Gender and the Misallocation of Labor Across Countries*

Nava Ashraf† Oriana Bandiera‡ Virginia Minni§
Víctor Quintas-Martínez¶

June, 2023

Abstract

We study how the gendered division of labor inside and outside the home affects productivity by restricting the pool of talent firms hire from. Using the personnel records of a multinational firm, we show that the performance of female employees is higher in countries and cohorts where the ratio of women to men working outside the home is low. This is in line with gender barriers generating positive selection: the higher the barriers, the higher the productivity that women need to overcome them. This implies that gender specific barriers prevent the firm from hiring women who would increase productivity. We estimate that without these barriers productivity would be 32% higher for the same level of employment and wage bill. The findings suggest that selection is a powerful lens to understand the link between diversity and productivity.

*We thank our partners at the MNE for sharing the data, Sam Bowles, Stefano Caria, Anton Heil, Jonas Hjort, Erik Hurst, Rocco Macchiavello, Andrea Prat, Imran Rasul and seminar participants at CIVOE, Harvard, HBS, IIES, IZA, LSE, MIT Sloan, NBER Summer Institute, Northwestern, Princeton, Toulouse, UCL, UMichigan for useful feedback. We thank Rishabh Malhotra, Donata Schilling, Misha Seong for great research assistance. All errors are our own.

†STICERD and Department of Economics, LSE. n.ashraf1@lse.ac.uk
‡STICERD and Department of Economics, LSE. o.bandiera@lse.ac.uk
§STICERD and Department of Economics, LSE. v.m.minni@lse.ac.uk
¶Department of Economics, MIT. vquintas@mit.edu

1 Introduction

The division of labor inside and outside the home varies across countries and time, but it is always within the confines of norms that assign the largest share of housework to women (Jayachandran (2015); Fernández et al. (2021); Bursztyn et al. (2023)). This under-representation of women in spheres of influence and employment has led to significant efforts in both the private and public sectors to address the gap through extensive diversity initiatives (Bertrand (2011); Olivetti and Petrongolo (2016); Bertrand (2020)). Critics of these initiatives argue that this can encourage a lower quality bar for minority candidates and, in the corporate world, be ultimately worse for business by hiring or promoting lower productivity candidates. Supporters argue that diversity per se could be beneficial for productivity and profits due, for example, to the nature of the production function or role model effects (Lazear (1999); Athey et al. (2000); Hong and Page (2001)).

If we take as given that the distribution of innate talent is orthogonal to group identity, we cannot understand the effect of diversity on productivity without understanding how under-represented groups select into the labor force. On the one hand, by definition, members of under-represented groups who seek employment likely had to overcome higher barriers, which implies they would be positively selected (Lazear, 2021). If so, hiring from under-represented groups brings in more able workers and increases average productivity directly. However, ability is not the only factor that affects selection: differences in other factors unrelated to individual ability- for example, family wealth- could be the reason why individuals are able to overcome barriers and join the labor force.

In this paper, we examine the relationship between gender diversity and performance in the workplace through the lens of selection. To do so, we combine micro-level variation in earnings and career paths of approximately 100,000 employees at a large multinational company operating in 101 countries, with variation in the barriers that current employees faced when deciding whether to work outside the home. We proxy these barriers with the ratio of women to men in the labor force (henceforth LFPR) in the decade when the choice was made. Since we observe employees of different ages in the firm, we can exploit both cross-country and cross-cohort variation in barriers. Importantly, and as is well known, we observe variation in LFPR even among countries with similar levels of income, thus our measure of barriers does not solely capture the level of economic development.[1]

A simple Roy model of occupational choice illustrates what the correlation between performance in the firm and LFPR at the point of entry into the labor market tells us about selection. The model shows that when gender roles are strong so that women face a high cost to leave the home, only women whose returns are high enough will do so. This implies that the marginal, and hence the average, productivity of female employees is decreasing in the share of women working outside the home. The empirical counterpart of this result is that a negative correlation between women's performance in the firm and LFPR is evidence of positive selection, under the assumption that the relationship between performance (which we do not observe) and pay and promotion (which we do) is orthogonal to LFPR. This assumption is supported by the accounts of the firm's HR managers and is consistent with similarly centralized policies in the sample of 1,213 multinationals analyzed by Hjort et al. (2020).

We establish three facts. First, the variation in the LFPR is correlated with the female share of employees in the firm. In other words, in countries and cohorts with low LFPR, the firm hires fewer women. This is in line with the firm using the same selection process in all countries; that is, they do not employ more pro-women policies in countries where norms keep women inside the home (or vice versa). This is also essential for our purposes because if there were no correlation, the variation in LFPR would be moot.

Second, in low LFPR countries and cohorts, women are over-represented in the highest rungs of the hierarchy and are more likely to be promoted relative to their counterparts in high LFPR countries and cohorts. Relatedly, women are over-represented in the top decile of the wage distribution, and under-represented in the bottom decile, when LFPR is low. This suggests that the women who we observe in the firm in low LFPR countries had the ability to overcome higher barriers, and are thus positively selected.

Third, there is a negative correlation between women's average performance and the LFPR. Moreover, it comes from the bottom percentiles of the wage distribution: the wages of the women at the bottom decile of the wage distribution decrease as the

[1] For instance, Algeria and Ecuador had a similar level of GDP per capita in the 1980s but the former's LFPR was less than 1/4 of the latter's LFPR.

LFPR increases, while the wages of the women at the top decile remain constant. We interpret this as evidence that, as the LFPR increases, the ability of the marginal female employees falls, in line with the existence of an ability threshold below which women work inside the home. The fact that the negative correlation is driven by the left tail of the wage distribution rules out alternative explanations that would also generate a negative correlation between average female performance and LFPR. For instance, this pattern rules out that the high productivity of women in low LFPR countries, and hence their high wages, are due to the fact that women bring different valuable inputs to the firm and therefore the marginal value of these inputs is high when the share of women is low. If this were the case, we would find that the productivity of the top percentiles of women would decrease as LFPR increased. It also rules out the argument that women's performance increases relative to men's because men in low LFPR countries have better outside options. Besides being inconsistent with the distributional changes, the assumption that men's outside option is decreasing in LFPR is inconsistent with the fact that multinationals pay higher wages and offer more amenities across the board (e.g., Verhoogen (2008); Javorcik (2015); Hale and Xu (2016); Alfaro-Urena et al. (2019)).

The reduced form analysis aggregates individual-level data on outcomes at the level of variation of LFPR, that is at the cohort-country-gender level. While the findings are consistent with the simple Roy framework, unobservables at the cohort-country-gender level correlated with LFPR would produce observationally equivalent results. For instance, a negative correlation between men's outside option and the LFPR would also explain the patterns found. To address this, in the second part of the paper, we use the Roy model to estimate individual-level ability as well as the parameters of the firm's pay policy and the utility of working in the home. The structural estimates leverage the significant advantage of our data: we observe several employees in each country, gender, tenure, and cohort, allowing us to estimate both a fixed parameter common to all employees in the same cell (e.g. discrimination based on gender) and, using the variation in wages within cell, differences in individual productivity. Individual-level data allow us to separately identify gender differences in fixed pay from differences in pay due to differences in productivity. Moreover, this allows us to quantify the productivity cost of barriers faced by women in different countries and to study policy counterfactuals.

We assume that pay is equal to an average wage plus a reward proportional to marginal productivity. This gives us a relationship that links salary in the MNE to the cost of working outside the home. Both variables have a component that is common to all people of the same gender in the same cohort-country-tenure group (for instance, country-specific social norms about gender) and a component that is specific to the individual (their own productivity and preferences). Because we have individual-level pay data, we can separate the common component of salary from the reward for individual-specific performance. We use this to identify productivity thresholds that split the population into workers and homemakers.

Three findings are of note. First, the ability of the average female employee ranges from .9 SD when LFPR is at its lowest to .3 at its highest. The estimate is larger in countries with weaker gender equity labor laws and with more conservative gender norms. Men's ability, in contrast, does not change with LFPR. This is important for the interpretation of our results because it rules out that men are differentially selected, for instance, because in low LFPR countries their outside option is higher outside the MNE. Second, we can relax the assumption of equal preferences for work in and outside the home across genders, and use the model to back out the difference in preferences that would make the current LFPR gap optimal. The implied preference gap is several orders of magnitude larger than any other gender gap in preferences estimated in the literature to date. Third, we show that our estimated gender productivity gap in the MNE is correlated with the productivity of other firms in the economy, and more so with those that operate in the same sector as the MNE. To do this, we extract balance sheet data from ORBIS to cover all manufacturing firms in the same countries where the MNE operates, yielding a sample of 2 million firms in 158 SIC3 sectors.

Finally, we draw implications for the firm and for policy. We first show that, given the productivity differences between men and women, the firm could increase productivity for the same wage bill if they were to change the terms of the wage contract to attract more women. We find that the optimal contract has a lower base pay and a steeper performance gradient than the observed contract. This brings the firm's gender ratio close to one and increases productivity by 22%. However, we note that such a contract would significantly increase inequality within and between genders;

most notably, the difference in pay between women and men would go up by 78%. This captures both differences in performance for the same job and differences in jobs as more able women climb the corporate ladder faster. Thus, whilst it is theoretically possible for the firm to adopt policies that compensate for societal norms, such a steep performance gradient would create a high level of inequality among employees. The highest inequality within the firm would occur where there are the most restrictive gender norms. It may quite possibly also be unsustainable for the firm: in order to hire more women without excessively increasing inequality, they would have to increase women's pay without decreasing men's pay.

A direct implication of these results is that the MNE would benefit from the elimination of gender norms. Indeed, we show that equalizing barriers between genders would bring the pay gap to zero and would increase productivity by 32%, while keeping the wage bill and employment constant. The productivity gains result from both high-productivity women joining and low-productivity men leaving. This echoes the results in Hsieh et al. (2019), which finds that reducing misallocation by lowering barriers across gender and race groups accounted for 41.5% of the increase in GDP per capita in the United States between 1960 and 2010. We note that the mirror image of the gender tax that women have to pay to work outside the home is the tax that men have to pay to work inside the home, independent of their skills and preferences. Thus, eliminating gender norms will also eliminate misallocation in work inside the home, by allowing the men who wish to do so to specialize in home production.

In terms of the implications for public policy, the results cast new light on pro-worker labor policies and the gender earnings gap. We show that more stringent labor regulations may hurt women more than men, especially when the barriers to female work outside the home are higher. Intuitively, most pro-worker measures, such as restrictions on hiring and firing, make it harder to link pay to performance and this leads to a larger intake of lower-productivity workers (Propper and van Reenen, 2010) who, by selection, are more likely to be male in places with more restrictive gender norms.

Second, policies that aim to close the gender earnings gap are not sufficient to restore gender equality, due to the fact that the measures of earnings gaps do not account for differential selection. This is especially true when barriers to women's work outside the home are high because women's productivity is higher in these cases. In line with this, we find that the gap between female and male earnings monotonically decreases as LFPR increases. Like Olivetti and Petrongolo (2008), we show a similar pattern in aggregate wage data from the International Labor Organization (ILO). We use our estimates to "correct" the gap for differences in ability and show that the ability-adjusted gap is always larger and up to four times so than what is reported in official statistics. Restoring equality based on merit would result in an inverted gender pay gap in most countries.

Our findings show that selection creates a link between the size of a group and the productivity of its members, thus connecting the literature on the barriers to female labor force participation (Goldin (1995); Fernández et al. (2004); Jayachandran (2015); Olivetti and Petrongolo (2016)) to the literature on the impact of diversity for firm productivity (Alesina and La Ferrara (2005); Hamilton et al. (2012); Hjort (2014); Bertrand and Duflo (2017); Marx et al. (2021)). Seen through the lens of selection, the link between diversity and productivity is underpinned by the traits of the minority due to the barriers they had to overcome rather than a direct "treatment" effect of diversity on productivity through, for instance, role model effects or changes in culture. Via selection, the productivity of the firm increases when there are more women in the applicant pool because the firm does not need to hire from the left tail of the distribution of men; a similar pattern is seen in the selection of political candidates in Sweden following the introduction of a gender quota (Besley et al., 2017).

Our results also inform the literature on the evolution of the gender earnings gap (Goldin, 1990; Blau, 2012; Olivetti and Petrongolo, 2017). A key advantage of our data is that we observe women and men with the same education, same tenure, and working in the same function in the same firm. Hence the earnings gap is not influenced by differences in occupational choices that make comparisons between genders difficult (e.g., Blau (1977); Goldin (2014); Card et al. (2015); Wiswall and Zafar (2017); Andrew et al. (2021)). We can thus use it to study whether differential labor force participation rates lead to underlying differences in ability across genders in the working population.

The paper is organized as follows. Section 2 presents the institutional context of the MNE, and describes the data sources. Section 3 introduces the model and provides reduced form evidence on the link between the LFPR and differential performance in the firm between genders. In Section 4, we calibrate the parameters of the model from the firm's personnel data and country-cohort level LFPR and Section 5 uses our estimates to evaluate the effects of different counterfactuals. Section 6 concludes by discussing welfare implications and other issues for further research.

2 Context and data

2.1 Context

We collaborate with an MNE with headquarters in Europe and offices in more than 100 countries worldwide as illustrated in Figure A.1. The MNE produces consumer goods; in 2019, it had a turnover of €20+ billion and employed over 120,000 workers, of which approximately 55% were white collars. We focus on white-collar workers because blue-collar workers are only observed in two-thirds of countries where the MNE has production activities. Typical white-collar jobs in the MNE involve sales, engineering, marketing, HR, R&D for product development, and general managerial activities. The workers have homogeneous levels of human capital as applications require a college degree, and most employees have degrees in either business administration (50%) or engineering (20%).

2.2 Data

Personnel records: Our sample covers the universe of employees between 2015 and 2019. We focus our analysis on local employees (non-expats), resulting in 100,819 distinct regular full-time workers over 2015-2019 in 101 countries (303,756 employee-year observations). The company is organized into a hierarchy of work levels that goes from work level 1 to 6 (C-Suite). For each employee, we observe: (1) their work level; (2) their tenure in the firm and job; and (3) total compensation (fixed plus variable pay in euros). Pay captures differences in performance between employees, encompassing off-peak salary increases as well as promotions. We look at four 10-year age cohorts within the company, 18-29, 30-39, 40-49, and 50-59[2].

[2]We do not have more granular data on age cohorts because of data privacy clauses. Due to a limited sample size of workers in age cohorts above 50-59, we only consider workers up to the 50-59 cohort.

A recent strand of evidence shows that multinationals pay higher wages in the countries where they operate (Verhoogen, 2008; Javorcik, 2015; Hale and Xu, 2016; Alfaro-Urena et al., 2019; Hjort et al., 2020). We confirm that this is the case for our firm in Appendix Figure A.2 which shows that the firm's average wages are usually well above the countries' average wages, using both the average wages in the manufacturing sector from the ORBIS database and the ILO estimates for white-collar employees. Table 1 presents summary statistics separately by gender at the gender-cohort-country-tenure level, which is the relevant unit of analysis used in the structural estimation.[3] The female cells show lower average pay, age, tenure, probability of being in managerial positions, and probability of experiencing fast promotions. Overall in the company, 40% of workers are aged between 30-39 and the majority of workers are in WL1 (>70%).

Country-cohort level data: We combine the firm's administrative records with country-cohort data on labor force participation rates of men and women from the World Bank. In particular, we match the age cohorts in the firm with the average LFP rate in the country in the decade of labor market entry, separately by gender. For example, employees of age 18-29 are associated with the LFP rates of the 2010-2020 decade while employees of age 30-39 are associated with the LFP rates of the 2000-2010 decade, and so on.[4]

Figure 1 plots moments of the distribution of the labor force participation ratio of women's LFP against men's LFP (the LFPR)[5] at different deciles of GDP per capita. It shows that while the mean exhibits the well-known U-shape pattern (Goldin, 1995), the distributions largely overlap: there is variation in LFPR at every level of GDP per capita. For instance, the interquartile range of LFPR is broadly similar across the deciles of GDP per capita. Moreover, the 90th percentile of LFPR only ranges between 1 and 0.9. This indicates that there are country-cohort cells with high LFPR at every level of economic development, instead of being solely concentrated in high-income countries. This is essential for the analysis that follows because it allows us to partial

[3]As we explain in section 4, we let the structural parameters vary by gender, cohort, country, and tenure in the firm to control for confounders and allow for the wage policy to differ across countries, as well as take into account factors such as on-the-job learning.

[4]Since data from the World Bank on the LFP for the 1980-1989 decade is not available, we use a linear interpolation of the LFP for that decade.

[5]We rescale women's LFP by men's LFP in order to account for country-cohort level differences in the probability of individuals to work. However, we note that the variation in men's LFP is tiny compared to the variation in women's LFP. In other words, the variation in LFPR is mostly driven by the variation in the female labor force participation, the numerator in the ratio.

out the differences in national income among the countries where the MNE operates.

3 Framework and facts

3.1 Framework

To formalize our intuition about differential selection into the labor force, consider a basic two-sector Roy (1951) model (as formalized by Borjas (1987)). Suppose that the utility from working outside the home is equal to pay and that these are a (log-)linear function of individual i's productivity, A_i:[6]

$$y_i^1 = \alpha^1 + \beta^1 A_i, \tag{1}$$

The term α^1 is the unconditional average wage and β^1 is the return to productivity. We interpret deviations from that average wage as arising from individual differences in productivity.

Similarly, we model worker i's value of housework as:

$$y_i^0 = \alpha^0 + \nu^0 N_i, \tag{2}$$

where N_i captures sources of individual heterogeneity in the value of staying out of the labor force that are independent of productivity. Here, α^0 captures the unconditional average value of staying out of the labor force (e.g. social norms that affect all women).

We make the distributional assumption:

$$\begin{bmatrix} A_i \\ N_i \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right)$$

equally distributed in all our sample. A difference with the canonical Roy-Borjas model is that, as a result, y_i^1 and y_i^0 are independent. This assumption will allow us, in section 4, to identify the parameters of the model with individual-level data only on those who are working at the firm. It has also become a standard assumption in the

[6]This could be micro-founded, for example, by assuming that workers are paid their marginal product of labor, that all workers supply the same amount of hours (full-time) and that the individual production function is Cobb-Douglas, $F(K, l_i) = e^{z_i} K^{\alpha} l_i^{1-\alpha}$, where $z_i \sim \mathcal{N}(\mu, \sigma^2)$.

literature on misallocation with structural models of sectoral choice (e.g. Hsieh et al. (2019)).

Individual i self-selects to work outside the home if and only if:

$$y_i^1 \geq y_i^0$$

$$\Leftrightarrow \quad \eta_i \equiv \frac{\beta^1 A_i - \nu^0 N_i}{\sqrt{(\beta^1)^2 + (\nu^0)^2}} \geq \frac{\alpha^0 - \alpha^1}{\sqrt{(\beta^1)^2 + (\nu^0)^2}} \equiv \xi.$$

Because $\eta_i \sim \mathcal{N}(0,1)$, that happens with probability $1 - \Phi(\xi)$. It is straightforward to show that:

Result 1: *Stronger norms about gender roles, that is a higher α_0, increase ξ and therefore reduce the probability of working outside the home $1 - \Phi(\xi)$.*

Moreover, we have that:[7]

$$\begin{aligned}\mathbb{E}[A_i \mid \eta_i \geq \xi] &= \frac{\text{cov}(A_i, \eta_i)}{\text{var}(\eta_i)} \mathbb{E}[\eta_i \mid \eta_i \geq \xi] \\ &= \frac{\beta^1}{\sqrt{(\beta^1)^2 + (\nu^0)^2}} \frac{\phi(\xi)}{1 - \Phi(\xi)}.\end{aligned}$$

This implies that:

Result 2: *The average ability of those who choose to work outside the home is increasing in ξ.* Taken together with the previous result this underpins the link between productivity and diversity: higher barriers lead to fewer individuals joining the labor force and their average ability to be higher, all else equal.

The framework guides our analysis in the next section. First, we show that the female share in the firm closely follows the LFPR across countries and cohorts, providing empirical support that barriers at the country-cohort level constrain the firm's talent pool. Given this, we bring Result 2 to the data and analyze the correlation between women's performance and the LFPR.

[7] To see why, realize that we can write $A_i = \text{cov}(A_i, \eta_i)/\text{var}(\eta_i) \times \eta_i + u_i$, where u_i is the population OLS projection error. By the normality assumption, u_i and η_i are not just uncorrelated, but also stochastically independent.

3.2 Facts

Figure 2 documents a large variation in female shares of employment within the MNE across countries, which closely matches the variation in LFPR across countries: the female to male ratio in the firm closely follows the same ratio in the labor force. The relationship between the firm and country ratios does not vary with the level of LFPR, which is confirmed by a formal test of differential slopes by above/below median LFPR. This suggests that the selection of men and women in the MNE follows closely their LFP decisions and hence that the countries' LFPRs "bind".

Under our framework, an increase in the proportion of women in the firm, keeping firm size constant, would imply that the marginal woman hired is less productive than the average woman while the man she replaces is less productive than the average man. Hence, we would expect the average performance of women to fall and the average performance of men to rise as the LFPR increases.

Figure 3 displays plots of women's performance measures compared to men's against the labor force participation ratio. In line with this intuition, Panels (a) and (b) of Figure 3 show that women are over-represented at the top decile of the wage distribution and under-represented at the bottom decile when overall LFPR is low and converge when LFPR is close to 1. Moreover, Panels (c) and (d) show that when LFPR is low women are over-represented in managerial positions and among those who get promoted quickly, but the two converge as participation rates get more equal. We only show the plots with the women share as the y-axis, since, by definition, the plots with the men share would be their symmetric equivalents.

Our framework indicates that as the LFPR increases, women with lower ability enter the labor force while high-ability women already in the labor force are unaffected. We test this result empirically in Figure 4 by looking at how women's wages at different points of the distribution change as the LFPR increases. Since overall wage levels change across countries, we control for the respective wage measure for men. The first panel from the left shows that women's average wages decrease as the LFPR increases. The remaining two panels indicate whether this decrease is driven by the bottom or the top of women's wage distribution. It comes from the bottom of the wage distribution: the 10th percentile of women's wages decreases with LFPR (second panel), while there is no impact on the 90th percentile (third panel).[8]

These differential patterns at different levels of the wage distribution are consistent with women facing a higher bar for participating in the labor force in countries with lower LFPR. As the LFPR increases, lower-ability women start to enter, hence decreasing the wages at the bottom of the distribution, while high-ability women are not affected, hence leaving the wages at the top of the distribution unchanged.

Figure 4 is also useful to rule out an alternative interpretation that posits that women's numerical scarcity leads to women earning wage premia in environments where women are scant. According to this interpretation, women receive high wages in low LFPR settings not because employed women are more able in general, but rather because female-specific skills are scarce and hence command a higher price. Specifically, women's contribution would be particularly valuable in a low LFPR environment, where most employees are men. If this were true, women's wages should be negatively correlated with the LFPR, and even more so for the wages of high-ability women compared to low-ability women, which is not what we find in Figure 4.

4 Analysis

In this section, we take the model of worker selection into the labor force that we introduced in section 3 to the data. The value added of the model, relative to the reduced form estimates and to the existing literature, is that we can leverage our individual-level data to separately identify gender differences in fixed pay (due to, for instance, discrimination) from differences due to productivity. Typically, this cannot be estimated with aggregate data that only contain the average wage, whereas individual data can be used to compute the variance which maps directly to the variation in pay due to differences in productivity. By using individual-level data, we can account for unobservables at the cohort-country-gender level which are correlated with the LFPR, for e.g., if there is a correlation between LFPR and men's outside option or between LFPR and women's wage contracts because of discrimination. We first discuss how to calibrate the parameters of the model using the firm's personnel data and country-cohort level LFP. Next, we present the results of our calibration and validate them

[8] Results are robust to using other percentiles, for example, the 25th percentile.

against different variables and independent datasets.

4.1 Method

For our structural exercise, we will use the two-sector Roy-Borjas model presented in section 3. We let the parameters vary by country (c), cohort (a), and tenure in the firm (t)[9] and gender (g) cells: $\alpha^1_{gtac}, \beta^1_{gtac}, \alpha^0_{gtac}, \nu^0_{gtac}$. This flexibility allows us to control for a variety of confounders: for example, we do not make assumptions as to whether the firm has the same wage policy for men and women or across countries, and the fact that we estimate β^1 by tenure group allows for on-the-job learning in a way that is possibly correlated with productivity. As in section 3, we keep assuming that all *gtac* cells have the same underlying distribution of (A_i, N_i) at birth, jointly standard normal, but we allow the parameters in the wage and value of housework to be cell-specific. In our calibration exercise, we cannot distinguish differences in distribution from differences in $\alpha^1_{gtac}, \beta^1_{gtac}, \alpha^0_{gtac}$ and ν^0_{gtac}.[10]

In what follows, we make three additional assumptions. First, as with any other binary choice models, the scale parameter of the selection Probit of our Roy model is not identified. Hence, we normalize it to 1 for calibration purposes, so that $\text{Var}(\eta_{igtac}) = \text{Var}(\beta^1_{gtac}A_i - \nu^0_{gtac}N_i) = (\beta^1_{gtac})^2 + (\nu^0_{gtac})^2 = 1$.

Second, we do not have individual-level data for those not working at the firm, so we are not able to identify the correlation between A_i and η_{igtac} (as one would do with a standard Heckman (1976), selection model). Hence, we continue to assume that A_i and N_i are independent.[11]

Finally, and again because we only have data on those in our firm, we need to obtain a measure of the probability that a worker will want to work for our firm, i.e. $\Pr(y^1_{igtac} \geq y^0_{igtac})$. We proxy that by country-cohort-level LFP data. In other words, we are identifying the probability that an individual chooses to work for our firm with the probability that he or she chooses to work for *any* other firm. For that, we need to assume that working at our firm is weakly preferred to working at other firms

[9] We aggregate tenure in groups of two years each.

[10] Goodness of fit plots of the wage distribution against the distribution implied by the model (Figure A.4) show that the distributional assumption of normality fits the data reasonably well.

[11] This is a standard assumption in the literature on misallocation with structural models of sectoral choice. For example Hsieh et al. (2019) assume that productivity draws in each sector (including housework) are i.i.d. Fréchet.

for all individuals. We provide supporting evidence of this in Appendix Figure A.2 (discussed in section 2) and in Figure 2 (discussed in section 3).

Under our distributional assumptions, labor force participation and the observed wage satisfy the following moment conditions:

$$\Pr(\text{employed}) = 1 - \Phi(\xi_{gac}) \tag{3}$$

$$\mathbb{E}[y^1_{igtac} \,|\, \text{employed}] = \alpha^1_{gtac} + (\beta^1_{gtac})^2 \lambda(\xi_{gac}) \tag{4}$$

$$\text{Var}(y^1_{igtac} \,|\, \text{employed}) = (\beta^1_{gtac})^2 + (\beta^1_{gtac})^4 \left[\xi_{gac}\lambda(\xi_{gac}) - \lambda(\xi_{gac})^2\right] \tag{5}$$

where $\lambda(\cdot) \equiv \phi(\cdot)/(1 - \Phi(\cdot))$ is the inverse Mills ratio and $\xi_{gac} \equiv \alpha^0_{gtac} - \alpha^1_{gtac}$ are the participation thresholds. Together with the restriction $(\beta^1_{gtac})^2 + (\nu^0_{gtac})^2 = 1$, we have four equations in four unknown parameters $(\alpha^1_{gtac}, \beta^1_{gtac}, \alpha^0_{gtac}, \nu^0_{gtac})$ for each gender, tenure level, age cohort, and country cell.

In order to calibrate those parameters, we match the moments above to their empirical counterparts. Since we observe the wage for those working in the firm, we can use the sample average and variance as the empirical analogs for $\mathbb{E}[y^1_{igtac} \,|\, \text{employed}]$ and $\text{Var}(y^1_{igtac} \,|\, \text{employed})$. To eliminate the effect of observables, we use the residuals of a regression of log(base pay + bonus) on year and function dummies as our measure of y^1_{igtac}. To calibrate the parameters in the participation decision, we use World Bank LFP data in each gender, cohort, and country cell, LFP_{gac}, as our empirical analog for Pr(employed). Table 2 provides a summary of the parameters of the model and the empirical target that each parameter tries to match in our calibration strategy. Since we have three equations and three unknown parameters to calibrate, these are solved exactly and thus, these three moments are fit exactly. Figure A.4 shows that the distributional assumption of normality fits the data reasonably well beyond the moments that we targeted for calibration.

4.2 Results

Once we have obtained the parameters in the firm's wage policy, we can recover productivity as:

$$\widehat{A}_i = \frac{y^1_{igtac} - \alpha^1_{gtac}}{\beta^1_{gtac}},$$

where y^1_{igtac} is the residualized log-wage described before. Figure 5 plots the average calibrated productivity for our sample of firm workers by LFPR. Because we normalized $A_i \sim N(0,1)$, differences in productivity will be in units of standard deviations of the underlying A_i. Average productivity is approximately constant for male workers in our countries, whereas for female workers, average productivity is very high when FLFP is much smaller than MLFP (about 0.8 standard deviations higher than for men) and decreases as the LFPR approaches 1. This result is consistent with selection with respect to productivity, so the lower a group's LFP, the more positively selected they are.

Figure 6 further shows that this is due to a shift in the whole productivity distribution of women. As the LFPR increases, there is a downward shift of the entire distribution but the right tail.

If we believe that the underlying distributions of A_i and N_i in the population are the same for men and women, differences in LFP across genders must be due to different payoffs for working outside and inside the home, for e.g. due to gendered social norms. In other words, if men and women have the same distribution of productivity[12] and preferences for staying at home at birth, the only reason for their LFP to differ is that they must face different payoffs from working outside the home.

Alternatively, we can relax the assumption of equal preferences for housework and use the model to back out the difference in preferences (N_i) that would make the current LFP gap optimal. Suppose that all parameters of the model are the same for men and women except the average value of staying at home α^0. Further assume that the observed LFPs are optimal, i.e. the marginal man and woman have the same productivity. In that case, $\alpha^0_{Ftac} - \alpha^0_{Mtac} = \Phi^{-1}(MLFP) - \Phi^{-1}(FLFP)$, so we can back out the average difference in preferences from the difference in LFPs. Because we normalized $N_i \sim N(0,1)$, these will be in units of standard deviations of the underlying N_i.

Figure 7 plots these average preference gaps by country. Two points are of note: first, differences between genders are large - well over $1 S.D.$ for the top decile of countries, and at least $0.5 S.D.$ for most of the sample. Second, differences across countries are also large, for instance, the interquartile and the interdecile ranges are

[12] Gender differences in educational attainment have been drastically reduced over the last decades and are much smaller compared to the LFP gap (see Appendix Figure A.8).

0.45$S.D.$ and 1.18$S.D.$, respectively. These patterns are in sharp contrast with the existing evidence on gender differences in economic preferences across countries. Figure A.3 in the appendix shows the distribution of gender differences in risk aversion, altruism, trust, patience, and positive and negative reciprocity from Falk et al. (2018) for the same set of countries. These differences between genders are much smaller than what we find as well as the variation across countries. Overall, considering the interquartile range, for differences in preferences to rationalize our results we would need these to be at least three times larger than any other economic preference.

4.3 Validation

4.3.1 Comparison with external variables

To validate our estimates we show that they correlate with other data not used to calibrate the model. We use three sets of external variables to validate our three key estimates: (i) individual performance data from the firm's records to validate calibrated productivity; (ii) country-level labor laws that constrain the firm personnel policy to validate the parameters of the wage policy and (iii) country-level social norms to validate our estimate of the difference in home payoffs.

The results for the first exercise are in Appendix Figure A.5. Panel (a) validates calibrated productivity against the firm's performance appraisals that a manager gives every year, panel (b) against pay growth in the first year (for new hires), and panel (c) against an objective productivity measure for the sales department (based on reaching specific targets). The correlation of all three indicators with our calibrated productivity is positive and strong.

Appendix Figure A.6 panel (a) shows the results for the second exercise. It plots our calibrated β^1_{gtac} (which represents returns to productivity, and, in our model, is what generates dispersion in pay within gender-country-cohort-tenure cells) against the Restrictive Labor Regulations Index from the World Economic Forum.[13] Consistent with stricter labor regulation limiting performance pay, we find that our calibrated

[13]The WEF Restrictive Labor Regulations Index is available for the period 2008–2020 and it is based on an annual survey on the most problematic factors for doing business (e.g. corruption, taxes, inflation, etc.). The survey is administered to a representative sample of around 15,000 business executives in 150 countries. The Restrictive Labor Regulations Index includes measures related to labor-employer relations, wage flexibility, hiring and firing practices, performance pay, labor taxes, attraction and retention of talent.

β^1_{gtac} is lower in countries with a higher value of the index.

Appendix Figure A.6, panel (b) plots the gap in our calibrated $\alpha^0_{Ftac} - \alpha^0_{Mtac}$ (that we interpret as the cost of gender norms for the average woman) against the Women, Business and the Law Index from the World Bank.[14] The figure shows that the gap $\alpha^0_{Ftac} - \alpha^0_{Mtac}$ is strongly negatively correlated with laws allowing or facilitating women's labor. Therefore, part of the restrictions to FLFP due to gender norms may actually be embedded in the laws of certain countries.

Finally, Appendix Figure A.7 shows the gap in our calibrated average value of staying at home, $\alpha^0_{Ftac} - \alpha^0_{Mtac}$, against the responses to four questions in the World Value Survey that relate to gender norms: (1) "Men make better business executives than women do," (2) "Pre-school child suffers with working mother," (3) "Being a housewife is just as fulfilling as working," (4) "When jobs are scarce, men should have more right to a job than women." For all four questions, we see a strong positive correlation between agreement with the statement and our gap in the average value of staying at home.

4.3.2 Comparison with other firms

To gauge the external validity of our estimates, we extract balance sheet data for all manufacturing firms from the ORBIS database for our sample countries and years, and we test whether our estimate of the gender productivity gap (the average productivity of women minus the average productivity of men at the country level) correlates with the productivity of other firms in the economy. We obtain a cross-section of firms based on the latest year of balance sheet reporting between 2012 and 2019. The sample contains two million firms in 158 SIC3 sectors. Intuitively, our estimate of the gender productivity gap captures differences in productivity due to differences in barriers at the societal level, as well as idiosyncratic differences due to the MNE pay and recruitment policies. If the latter dominates, the estimated productivity gap will not be correlated with the productivity of other firms. Table 3 reports the estimates of the

[14] The WB Women, Business and the Law Index covers 190 countries through the period 1971–2020 and is structured around the life cycle of a working woman. It consists of eight indicators constructed around women's interactions with the law — mobility, pay, workplace, marriage, parenthood, entrepreneurship, assets, and pensions — for current laws and regulations (i.e. religious and customary laws are not considered unless they are coded). Hyland et al. (2021) provides an overview of the data documenting how gender discrimination by law affects women's economic opportunities.

following model:

$$y_{isc} = \alpha l_{isc} + \beta k_{isc} + \gamma(A_F - A_M)_c + \delta X_c + \epsilon_{isc} \tag{6}$$

where y_{isc}, l_{isc}, k_{isc}, are, respectively, log operating revenue, log employment, and log capital of firm i in sector s in country c. Since our estimated productivity gap varies at the country level, we control for GDP and the LFPR in X_c and cluster standard errors by country. The coefficient of our estimated productivity gap is negative, precisely estimated, and quite large, likely due to correlated unobservables at the country level. To estimate correlations within country, we split sectors into two groups: those where the MNE operates and those where it does not. In column 3, we estimate a model with country fixed effects that identifies the correlation between our productivity gap and the productivity of firms in the same sector minus the correlation with the productivity of firms in different sectors. The estimated elasticity in column 3 is about half of that of capital.

Table 4 estimates the model in Equation 6 at the sector level, on both the mean and the variance of the productivity gap. In line with the findings at the firm level, average productivity is negatively correlated with the calibrated productivity gap. Moreover, the dispersion of productivity between firms in the same sector and country (a rough measure of misallocation) is higher in countries where the productivity gap is higher.

5 Implications

We use the model estimates to evaluate the effect of different counterfactuals on the LFPR, the pay gap, firm productivity and welfare. To do this, we need to take a stance on how the firm responds to changes in the environment. Since we do not observe the production function of the firm nor the elasticity of demand they face, we cannot use profit maximization as the guiding criterion. Rather, we take the observed level of employment and the wage bill in each country-cohort-tenure group as binding constraints.[15] Our first counterfactual asks whether, under these constraints, the firm

[15] In practice these are determined by a maximization problem that we do not observe. This is equivalent to assuming that the firm sets the optimal scale of operation and then decides who to hire to maximize productivity.

maximizes productivity. Second, we quantify the effect of misallocation on the firm's productivity. In the third counterfactual, we simulate the effect of stricter labor laws. Lastly, we analyze the implications that our empirical findings have for gender pay gaps using both the micro-data from the firm and aggregate wage data by gender and occupation from the ILO.

5.1 For the firm's wage policy

To compute the optimal wage policy, we let the firm choose α^1_{gtac}, β^1_{gtac} to maximize the productivity of its employees, subject to two constraints: (i) keeping total employment (or the labor force) fixed; and (ii) keeping the total wage bill constant.[16] The exact details of the maximization problems are discussed in Appendix A.3.1.

Figure 8, panels (a) and (b) compare the calibrated wage policy parameters to the solution of the optimization problem described above. We can see that these do not coincide, and for some countries, they are quite far apart. The difference between the optimal and the observed parameters follows the same pattern in most countries: to maximize productivity the firm should increase the fixed pay of women α^1_F, decrease the fixed pay of men α^1_M and increase variable pay β^1 for both genders.

Under the optimal policy, the firm equates the participation thresholds of men and women, so that the LFPR gets close to 1 as shown in Figure 8 panel (e). While equating marginal productivity across groups is obviously the solution to the unconstrained problem, this exercise tells us that it is possible to do so while simultaneously keeping employment and the wage bill constant. The optimal policy effectively undoes differences in LFP and leads to higher productivity in every country as shown in Figure 8 panel (c). On average, the firm could increase productivity by 22%.

Why is the firm not setting the optimal α^1_{gtac}, β^1_{gtac}? A possible answer is that while the average pay is constant by assumption, the optimal policy generates a stark increase in inequality between genders. Indeed, because women in the labor force are more positively selected in most countries, it would be productivity-maximizing to pay women more, both in terms of fixed and variable pay, so that the pay gap between

[16]Both constraints are necessary to make sure that we obtain a sensible solution. Without the employment constraint, the firm can increase average productivity just by hiring fewer people (because of positive selection). Without the wage bill constraint, the firm can raise both α^1_{gtac} and β^1_{gtac} in a way that increases productivity without changing employment, at the cost of paying much higher wages.

females and males would be even larger than what we observe. Without any change in norms, on average over all countries, the gender pay gap (Female – Male) would have to increase by 78% (Figure 8, panel (d)). Since individual productivity is not directly observable, such an increase might not be acceptable.

Another important reason may be labor regulations, which limit variable pay and pay inequality even within gender. We address this in subsection 5.3.

Finally, the fact that the firm would be better off hiring more women (in most cases, hiring as many women as men) suggests that quotas would not bind. However, meeting them would require a large increase in inequality between genders, with steep rewards for talent, in order to sufficiently attract women into the labor force. This could be as stark as, for example, most leadership positions being held by women with all men working under them. Note that this would be a very different scenario than that of many policies which prescribe equality in pay and rewards between genders, and would imply an equal number of men and women in top-level positions, as well as in lower-level positions.[17]

5.2 For the firm's productivity

Our next counterfactual quantifies the effect of misallocation on the firm's productivity. In our framework, the true value of staying at home for a woman equals that of a man with the same observable characteristics and (A_i, N_i) type, $y^{0,*}_{iFtac} = y^0_{iMtac}$. However, when entering the labor force, women must pay a "gender norms tax", τ_{tac}, (as in Hsieh et al. (2019)) proportional to their true value of staying at home, so that when making the decision of whether to enter the labor force or not they take into account $y^0_{iFtac} = (1 + \tau_{tac})y^{0,*}_{iFtac}$. In this counterfactual, we eliminate the gender norms tax by setting the value of the staying-at-home parameters of women (α^0, ν^0) equal to those of men (within the same country, cohort, and tenure cells).

We discuss effects both in the short run, that is, keeping the pay policy of the firm fixed, and in the long run, when the firm can optimally adjust its policy to the new environment. To do this, we must take a stance on the firm's optimization procedure. In light of the results in subsection 5.1 we let the firm choose α^1_{gtac}, β^1_{gtac} to maximize the

[17] If the policy were to combine gender quotas with the imposition of equal pay, rather than rewarding talent regardless of gender, the firm will reduce mean productivity to minimize labor costs. Results are available on request.

productivity of its employees, subject to three constraints: (i) keeping total employment (or labor force) fixed; (ii) keeping the total wage bill constant; and (iii) a bound on the pay gap in the firm. This bound is chosen to maximize the goodness of fit between the constrained optimal and calibrated wage policy parameters at baseline. We refer to this wage policy as the constrained optimal wage policy. The exact details of the maximization problems are discussed in Appendix A.3.1.

Figure 9 plots the short-run effects of eliminating gender norms on the LFPR, the pay gap, and average productivity. In the short run, MLFP does not change (because men's value of staying at home and wages stay the same), and FLFP increases, so LFP increases overall, which means that the ratio increases as well. These changes in LFP reduce the pay gap because the women that enter the labor force are less able on average than those who were already working at baseline, so women's average wages decrease. For the same reason, average productivity decreases, since MLFP does not change, but FLFP increases, and the entrants are, on average, lower productivity than those who were already in the labor force. These findings mirror those in subsection 5.1: the baseline policy sacrifices productivity to bound pay inequality when entry barriers vary by gender, and hence, once barriers are equalized, the productivity cost rises.

The trade-off between productivity and inequality when the gender tax is set to zero, however, disappears in the long run (Figure 10). If the firm can adjust its policy, it will equate the participation thresholds so that the LFPR becomes 1. The fixed total employment constraint means that MLFP decreases by the same amount as FLFP increases. In the long run, because we equate the parameters of staying at home, it is optimal for the firm to have the same wage policy for men and women, which eliminates the pay gap. Again, the constant wage bill constraint means that this is just a redistribution from men to women. Finally, in a world without norms tax, the firm replaces less able men with more able women, and average productivity increases.

Figure 11 shows, by baseline LFPR, the average productivity of LFP entrants and leavers. The average productivity of LFP entrants is, because of positive selection, lower in those countries where the LFPR changes are the largest. However, those entrants are on average much more productive than the leavers they replace. Figure 12 shows the productivity gains of eliminating gender norms by countries. The average across countries is 32%, although there is substantial heterogeneity, with some countries, such as Pakistan or Sri Lanka, potentially obtaining an increase of up to 90% in productivity. To benchmark the magnitude of our results, Hsieh et al. (2019) finds that roughly 40% of growth in aggregate market output per person can be explained by the convergence in the occupational distribution across gender and race between 1960 and 2010 in the U.S. economy.

Our data is well suited to quantify the cumulative productivity loss due to the gendered division of labor, but it does not shed light on the individual components that make up gendered labor division, such as barriers to education, how work affects marriage prospects, child penalties, and so on. We calibrate gains from eliminating gender norms at every stage; if we eliminated barriers to LFP without eliminating barriers to education, the gains would of course be smaller because the pool of qualified women applicants would be smaller. We thus do not know ex-ante what the size of the pool of qualified women would be if certain gender barriers were eliminated but others were not. This precludes us from analyzing the trade-off between diversity and quality that is at the core of popular policy tools such as quotas and affirmative action.

In addition, while we can assess the impact on firm productivity keeping the wage bill and total employment constant, we do not have enough information to evaluate general equilibrium effects, which are needed to evaluate welfare impacts. Intuitively, eliminating the norms that lead to higher barriers for women has two consequences for welfare. First, the reallocation of labor inside and outside the home creates winners and losers. The former consists of the women who are strongly suited for work outside the home and transition from inside to outside, as well as of the men who are strongly suited for work inside the home who move in the opposite direction. However, men who get crowded out by women in the workplace might lose.[18] The extent to which this happens depends on the second consequence of eliminating misallocation: the increase in efficiency and productivity overall, which generates a virtuous circle through increases in demand. We cannot assess these effects because we do not know the production function of the firm, so we cannot say how the firm would adjust employment in response to a convergence of gender norms.

[18]We discuss the effects on workers' welfare in more detail in Appendix A.3.2.

5.3 For labor regulations

In a third counterfactual, we ask what the effect of stricter labor laws, that limit performance pay, would be on average productivity in the firm. To answer this question, we consider the constrained optimal wage policy and add an additional cap on returns to productivity: $\beta^1_{gtac} \leq \max\{\beta^1_{Fta,FRA}, \beta^1_{Mta,FRA}\}$, where $\beta^1_{Fta,FRA}, \beta^1_{Mta,FRA}$ denote the corresponding parameters for France (since France is the 95th percentile of the WEF Restrictive Labor Regulations Index). We leave the value of staying-at-home parameters $(\alpha^0_{gtac}, \nu^0_{gtac})$ unchanged at baseline, but we allow for the LFP of each group to respond optimally to the change in incentives.

The results are plotted in Figure 13. The average productivity of women does not change, whereas the average productivity of men decreases substantially. Intuitively, limiting performance pay hurts the firm because they cannot screen out lower-productivity men, whereas it does not affect women because lower-productivity women are out of the labor force even at baseline.

5.4 For the gender pay gap

5.4.1 Gender pay gap in the firm

The differences in selection on ability across genders that we have documented have implications for how the gender pay gap changes with the LFPR.[19] Namely, the correlation between the gender pay gap and the LFPR should have a negative sign as the lower the LFPR is, the higher the ability of the women working would be.

Table 5 combines women and men's wages together and shows that the correlation between the gender pay gap and the LFPR is indeed negative. It looks at the gender pay gap using different sources of variation, thus also allowing us to check whether the sign of the correlation is consistent regardless of the source of variation that we use to identify it. We estimate the following model:

$$w_{iact} = \alpha LFPR_{ac} + \beta Female_i + \gamma LFPR_{ac} * Female_i + \mathbf{X'}_{iact}\mathbf{\Lambda} + \boldsymbol{\psi}_t + \epsilon_{iact} \qquad (7)$$

where w is log wage of employee i in country c and year t for age group a. $\boldsymbol{\psi}_t$

[19] We define the gender pay gap as the ratio of women's wages over men's wages.

represents year fixed effects to take out year-level macro shocks and $\mathbf{X}_{iact}$ is a vector of controls. We cluster standard errors at the same level as the RHS variable, that is country-cohort. The coefficient of interest is γ which measures the change in the pay gap as LFPR increases and $\gamma < 0$ would indicate the presence of differential barriers to entry for women.

We include different controls in $\mathbf{X}_{iact}$: column 2 controls for a quadratic function of tenure and function fixed effects, column 3 adds GDP per capita in logs so to show that the variation in the LFPR is not only a function of country income, column 4 adds cohort fixed effects so to only exploit the variation across countries; column 5 replaces the cohort fixed effects with country fixed effects hence only exploiting the variation within countries. The comparison between columns 4 and 5 is particularly informative as it uses one source of variation at a time.

The estimates of γ are negative and precisely estimated in all specifications. In Appendix Table A.1, we use the LFP data for individuals with advanced education only.[20] The results are nearly unchanged when we adopt this measure. However, we lose almost 20% of the sample and particularly from countries with low female labor force participation (FLFP). Hence, we employ the overall LFP as our default measure.[21] The fact that the γ coefficient is stable does reassure us that it is, in fact, measuring selection.[22]

Finally, columns 6 and 7 estimate the same specification as in column 1 to look at the pay progression for new hires. They show that women in low LFPR country-cohorts display faster pay growth and a higher probability of promotion and that this positive gender gap in realized performance decreases as the LFPR increases. This

[20] As defined by the World Bank, an individual with advanced education has completed a short-cycle tertiary education or a college degree and/or above.

[21] The correlation between overall LFPR and LFPR for individuals with advanced education is 68%.

[22] We conduct a number of additional robustness checks. Table A.2 in the Appendix shows that the patterns in Table 5 hold if we add fixed effects for the geographical region and if we split the sample by lower income and higher income countries (as defined by the World Bank). In Appendix Table A.3 we report the results when converting wages from euros into PPP 2017 $, using the PPP conversion rates of the ICP at the World Bank. The gender gap is unaffected and the only change is the magnitude of the coefficient on LFPR which shrinks to the level found when controlling for country fixed effects (column 5 in Table 5). This is what we would expect as differences in PPP exchange rates would not affect cross-country comparisons of the gender gap. Finally, we note that the results are driven by differences in fixed pay rather than in variable bonus, which constitutes a much lower proportion of overall salary (the median ratio of bonus to fixed pay is 13%) — see Appendix Table A.4. In the company, pay summarizes altogether most differences in performance between employees, encompassing off-peak salary increases as well as promotions.

evidence on the sample of new hires also supports the interpretation that the firm is not taking into account the implications on productivity that different LFP rates by gender might entail when making decisions on hiring and wage offers. It is only after some time at the firm that women and men *ex-post* pay gaps start to diverge the lower the LFPR is, indicating positive selection of women into the firm due to ex-ante unobservable characteristics.

5.4.2 Adjusting country-level pay gaps

A direct implication of our findings is that gender pay gaps in aggregate data should be larger once we account for differential selection. If women are more positively selected into the labor force compared to men, the observed pay gap will underestimate the true average pay gap (i.e. the pay gap that the average woman in the population would face vis-à-vis the average man). Table 5 shows that there is a negative correlation between the gender pay gap and the LFPR (-0.187). We also find a negative correlation between the gender pay gap and the LFPR when using aggregate data from the ILO for ISCO-08 categories 1–5 (white collars).[23] However, as anticipated, the correlation is significantly weaker (-0.087).

We use the estimates from our model to adjust the ILO estimates of the gender pay gap to take into account selection. Our model implies the following relationship between the moments of the observed and population pay distributions:

$$\mathbb{E}[y^1_{igtac} \mid \text{empl'd}\,] = \mathbb{E}[y^1_{igtac}] + (\beta^1_{gtac})^2 \lambda(\xi_{gac}).$$

Given external estimates of the observed pay gap, we could use our calibrated parameters to adjust them for selection as follows:

$$\underbrace{\mathbb{E}[y^1_{iFtac}] - \mathbb{E}[y^1_{iMtac}]}_{\text{Adjusted Gender Gap}} = \underbrace{\mathbb{E}[y^1_{iFtac} \mid \text{empl'd}\,] - \mathbb{E}[y^1_{iMtac} \mid \text{empl'd}\,]}_{\text{Unadjusted Gender Gap}} - \underbrace{[(\beta^1_{Ftac})^2 \lambda(\xi_{Fac}) - (\beta^1_{Mtac})^2 \lambda(\xi_{Mac})]}_{\text{Adjustment term}}$$

We use the above procedure to adjust the ILO estimates of the gender pay gap for

[23] Olivetti and Petrongolo (2008) documents a similarly negative correlation in OECD countries.

white-collar workers.[24] The results are in Figure 14. Once we adjust gender gaps for selection, we get much larger magnitudes (in absolute value), which are up to four times the size of the unadjusted gap. As expected, the difference is larger for countries with low LFPR, since those are the ones where the adjustment term is larger. The coefficient for the LFPR in a regression with the difference in gaps in absolute value as the outcome takes the value of -0.37 and is unchanged if we also control for GDP per capita in logs (p-value = 0.28).[25]

6 Conclusion

We have documented that the gendered division of labor between the market and the home creates a direct link between diversity and productivity through selection. That is, when most women stay at home, the few that work outside the home are those with the highest productivity. We note that the mirror image of the gender tax that women have to pay to work outside the home is the tax that men have to pay to work inside the home, independent of their skills and preferences. Thus, eliminating gender norms will also eliminate misallocation in work inside the home, by allowing the men who wish to do so to specialize in home production. Understanding differential selection by gender, or indeed by any under-represented group, is key to informing personnel policy as well as broader labor market policies, such as the design of labor laws and addressing the gender pay gap.

Our findings have three implications for the design of firms' diversity policies at the hiring and compensation stages. The first implication is that even if the firm does not observe potential productivity at the point of hiring, the logic of selection indicates that under-represented groups, other things equal, will likely have higher productivity. This implies that between two potential hires with the exact same observable qualities, the minority candidate will have the better unobservables on average. This leads to the second implication, namely that diversity is justifiable on the grounds

[24] For 68 countries not in our firm sample, we impute the selection adjustment term based on a cubic regression on the LFPR.

[25] As already noted, differences in the LFPR cannot be solely explained by differences in GDP. Hence, we find larger differences between the adjusted and unadjusted gender pay gaps in countries that share a low LFPR but that may be at different levels of economic development, for example, Sri Lanka and Italy or Guatemala and Turkey.

of productivity, regardless of whether diversity has an independent effect of its own on productivity. The third implication is that aiming for gender equity — in pay, promotions, and dismissals — can turn out to be inequitable because selection generates different distributions of productivity between genders. Perhaps counter-intuitively, gender equity policies might end up hurting women as they limit the firm's ability to reward performance. Under current circumstances, greater rewards for objective performance could induce more women to enter the labor market (and stay there). Rewards need not be monetary. Indeed, in the presence of social norms where women may bear a disproportionate burden of childcare and there are transaction costs, resources that would enable women to better manage childcare could have equally large productivity gains for the firm through the selection margin.

Awareness of the evidence of positive selection of under-represented groups changes how quality can be inferred, particularly if quality is not perfectly observable or objectively measured: the very presence of a member of an under-represented group should change one's prior on the talent of that member, simply through the logic of positive selection. This could lead to greater expectations, possibly in the form of greater compensation, placed on minority candidates, which could have opposing effects: on the one side, it could put greater pressure on minority candidates to perform or de-motivate non-minority employees and, on the other side, it could provide greater motivation for diverse employees by being received with positive expectations by management. The implications of increased awareness of the positive selection effect of diverse candidates by leaders for the recruitment, promotion, and productivity of under-represented groups are left to future research.

References

Alesina, Alberto and Eliana La Ferrara, "Ethnic diversity and economic performance," *Journal of economic literature*, 2005, *43* (3), 762–800.

Alfaro-Urena, Alonso, Isabela Manelici, and Jose P Vasquez, "The Effects of Joining Multinational Supply Chains: New Evidence from Firm-to-firm Linkages," *Working paper*, 2019.

Andrew, Alison, Oriana Bandiera, Monica Costa-Dias, and Camille Landais, "Women and men at work," *IFS Deaton Review of Inequalities*, 2021.

Athey, Susan, Christopher Avery, and Peter Zemsky, "Mentoring and diversity," *American Economic Review*, 2000, *90* (4), 765–786.

Bertrand, Marianne, "New perspectives on gender," in "Handbook of labor economics," Vol. 4, Elsevier, 2011, pp. 1543–1590.

_ , "Gender in the twenty-first century," in "AEA Papers and proceedings," Vol. 110 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2020, pp. 1–24.

_ **and Esther Duflo**, "Field experiments on discrimination," *Handbook of economic field experiments*, 2017, *1*, 309–393.

Besley, Timothy, Olle Folke, Torsten Persson, and Johanna Rickne, "Gender Quotas and the Crisis of the Mediocre Man: Theory and Evidence from Sweden," *American Economic Review*, August 2017, *107* (8), 2204–42.

Blau, Francine D, *Equal pay in the office*, Lexington Books, 1977.

_ , *Gender, inequality, and wages*, Oxford University Press, 2012.

Borjas, George J., "Self-Selection and the Earnings of Immigrants," *The American Economic Review*, 1987, *77* (4), 531–553.

Bursztyn, Leonardo, Alexander W Cappelen, Bertil Tungodden, Alessandra Voena, and David H Yanagizawa-Drott, "How Are Gender Norms Perceived?," *National Bureau of Economic Research*, 2023.

Card, David, Ana Rute Cardoso, and Patrick Kline, "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women*," *The Quarterly Journal of Economics*, 10 2015, *131* (2), 633–686.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde, "Global evidence on economic preferences," *The Quarterly Journal of Economics*, 2018, *133* (4), 1645–1692.

_ , _ , Thomas J Dohmen, David Huffman, and Uwe Sunde, "The preference survey module: A validated instrument for measuring risk, time, and social preferences," 2016.

Fernández, Raquel, Alessandra Fogli, and Claudia Olivetti, "Mothers and sons: Preference formation and female labor force dynamics," *The Quarterly Journal of Economics*, 2004, *119* (4), 1249–1299.

_ , Asel Isakova, Francesco Luna, and Barbara Rambousek, *Gender equality and inclusive growth*, International Monetary Fund, 2021.

Goldin, Claudia, *Understanding the gender gap: An economic history of American women* number gold90-1, National Bureau of Economic Research, 1990.

_ , "The U-shaped female labor force function in economic development and economic history," 1995.

_ , "A grand gender convergence: Its last chapter," *American Economic Review*, 2014, *104* (4), 1091–1119.

Hale, Galina and Mingzhi Xu, "FDI effects on the labor market of host countries," *2016-25*, 2016.

Hamilton, Barton H, Jack A Nickerson, and Hideo Owan, "Diversity and productivity in production teams," in "Advances in the Economic Analysis of participatory and Labor-managed Firms," Vol. 13, Emerald Group Publishing Limited, 2012, pp. 99–138.

Heckman, James, "The Common Structure of Statistical Models of Truncation, Sample Selection and Limited Dependent Variables and a Simple Estimator for Such

Models," in "Annals of Economic and Social Measurement," Vol. 5, National Bureau of Economic Research, Inc, 1976, pp. 475–492.

Hjort, Jonas, "Ethnic divisions and production in firms," *The Quarterly Journal of Economics*, 2014, *129* (4), 1899–1946.

_ , Xuan Li, and Heather Sarsons, "Across-Country Wage Compression in Multinationals," Technical Report, National Bureau of Economic Research 2020.

Hong, Lu and Scott E Page, "Problem solving by heterogeneous agents," *Journal of economic theory*, 2001, *97* (1), 123–163.

Hsieh, Chang-Tai, Erik Hurst, Charles I Jones, and Peter J Klenow, "The Allocation of Talent and US Economic Growth," *Econometrica*, 2019, *87* (5), 1439–1474.

Hyland, Marie, Simeon Djankov, and Pinelopi Koujianou Goldberg, "Do Gendered Laws Matter?," 2021.

Javorcik, Beata S, "Does FDI bring good jobs to host countries?," *The World Bank Research Observer*, 2015, *30* (1), 74–94.

Jayachandran, Seema, "The roots of gender inequality in developing countries," *economics*, 2015, *7* (1), 63–88.

Lazear, Edward P, "Globalisation and the market for team-mates," *The Economic Journal*, 1999, *109* (454), 15–40.

_ , "Why are some immigrant groups more successful than others?," *Journal of Labor Economics*, 2021, *39* (1), 115–133.

Marx, Benjamin, Vincent Pons, and Tavneet Suri, "Diversity and team performance in a Kenyan organization," *Journal of Public Economics*, 2021, *197*, 104332.

Minni, Virginia, "Making the Invisible Hand Visible: Managers and the Allocation of Workers to Jobs," *Working Paper*, 2022.

Olivetti, Claudia and Barbara Petrongolo, "Unequal Pay or Unequal Employment? A CrossCountry Analysis of Gender Gaps," *Journal of Labor Economics*, 2008, *26* (4), 621–654.

_ **and** _ , "The evolution of gender gaps in industrialized countries," *Annual review of Economics*, 2016, *8*, 405–434.

_ **and** _ , "The economic consequences of family policies: lessons from a century of legislation in high-income countries," *Journal of Economic Perspectives*, 2017, *31* (1), 205–230.

Propper, Carol and John van Reenen, "Can Pay Regulation Kill? Panel Data Evidence on the Effect of Labor Markets on Hospital Performance," *Journal of Political Economy*, 2010, *118* (2), 222–273.

Roy, Andrew Donald, "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 1951, *3* (2), 135–146.

Verhoogen, Eric A, "Trade, quality upgrading, and wage inequality in the Mexican manufacturing sector," *The Quarterly Journal of Economics*, 2008, *123* (2), 489–530.

Wiswall, Matthew and Basit Zafar, "Preference for the Workplace, Investment in Human Capital, and Gender*," *The Quarterly Journal of Economics*, 08 2017, *133* (1), 457–507.

7 Figures

Figure 1: LFPR and GDP per capita

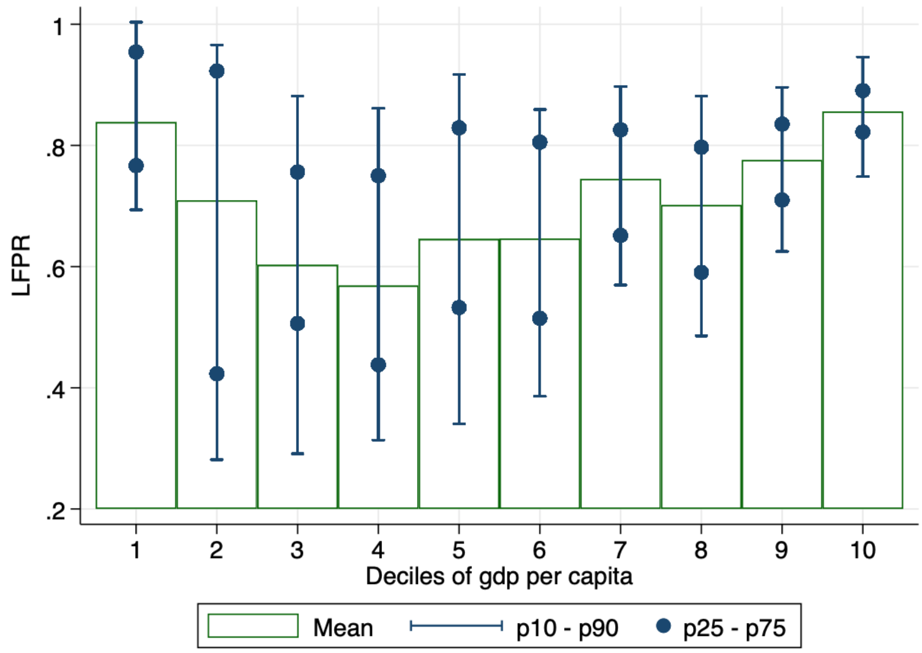

Notes. The figure plots the average, the interdecile range, and interquartile range of the LFPR across deciles of GDP per capita across countries and cohorts.

Figure 2: Female/Male in MNE vs LFPR

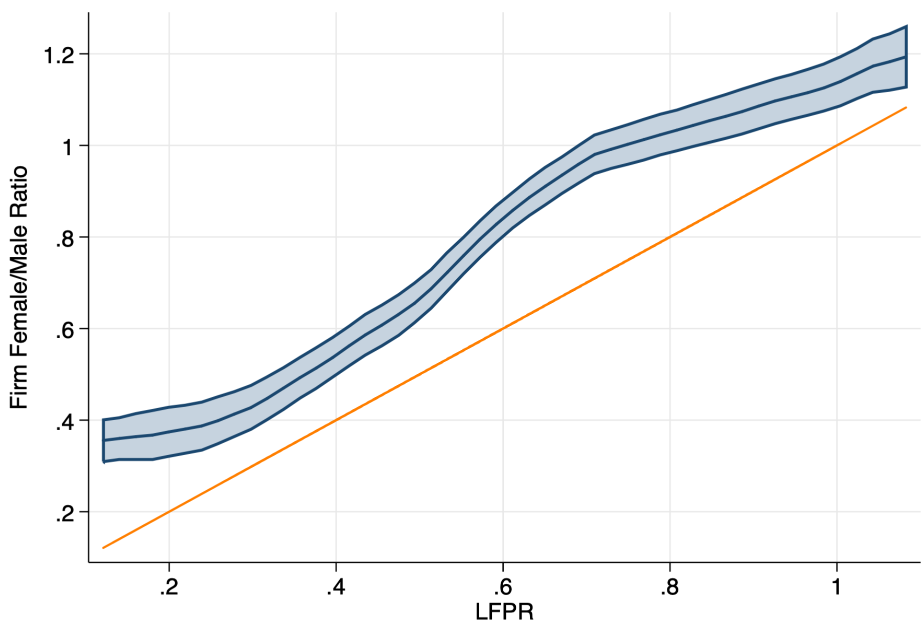

Notes. The y-axis corresponds to the female/male employment ratio in the MNE while the x-axis corresponds to the LFPR in the countries. The blue line shows the relationship between the two measures smoothed through a local linear regression. The orange line represents the 45 degree line.

Figure 3: Share of women and women's performance

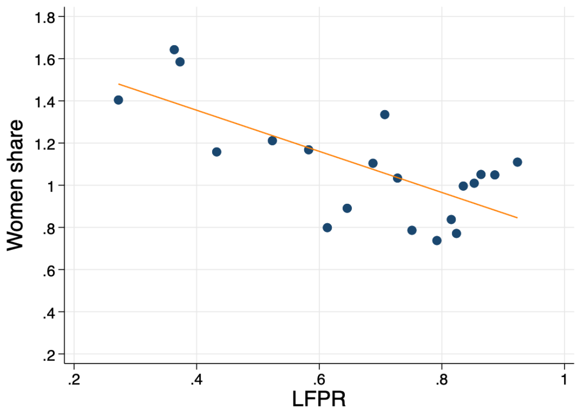

(a) Share of women, top decile of wages

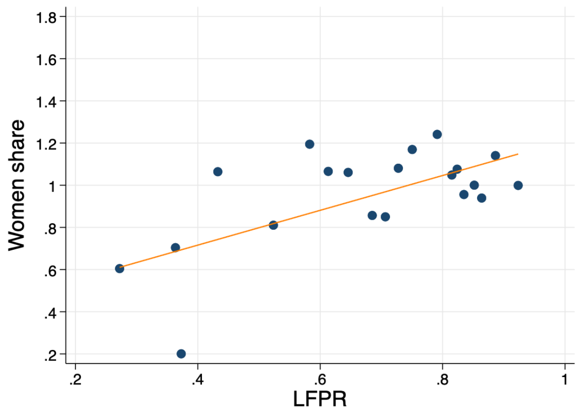

(b) Share of women, bottom decile of wages

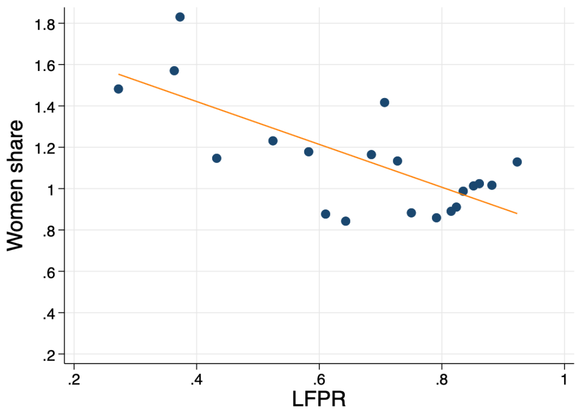

(c) Share of women, managerial positions

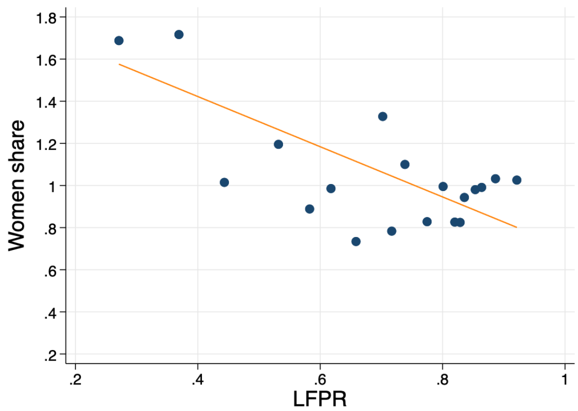

(d) Share of women, high promotion speed

Notes. The figures are binned scatterplots and a linear fit of the share of women with different performance characteristics as a proportion of the gender share in each country-cohort cell against the LFPR. Panel (a) looks at the top decile of wages; Panel (b) at the bottom decile; Panel (c) looks at the share of women in managerial positions; and Panel (d) at promotion speed to managerial positions (we apply the method developed in Minni (2022) to identify fast managerial promotions). In the regressions, we use analytical weights by employee size of each cohort-country cell.

Figure 4: Women's wages and the LFPR

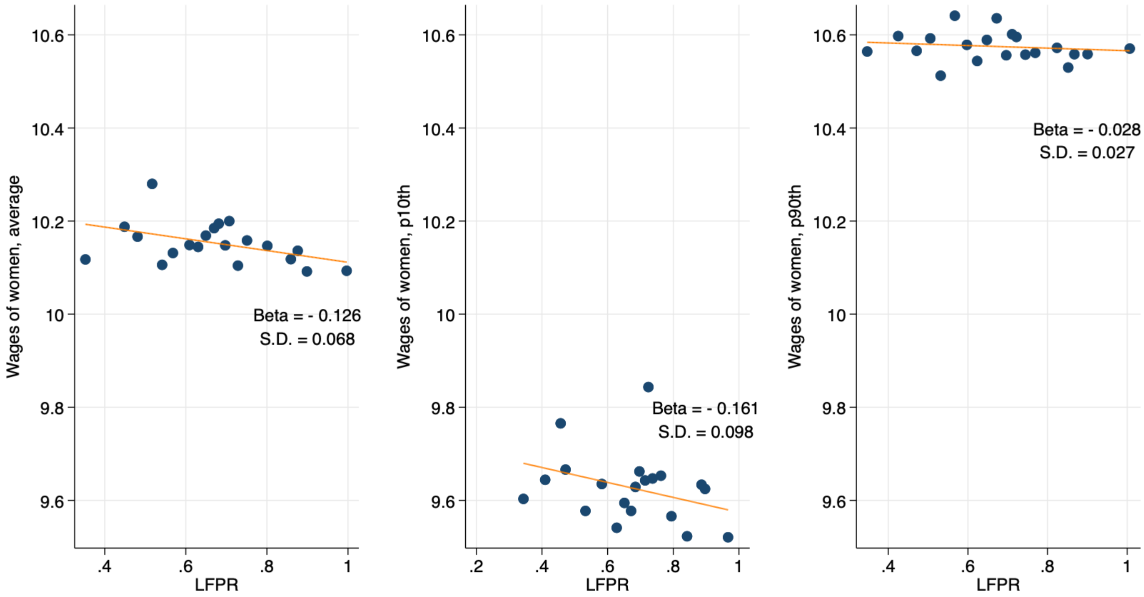

Notes. The figures are binned scatterplots and a linear fit of women's wages against the LFPR. The first figure from the left plots women's average wages, the middle one plots the 10th percentile, and the last one plots the 90th percentile. In the regressions, we control for the respective measure of men's wages, and we use analytical weights by employee size of each cohort-country cell and robust standard errors.

Figure 5: Calibrated productivity: average by LFPR

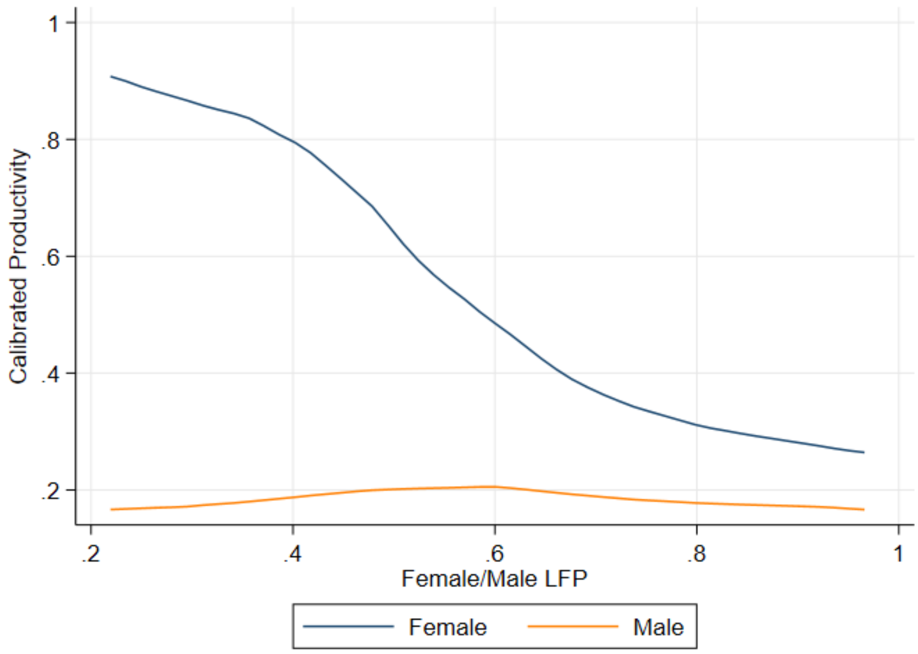

Notes. The figure plots average calibrated productivity for our sample of firm workers by Female/Male LFP, smoothed through a local linear regression.

Figure 6: Calibrated productivity: density by LFPR group

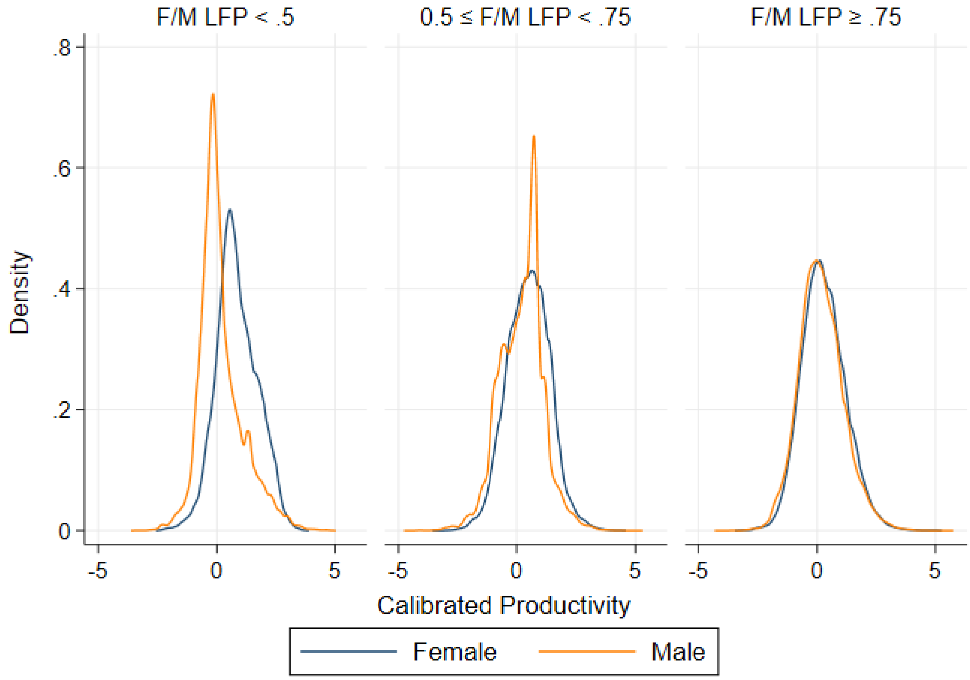

Notes. The figure plots a kernel density estimate of calibrated productivity for our sample of firm workers by three LFPR groups: $[0, .5)$, $[.5, .75)$ and $[.75, 1]$.

Figure 7: Counterfactual preference gap

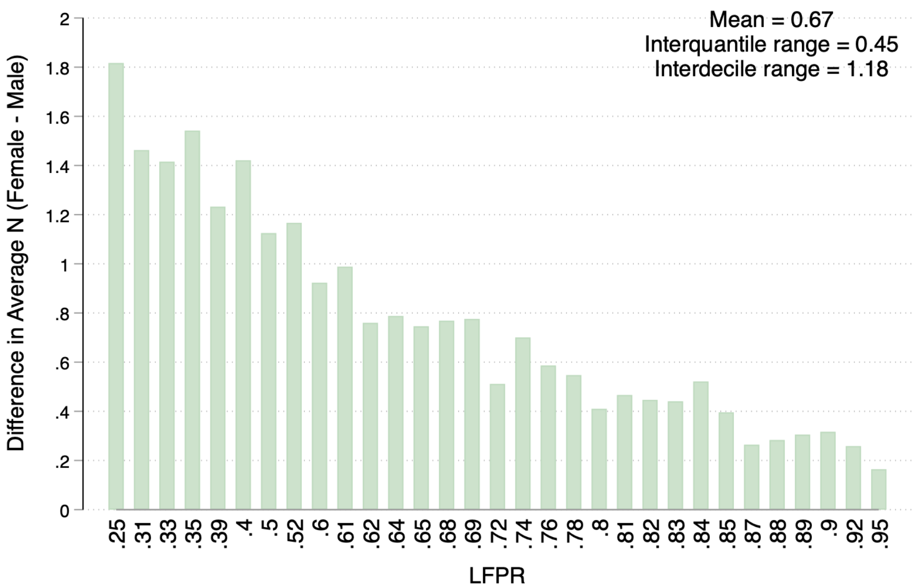

Notes. The figure plots the difference in means of N between genders that make the observed LFP gap optimal, computed as explained in the main text. The x-axis is the LFPR in each country.

Figure 8: Baseline vs. optimal wage policy

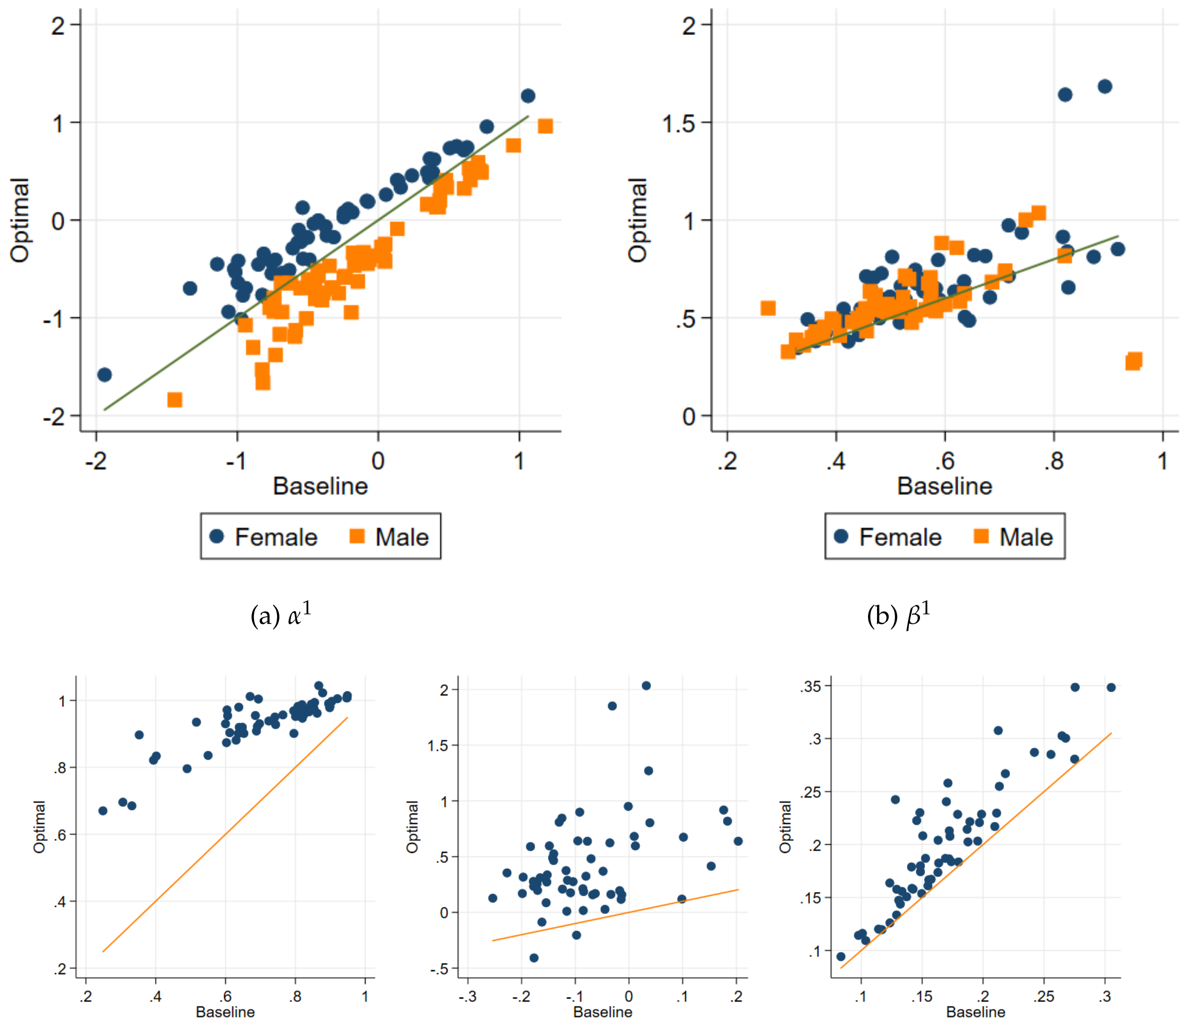

(a) α^1

(b) β^1

(c) F/M LFPR

(d) Pay gap

(e) Average productivity

Notes. The figures compare different outcomes (female to male LFP ratio, pay gap, and average productivity) and the wage policy parameters (α^1, β^1) at baseline vs. the optimal wage policy (see main text for details). Each dot represents a country and the 45-degree line is included.

Figure 9: Baseline vs. equating norms (short run)

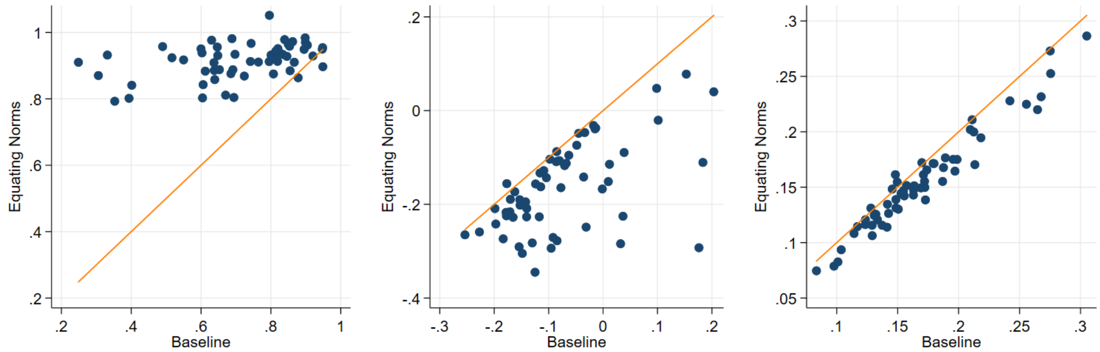

(a) F/M LFPR (b) Pay gap (c) Average productivity

Notes. The figures compare different outcomes (female to male LFP ratio, pay gap, and average productivity) under the baseline norm parameters (α^0, ν^0) to the counterfactual where these are equalized at the male levels. The "short-run" wage policy scenario keeps the wage policy of the firm fixed at the calibrated baseline parameters. Each dot represents a country and the 45-degree line is included.

Figure 10: Baseline vs. equating norms (long run)

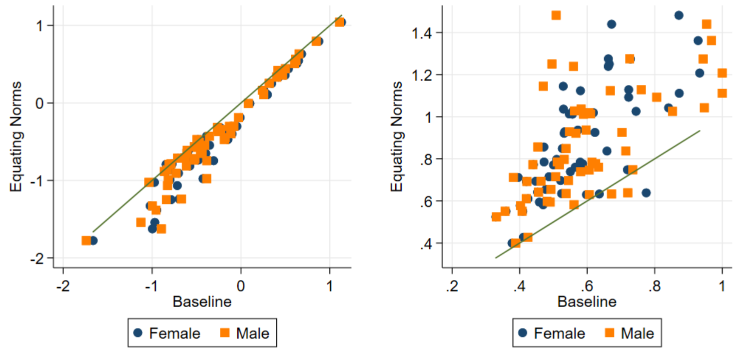

(a) α^1 (b) β^1

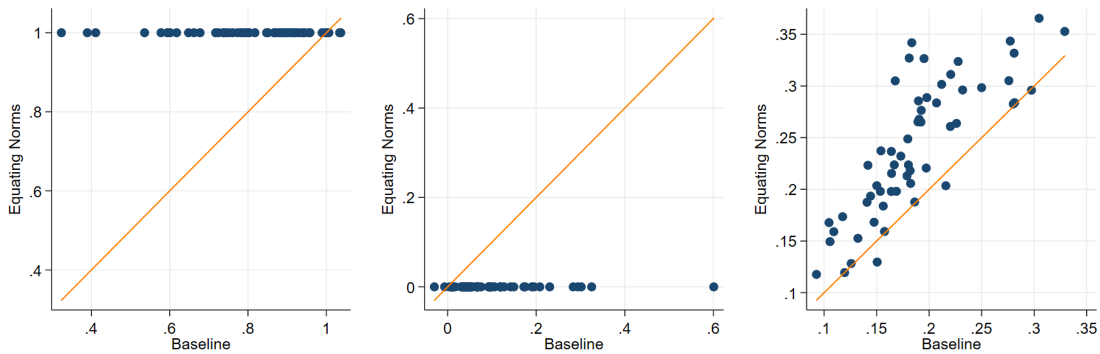

(c) F/M LFPR (d) Pay gap (e) Average productivity

Notes. The figures compare different outcomes (female to male LFP ratio, pay gap, and average productivity) and the wage policy parameters (α^1, β^1) under the baseline norm parameters (α^0, ν^0) to the counterfactual where these are equalized at the male levels. The "long-run" wage policy scenario lets the firm optimize the wage policy to maximize productivity under certain constraints (see the main text for details). Each dot represents a country and the 45-degree line is included.

Figure 11: Average productivity of entrants and leavers by country (long run)

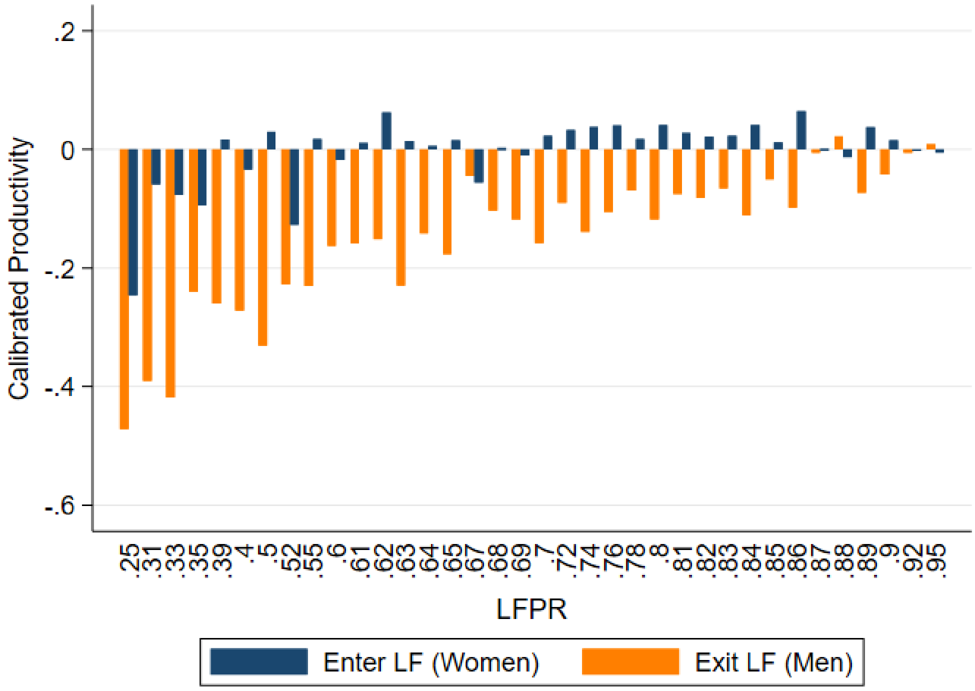

Notes. The figure plots the average productivity of LFP entrants against leavers in the long-run wage policy response. The x-axis is the LFPR in each country.

Figure 12: Increase in average productivity by LFPR (long run)

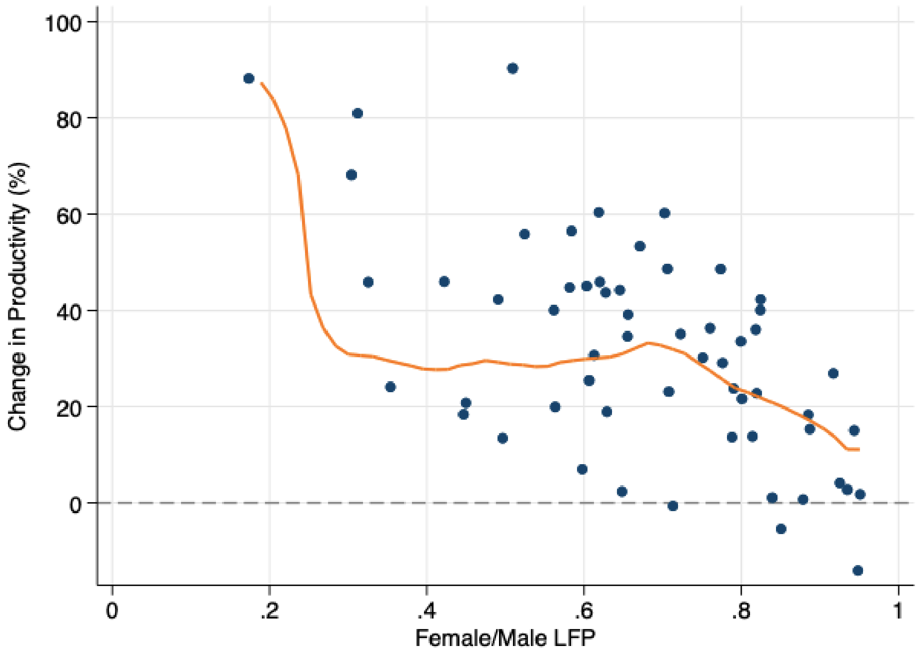

Notes. The figure plots the percentage increase in average productivity when eliminating social norms. Each dot is a country. The average across countries is 32%.

Figure 13: Counterfactual average productivity with strict labor regulation

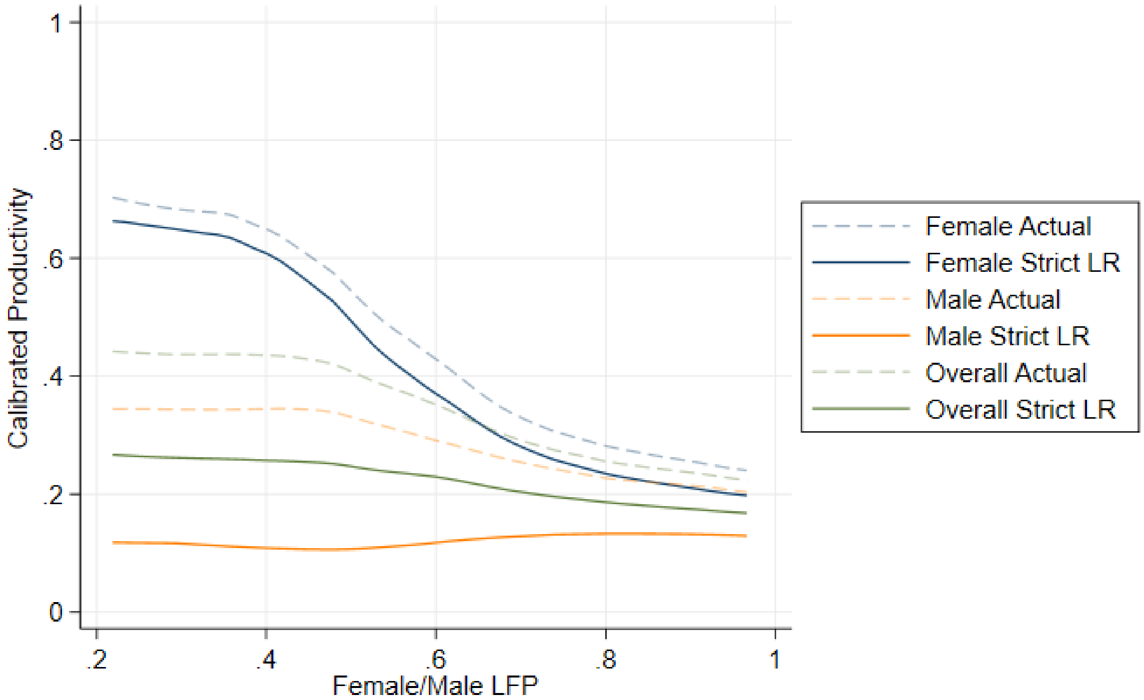

Notes. The figure plots a local polynomial regression of average productivity at baseline and under the strict labor regulation counterfactual, both estimated under the constrained optimal policy.

Figure 14: ILO gender pay gaps adjusted for selection

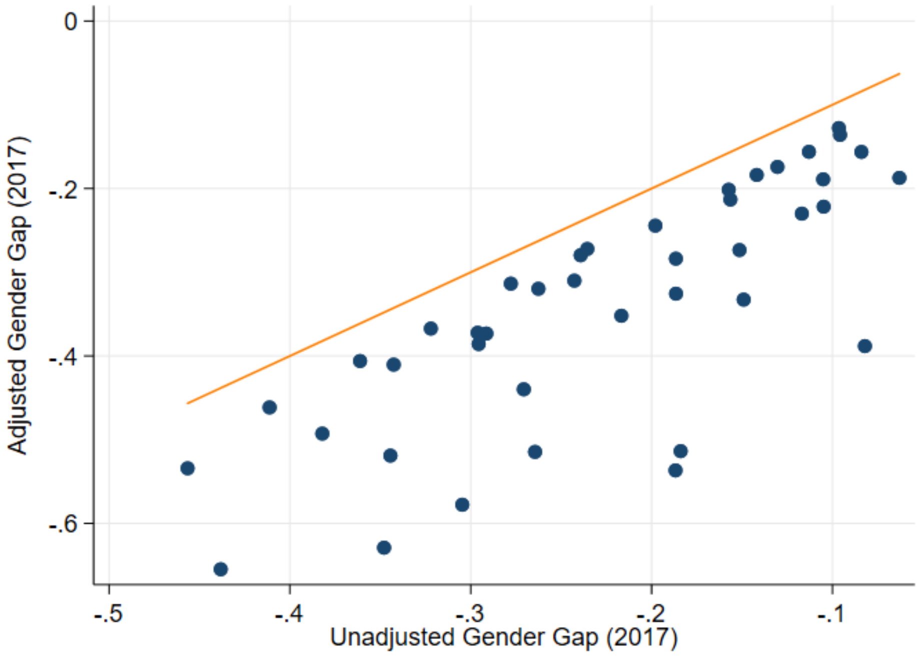

Notes. The figure is a scatterplot of the adjusted vs. unadjusted gender pay gaps. Each dot is a country. The original data (unadjusted) is from ILO, for ISCO-08 categories 1–5 (white collars). We adjust those for selection using our calibrated parameters by the procedure described in the main text.

8 Tables

Table 1: Summary statistics

	(1)	(2)
Variable	Male	Female
Pay + Bonus (logs)	10.467	10.426
	(0.695)	(0.648)
Age	42.859	42.404
	(10.074)	(10.159)
Tenure	12.027	11.869
	(8.605)	(8.609)
Share in Work-level 2+	0.243	0.214
	(0.286)	(0.272)
Share with fast promotions	0.186	0.180
	(0.305)	(0.301)
Share top performers	0.137	0.124
	(0.205)	(0.189)
Econ, Business, and Admin	0.506	0.526
	(0.381)	(0.371)
Share in Sales Function	0.461	0.335
	(0.337)	(0.316)
Observations	3,338	3,103

Notes. This table reports summary statistics for the relevant sample of workers used in the analysis. An observation is a gender-cohort-country-tenure cell (tenure is binned in groups of 2 years each). This is the relevant unit in the structural estimation. Tenure is measured in years. Work level denotes the hierarchical tier (from level 1 at the bottom to level 6). The share of fast promotions only considers workers that achieve at least work-level 2 or higher. The sales function is the most common function (39%). Top performers based on firm performance appraisals system (top 10%).

Table 2: Summary of model parameters and empirical targets

Param.	Interpretation	Empirical Target
α^0_{gtac}	Unconditional average value of staying at home	LFP
α^1_{gtac}	Unconditional average log-wage	Average observed log-wage (controlling for selection)
β^1_{gtac}	Returns to productivity in the firm	Variance of the observed log-wage (controlling for selection)
ν^0_{gtac}	Dispersion of the idiosyncratic taste for staying at home	Not identified separately (Normalize $(\beta^1_{gtac})^2 + (\nu^0_{gtac})^2 = 1$)

Table 3: $\overline{A}_F - \overline{A}_M$ and log(operating revenue), ORBIS

	Log(OpRev)		
	(1)	(2)	(3)
Productivity gap	-9.151	-12.975	
	(1.281)	(1.879)	
Same SIC 3=1 × Productivity gap			-0.165
			(0.039)
Log(employment)	0.607	0.616	0.610
	(0.212)	(0.215)	(0.001)
Same SIC 3=1 × Log(employment)			0.148
			(0.002)
Log(capital)	0.425	0.420	0.372
	(0.061)	(0.059)	(0.001)
Same SIC 3=1 × Log(capital)			0.016
			(0.001)
Log(GDP)	0.532	0.517	
	(0.156)	(0.136)	
LFPR		-6.530	
		(3.357)	
Same SIC 3=1			-0.363
			(0.014)
Country FE	No	No	Yes
R-squared	0.665	0.672	0.714
N	2239881	2239881	2239881

Notes. An observation is a firm in the Orbis database. Cross-section based on latest year up to 2019, sample restricted to firms whose latest year is after 2011. Standard errors clustered at the country level in cols. 1 and 2 and robust in col. 3.

Table 4: $\overline{A}_F - \overline{A}_M$ and productivity dispersion, ORBIS

	Log(OpRev/emp.), Mean			**Log(OpRev/emp.), CV**		
	(1)	(2)	(3)	(4)	(5)	(6)
Productivity gap	-17.549	-11.390	-11.310	0.577	0.556	0.529
	(3.757)	(2.422)	(2.492)	(0.168)	(0.144)	(0.145)
Log(GDP)	0.868	0.483	0.483	0.007	0.011	0.011
	(0.248)	(0.173)	(0.173)	(0.008)	(0.009)	(0.009)
LFPR	-7.750	-6.892	-6.907	0.365	0.388	0.386
	(5.438)	(3.493)	(3.501)	(0.278)	(0.242)	(0.242)
Log(capital), Mean		0.458	0.449			
		(0.065)	(0.065)			
Same SIC 3=1 × Productivity gap			-0.379			0.143
			(0.613)			(0.039)
Same SIC 3=1 × Log(capital), Mean			0.062			
			(0.019)			
Log(capital), CV					0.148	0.143
					(0.085)	(0.086)
Same SIC 3=1 × Log(capital), CV						0.037
						(0.032)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.751	0.890	0.891	0.485	0.523	0.529
N	2418	2418	2418	2418	2418	2418
Outcome mean	10.822	10.822	10.822	0.145	0.145	0.145
Outcome sd	2.006	2.006	2.006	0.059	0.059	0.059

Notes. An observation is an industry (US SIC 3) -country cell in the Orbis database. Analytics weights used. Measures based on cross-section of firms based on latest year up to 2019, sample restricted to firms whose latest year is after 2011. Standard errors clustered at the country level.

Table 5: Gender pay gap and LFPR

	Full Sample					New Hires	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			Pay + Bonus (logs)			**Pay Growth**	**Major promotion**
Female	0.377	0.256	0.028	0.253	0.197	0.159	0.107
	(0.142)	(0.096)	(0.173)	(0.094)	(0.077)	(0.067)	(0.047)
LFPR	1.640	1.641	0.596	1.662	0.138	0.514	-0.028
	(0.282)	(0.212)	(0.153)	(0.204)	(0.235)	(0.151)	(0.080)
Female × LFPR	-0.564	-0.471	-0.274	-0.451	-0.376	-0.142	-0.112
	(0.194)	(0.135)	(0.116)	(0.132)	(0.103)	(0.087)	(0.062)
GDP per capita (logs)			0.275				
			(0.021)				
Female × GDP per capita (logs)			0.008				
			(0.016)				
Controls	No	Yes	Yes	Yes	Yes	No	No
Cohort FE	No	No	No	Yes	No	No	No
Country FE	No	No	No	No	Yes	No	No
N	303756	303756	302567	303756	303756	8274	8274
R-squared	0.116	0.285	0.435	0.307	0.540	0.103	0.056

Notes. An observation is a worker-year. Controls include: tenure, tenure squared, year FE and function FE. The last two columns report estimates when restricting the sample to new hires at the entry level observed for at least four years. *Pay growth* is computed as the difference in log pay between the last year a worker is observed and the first year a worker is observed. *Probability of promotion* equals 1 if the worker was promoted to work-level 2 during the sample period. Controlling for starting salary in columns 6 and 7. Standard errors clustered at the country-cohort level.

A Appendix

A.1 Descriptives

Figure A.1: The countries where the MNE operates and female to male LFP

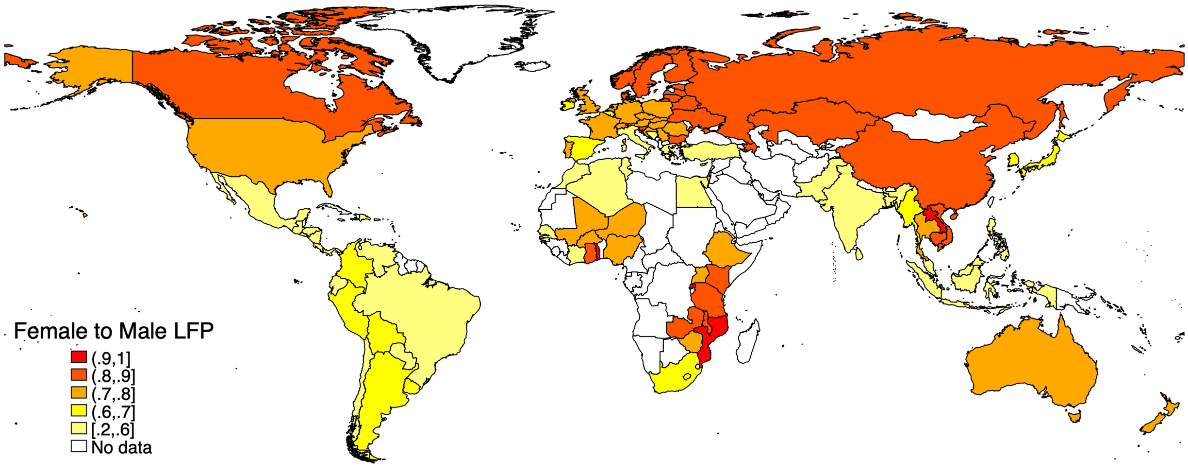

Figure A.2: Average wages in the firm and in the country overall: a) ILO, white collar occupations only; and b) ORBIS, manufacturing sector only

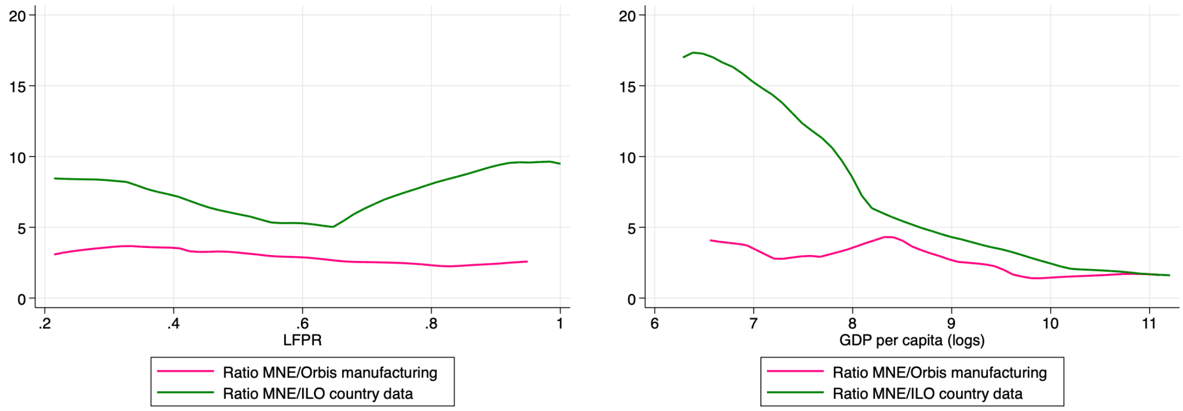

(a) Wage ratios against LFPR (b) Wage ratios against GDP per capita in logs

Notes. This figure plots the ratio of the average wage in the MNE and in the country overall: a) from the ORBIS database, considering the manufacturing sector only, and b) from the International Labor Organization (ILO), considering white-collar occupations only. Wages are measured in 2017 PPP $. The x-axis is the LFPR (panel a) and the GDP per capita in logs (panel b) in each country.

A.2 Analysis

Figure A.3: Gender differences in preferences

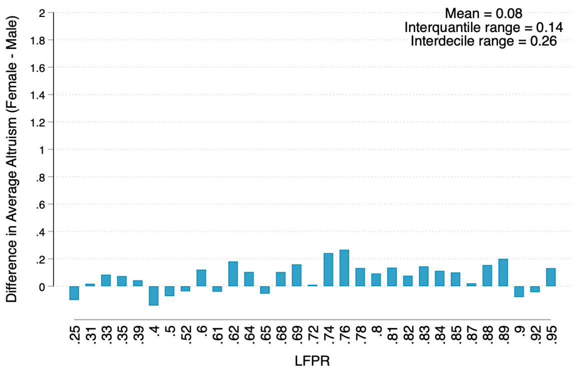

(a) Altruism

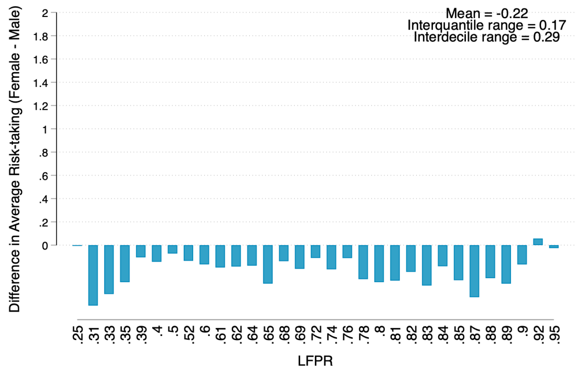

(b) Risk-taking

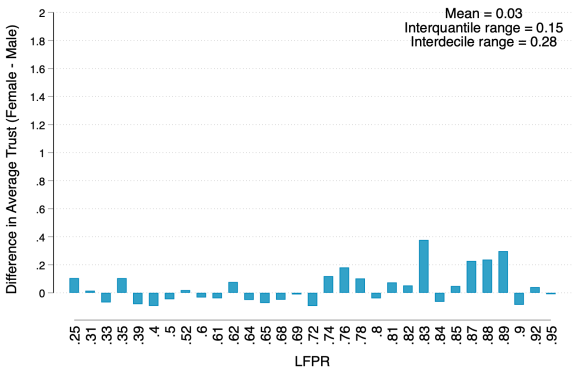

(c) Trust

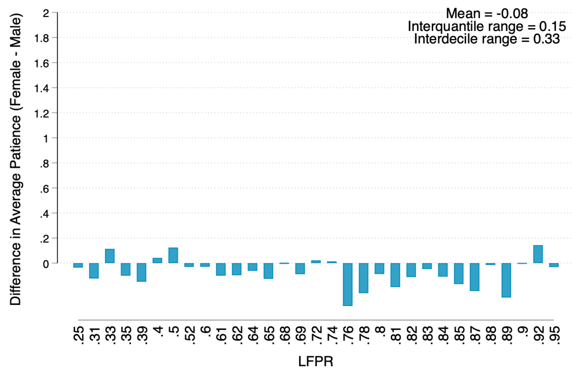

(d) Patience

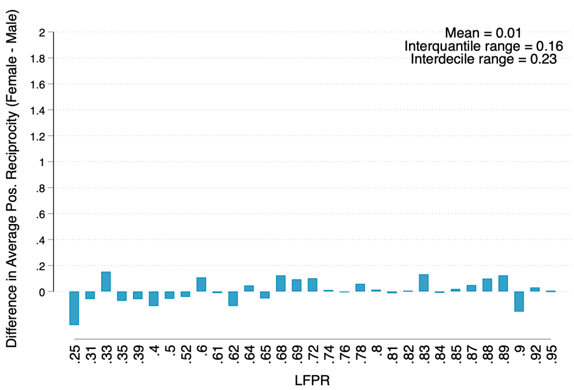

(e) Positive reciprocity

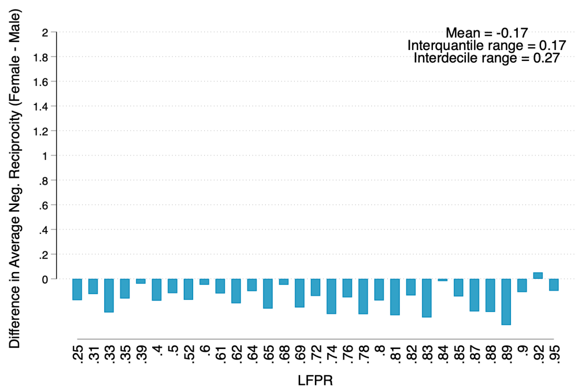

(f) Negative reciprocity

Notes. The figure plots the difference in means of economic preferences between genders. Data taken from the Global Preferences Survey (Falk et al. (2016), Falk et al. (2018)). The x-axis is the LFPR in each country.

Figure A.4: Goodness of fit: Wage distribution in India and Sweden

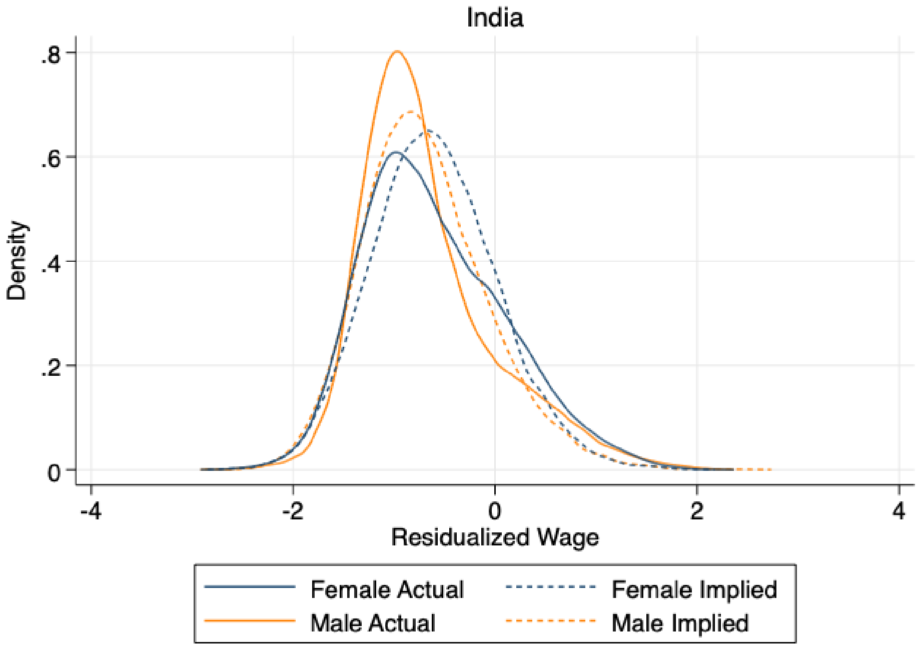

(a) India

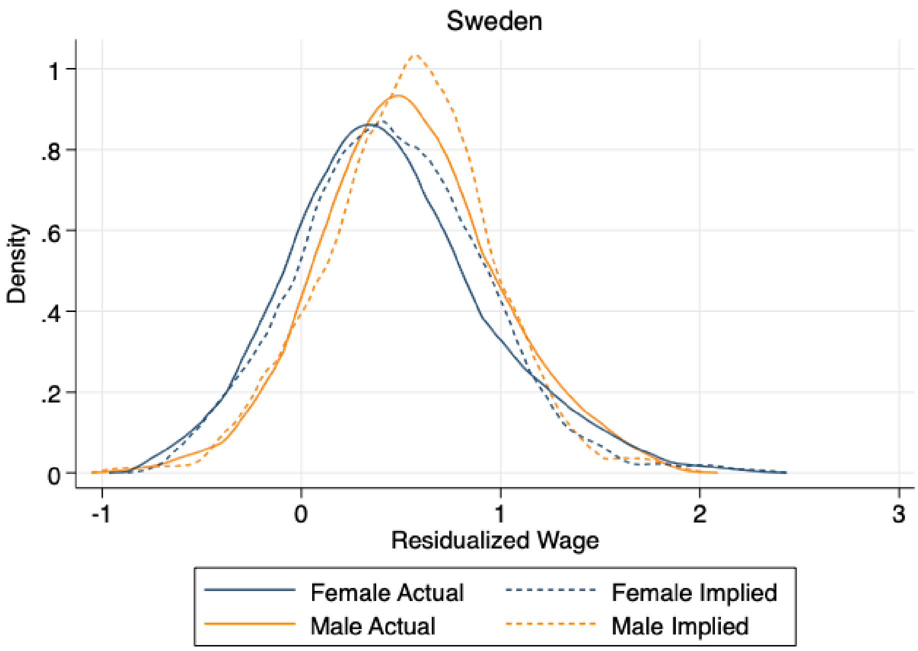

(b) Sweden

Notes. The figures compare the actual distribution of residualized log wages in India and Sweden to the distributions implied by our calibrated model.

Figure A.5: Validation of our calibrated productivity against other performance measures

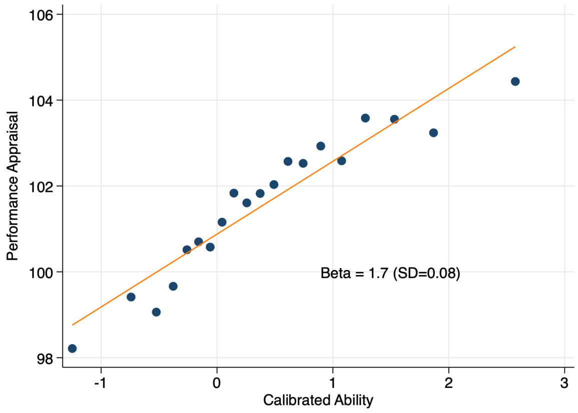

(a) Performance appraisals

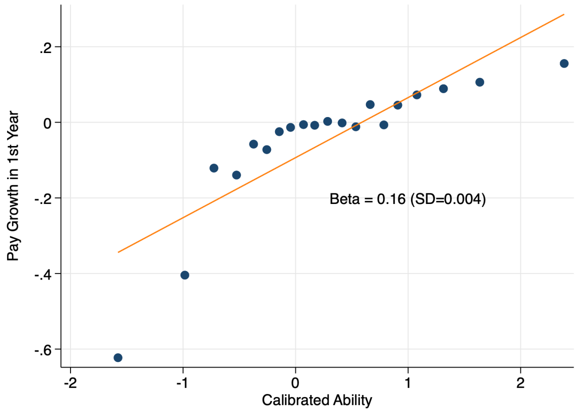

(b) Pay growth in the first year (new hires)

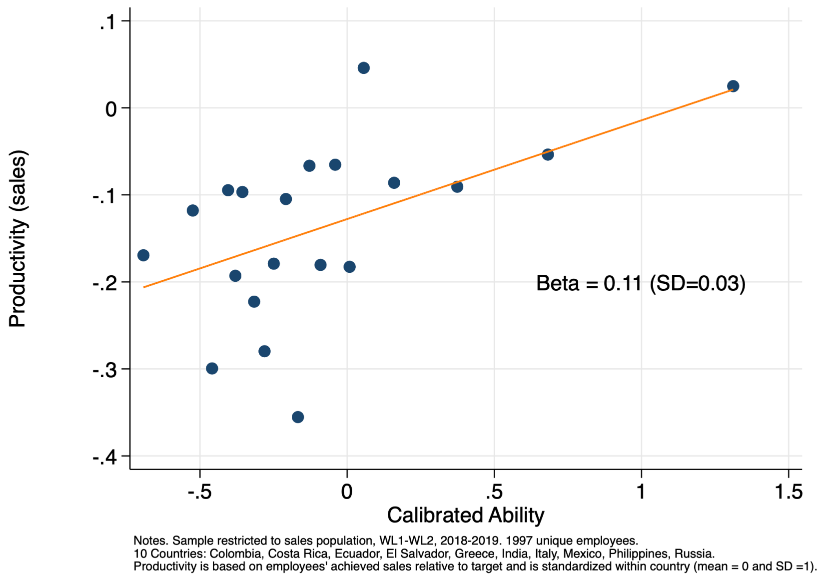

(c) Productivity (sales)

Notes. The figures are binned scatterplots and a linear fit of other performance measures (performance appraisals, pay growth for new hires, and objective productivity) against our calibrated productivity. The objective productivity measure is available only for the sales function in 10 countries and is based on reaching set targets.

Figure A.6: Validation of our calibrated parameters against labor regulations

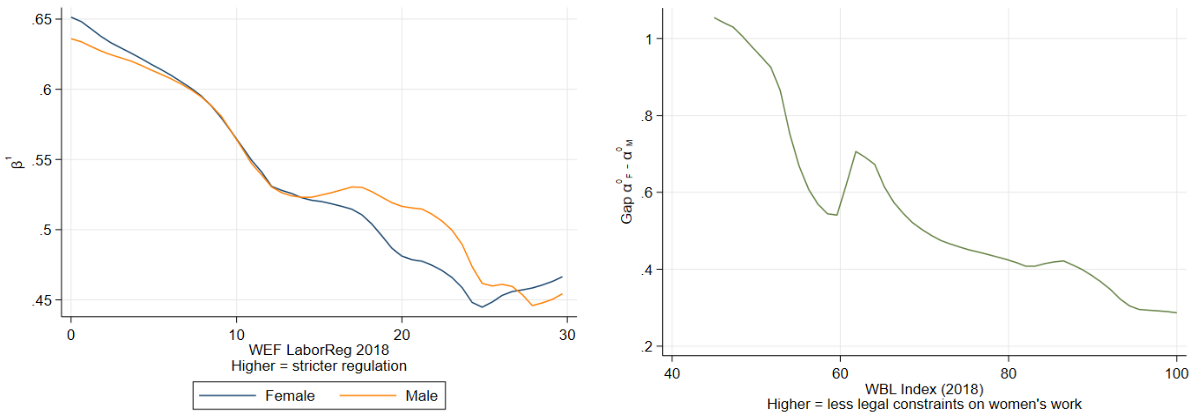

(a) Performance rewards (β^1) against the Restrictive Labor Regulations Index

(b) Value of staying at home ($\alpha_F^0 - \alpha_M^0$) against the Women, Business and the Law Index

Notes. Both panels show local polynomials of our calibrated parameters against two indices related to labor regulations. Panel (a) plots the calibrated β^1_{gtac}, for men and women separately, against the WEF Restrictive Labor Regulations Index. Panel (b) plots the gap in calibrated $\alpha_F^0 - \alpha_M^0$ against the WB Women, Business, and the Law index. Details about these indices are in the main text.

Figure A.7: Validation of our calibrated parameters against values

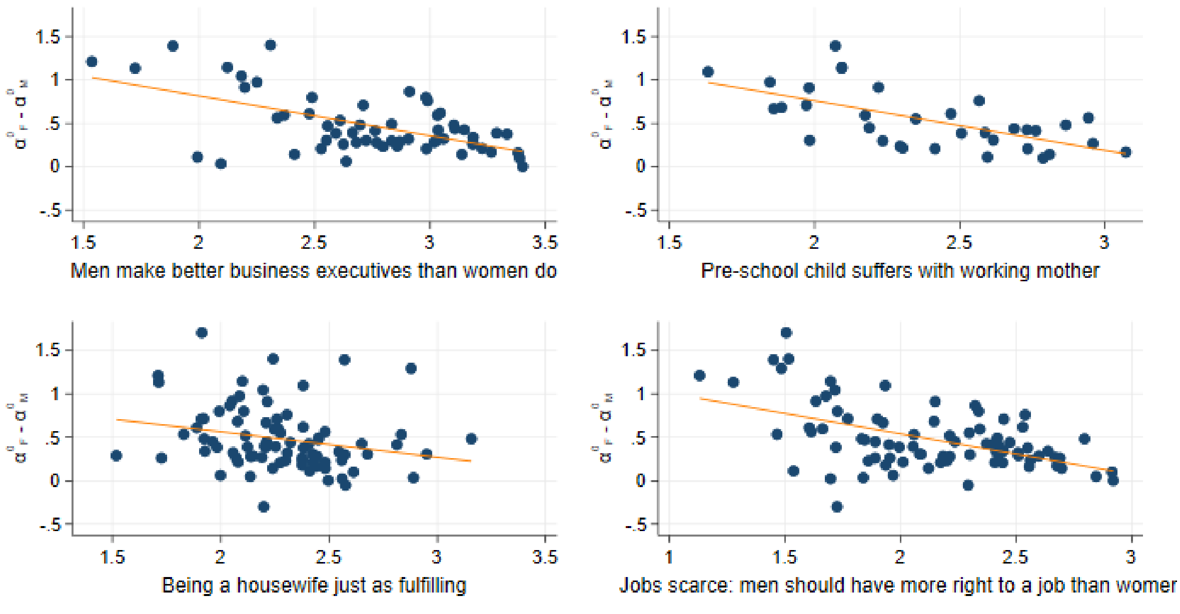

Notes. The figure shows scatterplots and fitted linear regressions of the gap in calibrated $\alpha^0_{Ftac} - \alpha^0_{Mtac}$ against four questions in the World Value Survey: (1) "Men make better business executives than women do," (2) "Pre-school child suffers with working mother," (3) "Being a housewife is just as fulfilling as working," (4) "When jobs are scarce, men should have more right to a job than women." For all questions, lower values of the index denote more agreement with the statement. Each dot is a country-cohort pair.

Figure A.8: The gender education gap versus the LFP gap

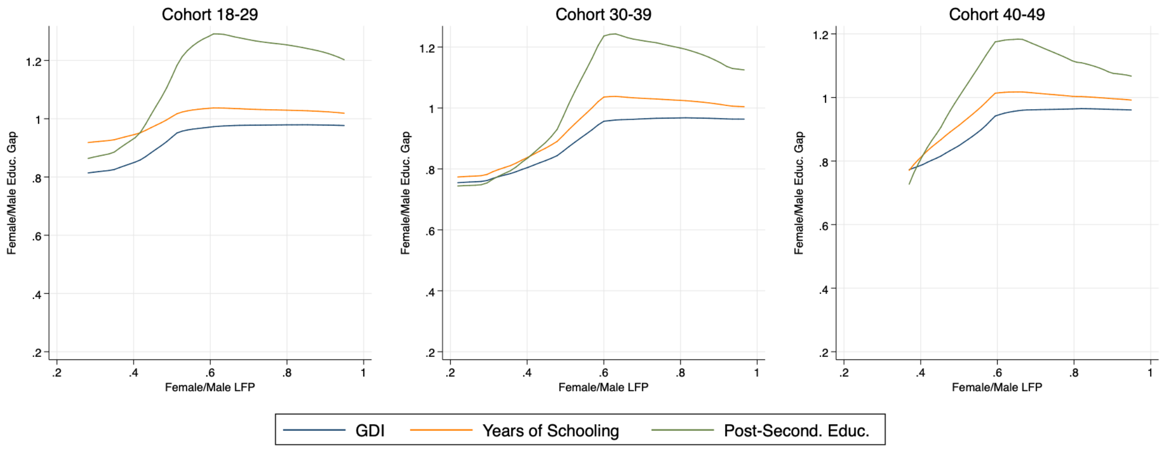

Notes. The figure plots the gender gap in education against the gender gap in LFP. We use a number of education measures: the gender development index (GDI, the ratio of female/male Human Development Index); years of schooling (female to male ratio), and the percentage in post-secondary education (female to male ratio). The GDI data is from the UNDP and the educational attainment data is from the World Bank.

A.3 Counterfactuals

A.3.1 Optimization problems for the firm's wage policy under the counterfactuals

The optimal policy we consider in subsection 5.1 solves, for each country-cohort-tenure cell, the following program:

$$\max_{(\alpha^1_{gtac},\beta^1_{gtac})_{g\in\{F,M\}}} \left[1-\Phi\left(\tilde{\xi}_{Ftac}\right)\right]\frac{\beta^1_{Ftac}}{\sigma_{Ftac}}\lambda\left(\tilde{\xi}_{Ftac}\right)+\left[1-\Phi\left(\tilde{\xi}_{Mtac}\right)\right]\frac{\beta^1_{Mtac}}{\sigma_{Mtac}}\lambda\left(\tilde{\xi}_{Mtac}\right)$$

$$\text{subj. to:}\quad 1-\Phi\left(\tilde{\xi}_{Ftac}\right)+1-\Phi\left(\tilde{\xi}_{Mtac}\right)=FLFP_{tac}+MLFP_{tac} \qquad \text{(i)}$$

$$\left[1-\Phi\left(\tilde{\xi}_{Ftac}\right)\right]\left[\alpha^1_{Ftac}+\frac{(\beta^1_{Ftac})^2}{\sigma_{Ftac}}\lambda\left(\tilde{\xi}_{Ftac}\right)\right]+$$

$$+\left[1-\Phi\left(\tilde{\xi}_{Mtac}\right)\right]\left[\alpha^1_{Mtac}+\frac{(\beta^1_{Ftac})^2}{\sigma_{Mtac}}\lambda\left(\tilde{\xi}_{Mtac}\right)\right]=\overline{y}^1_{Ftac}+\overline{y}^1_{Mtac} \qquad \text{(ii)}$$

where $\tilde{\xi}_{gtac}=(\alpha^0_{gtac}-\alpha^1_{gtac})/\sigma_{gtac}$, $\sigma_{gtac}=\sqrt{(\beta^1_{gtac})^2+(\nu^0_{gtac})^2}$. The objective is average productivity in the firm ($LFP\cdot\mathbb{E}[A\,|\,\text{empl'd}]$). Constraint (i) states that total employment (or LFP) should stay constant. Constraint (ii) states that the total wage bill should be unchanged.

The long-run constrained optimal policy considered in the counterfactuals in subsection 5.2

solves the program above, with the additional constraint:

$$\left| \left[\alpha^1_{Ftac} + \frac{(\beta^1_{Ftac})^2}{\sigma_{Ftac}} \lambda \left(\widetilde{\xi}_{Ftac} \right) \right] - \left[\alpha^1_{Mtac} + \frac{(\beta^1_{Mtac})^2}{\sigma_{Mtac}} \lambda \left(\widetilde{\xi}_{Mtac} \right) \right] \right| \leq B \qquad \text{(iii)}$$

Constraint (iii) places a limit B on how unequal the wage of the average men and women in the firm can be. The bound B is chosen to maximize the goodness of fit to the data, measured by the distance $\|\theta^*_B - \widehat{\theta}\|$ (where $\widehat{\theta}$ is the vector of calibrated $\alpha^1_{gtac}, \beta^1_{gtac}$ and θ^*_B is the solution of the constrained optimal policy problem with bound B).

A.3.2 Effects on workers' welfare

The effect of eliminating gender norms on wages (and hence on workers' utility) depends on the assumptions about the wage policy of the firm. If the wage parameters stay fixed at baseline levels, the only welfare effects are on the women who enter the labor force. By a revealed preference argument, those women gain because their wage in the firm is higher than their underlying value of staying at home (which is assumed to be the same as men's); the only reason why they were not in the labor force before at the current wages is the gender norm tax. There is no welfare effect on male workers nor on the women who are already in the labor force since their wages are unchanged. Hence, the net welfare effect on workers is positive in the short run.

When the firm adjusts, in contrast, because we are imposing the constraint that the wage bill stays fixed, there are winners and losers of the policy, and men can be affected by the elimination of the gender norm tax too. In particular, by the same revealed preference argument as above, those who enter the labor force will gain, and those who exit the labor force will lose. The sign of the welfare change for those who are always in the labor force depends on how the firm optimally reacts to the elimination of gender norms: in particular, if β_1 increases, as we find in Figure 10 panel (b), higher-productivity workers will benefit from a wage increase and lower-productivity workers will suffer from a wage decrease.

To illustrate, we focus on the two countries on opposite ends of the LFPR distribution: India and Sweden. The solid lines in Figure A.9 represent the participation frontier, $\beta^1 A_i - \nu^0 N_i \geq \alpha^0 - \alpha^1$, for India and Sweden in the long-run horizon (averaging across cohort and tenure groups, weighted by size). The dashed lines in the same figure show the change after the gender norm tax is eliminated. LFP and the participation frontier are very equal in Sweden, to begin with, hence the effect of equating α^0, ν^0 for men and women will be very small, and effectively there will be no welfare change. In contrast, in India, the participation frontiers are very different, to begin with, and equating the value of staying at home leaves room to attract high-productivity females into the labor force. Under the long-run wage policy we consider, once the gender norm tax is eliminated, the firm can maximize productivity by equating the wage parameters for men and women (which results, on average, in a wage increase for women and a wage decrease for men). The magnitude of this change depends on how constrained the firm was originally, as shown in Figure A.10.

The results for all countries are in Figure A.11, which plots the change in the value of chosen option y^* (which can be interpreted as a percentage change, because it is in units of log-wages). For men, $y^* = \max\{y^0, y^1\}$. For women, $y^* = y_F^1 \mathbf{1}\{y_F^1 \geq y_F^0\} + y_M^0 \mathbf{1}\{y_F^1 < y_F^0\}$ (i.e. they make a decision based on y_F^0 but the true value of staying at home is y_M^0, since we interpret the difference $y_F^0 - y_M^0$ as the gender norm tax).

Whether it is men or women who lose depends strongly on how much flexibility the firm had to react to differences in gender norms in the first place, and it is best understood by looking at the parameters under the counterfactual policies in Figure 10. For women, the increase in LFP and in variable pay is big enough to offset the decrease in fixed pay, so they are net winners. Men's LFP decreases and fixed pay falls enough that their overall welfare change is negative. Because we impose the constraint that the wage bill is kept constant, welfare changes are essentially a redistribution between female and male workers. It is important to note that this result is driven by the assumption that the firm cannot adjust employment levels. This is convenient for quantifying the pure effect of misallocation but it does not capture all the possible benefits of expanding women's access to the labor force. In a more general model where the firm can adjust employment in response to the change in gender norms, the effect is unlikely to be zero-sum.

Figure A.9: Examples: India and Sweden before and after equating norms

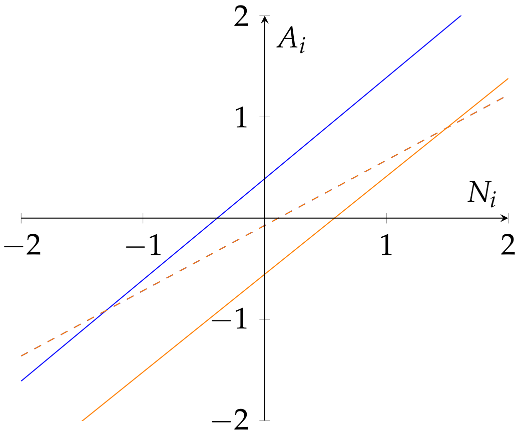

(a) India, constrained opt. policy

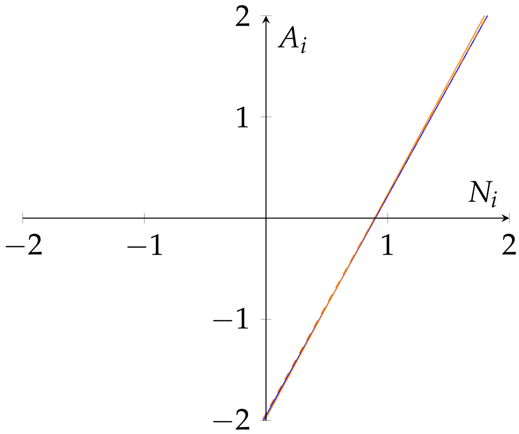

(b) Sweden, constrained opt. policy

Notes. The figures show the participation frontier, $\beta^1 A_i - \nu^0 N_i \geq \alpha^0 - \alpha^1$, for India and Sweden in the long-run horizon (averaging across cohort and tenure groups, weighted by size). Blue represents females, orange represents males. The solid line is the participation frontier with baseline α^0, ν^0; the dashed line is the counterfactual participation frontier after equating gender norms (and letting the wage policy respond). Those with (A_i, N_i) above and to the left of the participation frontier are in the labor force.

Figure A.10: Example: India, winners, and losers under different wage policies

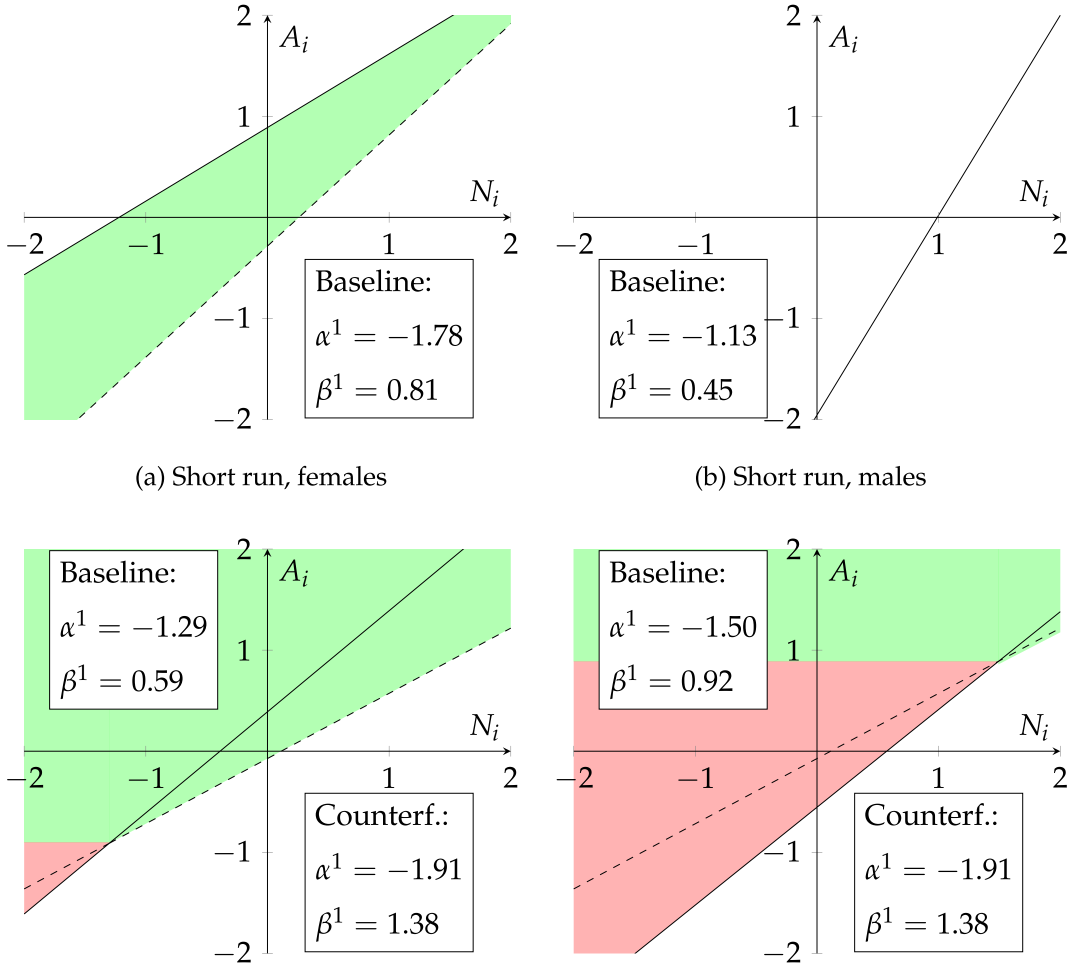

(a) Short run, females

(b) Short run, males

(c) Long run, females

(d) Long run, males

Notes. The figures show the participation frontier, $\beta^1 A_i - \nu^0 N_i \geq \alpha^0 - \alpha^1$, as well as the welfare effects for India under the two wage policy horizons we consider. The solid line is the participation frontier with baseline α^0, ν^0; the dashed line is the counterfactual participation frontier after equating gender norms (and letting the wage policy respond under the optimal policies). The area shaded green represents the population with (A_i, N_i) that benefits from equating gender norms, and the area shaded red is the population that is harmed.

Figure A.11: Welfare effects of eliminating the gender norms tax

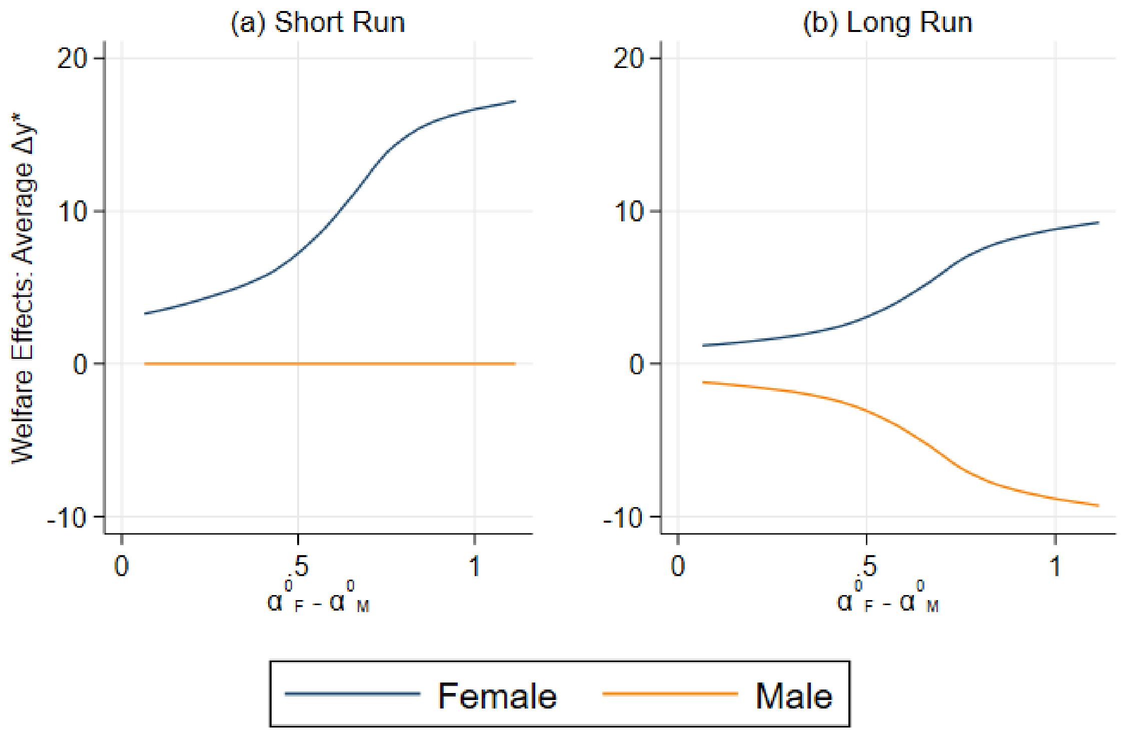

Notes. The figure shows the change in the value of chosen option y^* (which can be interpreted as a percentage change, because it is in units of log-wages) by the difference $\alpha_F^0 - \alpha_M^0$. For men, $y^* = \max\{y^0, y^1\}$. For women, $y^* = y_F^1 \mathbf{1}\{y_F^1 \geq y_F^0\} + y_M^0 \mathbf{1}\{y_F^1 < y_F^0\}$ (i.e. they make a decision based on y_F^0 but the true value of staying at home is y_M^0, since we interpret the difference $y_F^0 - y_M^0$ as the gender norm tax).

A.4 Implications for the gender pay gap - robustness

Table A.1: Gender pay gap and LFPR — LFP for population with advanced education

	Pay + Bonus (logs)				
	(1)	(2)	(3)	(4)	(5)
Female	0.427	0.269	0.161	0.303	0.381
	(0.173)	(0.082)	(0.150)	(0.079)	(0.073)
LFPR, advanced education	1.776	1.692	-0.138	1.684	0.073
	(0.188)	(0.113)	(0.222)	(0.116)	(0.273)
Female × LFPR, advanced education	-0.496	-0.375	-0.484	-0.401	-0.522
	(0.198)	(0.100)	(0.140)	(0.096)	(0.088)
GDP per capita (logs)			0.358		
			(0.033)		
Female × GDP per capita (logs)			0.020		
			(0.016)		
Controls	No	Yes	Yes	Yes	Yes
Cohort FE	No	No	No	Yes	Yes
Country FE	No	No	No	No	No
N	251336	251336	250152	251336	251336
R-squared	0.155	0.315	0.469	0.332	0.570

Notes. An observation is a worker-year. The LFPR is computed using the LFP for individuals with advanced education (short-cycle tertiary education or college degree and/or above). Controls include: tenure, tenure squared, year FE and function FE. Standard errors clustered at the country-cohort level.

Table A.2: Gender pay gap and LFPR — by region and income group

	All	Region FE	Lower inc.	Higher inc.
	(1)	(2)	(3)	(4)
	Pay + Bonus (logs)			
Female	0.256	0.208	0.265	-0.032
	(0.096)	(0.070)	(0.050)	(0.111)
LFPR	1.641	0.698	0.547	1.384
	(0.212)	(0.186)	(0.134)	(0.338)
Female × LFPR	-0.471	-0.424	-0.408	-0.129
	(0.135)	(0.100)	(0.068)	(0.156)
Controls	Yes	Yes	Yes	Yes
N	303756	303756	71658	232098
R-squared	0.285	0.387	0.218	0.239

Notes. An observation is a worker-year. Controls include: tenure, tenure squared, year FE and function FE. Standard errors clustered at the country-cohort level. Income group and geographical region are obtained from the World Bank.

Table A.3: Gender pay gap and LFPR — PPP conversion

	Pay + Bonus (logs), PPP 2017 USD				
	(1)	(2)	(3)	(4)	(5)
Female	0.318	0.195	0.312	0.197	0.198
	(0.127)	(0.080)	(0.141)	(0.080)	(0.076)
LFPR	0.266	0.249	0.085	0.259	0.129
	(0.200)	(0.136)	(0.155)	(0.130)	(0.239)
Female × LFPR	-0.463	-0.372	-0.275	-0.364	-0.381
	(0.165)	(0.104)	(0.083)	(0.105)	(0.102)
GDP per capita (logs)			0.044		
			(0.018)		
Female × GDP per capita (logs)			-0.021		
			(0.013)		
Controls	No	Yes	Yes	Yes	Yes
Cohort FE	No	No	No	Yes	Yes
Country FE	No	No	No	No	No
N	302789	302789	301600	302789	302789
R-squared	0.014	0.164	0.169	0.173	0.339

Notes. An observation is a worker-year. Controls include: tenure, tenure squared, year FE and function FE. Standard errors clustered at the country-cohort level. Wages are measured in PPP 2017 USD. Purchasing power parity (PPP) exchange rates are taken from the ICP (World Bank).

Table A.4: Gender pay gap and LFPR — fixed pay only

	Pay (logs)				
	(1)	(2)	(3)	(4)	(5)
Female	0.358	0.232	-0.023	0.230	0.167
	(0.135)	(0.092)	(0.162)	(0.091)	(0.074)
LFPR	1.582	1.578	0.501	1.599	0.144
	(0.282)	(0.212)	(0.145)	(0.204)	(0.234)
Female × LFPR	-0.533	-0.436	-0.241	-0.415	-0.328
	(0.185)	(0.131)	(0.114)	(0.129)	(0.100)
GDP per capita (logs)			0.283		
			(0.020)		
Female × GDP per capita (logs)			0.011		
			(0.015)		
Controls	No	Yes	Yes	Yes	Yes
Cohort FE	No	No	No	Yes	Yes
Country FE	No	No	No	No	No
N	303756	303756	302567	303756	303756
R-squared	0.110	0.285	0.448	0.309	0.552

Notes. An observation is a worker-year. Controls include: tenure, tenure squared, year FE and function FE. Standard errors clustered at the country-cohort level.