STAFF REPORTS

NO. 1092
MARCH 2024

# Miss-Allocation: The Value of Workplace Gender Composition and Occupational Segregation

Rachel Schuh

FEDERAL RESERVE BANK *of* NEW YORK

**Miss-Allocation: The Value of Workplace Gender Composition and Occupational Segregation**
Rachel Schuh
*Federal Reserve Bank of New York Staff Reports*, no. 1092
March 2024
https://doi.org/10.59576/sr.1092

## Abstract

I analyze the value workers ascribe to the gender composition of their workplace and the consequences of these valuations for occupational segregation, tipping, and welfare. To elicit these valuations, I survey 9,000 U.S. adults using a hypothetical job choice experiment. This reveals that on average women and men value gender diversity, but these average preferences mask substantial heterogeneity. Older female workers are more likely to value gender homophily. This suggests that gender norms and discrimination, which have declined over time, may help explain some women's desire for homophily. Using these results, I estimate a structural model of occupation choice to assess the influence of gender composition preferences on gender sorting and welfare. I find that workers' composition valuations are not large enough to create tipping points, but they do reduce female employment in male-dominated occupations substantially. Reducing segregation could improve welfare: making all occupations evenly gender balanced improves utility as much as a 0.4 percent wage increase for women and a 1 percent wage increase for men, on average.

JEL classification: J16, J24, J71
Key words: gender, labor, occupation choice

---

Schuh: Federal Reserve Bank of New York (email: rachel.schuh@ny.frb.org). The author thanks her advisors Nicholas Bloom, Isaac Sorkin, Peter Klenow, and Alessandra Voena for advice and support. She also thanks Pascaline Dupas, Muriel Niederle, Caroline Hoxby, Adrien Auclert, Basit Zafar, and seminar participants at Stanford, the All-California Labor Economics Conference, and the Young Economists Symposium for valuable feedback, and Emily Eisner for early collaboration on this project. Varanya Chaubey and Lisa Modifica provided excellent editorial assistance. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship grant no. DGE-1656518. The author also acknowledges financial support from the Stanford Institute for Economic Policy Research George P. Shultz Dissertation Fund and the Stanford University Graduate Research Opportunity Fund.

# I Introduction

Although women's labor market outcomes have improved markedly relative to men's in the last several decades, women and men still do very different jobs. In fact, 40% of men perform occupations where less than 20% of workers are female, 30% of women perform occupations where less than 20% of workers are male, and gender sorting across occupations explains around one-third of the mean gender earnings gap.[1] Occupational gender segregation has been studied in both economics and sociology for several decades, but its causes are still an active area of research.[2] In a standard model of occupation choice, the reasons men and women choose different jobs can be divided into a few broad categories: gender differences in skills, gender differences in preferences for tasks or amenities, discrimination, and social norms.[3] This paper focuses on a particular occupational characteristic that men and women may value differently: the gender composition of an occupation itself.

Preferences for gender composition are particularly interesting because they can amplify existing segregation, create tipping points, and create sorting externalities and thus a role for policy intervention. If women value gender homophily, it will be difficult to desegregate a mostly male occupation without central coordination. Pan (2015) and Henry and Sidorov (2020) show that the level and dynamics of occupational gender segregation in the US are consistent with homophilic gender composition valuations in a Schelling (1971)-style model. Yet there is little evidence that causally identifies and quantifies these valuations. Because the gender composition of an occupation is endogenous and tends to move slowly over time, it is difficult to isolate the value of gender composition separately from other occupational characteristics that vary by gender, like tasks, flexibility, and skill requirements. Delfino (2021) and Larson-Koester (2017) provide evidence that women prefer more female occupations, but whether these preferences are large enough to explain aggregate patterns of segregation is unclear.

In this paper, I causally identify how much men and women value the gender composition of their workplace and quantify the importance of these preferences for occupational gender segregation and welfare. Using a novel online survey experiment, I demonstrate that women and men prefer gender-diverse workplaces on average, but composition preferences have substantial heterogeneity across individuals. I then build a structural model of occupation choice using my survey-estimated composition preferences. Using this model, I show

[1] See Tables C.4 and C.5. Blau and Kahn (2017) and Sloane et al. (2021) estimate similar results.

[2] For a thorough review of the literature, see Cortés and Pan (2018).

[3] Gender differences in skills: see, e.g., Rendall (2018); Yamaguchi (2018); Cortes et al. (2021). Gender differences in preferences: see, e.g., Wiswall and Zafar (2018); Gelblum (2020); Lordan and Pischke (2022). Discrimination: see, e.g., Bertrand and Duflo (2017), Kessler et al. (2019). Social norms: see, e.g., Goldin (2014), Bursztyn and Jensen (2017), Bursztyn et al. (2017), Bertrand et al. (2021).

that the gender composition valuations I estimate are not large enough to produce multiple equilibria and tipping points in gender sorting unless other amenities are endogenous to the gender composition of a job. However, these valuations matter for sorting: if workers did not value the gender composition of their job, female shares would be higher in mostly male occupations. Finally, reducing segregation could improve welfare. I show that if all occupations became 50% female, both women and men would be better off, and a welfare-maximizing social planner would reallocate workers to reduce segregation

The reasons why men and women perform different occupations are often divided into "preferences" and "constraints." In my setting, I hold constraints fixed and measure preferences for gender composition directly. The gender composition of an occupation or workplace is a bundled treatment that may include attributes like expected discrimination or harassment (Folke and Rickne, 2022; Adams-Prassl et al., 2023; Exley and Nielsen, 2024); gender differences in behavior, e.g., competitiveness (Gneezy et al., 2003; Niederle and Vesterlund, 2007); non-promotable tasks (Babcock et al., 2017); the likelihood of friendship or romantic attachments with coworkers; and differences in information sharing within and across genders (Cullen and Perez-Truglia, 2023; Hampole et al., 2022; Bostwick and Weinberg, 2022; Gallen and Wasserman, 2021, 2022; Breen and García-Peñalosa, 2002). Gender differences in desired occupational and workplace gender composition are not necessarily inherent, biological, or constant over time. I ask, then, given the current equilibrium differences in workplace behavior between men and women, how much do workers value gender composition?

I begin my analysis by adapting a simple model of occupation choice based on Pan (2015), Henry and Sidorov (2020), and Schelling (1971) to build intuition for the key parameters driving the effects of gender composition preferences on gender sorting in equilibrium. In the model, male and female workers choose occupations based solely on occupational wages and gender compositions. If workers do not value the gender composition of their occupation, only exogenous gender wage gaps create gender sorting across occupations. If workers value the gender composition of their occupation, homophilic gender composition valuations amplify, and heterophilic gender composition valuations dampen, underlying gender sorting. The key parameters that determine the effect of gender sorting preferences are the willingness-to-pay for a particular level of the female share of an occupation and the wage elasticity of occupational labor supply.

To measure the size and shape of gender composition valuations, I design a hypothetical job choice survey experiment to elicit the willingness-to-pay for workplace gender composition. In the vein of Wiswall and Zafar (2018), Mas and Pallais (2017), and Maestas et al. (2023), in my survey respondents make several choices between pairs of hypothetical job offers that vary randomly in their pay and demographic composition at different workplaces within a fixed occupation and firm. For example, in one question respondents choose between two jobs as a sales associate at two different locations of the same retail store chain. I randomly vary the demographics of each job to address the endogeneity of gender composition. I give choices between workplaces within a fixed occupation and firm to hold constant other characteristics that might vary with (but are not necessarily caused by) the gender composition of a job, including occupational tasks, amenities contracted by the firm, and industry.[4]

I find that both men and women assign non-negligible value to the gender composition of their workplace, but the shape and size of these preferences differ by gender. Women prefer workplaces that are at least half, but not entirely, female. I estimate that women are willing to trade off 4.5% of their wages to avoid a workplace that is all male in favor of a workplace that is evenly gender balanced, but they are not willing to trade off more wages to access a workplace that is more than half female. Men, however, prefer workplaces that are half female and dislike segregation in either direction. I estimate that men are willing to trade off 3% of their wages to avoid a workplace that is all male and 2% of their wages to avoid a workplace that is all female in favor of a gender balanced workplace. By causally identifying the willingness-to-pay for workplace gender composition while allowing for non-linearities in its shape over the spectrum of female shares, these estimates enrich prior findings by Delfino (2021) and Larson-Koester (2017) that suggest women are more likely to choose more female-dominated occupations.

The average estimates of gender composition valuation mask substantial heterogeneity across individuals. To detect heterogeneity that is not correlated with observable characteristics, I estimate a latent-class logit model (Greene and Hensher, 2003) using the survey data. Among both men and women, I find that around half of respondents choose jobs purely based on the wage and assign zero value to the female share. The remaining women prefer majority female jobs and are willing to trade off over 10% of wages to avoid predominantly male workplaces. Among the remaining men, around half prefer majority male jobs, and half prefer majority female jobs.

Older men and women are substantially more likely to value gender homophily in the workplace, suggesting that the desire to segregate by gender may have declined as women's labor market outcomes converged toward men's over the 20th century. For women over 55, the likelihood of valuing a more female workplace increases with age, but for younger

[4] My design is not incentivized and relies on respondents answering questions in a way that reflects their true preferences. This style of survey experiment has been incentivized by offering customized job postings (Kessler et al., 2019; Bustelo et al., 2022) or information on survey results to respondents (Drake, Marshall et al., 2022). In my setting where workers are highly heterogeneous in employment status, occupation, and location, it was not clear that these incentives would be effective, but I have piloted randomly offering respondents job suggestions that depend on their responses.

women the age profile is flat. Although I cannot directly distinguish whether this is an age or cohort effect, a cohort effect tells an interesting story. These younger workers would have joined the labor force during or after the 1990s, when the increase in women's labor force participation began to slow. This suggests that gendered social norms surrounding work and occupation choice, which have likely lessened over time, may be a driver of homophilic gender composition valuations, as suggested by Akerlof and Kranton (2000), Brock and Durlauf (2001), and Pan (2015).[5]

Additional results from my survey suggest that men value female coworkers specifically for social interactions, while women expect many non-wage amenities to be better in a majority female environment. I find that men have a higher willingness-to-pay for a majority female workplace in occupations that require more interaction with coworkers (such as teaching or working in a retail store) than in occupations that rely more on solo work (such as working as an insurance sales agent in an office). I find no such differences among women. In addition, men report they would prefer the coworkers in a mostly female job, but they report they would prefer the work environment, schedule, tasks, and promotion ability in a mostly male job. Women, however, report they would prefer a mostly female job to a mostly male job in all of these attributes. This suggests not only that women value having female coworkers but also that a more female workplace serves as a signal of other workplace amenities women value—for instance, family-friendliness (Goldin and Katz, 2016; Mas and Pallais, 2017; Cortés and Pan, 2020; Morchio and Moser, 2020). Since most of these attributes are held constant in my hypothetical choice experiment, my estimates likely provide a lower bound to the total value women assign to a more female work environment.

I next expand my model of occupation choice to quantify the implications of these gender composition valuations for aggregate gender sorting and welfare in the US. In the model, male and female workers choose occupations to maximize their utility from wages, amenities, gender composition valuations, and random occupational preference draws while wages adjust to labor supply in equilibrium. I estimate the model using my survey results and Current Population Survey data on wages and employment across over 400 occupations. A model is necessary to aggregate my reduced-form results because understanding the implications of composition preferences across multiple occupations requires solving for a sorting equilibrium.

I use this model to perform four exercises. First, I demonstrate that the average gender composition valuations I estimate are not large enough to create multiple sorting equilibria and tipping points in gender segregation. Even for the group of women with the strongest

[5] Fernández et al. (2004), Goldin (2014), Bursztyn et al. (2017), Olivetti et al. (2020), Bertrand et al. (2021), and Cortés et al. (2022), among others, also document the importance of social norms for women and men's career choices, particularly as they relate to marriage and children.

preference for majority-female jobs, composition valuations would need to be seven times larger than those I measure to create a tipping point. However, tipping could occur if other amenities respond endogenously to the gender composition of an occupation. Coupled with my results suggesting that the value of gender homophily may have declined over time, this suggests that homophilic preferences may have contributed to the tipping points in segregation demonstrated by Pan (2015), but changes in amenities and social norms were likely also necessary.

Second, I estimate the aggregate utility cost of segregation by calculating the equivalent variation of eliminating gender segregation across occupations. I find that, on average, both men and women would be better off if their occupation were evenly gender balanced. The average utility gain from making an occupation exactly half female is equivalent to a 0.7% wage increase for women and a 2.8% wage increase for men. The utility gain is even larger for the men and women who most dislike male-dominated occupations.

Next, I measure the aggregate consequences of gender composition valuations for segregation by comparing wages and allocations in reality with those that would result if workers assigned no value to workplace gender composition. I find that absent gender composition valuations, female employment would increase by up to 10% in the most male occupations. I also find that the part of the gender wage gap that is attributable to occupational gender segregation would fall by 10% if workers did not value gender composition.

Finally, I study a social planner's problem in the model and show that reducing segregation would improve welfare modestly. This occurs because workers' valuations of gender composition create a sorting externality: women (or men) do not take into account the effect their entry into an occupation will have on the utility of men (or women) in that occupation. The social planner, then, allocates more women to male-dominated occupations and more men to female-dominated occupations so all occupations are more gender-mixed, which both genders value.

This paper contributes to the literature in three main areas. First, I estimate credibly and causally women's and men's preferences for workplace gender composition. This provides direct quantitative evidence for the causes of tipping points posited by Schelling (1971), Brock and Durlauf (2001), and Pan (2015) and further understanding of the tendency for women to choose more female jobs shown by Delfino (2021), Larson-Koester (2017), and Engel et al. (2022). This also adds to the literature on the importance of non-pecuniary amenities in job choice (Sorkin, 2018; Mas and Pallais, 2017; Wiswall and Zafar, 2018; Bustelo et al., 2022) by showing that the composition of the workplace itself, in addition to amenities provided directly by the firm, may be an important driver of job choice. This has important implications for understanding the causes of inequality across demographic groups. I also show that worker's preferences for gender composition can reduce women's presence in male-dominated jobs and that policies to reduce segregation could improve welfare, which contributes more broadly to the literature on the causes of gender segregation across occupations and the consequences of policies to reduce segregation (Hsieh et al., 2018; Kaplan and Schulhofer-Wohl, 2018; Sloane et al., 2021; Gelblum, 2020; Cortés and Pan, 2018; Cortes et al., 2021).

This paper is organized as follows. In Section II, I outline a basic model of occupation choice with gender composition preferences. I show that composition preferences can act to amplify or dampen existing sorting caused by, for instance, differences in skills or discrimination. In Section III, I discuss the design and results of the survey I administered to estimate gender composition preferences and their distribution across individuals. Section IV brings my survey-estimated preferences into a quantitative model of occupation choice, where I measure the aggregate importance of composition preferences and discuss the utility of alternative allocations. Section V concludes.

## II Occupation Choice Model

To fix concepts, I present a toy model of occupation choice in which workers value both the wage and gender composition of their occupation, in the spirit of Pan (2015) and Henry and Sidorov (2020). I show that underlying patterns of gender sorting will be amplified by homophilic preferences but dampened by heterophilic preferences. In Section III, I measure these preferences directly using a survey experiment.

### II.A Model Environment and Sorting Equilibrium

In my toy model, which is adapted from Henry and Sidorov (2020), male and female workers choose the occupation providing the highest utility from its gender composition and wage. If wage differences across occupations cause one occupation to have a higher female share, homophilic composition preferences cause even more women to choose that occupation. Heterophilic preferences have the opposite effect.

In the model environment, there are two genders of workers, $g \in \{f, m\}$, and two occupations, $k \in \{1, 2\}$. Workers of each gender are present in measure 1. Each occupation is characterized by a pair of gender-specific wages $w_{f,k}$ and $w_{m,k}$ and a female share, $f_o = \frac{l_{k,f}}{(l_{k,f}+l_{k,m})}$, where $l_{k,g}$ is the total number of workers of gender $g$ in occupation $o$. Each gender of worker has some preference profile over the female share of their occupation represented by $h_g(f)$. If preferences are homophilic, $h_g(f)$ is increasing in $f$ for women and decreasing in $f$ for men.

Each individual chooses an occupation to maximize their utility from the occupational wage, the occupational female share, and an idiosyncratic preference draw. The utility function for an individual $i$ of gender $g$ is given by

$$U_i = \max_{k \in \{1,2\}} \left[ \log(w_{g,k}) + h_g(f_k) + \varepsilon_{i,k} \right], \tag{1}$$

where $\varepsilon_{i,k}$ is an idiosyncratic preference shock for individual $i$ in occupation $k$. The idiosyncratic shock is added to the utility function to create dispersion in choices across individuals. I define the function $F_{\varepsilon_2-\varepsilon_1}(x)$ as the CDF of the difference between the preference shocks in occupation 2 and occupation 1.

An individual $i$ will choose occupation 1 if the utility they obtain in occupation 1 is greater than the utility they obtain in occupation 2:

$$log(w_{1,g}) + h_g(\ell_{1,f}/\ell_1) + \varepsilon_{i,1} > log(w_{2,g}) + h_g(\ell_{2,f}/\ell_2) + \varepsilon_{i,2}. \tag{2}$$

Thus, if wage and gender composition utility is higher in occupation 1, only individuals with a large idiosyncratic preference for occupation 2 will choose this occupation, and vice versa.

Following Henry and Sidorov (2020), I define a sorting equilibrium as an allocation of workers of each gender to each occupation such that no small mass of workers could improve their utility by switching occupations. In a sorting equilibrium, the share of workers of each gender in occupation 1 is equal to the probability that utility is higher in occupation 1, or, equivalently, the probability that the relative preference shock in occupation 2 is less than the relative wage and gender composition utility in occupation 1.[6]

$$\ell^*_{1,g} = Pr\left[ log(w_{1,g}) + h_g(\ell_{1,f}/\ell_1) + \varepsilon_1 > log(w_{2,g}) + h_g(\ell_{2,f}/\ell_2) + \varepsilon_2 \right], \qquad g = \{f, m\}. \tag{3}$$

Multiple equilibria will occur when this system has multiple solutions. Thus, tipping, which I define as a shift from one equlibrium to another, will only be possible if multiple equilibria exist. In the next section, I use Equation 3 to understand which conditions create the possibility for tipping and characterize the effect of composition preferences on sorting.

### II.B The Partial Equilibrium Effect of Composition Preferences

In this section, I use the model outlined above to evaluate which parameters determine the effect of composition preferences on gender sorting when wages are fixed exogenously. I find that the effect of composition preferences on sorting and wages will be larger, and tipping more likely, if gender composition preferences are strong relative to wages and if occupation choice is more responsive to changing wages.

[6] Henry and Sidorov (2020) establish conditions for existence and uniqueness of a sorting equilibrium: essentially, an equilibrium generally exists and will be unique if composition preferences are smooth and of small enough scale.

We can invert Equation 3 to find the labor supply function:

$$\log(w_{1g}) - \log(w_{2g}) = \underbrace{F^{-1}_{\varepsilon_2-\varepsilon_1}(\ell_{1g})}_{\text{marginal preference shock}} + \underbrace{\left[h_g\left(\frac{1-\ell_{1f}}{2-\ell_{1f}-\ell_{1m}}\right) - h_g\left(\frac{\ell_{1f}}{\ell_{1f}+\ell_{1m}}\right)\right]}_{\text{difference in composition utility}}.$$

To understand this labor supply function more concretely, I assume that the idiosyncratic preference draw follows a Type I Extreme Value (TIEV) distribution with shape parameter $\lambda$. The parameter $\lambda$ determines how responsive workers are to changing occupational conditions: if $\lambda$ is large, the variance of the preference shocks is high, which makes occupation choices relatively sticky and workers less responsive to changing occupational conditions. If $\lambda$ is small, the variance of the preference shock is low, which makes occupation choices relatively flexible and workers more responsive to changing occupational conditions.

Under this assumption, we can solve the inverse CDF of the preference shock to write the following labor supply function:

$$\underbrace{[\log(w_{1g}) - \log(w_{2g})]}_{\text{wage utility diff.}} = \underbrace{\lambda \cdot \log\left(\frac{\ell_{1g}}{1-\ell_{1g}}\right)}_{\text{pref shock utility diff.}} + \underbrace{\left[h_g\left(\frac{1-\ell_{1f}}{2-\ell_{1f}-\ell_{1m}}\right) - h_g\left(\frac{\ell_{1f}}{\ell_{1f}+\ell_{1m}}\right)\right]}_{\text{composition utility diff.}}, \quad g = f, m. \tag{4}$$

The first part of Equation 4 is standard: as the relative wage in occupation 1 increases, the relative labor supply to occupation 1 increases; the labor supply is more responsive to wages when the variance of the preference shock is smaller. The second part of this equation illustrates that labor supply will also depend on the relative composition utility of the two occupations: if composition utility is higher in occupation 1, more workers of gender $g$ will choose occupation 1 than they would in the absence of composition preferences.

Homophilic composition preferences will exacerbate gender sorting and reduce female wages in more female jobs. This is illustrated in Figure 1, which plots the female labor supply function with and without composition preferences given a fixed male labor supply to occupation 1. The solid line shows the labor supply function with no gender composition preference, i.e. $h_g(f) = 0$. The dashed line shows the labor supply function with homophilic composition preferences. This figure illustrates the two key effects of homophilic composition preferences. First, when female wages are relatively higher in occupation 1, at a fixed wage difference, homophilic preferences will induce more women to choose occupation 1 than otherwise would. This is illustrated by the fact that point B is to the right of point A.

Second, if occupation 1 is more than half female, the wage required to attract a fixed number of women to occupation 1 will be lower in the presence of homophilic composition preferences than without composition preferences. This is illustrated by the fact that point C is below point A.

The effect of composition preferences on the sorting equilibrium will be larger when composition preferences are steeper relative to preference shocks and wage differences. As the composition preference function increases in slope, the labor supply function will become flatter. This is illustrated in Figure 2 Panel A. Stronger composition preferences will cause the equilibrium to shift to an even higher level of female employment in occupation 1 than would occur with weaker composition preferences given a fixed wage difference.

The effect of composition preferences will also be larger when workers are more likely to switch occupations; that is, when the variance of the preference shock is smaller. In Figure 2 Panel B, I plot the labor supply function with and without composition preferences for two values of the TIEV preference shock variance $\lambda$. When $\lambda$ is relatively large, workers' occupation choices are less responsive to changing wages, so the labor supply function is steeper both with and without composition preferences (grey lines). When $\lambda$ is relatively small, workers' occupation choices are more responsive to changing wages, so the labor supply function is flatter. This means that with the same composition preference function (normalized to wages, i.e., equal willingness-to-pay for different gender compositions), the effect of composition preferences on wages and allocations will be larger when the variance of the preference shock is smaller.

When are tipping points possible in this model? I define a tipping point as a rapid shift between multiple sorting equilibria, either from a gender mixed to a gender segregated equilibrium or vice versa. In order for tipping to occur, either the male or female labor supply function must bend backward at some point so that the same relative wages can result in multiple possible allocations. This is illustrated in Figure 3. When the female labor supply function is backward bending, a change in relative wages (possibly caused by a shift in demand) or a change in male labor supply can remove or create a mixed gender equilibrium.

Thus, tipping will be more likely when either gender composition preferences are very strong or occupation choices are very flexible, because both of these conditions will create a flatter labor supply function. Therefore, the key parameters that determine the effect of composition preferences are the slope of the composition preference function relative to wages and the variance of the occupational preference shocks.

## II.C General Equilibrium Extension

In this section, I show that the effect of composition valuations will be dampened, but not nullified, when wages adjust to changing allocations in general equilibrium.

To allow for equilibrium wage adjustments, I close the model with a production function that is CES across occupations:

$$Y = \left( (A_1 \ell_1)^{\frac{\nu-1}{\nu}} + (A_2 \ell_2)^{\frac{\nu-1}{\nu}} \right)^{\frac{\nu}{\nu-1}}.$$

The two occupations are gross substitutes when $\nu > 1$, and gross complements when $\nu < 1$. $A_k$ is occupation-specific productivity.

To account for gender wage gaps within occupations, I assume that within occupations, the production function is CES across genders, as in Ngai and Petrongolo (2017):

$$\ell_k = \left( q_k (\ell_{k,m})^{\frac{\alpha-1}{\alpha}} + (1 - q_k)(\ell_{k,f})^{\frac{\alpha-1}{\alpha}} \right)^{\frac{\alpha}{\alpha-1}},$$

where $\alpha$ is the elasticity of substitution across genders, and $q_k$ represents the relative productivity of men. If $q_k > .5$, men are more productive in the occupation, and if $q_k < .5$, women are more productive. This wedge between male and female productivity can encompass both actual productivity differences (e.g., women are less productive in tasks requiring physical strength or work fewer hours) and pure wage discrimination.

Taking the first-order conditions of the profit function gives the following wage equations

$$w_{k,m} = q_k \left( \frac{\ell_k}{\ell_{k,m}} \right)^{1/\alpha} \frac{Y^{\frac{1}{\nu}} A_k^{\frac{\nu-1}{\nu}} p}{\ell_k^{1/\nu}}, \qquad w_{k,f} = (1 - q_k) \left( \frac{\ell_k}{\ell_{k,f}} \right)^{1/\alpha} \frac{Y^{\frac{1}{\nu}} A_k^{\frac{\nu-1}{\nu}} p}{\ell_k^{1/\nu}} \quad k = 1, 2,$$

where $p$ is the final good price. The wage for each gender increases in own-gender occupational productivity ($q_k$ for men and $(1-q_k)$ for women) and overall occupational productivity $A_k$ and decreases in the own-gender share of labor.

Allowing for equilibrium wage adjustments will dampen the effect of composition valuations on sorting and make tipping less likely. This is illustrated in Figure 4, which adds an equilibrium relative occupational demand curve to the previous labor supply plots. As more women enter an occupation, their relative wage in that occupation will fall, dampening the effect of composition preferences. It is notable, however, that accounting for equilibrium wage adjustments does not erase the effects of gender composition preferences, because the wage does not fully internalize the effect of a worker's entry into an occupation on the utility of the other gender.

## III Measuring the Value of Workplace Gender Composition

To measure the amenity value of workplace gender composition, I design and administer a survey with an embedded hypothetical job choice experiment to causally identify the amount respondents are willing to pay to be in a job with their preferred female share. This is meant to simulate an ideal experiment in which occupational gender compositions are changed at random. I find that on average both women and men value gender diversity, but within both genders there is substantial individual preference heterogeneity. Women are willing to trade off 4.5% of their wages to avoid an all-male workplace in favor of a gender-balanced workplace, but they are not willing to trade off additional wages to be in a majority female workplace rather than a gender-balanced workplace. Men are willing to trade off 2-3% of their wages to avoid both all-female and all-male workplaces in favor of a gender-balanced workplace. In Section IV, I use these estimates in a quantitative model to determine their consequences for gender sorting in the aggregate.

### III.A Survey Design

The main section of my survey is a hypothetical job choice conjoint experiment in which respondents choose between several pairs of job offers which vary randomly in pay and workplace demographic composition. Through this design, I can causally identify the respondents' willingness-to-pay for the female share of a job.

#### III.A.1 Hypothetical Job Choice Survey Experiment

My survey instrument is designed to elicit the willingness-to-pay for the gender composition of one's workplace within a fixed occupation and firm. I focus on a choice between workplaces, rather than occupations, for three main reasons. First, the estimand of interest is the effect of gender composition on utility within an occupation, which requires holding the occupation itself fixed. Second, a choice in an online survey can more easily resemble a choice between workplaces than a choice between occupations, which may involve long-term human capital investments. Finally, this design holds fixed other amenities that may on average vary with the gender composition of an occupation, but are not necessarily causally influenced by gender composition.

In each hypothetical choice, respondents choose between two jobs which are in the same occupation but are located at two different workplaces which vary in their wages and demographics. For each pair of jobs, respondents are instructed that in each job they will be performing the same occupation at different locations of the same firm, which are a similar distance from their home. They are asked to select the job they would be more likely to choose if offered both within each pair. For example, respondents may see the following instructions:

> *You are choosing between two jobs as a sales associate at a retail store. Both stores are locations of the same chain and are a similar distance from your home. Please select the store at which you would prefer to work.*[7]

For each individual job, respondents see the pay (hourly and annual equivalent), hours (for this section, all are full-time), and three characteristics of the demographic composition of their coworkers: gender (share female or share male); age (share under 40); and parental status (share with children). Respondents see nine to eleven of these choices, within five occupations: high school teacher, insurance sales agent, retail sales associate, nurse, and software developer.[8]

One reason the survey presents choices between workplaces, rather than occupations, is that the choice over a specific job rather than an entire occupation is easier to simulate in an online survey setting. Conjoint designs for hypothetical job choices are used frequently in the labor literature, and choices in these surveys have been shown to be correlated with realized job choices (Wiswall and Zafar, 2018; Mas and Pallais, 2017; Maestas et al., 2023; Folke and Rickne, 2022).[9] A choice in an online survey between two specific job offers closely resembles a choice applicants make in real life when they, for instance, select which jobs to apply to on a website. However, it is more difficult to replicate a choice between occupations which may involve retraining or other long-term human capital investments.

A choice between workplaces within a specific occupation and firm also holds fixed many amenities of a job that might vary with the female share. These may include formally contracted amenities, like schedule flexibility or paid parental leave, and job tasks. In the ideal experiment, in which the gender composition of an occupation changes randomly and rapidly, these characteristics may not change in the short run. On the whole, the workplace choice presents a cleaner exercise that better isolates the amenity value of gender composition itself rather than correlated attributes.

To incorporate the results from the survey of workplace choices into a model of occu-

[7] A fully formatted example of a question respondents see in this section is included in Appendix A, Figure A.1.

[8] The first three occupations are chosen because they are gender-balanced in reality and thus realistically could vary widely in workplace gender composition. The last two occupations are chosen to test for heterogeneity based on actual occupational gender composition because they are very segregated in reality.

[9] Gender composition differs from other amenities whose values are elicited in conjoint surveys because it is unlikely to be communicated directly to applicants. This is similar to the prevalence of workplace sexual harassment, which Folke and Rickne (2022) elicit valuations for in a conjoint design that includes vignettes about employee experiences of sexual harassment. In both cases, although the job characteristic is unlikely to be directly advertised by the firm, it may be conveyed to the applicant through informal channels such as conversations with current or past employees.

pation choice, we must understand how the gender composition of an occupation is related to the gender composition of a workplace. At the extreme end, if there is no gender sorting across firms within an occupation, workplace composition preferences won't matter at all for occupation choice because workers can choose the firm with their preferred gender composition within any occupation. However, this is unlikely to be the case. Hellerstein et al. (2008) estimate that as of 2000, 61% of gender segregation across firms was explained by gender segregation across occupations, which suggests that adjusting for the difference between occupations and firms would undo about 40% of the workplace-level valuations. In addition, I consider a static model of occupation choice, and when initial human capital investments are made, it is much more likely that individuals know the approximate gender composition of an occupation than the distribution of firm gender compositions within occupations. Finally, even within a firm, one's closest coworkers are most likely those who perform the same occupation, so differences in segregation across firms do not necessarily undo segregation across occupations.

### III.A.2 Variations in Job Choice Experiment

I include several variations on the hypothetical job choice conjoint across iterations of the survey to assess the robustness of my results and the reasons respondents value workplace gender composition.

In the first variation of the survey, respondents only see the pay and demographic composition of each job and no other information. This design eliminates possible effects of prior perceptions of specific occupations, which may affect the fixed-occupation workplace choice described above. It also allows for social norms to affect these job choices: without listing a specific occupation, respondents' perception of the societal view of the job may depend on the reported gender composition. However, with such a simple design it is also possible that respondents infer characteristics of the job that I do not intend to evaluate, such as job tasks and contracted amenities.

In the second variation, respondents again do not see a specific occupation but also see additional amenities of each job. This design allows me to analyze the mechanisms for the measured gender composition valuations by holding fixed other amenities that might vary with occupational gender composition. I separately control for two additional amenities: the probability of promotion or firing, and the job hours and schedule flexibility.

In the final variation, respondents choose between two specific occupations and I randomly vary whether they see the true gender composition of the occupation. In this setting, I can evaluate whether respondents still consider gender composition when they are choosing between two different occupations, which more closely matches the model exercise in Section

IV. This survey exercise is slightly more complex to analyze because the effect of information on the female share of an occupation will depend on the respondent's prior perception of the female share. Thus, in this exercise I do not include wage information, and instead I use this as a qualitative robustness check on the more standard conjoint designs.

#### III.A.3 Remaining Survey Structure

The survey is structured as follows. First, respondents answer a standard set of demographic questions regarding their age, gender, education, marital status, and family structure. I then ask respondents about the occupation, industry, hours, and wages of their current or most recent job. The hypothetical job choice conjoint experiment follows. After the conjoint experiment, I ask respondents to report their expectations about how their experience in a mostly female occupation and a mostly male occupation might differ. The purpose of this section is to disentangle the possible reasons for gender composition preferences. Finally, I ask respondents for their opinions on three gender attitudes using questions adapted from the General Social Survey and Pew American Trends Panel.[10]

### III.B Survey Data

I fielded my survey online to a sample of 8,850 US adults in multiple waves from October 2021 to October 2022. I designed the survey in Qualtrics and recruited participants using Lucid Theorem, a service that connects survey participants with academic researchers.[11]

An online survey is useful because it provides a controlled setting in which to run randomized experiments at a low cost, but there are several potential drawbacks. The main possible problems with an online survey are non-representative samples, inattentive respondents, and lack of external relevance. In this section, I discuss the first two issues and address their relevance in my setting. I address external relevance in Section III.D.5. Reassuringly, I find that my sample is overall demographically similarly to the US population and mostly quite attentive.

My survey sample is broadly similar to the population-representative March CPS, as

[10]The full survey is available upon request.

[11]Lucid, like many other online survey panel providers used by academic researchers, sources participants from multiple companies that recruit individuals to take online surveys. In contrast with a service like MTurk, I do not directly solicit participants; rather, I send an order for participants to Lucid, which funnels my order to a third-party survey respondent recruiter. Respondents may have been recruited through emails, push notifications, in-app pop-ups, or through sites offering multiple survey opportunities. (See `https://lucidtheorem.com/faq`.) Coppock and McClellan (2019) compare results from several survey experiments conducted through Lucid, MTurk, and probability samples. For all but one of five experiments, they find experimental effects that matched the sign and significance of the original estimates on the probability sample.

shown in Table 1.[12] My survey respondents are relatively younger and more educated than the US population, but there is sufficient presence of the under-represented groups to enable re-weighting.[13] In Appendix A, I compare my survey respondents to the CPS ASEC respondents on other characteristics that one might expect to vary with gender composition preferences, including income, realized occupational gender composition, and attitudes toward gender and work. I find that my sample is similar to the CPS on these characteristics.

To ensure my data is of high quality, I include several attention checks within my survey and limit my analysis to the subset of respondents who pass these attention checks. Overall, 78% of respondents pass all attention checks.[14] Most importantly, I include an attention check within the hypothetical job choice conjoint, where respondents see a choice between two jobs that differ only in their wages. As we see in Figure A.15, 85% of respondents choose the higher wage job. This is a similar rate of inattention to that found by Mas and Pallais (2017), who found that 85.5% of individuals chose the higher-paying job given a choice between two otherwise identical jobs to apply for. It is reassuring that the rates of inattention are similar, given Mas and Pallais (2017) had an incentivized choice between actual jobs and my survey is purely hypothetical and not incentivized. Results of my survey are similar whether or not I include respondents who failed one or more attention checks.

I also check that respondents' answers are consistent across questions within the hypothetical job choice module and find that in the vast majority of cases they are. First, I check that if respondents see the exact same pair of jobs twice, they make the same choice in both scenarios. As shown in Appendix A Table A.1, in 87% of such instances respondents choose the same job in both cases. Because it is fairly rare for respondents to see the exact same choice twice (it only happens 60 times), I also check that if respondents see the same pair of female shares with possibly differing wages, they either choose the same female share in both cases or make choices that represent internally consistent bounds on their willingness-to-pay for a particular female share. As detailed in Appendix A Table A.2, I find that respondents make internally consistent choices in 98% of these cases. My final estimation sample for the fixed-occupation hypothetical choice consists of 2,772 individuals and 29,972 total choices.

[12]CPS ASEC data accessed via Flood et al. (2018). I pool data from 2014 through 2019 and limit to individuals aged 18 and up.

[13]Among my survey respondents, Black people are slightly over-represented and Hispanic under-represented relative to the CPS. However, this may be due the fact that Hispanic ethnicity is a separate question in the CPS, whereas I ask for race and ethnicity in one question.

[14]In Appendix A, I provide further details on the attention checks and compare the demographics of attention check passers to all survey respondents.

### III.C Estimation Strategy

Given the randomized design of the hypothetical job choice experiment, I can use a simple conditional logit model to estimate the willingness-to-pay for the female share of a job. I also estimate a latent class logit model to determine the degree of preference heterogeneity, which allows me to determine the share of the population for which gender composition preferences will be an important determinant of job choice.

#### III.C.1 Average Preferences by Gender

To estimate the average willingness-to-pay (WTP) for the gender composition of an occupation, I use a standard multinomial logit model. Similar to Equation 1, the utility of an individual choosing among $J$ jobs indexed by $j$ is

$$U_{i,J} = \max_{j \in J} \left[ \beta_w \log(w_{g,j}) + h_g(f_j) + \kappa X_j + \varepsilon_{i,j} \right],$$

where $X_j$ is a vector of additional job attributes. To estimate the composition preference function $h_g(f)$, I take a discrete approximation

$$h_g(f) \approx \sum_{f_n = 0,.1,\ldots,.9,1} \beta_{f_n} \mathbb{1}(f_n = f), \tag{5}$$

where $f_n = 0, .1, \ldots, .9, 1$ are the eleven (minus one) possibilities for the female share in the survey. As in the model in Section II, I assume the preference shocks are Type I Extreme Value with shape parameter $\lambda$. Then, the probability that individual $i$ chooses alternative $j$ among the choices in $J$ takes the conditional logit form

$$P_{ij,J} = \frac{\exp((\beta_w \log(w_{g,j}) + \sum_{f_n} \beta_{f_n} \mathbb{1}(f_n = f_j) + \kappa X_j)/\lambda)}{\sum_{j' \in J} \exp((\beta_w \log(w_{g,j'}) + \sum_{f_n} \beta_{f_n} \mathbb{1}(f_n = f_{j'}) + \kappa X_{j'})/\lambda)} \tag{6}$$

Then, given each respondent's choices and available alternatives, I can estimate the utility parameters using a logit regression.

Importantly, the "pure" wage elasticity $\beta_w$ and gender composition elasticities $\beta_{f_n}$, that is, the responsiveness to wages and gender composition holding the preference draw fixed, are not level identified separately from the variance of the preference draw $\lambda$. The parameters I identify are $\hat{\beta_w} = \frac{\beta_w}{\lambda}$ and $\hat{\beta_f} = \frac{\beta_f}{\lambda}$. For my willingness-to-pay estimates this is not important because the variance of the preference draw is differenced out when I calculate the elasticity to gender composition relative to the wage elasticity. In the quantitative model in Section IV, I will normalize the pure wage elasticity to one, use the willingness-to-pay to measure composition preferences, and calibrate the preference draw variance.

Given this estimation, I can calculate the willingness-to-pay for a given gender composition relative to a baseline gender composition by taking the exponential of the ratio of the estimated gender composition coefficient to the wage coefficient:[15]

$$WTP_f = 1 - \exp\left(\frac{-\hat{\beta}_f}{\hat{\beta}_w}\right). \tag{7}$$

The WTP is positive for a good amenity and negative for a bad amenity: if an amenity is good, the worker will be willing to accept a lower wage in exchange for access to the amenity, and if an amenity is bad, the worker will need a higher wage in exchange for accepting the amenity.

### III.C.2 Heterogeneous Preferences

The logit model described above allows me to estimate average demographic valuations within each gender. However, these estimates may mask substantial heterogeneity. Understanding heterogeneity can both improve our understanding of what underlying attributes contribute to gender composition valuations and inform firm- and economy-wide policies to reduce segregation.

I estimate heterogenous composition valuations using a latent class logit model that allows me to detect heterogeneity that may not correlate with observable characteristics, following Greene and Hensher (2003). In this model, rather than assuming that all individuals of a given gender have the same preference parameters, I assume there are $Q$ preference types within each gender indexed by $q$. Then, the probability that an individual of class $q$ selects alternative $j$ in choice $t$ is

$$P[\text{alternative } j \text{ chosen by individual } i \text{ in choice } J| \text{ class } q] = \frac{\exp(x'_{iJ,j}\beta_q)}{\sum_{j=1}^{J}\exp(x'_{iJ,j}\beta_q)} = P_{ij,J|q}, \tag{8}$$

where $x_{iJ,j}$ are the characteristics of choice $j$ in situation $J$ and $\beta_q$ is a vector of class-specific preference parameters. This means that the probability that an individual of type $q$ makes some choice is fixed within class $q$ but varies across classes. I estimate this model using maximum likelihood, further details of which are provided in Appendix B.2.

I use a latent class model with a discrete number of types rather than estimating preferences at the individual level for two reasons. First, estimating preferences across the entire spectrum of gender compositions would require observing a large number of choices per per-

[15] Details on the steps to this willingness-to-pay formula are described in B.1.

son.[16] In the context of an online survey, respondent fatigue and inattention are relevant concerns, so asking respondents to make twenty or more choices of the same format is not desirable. Second, inputting preferences into a numerical model requires a discrete number of preference types, so discretization is required regardless. Thus, although I assess individual-level preference estimates to check validity and robustness, I rely on a discrete number of classes for most heterogeneity analysis.

### III.D Survey Results

I find that, on average, women are willing to trade off 4.5% of their wages to avoid an all-male workplace, but they do not value additional women once a job is at least 50% female. Men are willing to trade off 2% of their wages to avoid an all-female workplace and 3% of their wages to avoid an all-male workplace. However, these average estimates belie substantial heterogeneity.

#### III.D.1 Average Preference Results

Women and men both prefer gender-mixed workplaces to majority male or all female workplaces, but the shape and scale of these preferences differ by gender. Figure 5 shows the estimated WTPs for each possible female share separately for men and women, and Table 2 shows the estimated willingness-to-pay for each gender, age, and education share by respondent gender, age, and education.[17] In the analysis that follows, I use a 50% female workplace as the baseline so the WTP for a 50% female job is always zero, and other WTPs are relative to a 50% female job.

Women, on average, prefer workplaces that are at least half female to majority male workplaces, but they do not desire complete gender segregation, as shown in red in Figure 5. Women's composition valuations can be split into three regions. For workplaces that are less than half female, women have a negative willingness-to-pay, meaning they dislike mostly male workplaces. Women are willing to pay 4.5% of wages to avoid an all-male workplace. For workplaces that are 50-90% female, women's willingness-to-pay is approximately zero, meaning that there is no additional utility for having more women once a job is at least half female. Finally, for workplaces above 90% female, women have a small negative willingness-to-pay, suggesting that complete segregation is not desirable. Overall, women dislike being a gender minority, but do not want complete gender segregation.

[16] Drake, Marshall et al. (2022) develop a procedure to elicit individual-level WTPs using a Bayesian Adaptive Choice Experiment in which the questions that respondents see depend on their prior answers in such a way to maximize the information in each answer. I did not implement this method due to computational limitations, but this provides a promising avenue for further study.

[17] I show all coefficient estimates in Appendix Table B.1. Notably, the coefficient on the log wages for men and women is nearly identical.

Men have a strikingly similar composition valuation profile to women: they prefer gender-mixed to majority male workplaces. However, they have slightly more distaste for majority female workplaces than women do. For both workplaces that are less than 30 percent male and more than 70 percent female, men have a negative willingness-to-pay, meaning they dislike both predominantly female and predominantly male workplaces. Among workplaces that are 40-60% female, men appear largely indifferent. On the whole, men dislike gender-segregated workplaces with either gender in a large majority and prefer gender-mixed workplaces. Notably, though, they have greater distaste for majority male than majority female workplaces.

My WTP estimates enrich prior research suggesting women are more likely to choose more female jobs or workplaces by pinpointing their particularly strong distaste for predominantly male workplaces. They also complement prior estimates suggesting that men do not value gender composition by clarifying that men prefer gender diversity, but are largely indifferent to the exact composition in workplaces that are fairly gender-balanced. In particular, Wiswall and Zafar (2018) estimate that women have a small positive WTP for more female workplaces for jobs in the range of 26-72% female, and Delfino (2021) finds that women are more likely to apply for a social work job when the advertisement contains a female rather than a male picture. [18] Larson-Koester (2017) also estimates that women value more female occupations and men do not value gender composition using observational data on occupation choice.[19] Folke and Rickne (2022) find a 10% WTP to avoid a job with a reported incident of sexual harassment, and an increasing likelihood of sexual harassment for workers who do not match the gender majority of their firm, corroborating my results for women.

Relative to estimates of valuations for other job characteristics in the literature, the 4.5% WTP I find among women to avoid an all-male workplace is about half as large as Mas and Pallais (2017)'s estimate of the WTP for the ability to work from home (7-10%)[20] and much smaller than their estimate of the WTP to avoid employer discretion in scheduling (around

[18] Wiswall and Zafar (2018), using a similar survey, estimate an average WTP of -.08% among men and -.04% among women for the percent of men at a job, but the estimates are not significantly different from zero. If we multiply their female WTP by fifty for a move from a 70% to 20% male job, the resulting 2% is actually quite similar to my estimate of 3% for women's WTP for a similar move. For men, their estimate would translate into a 4% WTP for the same move compared to about 2% in my sample. Notably, the male shares in their survey range from only 26% to 72%, so we cannot extrapolate these estimates to the entire range I measure. Delfino (2021) finds that showing applicants a picture of a male worker for a social work recruitment advertisement has no effect on men's likelihood of applying but decreases women's likelihood of applying, which fits with women's desire to avoid mostly male workplaces and men's relatively symmetric preferences.

[19] Larson-Koester (2017) estimates a discrete choice model of occupation choice using data from the SIPP, ACS, and CPS and finds that women have increasing but concave utility in the female share of an occupation and men have no significant preference over the female share of an occupation.

[20] Notably, this experiment was run before the COVID-19 pandemic.

25-30%). Relative to the WTPs for other demographic characteristics I measure, the WTPs for gender composition are similar or larger.

Overall, my results suggest that both men and women care about the gender composition of their job, but their valuations are not symmetric. Women prefer jobs that are at least half, but not all female, indicating that women primarily have a distaste for the most male jobs. Men, on the other hand, most prefer gender-mixed jobs and dislike jobs that are very gender segregated in either direction. Even without formally aggregating these preferences using the structural model, these estimates tell us that women's average composition preferences may help explain gender segregation because women have mostly homophilic valuations, but men's average composition valuations will not help explain segregation because men prefer gender-mixed jobs.

### III.D.2 Heterogeneous Valuation Results

I estimate a latent class logit model to detect underlying heterogeneity in gender composition valuations and find that the average preference estimates among women are a mix of one type of respondent that has large valuations for more female jobs and one type that does not value the female share of a job at all. The average preference for diversity among men is a mix of men who do not value the gender composition of a job, men who prefer more female jobs, and men who prefer more male jobs.

I select the number of classes in the latent class logit model using cross-validation.[21] Through these exercises, I choose $Q = 2$ classes for women and $Q = 3$ classes for men.

I find that among women, there is a class that only values higher wages and a class that prefers more female workplaces, as illustrated in Figure 6. The wage-preferring class, shown in gold, has a large coefficient on the wage but a close to zero coefficient on all demographic characteristics. The female-preferring class, shown in red, also tends to choose higher wage jobs but has a large willingness-to-pay to avoid the most male jobs.[22]

The wage-preferring class, who value higher wages but assign little to no value to demographic composition, constitute 52% of women in my sample. This class has an extremely large coefficient (181) on the log wage, implying a member of this class would have over a 99.9% probability of choosing a job where the wage is 10% higher, all else equal. This class also has marginally negative, but small, WTPs for jobs less than 20% female. The WTPs

[21]See Appendix A for details. I run a 10-fold cross validation exercise on the log-likelihood for models with 1-10 classes. I also evaluate the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) for each number of classes. The details of these exercises are in Appendix A, Figures B.1 and B.2. I also assess the estimated preference classes for 1-5 classes, and among those ranked highly by the cross-validation and AIC/BIC, select a number of classes where the WTPs across classes are clearly distinguishable. Appendix A Figures B.3 and B.4 show the estimated WTPs for models with 1-5 classes.

[22]Estimates of all coefficients and WTPs are included in Appendix A Tables B.2 and B.3.

for other demographic characteristics in this class are small and not significantly different from zero. All of this suggests that the wage-preferring class is not particularly motivated by gender composition.

The female-preferring class, who assign large value to majority female workplaces, make up the remaining 48% of my female sample.[23] This class has large, negative WTPs for all female shares below 50%. This class also has a slight preference for younger workplaces where more coworkers have children. Members of this class still prefer higher wage jobs, but are willing to trade off a significant fraction of their wages for jobs that have more women.

Overall, the WTP estimates in Figure 6 and Table B.3 suggest that the latent class logit picks up on important heterogeneity in gender composition preferences among female respondents. The negative average WTP for mostly male workplaces among women shown in Figure 5 is in fact an amalgam of the WTPs of two groups of women: one that highly values more female jobs and one that does not value more female jobs.[24]

I find that among men, there is a class that only values higher wages, a class that prefers more female workplaces, and a class that prefers more male workplaces, as illustrated in Figure 7. As among women, the wage-preferring class, shown in gold, has a large coefficient on the wage but a close to zero coefficient on all demographic characteristics. The female-preferring class, shown in red, dislikes mostly male workplaces but is mostly indifferent between workplaces over 50% female. Finally, the male-preferring class, shown in light blue, dislikes mostly female workplaces but is mostly indifferent between workplaces under 50% female.

Similarly to the female sample, I estimate that 52% of the male sample comes from a wage-preferring class. The male wage-preferring class again has a large coefficient (181) on the log wage, which is almost identical to the log wage coefficient for the female wage-preferring class. This class also has a small negative WTP for 0-10% female workplaces, but the WTPs for all other demographic characteristics are small and not significantly different from zero. As in the female sample, this class seems to be motivated mostly if not entirely by higher wages.

In contrast to the female sample, in the male sample, the group with significant preferences for demographic composition comes from two distinct classes. I denote these as the female-preferring class, which makes up 29% of the male sample, and the male-preferring class, which makes up 19% of the male sample.

[23] This class has a coefficient on the wage of 15, which implies that given a choice between two jobs, one of which has a 10% lower wage, a woman from the female-preferring class will choose the higher wage job with a probability of 81%.

[24] See Table B.5 and B.4 for all WTP and coefficient estimates by class.

The female-preferring class among men assigns the most value to workplaces that are 50-90% female. For this class, the willingess-to-pay to avoid an all-male workplace is around 15% of wages. Members of this class still prefer higher wage jobs, but are willing to trade off a significant fraction of their wages for jobs that have more women.

Finally, the male-preferring class most prefers gender-mixed to male-dominated workplaces and dislikes majority female workplaces. This class has a willingness-to-pay of around 10% to avoid an all female workplace. As shown in Table B.4, the coefficient on the log wage for this class is only 9.295, implying a member of this class would choose a job with a 10% higher wage with a 71% probability. This causes the estimated WTPs to be noisier and suggests this class may be slightly less attentive.

As a whole, the latent class logit also picks up on important gender composition preference heterogeneity among men. The average hump-shaped preference profile for men, which suggested an average preference for gender-mixed jobs, is in fact an amalgam of wage-preferring, male-preferring, and female-preferring classes. Rather than preferring gender diversity, then, the latent class estimates suggest that men also prefer gender sorting to some extent but they do not all prefer to sort with their own gender. Notably, however, all preference classes are happy in a gender-balanced job.

#### III.D.3 Covariates of Valuation Heterogeneity

I find that older men and women are more likely to value gender homophily, suggesting that the value of gender sorting may have declined over time as women's participation and labor market options have increased. However, most heterogeneity in composition valuations is not explained by observable characteristics.

Older women are significantly more likely to belong to the female-preferring class, while older men are somewhat more likely to belong to the male-preferring class. Figure 8 plots the posterior probability of belonging to the class that prefers own-gender workplaces by age separately for women and men. Women over 55 are more likely than women under 55 to belong to the female-preferring class, and this probability continues to rise for older ages. Older men are also somewhat more likely to belong to the class that prefers more male workplaces, although the pattern is less clear than that for women.

If we assume that these age patterns are at least in part due to cohort effects, the age profiles suggest that composition valuations of men and women may be converging over time, matching a broader trend of gender convergence in the labor market. Notably, the age profile of women's composition values mirrors the time trend of women's labor force participation, which rose through the 20th century before flattening out in the mid-1990s. Women under 55 were likely to enter the labor market during or after the mid-1990s, when women's labor force participation began to flatten. One reason that older women are more likely to prefer gender homophily may be that these women were more likely to experience gender discrimination in the labor market. Similarly, older men may be more likely to subscribe to more traditional gender norms. However, in this cross-sectional data it is not possible to distinguish cohort from age effects, so some caution is required in interpreting these results.

I also find that younger, unmarried men are more likely to belong to the class that prefers more female workplaces. This suggests that marriage market matching could be a motivation for men who prefer more female workplaces, because they might be more likely to meet potential romantic partners at a workplace with more women. This could also be evidence that younger men with more progressive gender attitudes prefer to work around more women. All demographic covariates of the valuation classes are summarized in Appendix Figures B.5 and B.6.

### III.D.4 Channels for Composition Valuation

Evidence from additional survey questions and robustness checks suggests that men value mixed to more female workplaces primarily for the coworker interactions, while women value many aspects of more female workplaces.

Men are more likely to choose more female workplaces when an occupation involves more direct interaction with coworkers, suggesting coworker interactions are an important driver of men's composition valuations. I can detect this heterogeneity by estimating WTPs separately for the specific occupations used in the hypothetical workplace choice. Figures B.14 and B.15 illustrate this difference. Men value female workplaces more when the occupation in question is "retail store worker," "teacher," or "nurse," occupations that typically involve frequent face-to-face interaction with coworkers. In contrast, men's valuations of mostly female workplaces are lower when the occupation in question is "software developer" or "insurance sales agent," office-based occupations which may involve less face time with coworkers. Women's valuations display much less heterogeneity across occupations.

Additionally, in an alternative design in which respondents choose between jobs without a specific occupation listed, women's composition valuations are similar to the fixed-occupation workplace choice, but men value mostly male jobs more, suggesting that men's desire for mixed to more female jobs is specifically about the workplace. As shown in Appendix Figure B.16, men are willing to pay less to avoid mostly male and more to avoid mostly female jobs when the occupation is unknown. This suggests that if men infer the specific occupation being performed from the gender composition, they are somewhat more likely to choose mostly male occupations, and therefore they may prefer the tasks or amenities in mostly male occupations. However, if the occupation is fixed, men have more desire for more gender diversity among their coworkers.

Finally, in a section of the survey in which respondents report how satisfied they expect to be with various attributes of a mostly female and a mostly male job, women expect to be more satisfied with most aspects of a mostly female job, while men expect to be more satisfied with the coworkers in a mostly female job but all other attributes in a mostly male job. This is illustrated in Figure 9, which shows the share of respondents who report they would be more satisfied with the listed attribute in a mostly female job minus the share who report they would be more satisfied in a mostly male job. Women report they would be more satisfied with the coworkers, schedule, work environment, tasks, and promotion probability in a mostly female job, even though they expect to earn less. Men, on the other hand, are only more likely to prefer the coworkers in a mostly female job, and appear indifferent or prefer a mostly male job on all other listed attributes.

Overall, it seems that women prefer many attributes of majority female jobs, but men particularly like the coworkers in majority female jobs. However, this does not necessarily imply that for women female coworkers are only valuable as a signal of other job qualities. Older women are more likely to value more female workplaces, and women with children are no more likely to value more female workplaces. This suggests that women are willing to pay for a higher female share for something other than family-friendly amenities. Also, adding additional amenities to the hypothetical job choice changes the value of the female share relatively little. This suggests that female coworkers have some inherent value or at least that they signal some workplace qualities that are not easily measured or advertised.

#### III.D.5 Validation and External Relevance

The survey results shown thus far illustrate that women and men do value the gender composition of their workplaces. The model of occupation choice in Section II shows that workers who value more female workplaces will choose more female occupations, all else equal. In this section, I evaluate the external relevance of the survey-estimated valuations by measuring whether they correspond to real-life occupation choices.

The survey includes multiple measures of the gender composition of a respondent's occupation, job, and employer. First, early in the survey, respondents report their occupation from a drop-down list of 26 Census occupation groups. I then calculate the female share of the reported occupation using the March CPS sample. I also ask respondents to report the gender composition of their employer and their coworkers who perform the same job as they do. Finally, I ask respondents to report whether they think their job is viewed by others as more likely to be done by a man or a woman.

Women who report working for employers with a higher female share are more likely to belong to the female-preferring class, but men who report working for employers with a higher female share are somewhat more likely to belong to the male-preferring class. This is illustrated in Figure 10, which shows a binned scatter plot of the probability of belonging to each preference class against the reported female share of the respondent's employer, separately for men and women.

It is possible that women's preference class correlates more with their job choice for several reasons. First, for the female-preferring group among men, the estimated WTPs are large but imprecisely estimated, so it is possible that in reality the WTPs are much smaller than the point estimates for this group. Second, it is possible that survey estimated composition valuations are simply more relevant to job choice for women than they are for men. I find that the male valuation results are noisier overall, and are less stable across variations and survey waves, suggesting gender composition valuations are perhaps less stable or less well-formed among men.

Finally, men might choose jobs that do not align with their valuations for gender composition because these jobs are preferable for other reasons. As demonstrated in the previous section, women seem to prefer many attributes of more female workplaces, while men only prefer female coworkers. Thus, although men in the female-preferring class report valuing workplaces with more female coworkers, other attributes of these jobs may be unappealing to men. This demonstrates the importance of holding "all else equal" in my survey design—in reality, these jobs are not otherwise identical, so the fact that men's valuations do not seem to correlate with their actual job choices does not mean they wouldn't choose jobs more aligned with their composition valuations given the chance to do so "all else equal." For instance, male software engineers might prefer firms with more female engineers, but in reality these firms may be hard to find. In addition, for both men and women, employer-side frictions like discrimination may prevent them from choosing a job that aligns with their gender composition valuation.

## IV Quantifying the Aggregate Effects of Gender Composition Preferences

With estimates of the value of gender composition in hand, I can now quantify the aggregate consequences of these preferences. Using a structural model of occupation choice that expands on the model in Section II, I show that composition preferences matter for sorting but are unlikely to create tipping points and that reducing segregation could improve welfare.

### IV.A Quantitative Model Environment

To perform quantitative exercises using the survey-estimated preferences, I expand the toy model from Section II to include multiple occupations, unobserved amenities, and a nested occupation choice structure.

There are three additions to the quantitative model relative to the toy model. First, there are now $K$ occupations indexed by $k$ rather than only two occupations. Non-participation is considered its own occupation. Second, each occupation has a gender- and occupation-specific residual amenity $a_{k,g}$. Mathematically, this term allows me to match observed employment shares across occupations, which in reality depend on other characteristics in addition to gender composition and wages.[25] Practically, this term will include both positive and negative amenities, like flexibility and danger, as well as barriers to entry, like education. Third, the structure of the individual preference shocks is now nested. I apply a nested logit structure to the model so that preference shocks are more correlated within an occupation nest than across nests, and therefore in counterfactuals when workers switch occupations they will be more likely to move to occupations that are more closely related to their own.

The utility function in the quantitative model is as follows:

$$U_{i,g} = \max_{s,k \in s} \left\{ Z_{s,g} + \log(w_{k,g}) + a_{k,g} + h_g \left( \frac{\ell_{kf}}{\ell_k} \right) + \varepsilon_{i,k,s} \right\}. \tag{9}$$

As in Section II, $w_{k,g}$ is an occupation- and gender-specific wage, and $h_g\left(\frac{\ell_{kf}}{\ell_k}\right)$ is the gender-specific gender composition utility function. In addition, $a_{k,g}$ is an occupation- and gender-specific amenity, $Z_{s,g}$ is a nest- and gender-specific amenity, and $\varepsilon_{i,k,s}$ is an individual-specific occupation-level preference shock that follows a generalized extreme value distribution with correlation $1 - \lambda$ within nests and no correlation across nests. The parameter $\lambda$ serves a similar function to the shape parameter of the taste shock in the standard logit model—the larger $\lambda$ is, the more important taste shocks are for occupation choice within nest and the less important wages and other amenities are.[26]

A nested occupation choice accounts for the fact that some occupations are more closely related than others. Formally, the nested logit model breaks the independence of irrelevant alternatives that is inherent to the standard multinomial logit model: it is no longer the case that a utility change in a single occupation will have equal effects on all other occupations.

With the nested logit structure on preference shocks, the share of workers in occupation $k$ within nest $s$ is

$$\frac{\ell_{k,g}}{\sum_{j \in s} \ell_{j,g}} = \frac{\exp(V_{k,g}/\lambda)}{\sum_{j \in B_s} \exp(V_{j,g}/\lambda)}, \tag{10}$$

[25] See Grigsby (2024).

[26] If $\lambda = 1$, the model collapses to a standard logit.

where $V_{k,g} = \log(w_{k,g}) + a_{k,g} + h_g\left(\frac{\ell_{kf}}{\ell_k}\right)$. Within each nest, the model behaves similarly to the standard logit model: employment shares are proportional to the relative utility of each occupation.

The shares of each nest in overall employment are given by

$$P(s) = \frac{\exp(Z_s'\alpha + \lambda IV_s)}{\sum_l \exp(Z_l'\alpha + \lambda IV_l)}$$

where

$$IV_s = \log\left[\sum_{k\in S} \exp(V_k/\lambda)\right].$$

That is, the employment shares across nests are proportional to the total utility in each nest, which includes an aggregate of the utility in each occupation in the nest and a constant utility $Z_s$ which is equal across occupations within the nest. This structure ensures that in counterfactual exercises where utilities shift across occupations, workers will be more likely to switch to an occupation that is more closely related to their initial occupation.

As in the simple model, production is CES across occupations and across genders within an occupation.

$$Y = \left(\sum_k (A_k\ell_k)^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}},$$
$$\ell_j = \left(q_k(\ell_{k,m})^{\frac{\alpha-1}{\alpha}} + (1-q_k)(\ell_{k,f})^{\frac{\alpha-1}{\alpha}}\right)^{\frac{\alpha}{\alpha-1}}.$$

### IV.B From Survey to Model

Equating occupational gender composition utility, which I measure in the model, to workplace gender composition utility, which I measure in the survey, requires some assumptions on the structure of occupation choice. First, I assume that when workers in the model choose their occupation, they know the gender composition of the occupation as a whole but not the distribution of gender compositions across workplaces. Second, I assume that labor market frictions imply they will not be able to freely choose a workplace with their preferred gender composition within their occupation.

Even if these assumptions seem rather stringent, equating occupational composition utility to workplace composition utility may have limited importance for the overall results of the model. Preferences across gender compositions are strongest at the extreme ends of the distribution, and within extremely segregated occupations, workplaces whose gender compositions diverge greatly from the occupational average are likely to be rare. In addition, an individual's closest colleagues in a workplace are likely those with the same or very similar occupations. For the model, this means that if we estimate large movements out of extremely segregated occupations in counterfactuals, these movements would not be possible across workplaces without changing the gender composition of the occupation.[27]

### IV.C Model Estimation

I estimate the model using my survey results and data from the March CPS. My data on allocations ($\ell_{g,k}$) and wages ($w_{g,k}$) across occupations by gender come from the March CPS. I use the 2014-2019 data on workers age 18 and up. I divide workers into 434 occupations using 2010 Census occupation codes. For occupation- and gender-specific wages, I take the average annual earnings from the prior year for each occupation, residualized on education, age, age squared, usual weekly hours, and year fixed effects. I calculate the allocations across occupations as the share of workers of each gender with a given occupation listed as their main occupation in the previous year. I define non-participation or unemployment as its own occupation, whose wage and amenity I normalize to zero.

To apply my survey-estimated composition valuations in the model, I fit quadratic functions to the average composition value (measured as a percent of wages) separately for men and women. These functions are plotted with the composition value estimates in Figure 5. This approximation allows me to estimate gender composition utility for compositions I do not directly elicit values for in the survey. I calibrate $\lambda$, the parameter governing the correlation of extreme value preference shocks within nests, externally to 0.5 from estimates of occupational and firm labor supply elasticities in the literature.[28]

Next, I calibrate the production functions elasticities externally. I set the elasticity of substitution across occupations $\eta$ in the aggregate production function equal to 1.5, based on estimates of the elasticity of substitution across skilled and unskilled occupations.[29] This implies that occupations are gross substitutes. I set the elasticity of substitution across genders $\alpha$ equal to 2.5, following estimates in Ngai and Petrongolo (2017). This implies that within occupation, genders are also gross substitutes, but genders are more substitutable within occupations than occupations are. A non-infinite value for the elasticity of substi-

[27] I discuss the issue of comparing workplace and occupation gender composition further in section III.A.1.

[28] This parameter is roughly equivalent to the variance of the type I EV preference shock in a standard logit model. I use a value of $\lambda = 0.5$, corresponding to a wage elasticity of 2. This lies between estimates of the wage elasticity across a smaller number of occupations (Dix-Carneiro, 2014; Hsieh et al., 2018) and across firms (Bassier et al., 2022), because I estimate my model on a large number of occupations and incorporate a nested structure.

[29] Acemoglu and Autor (2011) note that most estimates for this elasticity lie between 1.4 and 2.

tution across genders can be due to differences in production stemming from, for example, differences in hours or differences in the specific tasks performed within an occupation. Table 3 summarizes the values and sources of parameters that enter the model estimation.

Finally, the nested logit model requires occupations be divided into nests before estimation. To do this, I follow Nimczik (2022) and Schubert et al. (2021) by defining a network of occupations based on observed worker flows. Intuitively, I combine occupations that workers frequently flow between into the same nest. The full details of the nest estimation process, and the estimated nests, are described in Appendix C.

For non-participation, I normalize the values of the amenity, wage, and composition utility to zero. This pins down the overall level of the amenities $a_{k,g}$ for each gender and occupation. Given wages and composition preferences, the amenity $a_{k,g}$ across occupations is the residual utility needed to rationalize allocations across occupations.

The overall occupational productivities $A_k$ and relative gender productivities $q_k$ are determined by gender-specific wages and occupational allocations. Taking the first-order conditions of the profit function gives the wage for each gender as a function of labor supply, $A_k$, and $q_k$, so given wages and allocations, I can back out $A_k$ and $q_k$ for each occupation. Intuitively, $A_k$ depends on the level of wages relative to total labor, and $q_k$ depends on the gender wage gap. I normalize the level of output so that wages sum to total output.

### IV.D Tipping and Labor Supply to a Single Occupation

In this section, I use my survey-estimated preferences to analyze the possibility of tipping across occupations. I find that tipping is unlikely to occur unless occupational wage elasticities are significantly larger than those typically estimated in the literature. However, allowing for amenity provision to respond endogenously to gender composition increases the likelihood of tipping.

Tipping occurs when a small perturbation in labor demand or supply induces a movement from a gender-mixed to a gender-segregated equilibrium, or vice versa. In order for tipping to occur, the occupational labor supply function must bend backwards—that is, at a given wage level, more than one sorting equilibrium must be possible.

Taking total occupational labor as fixed and normalized to one, I can write the female labor supply function to a fixed occupation $k$ relative to some outside occupation $j$ as

$$\begin{aligned}&\log(w_{k,g}) - \log(w_{j,g}) = \\ &\lambda\left[\log(f_k) - \log(f_j)\right] + \left[a_{j,g} - a_{k,g}\right] + \left[h_g\left(f_j\right) - h_g\left(f_k\right)\right],\end{aligned} \tag{11}$$

where $f_k$ is the female share of occupation $k$, and $\lambda$ is the inverse wage elasticity.[30] This function implies that the labor supply of women to an occupation $k$ will increase as its wage, amenity value, and gender composition value increase relative to occupation $j$.

From the labor supply function, I can derive a condition for tipping: at some point, the derivative of wages with respect to female labor supply must be negative. That is, there must be some point at which increasing the number of women in the occupation *lowers* the wage required to induce that number of women to work in the occupation. Setting the derivative of wages with respect to labor supply greater than zero gives

$$\frac{\lambda}{f_k} < \frac{\partial}{\partial f_k}\left[h_f(f_k)\right].$$

That is, tipping can occur when the slope of the composition preference function exceeds the slope of the labor supply function absent composition preferences.

In the case of composition preferences that are quadratic in the female share, I can derive a simple condition for tipping relating the slope of composition preferences to the wage elasticity. With composition preferences of the form

$$h_g(f_k) = a + bf_k - cf_k^2,$$

we must have

$$\lambda < \frac{b^2}{8c}$$

to create the possibility of a tipping point.[31]

Using reasonable estimates of the occupational wage elasticity and my survey estimates of composition preferences, it is unlikely that tipping will occur due to workplace composition preferences alone. In Table 4, I show the estimated composition utility function from the survey for women (average and female-preferring group) along with the maximum inverse wage elasticity $\lambda$ necessary to create a tipping point. I also show the multiplier on composition preferences required to create tipping with a standard estimate of the wage elasticity ($\lambda = 0.5$). I find that even for the women with a strong preference for a more female workplace, an extremely large wage elasticity (around 14) would be necessary to create a tipping point. With a more standard wage elasticity of 2 ($\lambda = 0.5$), even for the female-preferring group the composition preferences would need to be about 7 times as large as those which I

[30] When total labor is fixed and normalized to one, female labor supply in an occupation is equal to the female share.

[31] Appendix C.2 provides the derivation of this condition.

estimate to create a tipping point.[32]

Alternatively, it is possible that tipping could be generated not only from composition preferences but from downstream effects of changing composition on other job characteristics. For instance, as the female share increases, an occupation might increase its provision of female-friendly amenities, or it might become seen as less socially acceptable for men to perform. I generalize these possibilities by considering an extension of the job choice model where the gender-specific amenity depends on the female share:

$$a_{k,g} = a_g(f_k).$$

If I assume that the endogenous amenity is linear in the female share in the form $a_f(f_k) = \alpha f_k$, I calculate that to generate tipping with endogenous amenities, it must be the case that

$$\lambda < \frac{(b+\alpha)^2}{8c}.$$

As above, I can calculate the minimum derivative of amenities with respect to the female share, $\alpha$, such that tipping is possible for a fixed wage elasticity. I tabulate these estimates in Table 5 for both women (requiring a positive $\alpha$) and men (requiring a negative $\alpha$).

For the average woman and man, $\alpha$ would need to be around 0.5 to generate tipping. What does this mean? Consider a 10 percentage point increase in the female share $\Delta f = 0.1$. Then, $\alpha/10$ is the percent wage increase that is equivalent to the amenity increase brought about by a 10 percentage point increase in the female share. Then, $\alpha = 0.5$ implies that a 10 percentage point increase in the female share creates additional amenities with value equivalent to a 5% increase in wages. This is not unreasonable to imagine, particularly in the long run.[33] In addition, it is possible that discontinuous increases in female-friendly amenities could occur at some levels of the female share, making this possibility more likely. This result suggests that investigating the endogeneity of amenity provision and social norms to the demographic composition of a job could be a promising avenue for further research. In this vein, Corradini et al. (2022) find that the adoption by a Brazilian trade union of an agenda that increased female leadership and explicitly promoted female-friendly policies increased the provision of female-friendly amenities and increased female demand for these jobs.

[32] For men, the condition for tipping is slightly more complex because male labor contributes negatively to the female share, and on average men prefer a higher female share, at least at low male shares. I leave the derivation to the appendix, but I find tipping is unlikely for any group of men.

[33] Mas and Pallais (2017) find that the average worker is willing to give up 20% of wages to avoid a schedule that is set on short notice, and Wiswall and Zafar (2018) find that women are willing to pay about 7% of wages for an option to work part-time, for example.

On the whole, I conclude that the gender composition preferences I estimate are, on their own, unlikely to create tipping points. Increasing the occupational wage elasticity would make tipping more likely, suggesting that tipping is most likely to occur between occupations that are closely related, or over long periods of time in which workers are more flexible over their occupation choices. In addition, increasing the size of the composition preference increases the likelihood of tipping. Given that I find that older women are significantly more likely to have a large preference for a higher female share, this suggests that tipping between occupations may have been possible in the past. Finally, allowing for the provision of amenities to be endogenous to the female share of an occupation could make tipping much more likely.[34]

### IV.E The Utility Cost of Segregation

How much better off would workers be if there were no gender segregation across occupations? Given my estimates of preferences for gender composition, I can directly measure how much utility would be gained (purely due to composition preferences) from desegregating all occupations. To do this, I calculate the compensation required to make workers indifferent between their current occupational gender composition and an occupational female share of 50%. I find that on average for both men and women the equivalent variation (EV) for moving to total desegregation is positive, but there is wide variation across individuals.

The average worker would be slightly better off if their occupation were evenly gender balanced. Table 6 shows the average equivalent variation for one's occupation becoming 50% female by sex and preference class. In general, the average EV is positive, around 0.4% for women and around 1% for men. Splitting up these results by preference classes shows that the average EV is much larger for the female-preferring group of both men and women—on average, moving to a 50% female share is equivalent to a 0.7% wage increase for female-preferring women and an 2.8% wage increase for female-preferring men. Men's EV is larger because, although their WTP for an evenly mixed job is smaller, they are much more likely than women to be in a mostly male occupation.

Looking at the distribution of equivalent variations among women shows most respondents would be indifferent to their job becoming evenly gender balanced, but a sizeable minority would experience large gains in utility. Figure 11 shows the distribution of the equivalent variation for a 50% female occupation among employed female survey respondents. About 55% of the sample has an EV of essentially zero for a move to a 50% female

[34] These results suggest that the evidence for tipping found in Pan (2015) may come from a combination of larger composition preferences in the past and amenities or social norms that are dependent on the female share of an occupation. Also, although Schelling (1971) suggests that tipping can happen even with very small composition preferences, this result assumes that items in the choice set are otherwise identical.

occupation. Most of this group is the wage-preferring group, among whom the WTP for any change in female share is close to zero. About 6% of the sample, though, belongs to the female-preferring group but still has an EV of zero for moving to a 50% female occupation, because these respondents are already in an occupation with a female share very close to 50%. Another 23% of the sample has a negative, but small, EV—these women would actually be *less* happy if their occupation were 50% female. These women are in jobs that are 50-90% female, which are slightly preferred to half female jobs. Notably, though, the point estimates for WTPs in this range are not significantly different from zero, so these small utility losses may actually be negligible. Finally, 22% of the sample has a positive EV—these women would be better off if their occupation were half female. Although this group is in the minority, their EVs are large and widely varied—the maximum EV is 12%, and the average EV among the positive EV group is 2%.

Similarly, most men are nearly indifferent to their job becoming evenly gender balanced, but some men would experience large utility gains from such a change. Figure 12 shows the distribution of the equivalent variation for a change to a 50% female job among men, split by preference class. About 40% of the sample would experience basically no change in utility if their occupation became 50% female, most of which come from the wage-preferring group. Another 17% of men would be slightly worse off if their job became half female. Finally, 43% of the sample would be better off if their job became half female. Among this group, utility gains are widely dispersed, with an average EV of 3% and a maximum of 15%, which occurs for men who belong to the female-preferring group but are in extremely male occupations such as natural resource extraction and construction.

Overall, these results suggest that although most people would be indifferent to their occupation becoming desegregated, a small group of people would experience large utility gains if their occupation became evenly gender balanced. In fact, men would actually gain more utility than women, because they are less likely to be sorted in accordance with their composition preferences. This suggests that policies to reduce segregation could be desirable from a utility perspective, but they are best targeted at the individuals who care the most about gender composition.

### IV.F Aggregate Effect of Composition Valuations

Another way of assessing the aggregate importance of gender composition preferences is to estimate their effect on overall labor allocation across occupations. To do this, I use my model to construct counterfactual allocations and wages across occupations in a world with no gender composition preferences.

Formally, to remove gender composition valuations, I take the existing model estimates of occupational productivity and preference parameters and set the composition value $h_g(\ell_{g,k}/\ell_k)$ to zero. I then solve for the counterfactual allocations $\ell'_{k,g}$ that would occur in the absence of gender composition valuations:

$$\frac{\ell'_{k,g}}{\ell_g} = \frac{\exp[\log(w_{k,g}) + a_{k,g}]}{\sum_j \exp\left[log(w_{j,g}) + a_{j,g}\right]} \quad j = 1, ..., J.$$

I first estimate the new allocations in partial equilibrium. I use the heterogeneous valuations estimated using the latent class logit model. I observe the 2-digit occupations of survey respondents and impute allocations by gender and preference class across 6-digit Census occupations using the observed allocations of survey respondents to 2-digit occupations and allocations within 2-digit occupations across 6-digit Census occupations in the CPS.

Figure 13 illustrates that absent composition valuations, female employment would increase in mostly male occupations. I estimate that if workers did not value gender composition, the number of women working in occupations that are less than 9% female in reality would be on average 10% higher. Conversely, if workers did not value gender composition, female employment would be 0.5-1% lower in occupations that are 50-80% female in reality. The changes in men's employment have similar patterns, but are more muted. Thus, it does appear that workers' valuations of gender composition have nontrivial effects on occupational employment, particularly for women in very male occupations.[35]

In Appendix Tables C.5 and C.4, I include results from this counterfactual exercise in general equilibrium. In partial equilibrium, I keep wages fixed to their true values at the occupation-gender level. In general equilibrium, I allow wages to adjust in response to changes in worker allocations across occupations by gender. In general equilibrium, allocations move less but wages change more: the gender wage gap due to sorting falls from 4.7% to 4.1% when composition preferences are set to zero.

### IV.G Correcting Sorting Externalities

In this section, I investigate how we could optimally reallocate workers across occupations to improve utility. Reallocation may improve welfare in this context because composition valuations create a sorting externality: individuals do not take into account the effects of their entry into an occupation on the utility of others in that occupation. I solve a modified social planner's problem and conclude that utility could be improved by reallocating workers to

[35] I include additional statistics on the changes in wages and overall gender segregation absent composition preferences in the appendix. Overall, the changes in employment are not large enough to have a substantial effect on the average degree of segregation or gender wage gaps–the number of women in extremely male occupations is so small that female employment in these occupations would need to double or triple to see significant movement in the degree of segregation.

reduce segregation. This differs from the equivalent variation calculation because it accounts for differences in amenities and wages across occupations. A minority-gender-specific wage or amenity subsidy in very segregated occupations could achieve similar utility gains.

To solve for optimal allocations in a modified social planner's problem, I augment individual utility to account for each individual's effect on the utility of other workers in their occupation. Recall that in the individual utility function,

$$U_{i,g} = \max_k \left\{ \log(w_{k,g}) + a_{k,g} + h_g \left( \frac{\ell_{kf}}{\ell_k} \right) + \varepsilon_{i,k} \right\},$$

workers receive utility from wages, amenities, gender composition, and idiosyncratic shocks in their occupation.[36] In a true social planner's problem, to maximize total utility, holding wages fixed, we would choose an allocation of workers to occupations $\varphi : i \to k(i, \varphi)$ that maximizes the sum of occupational utility across individuals:

$$\max_{\varphi} \sum_i \left\{ \log(w_{k(i,\varphi),g}) + a_{k(i,\varphi),g} + h_g \left( \frac{\ell_{k(i,\varphi)f}}{\ell_{k(\varphi)}} \right) + \varepsilon_{i,k(i,\varphi)} \right\}.$$

Solving for this allocation in closed form is not straightforward due to the sum of extreme value draws (Davis and Gregory, 2021). Thus, I approximate the solution by solving for a decentralized equilibrium using an individual utility function which I augment to directly account for sorting externalities:

$$\begin{aligned} U'_{i,g} = \max_k \Bigg\{ & \log(w_{k,g}) + a_{k,g} + h_g \left( \frac{\ell_{kf}}{\ell_k} \right) + \varepsilon_{i,k} + \\ & \ell_{k,f} \frac{\partial}{\partial \ell_{k,g}} h_f \left( \frac{\ell_{kf}}{\ell_k} \right) + \ell_{k,m} \frac{\partial}{\partial \ell_{k,g}} h_m \left( \frac{\ell_{kf}}{\ell_k} \right) \Bigg\}. \end{aligned}$$

The above utility function adds to the original utility function the terms $\ell_{k,f} \frac{\partial}{\partial \ell_{k,g}} h_f \left( \frac{\ell_{kf}}{\ell_k} \right)$ and $\ell_{k,m} \frac{\partial}{\partial \ell_{k,g}} h_m \left( \frac{\ell_{kf}}{\ell_k} \right)$ which account for the effect of individual $i$ of gender $g$ joining occupation $k$ on the utility of other women and men, respectively, in occupation $k$. Given this modified utility function, I can simply solve for a new decentralized equilibrium which closely approximates the true social planner's solution. In this exercise, I use a simple logit model rather than the nested logit model for simplicity.

I find that a social planner would reallocate workers to reduce segregation by increasing the female share in mostly male occupations and reducing the female share in mostly female occupations. Figure 14 illustrates that the social planner would increase the female share of occupations that are less than half female by up to 2 percentage points, and decrease that

[36] In this section, I abstract from the nested structure used previously and revert to a basic logit model.

of occupations that are more than half female by up to 1.5 percentage points.

The social planner would create this decrease in segregation by substantially increasing female employment in mostly male occupations and slightly increasing male employment in mostly female occupations, and vice versa for own-gender-dominated occupations. Figure 15 illustrates that the social planner would increase female employment in the most male occupations by up to 20%, and increase male employment in the most female occupations by up to 10%.

I find that under the social planner's solution, utility improves slightly on average for both men and women, with larger improvements for individuals who begin in segregated occupations. I calculate an equivalent variation for the change in utility under the social planner's solution by finding the wage subsidy that equates utility in the decentralized equilibrium with utility in the social planner's solution. On average, we would need to raise women's wages by 0.02% and men's wages by .05% in the decentralized equilibrium to equate their utility to that in the social planner's solution. These numbers are rather small, but they are widely dispersed. Women in occupations that are initially 10-30% female would have welfare gains equivalent to a 0.2% percent wage increase, on average, in the social planner's solution. Notably, these gains may be longer in the long run if amenities respond endogenously to the changes in composition.

The social planner will move toward desegregation because both genders are the happiest in jobs that are relatively gender balanced. Given that both men and women, on average, dislike majority male workplaces, the social marginal utility of increasing female employment, and reducing male employment, in mostly male occupations is large. Men and women also dislike extremely female dominated workplaces, but less than they dislike male dominated workplaces, so the social planner will decrease female shares in these occupations by less than they change the female shares in the mostly male occupations.[37]

The social planner reduces segregation by overcoming the problem of the marginal gender minority worker. Although both genders will be better off as an occupation becomes more gender mixed, one woman (or man) will be reluctant to enter a majority male (or female) occupation because one person cannot guarantee that others will follow. If individuals could collude to create a large shift in the gender share of an occupation, desegregating would be easier, but in a decentralized equilibrium individuals will be more likely to choose occupations that already fit their preferred gender composition.

[37]Notably, the female shares change slightly less in the most segregated occupations, because in these occupations the gender difference in the residual amenity $a_{k,g}$ is so large that even with a higher female share, these occupations still offer low utility for the gender minority. In addition, because the initial number of women is so small, even increasing female employment by 20% does not have large effects on the female share.

The solution to the planner's problem suggests that desegregation is a desirable policy goal because both men and women would potentially be better off if occupations, particularly those with very high or very low female shares, were less segregated. This exercise demonstrates that policies that reduce the cost of entry into an occupation for the gender minority could be fruitful. For instance, policies that improve amenities that women value highly or subsidize education for women in male-dominated occupations could help reduce segregation and benefit workers of all genders.

Subsidizing wages or amenities for workers in the minority gender in segregated occupations could feasibly provide similar reductions in segregation to the social planner's solution. I evaluate the effects of a simple policy where women receive a 10% wage subsidy in occupations that are less than 30% female, and men receive a 10% wage subsidy in occupations that are more than 70% female. This subsidy on average increases utility by a very small amount (the equivalent variation is 0.01% for women and 0.05% for men, excluding the subsidy), but improves utility for women and men in segregated occupations. Women in occupations that are 10-30% female would have utility gains equivalent to a 0.3% wage increase, and men in occupations above 70% female would have utility gains equivalent to a 0.1% wage increase. In Appendix Figure C.2, I plot the changes in female shares under this policy.

## V Conclusion

In this paper, I study whether workers value the gender composition of their workplace and how this affects occupational gender segregation and worker welfare. Using a novel online survey, I estimate that both men and women are willing to trade off a nontrivial portion of their wages for a job that has their preferred gender composition. I find that on average, both women and men prefer gender-mixed jobs. Importantly, these preferences are heterogeneous across individuals. I estimate that about half of women and men do not care about the gender composition of their jobs, but the other half of women and men are willing to trade off a nontrivial portion of their wages for a job with their preferred gender composition. Importantly, older workers are more likely to value workplace gender homophily, suggesting that homophily has become less valuable as men and women's labor market outcomes converged over the 20th century.

I use these estimated valuations in a structural model of occupation choice to assess their implications for occupational gender segregation in the aggregate. I find that if workers did not value the gender composition of their occupation, the number of women working in majority male occupations would increase by up to 10%. Additionally, because mostly female occupations tend to have lower wages, removing gender composition preferences would reduce the portion of the gender wage gap due to occupational sorting by nearly 10% in general equilibrium. Finally, I show that reducing segregation could improve welfare. Making all occupations evenly gender balanced would improve utility by the equivalent of a 0.7% wage increase for women and a 2.8% wage increase for men, on average. In addition, a welfare-maximizing social planner would reduce gender segregation, which would improve welfare.

In all, this paper shows that the gender composition of a workplace is an important amenity that can have large consequences for occupation choice and worker welfare. This complements our understanding of the importance of non-wage amenities for job choice by showing that who performs a job may be as important as the benefits a workplace provides.

## References

**Acemoglu, Daron and David Autor**, "Skills, Tasks and Technologies: Implications for Employment and Earnings," *Handbook of Labor Economics*, 2011, *4*, 1043–1171.

**Adams-Prassl, Abi, Kristiina Huttunen, Emily Nix, and Ning Zhang**, "Violence Against Women at Work," *The Quarterly Journal of Economics*, 09 2023.

**Akerlof, George A and Rachel E Kranton**, "Economics and identity," *The quarterly journal of economics*, 2000, *115* (3), 715–753. Publisher: MIT Press.

**Alvarez, R. Michael, Lonna Rae Atkeson, Ines Levin, and Yimeng Li**, "Paying Attention to Inattentive Survey Respondents," *Political Analysis*, April 2019, *27* (2), 145–162.

**Babcock, Linda, Maria P Recalde, Lise Vesterlund, and Laurie Weingart**, "Gender differences in accepting and receiving requests for tasks with low promotability," *American Economic Review*, 2017, *107* (3), 714–747.

**Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu**, "Monopsony in movers: The elasticity of labor supply to firm wage policies," *Journal of Human Resources*, 2022, *57* (S), S50–s86.

**Bertrand, Marianne and Esther Duflo**, "Field experiments on discrimination," *Handbook of economic field experiments*, 2017, *1*, 309–393. Publisher: Elsevier.

_ **, Patricia Cortés, Claudia Olivetti, and Jessica Pan**, "Social Norms, Labour Market Opportunities, and the Marriage Gap Between Skilled and Unskilled Women," *The Review of Economic Studies*, July 2021, *88* (4), 1936–1978.

**Blau, Francine D. and Lawrence M. Kahn**, "The Gender Wage Gap: Extent, Trends, and Explanations," *Journal of Economic Literature*, September 2017, *55* (3), 789–865.

**Boas, Taylor C., Dino P. Christenson, and David M. Glick**, "Recruiting large online samples in the United States and India: Facebook, Mechanical Turk, and Qualtrics," *Political Science Research and Methods*, April 2020, *8* (2), 232–250.

**Bostwick, Valerie K. and Bruce A. Weinberg**, "Nevertheless She Persisted? Gender Peer Effects in Doctoral STEM Programs," *Journal of Labor Economics*, 2022, *40* (2), 397–436.

**Breen, Richard and Cecilia García-Pe nalosa**, "Bayesian Learning and Gender Segregation," *Journal of Labor Economics*, 2002, *20* (4), 899–922.

**Brock, W. A. and S. N. Durlauf**, "Discrete Choice with Social Interactions," *The Review of Economic Studies*, April 2001, *68* (2), 235–260.

**Bursztyn, Leonardo and Robert Jensen**, "Social Image and Economic Behavior in the Field: Identifying, Understanding, and Shaping Social Pressure," *Annual Review of Economics*, August 2017, *9* (1), 131–153.

— , **Thomas Fujiwara, and Amanda Pallais**, "'Acting Wife': Marriage Market Incentives and Labor Market Investments," *American Economic Review*, November 2017, *107* (11), 3288–3319.

**Bustelo, Monserrat, Ana Maria Diaz, Jeanne Lafortune, Claudia Piras, Luz Magdalena Salas, and Jose Tessada**, "What is the price of freedom? Estimating women's willingness to pay for job schedule flexibility," *Economic Development and Cultural Change*, January 2022, p. 718645.

**Coppock, Alexander and Oliver A. McClellan**, "Validating the demographic, political, psychological, and experimental results obtained from a new source of online survey respondents," *Research & Politics*, January 2019, *6* (1), 205316801882217.

**Corradini, Viola, Lorenzo Lagos, and Garima Sharma**, "Collective bargaining for women: How unions can create female-friendly jobs," 2022.

**Cortes, Guido Matias, Nir Jaimovich, and Henry E Siu**, "The growing importance of social tasks in high-paying occupations: implications for sorting," *Journal of Human Resources*, 2021, pp. 0121–11455R1. Publisher: University of Wisconsin Press.

**Cortés, Patricia and Jessica Pan**, "Occupation and gender," *The Oxford handbook of women and the economy*, 2018, pp. 425–452. Publisher: Oxford University Press New York, NY.

— **and** — , "Children and the remaining gender gaps in the labor market," *Journal of Economic Literature, forthcoming*, 2020.

— , **Gizem Koşar, Jessica Pan, and Basit Zafar**, "Should Mothers Work? How Perceptions of the Social Norm Affect Individual Attitudes Toward Work in the US," Technical Report, National Bureau of Economic Research 2022.

**Cullen, Zoë and Ricardo Perez-Truglia**, "The Old Boys' Club: Schmoozing and the Gender Gap," *American Economic Review*, July 2023, *113* (7), 1703–40.

**Davis, Morris and Jesse M Gregory**, "Place-Based Redistribution in Location Choice Models," Technical Report, National Bureau of Economic Research 2021.

**Delfino, Alexia**, "Breaking gender barriers: Experimental evidence on men in pink-collar jobs," 2021. Publisher: IZA Discussion Paper.

**Dix-Carneiro, Rafael**, "Trade Liberalization and Labor Market Dynamics," *Econometrica*, 2014, *82* (3), 825–885.

**Drake, Marshall, Thakral, Neil, and Tô, Linh**, "Wage Differentials and the Price of Workplace Flexibility," 2022.

**Engel, Yuval, Trey Lewis, Melissa S Cardon, and Tanja Hentschel**, "Signaling Diversity Debt: Startup Gender Composition and the Gender Gap in Joiners' Interest," *Academy of Management Journal*, 2022, (ja).

**Exley, Christine L. and Kirby Nielsen**, "The Gender Gap in Confidence: Expected but Not Accounted For," *American Economic Review*, March 2024, *114* (3), 851–85.

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

**Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, and J. Robert Warren**, "Integrated Public Use Microdata Series, Current Population Survey: Version 6.0," 2018. type: dataset.

**Folke, Olle and Johanna Rickne**, "Sexual harassment and gender inequality in the labor market," *The Quarterly Journal of Economics*, 2022, *137* (4), 2163–2212. Publisher: Oxford University Press.

**Gallen, Yana and Melanie Wasserman**, "Informed Choices: Gender Gaps in Career Advice," 2021. Publisher: CEPR Discussion Paper No. DP15728.

**_ and _** , "Does Information Affect Homophily?," Technical Report 2022.

**Gelblum, Madeleine**, "Preferences for Job Tasks and Gender Gaps in the Labor Market," 2020.

**Gneezy, Uri, Muriel Niederle, and Aldo Rustichini**, "Performance in competitive environments: Gender differences," *The Quarterly Journal of Economics*, 2003, *118* (3), 1049–1074. Publisher: MIT Press.

**Goldin, Claudia**, "A pollution theory of discrimination: male and female differences in occupations and earnings," in "Human capital in history: The American record," University of Chicago Press, 2014, pp. 313–348.

**_ and Lawrence F Katz**, "A most egalitarian profession: pharmacy and the evolution of a family-friendly occupation," *Journal of Labor Economics*, 2016, *34* (3), 705–746. Publisher: University of Chicago Press Chicago, IL.

**Greene, William H and David A Hensher**, "A latent class model for discrete choice analysis: contrasts with mixed logit," *Transportation Research Part B: Methodological*, 2003, *37* (8), 681–698. Publisher: Elsevier.

**Grigsby, John**, "Skill Heterogeneity and Aggregate Labor Market Dynamics," 2024.

**Hampole, Menaka, Francesca Truffa, and Ashley Wong**, "Peer effects and the gender gap in corporate leadership: Evidence from MBA students," Technical Report, Working Paper 2022.

**Hellerstein, Judith, David Neumark, and Melissa McInerney**, "Changes in workplace segregation in the United States between 1990 and 2000: Evidence from matched employer-employee data," in "The analysis of firms and employees: Quantitative and qualitative approaches," University of Chicago Press, 2008, pp. 163–195.

**Henry, Marc and Ivan Sidorov**, "Occupational segregation in a Roy model with composition preferences," *arXiv preprint arXiv:2012.04485*, 2020.

**Hsieh, Chang-Tai, Erik Hurst, Chad Jones, and Pete Klenow**, "The Allocation of Talent and US Economic Growth," 2018.

**Kaplan, Greg and Sam Schulhofer-Wohl**, "The Changing (Dis-)Utility of Work," *Journal of Economic Perspectives*, August 2018, *32* (3), 239–258.

**Kessler, Judd B., Corinne Low, and Colin D. Sullivan**, "Incentivized Resume Rating: Eliciting Employer Preferences without Deception," *American Economic Review*, November 2019, *109* (11), 3713–3744.

**Larson-Koester, Miriam**, "Occupation Gender Segregation: Empirical Evidence from a Matching Model with Transfers," 2017.

**Lordan, Grace and Jörn-Steffen Pischke**, "Does Rosie Like Riveting? Male and Female Occupational Choices," *Economica*, January 2022, *89* (353), 110–130.

**Maestas, Nicole, Kathleen J Mullen, David Powell, Till Von Wachter, and Jeffrey B Wenger**, "The value of working conditions in the United States and implications for the structure of wages," *American Economic Review*, 2023, *113* (7), 2007–2047.

**Mas, Alexandre and Amanda Pallais**, "Valuing Alternative Work Arrangements," *American Economic Review*, December 2017, *107* (12), 3722–3759.

**Morchio, Iacopo and Christian Moser**, "The Gender Pay Gap: Micro Sources and Macro Consequences," *Available at SSRN 3176868*, 2020.

**Ngai, L. Rachel and Barbara Petrongolo**, "Gender Gaps and the Rise of the Service Economy," *American Economic Journal: Macroeconomics*, October 2017, *9* (4), 1–44.

**Niederle, Muriel and Lise Vesterlund**, "Do women shy away from competition? Do men compete too much?," *The quarterly journal of economics*, 2007, *122* (3), 1067–1101. Publisher: MIT Press.

**Nimczik, Jan Sebastian**, "Job mobility networks and data-driven labor markets," Technical Report, Working Paper 2022.

**Olivetti, Claudia, Eleonora Patacchini, and Yves Zenou**, "Mothers, peers, and gender-role identity," *Journal of the European Economic Association*, 2020, *18* (1), 266–301. Publisher: Oxford University Press.

**Pan, Jessica**, "Gender Segregation in Occupations: The Role of Tipping and Social Interactions," *Journal of Labor Economics*, April 2015, *33* (2), 365–408.

**Rendall, Michelle**, "Brain Versus Brawn: The Realization of Women's Comparative Advantage," *SSRN Electronic Journal*, 2018.

**Schelling, Thomas C.**, "Dynamic models of segregation†," *The Journal of Mathematical Sociology*, July 1971, *1* (2), 143–186.

**Schubert, Gregor, Anna Stansbury, and Bledi Taska**, "Employer concentration and outside options," *Available at SSRN 3599454*, 2021.

**Sloane, Carolyn M., Erik G. Hurst, and Dan A. Black**, "College Majors, Occupations, and the Gender Wage Gap," *Journal of Economic Perspectives*, November 2021, *35* (4), 223–248.

**Sorkin, Isaac**, "Ranking Firms Using Revealed Preference," *The Quarterly Journal of Economics*, August 2018, *133* (3), 1331–1393.

**Wiswall, Matthew and Basit Zafar**, "Preference for the Workplace, Investment in Human Capital, and Gender*," *The Quarterly Journal of Economics*, February 2018, *133* (1), 457–507.

**Yamaguchi, Shintaro**, "Changes in Returns to Task-Specific Skills and Gender Wage Gap," *Journal of Human Resources*, 2018, *53* (1), 32–70.

**Figures**

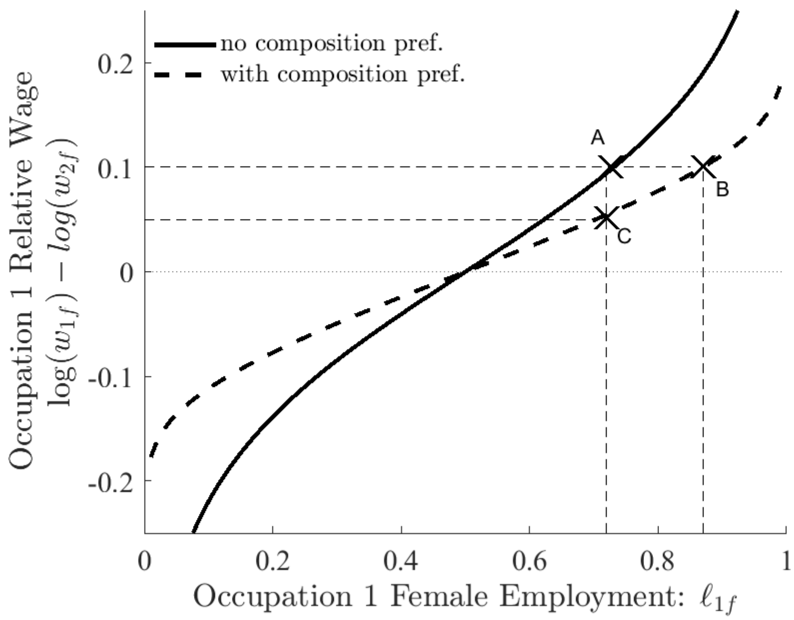


Figure 1: Equilibrium condition with composition preferences

Note: This plot displays the female labor supply function with and without gender composition preference. Here, I set $\eta = 0.1$, $\ell_{1m} = 0.5$, and $h_f(f) = 0.06 \cdot \log(f) - 0.5$.

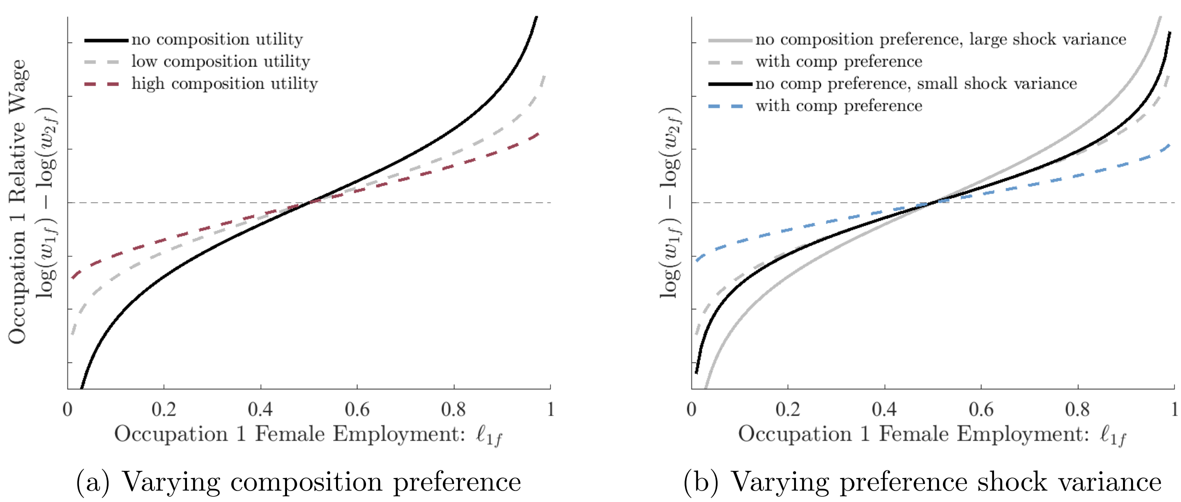


Figure 2: Equilibrium labor supply, varying composition preference, preference shock variance

Note: These plots display the female labor supply function with and without gender composition preferences of varying scale (panel a) or with varying preference shock variance $\eta$ (panel b). Here, I set $\eta = 0.1$, $\ell_{1m} = 0.5$, and $h_f(f) = 0.06 \cdot \log(f) - 0.5$ at baseline.

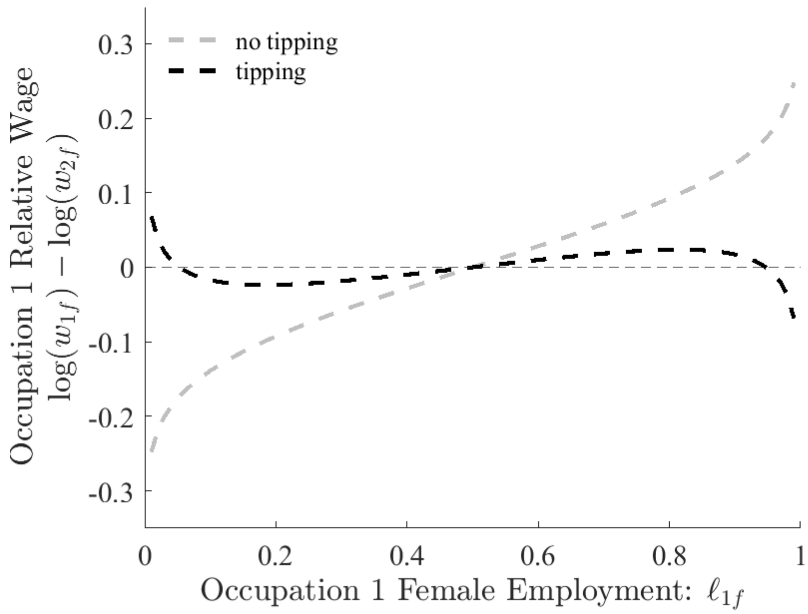


Figure 3: Tipping points in gender composition

Note: This plot displays the female labor supply function with and without gender composition preference. Here, I set $\eta = 0.1$, $\ell_{1m} = 0.5$, and $h_f(f) = 0.06 \cdot \log(f) - 0.5$.

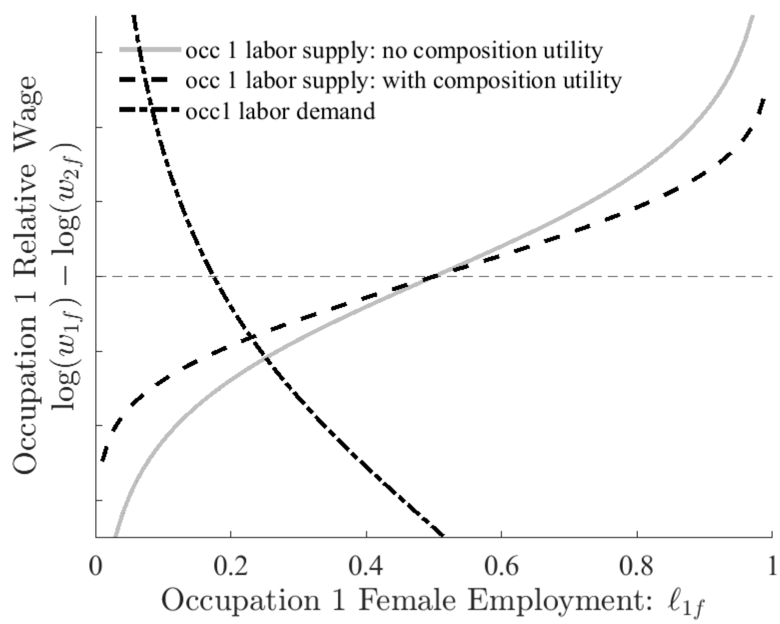


Figure 4: Equilibrium labor supply and demand

Note: This plot displays the occupation 1 female labor supply function with and without gender composition preferences and the occupation 1 equilibrium relative wage function. Here, I set $\nu = 0.1$, $\ell_{1m} = 0.5$, and $h_f(f) = 0.06 \cdot \log(f) - 0.5$ at baseline.

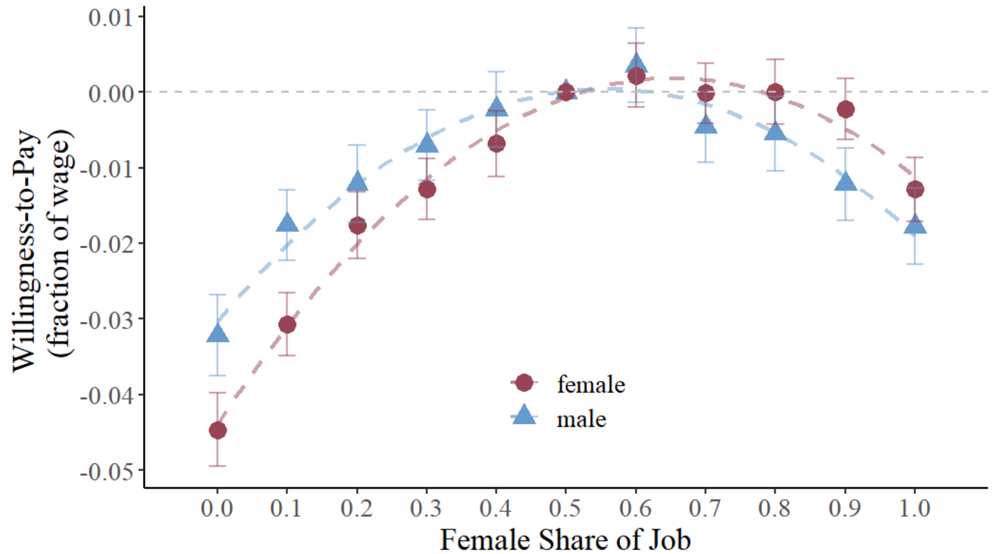


Figure 5: Average Willingness-to-Pay for Gender Composition

Note: This figure shows the willingness-to-pay, as a fraction of the wage, for each possible female share estimated on data from the conjoint job choice in my survey. Bars show 95% confidence intervals estimated using the delta method.

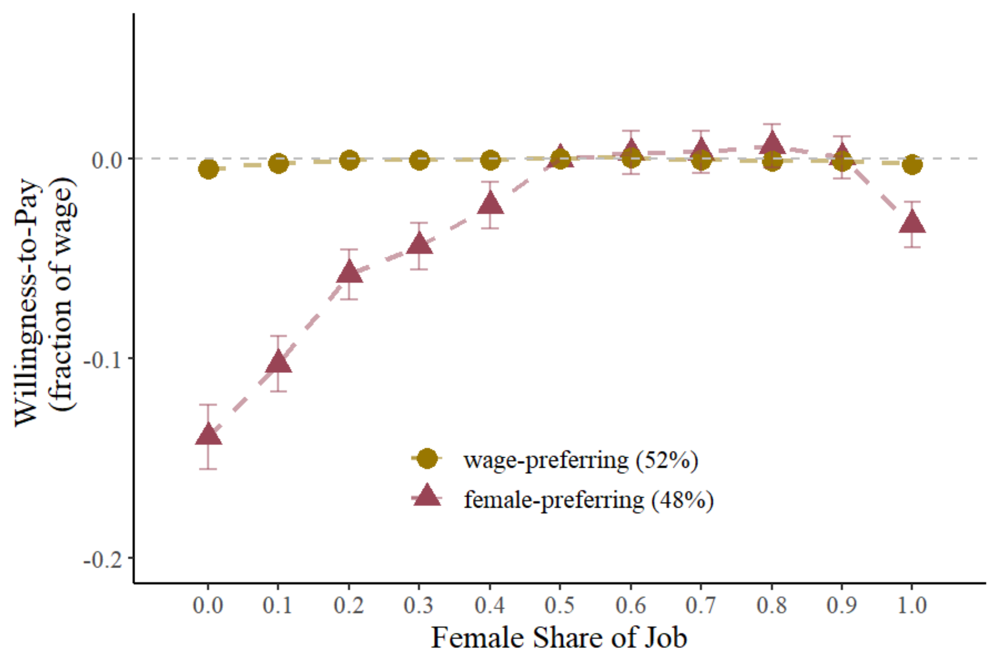


Figure 6: WTP for Gender Composition by Class, Female

Note: This figure shows the willingness-to-pay, as a fraction of the wage, for the two classes among women estimated using the latent class logit model described in Section III.C.2. Bars show 95% confidence intervals estimated using the delta method.

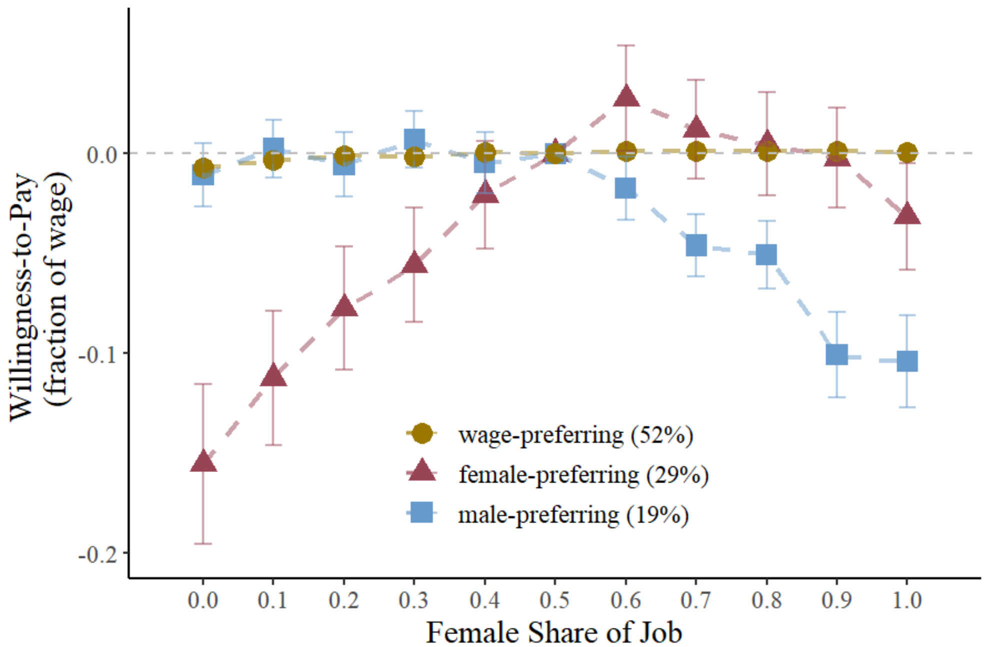


Figure 7: WTP for Gender Composition by Class, Male

Note: This figure shows the willingness-to-pay, as a fraction of the wage, for the three classes among men estimated using the latent class logit model described in Section III.C.2. Bars show 95% confidence intervals estimated using the delta method.

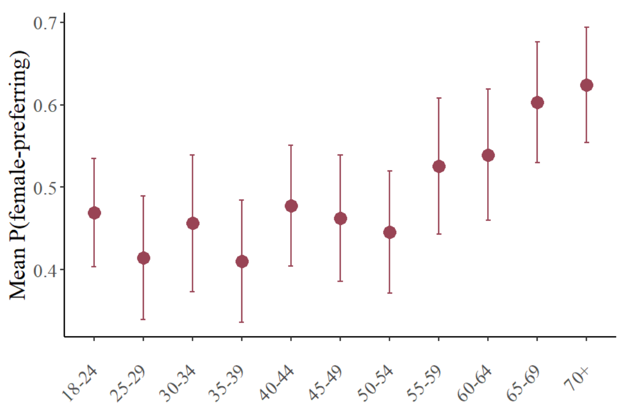


(a) Female

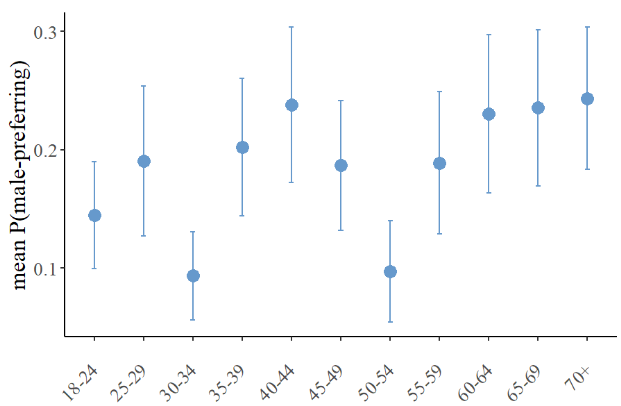


(b) Male

Figure 8: Age and Probability of Belonging to Own-Gender-Preferring Class

Note: This figure plots the average posterior probability of belonging to the own-gender-preferring preference class by age bin.

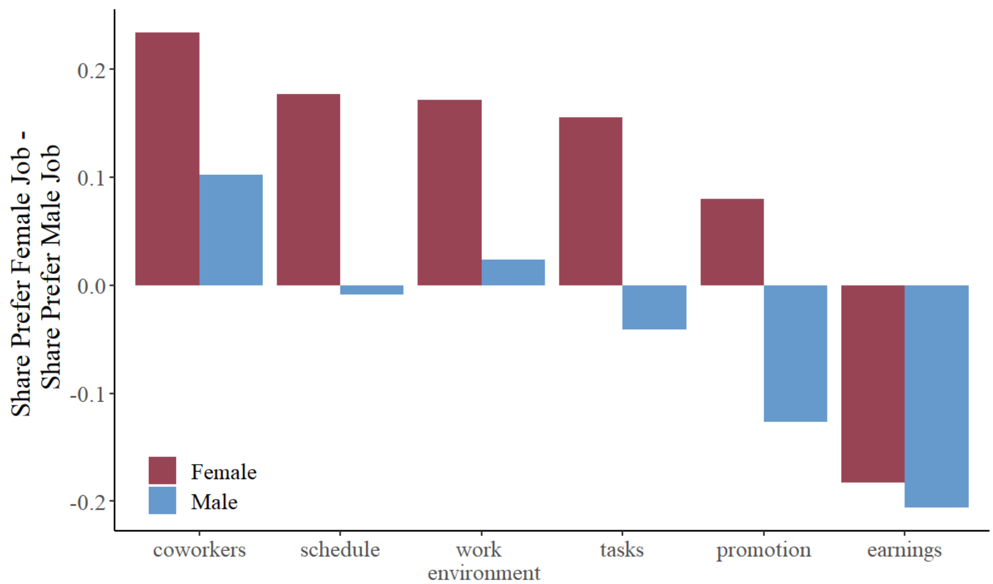


Figure 9: Share Preferring Female Job Minus Share Preferring Male Job on Each Attribute

Note: This figure plots the share of respondents who report that they would be more satisfied with a female job minus the share who report they would be more satisfied with a male job on the listed attribute, separately for men and women.

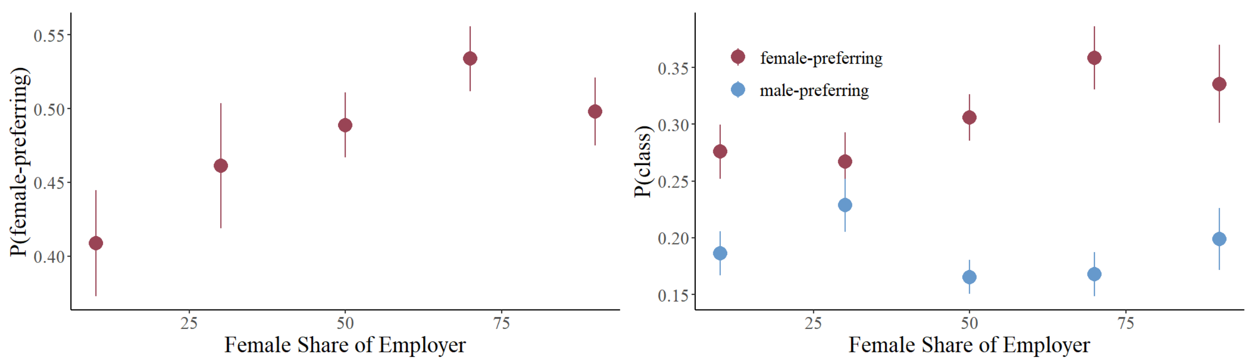


(a) Female

(b) Male

Figure 10: Most Likely Preference Class and Female Share of Employer

Note: This figure plots probability of belonging to each composition valuation class (omitting the wage-preferring class) against bins of the reported female share of the respondent's employer.

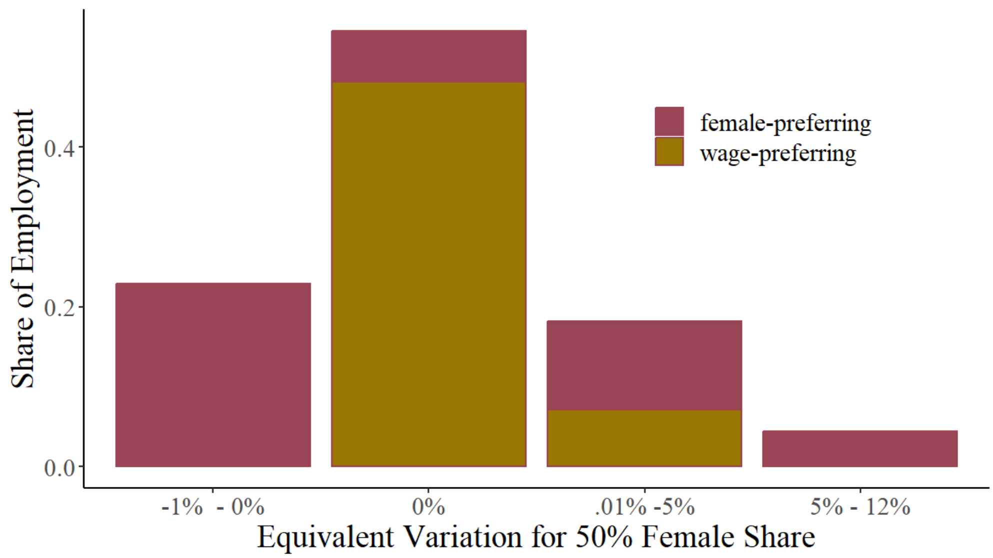


Figure 11: Distribution of Equivalent Variation for 50% Female Share - Women

Note: This figure shows the distribution of the wage change that is equivalent to one's occupation becoming 50% female for survey respondents, by preference class. I limit the sample to employed respondents.

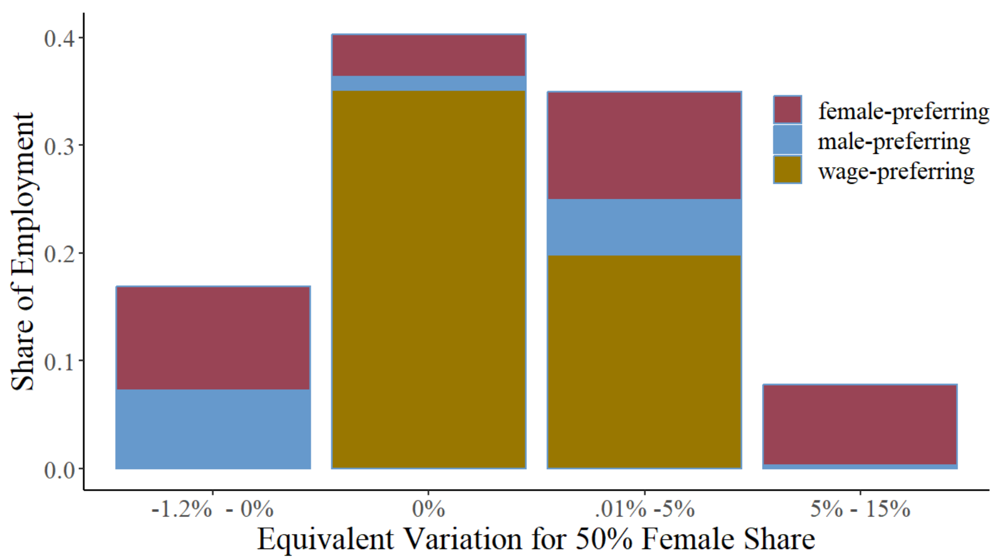


Figure 12: Distribution of Equivalent Variation for 50% Female Share - Men

Note: This figure shows the distribution of the wage change that is equivalent to one's occupation becoming 50% female for survey respondents, by preference class. I limit the sample to employed respondents.

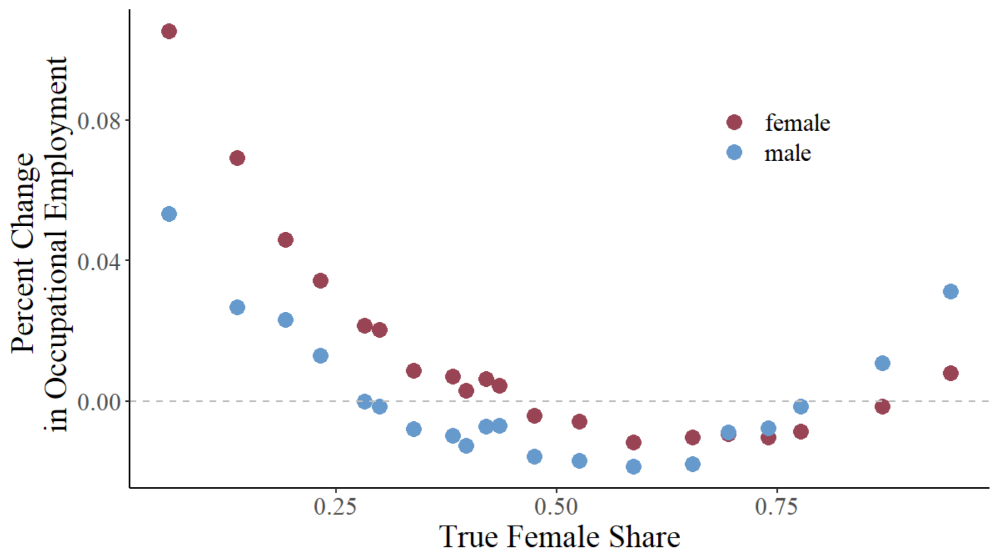


Figure 13: Change in Occupational Employment if Composition Preferences are Removed

Note: This figure shows a binscatter of the average percent change in occupational employment in a counterfactual allocation without composition preferences (counterfactual employment - true employment/true employment) against the average true female share of the occupations in that bin. The y-axis value of 0.01 corresponds to a 1% change. This model uses Census 2010 occupations and heterogeneous composition preferences. I impute employment in each Census occupation by class using the survey-reported aggregate occupations crosswalked with the distribution of Census occupations within aggregated occupations in the CPS.


Figure 14: Changes in Occupational Female Shares in Social Planner's Solution

Note: This figure shows results from the social planner's solution to maximize welfare by reallocating workers across occupations. Here, I use average valuations by gender. This figure displays a binscatter of the difference in the female share of an occupation (social planner - true) against the true female share of the occupation. Dot size corresponds to occupational employment.

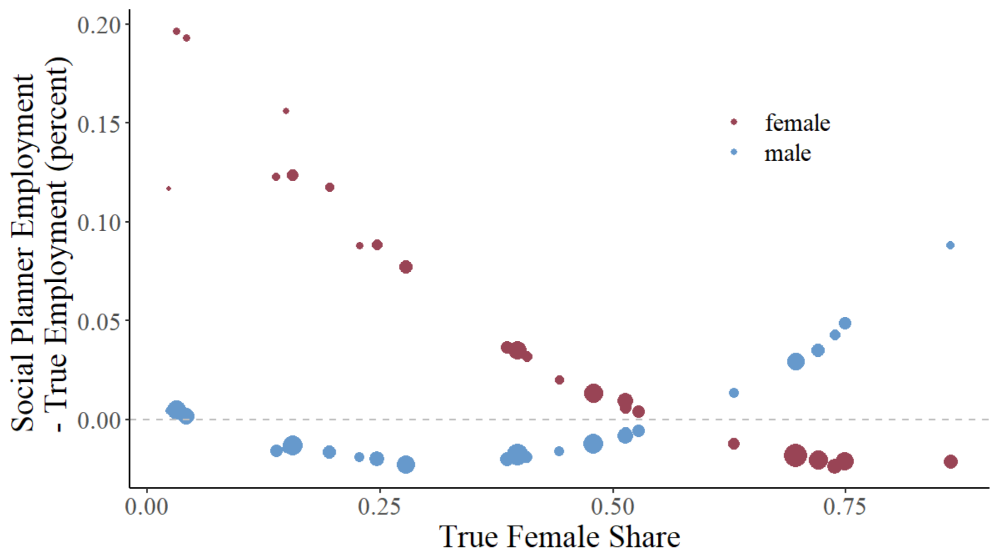


Figure 15: Changes in Occupational Employment in Social Planner's Solution

Note: This figure shows results from the social planner's solution to maximize welfare by reallocating workers across occupations. Here, I use average valuations by gender. This figure displays a binscatter of the difference in the employment of an occupation (social planner - true) against the true female share of the occupation. Dot size corresponds to occupational employment.

## Tables

Table 1: Survey Demographics vs. March CPS

| | CPS Share | Survey Share | Difference | t-stat |
|---|---|---|---|---|
| *Age* | | | | |
| <18 | .000 | .003 | .003 | |
| 18-24 | .112 | .123 | .011 | 3.281 |
| 25-34 | .170 | .195 | .025 | 5.835 |
| 35-44 | .138 | .197 | .059 | 13.928 |
| 45-54 | .158 | .166 | .008 | 2.000 |
| 55-64 | .182 | .151 | -.031 | -8.246 |
| 65-74 | .139 | .127 | -.012 | -3.427 |
| 75-84 | .072 | .034 | -.037 | -19.349 |
| 85+ | .029 | .004 | -.025 | -35.484 |
| *Education* | | | | |
| less than high school | .096 | .036 | -.060 | -29.968 |
| bachelor's degree or greater | .329 | .319 | -.010 | -2.077 |
| high school diploma | .292 | .279 | -.013 | -2.817 |
| some college | .283 | .363 | .080 | 15.670 |
| *Race* | | | | |
| White or Caucasian | .733 | .706 | -.027 | -5.505 |
| Hispanic or Latino | .108 | .081 | -.027 | -9.444 |
| Black or African American | .094 | .124 | .030 | 8.465 |
| Asian or Pacific Islander | .049 | .049 | .000 | .110 |
| Multiracial or Biracial | .010 | .018 | .008 | 5.726 |
| Native American or Alaskan Native | .006 | .014 | .008 | 6.389 |
| other race | .000 | .008 | .008 | |
| *Gender* | | | | |
| female | .521 | .520 | -.001 | -.221 |
| male | .479 | .476 | -.003 | -.588 |
| other | .000 | .004 | .004 | |

Note: This table compares the distribution of demographic characteristics of the sample for my survey conducted via Lucid in October 2021 and the March CPS samples from 2014 through 2019. Lucid targets the distributions of these characteristics, although not necessarily at the level of detail shown here.

Table 2: Willingness-to-Pay for Demographic Composition

| | all | female | male | < 40 | > 40 | no kids | has kids |
|---|---|---|---|---|---|---|---|
| Share Female (baseline: 50%) | | | | | | | |
| 0% | -.039 | -.045 | -.032 | -.031 | -.044 | -.042 | -.033 |
| | (-.043,-.036) | (-.049,-.04) | (-.038,-.027) | (-.037,-.026) | (-.049,-.04) | (-.047,-.038) | (-.039,-.027) |
| 10% | -.025 | -.031 | -.018 | -.021 | -.027 | -.027 | -.02 |
| | (-.028,-.022) | (-.035,-.027) | (-.022,-.013) | (-.026,-.016) | (-.031,-.023) | (-.031,-.023) | (-.025,-.014) |
| 20% | -.015 | -.018 | -.012 | -.015 | -.015 | -.016 | -.013 |
| | (-.019,-.012) | (-.022,-.013) | (-.017,-.007) | (-.021,-.01) | (-.019,-.011) | (-.021,-.012) | (-.019,-.007) |
| 30% | -.01 | -.013 | -.007 | -.011 | -.01 | -.011 | -.01 |
| | (-.014,-.007) | (-.017,-.009) | (-.012,-.002) | (-.016,-.007) | (-.014,-.006) | (-.015,-.007) | (-.015,-.004) |
| 40% | -.005 | -.007 | -.002 | -.006 | -.004 | -.005 | -.004 |
| | (-.008,-.001) | (-.011,-.002) | (-.007,.003) | (-.011,-.001) | (-.008,0) | (-.009,-.001) | (-.01,.001) |
| 60% | .003 | .002 | .004 | -.001 | .005 | .004 | .001 |
| | (0,.006) | (-.002,.006) | (-.001,.008) | (-.006,.004) | (.001,.009) | (0,.008) | (-.005,.006) |
| 70% | -.002 | 0 | -.005 | -.002 | -.002 | -.002 | -.002 |
| | (-.005,.001) | (-.004,.004) | (-.009,0) | (-.007,.003) | (-.006,.002) | (-.006,.002) | (-.007,.003) |
| 80% | -.002 | 0 | -.006 | 0 | -.004 | -.002 | -.002 |
| | (-.006,.001) | (-.004,.004) | (-.01,-.001) | (-.005,.006) | (-.008,0) | (-.006,.001) | (-.008,.004) |
| 90% | -.007 | -.002 | -.012 | -.004 | -.009 | -.008 | -.005 |
| | (-.01,-.004) | (-.006,.002) | (-.017,-.007) | (-.008,.001) | (-.013,-.005) | (-.012,-.004) | (-.01,.001) |
| 100% | -.015 | -.013 | -.018 | -.008 | -.019 | -.018 | -.009 |
| | (-.018,-.012) | (-.017,-.009) | (-.023,-.013) | (-.013,-.003) | (-.024,-.015) | (-.022,-.014) | (-.015,-.004) |
| Share with Kids (baseline: 50%) | | | | | | | |
| 30% | .001 | 0 | .002 | .001 | .001 | .004 | -.005 |
| | (-.001,.004) | (-.003,.004) | (-.001,.006) | (-.003,.005) | (-.002,.004) | (.001,.007) | (-.009,-.001) |
| 70% | .002 | .005 | -.002 | .002 | .001 | 0 | .006 |
| | (-.001,.004) | (.002,.008) | (-.005,.002) | (-.002,.006) | (-.002,.005) | (-.003,.003) | (.002,.01) |
| Share Under 40 Years Old (baseline: 50%) | | | | | | | |
| 30% | .003 | .003 | .004 | .002 | .004 | .004 | .002 |
| | (.001,.006) | (0,.006) | (0,.008) | (-.001,.006) | (.001,.007) | (.001,.007) | (-.002,.006) |
| 70% | -.006 | -.007 | -.006 | .001 | -.011 | -.008 | -.003 |
| | (-.009,-.004) | (-.01,-.004) | (-.009,-.002) | (-.003,.004) | (-.014,-.007) | (-.011,-.005) | (-.007,.001) |
| lefthand job | .005 | .005 | .006 | .005 | .005 | .005 | .007 |
| | (.004,.006) | (.004,.006) | (.004,.007) | (.004,.007) | (.004,.007) | (.003,.006) | (.005,.009) |
| num. obs. | 29972 | 16495 | 13334 | 11215 | 18757 | 20355 | 9617 |
| num. indiv. | 2772 | 1525 | 1234 | 1037 | 1735 | 1889 | 883 |

Note: this table shows the willingness-to-pay for each share of each demographic group estimated using data from the conjoint job choice in my survey. The logit coefficients that underlie these estimates are shown in Appendix Table B.1, which also includes the number of observations and individuals in each regression.

Table 3: Model Parameter Sources

| **Parameter** | **Meaning** | Value | Source |
|---|---|---|---|
| | Occupation Characteristics | | |
| $w_{g,k}$ | Occupation-Gender wage | | March CPS |
| $\ell_{g,k}$ | Occupation-Gender allocation | | March CPS |
| | Worker Utility Parameters | | |
| $h_g(\ell_{f,k}/\ell_k)$ | Gender composition valuation function | | Survey |
| $1-\lambda$ | Correlation of EV shocks within nest | .5 | Dix-Carneiro (2014), Bassier et al. (2022) |
| | Externally Calibrated Parameters | | |
| $\eta$ | Sub. Elasticity Across Occupations | 1.5 | Acemoglu and Autor (2011) |
| $\alpha$ | Sub. Elasticity Across Genders | 2.5 | Ngai and Petrongolo (2017) |

Note: This table shows the sources of parameters that are determined outside the model.

Table 4: Conditions for Tipping with Estimated Composition Preference, Female

| | female | |
|---|---|---|
| | average | female preferring |
| composition pref. $h_g(f)$ | $-.04 + .14f - .11f^2$ | $-.13 + .41f - .30f^2$ |
| max $\lambda$ | 0.0224 | 0.0693 |
| $\lambda = 0.5$, min wtp multiplier | 22.3 | 7.2 |

Note: This table shows the fitted polynomials for female composition preferences, and the conditions on preferences and the inverse wage elasticity $\lambda$ necessary to create a tipping point. The inverse of the maximum $\lambda$ is the smallest wage elasticity needed to create a tipping point.

Table 5: Conditions for Tipping with Endogenous Amenity

| | female | male |
|---|---|---|
| composition pref. $h_g(f)$ | $-.04 + .14f - .11f^2$ | $-.02 + .11f - .10f^2$ |
| min $\lvert\alpha\rvert$, $\lambda = 0.5$ | 0.5123 | .5059 |

Note: This table shows the fitted polynomials for female and male average composition preferences, and the conditions on the endogenous amenity function necessary to create a tipping point. The parameter $\alpha$ measures the responsiveness of amenity levels to the level of the female share. A value of $\alpha = 0.5$ indicates that a 10 percentage point increase in the female share would increase the value of amenities by the same amount as a 5% wage increase.

Table 6: Average Equivalent Variation for 50% Female Share in All Occupations

| | | female | male |
|---|---|---|---|
| aggregate | average preference | .0034 | .0111 |
| | heterogeneous preferences | .0036 | .0098 |
| by class | wage-preferring | .0004 | .0014 |
| | female-preferring | .0074 | .0278 |
| | male-preferring | – | .0026 |

Note: This table shows the equivalent variation for one's occupation becoming half female. The equivalent variation is the wage increase in the case with the true female share that would make utility equal to the case with a 50% female share.

# Appendix A Data

## A.1 Survey Design

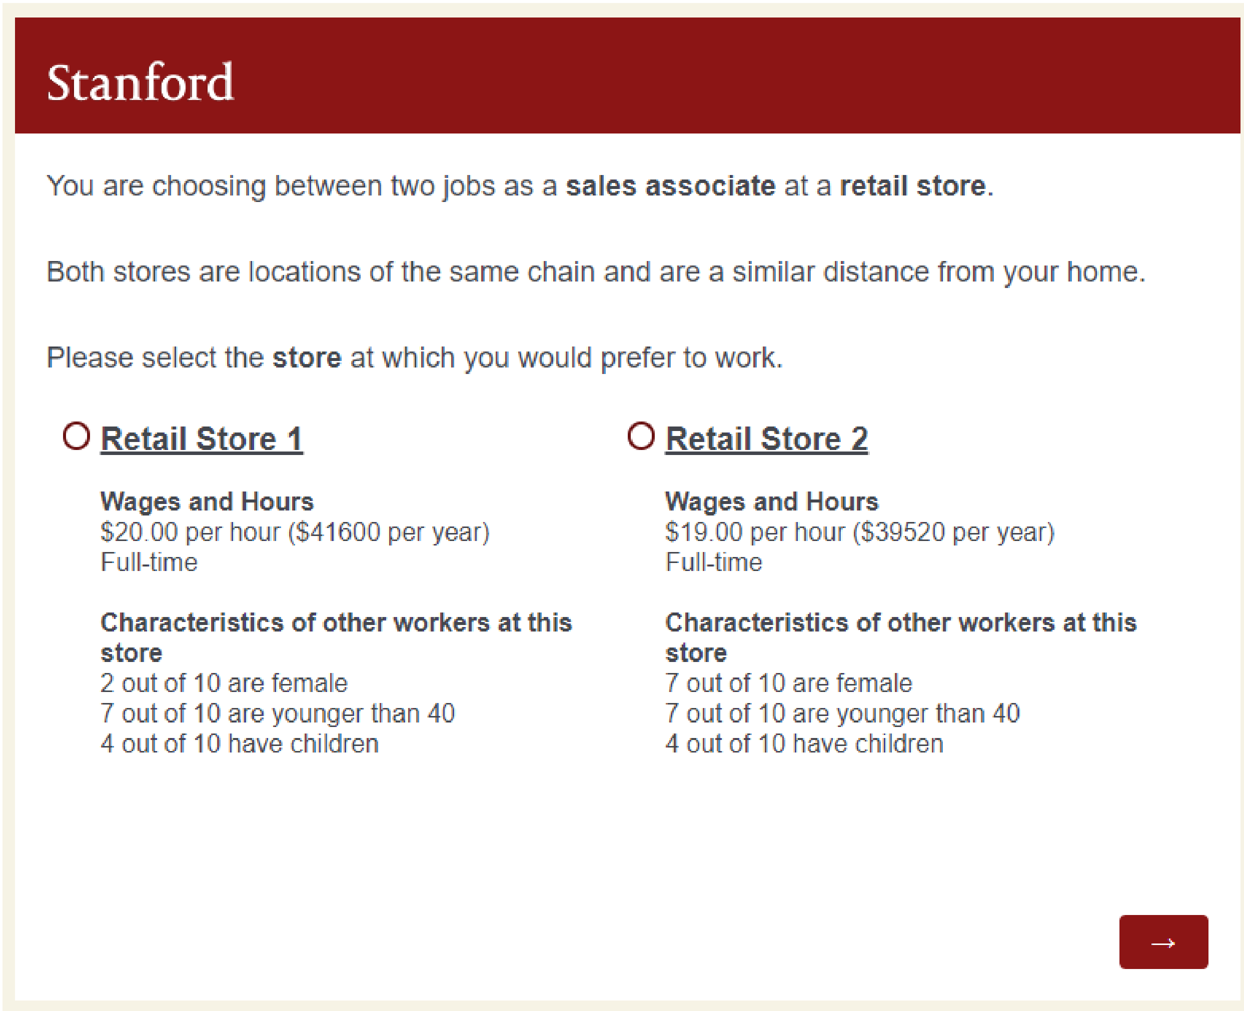


Figure A.1: Example Hypothetical Job Choice

## A.2 Survey Descriptive Statistics

A primary concern with online survey samples is representativeness. The respondents are a convenience sample of participants who have opted in to an online survey panel or website. This sample is inherently non-random. The primary concern with such a sample is whether survey results can be generalized to a broader population, i.e., working adults in the United States. Reassuringly, several papers in political science have conducted survey experiments using both online convenience samples from Lucid and similar providers and traditional random samples and found similar results.[38]

Whether survey results can be generalized to a larger population relies on both the presence of different types of people in the survey sample and the heterogeneity of the treatment effect of interest, here, gender composition preferences. If there is no population-level heterogeneity in gender composition preferences, the survey sample does not need to be demographically representative to make inferences about the general population. Either way, the average preference estimates for the survey and the population will be the same. If, on the other had, gender composition preferences are highly variable in the population, we must understand which individual characteristics co-vary with these preferences to re-weight and estimate population-level statistics. In this case, it is again not necessary that the survey sample exactly matches the population distribution of relevant traits, only that there is sufficient presence of each group in the sample to re-weight to match the population.

To assess whether my survey has sufficient representation across demographic groups, I compare my survey sample to the CPS Annual Social and Economic Supplement (CPS ASEC, commonly referred to as the March CPS) pooled from 2014-2019. I first compare my sample to the CPS ASEC by sex, age, race, and education in Table 1. My conjoint experiment measures preferences for sex, age, and educational composition of an occupation, so these characteristics are likely to be important covariates of composition preferences. These characteristics, in addition to race, are targeted by Lucid Theorem.

[38]Coppock and McClellan (2019) compare results from several survey experiments conducted through Lucid, MTurk, and probability samples. For all but one of five experiments, they find experimental effects that matched the sign and significance of the original estimates on the probability sample. Boas et al. (2020) compare Qualtrics, which recruits survey participants in a similar manner to Lucid, with probability samples and samples recruited through Facebook and MTurk and found that the Qualtrics sample was the most demographically and politically representative. They also find that treatment effects and in particular treatment effect heterogeneity by partisanship are similar in the benchmark probability samples (from the General Social Survey and YouGov Polimetrix) and the Qualtrics sample.

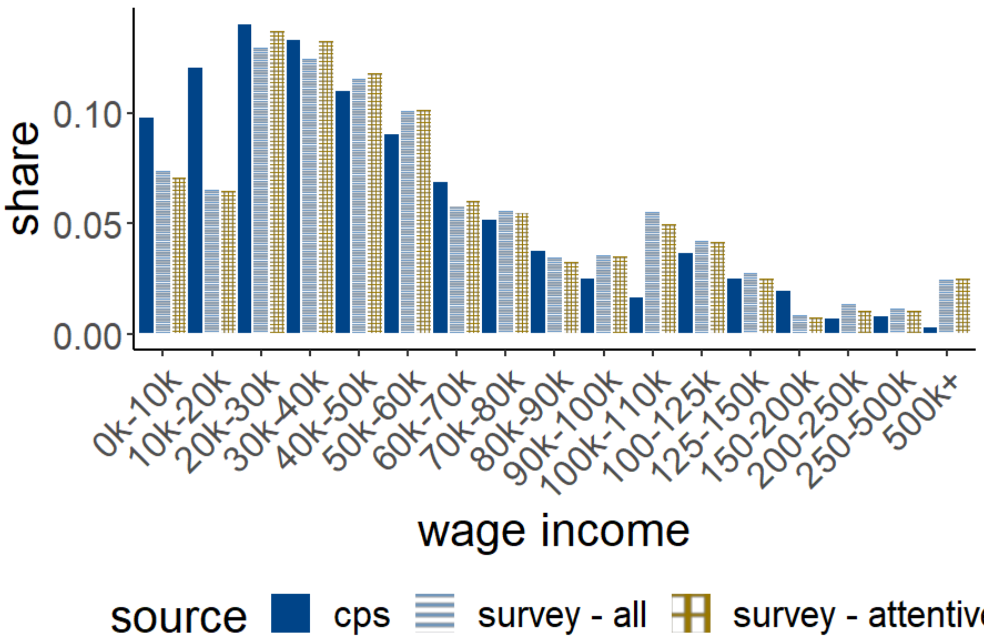


Figure A.2: Income Distribution: Survey vs. March CPS

This plot shows the distribution of individual annual wage income in the CPS ASEC from 2014 through 2019 and my survey. The distribution of incomes is overall similar, although my survey sample has less mass at the lowest incomes and more mass at the highest incomes.

Figure A.2 compares the distribution of reported annual income in my survey to that in the CPS ASEC. Relative to the CPS, respondents to my survey reported incomes that were higher on average; in particular, the lowest income groups were somewhat under-represented. On the whole, however, the distributions are relatively similar, which is reassuring given that one might expect income to be mis-reported in surveys. One might expect income to co-vary with gender composition preferences because female-dominated occupations have lower income on average. Thus, workers with a strong preference for those jobs might tend to have lower incomes.

Next, I compare the distribution of occupational female shares among survey respondents to that in the CPS. Naturally, I expect people who prefer more female occupations to work in more female occupations and vice versa. Figure A.3 compares the distribution of employment by occupational female shares in my survey and the CPS. Here, I calculated the survey respondents' occupational female shares by linking their reported occupation to that occupation in the CPS and calculating the occupation's female share in the CPS.[39] The distribution of female shares in survey-reported occupations is overall relatively similar to

[39]Specifically, I use the groups of the IPUMS CPS variable occ10ly, which are approximately equivalent to 2-digit SOC codes.

that in the CPS with representation across the spectrum of female shares. In Appendix A figure A.5, I compare the CPS occupational female share distribution to survey respondents' reported female shares of their employers and coworkers and the reported gender perception of their jobs and find similar results.

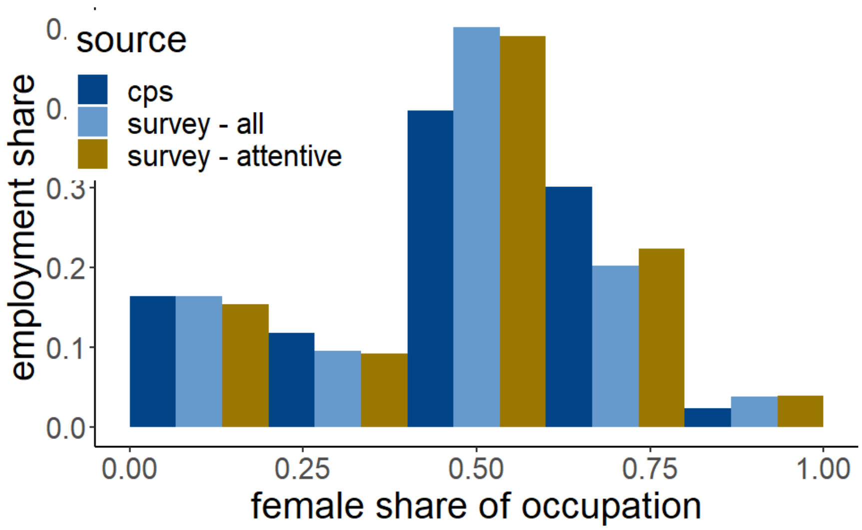


Figure A.3: Female shares: Survey vs. CPS

This plot shows the distribution of occupational female shares in the CPS ASEC from 2014 through 2019 and my survey. I calculate the female share of a respondent's occupation by linking their reported occupation to that occupation in the CPS. I group occupations at the level of the group headings of the IPUMS CPS variable *occ10ly*.

Finally, I compared the distribution of attitudes towards gender and work in my survey to surveys with population-representative random samples. Gender composition preferences may be related to an individual's view on gender roles and whether men and women should perform different tasks on the job and at home. Thus, I asked my survey respondents two questions from the General Social Survey (GSS) on gender roles and work and gender-based affirmative action and one question from the Pew American Trends Panel about gender attitudes.

In the first question from the GSS, shown in Figure A.4 Panel a, I asked survey respondents whether they agree or disagree with the following statement:

> *"It is much better for everyone involved if the man is the achiever outside the home and the woman takes care of the home and family."*

Relative to the GSS respondents, my survey respondents were more likely to agree or strongly agree with this statement, potentially indicating more conservative attitudes towards gender roles and work.

In the second question from the GSS, shown in Figure A.4 Panel b, I asked respondents the following question about gender-based affirmative action:

> *"Some people say that because of past discrimination, women should be given preference in hiring and promotion. Others say that such preference in hiring and promotion of women is wrong because it discriminates against men. What about your opinion - are you for or against preferential hiring and promotion of women?"*

Here, respondents to my survey were more likely to select the less-extreme options relative to GSS respondents, suggesting weaker preferences in either direction regarding gender-based affirmative action.

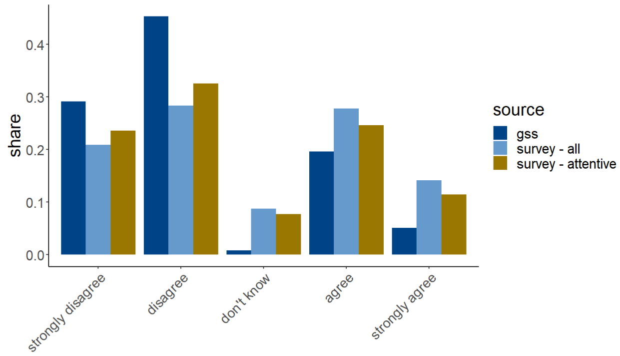


(a) GSS: Man Should Work, Woman Stay at Home

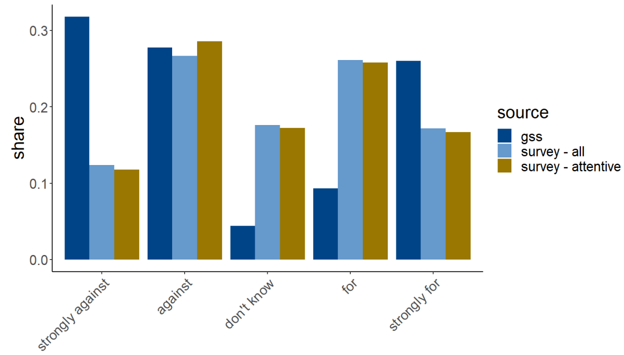


(b) GSS: Affirmative action for women in jobs

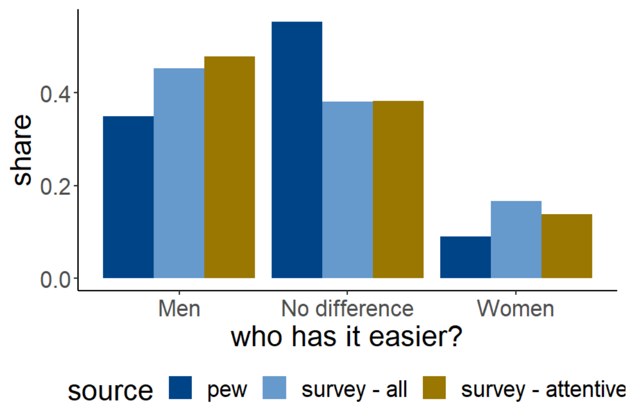


(c) Pew: Which Gender has it Easier?

Figure A.4: Gender Attitudes vs. Representative Surveys

These plots show the distribution of answers to three questions about gender attitudes in population representative surveys in green and my survey in coral. The first two questions come from the General Social Survey and the third from the Pew American Trends Panel. Descriptions of each question are detailed in the text.

In the last question about gender attitudes, from the Pew American Trends Panel, shown in Figure A.4 Panel c, I asked survey respondents the following question:

*"All things considered, who do you think has it easier in our country these days?"*

My survey respondents were slightly more likely to answer that women or men have it easier and less likely to answer no difference.

The answers to these questions suggest that the attitudes toward gender and work in my survey sample might differ from the general population, but not in straightforward ways. Respondents report more conservative attitudes toward gender roles, but more neutral stances toward affirmative action, than GSS respondents. Relative to the Pew survey, my respondents have more extreme attitudes toward gender-based privilege. Nonetheless, it is reassuring that there is representation across the spectrum of gender attitudes among my survey respondents.

Overall, comparing the composition of my survey sample to surveys with population-representative random samples indicates that the survey sample is representative enough of the general population. My survey sample has similar distributions of age, education, race, sex, income, and occupational female shares to the general population. Attitudes toward gender and work in my sample seem to differ somewhat from the population, but there is large variation in reported gender attitudes, which makes re-weighting possible. I am currently working on providing estimates of gender composition preferences that are re-weighted by demographic characteristics.


(a) Female Share of Employer

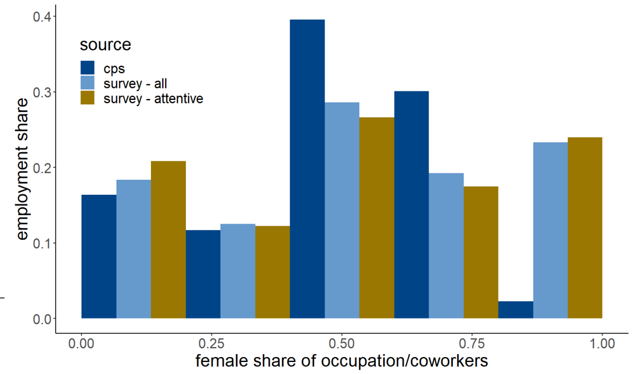


(b) Female Share of Coworkers with Same Job

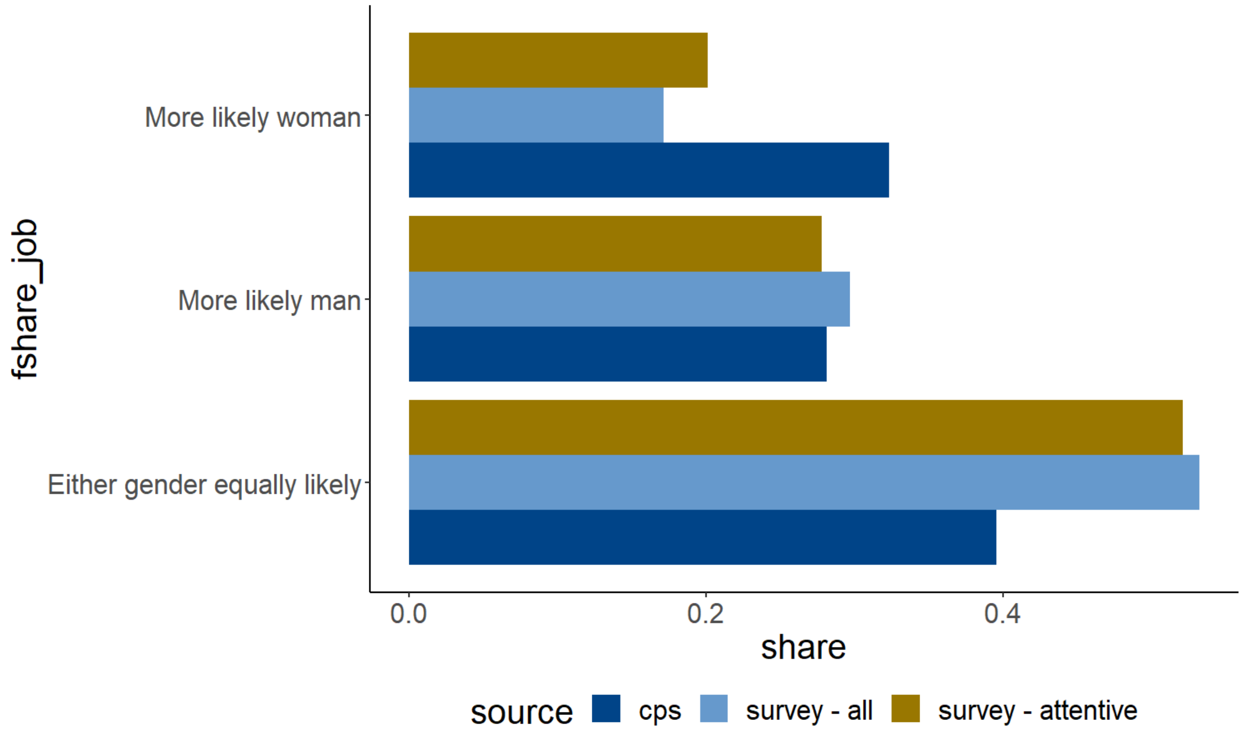


(c) Reported Gender Norm of Job

Figure A.5: Reported Female shares: Survey vs. CPS

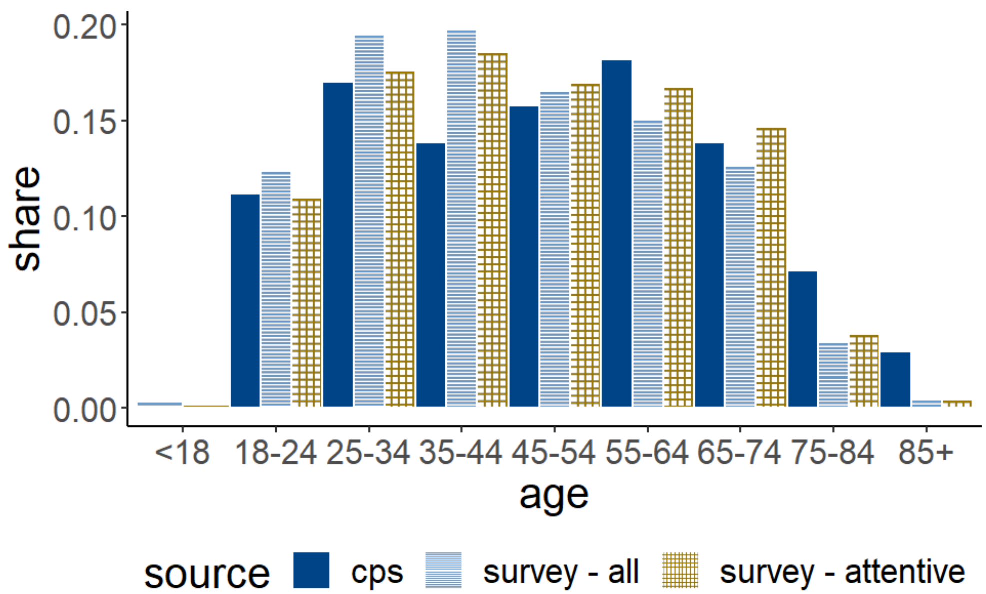


Figure A.6: Age Distribution: Survey vs. March CPS

This plot shows the distribution of ages in the CPS ASEC from 2014 through 2019 and my survey.

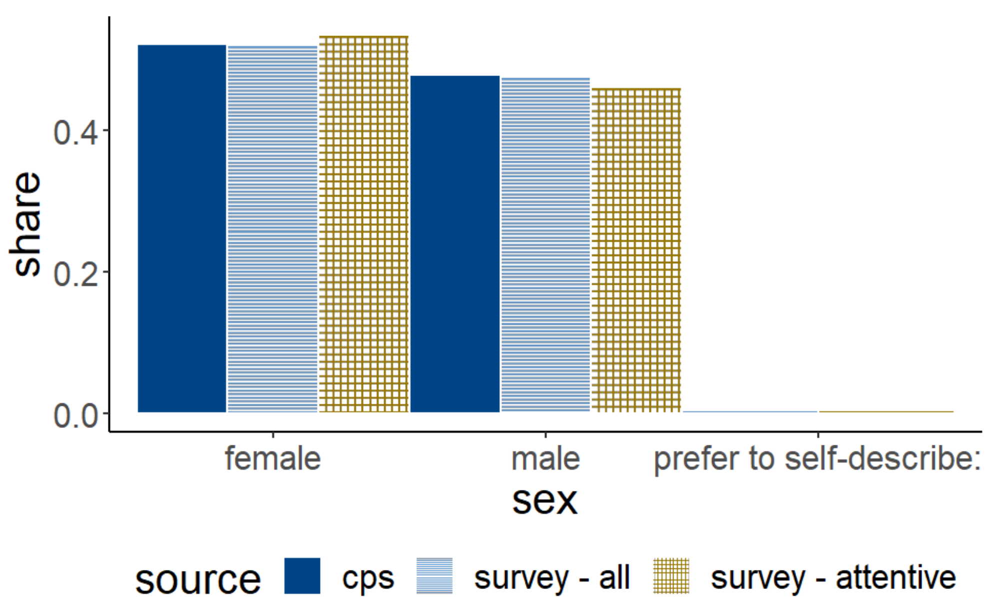


Figure A.7: Sex Distribution: Survey vs. March CPS

This plot shows the distribution of sex in the CPS ASEC from 2014 through 2019 and my survey.

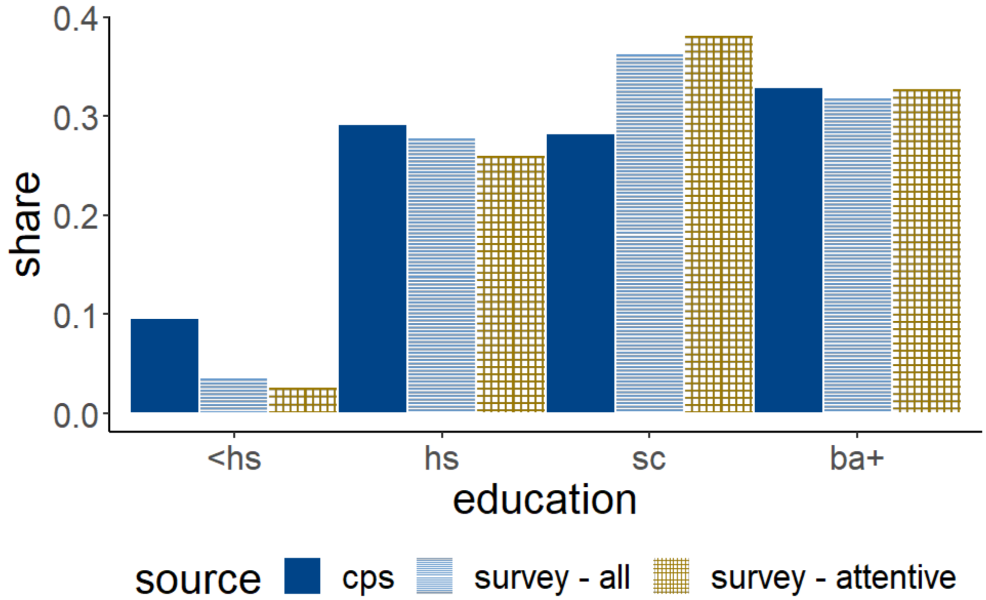


Figure A.8: Education Distribution: Survey vs. March CPS

This plot shows the distribution of education in the CPS ASEC from 2014 through 2019 and my survey.

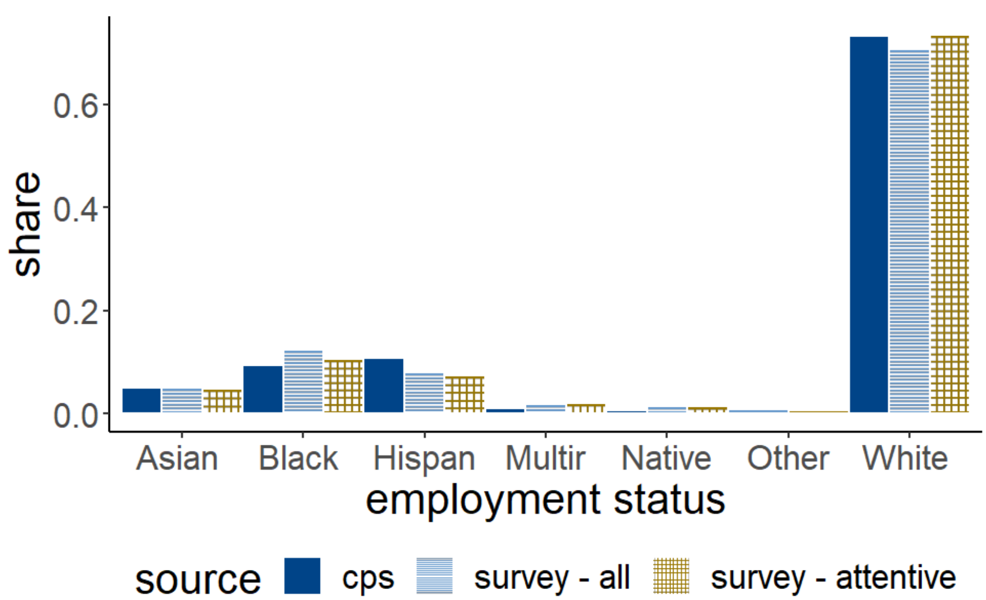


Figure A.9: Race Distribution: Survey vs. March CPS

This plot shows the distribution of race in the CPS ASEC from 2014 through 2019 and my survey.

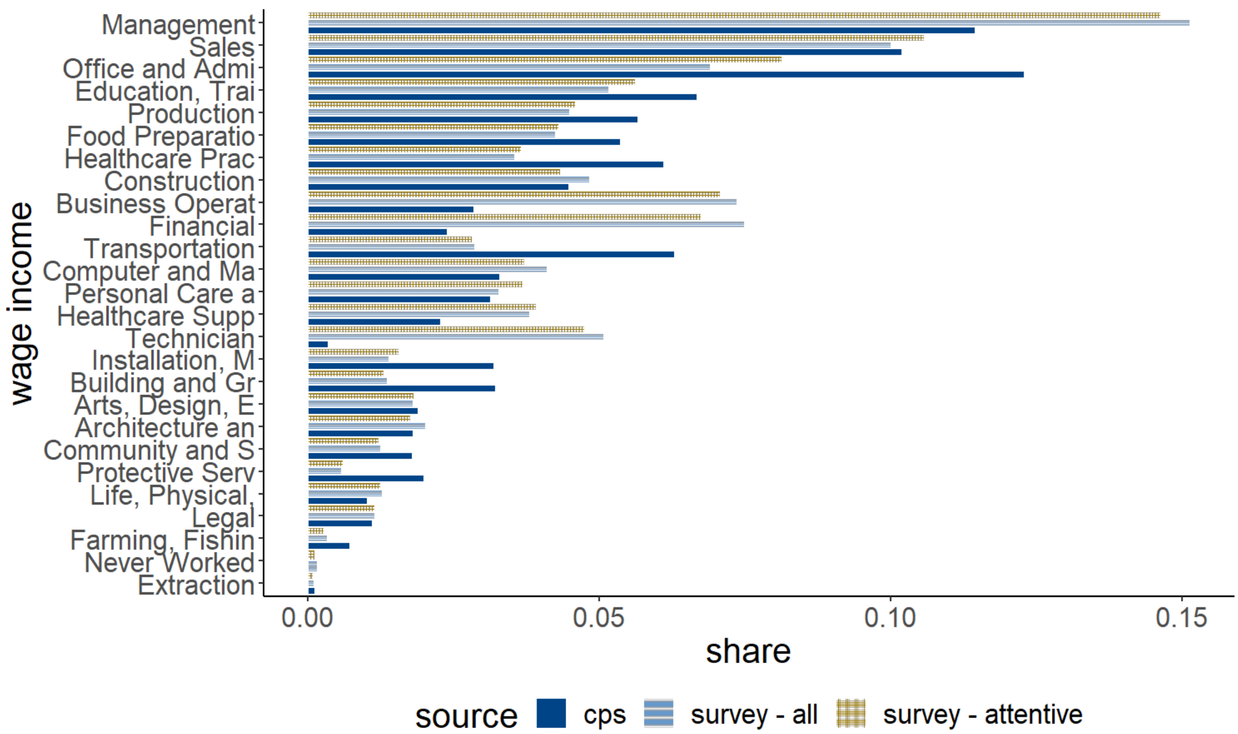


Figure A.10: Occupation Distribution: Survey vs. March CPS

This plot shows the distribution of occupations among employed respondents in the CPS ASEC from 2014 through 2019 and my survey.

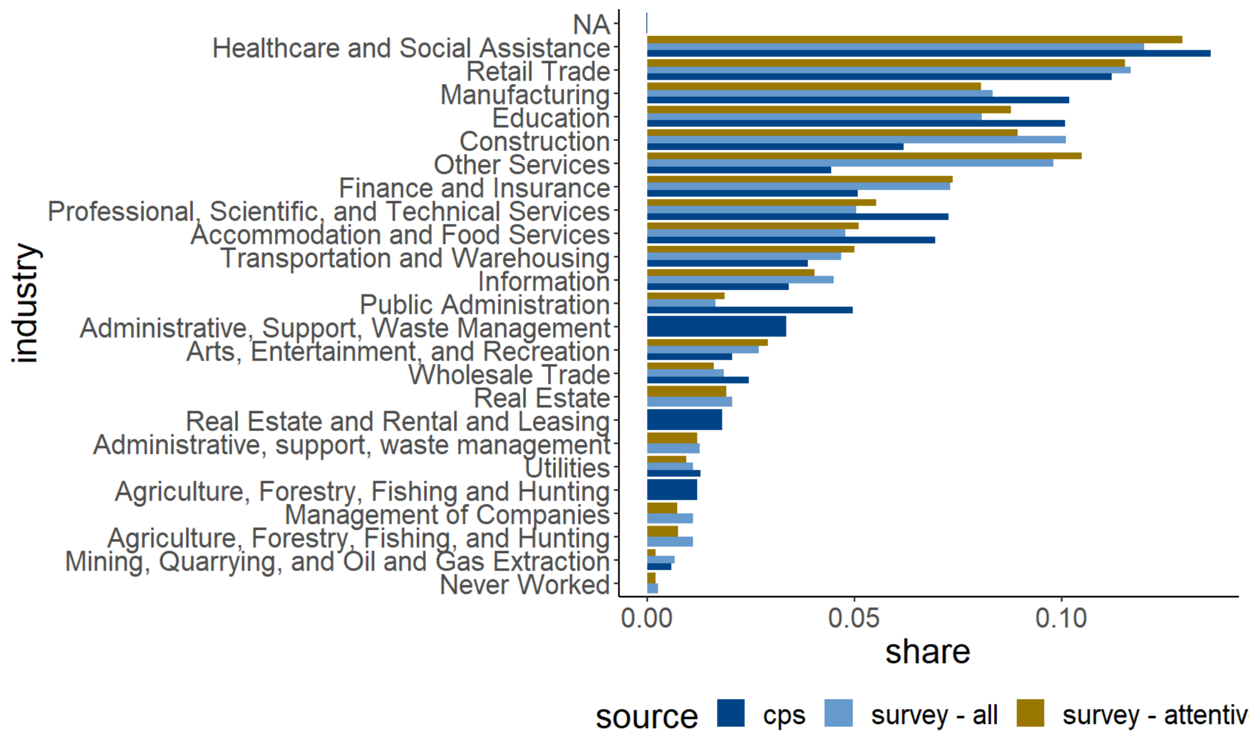


Figure A.11: Industry Distribution: Survey vs. March CPS

This plot shows the distribution of industries among employed respondents in the CPS ASEC from 2014 through 2019 and my survey.

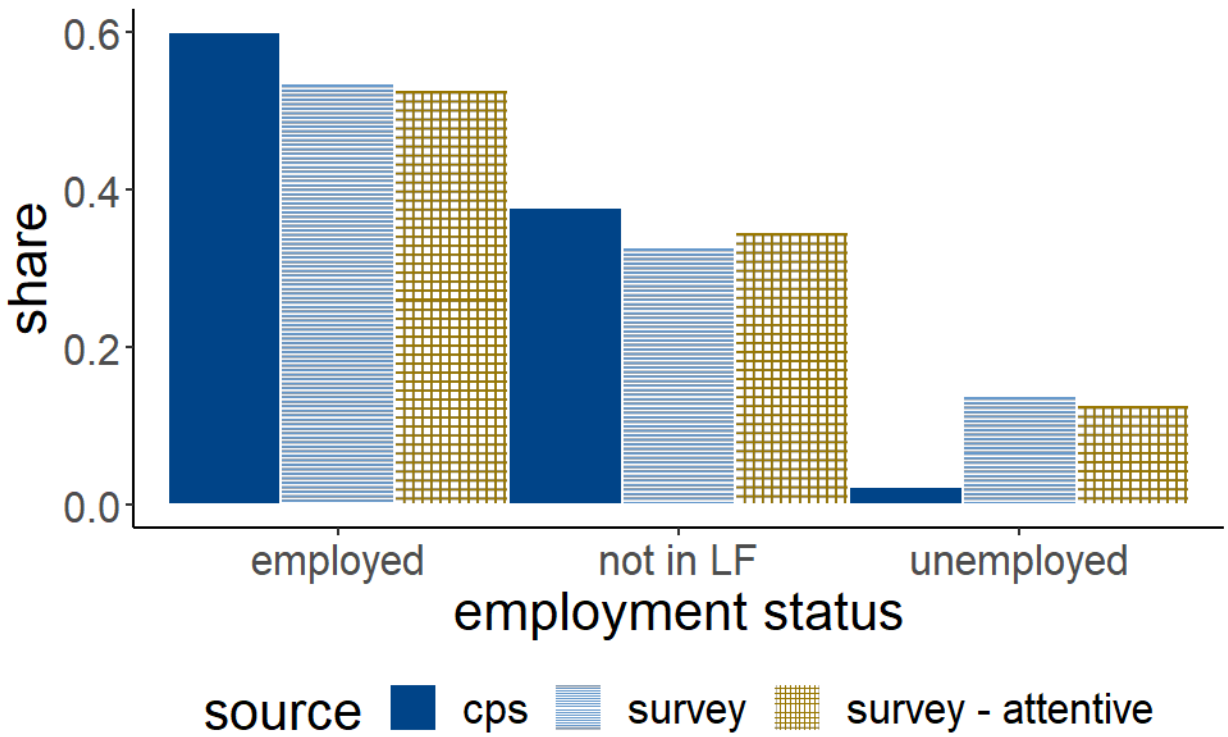


Figure A.12: Employment Status Distribution: Survey vs. March CPS

This plot shows the distribution of employment status among respondents in the monthly CPS averaged over 2022.

### A.3 Data Quality

A primary concern with online surveys is that respondents may be inattentive and not provide answers based on true preferences. I include several attention checks within my survey and find that the majority of respondents are attentive, and rates of attention are similar to incentivized job choice experiments found in the literature.

The simplest check for inattentive respondents is to look for people who complete the survey too quickly. Figure A.13 shows the distribution of the time taken to complete the survey in minutes. The average respondent completed the survey in 10.6 minutes, and the median respondent completed the survey in 7.6 minutes. There is a long right tail in the distribution of response times, but this is not concerning, as respondents may stop the survey midway and return to it later.

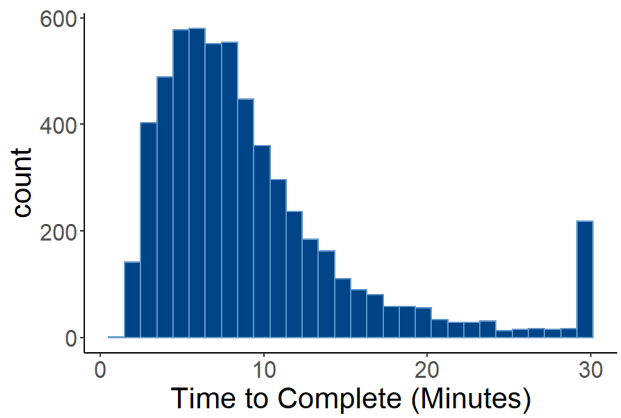


Figure A.13: Distribution of Survey Completion Time

This plot shows the distribution of time taken to complete my survey in minutes. Completion times are truncated at 30 minutes. The average completion time is 10.6 minutes, and the median is 7.6 minutes.

My survey included three attention checks, which were placed at the beginning, middle, and end of the survey instrument. Attention checks are a common method in survey design to distinguish low-effort respondents, whose responses may be discarded or down-weighted in the ultimate analysis.[40]

At the beginning and end of the survey, I had two basic attention checks. During the demographic information section, I included a question that simply asks respondents to write the number "13" in a text box. Responses to this question are shown in Figure A.14 Panel a. At the end of the survey, among the questions about gender attitudes, I included a multiple choice question with an agreement scale, where the question asks respondents to select "disagree." Responses to this question are shown in Figure A.14 Panel b. For both questions, the vast majority of respondents passed the attention check. As is expected given respondent fatigue, the rate of correct responses was lower for the attention check at the end of the survey.

[40] Alvarez et al. (2019) found that respondents who failed attention checks completed surveys faster, had a higher incidence of "don't know" responses, and provided lower-intensity responses regarding attitudes and behavior.

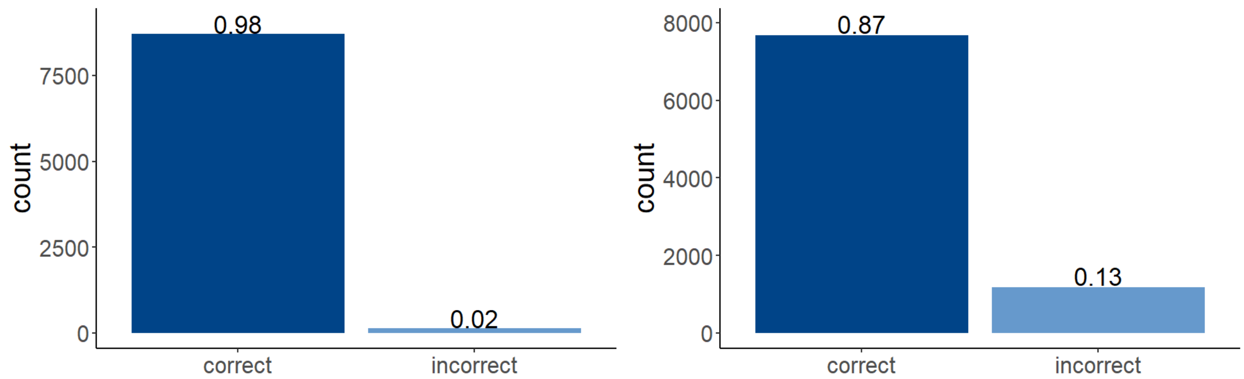


(a) Text Entry (Beginning)

(b) Multiple Choice (End)

Figure A.14: Attention Checks

These plots show the share of survey respondents answering the attention check questions correctly and incorrectly. In the text entry question, shown in Panel a, respondents are given a text entry box and instructed *"As a data quality check, please type the number 13 in the box."*. In the multiple choice question, shown in Panel b, respondents are given a multiple choice scale from strongly agree to strongly disagree and instructed *"Please select 'disagree' for data quality control."*

I also included an attention check within the hypothetical job choice conjoint, where respondents saw a choice between two jobs that differed only in their wages. As we see in Figure A.15, 85% of respondents choose the higher wage job. This is a similar rate of inattention to that found by Mas and Pallais (2017), who found that 14.5% of individuals chose the lower-paying job given a choice between two otherwise identical jobs to apply for. It is reassuring that the rates of inattention are similar, given Mas and Pallais (2017) had an incentivized choice between actual jobs and my survey is purely hypothetical and not incentivized. Results of my survey are similar whether or not I include respondents that failed one or more attention check.

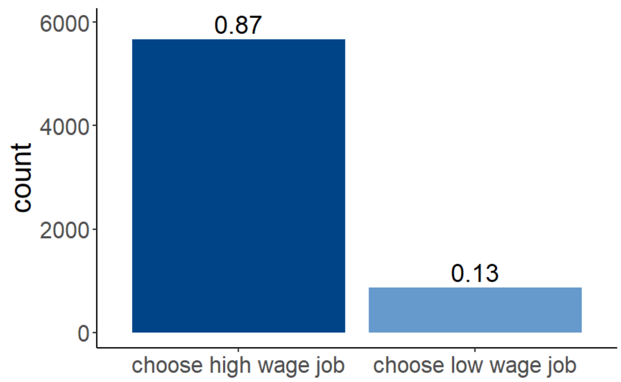


Figure A.15: Hypothetical Job Choice Attention Check

This plots show the share of survey respondents who choose the higher and lower wage job in a choice where the two jobs are identical except for the wage. This attention check occurs amid the other job choice questions at a randomized position.

I further assessed respondent attentiveness by evaluating the internal consistency of choices in the hypothetical job choice conjoint. First, I compared the choices of individual respondents who saw the same question twice. Because the job attributes and wages are randomized for five of the six questions that respondents saw, a small fraction of respondents saw the same question twice. If respondents were paying attention and answering according to a well-behaved utility function, they should have chosen the same alternative when faced with the same choice multiple times. In Table A.1, we see that 53 respondents saw the exact same choice more than once, and of those, 41 respondents (77%) made the same choice both times, while 12 respondents (23%) made a different choice.

Table A.1: Repeated Choice Consistency

| | Count | Share |
|---|---|---|
| different choice | 8 | 0.13 |
| same choice | 52 | 0.87 |
| Total | 60 | 1.00 |

This table shows the number and shares of respondents who made the same choice when they saw an identical pair of options in the job choice conjoint multiple times.

To include more data points, I also check for consistency across questions where individuals chose between jobs with the same attributes twice, but possibly with different wages. I present a brief example below to show how to check for consistency when wage pairs differ.

Suppose an individual is choosing between the jobs below, and denote by $WTP_a$ the willingness to pay for attribute $a$. If $WTP_a > 0$, a job with attribute $a$ is preferred, and if

$WTP_a < 0$, a job without attribute $a$ is preferred.

| | job 1 | job 2 |
|---|---|---|
| attribute a | 1 | 0 |
| wage | w1 | w2 |

Below, I enumerate what each possible choice and wage difference implies about $WTP_a$:

| job chosen | wage difference | implication |
|---|---|---|
| 1 | $w1 < w2$ | $WTP_a \geq w2 - w1 > 0$ |
| 1 | $w2 < w1$ | $WTP_a \geq w2 - w1 < 0$ |
| 2 | $w1 < w2$ | $WTP_a \leq w2 - w1 > 0$ |
| 2 | $w2 < w1$ | $WTP_a \leq w2 - w1 < 0$ |

We see that each choice implies an upper or lower bound for $WTP_a$. If the observed choices reveal multiple lower or multiple upper bounds, they are consistent (this means the respondent chose the job with the same attribute both times). If the observed choices reveal a lower and an upper bound, we can test for consistency by checking whether the lower bound is less than or equal to the upper bound.

For instance, if we found that an individual chose job 1 when $w_1 = 1$ and $w_2 = 2$, we would conclude that $WTP_a \geq 1$. If that individual later chose job 2 when $w_1 = 2$ and $w_2 = 1.5$, we would conclude $WTP_a \leq .5$. These choices would be inconsistent, then, as they would imply that $WTP_a \geq 1$ and $WTP_a \leq .5$.

Table A.2 shows the results of this consistency check. Of the 581 respondents who saw two pairs of jobs with the same choice between female shares, but possibly different wages, 360 chose the same female share both times, giving either two lower or two upper bounds for their WTP for that female share. Of the 221 respondents who chose different alternatives and thus produced a lower and upper bound on the WTP, 85% had choices that produced non-overlapping bounds.

Table A.2: Consistency Test for Job Choice Conjoint

| Choice | Bounds | Count | Share | Share of Different Job Choosers |
|---|---|---|---|---|
| chose same job | | 697 | 0.74 | |
| chose different jobs | bounds not consistent | 21 | 0.02 | 0.09 |
| chose different jobs | bounds consistent | 219 | 0.23 | 0.91 |

This table enumerates all choices made by respondents who saw the same pair of jobs multiple times possibly with different wages. Choosing the same job both times is consistent. When choosing a different job in each case, a lower and upper bound for the WTP for that pair of attributes is produced. These bounds are consistent if the lower bound is less than the upper bound.

Overall, the survey respondents appear to be relatively attentive. The majority of respondents took a reasonable amount of time to complete the survey and passed most attention checks. In most cases, respondents provided internally consistent answers within the job choice conjoint experiment. The results in the next section include respondents who failed attention checks, but results are nearly identical when they are excluded from the sample. In addition, through the latent preference type model, I will be able to distinguish inattentive respondents that these attention checks might not pick up.

# Appendix B Estimation of Gender Composition Valuations

## B.1 Calculating the Willingness-to-Pay

Below, I discuss how to convert the regression coefficients into willingness-to-pay estimates. Consider this example: a worker is choosing between jobs 1 and 2, with wages $w_1$ and $w_2$, respectively. Job 2 has some attribute $X$ that job 1 does not have. The worker's utility function is

$$U_i = \beta_w \ln(w_j) + \beta X_j + \varepsilon_{ij}, \tag{12}$$

where $\varepsilon_{ij}$ follows a type I extreme value distribution.

Then, the probability of choosing job 1 is

$$P(\text{choose job 1}|\text{job 1, job 2}) = \frac{1}{1 + \exp(X\beta + ln(w_2/w_1)\beta_w)}. \tag{13}$$

The worker is then indifferent between the two jobs when this probability is .5, and the difference between the wages exactly cancels out the utility from the amenity

$$.5 = \frac{1}{1 + \exp(X\beta + ln(w_2^*/w_1^*)\beta_w)}. \tag{14}$$

When we rearrange, we get

$$exp(-\beta/\beta_w) = w_2^*/w_1^*. \tag{15}$$

This means that when the ratio between the wages in job 2 and job 1 is exactly equal to $exp(-\beta/\beta_w)$, the worker will be indifferent between the job with the amenity and the job without the amenity. To express this as a percent of the wage in job 1, I take $1 - w_2^*/w_1^*$, so

$$WTP_f = 1 - \exp\left(\frac{-\beta_f}{\beta_w}\right). \tag{16}$$

## B.2 Maximum Likelihood Estimation of Latent Class Logit Model

Let $H_{iq}$ denote the prior probability of class $q$ for individual $i$. Then the likelihood for individual $i$ is

$$P_i = \sum_{q=1}^{Q} H_{iq} P_{i|q}, \tag{17}$$

and the log-likelihood for the sample is

$$\ln L = \sum_{i=1}^{N} \ln P_i = \sum_{i=1}^{N} \ln \left[ \sum_{q=1}^{Q} H_{iq} \left( \prod_{t=1}^{T_i} P_{it|q} \right) \right]. \tag{18}$$

I can estimate the parameters either using an EM-type algorithm, where the class probabilities $H_{iq}$ and the parameters $\beta$ are estimated iteratively, or estimate them jointly in a one-step maximum likelihood. Here, I estimate all parameters jointly in one step. I select the number of classes $Q$ by evaluating the Aikake Information Criterion (AIC) and Bayesian Information Criterion (BIC) and running a 10-fold cross validation on the log-likelihood.

## B.3 Demographic Conjoint Results

Table B.1: Conjoint Coefficients

| | All | Female | Male |
|---|---|---|---|
| log wage | 27.56* | 28.54* | 26.77* |
| | [26.78; 28.35] | [27.45; 29.64] | [25.62; 27.91] |
| 0% female | −1.06* | −1.25* | −0.85* |
| | [−1.15; −0.97] | [−1.37; −1.12] | [−0.98; −0.71] |
| 10% female | −0.68* | −0.86* | −0.47* |
| | [−0.76; −0.59] | [−0.98; −0.75] | [−0.59; −0.34] |
| 20% female | −0.42* | −0.50* | −0.32* |
| | [−0.51; −0.33] | [−0.62; −0.38] | [−0.46; −0.19] |
| 30% female | −0.29* | −0.36* | −0.19* |
| | [−0.37; −0.20] | [−0.48; −0.25] | [−0.31; −0.06] |
| 40% female | −0.13* | −0.19* | −0.06 |
| | [−0.22; −0.04] | [−0.32; −0.07] | [−0.19; 0.07] |
| 60% female | 0.08 | 0.06 | 0.09 |
| | [−0.01; 0.17] | [−0.06; 0.19] | [−0.04; 0.23] |
| 70% female | −0.06 | −0.00 | −0.12 |
| | [−0.14; 0.03] | [−0.12; 0.11] | [−0.25; 0.00] |
| 80% female | −0.06 | 0.00 | −0.15* |
| | [−0.15; 0.02] | [−0.12; 0.12] | [−0.28; −0.02] |
| 90% female | −0.19* | −0.06 | −0.32* |
| | [−0.27; −0.10] | [−0.18; 0.05] | [−0.45; −0.20] |
| 100% female | −0.41* | −0.37* | −0.47* |
| | [−0.50; −0.32] | [−0.49; −0.25] | [−0.60; −0.34] |
| 30% have kids | 0.03 | 0.01 | 0.07 |
| | [−0.03; 0.10] | [−0.08; 0.10] | [−0.03; 0.16] |
| 70% have kids | 0.05 | 0.13* | −0.04 |
| | [−0.02; 0.12] | [0.04; 0.23] | [−0.15; 0.06] |
| 30% under 40 | 0.09* | 0.08 | 0.10* |
| | [0.02; 0.16] | [−0.01; 0.17] | [0.00; 0.20] |
| 70% under 40 | −0.17* | −0.20* | −0.15* |
| | [−0.24; −0.10] | [−0.29; −0.11] | [−0.25; −0.05] |
| lefthand job | 0.15* | 0.14* | 0.16* |
| | [0.12; 0.17] | [0.10; 0.18] | [0.12; 0.20] |
| Num. obs. | 29972 | 16495 | 13334 |
| Num. indiv. | 2772 | 1525 | 1234 |

* Null hypothesis value outside 95% credible interval.
Omitted baseline categories are 50% female, 50% have kids, and 50% under 40.
Each individual sees 7-10 choices. Thirteen individuals chose to self-describe gender.

## B.4 Heterogeneous Preferences

For both genders, I find that a portion of the sample comes from what I term an inattentive class. The inattentive class is distinguished by very small (1.75 for women and -2.45 for men) coefficients on the wage and noisy coefficients on the demographic characteristics. These wage coefficients imply that given a choice between two jobs with one job having a 10% lower wage, a woman from the inattentive class will choose the higher wage job with a probability of 54%, and a man from the inattentive class will choose the higher wage job with a probability of 44%.

These estimates suggest that the respondents classified as inattentive are essentially answering the hypothetical job choice questions at random: we would expect that even if respondents do not care about the demographic composition of the job, they should choose higher wage jobs on average, while this group does not. I estimate that 28.5% of women and 16.8% of men belong to the inattentive class. The estimated logit coefficients for every class, including the inattentive class, are shown in Appendix A Tables B.2 and B.4. In the remainder of this analysis, I omit the inattentive group for both genders.

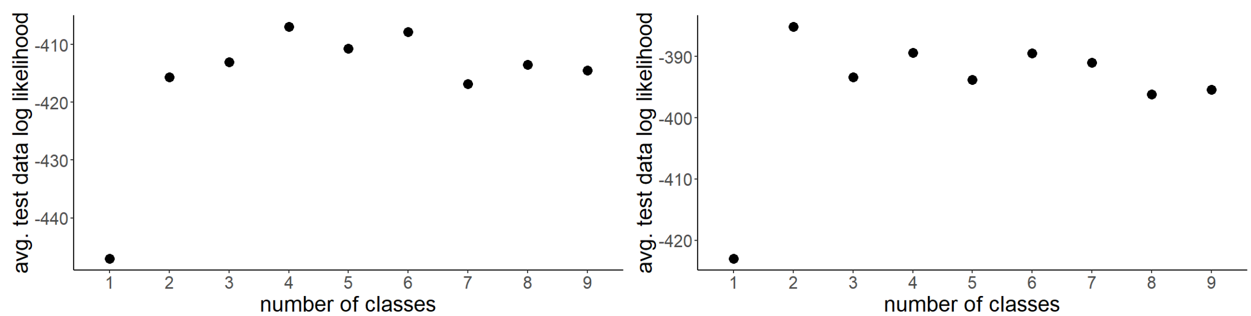


(a) Female (b) Male

Figure B.1: Avg. Test Log Likelihood in 10-fold Cross-Validation

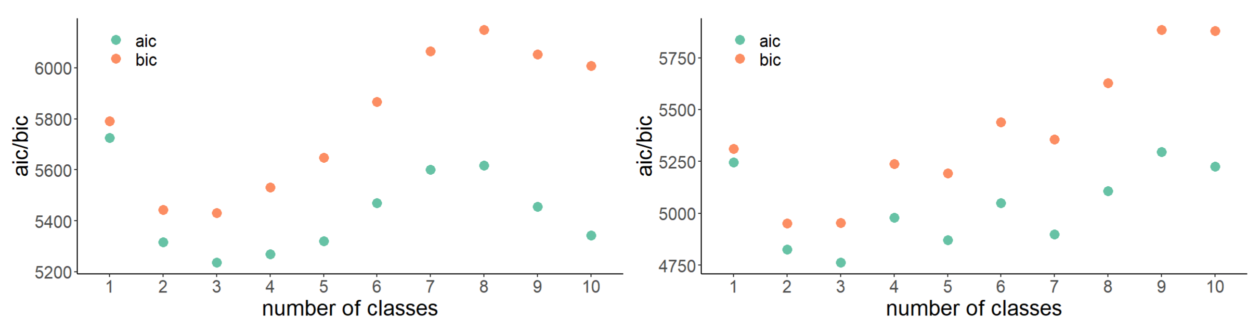


Female Male

Figure B.2: Akaike Information Criterion and Bayesian Information Criterion

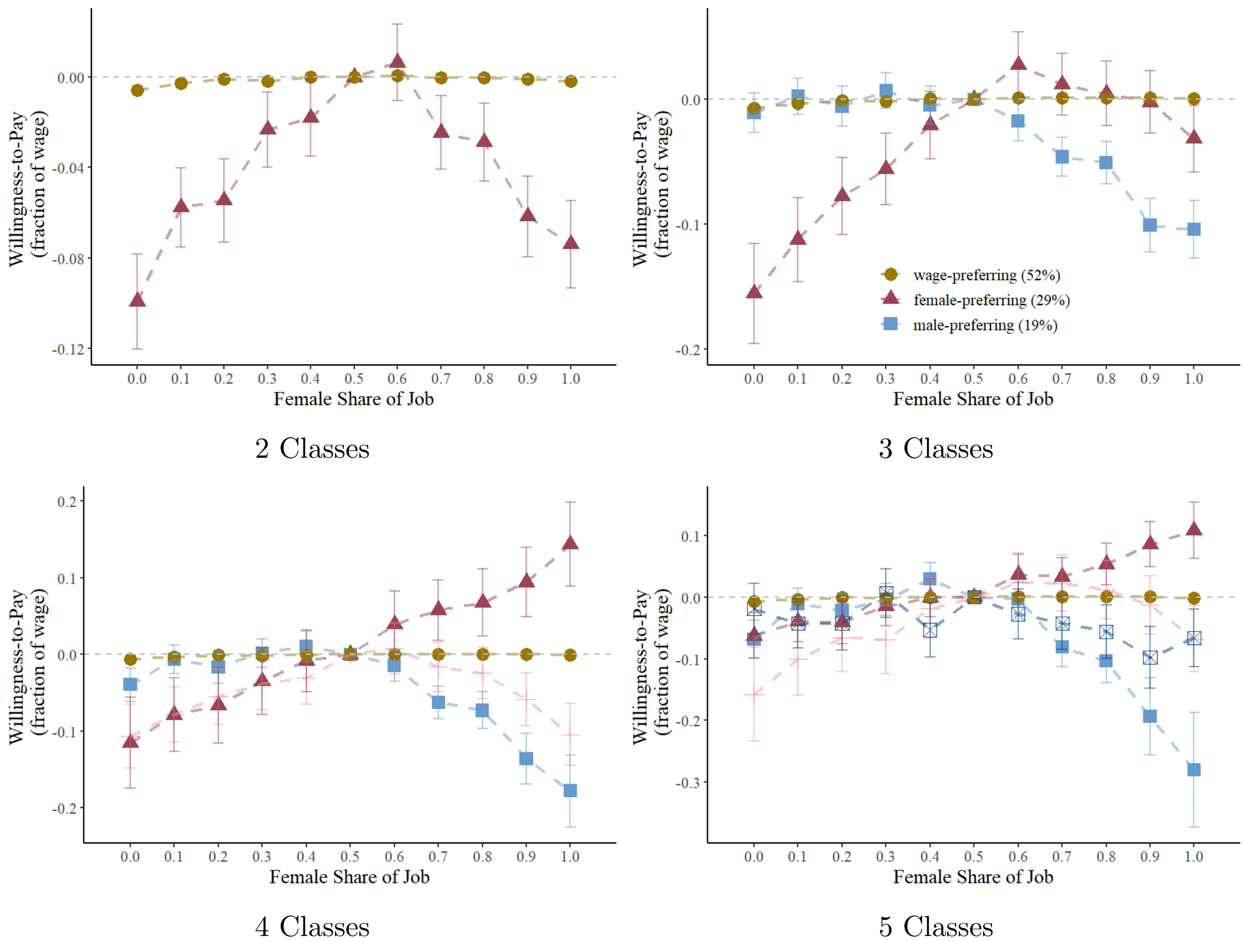


Figure B.3: Heterogeneous WTPs for Men, 2 to 5 Classes

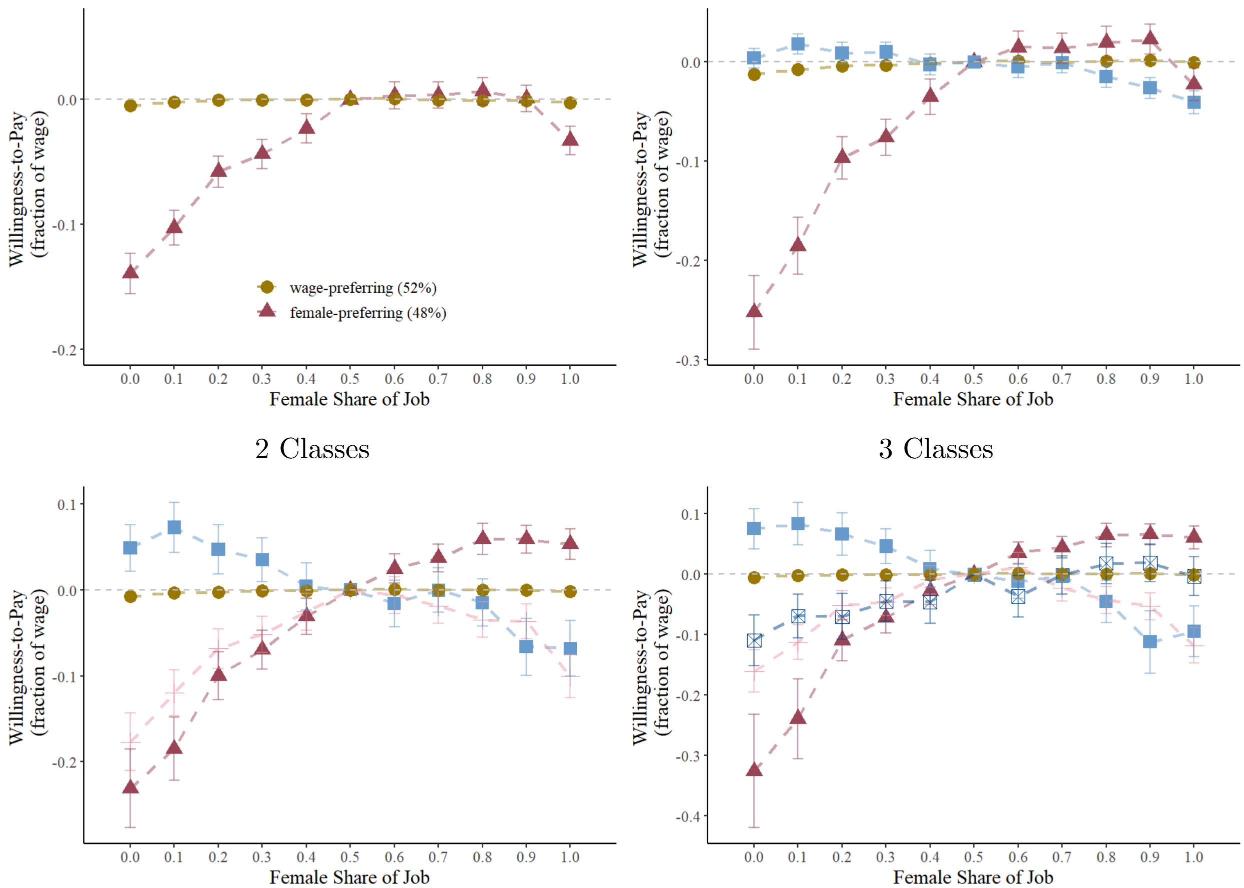


2 Classes

3 Classes

4 Classes

4 Classes

Figure B.4: Heterogeneous WTPs for Women, 2 to 5 Classes

Table B.2: Female Latent Class Logit Coefficients

| | wage-preferring | female-preferring |
|---|---|---|
| log(w) | 181.255 | 15.283 |
| | (168.747,193.763) | (14.229,16.338) |
| 0% female | -0.842 | -1.993 |
| | (-1.126,-0.558) | (-2.171,-1.815) |
| 10% f | -0.423 | -1.498 |
| | (-0.681,-0.165) | (-1.665,-1.331) |
| 20% f | -0.09 | -0.861 |
| | (-0.384,0.205) | (-1.032,-0.69) |
| 30% f | -0.129 | -0.656 |
| | (-0.389,0.132) | (-0.819,-0.492) |
| 40% f | -0.082 | -0.352 |
| | (-0.377,0.213) | (-0.525,-0.178) |
| 60% f | 0.106 | 0.048 |
| | (-0.18,0.392) | (-0.121,0.217) |
| 70% f | -0.065 | 0.051 |
| | (-0.329,0.199) | (-0.109,0.211) |
| 80% f | -0.137 | 0.1 |
| | (-0.433,0.16) | (-0.067,0.267) |
| 90% f | -0.202 | 0.012 |
| | (-0.458,0.054) | (-0.148,0.171) |
| 100% f | -0.463 | -0.495 |
| | (-0.754,-0.173) | (-0.662,-0.329) |
| 30% have kids | -0.062 | 0.078 |
| | (-0.267,0.144) | (-0.049,0.205) |
| 60% have kids | -0.043 | 0.25 |
| | (-0.253,0.167) | (0.122,0.378) |
| 30% <40 | -0.079 | 0.194 |
| | (-0.282,0.124) | (0.068,0.319) |
| 70% <40 | -0.123 | -0.295 |
| | (-0.326,0.08) | (-0.42,-0.17) |
| class share | 0.52 | 0.48 |

Table B.3: Willingness-to-Pay by Class, Female Latent Class Logit

| coefs | wage-preferring | female-preferring |
|---|---|---|
| 0% female | -0.005 | -0.139 |
| | (-0.006,-0.003) | (-0.155,-0.123) |
| 10% female | -0.002 | -0.103 |
| | (-0.004,-0.001) | (-0.117,-0.089) |
| 20% female | 0 | -0.058 |
| | (-0.002,0.001) | (-0.07,-0.046) |
| 30% female | -0.001 | -0.044 |
| | (-0.002,0.001) | (-0.055,-0.032) |
| 40% female | 0 | -0.023 |
| | (-0.002,0.001) | (-0.035,-0.012) |
| 50% female | | |
| | | |
| 60% female | 0.001 | 0.003 |
| | (-0.001,0.002) | (-0.008,0.014) |
| 70% female | 0 | 0.003 |
| | (-0.002,0.001) | (-0.007,0.014) |
| 80% female | -0.001 | 0.006 |
| | (-0.002,0.001) | (-0.004,0.017) |
| 90% female | -0.001 | 0.001 |
| | (-0.003,0) | (-0.01,0.011) |
| 100% female | -0.003 | -0.033 |
| | (-0.004,-0.001) | (-0.044,-0.022) |
| 30% kids | 0 | 0.005 |
| | (-0.001,0.001) | (-0.003,0.013) |
| 50% kids | | |
| | | |
| 70% kids | 0 | 0.016 |
| | (-0.001,0.001) | (0.008,0.025) |
| 30% <40 | 0 | 0.013 |
| | (-0.002,0.001) | (0.004,0.021) |
| 50% <40 | | |
| | | |
| 70% <40 | -0.001 | -0.019 |
| | (-0.002,0) | (-0.028,-0.011) |

Note: This table shows the willingness-to-pay, as a fraction of the wage, for each attribute for the two preference classes among women estimated using the latent class logit model described in Section III.C.2. 95% confidence intervals are in parentheses. The logit model coefficients are shown in Appendix Table B.2.

Table B.4: Male Latent Class Logit Coefficients

| | wage-preferring | female-preferring | male-preferring |
|---|---|---|---|
| log(w) | 181.212 | 26.159 | 9.295 |
| | (165.782,196.642) | (23.016,29.302) | (7.898,10.692) |
| 0% female | -1.336 | -0.28 | -1.344 |
| | (-1.659,-1.014) | (-0.69,0.13) | (-1.604,-1.084) |
| 10% f | -0.59 | 0.062 | -0.99 |
| | (-0.884,-0.297) | (-0.315,0.439) | (-1.231,-0.75) |
| 20% f | -0.223 | -0.144 | -0.694 |
| | (-0.55,0.104) | (-0.559,0.271) | (-0.941,-0.447) |
| 30% f | -0.287 | 0.183 | -0.505 |
| | (-0.591,0.018) | (-0.195,0.561) | (-0.749,-0.262) |
| 40% f | 0.081 | -0.124 | -0.19 |
| | (-0.254,0.416) | (-0.522,0.275) | (-0.434,0.054) |
| 60% f | 0.245 | -0.451 | 0.26 |
| | (-0.079,0.569) | (-0.866,-0.036) | (0.01,0.511) |
| 70% f | 0.184 | -1.178 | 0.112 |
| | (-0.122,0.491) | (-1.569,-0.787) | (-0.12,0.344) |
| 80% f | 0.186 | -1.296 | 0.045 |
| | (-0.126,0.497) | (-1.72,-0.872) | (-0.199,0.288) |
| 90% f | 0.207 | -2.514 | -0.021 |
| | (-0.1,0.513) | (-3.01,-2.018) | (-0.253,0.212) |
| 100% f | 0.093 | -2.593 | -0.289 |
| | (-0.218,0.405) | (-3.132,-2.055) | (-0.527,-0.051) |
| 50% have kids | -0.064 | 0.147 | -0.206 |
| | (-0.301,0.173) | (-0.17,0.464) | (-0.396,-0.015) |
| 60% have kids | 0.017 | 0.087 | -0.392 |
| | (-0.228,0.262) | (-0.204,0.378) | (-0.577,-0.206) |
| 50% <40 | -0.192 | -0.115 | -0.087 |
| | (-0.427,0.044) | (-0.427,0.197) | (-0.274,0.099) |
| 70% <40 | -0.226 | -0.301 | -0.374 |
| | (-0.451,-0.001) | (-0.596,-0.006) | (-0.562,-0.185) |
| class share | 0.518 | 0.29 | 0.192 |

Table B.5: Willingness-to-Pay by Class, Male Latent Class Logit

| coefs | wage-preferring | female-preferring | male-preferring |
|---|---|---|---|
| 0% female | -0.007 | -0.156 | -0.011 |
| | (-0.009,-0.006) | (-0.196,-0.115) | (-0.027,0.005) |
| 10% female | -0.003 | -0.112 | 0.002 |
| | (-0.005,-0.002) | (-0.146,-0.079) | (-0.012,0.017) |
| 20% female | -0.001 | -0.078 | -0.006 |
| | (-0.003,0.001) | (-0.109,-0.047) | (-0.021,0.01) |
| 30% female | -0.002 | -0.056 | 0.007 |
| | (-0.003,0) | (-0.085,-0.027) | (-0.007,0.021) |
| 40% female | 0 | -0.021 | -0.005 |
| | (-0.001,0.002) | (-0.048,0.006) | (-0.02,0.011) |
| 50% female | | | |
| 60% female | 0.001 | 0.028 | -0.017 |
| | (0,0.003) | (0.001,0.054) | (-0.033,-0.001) |
| 70% female | 0.001 | 0.012 | -0.046 |
| | (-0.001,0.003) | (-0.013,0.037) | (-0.062,-0.03) |
| 80% female | 0.001 | 0.005 | -0.051 |
| | (-0.001,0.003) | (-0.021,0.031) | (-0.068,-0.034) |
| 90% female | 0.001 | -0.002 | -0.101 |
| | (-0.001,0.003) | (-0.027,0.023) | (-0.122,-0.079) |
| 100% female | 0.001 | -0.032 | -0.104 |
| | (-0.001,0.002) | (-0.058,-0.005) | (-0.127,-0.081) |
| 30% kids | | | |
| 50% kids | 0 | -0.022 | 0.006 |
| | (-0.002,0.001) | (-0.044,-0.001) | (-0.006,0.018) |
| 70% kids | 0 | -0.043 | 0.003 |
| | (-0.001,0.001) | (-0.065,-0.021) | (-0.008,0.014) |
| 30% <40 | | | |
| 50% <40 | -0.001 | -0.009 | -0.004 |
| | (-0.002,0) | (-0.03,0.011) | (-0.016,0.008) |
| 70% <40 | -0.001 | -0.041 | -0.012 |
| | (-0.002,0) | (-0.063,-0.019) | (-0.023,0) |

Note: This table shows the willingness-to-pay, as a fraction of the wage, for each attribute for the two preference classes among men estimated using the latent class logit model described in Section III.C.2. 95% confidence intervals are in parentheses. The logit model coefficients are shown in Appendix Table B.4.

## B.5 Observable Heterogeneity

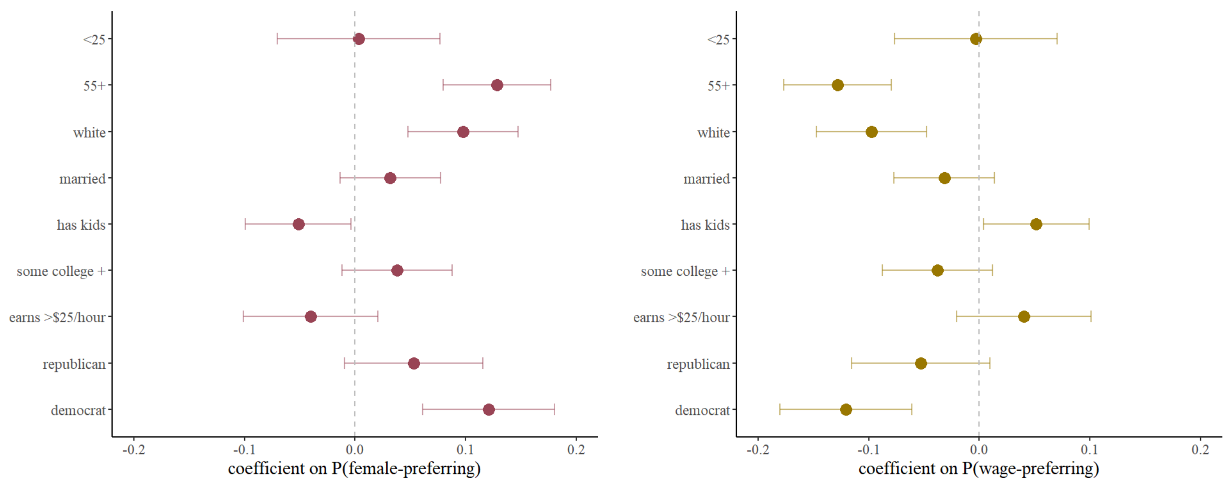


(a) Female-Preferring Class

(b) Wage-Preferring Class

Figure B.5: Correlates of Preference Class: Women

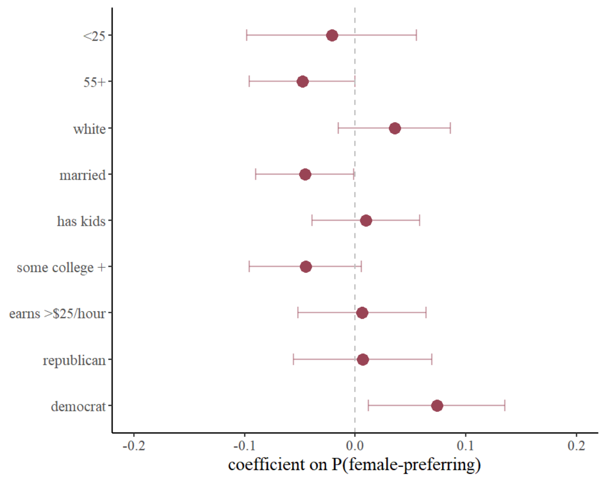


(a) Female-Preferring Class

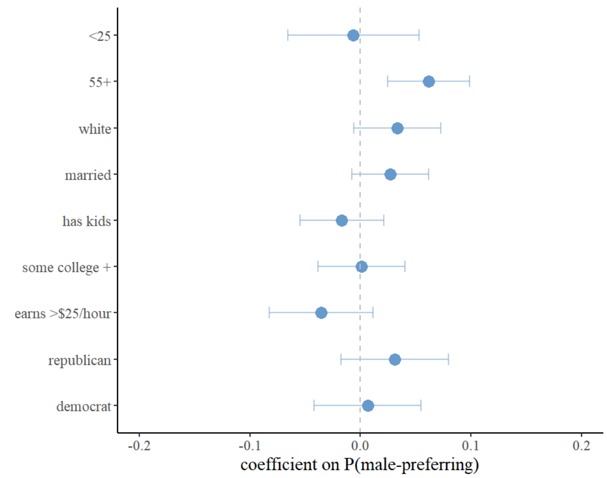


(b) Male-Preferring Class

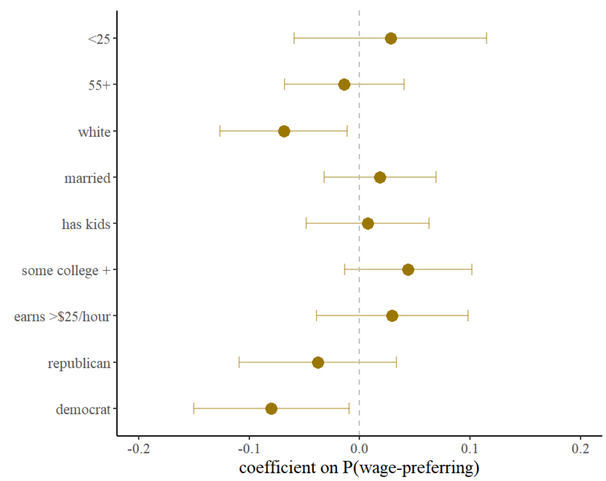


(c) Wage-Preferring Class

Figure B.6: Correlates of Preference Class: Men

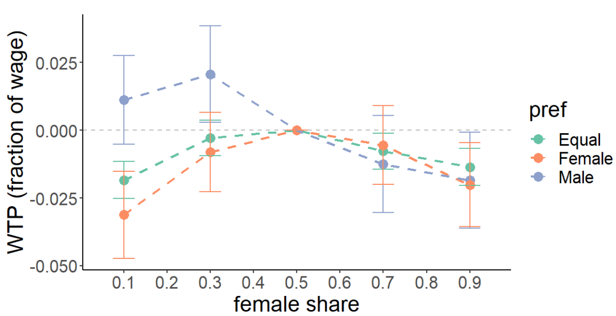


(a) Men

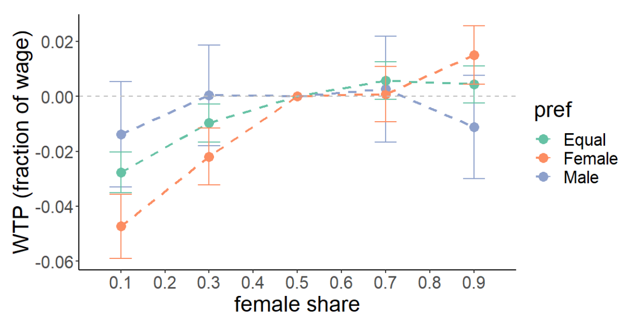


(b) Women

Figure B.7: More Satisfied with Coworkers: Male or Female

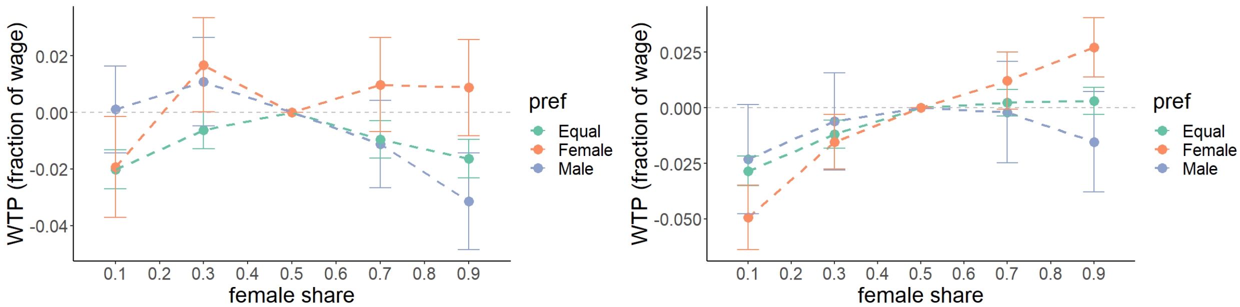


(a) Men (b) Women

Figure B.8: More Satisfied with Work Environment: Male or Female

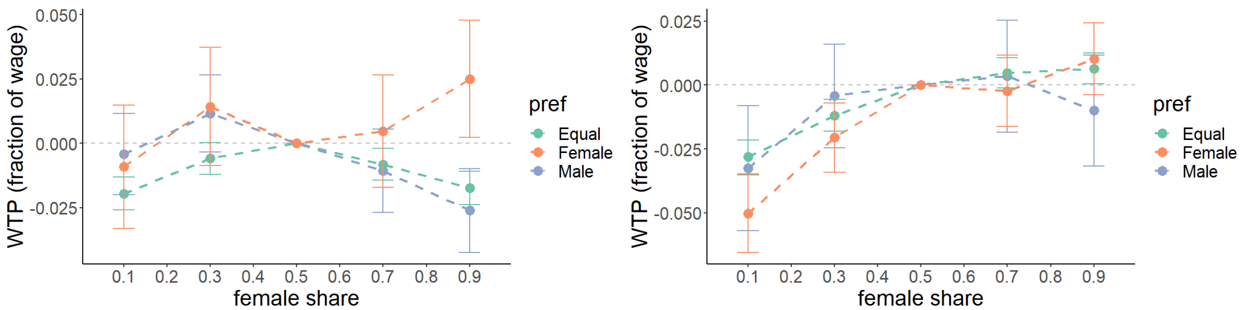


(a) Men (b) Women

Figure B.9: More Satisfied with Tasks: Male or Female

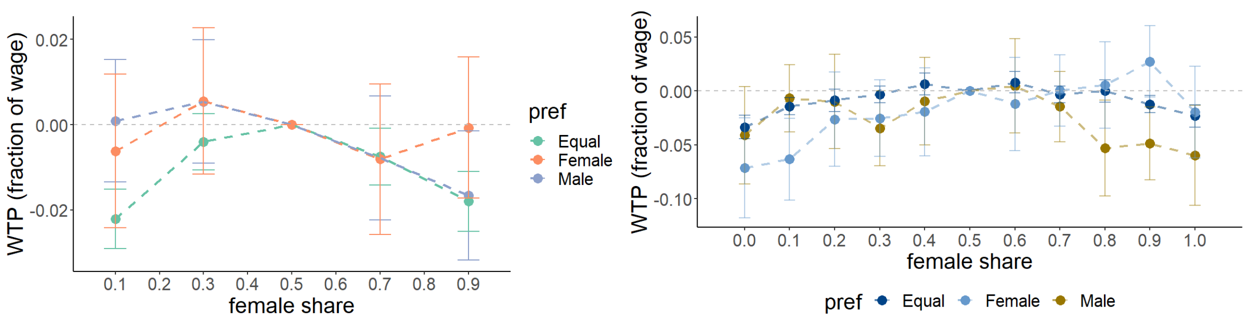


(a) Men (b) Women

Figure B.10: More Satisfied with Schedule: Male or Female

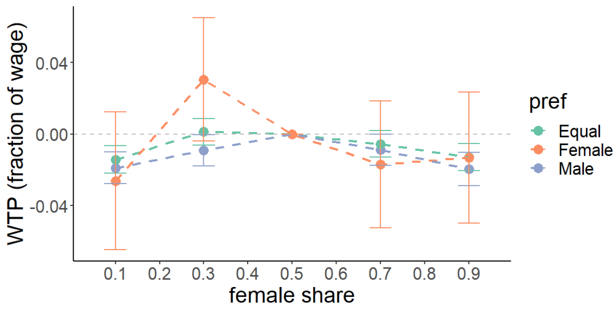


(a) Men

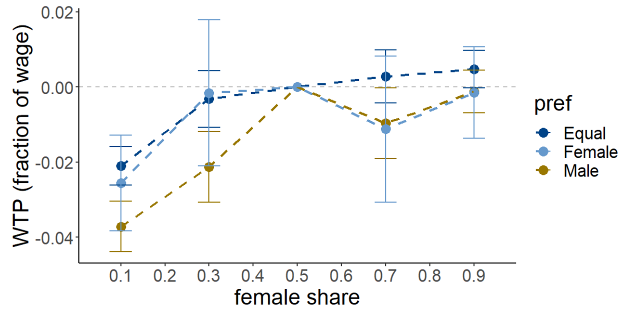


(b) Women

Figure B.11: Earn More: Male or Female

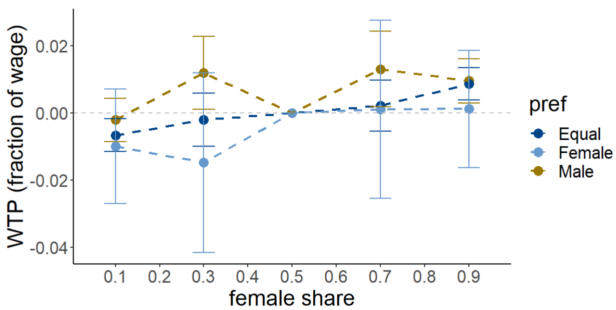


(a) Men

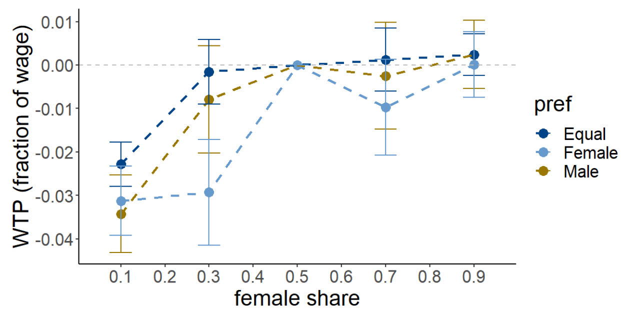


(b) Women

Figure B.12: More Likely to be Promoted: Male or Female

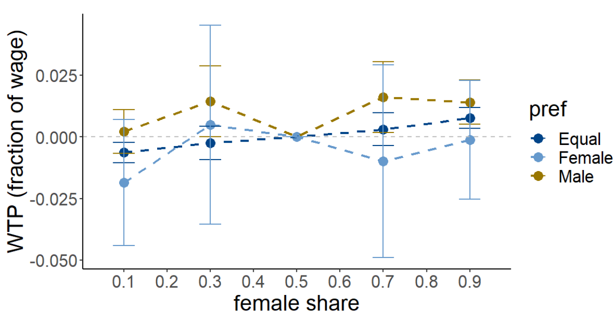


(a) Men

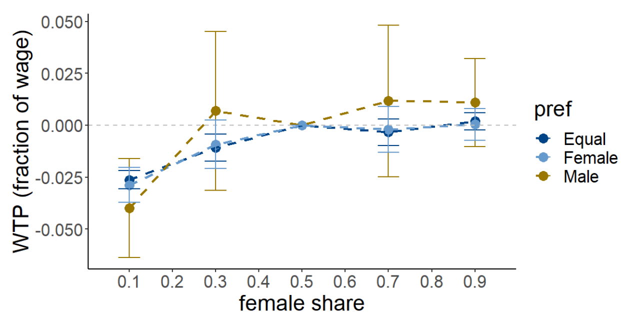


(b) Women

Figure B.13: Family Would Prefer: Male or Female

## B.6 Channels and Robustness

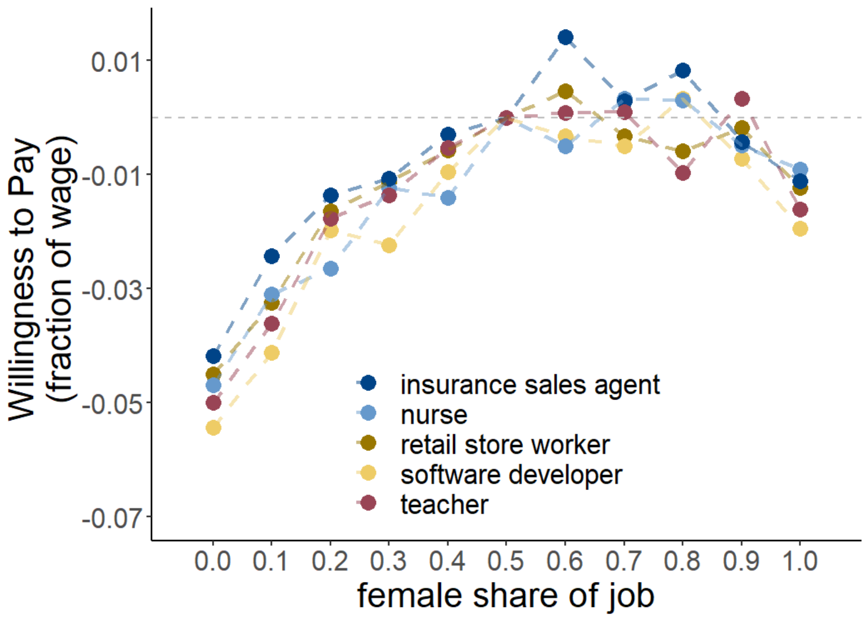


Figure B.14: WTP for Female Share by Occupation in Workplace Choice, Female

Note: This figure plots the willingness-to-pay for each possible female share for women, split by the occupation listed in the hypothetical workplace choice.

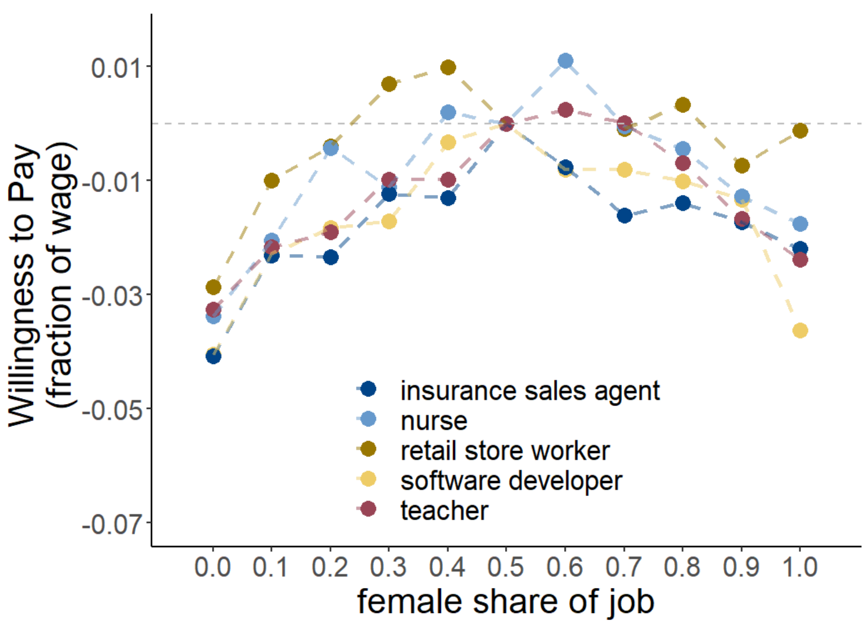


Figure B.15: WTP for Female Share by Occupation in Workplace Choice, Male

Note: This figure plots the willingness-to-pay for each possible female share for men, split by the occupation listed in the hypothetical workplace choice.

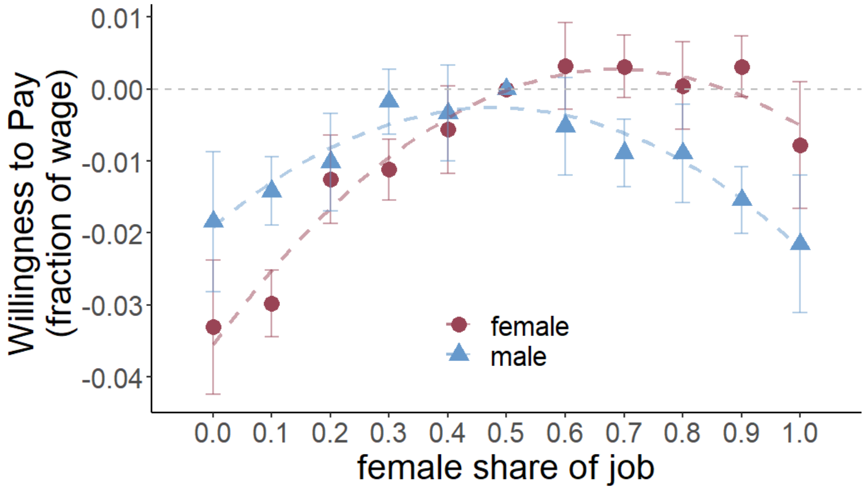


Figure B.16: WTP for Female Share, No Occupation Listed in Choice

Note: This figure plots the willingness-to-pay for each possible female share for men and women in the hypothetical choice with no occupation listed.

# Appendix C Quantitative Model

## C.1 Occupation Nests

I create occupation nests using data from the March CPS and the fast-greedy clustering algorithm. Using the same CPS sample described in Section IV.C, I limit to workers who are observed over two years. Using this sample, I create an adjacency matrix which counts observed flows for each pair of occupations. The fast-greedy algorithm creates a subnetwork of connections from this adjacency matrix, and then randomly samples connections that improve the modularity of the network. Essentially, the algorithm groups together the occupations that workers flow between most frequently. The 15 nests selected by the algorithm are presented in the following table, ordered from the smallest to the largest number of occupation codes in the nest.

Table C.1: Occupation Nests

| **Real Estate** |
|---|
| Property, real estate, and community association managers; Real estate brokers and sales agents |
| **Law and Economics** |
| Economists; Lawyers |
| **Air Travel** |
| Aircraft pilots and flight engineers; Air traffic controllers and airfield operations specialists; Flight attendants |
| **Postal Service** |
| Postal service clerks; Postal service mail carriers; Postal service mail sorters, processors, and processing machine operators |
| **Agriculture** |
| Farmers, ranchers, and other agricultural managers; First-line supervisors of farming, fishing, and forestry workers; Agricultural inspectors; Miscellaneous agricultural workers ; Fishers and related fishing workers |
| **Medicine** |
| Chiropractors; Dentists; Dietitians and nutritionists; Optometrists; Pharmacists; Physicians and surgeons; Podiatrists; Veterinarians |

| Transportation |
| --- |
| Couriers and messengers; Driver/sales workers and truck drivers; Taxi drivers and chauffeurs; Motor vehicle operators, all other; Locomotive engineers and operators; Railroad conductors and yardmasters; Sailors and marine oilers; Ship and boat captains and operators; Parking lot attendants; Pumping station operators; Refuse and recyclable material collectors |
| **Law Enforcement** |
| First-line supervisors of correctional officers; First-line supervisors of police and detectives; First-line supervisors of fire fighting and prevention workers; First-line supervisors of protective service workers, all other; Firefighters; Fire inspectors; Bailiffs, correctional officers, and jailers; Detectives and criminal investigators; Private detectives and investigators; Security guards and gaming surveillance officers; Crossing guards; Furnace, kiln, oven, drier, and kettle operators and tenders |
| **Food Service** |
| Food service managers; Chefs and head cooks; First-line supervisors of food preparation and serving workers; Food preparation workers; Bartenders; Combined food preparation and serving workers, including fast food; Counter attendants, cafeteria, food concession, and coffee shop; Waiters and waitresses; Food servers, nonrestaurant; Dining room and cafeteria attendants and bartender helpers; Dishwashers; Hosts and hostesses, restaurant, lounge, and coffee shop; Miscellaneous entertainment attendants and related workers; Bakers; Food and tobacco roasting, baking, and drying machine operators and tenders; Food cooking machine operators and tenders |
| **Science and Education** |
| Education administrators; Agricultural and food scientists; Biological scientists; Medical scientists; Astronomers and physicists; Atmospheric and space scientists; Chemists and materials scientists; Environmental scientists and geoscientists; Physical scientists, all other; Biological technicians; Postsecondary teachers; Preschool and kindergarten teachers; Elementary and middle school teachers; Secondary school teachers; Special education teachers; Other teachers and instructors; Archivists, curators, and museum technicians; Librarians; Library technicians; Teacher assistants; Other education, training, and library workers; Athletes, coaches, umpires, and related workers; Dancers and choreographers; Musicians, singers, and related workers; Childcare workers; Library assistants, clerical |
| **Misc. Services** |
| Legislators; Animal control workers; Ushers, lobby attendants, and ticket takers; First-line supervisors of retail sales workers; Cashiers; Counter and rental clerks; Parts salespersons; Retail salespersons; Advertising sales agents; Travel agents; Sales representatives, services, all other; Sales representatives, wholesale and manufacturing; Models, demonstrators, and product promoters; Sales engineers; Telemarketers; Door-to-door sales workers, news and street vendors, and related workers; Sales and related workers, all other; Communications equipment operators, all other; Gaming cage workers; Brokerage clerks; |

## Misc. Services Continued

Customer service representatives; Hotel, motel, and resort desk clerks; Order clerks; Reservation and transportation ticket agents and travel clerks; Cargo and freight agents; Shipping, receiving, and traffic clerks; Stock clerks and order fillers; Graders and sorters, agricultural products; Butchers and other meat, poultry, and fish processing workers; Food processing workers, all other; Textile cutting machine setters, operators, and tenders; Textile winding, twisting, and drawing out machine setters, operators, and tenders; Packaging and filling machine operators and tenders; Automotive and watercraft service attendants ; Other transportation workers ; Industrial truck and tractor operators; Laborers and freight, stock, and material movers, hand; Packers and packagers, hand

## Health and Personal Care

Medical and health services managers; Social and community service managers; Psychologists; Counselors; Social workers; Clergy; Directors, religious activities and education; Religious workers, all other; Physician assistants; Audiologists; Occupational therapists; Physical therapists; Radiation therapists; Recreational therapists; Respiratory therapists; Speech-language pathologists; Health diagnosing and treating practitioners, all other; Clinical laboratory technologists and technicians; Dental hygienists; Diagnostic related technologists and technicians; Emergency medical technicians and paramedics; Licensed practical and licensed vocational nurses; Opticians, dispensing; Other healthcare practitioners and technical occupations; Nursing, psychiatric, and home health aides; Occupational therapy assistants and aides; Physical therapist assistants and aides; Massage therapists; Dental assistants; First-line supervisors of personal service workers; Animal trainers; Nonfarm animal caretakers; Gaming services workers; Barbers; Hairdressers, hairstylists, and cosmetologists; Miscellaneous personal appearance workers; Baggage porters, bellhops, and concierges; Personal care aides; Recreation and fitness workers; Residential advisors; Personal care and service workers, all other

## Finance and Insurance

Financial managers; Buyers and purchasing agents, farm products; Wholesale and retail buyers, except farm products; Purchasing agents, except wholesale, retail, and farm products; Claims adjusters, appraisers, examiners, and investigators; Accountants and auditors; Appraisers and assessors of real estate; Budget analysts; Credit analysts; Financial analysts; Personal financial advisors; Insurance underwriters; Financial examiners; Credit counselors and loan officers; Tax examiners and collectors, and revenue agents; Tax preparers; Financial specialists, all other; Medical records and health information technicians; Insurance sales agents; Securities, commodities, and financial services sales agents; First-line supervisors of office and administrative support workers; Switchboard operators, including answering service; Telephone operators; Bill and account collectors; Billing and posting clerks ; Bookkeeping, accounting, and auditing clerks; Payroll and timekeeping clerks; Procurement clerks; Tellers; Financial clerks, all other; Court, municipal, and license clerks; Credit authorizers, checkers, and clerks; Eligibility interviewers, government programs; File clerks; Interviewers, except eligibility and loan; Loan interviewers and clerks; New accounts clerks; Human resources assistants, except payroll and timekeeping; Receptionists and information clerks; Information and record clerks, all other; Dispatchers; Production, planning, and expediting clerks; Secretaries and administrative assistants; Data entry keyers; Word processors and typists; Insurance claims and policy processing clerks; Mail clerks and mail machine operators, except postal service; Office clerks, general; Office machine operators, except computer; Proofreaders and copy markers; Office and administrative support workers, all other

| Management and Engineering |
|---|
| Chief executives; General and operations managers; Administrative services managers; Computer and information systems managers; Industrial production managers; Purchasing managers; Transportation, storage, and distribution managers; Construction managers; Architectural and engineering managers; Gaming managers; Natural sciences managers; Managers, all other; Agents and business managers of artists, performers, and athletes; Cost estimators; Logisticians; Management analysts; Computer programmers; Software developers, applications and systems software; Computer support specialists; Database administrators; Actuaries; Operations research analysts; Miscellaneous mathematical science occupations; Architects, except naval; Surveyors, cartographers, and photogrammetrists; Aerospace engineers; Chemical engineers; Civil engineers; Computer hardware engineers; Electrical and electronics engineers; Environmental engineers; Industrial engineers, including health and safety; Marine engineers and naval architects; Materials engineers; Mechanical engineers; Petroleum engineers; Engineers, all other; Drafters; Surveying and mapping technicians; Conservation scientists and foresters; Sociologists; Urban and regional planners; Artists and related workers; Designers; Actors; Entertainers and performers, sports and related workers, all other; Announcers; News analysts, reporters and correspondents; Public relations specialists; Technical writers; Writers and authors; Miscellaneous media and communication workers; Broadcast and sound engineering technicians and radio operators; Photographers; Television, video, and motion picture camera operators and editors; First-line supervisors of gaming workers; Embalmers and funeral attendants; Tour and travel guides; Computer operators; Statistical assistants; First-line supervisors of construction trades and extraction workers; First-line supervisors of mechanics, installers, and repairers; Computer, automated teller, and office machine repairers; First-line supervisors of production and operating workers; Photographic process workers and processing machine operators; Supervisors of transportation and material moving workers |
| **Misc. Production** |
| Engineering technicians, except drafters; Agricultural and food science technicians; Chemical technicians; Geological and petroleum technicians; First-line supervisors of housekeeping and janitorial workers; First-line supervisors of landscaping, lawn service, and groundskeeping workers; Janitors and building cleaners; Maids and housekeeping cleaners; Pest control workers; Grounds maintenance workers; Meter readers, utilities; Weighers, measurers, checkers, and samplers, recordkeeping; Forest and conservation workers; Logging workers; Boilermakers; Brickmasons, blockmasons, and stonemasons; Carpenters; Carpet, floor, and tile installers and finishers; Cement masons, concrete finishers, and terrazzo workers; Construction laborers; Paving, surfacing, and tamping equipment operators; Operating engineers and other construction equipment operators; Drywall installers, ceiling tile installers, and tapers; Electricians; Glaziers; Insulation workers; Painters, construction and maintenance; Pipelayers, plumbers, pipefitters, and steamfitters; Plasterers and stucco masons; Reinforcing iron and rebar workers; Roofers; Sheet metal workers; Structural iron and steel workers; Helpers, construction trades; Construction and building inspectors; Elevator installers and repairers; Fence erectors; Hazardous materials removal workers; Highway maintenance workers; Rail-track laying and maintenance equipment operators; Miscellaneous construction and related workers; Derrick, rotary drill, and service unit operators, oil, gas, and mining; Earth drillers, except oil and gas; Explosives workers, ordnance handling experts, and blasters; Mining machine operators; Other extraction workers; Radio and telecommunications equipment installers and repairers; Avionics technicians; Electric motor, power tool, and related repairers |

## Misc. Production Continued

Electrical and electronics repairers, industrial and utility; Electronic equipment installers and repairers, motor vehicles; Electronic home entertainment equipment installers and repairers; Security and fire alarm systems installers; Aircraft mechanics and service technicians; Automotive body and related repairers; Automotive glass installers and repairers; Automotive service technicians and mechanics; Bus and truck mechanics and diesel engine specialists; Heavy vehicle and mobile equipment service technicians and mechanics; Small engine mechanics; Miscellaneous vehicle and mobile equipment mechanics, installers, and repairers; Control and valve installers and repairers; Heating, air conditioning, and refrigeration mechanics and installers; Home appliance repairers; Industrial and refractory machinery mechanics; Maintenance and repair workers, general; Maintenance workers, machinery; Millwrights; Electrical power-line installers and repairers; Telecommunications line installers and repairers; Precision instrument and equipment repairers; Coin, vending, and amusement machine servicers and repairers; Locksmiths and safe repairers; Manufactured building and mobile home installers; Riggers; Helpers–installation, maintenance, and repair workers ; Other installation, maintenance, and repair workers; Aircraft structure, surfaces, rigging, and systems assemblers; Electrical, electronics, and electromechanical assemblers; Engine and other machine assemblers; Structural metal fabricators and fitters; Miscellaneous assemblers and fabricators; Food batchmakers; Computer control programmers and operators; Extruding and drawing machine setters, operators, and tenders, metal and plastic; Rolling machine setters, operators, and tenders, metal and plastic; Cutting, punching, and press machine setters, operators, and tenders, metal and plastic; Grinding, lapping, polishing, and buffing machine tool setters, operators, and tenders, metal and plastic; Lathe and turning machine tool setters, operators, and tenders, metal and plastic; Machinists; Metal furnace operators, tenders, pourers, and casters; Molders and molding machine setters, operators, and tenders, metal and plastic; Tool and die makers; Welding, soldering, and brazing workers; Plating and coating machine setters, operators, and tenders, metal and plastic; Tool grinders, filers, and sharpeners; Metal workers and plastic workers, all other; Prepress technicians and workers; Laundry and dry-cleaning workers; Pressers, textile, garment, and related materials; Sewing machine operators; Shoe and leather workers and repairers; Tailors, dressmakers, and sewers; Textile knitting and weaving machine setters, operators, and tenders; Upholsterers; Textile, apparel, and furnishings workers, all other; Cabinetmakers and bench carpenters; Furniture finishers; Sawing machine setters, operators, and tenders, wood; Woodworking machine setters, operators, and tenders, except sawing; Woodworkers, all other; Power plant operators, distributors, and dispatchers; Stationary engineers and boiler operators; Water and wastewater treatment plant and system operators; Miscellaneous plant and system operators; Chemical processing machine setters, operators, and tenders; Crushing, grinding, polishing, mixing, and blending workers; Cutting workers; Extruding, forming, pressing, and compacting machine setters, operators, and tenders; Inspectors, testers, sorters, samplers, and weighers; Jewelers and precious stone and metal workers; Medical, dental, and ophthalmic laboratory technicians; Painting workers; Adhesive bonding machine operators and tenders; Cleaning, washing, and metal pickling equipment operators and tenders; Etchers and engravers; Molders, shapers, and casters, except metal and plastic; Paper goods machine setters, operators, and tenders; Tire builders; Helpers–production workers; Production workers, all other; Subway, streetcar, and other rail transportation workers; Transportation inspectors; Crane and tower operators; Dredge, excavating, and loading machine operators; Hoist and winch operators; Cleaners of vehicles and equipment; Machine feeders and offbearers; Material moving workers, all other; Military, rank not specified

### C.2 Tipping math derivation

As noted in the section, we need

$$\frac{\lambda}{f_k} < \frac{\partial}{\partial f_k} h_f(f_k)$$

to generate tipping.

If I assume that the preference function has the form $h_f(f) = -a + bf - cf^2$ (as it does in the survey results), I can simplify the condition to

$$\lambda < (b - 2cf)f.$$

To create tipping, this condition need not be met for all values of the occupational female share $f$, but only for some value of $f$. Thus, to see if we can get a tipping point at the easiest possible spot, we can simply check to see if the maximum of the righthand side is greater than $\lambda$.

The righthand side $bf - 2cf^2$ is maximized at $f = \frac{b}{4c}$. Evaluating at this point gives us the minimal condition that

$$\lambda < \frac{b^2}{8c}.$$

Where, for this to work, it must be the case that $0 < \frac{b}{4c} < 1$.

I can do the same exercise for men, but now in the labor supply function we take derivatives relative to male labor supply, or $1 - f$, so the original tipping condition becomes

$$\frac{\lambda}{1-f} < \frac{\partial}{\partial 1 - f_k} h_f(f_k)$$

Again assuming the quadratic form above for the composition preference function, I simplify the condition for men to

$$\lambda < (-b + 2cf)(1 - f).$$

The righthand side is maximized at $f = \frac{2c+b}{4c}$. Evaluating at this point gives us the minimal

Table C.2: Tipping Conditions: Men

| | average | male female preferring | male preferring |
|---|---|---|---|
| composition pref. $h_g(f)$ | $-.02 + .11f - .10f^2$ | $-.15 + .48f - .35f^2$ | $-.01 + .10f - .20f^2$ |
| max $\lambda$ | .0096 | 0.0180 | .0559 |
| $\lambda = 0.5$, min wtp multiplier | 52 | 27.8 | 8.9 |

Table C.3: Tipping Conditions with Endogenous Amenity: Heterogeneous preference

| | female | | male | | |
|---|---|---|---|---|---|
| | average | female preferring | average | female preferring | male preferring |
| min $\alpha$ | 0.5123 | 0.6857 | .5391 | .9632 | .5977 |

condition for tipping among men that

$$\lambda < -b + \frac{(2c+b)^2}{8c},$$

where, again, it must be the case that $0 < \frac{(2c-b)^2}{8c} < 1$.

Repeating the endogenous amenity exercise for men, I find that we require

$$\lambda < (\alpha - b) + \frac{(b - \alpha + 2c)^2}{8c}.$$

## C.3 Counterfactual Model Results

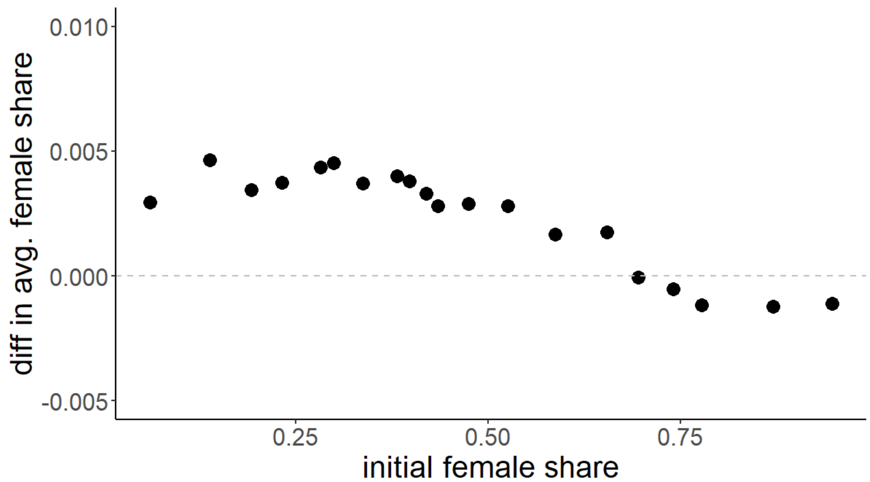


Figure C.1: Changes in Occupational Female Shares with No Composition Valuation

Note: This figure shows results from a counterfactual exercise where I eliminate composition valuations and re-calculate allocations and wages across occupation and gender. Here, I use average valuations by gender and do not allow wages to adjust in general equilibrium. This figure displays a binscatter of the difference in the female share of an occupation (counterfactual - true) against the true female share of the occupation. Bins are weighted by employment.

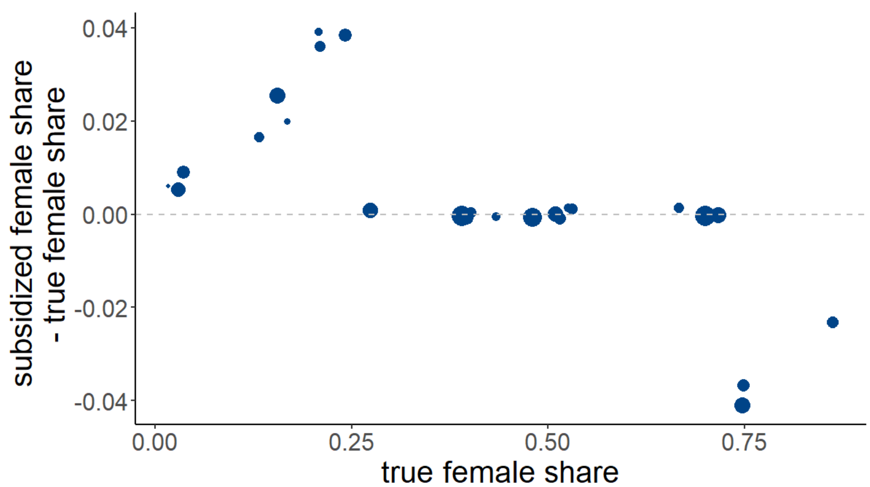


Figure C.2: Changes in Occupational Female Shares with Wage Subsidy for Gender Minorities

Note: This figure shows results from a counterfactual exercise where subsidize wages by 10 percent for women in occupations less than 30 percent female and men in occupations more than 70 percent female. Here, I use average valuations by gender and do not allow wages to adjust in general equilibrium. This figure displays a scatter of the difference in the female share of an occupation (counterfactual - true) against the true female share of the occupation. Dot size scales with employment. For this exercise, I use an aggregated group of 25 occupations.

Table C.4: Changes in Allocations: No Gender Composition Preference

| | | nested logit heterogeneous preference survey wages and shares | | | standard logit average preference CPS wages and shares | | |
|---|---|---|---|---|---|---|---|
| | | | no composition preference | | | no composition preference | |
| | | TRUE | all | female only | TRUE | PE | GE |
| share switching occupations | female | — | .008 | .008 | — | .006 | .005 |
| | male | — | .012 | .000 | — | .012 | .007 |
| avg. female share | female | .592 | .591 | .590 | .630 | .630 | .629 |
| | male | .358 | .358 | .361 | .308 | .307 | .308 |
| share in occs under 20 pct. female | female | .037 | .036 | .036 | .046 | .046 | .045 |
| | male | .233 | .227 | .220 | .396 | .398 | .395 |
| | total | .141 | .138 | .134 | .237 | .238 | .236 |
| share in occs over 80 pct. female | female | .205 | .205 | .205 | .296 | .296 | .296 |
| | male | .020 | .021 | .020 | .032 | .031 | .032 |
| | total | .107 | .107 | .107 | .152 | .152 | .152 |
| duncan-duncan seg. Index | | .405 | .404 | .400 | .490 | .490 | .490 |

Note: This table shows results from a counterfactual exercise where I eliminate composition preferences and re-calculate allocations and wages across occupation and gender. I use the WTPs for gender composition estimated in Section III.D.1. In the PE columns, I do not allow wages to adjust in response to changing allocations. I calculate the GE counterfactual only in the average preference case because distinguishing preference types in wage setting is not straightforward.

Table C.5: Changes in Gender Wage Gaps: No Gender Composition Preference

| | nested logit heterogeneous preference survey wages and shares | | | standard logit average preference CPS wages and shares | | |
|---|---|---|---|---|---|---|
| | | no composition preference | | | no composition preference | |
| | TRUE | all | female only | TRUE | PE | GE |
| Wage Gap (Female-Male) | -.259 | -.258 | -.258 | -.225 | -.224 | -.223 |
| Wage Gap (Male Wage) | -.110 | -.109 | -.108 | -.080 | -.079 | -.077 |
| Wage Gap (Female Wage) | -.081 | -.081 | -.080 | -.047 | -.046 | -.041 |

Note: The wage gap is measured as the mean male minus female wage divided by the male wage, $\frac{w_m - w_f}{w_m}$. This table shows results from a counterfactual exercise where I eliminate composition preferences and re-calculate allocations and wages across occupation and gender. I use the WTPs estimated in the latent class logit model in Section III.D.2.In the "male wage" row, I set $w_{f,j} = w_{m,j}$ and calculate overall wage gaps given true allocations, and vice versa in the "female wage" row. In partial equilibrium (PE) wages are fixed, in general equilibrium (GE) I allow wages to adjust as allocations change. I calculate the GE counterfactual only in the average preference case because distinguishing preference types in wage setting is not straightforward.