Discussion Paper Series

IZA DP No. 18843

July 2026

Preferences, Discrimination, and Occupational Segregation

Girish Bahal
University of Western Australia

Jan Feld
Victoria University of Wellington
and IZA@LISER

Anand Shrivastava
Azim Premji University

IZA@LISER NETWORK
ISSN: 2365-9793

IZA DP No. 18843 July 2026

Preferences, Discrimination, and Occupational Segregation*

Abstract

Preferences and discrimination are often treated as competing explanations for occupational segregation. This framing misses an important feedback. Preference-based sorting changes the sex composition of occupations; composition changes expected discrimination; and expected discrimination changes sorting. We formalize this mechanism through a general model, where occupations differ in flexibility and workers face more discrimination when their sex is underrepresented. We show two results. First, discrimination amplifies segregation caused by preferences. Second, preferences and discrimination are negatively correlated across occupations. We use the model to study anti-discrimination policies, preference-equalizing policies, and policies that make high-wage occupations more flexible.

JEL classification

J16, J24, J71

Keywords

labor market segregation, preferences, discrimination, sex, gender

Corresponding author

Jan Feld

jan.feld@vuw.ac.nz

* We are grateful to seminar participants at the University of Cambridge, New Zealand Economics eSeminar series, Alison Andrew, Erika Deserranno, Joseph-Simon Görlach, Rajiv Sethi, and Travis Lybbert for helpful discussions and comments.

1 Introduction

Why do women and men work in different occupations? Two common answers are preferences and discrimination (Blau and Kahn, 2017; Cortes and Pan, 2018). The preferences answer says that women and men value different job characteristics. The discrimination answer says that workers avoid occupations where they expect to be treated poorly. These explanations are often treated as independent. We argue that they are connected through an important feedback mechanism.

Suppose women value flexibility more than men. Even without discrimination, women will be more likely to choose more flexible occupations, while men will be more likely to choose less flexible occupations. This sorting changes the sex composition of occupations. Less flexible occupations have a higher share of men. If women expect worse treatment as the share of men in an occupation rises, those occupations become even less attractive to women. The initial sorting due to preference differences is therefore amplified by discrimination.[1]

We formalize this feedback in a model of occupational choice. Occupations differ in wages and flexibility. Women and men value wages equally, but women have a higher willingness to pay for flexibility. Workers also care about discrimination, which depends on the sex composition of an occupation. Women face more discrimination as the share of men in an occupation rises. Men face more discrimination as the share of women rises.

We study the feedback mechanism first in a model with two occupations. One occupation offers more flexibility and pays less. The other offers less flexibility and pays more. Workers evaluate occupations based on wages, flexibility, and expected discrimination. Expected discrimination is determined by the sex composition inherited from the previous cohort. We then show that the same logic extends to a setting with many occupations.

The model has two main results. First, discrimination increases the segregation caused by preferences. In the preferences-only benchmark, sex differences in preferences for flexibility generate occupational segregation. Adding discrimination amplifies this segregation: occupations that are less attractive to women because they offer less flexibility also expose women to more discrimination because they have a higher share of men. The same logic applies to men in occupations with higher shares of women.

Second, preferences and discrimination are negatively correlated across occupations. Occupations that women find less attractive because they offer less flexibility are also occupations where women face more discrimination, because they have a higher share of men. Occupations that men find less attractive are also occupations where men face more

[1] This mechanism is similar to that in Schelling (1971) and Sethi and Somanathan (2004), where preferences over the composition of neighborhoods lead to residential segregation. Pan (2015) provides related evidence for occupations: once the share of women in a US occupation crosses a threshold, male employment growth falls sharply and the occupation becomes predominantly female.

discrimination, because they have a higher share of women.[2]

This result matters for empirical work. Several studies ask how much of the gender wage gap or occupational segregation can be explained by preferences for job characteristics (e.g. Cortes and Pan, 2018; Wiswall and Zafar, 2018). These studies compare gaps before and after controlling for preference measures. Our model suggests that this approach can overstate the role of preferences. Because preferences shape occupational composition, and occupational composition shapes discrimination, preference measures partly proxy for discrimination. Controlling for preferences can attribute to preferences some of the segregation caused by discrimination.

The model rests on two assumptions. The first is that women and men differ in how much they value flexibility, which is well documented in previous work (e.g. Wiswall and Zafar, 2018; Mas and Pallais, 2017; Bolotnyy and Emanuel, 2022; Wasserman, 2023). The second assumption is that expected discrimination depends on the sex composition of an occupation. Women face more discrimination as the share of men in an occupation rises. Men face more discrimination as the share of women rises. We use discrimination to mean poor treatment at work, rather than unequal hiring or pay by employers. This includes sexual harassment, violence, bullying, exclusion, and other forms of poor treatment. It also includes workplace cultures that make it unpleasant to be in the minority sex: who gets heard in meetings, what behavior is tolerated, what people joke about, and who feels like they belong.

Existing evidence supports the link between occupational sex composition and sexual harassment. Folke and Rickne (2022) use data from a large Swedish survey on working conditions to show that both women and men report more sexual harassment when they work with more opposite-sex coworkers. One explanation for this result is exposure: perpetrators of sexual harassment are usually of the opposite sex from their victims (Cortina and Berdahl, 2008), so harassment can rise mechanically with the share of opposite-sex coworkers (see also Gutek, Cohen and Konrad, 1990; Folke and Rickne, 2022 on the contact hypothesis). A second explanation is workplace culture: as the share of one sex rises, norms may shift in ways that make poor treatment of the minority sex more likely or more tolerated.[3] Mistreatment also feeds back into composition. Using Finnish administrative data, Adams-Prassl et al. (2024) show that violence by men against female colleagues causes women to leave the firm and fewer women to be hired. This evidence covers both directions of the feedback in our model: composition shapes discrimination,

[2] See Delfino (2024) for experimental evidence on barriers to men's entry into female-dominated occupations.

[3] See Folke and Rickne (2022) for a review of the literature on sexual harassment. There is also more indirect evidence that women's poor treatment increases with the share of men. Lordan and Pischke (2022) use data from the US, UK and Russia to show that women's job satisfaction is negatively correlated with the share of men in an occupation. Zafar (2013) shows that female university students expect to be treated more poorly in jobs associated with majors that have a higher share of male students (see Appendix 3 in Zafar, 2013).

and discrimination reshapes composition.

To see whether poor treatment at work more generally also rises with the share of opposite-sex workers, we ran a short survey on Amazon Mechanical Turk. We asked 295 respondents to guess, for each occupation, the share of female workers and the share of male workers who had been treated poorly at work because of their sex in the last 12 months. Each respondent saw 10 occupations drawn from a list of 60. We then linked these responses to the share of women in each occupation in the 2018 American Community Survey. The results support our main assumption. Female respondents expect that more women are treated poorly in occupations with larger shares of men. Male respondents expect that more men are treated poorly in occupations with larger shares of women (see Appendix A for more details).

We use the model to study three policy levers. First, policies that reduce poor treatment at work weaken the feedback between occupational composition and expected discrimination. These policies reduce segregation by making it less costly for workers to enter occupations where their sex is underrepresented. In the limit, when discrimination is removed, segregation reflects only sex differences in preferences for flexibility.

Second, policies that narrow sex differences in preferences for flexibility also reduce segregation. Examples include childcare subsidies, paternity leave, and policies that shift caregiving norms. In the model, these policies move women's willingness to pay for flexibility closer to men's. Even without discrimination, this reduces segregation. With discrimination, the effect is larger because the amplification mechanism also works in reverse: reducing preference-based segregation reduces the discrimination that would otherwise reinforce it.

Third, we study changes in the wage-flexibility trade-off. Technologies or workplace reforms, such as remote work or more flexible scheduling, can make high-wage occupations more flexible. The effect on segregation is ambiguous. If the workers most attracted by the additional flexibility are women, segregation falls. If they are men, segregation can rise. This ambiguity does not change the main result: for any given wage-flexibility trade-off, discrimination amplifies the segregation generated by preferences.

Our paper contributes to a large literature on how preferences and discrimination affect sex segregation in the labor market. The empirical part of this literature documents sex differences in preferences over job attributes, including risk, workplace flexibility, and work tasks, as well as employer discrimination and sexual harassment (see Cortes and Pan, 2018 for a review). The theoretical link from preferences and discrimination to segregation is straightforward: workers sort into occupations that match their preferences and avoid occupations where they expect discrimination. This logic is central to models of compensating differentials, taste-based discrimination, and statistical discrimination (Becker, 1971; Arrow, 1973; Rosen, 1986). These models are not mainly models of occupational segregation by sex, but they have been used to study it (e.g. Goldin, 2014;

Goldin and Katz, 2016; Bergmann, 1974). Our paper focuses on a specific feedback between the two channels: preference-based sorting changes occupational sex composition, sex composition changes expected discrimination, and expected discrimination changes sorting.

The closest study to ours is Folke and Rickne (2022). They document that sexual harassment rises with the share of opposite-sex coworkers. They use a two-workplace model, adapted from the standard theory of compensating differentials, to organize their empirical analysis. In the model, harassment amplifies segregation initially generated by sex differences in preferences for workplace amenities.[4]

We build on their analysis by studying this feedback in a labor market with many occupations and tracing its consequences for aggregate segregation, empirical attempts to distinguish preferences from discrimination, and policies to reduce segregation. With many occupations, the aggregate effect is not obvious: the feedback can make some occupations more balanced and others less so. We show that it nevertheless raises aggregate occupational segregation, measured by the Duncan index. The feedback also complicates empirical attempts to separate preferences from discrimination. Because preferences shape occupational composition and composition shapes discrimination, estimates of the contribution of preferences may partly capture segregation caused by discrimination. Finally, we derive comparative statics showing how occupational segregation responds to three policies intended to reduce it: policies that reduce workplace discrimination, policies that narrow sex differences in preferences for flexibility, and policies that make high-wage occupations more flexible.

2 Model

2.1 Preferences for occupations

In our model, people differ in how they value job characteristics and in the discrimination they expect to face. For simplicity, we assume two sexes and equal numbers of women and men in the labor force. Occupations differ in wages and flexibility. Women and men value wages equally, but women have a higher willingness to pay for flexibility. These labels are concrete examples of a more general distinction: some job characteristics are valued similarly by women and men, while others are valued differently.

[4] Two other papers study related amplification mechanisms. In a two-sector Roy model, Chen et al. (2025) show that preferences over group composition amplify segregation generated by group differences in earnings opportunities. Their model treats the composition effect in reduced form and does not include a separate role for sex differences in preferences over job attributes. Using simulations of a search model, Usui (2015) finds that women's aversion to male-dominated jobs reinforces segregation, both directly and by leading firms to tailor their salary-hours offers toward men's working-hours preferences.

The utility derived by person i of sex j in occupation k is given by

$$U_{ijk} = \text{wage}_k + \beta_i \text{flex}_k - \gamma \text{disc}_{jk}, \tag{1}$$

where wage_k denotes the wage in occupation k, flex_k shows the job flexibility in occupation k, disc_{jk} shows how much discrimination people of sex j experience in occupation k, β_i shows how much person i values flexibility, γ shows the dislike for discrimination. We normalize the coefficient of wage_k to 1. We assume that everyone is either indifferent to flexibility or prefers more flexibility ($\beta_i \geq 0$) and that everyone dislikes discrimination equally ($\gamma > 0$). This setup allows us to interpret β_i as willingness to pay (WTP) for flexibility, which is greater than or equal to zero for everyone. Similarly, we interpret γ as WTP to avoid discrimination, which is positive and the same for everyone.[5]

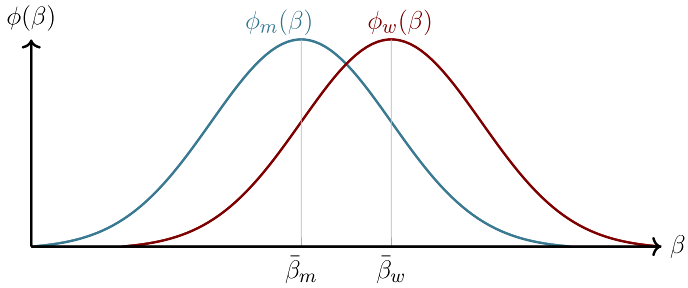

(a) PDFs of WTP

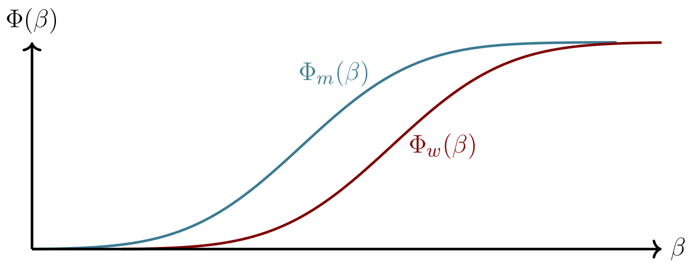

(b) CDFs of WTP

Figure 1: WTP for flexibility for men and women

To capture that women care more about flexibility than men, we assume that the cumulative distribution function (CDF) of women's WTP for flexibility, $\Phi_w(\beta)$, first-order stochastically dominates men's CDF, $\Phi_m(\beta)$.

Assumption 1. $\Phi_w(\beta) \leq \Phi_m(\beta)$ *for all* β*, with strict inequality for some* β*.*

Figure 1 shows distributions of women's and men's WTP for flexibility that are consistent with this assumption. Figure 1a shows that the probability density function (PDF) of women's WTP for flexibility, $\phi_w(\beta)$, is shifted to the right relative to men's PDF

[5]We assume γ to be the same for everyone to keep the algebra simple. Making it vary over the population will not affect our results as long as its distribution is independent of that of β_i.

of WTP for flexibility, $\phi_m(\beta)$. The corresponding CDF in Figure 1b shows that women's WTP for flexibility first-order stochastically dominates men's.

The key assumption in our model is that discrimination depends on occupational sex imbalance. Rather than designating particular occupations as hostile to one sex, we let discrimination arise endogenously from the distribution of workers across occupations. We normalize the mass of employed men to one and the mass of employed women to one. Let m_k denote the share of employed men who work in occupation k, and let w_k denote the corresponding share of employed women. We summarize occupational sex imbalance by

$$x_k = m_k - w_k.$$

The sign of x_k indicates which sex is overrepresented in occupation k: men when $x_k > 0$, women when $x_k < 0$.[6] Discrimination against women (disc_{wk}) is equal to $f(x_k)$ and discrimination against men (disc_{mk}) is equal to $f(-x_k)$. We are agnostic about the exact functional form of $f(\cdot)$, but assume that it is increasing through zero. Because $f(0) = 0$ and $f' > 0$, disc_{jk} is positive when sex j is underrepresented in occupation k (reducing utility through the $-\gamma\text{disc}_{jk}$ term) and negative when overrepresented (raising utility). The underrepresented sex therefore faces more discrimination in occupations that are more imbalanced against it.

Assumption 2. $\text{disc}_{wk} = f(x_k)$, $\text{disc}_{mk} = f(-x_k)$, *where* $x_k = m_k - w_k$, $f(0) = 0$, *and* $f'(.) > 0$.

Together, Assumptions 1 and 2 treat preferences and discrimination as analytically independent primitives, and we show that even under this orthogonality, the two channels interact through occupational composition to amplify segregation.[7]

We assume that there are only two occupations, which allows us to characterize the trade-off between wages and flexibility. Occupation 1 has a higher wage and is less flexible. Occupation 2 has a lower wage and is more flexible. Figure 2 shows both occupations on a wage-flexibility plane. The line connecting them (the wage-flexibility frontier) slopes downward; we define the absolute value of its slope as the wage-flexibility trade-off:

$$\mu = \frac{\text{wage}_1 - \text{wage}_2}{\text{flex}_2 - \text{flex}_1} \tag{2}$$

[6] The object $x_k = m_k - w_k$ is the contribution of occupation k to the Duncan index. It is not identical to the within-occupation male share, but it has the same sign as the occupation's deviation from sex parity. If π_k denotes the male share within occupation k, then, given our normalization of the employed male and female workforces, $\pi_k = m_k/(m_k + w_k)$ and $\pi_k - \frac{1}{2} = x_k/[2(m_k + w_k)]$. Thus $x_k > 0$ if and only if the occupation is male-overrepresented. We use x_k rather than a ratio measure because it aligns directly with the Duncan index and, in the two-occupation case, implies $x_2 = -x_1$, yielding a one-dimensional law of motion.

[7] This is a modeling choice, not an empirical claim: in practice, the two forces may shape each other. Endogenizing one with respect to the other would only add feedback and strengthen the interaction we study.

Unlike a worker's willingness to pay for flexibility β_i, the trade-off μ is a feature of the job environment. It reflects technological, organizational, and institutional constraints and is the same for all workers, regardless of sex or preferences.

Figure 2: Two occupations

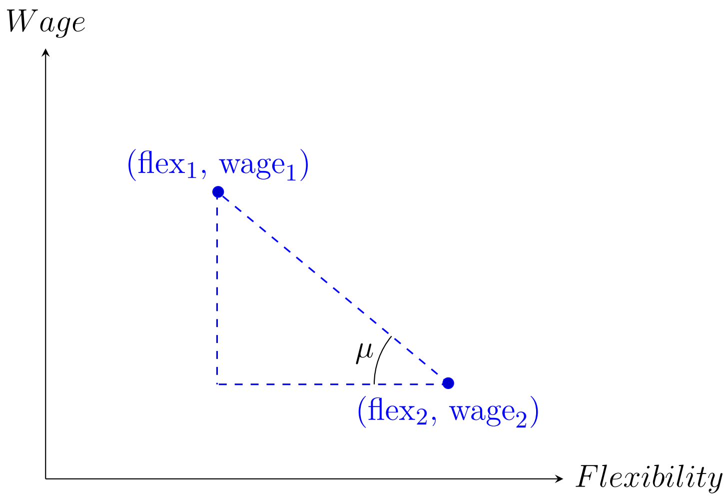

Women prefer occupation 1 over occupation 2 iff

$$U_{iw1} > U_{iw2} \Rightarrow \text{wage}_1 + \beta_i \text{flex}_1 - \gamma \text{disc}_{w1} > \text{wage}_2 + \beta_i \text{flex}_2 - \gamma \text{disc}_{w2} \tag{3}$$

$$\beta_i < \mu \left(1 - \gamma \frac{\text{disc}_{w1} - \text{disc}_{w2}}{\text{wage}_1 - \text{wage}_2}\right). \tag{4}$$

Analogously, men prefer occupation 1 over occupation 2 iff

$$\beta_i < \mu \left(1 - \gamma \frac{\text{disc}_{m1} - \text{disc}_{m2}}{\text{wage}_1 - \text{wage}_2}\right). \tag{5}$$

Equations 4 and 5 show that discrimination shifts the cutoff for preferring occupation 1 over occupation 2. This shift differs for women and men because discrimination depends on occupational sex imbalance. With two occupations, the imbalance in occupation 2 is the mirror image of the imbalance in occupation 1: $x_2 = -x_1$. This follows from $m_2 = 1 - m_1$ and $w_2 = 1 - w_1$.

To simplify notation, define women's normalized discrimination gap as

$$\delta \equiv \frac{\text{disc}_{w1} - \text{disc}_{w2}}{\text{wage}_1 - \text{wage}_2} = \frac{f(x_1) - f(-x_1)}{\text{wage}_1 - \text{wage}_2}. \tag{6}$$

The parameter δ measures women's discrimination gap between the two occupations in units of the wage gap. If $\delta > 0$, women face more discrimination in occupation 1 than in occupation 2. Because the sex imbalance in occupation 2 is the mirror image of the imbalance in occupation 1, the corresponding normalized discrimination gap for men is $-\delta$.

Therefore, a woman prefers the higher-paying occupation 1 over the more flexible occupation 2 iff

$$\beta_i < \mu(1 - \gamma\delta). \tag{7}$$

A man prefers occupation 1 over occupation 2 iff

$$\beta_i < \mu(1 + \gamma\delta). \tag{8}$$

Equations 7 and 8 show cut-off values for WTP for flexibility for women and men. People whose WTP for flexibility is below their sex-specific cut-off prefer the higher-paying occupation 1. Everyone else prefers the more flexible occupation 2. Overall, the preference for an occupation depends on a person's WTP for flexibility (β_i), the trade-off between wage and flexibility between the two occupations (μ), the WTP to avoid discrimination (γ), and the sex-specific trade-off between discrimination and wage (δ or $-\delta$).

Preview: model intuition

Below, we model the matching of people to occupations as a dynamic process in which people's expectations about discrimination, and their resulting occupational choices, are driven by the occupational segregation in the previous cohort. In this section, we preview the key insight from this dynamic model. We show how preferences for flexibility, together with discrimination, cause occupational segregation.

Figure 3 shows the PDFs of women's and men's WTP for flexibility for two scenarios. The top part of Figure 3 shows a world without discrimination ($\gamma\delta = 0$). In this scenario, the cutoff value for preferring occupation 2 is the same for women and men and solely driven by the trade-off between wages and flexibility of the two occupations shown (μ). With the same cutoff, more women prefer the more flexible occupation 2 because women's WTP for flexibility first-order stochastically dominates men's. Sex differences in preferences for occupations are entirely driven by sex differences in WTP for flexibility. We show below that these sex differences in preferences straightforwardly translate into occupational segregation in our dynamic model.

The bottom part of Figure 3 shows a scenario with discrimination, in the intuitive case where the previous cohort had more men than women in higher-paying occupation 1 ($x_1 > 0$), in line with sex differences in preferences. Because women now expect more discrimination in occupation 1, their cut-off for preferring occupation 2 shifts left to

Figure 3: PDFs of WTP without and with discrimination

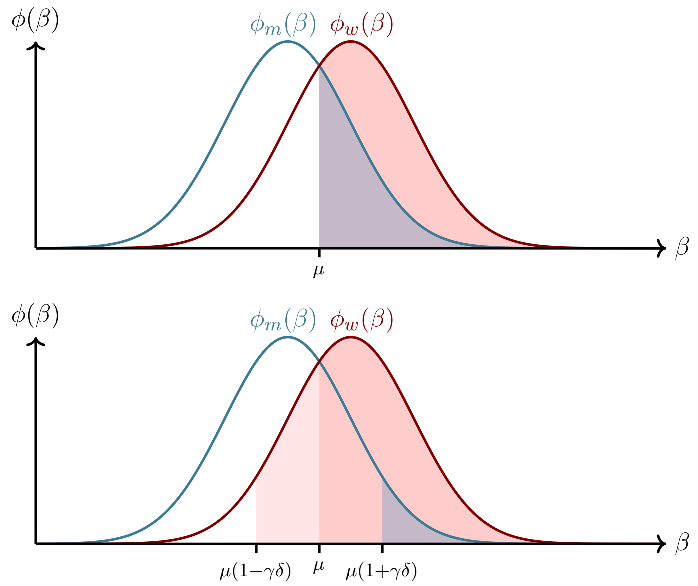

Note: Shaded areas in red and gray show the proportion of women and men, respectively, who prefer the more flexible occupation 2 over the higher-paying occupation 1.

$\mu(1-\gamma\delta)$, and more women prefer occupation 2 than in the no-discrimination case. At the same time, the cut-off shifts right for men to $\mu(1+\gamma\delta)$. Because there were more women than men in occupation 2 in the previous cohort, men expect more discrimination there, which makes occupation 2 less attractive for them. We show below that this result also holds in the dynamic model. Discrimination magnifies the occupational segregation caused by sex differences in preferences.[8]

2.2 Dynamic model

In this section, we model how sex differences in preferences and discrimination lead to occupational segregation as a dynamic process.

2.2.1 From preferred to realized occupations

Each person lives two periods. In the first period, they observe the wages and flexibility of the two occupations, along with the discrimination faced by the previous cohort of their sex. They expect the same level of discrimination, which determines which occupation they prefer. In the second period, they are allocated across the two occupations.

[8]Note that our model also predicts a gender wage gap, but the direction depends on an arbitrary modeling choice. In this version, the key trade-off is between wages and a characteristic women value more than men (flexibility), so women earn less on average. Had we instead modeled the trade-off as between wages and a characteristic men value more than women (e.g. working with things), the wage gap would run in the opposite direction. Even in that version, discrimination would still magnify the segregation caused by sex differences in preferences.

Let $s_{1w,t}$ denote the share of women who prefer the higher-paying occupation 1 at time t, and let $s_{1m,t}$ denote the corresponding share of men. From the cutoff rules above,

$$s_{1w,t} = \Phi_w[\mu(1-\gamma\delta(x_{1,t}))], \tag{9}$$

$$s_{1m,t} = \Phi_m[\mu(1+\gamma\delta(x_{1,t}))], \tag{10}$$

where

$$\delta(x_{1,t}) = \frac{f(x_{1,t}) - f(-x_{1,t})}{\text{wage}_1 - \text{wage}_2}.$$

The objects $s_{1w,t}$ and $s_{1m,t}$ are preference shares, not realized employment shares.

To allow for the fact that workers may not always enter their preferred occupation, we write the allocation as a convex combination of preferred and fallback allocations. With probability $\lambda \in (0,1)$, a worker enters their preferred occupation; with probability $1-\lambda$, they are assigned to a preference-neutral fallback allocation that places them in occupation 1 with probability q and occupation 2 with probability $1-q$. Hence,

$$m_{1,t+1} = \lambda s_{1m,t} + (1-\lambda)q, \tag{11}$$

$$w_{1,t+1} = \lambda s_{1w,t} + (1-\lambda)q. \tag{12}$$

The corresponding shares in occupation 2 are $m_{2,t+1} = 1 - m_{1,t+1}$ and $w_{2,t+1} = 1 - w_{1,t+1}$. Thus all workers are allocated to one of the two occupations, and the preference-neutral fallback allocation does not itself create sex-based segregation.

The fallback weight $1-\lambda$ is a catch-all for capacity constraints, skill requirements, geography, vacancies, search frictions, networks, and other sex-neutral allocation frictions. These frictions do not generate unemployment, although the model can easily accommodate it. The frictionless case $\lambda = 1$, where everyone enters their preferred occupation, is susceptible to corner solutions with all men in one occupation and all women in the other. The other extreme $\lambda = 0$ shuts down the preference channel entirely, so there is no segregation from preferences. We consider the non-trivial case $0 < \lambda < 1$.

2.2.2 Segregation in equilibrium

We measure segregation with the Duncan index (Duncan and Duncan, 1955):

$$D_t = \frac{1}{2}\sum_{k=1}^{2} |m_{k,t} - w_{k,t}| = |x_{1,t}| = |x_{2,t}|, \tag{13}$$

where $x_{k,t} = m_{k,t} - w_{k,t}$. In the two-occupation case, the Duncan index is equal to the absolute value of the difference between men's and women's employment shares in occupation 1. This allows us to use $x_{1,t}$ as a measure of occupational segregation. Taking the difference between men's and women's realized employment shares in occupation 1 gives

$$
\begin{aligned}
x_{1,t+1} &= m_{1,t+1} - w_{1,t+1} \\
&= \lambda(s_{1m,t} - s_{1w,t}) \\
&= \lambda\left[\Phi_m(\mu(1+\gamma\delta(x_{1,t}))) - \Phi_w(\mu(1-\gamma\delta(x_{1,t})))\right]. \qquad (14)
\end{aligned}
$$

The preference-neutral fallback allocation q drops out because it applies equally to men and women. A steady state equilibrium exists whenever the difference between men's and women's employment shares in occupation 1 does not change over time, i.e. when $x_{1,t+1} = x_{1,t}$. We characterize the steady state under different scenarios below.

Proposition 1. *There is a positive level of segregation with sex differences in preferences but without discrimination.*

This result follows directly from Equation 14. Without discrimination, either because $f(x_k) = 0$ for all x_k or because $\gamma = 0$, the law of motion does not depend on occupational composition in the previous period. The level of segregation in occupation 1 is

$$\bar{x}_1 = \lambda\left[\Phi_m(\mu) - \Phi_w(\mu)\right] > 0. \qquad (15)$$

Because women's willingness to pay for flexibility first-order stochastically dominates men's, $\Phi_w(\mu) < \Phi_m(\mu)$ at the relevant cutoff. Hence occupation 1 has more men than women, while occupation 2 has more women than men. The level of segregation is proportional to the difference in preferences, $\Phi_m(\mu) - \Phi_w(\mu)$, and to λ, the extent to which preferred occupational choices are reflected in realized employment.[9]

Proposition 2 characterizes an equilibrium in the more plausible scenario with differences in preferences *and* discrimination.

Proposition 2. *Compared to the preferences-only equilibrium, adding discrimination increases the level of segregation in the steady state.*

Proof: Let x_1^* denote a steady state equilibrium with differences in preferences and discrimination. Define the dynamic mapping

$$F(x) = \lambda\left\{\Phi_m\left(\mu\left(1+\gamma\frac{f(x)-f(-x)}{\text{wage}_1-\text{wage}_2}\right)\right) - \Phi_w\left(\mu\left(1-\gamma\frac{f(x)-f(-x)}{\text{wage}_1-\text{wage}_2}\right)\right)\right\}. \qquad (16)$$

A steady state satisfies $x_1^* = F(x_1^*)$.

At $x = 0$, we have $f(0) = 0$ and therefore

$$F(0) = \lambda[\Phi_m(\mu) - \Phi_w(\mu)] = \bar{x}_1 > 0.$$

[9] A trivial result is that there would be no segregation ($D = 0$) if there is neither discrimination nor a difference in preferences, i.e. $\Phi_m(\cdot) = \Phi_w(\cdot)$.

Moreover, since the difference between two CDFs is bounded above by one and $\lambda < 1$, we have $F(x) \leq \lambda < 1$ for all $x \in [0,1]$. Hence $F(1) - 1 < 0$. Since F is continuous, the intermediate value theorem implies that at least one positive fixed point exists in $(0,1)$.

Now consider any positive x. Since $f'(\cdot) > 0$, $f(x) > f(-x)$, and therefore discrimination raises men's cutoff for occupation 1 and lowers women's cutoff for occupation 1. Hence

$$F(x) > \lambda[\Phi_m(\mu) - \Phi_w(\mu)] = \bar{x}_1.$$

Since any positive steady state satisfies $x_1^* = F(x_1^*)$, it follows that

$$x_1^* > \bar{x}_1.$$

Thus, discrimination amplifies the level of segregation generated by sex differences in preferences.[10]

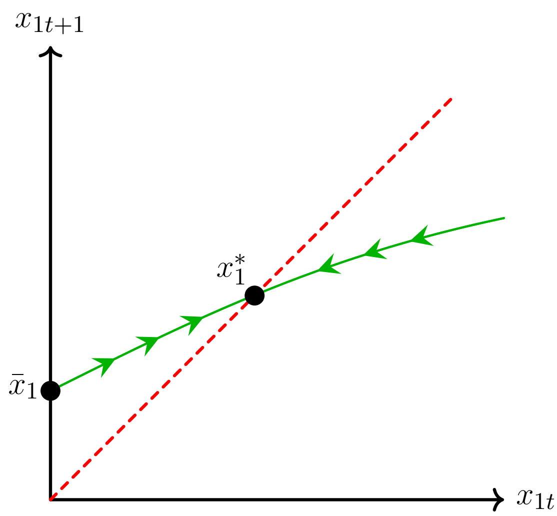

Figure 4: Steady State Equilibrium

Introducing discrimination causes future segregation to depend on current segregation. Given an initial occupational imbalance, discrimination shifts the sex-specific cutoffs, changing the shares of men and women who prefer occupation 1 in the next cohort. Figure 4 illustrates a locally stable positive steady state, denoted by x_1^*, where occupational composition no longer changes over time. At this equilibrium, more men work in the higher-paying occupation 1 and more women work in the more flexible occupation 2. In other words, discrimination exacerbates the segregation that already arises from differences in preferences.[11]

[10]It is theoretically possible to have a steady state equilibrium where x_1 is negative, i.e. where there are more women in the higher-paying occupation 1 and more men in the more-flexible occupation 2. We show in Appendix C that such an equilibrium will a) not *always* exist and b) even when it exists, it will have a smaller basin of attraction relative to the stable equilibrium where x_1 is positive.

[11]We discuss the scenario with discrimination but without differences in preferences in Appendix B.

The key reason discrimination generates dynamics is that it creates a feedback mechanism. When an occupation becomes even slightly imbalanced, workers in the minority sex anticipate worse treatment and become less willing to enter, causing the imbalance to widen over time. In terms of the dynamic mapping $x_{t+1} = F(x_t)$, discrimination makes this relationship steeper, so small imbalances today translate into larger imbalances tomorrow. By contrast, when discrimination is absent, the mapping is flat. Future composition does not depend on today's composition, and no amplification occurs.

Propositions 1 and 2 show how sex-based segregation depends on sex differences in preferences and on discrimination in a two-occupation labor market. Proposition 3 extends the result to a labor market with n occupations.

Proposition 3. *In a general n-occupation model, sex differences in preferences continue to generate a positive level of segregation, and discrimination continues to amplify segregation relative to the preferences-only benchmark.*

The proof is in Appendix D; here we describe the argument, beginning with why the two-occupation logic does not extend directly. With two occupations, discrimination acts on a single pair: it makes the male-dominated occupation more attractive to men and less attractive to women, driving that pair further apart, and this is the entire mechanism. With n occupations each occupation is compared with all $n-1$ others at once, and a property that holds for every pair need not hold for every occupation.

To see why, fix an occupation k. Its comparisons with the others pull it in opposite directions: comparisons with occupations more male-dominated than k pull it toward female, and comparisons with occupations more female-dominated than k pull it toward male. Discrimination strengthens every one of these comparisons, but it does so on both sides of k at once, so the direction in which k ultimately moves depends on which side dominates. An occupation that begins only mildly female-overrepresented may be driven further toward female by the strongly male occupations and back toward male by the strongly female ones; if the latter dominate, k becomes *less* female-overrepresented than before, even though discrimination has reinforced each of its pairwise comparisons.

The result therefore cannot be established occupation by occupation, since an individual occupation may move against its own initial imbalance. It instead requires an argument about the entire vector of imbalances at once, which shows that discrimination spreads the imbalances further apart and so raises the Duncan index. We develop the two ingredients of that argument in turn: a pairwise decomposition of preferences-only segregation, and the feedback through which discrimination acts on it. The n-occupation model retains

The key insight from this discussion is that discrimination alone does not necessarily lead to a segregated labor market, and the stable equilibrium is path dependent. Depending on the initial conditions, we can have a perfectly integrated labor market, a segregated labor market with more men in the higher-paying occupation and more women in the more flexible occupation, or a segregated labor market with more women in the higher-paying occupation and more men in the more flexible occupation.

the economic structure of the two-occupation model, with the wage–flexibility trade-off now generating a richer set of occupational rankings. We order occupations along a wage-flexibility frontier: occupation 1 is the highest-paid and least flexible, while occupation n is the lowest-paid and most flexible:

$$\text{wage}_1 > \text{wage}_2 > \cdots > \text{wage}_n, \qquad \text{flex}_1 < \text{flex}_2 < \cdots < \text{flex}_n.$$

Thus, no occupation dominates another in terms of both wage and flexibility. Workers with a low willingness to pay for flexibility rank high-wage, low-flexibility occupations more highly, while workers with a high willingness to pay for flexibility rank flexible, lower-wage occupations more highly. These n occupations generate at most $\binom{n}{2}$ unique pairwise slope thresholds, which correspond to at most $\binom{n}{2}+1$ unique preference rankings.

Occupation k's imbalance is $x_k = m_k - w_k$, as in the two-occupation model, and because each workforce has unit mass the imbalances sum to zero, $\sum_{k=1}^{n} x_k = 0$. Imbalance is therefore relative: no occupation can become more male-overrepresented without another becoming more female-overrepresented, so a rise in overall segregation means the distribution of men and women becomes more uneven, not that every occupation becomes more imbalanced. Segregation is measured by the Duncan index $D = \frac{1}{2}\sum_k |x_k|$. To track how preferences and discrimination shape these imbalances, write a worker's valuation of occupation k as $u_k(\beta) = a_k + \beta\,\text{flex}_k$, where β is her willingness to pay for flexibility and a_k is the occupation's intercept. This is the utility of Section 2 with discrimination folded into the intercept, so that $a_k = \text{wage}_k$ in the preferences-only case and discrimination enters by shifting a_k differently for men and women. A worker ranks occupations by u_k and is assigned to occupations according to the probability schedule below, a map we write R_Φ from the intercepts to realized employment shares.

As a worker's willingness to pay for flexibility rises, her ranking changes one adjacent pair at a time. Each pair of occupations $i < j$ has a single cutoff

$$\eta_{ij} = \frac{\text{wage}_i - \text{wage}_j}{\text{flex}_j - \text{flex}_i} > 0,$$

at which she switches between them: below the cutoff she prefers the higher-paying occupation i, above it the more flexible occupation j. These at most $\binom{n}{2}$ cutoffs carry the wage ranking, held at low willingness to pay, into the flexibility ranking, held at high willingness to pay. Figure 5 shows the construction: each occupation's score $a_k + \beta\,\text{flex}_k$ is a straight line in β, every pair of lines crosses exactly once at its cutoff η_{ij}, and the crossings, which need not occur in numerical order, carry the wage ranking into the flexibility ranking through a sequence of adjacent swaps.

To map rankings into realized employment, we use the same idea as in the two-occupation model: preferences affect realized employment, but not necessarily one-for-one.

Figure 5: Occupation scores and pairwise cutoffs

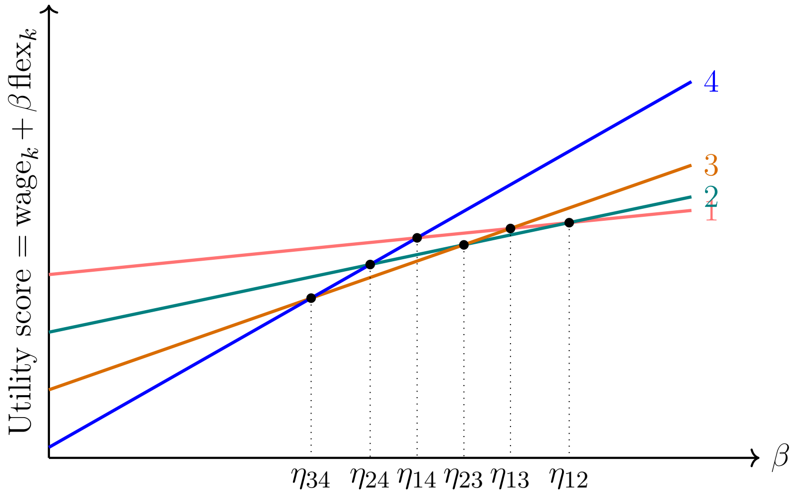

Note: Each line is an occupation's utility score wage$_k + \beta\,$flex$_k$; steeper lines represent more flexible occupations. The six dots mark the pairwise cutoffs η_{ij} where two occupations swap adjacent ranks.

Let λ_r denote the probability that a worker is assigned to her r-th ranked occupation, with

$$\lambda_r \geq 0, \qquad \sum_{r=1}^{n} \lambda_r = \Lambda < 1, \qquad \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n,$$

and at least one strict inequality, so that more preferred occupations are more likely to be realized. In the two-occupation model this intensity was a single scalar λ, the probability of entering one's preferred occupation; here it becomes a full schedule $(\lambda_1, \ldots, \lambda_n)$, so preferences shape realized employment throughout the ranking rather than only at the top. The remaining probability $1 - \Lambda$ is a preference-neutral fallback allocation $\mathbf{q} = (q_1, \ldots, q_n)$, with $q_i \geq 0$ and $\sum_{i=1}^{n} q_i = 1$, that is common to men and women and therefore drops out when we take the difference between their employment shares.

These pairwise crossings let us decompose preferences-only segregation cutoff by cutoff. Each crossing moves only the two occupations involved.

Lemma 1 (Adjacency at cutoffs)**.** *Assume the pairwise cutoffs are distinct, which holds for generic wages and flexibilities. At each cutoff η_{ij}, occupations i and j occupy adjacent ranks and exchange them, and the relative order of every other pair is unchanged.*

At a cutoff η_{ij}, let $\epsilon_{ij} = \Phi_m(\eta_{ij}) - \Phi_w(\eta_{ij}) \geq 0$ be the excess of men over women who cross it, nonnegative by stochastic dominance (Assumption 1), and let $\Delta\lambda_{ij} = \lambda_r - \lambda_{r+1} \geq 0$ be the gain in realization probability from moving up the adjacent rank the pair occupies. The nonnegative pairwise contribution

$$\sigma_{ij}^0 = \epsilon_{ij}\,\Delta\lambda_{ij} \geq 0$$

is the segregation this one crossing adds, and summing an occupation's crossings recovers its preferences-only imbalance.

Lemma 2 (Pairwise decomposition)**.** *For every occupation k,*

$$x_k^0 = \sum_{j>k} \sigma_{kj}^0 - \sum_{i<k} \sigma_{ik}^0. \tag{17}$$

Because every $\sigma_{ij}^0 \geq 0$, the cumulative imbalance in the top m occupations, $X_m = \sum_{k \leq m} x_k^0 = \sum_{i \leq m < j} \sigma_{ij}^0$, is nonnegative at every threshold: the highest-paying occupations are jointly male-overrepresented. As long as one crossing contributes ($\sigma_{ij}^0 > 0$ for some pair), the preferences-only Duncan index is strictly positive. This first result is established formally in Appendix D, and needs no assumption that the preference gap between men and women is constant across cutoffs.

Discrimination, in turn, amplifies this segregation. It enters by shifting each occupation's intercept, by sex: relative to the no-discrimination wages, men in occupation k perceive the intercept $a_k^m = \text{wage}_k - \gamma f(-x_k)$ and women $a_k^w = \text{wage}_k - \gamma f(x_k)$, where f is increasing with $f(0) = 0$ and $\gamma > 0$ is the disutility from discrimination. A discriminatory steady state is an imbalance vector $\boldsymbol{x}^*$ that reproduces itself once these shifted intercepts pass through the realized-employment map; such a state exists. Because each occupation's shift grows with its own imbalance, discrimination drives every pair apart: if occupation i is more male-overrepresented than j, it becomes still more attractive to men and less to women. This is the n-occupation counterpart of the two-occupation cutoff shifts $\mu(1 \pm \gamma\delta)$.

To turn these shifts into a change in segregation we need to know how realized employment responds when an intercept moves. For a sex s and a pair (k, ℓ), shifting the intercept gap slides the indifference cutoff between them, and the realized employment of sex s that crosses the margin per unit of gap is

$$c_{k\ell}^s = \underbrace{\phi_s(\eta_{k\ell})}_{\text{workers at the cutoff}} \times \underbrace{\Delta\lambda_{k\ell}}_{\text{employment per switch}} \times \underbrace{\frac{1}{|\text{flex}_\ell - \text{flex}_k|}}_{\text{cutoff movement per unit gap}} \geq 0:$$

the density of workers at the cutoff, times the realization gain from rising one rank, times how fast the cutoff moves. Writing h_k^s for the intercept shift of sex s in occupation k, these coefficients assemble the marginal response of realized employment.

Lemma 3 (Marginal effect of a small change in intercepts)**.** *For a sex with intercept shifts h_k^s, and at intercepts where the pairwise cutoffs are distinct, the change in its realized employment in occupation k is $\sum_{\ell \neq k} c_{k\ell}^s (h_k^s - h_\ell^s)$.*

Accumulating this marginal effect as discrimination is dialed on, from the wage intercepts to the steady-state intercepts, and differencing men and women relates the two benchmarks by

$$\boldsymbol{x}^0 = (I - L^*)\, \boldsymbol{x}^*,$$

where L^* is a symmetric operator, built from nonnegative pairwise weights, that averages each occupation's imbalance toward the others'. Its construction and the path argument behind this identity are in Appendix D; what matters here is that $I - L^*$ is an averaging map.

The amplification result follows. Under a stability condition (that the feedback does not overshoot, the n-occupation counterpart of $0 < F'(x^*) < 1$, namely that the largest eigenvalue of L^* is below one), the map $I - L^*$ is a genuine average, so $\boldsymbol{x}^0$ is a smoothed, less spread-out version of $\boldsymbol{x}^*$. Because the Duncan index is half the total spread of the imbalances, the discriminatory steady state is at least as segregated as the preferences-only benchmark,

$$D^* \geq D^0,$$

and strictly more so whenever the feedback links a male-overrepresented occupation to a female-overrepresented one with positive weight. The comparison holds even when discrimination is strong enough to change the order in which workers rank occupations. Proofs are in Appendix D.

A second feature of the model concerns where discrimination falls. Discrimination and preferences are negatively related across occupations: each sex faces the most discrimination in the occupations its members are least inclined to enter. The logic is direct. In any occupation k, women face discrimination $f(x_k)$ and men $f(-x_k)$, so the difference $f(x_k) - f(-x_k)$ carries the sign of x_k, and the underrepresented sex always bears the greater discrimination. Underrepresentation, in turn, reflects each sex sorting away from the characteristics it values less (women from the less flexible, higher-paying occupations and men from the more flexible ones), so the occupations a sex finds least attractive are exactly those where it faces the most discrimination. In the two-occupation model this is immediate. With many occupations it is no longer obvious, because an occupation's imbalance need not vary monotonically with its flexibility, so we verify it by simulating the n-occupation dynamics (construction in Appendix D). Figure 6 plots, for each occupation at the steady state, its flexibility against each sex's under-representation, which under Assumption 2 determines the discrimination it faces, and confirms the negative relationship.

The pattern in Figure 6 is intuitive. Women generally prefer more flexible occupations and face more discrimination in less flexible ones; men generally prefer less flexible, higher-paying occupations and face more discrimination in more flexible ones. Because discrimination is driven by occupational sex imbalance, the women and men who end up in "atypical" occupations are exactly those who face the most discrimination.

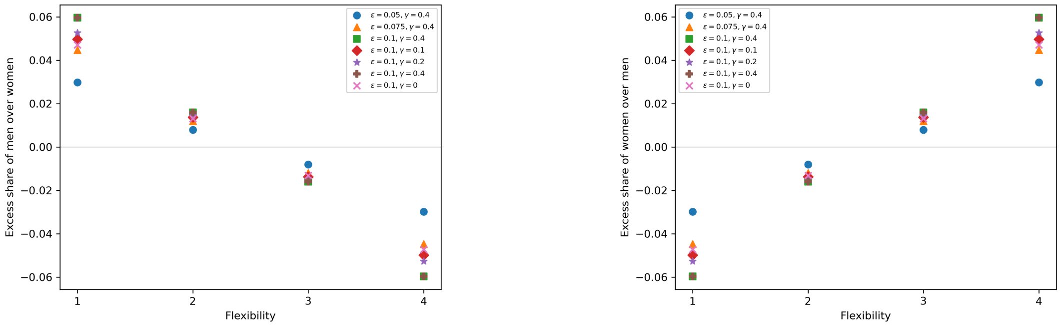

(a) Women

(b) Men

Figure 6: Relationship between occupation-level flexibility and discrimination for women and men

Note: The figure shows the relationship between occupational flexibility and the discrimination faced by women (men) at the steady state of the n-occupation simulation. The vertical axis plots the excess share of men (women) over women (men) in an occupation; under Assumption 2, discrimination is increasing in this quantity. The horizontal axis orders occupations in increasing order of flexibility. The code for these simulations is available in an OSF repository at https://osf.io/g2umr/.

3 Policy implications

Our model shows how different policy interventions affect occupational segregation. Broadly, three categories of policies correspond to three levers in the model: policies that reduce workplace discrimination, policies that reduce sex differences in preferences for flexibility, and policies or technologies that change the wage-flexibility trade-off. We examine the policies in turn, highlighting how they alter the steady-state equilibrium.

3.1 Policies that reduce workplace discrimination

Policies aimed at reducing on-the-job discrimination, such as stricter sexual harassment policies, bias training, or enforcement of equal treatment, can be interpreted in the model as a reduction in the slope of the discrimination function $f(x)$. Intuitively, if workplace culture improves such that being in the minority sex is less likely to invite hostility or exclusion, then a given imbalance x in the sex-composition generates a smaller disutility from discrimination. Exposure itself can reduce discrimination: randomly assigning women to traditionally male teams improved men's attitudes toward female colleagues (Dahl, Kotsadam and Rooth, 2021). A flatter $f(\cdot)$ means that $f'(x)$ is lower at all x. Graphically, this change pivots the best-response mapping $x_{t+1} = F(x_t)$ downwards (see Figure 7).

Reducing workplace discrimination lowers steady-state sex segregation. In the model, weaker discrimination leads to a lower equilibrium level of segregation, denoted by $\hat{x}_1$. This occurs because discrimination amplifies sex differences in preferences; when discrimination weakens, women and men sort less strongly into different occupations. The no-discrimination benchmark, $\bar{x}_1$, remains unchanged because it reflects segregation caused by preferences alone. Policies that reduce discrimination do not change workers' preferences; they reduce the costs of entering occupations where one's sex is underrepresented. As a result, the steady-state level of segregation with lower discrimination, $\hat{x}_1$, moves closer to the preferences-only benchmark, $\bar{x}_1$.

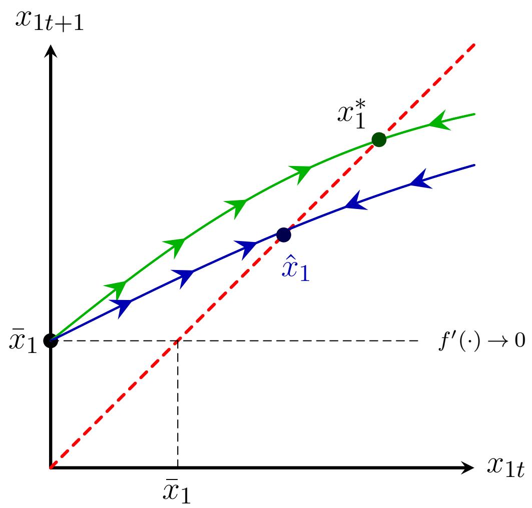

Figure 7: Reducing discrimination and steady-state segregation

In a labor market completely free from discrimination (where $f'(x) \to 0$), the equilibrium aligns with outcomes based solely on preferences. Any partial reduction in discrimination results in equilibrium segregation falling between the original x_1^* and $\bar{x}_1$. Therefore, discrimination-based policies unambiguously reduce occupational segregation in our model.

These theoretical results closely align with the empirical findings reported in the literature. For example, Folke and Rickne (2022) show that the incidence of sexual harassment for both women and men rises when they work in occupations where their sex is in the minority. It follows that mitigating such behavior through stronger anti-harassment policies will encourage more workers from the minority sex to enter that occupation.

3.2 Policies that narrow sex differences in preferences

While women tend to value flexibility more than men, this difference does not only reflect fixed or innate preferences. It also reflects constraints outside the workplace, such as childcare responsibilities (e.g. Kleven, Landais and Søgaard, 2019; Adda, Dustmann and Stevens, 2017), caregiving norms (Boelmann, Raute and Schönberg, 2025), and the division of household labor (e.g. Andrew et al., 2022). Policies that relax these constraints, such as childcare subsidies, paid parental leave for fathers, or efforts to shift caregiving norms, can therefore reduce sex differences in preferences for flexibility (see Olivetti and Petrongolo, 2017, for evidence on family policies).[12]

In our model, we represent policies that reduce sex differences in preferences for flexibility as shifts in the distribution of women's willingness to pay for flexibility. A policy that narrows the preference gap shifts Φ_w toward Φ_m, so that the average woman's willingness to pay for flexibility becomes more like the average man's.[13]

Let $\tau \in [0,1]$ index the strength of this policy, where higher values of τ correspond to a smaller sex difference in preferences for flexibility. We represent this shift by moving women's CDF toward men's CDF pointwise:

$$\Phi_w^\tau(\beta) \;\equiv\; (1-\tau)\,\Phi_w(\beta) \;+\; \tau\,\Phi_m(\beta). \tag{18}$$

When $\tau = 0$, women's willingness to pay is unchanged; when $\tau = 1$, women's CDF equals men's CDF. In the baseline without discrimination, the new steady-state level of segregation in occupation 1 is

$$\bar{x}_1(\tau) = \lambda\Big[\Phi_m(\mu) - \Phi_w^\tau(\mu)\Big] = (1-\tau)\,\bar{x}_1. \tag{19}$$

Hence $\bar{x}_1(\tau)$ falls linearly in τ and $\bar{x}_1(\tau) \to 0$ as $\tau \to 1$. With more similar preferences, women and men make more similar occupational choices. All else equal, fewer women prefer the low-wage, high-flexibility occupation 2, and more women consider the high-wage occupation 1. When we reintroduce discrimination into the model, and following Proposition 2, the next period's level of segregation is determined by

$$x_{1,t+1} = \lambda\left[\Phi_m\Big(\mu(1+\gamma\,\delta(x_{1,t}))\Big) - \Phi_w^\tau\Big(\mu(1-\gamma\,\delta(x_{1,t}))\Big)\right] \tag{20}$$

where $\delta(\cdot)$ is increasing in x_{1t} through the discrimination function $f(\cdot)$. As the policy brings women's preferences closer to men's, $\Phi_w^\tau(\cdot) \geq \Phi_w(\cdot)$ pointwise. Hence the segregation dynamics curve under the policy lies weakly below the baseline: $F^\tau(x) \leq F^0(x)$, with a strict inequality whenever $\tau > 0$. Under the usual stability condition (continuous, increasing F with $F'(x_1^*) \in (0,1)$ at the fixed point), the unique steady state $x_1^*(\tau)$ solves $x_1^*(\tau) = F^\tau\Big(x_1^*(\tau)\Big)$ and satisfies the comparative statics

$$x_1^*(\tau) \;<\; x_1^*(0) \quad \text{for all } \tau > 0, \qquad \text{and} \qquad \frac{\partial x_1^*(\tau)}{\partial \tau} \;<\; 0. \tag{21}$$

The effect of preference-changing policies is amplified by the discrimination mechanism. By reducing the preference gap, the policy first lowers the preferences-only level of segregation, $\bar{x}_1(\tau)$. This also weakens the discrimination response. As more women

[12] Fathers' quotas in parental leave, for instance, durably shift the household division of labor (Patnaik, 2019).

[13] One can equivalently think of men's distribution shifting toward women's, for example because caregiving norms change. What matters is that the distance between the two distributions narrows.

enter the high-wage occupation, women face less discrimination there, which makes the occupation still more attractive to them. The steady state with the policy, $x_1^*(\tau)$, is therefore lower than the original steady state, x_1^*. Policies that equalize preferences have a double benefit: they directly reduce preference-based segregation and indirectly reduce the discrimination that would otherwise increase segregation.

3.3 Policies that change the wage-flexibility trade-off

So far, we have treated the wage-flexibility trade-off, μ, as fixed. This trade-off measures how much wage workers give up for an extra unit of flexibility when moving from the high-wage, low-flexibility occupation to the low-wage, high-flexibility occupation. In practice, technologies and workplace reforms can change the wage-flexibility bundles that occupations offer. For example, remote work (Bloom et al., 2015; Barrero, Bloom and Davis, 2023), flexible scheduling, or shorter workweeks can make high-wage occupations more flexible. We study how such changes affect occupational choices and segregation.

Before we proceed, it is important to clarify that our results discussed above go through whether the wage-flexibility tradeoff is fixed or changes exogenously through technology or policies. Given different wage-flexibility bundles offered by different occupations, i.e. *for a given* μ (or a vector $\boldsymbol{\mu}$, in the n-occupations case), discrimination continues to increase the level of segregation in the steady state compared to the case with only difference in preferences. To keep the analysis tractable, we therefore study how changes in the wage-flexibility trade-off affect segregation in the model without discrimination.

Consider a change in technology such as the ability to work from home that raises the flexibility of the high-wage, low-flexibility occupation 1, while leaving wages unchanged. As Figure 8 shows, this causes the wage-flexibility frontier to become steeper. This change, in effect, increases the tradeoff (slope) μ between wage and flexibility, with the gain in flexibility being lower for the same reduction in wage.

As a consequence, some women and men switch from preferring occupation 2 to preferring occupation 1. If more women than men switch, then a technology that makes the high-wage occupation more flexible reduces segregation. If more men than women switch, segregation increases.

The effect of flexibility-enhancing technology depends on the sex composition of workers near the cutoff. Figure 9 illustrates this point. The dashed vertical line marks the value of μ at which the male and female densities are equal. To the left of this line, there are more men than women close to indifference between the two occupations. A rightward shift in the cutoff therefore changes the preferred occupation of more men than women. For example, if μ moves from μ_1 to μ_1', segregation rises. To the right of the dashed line, there are more women than men close to indifference. A rightward shift then changes the preferred occupation of more women than men. For example, if μ moves from μ_2 to μ_2',

Figure 8: Technology and wage-flexibility tradeoff

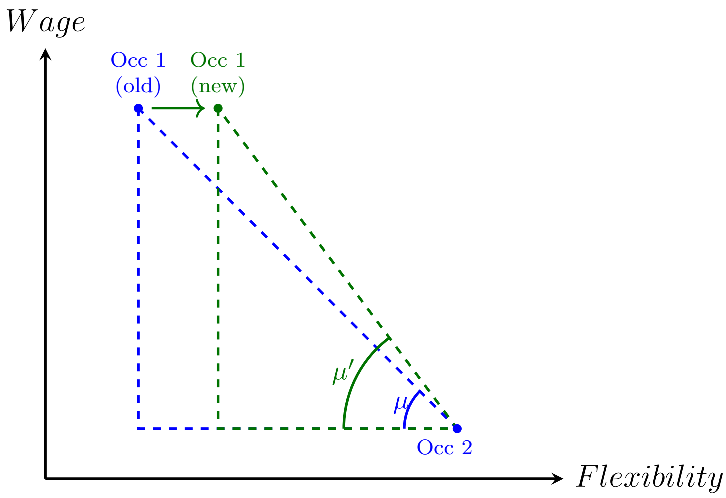

segregation falls.

Figure 9: Technology improvement and segregation

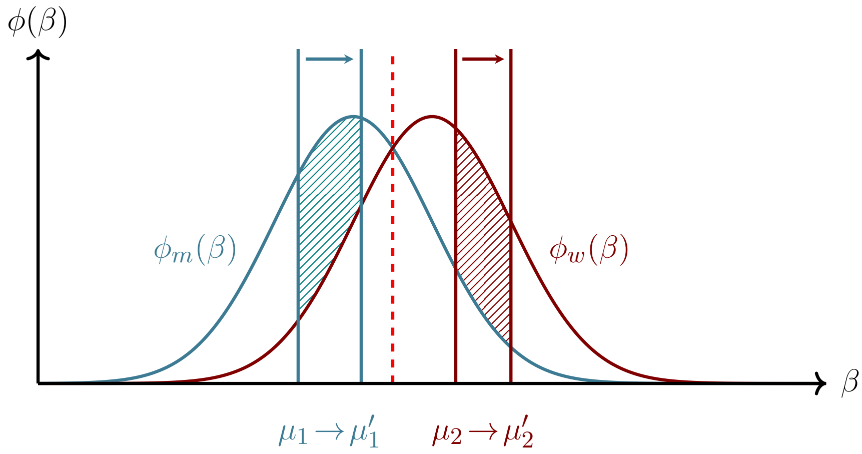

Note: The figure shows two possible effects of a technology that makes occupation 1 more flexible. In the left case, the cutoff moves from μ_1 to μ_1'; in the right case, it moves from μ_2 to μ_2'. The shaded regions show the difference between the male and female densities between the old and new cutoffs.

To see this analytically, consider the preferences-only benchmark from Equation 15, written as a function of the wage-flexibility trade-off:

$$\bar{x}_1(\mu) = \lambda \left[\Phi_m(\mu) - \Phi_w(\mu) \right].$$

Differentiating with respect to μ gives

$$\frac{d\bar{x}_1}{d\mu} = \lambda \left(\phi_m(\mu) - \phi_w(\mu) \right). \tag{22}$$

A change in technology that increases the wage-flexibility (μ) tradeoff will further increase segregation if the *pdf* of willingness to pay for men at the initial equilibrium is higher than that for women, and vice-versa. Hence, we have the following result.

$$\frac{d\bar{x}_1}{d\mu}\begin{cases} > 0 & \textit{if } \phi_m(\mu) > \phi_w(\mu), \\ = 0 & \textit{if } \phi_m(\mu) = \phi_w(\mu), \\ < 0 & \textit{if } \phi_m(\mu) < \phi_w(\mu). \end{cases}$$

The sign of the derivative depends on the relative densities at the cutoff. The derivative is positive when the male density is higher, negative when the female density is higher, and zero when the densities are equal. The zero case corresponds to the dashed vertical line in Figure 9.

The same logic extends to the n-occupation case. With many occupations, each pair of occupations has its own wage-flexibility trade-off and therefore its own cutoff. If all relevant cutoffs lie to the left of the point where the male and female densities are equal, then flexibility-enhancing technology increases segregation. If all relevant cutoffs lie to the right of that point, it decreases segregation. If some cutoffs lie on each side, the overall effect is ambiguous. The effect of technology therefore depends on the initial configuration of occupations. This ambiguity does not affect our main result: for any given wage-flexibility trade-off, discrimination amplifies the segregation generated by preferences.

4 Conclusion

Why do women and men choose different occupations? In particular, why do women remain underrepresented in male-dominated occupations like engineering? These questions are widely debated in academia, in the news, and on social media. Two of the most frequently discussed reasons are sex differences in preferences and discrimination. They are often pitched as competing explanations. Both sides of this debate are intuitively appealing. One can easily imagine that male-dominated occupations have job characteristics that fewer women find appealing. Similarly, one can easily imagine that women are treated more poorly in those occupations.

We contribute to this debate with a model that shows how both factors interact. The key ingredient in our model is how we link discrimination and preferences: we assume that discrimination increases with occupational sex imbalance, so that workers face worse treatment when their sex is underrepresented in an occupation.

We show two key results. First, discrimination amplifies preference-based segregation. If few women find the job characteristics of engineering appealing, engineering becomes male-dominated. Because women expect worse treatment as the share of men in an occupation rises, engineering then becomes even less appealing to women. Preferences change occupational composition, and occupational composition feeds back through discrimination.

This multiplier also works in reverse. Policies that reduce discrimination first make male-dominated occupations more attractive to women in the current cohort. As more women enter those occupations, their sex composition changes. They become less male-dominated, which further reduces the discrimination women expect to face there. The same logic applies to policies that narrow sex differences in preferences. These policies reduce preference-based segregation directly; by making occupations less segregated, they also reduce the discrimination that would otherwise amplify the initial sorting. However, even eliminating discrimination would not fully integrate the labor market as long as sex differences in preferences remain.

Policies that change the wage-flexibility trade-off are different. It may seem natural to expect remote work, flexible scheduling, or similar reforms to reduce segregation by making high-wage occupations more attractive to women. The model shows why this need not be true. Such reforms reduce segregation only if the workers drawn into the high-wage occupation are disproportionately women. If the marginal workers are mostly men, making the high-wage occupation more flexible can increase segregation.

Second, preferences and discrimination are negatively correlated. The occupations that few women find attractive because of their job characteristics are also the occupations in which women face more discrimination. This matters for empirical work on preferences and labor-market outcomes, such as occupational segregation or the gender wage gap (Cortes and Pan, 2018; Wiswall and Zafar, 2018). For example, Wiswall and Zafar (2018) carefully measure preferences for job characteristics of university students using hypothetical choice experiments. They then explore how the estimated gender wage gap changes once they control for these preferences and show that preferences "account for" at least a quarter of the gender wage gap. Wiswall and Zafar (2018) and careful readers of their paper likely recognize that preferences may be correlated with omitted variables. Our model makes this concern more specific. Because preferences shape occupational composition, and occupational composition shapes discrimination, preference measures partly proxy for discrimination. The bias therefore has a direction: controlling for flexibility can attribute to flexibility some of the wage gap that is actually caused by discrimination. Women may choose lower-paying occupations not only because they are more flexible, but also because they involve less discrimination.

The feedback mechanism we study is more general than the particular model we use to formalize it. We focus on sex differences in preferences for flexibility, discrimination at work, and occupational sorting, but none of these ingredients is essential. The same mechanism would arise with other preferences (e.g., competitiveness, Buser, Niederle and Oosterbeek, 2014), other groups, and other reasons why being underrepresented makes a setting less attractive. It applies whenever people sort into places (e.g., residential segregation, Card, Mas and Rothstein, 2008), and the composition of those places changes how attractive they are to different groups. Preferences not only sort people into places; by changing who is there, they also change what those places are like.

Appendix

A: Survey evidence for discrimination assumption

In this appendix, we present survey evidence for the assumption that the risk of being treated poorly increases with the share of opposite-sex workers in an occupation. The Qualtrics survey file (QSF file), the survey text, data, and analysis files for this section are available in the Open Science Framework (OSF) repository at https://osf.io/g2umr/.

Background: Survey and Occupation data

We ran the survey from November 3 to 6 of 2020 with US based participants on Amazon Mechanical Turk. We received ethics approval for this survey from the University of Western Australia. From this survey, we use three questions. The first one asked about respondents' sex. The second and third question asked respondents to guess the percentage of female workers (question 2) or male workers (question 3) who have been treated poorly at work because of their sex over the last 12 months in 10 occupations. We introduced questions 2 and 3 with four examples of poor treatment at work (being given bad work tasks, receiving inappropriate comments, being socially excluded, being ridiculed) and stated that *"People can be treated poorly because of their sex by, for example, co-workers, superiors, customers or clients."* (see OSF repository for full survey text). Respondents could state their guesses using a slider for 10 occupations (randomly selected from 60). Figure A.1 shows a screenshot of the question about female workers' poor treatment.

Tables A.1 and A.2 list the 60 occupations and the share of workers in those occupations that are women according to the 2018 American Community Survey (ACS). The 60 occupations are occupations which have a large number of workers and are easy to understand. For example, we included cooks because most people have a good idea of what a cook does. We did not include "First-Line Supervisors Of Non-Retail Sales Workers" because respondents would have found it harder to imagine the work environment for this occupation.

In total, 303 people started the survey. One did not consent and seven failed the attention check, leaving us with 295 people in the estimation sample (103 women and 192 men). The average age of respondents was 34.7 years.

Results: evidence for discrimination assumption

Figure A.2 shows the relationship between the share of women in an occupation and the average guess of the share of female workers (left) or male workers (right) who have been treated poorly in that occupation in the last 12 months. Each dot represents one occupation. The figure shows a negative relationship for women: women expect that women are, on average, treated less poorly because of their sex in an occupation with a higher share of women. The opposite holds for men. They expect more discrimination against men in occupations with larger shares of women.

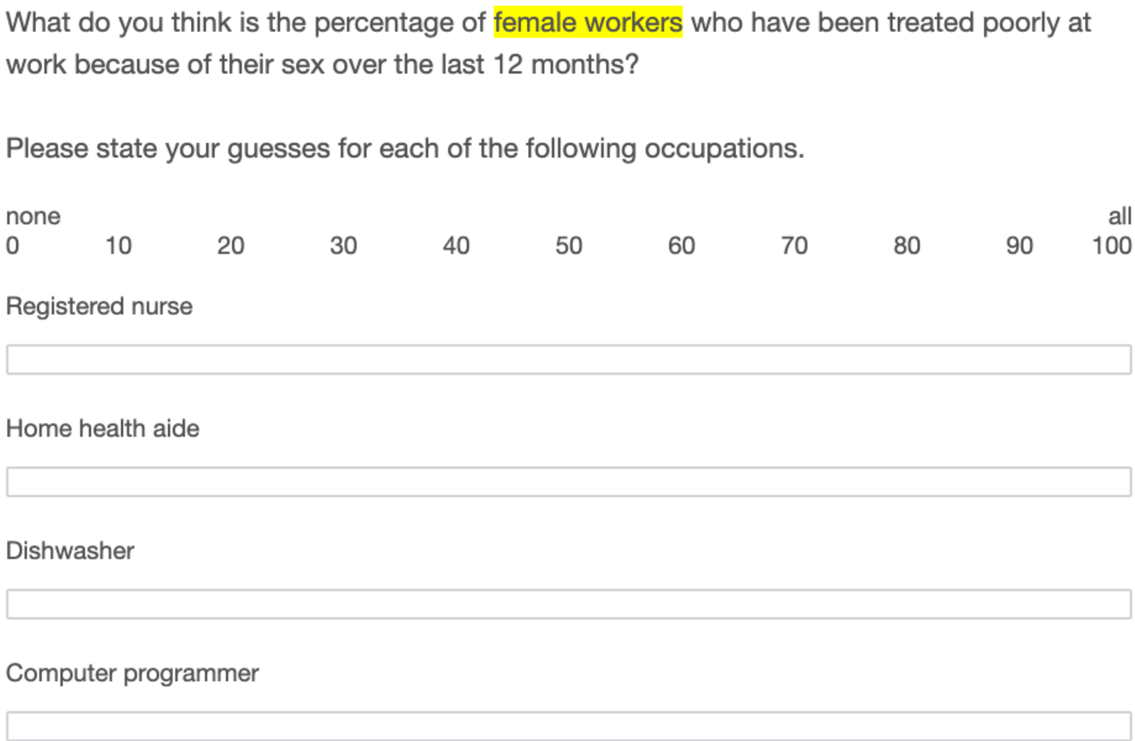

What do you think is the percentage of female workers who have been treated poorly at work because of their sex over the last 12 months?

Please state your guesses for each of the following occupations.

none | all
0 10 20 30 40 50 60 70 80 90 100

Registered nurse

Home health aide

Dishwasher

Computer programmer

Figure A.1: Screenshot question female workers treated poorly in occupation

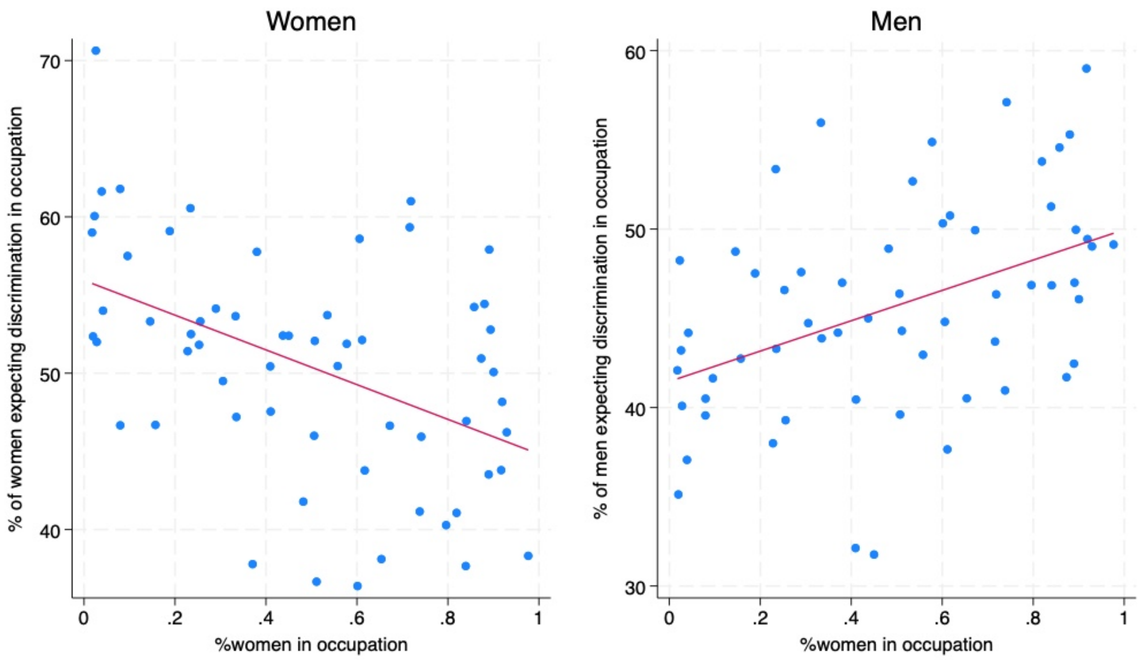

Figure A.2: Share of women/men expecting discrimination by share of women in occupation

Table A.1: Occupations and percent female (part 1).

Occupation (ACS label)	Percent Female
RPR-Automotive Service Technicians And Mechanics	1.8%
CON-Plumbers, Pipefitters, And Steamfitters	2.0%
CON-Carpenters	2.3%
CON-Construction Equipment Operators	2.6%
CON-Electricians	2.8%
CON-Construction Laborers	3.9%
RPR-Maintenance And Repair Workers, General	4.2%
MGR-Construction Managers	7.9%
CLN-Landscaping And Groundskeeping Workers	8.0%
CON-Painters And Paperhangers	9.6%
PRT-Police Officers	14.5%
MGR-Farmers, Ranchers, And Other Agricultural Managers	15.7%
CMM-Software Developers	18.9%
EAT-Dishwashers	22.8%
PRT-Security Guards And Gaming Surveillance Officers	23.4%
CMM-Computer Programmers	23.5%
CMM-Computer Support Specialists	25.3%
EAT-Chefs And Head Cooks	25.6%
SAL-Sales Representatives, Wholesale And Manufacturing	29.0%
FIN-Personal Financial Advisors	30.5%
CLN-Janitors And Building Cleaners	33.3%
MGR-Sales Managers	33.5%
TRN-Stockers And Order Fillers	37.1%
MED-Physicians	38.0%
OFF-Postal Service Mail Carriers	40.9%
BUS-Management Analysts	41.0%
BUS-Project Management Specialists	43.7%
EAT-Cooks	45.0%
SAL-Insurance Sales Agents	48.2%
EDU-Postsecondary Teachers	50.6%

Table A.2: Occupations and percent female (part 2).

Occupation (ACS label)	Percent Female
MGR-Property, Real Estate, And Community Association Managers	50.7%
MGR-Food Service Managers	51.1%
SAL-Retail Salespersons	53.5%
MGR-Financial Managers	55.8%
SAL-Real Estate Brokers And Sales Agents	57.8%
EDU-Secondary School Teachers	60.1%
EAT-Bartenders	60.6%
EAT-Food Preparation Workers	61.1%
FIN-Accountants And Auditors	61.7%
MGR-Education And Childcare Administrators	65.4%
OFF-Customer Service Representatives	67.2%
EAT-Waiters And Waitresses	71.6%
MGR-Medical And Health Services Managers	71.9%
SAL-Cashiers	73.8%
BUS-Human Resources Workers	74.1%
EDU-Elementary And Middle School Teachers	79.6%
OFF-Office Clerks, General	81.9%
HLS-Personal Care Aides	83.9%
LGL-Paralegals And Legal Assistants	84.1%
EDU-Teaching Assistants	85.8%
OFF-Bookkeeping, Accounting, And Auditing Clerks	87.3%
CLN-Maids And Housekeeping Cleaners	88.0%
OFF-Receptionists And Information Clerks	89.0%
HLS-Nursing Assistants	89.1%
MED-Registered Nurses	89.4%
HLS-Home Health Aides	90.1%
HLS-Medical Assistants	91.7%
PRS-Hairdressers, Hairstylists, And Cosmetologists	91.9%
PRS-Childcare Workers	92.9%
EDU-Preschool And Kindergarten Teachers	97.7%

We confirm these results with regressions of the expected share of people of one's own sex who are treated poorly in an occupation on the share of women in that occupation. We estimate those regressions separately for women and men. Because each respondent stated their guess for 10 occupations, we cluster our standard errors at the individual level. For both sexes, our results confirm the results shown in Figure A.2. For women, we estimate that as the share of women in an occupation increases by 10 percentage points, the expected probability of poor treatment of women in that occupation (as guessed by other women) decreases by 1.2 percentage points (p-value <0.001). For men, we estimate that as the share of women in an occupation increases by 10 percentage points the expected probability of poor treatment of men increases by 0.9 percentage points (p-value <0.001).

Our survey evidence supports our key assumption: both women and men expect more poor treatment in occupations with a higher share of opposite-sex workers.

B: Segregation with discrimination but no difference in preferences

The model in the main text discusses occupational segregation with 1) differences in preferences and 2) differences in preferences together with discrimination. An intermediate step is to understand the dynamics that discrimination introduces. Here we discuss the case where people get disutility from discrimination, but there are no sex differences in preferences ($\gamma > 0$ and $\Phi_m(\cdot) = \Phi_w(\cdot) = \Phi(\cdot)$). With no sex differences in preferences, the law of motion becomes

$$x_{1,t+1} = \lambda \left[\Phi(\mu(1+\gamma\delta(x_{1,t}))) - \Phi(\mu(1-\gamma\delta(x_{1,t})))\right], \tag{B.1}$$

where

$$\delta(x_{1,t}) = \frac{f(x_{1,t}) - f(-x_{1,t})}{\text{wage}_1 - \text{wage}_2}.$$

In this scenario there could be either one or two stable equilibria. There would be one stable equilibrium at a perfectly integrated labor market ($x_{1,t} = x_{1,t+1} = 0$) if changes in segregation from a perfectly integrated labor market lead to smaller changes in segregation in the next cohort, that is, when $\partial x_{1,t+1}/\partial x_{1,t} < 1$ at $x_{1,t} = 0$ (see Figure B.1a). In this case, any segregation disappears over time either because most people are far from the cutoffs for preferring occupation 1, $\mu(1+\gamma\delta(x_{1,t}))$ for men and $\mu(1-\gamma\delta(x_{1,t}))$ for women, or because the discrimination effect is weak.

There would be two stable equilibria if starting from a perfectly integrated labor market a change in segregation in the existing cohort led to an even larger change in segregation in the future cohort, that is, if $\partial x_{1,t+1}/\partial x_{1,t}\big|_{x_{1,t}=0} > 1$. One of those equilibria would be at a positive value of $x_{1,t}$ and the other at a negative value of $x_{1,t}$ (see Figure B.1b). That is, discrimination alone could lead to more men in the higher-paying occupation 1 and more women in the more flexible occupation 2, or the other way around. This could be the case if many people are close to being indifferent between the two occupations and small changes in sex composition push many workers in the next cohort to prefer a different occupation. Alternatively, $\partial x_{1,t+1}/\partial x_{1,t}\Big|_{x_{1,t}=0} > 1$ may arise because small changes in sex composition generate a large discrimination effect.

The distinction between the two cases becomes less pronounced when we have differences in preferences and discrimination, as in Proposition 2 of the main text. In both cases, the labor market is segregated even if we start from a perfectly integrated labor market, because $\bar{x}_1 > 0$ from Proposition 1. Any positive stable equilibrium satisfies $x_1^* > \bar{x}_1$.

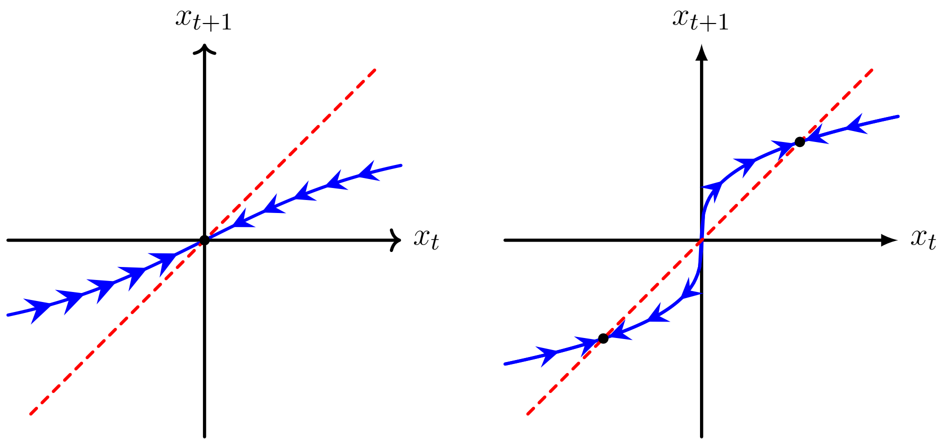

(a) $\partial x_{t+1}/\partial x_t < 1$ (b) $\partial x_{t+1}/\partial x_t > 1$

Figure B.1: Steady states with discrimination and no preference differences

C: Steady states of segregation with counter-intuitive initial condition

We are trying to find steady state equilibria of a dynamic system by finding the fixed points of the function $F(x_{1,t})$. The main text focuses on the intuitive case in which the high-wage, low-flexibility occupation 1 is male-overrepresented, so $x_{1,t} > 0$. There is, however, nothing in the law of motion that rules out counter-intuitive initial conditions with more women in the higher-paying occupation 1 and more men in the more flexible occupation 2, i.e. $x_{1,t} < 0$. Such starting points may arise for historical or institutional reasons.

When $x_{1,t} < 0$, we have $\delta(x_{1,t}) < 0$. Define

$$a(x_{1,t}) = -\gamma\delta(x_{1,t}) > 0.$$

The cutoff for women becomes

$$\beta_i < \mu(1 + a(x_{1,t})), \tag{C.1}$$

while the cutoff for men becomes

$$\beta_i < \mu(1 - a(x_{1,t})). \tag{C.2}$$

Thus, when occupation 1 is initially female-overrepresented, discrimination makes occupation 1 relatively more attractive to women and relatively less attractive to men.

In other words, for $x_{1,t} < 0$ we need to solve for $x_{1,t+1} = F(x_{1,t})$, where

$$F(x_{1,t}) = \lambda\left\{\Phi_m(\mu(1 - a(x_{1,t}))) - \Phi_w(\mu(1 + a(x_{1,t})))\right\}.$$

Although Assumption 1 implies $\Phi_m(\mu) \geq \Phi_w(\mu)$ at a common cutoff, the cutoffs are no longer common when $x_{1,t} < 0$. It is therefore possible that

$$\Phi_w(\mu(1 + a(x_{1,t}))) > \Phi_m(\mu(1 - a(x_{1,t}))),$$

so that $F(x_{1,t}) < 0$. That is, starting from $x_{1,t} < 0$, the next-period imbalance can also be negative. This does not by itself imply that a negative steady state exists. As Figure C.1(a) shows, if $F(x_{1,t})$ does not have a steep slope as $x_{1,t}$ approaches zero from the left, the system has only the positive stable steady state x_1^*.

When the slope of $F(x_{1,t})$ is steep as $x_{1,t}$ approaches zero from the left, there can also be a stable negative steady state. As Figure C.1(b) shows, in this case there are two stable equilibria: x_1^* and x_2^*. The basin of attraction of the positive steady state is larger because the positive vertical intercept, $\bar{x}_1$, reflects the underlying difference in preferences between men and women. Hence, the negative steady state equilibrium does not always exist, and even when it exists, it has a smaller basin of attraction than the positive steady state equilibrium.

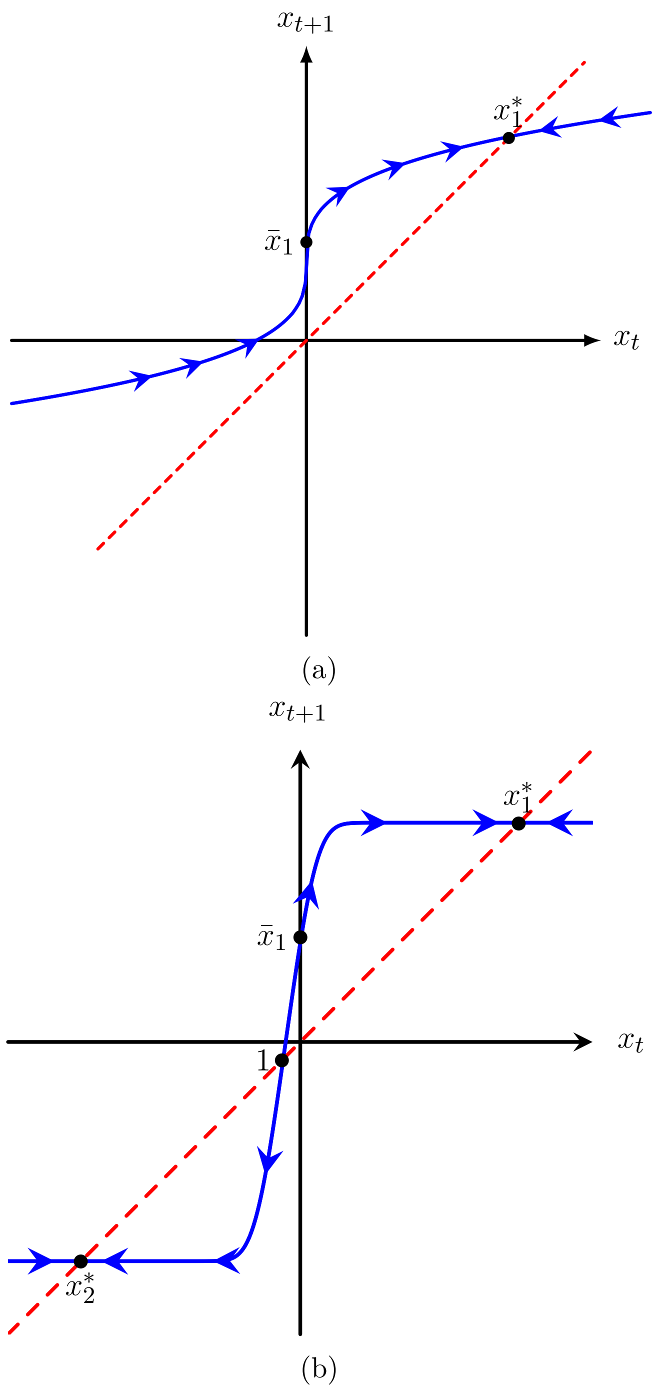

Figure C.1: Steady states from a counter-intuitive initial condition

D: Model with n occupations

This appendix proves Proposition 3. We show three things. First, sex differences in preferences alone generate positive occupational segregation in the n-occupation model (Sections D.2–D.3). Second, discrimination pushes each pair of occupations further apart in the direction of its existing imbalance (Section D.4). Third, and most importantly, the discriminatory steady state is *more* segregated than the preferences-only benchmark, as measured by the Duncan index. This holds *even when discrimination is strong enough to change the order in which workers rank occupations* (Section D.5). The only substantive assumption is a stability condition that is the n-occupation version of the local-stability condition already used in the two-occupation model. A numerical example in Section D.6 traces the dynamics to a steady state and exhibits the negative preference–discrimination correlation, which is not analytically transparent for $n > 2$.

D.1 Setup

We normalize the mass of employed men and of employed women each to one. Let m_k and w_k be the shares of employed men and women in occupation k, and let

$$x_k = m_k - w_k$$

be the male-minus-female imbalance in occupation k. Because both workforces have unit mass, the imbalances always sum to zero,

$$\sum_{k=1}^{n} x_k = 0. \tag{D.1}$$

The Duncan index is

$$D(\boldsymbol{x}) = \tfrac{1}{2}\sum_{k=1}^{n} |x_k| = \tfrac{1}{2}\|\boldsymbol{x}\|_1.$$

Occupations are ordered along a wage–flexibility frontier, so lower-indexed occupations pay more and are less flexible:

$$\text{wage}_1 > \text{wage}_2 > \cdots > \text{wage}_n, \qquad \text{flex}_1 < \text{flex}_2 < \cdots < \text{flex}_n. \tag{D.2}$$

A worker with willingness to pay for flexibility β values occupation k at the utility

$$u_k(\beta; \boldsymbol{a}) = a_k + \beta\,\text{flex}_k,$$

where a_k is the intercept of occupation k's utility. The worker ranks occupations by this utility, so the rank of occupation k is

$$r_k(\beta;\boldsymbol{a}) = 1 + \sum_{\ell \neq k} \mathbf{1}\{u_\ell(\beta;\boldsymbol{a}) > u_k(\beta;\boldsymbol{a})\},$$

and she is realized into her r-th ranked occupation with probability λ_r. We write Φ_m, Φ_w for the CDFs of men's and women's willingness to pay, with densities ϕ_m, ϕ_w; Assumption 1 (first-order stochastic dominance) states $\Phi_w(\beta) \leq \Phi_m(\beta)$ for all β, with strict inequality for some β. We assume throughout that Φ_m and Φ_w are continuous.

The preference-sensitive realized share of this population in occupation k is

$$R_\Phi(\boldsymbol{a})_k = \int \lambda_{r_k(\beta;\boldsymbol{a})} \, d\Phi(\beta), \tag{D.3}$$

where λ_r is the probability that a worker is realized into her r-th ranked occupation. Thus $R_\Phi(\boldsymbol{a})$ is not an additional behavioral assumption; it is only a compact way to convert a distribution of complete rankings into realized employment shares through the rank-probability schedule $(\lambda_1, \ldots, \lambda_n)$. Because Φ has a density, $R_\Phi(\boldsymbol{a})$ is continuous in $\boldsymbol{a}$ and differentiable wherever the pairwise cutoffs are distinct, a property we use in Section D.5.

The rank-probability schedule satisfies

$$\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n \geq 0, \qquad \sum_{r=1}^{n} \lambda_r = \Lambda < 1,$$

with at least one strict inequality. More-preferred occupations are therefore more likely to be realized. The residual probability $1 - \Lambda$ is assigned by a preference-neutral and sex-neutral fallback allocation $\boldsymbol{q} = (q_1, \ldots, q_n)$, where $q_k \geq 0$ and $\sum_k q_k = 1$. The full realized share of a population with CDF Φ in occupation k is therefore

$$R_\Phi(\boldsymbol{a})_k + (1 - \Lambda) q_k.$$

The fallback term can differ across occupations, because some occupations may have more capacity or lower entry barriers, but it is common to men and women. It therefore cancels in the imbalance $x_k = m_k - w_k$.

In the preference-only benchmark both sexes face the same intercept vector, namely $\mathbf{wage} = (\text{wage}_1, \ldots, \text{wage}_n)$. Hence

$$\boldsymbol{x}^0 = R_{\Phi_m}(\mathbf{wage}) - R_{\Phi_w}(\mathbf{wage}). \tag{D.4}$$

D.2 Pairwise decomposition

The body describes this decomposition and the figure that illustrates it; here we prove the two lemmas on which it rests.

Proof of Lemma 1. At $\beta = \eta_{ij}$ the scores of i and j are equal. Suppose some occupation h ranked strictly between them just below the cutoff. For i and j to reverse, one of them must pass h at this same β, which would require $\eta_{ih} = \eta_{ij}$ or $\eta_{jh} = \eta_{ij}$, contradicting distinctness. So no occupation lies between i and j at the cutoff: they are adjacent, and crossing η_{ij} interchanges the two neighbors and leaves all other ranks fixed. □

Proof of Lemma 2. Write ϵ_{ij}, $\Delta\lambda_{ij}$, and σ^0_{ij} for the quantities defined in the body. Intuitively, each cutoff belongs to exactly one pair, and by Lemma 1 it moves only that pair: the higher-wage occupation i loses a rank (a change of $+\sigma^0_{ij}$ in its imbalance) and the lower-wage occupation j gains it (a change of $-\sigma^0_{ij}$). Summation by parts makes this precise.

Order the distinct cutoffs as $0 = t_0 < t_1 < \cdots < t_M < t_{M+1} = \infty$, with $M = \binom{n}{2}$. On each interval (t_s, t_{s+1}) every worker holds the same ranking, so occupation k keeps a constant rank $r_k^{(s)}$, running from $r_k^{(0)} = k$ (the wage order) to $r_k^{(M)} = n - k + 1$ (the flexibility order). Writing $G = \Phi_m - \Phi_w$, equation (D.4) and the definition of R_Φ give

$$x^0_k = \sum_{s=0}^{M} \lambda_{r_k^{(s)}} \Big[G(t_{s+1}) - G(t_s) \Big].$$

Because Φ_m and Φ_w both integrate to one, $G(t_0) = G(t_{M+1}) = 0$, and summation by parts gives

$$x^0_k = -\sum_{s=1}^{M} \Big(\lambda_{r_k^{(s)}} - \lambda_{r_k^{(s-1)}} \Big) G(t_s).$$

Each t_s is a single cutoff η_{ij} at which, by Lemma 1, only the pair (i,j) swaps, so the increment vanishes unless $k \in \{i,j\}$. If $k = i$ (the higher-wage member, paired with some $j > k$), occupation k drops from rank r to $r+1$, the increment is $\lambda_{r+1} - \lambda_r = -\Delta\lambda_{kj}$, and $G(\eta_{kj}) = \epsilon_{kj}$, contributing $+\epsilon_{kj}\Delta\lambda_{kj} = +\sigma^0_{kj}$. If $k = j$ (the lower-wage member, paired with some $i < k$), occupation k rises from rank $r+1$ to r, the increment is $+\Delta\lambda_{ik}$, and the contribution is $-\sigma^0_{ik}$. Summing the nonzero terms yields (17). □

D.3 Preferences alone generate segregation

We now show that the pairwise contributions cannot all cancel: as long as at least one is strictly positive, the preferences-only allocation is segregated. The argument runs through the cumulative imbalance up the wage ladder. For each $m = 1, \dots, n$, let

$$X_m = \sum_{k=1}^{m} x_k^0$$

be the net male-minus-female share in the m highest-paying occupations. Lemma 2 writes each X_m as a sum of nonnegative pairwise contributions, so $X_m \geq 0$ at every threshold m: the highest-paying occupations are, taken together, male-overrepresented. The proposition below turns this into strictly positive segregation.

Proposition D.1 (Existence)**.** *If $\sigma_{ij}^0 > 0$ for at least one pair $i < j$, then $\boldsymbol{x}^0 \neq \mathbf{0}$ and $D^0 > 0$.*

Proof. Fix m and substitute the decomposition $x_k^0 = \sum_{j:k<j} \sigma_{kj}^0 - \sum_{i:i<k} \sigma_{ik}^0$ from Lemma 2 into $X_m = \sum_{k=1}^m x_k^0$. A pair (i,j) with $i < j$ enters x_i^0 with a $+$ sign and x_j^0 with a $-$ sign, so its coefficient in X_m is $\mathbf{1}\{i \leq m\} - \mathbf{1}\{j \leq m\} = \mathbf{1}\{i \leq m < j\}$. Hence

$$X_m = \sum_{i \leq m < j} \sigma_{ij}^0 \;\geq\; 0,$$

the total sorting that crosses the cut between the top m occupations and the rest. Every term is nonnegative, so $X_m \geq 0$ for each m.

Now suppose $\sigma_{ij}^0 > 0$ for some pair $i < j$. Taking the cut at $m = i$, the pair (i,j) crosses it, so

$$X_i = \sum_{i' \leq i < j'} \sigma_{i'j'}^0 \;\geq\; \sigma_{ij}^0 \;>\; 0.$$

Because $\sum_k x_k^0 = 0$, the positive entries of $\boldsymbol{x}^0$ sum to exactly $D^0 = \frac{1}{2}\sum_k |x_k^0|$. The partial sum $X_i = \sum_{k \leq i} x_k^0$ omits some positive entries and includes some negative ones, so $X_i \leq D^0$. Therefore

$$D^0 \;\geq\; X_i \;\geq\; \sigma_{ij}^0 \;>\; 0,$$

and in particular $\boldsymbol{x}^0 \neq \mathbf{0}$. $\square$

Remark (When the hypothesis holds)*.* The hypothesis holds whenever men's and women's willingness-to-pay distributions differ strictly at some cutoff η_{ij} whose two adjacent ranks carry distinct realization probabilities, so that $\epsilon_{ij} > 0$ and $\Delta\lambda_{ij} > 0$ there. In particular it is automatic when $\Phi_w < \Phi_m$ throughout the interior of the support and $\lambda_r > \lambda_{r+1}$ for at least one r, and it holds generically otherwise. The condition is exactly the statement that preferences segregate, $\boldsymbol{x}^0 \neq \mathbf{0}$, and is therefore slightly stronger than Assumption 1, which makes $\Phi_w < \Phi_m$ strict at some β but not necessarily at a cutoff.

D.4 Adding discrimination

Women in occupation k face discrimination $f(x_k)$ and men face $f(-x_k)$, where f is increasing with $f(0)=0$. Discrimination enters the model by shifting the intercepts: relative to the no-discrimination world, men and women in occupation k now perceive intercepts

$$a_k^m(\boldsymbol{x}) = \text{wage}_k - \gamma f(-x_k), \qquad a_k^w(\boldsymbol{x}) = \text{wage}_k - \gamma f(x_k). \tag{D.5}$$

A discriminatory steady state is a vector $\boldsymbol{x}^*$ that reproduces itself once we feed these shifted intercepts through the realized-employment map:

$$\boldsymbol{x}^* = R_{\Phi_m}\Big(\boldsymbol{a}^m(\boldsymbol{x}^*)\Big) - R_{\Phi_w}\Big(\boldsymbol{a}^w(\boldsymbol{x}^*)\Big). \tag{D.6}$$

Such a steady state exists. The right-hand side of (D.6) is continuous in $\boldsymbol{x}$ (the map R_Φ is continuous in the intercepts, and $\boldsymbol{a}^m, \boldsymbol{a}^w$ are continuous in $\boldsymbol{x}$ because f is), and it sends the compact, convex set of feasible imbalance vectors into itself, so Brouwer's theorem yields a fixed point.

Consider any two occupations i and j, and suppose i is more male-overrepresented than j, that is $x_i > x_j$. Then $f(x_i) > f(x_j)$ lowers women's relative valuation of i, while $f(-x_i) < f(-x_j)$ raises men's; both forces push this pair further toward i being male-overrepresented and j female-overrepresented. So discrimination reinforces whatever imbalance a pair already has. This is the n-occupation counterpart of the cutoff shifts $\mu(1 \pm \gamma\delta)$ in the two-occupation model.

The body explains why this pairwise reinforcement does not by itself imply that every occupation, or overall segregation, moves in the expected direction; the proof below therefore works with the entire vector of imbalances rather than occupation by occupation.

D.5 The steady state with discrimination

D.5.1 How discrimination moves realized employment

Discrimination acts by shifting the intercepts a_k, so to understand its effect we need to know how realized employment $R_\Phi(\boldsymbol{a})$ responds when the intercepts move. We compute this response for each sex separately. Throughout this subsection we hold one sex fixed, with intercepts $\boldsymbol{a}$, CDF Φ, and density ϕ, so that the cutoff $\eta_{k\ell}$, the density at it, and the coefficient $c_{k\ell}$ defined below all belong to that sex; in the notation of Lemma 3, this is $c_{k\ell}^s$ with the sex made implicit. Because men and women face different intercepts, they face different cutoffs; we combine the two in Section D.5.2. We build the response up from a single pair of occupations before assembling the full picture.

Consider two occupations k and ℓ, and take $k < \ell$, so that k is the higher-paying, less flexible of the two. As we saw in Section D.2, there is a single willingness-to-pay cutoff $\eta_{k\ell}$ at which a worker is indifferent between them: workers below the cutoff prefer k, workers above it prefer ℓ. The convention $k < \ell$ orients Figure D.1 and the discussion that follows; the coefficient $c_{k\ell}$ defined below is symmetric in k and ℓ, so the labeling is immaterial. Discrimination changes the intercepts a_k and a_ℓ, and so moves this cutoff. Since

$$\eta_{k\ell} = \frac{a_k - a_\ell}{\text{flex}_\ell - \text{flex}_k},$$

it is the *gap* between the two intercepts, $a_k - a_\ell$, that pins down the cutoff. If a change in intercepts widens this gap by one unit, the cutoff $\eta_{k\ell}$ shifts by $1/(\text{flex}_\ell - \text{flex}_k)$ in β. The flexibility difference is the conversion factor; the larger it is, the less the cutoff moves for a given change in intercepts.

When the cutoff moves, the workers it passes over change which of k and ℓ they prefer. How many change is set by how densely workers sit at the cutoff: each unit of cutoff movement flips the preferences of $\phi(\eta_{k\ell})$ workers. Which way preferences flip follows which way the cutoff moves: widening the gap $a_k - a_\ell$ makes more workers prefer k, narrowing it makes more prefer ℓ, and the number who flip is the same either way. Figure D.1 shows one such case.

Switching preference, however, is not yet a change in realized employment: a worker reaches her preferred occupation only probabilistically. Because k and ℓ sit at adjacent ranks at the cutoff (Lemma 1), a worker who switches merely swaps these two neighboring ranks, so the realized employment that moves with her is the rank gap $\Delta\lambda_{k\ell} = \lambda_r - \lambda_{r+1}$ introduced in Section D.2, the extra chance of being placed one rank higher.

Figure D.1: The marginal flow between one pair of occupations

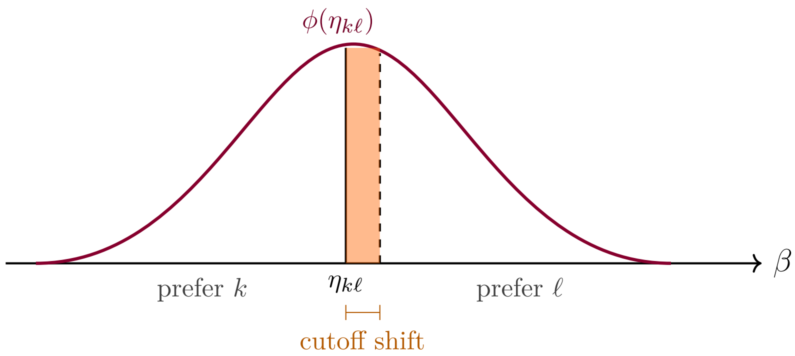

Note: The curve is the density of workers $\phi(\beta)$ over willingness to pay. The cutoff $\eta_{k\ell}$ separates workers who prefer k from those who prefer ℓ. Raising occupation k's intercept slides the cutoff to the right; only the shaded region of workers switch from preferring ℓ to k.

Collecting the pairs. The three pieces above combine into one coefficient per pair. For occupations k and ℓ, define

$$c_{k\ell} = \underbrace{\phi\big(\eta_{k\ell}\big)}_{\substack{\text{workers at} \\ \text{the cutoff}}} \times \underbrace{\Delta\lambda_{k\ell}}_{\substack{\text{realized employment} \\ \text{per switch}}} \times \underbrace{\frac{1}{|\text{flex}_\ell - \text{flex}_k|}}_{\substack{\text{cutoff movement per} \\ \text{unit intercept change}}} \geq 0.$$

This is the amount of realized employment that moves between k and ℓ per unit change in the intercept gap $a_k - a_\ell$. It measures the *rate* of flow along the k–ℓ margin, not its direction; the three factors are each nonnegative, so $c_{k\ell} \geq 0$. The coefficient belongs to the *pair*, not to either occupation: the cutoff $\eta_{k\ell}$, the density at it, the rank gap, and the size of the flexibility difference are all read the same whether we start from k or from ℓ, so $c_{k\ell} = c_{\ell k}$. A single occupation k thus sees $n-1$ such margins, one with each of the other occupations, but the economy has only $\binom{n}{2}$ distinct coefficients in all, one per cutoff, since the margin between k and ℓ is the same object viewed from either side.

The shift discrimination produces. What sets the direction is how far discrimination moves each occupation's intercept, and this differs by sex. From (D.5), relative to the no-discrimination wages, men's intercept in occupation k moves by $h_k^m = -\gamma f(-x_k)$ and women's by $h_k^w = -\gamma f(x_k)$. Each depends on the occupation's own imbalance x_k, so it differs across occupations. For men, h_k^m carries the sign of x_k: in a male-overrepresented occupation ($x_k > 0$) we have $f(-x_k) < 0$, so $h_k^m > 0$ and the occupation becomes more attractive to men; in a female-overrepresented one it is the reverse. For women the signs flip. On the k–ℓ margin, what matters is the change in the intercept gap; for men this is $h_k^m - h_\ell^m = -\gamma[f(-x_k) - f(-x_\ell)]$, which has the same sign as $x_k - x_\ell$ since f is increasing.

Proof of Lemma 3. Fix a sex $s \in \{m, w\}$ with continuous CDF Φ_s and density ϕ_s, and let h_k^s be its intercept shifts; at intercepts where the pairwise cutoffs are distinct, R_{Φ_s} is differentiable (Section D.1). Occupation k's share changes only through its margins with the other occupations. Along the margin with ℓ, what matters is the change in the intercept gap, $h_k^s - h_\ell^s$. A unit change in that gap moves $c_{k\ell}^s$ of realized employment, into k when $h_k^s > h_\ell^s$ and out of k when $h_k^s < h_\ell^s$, so this margin contributes $c_{k\ell}^s(h_k^s - h_\ell^s)$. Summing over the $n-1$ partners ℓ gives the stated change,

$$\sum_{\ell \neq k} c_{k\ell}^s (h_k^s - h_\ell^s). \tag{D.7}$$

If the change reorders some cutoffs, the $c_{k\ell}^s$ just take their current values; the statement is local and needs no fixed ordering. □

Picture the occupations as points, with an edge joining every pair. The weight $c_{k\ell}$ on an edge is how much realized employment can move along it: large where many workers sit near that pair's cutoff, zero where none do. The term $c_{k\ell}(h_k - h_\ell)$ is the flow along the edge, and employment moves toward whichever of the two occupations had the larger intercept rise. Two facts follow from accounting alone. First, the changes across occupations sum to zero: employment leaving ℓ for k is added to k and subtracted from ℓ, so nothing is created or lost. Second, each edge moves employment only *between* its two occupations, in proportion to the gap $h_k - h_\ell$, and always with a nonnegative weight. The full array of these edge weights is the matrix L^* we assemble at the steady state below. Both facts carry over to L^*: its rows sum to zero and its weights are nonnegative. In Section D.5.3 these properties make $I - L^*$ an averaging map, the step that links discrimination to the Duncan index.

D.5.2 Building the feedback

Comparing the no-discrimination world to the discriminatory steady state in a single step is not valid in general: if discrimination is strong, the cutoffs can reorder along the way, and the pairwise weights $c_{k\ell}$ of Lemma 3 change. We avoid this by moving gradually between the two worlds and accumulating the marginal changes. At the steady state $\boldsymbol{x}^*$, the intercept shifts discrimination induces, relative to the wage intercepts, are

$$h_k^m = -\gamma\, f(-x_k^*), \qquad h_k^w = -\gamma\, f(x_k^*),$$

so men face intercepts $\text{wage}_k + h_k^m$ and women $\text{wage}_k + h_k^w$. We reach these intercepts gradually. For each sex s, let occupation k's intercept travel along the straight-line path

$$a_k^s(t) = \text{wage}_k + t\, h_k^s, \qquad t \in [0,1],$$

from no discrimination at $t = 0$ to the steady state at $t = 1$; the parameter t records how far discrimination has been switched on, with the shifts h^s held at their steady-state values. Along the path a pair's cutoff and its coefficient both move with t,

$$\eta_{k\ell}^s(t) = \frac{(\text{wage}_k - \text{wage}_\ell) + t\,(h_k^s - h_\ell^s)}{\text{flex}_\ell - \text{flex}_k}, \qquad c_{k\ell}^s(t) = \frac{\phi_s\big(\eta_{k\ell}^s(t)\big)\, \Delta\lambda_{k\ell}(t)}{|\text{flex}_\ell - \text{flex}_k|}.$$

As t rises the cutoff slides, so the density of workers at it, $\phi_s(\eta_{k\ell}^s(t))$, changes continuously; the rank gap $\Delta\lambda_{k\ell}(t)$ is constant except at the isolated instants where a reordering changes which ranks the pair occupies. Thus $c_{k\ell}^s(t)$ is the *instantaneous* rate at which sex s's realized employment crosses the k–ℓ margin as discrimination is turned up at stage t.

Accumulating this rate over the transition gives the total change in sex s's realized employment in occupation k,

$$R_{\Phi_s}(\text{wage}+h^s)_k - R_{\Phi_s}(\text{wage})_k = \int_0^1 \sum_{\ell\neq k} c^s_{k\ell}(t)\,(h^s_k - h^s_\ell)\,dt = \sum_{\ell\neq k}(h^s_k - h^s_\ell)\,\bar{c}^s_{k\ell}, \qquad \bar{c}^s_{k\ell} = \int_0^1 c^s_{k\ell}(t)\,dt.$$

The gap change $h^s_k - h^s_\ell$ is fixed along the path, so it leaves the integral, and everything that varies as discrimination turns on—the sliding cutoff, the changing density, any reordering—collects into the single number $\bar{c}^s_{k\ell}$, the *cumulative capacity* of the k–ℓ margin: how much realized employment it carries over the whole transition. The integrand is piecewise continuous, its only jumps being in $\Delta\lambda_{k\ell}$ at finitely many instants, so $\bar{c}^s_{k\ell}$ is well defined; at no point did we need the ranking to stay fixed.

The imbalance is $\boldsymbol{x} = \boldsymbol{m} - \boldsymbol{w}$, so its change combines the two sexes,

$$x^*_k - x^0_k = \sum_{\ell\neq k}\Big[(h^m_k - h^m_\ell)\,\bar{c}^m_{k\ell} - (h^w_k - h^w_\ell)\,\bar{c}^w_{k\ell}\Big].$$

Both gap changes point the same way. Writing them through the secant slopes of the discrimination function f,

$$h^m_k - h^m_\ell = \gamma\, S^m_{k\ell}\,(x^*_k - x^*_\ell), \qquad -(h^w_k - h^w_\ell) = \gamma\, S^w_{k\ell}\,(x^*_k - x^*_\ell),$$

$$S^m_{k\ell} = \frac{f(-x^*_\ell) - f(-x^*_k)}{x^*_k - x^*_\ell} \geq 0, \qquad S^w_{k\ell} = \frac{f(x^*_k) - f(x^*_\ell)}{x^*_k - x^*_\ell} \geq 0,$$

both nonnegative because f is increasing. The slope $S^s_{k\ell}$ measures how sharply discrimination responds to the imbalance gap between the two occupations: it is large where f is steep over the relevant range, so that a difference in imbalance opens a wide gap in how the two occupations treat sex s, and it tends to f' at the common point as x^*_k and x^*_ℓ converge. The men's slope is taken over $-x^*_k$ and $-x^*_\ell$, since men face $f(-x)$; the women's over x^*_k and x^*_ℓ.

Substituting and factoring out the common gap $(x^*_k - x^*_\ell)$ writes occupation k's total change as a sum of pairwise terms,

$$x^*_k - x^0_k = \sum_{\ell\neq k}\omega^*_{k\ell}\,(x^*_k - x^*_\ell), \qquad \omega^*_{k\ell} = \gamma\Big(S^m_{k\ell}\,\bar{c}^m_{k\ell} + S^w_{k\ell}\,\bar{c}^w_{k\ell}\Big).$$

Each sex contributes the product of two nonnegative pieces: how sharply its discrimination responds to the imbalance gap, $S^s_{k\ell}$, and how much employment the margin carries over the transition, $\bar{c}^s_{k\ell}$. Scaled by the taste for avoiding discrimination γ, their sum is the realized employment discrimination moves across the k–ℓ margin per unit of imbalance gap. It is nonnegative and symmetric in k and ℓ because every factor is, and it vanishes when $x^*_k = x^*_\ell$: with no gap between the two occupations there is nothing to amplify. The construction uses only that f is increasing, not its shape. By conservation the same margin contributes $-\omega^*_{k\ell}(x^*_k - x^*_\ell)$ to occupation ℓ, equal and opposite.

Collect the pair weights into the symmetric matrix W, with $W_{k\ell} = \omega^*_{k\ell}$ and $W_{kk} = 0$, and let $D = \operatorname{diag}(d_1, \dots, d_n)$ record each occupation's total margin weight, $d_k = \sum_{\ell \neq k} \omega^*_{k\ell}$. Then

$$L^* = D - W$$

is the Laplacian of this weighted network, and the sum above is exactly $(L^* \boldsymbol{x}^*)_k$: the diagonal d_k multiplies x^*_k, and each off-diagonal $-\omega^*_{k\ell}$ multiplies x^*_ℓ. In vector form,

$$\boldsymbol{x}^* - \boldsymbol{x}^0 = L^* \boldsymbol{x}^*, \qquad \text{equivalently} \qquad \boldsymbol{x}^0 = (I - L^*)\, \boldsymbol{x}^*. \tag{D.8}$$

By construction L^* is symmetric, its rows and columns sum to zero, and its weights are nonnegative; the weights are accumulated along the actual path, so reordering changes them step by step but not these properties. Through the feedback $\boldsymbol{x}^* = \boldsymbol{x}^0 + L^* \boldsymbol{x}^*$, L^* pushes the steady-state imbalances further apart than the preferences-only ones, and Section D.5.3 turns this into the Duncan comparison.

D.5.3 Discrimination raises the Duncan index

Assumption D.1 (Stable feedback)**.** *The feedback satisfies $\rho(L^*) < 1$, where $\rho(L^*)$ is the largest eigenvalue of L^*.*

This is not a new economic assumption. In the two-occupation model the steady state is locally stable exactly when the best-response slope satisfies $0 < F'(x^*) < 1$; here L^* is built from the same marginal worker responses that govern how the system moves near $\boldsymbol{x}^*$, so $\rho(L^*) < 1$ is the n-occupation version of that condition. It says the discrimination feedback amplifies segregation but does not overshoot and destabilize the equilibrium. When γ is large enough that $\rho(L^*) \geq 1$, the steady state ceases to be stable and is no longer the object we compare to, the regime in which the two-occupation model develops the unstable, path-dependent equilibria of Appendix C.

Under Assumption D.1 the identity (D.8) also inverts: $I - L^*$ is nonsingular and

$$\boldsymbol{x}^* = (I - L^*)^{-1} \boldsymbol{x}^0 = \sum_{j \geq 0} (L^*)^j\, \boldsymbol{x}^0,$$

reading the discriminatory steady state as the preferences-only baseline amplified through successive rounds of feedback: $\boldsymbol{x}^0$ spread once by L^*, that spread again, and so on. The series converges precisely because $\rho(L^*) < 1$, which is the clearest reason to impose it. The expression relates $\boldsymbol{x}^*$ and $\boldsymbol{x}^0$ at the steady state rather than solving one from the other, since L^* is itself evaluated at $\boldsymbol{x}^*$; for the proof we keep the un-inverted form $\boldsymbol{x}^0 = (I - L^*)\boldsymbol{x}^*$, read as an average.

Proposition D.2 (Global Duncan amplification)**.** *Suppose the willingness-to-pay distributions are continuous, f is increasing, the realization probabilities are rank-monotone, and the steady state satisfies Assumption D.1. Then $D^* \geq D^0$, with strict inequality whenever the feedback links, with positive weight, a male-overrepresented occupation ($x_k^* > 0$) to a female-overrepresented one ($x_\ell^* < 0$).*

Proof. Let $P = I - L^*$. Since L^* has rows and columns summing to zero, those of P sum to one. Writing $d_k = \sum_{\ell \neq k} \omega_{k\ell}^* \geq 0$, we have $P_{kk} = 1 - d_k$ and $P_{k\ell} = \omega_{k\ell}^* \geq 0$ for $k \neq \ell$. Because its edge weights are nonnegative, L^* is positive semidefinite, so $d_k = \boldsymbol{e}_k^\top L^* \boldsymbol{e}_k \leq \rho(L^*) < 1$ and $P_{kk} > 0$. Hence P has nonnegative entries and rows summing to one: each x_k^0 is a weighted average of the entries of $\boldsymbol{x}^*$. An average lies between the values it mixes, so it cannot exceed the largest in magnitude; summing over occupations,

$$\|\boldsymbol{x}^0\|_1 = \sum_k \left| \sum_\ell P_{k\ell} x_\ell^* \right| \leq \sum_k \sum_\ell P_{k\ell} |x_\ell^*| = \sum_\ell |x_\ell^*| \sum_k P_{k\ell} = \|\boldsymbol{x}^*\|_1.$$

Dividing by two gives $D^0 \leq D^*$. The inequality is strict exactly when, for some row k, the terms $P_{k\ell} x_\ell^*$ do not all share a sign, that is, when row k puts positive weight on both a positive and a negative x_ℓ^*, the mixing condition in the statement. □

Identity (D.8) says the preferences-only vector is a weighted average of the discriminatory one, $\boldsymbol{x}^0 = P\boldsymbol{x}^*$. Averaging pulls numbers toward each other, so it can only shrink the total spread, and the Duncan index is exactly half that spread. Read backwards, the discriminatory vector is the un-averaged, more spread-out one, so it carries weakly more segregation. The argument never claims any single occupation moves in the intuitive direction; as (D.1) guarantees, some do not. It shows only that, taken together, discrimination spreads the imbalances apart, and that spread is segregation.

Remark (The n-occupation results)*.* With Propositions D.1 and D.2, the n-occupation results are analytic: preferences alone generate positive segregation, and discrimination weakly raises the Duncan index under stable feedback. Together they establish Proposition 3.

D.6 Simulation

To compute the discriminatory steady state numerically, and from it the flexibility–discrimination relationship plotted in Figure 6, we adopt a concrete, tractable parameterization. We take the preference gap to be constant across cutoffs, $\epsilon_{ij} \equiv \epsilon$, and willingness to pay to be locally uniform with support length $\bar{\beta}$, so the density at every cutoff is $1/\bar{\beta}$; we also use the linear discrimination function $f(x) = x$.

Under the constant gap, the pairwise decomposition of Section D.2 telescopes to a closed form for the preferences-only imbalance,

$$x_k^0 = \epsilon\,(\lambda_k - \lambda_{n-k+1}) \equiv \epsilon\, h_k,$$

so the baseline sorting vector $\mathbf{h}$ is fixed by the realization schedule (λ_r) alone. With the uniform density and linear f, the cutoff shift that discrimination induces is linear in the current imbalance, so the steady-state map of Section D.5 linearizes and the dynamics become

$$\mathbf{x}_{t+1} = A\,\mathbf{x}_t + \epsilon\,\mathbf{h}, \tag{D.9}$$

where the feedback matrix A collects the pairwise coefficients $c_{k\ell}$ of Section D.5 evaluated at the uniform density $1/\bar{\beta}$. It is the constant-gap, uniform, linear specialization of the averaging operator L^* (the discrimination intensity γ enters A through the size of the intercept shifts), and its largest eigenvalue is the $\rho(L^*)$ of Assumption D.1.

We set $n = 4$ occupations and $\Lambda = \sum_r \lambda_r = 0.95$, start from an integrated market $\mathbf{x}_0 = \mathbf{0}$, and iterate (D.9) to its steady state $\boldsymbol{x}^*$; every parameterization shown satisfies $\rho(A) < 1$ (Assumption D.1), so the iteration converges. At the steady state we read off each occupation's imbalance x_k^*, which under Assumption 2 determines the discrimination each sex faces ($f(x_k^*)$ for women and $f(-x_k^*)$ for men); Figure 6 plots these imbalances against occupational flexibility, confirming the negative preference–discrimination correlation across occupations.

References

Adams-Prassl, A., K. Huttunen, E. Nix, and N. Zhang (2024). Violence against women at work. *The Quarterly Journal of Economics 139*(2), 937–991.

Adda, J., C. Dustmann, and K. Stevens (2017). The career costs of children. *Journal of Political Economy 125*(2), 293–337.

Andrew, A., S. Cattan, M. Costa Dias, C. Farquharson, L. Kraftman, S. Krutikova, A. Phimister, and A. Sevilla (2022). The gendered division of paid and domestic work under lockdown. *Fiscal Studies 43*(4), 325–340.

Arrow, K. J. (1973). The theory of discrimination. In O. Ashenfelter and A. Rees (Eds.), *Discrimination in Labor Markets*, pp. 3–33. Princeton, NJ: Princeton University Press.

Barrero, J. M., N. Bloom, and S. J. Davis (2023). The evolution of work from home. *Journal of Economic Perspectives 37*(4), 23–49.

Becker, G. S. (1971). The economics of discrimination. *University of Chicago Press.*

Bergmann, B. R. (1974). Occupational segregation, wages and profits when employers discriminate by race or sex. *Eastern Economic Journal 1*(2), 103–110.

Blau, F. D. and L. M. Kahn (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature 55*(3), 789–865.

Bloom, N., J. Liang, J. Roberts, and Z. J. Ying (2015). Does working from home work? Evidence from a Chinese experiment. *The Quarterly Journal of Economics 130*(1), 165–218.

Boelmann, B., A. Raute, and U. Schönberg (2025). Wind of change? Cultural determinants of maternal labor supply. *American Economic Journal: Applied Economics 17*(2), 41–74.

Bolotnyy, V. and N. Emanuel (2022). Why do women earn less than men? Evidence from bus and train operators. *Journal of Labor Economics 40*(2), 283–323.

Buser, T., M. Niederle, and H. Oosterbeek (2014). Gender, competitiveness, and career choices. *The Quarterly Journal of Economics 129*(3), 1409–1447.

Card, D., A. Mas, and J. Rothstein (2008). Tipping and the dynamics of segregation. *The Quarterly Journal of Economics 123*(1), 177–218.

Chen, H., M. Dong, M. Henry, and I. Sidorov (2025). Occupational segregation in a Roy model with composition preferences. *Games and Economic Behavior 150*, 365–386.

Cortes, P. and J. Pan (2018). Occupation and gender. *The Oxford Handbook of Women and the Economy*, 425–452.

Cortina, L. M. and J. L. Berdahl (2008). Sexual harassment in organizations: A decade of research in review. *Handbook of Organizational Behavior 1*, 469–497.

Dahl, G. B., A. Kotsadam, and D.-O. Rooth (2021). Does integration change gender attitudes? The effect of randomly assigning women to traditionally male teams. *The Quarterly Journal of Economics 136*(2), 987–1030.

Delfino, A. (2024). Breaking gender barriers: Experimental evidence on men in pink-collar jobs. *American Economic Review 114*(6), 1816–1853.

Duncan, O. D. and B. Duncan (1955). A methodological analysis of segregation indexes. *American Sociological Review 20*(2), 210–217.

Folke, O. and J. Rickne (2022). Sexual Harassment and Gender Inequality in the Labor Market*. *The Quarterly Journal of Economics 137*(4), 2163–2212.

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

Goldin, C. and L. F. Katz (2016). A most egalitarian profession: pharmacy and the evolution of a family-friendly occupation. *Journal of Labor Economics 34*(3), 705–746.

Gutek, B. A., A. G. Cohen, and A. M. Konrad (1990). Predicting social-sexual behavior at work: A contact hypothesis. *Academy of Management Journal 33*(3), 560–577.

Kleven, H., C. Landais, and J. E. Søgaard (2019). Children and gender inequality: Evidence from Denmark. *American Economic Journal: Applied Economics 11*(4), 181–209.

Lordan, G. and J.-S. Pischke (2022). Does Rosie like riveting? Male and female occupational choices. *Economica 89*(353), 110–130.

Mas, A. and A. Pallais (2017). Valuing alternative work arrangements. *American Economic Review 107*(12), 3722–59.

Olivetti, C. and B. Petrongolo (2017). The economic consequences of family policies: Lessons from a century of legislation in high-income countries. *Journal of Economic Perspectives 31*(1), 205–230.

Pan, J. (2015). Gender segregation in occupations: The role of tipping and social interactions. *Journal of Labor Economics 33*(2), 365–408.

Patnaik, A. (2019). Reserving time for daddy: The consequences of fathers' quotas. *Journal of Labor Economics 37*(4), 1009–1059.

Rosen, S. (1986). The theory of equalizing differences. *Handbook of Labor Economics 1*, 641–692.

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology 1*(2), 143–186.

Sethi, R. and R. Somanathan (2004). Inequality and segregation. *Journal of Political Economy 112*(6), 1296–1321.

Usui, E. (2015). Occupational gender segregation in an equilibrium search model. *IZA Journal of Labor Economics 4*(1), 13.

Wasserman, M. (2023). Hours constraints, occupational choice, and gender: Evidence from medical residents. *The Review of Economic Studies 90*(3), 1535–1568.

Wiswall, M. and B. Zafar (2018). Preference for the workplace, investment in human capital, and gender. *The Quarterly Journal of Economics 133*(1), 457–507.

Zafar, B. (2013). College major choice and the gender gap. *Journal of Human Resources 48*(3), 545–595.