



Research
Education
Outreach

CCA

Gender Prescribed Occupations and the Wage Gap

Matteo Broso, Andrea Gallice, Caterina Muratori

No. 728
July 2024

Carlo Alberto Notebooks

www.carloalberto.org/research/workingpapers

Gender Prescribed Occupations and the Wage Gap

Matteo Broso*, Andrea Gallice†, Caterina Muratori‡

Abstract

Men and women often sort into different jobs, and male-dominated jobs typically pay more than female-dominated ones. Why is that the case? We propose a model where workers have heterogeneous attitudes with respect to the social norms that define gender prescribed occupations and face endogenous social costs when entering jobs deemed "appropriate" for the other gender. We show that: (i) workers trade off identity and wage considerations in deciding where to work; (ii) asymmetric social norms contribute to the gender pay gap by deterring women from entering higher-paying male-dominated sectors; (iii) breaking social norms generates positive externalities, reducing social stigma for everyone. Therefore, in equilibrium, there are too few social norm breakers.

JEL Classification: J24; J31

Keywords: Occupational Segregation; Wage Gap; Social Norms.

*Polytechnic University of Turin and Collegio Carlo Alberto

†University of Turin and Collegio Carlo Alberto

‡University of Barcelona and IEB; Center for Health Economics and Policy Studies, SDSU

1 Introduction

Gender differences in occupation and industry – i.e., *horizontal occupational segregation* – have been widely found to contribute to the gender wage gap [Groshe, 1991; Macpherson and Hirsch, 1995; Altonji and Blank, 1999; Blau and Kahn, 2017; Blau et al., 2022].¹ Blau and Kahn [2017] report that in 1980, horizontal occupational segregation accounted for 20% of the gap. By 2011, this percentage had increased to 50%, making horizontal occupational segregation the most relevant factor that accounts for male-female differences in wages. However, the literature has not yet provided a clear explanation of the reason why female occupations pay less than male occupations with similar measured characteristics [Blau and Kahn, 2017].

Multiple reasons have been advanced to explain why men and women sort into different jobs, social norms being one of them.² Gender social norms prescribe which jobs are "appropriate" for men and women. These societal prescriptions may align or clash with an individual's identity and thus influence their occupational choices, possibly contributing to occupational segregation. For instance, norms prescribing whether it is appropriate for women to work outside of the household appear to be key determinants of female labour force participation [Bursztyn et al., 2020].

In this paper, we build a sorting model in the labour market to show how social norms can generate occupational segregation and, in turn, contribute to the gender wage gap.

We consider a non-perfectly competitive labour market, where workers have market power. The lower the number of workers in a given sector, the higher their market power, and, therefore, their wage. Individuals have heterogeneous attitudes towards the norms that prescribe which jobs are more appropriate for men and which are more appropriate for women.

Inspired by the seminal paper by Akerlof and Kranton [2000], we let an individual's identity define how at ease the agent is in working in a masculine rather than in a feminine occupation. The agent discounts the monetary wage

¹The World Bank defines gender-based employment segregation as "the unequal distribution of female and male workers across and within job types". In particular, horizontal segregation is defined as "women and men concentrating in different sectors, industries, occupations, types of products, and business sizes" (Das and Kotikula [2019]).

²For a review, see Cortes and Pan [2018].

that they get in light of this self-perceived fit.³ Moreover, gender social norms impose a cost on those who decide to work in a sector that the society considers more appropriate for the other gender.⁴ This cost is endogenous and decreases in the share of norm-breaking workers. This feature of the model captures the idea that social norms can lead to "bad equilibria", where people do not coordinate on their most preferred action due to social stigma.

Within such a framework, and guided by the empirical evidence, we investigate the role of social norms in contributing to the gender wage gap. We start by noticing that gender social norms in the workplace are often stricter for women than men. [Usui \[2008\]](#) finds that women who work in male-dominated jobs report less job satisfaction than men who work in female-dominated jobs. [Lordan and Pischke \[2022\]](#) show that in occupations with a larger share of men, women are more likely to leave. [Morales and Marcén \[2023\]](#) report that where gender norms are stronger, women are less likely to enter male-dominated occupations. A key determinant of these results might be the risk of sexual harassment, which is higher in male-dominated occupations [[Hersch, 2011](#)] and is associated with occupational segregation and the resulting gender wage gap [[Folke and Rickne, 2022](#)]. Moreover, it is well documented that some male-dominated occupations are harder to reconcile with family commitments, as they display high rewards for long working hours and less flexibility. Such a job structure is often problematic for women, as they still carry a heavier load than men in the provision of housework and childcare [[Bertrand et al., 2015](#); [Schoonbroodt, 2018](#)]. In addition, women, but not men, are likely to suffer a penalty for parenthood in these workplaces [[Blau and Kahn, 2017](#); [Kleven et al., 2019](#)].

In light of this evidence, we explore the equilibrium predictions of our model under the assumption that the distributions of individual attitudes with respect to the occupational social norm differ across genders. More precisely, we postulate that the female distribution first-order stochastically dominates the male distribution. In other words, women are, on average, less at ease than men in working in sectors that are dominated by the other gender. This generates a wage gap. Because of the asymmetry in the incidence of social norms, the flow of women entering the male-dominated sector is smaller than the flow going in

³[Oh \[2023\]](#) provides empirical evidence about the connection between identity and job choice. They find that Indian workers are substantially less willing to take up job offers if the jobs require spending as little as ten minutes on tasks associated with castes other than their own.

⁴In this respect, we relate to papers that highlight the relevance of social concerns for individual and collective choices (see, for instance, [Bernheim \[1994\]](#), [Hopkins and Kornienko \[2004\]](#), [Levy and Razin \[2015\]](#), [Gallice and Grillo \[2020, 2024\]](#), and the references therein).

the opposite direction. Then, everything else equal, the male sector will pay higher wages and men on average will earn more than women. To the best of our knowledge, ours is the first paper that links (gender-asymmetric) social norms with the gender pay gap through this simple mechanism.

The model also highlights a perverse multiplier effect that amplifies the negative economic and social consequences on women: since they are less prone to switch sector than men, there will be fewer norm-breakers among women than among men. But since the social stigma that a norm-breaker bears is increasing in the number of norm compliers within their gender, fewer female norm-breakers imply larger social costs for women, further exacerbating the asymmetries between the two genders. The higher wage that the male sector pays countervails these forces, but only partially. Thus in equilibrium a gender wage gap survives and workers of both genders trade off their identity and wage considerations in deciding where to work.

Interestingly, breaking social norms generates positive externalities. If a worker decides to break the norm and enter a sector traditionally perceived as more appropriate for the other gender, they do it because they like that job so much that they are willing to sustain the social stigma. This choice turns out to be beneficial for other workers too. Indeed, all those other workers who are also inclined to break the norm now suffer lower social pressure from doing it. However, these positive externalities are non internalised at the individual level, and therefore in equilibrium there are too few social norm breakers. This suggests the desirability of policies aimed at boosting social norm breaking behaviors in the labour market, such as to dismantle the restraining effects of stereotypical gender prescribed occupations.⁵

The remainder of the paper is organised as follows. In Section 2 we describe the model setup. In Section 3, we solve the model and show how social norms can generate a gender wage gap. In Section 4 we provide a welfare analysis of equilibrium. Finally, Section 5 concludes.

⁵Let us consider the example of sports. Some sports, like football, are perceived as masculine. Some others, like dancing, are perceived as feminine. Our results show the desirability of policies subsidizing football camps for girls and dancing classes for boys. This will result in a lower social stigma for everybody, lead to a more efficient sorting, and thus ultimately increase welfare.

2 The Model

Consider an economy that features two occupational sectors: sector m , which is traditionally perceived as a "male" sector, and sector f , which is traditionally perceived as a "female" sector.

Players and Actions There is a continuum of workers with measure one. Let I be the set of workers. Each worker has a type (g, θ) , where $g \in \{m, f\}$ is the gender, and $\theta \in [0, 1]$ captures the alignment between the worker's identity and the occupational social norm. In other words, θ measures the individual's "self-perceived fit" in working in the sector to which, according to the social norm, they belong. Let $h(g, \theta)$ be the joint PDF, $h : I \rightarrow \{m, f\} \times [0, 1]$, and $h_g(\theta)$ be the PDF of θ conditional on g . We assume that gender is uniformly distributed, that is: $\mathbb{P}(g = f \mid \theta) = \mathbb{P}(g = m \mid \theta) = \frac{1}{2}$.⁶

Workers choose simultaneously in which sector to work: $\sigma \in \{m, f\}$. If $\sigma = g$, the worker works in the sector that is socially prescribed for their gender. If $\sigma \neq g$, the worker breaks the social norm.

Payoffs A worker of type (g, θ) that chooses action σ obtains the payoff:

$$u(\sigma \mid g, \theta) = \begin{cases} \theta w_\sigma & \text{if } \sigma = g \\ (1 - \theta) w_\sigma - c_g & \text{if } \sigma \neq g, \end{cases} \quad (2.1)$$

where $w_\sigma \geq 0$ is the monetary wage that sector σ pays.⁷ The worker discounts the monetary wage through θ . In particular, the worker enjoys a "fit-adjusted wage" equal to θw_σ if they work in the sector that, according to the social norm, is appropriate given their gender. The worker instead enjoys a fit-adjusted wage equal to $(1 - \theta) w_\sigma$ if they work in the sector that the social norm indicates as more appropriate for the other gender.

The labour market is not perfectly competitive. Let $B(g, \sigma)$ be the set of workers of gender g in sector σ . Let $\tilde{B}(g, \sigma) \in [0, \frac{1}{2}]$ be the measure of this set, that is, the share of workers of gender g in sector σ . The set of workers in sector σ is $B(\sigma) = B(m, \sigma) \cup B(f, \sigma)$, and the share of workers in this sector is

⁶To focus on the determinants of the wage gap that are unrelated to productivity differentials, we assume that all workers have the same productivity.

⁷Since workers choose in what sector to work, we use σ to denote both a player's action and the sector. Thus, for instance, a male that plays action σ gets the wage w_m if $\sigma = m$ and w_f if $\sigma = f$.

$\tilde{B}(\sigma) = \tilde{B}(m, \sigma) + \tilde{B}(f, \sigma)$, with $\tilde{B}(\sigma) \in [0, 1]$. Then, w_σ is given by:

$$w_\sigma = \frac{1}{\tilde{B}(\sigma)}, \quad (2.2)$$

so that the higher is the share of workers in the sector, the lower is the wage.

Finally, a worker of gender g who breaks the social norm and plays $\sigma \neq g$ pays a "social cost" equal to:

$$c_g = k \left(\frac{1}{2} - 2\tilde{B}(g, \sigma \neq g) \right), \quad (2.3)$$

where $k > 0$ is a proxy for the society's level of intolerance. The social cost is endogenous and gender-specific. It is positive and strictly decreasing for $\tilde{B}(g, \sigma \neq g) \in [0, \frac{1}{4}]$.⁸ Thus, a worker of gender g who breaks the social norm incurs a high social cost if they are one of the very few of their gender doing that. At the opposite, the cost approaches zero as the society is highly mobile and the sub-population of gender g equally splits between the two sectors, i.e., $\tilde{B}(g, \sigma \neq g) = \frac{1}{4}$.

3 Equilibrium

Our solution concept is Nash Equilibrium in Pure Strategies (equilibrium henceforth).

A worker of gender g is indifferent between working in the sector prescribed by the social norm or in the other sector whenever

$$u(\sigma = m \mid g, \theta) = u(\sigma = f \mid g, \theta). \quad (3.1)$$

Let $-g$ denote the complement of g in the set $\{m, f\}$. The indifferent type of gender g is given by $(g, \hat{\theta}_g)$, where:

$$\hat{\theta}_g = \max \left\{ 0, \frac{w_{-g}}{w_g + w_{-g}} - \frac{c_g}{w_g + w_{-g}} \right\}, \quad (3.2)$$

and $\hat{\theta}_g \in [0, 1]$. Payoff functions (2.1) are monotonic in θ , so the single crossing property holds. Then, all workers with $\theta > \hat{\theta}_g$ work in the sector prescribed by the social norm. On the contrary, all workers with type $\theta \leq \hat{\theta}_g$ violate the social

⁸When $\tilde{B}(g, \sigma \neq g) \in (\frac{1}{4}, \frac{1}{2}]$, the social cost turns into a benefit, as the majority of players of gender g choose to work in sector $\sigma \neq g$. We disregard this case as empirically not relevant, and we thus restrict our attention to equilibria where $\tilde{B}(g, \sigma \neq g) \in [0, \frac{1}{4}]$.

norm and work in the sector traditionally perceived as more appropriate for the other gender.⁹ Intuitively, and whenever strictly positive, $\hat{\theta}_g$ is increasing in the relative wage in sector $-g$ and decreasing in the social cost that a norm-breaker bears.

Let $H_g(x) = \mathbb{P}(\theta \leq x \mid g)$ be the CDF of θ conditional on g . Then, the share of workers of gender $g \in \{m, f\}$ who break the social norm and work in sector $\sigma = -g$ can be expressed as:

$$\tilde{B}(g, -g) = \frac{1}{2}H_g(\hat{\theta}_g), \quad (3.3)$$

so that the share of workers working in sector $\sigma = g$ is:

$$\tilde{B}(\sigma = g) = \frac{1}{2} \left((1 - H_g(\hat{\theta}_g)) + H_{-g}(\hat{\theta}_{-g}) \right). \quad (3.4)$$

Now let $\lambda \in [-1, 1]$ define the difference between the share of workers in the female sector and the share of workers in the male sector. Formally,

$$\lambda = \tilde{B}(\sigma = f) - \tilde{B}(\sigma = m) = H_m(\hat{\theta}_m) - H_f(\hat{\theta}_f). \quad (3.5)$$

Thus, when $\lambda > 0$ (respectively, $\lambda < 0$) the female sector is more (respectively, less) populated. By substituting (3.4) into (2.2), and using (3.5), the wages in the two sectors can be defined in terms of λ :

$$w_m = \frac{2}{1 - \lambda} \quad \text{and} \quad w_f = \frac{2}{1 + \lambda}. \quad (3.6)$$

Social costs that a worker who plays $\sigma \neq g$ suffers can instead be expressed as follows:

$$c_g = k \left(\frac{1}{2} - H_g(\hat{\theta}_g) \right) \quad (3.7)$$

Solving the system of equations given by (3.2), (3.5), (3.6), and (3.7) leads to the following equilibrium values for the indifferent types and the wages in the

⁹Without loss of generality, we break indifference in favor of the norm breaking behavior.

two sectors:

$$\begin{aligned}
\hat{\theta}_m^* &= \max \left\{ 0, \frac{1 - \lambda^*}{2} - \frac{1 - (\lambda^*)^2}{4} k \left(\frac{1}{2} - H_m(\hat{\theta}_m^*) \right) \right\} \\
\hat{\theta}_f^* &= \max \left\{ 0, \frac{1 + \lambda^*}{2} - \frac{1 - (\lambda^*)^2}{4} k \left(\frac{1}{2} - H_f(\hat{\theta}_f^*) \right) \right\} \\
w_m^* &= \frac{2}{1 - \lambda^*} \\
w_f^* &= \frac{2}{1 + \lambda^*}
\end{aligned} \tag{3.8}$$

with

$$\lambda^* = H_m(\hat{\theta}_m^*) - H_f(\hat{\theta}_f^*) . \tag{3.9}$$

We first briefly discuss the trivial equilibrium case in which there is no mobility across sectors. No mobility occurs when even the most inclined individual to break the norm (the worker of type $(g, \theta = 0)$) chooses to remain in sector g . Intuitively (see 2.1), this happens when social costs of changing sector are too high. We now solve for this threshold level of k . If no worker moves then $\hat{\theta}_g^* = 0$ and $H_g(\hat{\theta}_g^*) = 0$ for any $g \in \{m, f\}$, so that $\lambda^* = 0$. Then, for (3.8) to hold, it must necessarily be the case that:

$$\frac{1}{2} - \frac{1}{4} k \left(\frac{1}{2} - 0 \right) \leq 0, \tag{3.10}$$

which is satisfied for any $k \geq 4$. Thus, whenever $k \geq 4$, in equilibrium there is no mobility between sectors, $w_m^* = w_f^* = 2$, and there is no gender wage gap.

We now focus on the more interesting equilibria that feature some degree of mobility across sectors. These equilibria emerge when $k < 4$. In studying these equilibria, we assume that the social norm is asymmetric across genders and show how this asymmetry can generate a wage gap. In particular, and in line with the evidence discussed in the Introduction, we assume that the distribution of θ among females first-order stochastically dominates the distribution of θ among males.

Assumption 1. $H_f(\theta) < H_m(\theta)$ for all $\theta \in (0, 1)$.

Thus, under Assumption 1, females are, on average, less likely than males to experience low realizations of θ .




Figure 1: Equilibrium thresholds θ_m^* and θ_f^* .

Let $\bar{\sigma}$ denote a profile of players' actions. Formally, $\bar{\sigma}$ is a mapping from types (g, θ) to sectors σ . Then, we can define the gender wage gap as follows:

$$\Delta w = \mathbb{E}[w(\bar{\sigma} | m)] - \mathbb{E}[w(\bar{\sigma} | f)], \quad (3.11)$$

and thus state the following result.

Proposition 1. *Let Assumption 1 hold. Then, in equilibrium, $\Delta w^* > 0$. The gender wage gap Δw^* is strictly increasing in λ^* .*

Proof. See Appendix A. □

Proposition 1 shows that in equilibrium the average wage for male workers is higher than the average wage for female workers. In particular, the proof of Proposition 1 shows that, in spite of the fact that $\hat{\theta}_f^* > \hat{\theta}_m^*$, we have that $H_f(\hat{\theta}_f^*) < H_m(\hat{\theta}_m^*)$. This is shown in Figure 1. There is a larger fraction of social norm breakers among males, which in turn implies a more populated female sector and a lower wage w_f^* . Nonetheless, since sector m pays a higher wage, $\hat{\theta}_m^* < \hat{\theta}_f^*$. This force countervails, but only partially, the lower mobility of females, as well as the higher mobility of males.

It is then interesting to investigate how the gender wage gap could possibly depend on k , which we defined as a proxy for the society' level of intolerance towards norm-breakers. We know that if $k \geq 4$, then $\hat{\theta}_g^* = H_g(\hat{\theta}_g^*) = 0$ for any $g \in \{m, f\}$. Now, start from $k = 4$ and consider a marginal decrease in k . This reduction generates some mobility of workers across sectors. Some men (those

with types close to $\theta = 0$) will choose to work in sector $\sigma = f$. In the same way, some of the women with the lowest types will choose $\sigma = m$. The individual and gender specific incentives to change sector impact the equilibrium wage gap through two effects.

To disentangle these effects, suppose first that $\hat{\theta}_m^*$ and $\hat{\theta}_f^*$ increase by the same extent. By Assumption 1, this implies that more men than women break the social norm and choose to work in the sector perceived as more appropriate for the other gender. We name this the "Conformity Effect." As we assumed that men prefer working with women more than vice-versa, a reduction in social stigma makes men relatively more socially mobile than women, thereby positively contributing to the gender wage gap.

However, the marginal impact of a decrease in k on the threshold values $\hat{\theta}_m^*$ and $\hat{\theta}_f^*$ is actually asymmetric (see the equilibrium characterization in the proof of Proposition 1). This generates a second effect on the wage gap. To illustrate this effect, consider the expression for the endogenous social costs that a norm-breaker bears (see 3.7). Since in equilibrium $H_m(\hat{\theta}_m^*) > H_f(\hat{\theta}_f^*)$, then $c_f^* > c_m^*$. That is, female norm-breakers bear higher social costs than male norm-breakers. Moreover,

$$\frac{\partial (c_f^* - c_m^*)}{\partial k} > 0. \quad (3.12)$$

The intuition goes as follows. In equilibrium, more men than women break the social norm. Since social costs are endogenous and proportional to the mass of workers of the same gender that comply with the social norm, a reduction in k decreases social costs for women more than for men, thus making women marginally more inclined to mobility. That is, the difference $\hat{\theta}_f^* - \hat{\theta}_m^*$ can be increasing in k . This "Social Cost Effect" can thus have opposite sign with respect to the "Conformity Effect".

The net effect of k on social mobility and the wage gap is thus not univocal as it depends on the specific functional forms of the distributions $h_m(\theta)$ and $h_f(\theta)$. However, the trade-off that we discussed is general, and it follows from the dual nature of social norms. On the one hand, more men than women are inclined to break social norms. Then, as breaking social norms becomes cheaper, men initially benefit more, potentially resulting in a net positive flow of workers into the female sector, which negatively affects the wage rate in that sector and thus increases the wage gap. On the other hand, the cost of breaking social norms increases with the share of social norm compliers of a given gender. As

more women comply than men, a marginal reduction in the stringency of social norms benefits women more than men, positively contributing to their relative salary.

Notice finally that in a fully tolerant society where breaking social norms was not costly at all ($k = 0$), the gender wage gap in equilibrium would still be positive. We saw instead that the wage gap would be zero in a society that is extremely intolerant ($k \geq 4$). The comparison between these two polar cases already indicates that the minimization of the wage gap should not necessarily be the (unique) objective of economic policies. Indeed, in our model, such a goal could be reached by blocking mobility across the two sectors. However, such an outcome would perform very poorly from an aggregate welfare point of view, because it completely ignores workers' identities and desires. We thus now turn to a more careful investigation of the welfare implications of our model.

4 Welfare Analysis

Suppose that a benevolent social planner can allocate workers to sectors to maximise utilitarian welfare (the sum of workers' payoffs). How would this first best solution compare with the equilibrium solution obtained in Section 3?

Define the social planner's problem as follows.

$$\max_{\bar{\sigma}} \sum_{g \in \{m, f\}} \int_0^1 u(\bar{\sigma} | g, \theta) h_g(\theta) d\theta \quad (4.1)$$

Consider the sub-population of gender g . The planner chooses who to allocate to sector m and who to allocate to sector f . Formally, the planner chooses the welfare-maximizing sets $B^W(g, m)$ and $B^W(g, f)$. We have the following result.

Proposition 2. For any $g \in \{m, f\}$, the welfare-maximizing sets $B^W(g, \sigma)$ are such that

$$\begin{aligned} B^W(g, g) &= \left\{ (g, \theta) : g = g \text{ and } \theta \geq \hat{\theta}_g^W \right\}, \\ B^W(g, -g) &= \left\{ (g, \theta) : g = g \text{ and } \theta < \hat{\theta}_g^W \right\}, \end{aligned} \quad (4.2)$$

for some $\hat{\theta}_g^W \in [0, 1]$. Moreover, the relation $\hat{\theta}_g^W \geq \hat{\theta}_g^*$ holds so that in equilibrium there are too few social norm breakers.

Proof. See Appendix B. □

The proof of Proposition 2 follows in three steps. First, we show that the $B^W(g, \cdot)$ sets have a floor. If it is optimal to allocate worker with type (θ', g) to sector $\sigma = g$, it is also optimal to allocate workers with types (θ, g) for all $\theta > \theta'$. This follows from the single crossing property of agents' preferences. Second, we show that the allocation of workers across sectors does not impact the aggregate (and identity-adjusted) wage. This allows the social planner to select $\hat{\theta}_g^W$ by looking at social costs only. Third, we compare two cases: a marginal increase and a marginal decrease of equilibrium thresholds $\hat{\theta}_g^*$. Both marginal effects have no impact on identity-adjusted aggregate wages. For instance, consider a marginal increase in $\hat{\theta}_g^*$. This decreases the salary for workers in sector $-g$ and increases the salary for workers in sector g by the same extent. However, a marginal increase to $\hat{\theta}_g^*$ creates a positive externality: it reduces the social cost for all other norm-breakers with gender g without making any other worker worse off.

Proposition 2 shows that breaking social norms can be thought of as "good" with positive externality. If a worker decides to enter a sector traditionally perceived as more appropriate for the other gender, they do it because they find it optimal and, therefore, gain a private benefit. They like that job so much that they are willing to sustain the social cost associated with this switch. However, this behavior is beneficial for other workers too. All other workers who have a preference for that job will now suffer lower social cost from breaking the social norm. As any other good with positive externality, norm-breaking behavior is under-provided in equilibrium. From a policy perspective, Proposition 2 shows the possible desirability of policies subsidizing social mobility and norm-breaking behavior.

5 Conclusions

The extensive literature on the gender pay gap consistently highlights occupational segregation as a critical factor [Blau and Kahn, 2017], with gender norms significantly influencing individual career choices [Morales and Marcén, 2023]. Building on the seminal work by Akerlof and Kranton [2000], we construct a model of sorting into the labour market to show how gender norms can generate occupational segregation and, in turn, contribute to the gender wage gap. Our model contributes to the existing literature by demonstrating how the interaction between gender norms and the structural characteristics of male-dominated jobs drives disparities in gender-based occupational sorting and, subsequently, generates a wage gap. Specifically, long working hours, limited flexibility, substantial motherhood penalties, and a higher incidence of sexual harassment discourage women from entering traditionally male-dominated fields. When the flow of women entering the male-dominated sector is smaller than the flow going in the opposite direction, male-dominated sectors, all else being equal, tend to offer higher wages.

Our findings underscore the need for targeted interventions to (1) reform male-dominated job structures to make them more appealing to women and (2) diminish the influence of rigid gender norms to encourage nontraditional career paths for both men and women. Our model shows that policies encouraging norms-breaking behavior may help mitigate the pay gap issue by reducing the social stigma associated with gender non-conforming occupational choices.

Appendix

A Proof of Proposition 1

The system of equations given by (3.2), (3.4), (3.5) and the two wage-setting equations lead to the equilibrium solutions defined in (3.8). Consider now the difference between the threshold values $\hat{\theta}_f^*$ and $\hat{\theta}_m^*$. This difference is given by:

$$\Delta_{\hat{\theta}^*} = \hat{\theta}_f^* - \hat{\theta}_m^* = \lambda^* \left(1 - \frac{1 - (\lambda^*)^2}{4} k \right). \quad (\text{A-1})$$

The function $\Delta_{\hat{\theta}^*}$ is a cubic polynomial in λ^* and thus possibly displays three zeros:

$$\lambda^* \in \left\{ -\frac{1}{k} \sqrt{k(k-4)}, 0, \frac{1}{k} \sqrt{k(k-4)} \right\} \Rightarrow \lambda^* \left(1 - \frac{1 - (\lambda^*)^2}{4} k \right) = 0. \quad (\text{A-2})$$

However, we saw in the main text that any equilibrium that features some mobility of workers between sectors necessarily requires $k < 4$. With $k < 4$, the function $\Delta_{\hat{\theta}^*}$ has a unique zero, which occurs at $\lambda^* = 0$.




Figure A.2: The shape of $\Delta_{\hat{\theta}^*} = \hat{\theta}_f^* - \hat{\theta}_m^*$ as a function of λ^* .

Figure A.2 depicts the shape of $\Delta_{\hat{\theta}^*} = \hat{\theta}_f^* - \hat{\theta}_m^*$ as a function of λ^* and for different values of $k \in (0, 4)$. There are thus three possible cases to discuss: λ^* is negative and the function $\Delta_{\hat{\theta}^*}$ is negative (that is, $\Delta_{\hat{\theta}^*}$ is in quadrant III of the

Cartesian plane); λ^* is zero and the function $\Delta_{\hat{\theta}^*}$ is zero too; λ^* is positive and the function $\Delta_{\hat{\theta}^*}$ is positive (that is, $\Delta_{\hat{\theta}^*}$ is in quadrant I).

Formally, the three cases are:

- (a) $\lambda^* < 0$ and $\hat{\theta}_f^* < \hat{\theta}_m^*$
- (b) $\lambda^* = 0$ and $\hat{\theta}_f^* = \hat{\theta}_m^*$
- (c) $\lambda^* > 0$ and $\hat{\theta}_f^* > \hat{\theta}_m^*$

Cases (a) and (b) are impossible given Assumption 1: if $\hat{\theta}_f^* \leq \hat{\theta}_m^*$, then because of first order stochastic dominance it must be necessarily the case that $H_f(\hat{\theta}_f^*) < H_m(\hat{\theta}_m^*)$, i.e., $\lambda^* > 0$. The equilibrium thus necessarily belongs to case (c). Therefore, in equilibrium:

$$\begin{aligned}\hat{\theta}_f^* &> \hat{\theta}_m^* \\ H_f(\hat{\theta}_f^*) &< H_m(\hat{\theta}_m^*) \\ \lambda^* &\in (0, 1).\end{aligned}\tag{A-3}$$

Now consider average wages across genders. In equilibrium, $(1 - H_m(\hat{\theta}_m^*))$ of the male workforce works in the male sector and the remaining $H_m(\hat{\theta}_m^*)$ works in the female sector. Equilibrium wages are $w_m^* = \frac{2}{1 - \lambda^*}$ and $w_f^* = \frac{2}{1 + \lambda^*}$. Then, the male average wage is given by

$$\mathbb{E}[w(\bar{\sigma}^* | m)] = \frac{1 + \lambda^* (1 - 2H_m(\hat{\theta}_m^*))}{1 - (\lambda^*)^2}\tag{A-4}$$

Similarly, the female average wage is:

$$\mathbb{E}[w(\bar{\sigma}^* | f)] = \frac{1 - \lambda^* (1 - 2H_f(\hat{\theta}_f^*))}{1 - (\lambda^*)^2}\tag{A-5}$$

The wage gap is then given by:

$$\Delta w^* = 2\lambda^* \frac{(1 - H_m(\hat{\theta}_m^*) - H_f(\hat{\theta}_f^*))}{1 - (\lambda^*)^2}.\tag{A-6}$$

Since $\lambda^* \in (0, 1)$ and given that $H_g(\hat{\theta}_g^*) < \frac{1}{2}$ for any $g \in \{m, f\}$, the wage gap Δw^* is positive. Clearly, Δw^* is increasing in λ^* . \square

B Proof of Proposition 2

The proof involves three steps.

Step 1 We show that the sets $B^W(g, g)$ have a floor. Consider the subpopulation of workers of gender $g \in \{m, f\}$. Suppose it is optimal (in the sense of (4.1)) to allocate a worker of type θ' to the set $B^W(g, g)$. We want to prove that it is welfare-maximizing to allocate all workers with $\theta > \theta'$ to $B^W(g, g)$. We proceed by contradiction.

Define the interval $A = [\theta^{\underline{A}}, \theta^{\bar{A}}]$, where $\theta' < \theta^{\underline{A}} \leq \theta^{\bar{A}} \leq 1$. Suppose, towards a contradiction, that allocating workers with types $\theta \in A$ to the set $B(g, -g)$ is welfare maximizing.

Consider now the alternative interval $A' = [\theta', \theta^{\bar{A}'}]$, where $\theta^{\bar{A}'}$ is such that

$$\int_{\theta'}^{\theta^{\bar{A}'}} h_g(\theta) d\theta = \int_{\theta^{\underline{A}}}^{\theta^{\bar{A}}} h_g(\theta) d\theta. \quad (\text{A-7})$$

Allocating to the set $B(g, -g)$ the workers with $\theta \in A'$ achieves higher welfare with respect to allocating to the set those with $\theta \in A$. In fact, both changes have the same effect on wages and social costs, but workers in A' would enjoy higher utility (with respect to workers in A) from working in sector $\sigma \neq g$. Therefore, it is welfare-improving to allocate the worker with type θ' to sector $\sigma \neq g$. We reached a contradiction. Figure B.3 shows the sets A, A' , and the logic of this step.

This step allows us to simplify problem (4.1). In particular, for any genders $g \in m, f$, the social planner can simply choose the threshold $\hat{\theta}_g^W$, allocate all workers with $\theta \geq \hat{\theta}_g^W$ to the set $B^W(g, g)$, and all workers with $\theta < \hat{\theta}_g^W$ to the set $B^W(g, -g)$. Since the problem is concave, the thresholds $\hat{\theta}_g^W$ exist and are unique.¹⁰

Step 2 We now show that moving workers across sectors does not impact the aggregate (and preference-adjusted) wage. This will be useful in the next step

¹⁰Concavity of the aggregate welfare function follows from the concavity of (2.1). Since (2.1) is concave and bounded, the expectation exists, and it is finite and concave.




Figure B.3: The intervals A, A' are such that $\int_A h_g(\theta)d\theta = \int_{A'} h_g(\theta)d\theta$.

to show that the social planner can choose the optimal thresholds $(\hat{\theta}_f^W, \hat{\theta}_m^W)$ by considering social costs only.

The aggregate wage does not depend on the allocation of workers between sectors, i.e., it does not depend on λ :

$$\begin{aligned} \mathbb{E}[w_\sigma(\bar{\sigma} | g, \theta)] &= \\ \tilde{B}(m)w_m + \tilde{B}(f)w_f &= \\ \left(\frac{1}{2}(1-\lambda)\right)\frac{2}{1-\lambda} + \left(\frac{1}{2}(1+\lambda)\right)\frac{2}{1+\lambda} &= 2. \end{aligned} \tag{A-8}$$

The intuition for (A-8) is the following. When λ goes up, the female sector becomes relatively more populated with respect to the male sector, thus leading to a reduction of the wage in the female sector and an increase of the wage in the male sector. The two effects cancel out.

Workers discount wages according to their self-perceived fit in the sector. Then, we need to show that (in aggregate terms) preference-adjusted wages do not depend on λ . The average (or aggregate) preference-adjusted wage can be written as:

$$\mathbb{E}[\theta w_\sigma(\bar{\sigma} | g, \theta)] = \mathbb{E}[w_\sigma(\bar{\sigma} | g, \theta)]\mathbb{E}[\theta] + COV(\theta, w_\sigma(\bar{\sigma} | g, \theta)). \tag{A-9}$$

Since a worker's wage is not a function of θ , $COV(\theta, w_\sigma(\bar{\sigma} | g, \theta)) = 0$, so that $\mathbb{E}[\theta w_\sigma(\bar{\sigma} | g, \theta)]$ is also constant in λ (and $\bar{\sigma}$).

Step 3 In this last step, we show that $\hat{\theta}_g^* \leq \hat{\theta}_g^W$.

By the previous step, we know that moving workers between sectors does not impact aggregate preference-adjusted wages. We can thus compare $\hat{\theta}_g^*$ and

$\hat{\theta}_g^W$ by focusing on social costs only.

Consider the sub-population of gender g , and the corresponding equilibrium floor $\hat{\theta}_g^*$. Take an arbitrary $\epsilon \rightarrow 0_+$, and consider the following alternative floors: $\hat{\theta}_g^* + \epsilon$ and $\hat{\theta}_g^* - \epsilon$. The two marginal changes have the same (nil) impact on aggregate preference-adjusted wages. However, $\hat{\theta}_g^* + \epsilon$ generates a positive externality because increasing the share of norm-violators decreases the cost of breaking the norm for everybody. Therefore, $\hat{\theta}_g^* + \epsilon$ is (weakly) socially preferable to $\hat{\theta}_g^*$, whereas $\hat{\theta}_g^* - \epsilon$ is not. Therefore, $\hat{\theta}_g^* \leq \hat{\theta}_g^W$. \square

References

- Akerlof, G. A. and Kranton, R. E. [2000], 'Economics and identity', *The Quarterly Journal of Economics* **115**(3), 715–753.
- Altonji, J. G. and Blank, R. M. [1999], 'Race and gender in the labor market', *Handbook of Labor Economics* **3**, 3143–3259.
- Bernheim, B. D. [1994], 'A theory of conformity', *Journal of Political Economy* **102**(5), 841–877.
- Bertrand, M., Kamenica, E. and Pan, J. [2015], 'Gender identity and relative income within households', *The Quarterly Journal of Economics* **130**(2), 571–614.
- Blau, F. D., Ferber, M. A. and Winkler, A. E. [2022], *The economics of women, men, and work*, Oxford University Press.
- Blau, F. D. and Kahn, L. M. [2017], 'The gender wage gap: Extent, trends, and explanations', *Journal of Economic Literature* **55**(3), 789–865.
- Bursztyjn, L., González, A. L. and Yanagizawa-Drott, D. [2020], 'Misperceived social norms: Women working outside the home in Saudi Arabia', *American Economic Review* **110**(10), 2997–3029.
- Cortes, P. and Pan, J. [2018], 'Occupation and gender', *The Oxford handbook of women and the economy* pp. 425–452.
- Das, S. and Kotikula, A. [2019], *Gender-based employment segregation: Understanding causes and policy interventions*, World Bank.
- Folke, O. and Rickne, J. [2022], 'Sexual harassment and gender inequality in the labor market', *The Quarterly Journal of Economics* **137**(4), 2163–2212.
- Gallice, A. and Grillo, E. [2020], 'Economic and social-class voting in a model of redistribution with social concerns', *Journal of the European Economic Association* **18**(6), 3140–3172.
- Gallice, A. and Grillo, E. [2024], 'Shifting social norms through endorsements', *Mimeo*.
- Groshen, E. L. [1991], 'The structure of the female/male wage differential: Is it who you are, what you do, or where you work?', *Journal of Human Resources* pp. 457–472.

- Hersch, J. [2011], 'Compensating differentials for sexual harassment', *American Economic Review* **101**(3), 630–634.
- Hopkins, E. and Kornienko, T. [2004], 'Running to keep in the same place: Consumer choice as a game of status', *American Economic Review* **94**(4), 1085–1107.
- Kleven, H., Landais, C. and Søgaaard, J. E. [2019], 'Children and gender inequality: Evidence from denmark', *American Economic Journal: Applied Economics* **11**(4), 181–209.
- Levy, G. and Razin, R. [2015], 'Preferences over equality in the presence of costly income sorting', *American Economic Journal: Microeconomics* **7**(2), 308–37.
- Lordan, G. and Pischke, J.-S. [2022], 'Does rosie like riveting? male and female occupational choices', *Economica* **89**(353), 110–130.
- Macpherson, D. A. and Hirsch, B. T. [1995], 'Wages and gender composition: why do women's jobs pay less?', *Journal of Labor Economics* **13**(3), 426–471.
- Morales, M. and Marcén, M. [2023], 'The effect of gender norms on gender-based sorting across occupations', *Review of Economics of the Household* pp. 1–34.
- Oh, S. [2023], 'Does identity affect labor supply?', *American Economic Review* **113**(8), 2055–83.
- Schoonbroodt, A. [2018], 'Parental child care during and outside of typical work hours', *Review of Economics of the Household* **16**, 453–476.
- Usui, E. [2008], 'Job satisfaction and the gender composition of jobs', *Economics Letters* **99**(1), 23–26.