Gender-Targeted Job Ads in the Recruitment Process:

Evidence from China

June 28, 2016

Peter Kuhn, UC Santa Barbara

Kailing Shen, Australian National University

Explicit requests for applicants of a specific gender are widely used in emerging-economy labor markets. The extent to which these requests affect the allocation of workers to jobs depends on two factors: how much workers *comply* with firms' requests in their application decisions, and how much firms *enforce* their own requests when they encounter gender-inappropriate applicants. Using internal data from a Chinese job board, we show that both compliance and enforcement are substantial, but compliance accounts for the vast majority of the gender segregation associated with employers' explicit gender requests. Using firm*occupation fixed effects for our compliance analysis and worker fixed effects for our enforcement analysis, we argue that both effects are causal, in the sense that (a) changing the gender label on an ad leads to large differences in the gender mix of applications that arrive, and (b) the same worker is much less likely to be called back when applying to a job requesting the other gender than when applying to an observationally-identical ad with no stated gender preference.

This research is supported by the National Natural Science Foundation of China through Grant No. 71203188, "Impacts of *Hukou*, Education and Wage on Job Search and Match: Evidence Based on Online Job Board Microdata". We thank Austin Jones for careful research assistance

Statements in a job ad that either men or women are preferred for the job are widely used in many countries. This practice, which we call gender *profiling,* has been studied by Kuhn and Shen (KS, 2013) in China, and by Delgado Helleseter, Kuhn and Shen (DKS, 2016) in China and Mexico. Based on these studies, a number of empirical regularities have been established. For example, gender profiling appears to be relatively *symmetric*, with a roughly equal number of job ads requesting men and women, even within job skill levels. Profiling is also *job-specific* in the sense that a substantial share of the variation in requested gender occurs across jobs within the same firm. In addition, gender profiling (in both directions) is much more common in jobs requiring low skill levels, whether measured by education or experience requirements or the offered wage. This phenomenon, which we call the (negative) *skill-targeting relationship*, is dramatically reflected in the fact that almost three quarters of ads on a job board serving unskilled workers are explicitly gender-targeted (Delgado Helleseter, Kuhn and Shen 2016). Finally, DKS show that the direction of explicit gender profiling in both Chinese and Mexican job ads strongly favors women at young ages and men at higher ages (the *age twist* in gender targeting).

While the above results suggest that gendered job ads may play an important role in allocating labor in emerging-economy labor markets, direct evidence of such an effect does not yet exist. This is because all existing research on gender profiling we know of is based on samples of job *ads* only. Thus, while existing work has documented what employers *ask for* (in terms of employee gender) in different types of jobs, we do not yet know whether gendered job ads have real effects on where workers send their job applications, or on which workers ultimately get hired. At one extreme, advertised gender requests could be 'hard' requirements in the sense that gender-mismatched applications are always rejected. If so, one might expect workers' application behavior to strongly conform to firms' stated requests. At the other extreme, advertised gender requests could just be 'soft' suggestions that a particular gender is preferred, or even that a particular gender might prefer working in that job (for example due to the presence of same-sex co-workers). In this 'soft' case, gender-inappropriate applicants might fare quite well when they apply.

To measure how gendered job ads affect workers' application decisions and employers' callback behavior, this paper studies applicant and callback pools to job ads on a Chinese job board (XMRC.com) over a six-month period in 2010. A key advantage of this data is that --in addition to knowing the characteristics of all the ads (including the requested gender, if any)-- we know the gender and qualifications of every person who *applied* to each ad, and of persons who were *called back* to each ad. For gender-targeted ads, this allows us to measure the total amount of *gender-matching* that results from the recruiting process, i.e. the extent to which the pool of successful applicants (i.e. callbacks) matches the firm's explicit requests. We can also partition the total amount of gender-matching into portions attributable to self-selection by workers, i.e. to applicants' *compliance* with employers' gender requests when deciding where to send their applications, and to employers' *enforcement* of their own requests when choosing workers from the applicant pool. For our entire sample of ads (including the non-gendered ones) we can also measure the total amount of gender *segregation* across jobs, firms and occupations, and pose two questions, neither of which to our knowledge has been answered before. First, what is the contribution of explicit gender profiling to the amount of gender segregation among successful job applicants across jobs, firms and occupations? Second, does the observed level of segregation –whether linked to profiling or not-- result mostly from workers' self-sorting in deciding where to apply, or from employers' active selection among applicants that arrive? Existing studies of occupational segregation have not been able to address this question due to the absence of data on workers' application behavior.

Our main results are as follows. First, we find that total gender-matching is high: 95.1 percent of callbacks to gendered jobs are of the requested gender. Worker compliance is also high, since 92.5 percent of *applications* to gendered job ads are also of the requested gender. Firms' enforcement decisions reinforce these application patterns, and while enforcement is substantial it is far from complete: Among applicants to 'female' jobs, men are 75 percent as likely to get a callback as women. Among applicants to 'male' jobs, women are 43 percent as likely to get a callback as men. In an accounting sense, the vast majority of gender-matching between actual callback pools and firms' gender requests can be attributed to applicants' compliance, or *self-sorting* decisions. Second, again in an accounting sense, we calculate that 60 (60) [50] percent of the gender segregation across *all* jobs (firms) and [occupations] on this job board is associated with the explicit gender labels attached by employers to jobs. Like our results for gender matching, self-selection decisions by workers account for almost all of this label-linked segregation, and indeed for the vast majority of all gender segregation across the jobs, firms and occupations in our sample.

Third, regression analysis strongly suggests that the apparent effects of gender profiling in the above decompositions are mostly, if not exclusively causal consequences of the explicit gender label attached to a job ad. Concerning workers' application decisions, simple means indicate that 55 percent of applications to non-gendered jobs are from men, compared to just 7 percent of application to jobs requesting women, a reduction of 48 percentage points. Controlling for an extensive list of ad and job characteristics including fixed effects for occupation * firm cells, this reduction hardly changes at all: to 47 percentage points. Simple means also indicate that 45 percent of applications to nongendered jobs are from women, while only 8 percent of applications to 'male' jobs come from women, a decline of 37 percentage points. Adding the same set of controls reduces the magnitude of this decline to 24 percentage points. Thus, while about (.37 - .24)/.37 = 35 percent of the unadjusted 'effect' of a male job label on the gender mix of applications is a spurious consequence of the fact that certain occupations attract male applicants regardless of how they are labeled, none of the unadjusted effect of a female job label is attenuated by firm*occupation controls. We conclude that explicit gender requests in job ads convey information to applicants that is not otherwise available, and that this information *directs workers' applications* towards jobs where their gender is in greatest demand.

Concerning employers' callback decisions, simple means indicate that women's callback rate to nongendered job ads is 8.8 percent, compared to 4.1 percent in jobs requesting men, suggesting a mismatch penalty for women of 4.7 percentage points (or 53 percent). Adding controls for worker fixed effects –thus comparing the same woman applying to observationally identical jobs that differ only by their preferred gender-- reduces this gender-mismatch penalty only slightly, to 4.3 percentage points. Men's recall rate is 9.2 percent in nongendered jobs versus 5.4 percent in jobs requesting women, suggesting a male mismatch penalty of 3.8 percentage points (or 41 percent). Adding the same set of controls reduces this estimated penalty to 2.6 percentage points. Thus, women face a substantial callback penalty when applying to 'male' jobs, and essentially none of their apparent penalty in the raw data results from nonrandom selection of the women who apply to male jobs. Not only do men face a smaller 'raw' callback penalty when applying to female jobs, but about (.38 - .26)/.38 = 32 percent of this unadjusted callback penalty is spurious, in the sense that it is associated with negative selection of the men who apply to explicitly female jobs. To our knowledge, these callback penalties encountered by workers applying to 'wrong'-gender jobs are the first *within-worker* estimates of the effects of gender discrimination based on naturally-occurring job applications.

1. Related Literature

While there is a large literature on gender differentials in labor markets, almost none of it has focused on the explicit gender-typing of jobs that is widely used in emerging-economy labor markets. Thus our paper fills an important gap in the gender-differentials literature, especially in reference to the labor markets in which most of the world's workers are employed. In studying the effects of job ad content on application patterns, our paper also relates to a substantial theoretical literature on directed search in labor markets (e.g. Albrecht et al., 2006). With a few recent exceptions, this literature has not examined data on how workers direct their applications. These exceptions include Marinescu and Wolthoff (2015) who use Careerbuilder data to study the effects of the posted wage on the number and quality of applicants, and Marinescu and Rathelot (2015) who study the geographic scope of workers' search. Kudlyak, Lkhagvasuren and Sysuyev (2014) study how workers re-direct their search (according to the job's education requirements) over the course of a search spell.

Our paper also relates to an emerging literature that uses the contents of job ads to study labor markets. Such job-board studies include Herschbein and Kahn (2015) and Modestino, Shoag and Balance (2015) both of whom ask whether employers request higher qualifications for the same jobs ('upskilling') when local labor market conditions make workers "easier to get". Brencic and Norris (2009, 2010, 2012), and Brencic (2010, 2012) use the same type of data to study aspects of employer search strategies, including whether to post a wage and whether to adjust ad contents during the course of recruitment.

Finally, a large literature now studies which job applications receive callbacks using resume audit methods (Bertrand and Mullainathan 2004, Oreopoulos 2011, Kroft et al. 2013, Neumark et al. 2015). While studying how employers treat resumes from different applicants is important, our results suggest that the vast majority of gender segregation in labor markets is attributable, instead, to workers' decisions on where to send their resumes. While some sociologists have studied workers' application choices (Pager and Pedulla 2012), to our knowledge workers' application behavior has been almost completely unstudied by economists. Job-board-based studies like ours offer the potential to address this large gap in the literature.[1]

2. Data

As noted, our data consist of internal records of XMRC.com, an Internet job board serving the city of Xiamen. XMRC is a private firm, commissioned by the local government to serve private-sector employers seeking relatively skilled workers.[2] Its job board has a typical U.S. structure, with posted ads and resumes, on-line job applications and a facility for employers to contact workers via the site. It is nationally recognized in the job-board industry as dominant in Xiamen, possibly due to its close links with the local government and social security bureau.

To study the effect of gender profiling on application and callback patterns, we began with the universe of ads that received their first application between May 1 and October 30, 2010. We then matched those ads to all the resumes that applied to them, creating a complete sample of applications. Finally, for the subset of ads that used XMRC's internal messaging system to contact applicants, we have indicators for which applicants were contacted after the

[1] We also address the question of which applicants get callbacks, using worker fixed effects rather than random assignment to control for unobserved heterogeneity. Unlike audit studies which typically target a specific subset of jobs, our focus is on a population of naturally-occurring applications on a citywide job board. Also, most audit studies appear to only submit resumes that satisfy the employer's requested criteria, whereas our focus is on the effects of mismatches between employers' requests and workers' attributes. In earlier work on the XMRC data, we have used a similar approach to study employers' choices between under- and over-educated applicants (Kuhn and Shen 2013b) and between native and migrant workers (Kuhn and Shen 2015).

[2] The other major local job site, XMZYJS, is operated directly by the local government. It serves private sector firms seeking production and low-level service workers. Unlike XMRC, XMZYJS does not host resumes or provide a service for workers and firms to contact each other through the site.

application was submitted. This indicator serves as our measure of callbacks. Our primary dataset for the entire paper is this subset of ads for which callback information is available, which comprises 3489/41467 = 8.4 percent of all ads. Summary statistics for this sample are, however, very similar to the universe of ads, shown in Appendix Table A1. We also replicated our entire analysis of worker compliance (self-sorting) --which does not require callback information-- on the entire universe of ads in Appendix Table A2 with very similar results.

In all, our primary dataset comprises 221,135 applications made by 78,031 workers (resumes) to 3,489 ads, placed by 1,551 firms, resulting in 18,731 callbacks. Thus there was an average of 63 applications per ad and 5.4 callbacks per ad. One in twelve applications received a callback, while one in four resumes received a callback.[3] Descriptive statistics are provided in Tables 1 and 2 for ads and applications respectively. Table 1 shows that 840/3489 = 24 percent of ads requested female applicants, 18 percent requested male applicants and the remaining 58 percent did not specify a preferred gender.[4] The average number of education years requested was 12.2, and was more than a year higher in jobs requesting women than men. Forty-eight percent of ads specified a preferred worker age; the mean requested age was 28. Consistent with the age twist identified in DKS (2016), the requested age was considerably lower for jobs specifically requesting women. On average, one year of experience was requested. 58 percent of ads posted a wage; the mean posted wage was 2434 RMB per month overall but only 1983 RMB in jobs requesting women.

Table 2 shows that 120,172/221,135 = 54 percent of applications came from women. The typical *application* had 14.35 years of education, with women holding about half a year more education than men. Average applicant age was 24.7 years. Other applicant characteristics observed in our data (and used in the regression analysis) include experience, new graduate status, marital status, current wage (when provided), myopia, height, the number of experience and job spells listed, and whether an English version of the resume is available.

3. Descriptive Analysis

3.1 Total gender-matching (G) and its components

Descriptively, our first goal is to measure the extent to which the final pool of successful applicants to a job ad (i.e. the *callback pool*) reflects the employer's stated gender preferences. This concept of gender-matching, *G*, applies only to explicitly gendered ads. We also wish to measure the relative contributions of workers' *compliance* with firms' requests and employers'

[3] This analysis sample is identical to the one used in Kuhn and Shen (2013); additional details on its construction are available there.
[4] This compares to 20, 18 and 62 percent in the universe of job ads. See Table A1.

enforcement of their own stated requirements to the total amount of gender-matching that occurs. The analysis begins with some basic descriptive statistics on applications and callbacks in Table 3.

Starting with total gender-matching, row 1 of Table 3 shows the share of callbacks that are female (δ) for the three job types in our data: jobs requesting women (*F* jobs), jobs requesting men (*M* jobs) and jobs that do not state a gender preference (*N* jobs). These statistics indicate a high congruence of the callback pool with firm's stated requests. Specifically, 94.4 percent of callbacks to *F* jobs are female and 100 – 3.6 = 96.4 percent of callbacks to *M* jobs are male. Combining *F* and *M* jobs, 95.1 percent of callbacks to gendered job ads are of the requested gender. Row 2 shows the share of *applications* to the three job types that are female (α). It suggests that applicants' compliance with employers' gender requests plays a substantial role in accounting for this high level of gender-matching, since applicant pools are almost as highly sorted by gender as callback pools. Specifically, 92.7 percent of applications to *F* jobs are female and 100 – 8.0= 92.0 percent of applications to *M* jobs are male. Combining *F* and *M* jobs, 92.5 percent of applications to gendered job ads are of the requested gender.

The remaining rows of Table 3 show that employers' 'enforcement' of their own stated requests also helps to account for the overall amount of gender matching that occurs. Specifically, in jobs explicitly requesting female applicants, men who do apply are only 1/1.337 = 74.8 percent as likely to be called back as women. In jobs requesting men, female applicants are only 43.1 percent as likely to be called back as a man. Thus, at least in the raw data, employers' enforcement of their own gender requests is stronger against women applying to male jobs than men applying to women's jobs.

To get a better sense of the overall amount of gender-matching and its components, it is useful to define the following index of gender-matching:

$$G = \frac{g - g_0}{1 - g_0} \qquad (1)$$

where g is the share of gendered ads that are of the requested gender and g_0 is the share of gendered ads that *would be* of the requested gender if there was no gender-matching (i.e. if we re-allocated the total population of called-back workers across all jobs --whether *F*, *N* and *M*-- so that the total number of callbacks to each job remained the same, but the gender mix of callbacks was equalized across all jobs). Thus G=1 if all callbacks to gendered jobs match the employers' request, and G=0 if the female share of callbacks (δ) equals its population average in all jobs. In our data, g = .951 and g_0 = .502, so our overall index, G = .902. In other words, on a scale where zero indicates no gender matching and 10 indicates perfect matching, the total amount of matching equals 9.

With this index in hand, we can now assess the relative contributions of compliance and enforcement to gender-matching, G, using the identity:

$$\delta^J = \frac{\theta^J \alpha^J}{\theta^J \alpha^J + (1-\alpha^J)} \qquad (2)$$

where $J = F, N,$ or M and θ is women's relative risk of being chosen from the applicant pool, i.e. the ratio of callback rates (*f/m*). Equation (2) allows us to compute two counterfactual levels of g and G.[5] *Counterfactual 1* (no compliance) keeps enforcement, θ, at its actual level in each of the three job types, but sets α (the share of women in the *applicant* pool) at its population mean level in all jobs (i.e. at .543, from Table 3). *Counterfactual 2* (no enforcement) keeps compliance, α, at its actual level in each job type, but sets θ (women's relative risk of being picked from the applicant pool) at its population average (.863) in all jobs. The results are reported in Table 4.

According to row 2 of Table 4, eliminating worker compliance while maintaining actual levels of enforcement would reduce the share of callbacks that are of the requested gender, g, from .951 to .631. The corresponding decline in the gender-matching index, G, is from .902 to .259. Thus, workers' compliance with employers' gender requests accounts for (0.902-0.259)/0.902 = 71 percent of the gender-matching in our data. According to row 3, eliminating employers' enforcement while maintaining actual levels of worker compliance would have a much smaller impact, reducing g from .951 to .921 and G from .902 to .842. Thus, workers' compliance with employers' gender requests accounts for only (0.902-0.842)/0.902 = 7 percent of the gender-matching in our data. Because the decomposition in equation (2) is exact but nonlinear, the remaining 22 percent of gender matching is due to the interaction between compliance and enforcement.[6] We conclude that *compliance, i.e. applicants' self-sorting according to employers' gender requests in job ads, accounts for the vast majority of gender matching in gendered ads.* The intuition is straightforward: Because applicant pools are so highly gender-segregated, even completely equal treatment of male and female applicants in all job types would have only a small impact on the gender mix of callbacks to each job.

3.2 Gender profiling and gender segregation (S)

---

[5] Like other indices used in this paper, the G index depends on the relative sizes of the three job types (J), as well as on the overall share of workers who are called back to each job type. Throughout the paper, we design our counterfactual thought experiments to hold both of these quantities constant, varying only the gender *mix* of workers who apply to different job types (or firms, occupations, etc.) and the gender *mix* of callbacks.

[6] By 'exact' we mean that eliminating both compliance and enforcement would reduce G to zero.

In this second part of our descriptive analysis, we broaden our focus beyond the gendered jobs to all the jobs in our sample. Motivated by evidence of high levels of gender segregation across occupations (Blau et al. 2013), across firms (Card et al. 2016) and even across jobs within firms (Bielby and Baron 1984), we wish to assess the contribution of the explicit job labels (*F*, *N* and *M*) to gender segregation across all these partitions of the labor market. The contribution of explicit gender designations for jobs to gender segregation is analogous to the effects of a practice known as *red-lining* in urban residential segregation in the United States before the civil rights era. Red-lining refers to the explicit and often public designation of certain neighborhoods as, say, black, mixed, or white, analogous to employers' explicit designation of jobs as *F*, *N* or *M* in our data.[7] In the urban context, these labels presumably allocated home seekers to neighborhoods both by directing where homeseekers search for housing ('compliance') and via landlords' and homesellers' refusals to transact with 'race-inappropriate' persons who offer to purchase or rent a home. In addition to measuring the total contribution of job profiling to gender segregation, this section also decomposes that contribution into its compliance and enforcement components.

To accomplish these goals, we use Duncan and Duncan's (1955) segregation index, applied to the set of successful applicants (i.e. callbacks) in a unit, *i*, which can be a job ad, a firm, or an occupation. The index, *S*, can be calculated from the female shares, δ_i, in those units as:

$$S = \frac{\sum_i \gamma_i |\delta_i - \Delta|}{2\Delta(1 - \Delta)} \qquad (3)$$

where δ_i is the female share in unit *i*, Δ is the female share in the population, and γ_i is unit *i*'s share of the callback population. Thus, *S* is the population-weighted mean absolute deviation of the female share from its global mean, divided by its maximum attainable value, $2\Delta(1\text{-}\Delta)$.[8] Like our gender-matching index *G*, Duncan and Duncan's *S* index varies between 0 and 1. It is widely used in studies of residential segregation (Cutler et al. 1999, Logan et al. 2004). Duncan and Duncan's *S* also has a well-known, natural interpretation: In our context, it gives the share of

[7] For an example of officially sanctioned residential redlining, see Section 980 (3) of the Federal Housing Association's 1938 Underwriting Manual, which recommends "Prohibition of the occupancy of properties except by the race for which they are intended" in restrictive housing covenants. (Federal Housing Association, 1938).

[8] Equivalently, *S* can be calculated via the better known formula, $S = \frac{1}{2}\sum_j \left[\frac{\phi_i}{\Phi} - \frac{\mu_i}{\mathrm{M}}\right]$, where ϕ_i is the share of callbacks in unit *i* that go to women, $\mu_i = 1 - \phi_i$ is the share of callbacks in unit *i* that go to men, and Φ and $\mathrm{M} = 1 - \Phi$ are their population equivalents.

men (or women) who would have to be reassigned to a different unit (job, firm, occupation, etc.) in order for men and women to be distributed identically across units.[9]

To use Duncan and Duncan's index in our context, however, we need to address an issue that doesn't usually arise in the residential segregation context: the effect of small unit sizes. This effect is most important when we wish to measure and decompose segregation across individual job ads, since the average size of the callback pool to an ad in our data is 5.4 workers. Thus, purely random variation in where workers send their resumes and in which resumes are picked from the application pool could generate a considerable amount of *de facto* segregation.[10]

To quantify the importance of random variation in the gender mix of callback pools, we take as given the total number of applications and callbacks at every job ad. We then simulate the amount of segregation we would expect if the gender mix of applications to each ad, and of callbacks to each ad was the result of a random draw from binomial distribution with parameters given by the population mean levels of α and θ. The idea is to hold fixed the total number of applications men and women make, the number of applications arriving at each job, and the total number of 'interview slots' (callbacks) available for each job. With these 'structural' features of the labor market fixed, we then assume that workers direct their applications randomly and that firms select candidates randomly. How much gender segregation would we expect to see?

In more detail, recall that the overall mean of α, $\bar{\alpha}$ = .543 and consider an ad that received 80 applications and issued 5 callbacks. We first simulate the number of female and male applications to that ad (a^f and a^m) as a random draw of 80 applications from a pool with population parameter .543, i.e. $a^f \sim B(n,p)$ = B(80, .543), $a^m = 80 - a^f$, and B indicates the binomial distribution. Next, taking this randomly-generated application pool as given (say, 51 women and 29 men), we simulate the number of male and female callbacks (c^f and c^m) as a random draw of 5 callbacks from a pool with population parameter given by:

$$p^c = \frac{\bar{\theta}a^f}{\bar{\theta}a^f + a^m} \qquad (4)$$

where $\bar{\theta}$ = .863 is the overall mean of women's relative callback risk. Thus, $c^f \sim B(n,p)$ = B(5, p^c); $c^m = 5 - c^f$. Doing this for every job, then calculating the realized segregation index, S, completes a single iteration.

---

[9] This property is independent of which group is being re-allocated and of the relative size of the two groups (Zoloth 1976). Notably, however, the counterfactual reallocation of residents underlying this interpretation does not preserve the total populations of the units.

[10] To see the point, note that if each firm calls back only one worker, segregation will always be complete: every firm's callback pool will be entirely male or entirely female.

Figure 1 plots the distribution of realized *S* values from 1000 iterations in this baseline scenario where there is no systematic variation across jobs in either application or callback behavior. It shows a surprisingly concentrated distribution with a mean of .316 and all values falling between .30 and .34. Thus, while random matching can generate a high level of measured segregation, the amount of segregation it generates is tightly constrained by the distribution of applicant pool sizes and callback pool sizes and the overall share of men and women in the population.

To remove the effects of this randomness, we define a *noise-adjusted* segregation measure, $\tilde{S}$, as:

$$\tilde{S} = \frac{S - S_0}{1 - S_0} \qquad (5)$$

where *S* is the unadjusted segregation index from equation (3) and $S_0 = .316$ is the mean level of segregation expected from noise in matching. Since *S* = .731, the noise-adjusted index of gender segregation across jobs in our data equals $\tilde{S} = \frac{.731 - .316}{1 - .316} = .606$. Interestingly, this level essentially coincides with Cutler et al.'s (1999) threshold of 0.6 for defining a U.S city as having a residential ghetto.

Having developed a noise-adjusted measure of gender segregation across jobs, we next ask how much of this segregation is associated with gender profiling, and how much of that 'label-linked' segregation, in turn, is associated with compliance versus enforcement. To that end, we use different assumptions on α and θ to generate five *counterfactual* $\tilde{S}$ indices. Of these, *counterfactual A* measures the total contribution of the three job *types* (the equivalent of 'redlining' in the residential segregation context) to gender segregation across jobs. Here, instead of a common α and θ for all ads, we simulate *S* allowing both α and θ to take three levels, one for each job type. *Counterfactuals B and C* parse counterfactual A into portions related to active selection by employers versus self-selection by workers, by letting only one of α and θ vary across job types. Finally, *counterfactuals D and E* ignore the job labels (*F*, *N*, and *M*) and divide the total amount of gender segregation across jobs into a worker self-selection component and an active employer selection component. In the former case, we simulate *S* using the *actual* numbers of applications received by each ad, while imposing the same relative risk, θ, for all jobs. Counterfactual *E* is the mirror image of this case.

The mean levels (across 1000 iterations) of noise-adjusted segregation from all the above simulations are displayed in Table 5.[11] According to counterfactual A, if the parameters α and θ differ *only* across the three job categories (*F*, *N* and *M*), mean noise-adjusted segregation,

[11] The distributions of *S* values across these counterfactual simulations are also highly concentrated, similar to the baseline, 'noise-only' simulation.

$\tilde{S}$, would equal .361, which almost 59.5 percent of all the gender segregation across jobs. Thus, about 60 percent of the total gender-segregation in the populations of successful job applicants across individual job ads on this job board is associated with the explicit gender labels employers attach to ads. The remaining gender segregation (much of it likely within the *N* jobs) is not guided by employers' explicit requests, and is presumably similar to the type of segregation that prevails in countries that do not practice explicit gender profiling in jobs. Consistent with our findings for gender-matching, counterfactuals B and C show that essentially all of the 'label-linked' cross-job segregation is due to self-sorting: allowing only α to differ across the three categories leads to a level of noise-adjusted segregation that is 57.6 percent of the actual level, while allowing only θ to differ generates only 8.0 percent of actual noise-adjusted segregation. Thus, in the analogy to residential segregation, homeseekers' (jobseekers) compliance with the designations of three neighborhood (job) types accounts for 57.6 percent of the census-block level (job-level) segregation in the city (labor market).

Finally, counterfactuals D and E abstract completely from the gender labels attached to job ads and simply ask what share of noise-adjusted sex segregation in the successful applicant pools across individual job ads is associated with men's and women's differential application patterns, versus their differential success rates conditional on applying. Together these two counterfactuals show that *self-sorting* (both directed and undirected) *accounts for 96 percent of all the systematic gender segregation* across jobs in our data.[12]

The preceding methods for computing actual and counterfactual noise-adjusted segregation indices across jobs can also be applied to segregation across other labor market units, including occupations, firms, and occupation*firm cells. The results of these calculations are summarized in Table 6.[13] At $\tilde{S}$= .558, Table 6 shows that gender segregation is almost as high across firm*occupation cells as across individual job ads, and that explicit gender profiling accounts for just under 60 percent of that segregation. Segregation across firms and occupations is lower, though it is interesting to note that $\tilde{S}$ is slightly higher across the 35 occupation categories on the XMRC website than across the much larger number of firms in our sample.[14] Column 3 of Table 6 shows that explicit gender profiling accounts for 60 percent of

[12] Enforcement alone –without any self-selection—can account for as much as 18 percent. The enforcement and compliance shares now add up to more than 100 percent because (in contrast to Table 4) these counterfactuals add enforcement and compliance, in turn, into a baseline scenario where neither is present, rather than subtracting them from a scenario where both are present. Analogous to adding regressors to an equation sequentially, the share explained by the first factor considered is larger in the absence than in the presence of a control for the second.

[13] As one might expect, the impact of noise-adjustment on the estimated level of segregation diminishes as the unit size increases (from jobs through occupations) in Table 6. It is minimal in the case of occupations, of which there are only 35 categories.

[14] This is consistent with a long literature documenting the importance of occupational sex segregation, and with DKS's finding that a large share of the variance in gender profiling is *within* firms.

the gender segregation across firms, and for 50 percent of the gender segregation across occupations.

A final perspective on the contribution of explicit gender labels to gender segregation examines the amount of gender segregation that is present within the 60 percent of our sample of job ads are not explicitly gendered: If explicit labels are epiphenomena that do not affect the allocation of workers to firms, we might expect to see just as much segregation within the nongendered ads, as men and women choose to apply to, and succeed in receiving, different jobs, firms and occupations even when those jobs are not publicly labeled as 'male' or 'female'. We perform these calculations in Appendix Table A6, and find much less gender segregation within sample of nongendered ads than in our sample overall. Compared to .607 overall, noise-adjusted gender segregation within the sample of nongendered jobs is only .196. Occupational segregation by gender is .422 overall, but only .155 in nongendered jobs. These figures suggest, but do not prove, that gendered job ads have real effects on the allocation of labor in China.

Summarizing our descriptive analysis of sex segregation in callback pools, we find that jobs in China are highly gender-segregated, with a noise-adjusted Duncan and Duncan segregation index of .607. In other words, on average 60.7 percent of either men or women would have to change jobs to equate the gender ratio across all jobs. Explicit gender profiling in turn accounts for 60 percent of that job-level segregation. Gender profiling also accounts for 60 percent of the gender segregation across firms, and for 50 percent of the gender segregation across occupations. Finally, the vast majority of the effect of gender profiling on all the outcomes studied in this section works through workers' compliance with firms' advertised requests when deciding where to send their resumes, rather than through active denials of callbacks to gender-inappropriate applicants.

4. Isolating the Causal Effects of Gender Profiling

While the previous Section provides useful descriptive information about the relationship between gender profiling of job ads and the recruitment process, simple descriptive statistics like those in Table 3 do not necessarily reflect the causal effects of job profiling on either workers' application behavior or employers' selection behavior. In this Section, we first define what we mean by the causal effects of profiling, then attempt to estimate these effects on application and selection decisions in turn.

Starting with application behavior, our goal is to describe the results of the following thought experiment. Imagine that the observed patterns of job profiling and application decisions in our data constitute a labor market equilibrium in the sense that employers' advertising and selection decisions are optimal given workers' application behavior, and vice versa. In this equilibrium, row 2 of Table 3 indicates that *F*, *N* and *M* job ads attract applicant pools that are 92.7, 45.0 and 8.0 percent female respectively. Now imagine that we were to exogenously switch the explicit gender label attached to *one* of these *N* jobs to *F* or to *M*, keeping everything else unchanged. What will happen to the share of applicants to that job that are female? If this share does not change, then the large differences in the gender mix of these three job types in Table 3 is not causal, in the sense that the gender labels do not *direct* workers' application decisions. Instead, the labels may simply be standing in for other features of the job (such as the occupation) that tend to attract applicants of different genders.

Accordingly, our econometric attempts to isolate a causal effect of gender labels on application behavior will focus on controlling as tightly as possible for other characteristics of jobs (or job ads) that might also explain why different ads attract different mixes of men and women. In practice, the tightest controls available to us are fixed effects for firm*occupation cells. In effect we will therefore be comparing two observationally identical ads issued by the same firm for the same occupation (for example sales jobs at Dell computer), one of which specifically requests, say, a woman and the other which does not state a gender preference. Notably, these estimated causal effects of job profiling on workers' application decisions are *partial equilibrium* effects in the sense that they isolate the effect of changing the label assigned to one job holding all other jobs' labels fixed, and holding prospective applicants' *expectations* about their relative callback chances (θ) in the three different job types (*F*, *N* and *M*) fixed as well. Both of these quantities are likely to change when a market-wide policy is changed, such as a penalty or ban on gender profiling. We discuss possible general-equilibrium effects of such policy changes in Section 5.

Turning to employers' selection behavior, we again imagine a labor market initially in a search equilibrium, but now we consider the effects of exogenously re-directing a single worker's *application* from a non-gendered (*N*) job to a 'gender-inappropriate' job, i.e. to a job whose explicit label does not match the worker's gender. If little or nothing happens to the worker's chances of receiving a callback, then employers' advertised gender preferences are 'soft' preferences, in the sense that job labels are only suggestive and 'gender-mismatched' applicants are evaluated fully and fairly along with other applications that arrive. If instead there is a large gender mismatch penalty, gendered job labels are hard requirements.

Again, while the substantial differences in men's and women's callback rates across job types in rows 3 and 4 of Table 3 suggest that advertised gender preferences are fairly 'hard', those apparent mismatch penalties could either be under-or over-estimates of the causal effects described above. Consider for example, women's apparent 8.8 – 4.1 = 4.7 percentage point callback penalty when applying to *M* versus *N* jobs from row 3. If the sample of women who apply to male jobs is negatively selected relative to women who apply to nongendered jobs, this 4.7 percentage point penalty would underestimate the true effect of the *M* label on the callback chances of a woman of fixed ability. On the other hand, if women who apply to men's jobs are positively selected, the true effect of the 'male' label on a man of fixed ability is larger than the unadjusted gap in Table 3. Thus, to isolate the causal effect of job labels on workers' callback rates we try to control as tightly as possible for other aspects of worker quality that might affect callback rates. In practice we will address this problem by using worker fixed effects—i.e. we will compare the callback rates of the same worker who sends her resume to two observationally-identical jobs that differ only by their explicit gender label. In this sense, these are first estimates of gender discrimination that control for worker fixed effects using actual, rather than fictitious resumes.

4.1 Do "desired gender" labels have causal effects on workers' application (self-selection) decisions?

According to row 2 of Table 3, the unadjusted 'effect' of labeling a job as female (*F*) on the female share of applicants that arrive (relative to no label) is .927-.450 = .477. Similarly, the raw 'effect' of labeling a job as male on α is .080 - .450 = -.370. As discussed above, we now examine how these two gaps change when we add controls for job characteristics in order to isolate a more causal effect of gender profiling. If the unadjusted gaps attenuate, then the labels *F*, *N*, and *M* are simply 'standing in' for the types of jobs men and women would be applying to even in the absence of explicit gender profiling. Attenuation would suggest that gendered job ads do not actually *direct* men's and women's applications.

To this end, we run regressions in our sample of 3,489 ads where the dependent variable is the share of applications that are female (α).[15] The regressors of interest are the labels attached to the ad (*F*, *N* or *M*). In more detail, we estimate:

$$\alpha_j = a + bF_j + cM_j + dX_j + e_j \qquad (6)$$

where *j* indexes jobs (ads), *F* (*M*) is a dummy for whether the job requests women (men) and *N* is the omitted job type. Since the precision with which α_j is measured increases with the number of applications received, all regressions are weighted by that number. In column 1 of Table 7, we include no controls (*Xj*) and replicate the raw gaps in Table 3. Column 2 adds controls for the following job characteristics: requested education, experience, and age; the advertised wage; a dummy for whether a new graduate is requested; the number of positions advertised; plus dummies for missing education, age, wage and number of positions. Columns 3-5, in turn, add occupation fixed effects, occupation and firm fixed effects, and occupation*firm fixed effects. In the latter, most saturated specification, we are effectively

[15] We replicated these regressions on the universe of ads posted in our sampling period, which includes ads for which we have no callback information. The results, shown in Appendix Table A2, are very similar.

comparing the gender mix of applications to two observationally identical ads issued by the same employer for the same occupation, with different 'preferred gender' labels attached.

Table 7 shows that the unadjusted effects of both the *M* and *F* job labels do attenuate somewhat (from -.37 to -.24 for *M* jobs and .48 to .35 for *F* jobs) when controls for job characteristics and occupation fixed effects are added in columns 2 and 3. Thus, to some extent, the correlation between jobs' gender labels and the gender mix of applicants reflects the fact that men and women tend to apply to different types of jobs, regardless of whether those jobs are explicitly targeted at their gender. This is as we expect since, for example, men are probably more likely to apply to truck-driving jobs and women to social work jobs regardless of which gender is explicitly requested. Still, the estimated effects of the gender labels remain very large and highly statistically significant.

Interestingly, however, adding controls for occupation and firm fixed effects in column 4 and for their interaction in column 5 has no additional effect on the *M* coefficient and actually strengthens the *F* coefficient, leaving it essentially identical to the unadjusted effect in column 1. Thus, once we have controlled for 35 occupational categories in column (3), jobs explicitly labeled as 'female' are disproportionately located in firms and in firm*occupation cells that tend to attract *male* applicants. This suggests that, if anything, employers are using the "female" job label to communicate which positions in typically male firms and occupations are open to women.[16]

Summing up, our regression analysis of the effects of a job's gender label on the gender mix of applicants it attracts shows that even within occupation*firm cells, jobs that explicitly request male or female applicants strongly shift the gender mix of applicants towards the requested gender. Indeed, the estimated effect of a 'female' job label (relative to no gender label) in the presence of fixed effects for firm*occupation interactions is essentially identical to the very large unadjusted effect of 47 percentage points, indicating that in the end *none* of the 'raw' gap is an artifact of the tendency for certain types of jobs to be perceived as 'female' regardless of whether women are explicitly invited to apply. The adjusted effect of a 'male' job label on the other hand, at 24 percentage points is about 65 percent as large as the unadjusted effect, suggesting that about 35 percent of the unadjusted effect is a consequence of the fact that certain occupations consistently attract male applicants, even when the jobs are not explicitly labeled as male. Even here, however, the adjusted effect is economically very large and highly statistically significant. We conclude that the job labels *M* and *F* convey real information to prospective applicants that is not otherwise available. In other words, explicit

[16] This increase in the *FJ* coefficient is also present in the full-sample results shown in Table A1. The phenomenon should only be seen as suggestive, however, since even in the full sample the column 5 coefficient is not statistically different from the column 3 coefficient.

gender labels appear to *redirect workers' applications* towards jobs where their gender is most 'wanted'.

Finally, in interpreting these estimates it is important to recall that they represent the partial equilibrium effects of assigning a gender label to a job ad, holding all other aspects of the ad (and of other ads in the labor market) fixed. In addition these estimated effects are consistent with at least two mechanisms. On the one hand, job labels may simply communicate to applicants that they might enjoy working in the job (possibly due to the gender mix of their co-workers). On the other, job labels may also communicate the worker's relative chances of actually getting a callback, θ. If so, as noted Table 7's regression estimates should be interpreted as reflecting a fixed set of *applicant beliefs* (i.e. those prevailing in current market equilibrium) regarding how θ varies across the three types of jobs (*F*, *N* and *M*). Thus Table 7's estimates do not apply to policy changes that redirect large numbers of applications, or that otherwise affect how firms choose callbacks from their applicant pools.

4.2 Does applying to a 'gender-Inappropriate' job have a causal effect on a worker's callback probability?

Turning back to Table 3, our raw data suggests a substantial but by no means infinite callback penalty from applying to 'gender-mismatched' jobs. Specifically, compared to nongendered jobs, women's callback rate is .088 - .041 = .047 percentage points lower in *M* than *N* jobs. Men's is .092 - .054 = .038 percentage points lower in *F* than *N* jobs. But do these penalties exist because employers treat the same applicant differently when he or she applies to a gender-mismatched job, or because the relatively few workers who apply to gender-mismatched jobs are negatively selected? This could be the case, for example, if gender-inappropriate applications are mistakes made by careless or inattentive applicants.

To answer this question, we estimate callback regressions that control for all available measures of job requirements and applicant characteristics. If (say) women's callback penalty from applying to an *M* job attenuates when these controls are added, part of the unadjusted 'effect' of the *M* job label on women's success is a spurious consequence of negative selection into gender-inappropriate applications. If, on the other hand, the estimated penalty increases in magnitude, applicants to gender-inappropriate jobs would appear to be positively selected relative to applicants to nongendered jobs.

In more detail, we estimate the following linear probability model:

$$Callback_i = \alpha + \beta_1 FtoF_i + \beta_2 FtoM_i + \beta_3 MtoF_i + \beta_4 MtoM_i + \delta Mworker_i + \varphi X_i + \varepsilon_i \quad (7)$$

where *i* indexes *applications*. Of the six possible application types, women applying to nongendered jobs (*FtoN*) is the omitted type. In this specification, β_1 and β_2 give the effect on women of applying to *M* and *F* jobs (relative to nongendered jobs), while β_3 and β_4 give the effect on men of applying to *M* and *F* jobs (again, relative to nongendered jobs). The parameter δ gives the callback gap between men and women applying to nongendered jobs. As noted, the goal of these estimates is to isolate the effects of the following thought experiment: Take the same worker and send her application to two jobs that are observationally identical, except for the gender label (*F*, *N* or *M*) attached to the job. What will be the difference in callback rates? Our main focus will be on the *gender mismatch penalties* associated with applying to a job that is targeted at the 'other' gender, β_2 and β_3.

Column 1 of Table 8 estimates equation (7) without controls, reproducing the raw gaps reported in rows 3 and 4 of Table 3. Column 2 adds controls for the job's requested level of education, experience and age; the advertised wage; and an indicator for whether a new graduate is requested. Also included are a set of indicators of the match between the applicant's characteristics and those requirements, including indicators for whether the applicant's education, age and experience are below or above the requested level, the match between the advertised wage and the applicant's current or previous wage, and the match between requested and actual new-graduate status.[17] Column 3 adds controls for the following worker (CV) characteristics: whether he/she attended a technical school; the applicant's *zhicheng* rank; whether an English CV is available; the number of schools attended, experience spells and certifications reported.[18] Indicators for applicant height, myopia and marital status are also included, all interacted with the applicant's gender.[19]

Column 4 adds fixed effects for 35 occupation categories characterizing the advertised job. Column 5 adds two indicators of the amount of competition for the job: the number of positions advertised and the number of persons who applied to the ad, and column 6 adds firm fixed effects. Finally, our most saturated specification in column 7 adds applicant (resume) fixed effects. In this case, the effects of fixed applicant characteristics, including gender, are not identified. Interactions between applicant gender and job type, which are our main coefficients of interest, however, remain identified. In effect, the column 7 estimates compare the outcomes of *the same worker* who has applied to observationally identical jobs that differ only according to the gender label (*F*, *N* or *M*) attached to the job, while allowing for this effect to differ according to the applicant's gender.

[17] Indicators for missing (requested) age and wage information are also included.
[18] *Zhicheng* is a nationally-recognized worker certification system that assigns an official rank (from one through six) to workers in almost every occupation. Ranks are based on education, experience and in some cases nationwide or province-wide exams.
[19] The 'detailed CV controls' introduced in column 3 are not requested in job ads very often, so it is not practical to construct variables summarizing their match with the job's requirements.

Before discussing our main coefficients of interest (the mismatch penalties), it is worth noting that whenever they are statistically significant, observable indicators of the match between worker qualifications and job requirements are of the expected signs in Table 8: workers who have less education or experience than requested, or are older than requested are less likely to be called back. Also of some interest, workers whose current wage is above the job's posted wage are less likely to be called back.

Turning to the mismatch penalties, men's unadjusted mismatch penalty does attenuate somewhat across the columns of Table 8 (mostly when basic worker and match controls are added in column 2 and when occupation fixed effects are added in column 4). Specifically, men's unadjusted mismatch penalty of 3.8 percentage points (in column 1) falls to 2.6 percentage points in column 7, indicating that about (.038-.260)/0.038 = 32 percent of men's unadjusted mismatch penalty is a consequence of negative selection of the men who apply to 'female' jobs. Applied to men's mean callback rate of 9.2 percent in *N* jobs, men's 2.6 percentage point regression-adjusted mismatch penalty implies that the men who apply to nongendered jobs would have a 6.6 percent callback rate in female jobs, which is substantially lower than in nongendered jobs, but also economically and statistically much greater than zero.

In contrast, women's callback penalty when applying to 'male' jobs remains essentially unchanged across the columns of Table 8: it is 4.6 percentage points in the raw data and 4.3 percentage points in the presence of worker fixed effects. We conclude that women who apply to 'male' jobs are, on balance, neither positively nor negatively selected relative to women who apply to 'gender-appropriate' jobs. Thus, especially for women, the unadjusted callback penalties used in all of Section 3's decompositions are very similar to our best estimates of the causal effects of applying to a gender-inappropriate job. Applying women's 4.3 percentage point regression-adjusted mismatch penalty to their mean callback rate of 8.8 percent in *N* jobs implies that the women who apply to nongendered jobs would have a 4.5 percent callback rate in male jobs, which, like men's callback rate in women's jobs, is substantially less than in nongendered jobs, but well above zero. Interestingly, the above calculations also suggest that women fare considerably worse when applying to men's jobs than vice versa. One possible interpretation is that –at least in the jobs where gender-mismatched applications tend to occur-- employers have stronger preferences for keeping women out of 'men's' jobs than vice versa.[20]

[20] Such preferences could, of course, derive from the anticipated negative reactions of incumbent male workers to female applicants.

5. Discussion

We believe that this is the first paper to study the (proximate) *consequences* for workers and firms of a practice that is very common in the world's labor markets: job advertisements that specify the worker's desired sex. Using internal callback information from a Chinese Internet job board, we find that workers comply with firms' advertised gender requests when deciding where to send their job applications, and that firms penalize gender-inappropriate applicants. Overall, 60 (60) [50] percent of the gender segregation among successful (called-back) applicants across all jobs (firms) [occupations] is associated with explicit gender-profiling in job ads. In an accounting sense, virtually all of this effect is associated with applicants' *compliance* with employers' advertised gender preferences when deciding where to apply, with active selection by employers among applications contributing very little. Intuitively, since so few workers apply to gender-mismatched jobs, total gender segregation would change very little even if employers ignored gender in all their callback decisions.

While the above decompositions were conducted using simple population averages, regression estimates with firm*occupation fixed effects strongly suggest that gendered job ads have causal effects on application behavior that are similar in size for female jobs, and about a third smaller for male jobs than those suggested by the raw data. Regression estimates with worker fixed effects strongly suggest that something similar is true for employers' callback decisions: women's callback penalty in 'male' jobs is very similar in size to their estimated penalty in the raw data, and men's is about a third smaller. Thus, with minor changes, the decompositions we conducted in Tables 4-6 using the raw data yield very similar results using regression-adjusted estimates.[21]

While we believe that our analysis has increased our understanding of the role of gendered job ads in the recruitment process, a number of important questions remain unanswered. One such question concerns the size of our estimated 'enforcement' effects: while being of the 'wrong' gender for a job reduces both men's and women's callback rates substantially, it is perhaps surprising that these gender-mismatch penalties aren't larger: even according to Table 8's regression estimates with worker fixed effects, men (women) who apply to jobs requesting female (male) applicants still have a substantial chance of getting a callback. Indeed, given the substantial recall rates of gender-mismatched applications, it seems a little surprising that workers avoid applying to gender-mismatched jobs as strongly as they do.

Since our mismatch-penalty estimates survive worker fixed effects, we can be sure that mismatched applicants' perhaps unexpectedly-high callback rates are not a result of

[21] Appendix Tables A4 and A5 replicate Tables 4 and 5 using αs that hold ad characteristics fixed and θs that hold both ad and worker characteristics fixed.

nonrandom selection of workers who choose to apply for gender-mismatched jobs. Indeed, if gender-mismatched applicants were positively selected –which would raise their callback rates relative to other applicants-- the estimated mismatch penalty would *increase* when we add worker fixed effects, and that is not the case. The substantial callback rates of gender-mismatched applicants are also almost certainly not due to random miscoding of gender on workers' profiles. If worker gender was sometimes miscoded, we would expect to see a 'true' man who was mistakenly coded as a woman repeatedly apply for 'male' jobs. In Appendix Table A3, we exclude the very small number of workers who repeatedly apply to mismatched jobs from the sample, with very little change in the results.[22]

A third possible explanation is that much of the information conveyed by jobs' gender labels is not about whether a worker is likely to get a callback but about how much the worker is likely to enjoy the job. If most but not all workers prefer same-gender work environments, then those workers who do not mind a mixed-sex environment will apply to gender-mismatched jobs and not be heavily penalized in the callback process. Finally, there may be considerable heterogeneity in the amount of enforcement *within* explicitly gendered jobs. In this scenario, some types of gendered jobs never call back gender-mismatched applicants, while others (perhaps a minority) do not consider gender very strongly in their callback decisions. If applicants know which jobs these are, that could generate a substantial average callback rate for gender-mismatched applicants. This possibility seems to warrant further exploration in this and other data sets.

A final outstanding question concerns the effects of prohibiting gender-based job profiling, as is done in many industrialized economies. Can our estimates shed any light on these effects? According to Table 7's application regressions, banning gender profiling in China could lead to substantial changes in *where* workers send their applications. This is because the gender labels that are attached to four out of every ten jobs in our sample appear to convey information beyond what can be inferred from occupation*firm fixed effects and other detailed features of the job ad. Unless employers can create other, legal signals that direct workers to the formerly explicitly gendered jobs, banning gendered job ads should generate many more applications to jobs that would formerly have been labeled as gender-inappropriate, and fewer gender-appropriate applications to jobs explicitly requesting a particular gender.

Thus, if employers' choice patterns among applications that arrive, i.e. the θs in our notation, remain unchanged, banning profiling will lead to a decline in matching efficiency, with

---

[22] Miscoding of the *requested* gender is not a concern since our data are the exact record of requested gender that workers observe on the job board when deciding where to apply. See the Appendix for additional discussion of mismeasurement issues.

possible consequences like longer search spells or lower accepted wages. Importantly, most of the induced application mismatches would be *within* occupations and firms (where the marginal information conveyed by gender labels is most relevant), and at lower skill levels (since that is where most of the gendered ads occur). Among other effects, less-skilled female applicants would no longer know which of the jobs in typically-male occupations and firms would actually welcome their applications. Prohibiting gender-profiling could even hurt women more than men since (according to Table 8) women appear do somewhat worse in gender-mismatched jobs than men do. These efficiency losses will be mitigated if firms' selection decisions (the θs) become less gendered after profiling is banned, for example if hiring discrimination is also effectively prohibited or if employers examine *and* are pleasantly surprised by the new pool of gender-mismatched applicants that arrive. Thus, it is possible that banning profiling could raise efficiency by exposing employers to new, better- qualified applicants than they expected.

Finally, it is worth pointing out that the above efficiency results are conditional on viewing employers' gender preferences (i.e. the θs) as morally and ethically legitimate. Given any set of θs, gendered job ads may indeed increase matching efficiency by directing workers of both genders away from jobs in which they are not wanted. That said, gender-based (and indeed race-based) job ads have historically been used to indulge employer, co-worker and/or customer preferences that are not motivated by productivity differences and in retrospect are no longer seen as morally legitimate. Thus, efficiently serving a given set of tastes may not be the only criterion on which to judge the appropriateness of this particular labor market practice.

References

Albrecht, J., P. Gautier and S. Vroman. 2006. “Equilibrium directed search with multiple applications”. *Review of Economic Studies* 73: 869—91

Bielby, W. T., and J. N. Baron, ‘‘A Woman’s Place Is with Other Women: Sex Segregation within Organizations,’’ in Sex Segregation in the Workplace: Trends, Explanations, Remedies, B. F. Reskin, ed. (Washington, DC: National Academies Press, 1984).

Blau, Francine D., Brummund, Peter and Albert Yung-Hsu Liu (2013), “Trends in Occupational Segregation by Gender 1970–2009: Adjusting for the Impact of Changes in the Occupational Coding System”, *Demography* 50, 471–492

Brenčič, Vera. 2010. Do Employers Respond to the Costs of Continued Search? *Oxford Bulletin of Economics and Statistics*, 72(2), 211-245.

Brenčič, Vera. 2012. Wage Posting: Evidence from Job Ads, *Canadian Journal of Economics*, 45(4), 1529-1559.

Brenčič, Vera and John B. Norris. 2009. Employers’ Online Search: An Empirical Analysis, *Industrial Relations*, 48(4), 684-709.

Brenčič, Vera and John B. Norris. 2012. Employers’ Online Recruitment and Screening Practices, *Economic Inquiry*, 50(1), 94-111.

Brenčič, Vera and John B. Norris. 2010. Do Employers Change Job Offers in Their Online Job Ads to Facilitate Search? Economics Letters, 108(1), 46-48.

Card, David, Ana Rute Cardoso, and Patrick Kline. 2016 "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women" *Quarterly Journal of Economics*, forthcoming.

Cutler, David M., Edward L. Glaeser and Jacob L. Vigdor. 1999. “The rise and decline of the American ghetto” *Journal of Political Economy* 107(3): 455-506

Delgado Helleseter, Miguel, Peter Kuhn and Kailing Shen) “Age and Gender Profiling in the Chinese and Mexican Labor Markets: Evidence from Four Job Boards”, unpublished paper, UCSB, 2016.

Duncan, Otis Dudley, and Beverly Duncan. “A Methodological Analysis of Segregation Indices” *American Sociological Review* 20 (April 1955): 210-217.

Federal Housing Administration. 1938. *Underwriting Manual: Underwriting and Valuation Procedure Under Title II of the National Housing Act With Revisions to February, 1938* (Washington, D.C.), Part II, Section 9, Rating of Location. Available at: http://wbhsi.net/~wendyplotkin/DeedsWeb/fha38.html (visited May 10, 2016).

Hirschbein, Brad and Kahn, Lisa "Is College the New High School? Evidence from Vacancy Postings" unpublished paper, W. E. Upjohn Institute, 2015.

Kuhn, Peter and Kailing Shen. "Gender Discrimination in Job Ads: Evidence from China" *Quarterly Journal of Economics* 128(1) (February 2013a)

Kuhn, Peter and Kailing Shen. "Do Chinese Employers Avoid Hiring Overqualified Workers? Evidence from an Internet Job Board" *Research in Labor Economics* 37 (2013b): 1-30.

Kuhn, Peter and Kailing Shen. "Do Employers Prefer Migrant Workers? Evidence from a Chinese Job Board" *IZA Journal of Labor Economics* 4:22 (2015) DOI: 10.1186/s40172-015-0038-0

Kuhn, Peter and Kailing Shen "Gender Discrimination in Job Ads: Evidence from China" *Quarterly Journal of Economics*, February 2013a.

Kuhn, Peter and Kailing Shen. "Do Chinese Employers Avoid Hiring Overqualified Workers? Evidence from an Internet Job Board" *Research in Labor Economics* 37 (2013b): 1-30

Kuhn, Peter. "Employer Investments in Employee Assessment" unpublished paper, UC Santa Barbara, March 2016.

Kudlyak, Marianna, Damba Lkhagvasuren and Roman Sysuyev. "Systematic Job Search: New Evidence from Individual Job Application Data"

Logan, John R., Brian J. Stults and Reynolds Farley. "Segregation Of Minorities In The Metropolis: Two Decades Of Change" *Demography*, 41(1), February 2004: 1–22

Makinen, Julie. "China's women begin to confront blatant workplace bias" February 28, 2014. Available at: http://www.latimes.com/world/asia/la-fg-china-women-work-20140228-story.html#page=1. Visited April 8, 2015.

Marinescu, Ioana and Roland Rathelot "Mismatch Unemployment and the Geography of Job Search" unpublished paper, University of Chicago 2015. [pdf]

Marinescu, Ioana and Ronald Wolthoff "Opening the Black Box of the Matching Function: the Power of Words", unpublished paper, University of Chicago 2015. [pdf]

Modestino, Alicia Sasser, Daniel Shoag and Joshua Balance . "Upskilling: Do Employers Demand Greater Skill When Workers Are Plentiful?" unpublished paper, Northeastern University.

Neumark, David Ian Burn and Patrick Button. "Is It Harder for Older Workers to Find Jobs? New and Improved Evidence from a Field Experiment" NBER Working Paper No. 21669, October 2015

Pager, Devah and David Pedulla. "Race, Self-Selection, and the Job Search Process". Unpublished manuscript, Princeton University, 01/22/12.

Zoloth, B.S. 1976. *Alternative Measures of School Segregation*. Madison: Institute for Research on Poverty, University of Wisconsin

Table 1: Descriptive Statistics: Ad Sample

	Ad Requests Women (*F* jobs)	Gender not specified (*N* jobs)	Ad Requests Men (*M* jobs)	All Ads
Education specified?	0.961	0.899	0.923	0.919
Education Requested (years), if specified	12.68	12.30	11.27	12.20
Tech School Requested?	0.292	0.148	0.192	0.191
Desired Age Range specified?	0.635	0.389	0.566	0.481
Desired Age, if Requested (midpoint of interval)	25.96	28.74	29.40	28.00
Experience Requested (years)	0.781	0.995	1.236	0.987
New Graduate Requested?	0.068	0.023	0.031	0.035
Wage Advertised?	0.637	0.557	0.555	0.576
Wage, if advertised (yuan/month, midpoint of interval)	1983	2644	2454	2434
Number of positions specified?	0.963	0.923	0.972	0.942
Number of positions, if specified	1.875	2.222	2.006	2.099
Number of applicants	79.08	62.22	46.42	63.38
Sample Size	840	2009	640	3489

Table 2: Descriptive Statistics: Application Sample

	Applications from Women	Applications from Men	All Applications
Education (years)	14.56	14.10	14.35
Completed Tech School?	0.156	0.166	0.160
Age (years)	23.93	25.56	24.68
Experience (years)	2.679	3.895	3.234
New Graduate?	0.209	0.153	0.183
Current wage listed?	0.688	0.703	0.695
Current wage, if listed (yuan/month)	2090	2461	2261
Married (if marital status listed)	0.140	0.215	0.175
Occupational Qualification (*Zhicheng*)[1]	1.064	1.348	1.194
Myopic	0.328	0.268	0.300
Height (cm)	160.6	171.5	165.6
English CV available?	0.145	0.104	0.126
Number of Schools listed	0.848	0.699	0.780
Number of Experience Spells	2.454	2.432	2.444
Number of Certifications	1.448	0.879	1.188
Sample Size	120,172	100,963	221,135

Notes:

1. *Zhicheng* is a nationally-recognized worker certification system that assigns an official rank (from one through six) to workers in almost every occupation. Ranks are based on education, experience and in some cases nationwide or province-wide exams .

Table 3: Application and callback patterns by job type

	Ad Requests Women (*F* jobs)	Gender not specified (*N* jobs)	Ad Requests Men (*M* jobs)	All Ads
1. Share of callbacks that are female (δ)	0.944	0.437	0.036	0.507
2. Share of applications that are female (α)	0.927	0.450	0.080	0.543
3. women's callback rate (f)	0.072	0.088	0.041	0.079
4. men's callback rate (m)	0.054	0.092	0.096	0.092
5. ratio of callback rates ($\theta = f/m$)	1.337	0.949	0.431	0.863
N (ads)	840	2,009	640	3,489
N (callbacks)	4,726	11,281	2,724	18,731
N (applications)	66,425	125,003	29,707	221,135

Table 4: Actual and counterfactual gender-matching rates

	Share of callbacks that are of the requested gender (*g*) (1)	Gender-matching index (*G*) (2)
Baseline: Actual values	0.951	0.902
Counterfactual 1-- no compliance: Equal female share in applications, α, across all jobs[1]	0.631	0.259
Counterfactual 2:-- no enforcement: Equal female callback advantage (θ) in all jobs[2]	0.921	0.842

Notes:

1. Applies the population female applicant share (α) (.543) to all three job types.
2. Applies the population female risk ratio (θ) (.863) to all three job types.

Table 5: Actual and Simulated Noise-Adjusted Segregation Indices across Jobs (Ads)

	Noise-Adjusted Segregation Index ($\tilde{S}$)	Share of noise-adjusted segregation explained ($\tilde{S}$ simulated/$\tilde{S}$ actual)
ACTUAL	**0.607**	**1.000**
SIMULATIONS:		
Effects of job categories (*F*, *N* and *M*) on segregation:		
A. Total effect of job categories: both α and θ vary across job categories	0.361	0.595
B. Effect of self-sorting across the three job categories: α varies across job categories, θ is the same in all ads	0.349	0.576
C. Effect of enforcement in the three job categories: θ varies across job categories, α is the same in all ads	0.048	0.080
Effects of applicant self-sorting and employer choice on segregation:		
D. Effect of self-sorting across all jobs: each job has its own α, all jobs have the same θ	0.580	0.957
E. Effect of employer choice within all jobs: each job has its own θ, all jobs have the same α	0.111	0.183

Table 6: Actual and Counterfactual Segregation across Job Titles, Occupations and Firms

Gender Segregation across:	Actual, noise-adjusted segregation ($\tilde{S}$) (1)	Segregation associated with job profiling (Counterfactual A) (2)	Share associated with job profiling (2/1) (3)
Jobs (from Table 5)	.607	.361	.595
Firm*Occupation cells	.558	.327	.585
Firms	.391	.235	.600
Occupations	.422	.210	.500

Table 7: Effects of Employers' Gender Requests on the share of female applications received (α)

	(1)	(2)	(3)	(4)	(5)
Ad requests men (MJ)	-0.3692***	-0.3290***	-0.2368***	-0.2026***	-0.2405***
	(0.039)	(0.035)	(0.029)	(0.031)	(0.083)
Ad requests women (FJ)	0.4773***	0.4220***	0.3548***	0.4164***	0.4645***
	(0.043)	(0.046)	(0.043)	(0.038)	(0.120)
eduF1		0.0277	-0.0127	-0.0024	0.0901
		(0.025)	(0.024)	(0.033)	(0.077)
eduF2		-0.0736**	-0.0577**	-0.1132***	-0.0612
		(0.032)	(0.023)	(0.030)	(0.080)
eduF3		0.0554**	0.0301	-0.0011	0.0444
		(0.023)	(0.020)	(0.025)	(0.071)
eduF4		0.1211***	0.0586***	0.0319	0.0812
		(0.033)	(0.021)	(0.031)	(0.077)
eduF5		0.1027**	0.0340	0.0333	0.0807
		(0.048)	(0.034)	(0.034)	(0.062)
edumissF		-0.0110	-0.0360	-0.0571*	-0.0526
		(0.031)	(0.031)	(0.028)	(0.055)
Number of positions advertised		-1.2763**	-0.9107***	-0.0444	-0.0181
		(0.576)	(0.275)	(0.365)	(0.839)
Occupation Fixed Effects			Yes	Yes	
Firm Fixed Effects				Yes	
Occupation*Firm Fixed Effects					Yes
Observations (number of ads)	3,489	3,489	3,489	3,489	3,489
R-squared	0.570	0.620	0.739	0.853	0.960

Standard errors in parentheses, clustered by occupation. *** p<0.01, ** p<0.05, * p<0.1

Note: in addition to the covariates shown, columns 2-5 also control for the following job ad characteristics: requested experience level (quadratic), requested age level (quadratic in midpoint of range), -advertised wage (quadratic in midpoint of bin; 8 bins), dummy for whether new graduate requested, number of positions advertised, plus dummies for missing education, age, wage and number of positions. All regressions are weighted by the total number of applications received.

Table 8: Effects of Job Labels (*F*, *N* and *M*) on Callback Rates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Female Worker * Female Job	-0.0152	-0.0110	-0.0109	-0.0112	-0.0162*	-0.0076	-0.0121
	(0.009)	(0.010)	(0.010)	(0.009)	(0.009)	(0.008)	(0.010)
Female Worker * Male Job	-0.0462***	-0.0446***	-0.0441***	-0.0436***	-0.0437***	-0.0221**	-0.0432***
	(0.013)	(0.014)	(0.014)	(0.015)	(0.012)	(0.010)	(0.013)
Male Worker * Female Job	-0.0381***	-0.0323***	-0.0319***	-0.0268***	-0.0269***	-0.0226**	-0.0260***
	(0.010)	(0.010)	(0.010)	(0.010)	(0.009)	(0.009)	(0.010)
Male Worker * Male Job	0.0037	-0.0003	-0.0012	0.0007	0.0025	-0.0010	-0.0025
	(0.009)	(0.009)	(0.009)	(0.009)	(0.008)	(0.008)	(0.008)
Male Worker	0.0047	0.0012	-0.0024	-0.0048	-0.0123**	-0.0145***	
	(0.006)	(0.006)	(0.006)	(0.006)	(0.005)	(0.004)	
Education less than requested		-0.0056	-0.0117**	-0.0121**	-0.0116**	-0.0123***	-0.0050
		(0.005)	(0.006)	(0.005)	(0.005)	(0.003)	(0.009)
Education more than requested		-0.0048	0.0007	0.0008	0.0025	0.0006	0.0010
		(0.004)	(0.004)	(0.004)	(0.004)	(0.002)	(0.007)
Age less than requested		-0.0006	-0.0025	-0.0026	-0.0000	-0.0086**	0.0050
		(0.006)	(0.006)	(0.006)	(0.005)	(0.003)	(0.007)
Age more than requested		-0.0313***	-0.0288***	-0.0270***	-0.0172**	-0.0172**	-0.0175**
		(0.009)	(0.009)	(0.009)	(0.008)	(0.007)	(0.009)
Experience less than requested		-0.0062	-0.0067	-0.0076	-0.0108**	-0.0127***	-0.0075
		(0.006)	(0.006)	(0.005)	(0.005)	(0.003)	(0.006)
Experience more than requested		-0.0001	0.0004	-0.0001	0.0008	-0.0023	0.0048
		(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.005)
Wage below advertised		-0.0013	-0.0014	-0.0027	-0.0029	-0.0016	-0.0074
		(0.005)	(0.005)	(0.005)	(0.005)	(0.002)	(0.006)
Wage above advertised		0.0001	0.0004	-0.0004	-0.0017	-0.0061**	0.0004
		(0.004)	(0.004)	(0.004)	(0.005)	(0.003)	(0.007)
Detailed CV controls			Yes	Yes	Yes	Yes	Yes
Occupation Fixed Effects				Yes	Yes	Yes	
Job Competition Controls					Yes	Yes	
Firm Fixed Effects						Yes	
Worker Fixed Effects							Yes
Observations	221,135	221,135	221,135	221,135	221,135	221,135	221,135
R-squared	0.002	0.005	0.006	0.017	0.028	0.240	0.404

Standard errors in parentheses, clustered by ad. *** p<0.01, ** p<0.05, * p<0.1

Notes to Table 8:

In addition to the covariates shown, columns 2-7 include the following controls for *ad characteristics*: requested education (5 categories), experience (quadratic), age (quadratic), the advertised wage (quadratic in midpoint of bin; 8 bins) and a dummy for whether a new graduate is requested. Columns 2-7 also include a dummy for whether the applicant's new graduate status matches the requested status, plus indicators for missing age and wage information for either the ad or the worker

"Detailed CV controls" (used in columns 3-6) are an indicator for attending technical school; the applicant's *zhicheng* rank (6 categories); an English CV indicator; the number of schools attended, job experience spells and certifications reported; and the following characteristics interacted with gender: height, myopia, and marital status.
-marital status (interacted with applicant gender)

Occupation fixed effects control for the 35 categories used on the XMRC website.

Job competition controls are the number of positions advertised (plus a dummy for unspecified) and the number of persons who applied to the ad.

Appendix—for online publication

Figure A1: Simulated segregation indices with random allocation of applications to jobs, and random selection of callbacks from all applicant pools

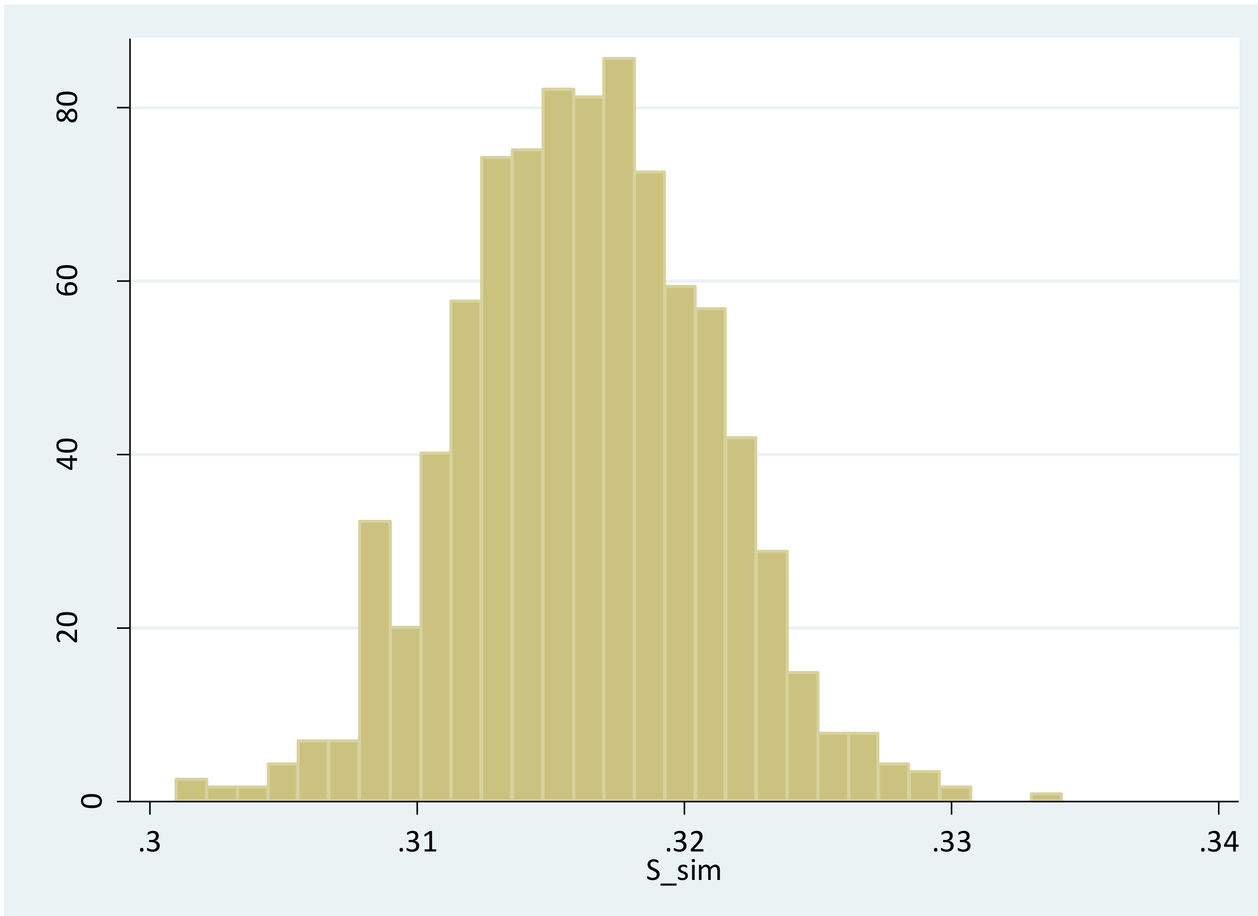

Table A1: Descriptive Statistics: Full Ad Sample

	Ad Requests Women (*F* jobs)	Gender not specified (*N* jobs)	Ad Requests Men (*M* jobs)	All Ads
Education specified?	0.945	0.886	0.931	0.905
Education Requested (years), if specified	12.84	12.73	11.69	12.56
Tech School Requested?	0.267	0.123	0.171	0.160
Desired Age Range specified?	0.573	0.319	0.529	0.407
Desired Age, if Requested (midpoint of interval)	26.41	29.53	30.31	28.85
Experience Requested (years)	0.839	1.150	1.344	1.124
New Graduate Requested?	0.035	0.017	0.019	0.021
Wage Advertised?	0.507	0.384	0.444	0.419
Wage, if advertised (yuan/month, midpoint of interval)	1999	2720	2508	2508
Number of positions specified?	0.961	0.935	0.965	0.945
Number of positions, if specified	1.580	1.808	1.687	1.742
Number of applicants	58.93	42.03	36.59	44.37
Sample Size	8,138	25,890	7,439	41,467

Table A2: Effects of Employers' Gender Requests on the share of Female applications received (α)—Full Ad Sample

	(1)	(2)	(3)	(4)	(5)
Ad requests men (MJ)	-0.3565***	-0.3242***	-0.2474***	-0.2289***	-0.2168***
	(0.042)	(0.035)	(0.027)	(0.023)	(0.039)
Ad requests women (FJ)	0.4939***	0.4502***	0.3704***	0.3858***	0.4368***
	(0.043)	(0.043)	(0.043)	(0.034)	(0.048)
eduF1		0.0251	0.0097	-0.0018	-0.0037
		(0.016)	(0.015)	(0.013)	(0.019)
eduF2		-0.0654***	-0.0527**	-0.0533***	-0.0490*
		(0.020)	(0.020)	(0.014)	(0.025)
eduF3		0.0670***	0.0481***	0.0460***	0.0510***
		(0.020)	(0.010)	(0.011)	(0.016)
eduF4		0.1161***	0.0649***	0.0636***	0.0564**
		(0.025)	(0.013)	(0.015)	(0.023)
eduF5		0.1198***	0.0494***	0.0579***	0.0395
		(0.028)	(0.016)	(0.017)	(0.027)
edumissF		0.0434**	0.0151	0.0013	-0.0030
		(0.020)	(0.012)	(0.012)	(0.015)
Number of positions advertised		-1.7570***	-0.9988***	-1.0455***	-1.0776***
		(0.500)	(0.119)	(0.137)	(0.271)
Occupation Fixed Effects			Yes	Yes	
Firm Fixed Effects				Yes	
Occupation*Firm Fixed Effects					Yes
Observations (number of ads)	41,467	41,467	41,467	41,467	41,467
R-squared	0.554	0.590	0.722	0.798	0.905

Standard errors in parentheses, clustered by occupation. *** p<0.01, ** p<0.05, * p<0.1

Notes: in addition to the covariates shown, columns 2-5 also control for the following job ad characteristics: requested experience level (quadratic), requested age level (quadratic in midpoint of range), -advertised wage (quadratic in midpoint of bin; 8 bins), dummy for whether new graduate requested, number of positions advertised, plus dummies for missing education, age, wage and number of positions. All regressions are weighted by the total number of applications received. s

Notes on gender misclassification

Miscoding of the *requested* gender is not a concern for our application analysis, since our data are the exact record of requested gender that workers observe on the job board when deciding where to apply. Miscoding of the requested gender could account for the relatively high success rates of gender-mismatched applicants if employers sometimes specify a gender requirement without intending to. If so, advertised gender requirements would be *de facto* rather soft. We view this as a possible interpretation of the relatively weak mismatch penalty in callbacks in our data.

Another possibility is that workers miscode their own gender when using the drop-down menu in the application process. The very high compliance rates we observe suggest that this is not a major concern. Nevertheless, we checked to see if miscoded applicant gender could account for the relatively weak enforcement in our data by re-running the main analysis on a restricted subsample for whom we are confident we have the right gender.[23]

To construct this sample, we first use the universe of applications, with no restrictions, to calculate the share of applications each CV in the sample sends to jobs which request the opposite gender. We then drop all the CVs in our sample for whom this share is 0.5 or higher. We also drop all CVs who submit fewer than 5 applications in the unrestricted data, because we may not have enough observations on them to reliably assess their application behavior. These restrictions only drop approximately 15,000 applications, leaving a sample size of 205,969.

We then re-run the application-level regressions from Table 8, and the results are very similar to those presented in the main analysis, which gives us confidence that the results are not being driven by misreported gender. They are reported in Table A3. Results for other cutoffs are not materially different.

[23] Note that miscoded applicant gender cannot explain weak enforcement if firms use resume-processing software to pre-screen resumes based on coded gender: such screens would eliminate both actual and false gender mismatches from consideration, generating a high level of *measured* enforcement. Miscoded applicant gender can only explain low compliance if employers can see that some apparently mismatched applicants are in fact of the requested gender (for example from the photo, name or other features of the resume.

Table A3: Effects of Job Labels (*F*, *N* and *M*) on Callback Rates for Gender Mismeasurement Robust Sub-Sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Female Worker * Female Job	-0.0142	-0.0098	-0.0097	-0.0104	-0.0153*	-0.0064	-0.0118
	(0.009)	(0.010)	(0.010)	(0.009)	(0.009)	(0.008)	(0.010)
Female Worker * Male Job	-0.0452***	-0.0438***	-0.0432***	-0.0429***	-0.0430***	-0.0220**	-0.0424***
	(0.013)	(0.014)	(0.014)	(0.015)	(0.012)	(0.010)	(0.013)
Male Worker * Female Job	-0.0390***	-0.0337***	-0.0333***	-0.0285***	-0.0283***	-0.0240***	-0.0264***
	(0.010)	(0.010)	(0.010)	(0.010)	(0.009)	(0.009)	(0.009)
Male Worker * Male Job	0.0028	-0.0007	-0.0014	0.0000	0.0017	-0.0015	-0.0021
	(0.009)	(0.009)	(0.009)	(0.009)	(0.008)	(0.007)	(0.007)
Male Worker	0.0044	0.0014	-0.0014	-0.0046	-0.0122**	-0.0135***	
	(0.006)	(0.006)	(0.006)	(0.006)	(0.005)	(0.004)	
Education less than requested		-0.0068	-0.0127**	-0.0130**	-0.0124**	-0.0128***	-0.0048
		(0.005)	(0.005)	(0.005)	(0.005)	(0.003)	(0.009)
Education more than requested		-0.0041	0.0010	0.0009	0.0027	0.0004	0.0006
		(0.004)	(0.004)	(0.004)	(0.004)	(0.002)	(0.007)
Age less than requested		-0.0008	-0.0026	-0.0028	-0.0003	-0.0089**	0.0053
		(0.006)	(0.006)	(0.006)	(0.006)	(0.004)	(0.007)
Age more than requested		-0.0301***	-0.0278***	-0.0265***	-0.0168**	-0.0171**	-0.0170**
		(0.009)	(0.009)	(0.009)	(0.008)	(0.007)	(0.008)
Experience less than requested		-0.0059	-0.0064	-0.0074	-0.0104**	-0.0124***	-0.0074
		(0.006)	(0.006)	(0.005)	(0.005)	(0.003)	(0.006)
Experience more than requested		-0.0001	0.0001	-0.0002	0.0007	-0.0023	0.0051
		(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.005)
Wage below advertised		-0.0027	-0.0028	-0.0039	-0.0041	-0.0024	-0.0070
		(0.005)	(0.005)	(0.005)	(0.005)	(0.002)	(0.006)
Wage above advertised		0.0001	0.0003	-0.0005	-0.0018	-0.0062**	0.0004
		(0.004)	(0.004)	(0.004)	(0.005)	(0.003)	(0.007)
Detailed CV controls			Yes	Yes	Yes	Yes	Yes
Occupation Fixed Effects				Yes	Yes	Yes	
Job Competition Controls					Yes	Yes	
Firm Fixed Effects						Yes	
Worker Fixed Effects							Yes
Observations	205,969	205,969	205,969	205,969	205,969	205,969	205,969
R-squared	0.001	0.005	0.005	0.015	0.026	0.239	0.353

Standard errors in parentheses, clustered by ad. *** p<0.01, ** p<0.05, * p<0.1

Table A4: Actual and counterfactual gender-matching rates

	Share of callbacks that are of the requested gender (*g*) (1)	Gender-matching index (*G*) (2)
Baseline: predictions for an average ad and applicant	.930	.862
Counterfactual 1-- no compliance: Equal female share in applications, α, across all jobs[1]	.613	.235
Counterfactual 2:-- no enforcement: Equal female callback advantage (θ) in all jobs[2]	.888	.778

The predicted α for the population of ads (.515) is calculated from the regression in column 5 of Table 7, with no controls for job type (*F*, *N* or *M*), at the mean of all observed characteristics. Predicted αs by job type (α^F=.912, α^N =.420 and α^M =.151) are based on the same regression *with* controls for job type. Thus (in contrast to Table 4), Counterfactual 1 in Table A4 calculates the effect on gender-matching of removing the gender labels from ads (thus setting α=.515 in all three job types), holding all ad characteristics fixed at their population mean.

The predicted θ for the population of ads (.863) is calculated from the predicted male and female callback rates (*f* and *m*) from the regression in column 3 of table 8 at the mean of all observed characteristics, with no controls for job type (*F*, *N* or *M*). Predicted θs by job type (θ^F=1.336, θ^N =.950 and θ^M =.432) are based on the same regression *with* controls for job type. Thus (in contrast to Table 4), Counterfactual 2 in Table A4 calculates the effect on gender-matching of eliminating the regression-adjusted mismatch penalties (thus setting θ =.863 in all job types) rather than eliminating the raw penalties.

Table A5: Actual and Simulated Noise-Adjusted Segregation Indices across Jobs (Ads), using regression-adjusted αs and θs.

	Noise-Adjusted Segregation Index ($\tilde{S}$)	Share of noise-adjusted segregation explained ($\tilde{S}$ simulated/$\tilde{S}$ actual)
ACTUAL	0.606	1.000
SIMULATIONS:		
Effects of job categories (*F*, *N* and *M*) on segregation:		
A. Total effect of job categories: both α and θ vary across job categories	0.365	0.602
B. Effect of self-sorting across the three job categories: α varies across job categories, θ is the same in all ads	0.345	0.569
C. Effect of enforcement in the three job categories: θ varies across job categories, α is the same in all ads	0.048	0.079

The predicted α for the population of ads (.515) is calculated from the regression in column 5 of Table 7, with no controls for job type (*F*, *N* or *M*), at the mean of all observed characteristics. Predicted αs by job type (α^F=.912, α^N =.420 and α^M =.151) are based on the same regression *with* controls for job type. Thus (in contrast to Table 5), Simulation B in Table A4 calculates the level of segregation attributable to compliance with the three job categories holding job characteristics fixed at the population mean.

The predicted θ for the population of ads (.863) is calculated from the predicted male and female callback rates (*f* and *m*) from the regression in column 3 of table 8 at the mean of all observed characteristics, with no controls for job type (*F*, *N* or *M*). Predicted θs by job type (θ^F=1.336, θ^N =.950 and θ^M =.432) are based on the same regression *with* controls for job type. Thus (in contrast to Table 5), Simulation C in Table A4 calculates the level of segregation attributable to enforcement of the three job categories for an application with fixed, mean characteristics.

Table A6: Actual and Counterfactual Segregation across Job Titles, Occupations and Firms—Nongendered Job Ads Only

	Actual, noise-adjusted segregation ($\tilde{S}$)	Segregation associated with self-sorting (Counterfactual *D*)	Segregation associated with employer choice (Counterfactual *E*)
Gender Segregation across:	(1)	(2)	(3)
Jobs:			
Segregation Index ($\tilde{S}$)	0.196	0.120	0.034
Share of Actual	1.000	0.612	0.172
Firm*Occupation cells:			
Segregation Index ($\tilde{S}$)	0.193	0.153	0.034
Share of Actual	1.000	0.791	0.174
Firms:			
Segregation Index ($\tilde{S}$)	0.151	0.138	0.037
Share of Actual	1.000	0.915	0.247
Occupations:			
Segregation Index ($\tilde{S}$)	0.155	0.179	0.008
Share of Actual	1.000	1.155	0.050