j \in \Omega_{t,t-1}} \log w_{jt} - \sum_{j \in \Omega_{t,t-1}} \log w_{jt-1} \right) = \mathbb{E}[\Delta \log w_{jt} | j \in \Omega_{t,t-1}]$$

where N_Ω denotes the number of firms in $\Omega_{t,t-1}$, that is, the number of comparison (non-CUT affiliated) firms. We estimate the component of welfare due to changes in \bar{w}_t^* as the average change in log wages across non-CUT establishments, captured by the coefficient β in the following regression:

$$\log w_{jt} = \alpha + \beta Post_t + \mu_j + \epsilon_{jt}, \quad j \in \Omega_{t,t-1}$$

where μ_j denotes establishment fixed effects and $Post_t$ is a dummy for the post-reform period (2015-2017). We estimate $\log\left(\frac{\bar{S}_t^*}{\bar{S}_{t-1}^*}\right)$ with a similar regression, using $\log(s_{jt})$ as dependent variable, where $s_{jt} = \frac{w_{jt}n_{jt}}{\sum_{k \in \Omega_{t,t-1}} w_{kt}n_{kt}}$.

To estimate the change in λ , we instead take a first order approximation around λ_{t-1}

$$\begin{aligned} \Delta \lambda_t = \lambda_t - \lambda_{t-1} &= \sum_{j \in \mathcal{J}} \frac{\partial}{\partial w_j} \lambda \cdot dw_j + \sum_{j \in \mathcal{J}} \frac{\partial}{\partial n_j} \lambda \cdot dn_j \Big|_{w_{t-1}, n_{t-1}} \\ &= \frac{\sum_{j \in (\mathcal{J} \setminus \Omega)} w_{jt-1} n_{jt-1}}{(\sum_{j \in \mathcal{J}} w_{jt-1} n_{jt-1})^2} \left(\sum_{j \in \Omega} n_{jt-1} \cdot dw_j + \sum_{j \in \Omega} w_{jt-1} \cdot dn_j \right) \\ &\quad - \frac{\sum_{j \in \Omega} w_{jt-1} n_{jt-1}}{(\sum_{j \in \mathcal{J}} w_{jt-1} n_{jt-1})^2} \cdot \left(\sum_{i \in \mathcal{J} \setminus \Omega} n_{jt-1} \cdot dw_j + \sum_{j \in \mathcal{J} \setminus \Omega} w_{jt-1} \cdot dn_j \right) \end{aligned}$$

where to simplify notation we use Ω in place of $\Omega_{t,t-1}$ to denote the set of non-CUT firms

(of measure N_Ω) and $\mathcal{J} \setminus \Omega$ to denote the set of CUT-affiliated firms (of measure $N_{\mathcal{J} \setminus \Omega}$).

We define $\tilde{s}_{jt} = \frac{w_{jt-1}n_{jt}}{\sum_{k \in \mathcal{J}} w_{kt-1}n_{kt-1}}$ and $\hat{s}_{jt} = \frac{w_{jt}n_{jt-1}}{\sum_{k \in \mathcal{J}} w_{kt-1}n_{kt-1}}$ and re-write the expression above as

$$\begin{aligned} \Delta\lambda_t &= N_\Omega(1 - \lambda_{t-1}) (\mathbb{E}[\Delta\tilde{s}_{jt}|j \in \Omega] + \mathbb{E}[\Delta\hat{s}_{jt}|j \in \Omega]) \\ &\quad - N_{\mathcal{J} \setminus \Omega}\lambda_{t-1} (\mathbb{E}[\Delta\tilde{s}_{jt}|j \in \mathcal{J} \setminus \Omega] + \mathbb{E}[\Delta\hat{s}_{jt}|j \in \mathcal{J} \setminus \Omega]) \end{aligned}$$

where $\mathbb{E}[\cdot]$ denotes an average across firms. Finally, because $\log\left(\frac{\lambda_t}{\lambda_{t-1}}\right) = \log\left(\frac{\Delta\lambda_t}{\lambda_{t-1}} + 1\right) \approx \frac{\Delta\lambda_t}{\lambda_{t-1}}$, we can write:

$$\begin{aligned} \log\left(\frac{\lambda_t}{\lambda_{t-1}}\right) &\approx \\ N_\Omega \frac{(1 - \lambda_{t-1})}{\lambda_{t-1}} (\mathbb{E}[\Delta\tilde{s}_{jt}|j \in \Omega] + \mathbb{E}[\Delta\hat{s}_{jt}|j \in \Omega]) &- N_{\mathcal{J} \setminus \Omega} (\mathbb{E}[\Delta\tilde{s}_{jt}|j \in \mathcal{J} \setminus \Omega] + \mathbb{E}[\Delta\hat{s}_{jt}|j \in \mathcal{J} \setminus \Omega]) \end{aligned}$$

We estimate the average change in \hat{s}_{jt} and \tilde{s}_{jt} across establishments with a within-establishment pre-post comparison. That is, we run four regressions of the form

$$y_{jt} = \alpha + \beta Post_t + \mu_j + \epsilon_{jt}$$

using as dependent variable \hat{s}_{jt} and \tilde{s}_{jt} —separately for CUT and non-CUT firms—and we combine these estimates with λ_{t-1} , N_Ω , and $N_{\mathcal{J} \setminus \Omega}$ which are directly computed from the data.

To obtain standard errors around total welfare and each one of the three welfare components, we bootstrap the entire estimation exercise 1000 times. In each bootstrap we re-draw with replacement a new sample of establishments from our initial sample and re-run the establishment-level regressions on the new sample.