

NBER WORKING PAPER SERIES

WHAT HAPPENS WHEN EMPLOYERS CAN
NO LONGER DISCRIMINATE IN JOB ADS?

Peter J. Kuhn
Kailing Shen

Working Paper 29116
<http://www.nber.org/papers/w29116>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2021

We thank David Card, Kory Kroft, Dan Kreisman, Rafael Lalive, David Ong, and seminar participants at the University of Illinois at Urbana-Champaign, Universidade do Minho, University of Toronto, Peking University HSBC Business School, Georgia State University, IESR (Jinan University), Melbourne Institute of University of Melbourne, East Asian Institute of National University of Singapore, Renmin University, East China Normal University, Nanyang Technological University for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Peter J. Kuhn and Kailing Shen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

What Happens When Employers Can No Longer Discriminate in Job Ads?

Peter J. Kuhn and Kailing Shen

NBER Working Paper No. 29116

July 2021

JEL No. J16,J63,J71

ABSTRACT

When employers' explicit gender requests were unexpectedly removed from a Chinese job board overnight, pools of successful applicants became more integrated: women's (men's) share of call-backs to jobs that had requested men (women) rose by 63 (146) percent. The removal 'worked' in this sense because it generated a large increase in gender-mismatched applications, and because those applications were treated surprisingly well by employers. The removal had little or no effect on aggregate matching frictions. The job titles that were integrated however, were not the most gendered ones, and were disproportionately lower-wage jobs.

Peter J. Kuhn

Department of Economics

University of California, Santa Barbara

2127 North Hall

Santa Barbara, CA 93106

and IZA

and also NBER

pjkuhn@econ.ucsb.edu

Kailing Shen

Research School of Economics

ANU College of Business & Economics

HW Arndt Building (25a)

The Australian National University

Canberra ACT 0200

Australia

and IZA

kailing.shen@anu.edu.au

1. Background and Motivation

Explicit requests for male or female applicants were once a common feature of job ads in the United States, Canada and Europe. While many of these countries prohibited such ads during the last 50 years, explicitly gendered job ads are still used in many nations, which together account for a substantial portion of the world's labor force.¹ Yet despite many countries' continued acceptance of explicitly gendered job ads, and despite this history of policy interventions, little is currently known about the effects of allowing versus prohibiting explicit gender requests on the firms and workers in a labor market.²

One of policymakers' motivations for banning gendered job ads has been a widely stated desire to eliminate a practice that is, on the surface, overtly discriminatory. Other possible motivations include reducing gender segregation in employment, increasing women's access to better jobs, and improving labor market matching efficiency by eliminating 'artificial' barriers to the employment of otherwise qualified workers. Gendered ad bans could however fail to achieve these goals for at least six reasons. First, if gendered ads were used mostly for affirmative-action purposes before the ban, banning them could increase gender segregation. Second, employers who previously requested a particular gender could simply ignore any new 'gender-mismatched' applications caused by the ban, leading to a null effect on segregation. Third, such employers could avoid even *receiving* new applications from the 'unwanted' gender by putting other signals of their gender preferences (such as code words) into their job ads. Fourth, qualifications for many predominately male and female jobs could be so different that few additional qualified applicants become available when employers 'open up' their jobs to workers of all genders: at least in the short run, signaling an openness to male nurses and female electricians may not yield any new, qualified applicants. Fifth, if men's application decisions respond more strongly to the ban than women's do, women may be crowded out of jobs that previously invited them, without gaining much access to jobs that previously invited men. Finally, banning gendered ads could raise labor market frictions by making it harder for

¹ For U.S. examples of gender- and race-specific job ads before 1974, see Darity and Mason (1998) and Walsh et al. (1975).

Appendix 1 documents gendered ad bans in various countries since 1970. Appendix 1 in Kuhn, Shen and Zhang (2020) gives recent examples of gendered ads in large labor markets -- including India, Indonesia, Brazil, Nigeria, Mexico, and Russia.

² In previous work, we have studied how and when employers use gendered job ads when such ads are permitted (Kuhn and Shen 2013; Delgado Helleseeter, Kuhn and Shen 2020). In the same policy environment, Kuhn, Shen and Zhang (2020) show that employers' gender requests direct workers' application decisions, and predict a worker's chances of success conditional on applying to a job. But apart from Card, Colella and Lalive (2021) we are not aware of any other research that estimates the effects of prohibiting explicitly discriminatory job ads.

both men and women to locate jobs where their applications are likely to succeed, and harder for them to avoid jobs where they are not wanted.

Motivated by these questions, this paper studies what happened when employers' 'preferred gender' field was removed without advance notice from all the job ads on a Chinese job board on March 1, 2019. Using data on ads whose content was changed by the job board overnight, we estimate the causal effects of removing this single piece of information on workers' application decisions and on the gender composition of successful applicant pools. Importantly, we can calculate these effects separately for jobs which originally requested men and for jobs which originally requested women, which allows us to estimate whether the ban increased women's access to 'men's' jobs, and men's access to women's. We also estimate the ban's effects on three indicators of search frictions: the arrival rate and quality of applications received by firms, and workers' call-back chances conditional on applying. We also characterize the types of jobs that were opened up by the ban in terms of the gender mix of the incumbent workforce and the wages offered.

Our first main finding is that the ban integrated call-back pools in the jobs that formerly specified a preferred gender: it raised the share of women in the pool of successful applicants to jobs that previously requested men (*M* jobs) and raised the share of men in the pool of successful applicants to jobs that previously requested women (*F* jobs). Notably, these effects were immediate—they are visible in the first week after the ban—and they persist for at least six months afterwards. The ban had these effects because (a) it generated a large number of additional gender-mismatched applications, and (b) these applications were treated surprisingly well by employers (both before and after the ban).

Second, these effects were not symmetric: women's share of call-backs to *M* jobs rose by 3.0 percentage points, or 63 percent, while men's share of call-backs to *F* jobs rose by 9.9 percentage points, or 146 percent. Proximate causes of this asymmetry include a larger increase in male *applications* to gender-mismatched jobs, and the fact that men did not reduce their applications to men's jobs after the ban, while women did reduce their applications to women's jobs. On a more fundamental level, women's jobs on this board appear to be less differentiated than men's and to require fewer industry-specific skills. This may have made it easier for men to enter women's jobs than vice versa, at least in the short run.

Third, the gender requests posted by employers before the ban were overwhelmingly *not* used for affirmative-action purposes (i.e. to request the under-represented gender in a job title or workplace). It follows that the ban acted to integrate jobs and workplaces, though it is important to note that these integrating effects mostly occurred in relatively low-wage jobs. This is primarily because explicit gender requests were disproportionately used in less-skilled

jobs before the ban—a tendency that appears to be widespread.³ Related, the ban’s integrating effects did not touch many of the most extremely gendered jobs on this board, such as drivers, electricians, and nurses. This could reflect deeply ingrained stereotypes, plus the fact that some of these jobs require specific skills and certifications that take time to acquire. In the longer run, banning explicit gender requests may impact jobs like nurses and electricians as well, by signaling that investing in these skills may pay off for gender-atypical workers.

Finally, while we do not detect any spillover effects of the ban on jobs that never made a gender request (*N* jobs), we document some interesting changes in the segments of the XMRC labor market that were directly treated by the ban: *F* and *M* jobs. Specifically, total application rates to these jobs increased substantially (by 14.1 and 8.6 percent respectively), mostly because gender-mismatched workers who were discouraged by employers’ gender requests began to apply in greater numbers. Perhaps because these new applicants were highly qualified on other dimensions (like education, experience and location), the quality of applications to *F* and *M* jobs appears to have increased as well. Notably, while this loosening of labor markets at *F* and *M* jobs should benefit employers, it did not necessarily hurt workers. This is because many of the (lower-yield) gender-mismatched applications that were made after the ban would not have been made before the ban. They were discouraged by employers’ gender requests, and thus had no chance of yielding a call-back.

Our analysis contributes to a large literature on gender segregation in employment and its causes (Bielby and Baron 1984, Sorensen 1990, Blau et al. 2013, Baker and Cornelson 2018), and to a more recent empirical literature on directed search by workers. The latter literature has considered the effects of posted wages (Belot et al. 2018, Banfi and Roldan 2019); affirmative action statements (Ibanez and Riener 2018, Leibbrandt and List 2019); job characteristics such as negotiated salaries, competitiveness, and flexible work hours (Leibbrandt and List 2014, Flory, Leibbrandt and List 2015, Mas and Pallais 2017); and information on the number of competing applicants (Gee 2018) on workers’ application behavior. We also contribute to a literature that studies the effects of prohibiting the use of various types of information in the recruiting process. This information includes the worker’s gender (Goldin and Rouse 2000, Krause et al 2012), ethnicity (Behagel et al 2015), criminal history (Agan and Starr 2018, Doleac and Hansen 2020) credit history (Bos et al. 2018, Ballance et al. 2020) and salary history (Agan, Cowgill and Gee 2021). Relative to these papers, a unique aspect of our paper is that the prohibition applies to what workers see when they apply for jobs, not to what employers see when they are deciding whom to interview or hire.

Previous work on explicit gender requests in job ads includes Kuhn and Shen (2013) who provide a model of why and when employers might choose to include requests for demographic characteristics in job ads. Consistent with that model (which is based on application processing

³ Delgado Hellesester, Kuhn and Shen (2020) document this fact on four different job boards.

costs), we show that profiling is more common in less-skilled jobs. Delgado Helleseeter, Kuhn and Shen (2020) and Ningrum et al. (2020) show that there is a strong interaction between gender-profiling and explicit employer requests for other demographic characteristics, including age, marital status and beauty. Specifically, employers ask women to be young, single and attractive, and men to be older and married. Finally, Kuhn, Shen and Zhang (2020) show that—in a policy environment where gendered job ads are permitted—employers’ gender requests direct workers’ application decisions, and predict a worker’s chances of success conditional on applying to a job.

The most closely related research to our paper is Card, Colella and Lalive (2021), who study the effects of Austria’s prohibition of stated gender preferences in worker recruitment in 2004. One advantage of their setting relative to ours is that they observe several hiring outcomes—specifically, which workers were hired, the wage they earned and their job durations. In contrast, while we have some measures of match quality and recruiting frictions, we cannot see what happens after the call-back stage of the recruitment process. Card et al. also study somewhat longer-term causal impacts than we do; we focus on a one-year period surrounding the ban. Advantages of our setting include the following. First, the removal of gender labels from the XMRC website was not accompanied by any other legislative or regulatory changes affecting discrimination in job advertising or recruitment. In contrast, the Austrian Equal Treatment Act also regulated other aspects of recruitment, promotion and pay. Second, the treatment in our case was sudden and unexpected, being implemented overnight. This, combined with our high-frequency data, allows us to use a sharp temporal discontinuity as our main identification approach.

Third, Card et al.’s setting is one where vacancies are recommended to jobseekers by caseworkers of the Austrian labor market service (*Arbeitsmarktservice* or *AMS*), based on employers’ descriptions of the job and the ideal applicant (including gender). Our environment is a more common one where jobseekers decide where to apply based on the content of posted job ads. Since we have application data, we can learn considerably more about the mechanisms via which the ad ban operates in this more typical environment. Fourth, our data and setting allow us to create an accurate counterfactual for gender-targeted job advertisements after the ban: We know exactly which jobs *would have contained* explicit gender requests after the ban because we can observe the same ad before and after its gender request was forcibly removed.⁴ A final difference between the papers’ approaches is Card et al.’s focus on gender segregation across firms as their main outcome of interest; inspired by a long literature on gender-stereotyping of different types of work, we focus more on the ban’s integrating effect at the job

⁴ Notably, the gender labels were attached to these ads when they were first posted, at which time there was no expectation of an imminent ban. Card et al. are forced to classify jobs that are likely to be more versus less impacted by the ban, by constructing a gender preference index using characteristics of the job advertisements and firm fixed effects.

title level. Despite these differences in setting, data and focus, several key results are common to both papers. Most notable of these is the fact that prohibiting explicit gender requests integrated labor markets. In both papers, this integration occurred despite the ban's segregating effect in the small number of cases where employers requested gender-atypical workers before the ban.

2. XMRC, the Policy Change, and the Analysis Sample

2.1 XMRC and the Policy Change

Our data consist of internal records of XMRC.com, an Internet job board serving the city of Xiamen and surrounding XiaZhangQuan region since 2000. Xiamen and its surroundings comprise a high-income coastal Chinese city with a population (including migrant workers) of about 14 million people. XMRC is a private firm, commissioned by the local government to serve private-sector employers seeking relatively skilled workers.⁵ Its job board has a traditional structure, with posted ads and resumes, on-line job applications, and a facility for employers to contact workers via the site.

Advantages of the XMRC data for the analysis of a gendered ad ban include the following. First, before the ban, ads on the site contained a simple, standardized indicator of employers' gender requests: when a job profile was created, the hiring agent selected the preferred gender from drop-down menu: *F*, *M*, or *N* (no request). Workers saw this selection when they read job ads and searched for jobs. On many other job boards (both in China and elsewhere), explicit and implicit gender requests must be inferred by parsing the text of the ad, a process which requires some judgment calls.⁶ Second, our data come from a period during which XMRC offered a very simple search technology: workers used keywords and menus to look for jobs, and (much less often) employers did the same to find workers. On some other job boards today, proprietary algorithms display suggested job matches to individual workers based on the worker's location, qualifications, employment history and recent searches. In these cases, unlike ours, workers' application decisions are jointly determined by the jobs that are suggested to them by the board's algorithms and their choices from that set.⁷

Third, the removal of gendered job ads on XMRC was sudden and unexpected: While China introduced fines for job boards posting such ads in 2016, enforcement (especially for

⁵ The only other major local job site, XCMHRSP (formerly XMZYJS) serves private sector firms seeking mostly production and low-level service workers. Public sector recruitment operates via separate channels.

⁶ For example, in Spanish one must decide whether the job titles "abogada" and "abogado" are explicit gender requests.

⁷ We do not observe which ads were viewed by workers; thus our estimated effects should be interpreted as incorporating workers' decisions regarding which types of jobs to search for.

regional boards like XMRC) was haphazard.⁸ Essentially, XMRC was able to ignore this regulation until late February 2019, when it was directed to remove the pre-configured ‘desired gender’ field from all its ads overnight, without any advance notice to recruiters or workers. To interpret our results, it is important to note that *only* these encoded gender labels were removed on March 1, 2019: expressions of gender preferences in the text of the job description were not removed from ads that were present on the site on March 1.⁹ Thus, our main estimates refer to the effects of removing only the most salient expressions of gender preferences from a job ad. In our robustness analysis, we show, as one might expect, that estimated treatment effects are modestly higher when we exclude ads which contained additional statements of gender preference in their text. These ‘embedded’ preferences were not removed on March 1, 2019.

Fourth, while XMRC ads and resumes are not representative of the employed population of Xiamen, we believe they are quite representative of the stock of skilled, private-sector *vacancies* and *job-seekers* in large Chinese cities (which are younger, better educated, better paid, more female, and much more likely to have non-local *hukou* than these cities’ employed populations).¹⁰ This is because (a) Xiamen has a similar demographic profile to other large, prosperous Chinese cities, and (b) XMRC is the dominant platform for skilled, private sector workers in Xiamen.¹¹ XMRC attributes this dominance, in part, to its close association with the local government, which provides complementary services (e.g. for social security and payroll taxation) that are electronically and physically linked to the job board, thus providing ‘one-stop shopping’ for local employers. Finally, while the XMRC board we study caters to the more skilled side of Xiamen’s labor market, we note that explicit gender requests are considerably more common at lower than higher skill levels (Kuhn and Shen, 2013). This suggests that our results could underestimate the impact of removing such labels in the labor market as a whole.

2.2 Sample Construction

Our main analysis sample is extracted from the population of 239,564 ads present on XMRC during the one-year period surrounding the ad ban, and the 4,338,694 applications that were made to those ads during that time period. From these, we retain only the ads that ‘span the ban’: they were posted before the ban, and received a positive number of applications both before and after the ban. As a result, every ad in our sample has an explicit gender label (*F*, *N*,

⁸ See Appendices 2 and 3 in Kuhn, Shen and Zhang (2020) for a recent history of regulation and practice regarding gendered job ads in China.

⁹ For ads newly posted or renewed by the employer after March 1, XMRC personnel were instructed to monitor the content and ask employers to remove any discriminatory content. By construction, these ads are not included in our estimation sample.

¹⁰ *Hukou* refers to the province of a Chinese person’s permanent residence permit, which in most cases is determined by their province or city of birth. See Table A4.4 in Kuhn, Shen and Zhang (2020).

¹¹ Table 2 in Delgado Helleseter, Kuhn and Shen (2020) provides descriptive information about Xiamen in comparison with other large Chinese cities. At the time of our study, the large national job boards like 51job and Zhaopin did not have a large presence in Xiamen.

or *M*) that was assigned when it was posted. In addition, recruitment for the ad was active (in the sense that the ad was receiving applications) both before and after the ban. This allows us to study how application and recruitment behavior changed within the same ad, before and after its explicit gender request was removed by the job board.

In all, our primary dataset comprises 3,130,317 applications made by 204,343 workers (resumes) to 116,725 ads, placed by 154,437 firms, resulting in 348,062 call-backs. Thus there was an average of 26.8 applications per ad and 2.98 call-backs per ad. 11.1 percent of applications (for which we have call-back information) received a call-back, suggesting that call-backs indicate a relatively high level of success in the recruiting process. Overall, 11.6 percent of job ads requested female applicants, 12.5 percent requested male applicants and the remaining 75.8 percent did not specify a preferred gender. 41.2 percent of applications came from women. Reflecting XMRC's focus on skilled jobseekers, the average ad requested 13.30 years of education and the mean requested worker age was 29.78. Among the 75.7 percent of ads that posted a wage, the mean posted wage was 5,794 and 4,432 RMB per month in jobs requesting men and women respectively, yielding a raw gender wage gap of 23.5 percent.¹²

In addition to this main analysis sample, we construct two additional samples to address specific questions and conduct robustness checks. The first of these—our *call-back sample*—addresses the fact that we do not observe all the call-backs that were made on XMRC. Specifically, we only see call-backs for the job ads where the recruiter used XMRC's internal messaging system to contact applicants, which account for 62 percent of all applications. Thus, while we use our full sample for most of our analysis, our analyses of call-back outcomes are limited to the 1,939,935 applications for which call-backs are observed. Fortunately, Appendix 2 shows that this sample has very similar observable characteristics to the full sample. In addition, Appendix 7 shows that replicating our main full-sample analyses that do not require call-back data on the call-back sample yields almost identical results.

Our second alternative sample replicates our main estimation sample—which comprises applications made between September 2018 and August 2019—on two different periods: January – August 2018 and January-August 2019. The latter period contains the date on which the 2019 ban occurred, and the former contains the date on which it *would have occurred* in 2018. Unfortunately, these two periods cannot be designed to exactly mimic our main analysis sample because we have no XMRC data from 2017. While this restricts the length of the pre-ban period in both years to just two months, it allows us to compare trends in our main outcomes between 2018 and 2019 on both sides of the ban date. It also allows us to conduct a difference-in-difference analysis of the ban's effect—which uses equivalent days or weeks from 2018 as controls for 2019-- as a robustness check of our main results. Notably, since important events affecting China's labor market—especially the Spring Festival-- are determined by the lunar calendar, this new *DiD sample* requires us to line up days and weeks between 2018 and

¹² Wages are calculated using the midpoint of each job's posted wage range. See Appendix 1 for additional descriptive statistics on the ads and applications in our main analysis sample.

2019 to that they represent the same days and weeks relative to the start of the Spring Festival in both years. Additional details are provided in Appendices 2 and 11.

3. Estimation Approach

Of the three types of jobs in our data -- M jobs, which requested men before the ban; F jobs, which requested women before the ban; and N jobs which never made a gender request—only F and M jobs are directly treated by the ad ban, in the sense that only these jobs had their contents changed. But even though these treated ads comprise only about one fourth of the ads in our sample, there are two main reasons why we might expect changes in their contents to affect the non-treated (N) ads. The first is treatment spillovers, or substitution effects: for example, the ban could re-direct some workers' applications from non-gendered ads into formerly gendered ads, because workers who once felt excluded now apply there. Second, the ban could cause changes in the equilibrium levels of job search and recruiting intensity. For example, if the ban made it harder for workers to find jobs (because it became harder for workers to direct their applications away from jobs where they are less likely to succeed), workers might compensate by applying to more jobs of all kinds.¹³ To incorporate these types of spillover and equilibrium effects, we therefore need to estimate the effect of the ban on the non-treated ads (N) as well as the directly treated ones. We cannot use N jobs as controls for F and M jobs.

To estimate the ban's effects on all jobs and workers, our main approach uses a full year of data surrounding the gendered ad ban, and estimates the discontinuity in each of our outcome measures on the first day of treatment--March 1, 2019. Two key choices in any such design are the bandwidth around the policy change, and the parametric modelling of trends in the running variable (in our case, time) within that interval. Motivated by the fact that continuity assumptions play an especially important role in RD-in-time contexts like ours (Hausman and Rapson 2018), our approach to these issues begins with an inspection of the raw trends in each of our outcome variables. For a key subset of the outcomes we examine -- specifically, the gender composition of applicant and call-back pools—these trends are quite smooth during the six months on either side of the ban. For these outcomes we analyze the data at the weekly level using the full year of data surrounding the ban, and fit a flexible polynomial to the time trend over this entire interval. Appealing features of this approach include the ability to estimate causal effects of the ad ban over a longer time horizon.

For our remaining outcome variables, the time trends during the year surrounding the ban are far from smooth. This is especially true for outcomes related to the overall level of recruiting activity, such as the total number of applications received by a job in a given day or week. Because the ad ban happened 24 days after the 2019 Chinese New Year, it occurred shortly after an annual spike in recruiting activity associated with the return to work of millions

¹³ Alternatively, workers could become discouraged and search less. Shimer (2004) shows that offer arrival rates can affect optimal search intensity in either direction.

of workers, many of whom travel home to rural areas for the New Year and Spring Festival holiday. To abstract from these trends, we turn to daily data and focus on a narrow window of 30 days surrounding the ban. This window starts after the peak of post-holiday recruiting activity and exhibits relatively smooth time trends within both the pre-and post-ban periods. In this sample, we estimate local linear regressions, fitting different linear trends on either side of the ban. While this approach sacrifices a lot of data, we believe it provides the most accurate estimates of the ban's causal effects on these outcome variables.

As an alternative approach to identification, Appendix 11 (discussed in Section 7) uses our *DiD* sample to produce difference-in-difference estimates that use weeks in 2018 as controls for 2019. This allows us to relax the smoothness assumptions we need for our main estimates by using a full set of calendar week fixed effects to capture sharp time trends that are common to both years. It also obviates the need to focus on just a 30-day window for the above outcomes: We can now use weekly data for the entire year for all our outcome variables. In addition to some data constraints (most importantly a shorter pre-ban period and poorer-quality call-back data), the cost, of course, is the need for a parallel trends assumption. Fortunately, the main results are very similar.

4. Did the Ad Ban Change the Gender Mix of Successful Applicant Pools?

4.1 Unadjusted Time Trends

To provide a first look at the how the ban affected the gender mix of applicant and called-back worker pools, Figure 1 displays aggregate, unadjusted time trends in these outcomes for non-gendered (*N*) jobs, and for jobs that (initially) requested women and men (*F* and *M* jobs). It shows, first of all, there was no perceptible change in these two outcomes at the time of the ban in *N* jobs-- the 75 percent of job ads whose contents were not altered overnight. In jobs that initially requested men, however, removing that request raised women's share of call-backs by about three percentage points-- from 4.86% to 7.93%. This shift was immediate, constant in magnitude, and persisted for the lifetime of these ads (up to six months after the ban). Figure 1(b) also shows that the main mechanism for this effect was application behavior: the share of women in *applicant* pools to male jobs increased from 5.57% to 10.52%. Interestingly, this surge in applications was attenuated by only a modest decline in women's relative chances of success in men's jobs. Specifically, conditional on applying, women's call-back rates in *M* jobs fell from 86.6 percent of men's to 73.3 percent. Similarly, Figure 1(c) shows that the ad ban raised men's share of call-backs to (previously) female jobs by almost 13 percentage points, from 6.81% to 19.44%. Again this effect was immediate, constant, and permanent; and the main mechanism was application behavior: men's share in the applicant pool rose from 10.04% to 25.95%. Interestingly, this surge in applications was reinforced by a small increase in men's relative chances of success in women's jobs, from 65.4 percent of women's to 68.8 percent.

Could the sharp jumps in application patterns on March 1, 2019 be associated with some other event at that time of year that we cannot observe? To check for this possibility, Figure 2 turns to our *DiD* sample, which lines up weeks in 2018 and 2019 such that the first week after 2018's pseudo-ban starts on the same lunar day as the actual ban. While the time trends for N jobs show no jumps on that date in either year, the time trends in applications and call-backs in both F and M jobs show a clear pattern: large jumps in 2019 and no change in 2018. We view this as compelling evidence that the 2019 jumps in the above outcomes were caused by the gendered-ad ban.

To summarize, Figures 1 and 2 and illustrate three main results. First, the ban caused a large upsurge in the share of 'gender-mismatched' applications to the ads whose contents were changed by the ban. Second and somewhat surprisingly, the rise in male applications to formerly female jobs was much larger than the increase in female applications to male jobs. We will explore possible causes for this asymmetry later in the paper. Third and perhaps most unexpected is *how well gender-mismatched applications were treated by employers* both before and after the ban, with success rates ranging from 65 to 87 percent of gender-matched applications. This fact explains why the upsurge in gender-mismatched applications changed the gender mix of call-back pools.

4.2 Regression Analysis

We next apply regression analysis to our main estimation sample to control for three potentially confounding factors affecting the time trends in Figure 1. The first of these factors is duration dependence within recruiting spells: since all our job ads were posted before the ban, the post-ban weeks will, on average, occur later in a recruiting spell. This could affect our estimates if, for example, men and women had different tendencies to apply to new versus old ads. Second, as recruiting spells become complete after the ban, the mix of job ads remaining in the estimation sample could change. For example, if ads for receptionists tend to be filled relatively quickly, then these stereotypically female jobs will be under-represented in the post-ban period relative to the pre-ban period. While we expect both these factors -the average age of ads and the mix of ads remaining active—to change relatively smoothly with calendar time, we nevertheless assess the effects of controlling for them on our main results. Finally, we are concerned that cyclical and seasonal factors could affect application and call-back patterns, for both macro-economic and cyclical reasons.

In more detail, we run regressions in which an observation is an ad-week cell, where j denotes ads and t denotes weeks. The outcome variable, Y_{jt} is the female share of applications to the job, or the female share of call-backs to the job occurring in week t .¹⁴ Our objective is to

¹⁴ Call-backs to an application can of course take place in a later week than the application. Our estimation approach in this section counts call-backs in the week they occurred; this allows the ad ban to affect call-backs to applications that were made before, as well as after the ban.

estimate the ban’s effect on the representation of women among successful applicants to jobs that (initially) requested men (M jobs), requested women (F jobs), or that never made an explicit gender request (N jobs). We parameterize the effects of the ad ban on these outcomes as follows:

$$Y_{jt} = \beta^1(Post_t \times F_j) + \beta^2(Post_t \times M_j) + \beta^3(Post_t) + \beta^4 F_j + \beta^5 M_j \quad (1)$$

where F_j and M_j indicate the job’s gender request *when it was first posted*. Thus, we treat non-gendered (N) jobs in the pre-ban period as our reference category. In this parameterization, β^4 measures how the female share of applications or call-backs differed between F and N jobs during the pre-ban period. Analogously, β^5 compares the female share between M and N jobs in the pre-ban period. β^3 measures the ban’s effect on the female share in ads that do not express a gender preference. Finally, β^1 and β^2 measure how banning gendered ads affected recruitment outcomes in jobs that requested men and women before the ban, *relative* to the ban’s effect on non-gendered jobs.

Column 1 of Table 1(a) reports regression results for female applicant shares exactly as specified in equation 1, with no additional controls.¹⁵ Importantly, because the female share is estimated much more precisely in weeks when, say, 20 applications arrived than when just one application arrived, all these regressions are weighted by the total number of applications received in that week.¹⁶ Column 1 suggests, first of all, that removing all the gender labels from this job board had a small positive effect on the female share of applications to previously non-gendered job ads: the coefficient of .0048 represents about a 0.5 percentage point increase on the pre-ban female share of 40.9 percent. Removing gender labels, however, raised the female share of applications to jobs that previously requested men by 4.95 percentage points, from a base of 5.58 percent.¹⁷ Even more dramatically, removing gender labels reduced the female share of applicants to previously female jobs by 15.91 percentage points. Thus, the ban raised the male share of applicants to ‘female’ jobs from a base of 10.04 to 25.95 percent. These results mirror the patterns observed in Figure 1.

To address possible contamination of the above results from vacancy duration dependence, column 2 of Table 1 controls for the age of each job ad using a quartic for the

¹⁵ Because only ad-week cells that contain a positive number of applications provide information on female shares, Table 1 is estimated on the subset of ad-weeks that received at least one application. This is considerably smaller than the total number of application-weeks in our data, because many of those weeks occur late in a recruitment spell, when few applications arrive.

¹⁶ We have verified that these application-weighted regressions of the female applicant share yield numerically identical estimates as unweighted linear probability models where an observation is a single application that is made in a job*week and the outcome is the probability the application came from a woman.

¹⁷ This effect (like all the results summarized in this Section) is the sum of the post-ban coefficient and the corresponding interaction term, i.e. $4.95 = 0.48 + 4.47$.

number of weeks since it was posted. Column 3 introduces a quartic in calendar time to control for secular and seasonal trends. Column 4, which is the most saturated specification that identifies the effect of the ad ban on all three job types (*F*, *M* and *N*) adds fixed effects for job ads to control for any selective re-filling of vacancies. Notably, these three treatment effects remain very similar to the ones in columns 1-3. Notably, column 4 implies that the ban raised the female applicant share to *M* jobs by $0.53 + 3.48 = 4.01$ percentage points, or $4.01/5.58 = 71.9$ percent. At the same time, the ban raised the male applicant share to *F* jobs by a much larger $13.92 - 0.53 = 13.39$ percentage points, or $13.39/10.04 = 133$ percent. The ban's effect on the gender mix of *N* jobs remains very small and precisely estimated.

Finally, we do two things to assess whether Table 1's results are robust to our use of a quartic to capture seasonal and cyclical effects. First, Appendix 5 shows that the ban's estimated effects are unchanged if we replace the quartic in calendar time by fifth- or sixth order polynomials, by separate quadratics on either side of the ban, or by separate quartics on either side. Second, column 5 of Table 1 adds a full set of calendar week fixed effects to the regression. Since the gender ad ban happened at a single point in time, the treatment effect on the non-gendered job (*N*) ads is no longer identified, but the ban's effects on the *F* and *M* jobs *relative* to the *N* jobs remain identified, and are essentially unchanged from the previous columns.

Table 2 performs a parallel analysis to Table 1 for call-backs, which exhibit very similar patterns. Focusing again on the column 4 results, Table 2 shows that the large increases in gender-mismatched applications caused by the policy also led to substantial increases in the gender mix of *successful* applicants at those jobs. Specifically, we estimate that banning gendered ads raised the representation of women among call-backs to men's jobs by $(2.46 + 0.49) = 2.95$ percentage points, or $2.95/4.77 = 61.8$ percent. It raised the share of men in the call-back pool to *F* jobs by $(10.39 - 0.49) = 9.90$ percentage points, or $9.90/6.79 = 145.8$ percent. The ban had no statistically significant effect on the share of women in the call-back pool to *N* jobs. All these results survive the same robustness checks that were applied to our application analysis (i.e. using a variety of polynomials in Appendix 5, and adding calendar week fixed effects in column 5 of Table 2).

In sum, Section 4 has demonstrated that (a) the ban increased the representation of gender-mismatched workers in call-backs to previously gendered jobs; (b) the main mechanism was via changes in application behavior, and (c) these effects are considerably stronger for men than women.

5. Which Types of Work, and Which Types of Workplaces were Most Affected by the Ban?

5.1 Historical Gender Mix of the Job Titles Most Affected by the Ban

Table 2 has shown that the gendered ad ban increased women’s representation in jobs that previously requested men, and increased men’s representation in jobs that previously requested women. Whether these changes reduced gender segregation, however, depends on how employers used their gender requests before the ban. If employers’ gender requests mostly reproduced the gender mix of the incumbent workers in a job or workplace, our estimates imply that the ad ban reduced gender segregation. However, if the explicit gender requests on XMRC were used primarily for affirmative action purposes (i.e. to attract gender-*atypical* applicants), the ban would be impeding employers’ good-faith attempts to integrate their workforces.

To distinguish between these scenarios, we turn to data on all the call-backs issued on XMRC between January and August 2018—a period that pre-dates our estimation sample. In Figure 3 we group these call-backs by their job title, and use the gender mix of call-backs to a job title as a proxy for the gender mix of the incumbent workers in that type of work (i.e. the workers who were previously hired into the title). Job titles are the short piece of text (typically 3-8 words) at the top of a job ad that summarize the main type of work involved; they are more detailed and descriptive than six-digit SOC codes.¹⁸ In Figure 3, the horizontal axis shows the incumbent female share in a job title, calculated in this way. The vertical axis shows the share of job ads for that title that explicitly requested men or women in our *estimation* sample. The message of Figure 3 is unmistakable: Overwhelmingly, employers on XMRC did not use explicit gender requests for affirmative action purposes (i.e. to ‘lean against the wind’). Instead, employers requested workers whose gender was already dominant in the job title being filled.

Given that employers’ gender requests reinforced existing stereotypes, just how historically male or female were the jobs that men and women entered (in the sense of increasing their share of call-backs) as a result of the ban? Motivated by Table 2’s evidence that the ban did not affect gender mix in non-gendered (*N*) jobs, we focus our attention on the substantial changes that occurred in *F* and *M* jobs. We begin in Table 3(a), by using the above historical title-level means to calculate the incumbent male share of all job ads that requested men (i.e., all *M* jobs) in our estimation sample. In a sense, these *M* jobs were the ones that were ‘opened up’ to women by the ban (since it removed their requests for men). This share is 0.878, indicating that the jobs *whose requests for men were removed* had a high incumbent male share. To measure the incumbent male share of the job titles where women’s call-back shares actually increased as a result of the ban, Appendix 3 decomposes the total increase in women’s

¹⁸ See Marinescu and Wolthoff (2020), who show that job titles are more predictive of wages and application decisions than are six-digit SOC codes.

share of call-backs to M jobs (i.e. the $2.46 + 0.49 = 2.95$ percentage point percentage point treatment effect identified in Table 2) into contributions from individual job titles: Each title's contribution is the product of the title's *prevalence* among M ads (α^T) and the pre- versus - post ban *increase* (within M ads) in women's share of call-backs to that title (δ^T).¹⁹ Using these contributions as weights, column 2 calculates the incumbent male share of the job titles that women *entered* as a result of the ban as 0.677. While this is still more historically male than an average job (0.499), it is much less male than a typical male-requesting job.

Table 3(b) repeats the preceding exercise for men's entry into F jobs, showing first that the jobs that requested women before the ban were historically 78.5 percent female (i.e. less female than the male-requesting jobs were male.) The jobs that men entered as a result of the ban, however, were 73.4 percent female, indicating that men made deeper inroads into historically female job titles compared to women's inroads into men's. This is starkly illustrated by column 3 of Tables 3(a) and (b): the gap between the historical male share of the jobs the ban 'opened up' to women and jobs they actually entered (20.1 percentage points) was almost four times as large as the analogous number for men (5.1 percentage points).

To paint a broader picture of the historical gender mix of the job titles that women and men entered because of the ban, Figure 4(a) plots the relationship between the incumbent male share in a job title and the (pre- versus post ban) *change* in women's share of call-backs to those jobs. Overall, consistent with Table 2, Figure 4(a) shows that women's access to jobs that previously requested men increased after the ban. In addition, Figure 4(a) shows that these increases were substantial in male-requesting jobs with incumbent male shares of up to about 80 percent, but small and statistically insignificant (1.87 percentage points, $p=0.130$) in job titles whose incumbent male share was above that level. Similarly, Figure 4(b) shows that men's inroads into jobs that previously requested women were smallest in the job titles that were the most historically female. In contrast to Figure 4(b) however, we still see substantial and statistically significant increases in men's call-back share even in job titles that were historically more than 80 percent female. When a request for female workers was removed from jobs with titles that were 80 percent or more female, men's call-back share still increased by 6.56 percentage points ($p = 0.000$).

5.2 Historical Gender Mix of the Positions, Workplaces and Firms Most Affected by the Ban

In Table 3(c) we apply the methods used for job titles in Tables 3(a) and (b) to measure the incumbent gender mix of the *positions*, *workplaces*, and firms that women and men entered as a result of the ban. A position (firm x job title cell) is a particular type of work at a given firm,

¹⁹ For simplicity and transparency, these decompositions were performed on unadjusted pre-post-ban changes in the gender mix within job titles. Thus, the total ban effects that they account for differ slightly from the regression coefficients in Tables 1 and 2.

such as all the HR managers at Dell Computer. A workplace (firm x location cell) refers to the operations of a firm at one of the 47 different job locations that are coded in XMRC's system.²⁰ Table 3(c) shows that—in addition to pulling women into *job titles* that were historically more male than average—the ban also pulled both men and women into positions, workplaces and firms where their gender had been under-represented. For both men and women, the ban's integrating effects were strongest at the *position* level, which is the most disaggregated unit we consider, and next-strongest at the title level. In these two categorizations (which both involve the type of work that is performed) the ban's integrating effect was stronger for men than women; we discuss possible reasons for this below. Interestingly, the ban also acted to integrate firms and workplaces, but these effects were considerably weaker, and slightly stronger for women than men.

5.3 Wage Levels in the Jobs Most Affected by the Ban

Finally, to characterize the level of wages in jobs that men and women entered because of the ban, Table 3(d) applies the same decomposition methods used in parts (a) and (b), but changes the outcome variable from a job title's incumbent gender mix to the mean posted wage in all (historical) ads for title.²¹ To provide some context for this exercise, column 1 reports the mean (historical) log wages of all the job ads requesting men, women and making no request. Column 1 shows, importantly, that jobs that make explicit gender requests are disproportionately low-wage jobs. Specifically, job ads requesting men (women) offer wages that are $8.858 - 8.738 = .120$ and $(8.858 - 8.634 = .224)$ log points lower than jobs that make no gender request. This is a common feature of gender requests on a number of different job boards (Delgado Hellesester et al. 2020), and is consistent with the idea that firms abstain from gender typing in the most skilled jobs because it is more important to identify the single best candidate (Kuhn and Shen 2013). Thus, simply because of the way employers tend to use gender requests in job ads, we should not necessarily expect the removal of gender requests to substantially raise either men's or women's access to highly paid positions. Next, column 2 of Table 3(d) shows that --compared to all jobs that had previously requested men—the *M* jobs that women entered after the ban paid about 11.7 percent less; for men this number was almost identical (11.0 percent less). Thus, while the ban integrated workplaces, its effects were concentrated in types of work that were not highly paid.²²

²⁰ The 47 locations available in the data pertain to a mix of geographic units. These include the six districts of the city of Xiamen, various cities (including Xiamen), provinces within mainland China, and outside mainland China. 93.2% of the applications are sent to job ads located in the city of Xiamen or one the six districts of Xiamen.

²¹ See Appendix A3.4 for details on how the shift-share analysis was applied to posted wages.

²² Table 3(d) is not directly informative of the ban's effects on either men's or women's wage-earning opportunities. This is because opportunities depend on all the call-backs a worker receives and how each individual worker's call-back set changes. For example, if all the additional women who got call-backs from *M* jobs previously only got offers from *N* jobs, their mean call-back wages probably fell (though the

5.4 Which Job Titles Mattered Most?

Another advantage of the shift-share decomposition underlying Table 3 and Figure 4 is that it lets us list the individual job titles that account for the main treatment effects we identified in Table 2. To that end, Appendices 3.1 and 3.2 list the 20 job titles that made the largest contributions to those effects (for women and men separately). Consistent with Table 3 and Figure 4, only 13 of top 20 job titles that accounted for women’s increased representation in male-requesting jobs had above-average historical male shares (for example, warehouse management). The remaining seven titles (like manager’s assistants) made substantial contributions to women’s increased representation in *M* jobs because they were common (i.e. with high α ’s) and sometimes requested men. In contrast, 18 of the 20 most important titles to which men gained access were female-dominated (like administration specialist); only two (like warehouse management) were historically male. Thus, the fact that the ban removed *some* gender requests that ‘leaned against the wind’ limited the magnitude of ban’s overall desegregating effect, especially for women.

As a point of reference for the job titles most impacted by the gendered ad ban, Appendix 3.3 lists the historically *most* male- and female-dominated jobs on XMRC. These extremely gendered titles include a substantial large number of driving titles (call-backs to “driver” positions were 99.66 percent male), and skilled trades (plumber, welder, electrician, mold fitter and internet management staff were each more than 99 percent male). By far the largest female-dominated title in this list was “administration and receptionist”, which was 98.25 female. Notably, only one (relatively small) title –civil construction worker—in these 80 extremely gendered job titles appears among the 40 titles that account for most of the ban’s effects. Thus, the ban left these extremely gendered jobs essentially untouched.

Taken together, the results in this Section shed additional light on the striking asymmetry of the ad ban’s effects on women and men documented in Section 2 of the paper. Aside from the fact that the ban raised men’s *applications* to *F* jobs much more than women’s applications to *M* jobs, one more fundamental potential cause of this asymmetry is the surprisingly large share of driving jobs among highly male jobs. Whether for personal preference and safety reasons (as documented in Cook et al.’s (2021) study of Uber drivers), strong cultural norms, or anticipated insider resistance, Chinese women did not move into driving jobs after the ban. Another potential contributor is the large role of skilled trades in extremely male-dominated job titles: in the short time horizon of our study, women interested in entering such jobs would not have the time to acquire the required credentials. A final possibility is a granular form of skill specificity: If we think of job titles as indicators of distinct skills (or skill bundles) that are needed to perform a job well, then the number of distinct titles describing the most

maximum may not have changed). The opposite is the case if these women previously only had call-backs from *F* jobs.

gendered jobs on XMRC would provide an indicator of how likely a random applicant is to be qualified for those jobs. Defining *highly male (female)* vacancies as vacancies that requested men (women) in *job titles* with a historical male (female) call-back share of at least 80 percent, Appendix A3.5 shows that 23, 33 and 50 percent of highly female vacancies were in the 5, 10 and 20 most common job titles, compared to 14, 23 and 33 percent highly male vacancies. Much of this difference is associated with the dominant role of female job titles related to administration, clerical, and reception tasks, which—in contrast to trades—seem less industry-specific. The predominance of such highly female titles may have made it easier for men to access women’s jobs (in the short run) than the converse.

6. Effects of the Ban on Firm and Worker Well-being

So far, our analysis has shown that the gendered ad ban integrated pools of successful applicants. As already noted, however, this integration may have come at a cost: the ban may have made it harder for workers to find good jobs, and for firms to hire qualified workers. To investigate these effects, this Section investigates the ban’s effects on two indicators of firms’ ability to recruit qualified workers—the arrival rate of applications to job ads and the match quality of those applications (application *arrivals* and *quality*, respectively)—and one indicator of workers’ ability to find good jobs: the call-back rate per application submitted (application *yield*).²³ The first stage of our analysis is conducted on an aggregate level that combines all the job ads in our sample, regardless of whether they made a gender request. The goal of this stage is to see whether the ad ban—which changed the contents of about one fourth of the ads on the job board—had any detectable effects on aggregate outcomes, allowing for the fact that treatment of some ads can have spillover effects on other ads, for example via changes in workers’ application behavior. The second stage of our analysis focuses only on the ads that were directly treated by the ban (i.e. the *F* and *M* jobs), where we expect to see the largest effects on firm and worker well-being. Specifically, how did removing the gender requests from an employer’s job ad affect the number and quality of applications that employer received? And how did workers’ chances of getting a call-back from those job ads change after applications were ‘officially’ opened up to both genders? In addition to answering these questions, this disaggregated analysis has two additional functions: to illuminate the mechanisms behind the ban’s integrating effects identified in Sections 4 and 5, and to shed additional light on the striking gender asymmetry in those effects.

²³ We explored using the posted wages of workers’ call-backs as an additional measure of worker well-being, but our estimates were very imprecise. We suspect this is because wages attached to individual jobs are posted in wide ranges and are not available for all jobs. These issues have less impact in Section 5, where we used mean posted wages of *job titles* to generate descriptive information on which titles that were most affected by the ban.

6.1 Aggregate Effects of the Ban

As a first look the ban's aggregate effects on firm and worker well-being, Figure 5 shows unadjusted trends in application arrivals, quality and yield during our year-long sample period surrounding the ban. Notably, all three outcomes show pronounced time trends shortly before the ad ban, associated with large changes in recruiting and job search activity before and during China's Spring Festival. Most dramatically, part (a) of Figure 5, on application arrival rates, shows a steady decline in applications per job ad between October 2018 and the end of January 2019. This reflects the fact that many employment spells in urban China end at the start of the Spring Festival, so there is little incentive to change jobs or workers as the holiday approaches. During the Spring Festival (the shaded area shortly before the ban), application activity rises dramatically. This suggests that many workers are looking for new job matches to start after the holiday. The sharp increase in application activity slows dramatically during the two weeks before the ban.

While the pronounced holiday effects in Figure 5(a) make it hard to estimate the ban's effects on outcomes like the overall application arrival rate, there is no obvious discontinuity in overall arrival rates in the week of the ban. Figure 5(a) does, however, show a pronounced 'fanning-out' of application rates between F , N and M jobs at the ban date, with daily applications to the previously gendered jobs (F and M) rising above applications to N jobs; we will study these differential trends by job type in Section 6.2. Parts (b) and (c) of Figure 5 show that the large increase in applications during the Spring Festival coincided with a substantial rise in application quality and yield, complicating our attempts to isolate the ban's effects on these outcomes as well. Still, in sharp contrast to Figures 1 and 2 (which showed large discontinuities in the gender mix of applications and call-backs) none of the three series in Figure 5 exhibit obvious discontinuities at the time of the ban.

Motivated by the trends in Figure 5, our main analyses of the ban's effects on firm and worker well-being fit local linear regression discontinuity models to *daily* data from the 30-day period surrounding the ban. During this window, the time trends in all these outcomes on either side of the ban are much less pronounced than during the Spring Festival. This gives us a better chance of isolating the ban's causal effects, but it does reduce the statistical power of some of our estimates, relative to the main effects identified in Section 4.²⁴

Starting with the ban's effects on firms' overall abilities to fill a vacancy, columns 1 and 2 of Table 4 study the daily arrival rate of applications to active vacancies.²⁵ Focusing on the 30-day period surrounding the ad ban, these columns estimate:

²⁴ In Appendix 11 we identify the ban's effects on application arrivals, quality and yield using an alternative approach that allows us to use data from a much longer period surrounding the ban. This approach uses 2018 data as a 'control' for 2019 outcomes, thus relying on a different set of identifying assumptions. The results are broadly similar.

²⁵ A vacancy is defined as active between the time it is posted and the time it receives its last application. We focus on the application arrival rate rather than the total number of applications eventually received, because the completed applicant pool size is, in part, a choice made by the firm.

$$Y_{jt} = \alpha + \beta(Post_t) + \delta^1 t + \delta^2 t(Post_t) + \theta X_{jt} \quad (2)$$

where Y_{jt} is now the *number* of applications (including zero) received by ad j on day t , and t indexes days relative to the gendered ad ban. The parameter β gives the size of the discontinuity on the first treatment day (March 1, 2019), while the δ s allow for different linear time trends on either side of the ban. In column 1, the controls, X_{jt} , include day-of-week fixed effects and a quartic in job age. Column 2 adds job ad fixed effects, giving us our most tightly controlled estimate of the ban's effects on the total number of applications received per day by a representative job ad on XMRC, essentially comparing the flow of applications within each individual job ad before versus after the ban. It is a fairly precisely estimated zero. Specifically, the estimated coefficient shows an increase in the daily arrival rate of .0017 on a baseline rate of .140. The 95 percent confidence interval around the coefficient runs from -.0008 (a 0.6 percent decline) to .0042 (a 3.8 percent increase).²⁶ Thus, while our estimates do not rule out a very small negative effect of the ban on overall application arrival rates, in combination with Figure 5(a) they suggest that the ban did not reduce the size of employers' applicant pools.

As a second indicator of whether the ad ban affected firms' ability to recruit, columns 3 and 4 of Table 4 use an index of the mean quality of applications received by an ad during a week. This quality index is based on an indicator developed by XMRC to measure the match between the job's requirements (from the job description) and the worker's previous working experience (from the resume).²⁷ To assess the validity of this score as a proxy for match quality, we regressed our call-back indicator on it, while controlling for job ad, applicant, and job day fixed effects.²⁸ In the sample of 3,023,024 applications with non-missing match scores, a one standard deviation increase in the match score raises the chance of a call-back by 1.028 percentage points ($p=.00000$). This represents a 9.26 percent increase relative to the mean call-back rate of 11.1 percent in the data.

Aside from the different outcome variable, the only change in how columns (3) and (4) were estimated (relative to columns 1 and 2) is in the unit of analysis, which is now at the *application* level (not the ad-day level). Thus an application subscript, i , replaces jt in equation (2). Notably, the *Post* indicator now indicates whether the application was made after the ban: This allows the ban to affect application quality via changes in workers' application behavior. For example, the ban might induce workers to send new, well-matched applications to jobs that had previously requested the 'other' gender. Once again, our most tightly controlled estimate is

²⁶ Regression-discontinuity plots corresponding to these estimates, and to the other main estimates in this Section, are provided in Appendix 4.

²⁷ Notably, the index does not incorporate the match between the job's requested gender and the worker's gender. It is based only on the types and amounts of skills required by the job and possessed by the worker.

²⁸ Job day fixed effects are indicators for every day in a job ad's 'life'; their role is to capture the effects of duration dependence in a recruiting spell in a very flexible way.

in column 4, which indicates that the ad ban improved the quality of a typical application by a very small but statistically significant amount (0.0168 standard deviations).²⁹

In Table 5 we shift our attention from firms' to workers' well-being, and use similar methods to estimate whether the ban made it harder for a typical worker to find a new job. Our estimate of the quality of workers' job-finding opportunities is application yield: the probability that an application results in a call-back. In addition to raising the chances a worker will find a new match, higher yields should also help workers find jobs with higher wages and better benefits and amenities, for the simple reason that the worker will have more offers to choose from. As in Table 4, Table 5 presents pooled estimates that combine all three job types (*F*, *N* or *M*). This is because one of the main mechanisms via which the gendered ad ban might reduce call-back rates is by *hiding* employers' gender preferences from workers. If workers try to maximize their call-back chances by avoiding jobs that request the 'other' gender, the ban could reduce call-back chances by making it harder for workers to target their applications in this way.

To see whether this occurred, Table 5 reports the results of local linear regressions similar in structure to the application quality regressions in Table 4: the unit of observation is again an *application*, and the regressor of interest is an indicator for whether the application was submitted before or after the ban. The only differences are the outcome --which equals one if the application *ever* received a call-back from the job to which it was submitted—and the fact that applicant fixed effects replace job ad fixed effects.³⁰ Thus we are effectively comparing the call-back chances of applications submitted by the same worker, before versus after the ban.

When applicant fixed effects are controlled for, Table 5 shows a fairly precisely estimated zero effect of the ban on workers' chances of getting a call-back, per application submitted.³¹ The estimates are almost identical for men, women and both genders combined. Specifically, for women the estimated coefficient shows a decline in the probability of (ever) getting a call-back of 0.34 percentage points, on a baseline call-back rate of 12.4 percent. The 95 percent confidence interval for the effect of the ban runs from -1.12 to 0.44 percentage points. For men, the estimated coefficient shows a decline in the probability of (ever) getting a call-back of 0.35 percentage points on a baseline call-back rate of 9.8 percent. The 95 percent confidence interval for the effect of the ban runs from -0.94 to 0.24 percentage points.

Combining our three indicators of firm and worker well-being, our results in this Section suggest that the ad ban had little or no detectable effects on firms' application arrival rates and workers' application yields, and a small positive effect on mean application quality. While these small or null effects may not be surprising in view of the fact that only one quarter of job ads

²⁹ Again, the RD graph corresponding to this estimate is provided in Appendix 4.

³⁰ The sample sizes differ because whether a call-back was received is observed for all applications; application quality is not always observed.

³¹ Again, Appendix 4 presents the associated RD plots.

were directly treated by the ban, they rule out the important possibility that the ban had large spillover effects on non-treated ads and the workers who applied to them. For example, the removal of gender requests could have signaled to men that all traditionally female jobs (not just those whose requests were removed) were now open to men. We do not see any evidence of such spillovers.

6.2 Effects of the Ban on Directly Treated Jobs, and on the Workers Who Applied to Them

We now turn our attention to the job ads whose gender requests were removed on March 1, 2019. How did this removal affect the number and quality of applications those job ads received? How did these changes affect the call-back chances of the workers who chose to apply to those ads, and were these effects different for men and women? Starting with the ban's effects on employers, Table 6 replicates the regressions in Tables 4 and 5 for the job ads that requested men and women only and shows that the ban led to substantial increases in the number of applications made to those jobs. Specifically, column 2 of panel (a) shows a $.0164/.157 = 10.4$ percent increase in the total number of applications to previously-female jobs, and a $.0093/.147 = 6.3$ percent increase in the total number of applications to previously-male jobs.

Columns 3 and 4 of Table 6's panel (a) disaggregate arrival rates by applicant gender to shed some additional light on the mechanisms underlying the large changes in applicant gender mix at formerly gendered jobs documented in Section 4. Interestingly, the large increase in men's share of applicant pools at formerly-female jobs was not driven solely by a rise in the number of applications coming from men. Instead, after the ban, women applied to formerly-female ads at a slower rate. Specifically, male applications to *F* jobs rose by $.0317$ per day, while women's dropped by $.0153$ per day. The pattern at formerly-male jobs was different, however: while women's applications to these jobs rose, men—unlike women—did not retreat from applying to their 'own' jobs after the ban. One possible explanation for this difference is related to gender differences in ambiguity aversion: when jobs are not clearly gender-stereotyped, Kuhn, Shen and Zhang (2020) found --loosely speaking-- that women tend to apply only when they are explicitly invited, while men *refrain* from applying only when they are explicitly disinvited. When this is the case, removing an *F* label from an ad will deter women from applying, but removing an *M* label will not deter men.³² Panel (b) of Table 6 replicates panel (a) for the mean match quality of applications. On average, the match quality of these applications (on dimensions other than gender) rose, though most estimates are imprecise (and insignificant

³² Gee (2018) and Roussille (2021) present evidence of similar ambiguity-averse behavior in other aspects of the job application process.

in the case of previously-female jobs). A quality increase is consistent with the idea that gender requests discouraged some otherwise highly qualified candidates from applying to jobs.³³

In sum, Table 6 shows that employers whose ads were forcibly modified had more applications to choose from after the ban; if anything, these applications were of higher quality than before. Further, since these additional applications to *F* and *M* jobs did not come at the expense of non-gendered (*N*) jobs, our analysis does not identify any harmful effects of the ban on employers.³⁴

Finally, if the ad ban raised the total number of applications to *M* and *F* jobs, it seems reasonable to ask whether it reduced the per-application call-back rates of the workers who chose to apply to those jobs. Table 7 addresses this issue, showing that the mean call-back rate fell by $.0141/.144 = 9.8$ percent in *F* jobs, and by $.0086/.099 = 8.7$ percent in *M* jobs. Thus, as one might expect, the inflow of applicants to previously gendered jobs *loosened* the micro-labor markets there. While this should benefit employers, it does not imply that the ban harmed workers. This is because—as we have documented—many of the applications to gendered jobs after the ban *would not have been made* before the ban (they were discouraged by employers' requests for the 'other' gender). Because they were never made, these 'missing' pre-ban applications had no chance of getting a call-back.³⁵

In sum, Section 6 shows that in the XMRC labor market *as a whole*, the ad ban had little or no detectable effects on firms' application arrival rates and workers' application yields, and a small positive effect on mean application quality. In the *F* and *M* jobs whose gender requests were removed by the ban, application rates increased substantially (by 14.1 and 8.6 percent respectively), mostly because gender-mismatched workers (who had been discouraged by employers' gender requests) began to apply in greater numbers. If anything, the match quality of these applications rose as well, suggesting an overall gain to employers from the ban. While the increase in applications reduced the chances that an applicant to *F* and *M* jobs received a call-back, this does not imply a reduction in worker well-being because mean call-back rates in the pre-ban period do not account for the fact that some pre-ban applications to *F* and *M* jobs were discouraged by employers' explicit requests for the 'other' gender.

³³ While we estimate a large increase in the quality of men's applications to women's jobs, we emphasize that this is very imprecisely measured, in part due to the very small number of male applications to women's jobs before the ban.

³⁴ The null effect of the ban on applications to *N* jobs is implied by Table 4's aggregate analysis. Direct estimates of the ban's effects on total applicant flows to *N* jobs show a precisely estimated zero effect of $-.0018$ ($.0014$).

³⁵ Columns (2) and (3) of Table 7 would seem to suggest that women bore the brunt of this loosening of labor markets, especially in previously-female jobs (which experienced a large influx of male applicants). This particular result is however not robust to alternative identification approaches (see Appendix 11).

Finally, Section 6 shed some additional light on the ban’s asymmetric effects on gender mix in *F* and *M* jobs. When explicit gender requests were removed from jobs requesting women, men responded with a large increase in applications while women *reduced* their applications. In contrast, when explicit gender requests were removed from jobs requesting men, women responded with a more modest increase in applications, while men *did not stop* applying. Together with the possibility of gender differences in skill specificity (discussed in Section 5.3), these shifts help make sense of why men gained more access to women’s jobs than vice versa.

7. Robustness

Appendices 5-11 explore the robustness of our results to a wide variety of specification changes. In Appendix 5, we assess the sensitivity of our main regression results (Tables 1 and 2) to the use of a quartic in time to control for time trends. We show that the estimated effects of the ban on the gender mix of applicant and call-back pools are unchanged if we replace the quartic by a fifth- or sixth-order polynomial, or by separate quadratics or quartics on either side of the ban. Appendix 6 addresses the fact that some ads contained statements of gender preferences in the text of their job descriptions, in addition to the standardized ‘preferred gender’ field whose removal we study here. As noted, these ‘embedded’ statements in the text were not removed on March 1, 2019. Consistent with expectations, we show that the effects of the ban on the gender mix of call-backs to *F* and *M* jobs (the jobs that had their gender requests removed) is 60 to 70 percent smaller in magnitude when an embedded gender request was present. Conversely, when we look only at the large majority of ads without embedded gender requests, the effect of the ban is 10-13 percent larger than our main estimates. We view the latter result as suggestive of what a gender ban would do if it was applied to the entire text of the job ad.

To address the concern that our call-back indicator is available for only 62 percent of applications, Appendix 7 replicates our regression analyses of *application* behavior (Tables 1, 4 and 5) --which do not require call-back data—on our full sample of ads. The results are nearly identical to the main analysis. Appendix 8 addresses the same concern by taking advantage of the fact that we have a second indicator of application success: whether the application was *read* by the hiring agent. While being read is not as selective as receiving a call-back, our ‘read’ indicator is available for essentially all applications.³⁶ Replicating Table 2 (effects of the ban on the female share of successful applications) using *reads* produces very similar results except for the scale of the coefficients, most likely because reads are more common than call-backs.

³⁶ We define ‘reads’ as being observed for a job *ad* if at least one application was read. By this measure, 99.86 percent of applications have valid read information.

Appendix 9 examines the robustness of our main regression results (the integration of application and call-back pools and zero or small effects on aggregate firm and worker well-being) to the width of the estimation windows used: a full year surrounding the ban in Tables 1 and 2, and a 30-day window in Tables 4 and 5. It shows that our results for the gender mix of application and call-back pools are highly stable to a large variety of application windows, ranging from 9 to 51 weeks in width. For our well-being results, the massive shifts in recruiting activity associated with the Spring Festival prevent us from expanding the estimation window before the ban, so our approach here is to gradually lengthen the post-ban estimation interval from 15 to 40 days (bringing us to the start of another holiday, the Qingming Festival). For match quality and call-back rates, these results are also very stable across all window lengths and very close to our main estimates. For application arrival rates, on the other hand, the estimated effect of the ban first shifts from zero to negative as the estimation window lengthens, then recovers back towards zero in longer windows.

Appendix 10 probes the robustness of the same set of results by estimating the effects of placebo bans in every possible week preceding the actual ban, and comparing these estimates to our main results.³⁷ For our gender mix outcomes, the estimates of the ban's effects on gender mix in *F* and *M* jobs are (a) largest in magnitude on the date of the actual ban, (b) only present during estimation windows that include the actual ban, and (c) approximately zero in estimation windows that exclude the actual ban. We see this is additional confirmation of the ban's causal, integrating effects on call-back pools. For our estimates of aggregate matching frictions, aside from large, spurious estimates that coincide with major holidays and April 2018 XMRC site upgrade, our estimates of the effects of placebo bans are very similar to the small estimated effects of the actual ban.

Finally, to see if our discontinuity-based results are robust to a fundamentally different set of identifying assumptions Appendix 11 exploits the fact that our data run back to the start of 2018.³⁸ Since China's labor market activity is highly synchronized with holidays based on its lunar calendar, we compare the 2019 outcomes with 2018 outcomes on the same lunar dates. Specifically, we construct a full set of week-of-the (lunar) year indicators for 2018 and 2019, which match equivalent weeks in the lunar calendar. We then conduct an augmented difference-in-difference analysis, using 2019 as the treated group and 2018 as the control, with only 2019 receiving treatment on the day of its ban. In addition to the control variables used in our main analysis, we also include a quartic in week of the year (WOY), interacted with year. This 'augmented' *DiD* specification allows the pattern of time trends to differ in a smooth

³⁷ Our analysis cannot be replicated for placebo bans that occur after the actual ban because ads posted on those dates will not contain the *F*, *N* and *M* labels needed to classify jobs. Our sample window includes the 2018 lunar equivalent date to the actual ban however.

³⁸ We also have some data from 2015 and 2016, which we do not use here. Aside from their greater distance in time, the sample sizes are much smaller, and call-back information was especially scarce in those years.

fashion between the two years, while still using the week fixed effects to capture high-frequency variation that is common to the two years.³⁹

Relative to our main estimation approach, advantages of this alternative approach include the fact that the full set of WOY effects relaxes our smoothness assumptions, letting us capture an arbitrary pattern of high-frequency temporal variation that is common across years. This allows us to use weekly data for the entire year for all our outcomes, not just the gender-composition outcomes. Limitations of the alternative approach include a considerably shorter pre-ban period and lower-quality call-back information. (Most of the 2018 data, including the period surrounding the ban, is from before a 2018 site upgrade that improved the company's counts of call-backs.) Most importantly, the alternative approach forces us to rely on a parallel trends assumption. Specifically, net of parameterized *differences* in time trends between the two years, the week fixed effects are assumed to be the same in both years. Because Appendix 11 explores the effects of a fundamentally different source of identification, it replicates all the regression tables in the paper (1, 2, and 4-7).

Appendix 11's estimates of the ban's effects on the gender mix of application and call-back pools are almost identical to our main estimates in Tables 1 and 2.⁴⁰ Once again, we conclude that those estimates are highly robust. Turning to the ban's aggregate effects on firms, Table A11.3 shows similar patterns to those in Table 4: small and statistically insignificant effects of the ad ban on application inflows, and a small but statistically significant positive effect on mean application quality. In contrast to Table 5, however, Table A11.4 shows a decline in aggregate application yields that is statistically significant for men and all workers combined (the effect was negative but insignificant in Table 5). While this suggests that the ban loosened the entire XMRC labor market, we are less confident in this result than our main estimates, primarily because of the DiD method's reliance on parallel trends assumptions. Inspection of the 2018 and 2019 trends for call-back yields shows a large pre-ban surge in 2019 that is absent in 2018. While we attempt to capture this by interacting the effects of the Spring Festival (and following week) with year, it is unclear that this accounts for all the relevant differences between the two years.

Turning to the ban's direct effects on the jobs whose gender requests were removed, Table A11.5 replicates Table 6's analysis of the ban's effects on employers. Aside from the higher magnitudes in panel (a), associated with the switch to weekly application counts, the results are very similar to Table 6: Both tables show an increase in total applications to jobs that previously requested women and men, with the former increase much larger than the latter.

³⁹ We also allow the three main national holidays to have different effects in each year. For call-back-related outcomes, we also add a control for XMRC's April 2018 system upgrade.

⁴⁰ The only difference of potential interest is that the ban's effect on the female share of *applications* to *N* jobs becomes statistically insignificant. The ban's effects on the female share of *call-backs* in *N* jobs is small and insignificant in both our main and DiD specifications.

Both tables attribute the increase in applications to *M* jobs solely to additional applications from women, while the increase in applications to *F* jobs is the net result of a large increase in applications from men and a decline in women's applications. Both Tables also show imprecisely estimated but largely positive effects of the ban on the quality of applications to *M* and *F* jobs. Importantly, all these patterns confirm our understanding of the mechanisms underlying the ban's asymmetric integrating effects on call-back pools: These effects were driven by a large surge in men's applications to *F* jobs, accompanied by a smaller increase in women's applications to *M* jobs and by a decline in women's applications to their 'own' (*F*) jobs.

Table A11.6 replicates Table 7's analysis of call-back yields in *F* and *M* jobs, with somewhat different results. Unlike Table 7, it does not find a disproportionate loosening of labor markets in the previously gendered jobs, with one exception: the call-back rate of women who applied to previously male jobs fell after the ban. While this strengthens the argument that the ban did not adversely affect workers' abilities to find new jobs (at least not disproportionately in directly affected jobs compared to other jobs), we place less weight on this estimate than on our main estimates due to its reliance on parallel trends assumptions.

8. Discussion

For a number of highly plausible reasons, a thoughtful observer could argue that prohibiting explicit gender requests in job ads could prove to be ineffective or counterproductive in integrating jobs and workplaces. For example, little or no change in application behavior might occur; employers could simply ignore any new 'gender-mismatched' applications that are generated; and a ban could restrict employers' good-faith efforts to diversify their workforces. Despite these and other possibilities, our study shows that an unexpected, overnight removal of explicit gender requests on a Chinese job board 'worked', in the sense that it raised the representation of gender-atypical workers in pools of successful applicants. In addition to integrating jobs and workplaces, we find that the ban did not increase aggregate matching frictions: application arrival rates to firms and aggregate application success rates for workers appear to have remained unchanged. Mean application match quality appears to have improved. In the jobs that were directly treated by the ban, application arrival rates and application match quality improved. While application success rates fell in the directly treated jobs, this is a direct consequence of the fact that workers were now applying to jobs they had previously avoided even trying to get.

While the XMRC ad ban integrated labor markets without apparently harming firms or workers, it also illustrates some limits of gendered ad bans as policy tools to reduce gender disparities in labor markets. One such limitation is the striking asymmetry of the ban's effects: men entered jobs that (previously) requested women at a much greater rate than women entered men's. If the goal of the ban is primarily to improve women's access to desirable jobs once reserved for men, this asymmetry may limit its potential. Closely related, --and simply because explicit gender requests are most frequently used in lower-paying types of work-- the

jobs that both women and men entered because of the ban were mostly low-wage positions. Thus, at least in the short run, banning gendered ads is not an effective way to increase, say, women's share in senior management positions. Finally, we also show that the historically most gender-stereotyped job titles on XMRC, such as drivers, electricians, nurses and receptionists—were not meaningfully integrated by the ban.

We also note two important limitations of our study, the first of which is the fact that we do not observe which workers are *hired* after a gendered ad ban—we only see call-backs. In this respect our work is complementary with Card, Colella and Lalive (2021) who study the effect of Austria's 2004 Equal Treatment Act, which also banned explicit gender requests in the hiring process. While Card et al.'s policy discontinuity is considerably fuzzier than ours, their administrative data allow them to show that the Act's integrating effect extends to the hiring stage. A second limitation is the short run nature of our analysis, which does not give workers time to acquire skills and qualifications that would qualify them for many of the jobs whose gender requests were removed. For this reason, our estimates could understate the harder-to-identify longer run integrating effects of a gendered ad ban. In the longer run, a ban might encourage both men and women to invest in skills that are currently highly gendered, such as electricians and nurses.

While our findings in this paper have direct implications for the many countries that still permit gender requests in job ads, they may also be of interest to countries that banned such requests long ago, for two specific reasons. First, recall that one of our most surprising findings is the high receptivity of employers who posted explicitly discriminatory job ads to the gender-mismatched applications that arrived when that discriminatory content was removed. Notably, this receptivity took place in a regulatory context where anti-discrimination regulations affecting other aspects of the recruitment process (such as selection from the applicant pool) were weak, and were not changed at the time of the ban. This suggests that even the *explicit* gender requests we study are in many cases relatively weak employer preferences, which are perhaps driven less by strong tastes or beliefs about productivity than by unconscious stereotypes that associate different jobs and tasks with different genders (Bohren, Imas, and Rosenberg 2018; Carlana 2019). If this is the case, recent attempts to remove subtler (and often unconscious) signals of gender preferences from job ads—such as those advocated by Glassdoor (2021)—may have significant potential to integrate labor markets.

Second, the explicit gender requests we study have an interesting parallel in the new opportunities to micro-target all types of ads that have been created by evolving Internet technologies. These new technologies appear to have allowed some employers to ensure that their job ads are seen only by the age- and gender groups they hope to hire (Angwin et al. 2017; Dwoskin 2018). Interestingly, compared to explicit gender requests, this emerging form of job ad targeting has the potential to cause even more segregation, because—in contrast to the situation we study—the excluded group is typically not even aware of the job's existence. While

such emerging forms of job ad targeting may be harder to regulate than the ads we study, we hope that some of the methods and lessons of our study may prove useful in studying them as well.

References

- Agan, Amanda, Bo Cowgill and Laura Gee 2021 "Salary Disclosure and Hiring: Field Experimental Evidence from a Two-Sided Audit Study". Unpublished paper, Rutgers University.
- Agan, Amanda and Sonja Starr. [Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment](#) (with Sonja Starr), *Quarterly Journal of Economics*, February 2018.
- Angwin, Julia, Noam Scheiber and Ariana Tobindec. 2017. "[Facebook Job Ads Raise Concerns About Age Discrimination.](#)" *New York Times*, December 20, 2017.
- Ballance, Joshua, Robert Clifford and Daniel Shoag "[No more credit score: Employer credit check bans and signal substitution](#)", *Labour Economics* Volume 63, April 2020.
- Banfi, Stefano, and Benjamin Villena-Roldan. 2019. "[Do High-Wage Jobs Attract more Applicants? Directed Search Evidence from the Online Labor Market.](#)" *Journal of Labor Economics*, 37(3), 715-746.
- Baker, Michael and Kirsten Cornelson. 2018. "[Gender Based Occupational Segregation and Sex Differences in Sensory, Motor and Spatial Aptitudes.](#)" *Demography* (2018) 55 (5): 1749–1775.
- Behagel, Luc, Bruno Crepon and Thomas Le Barbanchon. "[Unintended Effects of Anonymous Résumés](#)" *American Economic Journal: Applied Economics* 2015, 7(3): 1–27
- Belot, Michele, Philipp Kircher and Paul Muller. 2018. "[How Wage Announcements Affect Job Search Behaviour - A Field Experimental Investigation.](#)" IZA discussion paper no. 11814
- Bielby, W. T., and J. N. Baron. 1984. "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.
- 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(2), 471-492.
- Bohren, J. Aislinn, Alex Imas, and Michael Rosenberg. 2019. "[The Dynamics of Discrimination: Theory and Evidence.](#)" *American Economic Review*, 109 (10): 3395-3436.
- Bos, Marieke, Emily Breza and Andres Liberman. [The Labor Market Effects of Credit Market Information.](#) *The Review of Financial Studies*, Volume 31, Issue 6, June 2018, Pages 2005- 203.
- Card, David, Fabrizio Colella, and Rafael Lalive (2021) "Does Stated Gender Preference Lead to Workplace Gender Segregation?" unpublished manuscript, UC Berkeley.
- Carlana, Michela. "[Implicit Stereotypes: Evidence from Teachers' Gender Bias](#)" *Quarterly Journal of Economics*, Volume 134, Issue 3, August 2019, Pages 1163–1224.

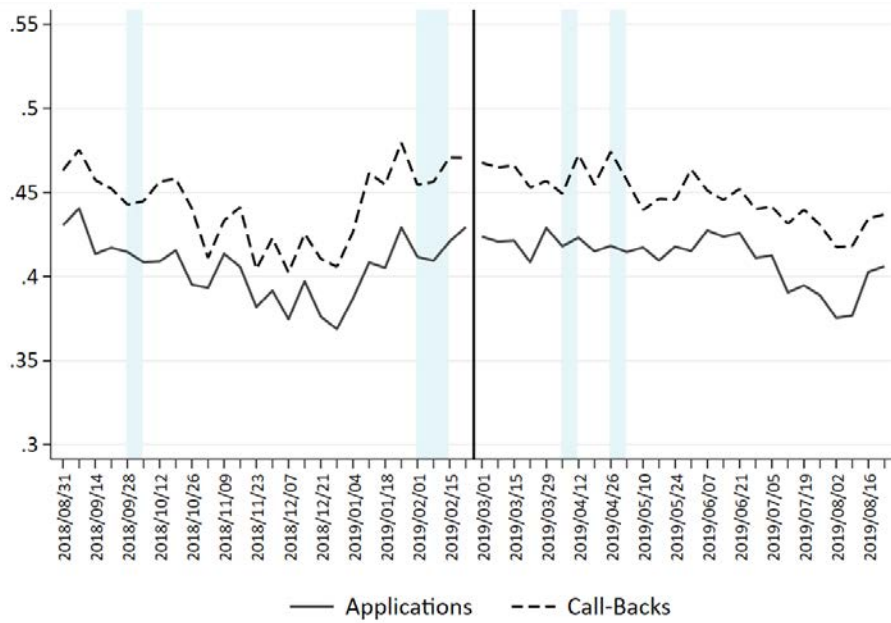
- Cook, Cody, Rebecca Diamond, Jonathan V Hall, John A List, and Paul Oyer. 2021 "[The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers](#)", *The Review of Economic Studies*, (forthcoming).
- Darity, William A. Jr, and Patrick L. Mason, "[Evidence on Discrimination in Employment: Codes of Color, Codes of Gender](#)," *Journal of Economic Perspectives*, 12 (1998), 63–90
- Delgado Hellesester, Miguel, Peter Kuhn and Kailing Shen. 2020. "[The Age Twist in Employers' Gender Requests: Evidence from Four Job Boards](#)", *Journal of Human Resources* vol. 55(2): 428-469
- Doleac, Jennifer, and Benjamin Hansen (2020). [The Unintended Consequences of 'Ban the Box': Statistical discrimination and employment outcomes when criminal histories are hidden.](#) *Journal of Labor Economics* 38(2): 321-374.
- Dwoskin, Elizabeth. "[Men \(only\) at work: Job ads for construction workers and truck drivers on Facebook discriminated on gender, ACLU alleges](#)" *Washington Post*, September 18, 2018.
- Flory, Jeffrey A., Andreas Leibbrandt and John A. List. 2015. "[Do Competitive Workplaces Deter Female Workers? A Large-Scale Natural Field Experiment on Job Entry Decisions.](#)" *The Review of Economic Studies*, 82(1), 122-155.
- Gee, Laura. 2018. "[The More You Know: Information Effects on Job Application Rates in a Large Field Experiment](#)", *Management Science*, 65(5), 2077-2094.
- Glassdoor Team. 2021. "[10 Ways to Remove Gender Bias from Job Descriptions](#)" (April 2).
- Goldin, Claudia and Cecilia Rouse. (2000). "[Orchestrating Impartiality: The Impact of 'Blind' Auditions on Female Musicians](#)" *American Economic Review* 90(4): 715-41
- Ibanez, Marcela and Gerhard Riener. 2018. "[Sorting through Affirmative Action: Three Field Experiments in Colombia.](#)" *Journal of Labor Economics*, 36(2), 437-478.
- Krause, Annabelle, Ulf Rinne and Klaus F Zimmermann. (2012) [Anonymous job applications in Europe.](#) *IZA Journal of European Labor Studies*, 1:5
- Kuhn, Peter and Kailing Shen. 2013. "[Gender Discrimination in Job Ads: Evidence from China.](#)" *Quarterly Journal of Economics*, 128(1), 287-366.
- Kuhn, Peter and Kailing Shen. 2015. "[Do Employers Prefer Migrant Workers? Evidence from a Chinese Job Board.](#)" *IZA Journal of Labor Economics* 4:22.
- Kuhn, Peter, Kailing Shen and Shuo Zhang. 2020. "[Gender-Targeted Job Ads in the Recruitment Process: Evidence from China](#)" *Journal of Development Economics* 147:102531.
- Leibbrandt, A. and John A. List. 2014. "[Do women avoid salary negotiations? Evidence from a large-scale natural field experiment.](#)" *Management Science*, 61(9), 2016-2024.

- Leibbrandt, A. and John. A. List. 2019. "[Do Equal Employment Opportunity Statements Backfire? Evidence from a Natural Field Experiment on Job-Entry Decisions.](#)" unpublished paper, University of Chicago.
- Marinescu, Ioana and Ronald Wolthoff. 2020. "[Opening the Black Box of the Matching Function: the Power of Words](#)", *Journal of Labor Economics* 38(2): 535-568.
- Mas, Alexandre and Amanda Pallais. 2017. "[Valuing Alternative Work Arrangements.](#)" *American Economic Review*, 107(12), 3722-3759.
- Ningrum PK, T. Pansombut and A. Ueranantasun (2020) "[Text mining of online job advertisements to identify direct discrimination during job hunting process: A case study in Indonesia](#)" *PLoS ONE* 15(6): e0233746
- Hausman, Catherine and David S. Rapson (2018) "[Regression Discontinuity in Time: Considerations for Empirical Applications](#)" NBER working paper no. 23602.
- Roussille, Nina. 2021. [The Central Role of The Ask Gap in Gender Pay Inequality](#), unpublished manuscript, UC Berkeley.
- Shimer, Robert. 2004 "[Search Intensity](#)", unpublished paper, University of Chicago.
- Sorensen, Elaine. 1990. "[The Crowding Hypothesis and Comparable Worth.](#)" *Journal of Human Resources*, 25(1), pp. 55-89.
- Walsh, John, Miriam Johnson and Marged Sugarman. 1975. "[Help Wanted: Case Studies of Classified Ads.](#)" Salt Lake City, Utah: Olympus Research Corporation.


Figures

Figure 1: Gender Mix of Applications and Call-Backs by Job Type, Before and After the 2019 Ban


a. Female Share of Applications and Call-Backs to Non-Gendered (N) Jobs:



b. Female Share of Applications and Call-Backs to Male (M) Jobs:



c. Male Share of Applications and Call-Backs to Female (F) Jobs:




Notes to Figure 1:

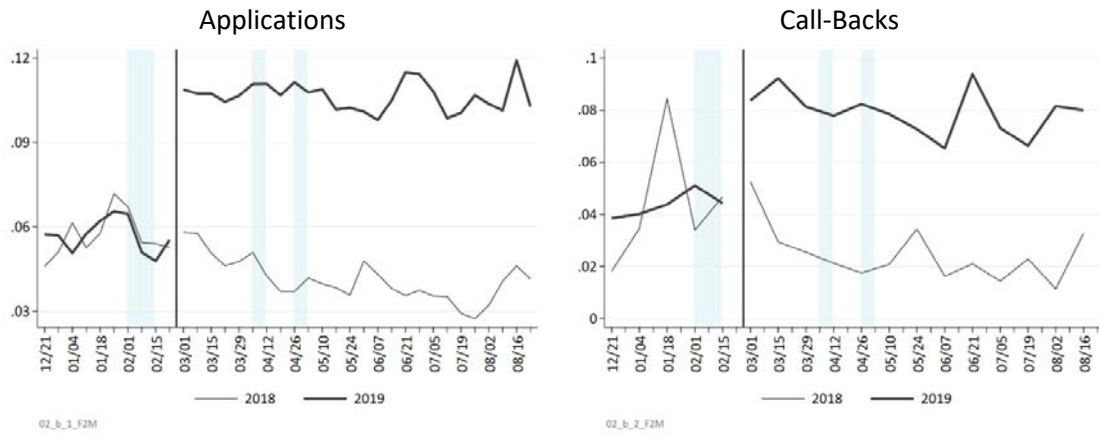
The four shaded areas correspond to the weeks covering four major holiday periods: National Day (October 1-7, 2018), the Spring Festival (February 4-10, 2019), the Qingming Festival (April 5-7) and Labor Day (May 1-4). Holidays have large, negative effects on the total volume of recruiting and application activity, but do not appear to affect the gender mix of applications and call-backs. The solid vertical line indicates the start of the gendered-ad ban, March 1, 2019.

Figure 2: Gender Mix of Applications and Call-backs by Job Type, Before and After the 2019 Ban and the 2018 Pseudo-Ban


a. Female Share of Applications and Call-Backs to Non-Gendered (*N*) Jobs



b. Female Share of Applications and Call-Backs to Male (*M*) Jobs




c. Male Share of Applications and Call-Backs to Female (*F*) Jobs



Notes to Figure 2:

1. Dates on the horizontal axis refer to 2019. Outcomes in 2018 are plotted for the same lunar dates, which are 11 calendar days later.
2. Time series have been adjusted to have the 2019 pre-ban mean in both years.
3. The three shaded areas correspond to the weeks containing three major holidays: the Spring Festival, the Qingming Festival and Labor Day. In 2019, these occurred in February 4-10, April 5-7, and May 1-4 respectively. The dashed vertical line indicates the first day of the gendered-ad ban, March 1, 2019.
4. Prior to XMRC's website upgrade in April 2018, XMRC reported many fewer call-backs than in 2019 (because fewer firms used its internal messaging system). To reduce the resulting noise in the time series, we plot all our call-backs in two-week bins, in contrast to the weekly bins in Figure 1.

Figure 3: Requests for Male and Female Applicants as a Function of Incumbent Gender Mix, by Job Title

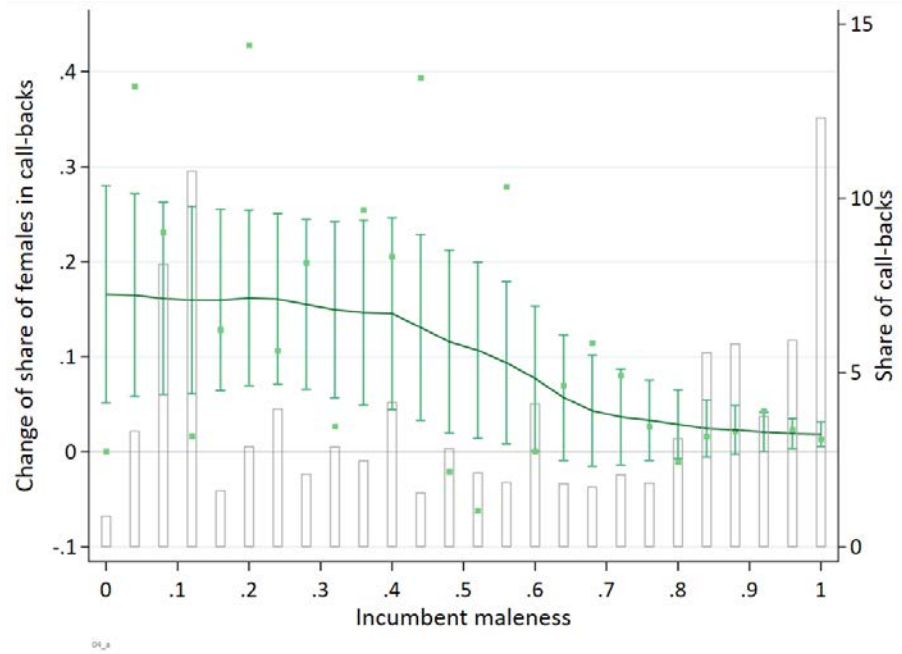


Notes:

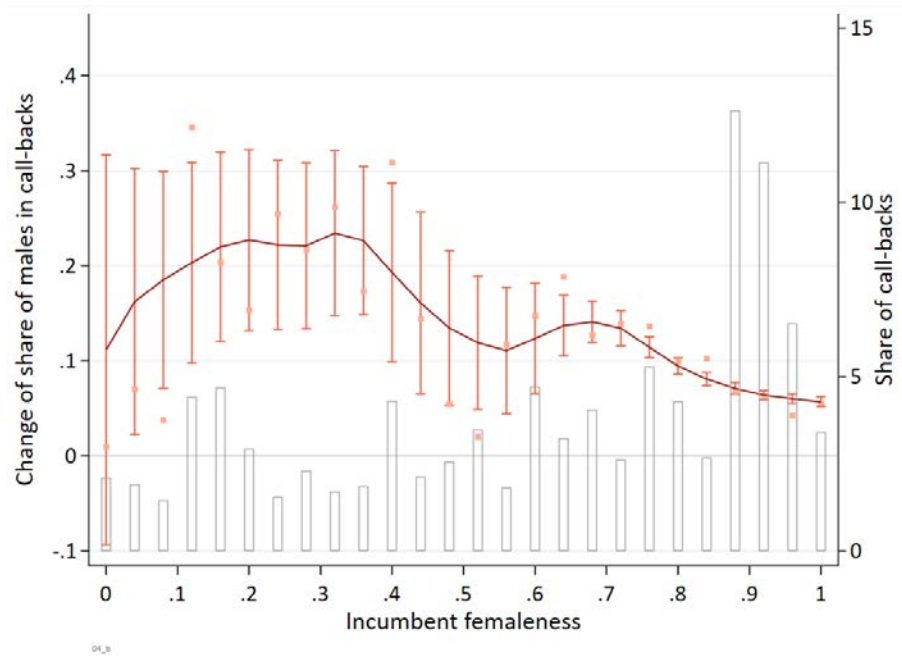
1. Figure shows two binned scatter plots and curves fitted to them. Underlying observations are job titles.
2. Incumbent femaleness (horizontal axis) is the female share of all the call-backs issued by a job title in the *pre-estimation period* (January through August 2018).
3. The square dots in each bin show the share of vacancies requesting men (women) in our *main estimation sample* (ads posted between September 2018 and February 2019). The curves are local polynomial fitted lines of these scatter plots, where the bars are the share of vacancies in each bin. The right axis is for the scatter plots and the curves. The left axis is for the bars.

Figure 4: Women’s and Men’s Increased Access to Job Titles, by Incumbent Gender Mix

a. Job Ads that Initially Requested Men: Pre-Post Ban Changes in Women’s Share of Call-backs by Incumbent Male Share of the Job Title



b. Job Ads that Initially Requested Women: Pre-Post Ban Changes in Men’s Share of Call-backs by the Incumbent Female Share of the Job Title



Notes to Figure 4:

1. The horizontal axis shows the male (female) share of call-backs in the pre-estimation period (between January and August 2018), by job title.
2. The left vertical axis (for the scatter plots and fitted curves) is the change in the female (male) share of call-backs to the relevant job titles (post ban minus pre-ban) in our main estimation sample. To fit the curves, we weight the observations by the number of call-backs they represent in the pre-estimation period to all job types. The vertical axis (for the bars) is the share of call-backs in each incumbent maleness (femaleness) group.
3. In jobs with incumbent maleness of 0.8 or greater, the increase in women's representation was 1.87 percentage points, which was not significantly different from zero ($p = .129$) (from part (a)).
4. In jobs with incumbent femaleness of 0.8 or greater, the increase in women's representation was 6.56 percentage points, which was significantly different from zero. ($p = .000$) (from part (b)).

Figure 5: Time Trends in Application Arrival Rates, Application Quality, and Call-back Rates

a. Weekly Flow of Applications per Active Vacancy, by Job's Gender Request (*F*, *N*, or *M*)



Shaded areas represent major holidays; vertical line is the date of the ban.

b. Average Weekly Application Match Quality, by Job's Gender Request (*F*, *N*, or *M*)




The standardized matching scores are based on raw scores with mean 72.79 and standard deviation 10.11.

S_A_Qual

Shaded areas represent major holidays; vertical line is the date of the ban.

c. Probability an Application Results in a Call-back, by Week Submitted and Applicant Gender



Shaded areas represent major holidays; vertical line is the date of the ban.

Tables

Table 1: Effects of the Gendered Ad Ban on the Female Share of Applications

	(1)	(2)	(3)	(4)	(5)
Post ban week * Female job	-0.1639*** (0.0042)	-0.1631*** (0.0041)	-0.1626*** (0.0041)	-0.1392*** (0.0023)	-0.1388*** (0.0023)
Post ban week * Male job	0.0447*** (0.0033)	0.0464*** (0.0033)	0.0465*** (0.0033)	0.0348*** (0.0018)	0.0350*** (0.0018)
Post ban week	0.0048*** (0.0018)	0.0166*** (0.0025)	-0.0009 (0.0021)	0.0053*** (0.0011)	
Female job	0.4906*** (0.0036)	0.4902*** (0.0036)	0.4898*** (0.0036)		
Male job	-0.3532*** (0.0036)	-0.3522*** (0.0036)	-0.3523*** (0.0036)		
Quartic in job weeks		Y	Y	Y	Y
Quartic in calendar weeks			Y	Y	
Job ad fixed effects				Y	Y
Calendar week fixed effects					Y
Effective # of obs	1,426,896	1,426,896	1,426,896	1,426,896	1,421,850
R^2	0.193	0.194	0.194	0.647	0.646

Notes:

1. All job ads are used here.
2. All regressions are weighted by the number of applications in each job*week cell, and clustered by firm ID. The numbers of observations used for all specifications are 1,426,896. As a result of fixed effects specifications, as singleton observations were dropped, the effective numbers of observations are smaller in column (5).
3. The dependent variable is the share of applications from female applicants; its weighted mean is 80.2%, 41.2%, 8.6% for female, non-gendered, and male jobs, respectively. The estimated constant term of 0.4090 in column 1 gives the raw, pre-ban female share of applications in N jobs. Thus, the raw female share of applications in male jobs pre-ban is $0.4090 - 0.3532 = 5.58\%$. The raw male share of applications in female jobs pre-ban is $1 - (0.4090 + 0.4906) = 10.04\%$.

Table 2: Effects of the Gendered Ad Ban on the Female Share of Call-backs

	(1)	(2)	(3)	(4)	(5)
Post ban week * Female job	-0.1269*** (0.0079)	-0.1263*** (0.0079)	-0.1254*** (0.0079)	-0.1039*** (0.0062)	-0.1037*** (0.0059)
Post ban week * Male job	0.0275*** (0.0064)	0.0282*** (0.0064)	0.0284*** (0.0064)	0.0246*** (0.0048)	0.0247*** (0.0046)
Post ban week	0.0037 (0.0043)	0.0152*** (0.0052)	0.0034 (0.0054)	0.0049 (0.0040)	
Female job	0.4847*** (0.0057)	0.4845*** (0.0057)	0.4837*** (0.0057)		
Male job	-0.3997*** (0.0070)	-0.3987*** (0.0071)	-0.3988*** (0.0071)		
Quartic in job weeks		Y	Y	Y	Y
Quartic in calendar weeks			Y	Y	
Job ad fixed effects				Y	Y
Calendar week fixed effects					Y
Effective # of obs	214,489	214,489	214,489	214,489	196,850
R^2	0.219	0.219	0.220	0.691	0.667

Notes:

1. Only job ads with a positive number of call-backs are used here.
2. All regressions are weighted by the number of call-backs in each job*week cell, and clustered by firm ID. The numbers of observations used for all specifications are 214,489. As a result of fixed effects specifications, singleton observations were dropped and the effective numbers of observations are smaller in column (5).
3. The dependent variable is the share of call-backs from female applicants at each job*week cell; its weighted mean is 85.7%, 45.0% and 6.7% for female, non-gendered, and male jobs, respectively. The estimated constant term of 0.4474 in column 1 gives the raw, pre-ban female share of applications in N jobs. Thus, the raw female share of call-backs in male jobs pre-ban is $0.4474 - 0.3997 = 4.77\%$. The raw male share of call-backs in female jobs pre-ban is $1 - (0.4474 + 0.4847) = 6.79\%$.

Table 3: Characteristics of the Jobs and Workplaces Women and Men Entered Because of the Ban

a. Incumbent Male Share of the Job Titles Women Gained Access to Because of the Ban

(1)	(2)	(3)
All M Job Titles (Opened to Women by the Ban)¹	M Titles Women Entered (were Called Back to) because of the ban²	(1)-(2)
0.878	0.677	0.201

1. The mean incumbent male share of all job titles among M jobs (jobs requesting men).
2. The mean incumbent male share of all job titles, weighted by their contribution to women's increased share in call-backs to M jobs (See Appendix 3).

b. Incumbent Female Share of the F Jobs Men Gained Access to Because of the Ban

(1)	(2)	(3)
All F Job Titles (Opened to Men by the Ban)¹	F jobs Men Entered (were Called Back to) because of the ban²	(1)-(2)
0.785	0.734	.051

1. The mean incumbent female share of all job titles among F jobs (jobs requesting women).
2. The mean incumbent female share of all job titles, weighted by their contribution to men's increased share in call-backs to F jobs (See Appendix 3).

c. Incumbent Gender Mix of the Job Titles, Positions, Workplaces and Firms Entered by Women and Men because of the Ban

	UNIT			
	Job Titles	Positions	Workplaces	Firms
Incumbent Male share of units entered by Women	0.677	0.813	0.610	0.579
Incumbent Female share of units entered by Men	0.734	0.855	0.576	0.540

1. Positions are firm x job title cells; workplaces are firm x job location cells.
2. Incumbent Male and Female Shares are calculated from all the call-backs issued by a unit (Position, Workplace, etc.) during the pre-estimation period only (January - August 2018). The mean incumbent female share of all call-backs on XMRC during the pre-estimation period is 0.501.
3. The incumbent gender mix of the units (e.g. job titles) that were entered is the product of the unit's prevalence (among M or F jobs) before the ban (α), and the change in the incumbent male or female share of call-backs between the pre- and post-ban periods (δ).

d. Mean Log (Wages) of the Job Titles Entered by Women and Men

	All Titles ¹	Titles Entered because of the Ban ²	Difference (2) – (1)
Mean log (wage) in:	(1)	(2)	(3)
Job Ads Requesting Men (<i>M</i> jobs)	8.738	8.621	-0.117
Job Ads Requesting Women (<i>F</i> jobs)	8.634	8.524	-0.110
Jobs Ads with no Gender request (<i>N</i> jobs)	8.858	--	--
All Jobs	8.826	--	--

1. The log (wage) of an individual job title is the number-of-vacancies weighted average of the log of the posted wage for all job ads in the pre-estimation sample for that title. The mean log (wage) for job ads requesting men is the average of the log (wage) of all *M* jobs in the estimation sample. Mean log (wages) for other job types are defined analogously.

2. The mean log (wage) of 'entered' titles is the mean across all job titles, weighted by their contribution to men's (women's) increased share in call-backs to *F* (*M*) jobs (See Appendix 3).

Table 4: Aggregate Effects of the Ad Ban on Employer Well-Being:**Local Linear Regressions**

	(1)	(2)	(3)	(4)
	Arrivals	Arrivals	Quality	Quality
Post ban	0.0011 (0.0013)	0.0017 (0.0013)	0.0218*** (0.0059)	0.0168*** (0.0052)
Day-of-week fixed effects	Y	Y	Y	Y
Quartic in job days	Y	Y	Y	Y
Job ad fixed effects		Y		Y
# of obs	3,494,575	3,494,575	439,173	439,173
R^2	0.021	0.360	0.005	0.610

Notes:

1. Observations are ad-day cells in columns 1 and 2, and applications in columns 3 and 4. Sample is restricted to days within 15 days of the ad ban.
2. All specifications include a linear trend, interacted with Post Ban. Days from the ban (t) is defined as the date of the application minus the date of the generated ad ban. Job days are the number of days elapsed since the job received its first application. All regressions are clustered by firm ID.
3. Application arrivals are the number of applications received per vacancy in each job*calendar day cell. The average daily number of applications received is .143.
4. Match quality is normalized to a mean of zero and standard deviation of one among all applications in our data. The average normalized match quality for the applications in the current estimation sample is .008.

Table 5: Aggregate Effects of the Ad Ban on Jobseekers:
Local Linear Regressions for Call-Back Yields

	(1)	(2)	(3)	(4)	(5)	(6)
	All		Applications from: Women		Men	
Post ban	-0.0060*** (0.0022)	-0.0035 (0.0026)	-0.0054 (0.0033)	-0.0034 (0.0040)	-0.0056** (0.0025)	-0.0035 (0.0030)
Day-of-week fixed effects	Y	Y	Y	Y	Y	Y
Quartic in job days	Y	Y	Y	Y	Y	Y
Applicant fixed effects		Y		Y		Y
Effective # of obs	637,711	637,711	270,994	270,994	366,717	366,717
R^2	0.001	0.133	0.001	0.130	0.001	0.132

Notes:

1. Observations are applications that were made within 15 days of the ban. 'Post ban' means the *application* was made after the ban. The dependent variable equals one if the application *ever* received a call-back.
2. All specifications include a linear trend and its interaction with the Post Ban dummy.
3. All regressions are clustered by firm ID.
4. The average call-back probability is 0.124 and 0.098 for applications from women and men, respectively, and 0.109 overall.

**Table 6: Effects of the Ad Ban on Employers who Posted Gendered Job Ads:
Local Linear Regressions**

a. Effects on the Daily Application Arrival Rate:

	(1)	(2)	(3)	(4)
	All	All	Women	Men
Applications from:				
<i>To Ads that Requested Women:</i>				
Post ban	0.0151*** (0.0037)	0.0164*** (0.0037)	-0.0153*** (0.0032)	0.0317*** (0.0016)
Job Ad Fixed Effects?		Y	Y	Y
# of obs	405,249	405,249	405,249	405,249
R ²	0.023	0.319	0.307	0.205
<i>To Ads that Requested Men:</i>				
Post ban	0.0088** (0.0038)	0.0093** (0.0038)	0.0102*** (0.0010)	-0.0009 (0.0036)
Job Ad Fixed Effects?		Y	Y	Y
# of obs	439,775	439,775	439,775	439,775
R ²	0.013	0.411	0.214	0.415

b. Effects on the Mean Match Quality of Applications:

	(1)	(2)	(3)	(4)
	All	All	Women	Men
Applications from:				
<i>To Ads that Requested Women:</i>				
Post ban	0.0073 (0.0163)	0.0185 (0.0153)	0.0264 (0.0165)	0.1036* (0.0560)
Job Ad Fixed Effects?		Y	Y	Y
# of obs	51,672	51,672	44,029	11,106
R ²	0.007	0.551	0.565	0.709
<i>To Ads that Requested Men:</i>				
Post ban	0.0294* (0.0153)	0.0260* (0.0141)	0.0236 (0.0658)	0.0301** (0.0149)
Job Ad Fixed Effects?		Y	Y	Y
# of obs	56,293	56,293	5,619	52,455
R ²	0.003	0.574	0.752	0.578

Notes:

1. Sample is restricted to days within 15 days of the ad ban.
2. All specifications include a linear trend, interacted with Post Ban. Days from the ban (t) is defined as the date of the application minus the date of the gendered ad ban. All regressions include day-of-week fixed effects and a quartic in job days (the number of days elapsed since the job received its first application). All regressions are clustered by firm ID.
3. In panel (a), observations are ad-day cells, and the dependent variable is the number of applications received per vacancy at each job*calendar day cell. Column 3 (4) only counts applications from women (men). The average daily number of applications received is 0.159 and 0.149 in ads requesting women and men respectively.
4. In panel (b), observations are applications, and the dependent variable is the application's match quality, normalized to a mean of zero and standard deviation of one among all applications in our data. The average normalized match quality for the applications in the current estimation sample is -0.050 and 0.089 in ads requesting women and men respectively.

Table 7: Effects of the Ad Ban on Workers who Applied to Gendered Job Ads:
Local Linear Regressions for the Call-Back Rate per Application

	Applications from:		
	All	Women	Men
	(1)	(2)	(3)
Post ban	-0.0008 (0.0027)	0.0003 (0.0042)	-0.0017 (0.0031)
Post ban * Female job	-0.0141*** (0.0045)	-0.0112** (0.0051)	-0.0076 (0.0085)
Post ban * Male job	-0.0086** (0.0040)	-0.0172* (0.0102)	-0.0048 (0.0043)
Day-of-week fixed effects	Y	Y	Y
Quartic in job days	Y	Y	Y
Applicant fixed effects	Y	Y	Y
Effective # of obs	637,711	270,994	366,717
R^2	0.133	0.131	0.132

- Notes:
1. Observations are applications that were made within 15 days of the ban. 'Post ban' means the *application* was made after the ban. The dependent variable equals one if the application *ever* received a call-back.
 2. All specifications include a linear trend, and its interaction with the Post Ban dummy.
 3. All regressions are clustered by firm ID. All columns include controls for Female Job and Male Jobs.
 4. The average call-back probabilities in female and male jobs were 0.144 and 0.099 respectively.

What Happens When Employers Can No Longer Discriminate in Job Ads?

Peter Kuhn and Kailing Shen

Appendix—For online publication

July 24, 2021

Appendix 1—Gendered Ad Bans in Selected Countries¹

1.1 The Decline of Gendered Job Ads in United States, Austria, and China

Information on legislation and advertising practices in these three countries is available from other research articles. The dates refer to the period during which explicit gender requests in job ads declined dramatically, which in most cases includes the date of related legislation.

United States (1968-1973)

A 1973 U.S. Supreme Court decision (*Pittsburgh Press Co. v Pittsburgh Commission on Human Relations*) declared that sex-segregated advertising was illegal, and that newspapers could not publish ads seeking applicants of a specific gender. Available evidence suggests, however, that explicit gender requests were already being abandoned by newspapers before that date, perhaps due to public pressure. In particular, Walsh et al.'s (1975) study of help wanted ads in San Francisco and Salt Lake City showed that such ads declined dramatically between 1968 and 1972. Still, 15 percent of job ads continued to designate sex in San Francisco in 1972, and 33 percent did so in Salt Lake City.

Austria (2004-2008)

In the early 2000s, the European Union published a number of directives that prohibited discrimination based on gender, ethnicity, religion, or disability in all EU member countries. These directives resulted in anti-discriminatory changes to labor laws in a number of countries, which included prohibitions of explicitly gendered job ads. In Austria's case, these laws took the form of the Austrian Equal Treatment Act (AETA), which became effective on July 1, 2004. Card, Colella and Lalive (2021) demonstrate that stated gender preferences declined precipitously during the three subsequent years, and were essentially absent by 2008.

China (2016-2019)

Appendices 2 and 3 in Kuhn, Shen and Zhang (2020) provide information on legislation and implementation of gendered ad bans in China. Essentially, China stipulated that job boards could be fined for posting job ads containing explicit gender requests in 2016. As in some other countries (such as Austria and Ireland, for example) enforcement was not immediate. While the large, national job boards eliminated most gendered ads quite early, XMRC was not ordered to do so until March 2019.

¹ This appendix was prepared with the energetic and capable assistance of Alice Liu and Billy Troutman.

1.2 The Decline of Gendered Job Ads-- Other English-Speaking Countries

To provide additional background for the current paper, we collected information on gendered job ads and associated legislation in these countries using a two-step process:

1. We identified one or two major newspapers in the country were that were active in 1960 and for several decades after. We then examined the classified job ads starting in 1960 to identify a period during which the prevalence of explicit gender requests declined dramatically.
2. We then searched news articles and other sources around that period to identify changes in public policy –especially anti-discrimination legislation-- that may have accounted for this sharp decline.

Our analysis was limited to English-speaking countries due to language barriers.

Unless otherwise stated, the source for most of our newspapers analysis is newspapers.com, which is a database of newspapers in the English language, some dating a few centuries back. Once again, the dates refer to the period during which explicit gender requests in job ads declined dramatically, which in most cases includes the date of related legislation.

Canada (1971-1978)

Since Canadian labor law is primarily under provincial jurisdiction, our newspaper analysis focused on the *Financial Post* and *Vancouver Sun*, prominent newspapers in Canada's two largest English-speaking provinces: Ontario and British Columbia.

Despite the passage of the Ontario Human Rights Code in 1962, job ads in the *Financial Post* contained large numbers of explicit gender requests between 1960 and 1970. In 1971, there was a significant decline in the prevalence of job ads containing gender requests; these were completely gone by February 5th, 1972. This was ten years after Ontario's Human Rights Code was passed, and five years before the Canadian Human Rights Act (affecting federal workers) was passed in 1977.

Between 1960 and 1973, job ads in the *Vancouver Sun* were separated into men's and women's sections. This practice ended in 1974, although explicit gender requests still persisted within individual ads. In 1977, editors of the newspaper began issuing a warning regarding discrimination in employment ads and telling applicants to disregard explicit requests and treat them as if they are requesting both genders. Explicit gender requests were completely absent from the April 7th, 1978 issue.

The 1974 disappearance of separate men's and women's sections could be related to the passage of the BC Human Rights Code in 1973; the 1977 change may be related to the Canadian Human Rights Act, though this Act only affected workers under federal jurisdiction.

United Kingdom (1972-1975)

Our newspaper analysis for the U.K. was based on job ads in *The Guardian*. Explicit gender requests were common in those ads between 1960 and 1974. A steep decline in such ads was observed by April 1975, and no examples of gender requests were found in April 1976.

This drop in discriminatory ads was likely due to the passage of the UK's *Sex Discrimination Act of 1975*, which was introduced to prevent discrimination on the basis of gender. The primary stated goal of the Act was to bring about gender equality.

New Zealand (1977-1980)

Our analysis of job ads in New Zealand used two newspapers, *The Press* and the *Manawatu Evening Standard*.² From 1965 to 1975, both papers' employment sections appear were saturated with gendered job ads. *The Press* exhibited many fewer of these ads by 1977, and the *Evening Standard* by 1978. By 1980, gendered ads were rare in *The Press* and practically nonexistent in the *Evening Standard*.

The disappearance of New Zealand's gendered job ads appears to be driven by New Zealand's Human Rights Commission Act of 1977, which created a commission to receive and mediate complaints of discrimination on bases that included marital status, sex, religion, and ethical beliefs. Gendered ads had already disappeared when New Zealand's 1993 Human Rights Act was passed.

Ireland (1977-1981)

Our analysis of job ads in Ireland used two newspapers, *The Irish Independent* and *The Evening Herald* (*The Herald* after 2013).³ Between 1960 and 1976, both papers' job ads were saturated with explicit requests for males and females.

By July of 1977, both papers had notices at the top of their employment sections regarding the *Employment Equality Act*, a 1977 law which made it illegal for employers to discriminate based on gender or marital status. The notices told readers and employers that all jobs under the jurisdiction of the law ("domestics wanted" ads were an exception) were open to both males and females, regardless of whether they expressed preferences for a specific gender. This did not appear to have any immediate effect, as both papers were still filled with gender explicit requests.

By June 1978, the volume of explicit gender requests had diminished substantially, and the wording of ads had become more gender neutral. For example, ads for "bar men" and "bar maids" were replaced by ads for "bar person". By 1982, we did not encounter a single explicit gender request in either newspaper.

In all, it appears that the job advertising changes in Ireland were driven by the 1977 *Employment Equality Act*, but full compliance took roughly 4 to 5 years afterwards.

Australia (1983-1985)

Our newspaper analysis was based on *The Sydney Morning Herald*. Between 1960 and 1983, its job ads were separated into "Men and Boys" & "Women and Girls" sections. In the June 4th, 1983 issue, employment ads were no longer separated by gender. Keywords like "man" and "woman" were also largely absent from the job ads after that date. Newspapers from 1984 and 1985 showed a similar pattern.

The timing of this change coincides with Australia's *Sex Discrimination Act of 1984*. This law made it illegal to discriminate against people because of their sex, intersex status, gender identity, sexual orientation, marital or relationship status, pregnancy or potential pregnancy status, and family responsibilities. The law also protects against sexual harassment, and covers most areas of public life, not just employment.

² Access was via scanned copies provided by librarians in New Zealand.

³ Access was via Irish Newspaper Archives, an online database of Irish newspapers.

Works Cited in Appendix 1

- Delgado Helleseeter, Miguel, Peter Kuhn and Kailing Shen. 2020. "The Age Twist in Employers' Gender Requests: Evidence from Four Job Boards", *Journal of Human Resources* vol. 55(2): 428-469
- Kuhn, Peter, Kailing Shen and Shuo Zhang. 2020. "[Gender-Targeted Job Ads in the Recruitment Process: Evidence from China](#)" *Journal of Development Economics* 147:102531.
- Pittsburgh Press Co. vs Pittsburgh Commission on Human Relations et al., U.S. Supreme Court decision argued March 20, 1973, decided on June 21, 1973.

Appendix 2: Sample Design and Descriptive Statistics

To construct our main analysis sample, we started with the universe of applications that were made on XMRC between January 1, 2018 and October 25, 2019, and the corresponding ads.

We then retain only:

- Job *ads* that received at least one application both before and after the ad ban (March 1, 2019)
- *Applications* to those ads that were made between August 31 2018 and August 29, 2019. This gives us a sample of 52 complete weeks (26 before and 26 after the ban), in which the first post-ban week begins on Friday March 1 2019—the first day of the ban.

In sum, we start with a wide window (almost two years) to make sure we capture all the job ads that ‘straddled’ the ban. We then retained only the ads that actually straddled the ban (received at least one application before and after it). Finally, our analysis sample comprises all applications to those ads that occurred during a one-year window surrounding the ban.

Table A2.1 shows the mean characteristics of the 166,725 *job ads* in our main analysis sample, and the 61,960 job ads for which we observe call-back information (our *call-back sample*). In both cases, the means are presented separately for jobs requesting women, jobs requesting men, and jobs with no explicit gender request (*F*, *M*, and *N* jobs respectively). Table A1.1 documents some noteworthy and perhaps surprising differences between these three types of jobs. One of these is the fact that both *M* and *F* jobs advertise lower wages than *N* jobs: At 5,794 yuan per month, *M* jobs pay 6.1 percent less than *N* jobs; at 4,432 yuan per month, *F* jobs pay 28.1 percent less. This appears to be a common feature of gender requests on a number of different job boards (Delgado Helleseer et al., 2020), and is consistent with the idea that firms abstain from gender typing in the most skilled jobs because it is more important to identify the single best candidate for those jobs (Kuhn and Shen, 2013).

Table A2.1 also shows that jobs requesting men and women have different education and experience requirements: jobs requesting men ask for more experience but less education. In part, this is because ads for women are highly focused on young, single new labor market entrants (who tend to be highly educated in China (Delgado et al. 2020)).

Table A2.2 shows the mean characteristics of the 3.1 million applications in our main analysis sample, and the 1.9 million applications for which we observe call-back information (our *call-back sample*). In both cases, the means are presented separately for applications made by women and by men. Part (a) of this table shows the characteristics of the applicant. Part (b) shows the characteristics of the match between the applicant’s characteristics and job’s requirements (for example, do the applicant’s age and experience fall into the ranges requested in the ad?).

Finally, to allow us to compare changes on XMRC around the ban with changes in 2018, we also create a *DiD* sample. To do so, we replicated our main estimation sample—which comprises applications that were made between September 2018 and August 2019—on two different periods: January – August 2018 and January–August 2019. The latter period contains the date on which the 2019 ban occurred, and the former contains the date on which it would have occurred in 2018. Unfortunately, these two periods cannot be designed to exactly mimic our main analysis sample because we have no XMRC data from 2017. While this restricts the length of the pre-ban period in both years to just two months, it allows us to compare trends in our main outcomes between 2018 and 2019 on both sides of the ban date. We use this sample in Appendix 11 to conduct a

difference-in-difference analysis of the ban's effect –which uses equivalent days or weeks from 2018 as controls for 2019-- as a robustness check of our main results. Notably, since important events affecting China's labor market –especially the Spring Festival-- are determined by the lunar calendar, this new DiD sample requires us to line up days and weeks between 2018 and 2019 so that they represent the same days and weeks relative to the start of the Spring Festival in both years.

Table A2.1: Sample Means (Job Postings)

Job Ad Characteristic	Full sample				Call-back sample			
	Female	None	Male	All	Female	None	Male	All
Education requirement (years)	13.04	13.59	11.81	13.30	13.08	13.40	11.60	13.13
Require technical school	0.214	0.103	0.156	0.123	0.218	0.117	0.171	0.137
Experience requirement (years)	0.832	1.281	1.616	1.271	0.828	1.156	1.518	1.157
Require new graduates?	0.029	0.020	0.020	0.021	0.030	0.021	0.018	0.021
Age requirement?	0.662	0.442	0.684	0.498	0.684	0.484	0.710	0.540
Age required (mean)	28.11	29.72	31.50	29.78	28.09	29.43	31.30	29.51
Explicit offered wages?	0.839	0.744	0.765	0.757	0.858	0.797	0.797	0.805
Wage offered (Yuan)	4,432	6,168	5,794	5,897	4,457	5,816	5,628	5,596
Bonus offered?	0.167	0.211	0.104	0.193	0.175	0.253	0.109	0.224
Bonus offered (Yuan)	6,701	8,443	7,729	8,219	6,704	8,409	7,658	8,182
Commission offered?	0.173	0.225	0.112	0.205	0.181	0.264	0.117	0.234
Explicit vacancy numbers?	0.969	0.949	0.956	0.952	0.971	0.952	0.960	0.956
Number of vacancies	2.026	2.548	2.323	2.458	2.139	2.735	2.368	2.607
Firm size stated?	0.998	0.999	0.998	0.998	0.999	0.999	0.998	0.999
Number of workers in the firm	347	800	607	724	307	521	451	483
Firm ownership								
State Owned Enterprises	0.018	0.041	0.046	0.039	0.012	0.021	0.028	0.021
Stock market listed	0.023	0.058	0.044	0.052	0.020	0.038	0.038	0.035
Private shared	0.662	0.611	0.575	0.612	0.675	0.661	0.605	0.656
Private	0.134	0.129	0.116	0.128	0.141	0.138	0.113	0.135
Taiwan/HK/Macau	0.061	0.045	0.080	0.051	0.059	0.046	0.087	0.053
FDI from US or European	0.012	0.015	0.021	0.015	0.011	0.009	0.018	0.011
Other FDI	0.017	0.017	0.025	0.018	0.016	0.017	0.027	0.018
Joint venture with US or other Joint venture	0.007	0.009	0.012	0.009	0.007	0.007	0.013	0.008
Foreign Representative	0.011	0.010	0.010	0.010	0.009	0.007	0.008	0.008
NGO	0.003	0.003	0.002	0.003	0.002	0.003	0.002	0.002
Missing	0.035	0.037	0.032	0.036	0.031	0.033	0.027	0.032
Average # of applications received per job ad	40.17	40.40	45.29	40.99	48.24	47.12	53.84	48.12
# of job postings	13,596	88,518	14,611	116,725	8,416	45,780	7,764	61,960
% of job postings	11.65	75.83	12.52	100.00	13.58	73.89	12.53	100

Table A2.2: Sample Means (Applications)

a. Applicants' Characteristics

	Full sample			Call-back sample		
	Females	Males	All	Females	Males	All
Age	29.20	31.56	30.59	29.13	31.38	30.41
Married?	0.448	0.487	0.471	0.456	0.476	0.468
Education	14.98	14.45	14.67	14.86	14.29	14.54
Technical school?	0.081	0.117	0.102	0.091	0.131	0.114
Experience (years)	8.06	10.30	9.38	8.09	10.22	9.30
New graduates?	0.001	0.000	0.000	0.001	0.000	0.000
Local <i>hukou</i>	0.355	0.269	0.305	0.343	0.253	0.292
Fujian <i>hukou</i>	0.817	0.748	0.776	0.806	0.727	0.761
Short sighted?	0.344	0.297	0.317	0.337	0.285	0.308
Valid height?	0.905	0.960	0.937	0.907	0.961	0.937
Height	160.7	171.9	167.5	160.6	171.8	167.1
Any photos?	0.436	0.395	0.412	0.423	0.379	0.398
# of photos	1.002	1.003	1.002	1.002	1.003	1.003
Valid current wage?	0.750	0.782	0.769	0.758	0.789	0.775
Current wage	5,657	7,150	6,549	5,444	6,741	6,192
Valid intended wage?	0.652	0.622	0.635	0.659	0.625	0.640
Intended wage	5,922	7,565	6,869	5,687	7,100	6,470
Current status?						
Employed, want to stay	0.078	0.066	0.071	0.081	0.069	0.074
Employed, will move if offered better	0.026	0.033	0.030	0.024	0.031	0.028
Employed, want to move	0.182	0.204	0.195	0.166	0.180	0.174
Unemployed	0.714	0.696	0.703	0.729	0.720	0.724
Chinese resume						
Complete	0.545	0.489	0.512	0.525	0.464	0.490
Mostly complete	0.451	0.506	0.483	0.472	0.531	0.505
Incomplete	0.004	0.005	0.004	0.004	0.005	0.004
English resume						
Complete	0.050	0.040	0.044	0.046	0.038	0.041
Mostly complete	0.016	0.016	0.016	0.015	0.015	0.015
Incomplete	0.934	0.944	0.940	0.939	0.947	0.944
Education info complete?	0.998	0.998	0.998	0.998	0.998	0.998
Experience info complete?	0.969	0.973	0.972	0.970	0.974	0.972

b. Characteristics of the Match between the Applicant and the Job

	Full sample			Call-back sample		
	Females	Males	All	Females	Males	All
Current wage match offered?						
Related data missing	0.384	0.390	0.387	0.344	0.343	0.344
Current wage lower	0.202	0.224	0.215	0.211	0.238	0.226
Current wage similar	0.278	0.239	0.255	0.300	0.261	0.278
Current wage higher	0.136	0.147	0.143	0.144	0.159	0.152
Intended wage match						
Related data missing	0.463	0.514	0.493	0.429	0.479	0.457
Intended wage lower	0.152	0.160	0.157	0.160	0.172	0.167
Intended wage similar	0.248	0.193	0.216	0.266	0.209	0.233
Intended wage higher	0.138	0.133	0.135	0.146	0.141	0.143
Age match required?						
Younger than required	0.036	0.033	0.034	0.038	0.035	0.036
Age proper	0.903	0.888	0.894	0.895	0.880	0.887
Older than required	0.061	0.079	0.071	0.067	0.085	0.077
Education match required?						
Less educated than required	0.126	0.135	0.131	0.122	0.126	0.124
Education proper	0.525	0.523	0.524	0.509	0.509	0.509
More educated than	0.350	0.342	0.345	0.370	0.365	0.367
Technical school educated and required	0.022	0.026	0.024	0.026	0.029	0.028
Experience match required?						
Less experienced than required	0.037	0.032	0.034	0.034	0.029	0.031
Experience proper	0.311	0.237	0.268	0.305	0.233	0.264
More experience than required	0.652	0.732	0.699	0.662	0.738	0.705
Gender match preferred?						
Gender mismatch	0.029	0.039	0.035	0.028	0.048	0.039
To F jobs	0.225	0.039	0.116	0.257	0.048	0.138
To N jobs	0.746	0.747	0.746	0.715	0.725	0.721
To M jobs	0.029	0.214	0.138	0.028	0.227	0.141
# of applications	1,290,206	1,840,111	3,130,317	839,101	1,100,834	1,939,935
% of applications	41.22	58.78	100	43.25	56.75	100

Note: Each panel of the table reports the extent to which applicant characteristic X matches the ad’s request for characteristic X. For example, in the full sample 90.3 percent of applications from women were in the age range requested by the job ad. 3.6 percent of female applicants were younger than requested and 6.1 percent were older than requested.

Appendix 3: Which Job Titles Account for the Ban’s Main Effects?

Our main empirical findings in this paper are that (a) removing explicit requests for men from job ads raises women’s share in call-backs to those ads; and (b) removing explicit requests for women raises men’s call-back share by a substantially larger amount. In Sections A3.1 and A3.2 of this Appendix we use a simple shift-share approach to identify which specific job titles account for most of these effects. To put these ‘impacted’ job titles in context, Section A3.3 uses data from before our estimation sample to identify the historically *most* male- and female-dominated job titles on XMRC, and discusses why these extremely male and female job titles were not integrated by the ban. Section A3.4 describes how we modify our shift-share analysis to calculate the wages associated with these impacted titles. Section A3.5 explores the extent to which title-specific skills might help account for the asymmetry in the ban’s effects on men versus women.

A3.1 Women’s Increased Access to *M* jobs—Which job titles contributed most?

In this Section our goal is to identify the exact job titles that accounted for women’s increased representation in *M* jobs. To illustrate our approach, consider all the call-backs to job ads that requested men in the pre-ban period (*M* ads), and categorize them by their job titles, T . Let α^T be the share of call-backs with title T among call-backs to *M* ads; thus α^T measures the *prevalence* of title T among jobs requesting men. Let δ^T be the *change* in the female share of call-backs in each of these job titles (within the *M* jobs) between the pre- and post-ban periods. Then the overall change in the female share of call-backs to *M* jobs is given by $\delta = \sum_T \alpha^T \delta^T$, and title T ’s contribution to this change is just $\alpha^T \delta^T$. Thus, a title’s contribution depends both on its prevalence among *M* jobs, and on the increase in women’s representation it experienced.

Table A3.1 lists the twenty job titles that made the largest contributions to the increased representation of women in jobs that previously requested men. Together, these 20 titles accounted for 1.434 percentage points (just under half) of the 3.038 percentage point increase in the share of women called back to *M* job ads in our data.⁴ To measure of how *stereotypically* male these job titles were, column 4 takes advantage of the fact that we have access to call-back data from an eight-month period that precedes our estimation sample (January - August 2018), and uses this data to estimate every job title’s *incumbent* gender mix.⁵ Under this definition, notice that job ads that request men can occur in job titles that are stereotypically male *or* female, and that, say, a stereotypically female job title could be predominantly male in some individual firms.

According to column 4 of Table A3.1, only 13 of the twenty titles that contributed most to women’s increased representation in *M* jobs had an incumbent male share that was above the overall share of call-backs that went to men (.499). The most important of these were three warehouse management titles plus “general

⁴ Because of sampling differences and regression controls, this 3.038 percentage point estimate in these Appendix calculations differs slightly from our main, regression- adjusted estimate of $2.46 + 0.49 = 3.05$ percentage points (Table 2, column 4). We use raw means to perform the decompositions in this Appendix to maximize simplicity and transparency.

⁵ Our estimates of a job title’s historical gender mix combine all three job types (*F*, *N*, and *M* together). Calculating incumbent gender mix from pre-estimation sample data ensures that pre-post ban changes in the gender mix of call-backs in our estimation sample are not affected by any dynamic statistical processes, such as mean reversion.

labor". In these job titles, the ad ban reduced gender-stereotyping because it brought more women into job titles in which women were historically underrepresented. Somewhat more surprisingly, the ad ban also had a positive effect on women's representation in some job titles where women were modestly or highly over-represented: manager's assistants and procurement specialists, respectively. Within these more-female titles, the ad ban raised women's success rate *in the subset of job ads that had previously requested men*. This change works to *increase* gender stereotyping at the job title level (making stereotypically female titles more female), but its effects must be weighed against the effect of removing requests for women in these female titles.

Taken together, all the 20 job titles listed in Table A3.1 had a (contribution-weighted) incumbent male share of .660, which is higher than the pre-ban mean male share of call-backs across all ads in our data (.499), but not 'extremely' male. Finally, going beyond the top 20 contributing titles, the contribution-weighted mean incumbent male share of all the job titles that account for the women's increased representation in jobs that previously requested men was very similar at 0.677. Thus, the gendered ad ban increased women's access to *types of work* (job titles) that historically were more male than average, but not dramatically so.

Table A3.1: Job Titles Accounting for Women’s Increased Representation in *M* jobs

	Job Title (7)	Translation	Contribution to women’s increased call-back share in <i>M</i> jobs ($\alpha^T \delta^T$)¹ (percentage points)	Incumbent male share²
	(1)	(2)	(3)	(4)
T1	仓管员	Warehouse management staff	0.251	0.852*
T2	经理助理	Manager assistant	0.201	0.356
T3	仓管理	Warehouse management staff	0.129	0.874*
T4	交易员	Trader	0.083	0.967*
T5	采购专员	Procurement specialist	0.078	0.395
T6	普工	General labor	0.072	0.853*
T7	电商仓库仓管 配货员打包员	E-commerce warehouse management staff, picker, packer	0.058	0.966*
T8	市场专员	Marketing specialist	0.055	0.684*
T9	市政作业员	Civil construction worker	0.051	1.000*
T10	仓库管理员	Warehouse management staff	0.046	0.861*
T11	行政主管	Administrative chief	0.045	0.302
T12	前端开发	Client-side developer	0.043	0.830*
T13	工程助理	Engineer assistant	0.043	0.635*
T14	行政专员	Administration specialist	0.043	0.134
T15	行政助理	Administration assistant	0.043	0.078
T16	审计税审专员	Auditing specialist (taxation)	0.043	0.237
T17	财务经理	Financial manager	0.042	0.488
T18	平面设计	Graphic design	0.038	0.559*
T19	活动执行	Event implementation	0.037	0.630*
T20	Java 工程师	Java engineer	0.036	0.920*
Top 20 titles, combined (T1 to T20) ³			1.434	0.660*
All contributing job titles (TOTAL) ³			3.038	0.677*
Contribution of the top twenty titles [(T1 to T20)/TOTAL]			47.2%	

* Incumbent male share exceeds the sample mean (.499).

Notes to Table A3.1:

1. The *prevalence* of each title, α^T is given by the title's share in male-requesting ads in the pre-ban portion of our main estimation sample. The increase in women's representation in each title, δ^T is the change in women's share of call-backs from male-requesting ads between the pre- and post-ban periods.
2. The incumbent male shares are calculated from our pre-estimation sample only (January – August 2018). Incumbent male shares are men's share in all call-backs to each job title (regardless of the gender requested) during that period
3. Incumbent male shares for the top 20 and all contributing titles are weighted by their relative contributions ($\alpha^T \delta^T$).

A3.2 Men’s Increased Access to *F* jobs—which job titles contributed most?

Turning now to the much larger effects of removing requests for women on the gender mix of call-backs, Table A3.2 lists the twenty job titles making the largest contribution to the increased representation of men in female-requesting jobs. Together, these 20 titles accounted for 3.572 percentage points (or about 37 percent) of the 9.608 percentage point increase in the share of men called back to formerly *F* job ads in our data.⁶ In addition, Table A3.2 shows that all these titles except one –warehouse management staff-- had an incumbent female share that was above the mean female share of .501. Thus, considerably more than was the case for women, the ban opened up a long list of highly female job titles to men.

Taken together, the 20 job titles listed in Table A3.2 had a (contribution-weighted) incumbent female share of .736, which is considerably higher than the mean female call-back share of .501. Finally, going beyond these 20 most important titles, the contribution-weighted mean incumbent female share of the job titles that account for the men’s increased representation in *F* jobs was almost identical at .734.

⁶ As was the case for *M* jobs, the 9.608 percentage point impact used in these Appendix calculations differs slightly from our main, regression- adjusted estimate of $(-0.1039 + 0.49 =)$ 9.9 percentage points (from Table 2, column 4).

Table A3.2: Job Titles Accounting for Men's Increased Representation in *F* jobs

	Job Title (<i>T</i>)	Translation	Contribution to men's increased call-back share in <i>F</i> jobs ($\alpha^T \delta^T$) ¹ (percentage points)	Incumbent female share ²
	(1)	(2)	(3)	(4)
T1	经理助理	Manager assistant	0.414	0.644*
T2	仓管员	Warehouse management staff	0.290	0.148
T3	行政专员	Administration specialist	0.266	0.866*
T4	采购员	Procurement officer	0.242	0.588*
T5	会计	Accountant	0.241	0.887*
T6	文员	Office clerk	0.216	0.940*
T7	财务会计	Financial accountant	0.197	0.889*
T8	人事行政专员	HR administrative specialist	0.174	0.922*
T9	外贸业务员	International trade sales staff	0.173	0.753*
T10	主办会计	Chief accountant	0.144	0.865*
T11	幼儿园出纳	Kindergarten cashier	0.142	0.957*
T12	运营助理	Operation assistant	0.141	0.409
T13	商务助理	Business assistant	0.139	0.822*
T14	人事专员	HR assistant	0.129	0.912*
T15	客服专员	Customer service specialist	0.119	0.788*
T16	销售助理	Sales assistant	0.115	0.746*
T17	客服	Customer service	0.110	0.677*
T18	采购	Procurement officer	0.109	0.591*
T19	仓管文员	Warehouse management clerk	0.107	0.925*
T20	财务助理	Financial assistant	0.104	0.916*
Top 20 titles, combined (T1 to T20) ³			3.572	0.736*
All contributing job titles (TOTAL) ³			9.608	0.734*
Contribution of the top 20 titles [(T1 to T20)/TOTAL]			37.2%	

* Incumbent female share exceeds sample average (.501).

Notes to Table A3.2:

1. The *prevalence* of each title, α^T is given by the title's share in female-requesting ads in the pre-ban portion of our main estimation sample. The increase in men's representation in each title, δ^T is the change in men's share of call-backs from female-requesting ads between the pre- and post-ban periods.
2. The incumbent female shares are calculated from our pre-estimation sample only (January – August 2018). Incumbent female shares are women's share in all call-backs to each job title (regardless of the gender requested) during that period.
3. Incumbent female shares for the top 20 and all contributing titles are weighted by their relative contributions ($\alpha^T \delta^T$).

A3.3 Which Job Titles were the Most Male- and Female Dominated before the Ban?

As a point of reference for the job titles identified in Sections A3.1 and A3.2 (where the ad ban had its largest effects), this Section again draws on historical call-back data to identify the job titles that were the *most* male- and female-dominated on XMRC prior to our estimation period. Turning first to male-dominated job titles, Appendix Table A3.3a lists the 20 largest job titles that *never* called back a woman during our pre-estimation period. Thus for example, none of the 330 call-backs issued to the title “driver for the manager” went to women. Related, Table A3.3b lists the 20 most-male job titles that called back at least one woman. To illustrate, 99.66 percent of the 1480 call-backs to the title ‘driver’ went to men. Notably, only one of the 40 job titles that appear in Table A3.3 (civil construction worker) appears in Table A3.1, which shows the 20 titles accounting for most of the increase in women’s representation in jobs that requested men. We conclude that the gendered ad ban did not integrate the most-male jobs on XMRC. Tables A3.4a and A3.4b perform the same exercise for extremely female jobs, with patterns that are only slightly less extreme. None of the 40 job titles listed here appear in Table A3.2—the job titles accounting for most of men’s increased representation in *F* jobs due to the ban. Thus, the gendered ad ban did not integrate the most-female jobs on XMRC either.

Closer inspection of Table A3.3’s list of intensely-male job titles reveals that two broader types of work seem to dominate here. One is a long list of trades (plumber, welder, electrician) and production jobs (master mold fitter, fitter, CNC operator) that would seem to require occupation- or industry-specific training. The other—even more dominant—is driving. For example, *none* of the $330 + 286 + 143 + 135 = 914$ call-backs to the following job titles went to women: driver for the manager, driver for chairperson, driver for president, and sales driver. In addition, by far the largest job title among extremely male titles is simply “driver”. Thus, a combination of occupation or industry-specific training requirements and a lack of female interest in driving jobs could explain why women didn’t access these extremely male jobs.⁷

Inspection of Table A3.4’s list of intensely-female jobs titles also reveals one dominant type of work: by far the largest single title is “administration and receptionist” and the words administration, reception, and clerk appear in a substantial number of titles. While none of Table A3.4’s 40 extremely female jobs appear in Table A3.2 (the 20 titles accounting for most of men’s increased representation in *F* jobs) the following Section shows that men did make substantial inroads into other job titles that were more than 80 percent female. This— together with the notion that many office tasks are not obviously specific to an occupation or industry—may help explain why the ban raised men’s access to women’s jobs considerably more than the converse.

⁷ Cook et al. (2021) suggest that a desire to be located closer to home and higher concerns for (or risks to) personal safety make driving for Uber less attractive to women than for men. Other factors could be the expectation of hostile reception by incumbent male workers, or very strong gender norms against having women in driving jobs.

Table A3.3: Job Titles with the Highest Historical Male Call-Back Shares

a) 100% of call-backs went to men				
	Job Title (T)	Translation	Number of Call-Backs Issued	Share of Call-Backs Issued to Men
1	经理司机	Driver for the manager	330	1
2	董事长司机	Driver for chair person	286	1
3	注塑技术员	Injection molding technician	228	1
4	装配钳工	Assembly fitter	181	1
5	生产课长	Production section chief	156	1
6	总裁司机	Driver for president	143	1
7	业务司机	Sales driver	135	1
8	CNC 编程	Programmer for computer numerical control machine	125	1
9	土建施工员	Civil construction worker	119	1
10	机修电工	Mechanic electrician	113	1
11	车间主管	Workshop supervisor	97	1
12	网管	Network management staff	96	1
13	机修	Machine repair	92	1
14	仓储课长	Chief of warehouse	78	1
15	CNC 操机员	Operator of computer numerical control machine	75	1
16	CNC 操机	CNC operator	75	1
17	水电工	Plumber	75	1
18	工程监理	construction project supervisor	75	1
19	装修施工员	Decoration builder	72	1
20	注塑主管	Injection supervisor	72	1

Note: Statistics refer to all the call-backs that occurred during our pre-estimation sample (January – August 2018). Panel (a) lists the twenty largest job titles (by call-backs) that sent all their call-backs to men.

Table A3.3 (continued)

b) At least one woman received a call-back				
	Job Title (T)	Translation	Number of Call-Backs Issued	Share of Call-Backs Issued to Men
1	行政司机	Driver for administration tasks	469	.9979
2	司机	Driver	1480	.9966
3	注塑领班	Injection molding team leader	170	.9941
4	网络管理员	Network management staff	169	.9941
5	货车司机	Truck driver	334	.994
6	生产厂长	production CEO	164	.9939
7	焊工	Welder	146	.9932
8	电工	Electrician	270	.9926
9	设备工程师	Equipment Engineer	267	.9925
10	模具钳工	Mold fitter	126	.9921
11	生产经理	Production manager	222	.991
12	送货司机	Delivery driver	393	.9898
13	保安	Security	362	.989
14	钳工	Fitter	87	.9885
15	CNC 操作员	CNC operator	77	.987
16	模具钳工师傅	Master mold fitter	74	.9865
17	生产部经理	production sector manager	142	.9859
18	模具设计	Mold design	67	.9851
19	设备技术员	Equipment technician	66	.9848
20	机械研发工程师	Mechanical R&D engineer	63	.9841

Note: Statistics refer to all the call-backs that occurred during our pre-estimation sample (January – August 2018). Panel (b) lists the twenty job titles with the highest male shares that were strictly less than 100 percent.

Table A3.4: Job Titles with the Highest Historical Female Call-Back Shares

a) 100% of call-backs went to women				
	Job Title (T)	Translation	Number of Call-Backs Issued	Share of Call-Backs Issued to Females
1	护士	Nurse	61	1
2	人事前台	HR & receptionist	57	1
3	前台行政助理	Front desk, administrative assistant	40	1
4	幼教老师	Childcare teacher	40	1
5	打字文员	Typing clerk	39	1
6	行政前台人事助理	Admin, front desk, HR assistant	37	1
7	人事行政出纳	HR, admin and cashier	35	1
8	部门文员	Department Clerk	31	1
9	销售行政助理	Sales admin assistant	31	1
10	财务助理会计助理出纳员	Financial assistant, accounting assistant, cashier	28	1
11	数据整理文员	Data management clerk	26	1
12	财务会计助理	Financial accountant assistant	24	1
13	保育员	Childcare teacher	24	1
14	人资行政专员	HR payroll admin specialist	24	1
15	出纳文员	Cashier clerk	24	1
16	行政专员前台	Admin specialist and front desk	23	1
17	美妆在线销售客服	Online sales & customer service for make-up products	23	1
18	国际展览业务员客服方向	International conference business customer service	23	1
19	跟单采购助理	Procurement assistant	23	1
20	现场咨询师	Onsite consultant	22	1

Note: Statistics refer to all the call-backs that occurred during our pre-estimation sample (January – August 2018). Panel (a) lists the twenty largest job titles (by call-backs) that sent all their call-backs to women.

Table A3.4 (continued)

b) At least one call-back went to men				
	Job Title (T)	Translation	Number of Call-Backs Issued	Share of Call-Backs Issued to Females
1	前台文员	Receptionist and clerk	397	.9950
2	出纳人事	Cashier & HR	91	.9890
3	行政前台文员	Administration, receptionist, and clerk	85	.9882
4	财务行政	Financial admin	68	.9853
5	人事行政文员	HR, administration clerk	123	.9837
6	行政前台	Administration and receptionist	1031	.9825
7	母婴在线销售客服	Mom-Baby online sales customer service specialist	48	.9792
8	生产统计	Production statistic	46	.9783
9	平面模特	Still model	44	.9773
10	小学英语教师	English teacher for primary school	86	.9767
11	生产统计文员	Production statistic clerk	42	.9762
12	业务助理业务跟单	Business assistant, merchandiser	82	.9756
13	财务人事	Finance & HR	39	.9744
14	助理文员	Assistant clerk	78	.9744
15	财务会计商务	Financial accountant, business	38	.9737
16	会计员	Accountant	75	.9733
17	内勤	Internal clerk	73	.9726
18	办公室文员打单人员文秘	Office clerk, documentary handler, secretary	36	.9722
19	外贸单证	International trade document clerk	71	.9718
20	客服代表	Customer representative	35	.9714

Note: Statistics refer to all the call-backs that occurred during our pre-estimation sample (January – August 2018). Panel (b) lists the twenty job titles with the highest female shares that were strictly less than 100 percent.

A3.4 Calculating Mean Posted Wages of the Job Titles Women and Men Accessed because of the Ban

To characterize the wage levels of the job titles workers gained access to, we first calculate the mean log (advertised) wage of every job title using data from the pre-estimation sample. This is done separately for F and M jobs (w^{FT} and w^{MT}). Letting, for example, α^{MT} be the share of male-requesting ads in the *estimation* sample with title T we can then express the (historical) mean wages of *all* M jobs in the estimation sample as $\sum_T \alpha^{MT} w^{MT}$. To calculate the wages of the M jobs that women entered after the ban, we normalize each title's *contribution* to women's increased representation in M jobs ($\alpha^{MT} \delta^{MT}$ from Appendix 3.1) by the total contribution of all M titles ($\sum_T \alpha^{MT} \delta^{MT}$) to obtain a new set of weights, $\theta^{MT} = \frac{\alpha^{MT} \delta^{MT}}{\sum_T \alpha^{MT} \delta^{MT}}$. The mean wage of all the job titles women entered because of the ban is then given by $\sum_T \theta^{MT} w^{MT}$. A parallel procedure yields the mean log wage of the F titles that men accessed as a result of the ban.

A3.5 Concentration of Job Titles— Can Skill Specificity Help Explain Men’s Greater Access to Historically Female Job Titles?

In Section A3.3 we suggested that skills or certifications that are specific to an industry or occupation might help explain why women made fewer inroads into men’s jobs after the gendered ad ban. In this Section we explore this hypothesis in one additional way, by using job titles themselves as indicators of distinct skills or skill bundles that are needed to perform a job well. If titles represent required skill bundles that differentiate jobs from one another, then the number of distinct titles among the most gendered jobs on XMRC would provide a rough indicator of the extent to which specific skills or certifications are needed to enter those jobs.

To assess this hypothesis (and inspired by Figures 4.a and 4.b), we defined *highly male* jobs as ads that requested men during our sample period, in *titles* with a historical male call-back share of at least 80 percent. We did the same for *highly female* jobs, then listed the 20 most common highly-male and highly-female job titles among the vacancies in our estimation sample. The results are presented in Tables A3.5 and A3.6. Overall, we find considerable support for the skill-specificity hypothesis: 23, 34 and 50 percent of the vacancies requesting women were in the 5, 10 and 20 most common job titles, compared to 14, 23 and 33 percent of the vacancies requesting men.

Table A3.5 Ads Requesting Men with High Incumbent Male Shares: Titles with the Most Vacancies

	Job Title (7)	Translation	Number of Vacancies	Share of Vacancies	Cumulative Share of Vacancies
1	保安	Security	544	3.72	3.72
2	普工	General labor	468	3.20	6.92
3	仓管员	Warehouse management staff	447	3.06	9.98
4	施工员	Construction worker	343	2.35	12.33
5	焊工	Welder	302	2.07	14.39
6	电工	Electrician	288	1.97	16.36
7	仓管理	Warehouse management staff	283	1.94	18.30
8	司机	Driver	273	1.87	20.16
9	项目经理	Project manager	234	1.60	21.76
10	装配钳工	Assembly fitter	209	1.43	23.19
11	保安员	Security	200	1.37	24.56
12	电气工程师	Electronic engineer	192	1.31	25.88
13	技术员	Technician	168	1.15	27.02
14	销售工程师	Sales engineer	162	1.11	28.13
15	机械工程师	Mechanic engineer	143	0.98	29.11
16	注塑技术员	Injection molding technician	126	0.86	29.97
17	注塑领班	Injection molding team leader	125	0.85	30.83
18	操作工	Operator	124	0.85	31.68
19	土建施工员	Civil construction worker	107	0.73	32.41
20	模具钳工	Mold fitter	106	0.73	33.13

Notes:

1. Sample is job ads that requested men in the estimation sample period, with job titles that with an incumbent male share of at least 80 percent (calculated from the pre-estimation sample). There are altogether 14,620 vacancies here.
2. Calculations reflect the fact that job ads may have multiple vacancies.

Table A3.6 Ads Requesting Women with High Incumbent Female Shares: Titles with the Most Vacancies

	Job Title (7)	Translation	Number of Vacancies	Share of Vacancies	Cumulative Share of Vacancies
1	文员	Office clerk	542	6.12	6.12
2	会计	Accountant	451	5.09	11.21
3	出纳	Cashier	416	4.70	15.91
4	行政前台	Administration and receptionist	346	3.91	19.81
5	人事专员	HR assistant	279	3.15	22.96
6	人事行政专员	HR administrative specialist	248	2.80	25.76
7	财务	Finance	203	2.29	28.05
8	行政助理	Administration assistant	183	2.07	30.12
9	行政专员	Administration specialist	173	1.95	32.07
10	主办会计	Chief accountant	157	1.77	33.85
11	行政文员	Administrative clerk	155	1.75	35.59
12	人事行政	HR administrative	150	1.69	37.29
13	前台文员	Receptionist and clerk	150	1.69	38.98
14	财务助理	Financial assistant	144	1.63	40.61
15	护士	Nurse	141	1.59	42.20
16	外贸业务助理	International trader	137	1.55	43.75
17	商务助理	Business assistant	132	1.49	45.24
18	收银员	Cashier	129	1.46	46.69
19	前台	Front desk	129	1.46	48.15
20	育婴师	Baby sitter	121	1.37	49.51

Notes:

1. Sample is job ads that requested women in the estimation sample period, with job titles that with an incumbent female share of at least 80 percent (calculated from the pre-estimation sample). There are altogether 8,858 vacancies here.
2. Calculations reflect the fact that job ads may have multiple vacancies.


Reference for Appendix 3:

Cook, Cody, Rebecca Diamond, Jonathan V Hall, John A List, and Paul Oyer. 2021 “The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers”, *The Review of Economic Studies*, (forthcoming).

Appendix 4: Local Linear Regression Plots

The figures in this Appendix graph the main results of the local linear regression analyses of aggregate matching frictions in Tables 4 and 5. Specifically, they display the regressions estimated in columns 2 and 4 of Table 4 and columns 4 and 6 of Table 5. These regressions are our most tightly controlled estimates of the effect of the ban on the size and quality of firms' applicant pools, and on workers' call-back chances, pooling all three job types (*F*, *N*, and *M*) together. The results are discussed in the text of the paper.

Figure A4.1: Daily Flow of Applications per Active Vacancy, Local Linear Regressions



Notes:

1. The vertical line separates the pre and post gendered-ad ban periods, that is just before March 1, 2019.
2. Time series are adjusted for day-of-week and ad fixed effects, plus a quartic in the age of the job ad.
3. Figure is based on the local linear regression coefficients in column 2 of Table 4, where the post-ban effect is 0.0017 (0.0013).

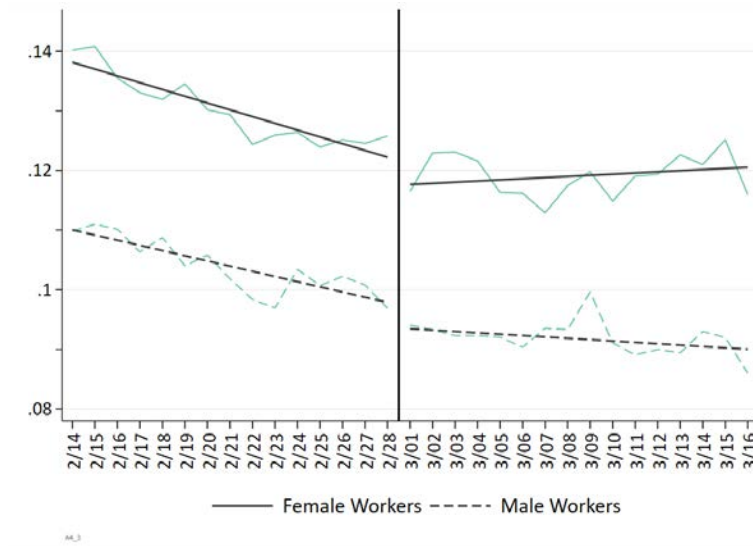
Figure A4.2: Daily Average Match Quality Scores, Local Linear Regressions



Notes:

1. Sample includes all applications to ads that have a matching score (98% of applications).
2. The vertical line separates the pre and post gendered-ad ban periods, that is just before March 1, 2019.
3. Time series are adjusted for day-of-week and ad fixed effects, plus a quartic in the age of the job ad.
4. Local Linear Regression coefficients correspond to column 4 of Table 4, where the post-ban coefficient is 0.0168*** (0.0052).

Figure A4.3: Call-Back Chances per Application by Day of Submission, Local Linear Regressions



Notes:

1. The vertical line separates the pre and post gendered-ad ban periods, that is just before March 1, 2019.
2. Time series are adjusted for day-of-week and worker fixed effects, plus a quartic in age of the job ad.
3. Figure is based on the local linear regression coefficients in columns 4 and 6 of Table 5, which are:

For female workers: Post-ban effect = -0.0034 (0.0040)

For male workers: Post-ban effect = -0.0035(0.0030).

Appendix 5: Using Other Polynomials to Absorb the Time Trend

Our main analysis of the ban’s effects on the gender mix of applications and call-backs (Tables 1 and 2) used a single quartic in weeks to model the time trends in these outcomes during the one-year period surrounding the ban. In this Appendix we explore the sensitivity of those results to other functional forms. Table A5.1 focuses on the female share of *applications*, and replicates column 4 of Table 1 (reproduced in column 1 here) four different ways. Columns 2 and 3 replace the quartic by a fifth- and sixth order polynomials respectively. Column 4 estimates separate *quadratics* on either side of the ban, column 5 estimates separate *quartics* on either side. Table A5.2 repeats this entire analysis for the female share of *call-backs*, rather than applications. In all cases, the ban’s estimated effects on the gender mix of applications and call-backs to previously gendered jobs (*F* and *M*) are essentially unchanged. The ban’s small, positive effect on the female share of applications to *N* jobs (in Table 1) becomes a small negative effect in some specifications. The ban’s (null) effect on the gender mix of *call-backs* to *N* jobs is, however, confirmed in all cases.

Table A5.1: Effects of the Gendered Ad Ban on the Female Share of Applications

	(1)	(2)	(3)	(4)	(5)
	Quartic in calendar weeks	5 th -order polynomial in calendar weeks	6 th -order polynomial in calendar weeks	Separate quadratics before and after ban	Separate quartics before and after ban
Post ban week * Female job	-0.1392*** (0.0023)	-0.1390*** (0.0023)	-0.1390*** (0.0023)	-0.1390*** (0.0023)	-0.1389*** (0.0023)
Post ban week * Male job	0.0348*** (0.0018)	0.0349*** (0.0018)	0.0349*** (0.0018)	0.0349*** (0.0018)	0.0350*** (0.0018)
Post ban week	0.0053*** (0.0011)	-0.0026** (0.0012)	-0.0029** (0.0012)	0.0000 (0.0012)	-0.0050** (0.0020)
Quartic in job weeks	Y	Y	Y	Y	Y
Job ad fixed effects	Y	Y	Y	Y	Y
Effective # of obs	1,426,896	1,426,896	1,426,896	1,426,896	1,426,896
R^2	0.647	0.647	0.647	0.647	0.647

Note: Column 1 reproduces column 4 of Table 1. The remaining columns change only the modeling of the calendar time trend, as described.

Table A5.2: Effects of the Gendered Ad Ban on the Female Share of Call-Backs

	(1)	(2)	(3)	(4)	(5)
	Quartic in calendar weeks	5 th -order polynomial in calendar weeks	6 th -order polynomial in calendar weeks	Separate quadratics before and after ban	Separate quartics before and after ban
Post ban week * Female job	-0.1039*** (0.0062)	-0.1039*** (0.0062)	-0.1037*** (0.0061)	-0.1038*** (0.0061)	-0.1037*** (0.0061)
Post ban week * Male job	0.0246*** (0.0048)	0.0246*** (0.0048)	0.0247*** (0.0048)	0.0247*** (0.0048)	0.0248*** (0.0048)
Post ban week	0.0049 (0.0040)	0.0032 (0.0043)	-0.0014 (0.0043)	0.0002 (0.0042)	-0.0035 (0.0071)
Quartic in job weeks	Y	Y	Y	Y	Y
Job ad fixed effects	Y	Y	Y	Y	Y
Effective # of obs	214,489	214,489	214,489	214,489	214,489
R^2	0.691	0.691	0.691	0.691	0.691

Note: Column 1 reproduces column 4 of Table 2. The remaining columns change only the modeling of the calendar time trend, as described.

Appendix 6: Comparing Ads with and without ‘Embedded’ Gender Requests

In this Appendix, we use the fact that some of the ads in our data have gender requests that are ‘embedded’ in the text of their job descriptions. Specifically, among the 116,725 job ads, 3,529 or 3.023 percent had an embedded gender request. These embedded requests were evenly split between requests for men and women (1.505 percent and 1.518 percent respectively.) Unlike the gender preference fields whose removal we study in this paper, these embedded requests were not removed by XMRC’s employees overnight on March 1, 2019.⁸ Thus, ads that contained embedded requests essentially received a smaller dose of treatment, since some but not all of their explicit gender preference statements were removed by the ban. If so, then our main estimates in the paper may understate the effects of longer-term effects of the ban, since XMRC was directed to remove all forms of gender requests from ads that were newly posted after March 1, 2019.

To test this idea, Tables A6.1 and A6.2 add two interaction terms to Tables 1 and 2: Post ban x Female job x Embedded, and Post ban x Male job x Embedded, where “Embedded” means the ad’s text also contained an explicit gender request. In this specification, the Post ban x Female and Post ban x Male coefficients show the effect of the ban in ads that had no embedded gender requests. These are fully treated ads (because all the gender requesting information was removed), which provide a better estimate of the longer-run effects of the ban. The new interaction terms tell us how much this effect is attenuated when only *some* of the gender-requesting information is removed. To simplify the discussion of the results, we focus on the ban’s effects on *M* and *F* jobs *relative* to *N* jobs, using the most saturated specification of both Tables (column 5).

According to Table A6.1, the ban raised the female share of applications to *M* jobs that did *not* contain any embedded gender requests by 3.84 percentage points. Consistent with the idea that these jobs were more intensely treated, this effect is about ten percent larger than the ban’s effects in all *M* job ads combined (3.50 percentage points, from Table 1). More dramatically, the ban’s effect on the female share of applications to *M* jobs that contained an embedded gender request (that was not removed by the ban) was only $3.84 - 2.29 = 1.55$ percentage points. As Table A6.1 indicates, this 2.29 percentage point (or 60 percent) reduction in the ban’s causal effect is highly statistically significant.

Table A6.1 also indicates that the ban reduced the female share of applications to *F* jobs that did *not* contain any embedded gender requests by 15.06 percentage points. Consistent with the idea that these jobs were more intensely treated, this effect is about nine percent larger than the ban’s effects in *F* job ads combined (13.88 percentage points, also from Table 1). More dramatically, the ban’s effect on the female share of applications to *F* jobs that contained an embedded gender request (that was not removed by the ban) was a reduction of only $15.06 - 8.09 = 6.97$ percentage points. As Table A6.1 indicates, this 8.09 percentage point (or 54 percent) reduction in the magnitude of the ban’s effect is highly statistically significant.

Table A6.2 replicates Table A6.2 for call-backs instead of applications, showing very similar patterns. Compared to our main estimates (that combine ads with and without embedded requests) the ban’s effects on ads without embedded requests are 11 to 13 percent greater in magnitude. In *F* jobs, the ban’s effects on the

⁸ XMRC was, however, directed to inspect the text of new ads posted after March 1, and to request the removal of any discriminatory content. By construction, these new ads do not appear in our estimation sample.

female share of call-backs is about 61 percent smaller when an embedded request is present in the ad. In *M* jobs, the magnitude of the ban's effect is reduced by 74 percent when an embedded request is present, compared to when one is not present.

Taken together, the results in this Appendix strongly confirm the hypothesis that job seekers respond to the content of job ads when deciding where to send their applications: The gender mix of the workers who apply to an ad changes much more dramatically when *all* the gender-relevant information is reduced from the ad, compared to when only some of that information is removed. In addition, the results suggest that a more complete ban on gendered content in job ads would have an integrating effect that is between 9 and 13 percent greater than the partial ban we study here.

Table A6.1: Effects of the Gendered Ad Ban on the Female Share of Applications, Accounting for Embedded Gender Requests

	(1)	(2)	(3)	(4)	(5)
Post ban * Female job *Embedded	0.0973*** (0.0082)	0.0962*** (0.0081)	0.0968*** (0.0081)	0.0808*** (0.0051)	0.0809*** (0.0051)
Post ban * Male job *Embedded	-0.0271*** (0.0076)	-0.0298*** (0.0077)	-0.0289*** (0.0077)	-0.0232*** (0.0044)	-0.0229*** (0.0044)
Post ban week * Female job	-0.1750*** (0.0044)	-0.1741*** (0.0044)	-0.1736*** (0.0044)	-0.1510*** (0.0025)	-0.1506*** (0.0025)
Post ban week * Male job	0.0483*** (0.0034)	0.0503*** (0.0035)	0.0504*** (0.0035)	0.0382*** (0.0019)	0.0384*** (0.0019)
Post ban week	0.0048*** (0.0018)	0.0167*** (0.0025)	-0.0010 (0.0021)	0.0053*** (0.0011)	
Female job	0.4864*** (0.0037)	0.4858*** (0.0037)	0.4854*** (0.0037)		
Male job	-0.3509*** (0.0036)	-0.3500*** (0.0036)	-0.3499*** (0.0036)		
Quartic in job weeks		Y	Y	Y	Y
Quartic in calendar weeks			Y	Y	
Job ad fixed effects				Y	Y
Calendar week fixed effects					Y
Effective # of obs	1,426,896	1,426,896	1,426,896	1,426,896	1,421,850
R^2	0.194	0.195	0.195	0.647	0.646

Notes:

1. All regressions also include controls for Embedded, Embedded * Male Job, and Embedded * Female Job.
2. All regressions are weighted by the number of applications in each job*week cell, and clustered by firm ID. The numbers of observations used for all specifications are 1,426,896. As a result of fixed effects specifications, as singleton observations were dropped, the effective numbers of observations are smaller in column (5).
3. The dependent variable is the share of applications from female applicants; its weighted mean is 80.2%, 41.2%, 8.6% for female, non-gendered, and male jobs, respectively. The estimated constant term of 0.4090 in column 1 gives the raw, pre-ban female share of applications in N jobs. Thus, the raw female share of applications in male jobs pre-ban is $0.4090 - 0.3509 = 5.81\%$. The raw male share of applications in female jobs pre-ban is $1 - (0.4090 + 0.4864) = 10.46\%$.

Table A6.2: Effects of the Gendered Ad Ban on the Female Share of Call-Backs, Accounting for Embedded Gender Requests

	(1)	(2)	(3)	(4)	(5)
Post ban * Female job *Embedded	0.0803*** (0.0169)	0.0786*** (0.0171)	0.0784*** (0.0171)	0.0699*** (0.0126)	0.0703*** (0.0121)
Post ban * Male job *Embedded	-0.0249** (0.0115)	-0.0265** (0.0115)	-0.0264** (0.0117)	-0.0203* (0.0113)	-0.0206* (0.0109)
Post ban week * Female job	-0.1379*** (0.0081)	-0.1370*** (0.0080)	-0.1361*** (0.0080)	-0.1152*** (0.0066)	-0.1151*** (0.0063)
Post ban week * Male job	0.0308*** (0.0066)	0.0318*** (0.0066)	0.0319*** (0.0066)	0.0276*** (0.0051)	0.0278*** (0.0049)
Post ban week	0.0037 (0.0043)	0.0153*** (0.0052)	0.0033 (0.0054)	0.0049 (0.0040)	
Female job	0.4840*** (0.0057)	0.4835*** (0.0057)	0.4827*** (0.0058)		
Male job	-0.3996*** (0.0071)	-0.3987*** (0.0071)	-0.3987*** (0.0071)		
Quartic in job weeks		Y	Y	Y	Y
Quartic in calendar weeks			Y	Y	
Job ad fixed effects				Y	Y
Calendar week fixed effects					Y
Effective obs	214,489	214,489	214,489	214,489	196,850
R^2	0.219	0.220	0.220	0.691	0.667

Notes:

1. All regressions also include controls for Embedded, Embedded * Male Job, and Embedded * Female Job.
2. Only job ads with a positive number of call-backs are used here.
3. All regressions are weighted by the number of call-backs in each job*week cell, and clustered by firm ID. The numbers of observations used for all specifications are 202,607. As a result of fixed effects specifications, singleton observations were dropped and the effective numbers of observations are smaller in column (5).
4. The dependent variable is the share of call-backs from female applicants at each job*week cell; its weighted mean is 85.7%, 45.0% and 6.7% for female, non-gendered, and male jobs, respectively. The estimated constant term of 0.4474 in column 1 gives the raw, pre-ban female share of applications in N jobs. Thus, the raw female share of call-backs in male jobs pre-ban is $0.4474 - 0.3996 = 4.78\%$. The raw male share of call-backs in female jobs pre-ban is $1 - (0.4474 + 0.4840) = 6.86\%$.

Appendix 7: Robustness to using the Call-Back Sample

As noted, the outcomes in this paper that refer to call-backs can only be estimated on the sub-samples of ads and applications for which call-backs are observed (the *call-back sample*). Appendix 1 has already shown that the observed characteristics of the call-back sample are very similar to the full sample. In this Appendix we provide additional evidence on the representativeness of the call-back sample by replicating the main parts of our analysis that do not require call-back data on the call-back sample. These analyses study the ban's effects on the female share of *applications* (Table 1); and on application arrivals and quality (Tables 4 and 6)

Our main findings are as follows. Table A7.1 replicates Table 1 (female share of applications) on the call-back sample only. The results are almost identical. Table A7.2 replicates Table 4 (aggregate application arrival rates and quality) on the call-back sample only. The main results are very similar: zero aggregate effect on arrival rates, and an increase in quality. Finally, Table A7.3 replicates Table 6 (application arrivals and quality in *F* and *M* jobs) on the call-back sample only. Again, the main results are similar. For arrival rates we again see more men applying to *F* jobs, more women applying to *M* jobs, women reducing their applications to *F* jobs; and men *not* reducing their applications to *M* jobs. One difference is that the total increase in applications to *M* jobs loses statistical significance. For quality, all but one of the post-ban coefficients is positive but our power is low: most coefficients are statistically insignificant.

Table A7.1: Effects of the Gendered Ad Ban on the Female Share of Applications, Call-Back Sample only

	(1)	(2)	(3)	(4)	(5)
Post ban week * Female job	-0.1625*** (0.0053)	-0.1619*** (0.0052)	-0.1612*** (0.0052)	-0.1387*** (0.0028)	-0.1383*** (0.0028)
Post ban week * Male job	0.0459*** (0.0041)	0.0472*** (0.0042)	0.0473*** (0.0041)	0.0342*** (0.0022)	0.0345*** (0.0022)
Post ban week	0.0026 (0.0023)	0.0169*** (0.0032)	-0.0025 (0.0026)	0.0047*** (0.0014)	
Female job	0.4751*** (0.0044)	0.4743*** (0.0043)	0.4738*** (0.0044)		
Male job	-0.3725*** (0.0043)	-0.3713*** (0.0043)	-0.3714*** (0.0043)		
Quartic in job weeks		Y	Y	Y	Y
Quartic in calendar weeks			Y	Y	
Job ad fixed effects				Y	Y
Calendar week fixed effects					Y
Effective # of obs	851,759	851,759	851,759	851,759	850,232
R ²	0.215	0.216	0.216	0.656	0.656

Notes:

1. This Table replicates Table 1 on the sub-sample of job ads for which call-backs are observed.
2. All regressions are weighted by the number of applications in each job*week cell, and clustered by firm ID. The numbers of observations used for all specifications are 851,759. As a result of fixed effects specifications, as singleton observations were dropped, the effective numbers of observations are smaller in column (5).
3. The dependent variable is the share of applications from female applicants; its weighted mean is 80.4%, 42.9%, 8.5% for female, non-gendered, and male jobs, respectively. The estimated constant term of 0.4276 in column 1 gives the raw, pre-ban female share of applications in *N* jobs. Thus, the raw female share of applications in male jobs pre-ban is $0.4276 - 0.3725 = 5.51\%$. The raw male share of applications in female jobs pre-ban is $1 - (0.4276 + 0.4751) = 9.73\%$.

Table A7.2: Aggregate Effects of the Ad Ban on Employer Well-being: Call-back Sample Only

	(1)	(2)	(3)	(4)
	Dependent Variable:			
	Arrivals	Arrivals	Quality	Quality
Post ban	-0.0026 (0.0018)	-0.0020 (0.0018)	0.0189** (0.0076)	0.0215*** (0.0067)
Day-of-week fixed effects	Y	Y	Y	Y
Quartic in job days	Y	Y	Y	Y
Job ad fixed effects		Y		Y
# of obs	1,868,265	1,868,265	257,151	257,151
R^2	0.020	0.352	0.007	0.599

Notes:

5. This table replicates Table 4 on the sub-sample of job ads for which call-backs are observed.
6. Observations are ad-day cells in columns 1 and 2, and applications in columns 3 and 4. Sample is restricted to days within 15 days of the ad ban.
7. All specifications include a linear trend, interacted with Post Ban. Days from the ban (t) is defined as the date of the application minus the date of the gendered ad ban. Job days are the number of days elapsed since the job received its first application. All regressions are clustered by firm ID.
8. Application arrivals are the number of applications received per vacancy in each job*calendar day cell. The average daily number of applications received is .154.
9. Match quality is normalized to a mean of zero and standard deviation of one among all applications in our data. The average normalized match quality for the applications in the current estimation sample is -.007.

Table A7.3: Effects of the Ad Ban on Employers who Posted Gendered Job Ads: Call-Back Sample Only

a. Effects on the Daily Application Arrival Rate:

	(1)	(2)	(3)	(4)
	All	Applications from:		
		All	Women	Men
<i>To Ads that Requested Women:</i>				
Post ban	0.0139*** (0.0049)	0.0143*** (0.0049)	-0.0203*** (0.0043)	0.0346*** (0.0021)
Job Ad Fixed Effects?		Y	Y	Y
# of obs	252,484	252,484	252,484	252,484
R ²	0.023	0.319	0.305	0.215
<i>To Ads that Requested Men:</i>				
Post ban	0.0065 (0.0055)	0.0082 (0.0054)	0.0106*** (0.0014)	-0.0023 (0.0051)
Job Ad Fixed Effects?		Y	Y	Y
# of obs	235,249	235,249	235,249	235,249
R ²	0.013	0.392	0.215	0.395

b. Effects on the Mean Match Quality of Applications:

	(1)	(2)	(3)	(4)
	All	Applications from:		
		All	Women	Men
<i>To Ads that Requested Women:</i>				
Post ban	0.0040 (0.0195)	0.0265 (0.0180)	0.0369* (0.0190)	0.1225* (0.0674)
Job Ad Fixed Effects?		Y	Y	Y
# of obs	36,211	36,211	31,098	7,670
R ²	0.008	0.542	0.555	0.696
<i>To Ads that Requested Men:</i>				
Post ban	0.0167 (0.0201)	0.0217 (0.0183)	-0.0056 (0.0899)	0.0263 (0.0192)
Job Ad Fixed Effects?		Y	Y	Y
# of obs	33,647	33,647	3,375	31,425
R ²	0.004	0.565	0.735	0.569

Notes:

5. This table replicates Table 6 on the sub-sample of job ads for which call-backs are observed.
6. Sample is restricted to days within 15 days of the ad ban.
7. All specifications include a linear trend, interacted with Post Ban. Days from the ban (t) is defined as the date of the application minus the date of the gendered ad ban. All regressions include day-of-week fixed effects and a quartic in job days (the number of days elapsed since the job received its first application). All regressions are clustered by firm ID.
8. In panel (a), observations are ad-day cells, and the dependent variable is the number of applications received per vacancy at each job*calendar day cell. The average daily number of applications received is 0.180 and 0.168 in ads requesting women and men respectively.
9. In panel (b), observations are applications, and the dependent variable is the application's match quality, normalized to a mean of zero and standard deviation of one among all applications in our data. The average normalized match quality for the applications in the current estimation sample is -0.038, and 0.095 in ads requesting women and men respectively.

Appendix 8: Using Application Reads to Measure Jobseekers' Success

As noted, our primary measure of an application's success in this paper --whether it received a call-back—is only observed for 62 percent of applications. In this Appendix we provide additional evidence of our results' robustness by measuring the success of an application by whether it was read by the employer's HR agent. While not as selective as receiving a call-back, this indicator of success is observed for essentially all of the applications (99.86 percent) in our data.⁹ Accordingly, this Appendix replicates our estimates of the ban's effects on women's share of call-backs (Table 2) using our indicator for whether the application was read as our measure of applicant success. Aside from a difference in scale (the main coefficients are about 25 to 30 percent larger in magnitude) the results are highly similar.

⁹ We define reads are being observed for an ad if at least one application was read.

Table A8.1: Effects of the Gendered Ad Ban on the Female Share of Reads

	(1)	(2)	(3)	(4)	(5)
Post ban week * Female job	-0.1538*** (0.0041)	-0.1534*** (0.0041)	-0.1531*** (0.0041)	-0.1327*** (0.0025)	-0.1323*** (0.0025)
Post ban week * Male job	0.0406*** (0.0034)	0.0424*** (0.0034)	0.0425*** (0.0034)	0.0317*** (0.0019)	0.0320*** (0.0019)
Post ban week	0.0074*** (0.0020)	0.0182*** (0.0025)	0.0013 (0.0024)	0.0077*** (0.0013)	
Female job	0.4898*** (0.0036)	0.4894*** (0.0036)	0.4892*** (0.0036)		
Male job	-0.3596*** (0.0036)	-0.3590*** (0.0036)	-0.3590*** (0.0036)		
Quartic in job weeks		Y	Y	Y	Y
Quartic in calendar weeks			Y	Y	
Job ad fixed effects				Y	Y
Calendar week fixed effects					Y
Effective # of obs	1,136,908	1,136,908	1,136,908	1,136,908	1,129,833
R ²	0.201	0.202	0.202	0.652	0.651

Notes:

1. This Table replicates Table 2, replacing call-backs by whether an application was read as an indicator of success. Reads are observed for a larger sample of ads (1,136,908) compared to call-backs (851,759).
2. All regressions are weighted by the number of call-backs in each job*week cell, and clustered by firm ID. The numbers of observations used for all specifications are 1,136,908. As a result of fixed effects specifications, singleton observations were dropped and the effective numbers of observations are smaller in column (5).
3. The dependent variable is the share of reads from female applicants in each job*week cell; its weighted mean is 81.2%, 41.9% and 8.5% for female, non-gendered, and male jobs, respectively. The estimated constant term of 0.4147 in column 1 gives the raw, pre-ban female share of reads in *N* jobs. Thus, the raw female share of reads in male jobs pre-ban is 0.4147 - 0.3596 = 5.51%. The raw male share of reads in female jobs pre-ban is 1 - (0.4147 + 0.4898) = 9.55%.

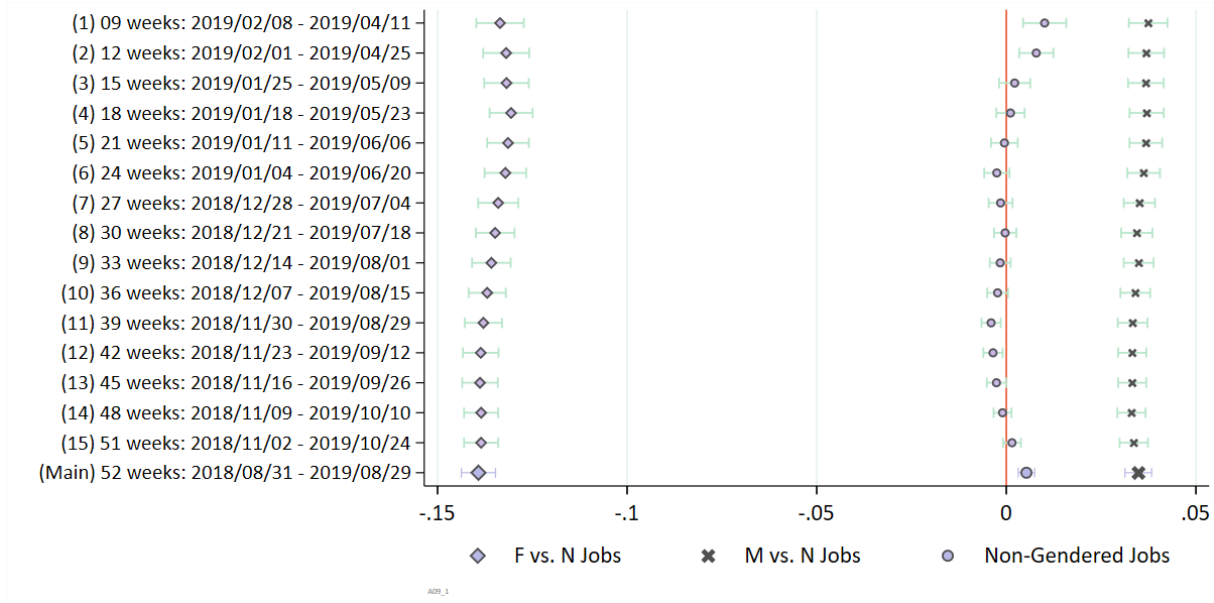
Appendix 9: Robustness to the Length of the Estimation Window

9.1 Gender Mix of Applicant and Call-back Pools

In Appendix 9.1 we explore the sensitivity of our results for the gender composition of applicant and call-back pools (Tables 1 and 2) to the length of the estimation window, which was one year surrounding the ban. To allow for a larger number of alternative windows despite the absence of data from 2018, all our alternative estimation windows put the ban one third of the way between the start and end of the window. For example, if the data window is 30 weeks, then it covers 10 weeks before the ban date and 20 weeks after the ban date. The window lengths used in Appendix 9.1 range from 9 weeks to 51 weeks, each of which covers the pre-ban period for one third of its duration and the post-ban period for two thirds of its duration. At the bottom of each figure, we also report the results from our main specification, which used a 52-week window, half of which is before the ban.

Figure A9.1 shows that –despite this wide range of estimation windows-- all of our alternative estimates of the ban’s effects on the female share of applications to jobs that previously requested men (*M* jobs) are very close to Table 1’s estimate of $-.1392$. The same is true for the other two coefficients of interest in Table 1. Similarly, Figure A9.2 shows that our estimates of the main parameters in Table 2 are remarkably insensitive to the estimation window length that is used. Given the clear visual message of Figure 1 in the paper, this robustness is perhaps not surprising.


Figure A9.1 Effects of the Ban on the Female Share of Applications: Robustness to Estimation Window Length



Notes:

1. The vertical axis shows the various data windows used. In each of the first 15 windows, the length of the data window after the actual gender ad ban is twice the length of the data window before the ban date. Only job ads that received applications both pre and post the gendered ad ban date within each data window are used.
2. The horizontal axis shows the estimated coefficients and 95% confidence intervals.
3. The regression specification of column 4 in Table 1 is used.

Figure A9.2 Effects of the Ban on the Female Share of Call-Backs: Robustness to Estimation Window Length



Notes:


1. The vertical axis shows the various data windows used. In each of the first 15 windows, the length of the data window after the actual gender ad ban is twice the length of the data window before the ban date. Only job ads that have received applications both pre and post the gendered ad ban date within each data window are used.
2. The horizontal axis shows the estimated coefficients and 95% confidence intervals.
3. The regression specification of column 4 of Table 2 is used.

9.2 Employer and Worker Well-Being

In Appendix 9.2 we explore the sensitivity of our results for our two measures of employer well-being—application arrival rates and application match quality, and our measure of worker well-being—application yield—to the length of the estimation window. In our main analysis this window was the 30 days surrounding the ban. Since the massive shifts in recruiting activity associated with the Spring Festival prevent us from expanding this estimation window before the ban, our approach here is simply to explore how the results change as we gradually lengthen the post-ban interval. Specifically, we add one day at a time to the post-ban interval till we reach April 19, 2019 (40 days after the ban). This is just before the steep decline in recruiting activity that precedes the Qingming Festival. In all cases we limit our attention to the effects of the ban on aggregate effects of the ban (in Tables 4 and 5).

For application arrival rates, the estimated effect of the ban first shifts from zero to negative as the estimation window lengthens, then recovers back towards zero in longer windows. For application quality and application yield the results are very stable across all window lengths and very close to our main estimates.


Figure A9.3 Effects of the Ban on the Daily Flow of Applications: Robustness to Estimation Window Length



Notes:

1. All the specifications here use the same data window as the main estimation sample. That is, only job ads that received applications both before and after the gendered ad ban date, and between 2018/8/31 to 2019/8/29 are used.
2. The vertical axis shows the various estimation windows, all of which start on 2019/2/14, 15 days before the gendered ad ban.
3. The horizontal axis shows the estimated coefficients and 95% confidence intervals.
4. The regression specification of Table 4, column 2 is used.


Figure A9.4 Effects of the Ban on Mean Application Quality: Robustness to Estimation Window Length



Notes:

1. All the specifications here use the same data window as the main estimation sample. That is, only job ads that received applications both before and after the gendered ad ban date, and between 2018/8/31 to 2019/8/29 are used.
2. The vertical axis shows the various estimation windows, all of which start on 2019/2/14, 15 days before the gendered ad ban.
3. The horizontal axis shows the estimated coefficients and 95% confidence intervals.
4. The regression specification of Table 4, column 4 is used.

Figure A9.5 Effects of the Ban on the Probability an Application Yields a Call-Back: Robustness to Estimation Window Length



Notes:

1. This figure replicates the coefficient estimates in columns (2), (4) and (6) of Table 5 for 25 different estimation windows. Each window adds one day to the original 30-day window.
2. All the specifications here use the same *data window* as the main estimation sample. That is, only job ads that received applications both before and after the gendered ad ban date and between 2018/8/31 to 2019/8/29 are used.
3. The vertical axis shows the various *estimation windows*, all of which start on 2019/2/14, 15 days before the gendered ad ban.
4. The horizontal axis shows the estimated coefficients and 95% confidence intervals.

Appendix 10: Placebo Bans

10.1 Gender Mix of Applicant and Call-back Pools

To further confirm that our main estimates of the ban's effects on the composition of applicant and call-back pools are not a spurious result of random temporal variation in these outcomes, this Section replicates Tables 1 and 2 71 times, in each case assuming the ban occurred in a different week ranging from March 2, 2018 through July 5, 2019. This is the largest number of placebo ban weeks for which our data allow us to construct a 24-week window surrounding each placebo ban, comprising eight weeks before the ban and 16 weeks afterwards.¹⁰ The goal is to see if estimates that use the actual ban (or a week close to it) are noticeably different from other weeks further away from the ban.

Our results are summarized in Figures A10.1 and A10.2, where the horizontal axis corresponds to the date of the placebo ban. The vertical axis shows coefficient estimates for the ban's effect, from the specification in column 4 of Tables 1 or 2. The solid black vertical line shows the actual ban date, and light blue vertical lines indicate that the estimation window included the actual ban. The dashed black vertical line indicates the week that contained the 2018 lunar equivalent of the actual ban. Three coefficient estimates are reported for each placebo ban date: β^3 (the effects of the ban on non-gendered (N) jobs), and β^1 and β^2 (the effects of the ban on F and M jobs, *relative* to N jobs) as described in equation 1 of the paper.

Figures A10.1 and A10.2 show that our estimates of β^1 and β^2 are (a) larger in magnitude on the date of the actual ban than on any other date, (b) larger in estimation windows that include the actual ban than other windows, and (c) hover around zero in estimation windows that exclude the actual ban. We see this as strong, additional confirmation of a causal effect of the actual ban.

¹⁰ In all cases we also restrict our sample to job ads that received applications both before and after the *actual* ban date. As in our main estimation sample, this ensures that all ads have a known, pre-ban gender request (F , N or M).

Figure A10.1 Effects of Placebo Ad Bans on the Female Share of Applications


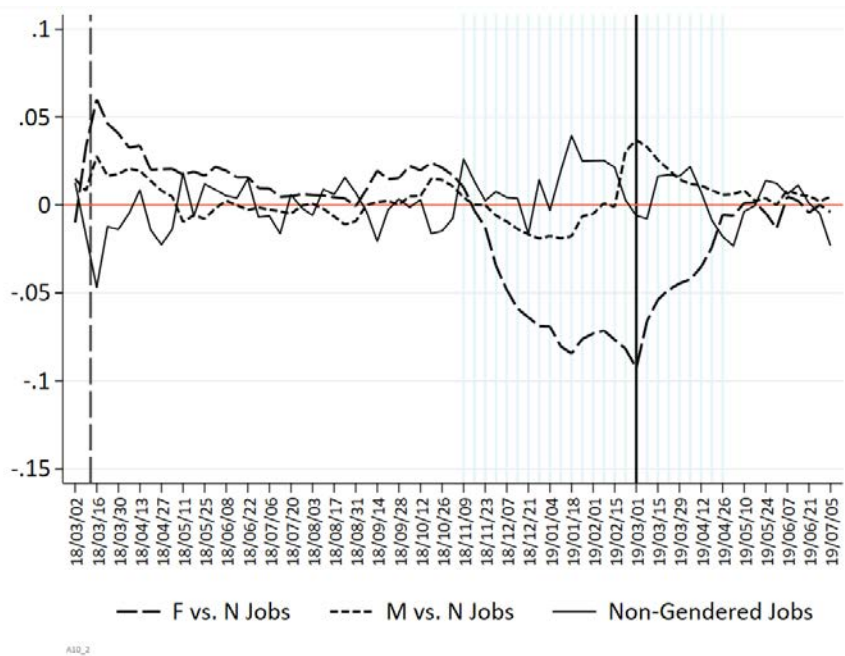


Figure A10.2 Effects of Placebo Ad Bans on the Female Share of Call-Backs



Notes to Figures A10.1 and A10.2:

1. The horizontal axis shows the date of the placebo ban; the vertical axis reports estimated coefficients. Regressions use the specification in column 4 of Tables 1 and 2.
2. The solid black vertical line shows the actual ban date; the dashed black vertical line shows its 2018 lunar equivalent (March 12). The blue vertical lines indicate that the estimation window included the actual ban.

10.2 Employer and Worker Well-Being

In this Section we replicate our estimates of the ban’s effects on our two measures of employer well-being—application arrival rates and application match quality—and on our estimate of well-being --application yield— for placebo ban dates ranging from March 2, 2018 through July 5, 2019. In all cases (as in Tables 4-7) the estimation window is the 30 days surrounding each placebo ban date.¹¹ As in Appendix 9.2, we restrict our attention to the ban’s effects on aggregate matching frictions.

Our results for firm well-being are summarized in Figures A10.3 and A10.4. These estimates use the specification in columns 2 and 4 of Table 4. Our results for worker well-being (application yield) are summarized in Figure A10.5; these use the specification in column 2 of Table 5. As in Figures A10.1 and A10.2, the horizontal axis in all these figures shows the date of the placebo ban, and the vertical axis shows coefficient estimates for the ban’s effect. The solid black vertical line shows the actual ban (March 1, 2019); the five solid blue vertical lines correspond to estimation windows that include the actual ban. The dashed black vertical line shows the 2018 lunar equivalent of the actual ban. The blue shaded areas show major holidays. Finally, the dashed blue-green vertical line indicates the XMRC site upgrade on April 20 2018, which causes a temporary decline in the level of recruiting activity.

Apart from some large, spurious estimates that coincide with major holidays and April 2018 XMRC site upgrade, Figures A10.3 - A10.5 show that our (mostly small or zero) estimated effects of the actual ban on firms’ and workers’ well-being are very similar to effects that would be estimated from a randomly selected placebo ban at other (non-holiday) times. Interestingly, this also applies to match quality, which is the only measure of matching frictions that was significantly affected by the ban in our main estimates. In consequence, we view our positive estimate of the ban’s effects on match quality with less confidence than our other estimates of aggregate matching frictions, even though the positive effect is robust to estimation window length (Appendix 9) and to a different identification strategy in Appendix 11.

¹¹ To construct estimation samples for these placebo regressions, we started with the 24-week window for each placebo ban from Appendix A10.1, then kept only ads or applications that were made within the narrower 30-day window. As in all our analyses, we use only job ads that received applications both before and after the *actual* ban date, to ensure that each ad have a known, pre-ban gender request (*F*, *N* or *M*).

Figure A10.3 Effects of Placebo Ad Bans on the Daily Flow of Applications





Figure A10.4 Effects of Placebo Ad Bans on the Quality of Applications Received




Notes to Figures A10.3 and A10.4:


1. The horizontal axis corresponds to the date of the placebo ban; the vertical axis corresponds to the estimated coefficients. All estimation samples use the 30 days surrounding each placebo ban; these samples were constructed from the 24-week windows used in Figures A10.1 and A10.2. To ensure that all ads have a gender label, we include only ads that received applications both before and after the actual ban date. The regression specification is from columns 2 and 4 of Table 4.
2. The solid black vertical line shows the actual ban (March 1, 2019); the five solid blue vertical lines correspond to estimation windows that include the actual ban. The dashed black vertical line shows the 2018 lunar equivalent of the actual ban.
3. The blue shaded areas correspond to the weeks covering four major holiday periods: National Day (October 1-7, 2018), the Spring Festival (February 4-10, 2019), the Qingming Festival (April 5-7, 2019) and the Labor Day (May 1-4, 2019). The dashed blue-green vertical line corresponds to the XMRC site upgrade on April 20 2018.

Figure A10.5 Effects of Placebo Ad Bans on the Probability an Application Yields a Call-Back


a. Applications from All



b. Applications from Women



c. Applications from Men



Notes to Figure A10.5:

1. These figures replicate columns (2), (4) and (6) from Table 5 for placebo ban dates between March 2018 and July 2019.
2. The horizontal axis corresponds to the date of the placebo ban; the vertical axis corresponds to the estimated coefficients. All estimation samples use the 30 days surrounding each placebo ban; these samples were constructed from the 24-week data windows used in Figures A10.1 and A10.2. To ensure that all ads have a gender label, we include only ads that received applications both before and after the actual ban date.
3. The solid black vertical line shows the actual ban (March 1, 2019); the five solid blue vertical lines correspond to estimation windows that include the actual ban. The dashed black vertical line shows the 2019 lunar equivalent of the actual ban.
4. The blue shaded areas correspond to the weeks covering four major holiday periods: National Day (October 1-7, 2018), the Spring Festival (February 4-10, 2019), the Qingming Festival (April 5-7, 2019) and the Labor Day (May 1-4, 2019). The dashed blue-green vertical line corresponds to the XMRC site upgrade on April 20 2018.

Appendix 11: Difference-in-Difference Analysis

In this Appendix, we replace our single-year, discontinuity-based identification approach by a difference-in-difference approach that uses 2018 as a control for 2019. This approach allows us to include richer controls for time trends, including week-of-the-year (WOY) fixed effects, thus relaxing the smoothness assumptions needed to identify the discontinuity-based estimates. The cost is that we now require a parallel trends assumption: Specifically, net of parameterized *differences* in time trends between the two years, we assume that both years a common set of WOY fixed effects. Other costs include a shorter pre-ban period (since we have no data from 2017) and a noisier measure of call-backs (because some of the data pre-date a 2018 site upgrade that increased the reporting of call-backs).

Since this approach compares the ‘same’ weeks in 2018 and 2019, and important events in China are determined by the lunar calendar, the way we line up weeks between 2018 and 2019 requires some discussion. In 2019, Chinese New Year occurred on February 5. The ad ban happened 24 days later, on Friday March 1. In 2018 Chinese New Year occurred on February 16. Thus, relative to the Chinese lunar calendar, a comparable ad ban in 2018 *would have occurred* 24 days after that, on Monday, March 12. Thus, to construct equivalent whole weeks around the ban (or ‘pseudo-ban’) in both years, we define the first post-ban week in both years as the seven days right after the ban (March 1-7 2019 and March 12-18 2018). Other whole weeks are constructed accordingly, up to six months after the ban, and as far as our data allow (10 weeks) before the ban.

11.1 Gender Mix of Applicant and Call-back Pools

In this Section we explore the robustness of our main results for the gender mix of applicant and call-back pools by replicating Tables 1 and 2 using our alternative, difference-in-difference approach that essentially uses 2018 as a ‘control year’ for 2019. Using the two-year *DiD* sample described in Appendix 2, columns 1 and 2 of Tables 11.1 and 11.2 replicate columns 1 and 2 of Tables 1 and 2, except that we add a fixed effect for all observations from 2019 (to allow for differences in outcome levels between the two years).¹² Only the weeks after the actual ban date, on March 1, 2019 are designated as treated by the ban. Columns 3 and 4 replicate columns 3 and 4 of Tables 1 and 2, replacing the quartic in calendar weeks by the following controls for time trends: (a) a full set of WOY fixed effects (to allow for an arbitrary, non-smooth, seasonal pattern that is common to both years), (b) a quartic in calendar weeks interacted with 2019 (to let lower-frequency and secular trends differ in a smooth way between the two years), and (c) dummies for the Spring Festival period interacted with year (to let the Spring Festival have a different effect in the two years).¹³ Notice that replicating column 5 of Tables 1 and 2 is unnecessary since the calendar week fixed effects added in column 5 of Tables 1 and 2 are replaced by our WOY fixed effects.

Notably, the estimates in Tables A11.1 and A11.2 are almost identical to their counterparts in Tables 1 and 2. The only difference of potential interest involves the effect of the ban on the female share of applicants to *N* jobs (the Post ban week *2019 coefficient): This effect is now statistically insignificant (it was small, positive and statistically significant in Table 1). The ban’s effects on the female share of *call-backs* in *N* jobs is small and insignificant in both our main and DiD specifications. Once again, we conclude that our main estimates of the ban’s effects on the gender composition on application and call-back pools are highly robust.

¹² All regressions in Appendix 11 also contain dummies that capture XMRC’s system transition during 2018.

¹³ Among other reasons, the Spring Festival’s labor market effects may differ between years because it occurs on different days of the week. For example, in 2018 it was Saturday – Wednesday and in 2019 from Tuesday – Sunday.

Table A11.1: Effects of the Gendered Ad Ban on the Female Share of Applications, DiD Approach

	(1)	(2)	(3)	(4)
Post ban week * Female job *2019	-0.1383*** (0.0056)	-0.1391*** (0.0056)	-0.1383*** (0.0056)	-0.1343*** (0.0026)
Post ban week * Male job *2019	0.0486*** (0.0055)	0.0486*** (0.0055)	0.0490*** (0.0055)	0.0359*** (0.0020)
Post ban week *2019	-0.0007 (0.0020)	0.0018 (0.0026)	-0.0071** (0.0034)	0.0025 (0.0023)
Female job	0.4741*** (0.0034)	0.4742*** (0.0034)	0.4734*** (0.0034)	
Male job	-0.3546*** (0.0032)	-0.3545*** (0.0032)	-0.3547*** (0.0032)	
Quartic in job weeks		Y	Y	Y
Time trends			Y	Y
Job ad fixed effects				Y
Effective # of obs	1,805,320	1,805,320	1,805,320	1,805,320
R^2	0.201	0.201	0.202	0.658

Notes:

1. All specifications include a dummy for 2019 and dummies to capture XMRC's system transition during 2018.
2. Time Trends include (lunar) calendar week fixed effects; a quartic in calendar weeks interacted with a dummy for 2019; plus interactions of 2019 with the Spring Festival week, and with the following week.
3. The dependent variable is the share of applications from female applicants; in 2018, its weighted mean is 89.0%, 42.1%, 6.7% for female, non-gendered, and male jobs, respectively; in 2019, its weighted mean is 78.8%, 41.3%, 9.3% for female, non-gendered, and male jobs, respectively.

Table A11.2: Effects of the Gendered Ad Ban on the Female Share of Call-Backs, DiD Approach

	(1)	(2)	(3)	(4)
Post ban week * Female job *2019	-0.0967*** (0.0101)	-0.0969*** (0.0101)	-0.0974*** (0.0101)	-0.1006*** (0.0075)
Post ban week * Male job *2019	0.0388*** (0.0086)	0.0389*** (0.0086)	0.0392*** (0.0086)	0.0311*** (0.0064)
Post ban week *2019	-0.0073 (0.0048)	-0.0041 (0.0059)	-0.0011 (0.0159)	-0.0026 (0.0139)
Female job	0.4629*** (0.0063)	0.4625*** (0.0063)	0.4630*** (0.0063)	
Male job	-0.4106*** (0.0055)	-0.4106*** (0.0055)	-0.4102*** (0.0055)	
Quartic in job weeks		Y	Y	Y
Time trends			Y	Y
Job ad fixed effects				Y
Effective # of obs	215,587	215,587	215,587	215,587
R^2	0.216	0.217	0.217	0.709

Notes:

1. All specifications include a dummy for 2019 and dummies to capture XMRC's system transition during 2018.
2. Time Trends include (lunar) calendar week fixed effects; a quartic in calendar weeks interacted with a dummy for 2019; plus interactions of 2019 with the Spring Festival week, and with the following week.
3. The dependent variable is the share of call-backs to female applicants; in 2018, its weighted mean is 91.6%, 45.9%, 5.1% for female, non-gendered, and male jobs, respectively; in 2019, its weighted mean is 84.9%, 45.3%, 7.0% for female, non-gendered, and male jobs, respectively.

11.2 Employer and Worker Well-Being

In this Section we use our DiD estimation approach to estimate the ban’s effects on firm and worker well-being (Tables 4-7 in our main analysis). A noteworthy benefit of the DiD approach for these outcomes in particular is that we can now abandon the short, 30-day estimation windows in Tables 4-7 and instead use weekly data for two comparable periods: January – August 2018 and January-August 2019. This longer estimation period allows us to use week-of-the-(lunar)-year fixed effects to capture an arbitrary, high-frequency time trend that is common to the two years.

Focusing first on the aggregate effects of the ban on employers, Table A11.3 replicates Table 4 of the paper using the same set of time trends controls as Tables A11.1 and A11.2 (which include week-of-the-year fixed effects). Consistent with Table 4, column 2 of Table A11.3 shows a statistically insignificant effect of the ad ban on the daily flow of applications per vacancy on XMRC.¹⁴ Also consistent with Table 4 we find a small but statistically significant increase in the mean match quality of applications. Table A11.4 replicates Table 5, which estimates the ban’s aggregate effects on workers. Here, the small, insignificant declines in call-back yields become slightly larger and are now statistically significant (though not for women). While this suggests that the ban loosened the entire XMRC labor market, we are less confident in this result than our main estimates, primarily because of DiD method’s reliance on parallel trends assumptions. Inspection of the 2018 and 2019 trends for call-back yields shows a large pre-ban surge in 2019 that is absent in 2018. While we attempt to capture this by interacting the effects of the Spring Festival (and following week) with year, it is unclear that this accounts for all the relevant differences between the two years.

Turning now to the ban’s effects on the jobs whose gender requests were removed, Table A11.5 replicates Table 6. Aside from the higher magnitudes in panel (a), associated with the switch from daily to weekly application counts, the results are very similar. Both Tables show an increase in total applications to jobs that previously requested women and men, with the former increase much larger than the latter. Both tables attribute the increase in applications to men’s jobs solely to additional applications from women, while the increase in applications to women’s jobs is the net result of a large increase in applications from men and a decline in women’s applications. Both Tables also show imprecisely estimated but largely positive effects of the ban on the quality of applications to *M* and *F* jobs. Finally, Table A11.6 replicates Table 7, with somewhat different results. Unlike Table 7, it does not find a disproportionate loosening of labor markets in the previously gendered jobs, with one exception: the call-back rate of women who applied to previously male jobs fell after the ban. Again, because of the need to rely on parallel trends assumptions, we place less weight on this estimate than on our main estimates.

¹⁴ Both the coefficient and standard error are larger, reflecting the fact that the dependent variable is now applications received per week, rather than per day.

Table A11.3: Aggregate Effects of the Ad Ban on Employers: DiD Approach

	(1)	(2)	(3)	(4)
	Dependent Variable:			
	Application Arrivals	Application Arrivals	Application Quality	Application Quality
Post ban	-0.0024 (0.0134)	0.0162 (0.0145)	-0.0076 (0.0067)	0.0094** (0.0037)
Quartic in job age	Y	Y	Y	Y
Job ad fixed effects		Y		Y
# of obs	4,122,080	4,122,080	3,891,632	3,891,632
R ²	0.056	0.453	0.003	0.494

Notes:

1. This table replicates Table 4 using a Difference-in-Difference approach, treating weeks from 2018 as a control group for 2019.
2. In columns 1 and 2, observations are ad-week cells, and the dependent variable is the number of applications received per vacancy in each cell. Post ban means the week was after the actual ban. The average weekly number of applications received is 0.83.
3. In columns 3 and 4, observations are applications, and the dependent variable is the application's match quality, normalized to a mean of zero and standard deviation of one among all applications in our data. Post ban means the application was submitted after the actual ban. The average normalized match quality for the applications in the current estimation sample is -0.01 (s.d. 0.99).
4. All specifications include a quartic in job age and a dummy for 2019. Time trends are modelled with (lunar) calendar week fixed effects; a quartic in time interacted with 2019; plus dummies for the Spring Festival week and the week following it, both interacted with 2019.
5. All regressions are clustered by firm ID.

Table A11.4: Aggregate Effects of the Ad Ban on Jobseekers: DiD Approach for Application Yield

	(1)	(2)	(3)	(4)	(5)	(6)
	All		Applications from: Women		Men	
Post ban	-0.0056** (0.0023)	-0.0056** (0.0024)	-0.0070** (0.0032)	-0.0051 (0.0035)	-0.0043* (0.0025)	-0.0061** (0.0027)
Quartic in job weeks	Y	Y	Y	Y	Y	Y
Applicant fixed effects		Y		Y		Y
Effective # of obs	3,986,179	3,986,179	1,668,960	1,668,960	2,317,219	2,317,219
R ²	0.018	0.129	0.021	0.134	0.016	0.123

Notes:

1. This table replicates Table 5 using a Difference-in-Difference approach, treating weeks from 2018 as a control group for 2019.
2. Observations are applications. Post ban means the application was made after the actual ban. The dependent variable equals one if the application *ever* received a call-back. The average call-back probability is 0.100 and 0.082 for applications from women and men respectively, and 0.089 overall.
3. All specifications include a quartic in job weeks and a dummy for 2019. Time trends are modelled with lunar calendar week fixed effects; a quartic in time interacted with 2019; plus dummies for the Spring Festival week and the week following it, both interacted with 2019.
4. All regressions are clustered by firm ID.

Table A11.5: Effects of the Ad Ban on Employers who Posted Gendered Job Ads: DiD Approach

a. Effects on the Weekly Application Arrival Rate:

	(1)	(2)	(3)	(4)
	All	Applications from:		
		All	Women	Men
<i>To Ads that Requested Women:</i>				
Post ban	0.1525*** (0.0388)	0.1761*** (0.0415)	-0.1152*** (0.0374)	0.2913*** (0.0138)
Job Ad Fixed Effects?		Y	Y	Y
# of obs	425,656	425,656	425,656	425,656
R ²	0.071	0.425	0.421	0.307
<i>To Ads that Requested Men:</i>				
Post ban	0.0469 (0.0377)	0.0670* (0.0407)	0.1039*** (0.0088)	-0.0369 (0.0393)
Job Ad Fixed Effects?		Y	Y	Y
# of obs	509,644	509,644	509,644	509,644
R ²	0.041	0.462	0.363	0.459

b. Effects on the Mean Match Quality of Applications:

	(1)	(2)	(3)	(4)
	All	Applications from:		
		All	Women	Men
<i>To Ads that Requested Women:</i>				
Post ban	-0.0164 (0.0182)	0.0147 (0.0106)	0.0279** (0.0112)	0.0528 (0.0362)
Job Ad Fixed Effects?		Y	Y	Y
# of obs	426,100	426,100	356,223	69,877
R ²	0.004	0.448	0.453	0.568
<i>To Ads that Requested Men:</i>				
Post ban	-0.0021 (0.0166)	0.0016 (0.0107)	0.0454 (0.0386)	0.0000 (0.0112)
Job Ad Fixed Effects?		Y	Y	Y
# of obs	511,557	511,557	41,065	470,492
R ²	0.003	0.424	0.634	0.422

Notes:

1. This table replicates Table 6 using a Difference-in-Difference approach, treating weeks from 2018 as a control group for 2019.
2. In panel (a), observations are ad-week cells, and the dependent variable is the number of applications received per vacancy in each job*week cell. Post ban means the week was after the actual ban. Column 3 (4) only counts applications from women (men). The average weekly number of applications received is 0.96 and 0.86 in ads requesting women and men respectively.
3. In panel (b), observations are applications, and the dependent variable is the application's match quality, normalized to a mean of zero and standard deviation of one among all applications in our data. Post ban means the application was submitted after the actual ban. The average normalized match quality for the applications in the current estimation sample is -0.06 (s.d. 0.98) and 0.07 (s.d. 0.96) in ads requesting women and men respectively. Please check numbers.
4. All specifications include a quartic in job age and a dummy for 2019. Time trends are modelled with (lunar) calendar week fixed effects; a quartic in time interacted with 2019; plus dummies for the Spring Festival week and the week following it, both interacted with 2019.
5. All regressions are clustered by firm ID.

Table A11.6: Effects of the Ad Ban on Workers who Applied to Gendered Job Ads: DiD Approach for Application Yield

	Applications from:		
	All	Women	Men
	(1)	(2)	(3)
Post ban	-0.0059** (0.0025)	-0.0052 (0.0037)	-0.0066** (0.0028)
Post ban * Female job	-0.0015 (0.0044)	0.0047 (0.0047)	-0.0073 (0.0071)
Post ban * Male job	0.0021 (0.0036)	-0.0121** (0.0058)	0.0048 (0.0038)
Day-of-week fixed effects	Y	Y	Y
Quartic in job weeks	Y	Y	Y
Applicant fixed effects	Y	Y	Y
Effective # of obs	3,986,179	1,668,960	2,317,219
R ²	0.129	0.135	0.123

Notes:

1. This Table replicates Table 7 using a Difference-in-Difference approach, treating weeks from 2018 as a control group for 2019.
2. Observations are applications. Post ban means the application was made after the actual ban. The dependent variable equals one if the application ever received a call-back. The average call-back probability is 0.100 and 0.082 for applications from women and men respectively, and 0.089 overall.
3. All specifications include a quartic in job weeks, applicant fixed effects, and a dummy for 2019. Time trends are modelled with lunar calendar week fixed effects; a quartic in time interacted with 2019; plus dummies for the Spring Festival week and the week following it, both interacted with 2019.
4. All regressions are clustered by firm ID.