

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

Viola Corradini Lorenzo Lagos Garima Sharma*

May 24, 2023

Abstract

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. Using a difference-in-differences design that exploits variation in affiliation to the federation, we find that "bargaining for women" increases female-centric amenities in collective bargaining agreements, which are then reflected in practice. These changes lead 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, wages, or firm profits. Our results suggest that changing institutional priorities can narrow the gender compensation gap.

Key words: gender gap; amenities; collective bargaining; unions

JEL Codes: J31, J33, J51, J52

*Corradini: Massachusetts Institute of Technology (email: vcor@mit.edu); Lagos: Brown University and IZA (email: lorenzo.lagos@brown.edu); Sharma: Massachusetts Institute of Technology (email: gsharma@mit.edu). We thank Daron Acemoglu, David Atkin, David Autor, Abhijit Banerjee, Esther Dufo, Brandon Enriquez, Henry Farber, and Simon Jäger, as well as seminar participants at MIT, Brown, Princeton, NBER Labor Studies 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, and Iacopo Morchio and Christian Moser for their 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 alike have argued that making workplaces more female-friendly is key to lowering gender disparities—for example, Goldin (2014) argues that changing the structure of jobs may cause all remaining gender earnings gaps to vanish—little is known about if and how labor market institutions can be redesigned from within to ameliorate the stark trade-offs faced by working women.

Per one view, making workplaces female-friendly—providing maternity leave, childcare, and flexible work schedules—is not worth the expense to employers since the marginal worker does not value them enough. This paper tests an alternate view: that, instead of the marginal worker’s preference, the priorities of those designing compensation determine workplace amenities. Because a few individuals typically decide workplace policies, their priorities take precedence and may not always feature women’s needs on top. When these priorities change, so too do workplaces. Unions provide a natural setting in which to test this hypothesis since, for nearly 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

This paper investigates how changing leaders’ priorities in women’s favor changes the workplace. The ideal experiment to study this question requires a top-down change in priorities that is uncorrelated with changes to a firm’s labor demand or workers’ preferences. We exploit such a natural experiment in Brazil, that spurred leaders of its 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, the CUT reserved half its leadership positions for women and emphasized the provision of female-focused policies, such as 6 months of paid maternity leave, flexible work schedules, and childcare. Because unions seldom change affiliation to a union central, and neither workers nor establishments choose their union, the CUT reform represents a top-down pro-women directive to union leaders that is unrelated to an establishment’s labor demand or supply. This motivates using a difference-in-differences design to compare amenities and costs (wages, employment) at establishments negotiating with CUT unions (treated) to non-CUT affiliates (comparison). The two sets

¹For example, nearly half of all workers but only 12% of union leaders in Brazil are women. In continental Europe, where collective bargaining covers a majority of workers, including Germany, Austria, and the Netherlands, less than 30% of union members are women (Skorge and Rasmussen, 2022).

²Union centrals are umbrella organizations that coordinate priorities among local unions. Over half of all formal workers in Brazil are covered by collective bargaining and 20% of unions affiliate with CUT.

of establishments closely resembled each other at baseline; together they comprise 19% of formal employment in Brazil, and employ 11.5 million workers across 80,000 establishments.

Unique to the Brazilian setting, our analysis relies on linking three rich sources of data: (i) establishment-level amenities from the text of all collective bargaining agreements (CBAs), (ii) worker outcomes from linked employer-employee records covering all formal employment (RAIS), and (iii) union affiliation and leadership covering all unions. CBAs offer uniquely high quality information on 137 different amenities offered by establishments, including maternity leave, workplace safety, absences, and work hours. The administrative data track workers over time and report their gender, wages, and instances of maternity leave.

We begin 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 enabling work-life balance, including maternity protections, childcare payments, absences, and workday reductions (“female-centric” amenities). In contrast, men value higher pay and safety, such as clauses governing profit sharing, hazard pay, life-insurance, and safety equipment (“male-centric” amenities).³ In an out-of-sample sense check, we find that female amenities increase—and male amenities decrease—with the share of women in an establishment’s workforce, providing the first clue that representation could influence amenities.

The second part of our analysis studies the causal effect of shifting union leaders’ priorities on female and male-centric amenities, and its downstream effect on workers and establishments, on their wages, retention, and employment.

Our first main takeaway is that female-centric amenities increase on paper and in practice. On paper, we find a 19% increase in female-centric amenities. This is a large improvement, equivalent to moving from the average baseline amenity count at a minority female establishment to one where over 80% of workers were women. Provisions governing leaves and childcare account for much of the gain, suggesting that the reform spurs benefits especially for childbearing women. The largest gains occur at establishments where women had limited voice at baseline, either by being in the minority among workers or among union leaders.

Amenity improvements on-paper translate into practice. Following the reform, women at treated establishments take longer maternity leaves, enjoy job protection following these leaves, and constitute a larger share of managers.

³We mitigate simultaneity bias, i.e., that employers increase female-centric amenities when wanting to hire more women, by using amenities from sectoral agreements negotiated with multiple employers in an industry instead of firm-level agreements negotiated with a single employer. Unlike the latter, sectoral CBAs are unlikely to be influenced by employer-level demand shocks.

Our second main takeaway is that women value changes to the work environment induced by the CUT reform, ruling out a pure compensating differences story for the amenity gains. Specifically, we find women separating less from and queuing for jobs at treated establishments, both of which are revealed preference measures of firm value (Krueger and Summers, 1988; Holzer et al., 1991). Retention among women increases by 10% and reflects a decline in voluntary separations. While we do not directly observe job queues, we proxy for them using probationary contracts that are commonly used by employers to screen applicants. Women’s share among probationary workers rises by 10%. In sum, higher female-centric amenities cause women to flock to CUT-affiliated establishments.

Our third main takeaway is that these improvements for women manifest without observed tradeoffs in wages or employment. Compensating differences would suggest that women’s wages fall to finance amenity improvements (Rosen, 1986). However, we find no effect on the earnings of either new or incumbent workers, ruling out even very small declines with a high degree of confidence. Given no wage change, establishments may reduce women’s employment because they are now more expensive to employ. We find no evidence of this; employment remains unchanged. Instead, women comprise a larger share of the workforce (by 0.2pp relative to 36% at baseline).

If women are not losing, perhaps men are. However, there is little evidence of this. We find no decline in the earnings or employment of incumbent male workers. Male amenities do not decline. If anything, there is a small positive treatment effect on retention among incumbent male workers, suggesting that men *value* the changes to the work environment spurred by the CUT reform. Overall, our findings are consistent with a model of the labor market wherein firms post utility offers for each gender (e.g. Card et al. (2018); Berger et al. (2022)). The reform causes this posted utility for women to rise without a corresponding decline in men’s utility.

If workers do not finance amenity improvements, perhaps firms finance them through lower profits. Both the empirical evidence and theoretical reasons point against this explanation. Empirically, we find no treatment effect on establishment exit—which is a non-trivial margin of adjustment in Brazil, with 8.7% of control establishments having exited within two years of the reform. For the subsample of establishments that report to Orbis, we find no evidence of a decline in measured profits. Theoretically, the CUT reform shifted union priorities rather than raising unions’ bargaining power. As such, 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, we find a precisely estimated zero effect.

How, then, are these amenity gains paid for? One explanation is that the union shifts

rents from men to women (albeit not observed in the data as wage or amenity declines), but men are not marginal to these rents as they would not obtain them elsewhere. A second possibility is that bargaining was inefficient ex-ante and changing union priorities led to a pareto improvement for workers and firms. At least two general models could explain these results. In one, frictions in the bargaining process or in aggregating workers’ interests to the union level (e.g., information or contracting frictions) yield the possibility for win-win situations once union attention is refocused on previously ignored issues. In another model, behavioral firms and unions did not conceive of providing female-centric amenities until changing union priorities put these issues front-and-center.

The final part of our analysis develops a revealed preference method to quantify the welfare effect of changing the work environment by drawing an analogy with consumer theory (Feenstra, 1994; Redding and Weinstein, 2016). Just as gains to consumer welfare from improving product varieties are quantifiable via an increase in the expenditure share on these improving varieties, gains to worker welfare from improving workplace amenities are quantifiable via an increase in the wage bill at these amenity-improving employers. In other words, workers vote with their feet toward desirable employers. A few sufficient statistics then quantify gains in welfare.⁴ This sufficient statistics approach allows us to remain agnostic regarding the precise functional form linking amenities to utility. Consistent with our reduced-form findings, we find that the CUT reform raises women’s welfare by 6% while leaving men’s welfare unchanged.

This paper contributes to four literatures. First, on unions and inequality. While firms care about the marginal worker, it is unclear who the union cares about (Farber, 1986). Unions have long struggled to organize workers with competing interests (Hill, 1996) 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’ priorities determine whose interests they represent. When unions prioritize women, they can lower within-firm gender inequality.

Second, on the importance of leaders’ priorities in how institutions function. Political

⁴For tractability, we assume that workers possess nested CES preferences over employers. Just as gains to consumer welfare from improving product varieties are quantifiable via changes to the price index, i.e., change in the cost of purchasing an additional util of utility, gains to worker welfare from improving workplace amenities are quantifiable using changes to a wage index, i.e., change in the wage for working one disutility-weighted hour. Under CES, only four sufficient statistics quantify these gains: an increase in the share of labor income at treated establishments (capturing workers flocking to amenity-improving employers), workers’ elasticity of substitution across establishments (capturing how difficult these moves were), change in the dispersion of labor income at comparison establishments (capturing where workers are drawn away from), and any change in the wage at comparison establishments (capturing pro-competitive responses).

leaders are found to 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, here too, leaders could step in on their behalf. While women on company boards have been found to have limited 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 to union leaders’ priorities alter the workplace, in this case making it better for women.

Third on whether providing female-focused amenities leads employers to lower women’s wages (Gruber, 1994) or reduce hiring when they cannot (Summers, 1989). We find no evidence of this: although the work environment improves for women, we cannot reject the null that their wages and employment do not suffer (and, indeed, rule out very small declines with high confidence). By way of benchmark, Lagos (2021) estimates that leave clauses—many of which emerge as female-focused in our revealed preference approach—are valued at 8.4% of a worker’s wage. Instead, although recent work demonstrates limited gains for workers from greater voice on corporate boards (Harju et al., 2021; Blandhol et al., 2020), we find substantial gains from elevating women’s voices on union boards. One exception is Boudreau (2023), who finds that elevating worker voice through Occupational Safety and Health (OSH) committees in Bangladeshi garment factories has a small, positive effect on workplace safety without detectable impacts on wages and employment.

Finally, our paper contributes to the revealed preference literature in three ways. First, we provide quasi-experimental evidence that workers move toward improving amenities, consistent with several 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). Second, we use worker moves and variation in amenities across establishments to identify what workers value, using a richer set of amenities and higher stakes environment than possible in experiments. Encouragingly, our results match this experimental work—in particular, women value flexibility (Mas and Pallais, 2017; Wiswall and Zafar, 2017; Maestas et al., 2018). Finally, we bring a revealed preference approach from consumer theory to the labor setting to quantify the welfare effect of changes to a firm’s work environment. This sufficient statistics approach remains agnostic regarding the precise functional form linking a rich set of amenities to utility; it can be similarly employed in future work to quantify the welfare effect of a change to the work environment.

The rest of the paper proceeds as follows. Section 1 discusses the institutional context and CUT reform. Section 2 describes the data and details our approach for classifying amenities as female- or male-centric. Section 3 presents our empirical strategy. Section 4

presents our main results on the causal effect of changing union priorities on amenities, revealed preference measures of firm value, and labor market outcomes. Section 5 quantifies the welfare impact of improving (female-centric) amenities on men and women. Section 6 concludes.

1 Institutional Context

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

1.1 Collective Bargaining and Union Centrals

Types of CBAs Brazil has two types of collective bargaining agreements (CBAs): sectoral and firm-level. In sectoral CBAs, unions negotiate with employer associations representing establishments in a specific industry and geography, for example, the car manufacturers in Curitiba. In firm-level CBAs, unions negotiate with individual employers, for example, Volkswagen. Given their wider coverage, sectoral agreements typically set general floors for wage and non-wage benefits. By contrast, firm-level agreements generally build on these floors to expand benefits for workers at individual employers (Horn, 2009). Our main analysis studies the impact of the CUT reform on firm-level CBAs. However, we leverage amenities contained in sectoral CBAs to identify the clauses that are highly valued by female and male workers (Section 2.2).

Union determination The union that negotiates CBAs on behalf of workers at any given employer is chosen neither by the workers nor by the employer. Rather, representation depends on two factors: 1) industry (or category); and 2) geographic location (municipality).⁵ Examples of unions include the bank workers’ union of São Paulo and the teachers’ union of Florianopolis.

Neither workers nor employers can change their union. As a legacy of Brazil’s corporatist past, the first union approved to represent a given category-geography cell enjoys a

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

lifetime monopoly.⁶ As such, workers can only influence their union priorities from within, for example, by voting in union elections, running for union leadership, or voicing their concerns to union leadership. At the same time, employers cannot avoid unions by virtue of this predetermined assignment of the same union to all employers in a category-geography cell. Naturally, union assignment by these cells produces an incredibly fragmented union landscape, with over 11 thousand labor unions operating in Brazil.⁷

CBA coverage Neither workers nor 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.⁸ 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. CBA clauses therefore build on top of these basic guarantees that are enjoyed by all workers.

Negotiation process Unions' priorities play a central role in determining the content of CBA negotiations. Before the expiration of an existing CBA, 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 begin when the union sends these demands to employers. They occur over several rounds. Most CBAs are signed for a duration of 12 months, giving rise to annual negotiations.⁹ The union board also decides which representatives sit at the bargaining table, which is not restricted exclusively to board members.

Union centrals Unions can affiliate with union centrals (or *centrais sindicais*), which are somewhat analogous to trade union federations such as the AFL-CIO in the United States. These centrals are national level, umbrella organizations that coordinate the activities of

⁶President Getúlio Vargas instituted this “monopoly union” framework, known as *unicidade sindical*, in the late 1930s as a means to co-opt the labor movement by enabling the federal government to control the union given the right to represent workers in collective bargaining.

⁷It'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.

⁸Despite universal coverage, CBA coverage in Brazil is around 50% partly because not every union has a CBAs covering all (sometimes any) of the municipalities they represent.

⁹In some cases negotiations occur every two years—the maximum possible duration for a CBA.

local unions and lobby for political favor (Liukkonen, 2019). While union centrals do not directly participate in collective bargaining, they are indirectly involved in coordinating union priorities *across* worker categories. For example, union centrals regularly organize general strikes, plan annual conferences attended by union representatives, provide support to local unions, participate in public discussion forums on behalf of constituent unions, and steer union attention toward broad priorities such as 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 organizations, representing 30.4% percent of all organized workers in Brazil as of 2016.¹⁰ CUT is Latin America’s largest union central, and among the largest in the world. It has close ties with the *Partido dos Trabalhadores* (PT), or 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 his move into politics—a common path for PT politicians (Lang and Gagnon, 2009).

CUT has vertically organized 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.¹¹ They meet once every three years to vote on CUT’s overarching priorities for the subsequent 3 years, recorded in a book of resolutions or “fight plan”. Executive boards comprise a smaller group of leaders elected by congresses to oversee CUT’s day-to-day functioning. They manage CUT’s finances, oversee the execution of the fight plan, organize meetings and training for local union leaders, and organize committees to tackle specific topics like gender equality within CUT.¹²

1.2 CUT Reform

The origin 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 its female presidential candidate, Dilma Rousseff, was elected as Brazil’s first woman president. Together these events spawned a demand for greater gender parity within the CUT. Prominent female CUT leaders authored op-eds demanding greater say for

¹⁰The 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 Classe Trabalhadora*, and *Central Sindical e Popular Conlutas*.

¹¹Elected 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](#).

¹²For 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.

women within CUT’s leadership and a similar quota for women in the union (Godinho Delgado, 2017). They were successful. CUT’s 2015 state and national congresses witnessed an unprecedented focus on women and instituted a pro-women reform that 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% 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 of around 21-25% (averaged across union centrals).¹³

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 employs a difference-in-differences strategy to compare the share of women on the local union boards of CUT and non-CUT affiliated unions. There is a small positive treatment effect, of 3% relative to baseline. This estimated effect is not mechanical as the quota only applied at the union central level and not also for its affiliates. Hence, this 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, coordinating the activities of the affiliated unions and their bargaining priorities.

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 mandate of 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 fight plan, compared with 46 occurrences in 2012 and 74 in 2009.

CUT’s 2015 fight plan also detailed a series of measures to promote gender parity within local unions. These included giving women chairmanship of important committees (like fi-

¹³The small increase in the share of women on the boards of non-CUT union centrals in 2015 is driven by *Conlutas*—an even more combative left-leaning union central with a very small number of affiliated unions. CUT’s main competitor union central is *Força Sindical*, which saw a small decline in the share of women on its national board in 2017 (Figure B1).

nance and communications) and involving women in the drafting of *pautas de reivindicações*, i.e., the list of union demands which are taken to employers for negotiation.¹⁴ Therefore, independent of any change in women’s representation on local union boards, these recommendations potentially translated into practices that elevated women’s voices within local union boards.

Summary In sum, starting in 2015, the CUT had more female leaders and vowed to advocate for women’s priorities at the bargaining table. It made its commitment to gender equality especially evident to the local union delegates who attended its congresses. Importantly, the CUT reform did not change the bargaining power of unions relative to employers but merely refocused union priorities toward women. Hence, any improvements for women realized due to the reform are likely to reflect these new priorities, as opposed to a change in the share of surplus accruing to workers.

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. This section first describes the data that satisfy these requirements. We then detail our data-driven approach to classifying amenities as male- or female-centric.

2.1 Data Sources

Our analysis relies on linking three rich sources of data: (i) amenities at the establishment-level from the text of all CBAs; (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 of each type as a measure of amenities offered to workers.

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

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

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 each employer to report key information regarding each worker employed in any given year. For each work spell, RAIS reports average monthly earnings, leaves taken, and (6-digit) occupation. It also reports worker characteristics like gender, age, and education; and establishment characteristics such as location (municipality) and industry (6-digit). We link RAIS to CBAs using an establishment identifier, known as CNPJ, that is common to both datasets.

For information regarding each union’s affiliation to a union central and its leadership composition over time, we use the national registry of unions, known as *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 in the Brazilian census are women (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 a comprehensive set of amenities provided at each job. However, whether a CBA clause is differentially valued by women relative to men (what we denote as a female-centric amenity) is not directly observed in these data. We adopt two approaches to classify clauses as female-centric. Here we describe the key steps of each approach, with details in Appendix C.

1) Intuitive approach In the intuitive approach, we classify 20 of the 137 pre-specified clause types in *Sistema Mediador* as disproportionately valued by female workers (Table 2, Column 1). They fall into four broad themes, 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) include clauses that one could reasonably associate with women. The last theme reflects the fact that women disproportionately value flexibility

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

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 a common ranking over establishments $j \in \mathcal{J}$. A worker’s utility from working at establishment j is rising in the wage and amenities that it offers to their group G . In particular, we assume that the gender-specific value of working at an establishment (denoted V_j^G) is a linear function of wages, amenities, and an unobserved component:

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

where Z denotes the set of all amenities. Our classification problem is then to find the set of amenities for which the difference $\beta_z^F - \beta_z^M$ is positive, which we denote as “female-centric”, as well as those for which this difference is negative, which we denote as “male-centric”.¹⁷

At a minimum, we must measure the value of employment, wages, and amenities provided at each establishment. For the value of employment, we estimate gender-specific PageRank values by leveraging worker flows across establishments (Sorkin, 2018; Morchio and Moser, 2020). This is a revealed preference measure of the value of working at an establishment, which relies on the idea that good employers attract more workers, especially from other good employers.¹⁸ For wages, we estimate the gender-specific wage premium at an establishment (ψ_j^G) using gender-specific AKM models.¹⁹ For amenities, we use the average annual count of clauses $a(z)_j$ for each of the 137 clause types $z \in Z$ included in CBAs covering establishment j .

Hence, while we measure the gender-specific value of employment and wage premia at each establishment, we only observe a proxy for amenities without knowing which clause types are disproportionately valued by women and which by men. To identify these clauses, we take the difference between the female and the male version of Equation (1) and estimate the following hedonic regression:

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

¹⁷An advantage of the data-driven approach relative to the intuitive approach is that it identifies male-centric clauses, allowing us to test for tradeoffs in male amenities following the CUT reform.

¹⁸Appendix D describes the approach in detail and Appendix C describes our implementation.

¹⁹AKM is the acronym for Abowd et al. (1999), which is the original paper estimating firm-specific wage premia using linked employer-employee data. Their underlying model also assumes a common job ladder among workers and identifies the firm effect using worker flows (see Appendix D for details and Appendix C for implementation).

where $\beta_z = \beta_z^F - \beta_z^M$ captures the value of the amenity for women relative to men. We estimate this regression 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 “female-centric”, and the bottom 20 with the lowest β_z “male-centric”. To the best of our knowledge, this is the first time that such a rich description of the work environment has been combined with administrative data on worker flows to uncover which features of the workplace are valued by different groups of workers.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 more hire women may redouble their recruitment efforts or provide other job features that are valued by women, in addition to increasing observed clauses. Because we do not directly observe recruitment intensity or perfectly observe the work environment, we may erroneously identify a clause as valuable because it covaries with these unobserved features.²¹ To mitigate this bias, we use amenities $a(z)_j$ from sectoral CBAs negotiated with several employers in an industry and geography instead of firm-level agreements negotiated with a single employer. Unlike the latter, sectoral CBAs are not influenced by demand shocks affecting individual employers.²² Using sectoral CBAs for classification is also important because we use firm-level CBAs to study the CUT reform’s causal effect. Using separate CBAs for classification and analysis prevents a mechanical relationship between clauses identified as female-centric and those that increase after the reform. Women flocking to treated establishments following the rise in female-centric amenities is then not a pre-determined result.

Estimation sample We estimate Equation (2) in the cross-section of establishments for which we can estimate V_j^G , ψ_j^G , and $a(z)_j$. First, because we must observe PageRank values for both genders, which can only be estimated for the largest super-connected set of employers (i.e., each establishments must hire from and lose a worker to another establishment in the set), our sample is restricted to the 2009-2016 intersection of these gender-specific super-connected sets. Second, AKM wage premia are only estimated for the largest connected set of establishments for which estimates are not noisy (average size ≥ 10 workers). The

²⁰Several papers elicit workers’ willingness-to-pay for a small set of workplace attributes such as 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 undistinguished by gender.

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

²²The results are not driven by industry-specific amenities and are similar when including industry fixed effects to leverage variation across geography; see footnote 23.

sample is thus also restricted to the 2006-2016 intersection of these largest connected sets between genders. Third, we reduce noise in $a(z)_j$ (i.e., the over-year average of clause type z), by restricting the sample to employers covered by at least four sectoral CBAs between 2009-2016.

Normalization Both PageRank values and AKM wage premiums must be normalized to make the gender difference in them interpretable. For AKM premiums, we normalize ψ_j^F and ψ_j^M relative to the restaurant sector—a fairly competitive industry where one can reasonably assume a zero wage premium for both genders. For PageRank values, V_j^F and V_j^M are unique up to unknown multiplicative factors. Our results are robust to three alternative methods for calculating $V_j^F - V_j^M$. The first chooses the establishment with the smallest gender gap in wage premiums as the normalizing establishment, and divides the female value of all other establishments by the ratio $\frac{V_j^F}{V_j^M}$ at this establishment. The second simply assumes the multiplicative factor is the same for both genders, i.e., no normalization is needed. The third method re-scales the values V_j^F and V_j^M to a scale from 0 to 100. The base method for identifying male and female-centric amenities in the data-driven classification uses a 50% random sample of establishments and the first method for normalizing PageRank values.

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

In line with the intuitive definition, the data-driven approach reveals that women disproportionately value clauses governing leaves (e.g., following adoption and miscarriage), childcare, and maternity (e.g., childcare assistance, maternity protections, and policies for dependents). In addition, they value 12 other provisions missing from the intuitive classification, including absences, extensions or reductions of the workday, medical exams, and health education campaigns.

On the male side, we also obtain sensible results. Men highly value additional pay, such as clauses governing on-call pay, profit sharing, hazard pay, workday compensation, life insurance and death or funeral assistance. They also disproportionately value workplace safety, such as protections for injured workers, machine and equipment maintenance, and safety equipment.²³

²³The clauses classified female-centric remain similar across various normalizations of PageRank values (Tables A4 and A5). Moreover, the classification is not driven by industry or geography-specific amenities, since it is invariant to including industry- and state- fixed effects. The rank correlation of the coefficient β_z on the selected clauses with and without these fixed effects is positive and statistically significant (0.56 with p -value < 0.01). Tables A2 and A3 offer specific examples of clauses identified as female and male-centric.

The fact that “female workforce” clauses appear among those disproportionately valued by men highlights the fact that our approach does not account for variation in the text of clauses. These “female workforce” clauses vary widely in content, including items that are clearly beneficial to women (e.g., free provision of sanitary pads), as well as those clearly beneficial to men (e.g., forbidding women to cast concrete or install scaffolding). It is likely, then, that our data-driven approach captures the latter. While the availability of pre-specified clause types allows us to have a simple measure of CBA content that avoids the drawbacks that plague more complicated topic models—such as text pre-processing, 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 identify clauses that women (or men) value disproportionately more than the other gender. Using firm-level CBAs signed in 2014—the year prior to the CUT reform—we find that female (male)-centric clauses increase with the share of women (men) at an establishment.²⁴ Figure 3a shows that intuitively classified female-centric clauses increase almost linearly with this share. Figure 3b shows a similar relationship for male and female-centric clauses defined using the data-driven method. Specifically, all-male workplaces offer ≈ 1.5 more male than female clauses, with this gap shrinking to almost zero at all-female workplaces. Interestingly, female clauses per the data-driven classification only begin to increase once women comprise the majority in an establishment (above the 50% threshold). This suggests either that women successfully advocate for these amenities once in the majority, or that establishments provide them to attract female workers—both implying higher value among women.

3 Empirical Strategy

We employ a difference-in-differences strategy to study the CUT reform’s effect on amenities and labor market outcomes. This section first describes the three analysis samples we use to study the reform’s effect on collective bargaining agreements, establishments, and workers. We then detail our empirical approach and identifying assumptions.

3.1 Analysis Samples

We construct three analysis samples to study the CUT reform’s effects on negotiated CBAs, establishments, and workers. Appendix C provides detail.

²⁴In addition, the number of female clauses is strongly positively correlated with the difference between women and men’s PageRank valuation of an establishment (Figure 3).

1) Amenities sample To study the evolution of amenities, we construct a balanced panel of each pair of establishment-and-negotiating union covered by firm-level collective bargaining between 2012 and 2017. Each of these pairs can be thought of as constituting a unique worker group, because the same union represents any category (usually industry) of workers in a given geography.²⁵ Our analysis focuses on clauses in firm-level CBAs because most improvements in amenities and working conditions are achieved through these agreements (Horn, 2009; Liukkunen, 2019).²⁶

While not every establishment-union pair renegotiates its contract every year, we obtain a balanced panel of contracts by exploiting the fact that, during our period of study, the coverage of old CBAs is automatically extended until a new agreement is negotiated (Lagos, 2021). Given both that all CBAs were required to be registered in *Sistema Mediador* beginning in 2009, and that they span at most 2 years, our panel paints an accurate picture of active CBAs between 2012 and 2017. Our results are robust to instead using an unbalanced panel that comprises only new contracts.

2) Establishment sample To study the possible downstream effects of changing amenities on labor market outcomes as well as wage and employment tradeoffs, we construct a sample of establishments signing CBAs in our *amenities sample*, and track their outcomes in RAIS. Outcomes include employment, the share of women among workers, and mean log wages. We make two additional restrictions to this sample. First, we restrict attention to establishments that employed both men and women at baseline (2014). Second, we only consider an establishment signing a contract as covered by its contents if it lies within the contract’s geographic coverage. This restriction allows us to exclude headquarters that sign contracts on behalf of their subsidiaries, and are hence outside the contract’s geographic coverage.

3) Incumbent worker sample To study individual worker-level outcomes such as 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 their baseline employer.

²⁵Most signing establishments (93%) negotiate with a single union over the entire study period, meaning that employers rarely negotiate with more than one worker category.

²⁶In an informal conversation, the President of the bankers’ union of São Paulo also confirmed that most amenity improvements are achieved through firm-level CBAs. Sector-level negotiations typically involve several tens (or even hundreds) of employers, making it difficult to reach consensus over a rich set of amenities. Unions therefore typically reserve these topics for negotiation with individual employers.

Treatment definition Following the 2015 reform, CUT-affiliated unions prioritized women in their collective bargaining strategy. While the reform was enacted in 2015, the gender quota was approved in 2012 (see Section 1.2), suggesting that CUT’s pro-women pivot may have been anticipated and spurred unions to switch affiliation to avoid or benefit from the pivot. Although unions rarely switch their 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 confirms that neither treated nor comparison unions systematically switched affiliation away from or toward the CUT following its 2012 announcement of a gender quota. Thus, there is no concern from endogenous selection even had we used a later affiliation year.

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 the CUT in 2012. In the *establishment sample*, a treated establishment is one belonging to treated pair.²⁷ Finally, in the *incumbent worker sample*, a worker is treated if employed at a treated establishment at baseline (2014).

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

Panel A reports sample sizes. Our sample comprises more than 211 thousand firm-level CBAs signed by 89,920 establishment-union pairs. These pairs cover 80,131 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). Of all pairs, 21% are treated and 79% are in the comparison group.

The amenities sample covers over 19% of total formal employment in Brazil, and 2.1% of establishments. These numbers highlight two points. First, only a select set of employers negotiate firm-level CBAs. Second, these establishments are substantially larger than the average establishment in Brazil, employing 143 workers on average compared to 16 among all establishments (Table A6).²⁸ The establishment sample, where establishments must additionally have been employing both men and women in 2014, covers 15% of the total 2014 workforce, and otherwise resembles the amenities sample in the size, sector and regional distribution of its establishments.

²⁷Over 93% of establishments negotiate with a single union and 98% with all unions with the same union central affiliation. For the remaining 2% of establishments, treatment is defined as negotiating with any treated union.

²⁸Compared to the average Brazilian establishment, an establishment signing firm-level CBAs is more likely to operate in manufacturing rather than commerce (difference of 16-19pp for each); these establishments are more likely to be located in the affluent Southeast and less in the poorer Northeast region of Brazil (Table A6).

Panel B of Table 1 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 classified “female-centric” per our data-driven definition (Section 2.2). On average, contracts feature 1.7 more male clauses than female clauses. These numbers are statistically indistinguishable across treated and control contracts. Although the share of female-specific clauses may appear to be small, this statistic may not accurately represent the value and importance of these clauses. For example, even a single contract provision extending maternity leave by 60 days may prove very valuable to young women. Thus, in addition to considering how the CUT reform affects amenities on paper, we will infer how valuable these changes are to women by studying revealed preference changes in their sorting behavior across establishments.

Panels C and D document establishment- and union-level characteristics, respectively, at baseline (2014). Our sample comprises large employers (especially in the treated group). The average establishment employs 143 workers, over a third of whom are women. A majority of establishments employ both men and women (82%). On the union side, treated unions have larger boards but with a similar share of women as comparison unions (around 23%), indicating no baseline difference between CUT and non-CUT affiliates. Only about 17% of unions have a female president.

Treated and comparison establishments exhibit substantial overlap along a number of observable dimensions, including their distribution of size, geography, industry, and share of women in the workforce (Figure B6). Appendix Table A7 statistically explores differences by treatment status. Treated establishments are larger than control establishments, but employ a similar share of women. They are more likely to be located in the North-East region (15% treated versus 11% control) and engage in manufacturing (32% treated versus 28% control). All analyses control for differences in industry and geography across treatment status through 2-digit-industry by year and geography by year fixed effects.

3.2 Differences-in-Differences Design

To measure the causal effect 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.^{29,30} Idiosyncratic errors are captured by ε_{it} and standard errors are clustered by establishment.³¹

The coefficients of interest, β^t , capture 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 treated and the comparison units, absorbed by unit fixed effects α_i . The identifying variation occurs within the same unit, comparing outcomes in any year relative to 2014, and within the same time period, comparing treated and comparison units. 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 for parallel trends in the pre-period.

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 with year-specific indicators δ_t with a single interaction for the post-period, $D_i \times \delta_{t \geq 2015}$. In addition, 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 change as a downstream consequence of changing amenities (e.g., wages and retention) 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 the CUT reform’s effect on amenities, finding disproportionate gains in women’s amenities on paper and in practice. We then explore whether women value these changes to CUT workplaces by studying the reform’s impact on two revealed preference measures of firm value—retention, and job queues. We conclude by evaluating potential tradeoffs from the improvement in female-centric amenities—in men and women’s employment, wages, and in firm profits.

²⁹For industry we use the first two digits of Brazil’s CNAE codes. There are 87 unique industries, including textile production, road transportation, and construction.

³⁰For geography we use either states (27 in total) or microregions, which are neighboring municipalities grouped into 543 units that capture local labor markets.

³¹Clustering by establishment assumes that establishments negotiate with unions that, as of 2012, were affiliated at random with a union central. Results are unchanged when clustering by union.

4.1 Amenities: On Paper and In Practice

Negotiated amenities Table 3 reports the CUT reform’s pooled DiD treatment effect on female and male-centric clauses and Figure 4 reports year-specific effects.³² Pre-reform female amenities evolve in parallel, supporting our identification assumption. Immediately following the reform, female clauses in treated contracts rise sharply in number (intensive margin), incidence (extensive margin), and as a share of all clauses. On the intensive margin, the number of intuitively defined female clauses increases by 0.157 (SE 0.013)—a 17% increase relative to baseline (Panel A). Data-driven female clauses rise by 0.301 (SE 0.021), a 19% increase. These effects represent a substantial improvement, equivalent to moving from the average baseline amenity count at an establishment with a minority female population to one with over 80% women. The effects do not reflect a mere increase in the number for clause types already being provided in contracts, for example, going from 1 to 5 maternity leave clauses. Rather, they represent the inclusion of *new* female-focused clauses, with the sum of unique clause types increasing by 12% over its baseline value (Panel B).

The CUT reform also increases the occurrence of any female-centric clause (extensive margin, Panel C) and these clauses as a share of all clauses in the contract (Panel D). On the extensive margin, the incidence of female-centric clauses increases by 1.7pp (SE 0.003)—a 5% gain over baseline. Using the data-driven classification, this effect is 3.4pp (SE 0.003), representing a 10% increase. As a share of all clauses, intuitive female clauses rise by 0.5pp (SE 0.001), a 10% increase over baseline, and data-driven clauses by 2.1pp (SE 0.001), denoting a 30% increase.

All four types of female clauses rise—leaves, childcare payments, anti-harassment, and flexibility (Column 2-5), with clauses governing leaves and childcare accounting for 76% of this gain. The CUT-driven improvement in amenities is thus likely to differentially impact workers of childbearing age, a fact that we will later exploit to zoom in on labor market outcomes among these workers.

There is 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 decline by small amounts: by 0.1pp (SE 0.003) relative to 46% at baseline for the former, and by 0.3pp point (SE 0.002) relative to 14% for the latter (Column 7). While the number of male-centric clauses increases, this gain is more than overshadowed by the gain in female-centric clauses. Moreover, while the treatment effect on female amenities occurs sharply in 2015, for male amenities it occurs in 2017, suggesting that the male effect is unlikely to be driven by the CUT reform (Figure 4).³³ Overall, the ratio of female-to-male clauses rises by 21%

³²Figure B7 plots the raw path of female-centric clauses in treated and comparison contracts.

³³The increase in male amenities is not robust to clustering by the union, whereas the gain in female amenities

over its baseline value in treated versus comparison contracts (Column 8). The CUT reform therefore increases the female-orientation of contracts, driven by an increase in female-centric clauses.³⁴

Turning to the question of *where* union priorities exert the greatest influence on female amenities, we find the largest gains at establishments where women could not already advocate for themselves either as workers or as union leaders (Table 4, Figure B9). Specifically, at establishments with a small baseline share of women in the workforce (below median, Column 2), in union leadership (Column 3), or without a female union president or vice-president (Column 4).

In terms of mechanisms, our findings could either reflect a change in the composition of union leadership through more female leaders, or a shift in broader union priorities without a direct increase in female leadership. Figure B10 shows a small positive treatment effect on the share of women among union leaders (0.7pp or 3%). However, while these newly elected female leaders may have been instrumental in implementing the CUT’s new priorities, they do not account for the reform’s full effect, as we also find large improvements in female amenities in contracts negotiated with unions without any new female leaders. We interpret this as evidence that the amenity increase stems from a broader shift in union priorities toward women, rather than simply changes enacted by the women themselves.

On a final note, it is worth highlighting that CBA clauses represent equilibrium outcomes resulting from negotiations between unions and employers. As such, our results show employers’ willingness to sign off on female-friendly amenities. Upcoming analyses explore whether this willingness reflects changes on paper not translating into practice, employers adjusting compensation along other dimensions (such as wages), a reallocation of surplus toward workers, or the proposed changes leading to pareto improvements for workers and employers.

Actual amenities To assess whether the change in amenities on paper translates into practice, we draw on the text of female-centric clauses to identify three outcomes that they can directly affect: (i) whether women are managers—corresponding with equal opportunity clauses; (ii) whether women take longer maternity leaves—corresponding with clauses that extend maternity leave; and (iii) if women enjoy job protection post maternity leave—corresponding with job protection clauses.

The reform positively affects outcomes along all three dimensions (Figure 5). The share

is (Table A8).

³⁴These results are robust to reasonable amendments to the data-driven definition of male- and female-centric amenities, the inclusion of more granular industry-geography-year fixed effects, and conditioning on establishment-union pairs with coverage in 2014 (Tables A9, A10, A11).

of women among managers at treated establishments increases by 2% relative to baseline. Women also take longer maternity leaves, with a 14% treatment effect on the share of mothers taking leaves longer than the state mandate of 120 days. Despite these longer leaves, mothers are no less likely to return to their employer following maternity, implying that they enjoy longer periods of job security. Thus, the new union priorities enacted by the CUT reform lead to actual improvements in the workplace for women.

We similarly draw on the text of male-centric clauses to study whether workplaces deteriorate for men. Per the data-driven approach, men value safety. We find no treatment effect on workplace safety as captured by the share of workers taking work-related injury leave. If anything, there is a -3% treatment effect on this outcome. Thus, at least on this dimension, the workplace does not deteriorate for men.

4.2 Revealed preference changes in firm value

Our analysis of improvements in actual amenities is limited to observables in the RAIS data. To more comprehensively understand whether workers actually value these changes to CUT workplaces, we study the reform’s impact on two revealed preference measures of job quality: employee retention and job queues.

Retention Retention serves as a revealed preference measure of an employer’s attractiveness relative to others (Krueger and Summers, 1988). We find a 1.8pp (SE 0.004) increase in retention among incumbent women, a 2.5% improvement over baseline.³⁵ The gender difference in this treatment effect is 0.08pp (SE 0.003), suggesting that incumbent women disproportionately value the reform over its value for incumbent men (Figure 6a). Since we find the largest improvement in amenities related to maternity and childcare, we also zoom in on retention among workers of childbearing age (20-35 years), finding a similar treatment effect (Figure 6b).

However, higher retention need not imply that women value these jobs more if it reflects fewer firings instead of fewer quits. To assess this possibility, we decompose the treatment effect on retention into a component explained by employer-to-employer transitions, likely reflecting quits, versus transitions into unemployment, more likely after a firing. Consistent with a revealed preference story, we find that the treatment effect on retention is explained by fewer voluntary employer-to-employer transitions as opposed to fewer firings into unemployment (Table A12).³⁶

³⁵The two-year baseline retention rate among women is 71%.

³⁶Voluntary transitions among incumbent women (men) decline by 1.1pp (0.8pp).

Since the share of male-centric clauses negligibly falls, men may value CUT employers less. However, we find a 1.0pp *increase* in retention among incumbent men (Table A12), representing a 1.5% increase over baseline. That men quit less provides strong evidence against the hypothesis that men are worse off due to the CUT reform. Thus, although the reform disproportionately improves working conditions for women, it does so without driving men to other jobs.

Job queues Job queues are a second revealed preference measure of value (Holzer et al., 1991). Because we do not directly observe applications, we use workers in the probationary period, i.e., the first 3 months of tenure, as a proxy measure. Since Brazilian labor law permits employers to terminate probationary workers without severance pay, such contracts are commonly used to screen workers.³⁷ We find a 0.6pp increase (SE 0.003) in women’s share among probationary workers (Figure 6c), a 1.7% improvement over baseline. This suggests that women are more likely to queue for jobs at treated establishments.

Although precise, the magnitude of this estimate is small. Three factors likely dampen the estimate of women’s queuing response at CUT establishments. The first (as previously discussed) is our inability to directly observe changes in amenity values using which to scale treatment effects.³⁸ The second is information frictions that may prevent workers from learning of newly instituted amenities at CUT establishments.³⁹ Finally, employers may potentially screen women out at the hiring stage, such that any change in composition among probationary workers is already muted.

In sum, we find that women flock to CUT establishments following the reform. Together their lower separation from, and higher likelihood of queuing for jobs at, CUT establishments translate into a 0.2pp increase in women’s share among employees. Section 5 uses these revealed preference changes in firm value to quantify the CUT reform’s effect on worker welfare.

³⁷For example, 25% of all separations occur between 3 months and 3 months and 1 day.

³⁸Since PageRank values can only be estimated for the super-connected set of firms, it is infeasible to separately estimate pre and post-period values covering a reasonably large sample of firms given only 3 years of data per period.

³⁹As an anecdotal example, an economics professor believed that she was eligible for extended maternity leave because a co-worker at the same institution had obtained such an extension. However, this professor’s location was not covered by the same CBA as her colleague, meaning that she was ineligible for the maternity leave extension.

4.3 Tradeoffs

How are the improvements in female-focused amenities paid for? Table 5 explores three potential explanations.⁴⁰ First, employers may finance amenity improvements by reducing women’s wages, as predicted by compensating differences (Rosen, 1986). Alternatively, if unable to pass the cost of amenities onto workers’ wages, employers may reduce employment or employ relatively more men or inexpensive workers/older women (Summers, 1989). Finally, firms may finance improvements through lower profits.

Wages If amenity improvements operate in a compensating differences world, women’s wages should disproportionately decline. Table 5, Panel A reports the treatment effect on wages and Figure B11 shows parallel pre-trends. Because Brazilian employers cannot cut nominal wages for existing workers without the union’s approval, wage adjustments may only realize for new workers. We therefore separately study the reform’s effect on the mean log wage of established workers, with tenure over 12 months, and new workers, with tenure under 12 months, separately by gender.

There is no treatment effect on the mean log wage of any worker group—established, new, men, or women. All point estimates are negative but very small and precise—the largest decline occurs for new male workers, whose wages fall by 0.6pp (SE 0.003).⁴¹ We rule out negative effects greater than 1.2-1.3pp for new workers, and 0.7-0.8pp for established workers, at a 95% confidence level.⁴² Given the similar point estimates for wage changes among men and women, there is no change in the gender wage gap. Overall, there is little evidence in favor of employers lowering wages to pay for higher female-centric amenities.

There are two important caveats to this finding. First, the average worker may not adequately represent workers whose wages are actually influenced by unions. For a more direct measure of union-negotiated wage changes, we extract the percentage wage adjustments negotiated in collective bargaining agreements. The treatment effect on these wage adjustments is 0.032pp (SE 0.021), allowing us to rule out a more than 0.009pp fall in wages with a high degree of confidence (95%). Second, employers may respond by changing the composition of their workforce, such that zero wage effects mask effective wage changes for new workers. However, we find an incredibly precise null treatment effect on the wages of incumbent workers, whose composition is unchanged (Table A12).

⁴⁰Using the establishment sample.

⁴¹This result is not robust to including industry-geography-year fixed effects.

⁴²By way of reference, Lagos (2021) finds that workers value leave clauses, many of which are classified as female-centric, at 8.4% of their wage. That paper pools men and women together.

Employment If employers cannot pass the cost of amenity improvements onto workers' wages, they may lower employment. Table 5, Panel B reports the treatment effect on employment and Figure B11 shows parallel pre-trends. Column 1 reports the effect on overall employment and Column 4 on hiring. We find no statistically significant impact on employment or new hiring among treated employers, and can rule out negative effects larger than 0.15pp with a high degree of confidence (95%). We also find no decline in the employment or hiring of female workers; if anything, as previously discussed, women's share among all workers increases by 0.2pp and among probationary workers by 0.6pp.⁴³

A second dimension of adjustment is worker composition—employers may hire more skilled or older workers. Table A13 provides evidence against this hypothesis. There is no change in the proportion of female workers poached from other employers (a measure of positive selection). Moreover, there is no treatment effect on the mean age, tenure, contracted hours, or schooling of female workers.

In sum, we find no evidence that employers hire fewer women, fewer workers, or different or more productive workers as a result of the CUT reform. Of course, we cannot rule out productivity gains among female workers *as a result of* the change in workplace environment. Indeed, this is a candidate explanation for our finding of no wage or employment tradeoffs due to the CUT reform.

Profits If workers do not finance the amenity improvement through lower wages or employment, perhaps firms finance it through lower profits. We provide empirical evidence and theoretical reasons against this explanation.

Table 5, Panel C shows no treatment effect on firm profits, measured in two different ways. First, we find no treatment effect on establishment exit. Exit is a non-trivial margin of adjustment in Brazil, with 8.7% of control group establishments exiting between 2014 and 2017. Second, we estimate a statistically insignificant 0.70pp (SE 1.17) treatment effect on profits among the sample of establishments that is observed in Orbis data during our study period. For this restricted sample, we rule out a higher than 1.59pp decline in profits with a high degree of confidence (95%).

Theoretically, profits could only fall if CUT unions were able to bargain away a larger share of surplus from employers. However, there is little reason to think that the CUT reform increased unions' bargaining power; rather, it merely shifted union priorities to favor women. If anything, the position of CUT-affiliated unions grew increasingly precarious around this time, following the 2015 impeachment of President Dilma Rousseff of the left-

⁴³There is also a small positive effect on the share of women among separators due to more women being hired and working at the firm. However, on net, the share of female workers increases.

wing Workers' Party with which the CUT has close ties. Moreover, while increasing union bargaining power generally predicts a change in employment—either moving right along a firm's upward-sloping labor supply curve, or left along its labor demand curve—we find a precisely estimated zero.

In sum, we find no evidence that profits decline to pay for the female-focused improvement in amenities.

4.4 Robustness

Brazil experienced a recession between 2014 and 2016. Our findings may be driven by the recession as opposed to a shift in union priorities if CUT unions either represent systematically different industries that are differently impacted by the recession, or if these unions differently respond to the recession. Several findings point against the differential impact of, or response to the recession as driving our findings. First, our results reflect an increase in female amenities in CUT contracts as opposed to a potentially-recession-induced-decline in amenities in non-CUT contracts (Figure B8). Second, there is little reason to expect the recession to have increased the CUT's demand for female-centric amenities such as maternity leave or childcare payments (as opposed to clauses that shield workers' wages, which may arguably constitute a more natural demand during a recession). Third, we find heterogeneous treatment effects, with the largest amenity gains occurring at establishments with a small baseline share of women; this heterogeneity counters the idea that the CUT in general responded differently to the recession. Finally, all specifications include 2-digit-industry and location-specific time varying shocks through industry by year and microregion by year fixed effects.

4.5 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. Women valued these changes, becoming less likely to separate from and more likely to queue for jobs at CUT establishments. Perhaps surprisingly, we find no evidence that these gains in female-focused amenities come at the expense of women or men's wages and employment, or of firm profits. While amenities for men may have fallen (in some unobserved way), men do not exit more.

Together our findings demonstrate that shifting union priorities can reduce the gender compensation gap. Just as in politics, where leaders' priorities determine policy design (Chattopadhyay and Duflo (2004); Pande and Ford (2012)), we show that unions' priorities

determine workplace design. We consider a broader definition of compensation than wages alone, including also amenities such as family allowances, leaves, and flexibility, and show that these are key levers through which unions influence inequality.

There are at least three models that could explain our results. In one model, men lose rents due to the reform—which are not observed in amenity or wage changes—but they are not marginal to this loss since these rents are not provided elsewhere. A second model is one in which the union was inefficiently aggregating workers’ preferences. Shifting union priorities caused it to focus on previously ignored female amenities that could be provided at net zero cost to employers. Finally, providing amenities may have increased firm profits and the total size of rents split between unions and employers. Behavioral unions and firms may have been leaving these gains on the table until the reform spurred a shift in focus. The last two explanations represent pareto improvements. Since male retention slightly improves, and male wages and amenities do not decline, our findings are most in line with one among the last two explanations.

5 Quantifying the Welfare Effect of the CUT Reform

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

Approach and Intuition We quantify the CUT reform’s effect on worker welfare through a revealed preference approach that (i) relies on a few sufficient statistics that are easily computable in the data; and, thus, (ii) takes no stance on the precise functional form linking amenities to worker 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 possess CES preferences over employers, as is common in the consumer setting (Feenstra, 1994; Atkin et al., 2015). As shown in 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 based on where they obtain the highest utility. This is a common way of modeling the labor market (in 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.

Then, just as gains to consumer welfare from improving product varieties can be measured through changes to the price index—i.e., the change in cost of purchasing one util worth of

utility—the gains to worker welfare from improving workplace amenities can be measured through changes to the wage index—i.e., how much more (or less) the representative worker earns to work one disutility-weighted hour.

Under CES preferences, only four sufficient statistics quantify the change in worker welfare, i.e., measure the change in the wage index. First, welfare increases with the share of total labor income found at treated establishments, which captures workers flocking to these employers after they improve amenities. Second, the same change in labor income at treated establishments corresponds with a higher increase in welfare if workers are less elastic to begin with, since it takes a larger improvement in amenities to draw them away. Third, welfare is higher if workers are drawn away from non-CUT firms with initially low value, capturing a bigger upgrade in employer quality across regimes. Finally, welfare increases with wages at non-CUT establishments, potentially capturing the pro-competitive spillover effects of the reform.

Model In each period, a representative household with CES preferences over employers is willing to work a fixed number of (dis)utility-weighted hours. It chooses labor supply to each firm to maximize total income, subject to this 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 denotes the set of firms operating at time t , n_j is the number of hours supplied to firm j , w_j is the wage at j , η is the elasticity of substitution across firms, and b_j represents the “taste-shifter” for firm j . b_j captures all non-wage attributes that commonly affect each worker’s utility at j . Worse amenities increase this disutility b_j . We assume a utility-posting world without job rationing, where a firm accepts any worker who wishes to work there. For simplicity, since worker welfare only depends on firms’ final wage and amenity offers, regardless of how firms arrive at them, we do not model the firm side.

The wage index measures how much the representative worker is paid to work a disutility-weighted hour, and serves as a measure of welfare:

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

Any change in the wage index across two periods captures changes to worker welfare, mea-

sured by the ratio:

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

The CUT reform changes amenities, or taste shifters b_{jt} , at treated establishments. The key challenge in estimating welfare changes is that these $\{b_{jt}\}_{j \in \mathcal{J}_t}$ are unobserved. However, assuming CES preferences allows us to overcome this challenge. Under CES, any welfare change depends only on the *observed* pre- and post-reform wages and employment at CUT and non-CUT employers.⁴⁴ Formally:

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

where λ_t is the share of total labor income in t at non-CUT firms, \overline{S}_t^* is a geometric average of the share of labor income at each non-CUT firm in t , and \overline{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 relative to before the reform. This ratio captures welfare changes through a revealed preference logic: workers substitute toward CUT firms once their amenities improve, 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 toward amenity-improving CUT-firms implies a larger welfare increase because it takes a bigger improvement in amenities to draw workers away.

The term $\left(\frac{\overline{S}_t^*}{\overline{S}_{t-1}^*} \right)^{-\frac{1}{1+\eta}}$ captures the heterogeneity in labor income at non-CUT firms: welfare increases by more if CUT firms draw workers away from less valued non-CUT firms, thereby increasing dispersion in and lowering the geometric mean of their wage bill share. As in the “variety-adjustment” term, the implied effects are larger as workers become more inelastic. The final term $\left(\frac{\overline{w}_t^*}{\overline{w}_{t-1}^*} \right)$ 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 We separately estimate Equation (5) for men and women. Our estimates employ the establishment sample from Section 4.3. Years 2012-2014 comprise the pre-reform period ($t - 1$) and 2015-2017 the post-reform period (t). We calibrate an estimate of the

⁴⁴Under CES, the relative (dis)utility of working at an employer is captured by its expenditure share, which depends exclusively on prices and quantities.

cross-firm elasticity of substitution (η) from Felix (2022), but assess robustness to other reasonable values.

We estimate the log change in \bar{w}^* and in \bar{S}^* using average changes across non-CUT establishments between $t - 1$ and t , estimated via the following regression:

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

where y_{jt} is either the average log earning at establishment j ($\log w_{jt}$) or the log of the share of labor income among non-CUT establishments at that establishment ($\log s_{jt}$). The specification includes establishment fixed effects μ_j . The coefficient of interest, β , 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 us to map the market-level change in the share of labor income at CUT establishments (the desired object) to changes in quantities that are estimable through establishment-level regressions as in Equation (6). We refer the reader to Appendix E for details.

Results Table 6 reports results. Women’s welfare increases by 0.059 log points (or 6.1%), consistent with our reduced form results that women are more likely to remain at, and comprise a larger share of new workers among, CUT establishments.⁴⁵ Worker moves following the reform account for over half of the increase in welfare. Women become more likely to work at CUT establishments, accounting for 15% of the welfare gain (a 1.8% rise in the share of CUT wage bill). In addition, the dispersion in the labor income across non-CUT firms rises (i.e., S^* falls), accounting for 48% of the increase in welfare.

The remaining 37% of the welfare gain is accounted for by higher wages among non-CUT employers. To the extent that these wage increases reflect pro-competitive responses to the CUT reform, any change in welfare from them can also be attributed to the reform. We recognize, however, that the increase in real wages at non-CUT employers following 2015 could be driven by a host of factors that are unrelated to the CUT reform. We therefore only view the change in welfare due to worker moves across firms—amounting to a 3.8% increase—as the credible estimate of the reform’s welfare impact for women workers.⁴⁶

For men, welfare is slightly higher (1.3%), but remains essentially unchanged if one only

⁴⁵As predicted by the model, workers’ elasticity of substitution across employers amplifies (or dampens) the welfare effect due to the shifts in employment across firms induced by the reform. For other reasonable values of η in the literature, ranging from 0.1 (Staiger et al., 2010) to 10.9 (Berger et al., 2022), women’s welfare increases by between 2.8% and 9.5%.

⁴⁶Table A14 computes the change in welfare separately for workers of child-bearing ages (i.e., between 20 and 35 years old), and finds qualitatively very similar results to those unrestricted by age.

considers the component due to worker moves across firms (0.2%). Thus, the CUT reform improves women’s welfare without reducing men’s welfare.

6 Conclusion

This study finds that one reason that workplaces do not provide job features valued by women is that decision-makers do not prioritize women’s preferences. Studying a top-down change in Brazil that led Latin America’s largest trade union federation, the *Central Única dos Trabalhadores*, to adopt a bargaining plan more attentive toward women’s needs, we find a sharp increase in female-centric amenities, without corresponding tradeoffs in women or men’s wages and employment, or in firm profits. The reform increases female-centric amenities on paper, such as those governing maternity leaves, job protection, childcare, and flexibility. These changes on paper translate into practice, with women taking longer maternity leaves and becoming more represented among managers. We find that women value these changes; they are less likely to separate from and more likely to queue for jobs at CUT establishments. Although the reform may have reduced male amenities, men do not exit more. Finally, we find no evidence that firm profits fall.

In sum, we provide causal evidence that union priorities importantly shape compensation, and, consequently, within-firm inequality. While gender gaps in virtually all labor market outcomes have narrowed at a fast pace in the last century, more recently reducing inequality has proven harder (Goldin, 2014; Blau and Kahn, 2017, 2006). Policies increasing women’s representation in the workplace, for example via quotas on firm boards, have had null effects (Bertrand et al., 2018; Maida and Weber, 2020). By contrast, we find an important role for representing women’s interests in collective bargaining.

In our setting, prioritizing women appears to usher in more efficient compensation for workers. One possible explanation for these findings is that the union was originally striking an inefficient bargain for workers. An alternative possibility is that the reform increased firm profits and the total size of rents split between unions and employers. Turnover is typically costly to the firm, and we find lower separations among women. Happier workers may also be more productive. Finally, the reform may have simply shifted worker rents from men to women, with no increase in male quits because men could not obtain similar rents elsewhere. This last explanation, while possible, is not supported by the evidence since we find a small increase in male retention and no observed changes in men’s wages or amenities.

Our findings raise several new questions. First, given that leaders’ priorities can influence compensation and inequality, how do these priorities emerge? A historical literature emphasizes the inherently political nature of labor unions, and argues that their objectives

are ultimately shaped by their internal organization (Farber, 1986; Ross, 1950). In light of our findings, this hypothesis is especially fruitful to revisit empirically. Second, if leaders influence workplace conditions, might they also influence investments that affect worker productivity? Studying how productivity endogenously evolves as a consequence of leadership decisions is an exciting area for future research.

References

- Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High Wage Workers and High Wage Firms. *Econometrica*, 67(2):251–333.
- 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., Herkenhoff, K., and Mongey, S. (2022). Labor market power. *American Economic Review*, 112(4):1147–93.
- 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. (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.
- Bolotnyy, V. and Emanuel, N. (2022). Why do women earn less than men? evidence from bus and train operators. *Journal of Labor Economics*, 40(2):283–323.
- Boudreau, L. (2023). Multinational enforcement of labor law: Experimental evidence from bangladesh’s apparel sector. *Mimeo*.
- 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.
- 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.
- 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 No. 28522*.
- 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 No. 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.
- Mas, A. and Pallais, A. (2017). Valuing Alternative Work Arrangements. *American Economic Review*, 107(12):3722–59.
- 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.
- Ross, A. M. (1950). *Trade Union Wage Policy*. Univ of California Press.
- Skorge, Ø. S. and Rasmussen, M. B. (2022). Volte-face on the welfare state: Social partners, knowledge economies, and the expansion of work-family policies. *Politics & Society*, 50(2):222–254.
- 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.
- 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.

Tables and Figures

Table 1: 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 | 80,131 | 18,103 | 62,028 |
| 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 | 198 | 127 |
| 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 | 0.64 | 0.63 | 0.64 |
| <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.

Table 2: Female- and Male-Centric Amenities

| Intuitive definition | Data-driven definition | | Rank |
|----------------------------|------------------------------------|--|------|
| | Female clauses | Top 20 female 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.

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 unique clause types)</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. The unit of observation is a union-employer pair. Panel A reports effects on the total number of clauses, an intensive margin measure of amenities. Panel B reports effects on the sum of unique clause types in the corresponding categories exist in a contract, capturing changes to the *space* of female (male) clauses (as opposed to their number). For example, two anti-harassment clauses will raise the outcome value by two in Panel A of Column 6 but by one in Panel B. Panel C reports effects on a cumulative indicator for whether any clause of the corresponding type exists in a contract as an extensive margin measure of amenities. Panel D uses the share of clauses among all clauses in a contract. 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.

Table 4: Heterogeneity by Baseline 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 of 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.

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

| <i>Panel A: 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 B: 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 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 |

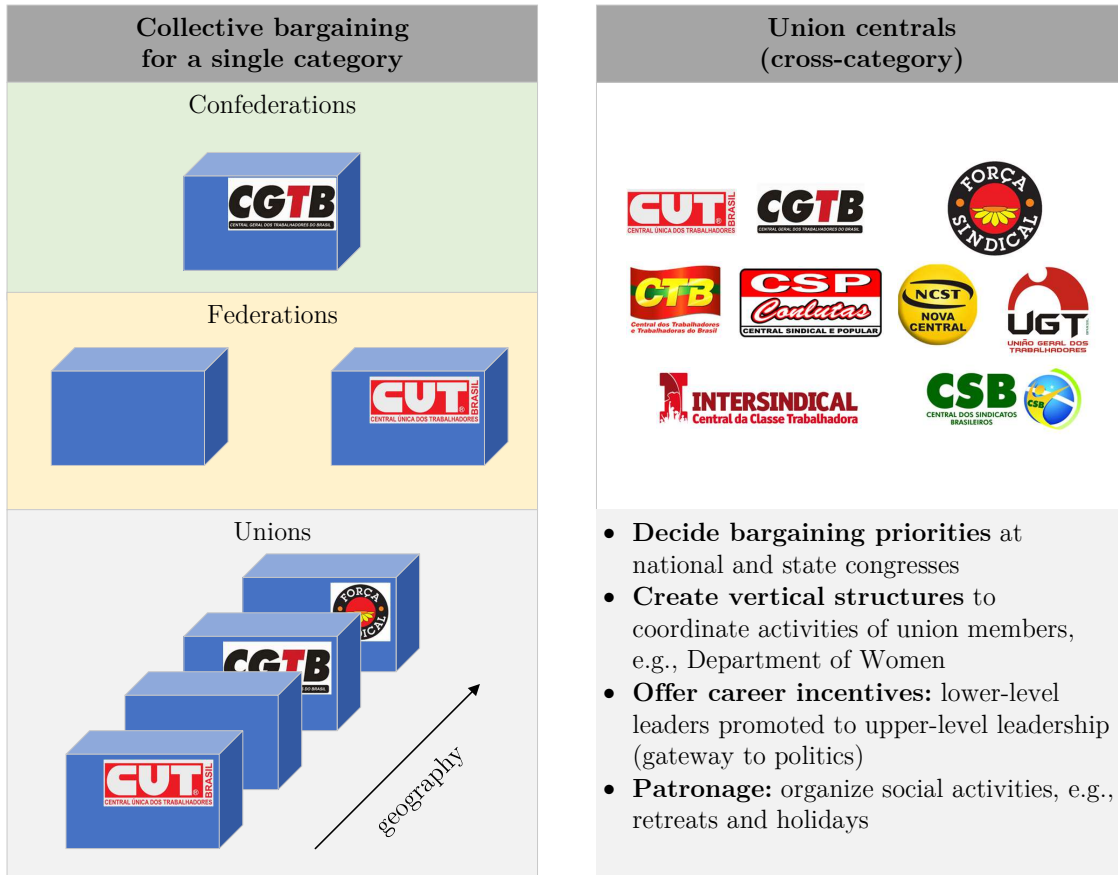
Notes: Table reports the coefficients for the establishment-level DID regression from Equation (3), comparing treated to comparison establishments on wage, employment, and profit outcomes. An establishment is treated if the union with which it negotiates is affiliated to CUT in 2012. Each regression includes establishment fixed effects, industry-year fixed effects, and microregion-year fixed effects. Panel A uses 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 B 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 C studies three imperfect measures of firm profits. Standard errors are clustered by establishment and reported in parentheses.

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.


Figure 1: Workers' Bargaining 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.

Figure 2: The 2015 CUT Reform

(a) Gender parity in national leadership




(b) Impact on local union boards



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.

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).

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

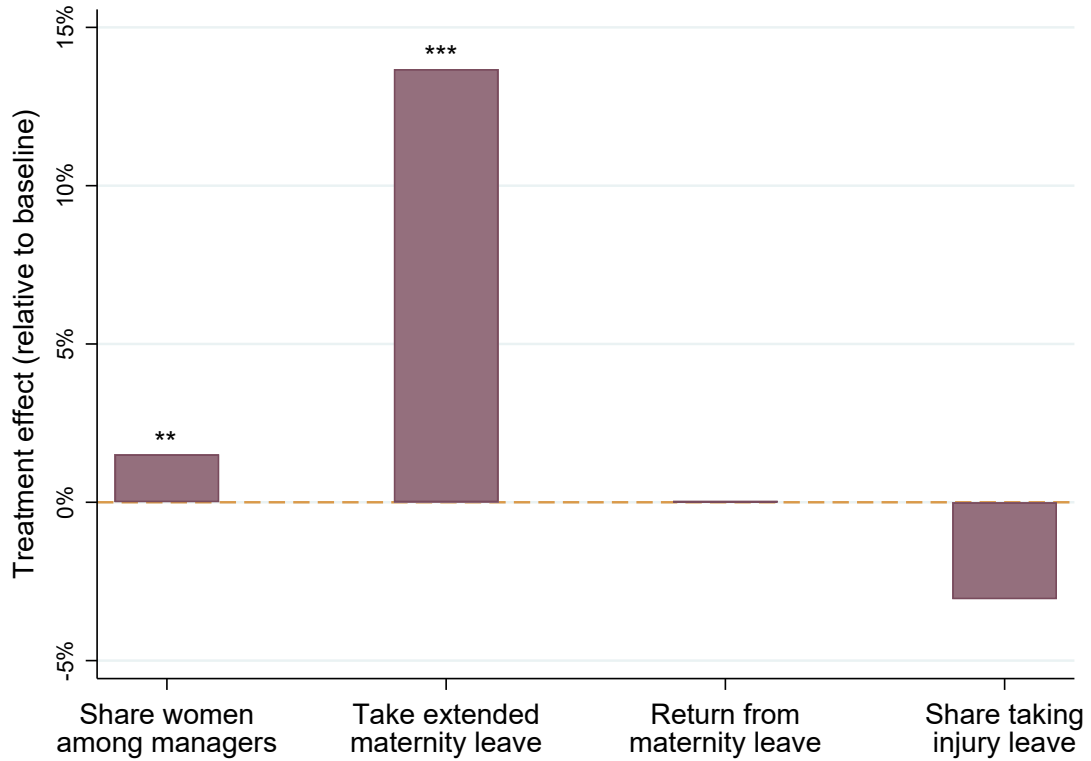


(d) Male clauses: as a share of clauses



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.


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.

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.

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: Examples of Female-Centric Amenities

| | |
|----------------------------|---|
| Childcare assistance | The company will reimburse all female employees, the monthly amount of R\$ 110, as a “day care allowance”, per child up to 6 years old. This benefit applies to any employee with custody of the child(ren). |
| Absences | The employee will receive full pay for absences upon proof of the following cases: a) bereavement (5 consecutive days); b) hospitalization of direct family or legal dependents; c) medical and dental consultations; d) marriage (5 working days) |
| Adoption leave | The employee who adopts or obtains legal custody for adoption will be granted maternity leave as follows: a) 120 days for children up to 1 year old; b) 60 days, for children from 1 to 4 years old; c) 30 days for children from 4 to 8 years old. |
| Other: holidays and leaves | The start of vacations cannot coincide with Saturdays, Sundays, holidays, or days already compensated. Vacations will start on the first working day of the week, communicated to the union within 10 working days by the company. |
| Seniority pay | The company will pay the employee who completes 5 uninterrupted years of work an additional 5% per length of service payable monthly, calculated on the monthly fixed base salary. |

Notes: Table lists examples of CBA clauses from the top 5 clause types selected as “female-centric” or “male-centric” based on our data-driven approach—refer to Section 2.2 for details on the data-driven approach. The clauses were selected based on the number of unique tokens appearing in the clause that are within the top 20 TF-IDF tokens of each specific clause type.

Table A3: Examples of Male-Centric Amenities

| | |
|--|--|
| On-call pay | The company will pay an additional 35% of the normal hours to employees, when scheduled to be on-call. This additional pay will not apply when the on-call becomes a service actually provided, in which case overtime will be due. |
| Life insurance | The company will maintain group life insurance, guaranteeing a single and total indemnity of at least R\$ 10,000 in the event of death or permanent disability of the employee resulting from an accident at work. |
| Strike procedures | The union assumes formal commitment not to promote or encourage stoppages, except in cases of non-compliance with clauses of this agreement or current laws, and even so, only after communicating the transgressions in writing to the employers. |
| Other: protections for injured workers | The company will communicate to Social Security, and subsequently to the union, injuries incurred by employees at the company or while commuting to/from work. |
| Profit sharing | The company will maintain a Profit Sharing Program with the amount made available for payment may be up to 1 nominal salary per employee. The payment period after the calculation of the results will be the month of February. |

Notes: Table lists examples of CBA clauses from the top 5 clause types selected as “female-centric” or “male-centric” based on our data-driven approach—refer to Section 2.2 for details on the data-driven approach. The clauses were selected based on the number of unique tokens appearing in the clause that are within the top 20 TF-IDF tokens of each specific clause type.

Table A4: 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 A5: 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 A6: Establishment Descriptives—RAIS vs. Analysis Samples

| | All RAIS (1) | Amenities sample (2) | Difference p-value (3) | RAIS: employ men and women (4) | Establishment sample (5) | Difference p-value (6) |
|--|-----------------|----------------------------|------------------------------|--------------------------------------|--------------------------------|------------------------------|
| <i>Employment and firm characteristics</i> | | | | | | |
| Size | 16.19 | 143.11 | 0.00 | 31.87 | 150.22 | 0.00 |
| Share women | 0.45 | 0.38 | 0.00 | 0.45 | 0.40 | 0.00 |
| Employs both men and women | 0.46 | 0.82 | 0.00 | 1.00 | 1.00 | 1.00 |
| Single person firm | 0.27 | 0.04 | 0.00 | 0.00 | 0.00 | 1.00 |
| Single establishment firm | 0.77 | 0.65 | 0.00 | 0.77 | 0.63 | 0.00 |
| <i>Sector</i> | | | | | | |
| Agriculture & extraction | 0.09 | 0.04 | 0.00 | 0.05 | 0.03 | 0.00 |
| Manufacturing | 0.09 | 0.28 | 0.00 | 0.11 | 0.30 | 0.00 |
| Construction & utilities | 0.05 | 0.06 | 0.00 | 0.04 | 0.05 | 0.00 |
| Commerce | 0.39 | 0.23 | 0.00 | 0.41 | 0.24 | 0.00 |
| Services | 0.38 | 0.39 | 0.00 | 0.38 | 0.39 | 0.00 |
| <i>Region</i> | | | | | | |
| North | 0.04 | 0.05 | 0.00 | 0.05 | 0.05 | 0.00 |
| Northeast | 0.16 | 0.12 | 0.00 | 0.16 | 0.12 | 0.00 |
| Central | 0.10 | 0.07 | 0.00 | 0.10 | 0.08 | 0.00 |
| South | 0.21 | 0.21 | 0.00 | 0.20 | 0.21 | 0.00 |
| Southeast | 0.49 | 0.56 | 0.00 | 0.49 | 0.54 | 0.00 |
| N establishments | 3,798,207 | 80,131 | | 1,739,255 | 61,752 | |
| N workers | 61,492,768 | 11,467,760 | | 48,564,436 | 9,276,475 | |
| % workforce | 100% | 19% | | 79% | 15% | |

Notes: Table compares descriptive statistics of establishments in Brazil’s formal sector (Column 1) and our analysis samples, i.e., the amenity sample (Column 2) and the establishment sample (Column 5). The p-values of the differences between these samples are reported in Columns 3 and 6. The bottom of the table includes the number of unique establishments and workers in each sample, as well as the percentage of the formal workforce present in the corresponding sample.

Table A7: Treated and Control Establishments Descriptives

| | Amenities sample | | | Establishment sample | | |
|--|------------------|----------------|----------------------|----------------------|----------------|----------------------|
| | Treatment (1) | Control (2) | Diff. p-value (3) | Treatment (4) | Control (5) | Diff. p-value (6) |
| <i>Employment and firm characteristics</i> | | | | | | |
| Size | 198.21 | 127.03 | 0.00 | 200.37 | 135.95 | 0.00 |
| Share women | 0.36 | 0.38 | 0.00 | 0.38 | 0.40 | 0.00 |
| Employs both men and women | 0.83 | 0.82 | 0.00 | 1.00 | 1.00 | 1.00 |
| Single person firm | 0.03 | 0.04 | 0.00 | 0.00 | 0.00 | 1.00 |
| Single establishment firm | 0.66 | 0.65 | 0.11 | 0.64 | 0.63 | 0.06 |
| <i>Sector</i> | | | | | | |
| Agriculture & extraction | 0.03 | 0.04 | 0.00 | 0.02 | 0.03 | 0.00 |
| Manufacturing | 0.32 | 0.27 | 0.00 | 0.33 | 0.29 | 0.00 |
| Construction & utilities | 0.08 | 0.06 | 0.00 | 0.06 | 0.04 | 0.00 |
| Commerce | 0.21 | 0.24 | 0.00 | 0.19 | 0.25 | 0.00 |
| Services | 0.37 | 0.39 | 0.00 | 0.39 | 0.38 | 0.04 |
| <i>Region</i> | | | | | | |
| North | 0.04 | 0.05 | 0.00 | 0.05 | 0.06 | 0.00 |
| Northeast | 0.15 | 0.11 | 0.00 | 0.16 | 0.11 | 0.00 |
| Central | 0.09 | 0.06 | 0.00 | 0.11 | 0.07 | 0.00 |
| South | 0.22 | 0.20 | 0.00 | 0.22 | 0.20 | 0.00 |
| Southeast | 0.50 | 0.58 | 0.00 | 0.46 | 0.56 | 0.00 |
| N establishments | 18,103 | 62,028 | | 13,677 | 48,075 | |
| N workers | 3,588,153 | 7,879,607 | | 2,740,517 | 6,535,958 | |

Notes: Table compares descriptive statistics of establishments between the treated (Columns 1 and 4) and comparison groups (Columns 2 and 5) in our analysis samples, i.e., the amenity sample and the establishment sample. The p-values of the differences between the treated and comparison groups are reported in Columns 3 and 6. The bottom of the table includes the number of unique establishments and workers in each group.

Table A8: 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 (unique clause types)</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 unique clause types, 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 A9: 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> | | | | | | | | |
| $D_i \times \delta_{year \geq 2015}$ | 0.096*** (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 (unique clause types)</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 unique clause types, 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 A10: Effect of CUT Reform on Female Amenities

| | Female-Centric Clauses: Intensive Margin | | | | | |
|--|--|---------------------|---------------------|---------------------|---------------------|---------------------|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| <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) | 0.297*** (0.019) | 0.096*** (0.015) |
| Mean outcome | 0.95 | 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) | 0.417*** (0.030) | 0.121*** (0.023) |
| Mean outcome | 1.58 | 2.05 | 1.17 | 1.58 | 1.58 | 2.71 |
| Data-driven clauses | baseline | any | all | baseline | baseline | baseline |
| Geography-year FEs | state | state | state | microregion | micro×ind | state |
| CBA coverage in 2014 | no | no | no | no | no | yes |
| Observations | 600,960 | 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 A4 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 microregion instead of state. Column 5 uses a microregion-industry time-varying fixed effect. Column 6 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 A11: Effect of CUT Reform on Female Amenities

| | Female-Centric Clauses: As a Share of All Clauses | | | | | |
|--|---|---------------------|---------------------|---------------------|---------------------|---------------------|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| <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) | 0.005*** (0.001) |
| Mean outcome | 0.05 | 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) | 0.011*** (0.001) | 0.004*** (0.001) |
| Mean outcome | 0.07 | 0.08 | 0.04 | 0.07 | 0.07 | 0.11 |
| Data-driven clauses | baseline | any | all | baseline | baseline | baseline |
| Geography-year FEs | state | state | state | microregion | micro×ind | state |
| CBA coverage in 2014 | no | no | no | no | no | yes |
| Observations | 600,960 | 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 A4 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 uses a microregion-industry time-varying fixed effect. Column 6 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 A12: 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 A13: Impact of CUT Reform on Worker Composition (Female)

| | 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 A14: 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 |
| | $\ln(\bar{w}_t^*) - \ln(\bar{w}_{t-1}^*)$ | 0.015 | 0.022 | -0.001 |
| | $\ln(\bar{S}_t^*) - \ln(\bar{S}_{t-1}^*)$ | -0.046 | -0.058 | 0.013 |
| 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



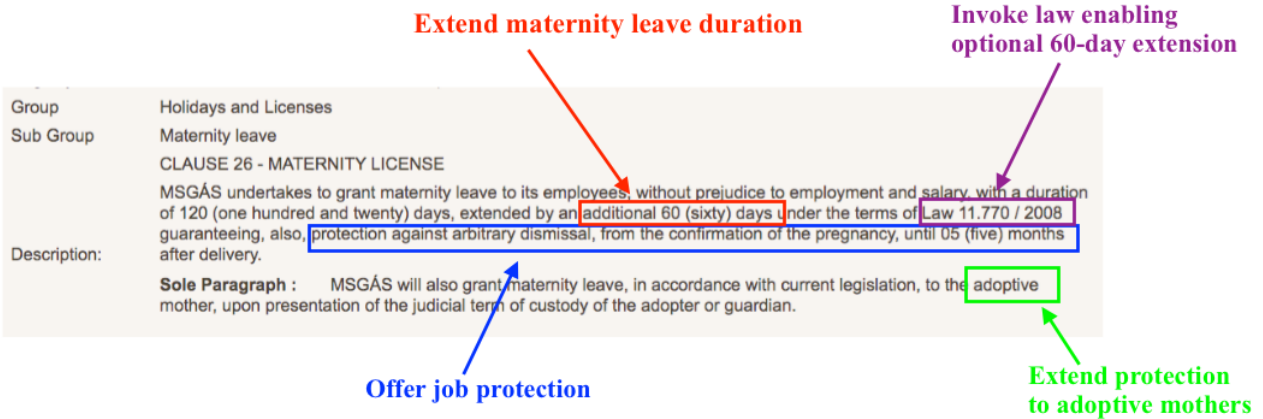
Notes: Figure plots the annual share of women on each union central’s national executive committee (*Inter-sindical* is dropped due to missing information on its board). The line for CUT is the same as in Figure 2a, while the unweighted average of all other union centrals make up the other line reported 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 (9 in total) and the “Maternity Leave” clause types (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 clause type classification of the different clauses, ignoring the variation in the text that may exist within each individual clause belonging to a specific type.


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



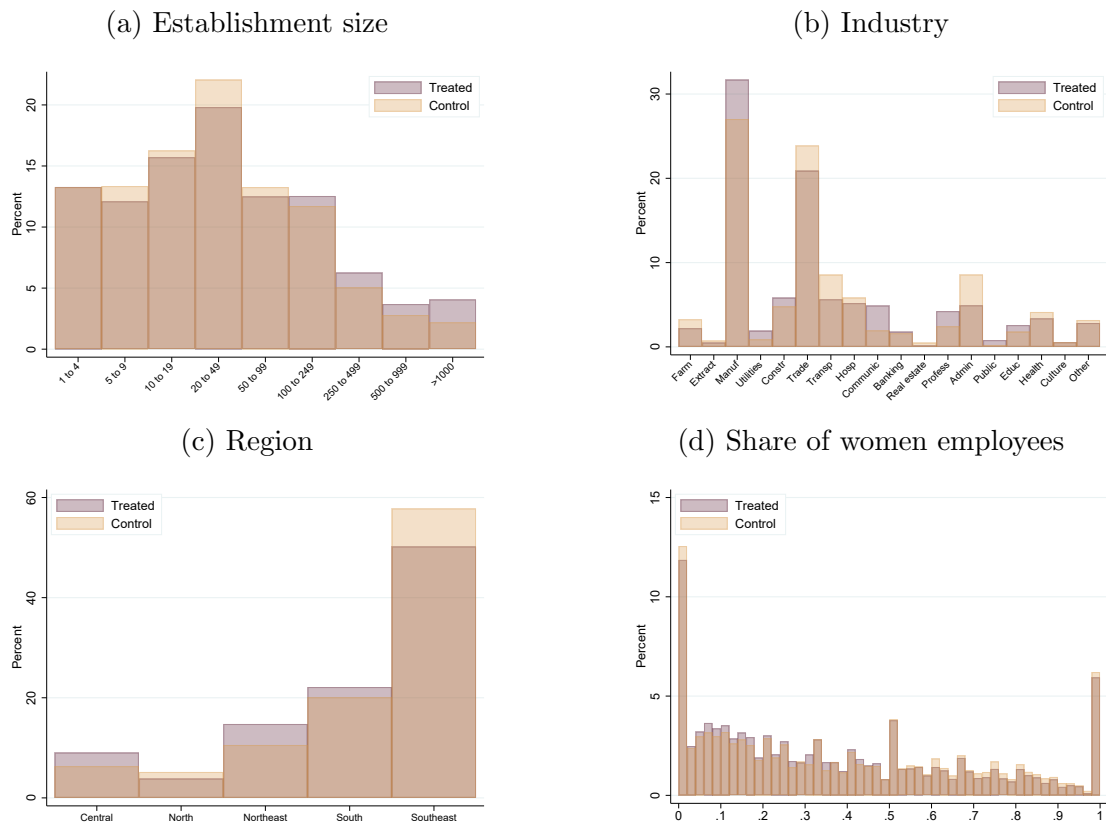
Notes: Figures depict binned scatterplots of the establishment-level gender gaps in PageRank values by the average female-centric clauses from sectoral CBAs applying to the establishment. Figure B4a uses the intuitive definition of female-centric amenities, while Figure B4b uses the data-driven approach. 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



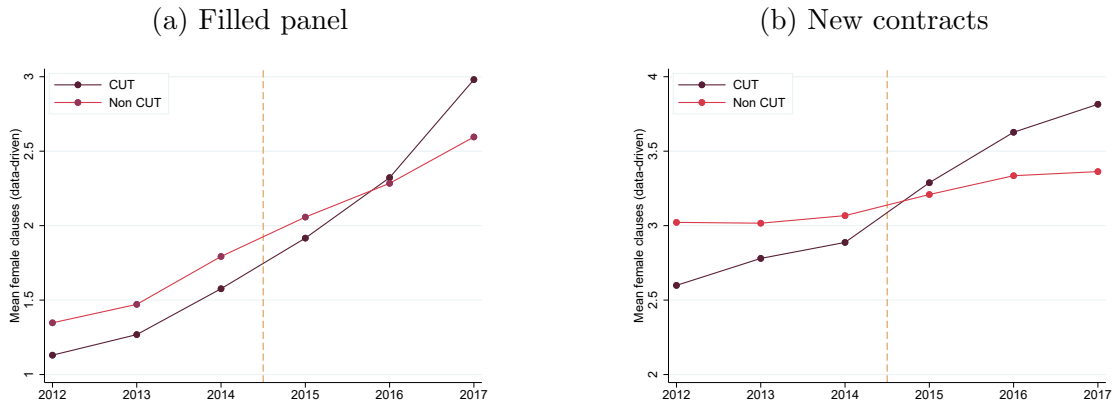
Notes: Figure plots changes in the probability of being affiliated to CUT between 2012 and 2016 separately for unions having either a high or a low share of women among the workers they represent (above or below the mean, i.e., 33% women). Coefficients represent the change with respect to 2012, in which 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



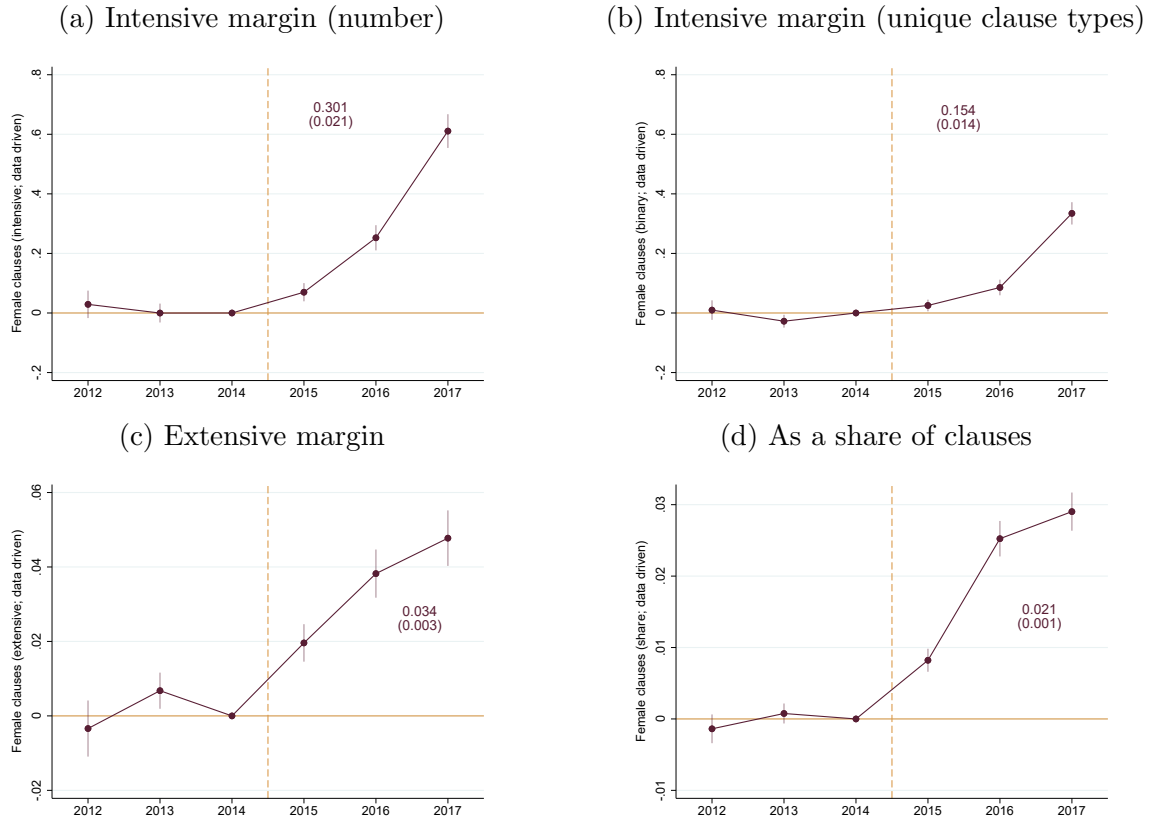
Notes: Figures show the treated and control establishments distributions of size, industry, regional location, and female share of employment at baseline. The establishments come from the starting sample detailed in Table 1.

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



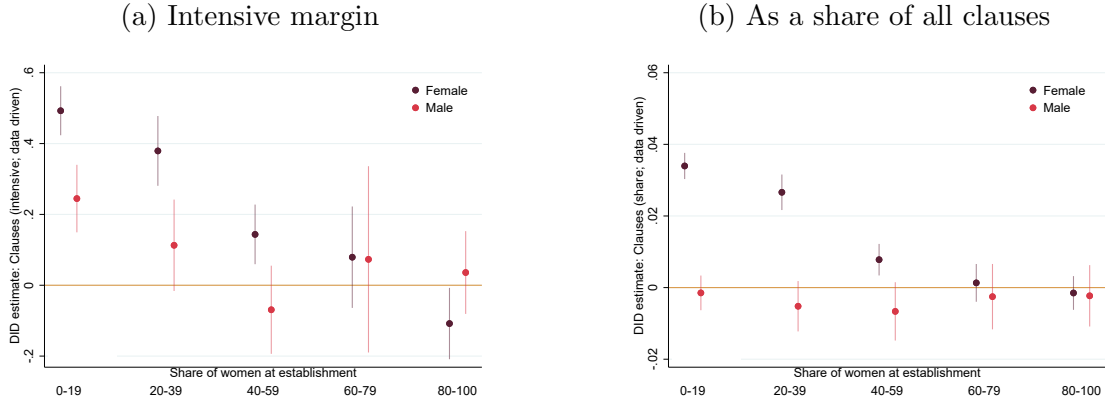
Notes: Figures plot the raw average number of female-centric clauses for treated (CUT) and control (Non CUT) establishment-union pairs over the years. Female-centric clauses are based on the data-driven classification. Figure B7a plots the average number of female-centric clauses for the filled panel, 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



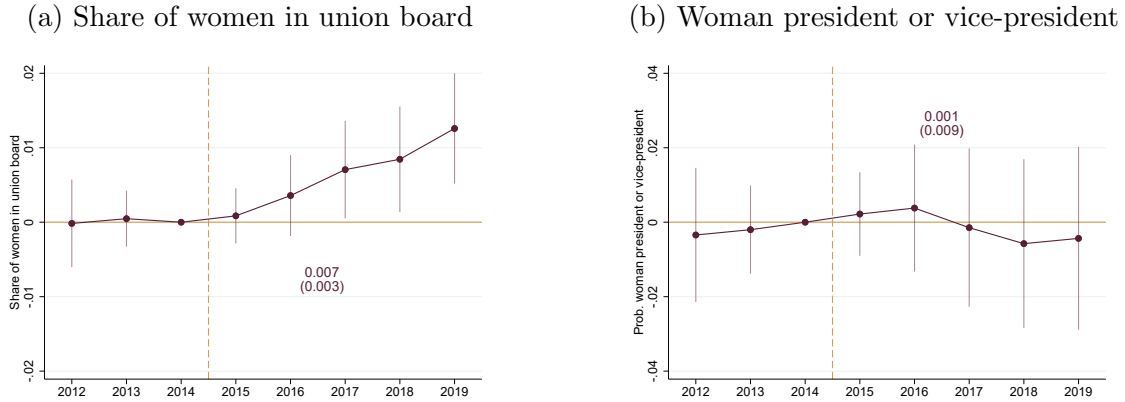
Notes: Figures show estimates of the δ_t coefficients for $t \in [2012, 2017]$ (with 2014 omitted) from the DID specification in Equation (3) on all margins considered for female-centric clauses, defined using the data-driven method. Confidence intervals at a 95% level are reported. Standard errors are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



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.⁴⁷

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).⁴⁸ We additionally drop contracts signed by multiple unions with different union central affiliations: this is fewer than 0.33% of CBAs.⁴⁹

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

⁴⁷*Sistema Mediador* classifies clauses into 137 categories, e.g., maternity assistance, overtime pay, life insurance, procedures in relation to strikes and strikers, etc.

⁴⁸Unions 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.

⁴⁹Of 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.⁵⁰ 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 1.

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.⁵¹ 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

⁵⁰We 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.

⁵¹All 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.⁵² 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

⁵²Firm-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).⁵³ 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.

⁵³Developed 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.⁵⁴ 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 2.

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

⁵⁴Section 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.⁵⁵

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 A4 and A5 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

⁵⁵Specifically, 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.⁵⁶ 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 $\{\varepsilon_{i,j}\}$.⁵⁷ 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

⁵⁶The 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.

⁵⁷This 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.⁵⁸

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 .

⁵⁸In 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.