

Monopsony and Gender

Garima Sharma*

January 22, 2023

[Updated frequently. Click here for the latest version.](#)

Abstract

I investigate the role of labor market power in driving the gender wage gap in Brazil. Exploiting firm-level shocks induced by the end of the Multi-Fiber Arrangement, I show that women are substantially less likely than men to separate from their employer following a wage cut. The ensuing gender difference in monopsony power would explain 50% of the gender wage gap (an 18pp difference). To probe the source of higher monopsony power over women, I develop and estimate a discrete choice model featuring two explanations: women strongly prefer their specific employer (horizontal difference) or have fewer good employers (vertical difference). Of the 18pp gender gap due to monopsony, I estimate 10pp as attributable to the former and 8pp to the concentration of good jobs for women in the textile sector. This concentration in turn reflects amenities/disamenities present in different sectors and not gender-specific comparative advantage: specifically, eliminating gender gaps in productivity across sectors erodes 4pp of the monopsony gender gap whereas leveling amenities entirely erodes the 8pp gap due to concentration. My findings demonstrate that although the textile industry provides women desirable jobs, this desirability confers its employers with higher monopsony power. By contrast, desirable jobs for men are not similarly concentrated.

*gsharma@mit.edu. I thank my advisors Esther Duflo, David Atkin, David Autor, and Abhijit Banerjee for their comments and encouragement. I also thank Daron Acemoglu, Sydnee Caldwell, Viola Corradini, Jonathan Cohen, Arindrajit Dube, Mayara Felix, Lisa Ho, Simon Jäger, Namrata Kala, Tavneet Suri, Diana Sverdlin-Lisker, Martina Uccioli, Sean Wang, Laura Weiwu, Samuel Young and attendees of MIT's development and labor lunches for helpful comments. I am incredibly grateful to Marc Muendler for sharing access to the SECEX customs data and Chris Poliquin for the crosswalk from CBO to O*NET occupations. Thanks to Joao Fernandes, Stanley Gacek, Juvandia Moreira, Juliana da Penha Thomaz, and Beatriz Santos for their help in better understanding the context. Thanks also to Berger, Herkenhoff, and Mongey (BHM'22) for their MATLAB code on which mine is based. Anika Bokil provided excellent research assistance. Access to Brazil's RAIS database is governed by the Data Use Agreement between MIT and Brazil's former Ministry of Labor. I 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. All errors are mine. First posted version: 19 November, 2022

Introduction

Across most labor markets, women are paid less than men. While these gender wage gaps could emerge purely from productivity differences or taste-based discrimination, they could also reflect lower competition in women’s labor market. Indeed, when Joan Robinson (1933) first introduced the concept of monopsony, she conjectured, “a type of discrimination may arise when men and women of equal efficiency are paid at different rates ... if their conditions of supply are different”. In developing countries, where a lack of safety, sparse job networks, or the notion that certain work is “inappropriate” for women can limit women’s mobility even more than in developed economies, we may naturally expect monopsony to be an even more prominent force generating their large observed gender wage gaps (35% in India, 28% in Brazil, and 22% in Mexico, ILO (2018)). There is, however, surprisingly little evidence of either the extent or the sources of differential monopsony by gender in any setting, but especially in developing countries.

To fill this gap, I study gender differences in monopsony power in the textile and clothing manufacturing industry in Brazil. This industry employs over 90 million workers across the developing world, over half of them women, making it among its largest industrial employers (ILO, 2022). I make three contributions. First, I quasi-experimentally show that employers possess higher monopsony power over women than over men. Exploiting negative, firm-specific demand shocks induced by the end of the Multi-Fiber Arrangement, I find that women are significantly less likely than men to separate from employers that lower their wage, with the resulting gender difference in monopsony power generating an 18pp wage gap among equally productive workers. Underlying women’s lower separation is their inability to leave the textile industry. Motivated by this result, I build a discrete choice model to probe three sources of differential monopsony by gender: (i) women strongly prefer their specific employer, all else equal (horizontal preference), (ii) women value amenities provided by fewer employers (vertical preference), and (iii) women and men have different schedules of industry-specific comparative advantage (vertical productivity). Of the 18pp gender wage gap due to monopsony, I find that 10pp is attributable to women’s stronger affinity for their specific employer and 8pp to the concentration of good jobs for women in the textile industry. Surprisingly, this concentration reflects the concentration of female-focused amenities in the textile industry and not gender-specific comparative advantage. Finally, I study how policy may remedy the monopsony-induced gender wage gap by calibrating the model with my estimates and counterfactually amending its determinants in turn (amenities, skills, and safety).

For data, I use rich employer-employee linked records covering the universe of formal employment in Brazil, merged with customs records detailing establishment-level exports, and the text of all collective bargaining agreements detailing establishment-level amenities.

I begin with the quasi-experiment. Empirically establishing gender differences in monopsony power requires estimating workers’ elasticity of labor supply to a single employer. This in turn requires a firm-level instrument for the wage. I exploit such variation due to the end of the Multi-Fiber Arrangement in 2005, which lifted decades-long export quotas on very specific textile and clothing products from China to the United States, Canada, and the EU. In a single year, Chinese

exports of quota-bound products grew 270% and competing Brazilian exports fell 20%. Because MFA rules deemed virtually identical Chinese products quota-bound and quota-free, the Brazilian workers producing these products were also virtually identical. For example, while “men’s shirts of cotton” and “women’s skirts of wool” were quota-bound, “men’s t-shirts of cotton” and “women’s trousers of wool” remained quota-free. Unsurprisingly, then, the workers engaged in manufacturing exposed and unexposed products were indistinguishable on baseline characteristics, in wage levels, 4-digit occupations, and geographies (of 557). I infer elasticities using a difference-in-differences strategy comparing their wages and employment.

My first takeaway is that employers possess higher monopsony power over women than over men. The MFA expiration causes a 6pp decline in both men and women’s wages, but men are substantially more likely to switch jobs as a result, such that their wages eventually recover whereas women’s remain persistently lower. Five years following the MFA, I find that treated men are 10pp more likely to quit their baseline employer compared to 5pp among women. Several distinct pieces of evidence point to monopsony as an explanation for these findings. An obvious objection is that the findings reveal gender differences in comparative advantage instead of monopsony, i.e., imperfect substitutability between treated and comparison women but not among men. I rule this out by showing the stability of estimates when leveraging variation in treatment among identically skilled workers — tailors. I rule out gender differences in forced separations by showing that most transitions are to new employers (as opposed to unemployment) and correspond with full wage recovery.

Although the MFA shock is firm-specific, estimating residual elasticities with respect to an employer’s own wage change requires ruling out wage spillovers to competing employers that in turn alter labor supply to the originally shocked employer (the exclusion restriction).¹ I rule out (strategic) wage spillovers via a novel test from the exchange rate pass-through literature (Amiti et al., 2019)². Its key insight is that spillovers operate by changing employers’ markdowns — as workers exit China-competing employers, non-China-competing employers can lower wages below marginal product. I show that for any structure of competition among employers (including oligopsony) and invertible labor supply system (including nested CES) I can thus estimate spillovers by regressing an employer’s wage change on a sufficient statistic for changes to competitor wages, controlling for the change in its own marginal product. Using the MFA shock to provide both a market-level instrument for the first and employer-level instrument for the second, I find no spillovers. I interpret this as evidencing the MFA’s small size — it affects less than 2% of establishments.³ Per my elasticity estimates, gender differences in monopsony power would generate an 18pp gender wage gap among equally productive workers, explaining nearly half the observed gender wage gap.⁴

¹The elasticity of residual labor supply is the partial equilibrium elasticity with respect to an employer’s own wage change, holding constant competitor wages. It governs markdowns in standard monopsony models.

²While its primary purpose here is to establish exclusion, the test can identify oligopsony in future work, an open question in the field (Card, 2022).

³My evidence remains consistent with oligopsony. As predicted in standard oligopsony models, own marginal product shocks have incomplete pass-through that falls with employer size.

⁴Average elasticity estimates mask heterogeneity. In both the model and the data, elasticities fall with employer

To probe the source of the gender difference in monopsony power, I next develop a model motivated by four empirical results. First, I find that workers in the MFA’s aftermath are most mobile across employers within the industry, then across industries, and finally across geography. Second, men exit industries substantially more than women. Third, women are disproportionately employed in fewer industries. Over 30% of women work in only two industries, 60% in their top five, and 80% in their top ten, compared to 15%, 35% and 40% among men. Industries thus seemingly differ in gender-specific amenities or comparative advantage. Finally, gender differences in observable skill (measured via O*NET or education) do not drive gender differences in exit, providing the first clue that gender differences in comparative advantage do not drive differential monopsony by gender.

Pursuant to these results, the model features three nests (location, industry, employer) and horizontal and vertical differences across industries as sources of monopsony power. Employers post wages and amenities and workers choose their highest utility employer subject to an idiosyncratic preference draw. Workers are most mobile within an industry, then across industries, and finally across geography, with mobilities governed by three elasticities of substitution: across employers within an industry (η_g), across industries (θ_g) and across geography (λ_g). Because workers flock to relatively desirable (high wage and amenity) employers, these employers are larger in their nest — good textile employers are large in the industry, and good industries large in their geography. Employers compete in a Cournot oligopsony.⁵

A key implication of the model is that a few sufficient statistics quantify the contribution of horizontal and vertical components to the markdown on the average worker’s wage. First, the three elasticities of substitution (η_g , θ_g , and λ_g), govern horizontal preferences. Second, within-industry concentration reflects vertical preferences and is high when only a few within-industry firms offer desirable jobs. It lowers markdowns because these firms compete tightly with less desirable employers within their industry ($\theta_g < \eta_g$). Finally, cross-industry concentration is high when only a few industries offer desirable jobs (vertical preference or productivity). It lowers markdowns because these industries compete tightly with less desirable within-geography industries ($\lambda_g < \theta_g$). Estimation employs moments derived from the labor supply system. I directly observe concentration. The data validate an important model prediction: estimated elasticities fall with employer size, especially when the textile industry is large in its geography.

My second takeaway is that both match-specific reasons and the concentration of women’s jobs in the textile industry generate gender differences in monopsony power. By themselves, match-specific reasons (η_g) prevent workers from exiting atomistic employers even when their local labor market abounds with textile jobs. This generates a 10pp gender wage gap, with the remaining 8pp attributable to gender differences in concentration. Concentration in turn reflects not within, but, rather, higher cross-industry concentration. Within the industry, markets are similarly concentrated

size. However, the model-consistent elasticities calculated later in the paper yield exactly the same gender wage gap: 18%.

⁵The model is similar to Berger et al. (2022), but departs in featuring three nests and amenities, both of which emerge as key features driving differential monopsony by gender.

(men’s Herfindahl Hirschmann Index across textile firms is, if anything, *higher*). However, the textile industry comprises a much larger share of women’s labor market (11% vs 3% for men). In textiles, this concentration contributes 10pp to the monopsony-induced gender wage gap. More broadly in the economy, women’s concentration in fewer industries contributes 12pp to the gender wage gap.

My third takeaway is that women’s higher concentration in textiles is almost entirely attributable to non-wage amenities as opposed to comparative advantage. To infer amenities I use the model’s insight that high wages or amenities draw workers to an industry. I use this structure to back out non-wage amenities given observed wages and industry size. I separately estimate gender-specific productivity across industries via standard production function methods (Ackerman et al., 2015).⁶ Remarkably, I find that productivity differences and gender wage gaps predict a much smaller fraction of women in the textile industry than observed (by 6 times). Non-wage amenities therefore explain textile’s prominence in women’s labor market.

Differential monopsony by gender thus has two intuitive sources: women prefer their specific employer even when their labor market abounds with opportunity; they are also tethered to relatively desirable textile jobs. But what drives these match-specific constraints and sector-specific amenities?

I find that safety importantly predicts the former and amenities in collective bargaining contracts importantly predict the latter. Given Brazil’s highest rate of violent crime in the world (UNODC), one might expect women’s labor supply to be less elastic if unsafe commutes make proximate employers appealing. Exploiting municipality-level data on the homicide rate, I show that low safety (below the 25th percentile) indeed predicts lower labor supply elasticities among women but not men. Exploiting the text of all collective bargaining agreements, I define “female-focused” amenities as those predicting women’s revealed preference value for an employer (following Corradini et al. (2022)). They include provisions governing maternity protections, childcare, and absences, among others. These female-focused contracted amenities strongly positively correlate with model-inferred amenities.⁷

Having uncovered the potential drivers of gender differences in monopsony power — amenities, skills, and safety — I finally study the prospect for policy to remedy the monopsony-induced gender wage gap via counterfactual changes to each. Counterfactuals account for general equilibrium effects, which crucially underpin equity and efficiency. Consider, for example, the effect of leveling cross-industry gender gaps in productivity: as non-textile industries raise compensation, women flock to them from textiles, which erodes the monopsony power of large textile employers. Women’s wages rise on account of lower monopsony, and because they reallocate to large and productive employers whose monopsony wedge disproportionately falls. Men’s wages fall as they reallocate

⁶Specifically, I assume a Cobb-Douglas production function in capital and labor, with labor a CES aggregation of male and female workers. I estimate women’s CES productivity shifter across industries. As a caveat, pending a Brazilian data application, I estimate productivity in two other developing countries, India and Chile, and use the average β estimate across the two.

⁷The correlation need not be causal. Corradini et al. (2022) quasi-experimentally evidence Brazilian women flocking to employers who improve amenities. The counterfactuals employ this quasi-experimental estimate.

downward.⁸ Productive efficiency rises as large employers expand (Baqaee and Farhi, 2020).

By way of benchmark, I estimate an 8pp reduction in the gender wage gap from leveling gender gaps in sector-specific amenities. By contrast, leveling gender gaps in productivity only achieves half this gain (4.1pp). In a policy-relevant counterfactual, I find that leveling gender gaps in contracted amenities reduces the gender wage gap by 4.5pp. Finally, improving safety to lower women’s match-specific preference for their employer reduces the monopsony gap by 3.6pp. Across scenarios, wage gains for women account for 70% of the effect whereas wage losses for men account for 30%.

My fourth and final takeaway is, therefore, that improving non-traditional jobs for women can have positive spillovers by reducing monopsony in women’s current jobs. Given the magnitudes reported above, these improvements must necessarily feature not only within, but, importantly, outside contracts — by combating sexual harassment in the workplace, instituting flexible work arrangements, or expanding maternity protections. While some interventions represent unambiguous improvements (combating sexual harassment), others likely entail costs (enhancing flexibility). I do not speak to the direct costs of these interventions, but evidence two offsetting gains. First, I show that equity begets efficiency. Because upgrading non-textile amenities disproportionately erodes the monopsony power of large employers, their expansion increases productive efficiency in textiles. Second, Corradini et al. (2022) evidence an inefficient underprovision of female-focused amenities in Brazil — the authors observe no tradeoffs (in wages, employment, or profits) following a union reform that institutes large improvements in female-focused amenities among 20% of the formal labor force. The costs of reducing differential monopsony by gender may therefore not be prohibitive.⁹

Related Literature This paper builds on a large literature studying imperfect competition in labor markets (reviewed in Manning (2003), Manning (2011), Card et al. (2018), Sokolova and Sorensen (2021), Manning (2021), Card (2022)), including in developing countries (Tortarolo and Zarate (2018), Felix (2022)). Most closely related are papers examining gender gaps due to imperfect competition (Card et al. (2016), Morchio and Moser (2021), Caldwell and Oehlsen (2022)), and especially those estimating men and women’s firm-specific labor supply elasticities using wage variation at neighboring firms (Barth and Dale-Olsen (2009), Hirsch et al. (2010), Ransom and Sims (2010), Webber (2016)). I make four contributions to this literature. First, I quasi-experimentally show the potential for gender differences in monopsony power to generate large gender wage gaps in an important industrial setting in a developing country (18% of the observed 35%). Relative to prior work, the advantage of a quasi-experiment is that I compare as-if-identical workers (independence condition) and rule out market-level changes in supply as driving results (exclusion

⁸Men reallocate to smaller and less productive employers whenever they are substitutes in production with women, as I assume.

⁹Two findings evidence employers exercising their higher monopsony power over women. First, gender wage gaps rise with employer size in industries that are large in women’s labor market. Second, shocks to marginal product have lower pass-through to women than men’s wages. This is exactly as predicted in oligopsony models where large textile employers are larger in women’s overall labor market.

restriction). The gap I estimate is substantially larger than found in developed economies (for example, Webber (2016) estimates a 3.3% gap in the US). Second, I identify employer differentiation as importantly driving differential monopsony by gender. Consistent with this, Caldwell and Oehlsen (2022) experimentally find no gender difference in monopsony power among US-based Uber drivers, perhaps owing to its similarity with Lyft. Third, I show that traditional measures of labor market concentration, the within-industry or within-occupation Herfindahl Hirschmann Indices can misdiagnose gender differences (e.g., as measured in Azar et al. (2022), Berger et al. (2022), Felix (2022), Rinz (2022)). While this HHI is similar by gender (or lower for women) across nearly all Brazilian industries, women’s labor market exhibits higher concentration because they are disproportionately clustered in fewer industries. New, data-driven methods may uncover more useful labor market boundaries within which to measure concentration to diagnose market power (e.g., Nimczik (2020), Schubert et al. (2021)). In this spirit, ongoing work in Appendix E uses the MFA shock to quasi-experimentally uncover the boundaries of men and women’s labor markets. Finally, I bring a test of strategic price spillovers from trade (Amiti et al., 2019) to labor markets. Valid for any structure of competition and invertible labor supply system, it can uncover oligopsony (or its lack thereof) in future work.

I link the study of gender wage gaps (reviewed in Blau and Kahn (2017)), especially in developing countries (Fletcher et al., 2017), to imperfect competition. Several papers decompose these gaps into a component “explained” by gender differences in observable characteristics, thought to reflect competitive explanations, and an “unexplained” component, potentially reflecting discrimination. I show that even among observably different workers, as men and women often are, gender wage gaps can reflect differential monopsony power in addition to any productivity difference. Finally, my findings add to a burgeoning literature examining the causes of low female labor force participation in developing countries. Recent work identifies the importance of limited autonomy over own earnings (Field et al., 2021), safety (ILO and Gallup, 2017), and behavioral biases (McKelway, 2018). I highlight the role of a different market failure — disproportionate monopsony power over women — in keeping female employment below its competitive market levels.

Outline The rest of the paper proceeds as follows. Section 1 describes the MFA shock and data. Section 2 provides evidence of employers’ higher monopsony power over women. Section 3 documents several empirical facts to motivate key ingredients of the model. Section 4 develops the model and derives sufficient statistics to probe the sources of differential monopsony by gender. I estimate these statistics in Section 5. Section 6 studies the effect of counterfactual policies on the monopsony-induced gender wage gap. Section 7 concludes.

1 Empirical Setting and Data

Estimating firm-specific labor supply elasticities requires a firm-specific instrument for the wage, which satisfies both independence (is as-if-randomly assigned) and exclusion (only alters labor

supply to the originally shocked firm through changes to its own and not competitor wages). This section provides a history of the Multi-Fiber Arrangement and argues that it plausibly satisfies both the independence and exclusion restrictions. I then describe my data sources and report summary statistics. Section 2 explicitly assesses strategic wage spillovers.

1.1 Multi-Fiber Arrangement

For three decades spanning 1974 to 2004, the Multi-Fiber Arrangement (MFA) restricted textile and clothing exports from developing to developed nations (the United States, EU, and Canada) by deeming these products outside the purview of multilateral negotiations at the World Trade Organization (WTO). In 1994, however, the WTO’s Uruguay Round abolished the MFA, with countries agreeing to lift quotas in four phases beginning on 1 January 1995, 1998, 2002 and 2005. Although this expiration was in principle gradual, in practice the most restrictive quotas were integrated in the final phase (2005), accounting for 60% of 10-digit textile and clothing HS codes and 49% of trade volume (Appendix Table C2, Brambilla et al. (2010), Khandelwal et al. (2013)).

MFA quotas historically bound Chinese exports over those from other developing countries. Whereas over 60% of Chinese products faced quotas, this fraction in Brazil was only 14%. Moreover, the “fill-rate”, or share of quota limits actually exported, averaged over 80% in China relative to 13% in Brazil.¹⁰

Upon the MFA’s expiration, Chinese exports of quota-bound products thus grew dramatically (over 270%) between 2005 and 2006 (Figure 1, Panel A), with competing Brazilian exports declining 20% (Panel B). This 2005 ending serves as a negative, firm-specific demand shock to the subset of Brazilian employers exporting hitherto-quota-bound products, and can be employed to trace out men and women’s firm-specific labor supply elasticities.

The empirical setting offered by the MFA is appealing for several reasons. First, because MFA rules required importers to retire quotas on very similar products in each phase, treated (quota-bound in China) and comparison (quota-free in China) products, and, thereby, the Brazilian workers manufacturing them were indistinguishable at baseline (independence condition). Specifically, in each MFA phase countries had to retire quotas on some products in each major textile and clothing category — yarn, fabrics, made-up textile products, and clothing, comprising a non-trivial fraction of 1990 volumes: 16%, 17%, 18%, and 49%. Thus, whereas “men’s shirts of cotton” and “women’s skirts of wool” were quota-bound in China, “men’s t-shirts of cotton” and “women’s trousers of wool” remained quota-free. The Brazilian workers manufacturing MFA-treated and comparison products were as-if-randomly assigned.

Second, the MFA shock plausibly satisfies the exclusion restriction by ushering a small shock in local textile and clothing labor markets — affecting fewer than 2% of establishments and 10% of employment in the average geography.¹¹ Importantly, because the MFA’s impact varies across

¹⁰Even prior to 2005, Chinese quota limits were eased at a much slower pace than those in other developing countries. In addition, Chinese exporters were restricted from shifting quota allocations across MFA good categories and years (Brambilla et al., 2010).

¹¹Brazil has 557 microregions.

geography, I can employ this market-level variation to directly examine spillovers (Section B).¹²

Finally, because the sheer magnitude of the MFA’s effect on Chinese exports was unanticipated, it provides sharp variation in 2005. Expert predictions of Chinese export growth ranged between 6 percent (Diao and Somwaru, 2001) and 104 percent (Rivera et al., 2003), still far below the realized 270 percent growth rate. Not anticipating the large positive shock to Chinese and adverse shock to their own exports, treated Brazilian exporters neither systematically changed their pre-period export composition nor their wages and employment (Figure 4).

1.2 Data

I link three rich sources of data: (i) bindingness of MFA quotas at the product-level from a database compiled by the US Office of Textiles and Apparel; (ii) exports at the establishment-level from customs records at Brazil’s foreign trade department; and (iii) worker outcomes from linked employer-employee records on the universe of formal sector workers.

For product-wise information on MFA quotas in China, I use two datasets from the US Office of Textiles and Apparel (OTEXA).¹³ The first report quota limits, “fill-rates” or how much of the limit was actually exported, and the date of quota removal for each country with which the United States had bilateral quota arrangements. These data span the period from 1984 and 2004. MFA quotas either fully or partially cover 194 different MFA product groups each lying in one of four possible textile and clothing segments: yarn, fabrics, made-ups, and clothing. Examples of MFA groups include “yarns of different colors” and “women’s and girls’ trousers, breeches, and shorts (cotton)”. A companion dataset links every 10-digit HS textile and clothing product with its corresponding MFA subgroup.

For establishment-level information on Brazilian exports I use customs data from the Brazilian foreign trade ministry known as *Secretaria Comércio Exterior* (SECEX). These data span the period between 2000 and 2009 and report annual information on each product exported by an establishment, including its value and destination. Product information is recorded at the 6-digit level, which encodes both product type and material, for e.g. “women’s or girls’ track suits of cotton” and “men’s or boys’ shirts of man-made fibers”. I link exports to MFA quotas using HS codes and to employer-employee records using unique firm identifiers.

For worker-level outcomes I use linked employer-employee records from Brazil’s labor ministry known as *Relação Anual de Informações Sociais* (RAIS), which cover the universe of formal sector workers. For each work spell, RAIS reports a worker’s average monthly earnings, (6-digit) occupation, tenure, hours; worker characteristics like gender, age, and education; and establishment characteristics like location (municipality) and industry (6-digit).

I supplement these datasets with two others: I use O*NET data to impute skills employed in an occupation. These data are based on interviews with thousands of US-based workers and report

¹²Intuitively, I test for strategic wage spillovers by comparing wages at non-MFA-treated establishments in geographies with many versus few MFA-treated establishments.

¹³These datasets were obtained from Brambilla et al. (2010) and are assembled using “U.S. trading partners’ Expired Performance Reports” used by OTEXA to monitor MFA quota compliance.

the skill level (on a 1-8 scale) required to execute each of 35 skills in an 8-digit SOC occupation. Examples of skills include coordination, operation and control, and equipment maintenance, among others. I follow Maciente 2013 in using job titles to map SOC codes to Brazilian occupation codes.

Second, I use the text of all collective bargaining agreements to uncover information on establishment-level amenities. These agreements are registered on the Sistema Mediador registry maintained by Brazil's Ministry of Labor (Lagos, 2019) and describe 137 different types of amenities offered by establishments, including maternity leave, childcare, worker safety, hazard pay, and work hours. I use the revealed preference procedure developed in (Corradini et al., 2022) (Table 2) to classify amenities as female or male-centric. Intuitively, amenities that predict women's (men's) revealed preference value for an employer are female (male)-centric.

1.3 Sample and Treatment Definitions

I define treatment using quotas imposed by the United States, as these were most binding on China. Treatment at the product, establishment, and incumbent worker-level is defined as follows:

Product A textile and clothing product is treated if it was quota-bound in China at baseline and in the comparison group if not. The sample comprises all HS codes with first two digits 61 (textile) or 62 (garments). MFA quotas were imposed at the subgroup level, comprising on average fifteen 10-digit-HS products. I consider a quota binding if its fill rate exceeds 90%, and assign each product the treatment status of its parent-subgroup. Because Brazilian export data is at the 6-digit level, I assign each 6-digit product the treatment status of its modal 10-digit product. Ultimately, this yields 181 treated and 680 comparison products of a total 861.

Establishment An establishment is treated if its highest sales value exported product was quota-bound in China and in the comparison-group if quota-free. This treatment definition is nearly unambiguous since over 50% of establishments export a single product. Even among multi-product exporters, the highest value product accounts for 80% of revenue.

Incumbent worker A worker is treated if employed at a treated establishment at baseline and in the comparison group if employed at a comparison establishment. The treatment definition is robust to reasonable changes along three dimensions: (i) using the share of an establishment's sales under binding quota as opposed to its highest value export, (ii) using an 85% fill rate to define quota bindingness, and (iii) using a higher threshold of treated 10-digit-HS products to define treatment for a 6-digit product.

1.4 Descriptive statistics

I show that Brazilian workers producing China-competing and China-non-competing exports were balanced on observable characteristics. I next show that the MFA shock had small effects on aggregate employment across local textile and clothing labor markets.

Treated and comparison products Table 1 demonstrates through examples that treated and comparison products closely resembled each other in material and good type. For example, while “men’s shirts of cotton” and “women’s skirts of wool” were quota-bound in China, “men’s t-shirts of cotton” and “women’s trousers of wool” were quota-free.

Treated and comparison workers Unsurprisingly, then, the workers employed at treated and comparison establishments were balanced on baseline characteristics. At baseline, they had similar levels of wages, tenure, education, and age (Table 2). The women earned on average 560 Brazilian real per month which was roughly equivalent to USD in PPP terms at the time (879 USD PPP in 2022). They had 4.3 years of tenure and were on average 33 years old. About 75% had not graduated from high school, 24% had high school degrees and only 1% had attended college. The men earned on average 1096 Brazilian real per month. They had 5.2 years of tenure and were on average 32 years old. They had obtained more years of schooling than women: 67% did not have high school degrees, 31% did, and 1% had some college education. Treated and comparison workers are statistically indistinguishable on these observable margins.

Figure 2 (A) shows that treated and comparison workers were employed in overlapping distributions of (4-digit) occupations. Over half of all women and about 20% of men were employed as tailors. The men were employed across a wider range of occupations; prominent ones include spinning operators (9%), production line feeders (8%), and machine operators (6%). Figure 2 (B) shows that treated and comparison workers were also employed in overlapping geographies or microregions (of which there are 557 across Brazil).¹⁴ Finally, Figure 2 (C) shows that treated and comparison workers had similar patterns of employer-to-employer transitions in the pre-period. Each dot on this figure denotes a 4-digit occupation: the x- and y-axes plot the share of transitions from treated and comparison employers into that occupation. The two are strongly positively correlated.

Size of shock Figure 3 and Appendix Figure C2 show that the MFA shock affected a small number of employers and small share of employment in the textile and clothing manufacturing industry. Figure 3 shows that the MFA shock treated on average 2% of establishments and 10% of employment in a geography. Appendix Figure C2 shows that this caused treated employers to lose on average 10-20% of workers over the following five years, but left employment at non-MFA-exporters and domestic producers unchanged. Therefore, the MFA ushered only a small change in aggregate demand in the textile and clothing industry and can be thought of as firm-specific. Section 2 formally establishes a lack of wage spillovers to competing employers.

2 Gender differences in monopsony power

To begin, I establish my main reduced form finding: the MFA spurs wage declines among treated relative to comparison workers, but men exit substantially more in response, such that their wages

¹⁴Although textile workers are concentrated in some hubs of manufacturing, the MFA shock is not. Appendix Figure ?? shows no correlation between the number of total and share of treated textile workers.

eventually recover over five years whereas women’s are persistently lower. I show that lower separations indicate higher monopsony power over women, by ruling out several competing explanations. I conclude by estimating men and women’s average separation elasticities, which govern markdowns across standard monopsony models (e.g., Manning (2011) and Card et al. (2018)). To do so, I rule out strategic wage spillovers to competitors, ensuring that I estimate the partial equilibrium residual elasticity that governs markdowns but does not also encode general equilibrium changes to supply as under oligopsony.

2.1 Reduced Form Effects on Wages and Employment

2.1.1 Empirical Strategy and Identification

I employ the following dynamic difference-in-differences (DiD) specification:

$$Y_{it} = \alpha_i + \gamma_{mt} + \sum_{t=-3}^5 \beta_t(D_i \times 1_{year=t}) + \lambda X_{it} + v_{it} \quad (1)$$

Where i indexes a worker and t a year. The unit of observation is an incumbent worker who is tracked wherever she or he goes. Y_{it} denotes the outcome, which can be either wages or exit. $1_{year=t}$ is an indicator equal to one in periods $t \in \{-3, 5\}$ relative to the MFA’s end in 2005, D_i is a treatment indicator equal to one if a worker’s baseline employer exported a treated product (quota-bound in China) and zero if it exported a comparison (quota-free) product. I control for location-specific shocks to the textile industry through microregion-year fixed effects γ_{mt} . v_{it} is an idiosyncratic error term. I compare treated women with their comparison counterparts and the same for men. In later specifications I compare as-if-identical workers by leveraging variation in treatment within the same geography and occupation through occupation-geography-year fixed effects (X_{it}). Standard errors are clustered by establishment.¹⁵ Because the MFA shock occurred at the same date for all workers and treatment is binary, the OLS twoway fixed-effects estimator is unbiased for the true ATT (De Chaisemartin and d’Haultfoeuille, 2020).¹⁶

The identifying variation in this regression occurs within the same worker, comparing outcomes in any year relative to $t = -1$, and within the same time period, comparing treated to comparison workers. The key identifying assumption is a parallel evolution in outcomes absent the MFA shock. The parallel evolution assumption does not require workers to possess similar baseline levels. However, it may be implausible if workers differ in pre-treatment characteristics that predict the *dynamics* of wage growth or labor supply. I therefore both check for parallel trends in the pre-period and show that treated and comparison workers resemble each other in levels (Section 2).

¹⁵The Bartik analogy for clustering standard errors by establishment is “exogenous shares”, i.e. that establishments are shocked at random. As discussed, the shock reflects as-if-random treatment assignment at the worker level. Clustering standard errors by worker leaves conclusions unchanged.

¹⁶Formally, each worker receives exactly the same weight and the DiD estimate is a weighted average of these heterogeneous treatment effects.

Gender differences I use the following pooled regression to evaluate gender differences in treatment effects:

$$Y_{it} = \alpha_i + \gamma_{mgt} + \sum_{p \in \{1,2\}} \delta_p(D_i \times Post_p) + \sum_{p \in \{1,2\}} \beta_p(D_i \times Post_p \times F_i) + \lambda X_{it} + \epsilon_{it} \quad (2)$$

Where F_i is an indicator for female, $Post_1$ is an indicator equal to one in the first three years following the MFA (years 0 to 2) and $Post_2$ an indicator equal to one in the following three years (years 3 to 5), γ_{mgt} represent microregion-gender-year fixed effects. δ_1 and δ_2 estimate the average treatment effect for men in the two post periods, whereas β_1 and β_2 estimate gender differences in treatment effects. I pool time into these two periods because the dynamic DiD reveals temporal changes in treatment effects between the first three (2005-2007) and subsequent three (2008-2010) post-MFA years.

2.1.2 Results

Wages Figure 4 plots the dynamic treatment effects on the log of workers' monthly earnings. Panel A shows results for women and Panel B for men. Table 3 reports gender differences. There are no pre-trends for either men or women. However, the MFA expiration causes both men and women's earnings to fall 6% points compared to the comparison group. Over the following five years, while men's wages fully recover, women's remain lower (3.6% points by year five). The difference in recovery by gender is statistically distinguishable at the 5% level. Appendix Figure C3 rules out declining hours as an explanation for treated workers' lower earnings (fall over 0.1pp). Appendix Table C3 shows that the negative treatment effect on earnings reflects both slower wage growth compared to the comparison group and nominal declines relative to workers' own $t = -1$ earnings. The likelihood of a nominal earnings drop is 6% points higher among treated men and 5.3% points among treated women.¹⁷

Overall, my findings reveal that the MFA causes both men and women's wages to decline but only men's recover over time. Below I argue that these MFA-induced wage drops signify imperfect competition in both men and women's labor market, and that wage recovery among men reflects higher competition in theirs.

Retention Figure 4, Panels C and D plot the dynamic treatment effects on the probability of retention at a worker's baseline (2004) employer. Both within the first few years following the MFA, and five years later, treated men are substantially more likely to exit treated employers compared to women. In the first three years following the MFA, the treatment effect on exit among men is 3pp, compared to no significant effect for women. By year five, this rate rises to 10pp among men and

¹⁷Brazilian labor law prohibits nominal wage cuts unless authorized by a union. My results reflect declines in the log of monthly nominal earnings as opposed to in contracted wages. In addition to wages, earnings include other forms of compensation such as a 13th month salary, which comprise an insignificant share of overall compensation. At baseline, over 30% of workers employed at textile exporters earn at least 10% over their contracted wage. Among all textile workers this fraction is 43%. I also observe a decline in workers' December earnings (Appendix Table C3), which is the traditional wage measure in papers using RAIS data (e.g. Gerard et al. 2021, Felix 2022).

5pp among women. I find that gender differences in exit entirely account for gender differences in wage recovery following the MFA – while both male and female stayers’ wages remain persistently lower at treated employers, leavers’ wages recover (Figure 5). The following section argues that lower retention among women following wage drops reflects lower competition in women’s labor market.

2.2 Establishing Monopsony Power

At face value, men’s substantially higher exit following similar wage declines indicate employers’ higher monopsony power over women. However, two potential concerns confound this interpretation, which I address in turn.

First, even in a perfectly competitive world, per a standard comparative advantage argument, imperfect substitutability between treated and comparison workers could cause their wages to diverge following a divergence in marginal products. To compare substitutable workers whose wages would not diverge under perfect competition, I first leverage variation in MFA treatment within the same narrow 4-digit-occupation (e.g., tailor, loom operator, supervisor) and geography, via occupation by geography by time fixed effects. I also zoom in on a single large occupation, tailoring, which employs over 50% of women and 20% of men at both treated and comparison employers. Table 3 shows no change in point estimates or conclusions from drawing these narrow comparisons, suggesting that the MFA-induced wage drop reflects imperfect competition and not imperfect substitutability between treated and comparison workers.

An analogous argument rules out gender differences in comparative advantage as driving women’s lower exit. Over 50% of women at both treated and comparison employers are tailors with balanced baseline characteristics (wages, tenure). In a competitive labor market, they relative wages would not diverge absent quits. Yet, I observe treated female tailors earning 3.6% less than comparison counterparts even five years post the MFA. Thus, women’s labor market is imperfectly competitive.

A second concern is that higher male exits indicate larger MFA-induced marginal product declines for men, and, hence, higher layoffs. Two findings counter this claim. First, the entire treatment effect on exit reflects moves to new employers as opposed to unemployment (Figure C4). Second, leavers’ wages fully recover (Figure 5). Both findings contrast with a generation of involuntary job loss studies, which find persistent unemployment and earnings losses among laid-off workers (Jacobson et al., 1993), including in Brazil outside of recessive periods (Hoek (2006), Kaplan et al. (2005)). Higher separations among men therefore suggest quits instead of layoffs.

In addition to ruling out alternatives (to gender differences in monopsony power), I identify two patterns consistent with static wage posting models. First, these models predict that monopsonistic employers with upward-sloping labor supply lose marginal workers off this supply curve when reducing wages, but retain inframarginal workers at the lower wage (Manning, 2003). As predicted, Figure 5 demonstrates persistent wage declines among retained workers (by $t = 4$), alongside full wage recovery among leavers. Second, wage posting models predict wage declines among new hires. As predicted, Appendix Figure C5 shows a 6pp decline in new workers’ wages. Of course, this

decline may reflect quality downgrading. Indeed, such concerns plague early, seminal studies of rent-sharing using establishment-level data (e.g. Van Reenen (1997), Hildreth (1998)). Without taking a strong stance on the topic, I demonstrate no treatment effect on new worker characteristics (Appendix Table C4).

In sum, the MFA expiration spurs wage declines among treated versus seemingly-identical comparison workers, reflecting imperfect competition. These wage drops spur substantially higher separations among men, indicating higher monopsony power over women. Over time, lower monopsony over men spurs full wage recovery whereas higher monopsony power over women manifests in persistently lower wages.

2.3 Estimating average separation elasticities

Empirical strategy and identification I next use an instrumental variables approach to estimate men and women’s average elasticity of labor supply to an employer. These elasticities measure the change in labor supply to an employer for a given change in its offered wage, holding fixed competitor wages. They govern monopsonistic markdowns across standard models, where the firm’s FOC is: $w_{gj} = \frac{e_{gj}}{1+e_{gj}} mrpl_{gj}$. w_{gj} is the wage received by group g at employer j , $mrpl_{gj}$ is their marginal revenue product, and e_{gj} is the elasticity of residual labor supply.¹⁸ It is twice the separation elasticity, which I estimate using the following IV system:

$$\begin{aligned} Exit_{igt} &= \alpha \Delta \ln w_{ig} + \gamma_{1mt} + v_{1igt} \\ \Delta \ln w_{ig} &= \beta D_i + \gamma_{2mt} + v_{2igt} \end{aligned} \tag{3}$$

where $\Delta \ln w_{ig}$ denotes the change in a worker’s offered wage, $Exit_{igt}$ equals one if a worker is at an employer different from her baseline employer as of t , D_i denotes MFA treatment. γ_{mt} capture microregion-year fixed effects. I estimate elasticities using data from years three to five post the MFA and interpret them as five-year elasticities. Because the correct measure of an offered wage is not obvious, I use two different measures which yield identical estimates. First, I use the worker’s own wage change between $t = -1$ and $t = 1$ (main) and, second, I use the change in stayer wages at one’s baseline establishment between $t = -1$ and $t = 1$ (appendix).¹⁹

The two key identifying assumptions are: (i) independence of potential outcomes (change in labor supply) and treatment assignment (change in offered wage) from instrument assignment, and (ii) exclusion (that the shock only affects labor supply through changes to an establishment’s own wage as opposed to changes in amenities or competitor wages). Baseline balance between treated and comparison workers indicates independence.

¹⁸The average elasticity masks heterogeneity: elasticities fall with employer size, in both the model and data. However, estimates of the average are still useful for comparability with a large empirical literature on monopsony (reviewed in Sokolova and Sorensen (2021)).

¹⁹The wage drop between $t = -1$ and $t = 1$ is a credible proxy for the change in offered wage because treated wages fall once and remain persistently lower.

However, the exclusion restriction may be violated if the MFA shock spurs wage spillovers. Intuitively, this is because when employers are large, labor supply to an employer can depend not only on its own wage but also the wages of competitors. A shock to one firm can spur best response wage changes among competitors, in turn changing labor supply to the originally shocked firm. I show that the elasticity in (3) would, then, estimate this composite effect as opposed to the Nash equilibrium e_{gj} governing markdowns (Appendix A.1.4).²⁰ In other words, the instrument would violate exclusion. Prior empirical work on monopsony assumes away wage spillovers, but we should expect them if employers are oligopsonistic (Berger et al., 2022).

Spillovers represent a potent concern in my setting because exporters are typically large and treated firms constitute 7% of the textile market in their geography. The industry is also concentrated: women’s Herfindahl Hirschmann Index (HHI) in the average geography is 0.28 and men’s is 0.31, Thus, there are only $1/0.28 = 3.5$ firms in women’s average geography and $1/0.31 = 3.22$ in men’s.

Testing for strategic wage responses I adapt a test of wage spillovers from the exchange rate pass-through literature (Amiti et al., 2019). Its key insight is that strategic wage spillovers alter markdowns: as workers exit China-competing employers, non-China-competing employers can lower wages below marginal product. I show in Appendix B that \forall competition structures (including oligopsony) and invertible labor supply systems (i.e., employers are not perfect substitutes) I can estimate wage spillovers by regressing an employer’s wage change on changes to its own marginal product and a sufficient statistic for competitor wage changes. This sufficient statistic is $\Delta \ln w_{-jt} = \sum_{j'} s_{j'} \Delta \ln w_{j'}$; it precludes accounting for which competitors change wages. I also derive the estimating equation, which is:

$$\Delta \ln w_{jt} = \delta \Delta \ln mrpl_{jt} + \gamma \Delta \ln w_{-jt} + \xi_{jt} \quad (4)$$

The coefficient δ estimates pass-through of a shock to employers’ own marginal product and γ estimates spillovers.²¹ The MFA shock provides an employer-level instrument $\Delta mrpl$ and a market-level instrument for $\Delta \ln w_{-jt}$ (Acemoglu & Angrist 2001). γ is identified off comparing employers with few and many MFA-competing firms. Table 4 evidences a strong first stage (F-stat > 60).

Table 4 shows that spillovers are absent. Column 1 demonstrates that the MFA had no spillovers on the wages of incumbent workers employed at any exporter in year $t = -1$, Column 2 demonstrates this for MFA-unaffected exporters, and Column 3 for workers at all untreated employers (including non-exporters) in the textile and clothing manufacturing industry. Across samples, $\hat{\gamma}$ is a tightly estimated zero, ruling out changes above 0.01% with a high degree of confidence (95%). In addition, I cannot reject the sharp null of zero spillovers for every employer using randomization inference p-values. Because new worker wages may more accurately capture wages treated workers’ expected

²⁰I show this under the three-nested-CES system and oligopsony I present in the following section.

²¹It’s controlling for changes to one’s own marginal product that makes this a test of strategic wage spillovers; even in a competitive market, a market-level shock spurs wage spillovers but those operate by changing marginal product and not markdowns.

wages, column (4) reports spillover effects on new hire wages. I reject changes over 0.01% with a high degree of confidence.²²

Appendix Section A.1.4 shows that my estimates therefore recover the elasticity governing markdowns.²³ I interpret a lack of spillovers as evidencing the MFA’s small size. Recent work highlights a number of frictions that could explain why non-MFA employers do not best respond by changing wages, including optimization errors (Dube et al., 2020) or nominal rigidities (Hazell and Taska, 2020). I do not interpret a lack of spillovers as ruling out oligopsony because, as predicted under oligopsony, shocks to employers’ own marginal product have incomplete pass-through that falls with employer size (Appendix Figure C8).

Results Table 5 reports estimates of the average separation elasticity among men and women. Panel A evidences a strong first stage: treated men and women’s wages decline by 5.3% points (F-stat 35) and 5.2% points (F-stat 21) respectively between $t = -1$ and $t = 1$ relative to the comparison group. Instrumenting this wage decline using the MFA shock to estimate (negative) separation elasticities yields an estimate of 0.767 among women and 1.921 among men (Panel B). The implied labor supply elasticities — double these separation elasticities — predict that the average woman is paid 61% the value of her marginal product and the average man is paid 79% and I can statistically reject equality by gender at the 90% level. Using weak IV-robust confidence intervals, I can also reject that women earn over 72% the value of their marginal product and men over 86% of theirs.

Gender differences in monopsony would thus generate a 18pp gender wage gap among equally productive men and women, explaining over half the observed 35pp gender wage gap in the textile industry.²⁴ This average estimate masks heterogeneity – large employers are more monopsonistic in both the model (Section 4) and data (Section 5). However, model-consistent measures of the average markdown (surprisingly) also predict an 18pp gender wage gap.

The Brazilian labor market is, therefore, substantially more frictional for women than men, more than in developed countries (for example, Webber (2016) estimates a 3% point gender gap due to monopsony in the United States and Caldwell and Oehlsen (2022) find no gender difference in monopsony among Uber drivers). A number of factors could limit women’s mobility in developing economies —commuting frictions, gender norms, or gendered patterns of sectoral comparative advantage. I next turn to exploring the sources of differential monopsony by gender.

3 Motivating the Model: Stylized Facts

Section 4 develops a discrete choice model to identify the source of differential monopsony by gender. This section presents empirical facts to motivate its key ingredients: (a) three-nested

²²I discuss measurement error in *mrpl* in detail in Appendix B.

²³I illustrate the identification argument for nested CES supply and oligopsony. However, the test enables estimating the partial equilibrium separation elasticity for any competition structure and invertible supply system.

²⁴Appendix A Proposition 2 shows that the share of the gender wage gap potentially attributable to monopsony is given by the ratio of the difference in log average wages and log average markdowns.

structure (location, industry, and employer within industry), (b) horizontal differences by industry, and (c) vertical differences across industries due to industry-specific amenities or productivity. Fact 1 tracks workers' post-MFA destinations to motivate (a). Together, Facts 2 and 3 demonstrate that two-digit industries delimit women's opportunity more than men's, motivating (b) and (c). Finally, Fact 4 shows that gender differences in observable skill do not account for gender differences in exit. Therefore, I do not explicitly model gender-specific comparative advantage at baseline, but revisit it in Section 5.

Fact 1: MFA-treated workers are most likely to switch to a new textile employer, then to switch industries, and finally to switch geography

Fact 1 demonstrates that workers are appropriately considered as first choosing locations, then industries, and finally within-industry employers. Table 6 reports treatment effects on workers' likelihood of being employed at a new (non-baseline) textile employer, non-textile employer, and in a new geography, three to five years following the MFA. Examples of two digit industries include the manufacturing of metal products, food and beverages, and leather. I find that five years following the MFA, treated men are 5.9% points more likely than the comparison group to switch to a new textile employer, 4% points more likely to switch to a non-textile employer, and no more likely to switch geography. Among women, these rates are 4.8% points, 0.1% points and 0% points respectively. Thus, in the MFA's aftermath workers are most likely to move to new textile employers, then to non-textile employers, and virtually none exit their geography.²⁵

Fact 2: Men exit industries substantially more than women

Fact 2 suggests that women's lack of mobility outside the textile industry governs their low responsiveness to wage drops. Table 6 shows that gender differences in cross-industry switches entirely account for gender differences in exit (Columns 2 and 3). While within-industry moves explain the entire treatment effect on exit for women, they account for only 5.9% points of the total 10% point effect on exit among men.

Therefore, the model introduces two forces that could tether women to the textile industry: idiosyncratic reasons that are specific to a worker and industry and horizontally differentiate employers, and industry-specific attributes that vertically differentiate textile employers by making the industry on average more attractive for all women. The first may encompass different proclivities for learning new skills. The second may encompass common amenities, such as longer maternity leaves, flexibility, safety, female coworkers, or childcare. Vertical differentiation may also arise if women are disproportionately more productive within than outside textiles (gender differences in comparative advantage).

Fact 3: Industries are large, especially for women

Fact 3 suggests that industries likely differ in amenities or productivities for women more than for men. Figure C13 shows that the textile

²⁵Limited geographic mobility in the face of the MFA shock is consistent with a number of studies finding the same result following large trade-induced negative shocks (Topalova 2010, Autor et al. 2013) including in Brazil (Dix-Carneiro & Kovak 2015).

industry in particular, and a small number of industries in general, employ a disproportionately large share of women. Among high school dropouts, over 11% of women but only 3% of men are employed in textile jobs.²⁶ More generally, over 40% of women compared to 20% of men are employed in their two most important industries; over 60% of women compared to 35% of men are employed in their five most important industries; and over 80% of women compared to 40% of men are employed in their ten most important industries. Together, these suggest that average attributes of an industry, such as amenities or industry-specific skills could systematically differ by gender.

Fact 4: Gender differences in observed skill do not drive gender differences in exit

Finally, Fact 4 suggests that gender differences in comparative advantage likely do not prevent women from leaving the textile industry. Table 7 shows that, following the MFA, women are as unlikely to (i) exit employers, and (ii) exit industries after controlling for gender differences in skill. I use three measures of skill. First, workers’ education, split into three categories: high school dropouts, high school graduates, and some college education. Second, workers’ 4-digit occupation, which proxies for currently employed skills. Finally, I create an O*NET-based measure of skill transferability by linking workers’ currently employed skills to those required in other jobs within their geography (Macaluso, 2022) (described in Appendix D.2.1). O*NET reports the skill level (from 1-8) on each of 35 skills employed in any occupation. Skill transferability is a weighted average of the skill distance (L1-norm) between a worker’s current occupation and other occupations in her geography, with weights equalling the share of local jobs in that occupation. I separate workers into 10 deciles of skill transferability.

I control for skill transferability through fixed effects that interact skill level with treatment status and indicators for each of two post periods (years 0-2 and years 3-5).²⁷ Controlling for a common average rate of exit among identically skilled workers, women remain as unlikely to exit their MFA-treated employer and industry as without controls. Therefore, gender differences in skill transferability do not by themselves explain differential rates of exit among men and women.

Next I turn to describing the model motivated by these facts.

4 Model: Sources of Gender Differences in Monopsony Power

In this section I develop a nested logit model to microfound gender differences in monopsony power due to both horizontal and vertical differences across employers and industries. I derive a set of sufficient statistics to quantify the relative contribution of each to differential monopsony by gender. I estimate these statistics in Section 5, and, in Section 6, I use the model’s structure to study the effect of counterfactual policies amending the horizontal and vertical determinants of monopsony power on the gender wage gap. Appendix A reports all omitted proofs and derivations. The model is similar to those in (Berger et al., 2022), which features nested logit supply and oligopsony among

²⁶Figure C13 plots the share of workers’ wage bill in various industries, which is the model-consistent measure of size.

²⁷Formally, the regression is $Y_{it} = \alpha_i + \gamma_{mgt} + \sum_{p \in \{1,2\}} \delta_p(D_i \times Post_p \times skill_{level}) + \sum_{p \in \{1,2\}} \beta_p(D_i \times Post_p \times F_i) + \epsilon_{it}$.

employers, but departs from it in featuring three nests and amenities, which emerge as important forces driving gender differences in monopsony power.

4.1 Setup

Market The economy features a continuum of geographies $r \in [0, 1]$ (microregions). Each geography features a discrete number of industries indexed by $k \in 1, \dots, M_r$, and firms within the industry $j \in 1, \dots, J_m$. Throughout my exposition I denote a geography as a “local labor market” and all same-gender jobs within an industry as the within-industry market.

Firms Firms compete in a Cournot oligopsony. Each posts gender-specific wages $\{w_m, w_f\}_j$ and chooses employment to maximize profits given the (inverse) labor supply it faces. The revenue function of each firm is differentiable and concave, with imperfect substitutability between male and female workers $F(f_j, m_j)$. Each firm is associated with exogenously-given industry-specific amenities (a_{gk}) and the firm’s deviation from this industry norm (a_{gj}). I abstract from competition in amenities for simplicity.²⁸

Workers Workers, indexed by group g (men or women), possess heterogeneous preferences over employers. They choose to work at their higher utility employer, and exhibit three-nested preferences, choosing first a location, then an industry, and finally an employer within the industry. Their utility from working at employer j has a common group-specific component (rising in wages and amenities), and an idiosyncratic preference shock specific to each employment relationship ϵ_{igjk} . Each worker must earn an income $y_i \sim F(y)$, where earnings are given by $y_i = w_{gj}h_{igj}$, with h_{igj} denoting hours of supplied work.

$$v_{igjk} = \ln w_{gj} + \ln a_{gk} + \ln a_{gj} + \epsilon_{igjk}$$

ϵ_{igjk} has a nested GEV Type I extreme value distribution. Its variance depends on three dispersion parameters η_g , θ_g , and λ_g .²⁹

$$F(\epsilon_{i1}, \dots, \epsilon_{NJ}) = \exp \left[- \sum_r \left(\sum_{k=1}^M \left(\sum_{j=1}^{J_m} e^{-(1+\eta_g)\epsilon_{igjk}} \right)^{\frac{1+\theta_g}{1+\eta_g}} \right)^{\frac{1+\lambda_g}{1+\theta_g}} \right]$$

²⁸Ex-post I observe that the amenities contributing to differential monopsony by gender are not directly contracted upon by employers. Therefore, one answer to the question of why non-textile industries do not provide these amenities if they confer monopsony in textiles is that they may represent features of the workplace that are difficult to change. One example is a majority-female workforce that improves safety.

²⁹As shown by Anderson et al. (1992), subject to functional form choices about the terms in the CES function, this discrete choice framework generates the same labor supply to firms as obtained in a nested CES setting with the outermost nest being regions, the middle nest being industries, and the bottom nest being employers within the industry. η_g then corresponds with the within-industry elasticity of substitution, θ_g is the cross-industry elasticity of substitution, and λ_g the cross-region.

Amenities Amenities a_{jg} and a_{kg} represent non-wage attributes that are commonly valued by all members of a group and vertically differentiate employers. They include tangible attributes of a workplace, such as flexible work or maternity leave policies, female coworkers, and female-friendly managers.³⁰ They may additionally constitute intangible gender norms that render certain work “inappropriate” for women. Amenities in this setup may also represent hiring discrimination.

Wages Wages rise in employer productivity and vertically differentiate employers. While the conclusions of this section remain valid absent assuming a functional form for the production function, Section 6 demonstrates how gender differences in industry-specific comparative advantage, by altering gender and industry-specific wage indices, may vertically differentiate employers across industries.

Dispersion parameters Dispersion parameters govern the distribution of idiosyncratic draws. The presence of idiosyncracies horizontally differentiates employers. Within an industry, η_g represents cross-employer mobility costs such as due to commuting frictions (Robinson, 1933).³¹ The higher the η_g , the lower the variance of idiosyncratic draws across employers within an industry and, thus, the higher workers’ cross-employer mobility. θ_g represent cross-industry mobility costs such as an aversion to learning new skills. The higher the θ_g , the lower the variance of idiosyncratic draws across industries and, thus, the higher workers’ cross-industry mobility. Finally, λ_g represent cross-location mobility costs, such as those associated with family relocation. A high λ_g lowers the overall variance of draws, and raises cross-location mobility. Pursuant to Fact 1 (Section 3), I assume that workers find it easiest to substitute across employers within an industry, then across industries, and finally across locations, $\eta_g > \theta_g > \lambda_g$. Although I do not impose this during estimation, I find it is true.

In this setup, η_g parameterizes match-specific reasons that prevent workers from switching to a new within-industry employer even when their own employer is atomistic. By contrast, workers’ cross-industry (θ_g) and cross-location (λ_g) parameters govern the contribution of concentration to monopsony, as I show below.

Labor supply I obtain labor supply to an employer by aggregating the preferences of individual workers. The probability of choosing employer j is rising in its wage and amenity and is given by the standard nested logit formula (McFadden, 1978):

³⁰A preference for female coworkers creates the possibility for multiple equilibria. While a fascinating question, modeling endogenous amenities is currently beyond the scope of this paper. Section 5 studies the sources of monopsony power in the current equilibrium. The prospect of multiple equilibria may alter the quantitative conclusions of counterfactuals, and they ought to be interpreted with this caveat.

³¹ η_g may also roughly represent loyalty, which is a commonly cited female attribute among employers in developing countries. The simple, static formulation of η_g developed here does not fully capture the dynamic nature of loyalty but serves as a rough proxy.

$$p_{gj} = \underbrace{\frac{(a_{gj}w_{gj})^{1+\eta_g}}{\sum_{j' \in k} (a_{gj'}w_{gj'})^{1+\eta_g}}}_{\text{prob of choosing firm } j \text{ in industry } k} \times \underbrace{\frac{a_{kg}^{1+\theta_g} \left(\sum_{j \in k} (a_{gj}w_{gj})^{1+\eta_g} \right)^{\frac{1+\theta_g}{1+\eta_g}}}{\sum_{k' \in R} a_{k'g}^{1+\theta_g} \left(\sum_{j \in k'} (a_{gj}w_{gj})^{1+\eta_g} \right)^{\frac{1+\theta_g}{1+\eta_g}}}}_{\text{prob of choosing industry } k} \times \underbrace{\frac{\bar{W}_{gr}^{1+\lambda_g}}{\sum_{R'} \bar{W}_{gr'}^{1+\lambda_g}}}_{\text{prob of choosing region } r}$$

Aggregating these probabilities over workers yields the upward-sloping labor supply to employer j :

$$n_{gjk} = \left(\frac{w_{gjk}}{\bar{W}_{kg}} \right)^{\eta_g} \left(\frac{\bar{W}_{kg}}{\bar{W}_{gr}} \right)^{\theta_g} \left(\frac{\bar{W}_{gr}}{\bar{W}_g} \right)^{\lambda_g} a_{gjk}^{1+\eta_g} a_{gk}^{1+\theta_g} N_g$$

Here $\bar{W}_{kg} = \left(\sum_{j \in k} a_{gjk} w_{gjk}^{1+\eta_g} \right)^{\frac{1}{1+\eta_g}}$ denotes the amenity-adjusted wage index of industry k , \bar{W}_{gr} denotes the wage index of region r , and \bar{W}_g the aggregate wage index of group g . The bars indicate that these expressions also include amenities, for example $(w_{gjk})^{\eta_g} = (a_{gjk} w_{gjk})^{\eta_g}$.

Workers flock to relatively desirable industries that offer relatively higher amenities or wages and, within them, to relatively desirable employers. These industries and employers are therefore larger, i.e., comprise a larger share of their within-nest wage bill.

$$s_{gkr} := \frac{\sum_{j' \in k} w_{gj'} n_{gj'}}{\sum_{k' \in R} \sum_{j' \in k'} w_{gj'} n_{gj'}} = \frac{a_{kg}^{1+\theta_g} \left(\sum_{j \in k} (a_{gj}w_{gj})^{1+\eta_g} \right)^{\frac{1+\theta_g}{1+\eta_g}}}{\sum_{k' \in R} a_{k'g}^{1+\theta_g} \left[\sum_{j' \in k'} (a_{gj'}w_{gj'})^{1+\eta_g} \right]^{\frac{1+\theta_g}{1+\eta_g}}}$$

$$s_{gjk} := \frac{w_{gj} n_{gj}}{\sum_{j' \in k} w_{gj'} n_{gj'}} = \frac{(a_{gj}w_{gj})^{1+\eta_g}}{\sum_{j' \in k} (a_{gj'}w_{gj'})^{1+\eta_g}}$$

Because the solution concept is Cournot, employers optimize taking into account their inverse labor supply curve, which is given by:

$$w_{gjk} = \left(\frac{n_{gjk}}{N_{gkr}} \right)^{\frac{1}{\eta_g}} \left(\frac{N_{gkr}}{N_{gr}} \right)^{\frac{1}{\theta_g}} \left(\frac{N_{gr}}{N_g} \right)^{\frac{1}{\lambda_g}} W_g a_{gk}^{\frac{-(1+\theta_g)}{\theta_g}} a_{gjk}^{\frac{-(1+\eta_g)}{\eta_g}}$$

Monopsony Power Employers face upward-sloping labor supply curves with residual labor supply elasticities that depend on shares as well as the three dispersion parameters:

$$e_{gj} = \left[\frac{1}{\eta_g} + \left(\frac{1}{\theta_g} - \frac{1}{\eta_g} \right) s_{gjk} + \left(\frac{1}{\lambda_g} - \frac{1}{\theta_g} \right) s_{gjk} s_{gkr} \right]^{-1} \quad (5)$$

These elasticities represent the partial equilibrium concept: how much does labor supply to an employer change with respect to its own wage, *holding fixed its competitors' responses*. Lower elasticities indicate higher monopsony power.

In this model, horizontal differences confer all employers with monopsony power and vertical differences confer large employers with higher monopsony power. Atomistic employers within an

industry face labor supply elasticity η_g because match-specific reasons represent the only force tethering workers to these small employers.³²

Large employers in any industry (high s_{gjk_r}) face less elastic labor supply because they compete more intensely with other, relatively less desirable, employers within their industry compared to those outside the industry ($\theta_g < \eta_g$). The condition ($\theta_g < \eta_g$), which I assume analytically but observe empirically, is key for this conclusion. When $\theta_g \rightarrow \eta_g$, employers compete in a geographically unified market such that large employers in an industry that are atomistic in their geography (small $s_{gjk_r}s_{gkr}$) face exactly the same elasticity as those that are also atomistic in the industry, η_g .

Large employers in large industries (high $s_{gjk_r}s_{gkr}$) are also larger in their geography and face less elastic labor supply on account of competing more intensely with relatively worse within-geography employers ($\lambda_g < \theta_g$). As above, the condition ($\lambda_g < \theta_g$) is key. When $\lambda_g \rightarrow \theta_g \rightarrow \eta_g$ employers compete in a nationally unified market and each employer is an atomistic monopsonist facing elasticity η_g .

Labor demand Each employer chooses employment to maximize profits given the inverse labor supply it faces, yielding the following FOC:

$$\frac{\partial R_j}{\partial n_{gj}} = w_{gj} \underbrace{\left(1 + \frac{1}{e_{gj}}\right)}_{\mu_{gj}^{-1}}$$

where e_{gj} denotes the elasticity of residual labor supply. Wages are marked down relative to workers' marginal revenue product of labor whenever $e_{gj} < \infty$.

4.2 Monopsony over the average worker

A key implication of the model is that market power, defined as the markdown on the average worker's wage, is determined by a few sufficient statistics.

Proposition 1 *The inverse markdown for the average worker in a group g in industry k in region r is:*

$$\bar{\mu}_{gkr}^{-1} = \frac{\bar{m}rpl_{gkr}}{\bar{w}_{gkr}} = \underbrace{1 + \frac{1}{\eta_g}}_{(i)} + \underbrace{\left(\frac{1}{\theta_g} - \frac{1}{\eta_g}\right) HHI_{gkr} + \left(\frac{1}{\lambda_g} - \frac{1}{\theta_g}\right) s_{gkr} HHI_{gkr}}_{(ii)}$$

Where the inverse average markdown is defined as the ratio of the average worker's marginal product ($\bar{m}rpl_{gkr} = \frac{\sum_{i \in gkr} mrpl_i n_i}{\sum_i n_i}$) and wage ($\bar{w}_{gkr} = \frac{\sum_{i \in gkr} w_i n_i}{\sum_i n_i}$). $HHI_{gkr} = \sum_{j \in k} s_{gjk_r}^2$ measures within-industry concentration through the payroll-weighted Herfindahl-Hirschmann Index and $s_{gkr} HHI_{gkr}$ weights this within-industry HHI_{gkr} by the share of the industry's wage bill in region r (s_{gkr}). Term (i) represents monopsony due to match-specific reasons and term (ii) represents monopsony due to concentration.

³²Changes in a firm's employment do not move the employment index in its industry or geography when it is infinitesimal.

Proof. The proof has two parts. First I show that the inverse of a group’s average markdown is the share-weighted average of individual employers’ inverse markdowns, i.e. $\bar{\mu}_{gkr}^{-1} = (\sum_{j \in k,r} s_{gj} \mu_{gj}^{-1})$. To obtain the expression from the proposition, I then calculate this inverse markdown ($\bar{\mu}_{gkr}^{-1}$) through the lens of the model by substituting in model-based inverse elasticities from equation (5). Proof in Appendix A. ■

The average worker faces a lower markdown (i.e., higher monopsony power) when match-specific reasons limit her mobility even as her local labor market abounds with textile jobs (low η_g). In addition, she faces a low markdown when fewer within-industry employers offer relatively desirable jobs, thereby raising concentration (HHI_{gkr}). The effect of fewer desirable within-industry jobs is exacerbated when the industry itself is relatively desirable, rendering each employer larger in its geography ($s_{gkr} HHI_{gkr}$).

4.3 Sufficient statistics: gender differences in monopsony power

The model yields two sufficient statistics to decompose the contribution of workers’ match-specific preference for their employer and concentration to gender differences in monopsony power. It yields four additional sufficient statistics to decompose the relative contribution of within and cross-industry concentration to the monopsony-induced gender wage gap.

Decomposition #1: Match-specific preference versus concentration Together, the elasticity of labor supply to atomistic employers, η_g , and the elasticity associated with the average inverse markdown, \bar{e}_g , quantify the relative contribution of match-specific and concentration to differential monopsony by gender.

Implication 1 Match-specific reasons differentially increase monopsony over women if women’s labor supply to atomistic employers is less elastic than men’s ($\eta_f < \eta_m$). In other words, $e_{fj} < e_{mj}$ even when $s_{gjk} \sim 0$. The share of the monopsony-induced gender wage gap that is attributable to worker-employer-specific match quality is therefore $\frac{\eta_m}{\eta_m+1} - \frac{\eta_f}{\eta_f+1}$.

Proof. Implication 1 comes from observing that $e_g = \eta_g$ when $s_{gjk} = 0$ in Equation (5). ■

Implication 2 Concentration differentially increases monopsony over women if women’s average elasticity (corresponding with the average inverse markdown) is further from η_f relative to this difference among men. In other words, $e_f(\bar{\mu}_f^{-1}) - \eta_f > e_m(\bar{\mu}_m^{-1}) - \eta_m$. The share of the monopsony-induced gender wage gap that is attributable to gender differences in concentration is therefore $\left(\frac{\bar{e}_m}{\bar{e}_m+1} - \frac{\bar{e}_f}{\bar{e}_f+1}\right) - \left(\frac{\eta_m}{\eta_m+1} - \frac{\eta_f}{\eta_f+1}\right)$.

Proof. Implication 2 comes from observing in Proposition 1 that concentration causes the average markdown to deviate from $1 + \frac{1}{\eta_g}$. Therefore, the average elasticity deviates from η_g . ■

Decomposition #2: Within or cross-industry concentration? Four sufficient statistics quantify the relative contribution of concentration within the industry and the concentration of

employment in the industry to the gender gap due to monopsony. First, HHI_{gkr} encodes the within-industry concentration of desirable employers. By itself, within-industry concentration increases monopsony whenever workers find it harder to substitute across than within industry, $\left(\frac{1}{\theta_g} - \frac{1}{\eta_g}\right) > 0$. Third, $s_{gkr}HHI_{gkr}$ encodes the concentration of desirable employers in a geography that are in industry k . This increases monopsony whenever workers find it harder to substitute across geography than across industries in a geography $\left(\frac{1}{\lambda_g} - \frac{1}{\theta_g}\right) > 0$

Before I proceed to estimation it is worth highlighting that the model distills a number of intuitive sources of differential monopsony over women to two forces: (i) match-specific reasons that prevent women from exiting although their local labor market abounds with opportunity (η_g), and (ii) concentration. For example, commuting frictions when homes are idiosyncratically distributed lowers η_g . Fewer employers paying women high wages w_{gj} or offering desirable amenities a_{gj} raises within-industry concentration (HHI_{gkr}). Hiring discrimination in an industry k' lowers its amenity $a_{k'g}$ and raises the size (s_{kg}) and monopsony power of employers in non-discriminating industry k . Wage discrimination similarly lowers an industry's wage index and raises the size (s_{kg}) and monopsony in non-discriminating industry k .³³ Higher comparative advantage in k raises its compensation, size, and, consequently, monopsony power.

5 Sources of Gender Differences in Monopsony Power

This section takes the model to the data. I first show that both match-quality and the concentration of women's employment drive differential monopsony by gender (Decomposition #1). Next, I show that concentration is itself driven by the concentration of women's jobs in the textile industry as opposed to gender differences in concentration *within* textiles (Decomposition #2). Finally, I show that women's concentration in textiles can be largely attributed to non-wage amenities and not gender differences in productivity. Therefore, although the textile industry draws women through desirable jobs, this in turn confers textile employers with higher monopsony power over women.³⁴

5.1 Match-specific constraints and concentration (Decomposition #1)

As shown in Section 4, η_g (the elasticity to atomistic employers) and its difference with \bar{e}_g (average elasticity of group g) capture the relative contribution of match-specific preferences or constraints and concentration to monopsony. To estimate each, I first estimate heterogeneous elasticities by employer size. The elasticity to atomistic employers yields an estimate of $\hat{\eta}_g$. Aggregating heterogeneous elasticities to an industry-level average by using the identity $\bar{\mu}_{gk}^{-1} = \sum_r s_{rk}^g \sum_{j \in k,r} s_{gj} \mu_{gj}^{-1}$ yields \hat{e}_g . Here s_{rk}^g is the share of the total gender-specific textile wage bill in region r . Appendix A shows that this sum aggregates to the average markdown of group g in industry k (closely linked to Step 1 of the proof of Proposition 1).

³³Discrimination could be a monopsonistic wedge between wages and marginal product or a non-monopsonistic wedge, which is a common way to model it e.g. Becker (1962), Altonji and Blank (1999), Hsieh et al. (2019).

³⁴All exercises in this section pertain to workers without high school degrees, as they constitute over 75% of textile workers.

I estimate heterogeneous elasticities using the following IV system, separately for regions with a large (over 10% of local wage bill share) and small (less than 10% of local wage bill share) textile industry, $L_{txt,j} \in 0, 1$:

$$\begin{aligned} Stay_{igt} &= \alpha_1 D_i + \alpha_2 D_i s_{gj} + \alpha_3 s_{gj} + \gamma_{1mt} + v_{igt} \\ \Delta \ln w_{igt} &= \beta_1 D_i + \beta_2 D_i s_{gj} + \beta_4 s_{gj} + \gamma_{2mt} + \psi_{igt} \end{aligned}$$

In this system, s_{gj} denotes an employer's share of the within-industry wage bill and γ_{mt} are microregion-year fixed effects. Other variables are defined as when estimating average elasticities. The coefficients α_1 (β_1), and α_2 (β_2) capture the average effect of treatment and its differential effect by employer share. The elasticity of labor supply to employer j is:

$$\frac{\Delta \ln n_{jgt}(s_{jg}, s_{gk})}{\Delta \ln w_{jgt}(s_{jg}, s_{gk})} = \frac{\hat{\alpha}_1 + \hat{\alpha}_2 s_{gj}}{\hat{\beta}_1 + \hat{\beta}_2 s_{gj}}$$

As with average elasticities, the exclusion restriction to estimate residual elasticities requires the MFA to not spur wage spillovers or changes to an establishment's amenities. In Appendix C.1.1 I formally show that the MFA satisfies both conditions.

Table 8 reports heterogeneous elasticities. As predicted, elasticities for both men and women fall as employers grow larger. They fall faster and are lower when the textile industry also comprises a large share of the overall wage bill in a worker's geography. I find that both match-specific reasons and concentration generate differential monopsony by gender. I estimate $\eta_f = 2.19$ and $\eta_m = 3.89$ (with p-value ≤ 0.05 for the difference). I estimate $\bar{e}_{fk} = 1.236$ and $\bar{e}_{mk} = 2.689$. These elasticity estimates imply that the average man earns 73% of the value of his marginal product whereas the average woman earns 55% the value of hers. If, instead, only match-specific reasons governed monopsony over workers then the average man would earn 79% the value of his marginal product and the average woman 69%. Therefore, match-specific reasons alone contribute $\frac{\eta_f}{\eta_f+1} - \frac{\eta_m}{\eta_m+1} = 10pp$ to the monopsony-induced gender wage gap. Concentration additionally contributes $\left(\frac{\bar{e}_f}{\bar{e}_f+1} - \frac{\bar{e}_m}{\bar{e}_m+1}\right) - \left(\frac{\eta_f}{\eta_f+1} - \frac{\eta_m}{\eta_m+1}\right) = 8pp$.

The decomposition would misestimate the relationship between concentration and monopsony if workers sort on elasticities. For example, if elastic women sort to locations that are abundant in non-textile jobs, I would wrongly infer their high elasticities as reflecting low concentration. Two facts mitigate this concern: first, most workers work where they are born (de Lima Amaral, 2013). About 6% of the population migrated between 1991 and 2000.; second, nearly 30% of women's non-textile jobs are in the public sector, whose size in a worker's birthplace is predicted by population density (calculated).

In sum, I find that both match-specific reasons and concentration of job opportunity among a smaller set of employers importantly generate gender differences in monopsony power. Below I turn to exploring the drivers of each.

Real-world determinants of η What governs women’s match-specific affinity for their employer in the real world? Appendix Table C8 explores heterogeneity along two dimensions: motherhood and safety. First, one might expect current or soon-to-be mothers to supply labor less elastically than older workers. I find no evidence of this when comparing older women with their younger counterparts of childbearing age (20-35 years) (Columns 1 and 2).

Second, given that Brazil features the highest rate of violent crime in the world (UNODC), one might expect women’s labor supply to be less elastic if unsafe commutes make proximate employers appealing. I implement this heterogeneity check using municipality-level data on homicide rates from (Dix-Carneiro et al., 2018). Defining a municipality lying above the 75th percentile of homicides in the pre-period (2000-2004) as unsafe, I find that low safety indeed predicts lower labor supply elasticities among women but not men.³⁵ While women’s labor supply elasticity to atomistic employers in safe municipalities is 2.278, in unsafe municipalities it is 1.727 (p-val of difference < 0.05). When translated to markdowns, employers’ relatively higher monopsony power over women in unsafe municipalities would spur a 6pp higher gender wage gap compared to safer municipalities.

5.2 Within or cross-industry concentration? (Decomposition #2)

Having demonstrated that higher concentration among women contributes 8pp to the monopsony-induced gender wage gap, I employ Decomposition#2 to quantify the relative contribution of within and cross-industry concentration. Of the four sufficient statistics underlying this decomposition, I estimate the gap between the inverse parameters $\left(\frac{1}{\theta_g} - \frac{1}{\eta_g}\right)$ and $\left(\frac{1}{\lambda_g} - \frac{1}{\theta_g}\right)$ via moments from the labor supply system. I directly measure concentration (HHI_{gkr}) and industry size (s_{gkr}) in the RAIS data. I aggregate each quantity to an industry-level average through a weighted sum over regions, i.e. $HHI_{gk} := \sum_r s_{kr}^g HHI_{gkr}$ and $s_{gk} HHI_{gk} := \sum_r s_{kr}^g s_{gkr} HHI_{gkr}$, where s_{kr}^g denotes the share of a group g ’s overall textile wage bill located in region r .³⁶

Gaps between elasticities of substitution The parameters θ_g and λ_g govern mobility across industries and geography. I estimate them using the following partial derivatives with respect to the wage at employer j :

$$\begin{aligned} \frac{\partial \ln n_{gkr,j}}{\partial \ln w_{gjr}} &= \theta_g (s_{gj})(1 - s_{gk}) + \lambda_g s_{gj} s_{gk} = \theta_g s_{gj} \text{ when } s_{gk} \sim 0 \\ \frac{\partial \ln n_{gr,j}}{\partial \ln w_{gjr}} &= \lambda_g s_{gj} s_{gk} \end{aligned}$$

Intuitively, η_g is the elasticity of labor supply to an atomistic employer, θ_g the elasticity of labor supply to an atomistic industry (from a large employer in that industry), and λ_g the elasticity of

³⁵Brazil has 2500 municipalities, which are therefore at a lower level of aggregation than its 557 microregions that comprise regions r in the model.

³⁶The weighted sums correspond with the inverse markdown facing the average worker from g employed in k : $\bar{\mu}_{gk}^{-1} = \sum_r s_{kr}^g \bar{\mu}_{gkr}^{-1}$ (Appendix A).

labor supply to an atomistic geography (from a large employer in that geography). Heterogeneous preferences are the only force tethering workers to these atomistic employers, industries, and geographies. These partial derivatives are not typically estimable in the presence of wage spillovers or amenity changes. However, I formally show that the MFA provides a credible instrument because it does not spur wage spillovers and because it is assumed orthogonal to amenity changes (Appendix C.1.1). The following establishment-level regressions estimate θ_g and λ_g :

$$\Delta \ln n_{gtxt,j} = \theta_g s_{gj} \Delta \ln w_{gtxt,j} + \epsilon_j$$

$$\Delta \ln n_{gr,j} = \lambda_g s_{gj} s_{gk} \Delta \ln w_{gtxt,j} + \epsilon_j$$

Table 9 reports estimates. Men and women are nearly identically mobile across industries (θ_g) and locations (λ_g). I estimate $\hat{\theta}_f = 0.89$, $\hat{\theta}_m = 0.87$, $\hat{\lambda}_f = 0.03$, and $\hat{\lambda}_m = 0.05$. An industry would therefore have to raise its wage by $1/0.87=114\%$ in order to double its workforce and a microregion by 2000%. Limited geographic mobility ($\lambda_g \sim 0$) in the face of a negative wage shock is consistent with a similar finding following other large, adverse trade shocks (Topalova (2010), Autor et al. (2014)), including in Brazil (Dix-Carneiro and Kovak, 2015). My estimate of θ_g is nearly identical to Felix (2022), who uses the same data but defines a market as an occupation and geography. We would expect this similarity given that each textile occupation is in the textile industry, and workers are immobile across geography. It is thus reassuring that they are. I use these estimates of θ_g and λ_g as my baseline values. In Appendix C I assess robustness to re-estimating them using a technique due to Costinot et al. (2016) and described in Felix (2022). This strategy requires shocks at the industry and location-levels to estimate θ_g , λ_g , not provided by the small MFA shock. I thus use variation due to Brazil’s trade liberalization in the 1990s.

Concentration To set the stage for Decomposition #2, Figure 8 plots the female/male ratio of within-industry concentration (HHI_{gk}) and what I term concentration in the industry ($s_{gk}HHI_{gk}$) across Brazil’s twelve largest industries by wage bill share. Together, these industries constitute over 60% of the formal sector wage bill. Across industries, within-industry HHI_{gk} is typically higher among men or equal across gender. However, women exhibit higher cross-industry concentration ($s_{gk}HHI_{gk}$) than men. Therefore, women’s disproportionate employment in fewer industries (Fact 3 of Section 3) spurs a higher model-consistent measure of concentration in women’s labor market. Together these patterns suggest, as substantiated below, that higher cross-(as opposed to within-) industry concentration drives gender differences in monopsony power.

Decomposition Table 10 (left panel) decomposes the contribution of within-industry concentration and cross-industry concentration to the monopsony-induced gender wage gap by progressively calculating the markdown from adding $\left(\frac{1}{\theta_g} - \frac{1}{\eta_g}\right) HHI_{gk}$ and $\left(\frac{1}{\lambda_g} - \frac{1}{\theta_g}\right) s_{gk}HHI_{gk}$ to the markdown generated by match-specific preferences alone $\left(1 + \frac{1}{\eta_g}\right)$. As noted above, different η_g generate

a 10% gender wage gap, with women paid 69% and men 79% the value of their marginal product.

Gender differences in within-industry concentration *reduce* this monopsony-induced gender wage gap by 2%, to 8%. Adding the $\left(\frac{1}{\theta_g} - \frac{1}{\eta_g}\right) HHI_{gk}$ term reduces women’s markdown to 66% and men’s to 74%. The decline in the gender gap reflects men’s higher within-industry concentration ($HHI_m = 0.11$, $HHI_f = 0.08$). In addition, it reflects higher monopsony power over men for the same HHI_{gk} due to their relatively higher captivity within industries: $\left(\frac{1}{\theta_m} - \frac{1}{\eta_m}\right) = 0.89$ and $\left(\frac{1}{\theta_f} - \frac{1}{\eta_f}\right) = 0.67$.³⁷

By contrast, women’s higher concentration in the textile industry adds 10% to the gender wage gap. Adding the $\left(\frac{1}{\lambda_g} - \frac{1}{\theta_g}\right) s_{gkr} HHI_{gkr}$ term reduces women’s markdown to 55% and men’s to 73%.³⁸ Together these findings demonstrate a prominent role for concentration in (but not within) the textile industry as driving gender differences in monopsony power.

Monopsony in the economy I next study a natural follow-up question: might employers in other industries possess disproportionate monopsony power over men? The answer is no, as portended by Figure 8. Table 10 (Panel B) shows that, akin to the textile industry, gender differences in within-industry concentration subtract 3pp from the economy-wide monopsony-induced gender wage gap, whereas women’s higher concentration in fewer industries adds 12pp. To generate economy-wide markdowns I combine η_g , θ_g , and λ_g estimates with a weighted sum of within-industry concentration $\sum_k s_k^g \sum_r s_{gk}^r HHI_{gkr}$ and its industry-size-weighted counterpart $\sum_k s_k^g \sum_r s_{gk}^r s_{gkr} HHI_{gkr}$. I take the outer sum over industries with weights s_k^g denoting the share of the group-specific wage bill in industry k . The weighted sum estimates the average markdown among all same-gender workers (Appendix A). Through the lens of the model, thus, women’s disproportionate employment in fewer industries confers employers with higher monopsony power over the average woman.

In sum, both match-specific reasons and concentration contribute to gender differences in monopsony power. Concentration is entirely driven by women’s disproportionate employment in the textile industry. Together these findings highlight the importance of market definition in diagnosing differential monopsony by gender. They show that traditional concentration measures, the within-industry or occupation HHIs (e.g. as measured in Azar et al. (2022), Berger et al. (2022), Felix (2022), Rinz (2022)) can misdiagnose gender differences in monopsony power by failing to account for women’s concentration in fewer industries.

5.3 What drives concentration in textiles? (non-wage amenities vs comparative advantage)

I next turn to probing the source of textile’s prominence in women’s labor market. Through the lens of the model, but also intuitively, this high share reflects either women’s relatively higher skills

³⁷The gender difference in inverse elasticities stems from η as opposed to θ differences.

³⁸Usefully, the 18pp model-generated gender wage gap exactly matches the non-targeted monopsony-induced gender gap estimated using heterogeneous elasticities in Section 5.1.

in the textile industry (comparative advantage) or its provision of relatively desirable non-wage amenities. I estimate both, gender-specific schedules of comparative advantage by industry, and gender-specific amenities by sector. I then evaluate their importance in driving the textile industry's high share.

Non-wage amenities Relatively desirable amenities raise market shares and therefore monopsony power. I estimate gender-specific non-wage amenities across industries by first constructing industry-specific wage indices from observed wages in RAIS and then inferring amenities given this wage index and industry size (the group-specific wage bill share in that industry). I employ the following three-step procedure and highlight objects computed at each stage in red:

Step 1. Estimate the wage index of industry k : $s_{gj} = \frac{(w_{gj})^{1+\eta_g}}{\sum_{j'} (w_{gj'})^{1+\eta_g}} \forall j$; re-arrange, take logs, and sum over all j : $W_{gk} = \tilde{w}_{gk} \tilde{s}_{gk}^{\frac{-1}{1+\eta_g}}$ where the tilde denote a geometric mean. I first calculate industry-region-specific wage indices and aggregate to the industry.³⁹

Step 2. Estimate men's amenities in k , normalizing $a_{m,txt} = 1$: $\frac{s_{mk}}{s_{m,txt}} = \frac{a_{mk}^{1+\theta_m} \frac{W_{mk}^{1+\theta_m}}{W_m^{1+\theta_m}}}{\frac{W_{m,txt}^{1+\theta_m}}{W_m^{1+\theta_m}}}$.

Step 3. Estimate women's amenities in k : $\frac{s_{fk}}{s_{mk}} = \frac{a_{fk}^{1+\theta_f} \frac{W_{fk}^{1+\theta_f}}{W_f^{1+\theta_f}}}{a_{mk}^{1+\theta_m} \frac{W_{mk}^{1+\theta_m}}{W_m^{1+\theta_m}}}$.

Figure 9 evidences the relatively prominent role of non-wage amenities compared to wages in drawing women to the textile industry. It plots the female/male ratio of industry-specific wage and amenity indices across Brazil's fifteen large industries by wage bill share (comprising 70% of the total wage bill). While all industries exhibit gender wage gaps, the standard deviation of these gaps (0.18) is substantially lower than of amenity gaps (0.35). Put another way, while wages alone would predict an identical female share in the textile and auto manufacturing industries, textiles employ 11% but automobiles only 2%.

Comparative advantage Relatively high productivity in textiles also raises compensation, shares, and, consequently, monopsony power. Consider a Cobb-Douglas production function $Y_j = z_j K_j^{\alpha_{k1}} l_j^{\alpha_{k2}}$ with labor a CES aggregation of male and female workers $l_j = [\beta_k f_j^\sigma + m_j^\sigma]^{\frac{1}{\sigma}}$. β_k represents women's relative productivity in k .⁴⁰ A high value raises k 's wage or amenity index, thereby raising its relative share. The marginal revenue product of female labor at j is:

$$mrpl_{jf} = \underbrace{\beta_k}_{\text{common to industry}} \underbrace{\sigma \alpha_{k2} z_j K_j^{\alpha_{k1}} l_j^{\alpha_{k2}-1} [\beta_k f_j^\sigma + m_j^\sigma]^{\frac{1}{\sigma}-1} f_j^{\sigma-1}}_{t_j}$$

β_k therefore aggregates to the industry-wide wage or amenity index, showing wages here:

³⁹Full derivation in Appendix A.

⁴⁰A large literature studying the consequences of changes in the relative supply of high and low skilled labor on wage inequality uses a similar functional form, as illustrated in Katz and Murphy (1992)

$$w_{gj} = \beta_k \frac{\epsilon_{gj}}{\epsilon_{gj} + 1} t_j$$

$$\sum_j (w_j^{1+\eta_g})^{\frac{1}{1+\eta_g}} = \beta_k \left[\sum_j \tilde{w}_j^{1+\eta_g} \right]^{\frac{1}{1+\eta_g}}$$

I estimate $\hat{\beta}_k$ across five key industries: the manufacturing of textile and clothing products, food and beverages, automobiles, machinery, and metal products using standard production function techniques (Akerberg et al., 2015). I choose these industries because each has at least 250 establishments observed in a panel between 2000 - 2010, as needed to obtain consistent estimates and tight standard errors (Demirer, 2020). Together they employ 10.4% of formal sector workers. A caveat to my conclusions is that, pending a data application in Brazil, I estimate production functions in two other developing countries: India and Chile. I obtain similar estimates across the two and use the mean value of $\hat{\beta}_k$. Appendix C.3 describes the data and procedure in detail. A second caveat is that the presence of market power potentially biases productivity estimates (see footnote).⁴¹

Figure 10 shows that gender productivity gaps predict a much smaller difference in female/male shares across industries than observed. They predict a ratio of female to male shares equal to 1.24 in textiles, 1.08 in food and beverages, 0.55 in automobiles, 0.58 in machinery, and 0.39 in the manufacturing of metal products. In reality, these shares are respectively 6, 0.76, 0.27, 0.2, and 0.18. Thus, the textile industry attracts a substantially higher share of women than predicted by gender differences in productivity alone, implying a prominent role for non-wage amenities in drawing women to the industry.

5.4 Summary

The findings of this section can be summarized in three takeaways. First, I find that a match-specific preference for their current employer and the concentration of good jobs for women in the textile industry both contribute to gender differences in monopsony power. Second, I find that the traditional measure of concentration, the within-industry Herfindahl-Hirshmann Index, can misdiagnose gender differences in labor market concentration. While men’s markets within industries typically exhibit higher or equivalent levels of concentration as women’s, women’s labor market is more concentrated because they are employed in fewer industries. Finally, I show that this concentration largely reflects desirable non-wage amenities in textiles and not gender differences in comparative advantage. These amenities could be those drawing women to the textile industry

⁴¹Market power in inputs potentially violates the “scalar unobservables” assumption of the proxy variables method, preventing me from inverting input demand to recover productivity shocks (Bond et al. (2021), De Loecker and Syverson (2021)). However, two facts mitigate the concern that bias in $\hat{\beta}_k$ leads to a false conclusion that amenities and not productivity drive women to textiles. First, Yeh et al. (2022) show using simulations that this method recovers the true parameters in its 95% confidence interval and delivers tight standard errors. Second, gender gaps in wages plus contracted amenities predict a much smaller difference in the relative female/male share in industries than observed — for example, they predict similar ratios in textiles and the manufacturing of metal products. In reality, the female share/male share in the former is 6 and in the latter is 0.2.

(e.g., female-focused contracts, female coworkers, safe work environments), or those driving them away from non-textile employers (discrimination, a lack of safety, or gender norms). In the following section I examine the counterfactual effect of altering sectoral gaps in amenities, both in residuals obtained from the model and in their real-world amenity analogs (such as maternity leave, flexibility, and childcare) observed in collective bargaining contracts.

6 Counterfactuals

This section studies the prospect for policy to remedy the monopsony-induced gender wage gap via a series of counterfactuals amending gender gaps in amenities, skills, and safety. To benchmark the contribution of sectoral differences in gender-specific amenities and comparative advantage, I first calculate the effect of leveling each. In two subsequent policy-relevant counterfactuals, I study the effect of improving female-centric amenities in contracts and safety.

6.1 Preliminaries

The counterfactual exercises require four key ingredients. First, I parameterize the labor supply system with estimates of the three elasticities of substitution η_g , θ_g , λ_g . Second, I assume that employers compete a la Cournot, choosing employment to maximize profits while taking as given the employment at other firms. Third, I assume a Cobb-Douglas production function with labor a CES aggregation of male and female workers (as above). I assume that capital is supplied in a competitive market.⁴² I calibrate other necessary parameters (summarized in Appendix Table C10). I calibrate productivity dispersion to match the observed firm size distribution in the textile industry. I calibrate the distribution of firm frequency in each industry and geography to a Pareto distribution with shape, scale, and location parameters set to minimize distance with the first three moments of the economy-wide distribution. Finally, I calibrate the elasticity of substitution between male and female workers (σ) from Gallen (2018).

I also make a few simplifying assumptions. I assume away establishment-specific amenities a_{gj} that deviate from industry norms. I also assume geographic immobility, $\lambda_g = 0$. I restrict counterfactuals to the five major industries in which I estimate gender-specific productivity (β_k above): the manufacturing of textiles and clothing, food and beverage products, metal products, machinery, and automobiles. Together these industries employ over 10.3% of the formal labor force (7.6% of women and 10.4% of men).

6.2 Mechanics

For each counterfactual, I quantify its effect on the monopsony-induced gender wage gap. Unlike estimation, counterfactuals require solving the model. I solve for two fixed points among

⁴²The firm’s optimization problem can be recast as optimizing over labor alone by substituting in its optimal capital demand (Berger et al., 2022).

women: an upper-level industry share and lower-level within-industry share. For counterfactuals that alter amenities or productivity, I start from the industry shares predicted by the new amenity/productivity distribution and old industry-specific wage part of the wage index (equation 5). As women substitute from textiles to amenity/productivity-enhancing industries, the textile industry’s share declines. Its employers face more elastic labor supply (equation 6), which causes them to alter wages and workers to reallocate across textile employers (lower-level fixed point). This yields a new wage index for each industry and upper-level industry share; I solve until a fixed point in this upper-level share. The last counterfactual increases mobility (η_f) to its value in the 75th percentile safest municipality, causing workers to reallocate across employers. As wage indices change, so too do industry shares. As before, I solve for the upper and lower level fixed points.

6.3 Counterfactuals

Level all gender-specific amenities I first study the effect of eliminating gender gaps in non-wage amenities by setting the ratio $a_{fk'}^{1+\theta_f}/a_{mk'}^{1+\theta_m} = 1$ across non-textile industries. Figure 11 shows a 8pp decline in the monopsony-induced gender wage gap. Women flock to amenity-improving non-textile industries, which lowers the textile share to 3.4% (from 10.5%) and increases the elasticity of labor supply to textile firms that are now smaller in women’s labor market. Women’s wages increase for two reasons: (i) they earn a higher share of marginal product, and (ii) they reallocate to larger, higher-wage employers whose monopsony wedge disproportionately falls. Simultaneously, men’s wages decline as their marginal product at large employers declines and they reallocate to smaller textile firms. The average woman’s markdown rises 6.6pp and the average man’s declines 1.4pp. In this way, leveling gender gaps in sector-specific amenities fully erodes the 8pp gender wage gap due to concentration.

Women in non-textile industries do not lose from their now-larger size and higher monopsony power. At the outset, each comprises 1-2% of women’s labor market and gains 1.4% following amenity improvements. The average woman’s markdown in these industries falls 3pp. However, women are better off in utility terms, as evidenced by net movements into these industries. Formally, utility gains are reflected in (1-10%) higher industry-specific wage indices.⁴³ Just as changes to the price index measure changes to consumer welfare from improving product varieties, changes to the wage index, i.e., how much workers are paid to work one disutility-weighted hour, measures changes to worker welfare from amenity improvements (Corradini et al., 2022).

How costly are amenity improvements? This paper does not speak to direct costs, but evidences offsetting gains. First, I find that greater equity is coupled with higher efficiency. Because large and productive textile employers disproportionately lose monopsony power, women reallocate to them from smaller ones, which raises productive efficiency in textiles. Aggregate productivity, defined as the ratio of realized to potential production $\Omega = Y_{realized}/Y_{potential}$, increases by 1.1pp.⁴⁴ Second,

⁴³ $\bar{W}_{kg} = a_{kg}(\sum_{j \in k} w_{gjk}^{-1+\eta_g})^{\frac{1}{1+\eta_g}}$

⁴⁴ $Y_{potential}$ is defined as the output if employers did not exercise their monopsony power and instead paid workers their marginal product.

recent work by Corradini et al. (2022) suggests that Brazilian employers inefficiently underprovide female-friendly amenities: they find no tradeoffs in employment, wages, or profits when a female-focused union reform institutes large gains (0.13SD) in these amenities among 20% of the formal labor market. Finally, I show below that non-contracted (likely unpaid) amenities importantly drive women to textiles. On the one hand, these amenities may be resilient to change. On the other hand, changing them may not entail expensive investment.

Level gender-specific comparative advantage I next study the effect of leveling gender gaps in productivity by equalizing a β_k fraction of each industry’s amenity advantage. I focus on amenity instead of wage gaps because the former explain textile’s prominence in women’s labor market (Figure 9). While this choice biases against finding an important role for amenities, they remain important.

Figure 11 shows a 4.12pp decline in the monopsony-induced gender wage gap from eliminating gender productivity gaps. The textile industry share falls from 10.5% to 8.7%. Women’s wages increase 3.4pp and men’s decline 0.72pp. Productive efficiency increases 0.3pp. Leveling productivity thus achieves roughly half the effect of leveling gender amenity gaps.

Level gender-specific amenities in contracts Because a policymaker would likely find it difficult to alter amenity “residuals” that she cannot observe, I study a policy-relevant counterfactual leveling female-focused amenities in collective bargaining contracts. I define amenities predicting women’s revealed preference value for an employer as female-focused. These include twenty (of 137) provisions governing maternity leave, flexibility, childcare, and overtime, among others (Table 2 in Corradini et al. (2022)). Textile contracts feature 5.8 female-focused amenities on average compared to 4.37 among non-textile employers. For the counterfactual, I let women vote-with-their-feet toward amenity-improving employers and calibrate this rate with Corradini et al.’s quasi-experimental estimate.

Figure 11 shows that equalizing gender gaps in contracted amenities has roughly half the effect of leveling gaps in all amenities (4.54pp). The textile industry’s share falls to 8.1%. Women’s wages increase 3.69pp and men’s decline 0.85pp. Productive efficiency increases 0.25pp. Because Corradini et al.’s three-year estimate may underestimate the longer-term draw of amenities, I assess robustness to doubling the rate: the gender wage gap falls 5.5pp. Overall, improving observed and actionable contracted amenities substantially, but does not fully, erode the monopsony-induced gender wage gap due to concentration.

Improve neighborhood safety Finally, we could make the streets safer. Figure 11 shows a 3.64pp decline in the monopsony-induced gender wage gap from improving safety across all municipalities to the 75th percentile. Improving safety raises η_g from 2.19 to 2.278, enhancing cross-employer substitution to high-wage/amenity textile employers. Overall, women’s wages rise 2.96pp and men’s decline 0.68pp.

Summary I have two main conclusions. First, the textile industry attracts a much higher relative share of women than predicted by productivity differences alone and must thus offer women higher non-wage amenities. These amenities are not fully accounted for by contracted amenities, and would instead comprise gender norms, safety in the workplace, or employer discrimination.

Second, improving non-textile jobs to be as desirable for women as textile jobs would substantially erode the monopsony-induced gender wage gap. The largest gains would accrue from eliminating differences in non-wage amenities, with smaller gains from upskilling women. Improving safety on the streets, by enhancing women’s mobility, also reduces the gender gap.

7 Conclusion

This paper highlights the role of an understudied market failure in developing countries, differential monopsony by gender, in effecting a potentially large share of the gender wage gap. Using quasi-experimental firm-level variation in wages, I document substantially lower separation elasticities among women than men, with resulting monopsony power generating an 18pp gender wage gap among equally productive workers. I show that differential monopsony by gender has two intuitive sources: women are tethered to their current employer even when their labor market abounds with opportunity; in addition, they are tethered to relatively good jobs in the textile industry. Combined with a model, I estimate positive spillovers from improving non-traditional work environments for women via a reduction in monopsony in women’s current jobs. Whereas leveling sectoral amenity gaps erodes 8pp of the monopsony-induced gender wage gap, upskilling women to level gender productivity gaps achieves only half this gain.

A surprising conclusion of my findings is that greater equity is coupled with efficiency. Reducing monopsony reallocates women from smaller to larger/more productive employers. This suggests that policies improving non-traditional workplaces, such as combating sexual harassment, instituting flexible work, or enhancing maternity protections, can be a key lever to remedy labor market distortions. While some of these policies (combating sexual harassment), are unambiguously good, others (increasing flexibility), entail costs. Weighing these costs against the estimated gains is a fruitful avenue for policy.

My findings also raise several new questions. First, do employers exercise their differential monopsony power over women? In ongoing work I study this question by examining the MFA shock’s differential pass-through to men and women’s wages.⁴⁵ Second, does concentration in fewer professions also spur differential monopsony by gender in other contexts? Just as in Brazil, women in other parts of the world are employed in fewer industries and occupations: in the United States, for example, 1 in 8 women is a teacher or nurse (Sokolova and Sorensen, 2021). Studying the consequences of such concentration for monopsony is an important area for future research.

⁴⁵Intuitively, in the nested CES and Cournot/Bertrand model, the pass-through of shocks to own marginal product declines with employer size when employers exercise monopsony power. The rate of decline therefore sheds light on its exercise. Qualitatively, as predicted, large employers exhibit lower pass-through to women than men’s wages (they are larger in women’s labor market).

Finally, my findings raise the puzzle of why exactly such few jobs are “women’s jobs” although women are similarly skilled in, and receive similar wages across, industries. Identifying the amenities/disamenities drawing women towards/away from industries (gender norms, a desire to work with other women, discrimination, or valuing other industry-specific amenities), and, to the extent that they misallocate women’s talent, identifying remedies, is an exciting area for future work.

References

- Akerberg, D. A., Caves, K., and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451.
- Almagro, M. and Manresa, E. (2021). Data-driven nests in discrete choice models. *Slides*.
- Altonji, J. G. and Blank, R. M. (1999). Race and gender in the labor market. *Handbook of labor economics*, 3:3143–3259.
- Amiti, M., Itskhoki, O., and Konings, J. (2019). International shocks, variable markups, and domestic prices. *The Review of Economic Studies*, 86(6):2356–2402.
- Anderson, S. P., De Palma, A., and Thisse, J.-F. (1992). *Discrete choice theory of product differentiation*. MIT press.
- Atkeson, A. and Burstein, A. (2008). Pricing-to-market, trade costs, and international relative prices. *American Economic Review*, 98(5):1998–2031.
- Autor, D. H., Dorn, D., Hanson, G. H., and Song, J. (2014). Trade adjustment: Worker-level evidence. *The Quarterly Journal of Economics*, 129(4):1799–1860.
- Azar, J., Marinescu, I., and Steinbaum, M. (2022). Labor market concentration. *Journal of Human Resources*, 57(S):S167–S199.
- Baqae, D. R. and Farhi, E. (2020). Productivity and misallocation in general equilibrium. *The Quarterly Journal of Economics*, 135(1):105–163.
- Barth, E. and Dale-Olsen, H. (2009). Monopsonistic discrimination, worker turnover, and the gender wage gap. *Labour Economics*, 16(5):589–597.
- Becker, G. S. (1962). *The economics of discrimination*. University of Chicago press.
- Berger, D., Herkenhoff, K., and Mongey, S. (2022). Labor market power. *American Economic Review*, 112(4):1147–93.
- Blau, F. D. and Kahn, L. M. (2017). The gender wage gap: Extent, trends, and explanations. *Journal of economic literature*, 55(3):789–865.
- Bond, S., Hashemi, A., Kaplan, G., and Zoch, P. (2021). Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data. *Journal of Monetary Economics*, 121:1–14.
- Brambilla, I., Khandelwal, A. K., and Schott, P. K. (2010). China’s experience under the multi-fiber arrangement (mfa) and the agreement on textiles and clothing (atc). In *China’s Growing Role in World Trade*, pages 345–387. University of Chicago Press.

- Burdett, K. and Mortensen, D. T. (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, pages 257–273.
- Caldwell, S. and Oehlsen, E. (2022). Gender differences in labor supply: Experimental evidence from the gig economy. *University of California, Berkeley Working Paper*.
- Card, D. (2022). Who set your wage? *American Economic Review*, 112(4):1075–90.
- 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., Cardoso, A. R., and Kline, P. (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly journal of economics*, 131(2):633–686.
- Corradini, V., Lagos, L., and Sharma, G. (2022). Collective bargaining for women: How unions can create female-friendly jobs.
- Costinot, A., Donaldson, D., and Smith, C. (2016). Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world. *Journal of Political Economy*, 124(1):205–248.
- De Chaisemartin, C. and d’Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–96.
- de Lima Amaral, E. F. (2013). Brazil: internal migration. *The encyclopedia of global human migration*.
- De Loecker, J. and Syverson, C. (2021). An industrial organization perspective on productivity. In *Handbook of Industrial Organization*, volume 4, pages 141–223. Elsevier.
- Demirer, M. (2020). Production function estimation with factor-augmenting technology: An application to markups. *Job Market Paper*.
- Diao, X. and Somwaru, A. (2001). Impact of the mfa phase-out on the world economy an intertemporal, globalgeneral equilibrium analysis. Technical report.
- Dix-Carneiro, R. and Kovak, B. K. (2015). Trade liberalization and the skill premium: A local labor markets approach. *American Economic Review*, 105(5):551–57.
- Dix-Carneiro, R., Soares, R. R., and Ulyssea, G. (2018). Economic shocks and crime: Evidence from the brazilian trade liberalization. *American Economic Journal: Applied Economics*, 10(4):158–95.
- Dube, A., Manning, A., and Naidu, S. (2020). Monopsony and employer mis-optimization explain why wages bunch at round numbers. Technical report, National Bureau of Economic Research.

- Felix, M. (2022). Trade, labor market concentration, and wages. *Job Market Paper*.
- Field, E., Pande, R., Rigol, N., Schaner, S., and Troyer Moore, C. (2021). On her own account: How strengthening women’s financial control impacts labor supply and gender norms. *American Economic Review*, 111(7):2342–75.
- Fletcher, E., Pande, R., and Moore, C. M. T. (2017). Women and work in india: Descriptive evidence and a review of potential policies.
- Forgy, E. W. (1965). Cluster analysis of multivariate data: efficiency versus interpretability of classifications. *biometrics*, 21:768–769.
- Gallen, Y. (2018). Motherhood and the gender productivity gap. *Becker Friedman Institute for Research in Economics Working Paper*, (2018-41).
- Hazell, J. and Taska, B. (2020). Downward rigidity in the wage for new hires. *Working paper*.
- Hildreth, A. K. (1998). Rent-sharing and wages: product demand or technology driven premia? *Economics of Innovation and New Technology*, 5(2-4):199–226.
- Hirsch, B., Schank, T., and Schnabel, C. (2010). Differences in labor supply to monopsonistic firms and the gender pay gap: An empirical analysis using linked employer-employee data from germany. *Journal of Labor Economics*, 28(2):291–330.
- Hoek, J. (2006). Life cycle effects of job displacement in brazil.
- Hsieh, C.-T., Hurst, E., Jones, C. I., and Klenow, P. J. (2019). The allocation of talent and us economic growth. *Econometrica*, 87(5):1439–1474.
- ILO (2018). Global wage report 2018/19: What lies behind gender pay gaps.
- ILO (2022). The state of the apparel and footwear industry: employment, automation and their gender dimensions. *ILO and European Commission, Geneva, Switzerland and Brussels, Belgium*.
- ILO and Gallup (2017). Towards a better future for women and work: Voices of women and men. *Gallup Inc. and the International Labour Organization, Washington*.
- Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings losses of displaced workers. *The American economic review*, pages 685–709.
- Jäger, S. and Heining, J. (2022). How substitutable are workers? evidence from worker deaths. Technical report, National Bureau of Economic Research.
- Kaplan, D. S., Gonzalez, G. M., Robertson, R., Menezes-Filho, N., and Arias, O. (2005). What happens to wages after displacement?[with comments]. *Economia*, 5(2):197–242.

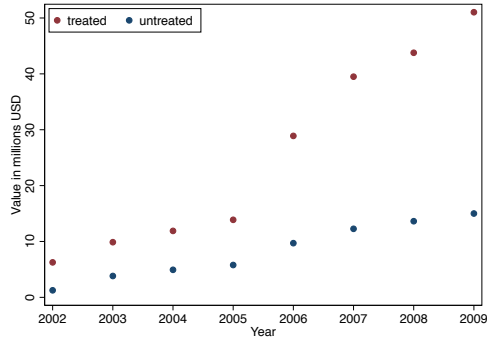
- Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: supply and demand factors. *The quarterly journal of economics*, 107(1):35–78.
- Khandelwal, A. K., Schott, P. K., and Wei, S.-J. (2013). Trade liberalization and embedded institutional reform: Evidence from chinese exporters. *American Economic Review*, 103(6):2169–95.
- Lagos, L. (2019). Labor market institutions and the composition of firm compensation: Evidence from brazilian collective bargaining. *Job Market Paper*.
- Lloyd, S. (1982). Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137.
- Macaluso, C. (2022). Skill remoteness and post-layoff labor market outcomes. *Working paper*.
- Manning, A. (2003). Monopsony in motion. In *Monopsony in Motion*. Princeton University Press.
- Manning, A. (2011). Chapter 11 - imperfect competition in the labor market. volume 4 of *Handbook of Labor Economics*, pages 973–1041. Elsevier.
- Manning, A. (2021). Monopsony in labor markets: A review. *ILR Review*, 74(1):3–26.
- Matsudaira, J. D. (2014). Monopsony in the low-wage labor market? evidence from minimum nurse staffing regulations. *Review of Economics and Statistics*, 96(1):92–102.
- McFadden, D. (1978). Modelling the choice of residential location. *Transportation Research Record*.
- McKelway, M. (2018). Women’s self-efficacy and women’s employment: Experimental evidence from india. *Unpublished Working Paper*.
- Morchio, I. and Moser, C. (2021). The gender pay gap: Micro sources and macro consequences.
- Nimczik, J. S. (2020). Job mobility networks and data-driven labor markets. Technical report, Working Paper.
- Organization, I. L. (2018). India wage report: Wage policies for decent work and inclusive growth. *India: International Labor Organization*.
- Ransom, M. R. and Sims, D. P. (2010). Estimating the firm’s labor supply curve in a “new monopsony” framework: Schoolteachers in missouri. *Journal of Labor Economics*, 28(2):331–355.
- Rinz, K. (2022). Labor market concentration, earnings, and inequality. *Journal of Human Resources*, 57(S):S251–S283.
- Rivera, S. A., Agama, L.-A., and Dean, J. (2003). Africa beyond 2005: Understanding the impact of eliminating ntbs and tariffs on textiles and clothing. *International Trade Commission, mimeo*.

- Robinson, J. (1933). *The economics of imperfect competition*. London: MacMillan Press.
- Schubert, G., Stansbury, A., and Taska, B. (2021). Employer concentration and outside options. *Working paper*.
- Sokolova, A. and Sorensen, T. (2021). Monopsony in labor markets: A meta-analysis. *ILR Review*, 74(1):27–55.
- Topalova, P. (2010). Factor immobility and regional impacts of trade liberalization: Evidence on poverty from india. *American Economic Journal: Applied Economics*, 2(4):1–41.
- Tortarolo, D. and Zarate, R. D. (2018). Measuring imperfect competition in product and labor markets. an empirical analysis using firm-level production data.
- Van Reenen, J. (1997). Employment and technological innovation: evidence from uk manufacturing firms. *Journal of labor economics*, 15(2):255–284.
- Webber, D. A. (2016). Firm-level monopsony and the gender pay gap. *Industrial Relations: A Journal of Economy and Society*, 55(2):323–345.
- Yeh, C., Macaluso, C., and Hershbein, B. (2022). Monopsony in the us labor market. *American Economic Review*, 112(4):1147–93.

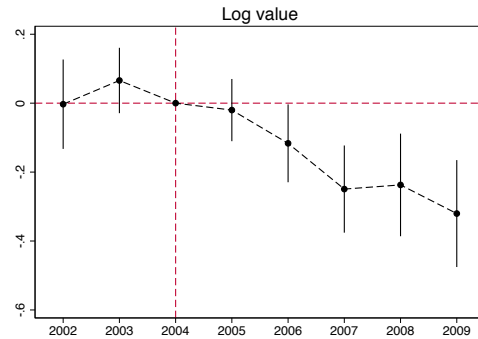
Figures

Figure 1: Effect of MFA's end on Chinese and Brazilian export values

(a) Chinese exports (value in millions USD)




(b) Brazilian exports (firm-level DiD, log value)



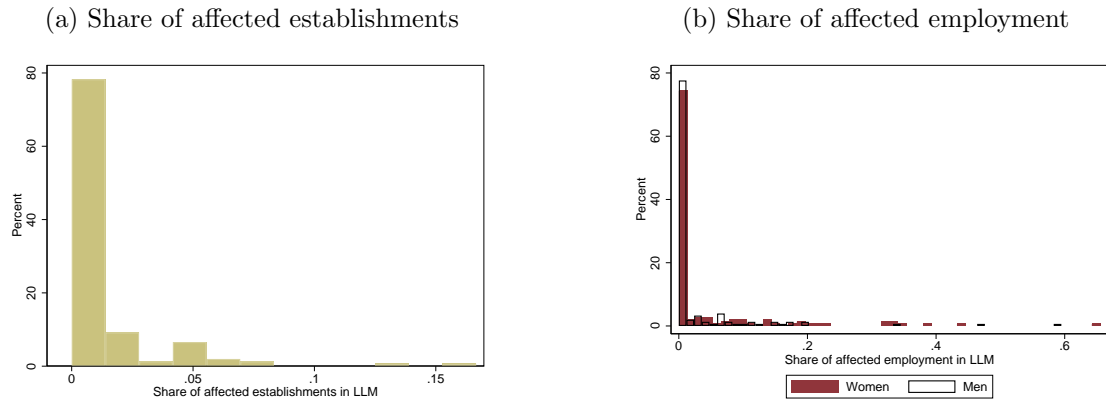
Notes: This figure plots the effect of the end of the MFA on Chinese (Panel a) and Brazilian (Panel b) exports of treated and comparison products. Treated products are those whose Chinese exporters faced a binding quota in from the US in 2004. Comparison products comprise all other products in the textile and clothing manufacturing industries. Panel a plots the total value of Chinese exports of the two sets of products between 2001 and 2008 (in millions of USD). Panel b plots an establishment-level DiD event study comparing the log of export value at Brazilian establishments exporting treated versus comparison products. The x-axis plots exports in the preceding year i.e. from year $t - 1$ to year t .

Figure 2: Worker characteristics



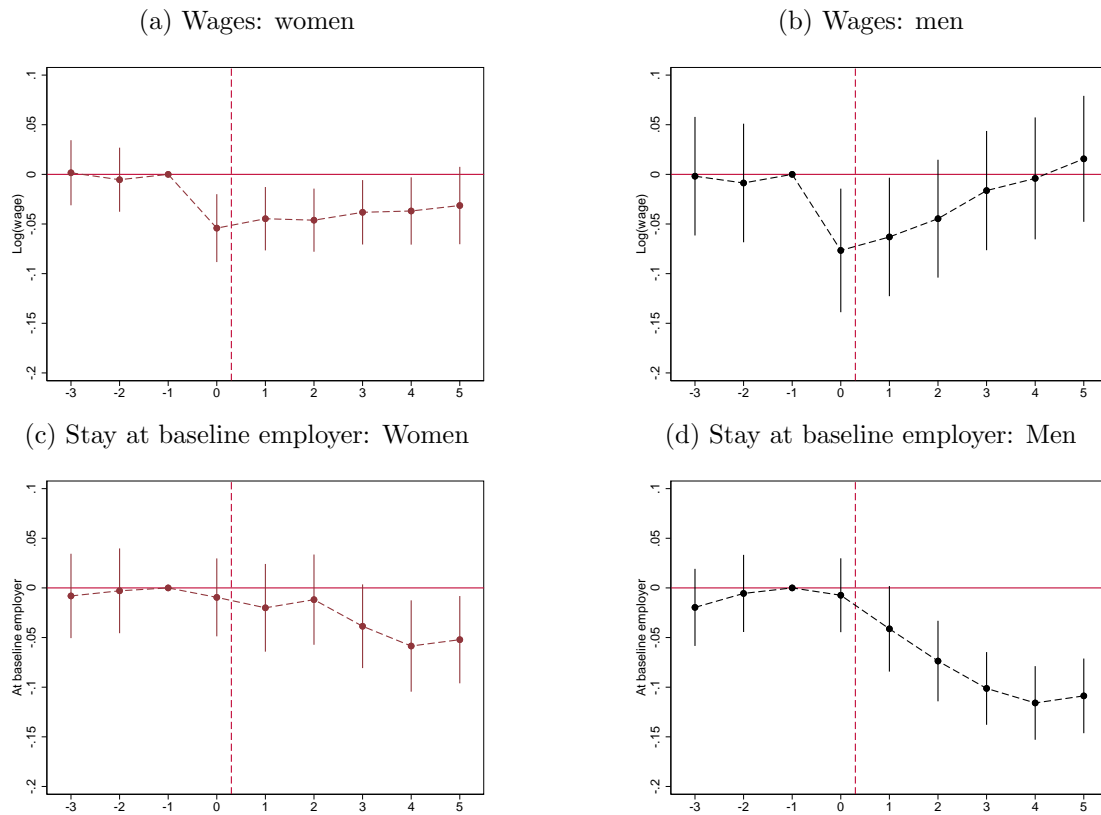
Notes: This figure plots baseline (2004) descriptive statistics of treated and comparison establishments. Panels a and b plot the distribution of 4-digit occupations for women and men at treated and comparison establishments. Panels c and d plot the correlation between the share of pre-period transitions to a given 4-digit occupation from treated establishments (x-axis) and comparison establishments (y-axis). Panels e and f plot the distribution of geography (microregions) for women and men at treated and comparison establishments.

Figure 3: Firm-specific shock in textile and clothing labor market



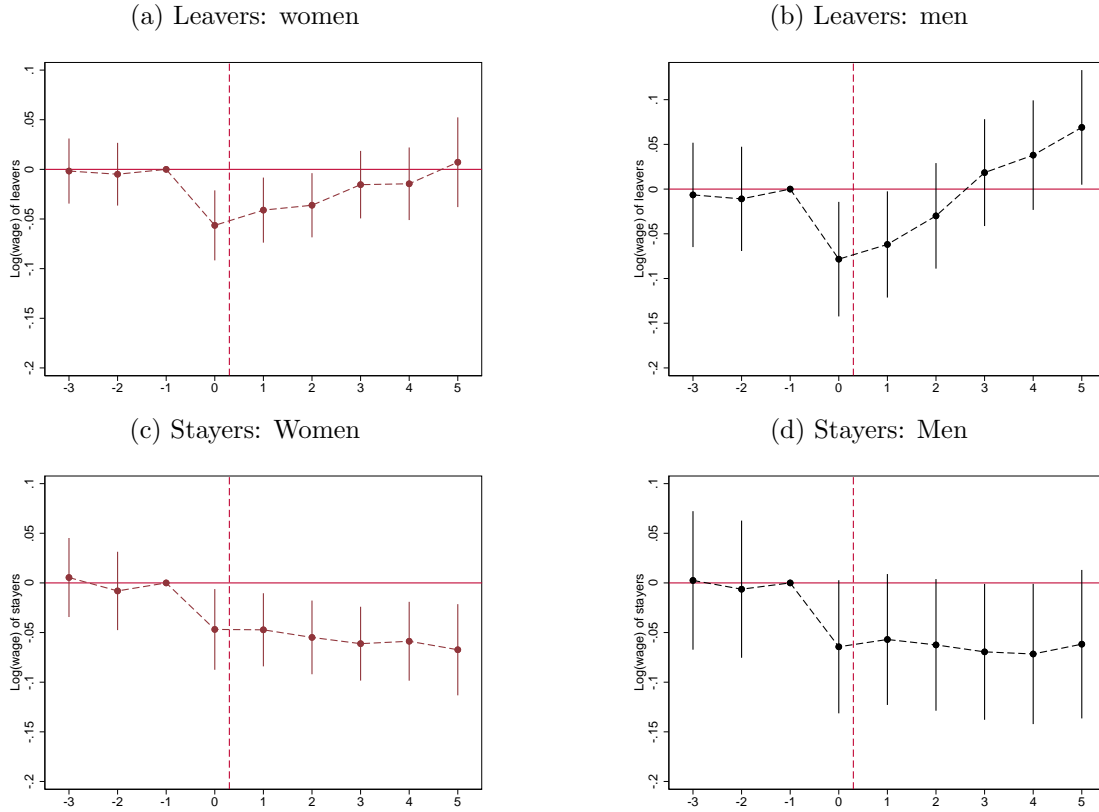
Notes: This figure shows that the MFA shock affected a small share of employment and establishments in the textile and clothing manufacturing industries. Panel a plots the share of treated establishments in a given geography (microregion). Panel b plots the share of treated male and female employment in a geography. The mean share of establishments is 2% and of employment is 10%.

Figure 4: Effect of the MFA's end on wages and employment



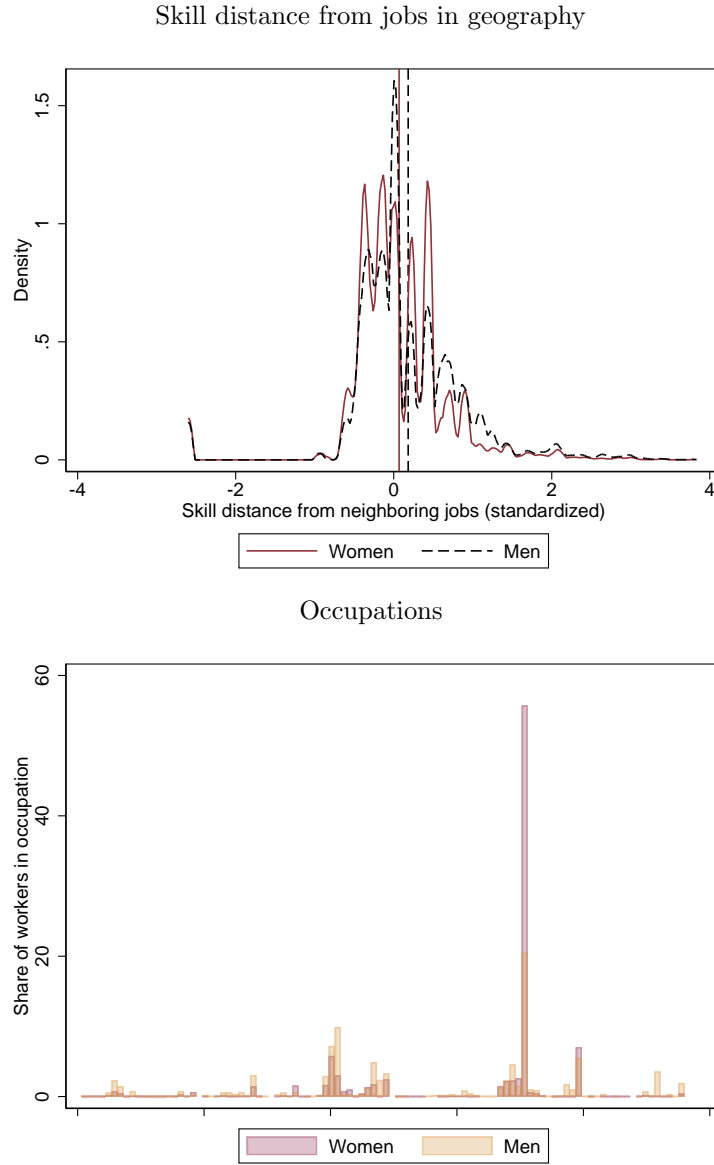
Notes: These figures plot estimates of the δ_t coefficients for $t \in [-3, 5]$ (with $t = -1$ omitted) from a DiD specification comparing treated men and women with their comparison counterparts. The outcome in Panels a and b is log wages and in Panels c and d is retention at one's baseline employer. $t = 0$ corresponds with 2005. MFA quotas were lifted in January 2005. Confidence intervals at a 95% level are reported. Standard errors are clustered at the establishment level.

Figure 5: Wages of stayers and leavers




Notes: These figures plot estimates of the δ_t coefficients for $t \in [-3, 5]$ (with $t = -1$ omitted) from the DiD specification on the log wages of leavers (Panels a and b) and stayers (Panels c and d), comparing them with all same-gender workers in the comparison group. $t = 0$ corresponds with 2005. MFA quotas were lifted in January 2005. Stayers are defined as those who remain at their baseline employer four years following the MFA's end. Standard errors are clustered at the establishment level.

Figure 6: Men and women’s skill transferability and occupations



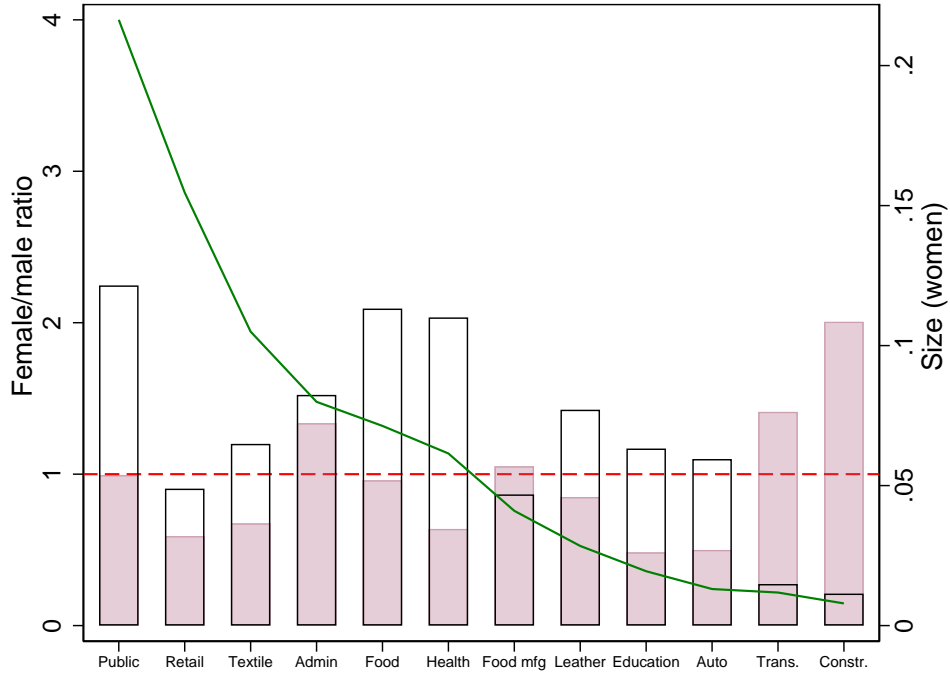
Notes: The top panel plots kernel density estimates of the distribution of skill distance from other jobs in a worker’s geography, separately for women (maroon, solid line) and men (black, dashed line) at treated employers. Vertical lines indicate the mean distance respectively for male and female. The skill level (1-8) on each of 35 skills required to do an occupation comes from O*NET. Skill distance is the weighted average of the distance (L1-norm) between a worker’s current occupation and all other occupations in her geography, with weights being that occupation’s share among jobs in the geography. I standardize this remoteness measure among all workers to have mean 0 and standard deviation 1. The bottom panel plots the distribution of occupations among male and female workers at treated employers. Occupations are 4-digit occupations.

Figure 7: Share of women’s and men’s wage bill in top industries



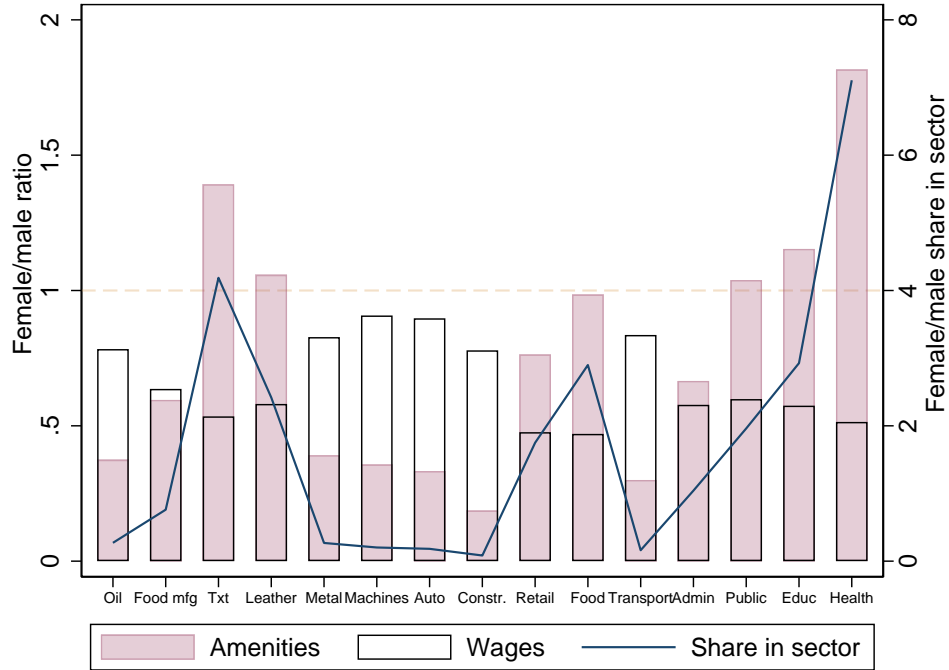
Notes: This figure plots the share of women and men’s overall wage bill in the textile industry as well as their two, five, and ten most important industries.

Figure 8: Within-industry HHI and industry-share-weighted HHI




Notes: This figure plots the ratio of female/male within-industry HHIs (term a) and concentration in industry (term b) in Proposition 1 for the twelve largest employers of workers who are high school dropouts. The line plots the share of women's wage bill in a given industry. The main message is that women's within-industry concentration resembles men's, but that women are disproportionately employed in a small number of industries raising overall concentration.

Figure 9: Amenities and wages across industries



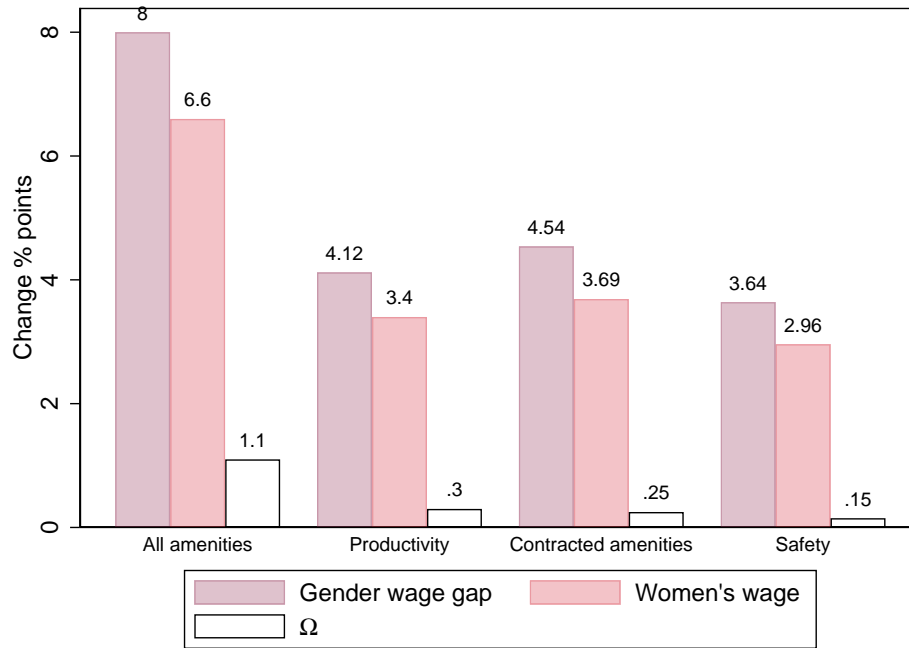
Notes: This figure plots the ratio of female/male amenity and wage indices across the fifteen largest 2-digit industries (sectors) by share of wage bill in Brazil. Together, these industries employ over 70% of workers. I observe wages across all formal sector employers in the RAIS data and aggregate them to a wage index at the sector level for men and women. I then use the discrete choice assumption of the model to infer amenities. Intuitively, a sector attracts many workers if it offers either high wages or high amenities. Appendix C.2 provides detail on constructing wage indices and inferring amenities.

Figure 10: Actual and predicted ratio of women to men across industries



Notes: The pink bars plot the actual ratio of the wage bill share of women’s labor market and men’s labor market in an industry. The black bars plot the implied ratio based on gender differences in productivity alone. I assume productivity differences are entirely compensated through non-wage amenities (using wages instead predicts an even smaller share of women in textiles). Appendix C.3 describes the procedure for estimating men and women’s relative productivity across industries.

Figure 11: Counterfactual exercises to alleviate gender differences in monopsony power



Notes: This figure plots the change in the gender wage gap, the average woman's wage, and productive efficiency (Ω) in the textile industry from various policies. The first equalizes all non-wage amenities across industries; the second equalizes gender productivity gaps across industries; the third equalizes non-wage amenities included in collective bargaining agreements; the final improves safety to the level of the 75th percentile municipality, thereby making it easier for women to substitute across employers within the industry η_g . Section 6 describes details.

Table 1: Examples of treated and comparison HS codes

Comparison		Treated	
Code	Product name	Code	Product name
620341	Men's or boys' trousers, bib and brace overalls, breeches and shorts of wool or fine animal hair	620461	Women's or girls' trousers, bib and brace overalls of wool or fine animal hair
620510	Men's or boys' shirts of cotton	610910	Men's or boys' t-shirts of cotton
620449	Women's or girls' dresses of other txt materials	620463	Women's or girls' trousers, bib and brace overalls of synthetic fibers
620451	Women's or girls' skirts of wool or fine animal hair	611011	Sweaters, pullovers, sweatshirts, waistcoats of wool or fine animal hair
610839	Women's or girls' nightdresses and pajamas of other textile materials	620811	Women's or girls' slips and petticoats of man-made fibers
621132	Men's or boys' track suits of cotton	620520	Men's or boys' dress shirts of cotton
621142	Women's or girls' track suits of cotton	620821	Women's or girls' nightdresses and pajamas of cotton
621050	Women's or girls' other garments of man-made fibers	620530	Men's or boys' shirts of man-made fibers
610729	Men's or boys' nightshirts and pajamas of other txt materials	610329	Men's or boys' ensembles of other txt materials
620332	Men's or boys' suit type jackets and blazers of cotton	620412	Women's or girls' suit type jackets and blazers of cotton
620333	Men's or boys' suit type jackets and blazers of synthetic fibers	620433	Women's or girls' suit-type jackets and blazers of synthetic fibers
620339	Men's or boys' suit type jackets and blazers of other txt materials	620419	Women's or girls' suits or other txt materials
610791	Men's or boys' nightshirts of cotton	620711	Men's or boys' underpants and briefs of cotton
621710	Other made-up clothing parts/accessories	621790	Other made-up clothing parts/accessories
620891	Women's or girls' singlets, slips, petticoats of cotton	620322	Men's or boys' ensembles of cotton
621210	Women or girls' brassieres of man-made fibers	620530	Men's or boys' shirts of man-made fibers
620341	Men's or boys' trousers, bib and brace overalls, breeches and shorts of wool or fine animal hair	620342	Men's or boys trousers, bib and brace overalls, breeches and shorts of cotton

Notes: This table depicts example 6-digit HS codes for products whose exports from China to the US did not face a binding quota under the MFA's phase IV (comparison) and whose exports from China to the US did face a binding quota under the MFA's phase IV (treated). A binding quota is defined as one whose fill rate at baseline (2004) is over 85%.

Table 2: Descriptive statistics

	<i>Full sector</i>			<i>Sample</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Estab characteristics</i>						
				All	Treated	Control
Avg. employment	29.855			152.070	150.675	152.914
Avg. # products exported	X	X	X	1.720	1.681	1.743
No. of estabs	15971			751	283	468
	Women			Men		
	All	Treated	Control	All	Treated	Control
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel B: Worker characteristics (Sample)</i>						
Avg. wage (per month)	559.769	550.856	567.697	1096.006	1050.331	1114.961
Avg. hours (per week)	43.885	43.961	43.737	43.731	43.647	43.745
Avg. tenure (years)	4.316	4.105	4.504	5.243	4.661	5.243
Age (years)	33.542	33.182	34.251	32.898	31.223	33.188
Education						
<i>Less than HS</i>	0.749	0.758	0.732	0.677	0.666	0.679
<i>HS grad</i>	0.242	0.232	0.262	0.315	0.319	0.315
<i>More than HS</i>	0.009	0.010	0.006	0.008	0.015	0.006
No. of workers	51533	24260	27273	62672	18381	44291

Notes: This table presents descriptive statistics for the sample of treated and comparison establishments and workers at baseline ($t = -1$). Treated establishments are those that were exporting treated products at baseline, i.e. those under binding quota in China under the MFA. Comparison establishments are those exporting other textile and clothing products. Panel A describes establishments. Panel B describes incumbent workers, i.e. those employed at treated or comparison establishments in ($t = -1$).

Table 3: Effects on wages and employment

	<i>All workers</i>						<i>Tailors</i>	
	Log earn (1)	Retention (2)	Log earn (3)	Retention (4)	Log earn (5)	Retention (6)	Log earn (7)	Retention (8)
D_i*Post_1	-0.059*** (0.013)	-0.032*** (0.011)	-0.058*** (0.013)	-0.018 (0.012)	x x	x x	-0.054*** (0.014)	-0.043* (0.024)
D_i*Post_1*F	0.012 (0.010)	0.022* (0.012)	0.011 (0.009)	0.015 (0.014)	0.003 (0.006)	0.019* (0.010)	-0.004 (0.014)	0.049* (0.026)
D_i*Post_2	0.001 (0.014)	-0.100*** (0.010)	-0.002 (0.013)	-0.099*** (0.011)	x x	x x	0.005 (0.016)	-0.178*** (0.022)
D_i*Post_2*F	-0.036*** (0.011)	0.054*** (0.012)	-0.029*** (0.011)	0.052*** (0.013)	-0.022*** (0.008)	0.045*** (0.010)	-0.057*** (0.016)	0.161*** (0.024)
Loc-gender-year FE	Yes	Yes	No	No	Yes	Yes	Yes	Yes
Loc-occ-gender-year FE	No	No	Yes	Yes	No	No	No	No
Est-year FE	No	No	No	No	Yes	Yes	No	No
N	765486	850646	765486	850646	765486	850646	236722	266139

Notes: This table estimates the MFA's treatment effects and gender differences in these effects. $Post_1$ is an indicator equal to one in years $t = 0$ to $t = 2$ and $Post_2$ equals one in years $t = 3$ to $t = 5$. F is a dummy variable indicating a female worker. The sample comprises all incumbent workers. These workers are tracked wherever they go. A worker is treated $D_i = 1$ if employed at an establishment exporting MFA-affected products at baseline and in the control group if exporting a different textile and clothing product. All specifications include worker fixed effects. Columns (1) and (2) are the baseline specification, including microregion x gender x year fixed effects. Columns (3) and (4) instead include 4 digit occupation x microregion x gender x year fixed effects. Columns (5) and (6) include establishment-year fixed effects, which absorb the main effect of treatment. Columns (7) and (8) present results narrowing in on workers in a specific but large 4-digit occupation, tailors, which employs 50% of women and 20% of men. Standard errors are clustered by establishment.

Table 4: Testing exclusion: pass-through estimates
(test of strategic wage spillovers)

	<i>Sample</i>			
<i>Panel A: Pass-through estimates</i>				
	Exporters	Untreated exporters	All unaffected employers	Establishments
	(1)	(2)	(3)	(4)
Δw_{-jt}	-0.003 (0.003)	0.000 (0.002)	0.005 (0.005)	0.034 (0.036)
$\Delta mrpl_{jt}$	0.145*** (0.052)			
<i>Panel B: First stage on Δw_{-jt}</i>				
Per 100 treated workers , excluding j	-0.100*** (0.013)	-0.099*** (0.014)	-0.089*** (0.021)	-0.036*** (0.003)
First stage F-stat	60.830	47.366	17.258	60.813
Avg. no. of treated workers , excluding j (hundreds)	25.833	20.018	30.239	38.284
Observations	147883	110595	426111	37674

Notes: This table estimates pass-through of own and competitor wage shocks. The outcome Δw_{jt} is the change in log wage at establishment j between $t = -1$ and $t \in 3, 5$. $\Delta mrpl_{jt}$ measures change in the log of own marginal product. Δw_{-jt} is the weighted sum of log wage changes at other textile and clothing establishments in j 's geography, excluding j , between $t = -1$ and $t = 1$, with weights equal to an establishment's baseline wage bill share in the labor market (textile and clothing industry x microregion). The coefficient on Δw_{-jt} estimates wage spillovers. The outcome in Columns (1)-(3) is the wages of incumbent workers whereas in Column (4) is the average wage of new workers. Own MFA status is the employer-level instrument for change in own marginal product. The share of treated employment in one's geography excluding one's own employer is the market-level instrument for competitor wage changes. Column (1) reports results for all establishments exporting textile or clothing products at baseline, column (2) for untreated exporters, column (3) for all textile and clothing establishments including non-exporters, and column (4) for textile and clothing exporters. Panel A reports pass-through estimates. Panel B reports the first stage on Δw_{jt} and the Kleibergen-Paap first stage F-stat. Standard errors are clustered at the establishment level.

Table 5: Average separation elasticities (negative)

	All workers		Tailors	
	Women	Men	Women	Men
	(1)	(2)	(3)	(4)
<i>Panel A: First stage</i>				
$\Delta \ln w_i$	-0.053*** (0.009)	-0.052*** (0.011)	-0.055*** (0.013)	-0.055*** (0.012)
F-stat	35.345	21.022	24.628	24.114
Observations	119754	135210	43905	26118
<i>Panel B: Elasticities</i>				
e_g	0.767 (0.279)	1.921 (0.431)	0.650 (0.358)	1.944 (0.426)
AR 90% CI	[0.334, 1.297]	[1.375, 2.982]	[0.087, 1.346]	[1.419, 2.639]
Observations	119754	135210	43905	26118

Notes: This table reports the negative of separation elasticities using the IV system from Section 2. The separation elasticity measures the change in separation for a given change in the offered wage. I use the change in a worker's wage between $t = -1$ and $t = 1$ to measure the change in offered wage. The MFA status of a worker's baseline establishment instruments for this change in offered wage. Columns (1) and (2) report results for all workers whereas Columns (3) and (4) report results for tailors (sewing machine operators) at baseline. All specifications include microregion x year fixed effects. I cluster standard errors by establishment. Panel A reports the first stage and the Kleibergen-Paap F-stat. Panel B reports average elasticity estimates, along with weak-identification robust confidence intervals. Appendix Table C5 reports elasticity estimates using the change in stayer wages to proxy for the change in offered wage, yielding very similar estimates.

Table 6: Gender differences in destination

	New txt. employer	New sector	New geography
	(1)	(2)	(3)
D_i*Post_1	-0.004 (0.008)	0.026*** (0.007)	0.010** (0.005)
D_i*Post_1*F	0.008 (0.008)	0.000 (0.008)	-0.006 (0.005)
D_i*Post_2	0.059** (0.008)	0.040*** (0.008)	0.006 (0.005)
D_i*Post_2*F	-0.011** (0.009)	-0.043*** (0.009)	-0.008 (0.006)
Observations	765486	765486	765486

Notes: This table estimates treatment effects and gender differences on worker destinations. $Post_1$ is an indicator equal to one in years $t = 0$ to $t = 2$ and $Post_2$ equals one in years $t = 3$ to $t = 5$. F is a dummy variable indicating female. The sample is all incumbent workers. They are tracked wherever they go. A worker is treated $D_i = 1$ if she worked at an establishment exporting MFA-treated products at baseline and in the control group if at a different textile exporter. All specifications include worker fixed effects. New is defined relative to baseline ($t = -1$). The outcome in Column (1) is a binary variable equal to one if a worker is employed at a new textile or garment industry employer. The textile and garment industries correspond with 2-digit 1995 CNAE industry codes 17 and 18. The outcome in column (2) is a binary variable equal to one if a worker is employed at a new non-textile employer. The outcome in column (3) equals one if the worker is employed in a new geography (microregion). Standard errors are clustered at the establishment level.

Table 7: Role of skills and occupations in explaining gender differences

	Role of skills			Role of occupations		
	Retention (1)	New sector (2)	New occupation (3)	Retention (4)	New sector (5)	New occupation (6)
D_i*Post_1*F	0.025** (0.012)	-0.009 (0.008)	-0.017* (0.009)	0.018 (0.012)	-0.001 (0.008)	-0.017* (0.009)
D_i*Post_2*F	0.057*** (0.012)	-0.051*** (0.009)	-0.062*** (0.010)	0.055*** (0.011)	-0.057*** (0.009)	-0.070*** (0.010)
Skill decile-treat-post FE	Yes	Yes	Yes	No	No	No
Occ-treat-post FE	No	No	No	Yes	Yes	Yes
Observations	850646	850646	850646	850646	850646	850646

Notes: This table explores the role of gender differences in skill and occupation in explaining gender differences in leaving after a wage drop. $Post_1$ is an indicator equal to one in years $t = 0$ to $t = 2$ and $Post_2$ equals one in years $t = 3$ to $t = 5$. F is a dummy for female. The skill level (1-8) on each of 35 skills required to do an occupation comes from O*NET. Skill distance is the weighted average of the distance (L1-norm) between a worker's current occupation and all other occupations in her geography, with weights being that occupation's share among jobs in the geography. Columns (1)-(3) control for a fixed effect for the post period x decile of skill distance from other jobs in one's geography, and treatment. This regression thus answers the question: controlling for some common rate of retention, switching industries, or switching to a new occupation among treated workers in the same skill decile in the post period, do women still leave less than men? Columns (4)-(6) do the analog, but with 4-digit occupation instead of the decile of skill transferability. The sample is all incumbent workers, tracked wherever they go. A worker is treated $D_i = 1$ if she worked at an establishment exporting MFA-affected products at baseline and in the comparison group if the establishment exported other textile and clothing products. All specifications include worker fixed effects. Standard errors are clustered by establishment.

Table 8: Elasticities fall as employers and industry grows larger

Panel A: Elasticities by industry and employer size

Firm share	Large txt sector		Small txt sector	
	Women	Men	Women	Men
	(1)	(2)	(3)	(4)
$s = 0.01$	2.113 (0.333)	x	2.294 (0.373)	3.597 (0.51)
$s = 0.05$	1.377 (0.192)	x	2.068 (0.287)	2.766 (0.380)
$s = 0.1$	0.767 (0.097)	x	1.8 (0.248)	2.105 (0.305)

Panel B: Average

	Women	Men
	(1)	(2)
e_g	1.236	2.689
μ_g	0.553	0.729

Notes: This table plots heterogeneity in labor supply elasticities by employer size (share of wage bill) and the size of the textile sector. The textile sector is considered large when its share is over 10% or more of a worker’s labor market (geography x education group). To estimate heterogeneous effects on the change in log wage and change in labor supply by share I interact a dummy for treatment with the share. I separately estimate this regression in locations with a large ($\geq 10\%$ wage bill at baseline) and small ($< 10\%$ wage bill at baseline) textile industry. Labor supply and wage variables are defined just as when estimating average elasticities. Standard errors are clustered by establishment in the regression. I use the delta method to create standard errors around elasticity estimates, since they are a ratio $\Delta \ln n / \Delta \ln w$.

Table 9: Estimates of η , θ , λ

Parameter	Women	Men
	(1)	(2)
η_g	2.19 (0.402)	3.89 (0.890)
θ_g	0.89 (0.355)	0.87 (0.421)
λ_g	0.03 (0.097)	0.05 (0.010)

Notes: This table reports estimates of η_g , θ_g , and λ_g . I estimate η_g as the elasticity facing atomistic employers in the textile and garment manufacturing industry, θ_g as the change in the rate at which workers exit an industry from a large versus small employer in the industry when the industry is small, and λ_g as the change in the rate at which workers exit a location from a large versus small employer in the location. Each estimate employs moments derived from the labor supply system, described in Section 5. In each specification I instrument for the change in wage using baseline MFA status. All specifications include microregion x year fixed effects. Standard errors are clustered at the establishment level.

Table 10: Sources: gender differences in monopsony power

	Textile		Economy-wide	
	(1)	(2)	(3)	(4)
	Women	Men	Women	Men
Match-specific preference				
$(1 + 1/\eta)$	69%	79%	69%	79%
Concentration				
Within-industry - $(1/\theta - 1/\eta) * HHI_{gk}$	66%	74%	66%	75%
Industry - $(1/\lambda - 1/\theta) * s_{gk} * HHI_{gk}$	55%	73%	45%	66%
ΔGWG				
Match-specific preference	10%		10%	
Within-industry concentration	-2%		-2%	
Concentration in industry	10%		12%	
Total monopsony-induced GWG	18%		20%	
p-val test: $1/\eta - 1/\theta < 0$	0.006	0.104	0.006	0.104
p-val test: $1/\lambda - 1/\theta < 0$	0.013	0.134	0.013	0.134

Notes: This table combines estimates of η_g , θ_g , and λ_g for each gender with measures of concentration and industry sizes in the RAIS data to calculate the gender wage gap from match-specific preferences, within-industry concentration, and concentration in some industries. As in Proposition 1.

C Appendix Tables

Table C1: Bindingness of MFA quota regime

	China	Brazil
6-digit products under quota	62%	14%
Fill rate (fr), conditional on quota	81%	13%
Binding (fr > 90%), conditional on quota	55%	5%

Notes: This table describes the share of 6-digit products facing quotas in China and Brazil at the MFA's end on January 1, 2005.

Table C2: MFA schedule of integration

Phase	Start date	No. of (10 digit) HS products integrated	Share of export volume
I	Jan 1, 1995	318	16%
II	Jan 1, 1998	744	17%
III	Jan 1, 2002	745	18%
IV	Jan 1, 2005	2978	49%

Notes: This table describes the schedule of quota removal on 10-digit HS products under the MFA, as implemented by the United States. It is from Brambilla et al. 2010.

Table C3: Nominal wage declines

	(1)	(2)	(3)	(4)
	Avg. monthly wage		December wage	
	Women	Men	Women	Men
D_i	0.053*** (0.012)	0.060** (0.029)	0.075** (0.033)	0.057** (0.027)
Constant	0.058*** (0.003)	0.072*** (0.003)	0.226*** (0.013)	0.225*** (0.013)
Observations	51533	62672	51533	62672

Notes: This table reports the MFA's effect on a dummy variable equal to one if a worker's nominal earnings fell between $t = -1$ and $t = 1$. The sample is incumbent workers. A worker is treated if s/he worked at an MFA-product-exporting establishment at baseline and comparison if at an exporter of a different textile and clothing product. Columns (1) and (2) report effects on monthly nominal earnings and Columns (3) and (4) on December earnings.

Table C4: New worker characteristics

	Women			Men		
	Age	No high school degree	Poached	Age	No high school degree	Poached
	(1)	(2)	(3)	(4)	(5)	(6)
Treated*post	-0.060 (0.294)	-0.010 (0.015)	-0.001 (0.011)	-0.540 (0.349)	-0.032 (0.020)	-0.004 (0.012)
Baseline mean	29.515	0.605	0.558	26.667	0.577	0.586
Observations	6759	6759	6759	6759	6759	6759

Notes: This table reports treatment effects on new worker characteristics before versus after the MFA. Each regression is at the establishment level. The outcome is the average characteristic of new workers in any given year. Each regression includes establishment and year fixed effects and clusters standard errors by establishment. Post is an indicator equal to one in and after 2005. Columns (1)-(3) report results for new female hires while Columns (4)-(6) report results for new male hires. Age is measured in years. High school dropout is the share of new workers without high school degrees. Poached is the share of workers poached in from another employer in the formal sector.

Table C5: Average separation elasticities (negative)
using change in stayer wage as change in offered wage

	All workers		Tailors	
	Women (1)	Men (2)	Women (3)	Men (4)
<i>Panel A: First stage</i>				
$\Delta \ln w_i$	-0.050*** (0.008)	-0.044*** (0.011)	-0.055*** (0.012)	-0.053*** (0.011)
F-stat	36.890	17.458	29.475	21.606
Observations	119754	135210	43905	26118
<i>Panel B: Elasticities</i>				
e_g	0.809 (0.299)	2.253 (0.562)	0.647 (0.352)	2.180 (0.523)
AR 90% CI	[0.334, 1.297]	[1.375, 2.982]	[0.087, 1.346]	[1.419, 2.639]
Observations	119754	135210	43905	26118
<i>Panel C: Gender differences</i>				
$e_m - e_f$	1.443		1.533	
90% CI	[0.399, 2.488]		[0.499, 2.567]	

Notes: This table is exactly as in 5, with the only difference being the measure of the change in a worker's offered log wage. Here I use the change in the average wage of stayers between $t = -1$ and $t = 1$ where a stayer is defined as any worker employed at an establishment in $t = -1$ who continues to remain employed in $t = 1$.

Table C6: Skills predictive of leaving for men

	Retention (1)	New sector (2)	New occupation (3)
$D_i * \text{Post}_1$	-0.027* (0.016)	0.020** (0.008)	0.015 (0.010)
$D_i * \text{Post}_2$	-0.069*** (0.014)	0.040*** (0.009)	0.045*** (0.010)
$D_i * \text{Post}_1 * \text{Transferable}$	-0.013 (0.019)	0.015 (0.011)	0.019 (0.012)
$D_i * \text{Post}_2 * \text{Transferable}$	-0.067*** (0.017)	0.039*** (0.013)	0.054*** (0.014)
Observations	397188	397188	397188

Notes: This table shows that skill transferability is predictive of leaving one's baseline employer, switching to a new sector, and switching to a new occupation for men. Skill transferability is defined as having above median transferable skills on the (negative) of the skill remoteness measure described in Appendix D.2. Post_1 is a dummy equal to 1 in years 0-2 after the MFA and Post_2 in years 3-5. The sample for these regressions is men. All specifications include geography x year fixed effects. Standard errors are clustered by establishment.

Table C7: Role of education and tenure

	Role of education			Role of tenure		
	Retention	New sector	New occupation	Retention	New sector	New occupation
	(1)	(2)	(3)	(4)	(5)	(6)
Treat*Post*F	0.023* (0.012)	0.000 (0.008)	-0.017* (0.009)	0.002 (0.013)	0.028*** (0.010)	-0.018* (0.011)
Treat*Post*F	0.054*** (0.012)	-0.042*** (0.009)	-0.059*** (0.010)	0.049*** (0.015)	-0.029** (0.011)	-0.063*** (0.012)
Education-treat-post FE	Yes	Yes	Yes	No	No	No
Tenure-treat-post FE	No	No	No	Yes	Yes	Yes
Observations	850646	850646	850646	850646	850646	850646

Notes: This table explores the role of gender differences in education and baseline tenure in explaining gender differences in leaving after a wage drop. $Post_1$ is an indicator equal to one in years $t = 0$ to $t = 2$ and $Post_2$ equals one in years $t = 3$ to $t = 5$. F is a dummy for female. Education is divided into three categories: high school dropouts, high school graduates, and those with some college education. Tenure is divided into ten deciles on baseline value. Columns (1)-(3) include controls for education level x treatment x post period and Columns (4)-(6) include controls for decile of tenure x treatment x post period. The sample is all incumbent workers, tracked wherever they go. A worker is treated $D_i = 1$ if she worked at an establishment exporting MFA-affected products at baseline and in the comparison group if the establishment exported other textile and clothing products. All specifications include worker fixed effects. Standard errors are clustered by establishment.

Table C8: Idiosyncratic constraints

	Women		Men	
	Childbearing	Unsafe municipality	Childbearing	Unsafe municipality
	(1)	(3)	(4)	(6)
$\Delta \ln w_i$	2.110*** (0.474)	2.278*** (0.486)	4.997*** (0.976)	3.104*** (0.925)
$\Delta \ln w_i \times X_i$	-0.542*** (0.070)	-0.551* (0.166)	-1.898*** (0.332)	-1.239 (1.381)
Observations	65913	65913	49482	49482

Notes: This table plots heterogeneous elasticities along two dimensions: being of childbearing age, defined as between 20 and 35 years; and employer location in an unsafe municipality. A municipality is designated unsafe if, in the four years between 2000 and 2004, its homicide rate was above the 75th percentile.

Table C9: Summary, sources of disproportionate monopsony

Name	Women	Men	Source	Specification
<i>Panel A: Horizontal differentiation</i>				
η_g Idiosyncratic ability to switch employers	2.19	3.89	Elasticity at atomistic employers ($s \sim 0$)	Section 5
θ_g Idiosyncratic ability to switch employers across sectors	0.89	0.87	Change in elasticity of leaving sector with employer size	Section 5
λ_g Idiosyncratic ability to switch employers across regions	0.03	0.05	Change in elasticity of leaving region with employer size	Section 5
<i>Panel B: Concentration within sector</i>				
HHI Payroll weighted Herfindahl	0.08	0.112	Calculated in RAIS	x
<i>Panel C: Share of textiles</i>				
s_{txt} Share of T&C in labor market	0.105	0.035	Calculated in RAIS	x

Notes: This table summarizes the sources of disproportionate monopsony. η_g , θ_g , λ_g are as estimated in Section 5. The payroll-weighted Herfindahl is calculated in RAIS as is the share of the textile and clothing industry in the labor market. The share of textiles is the weighed average over microregions, where each microregion is weighted by the share of the total textile wage bill in that microregion.

Table C10: Model parameters for counterfactuals

Parameter	Name	Value	Source	Component
<i>Estimated</i>				
η_g	Cross-employer	2.19 (w), 3.89 (m)	Elasticity estimate	LS
θ_g	Cross-industry elasticity of substitution	0.89 (w), 0.87 (m)	Elasticity estimate	LS
λ_g	Cross-location elasticity of substitution	0.03 (w), 0.05 (m)	Elasticity estimate	LS
s_{gk}	Share of industries	Varies	Data	Eqbm
a_{gk}	Industry-specific amenities	Varies	Match s_{gk} in data	Eqbm
W_{gk}	Industry-specific wages	Varies	Match s_{gk} in data	Eqbm
β_k	Women's productivity shifter in textiles	1.286	Estimated (Akerberg et al. 2015)	Prod
β_k	Food mfg	1.1	Estimated (Akerberg et al. 2015)	Prod
β_k	Metal	0.5	Estimated (Akerberg et al. 2015)	Prod
β_k	Machinery	0.54	Estimated (Akerberg et al. 2015)	Prod
β_k	Auto	0.34	Estimated (Akerberg et al. 2015)	Prod
ζ	Productivity dispersion	0.7	Firm size distribution in textiles	Prod
<i>Calibrated</i>				
σ	Elasticity of substitution between men and women	5.94	Gallen 2022	Prod

Notes: This table notes parameters needed to simulate the model, their source, and which feature of the environment they correspond with (LS = labor supply, Prod = production function, Eqbm = equilibrium object, LD = labor demand is determined under a Cournot oligopsony).

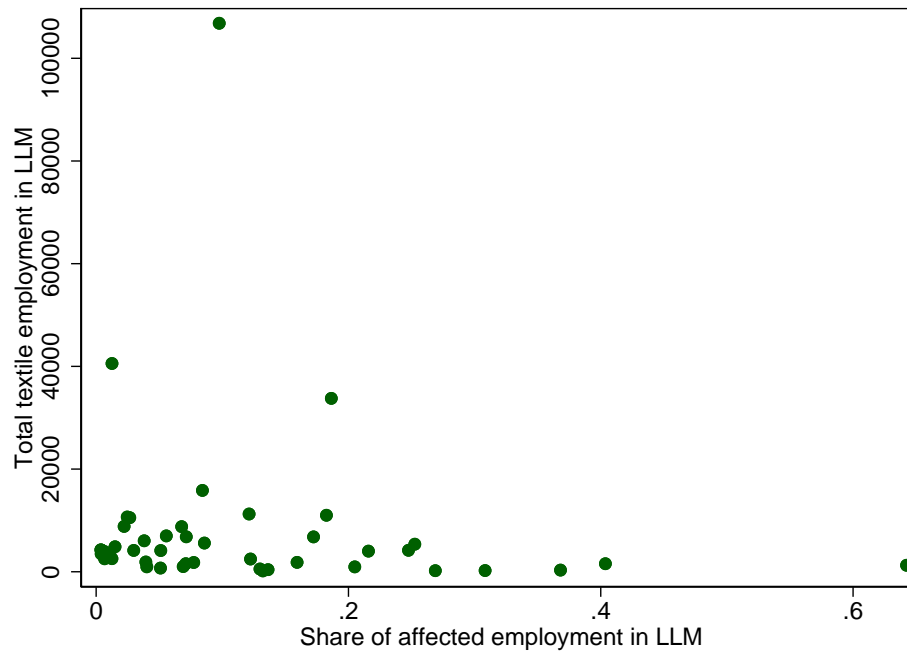
Table C11: Non-targeted moments

	Name	Model	Data
μ_f	Women's markdown	0.552	0.553
μ_m	Men's markdown	0.727	0.729
$\ln(\mu_m) - \ln(\mu_f)$	Monopsony-induced GWG	0.119	0.120

Notes: This table describes the average markdown and GWG simulated in the model using the baseline parameter values.

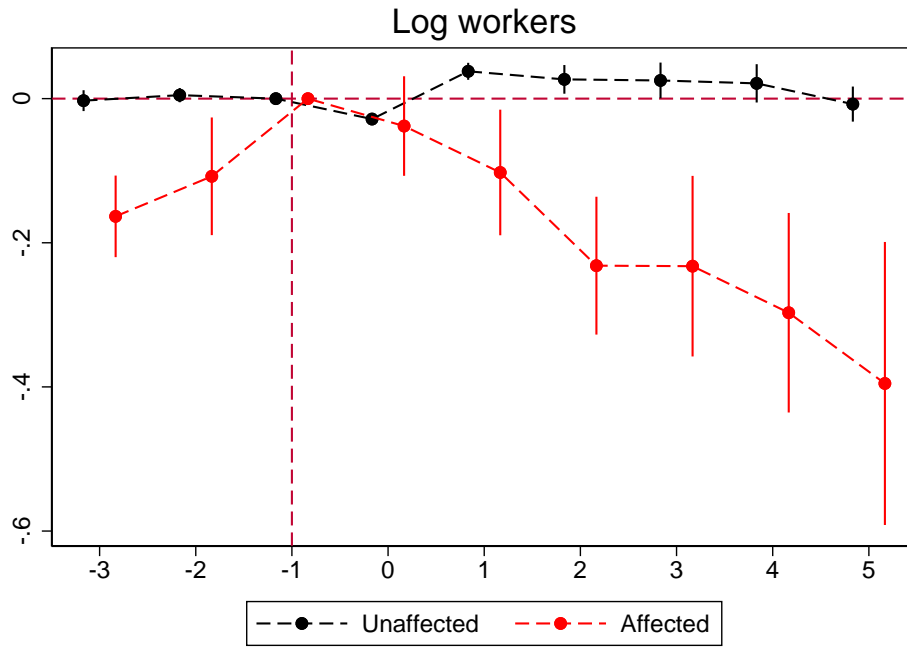
C Appendix Figures

Figure C1: Share of affected jobs over total number of textile jobs



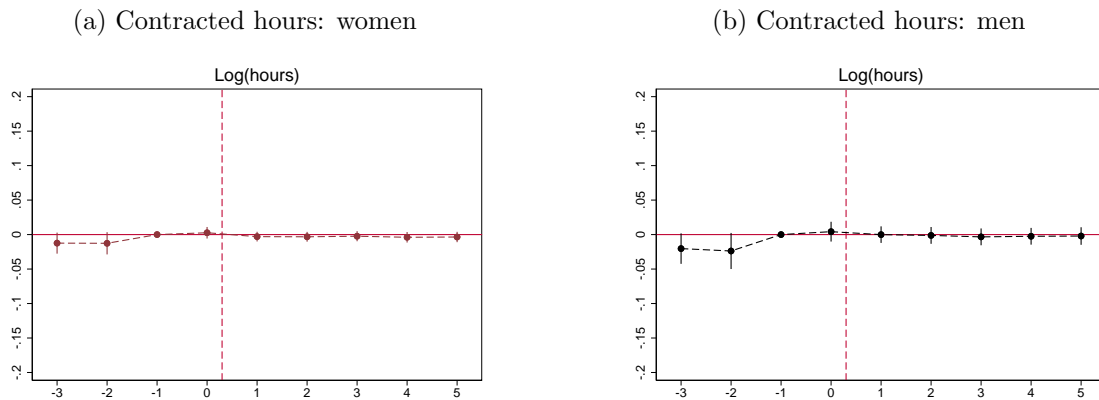
Notes: This figure shows that the MFA shock was small in all geographies, i.e. that, while textile and clothing industry may be clustered in some locations, treated employment still comprises a small share of total employment in these locations. In other words, the number of baseline T&C jobs in a microregion (y-axis) is uncorrelated with the share of treated T&C jobs in that microregion (x-axis).

Figure C2: Effect on aggregate employment



Notes: This figure plots the treatment effect on the log of employment at treated establishments and all other textile and clothing manufacturing establishments (including non-exporters). Standard errors are clustered by microregion


Figure C3: Effect on hours




Notes: This figure plots treatment effects on men and women's contracted hours.

Figure C4: Effect on employer-to-employer transitions

(a) Switching to a new employer: women




(b) Switching to a new employer: men




Notes: This figure plots treatment effects on employer-to-employer transitions, showing that the entire leaving response among men and women is driven by transitions to new employers as opposed to transitions into unemployment.

Figure C5: Effect on new worker log wages

(a) New workers: women




(b) New workers: men




Notes: This figure plots DiD treatment effects on the log of the average wage paid to new workers hired by an establishment in the past year. Each regression is at the establishment-level. Standard errors are clustered by establishment.

Figure C6: Testing exclusion: variation in instrument

(a) Share of treated wage bill in region, excluding self




(b) By gender




Notes: This figure plots the share of the baseline wage bill in a worker's microregion, excluding his or her own employer, that is treated by the MFA shock. The left panel plots this for all workers and the right separately by gender. This is the instrument that provides variation in how much share-weighted wages at other employers in the labor market change to test for wage spillovers.

Figure C7: Testing exclusion: pass-through of others' wage shocks

(a) First stage

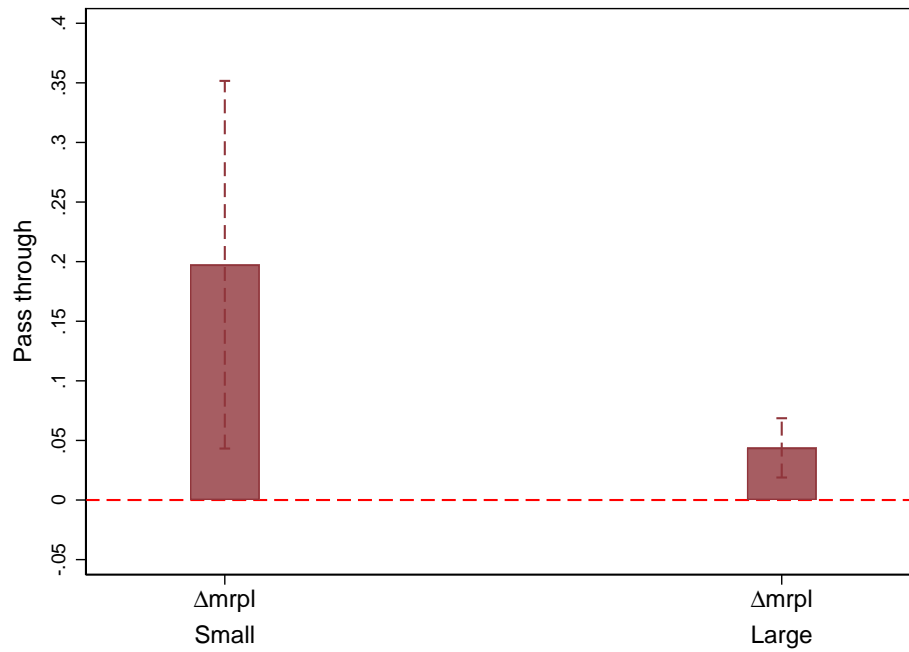


(b) Visual IV



Notes: This figure plots the visual IV version of Table 4. Here the treatment variable, instead of being the number of treated workers in one's industry and geography, is the share of treated wage bill in one's industry and geography (still excluding one's own employer).


Figure C8: Variation in own-mrpl pass-through by employer size




Notes: This figure plots the pass through of shocks to own mrpl on stayers' wages. Small employers are those with wage bill less than 1% of the local textile market whereas large employers are those with wage bill over 10%. The figure plots 95% confidence intervals.

Figure C9: Staying rates for childbearing and non-childbearing workers


(a) Childbearing: women




(b) Childbearing: men



(c) Non-childbearing: women

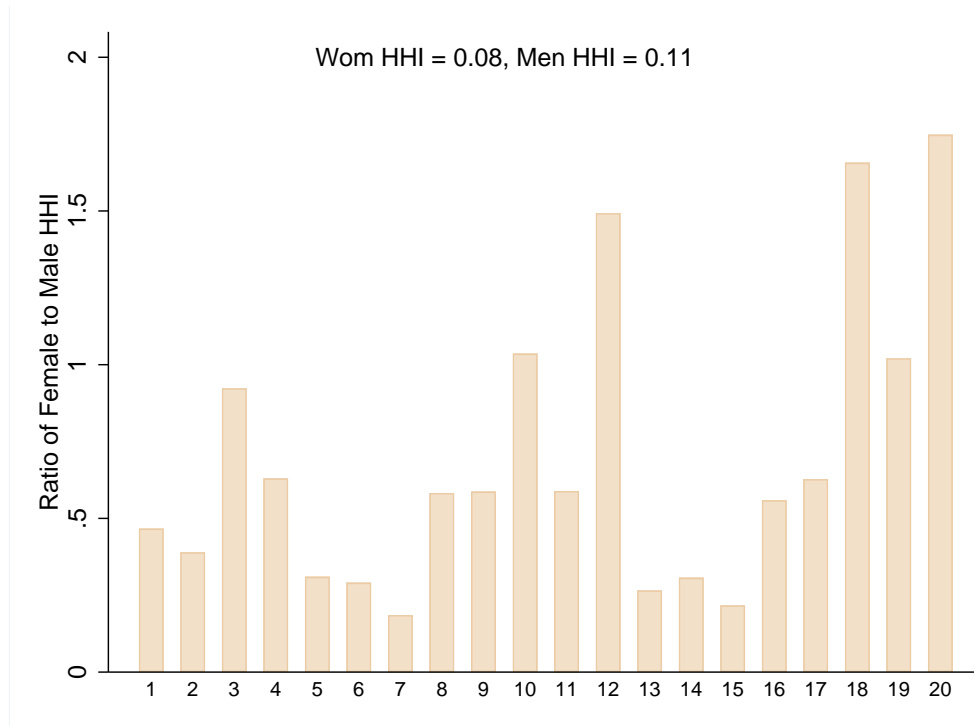


(d) Childbearing: men



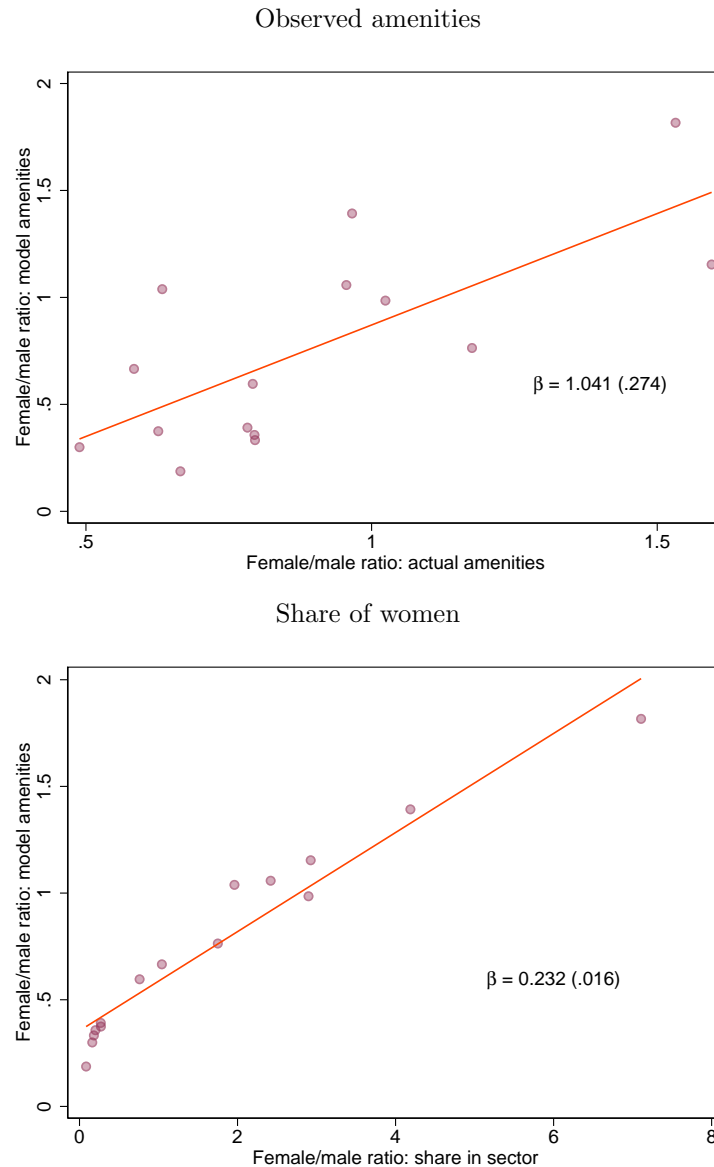
Notes: This figure plots treatment effects on retention for women and men of childbearing ages (20-35 years) and other ages. The figure plots 95% confidence intervals

Figure C10: Within-industry concentration in textiles



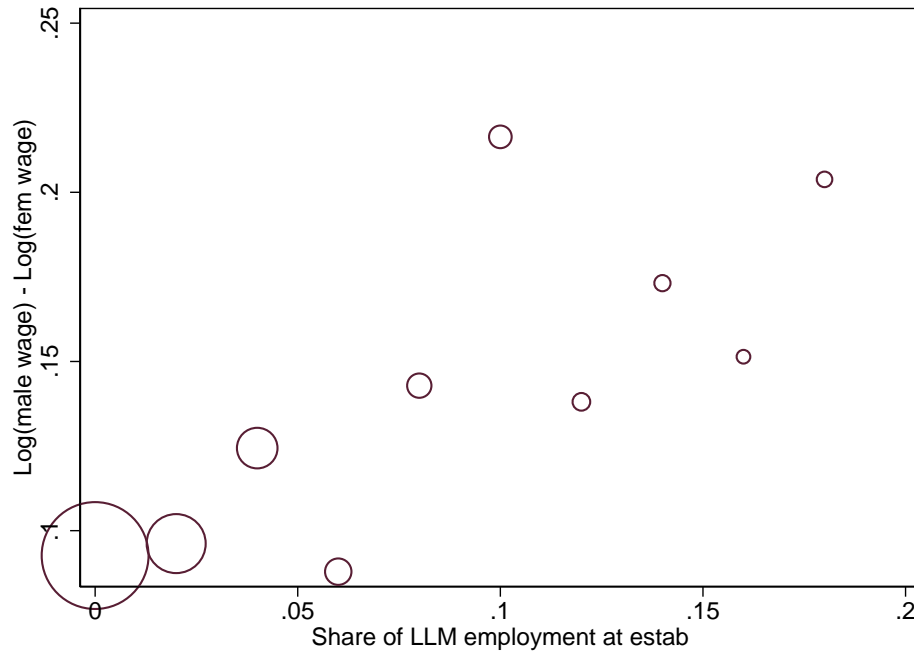
Notes: This figure plots the ratio of female to male payroll-weighted HHIs in the 20 microregions with the highest share of women’s total textile wage bill. About 60% of both men and women’s textile wage bill is in these microregions. Microregions are sorted from largest to smallest share, with 1 being the largest and 20 being the smallest. The text on top reports women and men’s aggregated payroll-weighted HHI, i.e., summing over microregions, where each microregion is weighted by the gender-specific share of the wage bill in that microregion. This is the right notion of an aggregate, as shown in Appendix A.2, where by “right” we mean the HHI facing the average women or the average man in the sector.

Figure C11: Correlation between inferred amenities and observed amenities, women’s share in industry



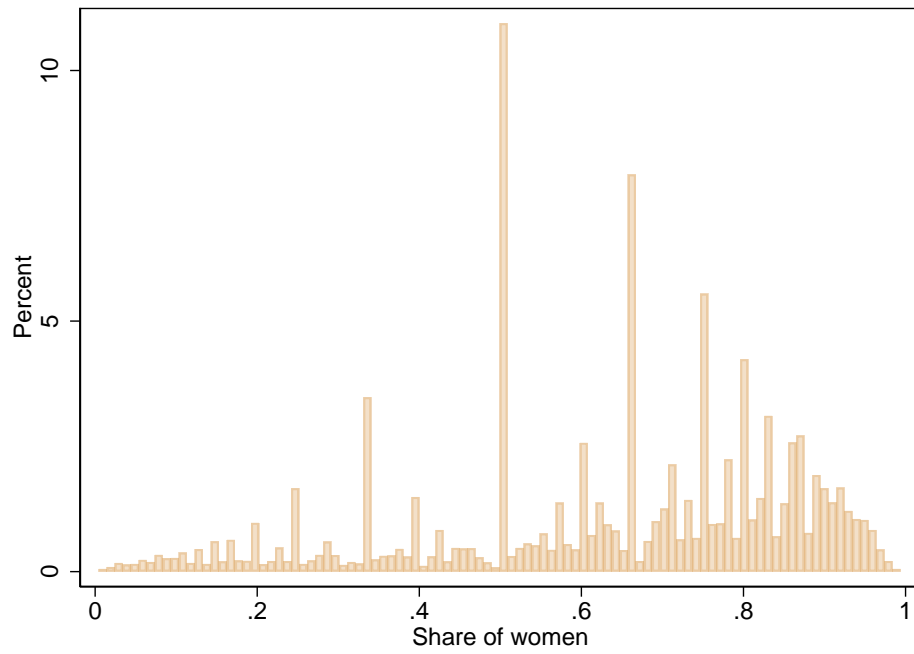
Notes: This figure plots, for each sector, observed and inferred female/male amenities. Observed amenities come from the text of all collective bargaining contracts, where we observe 137 different types of amenities at the establishment level, including maternity leave, childcare, and flexibility in work hours. I use the revealed preference approach developed in Corradini, Lagos & Sharma 2022 to define female and male-centric amenities: correlating gender-specific job ladders inferred from employer-to-employer moves with contracted amenities to infer which are valued. In each sector, average observed amenities are a weighted-average of female and male-centric amenities across all establishments, with the weights being that establishment’s share of the sector’s wage bill.

Figure C12: GWG increasing with firm size



Notes: This figure plots the gender wage gap at an establishment, i.e. the difference in the log of men and log of women's wages (y-axis) against the share of the T&C wage bill in that establishment's microregion at that establishment (x-axis). The size of bins corresponds with the size of employment in that bin. It shows that the gender wage gap rises with establishment size.

Figure C13: Share of women at treated firms



Notes: This figure plots the baseline distribution of the share of women in the workforce of treated establishments.

A Proofs and Derivations

A.1 Model Derivations

A.1.1 Setup

Markets The economy features a continuum of geographies $r \in [0, 1]$ (microregions). Each geography features a discrete number of industries indexed by $k \in 1, \dots, M_r$, and firms within the industry $j \in 1, \dots, J_m$. Throughout I describe a geography as a “local labor market” and all same-gender jobs within an industry as the within-industry market.

Firms Firms compete in a Cournot oligopsony. Each posts gender-specific wages $\{w_m, w_f\}_j$ and chooses employment to maximize profits given the (inverse) labor supply it faces. The production function of each firm is differentiable and concave, with imperfect substitutability between male and female workers $F(f_j, m_j)$. Each firm is associated with exogenously-given industry-specific amenities (a_{gk}) and the firm’s deviation from this industry norm (a_{gj}). I abstract from competition in amenities for simplicity.

Workers Workers, indexed by group g (men or women), possess heterogeneous preferences over employers. They choose to work at their higher utility employer, and exhibit three-nested preferences, choosing first a location, then an industry, and finally an employer within the industry. Their utility from working at employer j has a common group-specific component (rising in wages and amenities), and an idiosyncratic preference shock specific to each employment relationship ϵ_{igjk} . Each worker must earn $y_i \sim F(y)$, where earnings equal her wage multiplied by hours of supplied work $y_i = w_{gj}h_{igj}$.

$$u_{igjk} = \ln w_{gj} + \ln a_{gk} + \ln a_{gj} + \epsilon_{igjk}$$

ϵ_{igjk} has a nested GEV Type I extreme value distribution. Its variance depends on three dispersion parameters η_g , θ_g , and λ_g .⁴⁶

$$F(\epsilon_{i1}, \dots, \epsilon_{NJ}) = \exp \left[- \sum_r \left(\sum_{k=1}^M \left(\sum_{j=1}^{J_m} e^{-(1+\eta_g)\epsilon_{igjk}} \right)^{\frac{1+\theta_g}{1+\eta_g}} \right)^{\frac{1+\lambda_g}{1+\theta_g}} \right]$$

A.1.2 Labor supply

The probability of choosing an employer is given by the standard nested logit (McFadden, 1978):

⁴⁶As shown by Anderson et al. (1992) this discrete choice framework generates the same labor supply to firms as obtained in a nested CES setting with the outermost nest being regions, the middle nest being industries, and the bottom nest being employers within the industry. η_g then corresponds with the within-industry elasticity of substitution, θ_g is the cross-industry elasticity of substitution, and λ_g the cross-region.

$$p_{gjk r} = \underbrace{\frac{(a_{gj} w_{gj})^{1+\eta_g}}{\sum_{j' \in k} (a_{gj'} w_{gj'})^{1+\eta_g}}}_{\text{prob of choosing firm } j \text{ in industry } k} \times \underbrace{\frac{a_{kg}^{1+\theta_g} \left(\sum_{j' \in k} (a_{gj'} w_{gj'})^{1+\eta_g} \right)^{\frac{1+\theta_g}{1+\eta_g}}}{\sum_{k' \in r} a_{k'g}^{1+\theta_g} \left(\sum_{j'' \in k'} (a_{gj''} w_{gj''})^{1+\eta_g} \right)^{\frac{1+\theta_g}{1+\eta_g}}}}_{\text{prob of choosing industry } k} \times \underbrace{\frac{\bar{W}_{gr}^{1+\lambda_g}}{\sum_{r'} \bar{W}_{gr'}^{1+\lambda_g}}}_{\text{prob of choosing region } r}$$

I aggregate these probabilities over workers to obtain the labor supply to an employer j in industry k in region r :

$$n_{gjk r} = \int p_{gjk r} h_{igjk r} dF(y_i), \quad h_{igjk r} = \frac{y_i}{w_{gjk r}}$$

$$n_{gjk r} = \frac{(w_{gjk r})^{\eta_g}}{\sum_{j' \in k} (a_{gj'} w_{gj'})^{1+\eta_g}} \frac{\left(\left(\sum_{j' \in k} (a_{gj'} w_{gj'})^{1+\eta_g} \right)^{\frac{1+\theta_g}{1+\eta_g}} \right)^{\frac{1+\theta_g}{1+\eta_g}}}{\sum_{k' \in r} a_{k'g}^{1+\theta_g} \left(\sum_{j'' \in k'} (a_{gj''} w_{gj''})^{1+\eta_g} \right)^{\frac{1+\theta_g}{1+\eta_g}}} \times$$

$$\frac{\left(\sum_{k' \in r} a_{k'g}^{1+\theta_g} \left(\sum_{j'' \in k'} (a_{gj''} w_{gj''})^{1+\eta_g} \right)^{\frac{1+\theta_g}{1+\eta_g}} \right)^{\frac{1+\lambda_g}{1+\theta_g}}}{\sum_r \left(\sum_{k'' \in r} a_{gk''}^{1+\theta_g} \left(\sum_{j''' \in k''} (a_{gj''' } w_{gj'''})^{1+\eta_g} \right)^{\frac{1+\theta_g}{1+\eta_g}} \right)^{\frac{1+\lambda_g}{1+\theta_g}}} (a_{gjk r})^{1+\eta_g} a_{gk}^{1+\theta_g} Y_g$$

where $Y_g = \sum_n w_{gn} n_{gn}$ denotes the total labor income of the group summed over all employers in the economy. I define the following wage indices at the industry-region, region, and group levels:

$$\bar{W}_{gkr} = \left(\sum_{j' \in k, r} (a_{gj'} w_{gj'})^{1+\eta_g} \right)^{\frac{1}{1+\eta_g}}$$

$$\bar{W}_{gr} = \left(\sum_{k' \in r} a_{gk'}^{1+\theta_g} \left(\sum_{j' \in k} (a_{gj'} w_{gj'})^{1+\eta_g} \right)^{\frac{1+\theta_g}{1+\eta_g}} \right)^{\frac{1}{1+\theta_g}}$$

$$\bar{W}_g = \left(\sum_r \bar{W}_{gr}^{1+\lambda_g} \right)^{\frac{1}{1+\lambda_g}}$$

And the following employment indices:

$$N_{gkr} = \left(\sum_{j' \in k, r} (a_{gj'}^{-1} n_{gj'})^{\frac{1+\eta_g}{\eta_g}} \right)^{\frac{\eta_g}{1+\eta_g}}$$

$$N_{gr} = \left(\sum_{k \in r} (a_{gk}^{-1} N_{gkr})^{\frac{1+\theta_g}{\theta_g}} \right)^{\frac{\theta_g}{1+\theta_g}}$$

$$N_g = \left(\sum_r N_{gr}^{\frac{1+\lambda_g}{\lambda_g}} \right)^{\frac{\lambda_g}{1+\lambda_g}}$$

Along with (1), these indices imply $W_g N_g = Y_g$. To obtain the labor supply to an employer I plug these expressions back into (1), yielding the nested CES supply curve to j :

$$n_{gjdkr} = \left(\frac{w_{gjdkr}}{\bar{W}_{kgkr}} \right)^{\eta_g} \left(\frac{\bar{W}_{kgkr}}{\bar{W}_{gr}} \right)^{\theta_g} \left(\frac{\bar{W}_{gr}}{\bar{W}_g} \right)^{\lambda_g} a_{gjdkr}^{1+\eta_g} a_{gk}^{1+\theta_g} N_g$$

I invert the labor supply curve in three steps:

$$N_{gr} = \left(\frac{\bar{W}_{gr}}{\bar{W}_g} \right)^{\lambda_g} N_g$$

$$\bar{W}_{gr} = \left(\frac{N_{gr}}{N_g} \right)^{\frac{1}{\lambda_g}} W_g$$

Next:

$$N_{gkr} = \left(\frac{\bar{W}_{kgkr}}{\bar{W}_{gr}} \right)^{\theta_g} \left(\frac{\bar{W}_{gr}}{\bar{W}_g} \right)^{\lambda_g} a_{gk}^{1+\theta_g} N_g$$

$$\bar{W}_{kgkr} = \left(\frac{N_{gkr}}{N_{gr}} \right)^{\frac{1}{\theta_g}} a_{gk}^{-\left(\frac{1+\theta_g}{\theta_g}\right)} \bar{W}_{gr}$$

Finally:

$$n_{gjdkr} = \left(\frac{w_{gjdkr}}{\bar{W}_{kgkr}} \right)^{\eta_g} n_{gkr} a_{gjdkr}^{1+\eta_g}$$

$$w_{gjdkr} = \left(\frac{n_{gjdkr}}{n_{gkr}} \right) a_{gjdkr}^{-\left(\frac{1+\eta_g}{\eta_g}\right)} \bar{W}_{kgkr}$$

Together, these yield the inverse labor supply curve:

$$w_{gjdkr} = \left(\frac{n_{gjdkr}}{N_{gkr}} \right)^{\frac{1}{\eta_g}} \left(\frac{N_{gkr}}{N_{gr}} \right)^{\frac{1}{\theta_g}} \left(\frac{N_{gr}}{N_g} \right)^{\frac{1}{\lambda_g}} W_g a_{gk}^{\frac{-(1+\theta_g)}{\theta_g}} a_{gjdkr}^{\frac{-(1+\eta_g)}{\eta_g}}$$

Labor supply elasticity I obtain the inverse elasticity of residual labor supply to a single employer j by taking the derivative of its log wage wrt log employment:

$$\ln w_{gjdkr} = \frac{1}{\eta_g} \ln n_{gjdkr} + \left(\frac{1}{\theta_g} - \frac{1}{\eta_g} \right) \ln N_{gkr} + \left(\frac{1}{\lambda_g} - \frac{1}{\theta_g} \right) \ln N_{gr} + \text{Aggregates} + \text{Amenities}$$

Before doing so, I prove a useful lemma.

Lemma 1 $\frac{\partial \ln N_{gkr}}{\partial \ln n_{gjdkr}} = s_{gj}$

Proof. By definition, $N_{gkr} = (\sum_{j' \in k,r} a_{gj'}^{-1} n_{gj'}^{\frac{1+\eta_g}{\eta_g}})^{\frac{\eta_g}{1+\eta_g}}$. Thus:

$$\frac{\partial \ln N_{gkr}}{\partial \ln n_{gjk}} = \frac{\eta_g}{1 + \eta_g} \frac{\partial \ln (\sum_{j' \in k,r} (a_{gj'}^{-1} n_{gj'}^{\frac{1+\eta_g}{\eta_g}})^{\frac{1+\eta_g}{\eta_g}})}{\partial n_{gj}} n_{gj}$$

$$\frac{\partial \ln N_{gkr}}{\partial \ln n_{gjk}} = \frac{(a_{gj}^{-1} n_{gjk})^{\frac{1+\eta_g}{\eta_g}}}{\sum_{j' \in k,r} (a_{gj'}^{-1} n_{gj'})^{\frac{1+\eta_g}{\eta_g}}}$$

By definition, $s_{gjk} := \frac{w_{gjk} n_{gjk}}{\sum_{j' \in k,r} w_{j'} n_{j'}} = \frac{(a_{gj}^{-1} n_{gjk})^{\frac{1+\eta_g}{\eta_g}}}{\sum_{j' \in k,r} (a_{gj'}^{-1} n_{gj'})^{\frac{1+\eta_g}{\eta_g}}}$ (plugging in the inverse labor supply to j and $j' \in k, r$), thus proving the lemma. ■

By a similar argument, $\frac{\partial \ln N_{gr}}{\partial \ln n_{gjk}} = \frac{\partial \ln N_{gr}}{\partial \ln N_{gkr}} \frac{\partial \ln N_{gkr}}{\partial \ln n_{gjk}} = s_{gk} s_{gj}$. Therefore, the elasticity of residual labor supply to employer j in industry k in region r is:

$$e_{gj} = \left(\frac{\partial \ln w_{gj}}{\partial \ln n_{gj}} \right)^{-1} = \left[\frac{1}{\eta_g} + \left(\frac{1}{\theta_g} - \frac{1}{\eta_g} \right) s_{gj} + \left(\frac{1}{\lambda_g} - \frac{1}{\theta_g} \right) s_{gj} s_{gk} \right]^{-1} \quad (6)$$

A.1.3 Firm optimization: elasticities govern markdowns

Workers are paid only a fraction of their marginal product and markdowns are governed by employer-specific elasticities. Each employer chooses labor f_j and m_j to maximize profits, given the inverse labor supply it faces.

$$\max_{f_j, m_j} F(f_j, m_j) - w_{fj} f_j - w_{mj} m_j$$

with the FOC determining demand for each group:

$$MRPL_{gj} = \left(1 + \frac{1}{e_{gj}} \right) w_{gj}$$

$$MRPL_{gj} = \mu_{gj}^{-1} w_{gj}$$

When $e_{gj} = \infty$, in a competitive market, workers are paid exactly their marginal product. However, because firms have market power, workers are paid below their marginal product of labor with markdown μ_{gj} :

$$\mu_{gj} = \frac{e_{gj}}{1 + e_{gj}}$$

Workers with more elastic supply have higher markdowns, i.e. are paid closer to their marginal product.

A.1.4 Condition for estimated reduced form elasticity to equal structural elasticity

This section demonstrates why the elasticity of residual labor supply governing markdowns is typically inestimable with a firm specific shock. Intuitively, this is because such shocks typically uncover the total derivative with respect to j 's wage whereas the structural elasticity that governs markdowns as shown above is a partial derivative. It also yields the conditions under which this structural elasticity *is* estimable with firm-specific shocks. Starting from the labor supply system, where I omit time subscripts t for visual benefit but the changes are one period following time t :

$$\ln n_{gjjkr} = \eta_g \ln w_{gjjkr} + (\theta_g - \eta_g) \ln \bar{W}_{kgr} + (\lambda_g - \theta_g) \ln \bar{W}_{gr} + \text{Aggregates}$$

Consider a first-order approximation around the Nash equilibrium, following any change to firms in the region:

$$\begin{aligned} \Delta \ln n_{gjjkr} &= \eta_g \Delta \ln w_{gjjkr} + (\theta_g - \eta_g) \sum_{j' \in k, r} \frac{\partial \ln \bar{W}_{kgr}}{\partial \ln w_{gj'}} \Big|_{w^*_{j'=j}} \Delta \ln w_{j'} + (\lambda_g - \theta_g) \sum_{j'' \in r} \frac{\partial \ln \bar{W}_{gr}}{\partial \ln w_{gj''}} \frac{\partial \ln \bar{W}_{kgr}}{\partial \ln w_{gj''}} \Big|_{w^*_{j'=j''}} \Delta \ln w_{j''} \\ \frac{\Delta \ln n_{gjjkr}}{\Delta \ln w_{gjjkr}} &= (\eta_g + (\theta_g - \eta_g) s_{gj} + (\lambda_g - \theta_g) s_{gj} s_{gk}) + \frac{(\theta_g - \eta_g)}{\Delta \ln w_{gjjkr}} \sum_{j' \in k, r} s_{gj'} \Delta \ln w_{j'} + (\lambda_g - \theta_g) \sum_{j'' \in r} s_{gj''} s_{gk} \Delta \ln w_{j''} \end{aligned}$$

The estimated reduced form elasticity is:

$$\begin{aligned} \epsilon_{gjjkr} &= \frac{\Delta \ln n_{gjjkr}}{\Delta \ln w_{gjjkr}} \\ \epsilon_{gjjkr} &= e_{gjjkr} + \frac{1}{\Delta \ln w_{gjjkr}} ((\theta_g - \eta_g) \sum_{j' \in k, r} s_{gj'} \Delta \ln w_{j'} + (\lambda_g - \theta_g) \sum_{j'' \in r} s_{gj''} s_{gk} \Delta \ln w_{j''}) \end{aligned}$$

The test of strategic interaction I perform below argues that $\Delta \ln w_{j'} = 0 \forall j' \in k, r$. While I report results only for employers in the textile industry, I can also show that $\Delta \ln w_{j''} = 0 \forall j'' \in r \notin k$. This is intuitive, since employers in the textile industry compete most tightly with others that are also in the industry. Therefore, I estimate $\epsilon_{gjjkr} = e_{gjjkr}$.

A.2 Proofs of in-text propositions

Proposition 1 *Through the lens of the model, monopsony over the average worker in a group g in industry k is given by the following (inverse) average markdown:*

$$\bar{\mu}_{gkr}^{-1} = \frac{mr\bar{p}l_{gkr}}{w_{gkr}} = 1 + \frac{1}{\eta_g} + \left(\frac{1}{\theta_g} - \frac{1}{\eta_g} \right) HHI_{gkr} + \left(\frac{1}{\lambda_g} - \frac{1}{\theta_g} \right) HHI_{gkr} s_{gkr}$$

Proof. The proof has two steps. First I show that inverse of the average markdown for the average worker of group g in industry k in region r is just the share-weighted average of the inverse of individual employer markdowns: $\bar{\mu}_{gkr}^{-1} = (\sum_{j \in k, r} s_{gj} \mu_{gj}^{-1})$.

$$\begin{aligned}
\sum_{j \in k,r} s_{gj} \mu_{gj}^{-1} &= \sum_{j \in k,r} s_{gj} \left(1 + \frac{1}{e_{gj}} \right) \\
&= \sum_{j \in k,r} s_{gj} \left(\frac{mrpl_{gj}}{w_{gj}} \right) \\
&= \sum_{j \in k,r} \frac{n_{gj}}{\sum_{j' \in k,r} w_{gj'} n_{gj'}} \frac{mrpl_{gj}}{1} \\
&= \frac{\sum_{j \in k,r} mrpl_{gj} n_{gj}}{\sum_{j' \in k,r} w_{gj'} n_{gj'}} \times \frac{\sum_{j' \in k,r} n_{gj}}{\sum_{j' \in k,r} n_{gj}} \\
&= \frac{mrpl_{gkr}}{w_{gkr}} \\
&= \mu_{gkr}^{-1}
\end{aligned}$$

Next, I calculate the average group-specific markdown through the lens of the model:

$$\begin{aligned}
\sum_{j \in k,r} s_{gj} \mu_{gj}^{-1} &= \sum_{j \in k,r} s_{gj} \left(1 + \frac{1}{e_{gj}} \right) \\
\sum_{j \in k,r} s_{gj} \mu_{gj}^{-1} &= \sum_{j \in k,r} s_{gj} \left(1 + \frac{1}{\eta_g} (1 - s_{gj}) + \frac{1}{\theta_g} (s_{gj} (1 - s_{kr})) + \frac{1}{\lambda_g} (s_{gj} s_{kr}) \right) \\
\sum_{j \in k,r} s_{gj} \mu_{gj}^{-1} &= 1 + \frac{1}{\eta_g} + \left(\frac{1}{\theta_g} - \frac{1}{\eta_g} \right) HHI_{gk} + \left(\frac{1}{\lambda_g} - \frac{1}{\theta_g} \right) \sum_r s_{gr} HHI_{gkr} s_{gkr}
\end{aligned}$$

■

To obtain the average inverse markdown for a worker of group g in k I aggregate to the industry-level over regions r :

$$\begin{aligned}
\bar{\mu}_{gk}^{-1} &= \sum_r s_{rk}^g \sum_{j \in k,r} s_{gj} \mu_{gj}^{-1} \\
&= \sum_r s_{rk}^g \left[1 + \left(\frac{1}{\eta_g} \right) (1 - HHI_{gkr}) + \left(\frac{1}{\theta_g} \right) HHI_{gkr} (1 - s_{gkr}) + \frac{1}{\lambda_g} HHI_{gkr} s_{gkr} \right] \\
&= 1 + \frac{1}{\eta_g} + \left(\frac{1}{\theta_g} - \frac{1}{\eta_g} \right) HHI_{gk} + \left(\frac{1}{\lambda_g} - \frac{1}{\theta_g} \right) \sum_r s_{rk}^g HHI_{gkr} s_{gkr}
\end{aligned}$$

where s_{rk}^g is the share of the group's overall wage bill in industry k that is in microregion r . It is straightforward to see that this measures the average inverse markdown of a worker in industry k by closely paralleling Step 1 of the proof of the above proposition.

Proposition 2 *The share of the gender wage gap explained by monopsony over the average worker is.*

$$\frac{\ln(\bar{\mu}_m) - \ln(\bar{\mu}_f)}{\ln(\bar{w}_m) - \ln(\bar{w}_f)}$$

where \bar{w}_m is the average man's wage and \bar{w}_f is the average woman's wage and $\bar{\mu}_g = (\sum_j s_{gj} \mu_{gj}^{-1})^{-1}$ is the average markdown, i.e. how much the average worker is paid below marginal product with s_{gj} being the employer's share of the wage bill in textiles.

Proof. Start from the accounting identity $\bar{w}_g = \bar{mrpl}_g \bar{\mu}_g$. Taking logs and re-arranging yields the result:

$$\ln(\bar{w}_m) - \ln(\bar{w}_f) = \ln(\bar{\mu}_m) - \ln(\bar{\mu}_f) + \ln(\bar{mrpl}_m) - \ln(\bar{mrpl}_f)$$

Thus, the share potentially due to monopsony is $\frac{\ln(\bar{\mu}_m) - \ln(\bar{\mu}_f)}{\ln(\bar{w}_m) - \ln(\bar{w}_f)}$. ■

B Strategic Wage Setting

B.1 Test and Results

To test for spillover effects, I implement a test from the exchange rate pass-through literature (Amiti et al. 2019) in the labor setting. Its key insight is that strategic motives for wage setting operate by changing an employer's markdown: in other words, as workers flock from a China-competing employer that lowers their wage to a non-China-competing employer, the second can now pay them a smaller fraction of marginal product. I show that, for any structure of competition among employers (including Cournot and Bertrand oligopsony), and invertible labor supply system (where employers are not perfect substitutes), I can estimate wage spillovers by regressing an employer's wage change on a weighted average of wage changes at its competitors, controlling for changes to its own marginal product, without having to account for what spurred the change or who it comes from. While one main motivation for this spillover test is to assess the plausibility of the exclusion restriction, it also offers a general test of oligopsonistic wage setting that can be applied to other settings.

Estimating equation I start with an accounting identity that links an employer j 's log wage to the log of its marginal product and markdown:

$$\ln w_{jt} = \ln mrpl_{jt} + \ln \mu_{jt}$$

Proposition 3 For any competition structure among employers and any invertible labor supply system where firms are not perfect substitutes, \exists a log markdown function $\Lambda_j(w_{jt}, \mathbf{w}_{-jt}; a_{jt}, \mathbf{a}_{-jt})$, such that the firm's profit-max wage \tilde{w}_{jt} is the solution to the following fixed point, for any given wage vector among its competitors \mathbf{w}_{-jt} :

$$\ln \tilde{w}_{jt} = \ln mrpl_{jt} + \Lambda_j(\tilde{w}_{jt}, \mathbf{w}_{-jt}, a_{jt}, \mathbf{a}_{-jt}) \quad (7)$$

Proof. See below. ■ That j 's optimal markdown Λ_j can depend on other employer $-j$'s wages reflects the potential for wage spillovers. To derive the estimating equation, I totally differentiate and re-arrange equation (9):

$$dlnw_{jt} = \frac{1}{1 + \Gamma_{jt}} dlnmrpl_{jt} + \frac{\Gamma_{-jt}}{1 + \Gamma_{jt}} dlnw_{-jt} + \xi_{jt}$$

where $\Gamma_{jt} := -\frac{\partial \Lambda_j(w_t; a_t)}{\partial w_{jt}}$ captures how the optimal markdown changes with j 's own wage, $\Gamma_{-jt} := \sum_{j' \neq j} \frac{\partial \Lambda_j(w_t; a_t)}{\partial w_{j't}}$ captures how the optimal markdown changes with other employers' wages, $\xi_{jt} := \frac{1}{1 + \Gamma_{jt}} \sum_{j=1}^N \frac{\partial \Lambda_j(w_t; a_t)}{\partial a_{jt}} da_{jt}$ captures changes in employer-specific preferences or amenities, and $dlnw_{-jt} = \sum_{j' \neq j} \omega_{-j't} dlnw_{j't}$, is a weighted sum of wage changes at other employers $-j$, with weights equal to $\omega_{jj't} = \frac{\partial \Lambda_j(w_t; a_t) / \partial w_{j't}}{\sum_{k \neq j} \frac{\partial \Lambda_j(w_t; a_t)}{\partial w_{kt}}}$. I term $dlnw_{-jt}$ the competitor wage index. I show below that these weights ($\omega_{jj't}$) are proportional to competitor j' shares when the log expenditure function z_t is a sufficient statistic for competitor wages, i.e. labor supply to j can be denoted as $n_j(w_j, z_t; a_t)$ (this is true, for example, of nested CES). The empirical equation for estimating spillovers is:

$$\Delta lnw_{jt} = \delta \Delta lnmrpl_{jt} + \gamma \Delta lnw_{-jt} + \xi_{jt} \quad (8)$$

Therefore, under price or quantity competition when the expenditure function is a sufficient statistic for competitor wages (including Cournot or Bertrand oligopsony and nested CES supply), δ measures the pass-through of an employer's own marginal product shock, holding competitor wages constant, and γ measures the strength of strategic complementarities in wage setting. In typical monopsony models, where an atomistic monopsonist faces constantly elastic labor supply (e.g. Card et al. 2018), shocks to a firm's marginal product always exhibit complete pass-through since the wage is always proportional to marginal product ($\delta = 1$). However, when employers are large as opposed to atomistic, large employers face less elastic labor supply and need not pass-through all productivity gains to workers ($\delta < 1$). Typical models also abstract away from wage spillovers, $\gamma = 0$; however, if large employers alter markdowns in the face of competitor wage changes that alter worker flows to them, $\gamma > 0$.

Identification To isolate the causal effect of changes in an employer's own productivity (δ) and of competitor wages (γ) on its own wage, I need two demand-side instruments orthogonal to changes in its labor supply from changing preferences (ξ_{jt}). I use an employer's own MFA treatment status to provide variation in its marginal product ($\Delta lnmrpl_{jt}$) and the number of treated workers at other employers in its geography to provide variation in competitor wages (Δlnw_{-jt}). The identifying variation to estimate γ compares employers in places where the MFA shock spurs a large change in competitor wages to those in places where it spurs a small change. Appendix Figure 4 plots variation in the instrument, showing that for most workers less than 10% of their labor market is treated by the MFA shock, but for about 20% of workers this fraction is over 10%. Table 4, Panel B shows a strong first stage (F-stat 47): for every additional one hundred MFA-affected workers in

one's labor market, competitor wages fall by 0.02%. Appendix Figure 5 shows the first stage instead using the share of the MFA-affected wage bill (F-stat 50): as it grows from 0 to 10%, competitor wages fall by 1%.

Measurement As when estimating elasticities, $\Delta \ln w_{-jt}$ and $\Delta \ln mrpl_{jt}$ measure the change in wages and marginal products between $t = -1$ and $t = 1$, and $\Delta \ln w_{jt}$ measures wage changes between $t = -1$ and $t \in \{3, 5\}$. I use $t \in \{3, 5\}$ to account for spillovers over time as workers flock to other employers, as opposed to immediately, and because the labor supply response I am trying to transform into an own-wage elasticity occurs in years three to five. I assume that marginal product is proportional to average product, with the same change in log terms. This choice is consistent with a Cobb-Douglas production function, as commonly used to model production in textile and garment manufacturing industry (Cajal, Macchiavelo, & Rossi 2019 for Bangladeshi garment manufacturing, Atkin et al. 2019 for Egyptian carpet making), but also accommodates non-constant returns to scale.⁴⁷

Results Table 4 reports results. To gauge how the average worker's wage responds to competitors' wage changes, Columns (1) - (3) report results for incumbent workers employed in year $t=-1$. The MFA shock has no wage spillovers among exporters (Column 1), among MFA-unaffected exporters (Column 2) or among all unaffected employers in the textile and clothing manufacturing industry including non-exporters (Column 3). Across samples, $\hat{\gamma}$ is a tightly estimated zero, ruling out changes over 0.01% with a high degree of confidence (95%). In addition to the average zero effect, randomization inference p-values for the sharp null show that we cannot reject the null of zero effect for every employer.

Because new worker wages may more accurately capture the wage workers would receive when choosing a China-competing or non-China-competing employer, column (4) reports spillover effects on the wages of new hires. Again, I reject changes higher than 0.01% with a high degree of confidence.

A key concern is measurement error, which may attenuate estimates of γ and lead to a false conclusion of no strategic responses when in fact they exist ($\gamma > 0$). Measurement error comes from two possible sources. First, I may mismeasure marginal products because I have misspecified the production function and marginal product is not proportional to average product. Second, even if the production function is correctly specified, I may mismeasure marginal products because I only observe exports and not also revenues. To mitigate this concern, I omit $\Delta \ln mrpl_{jt}$ from the regression in equation 6 (Columns 2 and 3). Under strategic wage responses, the true $\Delta \ln mrpl_{jt}$ and $\Delta \ln w_{-jt}$ are positively related: as j 's competitors lower their wage, workers flock to j , who can hire more workers at lower wages, lowering marginal product. Eliminating $\Delta \ln mrpl_{jt}$ from the regression should yield upward biased estimates of γ . Here, too, I cannot reject either the null of

⁴⁷Specifically, I assume a production function of the form $PF = \frac{PL^{1+\lambda}}{1+\lambda}$ with returns to scale parameter $\frac{1}{1+\lambda}$ such that $\Delta \log APL = \Delta \log MPL$.

no average effect or the sharp null of no effect on any employer with a high degree of confidence (95%).

Discussion The small estimates of γ suggest that the exclusion restriction in my setting is satisfied, namely, that labor supply to MFA-affected employers does not change because of wage spillovers to other employers. I view this result as signifying that the shock was too small to compel a competitor response. No spillovers following a small shock might reflect optimization frictions (Dube et al., 2020), nominal rigidities (Hazell and Taska, 2020), or other institutional constraints that prevent employers from changing their wage when a small number of other employers do so. Nonetheless, the result is exciting because it offers the first test of oligopsonistic wage setting in labor markets, and provides a general rubric for such tests in future work.

While I interpret no spillovers as suggesting a small shock as opposed to a verdict that large employers do not fully exercise their market power, the latter is a possibility. My evidence points against it for two reasons: first, employers' own productivity shocks do not change workers' wages one-for-one, and, second, this pass-through falls as they become larger (Appendix Figure C8). Both findings are consistent with oligopsony.

B.2 Proofs

Proposition 3 *For any competition structure among employers and any invertible labor supply system where firms are not perfect substitutes, \exists a log markdown function $\Lambda_j(w_{jt}, \mathbf{w}_{-jt}; a_{jt}, a_{-jt})$, such that the firm's profit-max wage \tilde{w}_{jt} is the solution to the following fixed point, for any given wage vector among its competitors \mathbf{w}_{-jt} :*

$$\ln \tilde{w}_{jt} = \ln mrpl_{jt} + \Lambda_j(\tilde{w}_{jt}, \mathbf{w}_{-jt}, a_{jt}, \mathbf{a}_{-jt}) \quad (9)$$

Proof. This proof closely derives from Amiti et al. (2019). I omit subscripts t for brevity. Writing the profit maximization problem of a firm in conjectural variation form:

$$\max_{w_j, \mathbf{w}_{-j}} \{R_j(\exp(n_j(w_j, \mathbf{w}_{-j}; \mathbf{a}))) - \exp\{w_j + n_j(w_j, \mathbf{w}_{-j}; \mathbf{a})\} \mid s.t. h_{-j}(w_j, \mathbf{w}_{-j}; \mathbf{a}) = 0\} \quad (10)$$

where R_j is the revenue function⁴⁸, n_j is the **log** labor supply and w_j is the **log** wage in firm j ⁴⁹, and h_{-j} is the conjectural variation vector function with elements h_{jk} for $k \neq j$. For monopolistic and Bertrand competition, this vector takes the form:

$$h_{-j}(w_j, \mathbf{w}_{-j}; \mathbf{a}) = w_{-j} - w_{-j}^*$$

⁴⁸I assume firms are price-takers in the product market, which is reasonable for textile and clothing products.

⁴⁹Also for brevity, but different from the derivations above where they are levels.

and for Cournot competition it takes the form:

$$h_{-j}(w_j, \mathbf{w}_{-j}; \mathbf{a}) = -(n_{-j}(w_j, \mathbf{w}_{-j}; \mathbf{a}) - n_{-j}^*)$$

A conditional profit maximization with respect to wages therefore captures firm behavior under competition in prices and quantities. Introducing four pieces of notation:

1. $e^{w_j+n_j} \lambda_{jk}$ for $k \neq j$ represents the set of Lagrange multipliers for the constraints in (10).⁵⁰
2. $\zeta_{jkl}(\mathbf{w}, \mathbf{a}) := \frac{\partial h_{jk}(\mathbf{w}, \mathbf{a})}{\partial w_l}$, which is the elasticity of the conjectural variation function with respect to l 's wage (supply invertibility ensures invertibility of the matrix $\zeta_{jkl}(\cdot)\}_{k,j \neq l}$)
3. $\epsilon_j(\mathbf{w}, \mathbf{a}) := \frac{\partial n_j(\mathbf{w}, \mathbf{a})}{\partial w_j}$ and $\delta_{jk} := \frac{\partial n_j(\mathbf{w}, \mathbf{a})}{\partial w_k}$ are the own and cross-price elasticities of supply.

The FOCs of firm j 's profit maximization problem are then given by:

$$\begin{aligned} -1 - \epsilon_j + \epsilon_j e^{\mu_j} + \sum_{k \neq j} \lambda_{jk} \zeta_{jkj} &= 0 \\ -\delta_{jl} + \delta_{jl} e^{\mu_j} + \sum_{k \neq j} \lambda_{jk} \zeta_{jkl} &= 0 \quad \forall l \neq j \end{aligned}$$

where $\mu_j = mrpl_j - w_j$ is the log markdown. Solving out the Lagrange multipliers yields j 's optimal markdown:

$$\mu_j = \log \frac{\sigma_j}{\sigma_j + 1}$$

where σ_j is the perceived elasticity of labor supply and is given by:

$$\sigma_j := \epsilon_j - \zeta_j' \mathbf{Z}_j^{-1} \delta_j$$

where $\zeta_j := \{\zeta_{jkl}\}_{k \neq j}$, $\delta_j := \{\delta_{jkl}\}_{k \neq j}$, and $\mathbf{Z}_j := \{\zeta_{jkl}\}_{k \neq j, l \neq j}$ is an $(N-1) \times (N-1)$ matrix of cross-price elasticities (which has full rank because the labor supply system is invertible). Because all three are only functions of the wages and amenities at employers, it therefore follows that the optimal markdown is as well:

$$\Lambda_j(w_j, \mathbf{w}_{-j}, a_j, \mathbf{a}_{-j}) := \log \frac{\sigma_j(w_j, \mathbf{w}_{-j}, a_j, \mathbf{a}_{-j})}{\sigma_j(w_j, \mathbf{w}_{-j}, a_j, \mathbf{a}_{-j}) + 1}$$

and the firm's profit-maximizing wage solves the fixed point (adding back t subscripts):

$$\ln \tilde{w}_{jt} = \ln mrpl_{jt} + \Lambda_j(\tilde{w}_{jt}, \mathbf{w}_{-jt}, a_{jt}, \mathbf{a}_{-jt})$$

This proves the proposition. ■

⁵⁰This definition allows cancelling out the $e^{w_j+n_j}$ term from both sides of the FOC below.

Proposition 4 *If the log expenditure function z_t is a sufficient statistic for competitor wages, i.e. supply can be written as $n_{jt} = n_j(w_{jt}, z_t; \mathbf{a}_t)$ then the weights in the competitor wage index ($d \ln w_{-jt} = \sum_{j' \neq j} \omega_{-j't} d \ln w_{j't}$) are proportional to the competitor revenue market shares $S_{j't}$ for $j' \neq j$ and given by $\omega_{jj't} = \frac{S_{j't}}{1 - S_{jt}}$. Thus, the competitor wage index simplifies to:*

$$\sum_{j' \neq j} \frac{S_{j't}}{1 - S_{jt}} dw_{j't}$$

Proof. Following the steps from before, the optimal markdown is given by:

$$\Lambda_j(w_{jt}, z_t; \mathbf{a}_t) := \log \frac{\sigma_j(w_{jt}, z_t; \mathbf{a}_t)}{\sigma_j(w_{jt}, z_t; \mathbf{a}_t) + 1}$$

The weights in the competitor price index are therefore:

$$\begin{aligned} \omega_{jj't} &= \frac{\partial \Lambda_j(w_{jt}, z_t; \mathbf{a}_t) / \partial w_{j't}}{\sum_{k \neq j} \frac{\partial \Lambda_j(w_{jt}, z_t; \mathbf{a}_t)}{\partial w_{kt}}} \\ &= \frac{\partial \Lambda_j(w_{jt}, z_t; \mathbf{a}_t) / \partial z \frac{\partial z}{\partial w_{j't}}}{\partial \Lambda_j(w_{jt}, z_t; \mathbf{a}_t) / \partial z \sum_{k \neq j} \frac{\partial z}{\partial w_{kt}}} \\ &= \frac{S_{j't}}{\sum_{k \neq j} S_{kt}} = \frac{S_{j't}}{1 - S_{jt}} \end{aligned}$$

This relies on the fact that $\frac{\partial z_t}{\partial w_{jt}} = S_{jt} \forall j$ (Shepard's lemma). With (two nested) CES supply, for example, the log expenditure function is just the log wage index of the market, and $\frac{\partial \log W_{kt}}{\partial w_{jt}} = S_{jt}$ (where note that w_{jt} is still the log wage). ■

C Estimation

C.1 Parameters

C.1.1 Method 1

To estimate θ_g I take the partial derivative of labor supply to an industry with respect to an employer j 's wage:

$$n_{gkr} = \left(\frac{W_{gkr}}{W_{gr}} \right)^{\theta_g} \left(\frac{W_{gr}}{W_g} \right)^{\lambda_g} N_g$$

$$\frac{\partial \ln n_{gkr,j}}{\partial \ln w_{gj}} = \theta_g s_{gj} + (\lambda_g - \theta_g) s_{gj} s_{gk} = \theta_g s_{gj} \text{ when } s_{gk} \sim 0$$

Intuitively, η_g is the elasticity of labor supply to an atomistic employer, θ_g the elasticity of labor supply to an atomistic industry (from a large employer in that industry), and λ_g the elasticity of labor supply to an atomistic geography (from a large employer in that geography). Heterogeneous

preferences are the only force tethering workers to these atomistic employers, industries, and geographies. These partial derivatives are not typically estimable in the presence of wage spillovers or amenity changes. However, as above, this partial equilibrium elasticity is typically inestimable with a reduced form regression in the presence of spillovers. To derive the reduced form elasticity I start from the labor supply system in a geography where the industry is small:

$$\ln n_{gkr} = \theta_g \ln \bar{W}_{kgr} + \text{Aggregates}$$

Consider a first-order approximation around the Nash equilibrium, following any change to firms in the industry:

$$\begin{aligned} \sum_{j' \in k,r} \frac{\partial \ln n_{gkr}}{\partial \ln n_{gkr,j'}} \Big|_{n_{j'=j'}^*} \Delta \ln n_{gkr,j'} &= (\theta_g) \sum_{j' \in k,r} \frac{\partial \ln \bar{W}_{kgr}}{\partial \ln w_{gj'}} \Big|_{w_{j'=j'}^*} \Delta \ln w_{j'} + \Delta \text{Amenities} \\ \Delta \ln n_{gkr,j} + \sum_{j' \neq j \in k,r} \Delta \ln(n_{gkr,j'}) &= \theta_g s_{gj} \Delta \ln w_j + \theta_g \sum_{j' \neq j \in k,r} s_{gj'} \Delta \ln w_{j'} + \Delta \text{Amenities} \end{aligned}$$

I argue in the test of strategic interaction that $\Delta \ln w_{j'} = 0 \forall j' \neq j \in k, r$ and assume $\Delta \text{Amenities} = 0$. Similarly, Appendix Figure C2 shows roughly that $\sum_{j' \neq j \in k,r} \Delta \ln n_{gkr,j'} \sim 0$. Under both conditions, I can estimate the partial equilibrium elasticity through the reduced form $\epsilon_{gkr,j} = \frac{\Delta \ln n_{gkr,j}}{\Delta \ln w_{gj}} = \frac{\partial \ln n_{gkr,j}}{\partial \ln w_{gj}}$. The exclusion restriction may still be violated if there are multiple MFA employers in a geography. However, in half of all geographies there is only one MFA-treated employer and estimates closely resemble those in geographies with many.

Similarly, to derive the moment condition for λ_g I take the partial derivative of labor supply to a geography with respect to employer j 's wage:

$$\begin{aligned} n_{gr} &= \left(\frac{\bar{W}_{gr}}{W_g} \right)^{\lambda_g} N_g \\ \frac{\partial \ln n_{gr}}{\partial \ln w_{gj}} &= \lambda_g s_{gj} s_{gk} \end{aligned}$$

The rate of change of exiting a geography with employer size (in the geography) therefore estimates λ_g . Just as above, this can typically not be estimated under wage spillovers. Following the same steps as above:

$$\begin{aligned} \ln n_{gr} &= \lambda_g \ln \bar{W}_{gr} + \text{Aggregates} \\ \Delta \ln n_{gr,j} + \sum_{j' \neq j \in r} \Delta \ln n_{gr,j'} &= \lambda_g s_{gj} s_{gk} \Delta \ln w_j + \lambda_g \sum_{j' \neq j \in r} s_{gj'} s_{gk} \Delta \ln w_{j'} + \Delta \text{Amenities} \end{aligned}$$

The partial equilibrium elasticity is identified when $\sum_{j' \neq j \in r} s_{gj'} s_{gk} \Delta \ln w_{j'} \sim 0$ (no spillovers), $\Delta \text{Amenities} = 0$ (by assumption) and $\sum_{j' \neq j \in r} \Delta \ln n_{gr,j'}$ (small shock).

C.1.2 Method 2

I also adopt an alternate estimation strategy from Costinot, Donaldson, and Smith (2016) that requires shocks at the employer, industry and location levels. Because the MFA doesn't provide the requisite variation at industry and location levels, I instead use the shock from the Brazilian trade liberalization of the 1990s. Thanks to the lucid exposition of Felix (2022), who I follow almost exactly in this estimation, varying only the nesting structure, the estimating equations can be obtained by taking logs of the labor supply condition:

$$\log n_{gj} = \eta_g \log w_{gj} + \underbrace{(\theta_g - \eta_g) \log W_{gk,r}}_{\text{Industry} \times \text{Microregion} \times \text{Year FE}} + \underbrace{(\lambda_g - \theta_g) \log W_{gr}}_{\text{Microregion} \times \text{Year FE}} + \eta \log a_{gj}$$

Taking long differences and using an employer-level instrument yields an estimate of η_g . In other words:

$$\Delta \log n_{gj} = \eta_g \Delta \log w_{gj} + \Delta \delta_{im} + \Delta \delta_m + \epsilon_j$$

where a firm-level instrument can provide requisite variation in $\Delta \log w_{gj}$. To obtain an estimate of θ_g , observe:

$$\Delta \delta_{im} = \alpha + (\theta_g - \eta_g) \Delta \log W_{gk,r} + \epsilon_{gk,r}$$

an industry-level instrument provides the requisite variation in $\Delta \log W_{gk,r}$ to estimate θ_g . Similarly, to obtain an estimate of λ_g observe:

$$\Delta \delta_m = \alpha + (\lambda_g - \theta_g) \Delta \log W_{gr} + \epsilon_{gr}$$

a location-level instrument provides the requisite variation in $\Delta \log W_{gr}$.

I create instruments using the change in tariff exposure across industries, all of which were liberalized between 1991 and 1994. I instrument for $\Delta \log w_{gj}$ using $\Delta ICE_{j,k,r}$, instrument $\Delta \log W_{gk,r}$ using $\Delta ICE_{k,r}$, and $\Delta \log W_{gr}$ using ΔICE_r to obtain λ_g . The definitions are as follows, j is a firm, k is an industry, r is a region:

$$\begin{aligned} \Delta ICE_{j,k,r} &= \ln \left(\frac{1 + \tau_{j(k)1994}}{1 + \tau_{j(k)1991}} \right) \\ \Delta ICE_{k,r} &= \sum_{j \in k,r} \frac{s_{j,1991}^2}{\sum_j s_{j,1991}^2} \ln \left(\frac{1 + \tau_{j(k)1994}}{1 + \tau_{j(k)1991}} \right) \\ \Delta ICE_r &= \sum_{k \in r} \frac{s_{k,1991}^2}{\sum_{k \in r} s_{k,1991}^2} \sum_{j \in k,r} \frac{s_{j,1991}^2}{\sum_j s_{j,1991}^2} \ln \left(\frac{1 + \tau_{j(i)1994}}{1 + \tau_{j(i)1991}} \right) \end{aligned}$$

C.2 Estimating wage indices and amenities

I follow the following three steps to measure the industry-level wage index and infer amenity values in an industry for men and women. For this exercise I assume away establishment-specific amenities that deviate from industry norms.

Step 1. Estimating the wage index of an industry Re-arranging the expression for the wage bill share of a firm j in sector k :

$$s_{gj} = \frac{(w_{gj})^{1+\eta_g}}{\sum_{j'} (w_{gj'})^{1+\eta_g}}$$

$$\left[\sum_{j'} (w_{gj'})^{1+\eta_g} \right]^{\frac{1}{1+\eta_g}} = w_{gj} (s_{gj})^{\frac{-1}{1+\eta_g}} \forall j$$

Taking the log of both sides and summing over all j , the wage index of a industry k is:

$$W_{gk} = \tilde{w}_{gk} \tilde{s}_{gk}^{\frac{-1}{1+\eta_g}}$$

where \tilde{w}_{gk} is the geometric mean of wages and \tilde{s}_{gk} is the geometric mean of the share of the wage bill within textiles (i.e., exp of the mean of logs). The wage index is higher the higher the geometric mean of wages and the higher the dispersion in shares (the lower the geometric mean of shares).

Lemma 2

$$W_{gk} = \tilde{w}_{gk} \tilde{s}_{gk}^{\frac{-1}{1+\eta_g}}$$

where \tilde{w}_{gk} is the geometric average of wages for group g in sector k and \tilde{s}_{gk} is the geometric average of wage bill shares.

Proof. We know that:

$$s_{gj} = \frac{w_j n_j}{\sum_{j' \in k, r} w_{j'} n_{j'}} = \frac{(w_{gj})^{1+\eta_g}}{\sum_{j' \in k, r} (w_{gj'})^{1+\eta_g}}$$

Thus:

$$\left[\sum_{j' \in k} (w_{gj'})^{1+\eta_g} \right]^{\frac{1}{1+\eta_g}} = w_{gj} (s_{gj})^{\frac{-1}{1+\eta_g}} \forall j$$

Taking the log of both sides and summing over $j' \in k$:

$$\ln W_{gk} = \frac{1}{N} \sum_{j' \in k} \ln(w_{gj'}) + \frac{1}{N} \sum_{j' \in k} \ln(s_{gj'})^{\frac{-1}{1+\eta_g}} \ln W_{gk} = \ln((\Pi w_{gj'})^{\frac{1}{N}}) + \frac{-1}{1+\eta_g} \ln((\Pi s_{gj'})^{\frac{1}{N}})$$

Note that the geometric mean is defined:

$$(\Pi w_{gj'})^{\frac{1}{N}} = \exp\left(\frac{1}{N} \sum \ln w_{gj'}\right)$$

Therefore:

$$\ln W_{gk} = \ln(w_{gk}) - \frac{1}{1 + \eta_g} \ln(s_{gk})$$

Exponentiating both sides gives the result:

$$W_{gk} = w_{gk} s_{gk}^{\frac{-1}{1+\eta_g}}$$

■

Step 2. Estimating industry-specific amenity values for men Given estimates of the wage index by industry and geography and θ_m , the amenity values for men can be inferred from the share of men in each sector, normalizing the amenity value for men to 1 in textiles.

$$\frac{s_{mk}}{s_{mtxt}} = \frac{a_{mk}^{1+\theta_m} \frac{W_{mk}^{1+\theta_m}}{W_m^{1+\theta_m}}}{\frac{W_{mtxt}^{1+\theta_m}}{W_m^{1+\theta_m}}}$$

Step 3. Estimating women’s amenities relative to men Given estimates of the wage index by industry for each gender, (θ_m, θ_w) , and observed shares, the amenity values for women relative to men can then be inferred from the share of women in each industry relative to the share of men:

$$\frac{s_{wk}}{s_{mk}} = \frac{a_{wk}^{1+\theta_w} \frac{W_{wk}^{1+\theta_w}}{W_w^{1+\theta_w}}}{a_{mk}^{1+\theta_m} \frac{W_{mk}^{1+\theta_m}}{W_m^{1+\theta_m}}}$$

C.3 Estimating industry-specific productivity

I start with the following value added production function that is Cobb-Douglas in labor and capital with labor a CES aggregation of male and female workers: $Y_{jt} = z_{jt} K_{jt}^{\alpha_{k1}} l_{jt}^{\alpha_{k2}}$ with $l_{jt} = [\beta_k f_{jt}^\sigma + m_{jt}^\sigma]^{\frac{1}{\sigma}}$.⁵¹ I estimate the industry-specific parameters $\{\alpha_{k1}, \alpha_{k2}, \beta_k\}$ by employing standard production function techniques due to Akerberg et al. (2015) and calibrating σ to Gallen (2018). I describe my assumptions and procedure here.

A firm j picks capital for the subsequent period ($k_{j,t+1}$), labor in the current period (l_{jt}), and materials in the current period (x_{jt}) to maximize profits after observing its productivity shock ($\omega_{jt} := \log(z_{jt})$) that is unobservable to the econometrician. In other words, capital is a state variable (investment i_{jt-1} that determines k_{jt} is picked in $t - 1$) and supplied competitively in a rental market. Labor and material inputs are flexible (l_{jt} and m_{jt} are picked in period t). There is monopsony in the labor market but none in the market for materials. I make the following standard assumptions:

⁵¹One interpretation of this value-added production function is that gross output is Leontief in materials (Akerberg et al., 2015).

Assumption 1 Productivity evolves according to a first-order Markov process.

$$\omega_{jt} = f(\omega_{jt-1}) + \zeta_{jt}$$

which I take to be a polynomial of degree three in ω_{jt-1} .

Assumption 2 Scalar unobservable. For at least one flexible input, materials, the only unobservable factor in a firm's input demand function is productivity ω_{jt} .

Assumption 3 Strict monotonicity. A firm's input demand function for materials is strictly monotone in ω_{jt} .

Together, Assumption 2 and 3 imply that material input demand can be inverted to obtain ω_{jt} , which is necessary for Step 1 below. The estimation then proceeds in three steps:

Step 1 Purge output of measurement error, relying on the invertibility of input demand to obtain productivity. Let the log of output be observed with some measurement error τ_{jt} :

$$y_{jt} = f(v_{jt}; \beta) + \omega_{jt} + \epsilon_{jt}$$

where $v_{jt} := \{f_{jt}, m_{jt}, k_{jt}\}$ and $\beta := \{\alpha_{k1}, \alpha_{k2}, \beta_k\}$. By assumption 2:

$$x_{jt} = x_t(\omega_{jt}, k_{jt}, l_{jt})$$

By assumption 3:

$$\omega_{jt} = h_t(x_{jt}, k_{jt}, c_{jt})$$

where c_{jt} are controls that can influence input decisions⁵². Thus we have:

$$y_{jt} = f(v_{jt}; \beta) + h_t(x_{jt}, k_{jt}, c_{jt}) + \tau_{jt}$$

And I can obtain estimates for output net of measurement error by running a non-parametric regression (e.g., a high-order polynomial, in my case approximated by one of degree four) in only observables $\{k_{jt}, l_{jt}, x_{jt}\}$. Let this output be denoted $\phi_{jt}(v_{jt}, x_{jt}; \beta)$.

Step 2 Construct estimates of productivity, relying on timing assumption 1 and some guess of β (simultaneously determined in step 3):

$$\omega_{jt}(\tilde{\beta}) = \phi_{jt} - f(v_{jt}; \tilde{\beta})$$

⁵²In practice, I only use year fixed effects.

Step 3 Use GMM to recover the β parameters, relying on the timing of input choice to construct instruments:

$$E \left(\zeta_{jt}(\beta) \begin{pmatrix} f_{jt-1} \\ m_{jt-1} \\ k_t \end{pmatrix} \right) = 0$$

This yields three moments to estimate three parameters $\beta := \{\alpha_{k1}, \alpha_{k2}, \beta_k\}$.

Because the CES labor aggregation can be approximated by a translog function, we have for the functional form of f :

$$\begin{aligned} \ln Y_{jt} &= \ln z_{jt} + \alpha_{k1} \ln K_{jt} + \alpha_{k2} (\beta_k \ln f_{jt} + \ln m_{jt} + \sigma \beta_k (\ln f_{jt} + \ln m_{jt} - \ln f_{jt} \ln m_{jt})) \\ \ln Y_{jt} &= \ln z_{jt} + \alpha_{k1} \ln K_{jt} + (\alpha_{k2} \beta_k + \alpha_{k2} \sigma \beta_k) \ln f_{jt} + (\alpha_{k2} + \alpha_{k2} \sigma \beta_k) \ln m_{jt} - (\alpha_{k2} \sigma \beta_k) \ln f_{jt} \ln m_{jt} \end{aligned}$$

A caveat to the proxy variable method for estimating production functions is a potential violation of the “scalar unobservables” assumption under market power, preventing me from inverting input demand to recover productivity shocks (Bond et al. (2021), De Loecker and Syverson (2021)). Intuitively, this is because input demand cannot be inverted in Step 1 without additional controls since materials demand x_{jt} depends on establishment-specific input prices and shares (in c_{jt}). However, two facts mitigate the concern that bias in $\hat{\beta}_k$ leads me to falsely conclude that amenities and not productivity drive women to textiles. First, Yeh et al. (2022) show using simulations that this method recovers the true parameters in its 95% confidence interval and delivers tight standard errors. Second, gender gaps in wages plus contracted amenities predict a much smaller difference in the relative female/male share in industries than observed — for example, they predict similar ratios in textiles and the manufacturing of metal products. In reality, the female share/male share in the former is 6 and in the latter is 0.2.

I suffer an important data limitation because I do not observe workers separately by gender in the Brazilian production data. I thus estimate production functions in two other developing countries: India and Chile. For India, I use the Annual Survey of Industries between 2000 and 2010. The ASI is a manufacturing survey collected for a panel of large manufacturing establishments with over 100 workers. For Chile, I use the census of manufacturing plants covering all plants with over 10 employees between 1979 and 1996. The number of establishments used in estimation range between 295 and 2331. These are above the 250 threshold used in Demirer (2020) to obtain consistent estimates and tight standard errors.

D Data appendix

D.1 Sample construction

Establishment sample: MFA analysis To analyze the effect of the Multi-Fiber Arrangement on workers’ wage and employment outcomes I construct a sample of all establishments that exported

textile and clothing products in 2004 – with the first two digits 61 or 62. There are 751 such establishments in total, employing 114205 workers as of 2004.

Incumbent worker sample: MFA analysis Incumbent workers defined as those employed at a treated or comparison establishment as of 2004 (based on the establishment sample). Their treatment status depends on the treatment status of their baseline (2004) employer, as described above. Leveraging the linked employer-employee feature of RAIS, incumbent workers are tracked across jobs from 2000 to 2009 — that is, I do not restrict the sample to job spells at employers in the establishment sample. In constructing this sample, I only consider the “main job spell” for each worker in each year. I define the “main job spell” as a worker’s longest spell during the year. In case all job spells have the same duration, I break ties by keeping only one spell at random.

Sample for wage and amenity indices I use the full sample of workers without high school degrees in 2004 to calculate wage indices, shares, and infer amenity indices.

D.2 Outcomes

I briefly discuss how I define the outcomes used for analyses at the establishment-level and incumbent worker-level. While I use a worker’s main job spell for all worker-level outcomes, some establishment-level outcomes are constructed using all job spells. I first describe establishment-level outcomes derived from all job spells and then those derived using main job spells. Finally, I 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.
- 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.
- Education of new hires: Average years of schooling of workers newly hired at an establishment in a given year, separately by gender.
- Share of female new hires: Share of new hires that are female in any given year.

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.
- Wage bill share in industry. I divide an establishment’s wage bill by the overall textile wage bill in the geography in 2004. I do this separately by gender and only among workers without high school degrees.

Worker level outcomes - main job spell:

- 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 2004.
- 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.
- Employed at new textile employer A dummy that indicates whether the worker is observed working at a non-baseline employer in the textile industry in any given year, where the baseline employer is defined as the (main) establishment of employment in 2004.
- Employed at new non-textile employer A dummy that indicates whether the worker is observed working at a non-baseline non-textile employer in the textile industry in any given year, where the baseline employer is defined as the (main) establishment of employment in 2004.
- Employed in new geography A dummy that indicates whether the worker is observed working at a new employer in a different microregion in any given year, where the baseline employer is defined as the (main) establishment of employment in 2004.

D.2.1 Constructing O*NET-based skill transferability

I follow Macaluso (2022) to define O*NET-based measures of skill transferability for men and women. O*NET reports the skill level (from 1-8) on each of 35 skills required to do any job. These data are based on interviews with thousands of US-based workers and include a range of skills including coordination, operation and control, equipment maintenance, and materials management. O*NET data have been heavily employed to build indices of skill relatedness (Gathmann et al. 2010, Neffke et al. 2013). Most pertinently, Macaluso (2022) demonstrates that O*NET-based skill transferability predicts US workers’ wage recovery following job loss. The measure of skill transferability from occupation o in geography m is:

$$st_{om} = - \sum_{k \in O} \omega_{km} d_{ok}$$

where O is the set of all occupations, ω_{km} is the share of jobs in geography m in occupation k and d_{ok} is the “skill-distance” (L1 norm taken over 35 tasks) between o and k .

Figure 6, Panel A shows that treated men and women have overlapping distributions of skill transferability. This provides prima facie evidence that gender differences in skills do not explain gender differences in exit.

Appendix Table 4 shows that skill transferability is strongly predictive of switching industries among treated men: those in the bottom eight deciles exit their employer 4% points more compared to those in the comparison group. This rate, however, rises to 7.9% points among treated men with skills in the top two deciles of skill transferability.

Table 4 shows, however, that controlling for gender differences in skill transferability does not meaningfully change the gender difference in treatment effects on exiting one’s employer or industry.

E Causally inferring the boundaries of workers’ labor market

This section outlines a data-driven method to quasi-experimentally uncover the boundaries of men and women’s labor market following a grouping method analogous to Almagro and Manresa (2021). This analysis lets me answer two additional questions: (i) where do the men go? and, (ii) are men more likely than women to switch to jobs that use their skills? The authors observe that, in a discrete choice setting where consumers choose first a nest and then an alternative within a nest, the following holds for all alternatives within the same nest:

$$\log \frac{P_j}{P_0} = \beta^k x_j + \lambda_k$$

where x_j is an attribute of j , β^k and λ_k are a group (k)-specific slope and intercept, P_j is the probability of choosing alternative j , P_0 is the probability of choosing the outside option, and λ_k is the inclusive value of choosing nest k . Using kmeans clustering, they use this identity to group alternatives j into nests $k \in \{1, \dots, K\}$.

I use a slightly different but analogous argument. Consider women choosing over two occupation nests: a women tailor’s nest and a non-women tailor’s nest with elasticity of substitution η within each nest but no substitution across nests. The cross-price elasticity with respect to tailoring wages is equal for all occupations within a nest. For those in the nest of tailors it is:

$$\frac{\partial \ln n_{ok}}{\partial \ln w_{tk}} = -\eta s_{tk}$$

where s_{tk} is the share of tailoring jobs among all jobs in nest k .

I use a k-means algorithm to group together occupations with the same cross-price elasticity. The key challenge is obtaining an unbiased estimate of the cross-price elasticity. This is impossible absent an instrument for the change in tailoring wages. Simply calculating $\frac{d \ln n}{d \ln w_t}$ yields a biased estimate because switches into occupations are correlated with unobserved changes to these occupations. I use the MFA to provide instruments.

Formally, I estimate the following IV, just for the sample of tailors, separately by gender, for

each occupation o :

$$Inocc_{iot} = \alpha \Delta \ln w_{ig} + \gamma_{mt} + \phi_{igt} \quad (11)$$

$$\Delta \ln w_{ig} = \beta D_i + \gamma_{mt} + \psi_{igt} \quad (12)$$

which is analogous to the expression for estimating labor supply elasticities, with the outcome $Inocc_{iot}$ instead being 1 if worker i is observed in occupation o in year t . Just as there, $\Delta \ln w_{ig}$ is the change in a worker's offered wage between $t = -1$ and $t = 1$, D_i equals one for MFA-treated workers (instrument), and γ_{mt} are geography (microregion)-year fixed effects. I use three years of data (years three to five) and interpret these as five-year elasticities.

An ongoing econometric challenge is feasibly estimating this IV for a large number of occupations.

With these estimates, I can use kmeans clustering (Lloyd 1967, Forgy 1965) to group together occupations with the same cross-price elasticity into the same nest.

1. Start with a guess for each nest k : $dlnn_1^0, dlnn_2^0$. Practically, I set 0 and $dlnn$ for a random occupation with positive moves into it as the two initial values.
2. For all occupations o , and any step s with guesses $dlnn_1^s, dlnn_2^s$, compute the cluster that minimizes:

$$k(o)^{s+1} = \arg \min_{k \in \{1,2\}} (dlnn_o - dlnn_k)^2$$

3. Compute the mean $dlnn_1^1, dlnn_2^1$ using realized changes in the occupation.
4. Repeat until nest membership remains stable.