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

Rachel Schuh*

February 3, 2025

Abstract

I analyze individual preferences for workplace gender composition and their consequences for occupational segregation. To elicit these preferences, I survey 9,000 US 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 women, unmarried men, and people with more progressive gender views are more likely to prefer more female workplaces. Using these results, I estimate a structural model of occupation choice. I find that workers' composition preferences are not large enough to create tipping points, but they reduce female employment in male-dominated occupations substantially.

*Federal Reserve Bank of New York. Email: rachel.schuh@ny.frb.org. The views expressed in this paper are those of the author and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. I am grateful to my advisors Nicholas Bloom, Isaac Sorkin, Peter Klenow, and Alessandra Voena for advice and support. I would also like to thank 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. I also acknowledge financial support from the Stanford Institute for Economic Policy Research George P. Shultz Dissertation Fund and the Stanford University Graduate Research Opportunity Fund. The survey experiment was registered in the AEA RCT registry under number AEARCTR-0008414 and approved as exempt by Stanford University IRB under protocol 57557.

1 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] The causes and consequences of occupational gender segregation have long been an active area of research.[2] But one question that has received relatively little attention is whether workers actually care about gender segregation itself. That is, do they like being in segregated occupations and workplaces? Or do they prefer gender diversity?

Occupational gender segregation is often viewed as inherently harmful, but knowing its true welfare consequences requires understanding worker preferences for gender composition. Segregation resulting from discrimination certainly decreases aggregate productivity and welfare (Hsieh et al., 2018). But knowing whether gender segregation *itself* decreases welfare requires knowing what occupational gender composition workers *want*, and how much they value it. Even if policy to reduce segregation is inherently desirable, it may have limited effectiveness if workers prefer gender homophily. Pan (2015) shows that preferences for gender homophily can create tipping points where occupations rapidly shift from gender-mixed to majority female. If this is the case, policies intended to reduce segregation might have limited long-term effects—preferences for homophily will pull us back to segregation. There is some evidence that women prefer gender homophily (Larson-Koester, 2017; Wiswall and Zafar, 2018; Delfino, 2024), but we don't know what these preferences look like across the entire spectrum of gender composition or whether they can explain aggregate patterns of segregation.

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 and occupations on average, but this belies substantial heterogeneity across individuals. I then assess the aggregate consequences of my survey-estimated composition preferences in a structural model, and show that current levels of composition preferences are not large enough to create tipping points in segregation.

[1] See Tables III.3 and B.1. 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). Specific causes of occupational segregation that have been widely studied include gender differences in skills (Rendall, 2018; Yamaguchi, 2018; Cortes et al., 2021); gender differences in preferences (Wiswall and Zafar, 2018; Gelblum, 2020; Lordan and Pischke, 2022); discrimination (Bertrand and Duflo, 2017; Kessler et al., 2019); and social norms (Goldin, 2014; Bursztyn and Jensen, 2017; Bursztyn et al., 2017; Bertrand et al., 2021); gender-specific amenities (Goldin and Katz, 2016; Mas and Pallais, 2017; Morchio and Moser, 2020; Wiswall and Zafar, 2018).

However, composition preferences do reduce women's presence in majority-male occupations.

To measure the size and shape of gender composition preferences, I design a survey experiment to elicit the willingness-to-pay for workplace gender composition. In the experiment, which is designed in the vein of Wiswall and Zafar (2018), Mas and Pallais (2017), and Maestas et al. (2023), respondents make several choices between pairs of hypothetical job offers that are in the same occupation and firm but at different workplaces with randomly varying pay and demographic composition. For example, in one question respondents choose between two jobs as a sales associate at two different locations of the same retail store chain. This experimental design measures gender composition preferences at the workplace level, but it allows me to hold fixed attributes that may vary with the gender composition of an occupation such as amenity provision and tasks.[3][4]

I find that women and men both prefer gender-mixed occupations and workplaces on average, and they are willing to trade off a non-trivial portion of wages to be in a workplace with their preferred gender balance. 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 pay more for a workplace that is more than half female. 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. My results expand findings by Delfino (2024), Larson-Koester (2017), and Wiswall and Zafar (2018) by showing that *both* men and women care about gender composition and that composition preferences are highly non-linear, particularly at very segregated jobs, which is important for understanding the welfare consequences of segregation.

The average estimates of gender composition valuation mask substantial heterogeneity across individuals. I estimate a latent-class logit model (Greene and Hensher, 2003) using the survey data and find that around half of men and women choose jobs purely based on the wage and assign zero value to the female share. The remaining women prefer majority

[3]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.

[4]The gender composition preferences I measure are not necessarily inherent, biological, or constant over time, but take existing differences in behavior between men and women as given. Gender differences in behavior may include, e.g., expected discrimination or harassment (Folke and Rickne, 2022; Adams-Prassl et al., 2023; Exley and Nielsen, 2024); 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 (Macdonald et al., 2024); 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).

female workplaces and are willing to trade off over 10% of wages to avoid predominantly male workplaces. Among the remaining men, around half prefer majority male workplaces, and half prefer majority female workplaces.

Further analysis of preference heterogeneity suggests gender composition is not only a signal of other amenities but an important aspect of the social environment of a workplace. Women over 55 and women of all ages without children are the most likely to prefer female workplaces, suggesting a higher female share is not simply a signal of family-friendly amenities.[5] Young, single men are also more likely to prefer female workplaces, so the search for a romantic partner may be a motivating factor. Additionally, using follow-up questions on specific job attributes, I show that a preference for "coworker interactions" with a given gender is much more predictive of composition preferences than, e.g., tasks or promotion probability, suggesting that gender's influence on the social environment is key.

Finally, I build a model of occupation choice with composition preferences based on Pan (2015), Chen et al. (2023), and Schelling (1971) and quantify it with my survey results to study the implications of gender composition preferences for gender sorting and welfare in the US. I demonstrate that the gender composition valuations I estimate are not large enough to create tipping points in gender segregation, but they do decrease female employment by up to 10% in the most male occupations and increase the gender wage gap from occupational sorting by 10%. I also solve a social planner's problem in the model and show that reducing segregation would be optimal because composition preferences 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.

My paper's primary contribution is twofold. First, I enrich prior empirical evidence suggesting that women prefer gender homophily. Larson-Koester (2017) finds that women become more likely to choose an occupation when it becomes more female due to industrial and regional demand shifts, but we don't know if this is because of a preference for gender composition per se. Delfino (2024) finds showing a female (male) photo makes women (but not men) more likely to apply to a social work job, but we don't know how this translates to other occupations. And women are slightly more likely to choose a hypothetical job at a firm with a higher female share among mixed-gender jobs (Wiswall and Zafar, 2018), but we don't know how these preferences behave at jobs with more extreme gender compositions. My paper expands these results by causally estimating composition preferences over the entire spectrum of gender shares and showing that men and women actually both prefer

[5]More research is needed to determine if this age difference comes from an age or cohort effect, but several papers suggest that social norms which have changed over time can be an important driver of job choice across genders (Akerlof and Kranton, 2000; Brock and Durlauf, 2001; Pan, 2015; Fernández et al., 2004; Goldin, 2014; Bursztyn et al., 2017; Olivetti et al., 2020; Bertrand et al., 2021; Cortés et al., 2022).

gender diversity, on average, and that preferences are highly heterogeneous across individuals. My results for men also complement Greenberg et al. (2024), who find that integration of women into all-male military units did not affect performance but negatively affected men's perception of workplace quality.

Second, my detailed quantification of gender composition preferences allows me to directly test the theory that preferences for homophily can create tipping points in segregation (Schelling, 1971; Brock and Durlauf, 2001; Pan, 2015). I show that current levels of composition preferences are very unlikely to create tipping points. Tipping points in segregation would require either much larger composition preferences, extremely elastic labor supply across occupations, or additional amenities that are endogenous to occupational gender composition.

More broadly, my paper contributes to the literatures on the importance of amenities in job choice and the causes and consequences of occupational gender segregation. I contribute to our understanding of the importance of non-pecuniary amenities in job choice (Sorkin, 2018; Mas and Pallais, 2017; Wiswall and Zafar, 2018; Bustelo et al., 2022; Maestas et al., 2023) 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 and dynamics of inequality across demographic groups. I also show that 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 2, I discuss the design of the survey I administered to estimate gender composition preferences. In Section 3, I discuss the results of the survey and how the estimated preferences vary across individuals. In Section 4, I outline a basic model of occupation choice with gender composition preferences and show how composition preferences affect equilibrium sorting. Section 5 quantifies this model using my estimated composition preferences, which allows me to measure the aggregate importance of composition preferences and discuss the utility of alternative allocations. Section 6 concludes.

2 Survey Design

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.

In the experiment, respondents choose between several pairs of workplaces within a fixed occupation and firm, and I randomly vary the pay and gender composition of each workplace. This allows me to identify the willingness-to-pay for workplace gender composition while holding other attributes fixed.

2.1 Workplace Choice Experiment

The survey experiment is designed to quantify the willingness-to-pay for the female share of a workplace within a fixed occupation and firm. By randomly varying the pay and gender composition of hypothetical job offers, I can causally identify the dollar value of workplace gender composition across the entire spectrum of workplace female shares, from 0 to 100 percent female. By holding fixed the occupation, firm, and unlisted characteristics in each hypothetical choice, I can isolate the value of the experience with a particular set of coworkers and prevent inference about job amenities and tasks from the gender composition.

In this experiment, respondents make several choices between pairs of job offers which are in the same occupation and firm 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. Prior to making the choices, I instruct respondents that they will be performing the same tasks at both workplaces, the workplaces are identical in all but the listed characteristics, they have the necessary skills for both jobs, and they have offers for both jobs and can choose the workplace they prefer. This helps ensure that respondents hold fixed other job attributes while making their choices.

An example of a question respondents might see in this section is included in Figure 1. 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).[6] The level of each characteristic is varied randomly and the demographic shares are expressed as a number of coworkers out of ten. The wage offers are anchored to the average wage in the occupation, as estimated in the March CPS, and vary in increments of 0.5% to 10%.[7] Across respondents, I vary the order in which the demographic

[6]Gender and demographic 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.

[7]I round the average wage w to the nearest \$10 and randomly select from six possibly multipliers to

shares are presented and whether the gender composition is presented as the female share or the male share.

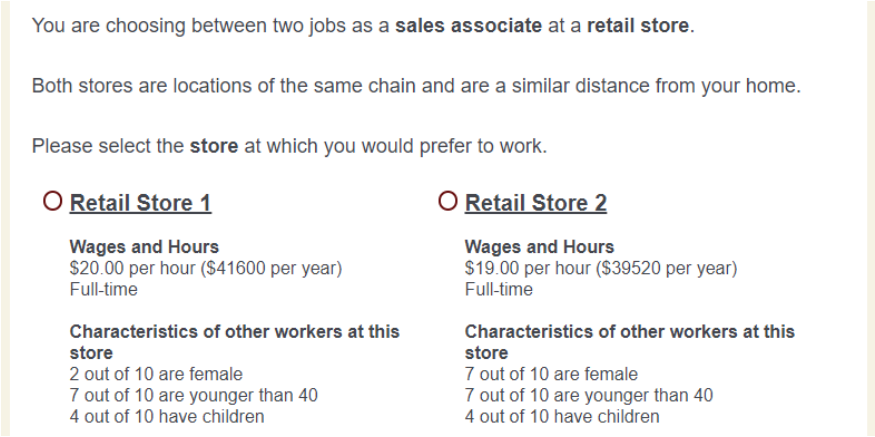


Figure 1: Example Hypothetical Workplace Choice

Respondents see nine to eleven job choices, within five occupations: high school teacher, insurance sales agent, retail sales associate, nurse, and software developer. 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. Each respondent sees two to three pairs of jobs within each occupation.[8]

In addition to enabling me to quantify the willingness-to-pay for different gender compositions across a range of occupations, this experiment holds fixed many attributes that respondents might expect to vary with the gender composition of an occupation. Respondents know that the two jobs are the same and even at the same firm, so they should not infer that the tasks are different. In addition, giving a choice between two jobs at different locations of the same firm implies that the contracted amenities should be the same. For instance, whether schedule flexibility and paid parental leave are offered should not differ across locations of the same firm within a local area.

Although the contracted amenities and tasks can reasonably be expected to be constant across the workplaces within an occupation and firm, changing the gender composition may still change a number of characteristics of the work experience, and thus I consider this a bundled treatment. First, the gender composition may change the social environment: how pleasant it is to interact with coworkers, how likely are friendships or romantic relationships

generate each wage pair. Thus, every wage is some value $m \cdot w$ where $m \in \{1, .995, .99, .975, .95, .9\}$.

[8] An example of the flow of hypothetical choices respondents might see in this section is included in Figure I.1 in the online appendix.

to form, how likely is harassment. In addition, the gender composition may change the opportunities: the likelihood of begin promoted or being assigned certain tasks within the job description. Finally, the gender composition may change the perception or availability of contracted amenities: if there are already many women with young children, a worker may expect less schedule flexibility for themselves; if most coworkers are male, taking parental leave even if offered may carry a stigma. In an additional follow up section of the survey, I attempt to parse in which way gender composition matters for job experience.

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 see the pay and demographic composition of each job but no information about which occupation they will perform. 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. 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, rather than choosing between two jobs with varying amenities in a fixed occupation, respondents choose between two different occupations, and I vary whether they see the true gender composition of each occupation.

2.2 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. Following this, respondents see two of the job choice experiments: the main workplace choice, or the other variants described above. I randomize the order of the two experiments and the order of questions within the experiment block. After the experimental questions, 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. Next, I ask respondents for their opinions on three gender attitudes using questions adapted from the General Social Survey and Pew American Trends Panel. Finally, I include a debrief question where I show respondents one of the choices they made in the workplace choice experiment

and ask them to explain in writing why they made that choice.[9]

2.3 Estimation Strategy

Given the randomized design of the hypothetical job choice experiment, the estimation process is straightforward. For the workplace choice, I use a 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.

2.3.1 Workplace Choice: Willingness-to-Pay for Gender Composition

To estimate the average willingness-to-pay (WTP) for the gender composition of an occupation, I use a standard multinomial logit model. 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

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

where $f_n = 0, .1, ..., .9, 1$ are the eleven (minus one) possibilities for the female share in the survey. As in the model in Section 4, I assume the preference shocks are Type I Extreme Value with shape parameter λ. Then, the probability that individual i chooses alternative j among the choices in J takes the conditional logit form

$$P_{ij,J} = \frac{\exp((\beta_w \log(w_{g,j}) + \sum_{f_n} \beta_{f_n} \mathbb{1}(f_n = f_j) + \kappa X_j)/\lambda)}{\sum_{j' \in J} \exp((\beta_w \log(w_{g,j'}) + \sum_{f_n} \beta_{f_n} \mathbb{1}(f_n = f_{j'}) + \kappa X_{j'})/\lambda)} \tag{2}$$

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

Importantly, the "pure" wage elasticity β_w and gender composition elasticities β_{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 λ. 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

[9] A diagram showing the survey flow is shown in Figure I.2 in the online appendix. The full survey is available upon request.

to gender composition relative to the wage elasticity.

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:[10]

$$WTP_f = 1 - \exp\left(\frac{-\hat{\beta}_f}{\hat{\beta}_w}\right). \tag{3}$$

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.

2.3.2 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 heterogeneous 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, 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 J is

$$P[\text{alternative } j \text{ chosen by individual } i \text{ in choice } J| \text{ class } q] = \frac{\exp(x'_{iJ,j}\beta_q)}{\sum_{j=1}^{J}\exp(x'_{iJ,j}\beta_q)} = P_{ij,J|q}, \tag{4}$$

where $x_{iJ,j}$ are the characteristics of choice j in situation J and β_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 A.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 individual preferences across the entire spectrum of gender compositions would require observing a large number of choices

[10] Details on the steps to this willingness-to-pay formula are described in Appendix A.1.

per person.[11] Asking respondents to make twenty or more choices of the same format is not desirable due to respondent fatigue and inattention. 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.

3 Survey Data and Results

I fielded my survey online to a sample of approximately 9,000 respondents. I find that, on average, both men and women prefer jobs and workplaces that are 50-70% female. However, women have slightly more distaste for majority male jobs and men slightly more distaste for majority female jobs. There is substantial preference heterogeneity across individuals.

3.1 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.[12] Of the total respondents, about half (4,604) see the main workplace choice experiment.[13]

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 3.5. 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 shown in Table 1.[14] My survey respondents are relatively younger and more educated than

[11] 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.

[12] Lucid, now owned by Cint, sources participants from multiple companies that recruit individuals to take online surveys. Respondents may have been recruited through emails, push notifications, in-app pop-ups, or through sites offering multiple survey opportunities. (See `https://theorem.cint.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.

[13] The online appendix shows the number of survey respondents who see each variation of the hypothetical job choice experiment.

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

the US population.[15] In Appendix A, I show that my sample is similar to representative samples on other characteristics that one might expect to vary with gender composition preferences, including realized occupational gender composition and gender attitudes.

Table 1: Survey Demographics vs. March CPS

	CPS Share	Survey Share	Difference	t-stat
Age				
18-24	.11	.12	.01	3.28
25-34	.17	.19	.02	5.84
35-44	.14	.20	.06	13.93
45-54	.16	.17	.01	2.00
55-64	.18	.15	-.03	-8.25
65+	.24	.17	-.07	-18.88
Education				
less than high school	.10	.04	-.06	-29.97
high school diploma	.29	.28	-.01	-2.82
some college	.28	.36	.08	15.67
bachelor's degree or greater	.33	.32	-.01	-2.08
Race				
White or Caucasian	.73	.71	-.03	-5.50
Hispanic or Latino	.11	.08	-.03	-9.44
Black or African American	.09	.12	.03	8.46
Asian or Pacific Islander	.05	.05	.00	.11
Native American or Alaskan Native	.01	.01	.01	6.39
Other	.01	.03	.02	9.46
Gender				
female	.52	.52	-.00	-.22
male	.48	.48	-.00	-.59

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.

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. Approximately 70% of respondents pass all attention checks. Most importantly, I include an attention check within the hypothetical job choice conjoint, where respondents see a choice

[15]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.

between two jobs that differ only in their wages. Depending on the exact design of the job choice question, 71% to 87% of respondents choose the higher wage job. This is similar to the 85.5% of respondents passing a similar attention check found by Mas and Pallais (2017), which is reassuring given theirs was an incentivized choice between two real jobs and mine is hypothetical. In the main analysis, I exclude respondents who failed one or more attention check, but results are similar if I include them. This leaves me with 2,772 total respondents for the workplace choice.[16]

3.2 Workplace Choice Results: Average Preferences

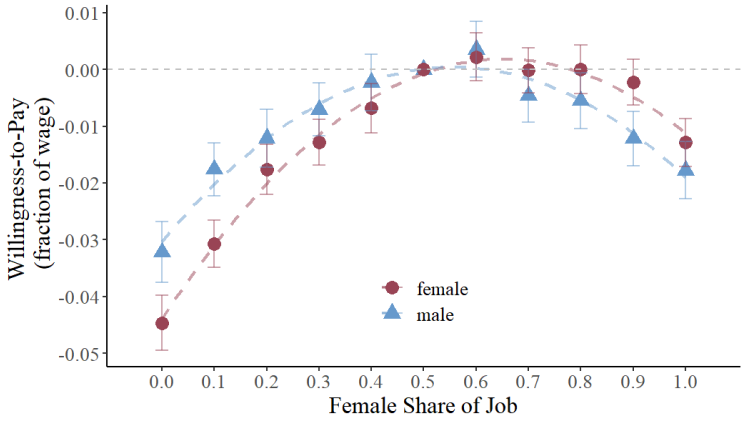


Figure 2: Average Willingness-to-Pay for Gender Composition

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.

The results of the workplace choice experiment show that men and women in fact prefer jobs that are gender mixed or moderately female-dominated, on average, and dislike workplaces that are very segregated in either direction. Figure 2 shows the estimated WTPs for each possible female share separately for men and women, and Table 2 adds the results for 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 2. Women dislike mostly male workplaces: they are willing to pay 4.5% of wages to avoid

[16] In Appendix A Section A.4, I detail additional results on the attention checks and the within-respondent consistency of responses within the job choice experiments.

[17] Notably, the coefficient on the log wages for men and women is very close at 28.54 for women and 26.77 for men, respectively.

an all-male workplace, and have a negative WTP for any female share below 50%. Women do not gain additional utility for female shares above 50%: the WTP is approximately zero for 50-90% female workplaces. However, complete segregation is not desirable: the WTP is negative for female shares above 90%. Overall, women dislike being a gender minority, but do not want complete gender segregation.

Men have a strikingly similar composition preference profile to women: they prefer gender-mixed workplaces. However, they have slightly more distaste for majority female workplaces than women do. Men's WTPs are negative for female shares below 40% and above 60%, and they are indifferent between workplaces that are 40-60% female. On the whole, men dislike gender-segregated workplaces, but they have greater distaste for majority male than majority female workplaces.

One might be concerned that the wide variation in female shares across jobs creates choices between unrealistic scenarios. I perform several robustness checks to allay these concerns. First, I limit the sample to questions where the difference in female shares between the two jobs is no more than 10 percentage points (e.g. 0% vs. 10% female). Second, I limit the sample to questions that present respondents with female shares that are within 30 percentage points of the reported female share of their own job. Finally, I limit the sample to questions that present female shares within 30 percentage points of the true female share of the occupation listed in the question. None of these sample restrictions change the results materially, seen in Appendix Figure A.2.

My WTP estimates enrich prior research in this area. I expand Wiswall and Zafar (2018)'s estimates of women's small positive WTP for more female workplaces between 26-72% female by showing that the distaste for extremely male jobs is even larger but adding more women above 50% makes little difference.[18] My estimates also help explain Delfino (2024)'s finding that women are more likely to apply for a social work job advertised with a female photo and Larson-Koester (2017)'s observational estimates of women's preference for more female occupations.[19] Further, my findings that men have smaller composition preferences than women and are largely indifferent among mixed-gender jobs clarifies the null results Larson-Koester (2017) and Delfino (2024) find for men. Folke and Rickne (2022) find a 10% WTP to avoid a job with a reported incident of sexual harassment, and an increasing

[18] Wiswall and Zafar (2018) 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. The male shares in their survey range from only 26% to 72%, so we cannot extrapolate these estimates to the entire range I measure.

[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.

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 coefficients on wages for women and men are 28.54 and 26.77, respectively.

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 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, and their average valuations are quite similar. Women prefer jobs that are at least half, but not all female, and dislike the most male jobs. Men also dislike male dominated jobs and like gender mixed jobs, but they have slightly more distaste for mostly female jobs. These results suggest that composition preferences are unlikely to be the primary driver of segregation since workers actually prefer diversity.

3.3 Heterogeneous Valuation Results

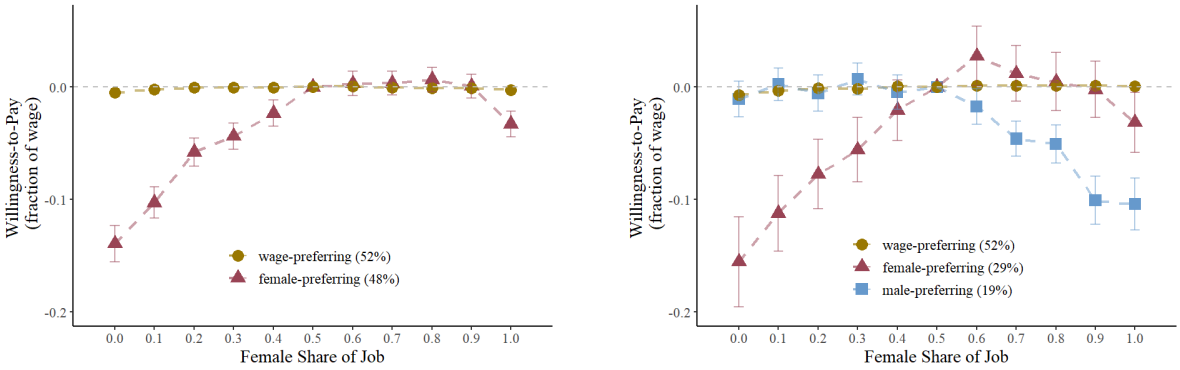
I estimate a latent class logit model to detect underlying heterogeneity in gender composition valuations. I find that some women have large valuations for majority female jobs and some do not value the female share of a job at all. I find that some men do not value the gender composition of a job, some prefer more female jobs, and some prefer more male jobs.[21]

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 3 Panel a. 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.

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: 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 has marginally negative, but small, WTPs for jobs less than 20% female. The WTPs for other demographic characteristics in this class are small and not significantly different from zero.

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

[21]I select the number of classes using cross-validation. I run a 10-fold cross validation exercise on the log-likelihood and evaluate the Akaike and Bayesian Information Criteria (AIC and BIC) for models with 1-10 classes. 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. Through these exercises, I choose $Q = 2$ classes for women and $Q = 3$ classes for men. Details are in the online appendix.



(a) Female (b) Male

Figure 3: WTP for Gender Composition by Class, Male

Note: This figure shows the willingness-to-pay, as a fraction of the wage, for the classes among women and men estimated using the latent class logit model. Bars show 95% confidence intervals estimated using the delta method.

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.[22] This class has large, negative WTPs for all female shares below 50%, and 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.

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 3 Panel b. 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

[22]The wage coefficient is 15, implying an 81% probability of choosing a job where the wage is 10% higher, all else equal.

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.

The female-preferring class among men assigns the most value to workplaces that are 50-90% female. For this class, the willingness-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. 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.

Overall, the latent class logit picks up on important heterogeneity in gender composition preferences among female and male respondents. Some women value more female workplaces and some do not care. Men's average preference for diversity masks 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.[23]

3.4 Who Values Gender Composition and Why?

Looking at how composition preference heterogeneity varies with observable characteristics can help us understand why men and women value gender composition. I find that older workers are more likely to prefer same-gender workplaces, unmarried men are more likely to prefer more female workplaces, and people with more progressive gender attitudes are more likely to prefer more female workplaces. Additionally, workers who say they prefer the interactions with coworkers (as opposed to, say, the schedule or tasks) at a majority female job are much more likely to belong to the female-preferring group. Together, this suggests that social factors, rather than correlated amenities, are important determinants of composition preferences.

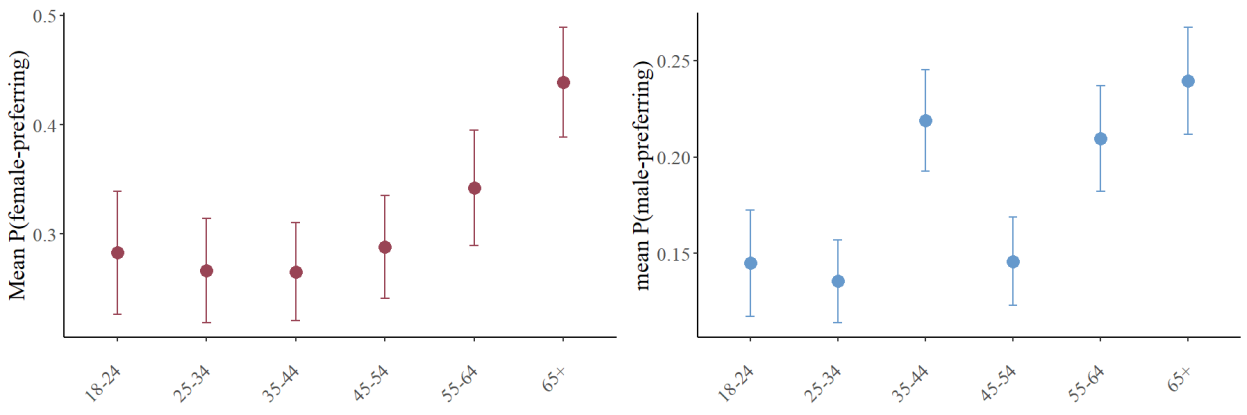
Older women and men are more likely to prefer a majority same-gender workplace. This pattern is more distinct for women—in Figure 4, I compare the probability of belonging

[23]See Table I.6 for all WTP estimates by class.

Table 3: Correlates of Preference Classes

	Probability of Preference Class		
	Female	Male	
	Female-preferring	Female-preferring	Male-preferring
Demographics			
Age: 55+	0.123	-0.045	0.063
	(0.022)	(0.024)	(0.018)
Race: white	0.108	0.036	0.033
	(0.023)	(0.026)	(0.020)
Married	0.026	-0.046	0.027
	(0.021)	(0.023)	(0.018)
Children at home	-0.061	0.009	-0.017
	(0.022)	(0.025)	(0.019)
Education: some college +	0.007	-0.045	0.001
	(0.023)	(0.026)	(0.020)
Income: >$25/hour	-0.020	0.006	-0.036
	(0.028)	(0.030)	(0.024)
Attitudes			
Democratic political views	0.056	0.043	-0.015
	(0.015)	(0.016)	(0.012)
Progressive gender norms	0.101	0.064	-0.055
	(0.024)	(0.025)	(0.019)
Prefer female job in listed attribute			
Coworker Interactions	0.137	0.118	-0.090
	(0.017)	(0.021)	(0.016)
Work Environment	0.155	0.095	-0.087
	(0.017)	(0.020)	(0.016)
Schedule	0.112	0.053	-0.072
	(0.019)	(0.023)	(0.018)
Tasks	0.090	0.067	-0.043
	(0.019)	(0.021)	(0.017)
Higher earnings in female job	-0.034	0.081	-0.040
	(0.016)	(0.020)	(0.015)
Family prefers female job	0.057	0.089	-0.067
	(0.021)	(0.022)	(0.017)
Promotion more likely in female job	0.009	0.104	-0.045
	(0.015)	(0.018)	(0.014)

Note: This table shows the coefficients from regressions of the posterior probability of belonging to each preference class on various questions. For women, estimates come from a 3-class model, shown in the online appendix. I run separate regressions for each variable. For democratic preference, progressive gender norms, and job attribute preferences, I transform variables to a -1 to 1 index with higher values representing, respectively, more democrat, more progressive norms, and more preference for a female job. The full gender attitude questions are shown in A.3. Observations: 1525 female, 1234 male.



(a) Female (b) Male

Figure 4: 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.

to the female-preferring group across age bins. The probability is relatively flat from ages 18 to 54, but increases at age 55 and linearly with age thereafter. From the survey data alone, I cannot determine if this is an age effect or a cohort effect. If it is an age effect, it may be that older workers generally value amenities more and wages less because they have higher savings and less time remaining in the labor market, or that preferences change or are learned with age. If it is a cohort effect, it may be the case that as gender norms and discrimination have lessened over time (Cortés and Pan, 2018; Hsieh et al., 2018), women and men have become more comfortable in mixed-gender work environments.

Further, Table 3 shows that women with children are actually less likely to belong to the female-preferring group.[24] As shown in Table 2, people with children have a slight preference for other coworkers with kids, but it is much smaller than women's preference for women coworkers. Along with the evidence that older women prefer female coworkers more, this suggests a high female share is not simply serving as a signal for the availability of family friendly amenities.

In addition, married men are less likely to belong to the female-preferring group, even after controlling for age. Perhaps unmarried men could be looking for a romantic partner in the workplace. In the context of recent research showing both the frequency and costs of romantic relationships in the workplace (Macdonald et al., 2024) and sexual harassment in the workplace (Folke and Rickne, 2022), this suggests that some men's desire for female coworkers may be at odds with women's desire to avoid sexual harassment, which could inform the policy implications of these results.

[24]This result holds when comparing instead women with kids at home to women with no kids at home.

Both men and women with more progressive gender attitudes are more likely to prefer majority female workplaces, as shown in Table 3.[25] Democrats are also more likely to belong to the female-preferring group.[26] This suggests that people who prefer more female workplaces may be in general more tolerant of women.

To get more information about the precise attributes that workers value at more and less female jobs, in an additional survey section I ask respondents whether they expect they would be more satisfied with various attributes in a majority female or majority male occupation. These attributes include interactions with coworkers, work environment, tasks, and schedule. I also ask in which job they think they would earn more, in which job they think they would be more likely to be promoted, and which job their family would prefer they do. Notably, this is slightly different from the workplace choice experiment in that the occupations may be different, and I do not instruct respondents that the jobs are identical other than the female share.

Figure 5 shows the share of respondents who report that they would prefer the more female job minus the share who report that they would prefer the more male job for each attribute, separated by respondent gender. Women on average are more likely to prefer the female occupation on all dimensions except pay. Men are more likely to prefer the more male job or are nearly indifferent on all dimensions except coworker interactions, where they prefer the more female job. Because the context of this question is different from the hypothetical workplace choice (focusing on different occupations, not necessarily at the same firm with otherwise identical characteristics) this does not imply that my composition preference estimates simply reflect gender composition as a signal of other amenities.

Comparing stated preferences for workplace characteristics with the results of the hypothetical workplace choice shows that coworker interactions are particularly important in explaining composition preferences. Table 3 shows that for both men and women, indicating a preference for the coworkers or work environment in a female (male) job is the largest predictor of belonging to the female (male) preferring group. This suggests that the preferences I measure are indicative of actual preferences for coworker gender, not just a preference for correlated amenities.

Additional results lend further credence to the idea that I am measuring actual preferences for coworkers rather than just correlated attributes. For one, men are more likely to choose more female workplaces in occupations that involve more direct interaction with coworkers, as illustrated in Appendix Figure A.3. Further, in an extension of the survey design where

[25] I construct a progressive gender norms index from three questions about gender norms. Full gender norms questions are described in Appendix A.3.

[26] Political affiliation is collected by Lucid when respondents sign up for the platform; it is not included in my survey.

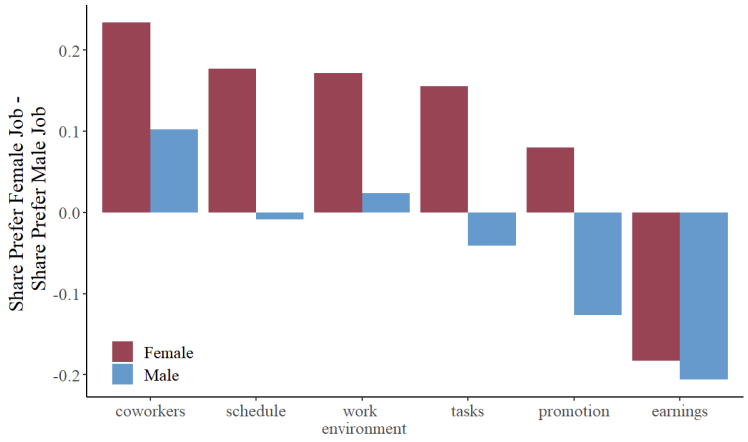


Figure 5: Share Preferring Female Job Minus Share Preferring Male Job on Each Attribute

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.

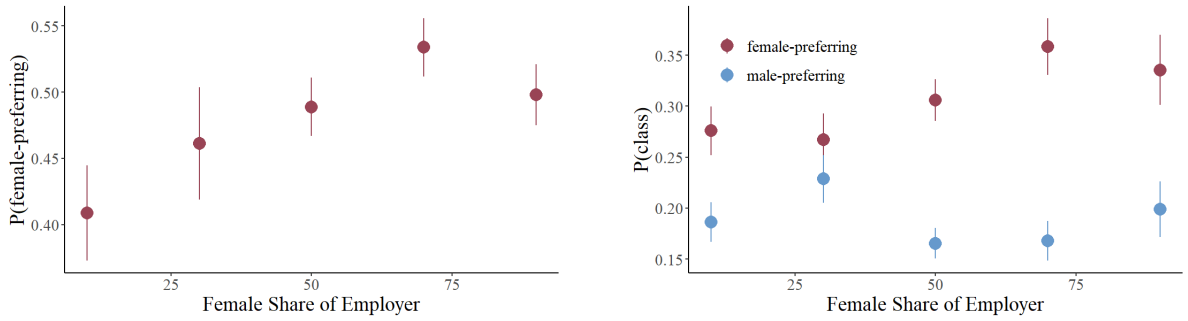
I add additional amenities, including schedule (fixed, variable, or flexible); hours (part-time or full-time); and probability of promotion or layoff; composition preferences retain the same overall shape.[27] Because these attributes are held fixed by design in the workplace choice experiment, it seems unlikely that a large share of the WTP is coming from inferred amenity provision.

Overall, these results suggest gender's effect on the social experience of work is a key driver of my measured composition preferences. Older women, women without children, single men, and individuals with more progressive gender attitudes value female coworkers more, suggesting that the social aspect of working with women rather than correlated family-friendly amenities is behind the preference for female coworkers. In addition, preferring the "coworker interactions" with a given gender is a strong predictor of a higher WTP to work with that gender, and adding additional amenities to the survey questions does not change the overall shape of composition preferences. It seems as a whole that gender composition may fundamentally affect the social environment of a workplace or occupation.

3.5 Do Composition Preferences Correspond to Job Choice?

In this section, I evaluate the external relevance of the survey-estimated valuations by measuring whether they correspond to real-life occupation choices.

[27]See Appendix Figure A.4. Although the magnitude of the estimates decreases somewhat, overall the preferences have the same shape. The sample size for this extension is about one-third of that for the main workplace choice, and the number of attributes is larger, so the estimates are much noisier.



(a) Female (b) Male

Figure 6: 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.

As shown in Figure 6, women and men who report working for employers with a higher female share are more likely to belong to the female-preferring class, although the probability declines somewhat for people who report the highest female shares at their employer. Men who report low female shares at their employers are also slightly more likely to belong to the male-preferring group. The relationship between preferences and reported female shares is stronger for women, perhaps indicative of more strongly-held preferences.

Of course, men and women might choose jobs that do not align with their valuations for gender composition because these jobs are preferable for other reasons. This demonstrates the importance of holding "all else equal" in my survey design—in reality, these jobs are not otherwise identical. For instance, male software engineers might prefer firms with more female engineers, but in reality these firms may be hard to find or specialize in different areas. 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.

4 Occupation Choice Model

To illustrate the intuition of how composition preferences affect gender sorting, I present a simple model of occupation choice in which workers value both the wage and gender composition of their occupation, in the spirit of Pan (2015) and Chen et al. (2023). I show that homophilic preferences amplify underlying patterns of gender sorting.

4.1 Model Environment and Sorting Equilibrium

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 $f_k = \frac{l_{k,f}}{(l_{k,f}+l_{k,m})}$, where $l_{k,g}$ is the

mass of workers of gender g in occupation k. Each gender has a 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, female share, and idiosyncratic preference draw $\varepsilon_{i,k}$. The utility function for an individual i of gender g is given by

$$U_i = \max_{k \in \{1,2\}} \left[\log(w_{g,k}) + h_g(f_k) + \varepsilon_{i,k} \right]. \tag{5}$$

The idiosyncratic shock creates 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.

Individual i will choose occupation 1 if it provides higher utility:

$$log(w_{1,g}) + h_g(\ell_{1,f}/\ell_1) + \varepsilon_{i,1} > log(w_{2,g}) + h_g(\ell_{2,f}/\ell_2) + \varepsilon_{i,2}. \tag{6}$$

Thus, if wage and gender composition utility is higher in occupation 1, individuals will choose this occupation unless their idiosyncratic utility draw is very high in occupation 2, and vice versa.

Following Chen et al. (2023), 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:[28]

$$\ell^*_{1,g} = Pr\left[log(w_{1,g}) + h_g(\ell_{1,f}/\ell_1) + \varepsilon_1 > log(w_{2,g}) + h_g(\ell_{2,f}/\ell_2) + \varepsilon_2 \right], \qquad g = \{f, m\}. \tag{7}$$

Tipping, which I define as a shift between two distant equilibria brought on by a small shock to the system, will only be possible if multiple equilibria exist. In the next section, I analyze which conditions create the possibility for tipping and characterize the effect of composition preferences on sorting.

4.2 The Partial Equilibrium Effect of Composition Preferences

In this section, I use the model to evaluate how composition preferences affect gender sorting when wages are fixed. 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

[28] Chen et al. (2023) establish conditions for existence and uniqueness of a sorting equilibrium: an equilibrium generally exists and will be unique if composition preferences are smooth and of small enough scale.

to wages and if occupation choice is more elastic.

We can invert Equation 7 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 λ. The parameter λ is the inverse of the occupational wage elasticity: if λ is large, the variance of the preference shocks is high, so occupation choices are sticky and workers are less responsive to changing occupational conditions. If λ is small, the variance of the preference shock is low, so occupation choices are flexible and workers are 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:

$$\underbrace{[\log(w_{1g}) - \log(w_{2g})]}_{\text{wage utility diff.}} = \underbrace{\lambda \cdot \log\left(\frac{\ell_{1g}}{1-\ell_{1g}}\right)}_{\text{pref shock utility diff.}} + \underbrace{\left[h_g\left(\frac{1-\ell_{1f}}{2-\ell_{1f}-\ell_{1m}}\right) - h_g\left(\frac{\ell_{1f}}{\ell_{1f}+\ell_{1m}}\right)\right]}_{\text{composition utility diff.}}, \quad g = f, m. \tag{8}$$

The first part of Equation 8 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 shows that 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. Figure 7 plots the female labor supply function with composition preferences of varying strength and varying wage elasticity, given a fixed male labor supply to occupation 1. This figure illustrates the key effects of homophilic composition preferences.

First, homophilic preferences flatten the labor supply curve and thus amplify sorting. We see this in Figure 7 by comparing the solid and dashed grey lines: if occupation 1 is already more female, composition preferences will increase female labor supply and depress female wages in this occupation, which can be seen on the right half of the figure.

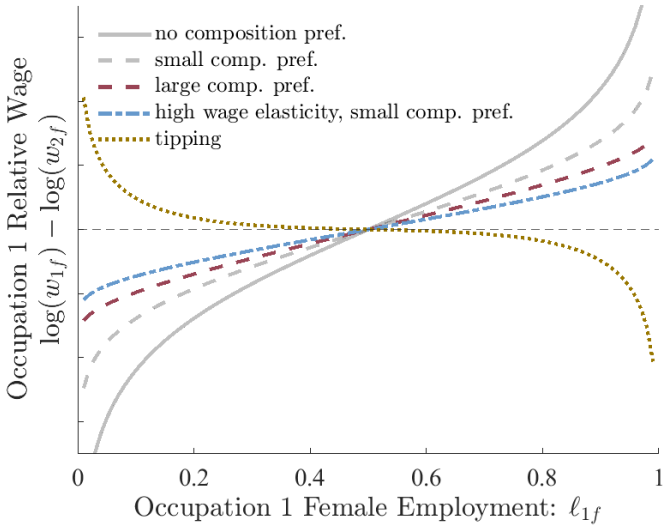


Figure 7: Labor supply with varying composition preferences and wage elasticity

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 η (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.

Second, the effect of composition preferences on sorting will be larger when composition preferences are steeper relative to preference shocks and wage differences. We see this in Figure 7 by comparing the dashed grey and dashed red lines: as the composition preference function increases in slope, the labor supply function becomes flatter.

Finally, the effect of composition preferences will also be larger when the occupational wage elasticity is large; that is, when the variance of the preference shock is small. We see this in Figure 7 by comparing the dashed grey and dashed blue lines: when λ is smaller, occupation choices are more responsive to changing wages, so the effect of the same composition preference function (normalized to wages, i.e., equal willingness-to-pay for gender composition) will be larger.[29]

When are tipping points possible in this model? I define a tipping point as a shift between multiple sorting equilibria. 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 by the yellow line in Figure 7. When the female labor supply function is backward bending, a change in relative wages 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

[29]In Online Appendix Section III.1, I show that the effect of composition preferences is dampened, but not nullified, in general equilibrium.

strong or occupation choices are very flexible, because both of these conditions will create a flatter labor supply function. 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.

5 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 an expanded version of the model in Section 4, I show that existing composition preferences are unlikely to create tipping points and reducing segregation would improve welfare.

5.1 Quantitative Model Environment

I make three additions to the quantitative model relative to the simple model: multiple occupations, unobserved amenities, and nested preference shocks. First, there are now K occupations indexed by k rather than only two occupations.[30] Second, each occupation has a gender- and occupation-specific residual amenity $a_{k,g}$, representing additional factors like other amenities and barriers to entry, and allowing me to match observed employment shares.[31] Third, the nested logit structure allows preference shocks to be more correlated within an occupation nest than across nests. Therefore, workers that switch occupations are more likely to move to occupations that are more closely related to their own.[32]

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 4, $w_{k,g}$ 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 λ serves the same function as the shape parameter of the taste shock in the standard logit model—the larger λ is, the smaller the occupational wage elasticity is within nests.[33]

[30] Non-participation is considered its own occupation.

[31] See Grigsby (2024).

[32] Formally, the nested logit model breaks the independence of irrelevant alternatives inherent to the standard multinomial logit model.

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

The share of workers in occupation k within nest s is

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

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

$$P(s) = \frac{\exp(Z_s'\alpha + \lambda IV_s)}{\sum_l \exp(Z_l'\alpha + \lambda IV_l)} \quad \text{where} \quad IV_s = \log\left[\sum_{k\in S}\exp(V_k/\lambda)\right].$$

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. This structure ensures that in counterfactuals, workers will be more likely to switch to an occupation that is more closely related to their initial occupation.

To close the model, I add a production function which is CES across occupations.

$$Y = \left(\sum_k (A_k\ell_k)^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}}.$$

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

To account for gender wage gaps within occupations, the within-occupation 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 α 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. This wedge between male and female productivity can encompass both actual productivity differences and pure wage discrimination.

5.2 From Survey to Model

Equating occupational gender composition utility 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. Because the model is static, we can think of it as a one-time, early-in-life occupation choice. Second, I assume that labor market frictions imply individuals will not be able to freely choose a workplace with their preferred gender composition within their occupation.[34]

Equating occupational composition utility to workplace composition utility may matter little for the overall results. Preferences across gender compositions are strongest at the extreme ends of the distribution. Within extremely segregated occupations, workplaces whose gender compositions diverge greatly from the occupational average may 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.

5.3 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.

To apply my survey-estimated composition valuations in the model, I fit quadratic functions to the heterogeneous composition values separately for men and women. I calibrate λ, 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.[35]

[34] How realistic this is depends on how gender shares are distributed across workplaces within an occupation. At the extreme end, firms could completely undo sorting across occupations if there are firms of every possible gender composition within all occupations. However, this is unlikely: Hellerstein et al. (2008) estimate that as of 2000, 61% of gender segregation across firms was explained by gender segregation across occupations.

[35] This parameter is roughly equivalent to the variance of the type I EV preference shock in a standard logit model. My value of $\lambda = 0.5$, corresponding to a wage elasticity of 2, 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).

Next, I calibrate the production functions elasticities externally. I set the elasticity of substitution across occupations η in the aggregate production function equal to 1.5, based on estimates of the elasticity across skilled and unskilled occupations.[36] Thus, occupations are gross substitutes. I set the elasticity of substitution across genders α equal to 2.5, following estimates in Ngai and Petrongolo (2017). Within occupation, genders are also gross substitutes, but genders are more substitutable than occupations.[37] Table 4 summarizes the values and sources of parameters that enter the model estimation.

Table 4: 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		
η	Sub. Elasticity Across Occupations	1.5	Acemoglu and Autor (2011)
α	Sub. Elasticity Across Genders	2.5	Ngai and Petrongolo (2017)

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

To divide occupations into nests, 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.[38]

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.

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

[37] A non-infinite elasticity of substitution across genders can be due to differences in production stemming from, e.g., differences in hours or specific tasks performed within an occupation.

[38] The estimated nests are described in the online appendix.

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.

5.4 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 would require very large occupational wage elasticities. 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 λ is the inverse wage elasticity.[39]

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 female employment 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

$$h_g(f_k) = a + bf_k - cf_k^2, \quad \text{we need} \quad \lambda < \frac{b^2}{8c}$$

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

to create the possibility of a tipping point.[40]

Using typical 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 5, I show the minimum wage elasticity (with estimated preferences) and minimum preference multiplier (with a standard wage elasticity) necessary for tipping. Even for the female-preferring women, a very large wage elasticity (around 14) would be necessary to create a tipping point. With a more standard wage elasticity of 2 ($\lambda = 0.5$), the composition preferences of the female-preferring group would need to be about 7 times as large as those I estimate to create tipping.

Table 5: Conditions for Tipping with Estimated Composition Preference

	female		male
	average	female preferring	average
composition pref. $h_g(f)$	$-.04 + .14f - .11f^2$	$-.13 + .41f - .30f^2$	$-.02 + .11f - .10f^2$
max λ	0.0224	0.0693	.0096
$\lambda = 0.5$, min wtp multiplier	22.3	7.2	52
min $\lvert\alpha\rvert$, $\lambda = 0.5$	0.5123	0.6857	.5059

Note: This table shows the fitted polynomials for composition preferences and the conditions on preferences, the inverse wage elasticity λ, and the endogenous amenity function necessary to create a tipping point. The inverse of the maximum λ is the smallest wage elasticity needed to create a tipping point. The parameter α measures the responsiveness of amenity levels to female share: $\alpha = 0.5$ indicates that a 10 percentage point increase in the female share increases the amenity value by the same amount as a 5% wage increase.

Alternatively, 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 with 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 can calculate the minimum α required to create tipping. I tabulate these estimates in Table 5 for both women (requiring a positive α) and men (requiring a negative α).

For the average woman and man, α would need to be around 0.5 to generate tipping. What does this mean? The value $\alpha = 0.5$ 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 and if allowing for discontinuous responses to the female share.[41] Thus, investigating the endogeneity of

[40] Appendix B.1 provides the derivation of this condition, and the tipping condition for men.

[41] Mas and Pallais (2017) find a 20% WTP to avoid a schedule that is set on short notice, and Wiswall

amenity provision and social norms to the demographic composition of a job could be a promising avenue for further research.[42]

On the whole, I conclude that the gender composition preferences I estimate are, on their own, unlikely to create broad-based tipping points. Tipping is more likely to occur when the occupational wage elasticity is high, which could occur between closely related occupations or over longer periods of time. Given that I find that older women have stronger preferences for higher female shares, 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.[43]

5.5 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.

To remove gender composition valuations, I take the model estimates of occupational productivity and preference parameters, set the composition value $h_g(\ell_{g,k}/\ell_k)$ to zero, and solve for the counterfactual allocations. I first estimate the new allocations in partial equilibrium, using heterogeneous preference groups.[44]

Figure 8 illustrates that absent composition valuations, female employment would increase in mostly male occupations. I estimate that if workers did not value gender composition, female employment would rise by 10% in the most male occupations and fall slightly in mixed gender occupations. The changes in men's employment have similar patterns, but are more muted. Thus, workers' valuations of gender composition have nontrivial effects on occupational employment, particularly for women in very male occupations.[45] If I decrease λ to increase the wage elasticity to 20 and make occupation choice more flexible, the no composition preference counterfactual produces larger movements across occupations, and

and Zafar (2018) find that women have a 7% WTP for an option to work part-time, e.g.

[42]Corradini et al. (2022) find a Brazilian trade union's female-centric agenda increased the provision of female-friendly amenities and increased female demand for these jobs.

[43]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 endogenous amenities or social norms. Although Schelling (1971) suggests that tipping can occur with very small composition preferences, this result assumes that items in the choice set are otherwise identical.

[44]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.

[45]I include additional statistics on the changes in wages and overall gender segregation absent composition preferences in the online appendix. Overall, the changes in employment are not large enough to have a substantial effect on the average degree of segregation.

we see female employment more than doubling in the most male occupations. However, this still doesn't change gender composition much—the female share is so low in the most male jobs that it only increases by about 5 percentage points on average even with these large employment shifts.

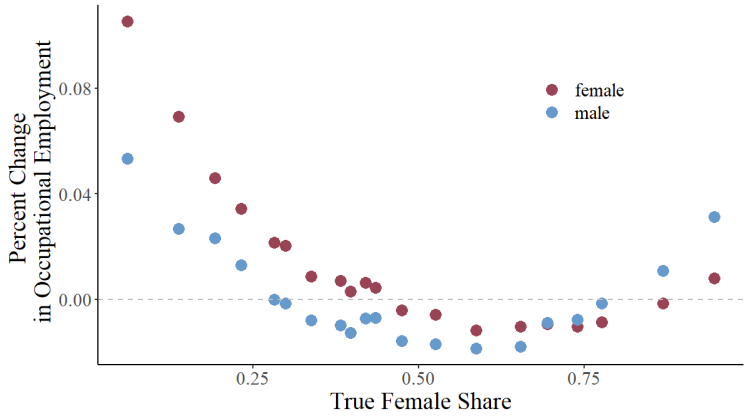


Figure 8: Change in Occupational Employment if Composition Preferences are Removed

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. 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.

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.[46] These results remain similar when I vary the production elasticities across occupations and genders.

A simple way to measure the cost of segregation is to calculate the equivalent variation (EV) for making one's occupation 50% female. The average EV is positive—0.4% for women and 1% for men—but it is larger for the female-preferring group of both men and women—0.7% for female-preferring women and 2.8% for female-preferring men. Notably, men gain more utility from desegregation than women, because they are less likely to be sorted in accordance with their composition preferences. The EVs vary widely across individuals, with those in very segregated jobs benefiting much more from desegregation.[47]

[46] In Appendix Table B.1 and Online Appendix Table III.3, I include full results from this exercise.
[47] The online appendix provides further detail of the EVs, including their distribution across individuals.

5.6 Correcting Sorting Externalities

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 reduce segregation. A minority-gender-specific wage or amenity subsidy in very segregated occupations could achieve similar utility gains.

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), and simulating is computationally infeasible with a large number of workers.

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:

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

The additional terms 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 solve for a new decentralized equilibrium which locally approximates the social planner's solution. In this exercise, I use a simple logit model rather than the nested logit model for simplicity.

Figure 9 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%.

The social planner reallocates 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 in mostly male occupations is large. Men and women also dislike extremely female dominated

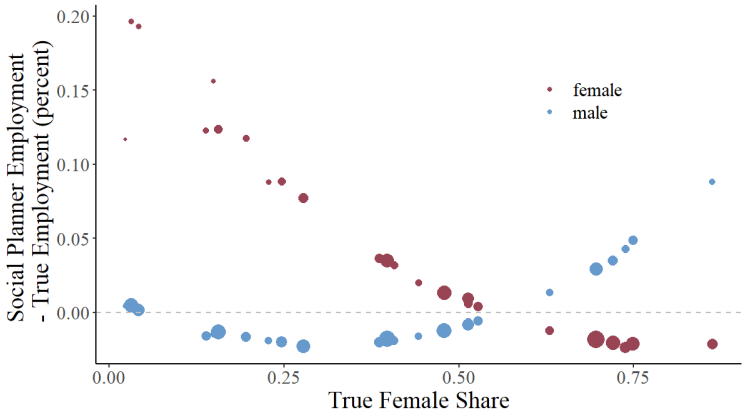


Figure 9: Changes in Occupational Employment in Social Planner's Solution

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.

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.[48]

The social planner overcomes 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.

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 demonstrates that policies that reduce the cost of entry into an occupation for the gender minority could be fruitful. For instance, policies that improve female-centric amenities or subsidize education for women in male-dominated occupations could help reduce segregation and benefit workers of all genders.

[48]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.

6 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. Older men and women are more likely to value gender homophily while men and women with more progressive gender views are more likely to value female coworkers, suggesting that gender norms may be important in explaining these preferences.

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 tipping points in segregation are unlikely to result from current composition preferences. However, composition preferences do matter for sorting. 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 0.4% wage increase for women and a 1% 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.

Chen, Haoning, Miaomiao Dong, Marc Henry, and Ivan Sidorov, "Occupational segregation in a Roy model with composition preferences," *arXiv preprint arXiv:2012.04485v2*, 2023.

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.

_ , 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," *American Economic Review*, 2024, *114* (6), 1816–1853.

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.

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.

Greenberg, Kyle, Melanie Wasserman, and E Anna Weber, "The Effects of Gender Integration on Men: Evidence from the US Military," Technical Report, National Bureau of Economic Research 2024.

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.

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.

Macdonald, David, Jerry Montonen, and Emily Nix, "DATING AND BREAKING UP WITH THE BOSS: BENEFITS, COSTS, AND SPILLOVERS," 2024.

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.

Appendix A Data

A.1 Calculating the Willingness-to-Pay

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 ε_{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

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

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 two jobs. To express this as a percent of the wage in job 1, I take $1 - w_2^*/w_1^*$, so

$$WTP = 1 - \exp\left(\frac{-\beta}{\beta_w}\right). \tag{15}$$

A.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

$$P_i = \sum_{q=1}^{Q} H_{iq} P_{i|q}, \tag{16}$$

and the log-likelihood for the sample is

$$\ln L = \sum_{i=1}^{N} \ln P_i = \sum_{i=1}^{N} \ln \left[\sum_{q=1}^{Q} H_{iq} \left(\prod_{t=1}^{T_i} P_{it|q} \right) \right]. \tag{17}$$

I estimate H_{iq} and the parameters β jointly in a one-step maximum likelihood.

A.3 Survey Descriptive Statistics

Whether survey results can be generalized depends on the heterogeneity of the treatment effect of interest and of the survey sample. Reassuringly, several papers have found similar results in experiments using online samples from Lucid and traditional random samples[49] In addition to the demographics in Table 1, I also compare the distribution of respondents across other variables that are expected to co-vary with composition preferences, and find that they are broadly consistent with the CPS. I compare the sample across other demographics in the Online Appendix.

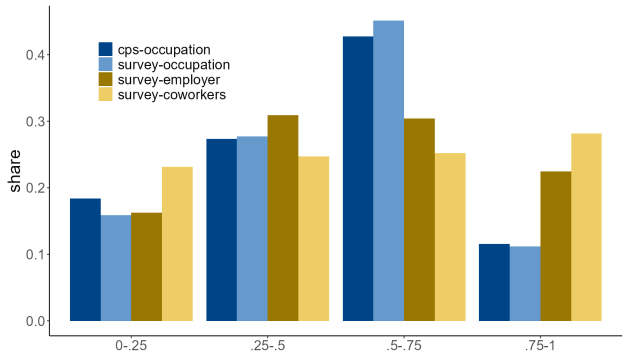


Figure A.1: Female shares: Survey vs. CPS

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 2010 Census occupation codes. Survey-employer and survey-coworkers refer to questions in which respondents report the female share of their employer and coworkers with their same occupation at their employer.

I compare gender attitudes in my survey to population-representative random samples from the General Social Survey (GSS) and the Pew American Trends Panel (Pew):

1. *"Do you 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."'* (GSS)

2. *"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?"* (GSS)

3. *"All things considered, who do you think has it easier in our country these days?"* (Pew)

[49]See, e.g., Coppock and McClellan (2019) and Boas et al. (2020).

Table A.1: Gender Attitudes

	Survey Mean All	Survey Mean Attentive	Rep. Sample Mean	Source
Against man achiever, woman home	0.066	0.146	0.369	GSS
For affirmative action for women	0.040	0.034	-0.249	GSS
Men have it easier	0.269	0.317	0.261	Pew

This table shows the average value of each gender norms response transformed to a numerical index. The index goes from -1 to 1 with higher values representing more progressive norms. Full distributions of the answers are shown in the online appendix.

A.4 Data Quality and Robustness

Using several attention checks, I find that the majority of respondents are attentive, and rates of attention are similar to incentivized job choice experiments found in the literature.[50] My attention checks are: (1) Type number "13" in a text box (beginning of survey, demographics section), (2) Select "disagree" in multiple choice (end, gender attitudes section), (3) Hypothetical job choice with identical jobs EXCEPT wages (middle, experimental section). Responses are shown in Table A.2. [51]

Table A.2: Attention check passing rates

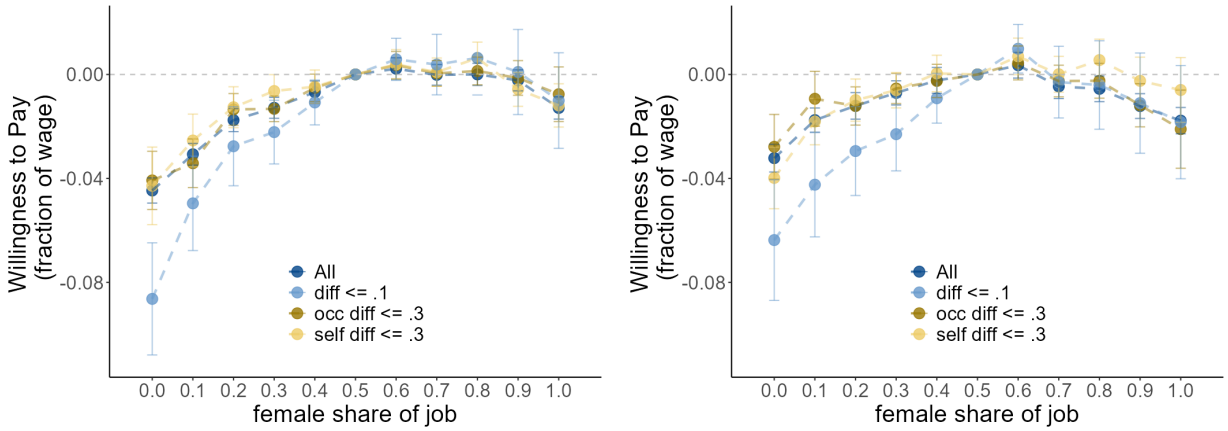
Question	Share correct	Total responses
text entry	0.98	8850
multiple choice	0.87	8850
workplace choice (choose higher wage)	0.71	4604
occupation choice (choose higher wage)	0.87	6538
all attention checks correct	0.69	8850

This table shows the share of respondents who answered each attention check question correctly. Workplace choice refers to the choice between two workplaces in a fixed occupation (here, teacher) and occupation choice refers to a choice between two jobs where the specific occupation is not listed.

I also evaluate the internal consistency of hypothetical choices. First, I compare the choices of individual respondents who saw the same question twice and find that 87% made the same choice both times. I also check for consistency across questions where individuals chose between jobs with the same attributes twice, but possibly with different wages. I use each answer to create bounds on the WTP estimate and find that the bounds are interally consistent in 98% of such cases. Further details are included in the online appendix.

[50] The average respondent completed the survey in 10.6 minutes, and the median in 7.6 minutes, details in online appendix.

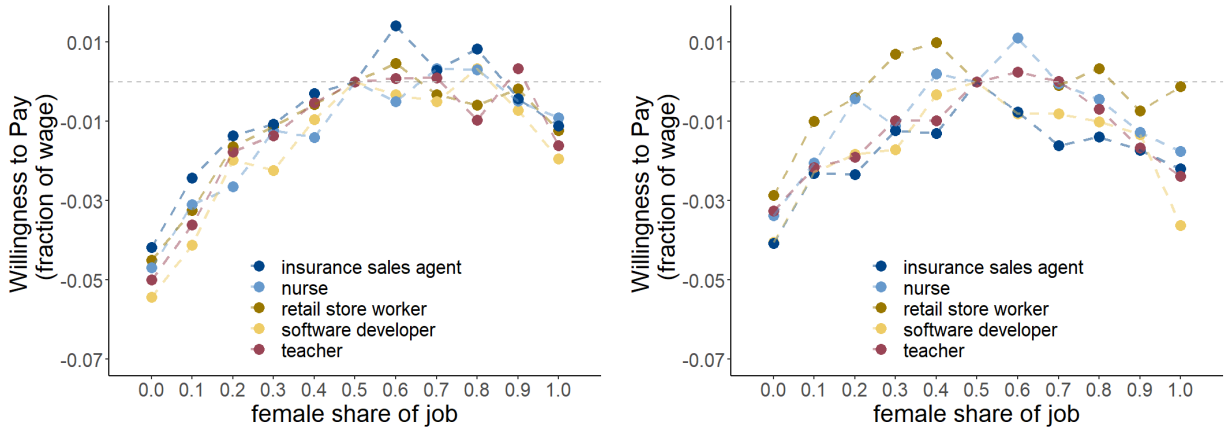
[51] Alvarez et al. (2019) found that respondents who failed attention checks completed surveys faster, had a higher incidence of "don't know" responses, and provided lower-intensity responses regarding attitudes and behavior.



(a) Female (b) Male

Figure A.2: WTP for Female Share, Robustness Checks

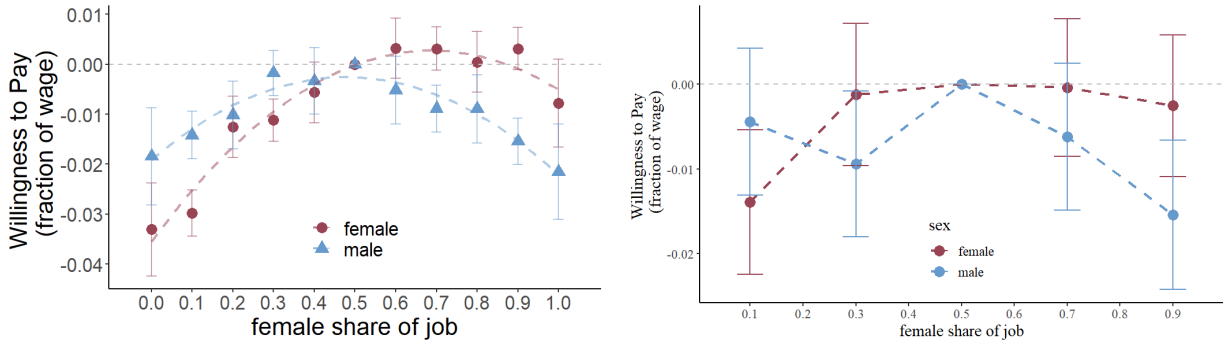
Note: This figure plots the willingness-to-pay for each possible female share for women and men. "All": full sample (16495 questions to 1525 women and 13334 questions to 1234 men). "diff $\leq$ .1": maximum difference of 10 ppt. between hypothetical female shares (4723 questions to women and 3713 questions to men, same individuals as full sample). "occ diff $\leq$ .3": difference between the self-reported female share of the respondent's employer and the hypothetical female share is less than 30 ppt. (4885 questions to 1404 women, 4005 questions to 1157 men). "self diff $\leq$.3": difference between the true female share in the occupation in the hypothetical question and the hypothetical female share is less than 30 ppt (6523 questions to women and 5227 questions to men, same number of respondents as full sample). True female shares: software engineer–20% female, insurance sales agent–50% female, retail sales–50% female, teacher–60% female, nurse–90% female. Calculated using the ACS from 2015-2019 via Flood et al. (2018).



(a) Female (b) Male

Figure A.3: 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.



(a) No Occupation Listed in Choice (b) Additional Amenities

Figure A.4: WTP for Female Share, Experiment Variants

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 on the left, and with no occupation listed where additional amenities are included on the right. Amenities include schedule (fixed, flexible, or variable); hours (part-time or full-time); probability of promotion; and probability of layoff; as well as the original demographic characteristics. Sample: left–14,400 questions to 2703 women and 12,385 questions to 2319 men; right–5,382 questions from 897 women and 4,962 questions from 827 men.

Appendix B Model

B.1 Tipping math derivation

As noted in Section 5.4, to generate tipping we need

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

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 only be met for some value of f, so we can simply check to see if the maximum of the righthand side is greater than λ. The term $bf - 2cf^2$ 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

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

which I simplify using the quadratic preference 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 condition for

tipping among men that

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

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

Including endogenous amenities gives the following conditions:

$$\lambda < \frac{(b+\alpha)^2}{8c} \text{ for women, and } \lambda < (\alpha - b) + \frac{(b-\alpha+2c)^2}{8c}, \text{ for men.}$$

B.2 Model Results

Table B.1: 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 model. 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.