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

Peter Kuhn and Kailing Shen

Appendix for online publication

February 15, 2023

Appendix 1: Gendered Ads, and Gendered Ad Bans in Selected Countries

A1.1 Estimated Prevalence of Explicit Gender Requests in Populous Countries[1]

To generate a rough estimate of the worldwide prevalence of explicitly gendered job ads, in Fall 2018 we started with a list of the world's 20 most populous countries. After excluding the U.S. and China, plus four countries that were not served by Indeed.com (Bangladesh, Ethiopia, Democratic Republic of Congo, and Iran), we searched the remaining countries' Indeed websites for job ads that made explicit gender requests.[2] We then excluded two countries (Egypt and Brazil) where fewer than two percent of the job ads appear to have made an explicit gender request. This left us with twelve countries in which gendered job ads were commonly used, ranging from 2.7 to 47.3 percent of all ads.

Table A1.1 reports the populations and estimated prevalence of gendered ads in these 12 countries. Together, these countries have a population of 2.8 billion people, or 37.4 percent of the world total. The population-weighted share of ads that made gender requests was 18.7 percent.

[1] The data on gender requests used in this Appendix were collected by the authors, with the assistance of undergraduate assistants Steve Li and Jia You. They can be accessed at Kuhn and Shen (2023).

[2] All the countries' Indeed job boards were accessed through Indeed's international portal, at https://www.indeed.com/worldwide.

Table A1.1.1: Estimated Share of Ads that are Gendered in Populous Countries, 2018

Country	Population (millions)	Share of ads that are gendered (estimated)
India	1339	0.1266
Indonesia	264	0.4732
Pakistan	197	0.3985
Nigeria	191	0.0885
Russia	147	0.0380
Mexico	129	0.0269
Japan	127	0.2384
Philippines	105	0.2988
Vietnam	95	0.2095
Germany	82	0.2987
Turkey	81	0.0851
Thailand	69	0.2475
Mean (weighted)		0.1872
Total:	2826	
World Pop:	7550	
Share of World:	0.374	

Notes:

1. Country and world population statistics are from United Nations Statistics Division (2017).

2. The countries are the 20 most populous, excluding the USA and China, excluding four countries without an Indeed site in 2018 (Bangladesh, Ethiopia, DR Congo, and Iran), and excluding two countries where fewer than 2 percent of Indeed ads made explicit gender requests (Egypt and Brazil).

3. The gendered share of ads is based on author-collected data derived from web searches of Indeed.com's international web portal. This data and the (Excel) code that constructs Table A1.1.1 are available at Kuhn and Shen (2023).

4. Ads are classified as gendered if they included at least one of four keywords (two male and two female) suggesting a gender request. Manual inspection of a random sample of ads indicated that this procedure omits some gender requests (that were made in other ways), and includes some keyword uses that do not constitute a stated gender preference (such as statements that welcome both genders). Thus, the numbers in Table A1.1 should be interpreted as rough estimates only. For a more exhaustive text analysis that identifies explicit gender discrimination in job ads in Indonesia, see Ningrum et al. (2020). They estimate that 38.8 percent of job ads on an Indonesian board discriminated on the basis of gender, which is somewhat lower than our estimate of 47.3 percent for that country.

A1.2 The Decline of Gendered Job Ads in United States, Austria, and China[3]

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.

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

A1.3 The Decline of Gendered Job Ads -- Other English-Speaking Countries[4]

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 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.

---

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

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*.[5] 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).[6] 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.

[5] Access was via scanned copies provided by librarians in New Zealand.
[6] Access was via Irish Newspaper Archives, an online database of Irish newspapers.

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.

References for Appendix 1

Card, David, Fabrizio Colella, and Rafael Lalive. 2021. "Gender Preferences in Job Vacancies and Workplace Gender Diversity." National Bureau of Economic Research Working Paper 29350.

Kuhn, Peter, and Kailing Shen. 2023. "Data and Code for: What Happens When Employers Can No Longer Discriminate in Job Ads?" *American Economic Association* [publisher], Inter-university Consortium for Political and Social Research [distributor]. http://doi.org/10.3886/E183021V1

Ningrum Panggih Kusuma, Tatdow Pansombut, and Attachai 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.

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.

United Nations Statistics Division. 2017 "Total Population - Both Sexes". *World Population Prospects, the 2017 Revision*. United Nations Department of Economic and Social Affairs, Population Division, Population Estimates and Projections Section. June 2017. *Quoted in Wikipedia, "List of countries by population (United Nations)*" (accessed November 8, 2018)

Walsh, John, Miriam Johnson, and Marged Sugarman. 1975. "*Help Wanted: Case Studies of Classified Ads*." Salt Lake City, Utah: Olympus Research Corporation.

Appendix 2: Setting, Sample Design and Representativeness

A2.1 A Sample XMRC Ad

A copy of a sample XMRC job ad (with translations) is available on Peter Kuhn's website, at the following link:

https://drive.google.com/file/d/1MfLUyBsma6KBM6re1167vtwaimhXV1Wv/view?usp=sharing

A2.2 Sample Design

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 A2.3.1 shows the mean characteristics of the 117,390 job ads in our main analysis sample, and the 62,116 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 A2.3.1 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 5,793 yuan per month, *M* jobs pay 6.1 percent less than *N* jobs; at 4,433 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 Helleseter 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 A2.3.1 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 Helleseter et al. 2020).

Table A2.3.2 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 to that they represent the same days and weeks relative to the start of the Spring Festival in both years.

A2.3 Descriptive Statistics

Table A2.3.1: 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	11.82	13.30	13.07	13.40	11.60	13.13
Require technical school education?	0.214	0.103	0.156	0.122	0.218	0.117	0.171	0.137
Experience requirement (years)	0.832	1.282	1.617	1.271	0.828	1.156	1.518	1.157
Require new graduates?	0.029	0.020	0.019	0.021	0.030	0.020	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	29.77	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)	4,433	6,168	5,793	5,897	4,457	5,814	5,629	5,596
Bonus offered?	0.167	0.211	0.104	0.192	0.175	0.252	0.109	0.224
Bonus offered (Yuan)	6,701	8,443	7,722	8,219	6,700	8,407	7,658	8,180
Commission offered?	0.173	0.225	0.112	0.204	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	2.024	2.546	2.322	2.456	2.138	2.734	2.367	2.605
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	342	801	610	724	305	521	453	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	0.037	0.037	0.035
Private shared	0.663	0.611	0.575	0.612	0.675	0.661	0.605	0.656
Private	0.134	0.129	0.115	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 European companies	0.007	0.009	0.012	0.009	0.006	0.007	0.013	0.008
Other joint venture	0.017	0.026	0.036	0.026	0.015	0.021	0.035	0.022
Foreign representative offices	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	0.033	0.036	0.031	0.033	0.027	0.032
Average # of applications received per job ad	40.09	40.27	45.17	40.86	48.19	47.07	53.81	48.06
# of job postings	13,663	89,053	14,674	117,390	8,436	45,905	7,775	62,116
% of job postings	11.64	75.86	12.50	100	13.58	73.90	12.52	100

Table A2.3.2: 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	10.31	9.38	8.09	10.22	9.30
New graduates?	0.001	0.000	0.000	0.001	0.000	0.000
Local *hukou*	0.355	0.269	0.305	0.343	0.253	0.292
Fujian *hukou*	0.817	0.748	0.776	0.806	0.727	0.761
Short sighted?	0.344	0.297	0.316	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	7,149	6,549	5,444	6,740	6,192
Valid intended wage?	0.652	0.622	0.635	0.659	0.625	0.640
Intended wage	5,922	7,563	6,868	5,687	7,099	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	0.035	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	0.349	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	0.267	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	0.747	0.715	0.725	0.721
To M jobs	0.029	0.213	0.137	0.028	0.227	0.141
# of applications	1,291,330	1,842,273	3,133,603	839,543	1,101,622	1,941,165
% of applications	41.21	58.79	100	43.25	56.75	100

Note:

1. 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.

A2.4 Did XMRC's Late Ban Adoption Affect Its Representativeness?

A potential concern about the representativeness of our XMRC setting is that XMRC was forced to remove its explicit gender requests considerably later than China's major national job boards. Enforcement actions against gendered job ads in China were absent before May 2016, when China's Ministry of Industry and Information Technology issued a regulation (henceforth MIIT 2016) that clarified the job boards' liability for posting such ads.[7] After that, enforcement actions against the job boards appears to have been rolled out quite gradually, not reaching XMRC till March 2019 when it was ordered to remove its gendered ads. This raises the possibility that the employers who were still posting gender requests on XMRC in 2019 were a highly selected group. For example, they could include a subset of employers who knowingly broke the law, or who were disproportionately uninformed.

To explore this possibility we first explored how the 2016 MIIT regulation was enforced in China. Based on our consultations with XMRC personnel and supported by our own search of media articles, enforcement actions (including fines) against *employers* posting gendered job ads were essentially non-existent between 2016 and our estimation period. In fact, they remain exceedingly rare even today. Essentially, throughout this period employers could safely conclude that they would not face any penalties for posting an explicit gender request, because MIIT focused all its enforcement actions on the job boards only. If true, this would suggest that the employers who were posting explicitly gendered ads on XMRC immediately before the March 2019 removal might not be appreciably different from the employers who posted such ads much earlier.

Second, to see if this was actually the case, we turned to our 'historical' XMRC data, which allows us to calculate the share of active job ads requesting men or women on a monthly basis for the periods January -- October 2016 and January 2018 -- February 2019. (A job ad is considered active if it received *any* applications in a month.) The results are shown in Figure A2.4.1. Notably, the share of ads requesting men and women exhibited no change in the months surrounding MIIT 2016. More importantly, the shares of XMRC job ads requesting men and women were unchanged during the fourteen months immediately preceding the overnight removal of gender requests from XMRC on March 1, 2019, and were essentially identical to the shares in 2016.

In sum, Figure A2.4.1 strongly suggests that employers on XMRC did *not* start to avoid making explicit gender requests during the period between the 2016 regulation and the 2019 removal of gendered ads from XMRC. While this lack of action might appear to constitute defiance of Chinese law, it can be explained by the fact that enforcement has never been aimed at employers (even to this day) and was only applied to individual job boards *after* they were directly ordered to remove such ads. Thus, we conclude that our estimates are not subject to the potential biases described at the start of this Section. Instead, Figure A2.4.1 suggests that our estimates are likely to be representative of what a ban would have done at any time between May 2016 and March 2019, and perhaps even earlier, before the Chinese government announced or started any enforcement actions against gendered ads.

---

[7] Additional details on this and earlier changes in Chinese policies and practices related to gender discrimination are available in Appendices 2 and 3 of Kuhn, Shen, and Zhang (2020).

Figure A2.4.1 Share of Job Ads with Explicit Gender Preferences in XMRC

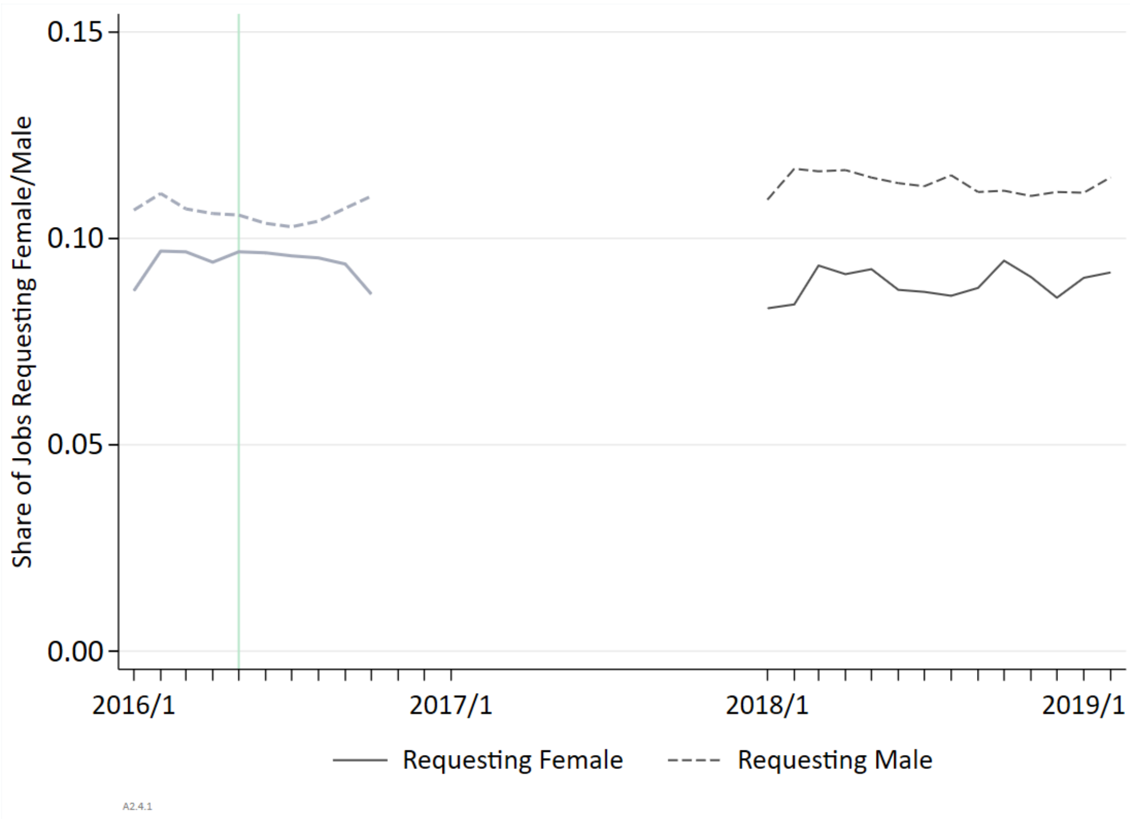

Notes:

1. For each month from 2016 January to 2016 October, and 2018 January to 2019 February, the vertical axis shows the share of job ads with explicit gender preferences among all active job ads in each period. Ads are weighted by the number of vacancies each ad represents.
2. The green vertical line represents May 2016, when the Ministry of Industry and Information Technology issued its regulation clarifying fines for posting gendered job ads.

A2.5 Understanding Changes in Relative Conditional Call-Back Rates After the Ban

In Section III.A of the paper, we noted an interesting pattern in the relative call-back rates experienced by applicants to previously gendered jobs: Men's relative call-back chances in previously female jobs improved modestly (from 65.4 percent of women's to 68.8 percent), while women's relative call-back chances in previously male jobs declined (from 86.6 percent of men's to 73.3 percent). In this Section we explore the relative importance of two possible causes of this change: changes in the match quality of applicant pools to gender-mismatched jobs, versus causal effects of the ban on employers' choices between equally qualified men and women.

Turning first to the match quality hypothesis, we first note that the ban caused a large upsurge in applications, and a large change in the gender mix of applications to previously gendered jobs. So it is quite possible that the quality of both the male and female applicants to these jobs changed after the ban. Next, we note that panels C and D of Table 5 present estimates of the ban's effects on the mean match quality of gender-mismatched applications. While some of these estimates are imprecise, they do suggest a large increase in the quality of men's applications to F jobs after the ban, and no change in the genders' relative match quality in *M* jobs -- a pattern that is strikingly consistent with the changes in relative call-back rates.

Turning now to the 'changing employer preferences' hypothesis, note first that -- since we are trying to identify *changes* in relative call-back rates, this hypothesis requires the gender preferences of recruiters for previously female and previously male jobs to *change* after the ban. While this might seem implausible, we note that the ban might, for example, have made recruiters more reluctant to indulge their gender preferences at the call-back stage, even though the ban imposed no direct constraints on call-back or hiring decisions. That noted, to explain the result in question, the preferences hypothesis requires a very particular type of preference change: the ban would need to shift employers' tastes *in favor of men in both F and M jobs*. This is because the observed decline in women's relative callback rates in *M* jobs, and the observed increase in men's relative call-back rates in *F* jobs both represent shifts in favor of men. This seems less plausible to us.

Finally, turning to the evidence on call-back rates, note that Table 6 estimates conditional call-back rates for both male and female applicants to previously gendered jobs. Notably, these estimates control for worker fixed effects, so they control for unmeasured aspects of applicant quality. As noted in the paper, the ban caused an overall reduction in call-back rates in previously gendered jobs, because the total number of competing applications increased. Comparing columns 3 and 4, however, we cannot rule out a null effect of the ban on the genders' relative success rates.

In sum, while our statistical power in Tables 5 and 6 is less than we would like, the evidence we have is more consistent with the application composition/match quality hypothesis than a preference-based one. The preference-based story also requires a specific type of change in employers' gender preferences, which we find implausible. We hasten to remind readers, however, that both these changes in relative, conditional call-back rates play very small roles in explaining the main results of our paper: the ban's effects on the gender mix of call-back pools. These effects are driven almost entirely by the dramatic changes in the gender mix of applications to gender-mismatched jobs, and by the fact that both men's and women's conditional call-back rates in those jobs remained surprisingly high after the ban.

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 modified 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. Finally, Section A3.6 measures the incidence of the ban's integrating effects in a different way, by calculating the number of distinct firms and job titles in which those effects occurred. We find that the effects were quite widespread, and not confined to a small number of firms or job titles.

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 A3.1.1 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 1.430 percentage points (just under half) of the 3.032 percentage point increase in the share of women called back to *M* job ads in our data.[8] 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.[9] Under this definition, notice that job ads that request men can occur in

[8] Because of sampling differences and regression controls, this 3.032 percentage point estimate in these Appendix calculations differs slightly from our main, regression- adjusted estimate of 2.46 + 0.48 = 2.94 percentage points (Table 2, column 4). We use raw means to perform the decompositions in this Appendix to maximize simplicity and transparency.

[9] 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

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 A3.1.1, 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 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 A3.1.1 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.

call-backs in our estimation sample are not affected by any dynamic statistical processes, such as mean reversion.

Table A3.1.1: Job Titles Accounting for Women's Increased Representation in *M* jobs

	(1)	(2)	(3)	(4)
	Job Title	English Translation	Contribution to Women's Increased Call-Back Share in *M* jobs ($\alpha^T \delta^T$)[1] (Percentage Points)	Incumbent Male Share[2]
1	仓管员	Warehouse management staff	0.247	0.852*
2	经理助理	Manager assistant	0.201	0.356
3	仓管理	Warehouse management staff	0.129	0.874*
4	交易员	Trader	0.083	0.967*
5	采购专员	Procurement specialist	0.078	0.395
6	普工	General labor	0.072	0.853*
7	电商仓库仓管配货员打包员	E-commerce warehouse management staff, picker, packer	0.058	0.966*
8	市场专员	Marketing specialist	0.055	0.684*
9	市政作业员	Civil construction worker	0.051	1.000*
10	仓库管理员	Warehouse management staff	0.046	0.861*
11	行政主管	Administrative chief	0.045	0.302
12	行政专员	Administration specialist	0.043	0.134
13	行政助理	Administration assistant	0.043	0.078
14	前端开发	Client-side developer	0.043	0.830*
15	工程助理	Engineer assistant	0.043	0.635*
16	审计税审专员	Auditing specialist (taxation)	0.043	0.237
17	财务经理	Financial manager	0.042	0.488
18	平面设计	Graphic design	0.038	0.559*
19	活动执行	Event implementation	0.037	0.630*
20	Java 工程师	Java engineer	0.036	0.920*
Top 20 titles, combined (T1 to T20)[3]			1.430	0.660*
All contributing job titles (TOTAL)[3]			3.032	0.677*
Contribution of the top twenty titles [(T1 to T20)/TOTAL]			47.2%	

* Incumbent male share exceeds the sample mean (0.499).

Notes:

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 A3.2.1 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 3.574 percentage points (or about 37 percent) of the 9.610 percentage point increase in the share of men called back to formerly *F* job ads in our data.[10] In addition, Table A3.2.1 shows that all these titles except two -- warehouse management staff and operation assistant -- 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 A3.2.1 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 .735.

[10] As was the case for *M* jobs, the 9.610 percentage point impact used in these Appendix calculations differs slightly from our main, regression- adjusted estimate of (–10.39 + 0.48 =) 9.91 percentage points (from Table 2, column 4).

Table A3.2.1: Job Titles Accounting for Men's Increased Representation in *F* jobs

	(1)	(2)	(3)	(4)
	Job Title	English Translation	Contribution to Men's increased Call-Back share in *F* jobs ($\alpha^T \delta^T$)[1] (Percentage Points)	Incumbent Female Share[2]
1	经理助理	Manager assistant	0.414	0.644*
2	仓管员	Warehouse management staff	0.290	0.148
3	行政专员	Administration specialist	0.266	0.866*
4	采购员	Procurement officer	0.242	0.588*
5	会计	Accountant	0.241	0.887*
6	文员	Office clerk	0.216	0.940*
7	财务会计	Financial accountant	0.197	0.889*
8	人事行政专员	HR administrative specialist	0.174	0.922*
9	外贸业务员	International trade sales staff	0.173	0.753*
10	主办会计	Chief accountant	0.147	0.865*
11	幼儿园出纳	Kindergarten cashier	0.142	0.957*
12	运营助理	Operation assistant	0.141	0.409
13	商务助理	Business assistant	0.139	0.822*
14	人事专员	HR assistant	0.128	0.912*
15	客服专员	Customer service specialist	0.119	0.788*
16	销售助理	Sales assistant	0.115	0.746*
17	客服	Customer service	0.110	0.677*
18	采购	Procurement officer	0.109	0.591*
19	仓管文员	Warehouse management clerk	0.107	0.925*
20	财务助理	Financial assistant	0.104	0.916*
Top 20 titles, combined (T1 to T20)[3]			3.574	0.736*
All contributing job titles (TOTAL)[3]			9.610	0.735*
Contribution of the top 20 titles [(T1 to T20)/TOTAL]			37.2%	

*Incumbent female share exceeds sample average (0.501).

Notes:

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, panel a of Table A3.3.1 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, panel b of Table A3.3.1 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, none of the 40 job titles that appear in Table A3.3.1 appears in Table A3.1.1, 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 A3.3.2(a) and A3.3.2(b) perform the same exercise for extremely female jobs, with patterns that are only slightly less pronounced. None of the 40 job titles listed here appear in Table A3.2.1 -- 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 A3.3.1’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 = 894 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.[11]

Inspection of Table A3.3.2’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 A3.3.2’s 40 extremely female jobs appear in Table A3.2.1 (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.

[11] 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 A3.3.1: Job Titles with the Highest Historical Male Call-Back Shares

a. 100% of Call-Backs Went to Men

	(1)	(2)	(3)	(4)
	Job Title	English 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 chairperson	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	CNC operator	75	1
16	CNC 操机员	Operator of computer numerical control (CNC) machine	75	1
17	水电工	Plumber	75	1
18	工程监理	construction project supervisor	75	1
19	注塑主管	Injection supervisor	72	1
20	装修施工员	Decoration builder	72	1

b. At Least One Women Received a Call-Back

	(1)	(2)	(3)	(4)
	Job Title	English Translation	Number of Call-Backs Issued	Share of Call-Backs Issued to Men
1	行政司机	Driver for administration tasks	469	0.9979
2	司机	Driver	1480	0.9966
3	注塑领班	Injection molding team leader	170	0.9941
4	网络管理员	Network management staff	169	0.9941
5	货车司机	Truck driver	334	0.9940
6	生产厂长	production CEO	164	0.9939
7	焊工	Welder	146	0.9932
8	电工	Electrician	270	0.9926
9	设备工程师	Equipment engineer	267	0.9925
10	模具钳工	Mold fitter	126	0.9921
11	生产经理	Production manager	222	0.9910
12	送货司机	Delivery driver	393	0.9898
13	保安	Security	362	0.9890
14	钳工	Fitter	87	0.9885
15	CNC 操作员	CNC operator	77	0.9870
16	模具钳工师傅	Master mold fitter	74	0.9865
17	生产部经理	Production sector manager	142	0.9859
18	模具设计	Mold design	67	0.9851
19	设备技术员	Equipment technician	66	0.9848
20	机械研发工程师	Mechanical R&D engineer	63	0.9841

Notes:

1. Statistics refer to all the call-backs that occurred during our pre-estimation sample (January - August 2018).
2. Panel a lists the twenty largest job titles (by call-backs) that sent all their call-backs to men.
3. Panel b lists the twenty job titles with the highest share of call-backs issued to men that were strictly less than 100 percent.

Table A3.3.2: Job Titles with the Highest Historical Female Call-Back Shares

a. 100% of Call-Backs Went to Women

	(1)	(2)	(3)	(4)
	Job Title	English Translation	Number of Call-Backs Issued	Share of Call-Backs Issued to Women
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	人资行政专员	HR payroll admin specialist	24	1
14	出纳文员	Cashier clerk	24	1
15	保育员	Childcare teacher	24	1
16	美妆在线销售客服	Online sales & customer service for make-up products	23	1
17	行政专员前台	Admin specialist and front desk	23	1
18	跟单采购助理	Procurement assistant	23	1
19	国际展览业务员客服方向	International conference business customer service	23	1
20	现场咨询师	Onsite consultant	22	1

b. At Least One Call-Back Went to Men

	(1)	(2)	(3)	(4)
	Job Title	English Translation	Number of Call-Backs Issued	Share of Call-Backs Issued to Women
1	前台文员	Receptionist and clerk	397	0.9950
2	出纳人事	Cashier & HR	91	0.9890
3	行政前台文员	Administration, receptionist, and clerk	85	0.9882
4	财务行政	Financial admin	68	0.9853
5	人事行政文员	HR, administration clerk	123	0.9837
6	行政前台	Administration and receptionist	1031	0.9825
7	母婴在线销售客服	Mom-Baby online sales customer service specialist	48	0.9792
8	生产统计	Production statistic	46	0.9783
9	平面模特	Still model	44	0.9773
10	小学英语教师	English teacher for primary school	86	0.9767
11	生产统计文员	Production statistic clerk	42	0.9762
12	业务助理业务跟单	Business assistant, merchandiser	82	0.9756
13	助理文员	Assistant clerk	78	0.9744
14	财务人事	Finance & HR	39	0.9744
15	财务会计商务	Financial accountant, business	38	0.9737
16	会计员	Accountant	75	0.9733
17	内勤	Internal clerk	73	0.9726
18	办公室文员打单人员文秘	Office clerk, documentary handler, secretary	36	0.9722
19	外贸单证	International trade document clerk	71	0.9718
20	客服代表	Customer representative	35	0.9714

Notes:

1. Statistics refer to all the call-backs that occurred during our pre-estimation sample (January - August 2018).
2. Panel a lists the twenty largest job titles (by call-backs) that sent all their call-backs to women.
3. Panel b lists the twenty job titles with the highest share of call-backs issued to women 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.

The results of these calculations are reported in Table 3.D of the paper.

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 A3.5.1 and A3.5.2. Overall, we find considerable support for the skill-specificity hypothesis: 23, 34 and 49 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 A3.5.1 Ads Requesting Men with High Incumbent Male Shares: Titles with the Most Vacancies

	(1)	(2)	(3)	(4)	(5)
	Job Title	English Translation	Number of Vacancies	Share of Vacancies	Cumulative Share of Vacancies
1	保安	Security	544	3.71	3.71
2	普工	General labor	471	3.21	6.93
3	仓管员	Warehouse management staff	448	3.06	9.98
4	施工员	Construction worker	343	2.34	12.32
5	焊工	Welder	302	2.06	**14.39**
6	电工	Electrician	288	1.97	16.35
7	仓管理	Warehouse management staff	283	1.93	18.28
8	司机	Driver	273	1.86	20.14
9	项目经理	Project manager	236	1.61	21.76
10	装配钳工	Assembly fitter	209	1.43	**23.18**
11	保安员	Security	200	1.36	24.55
12	电气工程师	Electronic engineer	195	1.33	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	29.10
16	注塑技术员	Injection molding technician	126	0.86	29.96
17	注塑领班	Injection molding team leader	125	0.85	30.82
18	操作工	Operator	124	0.85	31.66
19	土建施工员	Civil construction worker	107	0.73	32.39
20	模具钳工	Mold fitter	106	0.72	**33.12**

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 14,654 vacancies here.
2. Calculations reflect the fact that job ads may have multiple vacancies.

Table A3.5.2 Ads Requesting Women with High Incumbent Female Shares: Titles with the Most Vacancies

	(1)	(2)	(3)	(4)	(5)
	Job Title	English Translation	Number of Vacancies	Share of Vacancies	Cumulative Share of Vacancies
1	文员	Office clerk	543	6.10	6.10
2	会计	Accountant	451	5.07	11.17
3	出纳	Cashier	420	4.72	15.90
4	行政前台	Administration and receptionist	347	3.90	19.80
5	人事专员	HR assistant	279	3.14	**22.93**
6	人事行政专员	HR administrative specialist	249	2.80	25.73
7	财务	Finance	206	2.32	28.05
8	行政助理	Administration assistant	183	2.06	30.11
9	行政专员	Administration specialist	176	1.98	32.09
10	主办会计	Chief accountant	158	1.78	**33.86**
11	行政文员	Administrative clerk	155	1.74	35.60
12	前台文员	Receptionist and clerk	150	1.69	37.29
13	人事行政	HR administrative	150	1.69	38.98
14	财务助理	Financial assistant	144	1.62	40.61
15	护士	Nurse	141	1.59	42.18
16	外贸业务助理	International trader	137	1.54	43.72
17	商务助理	Business assistant	132	1.48	45.21
18	前台	Front desk	129	1.45	46.66
19	收银员	Cashier	129	1.45	48.11
20	育婴师	Baby sitter	121	1.36	**49.47**

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 8,895 vacancies here.
2. Calculations reflect the fact that job ads may have multiple vacancies.

A3.6 How Widespread Were the Ban's Integrating Effects?

If the integration of call-back pools after the gendered ad ban was mostly confined to a handful of the firms or job titles, we might be concerned that some idiosyncratic feature of our job board is responsible for the large and asymmetric effects we observed. To see if this was the case, Figure 3.6.1 shows cumulative distribution functions of the female share of call-backs before versus after the gendered ad ban, separately for three types of job ads: *F, N* and *M*. For each of these job types, we calculate the female shares at three different levels of aggregation: the job ad, the firm, and the job title. Consistent with our regression results (which showed a small and insignificant increase in the female share in *N* jobs), the post-ban cdfs of *N* jobs' female shares are slightly to the left of the pre-ban cdfs (panel a). Our main interest, however, is in *M* and *F* ads, where our statistical analysis demonstrated substantial causal effects of the ban in the direction of gender integration.

Turning next to *F* jobs (i.e. the jobs that were opened up to men by the ban; panel b), we first note that men's initial representation in these jobs was quite low: only 17.37, 21.03, and 32.94 percent of female-requesting job ads, firms, and job titles called back any men before the ban. After the ban, these cdfs shift dramatically: now 36.44, 41.98, and 55.63 percent of ads, firms and titles called back at least one man. Second, the pre-post differences are now much larger, with all the cdfs shifting far to the right. This dramatically illustrates the substantial inroads men made into 'women's jobs after the ban.

For *M* jobs (i.e. the jobs that were opened up to women by the ban), panel (c) of Figure 3.6.1 shows, first of all, that women's pre-ban representation here was much lower than men's in *F* jobs: only 10.07, 15.14, and 14.98 percent of job ads, firms, and job titles called back any women to jobs that requested men before the ban. Consistent with our regression results, after the ban these shares rose to 18.77, 29.07 and 24.72 percent.

To show the distribution of *changes* in men's and women's representation in *M* and *F* jobs associated with the ban more clearly, Figure 3.6.2 presents cdfs of the *changes* in women's representation at the title, firm and ad level. Focusing first on the *F* jobs, we see first that a large share of ads, firms and titles (56.36, 49.11 and 40.70 percent respectively) experienced no change in male representation at all.[12] That noted, 34.00, 40.99 and 47.51 percent of job ads, firms, and titles or 1,036, 827, 571 individual ads, firms, and titles -- experienced an increase in men's share of call-backs (i.e. a decline in women's share). The share experiencing declines in men's shares was much smaller and never more than 12 percent (right hand side of panel b). Thus, the increase in men's representation in jobs that previously requested women was widespread, across a large number of job ads, firms, and job titles.

Turning now to the *M* jobs (panel c) an even larger share of ads, firms, and titles (77.01, 64.25, and 71.54 percent) experienced no change in women's representation at all; this suggests that the very

[12] Especially at the ad level, this is because the absolute number of call-backs per unit is frequently quite small; thus is it not unusual for the share to be unchanged from one period to the next.

large share of *M* jobs that hired no women before the ban (Figure A3.6.1, panel c) continued to do so afterwards. Still, women did gain some access, with 16.24, 26.41 and 20.20 percent of ads, firms, and titles (or 452, 427, 318 individual ads, firms, and job titles) experiencing increases in women's representation. The shares experiencing decreases were much smaller and never more than 10 percent (left hand side of panel c). Thus, as was the case for men, women's representation in male-requesting ads was broadly based, across job, ads, firms and job titles.

In sum, Figures A3.6.1 and A3.6.2 show that the increases in men's representation in jobs that had requested women, and the increases in women's representation in jobs that had requested men, were quite broadly spread across the 21,757 job ads, 6,309 firms, and 8,624 job titles in our data.

Figure A3.6.1 Cumulative Distribution Functions of Female Share of Call-Backs

a. Non-Gendered (N) Jobs

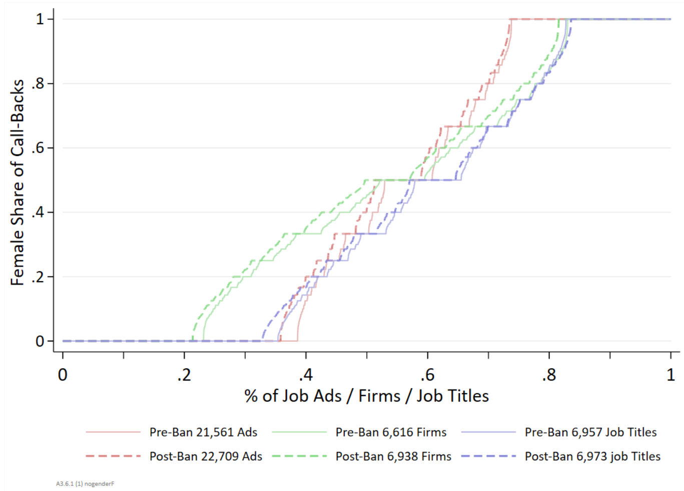

b. Female (F) Jobs

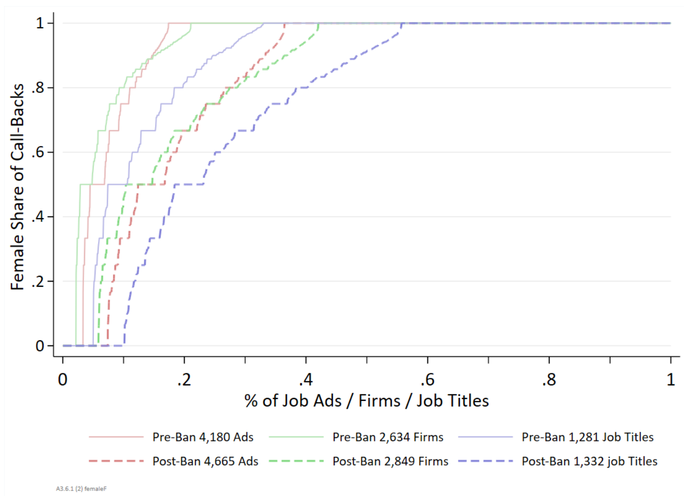

c. Male (M) Jobs

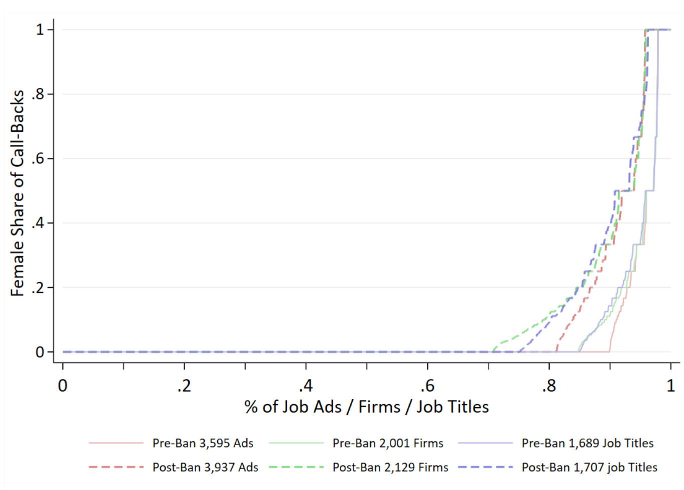

Figure A3.6.2 Cumulative Distribution Functions of the *Increase* in the Female Share of Call-Backs

a. Non-Gendered (N) Jobs

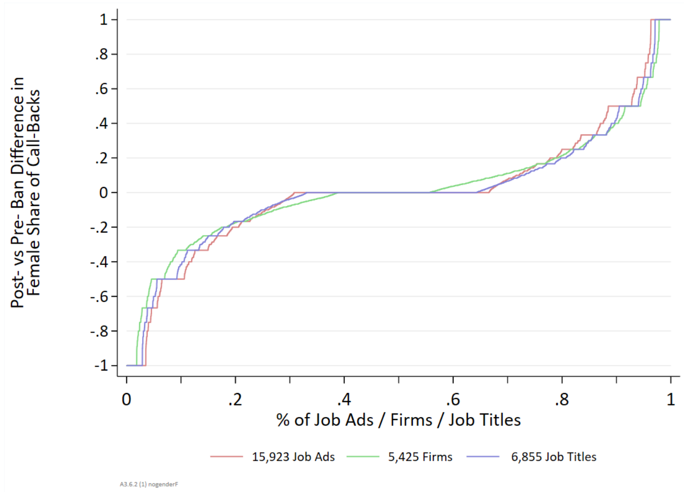

b. Female (F) Jobs

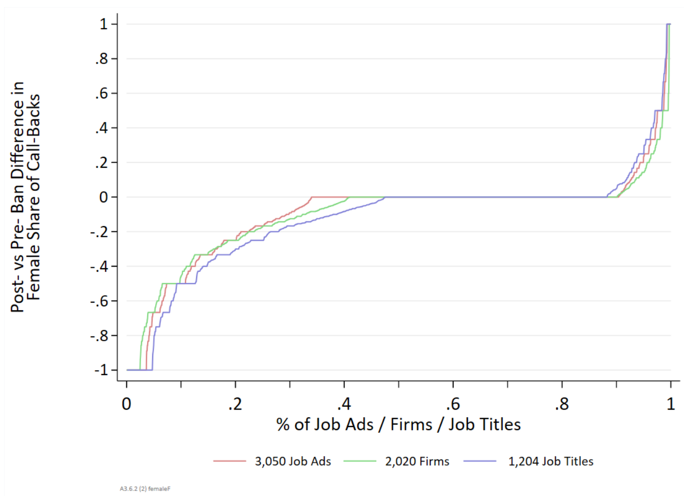

c. Male (M) Jobs

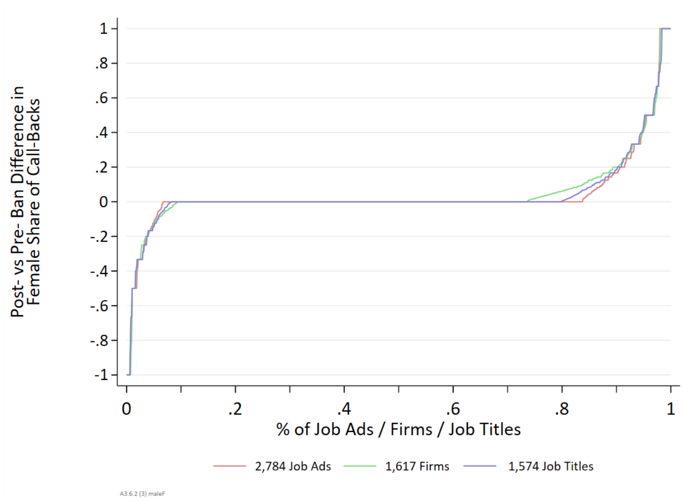

References for Appendix 3

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* 88(5): 2210-38.

Appendix 4: Local Linear Regression Plots

The figures in this Appendix graph the main results of the local linear regression analyses in column 2 of Tables 4 -- 6. 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

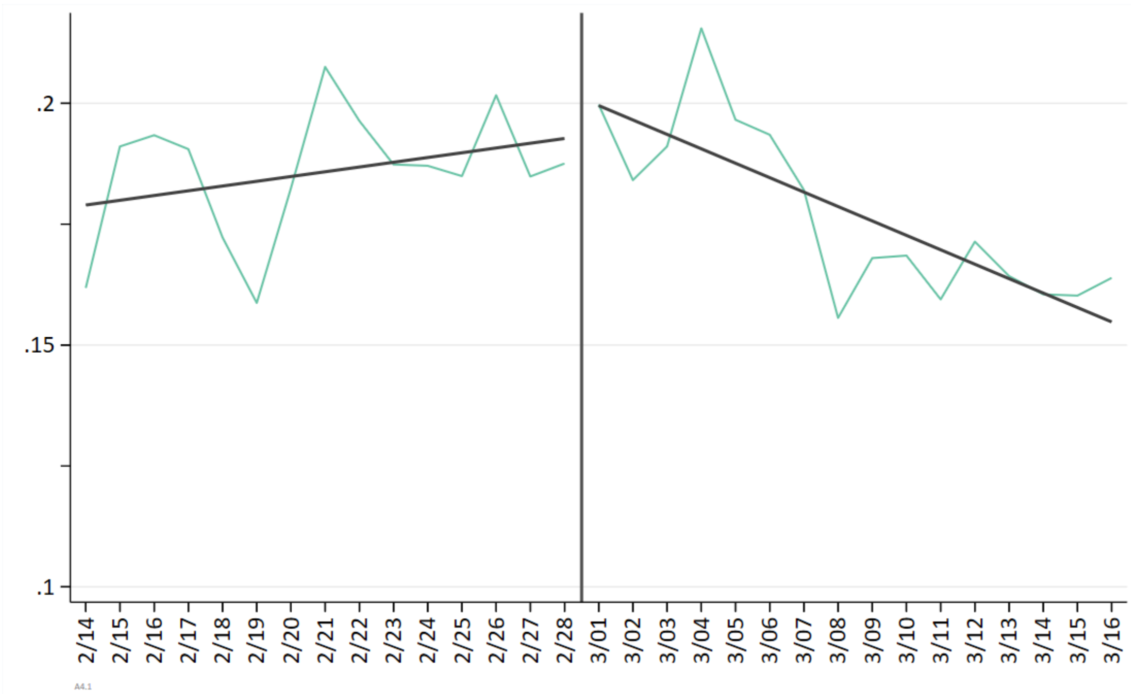

Notes:

1. The vertical line separates the pre and post gendered-ad ban periods, which is just before March 1, 2019.
2. Time series are adjusted for day-of-week and ad fixed effects, a linear trend in the age of the job ad, plus dummies for each of the first three days of an ad's life.
3. The figure is based on the local linear regression coefficients in column 2 of Table 4, where the post-ban effect is 0.0058*** (0.0015).

Figure A4.2: Daily 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, which is just before March 1, 2019.
3. Time series are adjusted for day-of-week and ad fixed effects, a linear trend in the age of the job ad, plus dummies for each of the first three days of an ad's life.
4. Local Linear Regression coefficients correspond to column 2 of Table 5, where the post-ban coefficient is 0.0169*** (0.0052).

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

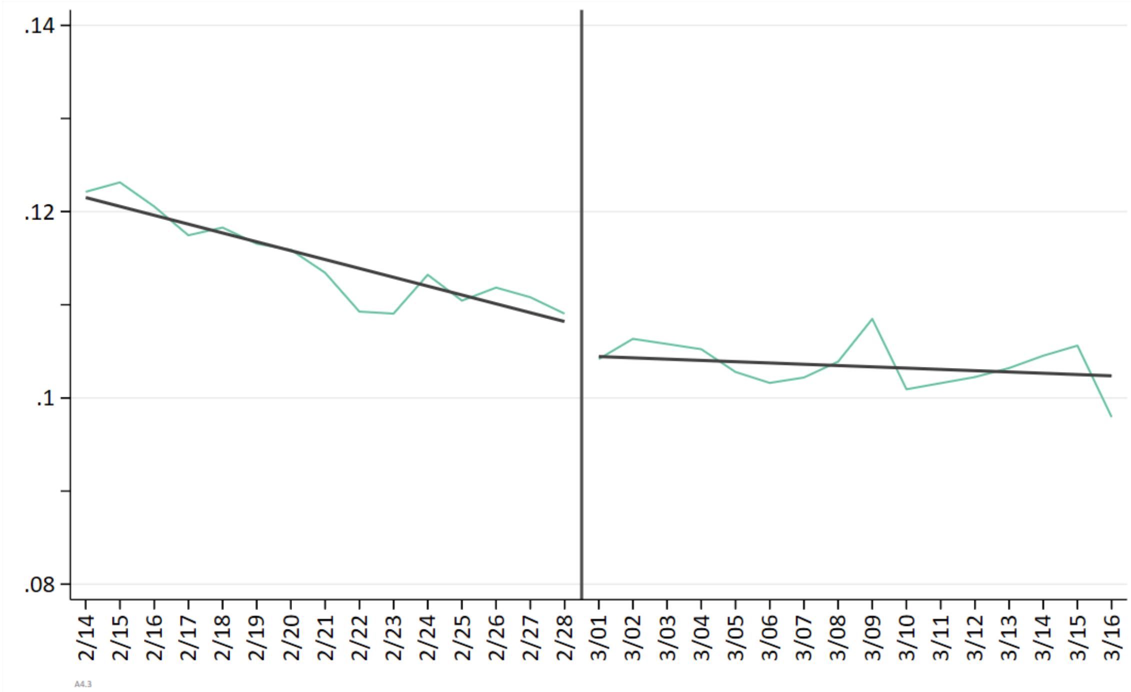

Notes:

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, a linear trend in the age of the job ad, plus dummies for each of the first three days of an ad's life.
3. Local Linear Regression coefficients correspond to column 2 of Table 6, where the post-ban coefficient is –0.0028 (0.0026).

Appendix 5: Robustness to Functional Form Assumptions

In Appendix 5, we assess the sensitivity of our main regression results (Tables 1 and 2) to the use of quartics to control for time trends and ad age effects. 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 in calendar time by a variety of other specifications, including separate quartics on either side of the ban. The results are also remarkably stable if we replace the quartic in *ad age* by fixed effects for each week of age. In Appendix A5.3 we provide descriptive evidence on the flow of applications during the first few days of an ad's life, to motivate our choice of functional form for the daily application flow regressions.

A5.1 Alternative Specifications of the *Calendar 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 A5.1.1 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 A5.1.2 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 A5.1.1: Effects of the Gendered Ad Ban on the Female Share of Applications

	(1)	(2)	(3)	(4)	(5)
	Quartic in calendar weeks	5th-order polynomial in calendar weeks	6th-order polynomial in calendar weeks	Separate quadratics before and after ban	Separate quartics before and after ban
Post ban week × Female job	–0.1393***	–0.1391***	–0.1390***	–0.1391***	–0.1390***
	(0.0023)	(0.0023)	(0.0023)	(0.0023)	(0.0023)
Post ban week × Male job	0.0348***	0.0349***	0.0349***	0.0349***	0.0350***
	(0.0018)	(0.0018)	(0.0018)	(0.0018)	(0.0018)
Post ban week	0.0053***	–0.0026**	–0.0029**	0.0000	–0.0050**
	(0.0011)	(0.0012)	(0.0012)	(0.0012)	(0.0020)
Quartic in job weeks	Y	Y	Y	Y	Y
Job ad fixed effects	Y	Y	Y	Y	Y
Effective # of obs	1,428,768	1,428,768	1,428,768	1,428,768	1,428,768
R^2	0.647	0.647	0.647	0.647	0.647

Note:

1. Column 1 reproduces column 4 of Table 1. The remaining columns change only the modeling of the calendar time trend, as described.
2. * p<.10, ** p<.05, *** p<.01.

Table A5.1.2: Effects of the Gendered Ad Ban on the Female Share of Call-Backs

	(1)	(2)	(3)	(4)	(5)
	Quartic in calendar weeks	5th-order polynomial in calendar weeks	6th-order polynomial in calendar weeks	Separate quadratics before and after ban	Separate quartics before and after ban
Post ban week × Female job	–0.1039***	–0.1039***	–0.1037***	–0.1038***	–0.1037***
	(0.0062)	(0.0062)	(0.0062)	(0.0062)	(0.0062)
Post ban week × Male job	0.0246***	0.0247***	0.0247***	0.0247***	0.0248***
	(0.0048)	(0.0048)	(0.0048)	(0.0048)	(0.0048)
Post ban week	0.0048	0.0031	–0.0014	0.0001	–0.0036
	(0.0040)	(0.0043)	(0.0043)	(0.0042)	(0.0071)
Quartic in job weeks	Y	Y	Y	Y	Y
Job ad fixed effects	Y	Y	Y	Y	Y
Effective # of obs	214,585	214,585	214,585	214,585	214,585
R^2	0.691	0.691	0.691	0.691	0.691

Note:

1. Column 1 reproduces column 4 of Table 2. The remaining columns change only the modeling of the calendar time trend, as described.
2. * p<.10, ** p<.05, *** p<.01.

A5.2 Alternative Specifications of the *Ad Age* Effect: Gender Mix Regressions

A well-known feature of job ads is that they tend to receive a flurry of attention right after they are posted, followed by a period of much more moderate activity. While it is not clear that such precipitous declines in activity should affect our main outcomes of interest -- the *gender mix* of applications and call-backs -- it still seems reasonable to ask whether the quartic we have used to control for ad age effects adequately captures this pattern. To answer this question, Tables A5.2.1 and A5.2.2 replicate our main regression tables (Tables 1 and 2), replacing the quartic in ad age by a full set of ad age fixed effects: one for each week of an ad's life.

Strikingly, the coefficients in these tables are almost identical to Tables 1 and 2. Intuitively, this occurs for two reasons. First, even though application and call-back activity decline sharply in the first few days of an ad's life, these changes in total activity do not affect the gender mix of applications or call-backs. Second, as we show in Appendix A5.3, all these sharp declines in activity occur during the first three *days* of an ad's life. In weekly data, they therefore manifest themselves as a much more modest decline between an ad's first and second week. These declines are easily captured by a quartic, and even simpler functional forms.

Table A5.2.1: Effects of the Gendered Ad Ban on the Female Share of Applications

	(1)	(2)	(3)	(4)	(5)
	Quartic in calendar weeks	5th-order polynomial in calendar weeks	6th-order polynomial in calendar weeks	Separate quadratics before and after ban	Separate quartics before and after ban
Post ban week × Female job	–0.1394***	–0.1392***	–0.1392***	–0.1392***	–0.1391***
	(0.0022)	(0.0022)	(0.0022)	(0.0022)	(0.0022)
Post ban week × Male job	0.0349***	0.0350***	0.0350***	0.0350***	0.0351***
	(0.0017)	(0.0017)	(0.0017)	(0.0017)	(0.0017)
Post ban week	0.0040***	–0.0037***	–0.0039***	–0.0016	–0.0058***
	(0.0011)	(0.0012)	(0.0012)	(0.0012)	(0.0019)
Job week fixed effects	Y	Y	Y	Y	Y
Job ad fixed effects	Y	Y	Y	Y	Y
Effective # of obs	1,423,319	1,423,319	1,423,319	1,423,319	1,423,319
R^2	0.646	0.646	0.646	0.646	0.646

Notes:

1. This table replaces the quartic in job ad age in Table 1 with a full set of ad age fixed effects (one for each week of an ad's life).
2. * *p*<.10, ** *p*<.05, *** *p*<.01.

Table A5.2.2: Effects of the Gendered Ad Ban on the Female Share of Call-Backs

	(1)	(2)	(3)	(4)	(5)
	Quartic in calendar weeks	5th-order polynomial in calendar weeks	6th-order polynomial in calendar weeks	Separate quadratics before and after ban	Separate quartics before and after ban
Post ban week × Female job	–0.1041***	–0.1041***	–0.1039***	–0.1040***	–0.1039***
	(0.0054)	(0.0054)	(0.0053)	(0.0054)	(0.0053)
Post ban week × Male job	0.0244***	0.0245***	0.0246***	0.0246***	0.0247***
	(0.0041)	(0.0041)	(0.0041)	(0.0041)	(0.0041)
Post ban week	0.0040	0.0014	–0.0029	–0.0013	–0.0051
	(0.0036)	(0.0038)	(0.0039)	(0.0038)	(0.0062)
Job week fixed effects	Y	Y	Y	Y	Y
Job ad fixed effects	Y	Y	Y	Y	Y
Effective # of obs	196,907	196,907	196,907	196,907	196,907
R^2	0.667	0.667	0.667	0.667	0.667

Notes:

1. This table replaces the quartic in job ad age in Table 2 with a full set of ad age fixed effects (one for each week of an ad's life).
2. * *p*<.10, ** *p*<.05, *** *p*<.01.

A5.3 Effects of *Ad Age* on Application Arrival Rates: Descriptive Statistics

While an ad's age may have only modest effects on the gender mix of applications in weekly data, ad age effects could be much more of a concern for some of our other analyses. Of greatest concern are the application arrival rate regressions in Tables 4, which are based on daily data surrounding the ban date. To explore this issue, Figure A5.3.1 displays the average daily number of applications (i.e. the outcome variable in Table 4) as a function of ad age, estimated from regressions on a full set of fixed effects for each day of an ad's age.

Importantly, Figure A5.3.1 suggests that after an ad has been posted for three days, the flow of applications is approximately linear, and in fact appears to be flat. For our daily arrival rate regressions (in Table 4), this suggests that fixed effects for the first three days of an ad's age (relative to all higher-order days) should be sufficient to capture the most important nonlinearities in application activity over the course of an ad's life. These considerations motivate our specification of ad age controls in the application arrival rate regressions of Table 4.

Figure A5.3.1: Number of Applications by Job Age, Calendar Date Controlled

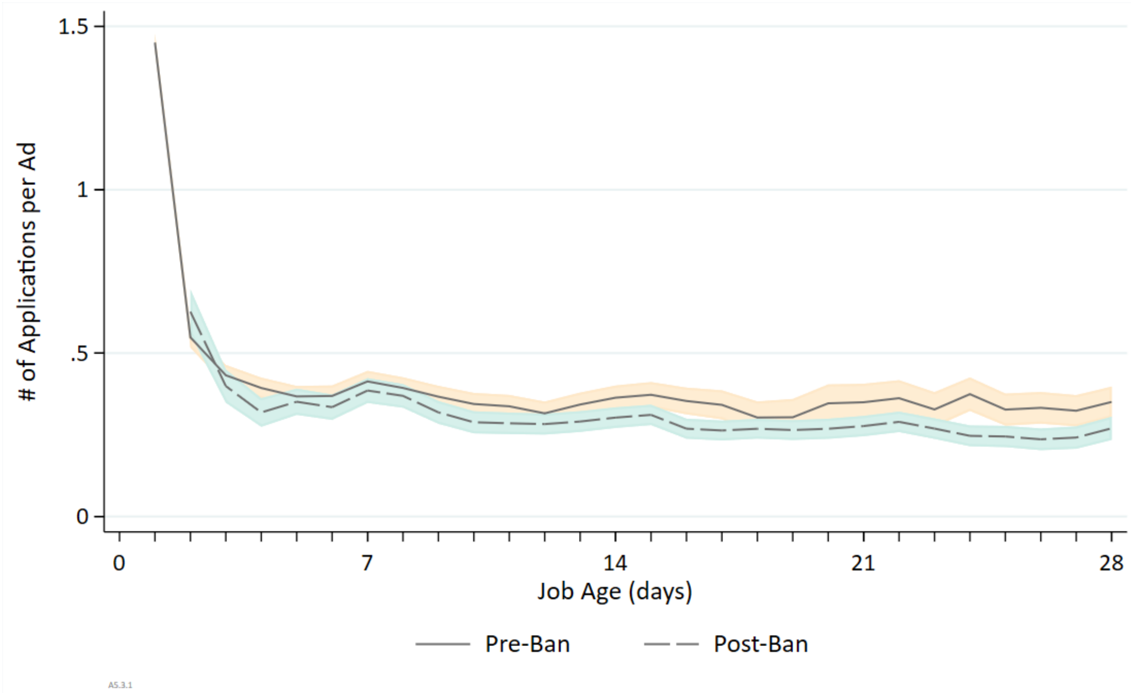

Notes:

1. Applications used are from Feb 14, 2019 to March 16, 2019, corresponding to our regression period.
2. The observation is calendar date × job age (days), with the number of active jobs as weight. A regression with calendar dates as fixed effect and weighted by active job postings and dummies for job age × post-ban dummy is estimated. The predicted number of applications by post-ban dummy is used for the graph.
3. The pre-ban data is adjusted to have the same mean as the post-ban data. This only affects the level but not the curvature of either line.
4. The maximum job age is 440 days in our data.

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 117,390 job ads, 3,529 or 3.006 percent had an embedded gender request. These embedded requests were evenly split between requests for men and women (1.509 percent and 1.497 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.[13] 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 nine percent larger than the ban's effects in all *M* job ads combined (3.51 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 15.07 percentage points. Consistent with the idea that these jobs were more intensely treated, this effect is about eight percent larger than the ban's effects in *F* job ads combined (13.89 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 15.07 – 8.12 = 6.95 percentage points. As Table A6.1

[13] 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.

indicates, this 8.12 percentage point (or 54 percent) reduction in the magnitude of the ban's effect is highly statistically significant.

Table A6.2 replicates Table A6.1 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 12 percent greater in magnitude. In *F* jobs, the ban's effects on the 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 8 and 12 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	0.0976***	0.0965***	0.0971***	0.0811***	0.0812***
	(0.0082)	(0.0081)	(0.0081)	(0.0051)	(0.0049)
Post ban × Male job × Embedded	–0.0271***	–0.0298***	–0.0290***	–0.0232***	–0.0229***
	(0.0076)	(0.0077)	(0.0077)	(0.0044)	(0.0042)
Post ban week × Female job	–0.1749***	–0.1739***	–0.1735***	–0.1511***	–0.1507***
	(0.0044)	(0.0044)	(0.0044)	(0.0025)	(0.0024)
Post ban week × Male job	0.0486***	0.0506***	0.0507***	0.0382***	0.0384***
	(0.0034)	(0.0035)	(0.0034)	(0.0019)	(0.0018)
Post ban week	0.0045***	0.0165***	–0.0012	0.0052***	
	(0.0018)	(0.0025)	(0.0021)	(0.0011)	
Female job	0.4862***	0.4855***	0.4852***		
	(0.0037)	(0.0037)	(0.0037)		
Male job	–0.3508***	–0.3499***	–0.3498***		
	(0.0036)	(0.0036)	(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	1,428,768	1,428,768	1,428,768	1,428,768	1,423,319
R^2	0.194	0.195	0.195	0.647	0.646

Notes:

1. See notes to Table 1 for regression specifications.
2. All regressions also include controls for Embedded, Embedded × Male Job, and Embedded × Female Job.
3. * p<.10, ** p<.05, *** p<.01.
4. 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 – 0.3508 = 5.82%. The raw male share of applications in female jobs pre-ban is 1 – (0.4090 + 0.4862) = 10.48%.

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***	0.0785***	0.0784***	0.0699***	0.0703***
	(0.0170)	(0.0171)	(0.0171)	(0.0128)	(0.0111)
Post ban × Male job × Embedded	–0.0248**	–0.0263**	–0.0262**	–0.0203*	–0.0206*
	(0.0115)	(0.0115)	(0.0117)	(0.0113)	(0.0099)
Post ban week × Female job	–0.1380***	–0.1370***	–0.1362***	–0.1152***	–0.1151***
	(0.0081)	(0.0080)	(0.0081)	(0.0066)	(0.0058)
Post ban week × Male job	0.0306***	0.0316***	0.0317***	0.0276***	0.0278***
	(0.0066)	(0.0066)	(0.0066)	(0.0051)	(0.0044)
Post ban week	0.0038	0.0154***	0.0034	0.0048	
	(0.0043)	(0.0053)	(0.0054)	(0.0040)	
Female job	0.4841***	0.4835***	0.4828***		
	(0.0057)	(0.0057)	(0.0058)		
Male job	–0.3995***	–0.3986***	-0.3985***		
	(0.0071)	(0.0071)	(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	214,585	214,585	214,585	214,585	196,907
R^2	0.219	0.220	0.220	0.691	0.667

Notes:

1. See notes to Table 2 for regression specifications.
2. All regressions also include controls for Embedded, Embedded × Male Job, and Embedded × Female Job.
3. * p<.10, ** p<.05, *** p<.01.
4. 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 –0.3995 = 4.79%. The raw male share of call-backs in female jobs pre-ban is 1 – (0.4474 + 0.4841) = 6.85%.

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 5)

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. Tables A7.2 and A7.3 replicate Tables 4 and 5 (application arrival rates and quality) on the call-back sample only. Again, the results are very similar, with one exception that reflects the difficulty of identifying spillover effects of the ban: The aggregate increase in the number of applications becomes insignificant, and we now see a negative 'spillover' effect of the ban on applications to *N* jobs. More importantly, the ban's effects on the number and quality of applications to directly treated (F and M) jobs are very similar to Tables 4 and 5.

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***	–0.1618***	–0.1612***	–0.1387***	–0.1384***
	(0.0053)	(0.0052)	(0.0052)	(0.0028)	(0.0027)
Post ban week × Male job	0.0460***	0.0472***	0.0474***	0.0342***	0.0345***
	(0.0041)	(0.0042)	(0.0042)	(0.0022)	(0.0021)
Post ban week	0.0024	0.0168***	–0.0026	0.0046***	
	(0.0023)	(0.0032)	(0.0026)	(0.0014)	
Female job	0.4749***	0.4742***	0.4736***		
	(0.0044)	(0.0044)	(0.0044)		
Male job	–0.3723***	–0.3712***	–0.3713***		
	(0.0043)	(0.0043)	(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	852,442	852,442	852,442	852,442	850,855
R^2	0.215	0.216	0.216	0.656	0.656

Notes:

1. See notes to Table 1 for regression specifications.
2. This table replicates Table 1 on the sub-sample of job ads for which call-backs are observed. 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.
3. The numbers of observations used for all specifications are 852,442. As a result of fixed effects specifications, as singleton observations were dropped, the effective numbers of observations are smaller in column (5).
4. * p<.10, ** p<.05, *** p<.01.
5. 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 – 0.3723 = 5.53%. The raw male share of applications in female jobs pre-ban is 1 – (0.4276 + 0.4749) = 9.75%.

Table A7.2: Effects of the Gendered Ad Ban on the Number of Applications Received, Call-Back Sample Only

	(1)	(2)	(3)	(4)
		Applications from:		
	All	All	Women	Men
1. To All Ads: [1,872,945 observations]				
Post ban	0.0011	0.0004	–0.0059***	0.0062***
	(0.0021)	(0.0021)	(0.0013)	(0.0015)
R^2	0.022	0.328	0.295	0.313
2. To Ads without Gender Request: [1,384,273 observations]				
Post ban	–0.0048**	–0.0055**	–0.0060***	0.0005
	(0.0023)	(0.0024)	(0.0015)	(0.0016)
R^2	0.023	0.325	0.293	0.287
3. To Ads that Requested Women: [253,082 observations]				
Post ban	0.0230***	0.0213***	–0.0221***	0.0434***
	(0.0055)	(0.0055)	(0.0048)	(0.0024)
R^2	0.024	0.300	0.288	0.201
4. To Ads that Requested Men: [235,590 observations]				
Post ban	0.0120*	0.0122**	0.0127***	–0.0005
	(0.0062)	(0.0061)	(0.0016)	(0.0058)
Job ad fixed effects		Y	Y	Y
R^2	0.016	0.363	0.202	0.366

Notes:

1. See notes to Table 4 for regression specifications.
2. This table replicates Table 4 on the sub-sample of job ads for which call-backs are observed. The dependent variable is the number of applications received. Its mean is 0.215, 0.196, and 0.221 for female, non-gendered, and male jobs, respectively.
3. * *p*<.10, ** *p*<.05, *** *p*<.01.

Table A7.3: Effects of the Gendered Ad Ban on the Match Quality of Applications, Call-Back Sample Only

	(1)	(2)	(3)	(4)
	Applications from:			
	All	All	Women	Men
1. To All Ads:				
Post ban	0.0154**	0.0219***	0.0042	0.0442***
	(0.0075)	(0.0067)	(0.0103)	(0.0097)
# of observations	257,274	257,274	123,280	155,176
R^2	0.018	0.599	0.649	0.638
2. To Ads without Gender Request:				
Post ban	0.0181**	0.0211***	–0.0078	0.0462***
	(0.0087)	(0.0079)	(0.0122)	(0.0114)
# of observations	187,385	187,385	88,796	116,060
R^2	0.021	0.614	0.675	0.649
3. To Ads that Requested Women:				
Post ban	–0.0006	0.0270	0.0370*	0.1224*
	(0.0194)	(0.0181)	(0.0190)	(0.0680)
# of observations	36,233	36,233	31,107	7,684
R^2	0.012	0.542	0.555	0.695
4. To Ads that Requested Men:				
Post ban	0.0165	0.0219	–0.0074	0.0265
	(0.0201)	(0.0184)	(0.0895)	(0.0192)
Job ad fixed effects?		Y	Y	Y
# of observations	33,656	33,656	3,377	31,432
R^2	0.027	0.565	0.736	0.569

Notes:

1. See notes to Table 5 for regression specification.
2. This table replicates Table 5 on the sub-sample of job ads for which call-backs are observed. The dependent variable is the mean match quality of applications. Its mean in the current estimation sample is –0.038, –0.019 and 0.096 for female, non-gendered, and male jobs, respectively.
3. * p<.10, ** p<.05, *** p<.01.

Appendix 8: Using Alternative Measures of Jobseekers' Success

In this Appendix we assess the robustness of our results to three alternative measures of an application's success. Appendix 8.1 replicates our main call-back regressions (Table 2) using an indicator for whether the application was read by the hiring agent as the outcome (Bartoš et al., 2016). While being read is not as selective as receiving a call-back, our 'read' indicator is available for essentially all applications. We find very similar results (except for the scale of the coefficients, which differs because reads are more common than call-backs).

To address the concern that receiving a call-back does not mean the worker received a job offer or was hired, Appendix 8.2 explores the sensitivity of our results to two other measures of workers' success. These two indicators use pauses in workers' application activity to infer likely hires. (For example, workers who stop looking shortly after a call-back are more likely to have been hired than workers who keep looking). Again, when Table 2 is replicated with these new measures, the results are very similar.

A8.1 Application Reads

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 Section 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.[14] 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 21 to 52 percent larger in magnitude) the results are highly similar.

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

Table A8.1.1: Effects of the Gendered Ad Ban on the Female Share of Reads

	(1)	(2)	(3)	(4)	(5)
Post ban week × Female job	–0.1537***	–0.1533***	–0.1530***	–0.1327***	–0.1323***
	(0.0041)	(0.0041)	(0.0041)	(0.0025)	(0.0024)
Post ban week × Male job	0.0407***	0.0426***	0.0426***	0.0317***	0.0321***
	(0.0034)	(0.0034)	(0.0034)	(0.0019)	(0.0018)
Post ban week	0.0073***	0.0181***	0.0012	0.0077***	
	(0.0020)	(0.0025)	(0.0024)	(0.0013)	
Female job	0.4897***	0.4893***	0.4890***		
	(0.0036)	(0.0036)	(0.0036)		
Male job	–0.3595***	–0.3590***	–0.3589***		
	(0.0036)	(0.0036)	(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	1,137,829	1,137,829	1,137,829	1,137,829	1,130,608
R^2	0.201	0.202	0.202	0.652	0.651

Notes:

1. See notes to Table 2 for regression specifications.
2. 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 (which correspond to 3,129,062 applications) compared to call-backs, the job ads of which correspond to 1,941,165 applications.
3. Observations are ad-week cells. All regressions are weighted by the number of reads. The dependent variable is the share of reads from female applicants. Its weighted mean is 81.2%, 41.9% and 8.5% for female, non-gendered, and male jobs, respectively.
4. The numbers of observations used for all specifications are 1,137,829. As a result of fixed effects specifications, singleton observations were dropped and the effective numbers of observations are smaller in column (5).
5. * p<.10, ** p<.05, *** p<.01.
6. 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 – 0.3595 = 5.52%. The raw male share of reads in female jobs pre-ban is 1 – (0.4147 + 0.4897) = 9.56%.

A8.2 Imputed Hires

Appendix 8.1 confirmed the validity of our call-back indicator by looking at a more expansive definition of workers' success (simply having one's application *read*). In this Section we construct two alternative measures of worker success ('imputed hires') from a completely different source of data -- time patterns in workers' application activity. One of these definitions (*Imputed Hires A*) is more restrictive than our main call-backs measure, because it defines 'hires' as the subset of call-backs whose corresponding applications occurred shortly before the end of a worker's search spell. The other (*Imputed Hires B*) -- while not necessarily more restrictive -- is of interest because it does not rely on our call-back data at all, by simply counting all applications that occur near the end of a search spell.

In more detail, both of our measures of imputed hires are based on a definition of workers' *search spells*, which we construct by connecting application dates together: any sequence of applications with less than *X* days between successive applications constitutes a search spell. Our first indicator of likely hires *(Imputed Hires A*) then defines an application as *successful* if it occurs less than *Y* days before the end of a search spell *and* received a call-back. In other words, *Imputed Hires A* are call-backs resulting from applications that were followed by a pause in application activity.

Imputed Hires A is, of course, an imperfect measure of hires because it could include some workers who are discouraged from searching, though we suspect this is unlikely because *all these imputed hires actually received a call-back*. *Imputed hires A* also assumes arbitrary values of *X* and *Y*, but Table A8.2.2 assesses the robustness of our results to alternative values of both and find little effect.

Imputed Hires B is the same as *A* except that it does not condition on receiving a call-back. Thus, all applications that occur within *Y* days of the end of a search spell are classified as hires. While this is more likely to include discouraged searchers in our count of hires, it has the countervailing advantage that *it does not rely on our indicator of call-backs at all*. Thus, it provides a more independent check on the validity of our call-back indicator.

As a point of reference, panel a of Table A8.2.1 re-displays Table 2 of the paper -- our main results for the effects of the ban on the female share of call-backs. Panel b then replicates Table 2 using *Imputed Hires A* as the outcome variable instead of call-backs. The results are highly similar. As noted, Table A8.2.2 tests the robustness of those panel b results, using different values of *X* and *Y* to construct search spells and define a successful application. Once again, the results are very similar.

Finally, Tables A8.2.3 and A8.2.4 replicate A8.2.1 and A8.2.1 using *Imputed Hires B* as the outcome variable instead of *Imputed Hires A.* Once again, the results are very similar, despite the fact that *Imputed Hires B* does not rely on any of our data on call-backs (which as noted are only available for 62 percent of applications).

Together, Tables A8.2.1 - A8.2.4 substantially increase our confidence that our call-back indicator is an accurate indicator of a worker's success in the recruiting process, and that our call-back indicator is strongly positively correlated with actual hires.

Table A8.2.1: Effects of the Gendered Ad Ban on the Female Share of *Imputed Hires A*

	(1)	(2)	(3)	(4)	(5)
a. Female Share of Call-Backs					
Post ban week × Female job	–0.1270***	–0.1263***	–0.1255***	–0.1039***	–0.1037***
	(0.0079)	(0.0079)	(0.0079)	(0.0062)	(0.0054)
Post ban week × Male job	0.0273***	0.0280***	0.0282***	0.0246***	0.0248***
	(0.0064)	(0.0064)	(0.0064)	(0.0048)	(0.0041)
Post ban week	0.0038	0.0152***	0.0035	0.0048	
	(0.0043)	(0.0052)	(0.0054)	(0.0040)	
Female job	0.4848***	0.4846***	0.4838***		
	(0.0057)	(0.0057)	(0.0057)		
Male job	–0.3995***	–0.3986***	–0.3986***		
	(0.0070)	(0.0071)	(0.0071)		
Effective # of obs	214,585	214,585	214,585	214,585	196,907
R^2	0.219	0.219	0.220	0.691	0.667
b. Female Share of *Imputed Hires A* hires = *call-backs* whose applications were at most **7** days before the end of an application spell defined by **X=7**.					
Post ban week × Female job	–0.1135***	–0.1128***	–0.1119***	–0.0925***	–0.0925***
	(0.0080)	(0.0079)	(0.0079)	(0.0071)	(0.0060)
Post ban week × Male job	0.0266***	0.0271***	0.0273***	0.0259***	0.0259***
	(0.0070)	(0.0071)	(0.0071)	(0.0060)	(0.0051)
Post ban week	0.0033	0.0138**	0.0042	0.0083	
	(0.0046)	(0.0055)	(0.0060)	(0.0051)	
Female job	0.4781***	0.4778***	0.4772***		
	(0.0059)	(0.0059)	(0.0059)		
Male job	–0.4126***	–0.4117***	–0.4117***		
	(0.0074)	(0.0075)	(0.0075)		
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	159,179	159,179	159,179	159,179	141,463
R^2	0.210	0.211	0.211	0.689	0.655

Notes:

1. See notes to Table 2 for regression specifications. Panel a reproduces Table 2, as a benchmark for panel b.
2. Panel b replicates Table 2, replacing call-backs by *Imputed Hires A* as an indicator of success. Only job ads with a positive number of imputed hires A are used here. *Imputed Hires A* are applications that result in a call-back, and that were made fewer than Y days before the end of an application spell. Application spells are defined as strings of applications with no more than X days between any pair of adjacent applications. Here panel b sets X = Y = 7.
3. Observations are ad-week cells in panel b. All regressions are weighted by the number of imputed hires A. The dependent variable in panel b is the share of imputed hires A that go to female applicants. Its weighted mean is 87.5%, 46.5% and 6.8% for female, non-gendered, and male jobs, respectively.
4. * p<.10, ** p<.05, *** p<.01.
5. The raw, pre-ban female share of imputed hires A in non-gendered jobs is 46.29%; the raw, pre-ban female share of imputed hires A in male jobs is 5.02%. The raw, pre-ban male share of imputed hires A in female jobs is 5.91%.

Table A8.2.2: Effects of the Gendered Ad Ban on the Female Share of *Imputed Hires A*, Robustness to Parameters

	(1)	(2)	(3)	(4)	(5)
a. hires = *call-backs* whose applications were at most **7** days before till the end of an application spell defined by **X=14**					
Post ban week × Female job	–0.1120***	–0.1113***	–0.1105***	–0.0898***	–0.0898***
	(0.0081)	(0.0081)	(0.0081)	(0.0074)	(0.0061)
Post ban week × Male job	0.0283***	0.0287***	0.0288***	0.0290***	0.0288***
	(0.0072)	(0.0072)	(0.0072)	(0.0070)	(0.0058)
Post ban week	0.0028	0.0139**	0.0060	0.0083	
	(0.0048)	(0.0057)	(0.0065)	(0.0060)	
Female job	0.4739***	0.4736***	0.4730***		
	(0.0061)	(0.0061)	(0.0061)		
Male job	–0.4190***	–0.4179***	–0.4178***		
	(0.0074)	(0.0074)	(0.0075)		
Effective # of obs	131,472	131,472	131,472	131,472	113,794
R^2	0.205	0.205	0.206	0.693	0.650
b. hires = *call-backs* whose applications were at most **7** days before till the end of an application spell defined by **X=28**					
Post ban week × Female job	–0.1142***	–0.1136***	–0.1128***	–0.0924***	–0.0928***
	(0.0080)	(0.0080)	(0.0080)	(0.0082)	(0.0066)
Post ban week × Male job	0.0258***	0.0262***	0.0264***	0.0264***	0.0262***
	(0.0074)	(0.0075)	(0.0075)	(0.0080)	(0.0065)
Post ban week	0.0066	0.0194***	0.0098	0.0107	
	(0.0050)	(0.0058)	(0.0067)	(0.0066)	
Female job	0.4723***	0.4720***	0.4714***		
	(0.0063)	(0.0063)	(0.0063)		
Male job	–0.4227***	–0.4215***	–0.4215***		
	(0.0076)	(0.0076)	(0.0076)		
Effective # of obs	112,587	112,587	112,587	112,587	95,210
R^2	0.200	0.201	0.201	0.699	0.648
c. hires = *call-backs* whose applications were at most **14** days before till the end of an application spell defined by **X=7**					
Post ban week × Female job	–0.1170***	–0.1164***	–0.1156***	–0.0976***	–0.0975***
	(0.0078)	(0.0078)	(0.0077)	(0.0066)	(0.0056)
Post ban week × Male job	0.0295***	0.0300***	0.0302***	0.0270***	0.0271***
	(0.0067)	(0.0067)	(0.0067)	(0.0052)	(0.0045)
Post ban week	0.0023	0.0137**	0.0023	0.0040	
	(0.0044)	(0.0054)	(0.0056)	(0.0045)	
Female job	0.4794***	0.4792***	0.4784***		
	(0.0058)	(0.0058)	(0.0058)		
Male job	–0.4098***	–0.4088***	–0.4088***		
	(0.0073)	(0.0073)	(0.0073)		
Effective # of obs	189,423	189,423	189,423	189,423	171,735
R^2	0.215	0.215	0.216	0.689	0.661

Table A8.2.2, continued:

d. hires = *call-backs* whose applications were at most **14** days before till the end of an application spell defined by **X=14**					
Post ban week × Female job	–0.1148***	–0.1142***	–0.1133***	–0.0919***	–0.0918***
	(0.0077)	(0.0077)	(0.0077)	(0.0067)	(0.0056)
Post ban week × Male job	0.0311***	0.0317***	0.0318***	0.0302***	0.0302***
	(0.0067)	(0.0067)	(0.0067)	(0.0059)	(0.0050)
Post ban week	0.0013	0.0139**	0.0009	0.0025	
	(0.0046)	(0.0055)	(0.0059)	(0.0050)	
Female job	0.4754***	0.4751***	0.4743***		
	(0.0059)	(0.0059)	(0.0059)		
Male job	–0.4171***	–0.4159***	–0.4158***		
	(0.0071)	(0.0072)	(0.0072)		
Effective # of obs	163,891	163,891	163,891	163,891	146,021
R^2	0.210	0.211	0.211	0.690	0.656
e. hires = *call-backs* whose applications were at most **14** days before till the end of an application spell defined by **X=28**					
Post ban week × Female job	–0.1165***	–0.1160***	–0.1152***	–0.0968***	–0.0969***
	(0.0077)	(0.0077)	(0.0077)	(0.0075)	(0.0062)
Post ban week × Male job	0.0284***	0.0290***	0.0292***	0.0266***	0.0268***
	(0.0069)	(0.0069)	(0.0069)	(0.0066)	(0.0055)
Post ban week	0.0053	0.0187***	0.0058	0.0076	
	(0.0047)	(0.0056)	(0.0061)	(0.0057)	
Female job	0.4739***	0.4737***	0.4729***		
	(0.0061)	(0.0061)	(0.0061)		
Male job	–0.4200***	–0.4187***	–0.4188***		
	(0.0074)	(0.0074)	(0.0074)		
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	140,438	140,438	140,438	140,438	122,699
R^2	0.204	0.205	0.205	0.692	0.652

Notes:

1. See notes to Table 2 and Table A8.2.1 for regression specifications.
2. * p<.10, ** p<.05, *** p<.01.
3. Panels a -- f explore two alternative values of Y (7 and 14) and three alternative values of X (7, 14, and 28). The dependent variables in panels a -- f is the share of imputed hires A that go to female applicants. Its weighted means and the pre-ban gender shares are listed in the following table,

Panel	Weighted mean of female shares of "hires" at job*week cells			Pre-ban "hires"		
	F jobs	N jobs	M jobs	Raw male share in F jobs	Raw female share in N jobs	Raw female share in M jobs
a	88.2%	47.7%	7.0%	5.42%	47.35%	5.08%
b	87.5%	46.8%	6.9%	5.77%	46.69%	4.98%
c	87.8%	47.1%	6.9%	5.66%	46.95%	5.05%
d	87.0%	46.1%	6.9%	6.11%	45.95%	4.97%
e	87.8%	47.4%	7.1%	5.53%	47.08%	5.08%

Table A8.2.3: Effects of the Gendered Ad Ban on the Female Share of *Imputed Hires B*

	(1)	(2)	(3)	(4)	(5)
(a) Female Share of Call-Backs					
Post ban week × Female job	–0.1270***	–0.1263***	–0.1255***	–0.1039***	–0.1037***
	(0.0079)	(0.0079)	(0.0079)	(0.0062)	(0.0054)
Post ban week × Male job	0.0273***	0.0280***	0.0282***	0.0246***	0.0248***
	(0.0064)	(0.0064)	(0.0064)	(0.0048)	(0.0041)
Post ban week	0.0038	0.0152***	0.0035	0.0048	
	(0.0043)	(0.0052)	(0.0054)	(0.0040)	
Female job	0.4848***	0.4846***	0.4838***		
	(0.0057)	(0.0057)	(0.0057)		
Male job	–0.3995***	–0.3986***	–0.3986***		
	(0.0070)	(0.0071)	(0.0071)		
Effective # of obs	214,585	214,585	214,585	214,585	196,907
R^2	0.219	0.219	0.220	0.691	0.667
(b) Female Share of Imputed Hires B hires = *applications* at most **7** days before till the end of an application spell defined by **X=7**.					
Post ban week × Female job	–0.1433***	–0.1426***	–0.1421***	–0.1229***	–0.1226***
	(0.0041)	(0.0040)	(0.0040)	(0.0025)	(0.0024)
Post ban week × Male job	0.0469***	0.0483***	0.0484***	0.0363***	0.0366***
	(0.0034)	(0.0035)	(0.0035)	(0.0020)	(0.0019)
Post ban week	0.0048***	0.0164***	–0.0012	0.0066***	
	(0.0019)	(0.0025)	(0.0022)	(0.0014)	
Female job	0.4880***	0.4875***	0.4872***		
	(0.0036)	(0.0036)	(0.0036)		
Male job	–0.3662***	–0.3652***	–0.3652***		
	(0.0037)	(0.0037)	(0.0037)		
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	1,116,929	1,116,929	1,116,929	1,116,929	1,108,451
R^2	0.183	0.184	0.184	0.619	0.617

Notes:

1. See notes to Table 2 and Table A8.2.1 for regression specifications.
2. This Table replicates Table A8.2.1, replacing imputed hires A by imputed hire B. An application is regarded as an imputed hire B if the application spell in which it is contained ends within less than Y days after the application. Application spells are defined as strings of applications with no more than X days between any pair of adjacent applications. Here, panel b sets X = Y = 7.
3. Observations are ad-week cells in panel b. All regressions are weighted by the number of imputed hires B. The dependent variable in panel b is the share of imputed hires B that go to female applicants. Its weighted mean is 82.9%, 42.7% and 8.9% for female, non-gendered, and male jobs, respectively.
4. * p<.10, ** p<.05, *** p<.01.
5. The raw, pre-ban female share of imputed hire B in non-gendered jobs is 42.39%; the raw, pre-ban female share of imputed hires B in male jobs is 5.78%. The raw, pre-ban male share of imputed hires B in female jobs is 8.80%.

Table A8.2.4: Effects of the Gendered Ad Ban on the Female Share of *Imputed Hires B*, Robustness to Parameters

	(1)	(2)	(3)	(4)	(5)
a. hires = *applications* at most **7** days before till the end of an application spell defined by **X=14**					
Post ban week × Female job	–0.1418***	–0.1411***	–0.1405***	–0.1220***	–0.1218***
	(0.0040)	(0.0040)	(0.0040)	(0.0027)	(0.0025)
Post ban week × Male job	0.0475***	0.0489***	0.0491***	0.0365***	0.0367***
	(0.0035)	(0.0035)	(0.0035)	(0.0022)	(0.0021)
Post ban week	0.0054***	0.0180***	–0.0009	0.0075***	
	(0.0019)	(0.0026)	(0.0024)	(0.0016)	
Female job	0.4869***	0.4863***	0.4859***		
	(0.0037)	(0.0037)	(0.0037)		
Male job	–0.3711***	–0.3700***	–0.3699***		
	(0.0037)	(0.0037)	(0.0037)		
Effective # of obs	941,604	941,604	941,604	941,604	930,736
R^2	0.176	0.177	0.178	0.608	0.604
b. hires = *applications* at most **7** days before till the end of an application spell defined by **X=28**					
Post ban week × Female job	–0.1426***	–0.1418***	–0.1411***	–0.1222***	–0.1220***
	(0.0040)	(0.0040)	(0.0040)	(0.0029)	(0.0027)
Post ban week × Male job	0.0476***	0.0491***	0.0492***	0.0373***	0.0375***
	(0.0036)	(0.0036)	(0.0036)	(0.0024)	(0.0022)
Post ban week	0.0075***	0.0213***	0.0010	0.0084***	
	(0.0020)	(0.0027)	(0.0025)	(0.0018)	
Female job	0.4859***	0.4851***	0.4846***		
	(0.0037)	(0.0037)	(0.0037)		
Male job	–0.3764***	–0.3751***	–0.3752***		
	(0.0038)	(0.0038)	(0.0038)		
Effective # of obs	811,949	811,949	811,949	811,949	798,906
R^2	0.172	0.173	0.173	0.603	0.598
c. hires = *applications* at most **14** days before till the end of an application spell defined by **X=7**					
Post ban week × Female job	–0.1501***	–0.1493***	–0.1489***	–0.1281***	–0.1278***
	(0.0040)	(0.0040)	(0.0040)	(0.0024)	(0.0023)
Post ban week × Male job	0.0469***	0.0483***	0.0485***	0.0365***	0.0367***
	(0.0034)	(0.0034)	(0.0034)	(0.0019)	(0.0018)
Post ban week	0.0048***	0.0172***	–0.0007	0.0051***	
	(0.0019)	(0.0025)	(0.0021)	(0.0012)	
Female job	0.4882***	0.4877***	0.4873***		
	(0.0036)	(0.0036)	(0.0036)		
Male job	–0.3634***	–0.3623***	–0.3623***		
	(0.0036)	(0.0037)	(0.0037)		
Effective # of obs	1,287,769	1,287,769	1,287,769	1,287,769	1,281,118
R^2	0.189	0.190	0.190	0.634	0.633

Table A8.2.4, continued:

d. hires = *applications* at most **14** days before till the end of an application spell defined by **X=14**					
Post ban week × Female job	–0.1485***	–0.1478***	–0.1472***	–0.1268***	–0.1264***
	(0.0039)	(0.0039)	(0.0039)	(0.0025)	(0.0024)
Post ban week × Male job	0.0486***	0.0500***	0.0502***	0.0378***	0.0380***
	(0.0034)	(0.0034)	(0.0034)	(0.0020)	(0.0019)
Post ban week	0.0045**	0.0174***	–0.0028	0.0038***	
	(0.0019)	(0.0025)	(0.0022)	(0.0014)	
Female job	0.4874***	0.4867***	0.4863***		
	(0.0036)	(0.0036)	(0.0036)		
Male job	–0.3690***	–0.3678***	–0.3678***		
	(0.0036)	(0.0037)	(0.0037)		
Effective # of obs	1,135,438	1,135,438	1,135,438	1,135,438	1,127,245
R^2	0.183	0.184	0.184	0.621	0.620
e. hires = *applications* at most **14** days before till the end of an application spell defined by **X=28**					
Post ban week × Female job	–0.1491***	–0.1484***	–0.1476***	–0.1273***	–0.1270***
	(0.0039)	(0.0039)	(0.0039)	(0.0026)	(0.0025)
Post ban week × Male job	0.0484***	0.0499***	0.0501***	0.0382***	0.0385***
	(0.0035)	(0.0035)	(0.0035)	(0.0021)	(0.0020)
Post ban week	0.0075***	0.0216***	0.0017	0.0069***	
	(0.0019)	(0.0026)	(0.0023)	(0.0015)	
Female job	0.4864***	0.4857***	0.4850***		
	(0.0037)	(0.0037)	(0.0037)		
Male job	–0.3738***	–0.3724***	–0.3725***		
	(0.0037)	(0.0037)	(0.0038)		
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	984,125	984,125	984,125	984,125	974,009
R^2	0.177	0.178	0.178	0.612	0.609

Notes:

1. See notes to Table 2 and Table A8.2.3 for regression specifications.
2. * p<.10, ** p<.05, *** p<.01.
3. Panels a -- f explore two alternative values of Y (7 and 14) and three alternative values of X (7, 14, and 28). The dependent variables in panels a -- f is the share of imputed hires B that go to female applicants. Its weighted means and the pre-ban gender shares are listed in the following table,

Panel	Weighted mean of female shares of "hires" at job*week cells			Pre-ban "hires"		
	F jobs	N jobs	M jobs	Raw male share in F jobs	Raw female share in N jobs	Raw female share in M jobs
a	84.1%	44.0%	9.2%	7.84%	43.57%	5.93%
b	82.9%	43.1%	9.1%	8.45%	42.81%	5.91%
c	83.5%	43.3%	9.1%	8.30%	43.01%	5.90%
d	82.2%	42.4%	8.9%	9.05%	42.13%	5.79%
e	83.4%	43.8%	9.3%	8.01%	43.35%	5.97%

Appendix 9: Robustness to the Length of the Estimation Window

Appendix 9 examines the robustness of our main regression results 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 -- 6. It shows that the former results (for the gender mix of applicant and call-back pools) are highly stable across application windows ranging from 3 to 51 weeks. We also apply Calonico, Cattaneo, and Titiunik's (CCT) selector to calculate MSE-optimal bandwidths for these regressions. The optimal windows fall roughly in the middle of the preceding grid searches and the coefficient estimates are similar.[15] We also explore the robustness of our results for application arrivals, quality and yield to increasing the post-ban estimation interval from 15 to 40 days.[16] While we find some instability in our estimates of aggregate application rates (to all jobs together), our results for aggregate quality and yield, and our estimates of the ban's effects on application rates to *M* and *F* jobs are also robust to window width.

A9.1 Exploring Alternative Estimation Windows: Gender Mix Regressions

In this Section 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 3 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 A9.1.1 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 women (*F* jobs) are very close to Table 1's estimate of –0.1393. The same remarkable stability holds for *F* jobs, though the coefficient for *N* jobs moves away from zero to become slightly positive in the very shortest windows. Figure A9.1.2 (for the female share of call-backs) exhibits the same remarkable stability for *F* and *M* jobs -- our main coefficients of interest -- and a slightly stronger trend towards positive estimates for the *N* jobs in the shortest windows. Given the clear visual message of Figure 1 in the paper, this robustness of our main results is perhaps not surprising.

[15] We were only able to estimate CCT-optimal bandwidths with a restricted set of control variables, relative to our preferred specification.

[16] We cannot increase the pre-ban window in these daily regressions due to the Spring Festival; we can't increase the post-ban window beyond 40 days due to the Qingming Festival.

Figure A9.1.1 Effects of the Ban on the Female Share of Applications, Robustness to Estimation Window Length

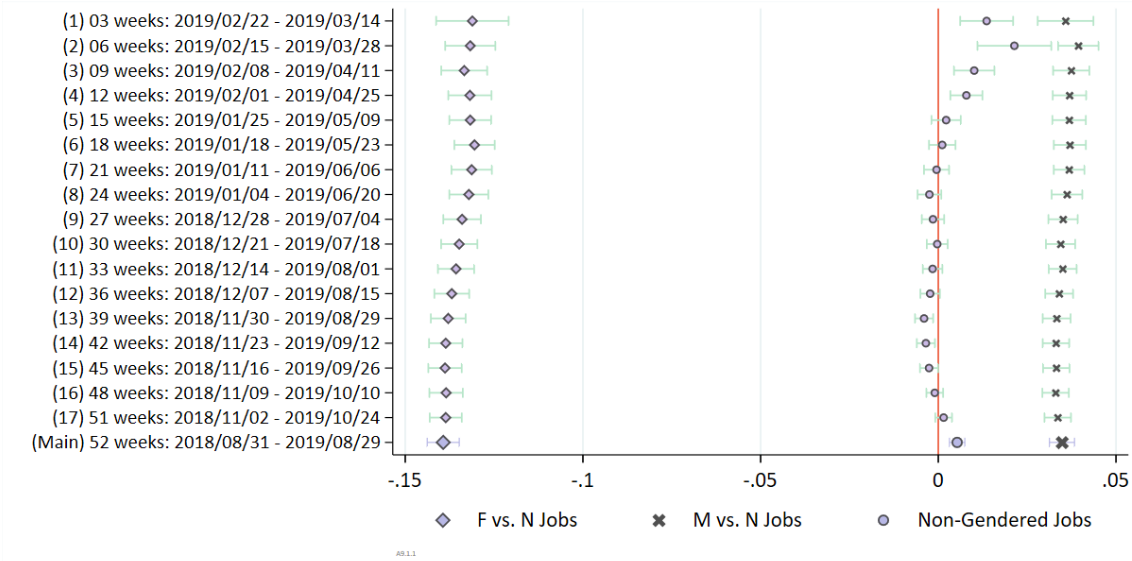

Figure A9.1.2 Effects of the Ban on the Female Share of Call-Backs, Robustness to Estimation Window Length

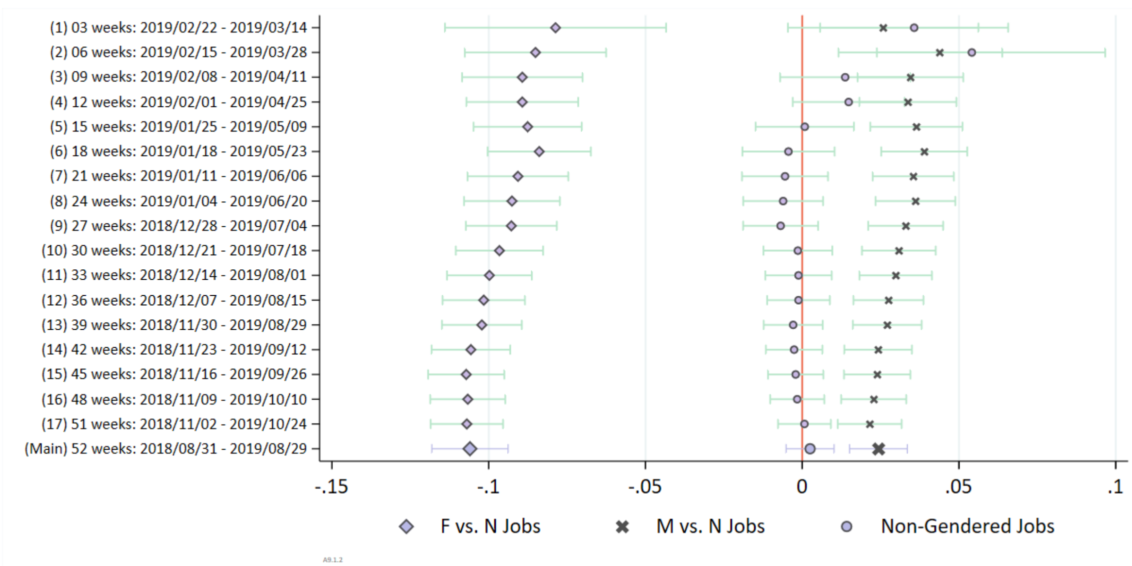

A9.2 Exploring Alternative Estimation Windows: Application Arrivals, Quality and Yield

In this Section we explore the sensitivity of our results for our three measures of search frictions -- application arrival rates, application match quality, and 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.

To keep the number of results manageable, we present two sets of results. First, Figures A9.2.1 -- A9.2.3 focus on the aggregate effects of the ban combining all jobs and workers together, i.e. on panel A of Tables 4 -- 6. Second, motivated by the insights gained about the mechanisms behind the ban's integrating effects from the changing arrival patterns to previously gendered jobs, Table A9.2.1 replicates panels C and D of Table 4 using alternative estimation windows.

For aggregate outcomes (Figures A9.2.1 -- A9.2.3), the results are somewhat mixed, but all are consistent with small or zero effects. For application arrival rates, the estimated effect of the ban first shifts from positive to negative as the estimation window lengthens, then recovers back to some zero and positive estimates in longer windows. For application quality, all the estimates show statistically significant increases that are very small in magnitude. For application yields, are the estimates are negative, small, and statistically insignificant.

In contrast, for disaggregated application flows (Table A9.2.1), all the new estimates are substantial in magnitude, highly statistically significant, and remarkably stable. Together, they dramatically confirm the following findings: The ban caused an increase in the number of applications to previously gendered jobs. These increases were driven exclusively by new applications from the previously excluded genders. Partially offsetting these increases, the ban also caused a decline in women's applications to previously female jobs. Because this decline was not paralleled by a decline in men's applications to previously male jobs, it may help explain why men gained so much more access to women's jobs than vice versa after the ban.

For disaggregated quality (Table A9.2.2), there are some changes in statistical significance as we lengthen the estimation window (the increase in quality of F-to-F applications becomes significant while the increase in M-to-F applications loses significance) but almost all of the estimated quality changes (21 of 24) are positive, and none are significantly negative. For application yield (Table A9.2.3), all of our new estimates, like the main ones, are small in magnitude and statistically insignificant. Together, these results confirm our conclusion that the ban did not reduce application quality or yield; if anything the mean match quality of applications received by previously gendered job ads rose.

Figure A9.2.1 Effects of the Ban on the Daily Flow of Applications, Robustness to Estimation Window Length

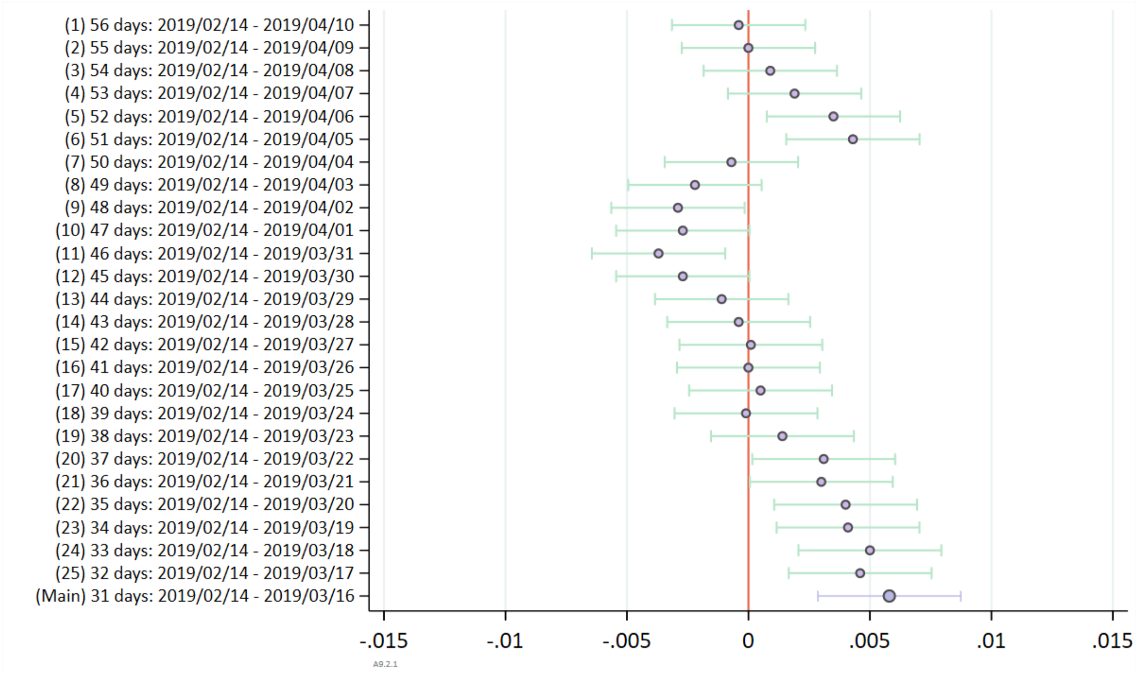

Notes:

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 A9.2.2 Effects of the Ban on the Match Quality of Applications, Robustness to Estimation Window Length

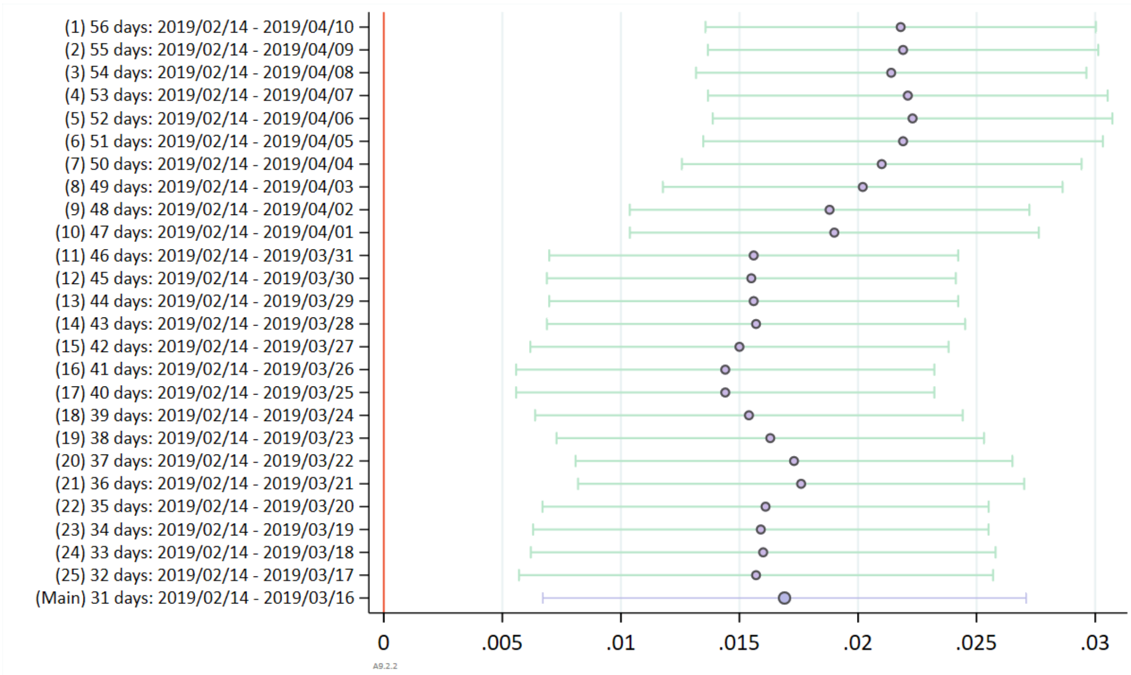

Notes:

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 5, column 2 is used.

Figure A9.2.3 Effects of the Ban on the Probability an Application Yields a Call-Back, Robustness to Estimation Window Length

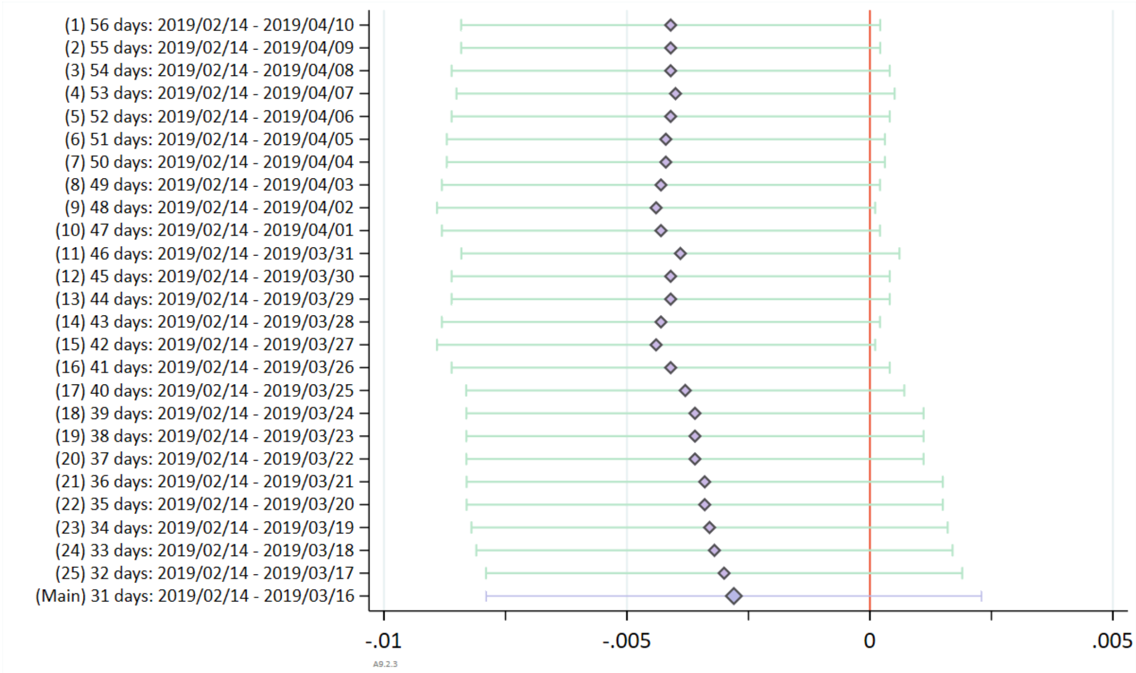

Notes:

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 6, column 2 is used.

Table A9.2.1 Effects of the Ban on Application Flows to *F* and *M* Jobs by Applicant Gender, Robustness to Estimation Window Width

	Applications from Women		Applications from Men	
	to *F* jobs	to M jobs	to *F* jobs	to *M* jobs
base (31 days)	–0.0160***	**0.0122*** **	**0.0388*** **	0.0043
35 days	–0.0164***	**0.0119*** **	**0.0384*** **	0.0012
40 days	–0.0163***	**0.0114*** **	**0.0370*** **	–0.0042
45 days	–0.0189***	**0.0111*** **	**0.0352*** **	–0.0079**
50 days	–0.0183***	**0.0114*** **	**0.0359*** **	–0.0058
55 days	–0.0175***	**0.0112*** **	**0.0365*** **	–0.0048

Notes:

1. This Table replicates panels C and D and columns 3 and 4 of Table 4 for alternative estimation window widths. See notes to Table 4 for regression specifications.
2. All windows include 15 days before the ban and 15 or more days after the ban.
3. * *p*<.10, ** *p*<.05, *** *p*<.01.
4. **Bolded** columns indicate gender-mismatched applications.

Table A9.2.2 Effects of the Ban on the Mean Quality of Applications to *F* and *M* Jobs by Applicant Gender, Robustness to Estimation Window Width

	Applications from Women		Applications from Men	
	to *F* jobs	to M jobs	to *F* jobs	to *M* jobs
base (31 days)	0.0268	**0.0217**	**0.1064***	0.0298**
35 days	0.0300*	**0.0075**	**0.0993***	0.0174
40 days	0.0403***	**0.0039**	**0.0974** **	0.0073
45 days	0.0425***	**0.0141**	**0.0798***	–0.0008
50 days	0.0407***	**–0.0019**	**0.0748**	0.0101
55 days	0.0430***	**–0.0034**	**0.0733***	0.0184

Notes:

1. This Table replicates panels C and D and columns 3 and 4 of Table 5 for alternative estimation window widths. See notes to Table 4 for regression specifications.
2. All windows include 15 days before the ban and 15 or more days after the ban.
3. * *p*<.10, ** *p*<.05, *** *p*<.01.
4. **Bolded** columns indicate gender-mismatched applications.

Table A9.2.3 Effects of the Ban on Application Yield in *F* and *M* Jobs by Applicant Gender, Robustness to Estimation Window Width

	Applications from Women		Applications from Men	
	to *F* jobs	to M jobs	to *F* jobs	to *M* jobs
base (31 days)	0.0017	**–0.0101**	**0.0004**	–0.0030
35 days	–0.0013	**–0.0195**	**–0.0078**	–0.0035
40 days	–0.0008	**–0.0281**	**0.0029**	–0.0032
45 days	–0.0027	**–0.0262**	**–0.0045**	–0.0016
50 days	–0.0031	**–0.0257**	**–0.0030**	–0.0006
55 days	–0.0038	**–0.0271**	**–0.0031**	–0.0001

Notes:
1. This Table replicates panels C and D and columns 3 and 4 of Table 6 for alternative estimation window widths. See notes to Table 4 for regression specifications.
2. All windows include 15 days before the ban and 15 or more days after the ban.
3. * p<.10, ** p<.05, *** p<.01.
4. **Bolded** columns indicate gender-mismatched applications.

A9.3 Estimates using CCT Optimal Bandwidths

Even though Appendix A9.1 has already shown that our main regression results in Tables 1 and 2 are remarkably stable to a grid search of the available range of estimation windows, it may be of interest to know which particular bandwidths are optimal in a well-defined sense. To answer this question, we applied Calonico, Cattaneo, and Titiunik's (CCT) selector to calculate MSE-optimal bandwidths to the regression framework underlying Tables 1 and 2. To accomplish this, we first disaggregate the week x ad level data we used in Tables 1 and 2 to the day x ad level. This is because RD methods in general require a continuous running variable (time in our context); reflecting this, our attempts to find an optimal bandwidth using weekly data did not yield a solution. As a result of this change, the dependent variable is now the female share of applications in each day x ad cell.

Next, since computations of optimal bandwidth become intractable while controlling for high-dimensional fixed effects, we restricted our search for optimal bandwidths to two specifications, neither of which controls for ad fixed effects. Finally, since the standard RD approach looks for discontinuities in a running variable without reference to a control group, we restricted our attention to the time series for *F* and *M* jobs and focused on identifying a discontinuity in those series on the ban date, using the MSE-optimal symmetric estimation window around the ban date.

The results (and the optimal bandwidths on which they are based) are shown in Table A9.3.1. For applications, the results are very similar to our estimates in columns 1-3 of Table 1 (the specifications that don't control for ad fixed effects), showing a decline in the female share in *F* jobs of about 15.6 percentage points in both cases. The increase in the female share in *M* jobs is 6.0 percentage points, compared to about 5.3 percentage points in Table 1. For call-backs, the results differ a little more, showing declines in the female share in *F* jobs of about 8 percentage points and increases in *M* jobs of about 1.6 percentage points (compared to 11.9 and 3.5 percentage points in Table 2). Some of the estimates for *M* jobs become statistically insignificant.

Since none of the estimates in Table A9.3.1 control for confounding effects we argued could be important, including changes in the mix of ads that are active during our estimation period, it is hard to directly compare Table A9.3.1 to our main estimates, and even to the 'grid search' results in Appendix 9.1, which do control for all those confounding effects. Nevertheless, we view the qualitative similarity of the results in Table A9.3.1 with our main estimates as additional confirmation of the main results' robustness.

Table A9.3.1 RD Treatment Effects on the Female Share of Applications and Call-Backs using CCT's MSE-Optimal Bandwidths

Covariates specified	Female jobs			Male jobs		
	Optimal bandwidth	RD treatment effect		Optimal bandwidth	RD treatment effect	
a. Applications						
(1) None	48.4	–0.1657***	(0.0038)	60.2	0.0587***	(0.0025)
(2) Day of the week dummies	47.6	–0.1656***	(0.0039)	73.1	0.0587***	(0.0023)
(3) (2)+ quartic in job days	47.0	–0.1657***	(0.0039)	73.2	0.0583***	(0.0023)
# of observations	276,970			315,021		
b. Call-Backs						
(1) None	40.5	-0.0769***	(0.0093)	31.9	0.0179*	(0.0093)
(2) Day of the week dummies	43.1	–0.0842***	(0.0090)	31.8	0.0145	(0.0093)
(3) (2)+ quartic in job days	53.8	–0.0902***	(0.0081)	31.8	0.0144	(0.0093)
# of observations	39,652			33,711		

* *p*<.10, ** *p*<.05, *** *p*<.01

Appendix 10: Placebo Bans

Appendix 10 probes the significance of our main results by estimating the effects of placebo bans in every possible week preceding the actual ban and comparing these estimates to our main results.[17] Turning first to our main results (Tables 1 and 2), our estimates of the ban's effects on gender mix in *F* and *M* jobs are dramatically different from all our placebo estimates. Mirroring the pattern in Appendix 9, our estimates of the ban's effects on aggregate application arrivals, quality and yield are not dramatically different from all the placebo estimates, suggesting that our statistical power to assess these outcomes is lower than our power to identify gender mix effects. That said, our estimates of the arrival rate of gender-mismatched applications differ dramatically from all the placebo estimates, suggesting strong power to identify the main mechanisms underlying the ban's effects.

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, this Section replicates Tables 1 and 2 thirty times, in each case assuming the ban occurred in a different week ranging from June 22, 2018 to November 2, 2018, or May 3, 2019 to 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, and where none of the estimation periods include the actual ban.[18] The goal is to see if our estimates of the actual ban's effects are noticeably different from the placebo estimates.

The results are shown in Figures A10.1.1 and A10.1.2, which show the frequency distribution of these estimates. For both applications and call-backs, the figures clearly show that the *F* and *M* coefficients for the actual ban are dramatically different from all the placebo estimates. The *N* coefficients, on the other hand, fall into the middle of the range of the placebo estimates, all of which are reasonably close to zero.

---

[17] 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.

[18] In all cases we continue to restrict our sample to job ads that received applications both before and after the *actual* ban date. As for our main estimation sample, this ensures that all ads have a known, pre-ban gender request (*F, N* or *M*). We also exclude placebo ban dates before June 22, 2018 to avoid the effects of an XMRC site upgrade that occurred just before the week of April 27, 2018.

Figure A10.1.1 Effects of Placebo Ad Bans on the Female Share of Applications

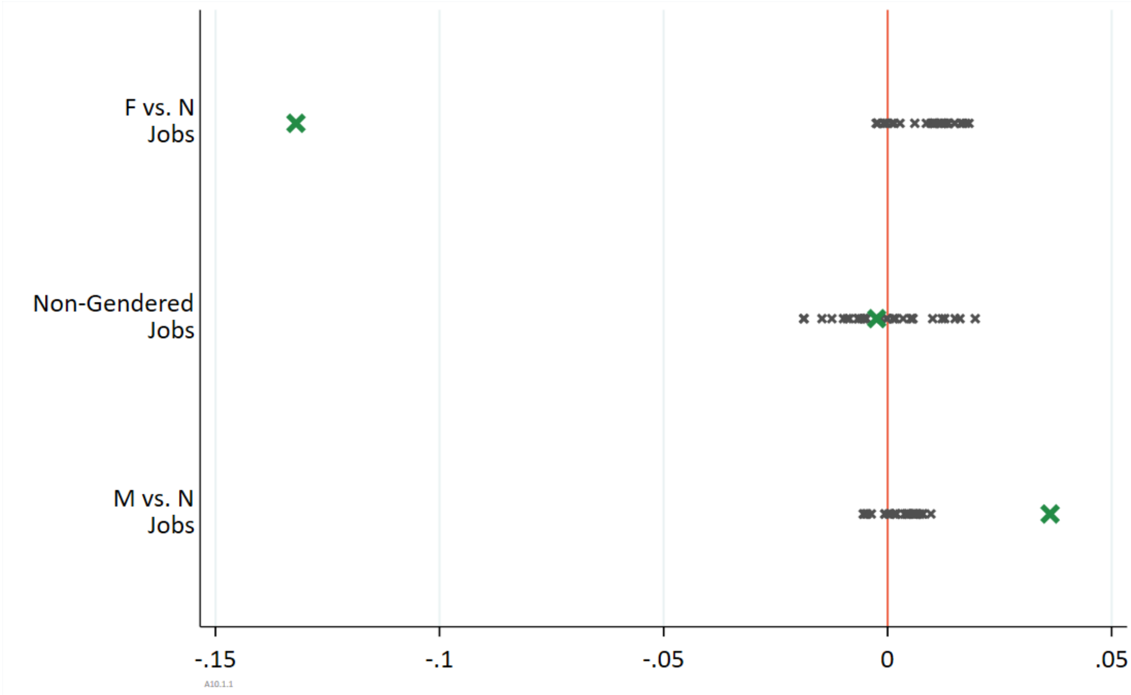

Figure A10.1.2 Effects of Placebo Ad Bans on the Female Share of Call-Backs

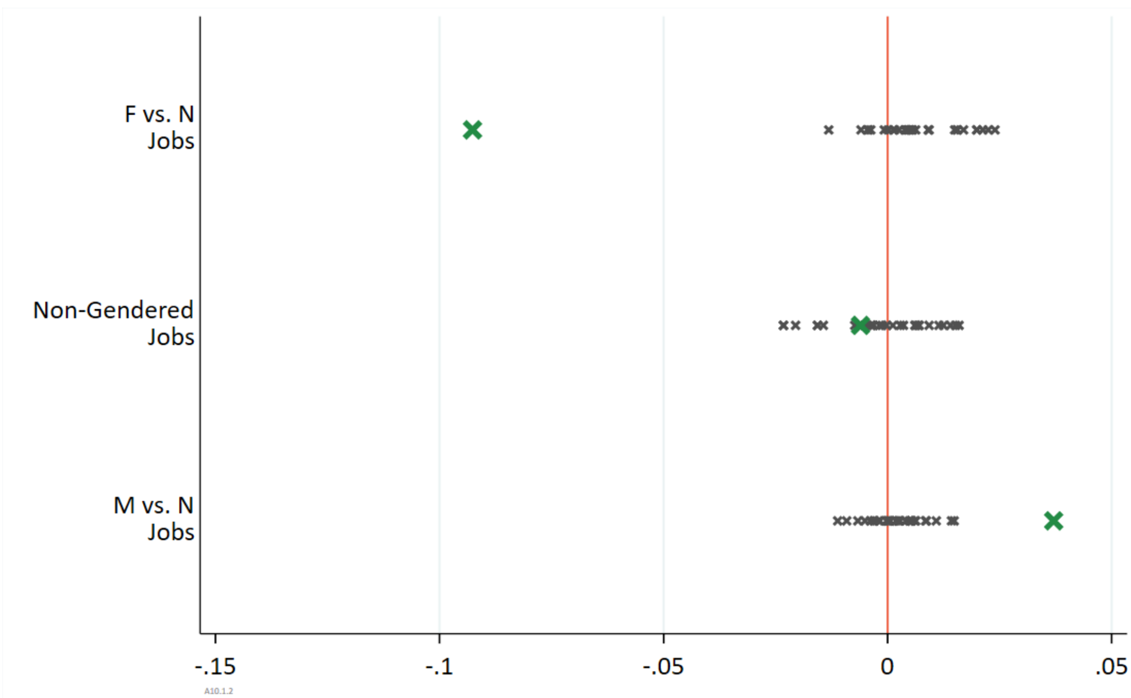

Note:

1. The dates used for placebo estimates are from 2018 June 22 to 2018 November 2nd, and 2019 May 3rd to 2019 July 5th. There are 30 dates used here.
2. The green symbols are estimates using the actual gender ban date, 2019 March 1st.

10.2 Search Frictions

In this Section we replicate our estimates of the ban's effects on our three measures of matching frictions: application arrival rates, match quality, and yield -- for placebo ban dates ranging from March 2, 2018 through July 5, 2019. In all cases (as in Tables 4 -- 6) the estimation window is the 30 days surrounding each placebo ban date.[19] As in Appendix 9.2, we restrict our attention to the ban's effects on aggregate frictions (i.e. for *F*, *N*, and *M* jobs combined), and on men's and women's application rates to *F* and *M* jobs.

Our results for aggregate search frictions affecting firms (application arrivals and quality) are summarized in Figure A10.2.1; results for workers' call-back yields are shown in Figure A10.2.2. All these estimates use the specification in panel A of Tables 4 -- 6. For application arrivals and quality, our positive main estimate is larger than most placebo estimates, but it is not the largest and does not stand out dramatically from most other estimates. For application yields our negative main estimate is in the middle of a large range of estimates. Overall, we conclude that it is difficult to identify the ban's effects on aggregate matching frictions using our natural experiment. This is primarily because the ban's effects spillover effects on these outcomes in non-treated (*N*) jobs are hard to identify, and because *N* jobs comprise about three fourths of the jobs on XMRC.

[19] 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 A10.2.1 Effects of Placebo Ad Bans on the Daily Flow of Applications and the Quality of Applications Received

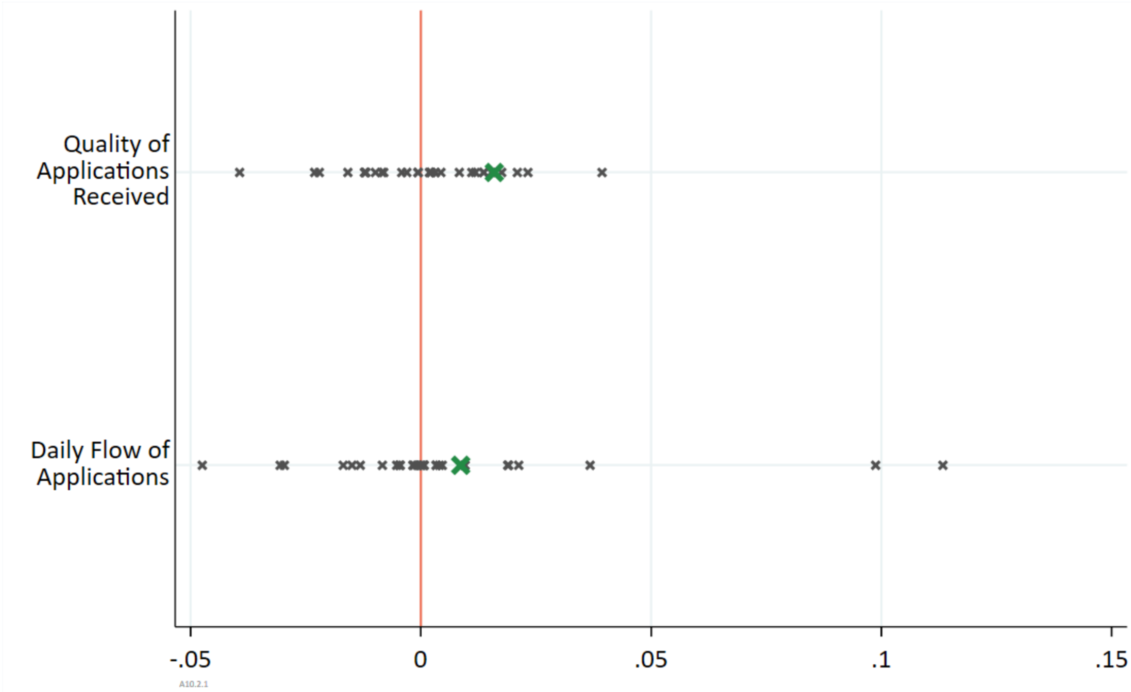

Notes:

1. The dates used for placebo estimates are from 2018 June 22 to 2018 November 2nd, and 2019 May 3rd to 2019 July 5th. There are 30 dates used here.
2. Each dot corresponds to an estimated coefficient. 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 column 2 of Tables 4 and 5.
3. The green symbols are estimates using the actual gender ban date, 2019 March 1st.

Figure A10.2.2 Effects of Placebo Ad Bans on the Probability an Application Yields a Call-Back

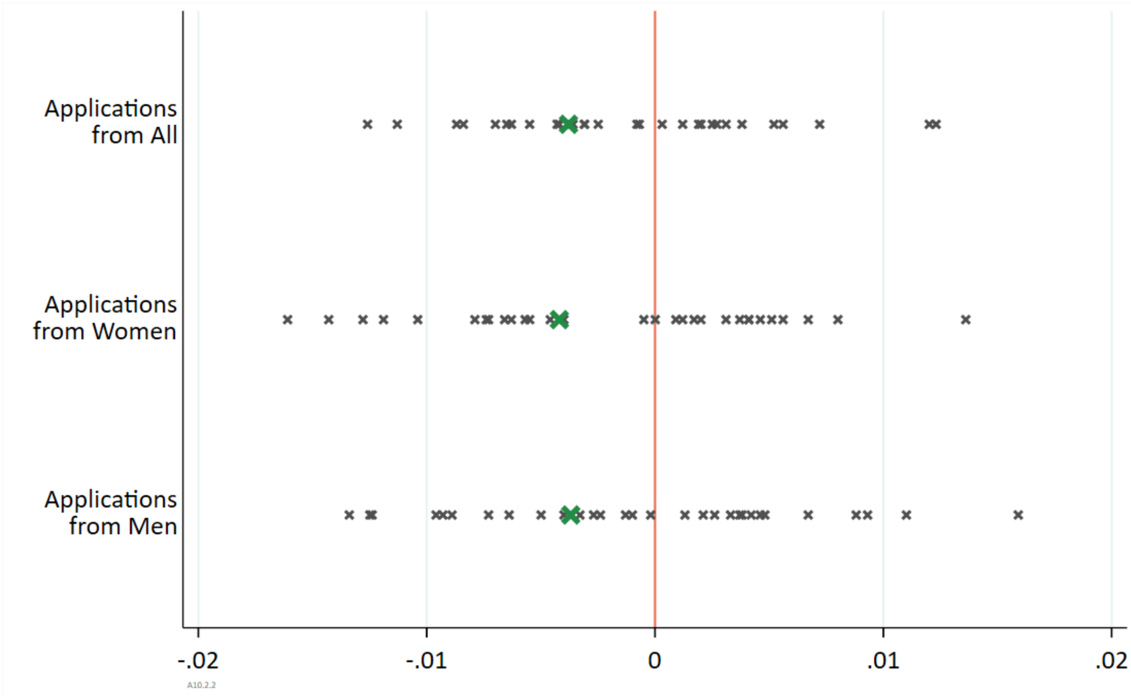

Notes:

1. The dates used for placebo estimates are from 2018 June 22 to 2018 November 2nd, and 2019 May 3rd to 2019 July 5th. There are 30 dates used here.
2. Each dot corresponds to an estimated coefficient. 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, 3 and 4 of Table 6.
3. The green symbols are estimates using the actual gender ban date, 2019 March 1st.

Figure A10.2.3 Effects of Placebo Ad Bans on Application Flows to *F* and *M* jobs, by Applicant Gender

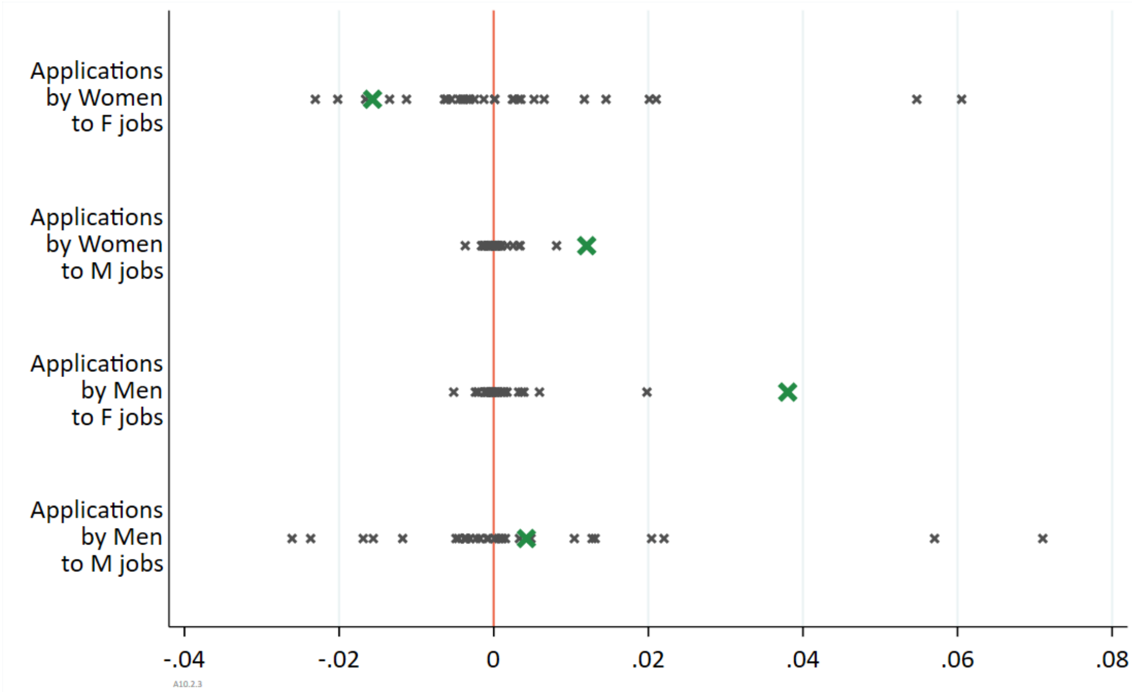

Notes:

1. The dates used for placebo estimates are from 2018 June 22 to 2018 November 2nd, and 2019 May 3rd to 2019 July 5th. There are 30 dates used here.
2. Each dot corresponds to an estimated coefficient. 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 3 and 4 of Table 4.
3. The green symbols are estimates using the actual gender ban date, 2019 March 1st.

Appendix 11: Difference-in-Difference Analysis

To see if our discontinuity-based results are robust to a fundamentally different set of identifying assumptions, this Appendix exploits the fact that our data run back to the start of 2018.[20] 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 fashion between the two years, while still using the week fixed effects to capture high-frequency variation that is common to the two years.[21]

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.[22] 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 our DiD approach explores the effects of a fundamentally different source of identification, we chose to replicate all the regression tables in the paper (1, 2, and 4 -- 6) here.

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.[23] Turning to application arrivals, quality and yield, the aggregate results are mostly confirmed: application arrivals and quality continue to increase after the ban, but the null effect on application yield becomes negative and statistically significant. Our DiD estimates of the ban's effects on men's and women's applications to *M* and *F* jobs, however, are again highly robust.

[20] 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.

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

[22] 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.

[23] The only difference of potential interest is that the ban's effect on the female share of *applications* to *N* jobs becomes statistically insignificant in three of four specifications. 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.

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 A11.1 and A11.2 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).[24] 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).[25] 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.

---

[24] All regressions in Appendix 11 also contain dummies that capture XMRC's system transition during 2018.

[25] 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 A11.1.1: 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***	–0.1391***	–0.1370***	–0.1343***
	(0.0056)	(0.0056)	(0.0055)	(0.0026)
Post ban week × Male job × 2019	0.0486***	0.0486***	0.0490***	0.0359***
	(0.0055)	(0.0055)	(0.0055)	(0.0020)
Post ban week × 2019	–0.0008	0.0017	–0.0070**	0.0025
	(0.0020)	(0.0026)	(0.0034)	(0.0023)
Female job	0.4740***	0.4741***	0.4706***	
	(0.0034)	(0.0034)	(0.0034)	
Male job	–0.3545***	–0.3544***	–0.3552***	
	(0.0032)	(0.0032)	(0.0031)	
Quartic in job weeks		Y	Y	Y
Time trends			Y	Y
Job ad fixed effects				Y
Effective # of obs	1,806,225	1,806,225	1,806,225	1,806,225
R^2	0.201	0.201	0.203	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.
4. * p<.10, ** p<.05, *** p<.01.
5. Since our DiD approach compares the 'same' weeks in 2018 and 2019, and important events in China are determined by the lunar calendar, we line up weeks between 2018 and 2019 relative to the Chinese New Year, which occurred on February 16, 2018 and February 5, 2019. Since ad ban happened 24 days after the 2019 New Year (on Friday March 1), a comparable ad ban in 2018 *would have occurred* on Monday, March 12. 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). The 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.

Table A11.1.2: 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***	–0.0969***	–0.0974***	–0.1006***
	(0.0102)	(0.0101)	(0.0101)	(0.0075)
Post ban week × Male job × 2019	0.0388***	0.0388***	0.0392***	0.0311***
	(0.0086)	(0.0086)	(0.0086)	(0.0064)
Post ban week × 2019	–0.0073	–0.0042	–0.0012	–0.0025
	(0.0048)	(0.0059)	(0.0159)	(0.0139)
Female job	0.4629***	0.4625***	0.4630***	
	(0.0063)	(0.0063)	(0.0063)	
Male job	–0.4106***	–0.4106***	–0.4102***	
	(0.0055)	(0.0055)	(0.0055)	
Quartic in job weeks		Y	Y	Y
Time trends			Y	Y
Job ad fixed effects				Y
Effective # of obs	215,635	215,635	215,635	215,635
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.
4. * p<.10, ** p<.05, *** p<.01.
5. 2018 and 2019 weeks are lined up according to China's lunar calendar. For details, see note 4 in Table A11.1.1.

11.2 Search Frictions

In this Section we use our DiD estimation approach to estimate the ban's effects on search frictions (Tables 4 -- 6 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 -- 6 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, Table A11.2.1 replicates Table 4 using the same DiD specification as Table A11.1.1; Table A11.2.2 replicates Table 5 the same way. Column 2 of both Tables shows that our original results -- increases in the number and quality of applications -- are reproduced.[26] Consistent with Table 5, Table A11.2.2 shows a small but statistically significant increase in the mean match quality of applications, both overall and in *N* jobs. In Table A11.2.3, which estimates the ban's effects on workers' application yields, the small, insignificant aggregate decline in call-back yields in Table 6 becomes slightly larger and statistically significant. 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.[27]

Turning now to the ban's effects on application flows to the jobs whose gender requests were removed, panels 3 and 4 of Table A11.2.1 show results that are strikingly similar to Table 4 (except for scale, since we are now modelling weekly application flows). 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. Thus, the regressions that illuminate how changes in application flows to *M* and *F* jobs explain the ban's effects are remarkably stable to our DiD identification approach.[28]

[26] Notice that all the coefficients and standard errors in Table A11.2.1 tend to be larger than in Table 4, reflecting the fact that the dependent variable is now applications received per week, rather than per day.

[27] 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.

[28] Like Tables 5 and 6, Tables A11.2.2 and A11.2.3 show mostly insignificant effects of the ban on the quality and yield of applications to *F* and *M* jobs.

Table A11.2.1: Effects of the Gendered Ad Ban on the Number of Applications Received: DiD Approach

	(1)	(2)	(3)	(4)
		Applications from:		
	All	All	Women	Men
1. To All Ads: [4,126,935 observations]				
Post ban	0.0393**	0.0678***	0.0071	0.0607***
	(0.0158)	(0.0170)	(0.0096)	(0.0114)
R^2	0.073	0.426	0.425	0.396
2. To Ads without Gender Request: [3,190,466 observations]				
Post ban	0.0037	0.0309*	0.0059	0.0249**
	(0.0169)	(0.0183)	(0.0103)	(0.0121)
R^2	0.076	0.427	0.424	0.374
3. To Ads that Requested Women: [426,174 observations]				
Post ban	0.2074***	0.2396***	–0.1195***	0.3591***
	(0.0426)	(0.0456)	(0.0407)	(0.0159)
R^2	0.086	0.401	0.399	0.293
4. To Ads that Requested Men: [510,295 observations]				
Post ban	0.1135***	0.1504***	0.1323***	0.0181
	(0.0423)	(0.0460)	(0.0102)	(0.0443)
Job Ad Fixed Effects?		Y	Y	Y
R^2	0.058	0.436	0.350	0.434

Notes:

1. This table replicates Table 4 using a Difference-in-Difference approach, treating weeks from 2018 as a control group for 2019.
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 1.164, 1.079 and 1.145 for female, non-gendered, and male jobs, respectively.
3. * p<.10, ** p<.05, *** p<.01.
4. All specifications include a quartic in job weeks. Time trends are modelled with (lunar) calendar week fixed effects; a quartic in time interacted with 2019; plus, dummies for the Spring Festival weeks and the week following it, both interacted with 2019. All regressions are clustered by firm ID.
5. 2018 and 2019 weeks are lined up according to China's lunar calendar. For details, see note 4 in Table A11.1.1.

Table A11.2.2: Effects of the Gendered Ad Ban on the Match Quality of Applications: DiD Approach

	(1)	(2)	(3)	(4)
	Applications from:			
	All	All	Women	Men
1. To All Ads:				
Post ban	–0.0077	0.0093**	–0.0067	0.0228***
	(0.0067)	(0.0037)	(0.0055)	(0.0051)
# of observations	3,893,360	3,893,360	1,626,568	2,266,792
R^2	0.003	0.494	0.537	0.508
2. To Ads without Gender Request:				
Post ban	–0.0066	0.0106**	–0.0190***	0.0282***
	(0.0077)	(0.0041)	(0.0063)	(0.0057)
# of observations	2,955,335	2,955,335	1,229,180	1,726,155
R^2	0.004	0.511	0.556	0.526
3. To Ads that Requested Women:				
Post ban	–0.0162	0.0149	0.0278**	0.0559
	(0.0181)	(0.0106)	(0.0112)	(0.0362)
# of observations	426,273	426,273	356,293	69,980
R^2	0.004	0.448	0.453	0.568
4. To Ads that Requested Men:				
Post ban	–0.0026	0.0013	0.0450	–0.0002
	(0.0166)	(0.0107)	(0.0386)	(0.0112)
Job Ad Fixed Effects?		Y	Y	Y
# of observations	511,752	511,752	41,095	470,657
R^2	0.003	0.424	0.634	0.422

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, and the dependent variable is the match quality. Match qualities are normalized to a mean of zero and standard deviation of one among all applications in our data.
3. Post ban means the week was after the actual ban.
4. * p<.10, ** p<.05, *** p<.01.
5. The average match quality for the applications in the current estimation sample is –0.056, –0.012 and 0.068 for female, non-gendered, and male jobs, respectively.
6. All specifications include a quartic in job weeks. Time trends are modelled with (lunar) calendar week fixed effects; a quartic in time interacted with 2019; plus, dummies for the Spring Festival weeks and the week following it, both interacted with 2019. All regressions are clustered by firm ID.
7. 2018 and 2019 weeks are lined up according to China's lunar calendar. For details, see note 4 in Table A11.1.1.

Table A11.2.3: Effects of the Gendered Ad Ban on Call-Back Chances Rate per Application Submitted: DiD Approach

	(1)	(2)	(3)	(4)
		Applications from:		
	All	All	Women	Men
1. To All Ads:				
Post ban	–0.0056**	–0.0055**	–0.0051	–0.0061**
	(0.0023)	(0.0024)	(0.0035)	(0.0027)
# of observations	3,987,960	3,987,960	1,669,543	2,318,417
R^2	0.018	0.129	0.134	0.122
2. To Ads without Gender Request:				
Post ban	–0.0029	–0.0035	–0.0012	–0.0053*
	(0.0025)	(0.0027)	(0.0040)	(0.0030)
# of observations	3,035,701	3,035,701	1,264,561	1,771,140
R^2	0.017	0.150	0.160	0.142
3. To Ads that Requested Women:				
Post ban	–0.0205***	–0.0104	–0.0122	0.0134
	(0.0062)	(0.0078)	(0.0084)	(0.0185)
# of observations	434,179	434,179	362,852	71,327
R^2	0.024	0.283	0.263	0.409
4. To Ads that Requested Men:				
Post ban	–0.0117**	–0.0084	–0.0377	–0.0068
	(0.0051)	(0.0065)	(0.0262)	(0.0066)
Applicant fixed effects?		Y	Y	Y
# of observations	518,080	518,080	42,130	475,950
R^2	0.015	0.272	0.596	0.249

Notes:

1. This table replicates Table 6 using a Difference-in-Difference approach, treating weeks from 2018 as a control group for 2019.
2. Observations are applications, and the dependent variable equals one if the application ever received a call-back. Post ban means the week was after the actual ban. The average call-back probability of applications is 0.117, 0.086, and 0.082 for female, non-gendered, and male jobs, respectively.
3. * p<.10, ** p<.05, *** p<.01.
4. All specifications include a quartic in job weeks. Time trends are modelled with (lunar) calendar week fixed effects; a quartic in time interacted with 2019; plus, dummies for the Spring Festival weeks and the week following it, both interacted with 2019; and dummies for the transition of XMRC's website upgrade in April 2018. All regressions are clustered by firm ID.
5. In Female jobs, we cannot reject that women's relative success rates (column 2 – column 3) were unchanged by the ban (p=.208).
6. In Male jobs, we also cannot reject that women's relative success rates (column 2 – column 3) were unchanged by the ban (p=.253).
7. 2018 and 2019 weeks are lined up according to China's lunar calendar. For details, see note 4 in Table A11.1.1.

Appendix 12: Re-Weighting the Estimates to Represent a National Job Board

If XMRC is an unusual job board relative to others in China, or to other countries that permit gendered job ads, the external validity of our results will be limited. Among other things, one might wonder whether Xiamen is an atypical city, or whether the surprising number of postings for jobs like drivers (at least among jobs that are extremely male dominated) is an idiosyncratic feature of XMRC. Could factors like these account for the striking asymmetry of our main results, where men gained much more access to women's jobs than vice versa?

This Appendix addresses these representativeness concerns in two ways. First, we replicate two of the paper's main regression results after re-weighting our XMRC data to be representative of a much larger national job board, Zhaopin. The results are quite similar. Second, we focus specifically on the role of driving jobs and show that they are indeed overrepresented on XMRC, relative to Zhaopin. This may be because XMRC is a local job board and driving jobs frequently require local hukou. However, our main results remain unchanged when we exclude all driving jobs from our estimation sample.

Our analysis in this Appendix draws on a comprehensive scraping of Chinese job ads we obtained in 2018 from a small Chinese company that has asked to remain anonymous, citing concerns with national security and the China-US trade war. Given the potential risks to this company from any disclosure of their data, we cannot provide the underlying data used to create the weights and descriptive statistics presented in this Section. However, the descriptive statistics and the weights that were calculated in this Appendix can be accessed for replication purposes from Kuhn and Shen (2023), specifically the datasets XMRC_ZP_Descriptives.dta and XMRC_ZP_wt.dta.

12.1 Re-Weighting Our Table 1 and Table 2 Results to Represent a National Job Board

In this Section, we replicate some of our main analyses after re-weighting the job ads on XMRC.com, which serves one city and hosted 403,208 job ads over almost two years (2018/1/1 to 2019/10/25) to represent a much larger, national job board, Zhaopin.com. Zhaopin serves all of China, with an especially large presence in mega-cities including Beijing, Shanghai, Guangzhou. It hosted almost 5.5 million job ads in a single month (2018/8/1 to 2019/9/6). Since job ads have many characteristics that are not standardized between platforms, we did this analysis using several alternative sets of weights, ranging from a simple version that uses only numerical variables that are comparably measured on both job boards (wage offered, education and experience required), to those variables plus a 20-dimensional feature vector derived from the unstructured text of the job ads using a Doc2Vec model. Additional details on the construction of these weights are provided in Appendix 12.2, which also demonstrates that all of these methods replicate the education, experience and wage distribution on Zhaopin very closely. Together, these weights adjust for a wide range of differences between XMRC and Zhaopin, including the fact that XMRC's job ads are more targeted at workers with middle levels of education (as opposed to less than high school or university), at workers with less experience, and offer wages that are 28 percent lower.

With these weights in hand, we re-estimated column 4 of Tables 1 and 2 of the paper -- our preferred estimates of the ban's effect on the gender mix of applications and call-backs --, re-weighting our XMRC sample to represent the ads posted on Zhaopin.com. The results are presented in Figures A12.1.1 and A12.1.2, for four different sets of weights. Figures A12.1.1 and A12.1.2 also display our original Table 1 and 2 estimates for comparison. Overall, the re-weighted estimates are very similar to our main estimates, and continue to support our result that men entered *F* jobs at a much higher rate than women entered *M* jobs. For call-backs (but not for applications, where our statistical power is considerably greater), some of the coefficients measuring women's increased access to *M* jobs become statistically insignificant. If anything, these estimates reinforce our conclusion that the ban had asymmetric effects, which were stronger in jobs that had previously requested women than in jobs that had requested men.

Figure A12.1.1: Effects of the Ban on the Female Share of Applications: Robustness to Compositional Adjustment to Mimic National Job Board Posting Samples

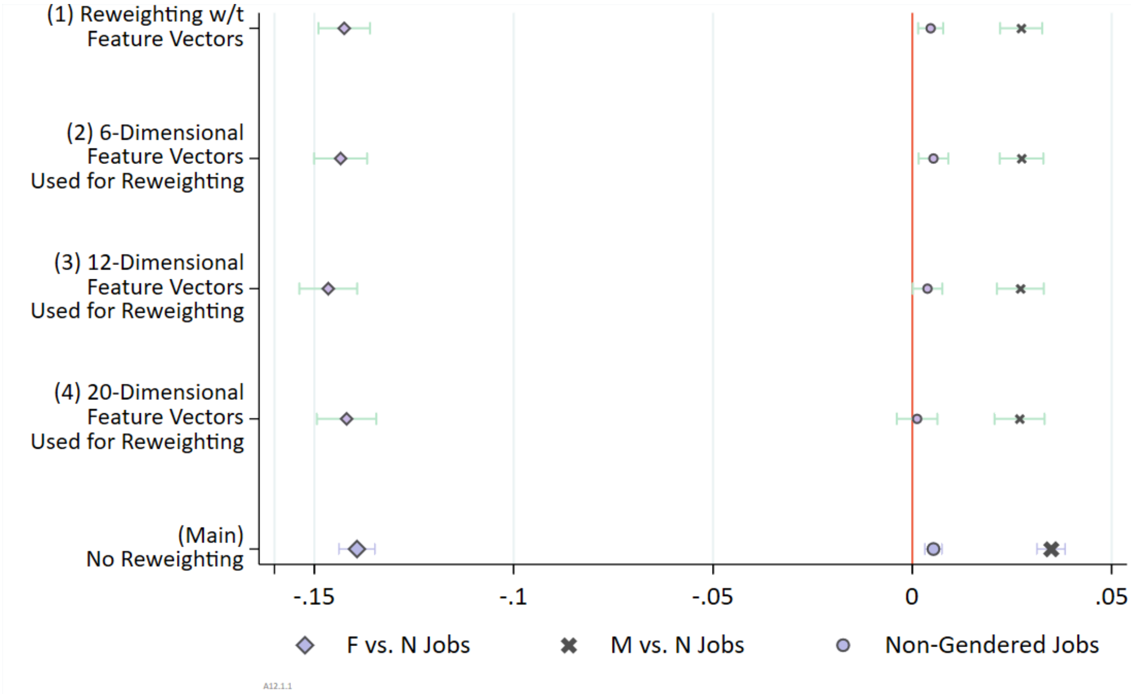

Notes:

1. The vertical axis shows the various composition adjustment designs used. Except for the last row, the XMRC job posting data is reweighted using Probit model predictions based on education, experience and offered wages in XMRC and Zhaopin job postings. Various dimensions of feature vectors are added in the Probit model in row 2, 3, and 4, respectively.
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 A12.1.2: Effects of the Ban on the Female Share of Call-Backs: Robustness to Compositional Adjustment to Mimic National Job Board Posting Samples

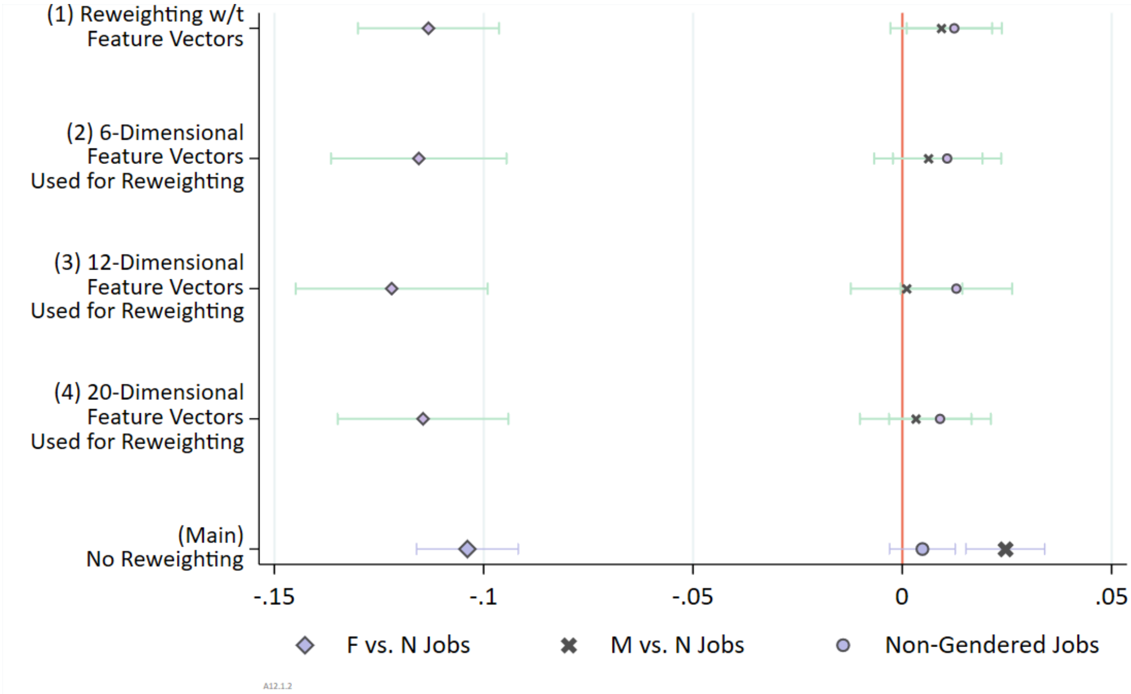

Notes:

1. The vertical axis shows the various composition adjustment designs used. Except for the last row, the XMRC job posting data is reweighted using Probit model predictions based on education, experience and offered wages in XMRC and Zhaopin job postings. Various dimensions of feature vectors are added in the Probit model in row 2, 3, and 4, respectively.
2. The horizontal axis shows the estimated coefficients and 95% confidence intervals.
3. The regression specification of column 4 in Table 2 is used.

12.2 Constructing the Weights: Additional Details

Our goal in this Section is to use the Zhaopin data to construct a set of weights that will allow us to re-weight the XMRC job titles according to their prominence in the broader Chinese labor market. Appendix 12 re-weights some of our main estimates to represent the population of job ads posted on a much larger, national job board (Zhaopin.com). As noted earlier, these weights were derived from a comprehensive scraping of Chinese job ads, which we obtained in 2018 from a small Chinese company that has asked to remain anonymous. The weights, but not the data used to create them, can be accessed for replication purposes at Kuhn and Shen (2023), specifically the datasets XMRC_ZP_Descriptives.dta and XMRC_ZP_wt.dta.

Since none of the occupation categories or job titles on any of China's major job boards (including XMRC) are standardized, we constructed several alternative sets of weights, ranging from a simple version that uses only numerical variables that are present on both sites (wage offered, education and experience required), to those plus a 20-dimensional feature vector derived from the unstructured text of the job ads using a Doc2Vec model (using the Python Gensim package). This model projects documents (sets of words) into vectors, so that a similarity score can be calculated using a cosine similarity formula.[29] The construction of weights involved four steps:

Step 1: [Model Training] we use the 910,323 Zhaopin job ads and the 403,208 XMRC ads that have detailed job description information to estimate (train) a Doc2Vec model using the Python Gensim package. This model essentially project documents (sets of words) into vectors, so that a similarity score can be calculated using cosine similarity formula of the vectors of any two documents. A key parameter in this estimation is the dimension of the document vector, *D*, we experiment with 6, 12, and 20 here.

Step 2: [Model Prediction] we then use the estimated Doc2Vec model to predict the vectors of the 929 job types of the full 5,499,015 job ads on Zhaopin -- the full sample for the period from 2018 August 1st to 2018 September 6th, as well as the job titles of 117,390 job ads of XMRC in our analysis sample.

Step 3: [Probit Estimation] we then estimate a Probit model using the pooled data of Zhaopin full sample and XMRC analysis sample as follows

$$\Pr(XM_i = 1) = \Phi(\alpha_0 + \alpha_1 \cdot EDU_i + \alpha_2 \cdot EXP_i + \alpha_3 \cdot lnWAGE_i + \alpha_4 \cdot WAGEMissing_i + \beta_1 \cdot DocV_i) \quad \text{(A12.2.1)}$$

[29] The choice of procedure was complicated by the fact that we only have rich unstructured job variables for a subsample of the Zhaopin data that we are targeting, and the fact that our XMRC analysis sample is a subset of the ads posted there as well. We decided to use the full XMRC sample and the subset of Zhaopin sample to generate a better language model.

Where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution, for each job posting i in this pooled data, EDU_i are the set of education requirement category dummies, EXP_i are the set of experience requirement category dummies, $lnWAGE_i$ is the offered wage in log form, $WAGEMissing_i$ is the dummy for no explicit offer wage, and $DocV_i$ is the feature vector predicted for the job ads in step 2.

Step 4: [Weight Generation] Every job posting in the XMRC analysis sample is assigned a weight based on the predicted likelihood of being belong to XMRC in the Probit model estimated in Step 3. The formula for these weights is provided in equation A12.1.1.

$$wt_i = \frac{\left(\frac{1-Pr\widehat{(XM_i)}}{Pr\widehat{(XM_i)}}\right)}{\frac{1}{N_{XMRCA}} \cdot \sum_{j \in XMRC\_Analysis\_Sample} \left(\frac{1-Pr\widehat{(XM_j)}}{Pr\widehat{(XM_j)}}\right)} \qquad \text{(A12.2.2)}$$

Where N_{XMRCA} is the number of job postings in XMRC analysis sample with non-zero estimated weights.

Given the dissimilarity of the Zhaopin and XMRC job postings, some job postings in XMRC analysis sample are dropped as they are predicted to be belong to XMRC for sure. As we increase the dimension of the feature vector, more dissimilarity is captured and more XMRC job postings are dropped.

Appendix Table 12.2.1 summarizes the effectiveness of our weighting schemes in reproducing the education, experience, and wage distribution on Zhaopin (the national board). It shows that the distribution of education, experience, and wages in all the re-weighted XMRC samples are similar to the targeted Zhaopin data.

Table A12.2.1 Characteristics of Job Postings on Zhaopin and XMRC, with and without Re-Weighting

	Zhaopin	XMRC				
Weight?	No	No	education, experiences, and wage offered	(3) + 6-dimensional feature vectors	(3) + 12-dimensional feature vectors	(3) + 20-dimensional feature vectors
	(1)	(2)	(3)	(4)	(5)	(6)
Education requirement						
Less than high school	.244	.119	.224	.221	.227	.218
High school	.007	.083	.007	.007	.007	.007
Technical school	.074	.122	.072	.070	.070	.079
College	.431	.481	.435	.429	.425	.421
University	.243	.195	.262	.272	.271	.275
Experience requirement						
None	.477	.487	.444	.432	.431	.433
1 year	.028	.169	.028	.028	.028	.027
2-3 years	.286	.263	.298	.296	.296	.289
4-5 years	.152	.064	.163	.169	.171	.173
6-10 years	.059	.018	.067	.075	.074	.077
No explicit wage	.020	.240	.020	.020	.020	.020
Wage	8,285	5,936	8,203	8,281	8,230	8,288
# of job postings	5,401,400	117,382	117,382	117,358	116,603	115,104

Notes:

1. Four sets of weights are constructed. Column (3) provides the benchmark weights where feature vectors derived from job titles and job types are not used. Column (4), (5) and (6) use the preferred weights where feature vectors derived from job titles and job types are used.
2. The feature vectors of job titles and job types are derived using Nature Language Process package Gensim's Doc2Vec algorithm trained on the pool of job descriptions of both Zhaopin and XMRC job postings. Column (4), (5) and (6) use 6-, 12- and 20-dimensional feature vectors, respectively.

12.3 Effects of Driving Jobs on our Results

When we compare the incidence of driving jobs on XMRC versus Zhaopin, we find that driving jobs are indeed overrepresented on XMRC, representing 1.1 percent versus 0.4 percent of the ads posted. This may be because XMRC is a local job board and driving jobs frequently require local *hukou*. To see how driving jobs affect our main results, we then replicated panels (a) and (b) of Table 3 after excluding all the driving jobs from the sample. The results, shown in Table A12.3.1, are very similar for women, and almost identical (up to three digits) for men. The intuition is straightforward: while driving jobs comprise a large share of the extremely male jobs on XMRC, they play almost no role in the jobs women gained access to (and none in the jobs men gained access to). As already noted in the paper, the marginal jobs women (men) gained access to were significantly more male (female) than average, but did not include the most extremely gendered jobs on XMRC, such as driving.

Table A12.3.1 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 ban[2]**	(1) – (2)
Main estimates (Table 3a)	0.878	0.677	0.201
Excluding all driving jobs	0.869	0.675	0.194

Notes:
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 ban[2]**	(1) – (2)
Main estimates (Table 3a)	0.785	0.735	.050
Excluding all driving jobs	0.785	0.735	.051

Notes:
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).

Reference for Appendix 12:

Kuhn, Peter, and Kailing Shen. 2023. "Data and Code for: What Happens When Employers Can No Longer Discriminate in Job Ads?" *American Economic Association* [publisher], Inter-university Consortium for Political and Social Research [distributor]. http://doi.org/10.3886/E183021V1

Appendix 13: Calculating Incumbent Gender Mix from a Longer Pre-Estimation Period

In Section 4 of the paper, we used the gender mix of call-backs issued on XMRC between January and August 2018 (a period immediately preceding our main estimation sample) as a proxy for the incumbent gender mix of job titles, positions, workplaces and firms before the ban.[30] To assess the sensitivity of our results to this decision, Appendix 13 replicates all our gender mix-related results using incumbent gender mix estimates derived from call-backs during 2015 (February 1 2015 to October 10 2015), 2016 (December 24 2015 to August 31 2016), and 2018, thereby adding 18 months of call-backs to the 8 months used in our main analysis.[31] All the main results were unchanged, both qualitatively and quantitatively. We attribute this stability to two factors: First, XMRC's dominance in Xiamen means that firms tend to remain attached to it for many years, making their past XMRC hires a good predictor of their more recent ones. Second, the gender stereotypes associated with job titles on XMRC appear to be quite stable over time.

In more detail, Figure A13.1 replicates Figure 3 of the paper, which shows the relation between a job title's incumbent gender mix prior to our estimation period and its explicit gender requests after that (specifically, during September 2018 -- February 2019, while gender requests were still allowed.) Our updated results in Figure A13.1 are very similar to Figure 3: Explicit gender requests strongly *reproduce* the incumbent gender mix in a job title.

Next, Figure A13.2 replicates Figure 4 with our new incumbent gender mix measures. Once again, the results are very similar: Women gained almost no access to jobs with incumbent male shares of more than 80 percent; men on the other hand made substantial inroads into jobs that had previously requested women regardless of the job title's incumbent gender mix. Finally, Table 13.1 reproduces Table 3 of the paper, again with very similar results: men advanced considerably further into women's jobs than women did into men's, but in both cases most of the increased access occurred in relatively low-paying jobs.

---

[30] We limited ourselves to this interval because we do not have XMRC data from 2017, and because we wanted a baseline that reflects relative hiring patterns. While one might argue that a firm's current *stock* of employees is the natural baseline against which to measure the ban's integrating effects, this stock could include hires made many years ago when gender norms were different. Thus, while both short- and long-memory proxies for past hires yield very similar results in our context, it is not clear which proxy is of greatest interest from a policy perspective.

[31] To adjust for the growth of XMRC over this period, and for the fact that call-back data are sparser in 2015 and 2016, we weight the observed call-backs so that each month over these three years contributes equally to the estimates. In additional analyses (not reported here) we also calculated incumbent gender mix using 2015 data only, and using 2016 data only. Once again, the results were very similar.

Figure A13.1: Share of Ads Requesting Women, Men or Making no Gender Request as a Function of Incumbent Female Share, by Job Title (robustness check)

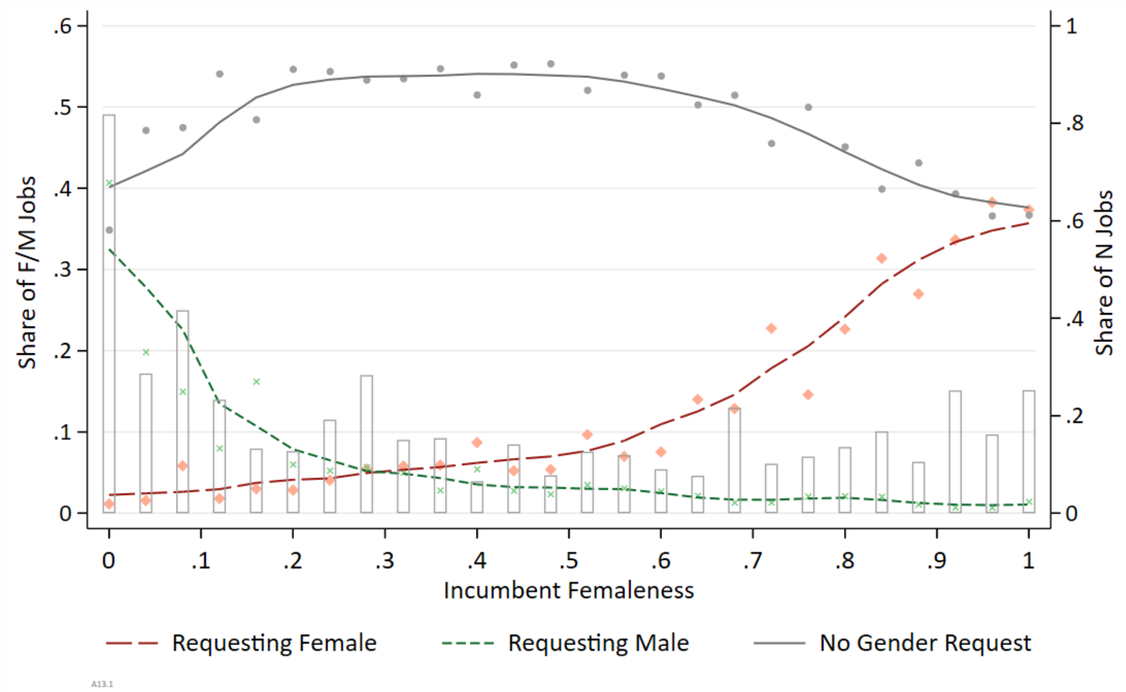

Note:

1. Figure 13.1 replicates Figure 3 using a longer period of time (2015, 2016, plus January -- August 2018) to estimate each job title's incumbent gender mix.

Figure A13.2 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

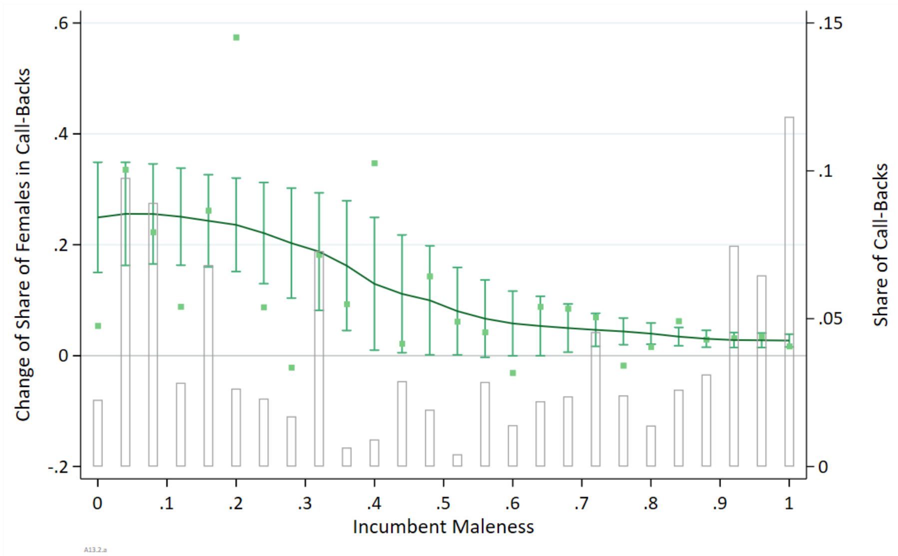

b. Job Ads that Initially Requested Women: Pre-Post Ban Changes in Men's Share of Call-Backs by Incumbent Female Share of the Job Title

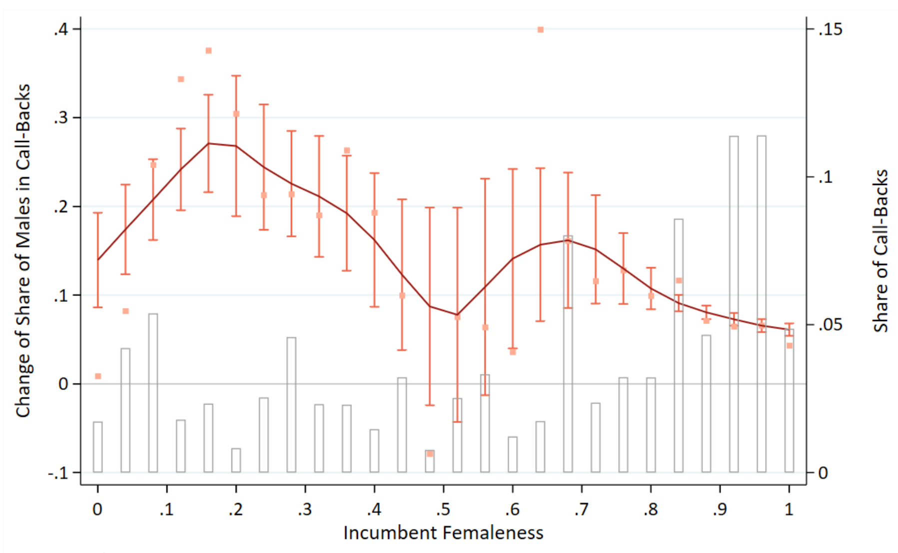

Note:

1. Figure 13.2 replicates Figure 4 using a longer period of time (2015, 2016, plus January -- August 2018) to estimate each job title's incumbent gender mix.

Table A13.1 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)	***M* Titles Women *Entered* (were Called Back to) Because of the Ban**	(1) – (2)
	original	
0.878	0.677	0.201
	with 2015, 2016 and 2018 jobs	
0.883	0.674	0.210

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)	**F jobs Men *Entered* (were Called Back to) Because of the Ban**	(1) – (2)
	original	
0.785	0.735	0.050
	with 2015, 2016 and 2018 jobs	
0.784	0.733	0.051

c. Incumbent Gender Mix of the Job Titles *Entered* by Women and Men because of the Ban

	original	*with 2015, 2016 and 2018 jobs*
Incumbent Male Share of Titles Entered by Women	0.677	0.674
Incumbent Female Share of Titles Entered by Men	0.735	0.733

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

original (Table 3)

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

with 2015, 2016 and 2018 jobs:

Mean log (wage) in:	All Titles	Titles Entered because of the Ban	Difference (2) – (1)
	(1)	(2)	(3)
Job Ads Requesting Men (*M* jobs)	8.739	8.598	–0.141
Job Ads Requesting Women (*F* jobs)	8.598	8.449	–0.149
Jobs Ads with no Gender request (*N* jobs)	8.859	--	--
All Jobs	8.826*	--	--

Notes:

1. Table A13.1 replicates Table 3 using a longer period of time (2015, 2016, plus 2018) to estimate each job title's incumbent gender mix.
2. *To account for inflation and real wage growth, log wages are adjusted so that for all jobs combined have the same mean across years. Wages were attached to job titles using all the ads that posted a wage in the relevant pre-estimation period.

Appendix 14: Did the Ban Redirect Active Search Spells or Initiate New Spells?

In principle, the additional gender-mismatched applications that were caused by the ban could have resulted from the re-direction of active worker search spells that were already in progress, from the revival of dormant worker profiles in response to newly opened opportunities, or from new signups of workers onto the job board. While it seems unlikely that an overnight change in the content of a minority of ads on XMRC caused an immediate increase in the creation of new profiles, distinguishing between the re-direction of ongoing search spells and the initiation of new spells by dormant profiles could shed some additional light on how the ban worked. For example, if previously gendered ads started receiving new, gender-mismatched applications from dormant profiles, this would suggest that removing gendered job ads encouraged some workers to start looking for new jobs.

While we do not have data on when workers first created their profiles on XRMC, this Appendix uses data on workers' recent search histories to distinguish between the mechanisms described above. Specifically, we categorize applications by number of days, *D*, since the worker last applied to another job. Thus, *D* > 7 means the worker has been inactive for at least 7 days. We then re-estimate our application flow regressions (Table 4), counting only applications from profiles that had been inactive for at least the last *D* days, where *D* = 7, 14, 21, …, 56. We also present results for "*D*=max", which counts only these applications for which we observe no previous activity in our entire estimation sample; these applications are the most likely to come from a new signup to the board. The results of this exercise are presented in Table 14.1. In our reading, three main patterns are evident.

First, looking at all job ads together, Table A14.1 shows that the ban was associated with a small increase in applications from jobseekers who had been inactive for a relatively short period: between 7 and 21 days. In the aggregate, there is no evidence that the ban re-activated the job search activity of workers who had been inactive for longer than that, or that it raised the number of new worker signups on the board.[32]

Second, when we restrict our attention to job ads that had previously requested men or women, these results change dramatically. Now we see increases in applications to *M* and *F* jobs, even among worker profiles that had been inactive on this job board for 56 days, or 8 weeks, and even among profiles for whom we observe no previous applications. Importantly, these applications from new or 'revived' profiles are entirely driven by *gender-mismatched* applications (men applying to jobs that had requested women and vice versa). Notably, while the estimated coefficients for these mismatched applications (especially in the *D > 56* and *D > max days observed* rows) are much smaller than the corresponding coefficients in the rows for active profiles (*D* ≤ 7) or for all profiles (Table 4), they are highly statistically significant. This is consistent with a scenario in which first applications from a

[32] In fact, the coefficient for "*D > max days observed"* for all jobs is negative and significant, though extremely small in magnitude.

dormant profile are rare, but are strongly affected by removing ad content that discourages one gender from applying.[33]

A final noteworthy feature of Table A14.1 is about the gender-*matched* applications to previously gendered jobs. While men's applications to previously male jobs either remained unchanged or increased slightly after the ban (depending on how long the profile had been active) women's applications to previously female jobs declined at all levels of previous application activity. As noted in the paper, this result mirrors an apparent gender gap in ambiguity aversion found by other authors, and in previous research of our own: Compared to men, women are much less likely to apply for a job when they are not explicitly invited. Table A14.1 extends this result by showing that the ban did more than cause actively-searching women to 'skip over' jobs that no longer explicitly invited women. The ban may also have discouraged some inactive female profiles from initiating new search spells on XMRC. As also argued earlier, these phenomena may help explain why women did not respond as strongly to the gendered ad ban as men did.

Overall, given this very specific way in which the ad ban re-activated dormant worker profiles (and possibly caused new signups on the board), we conclude that these two effects do not pose a threat to our main results. Instead, they illuminate a mechanism via which the ban worked of which we had previously been aware: By signaling that certain jobs were now open to a new gender, the ads re-activated the search activity of workers who had been discouraged from searching for quite some time.

[33] Note also that the effect of the ban on first applications from dormant profiles could dramatically understate the contribution of 'revived' profiles to gender integration. This is because all subsequent applications from such a revived profile will be counted as part of an active search spell.

Table A14.1: Ban Effects on the Daily Flow of Applications from Active versus Dormant Worker Profiles, Local Linear Regressions

Post ban coefficient for applications, by days since the worker's last application, *D*.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Applications from:									
	All	All	All	All	Women	Men	All	All	Women	Men
	To all Ads:		*To Ads that Requested Women:*				*To Ads that Requested Men*			
All profiles	0.0068***	0.0058***	0.0237***	0.0228***	*–0.0160****	**0.0388*****	0.0176***	0.0165***	**0.0122*****	*0.0043*
D ≤ 7 ("active" profiles)	0.0025**	0.0017	0.0185***	0.0177***	*–0.0142****	**0.0319*****	0.0071**	0.0061*	**0.0085*****	*–0.0024*
D > 7	0.0043***	0.0042***	0.0052***	0.0051***	*–0.0018*	**0.0069*****	0.0105***	0.0104***	**0.0037*****	*0.0067****
D > 14	0.0051***	0.0051***	0.0041***	0.0041***	*–0.0011*	**0.0052*****	0.0118***	0.0118***	**0.0028*****	*0.0091****
D > 21	0.0010*	0.0010*	0.0008	0.0008	*–0.0032***	**0.0041*****	0.0068***	0.0069***	**0.0025*****	*0.0044****
D > 28	–0.0001	–0.0001	0.0003	0.0003	*–0.0038****	**0.0040*****	0.0058***	0.0059***	**0.0024*****	*0.0034***
D > 35	–0.0001	–0.0001	0.0003	0.0003	*–0.0036****	**0.0039*****	0.0057***	0.0058***	**0.0024*****	*0.0033***
D > 42	–0.0005	–0.0005	0.0002	0.0002	*–0.0036****	**0.0038*****	0.0055***	0.0056***	**0.0023*****	*0.0033***
D > 49	0.0000	0.0000	0.0008	0.0008	*–0.0030***	**0.0038*****	0.0060***	0.0061***	**0.0023*****	*0.0038****
D > 56	–0.0002	–0.0002	0.0006	0.0006	*–0.0030***	**0.0037*****	0.0057***	0.0057***	**0.0023*****	*0.0035****
D > max days observed	–0.0008**	–0.0008**	–0.0004	–0.0005	*–0.0027****	**0.0022*****	0.0043***	0.0043***	**0.0015*****	*0.0028****
Job ad fixed effects		Y		Y	Y	Y		Y	Y	Y
Effective obs	3,514,552	3,514,552	407,240	407,240	407,240	407,240	441,722	441,722	441,722	441,722

Notes:

1. See notes to Table 4 for regression specifications.
2. This table replicates Table 4, disaggregating applications by the number of days, *D*, since the worker last applied to a job. *D > max* days *observed* means we do not observe any previous applications from that profile during our entire sample period. These applications are the most likely to have come from newly created profiles. *D* ≤ 7 counts only applications that had been *active* in the previous 7 days. Coefficients for **gender-mismatched applications** are **bolded**. Coefficients for *gender-matched applications (to previously gendered ads) are italicized.*
3. Applications sent within 15 days of the ad ban are used here. Observations are ad-day cells.
4. The dependent variable is the number of applications of each type received. Columns 5 and 9 only count applications from women. Columns 6 and 10 only count applications from men. For *D* ≤ 7, the average daily number of applications received is 0.135 and 0.135 in ads requesting women and men respectively. For *D* > 7, the average daily number of applications received is 0.051 and 0.060 in ads requesting women and men respectively. The average daily number of *all* applications received is 0.187 and 0.194 in ads requesting women and men respectively.
5. * *p*<.10, ** *p*<.05, *** *p*<.01.

Appendix 15: Using a Count Data Model for the Application Arrival Rate Regressions

Our outcome in Table 4's application flow regressions in Table 4 is the number of applications received by an ad on a given day. Even though these are count variables, we used linear probability models to estimate those regressions because of the practical difficulties in estimating count data models with high-dimensional fixed effects. To assess the likely impact of this choice, this Appendix replicates as many of the aggregate application flow regressions in Table 4 as possible using negative binomial count data regressions. We choose the negative binomial specification because ads frequently receive zero applications on a given day.

While we were not able to estimate a negative binomial model for our most saturated specification -- which includes job ad fixed effects -- Table A15.1 confirms Table 4's main result that the ad ban increased the overall flow of applications to job ads posted in XMRC for two other specifications: one with no controls, and one with a rich set of controls for the characteristics of the job ad. Interestingly, the estimated marginal estimated effects at the mean are about 49 percent larger in the count data regressions, suggesting that our linear probability models may underestimate the size of this effect. For both specifications, however, the estimated effects are relatively modest in size, corresponding to a 3.7 percent increase in the number of applications in the linear probability model and a 5.6 percent increase in the count data model. Unfortunately, we cannot confirm whether this larger effect persists in our preferred specification, which controls for job fixed effects. Still, the count data models confirm our main result of a modest rise in application flows: Overall, the ad ban gave employers more applications to choose from than before.

Table A15.1: Aggregate Effects of the Ad Ban on the Number of Applications Received: Linear Probability Models versus Count Data Models

	(1)	(2)	(3)	(4)	(5)	(6)
	Linear Probability Model			Negative Binomial Regressions		
Post ban	0.0068***	0.0068***	0.0058***	0.0107***	0.0101***	--
	(0.0015)	(0.0015)	(0.0015)	(0.0013)	(0.0013)	--
Job ad controls		Y			Y	
Job ad fixed effects			Y			Y
Effective # of obs	3,514,552	3,514,194	3,514,552	3,514,552	3,514,194	--
R^2	0.024	0.032	0.341	0.020	0.028	--

Notes:

1. Applications sent within 15 days of the ad ban are used here.
2. Observations are ad-date cells. All regressions are clustered by firm ID. The dependent variable is the number of applications received. Its mean is 0.1816.
3. * p<.10, ** p<.05, *** p<.01.
4. All specifications include day-of-week fixed effects, a linear time 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 specifications include a linear trend in the age of the job ad, plus dummies for each of the first three days of an ad's life.
5. Job ad controls include quadratics in worker age requirement, a dummy for valid worker age requirement, wage offered (ln), dummy for explicit wage offered, dummy for bonus offered, dummy for commission payment offered, dummy for requiring new graduates, 10 experience requirement dummies, seven education requirement dummies, 11 firm ownership dummies, 12 firm size (# of workers) dummies, the number of vacancies, and a dummy for whether the number of vacancies is listed.
6. Column 1 and 3 reproduces column 1 and 2 of Table 4.
7. Columns 4 and 5 report the marginal effects at the mean.

Table A15.2: Effects of the Ad Ban on Application Flows to *F* and *M* jobs, by Applicant Gender: Linear Probability Models versus Count Data Models

	(1)	(2)	(3)	(4)
	Applications from Women		Applications from Men	
	to *F* jobs	to M jobs	to *F* jobs	to *M* jobs
a. Linear Probability Model				
Post ban	–0.0156***	**0.0122*** **	**0.0393*** **	0.0053
	(0.0035)	**(0.0011)**	**(0.0018)**	(0.0040)
# of observations	407,209	**441,722**	**407,209**	441,722
R^2	0.035	**0.010**	**0.015**	0.024
b. Negative Binomial Regressions				
Post ban	–0.0092***	**0.0067*** **	**0.0261*** **	0.0087***
	(0.0028)	**(0.0006)**	**(0.0012)**	(0.0033)
# of observations	407,209	**441,722**	**407,209**	441,722
R^2	0.035	**0.066**	**0.060**	0.024

Notes:

1. See notes to Table 4 for regression specifications.
2. The dependent variable is the number of applications received. Its mean is 0.1541, 0.0155, 0.0326 and 0.1788 for columns 1 -- 4 respectively.
3. * p<.10, ** p<.05, *** p<.01.
4. Panel B reports the marginal effects at the mean.
5. **Bold** indicates gender-mismatched applications.

Appendix 16: Effects of the Ad Ban on Employers' Use of the Job Board

Did the prohibition of explicit gender requests on XMRC cause employers to shy away from using that job board? We think this is unlikely, in part because -- as Section 6 showed -- the ban did not reduce the quantity or quality of applications employers received, but also because XMRC is the dominant player in Xiamen's skilled regional labor market: No other platform hosts as many skilled resumes in Xiamen. Also, given the rollout of gendered ad bans across Chinese platforms, it would be reasonable for employers to expect that any alternative to XMRC would either already be covered by the ban, or would soon be covered. This fact should also mute employers' incentives to switch away from XMRC. Still, to investigate this question empirically, this Appendix calculates the changes in some indicators of employers' use of XMRC before versus after the ban, and compares those changes to the changes that occurred around a comparable pseudo-ban in 2018.

Notably, our main analysis sample does not allow us to estimate the ban's effects on employers' propensities to post new job ads on XMRC. This is because that sample, by design, includes only ads that 'span the ban'. This focus allows us to attach an employer's preferred gender request to every ad in our data, giving us a clean measure of exactly which ads were treated by the ban. As a result, our main analysis sample excludes new job ads that were posted, for example, the day after the ban went into effect.

To create estimates of the ban's effects on total employer activity on XMRC, this Appendix therefore uses a more comprehensive sample of job ads. Specifically, our sample now includes all job ads that received at least one application between January 1, 2018 and October 25, 2019, regardless of when the ad was first posted. Using this sample, Figure A16.1 shows unadjusted time trends in the total weekly number of active job ads, and of active firms, in 2018 and 2019, relative to the actual 2019 ban (on March 1) and relative to the lunar equivalent pseudo-ban in 2018 (March 12).[34] These time trends are strikingly similar in the two years.[35] Finally, Figure A16.2 uses the same sample to show time trends in the weekly number of new job ads, where an ad is defined as new in a given week if it received any applications, but no applications are observed for the ad in the preceding four weeks. Again, the time trends are strikingly similar in both years.

[34] A job ad is active in a given week if it received any applications; a firm is active if any one of its job ads received an application. Appendix 11 supplies additional details on how we constructed the 2018 pseudo-ban.

[35] Overall, new job ads are slightly less numerous in 2019 than 2018, but this difference is present both before and after the ban / pseudo ban. Notice also that the dips associated with the Labor Day and National holidays occur at different locations in the graph. This is because these holidays -- unlike the Spring Festival -- are determined by the solar calendar.

Figure A16.1 Counts of Active Jobs and Active Firms over time, 2018 and 2019

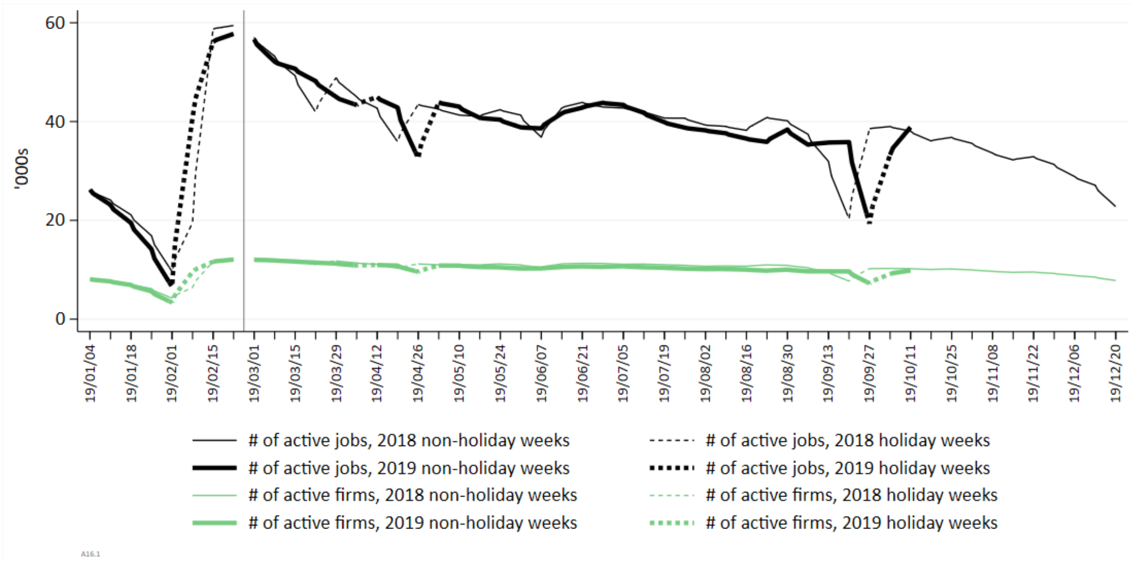

Notes:

1. The dashed parts of the lines correspond to weeks covering the Spring Festivals, Qingming festivals, Labor Day, and National Day for 2018 and 2019.
2. A job ad is considered active if it received any application during the week; A firm is considered active if one of the job ads it posted received any applications during the week.
3. The calendar weeks covered are the week starting on January 12 2018 up to the week starting on October 11 2019.

Figure A16.2 Counts of New Job Ads, 2018 and 2019

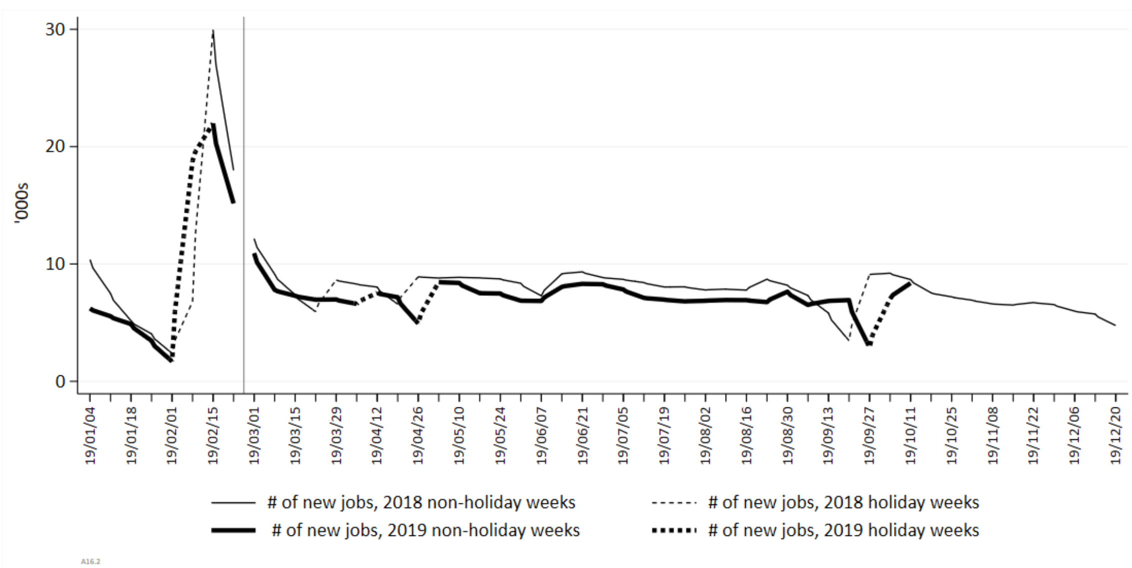

Note:

1. A job ad is “new” on a given week if it received any applications, but no applications are observed for the ad in the preceding four weeks.
2. The dashed parts of the lines correspond to weeks covering the Spring Festivals, Qingming festivals, Labor Day, and National Day for 2018 and 2019.

Appendix 17: Treatment Effect Heterogeneity by Job Characteristics

Job ads on XMRC frequently contain information on job characteristics, including work hours, work schedules, and employee benefits. In this Appendix we study how the ban's effects on applicant gender mix vary with some of these advertised job characteristics. For example, among the jobs whose request for male applicants was removed, we can ask whether the ones that advertised female-friendly work schedules experienced a larger increase in applications from women.

In more detail, we are able to construct two indicators of work schedules, and seven indicators of advertised employee benefits, for all the job ads in our sample.[36] Our indicators of work schedules are (i) the number of working days per week (such as 5, 5.5, or 6), and (ii) the type and number of shifts (such as regular day, regular night, two shifts, or 'irregular').[37] Our benefit indicators include, first, the five most common non-wage benefits (pension, health insurance, unemployment insurance, injury insurance, and fertility insurance), which are usually advertised together; plus two others: a housing provident fund, and private ("commercial") health insurance.

To estimate how the ban's effects on female application shares vary with the preceding job characteristics, we use the following regression specification:

$$Y_{jt} = \beta^1 Post_t \cdot F_j + \beta^2 Post_t \cdot N_j + \beta^3\, Post_t \cdot M_j$$

$$+ \sum\nolimits_{g \in G} \left[g_j \left(\alpha^{1g} Post_t \cdot F_j + \alpha^{2g} Post_t \cdot N_j + \alpha^{3g} Post_t \cdot M_j \right) \right] +$$

$$\textstyle\sum_{k=1,2,3,4} \gamma^k \left(JobWeek_{jt} \right)^k + \sum_{m=1,2,3,4} \theta^k (CalendarWeek_t)^m + \beta X_{jt} + \varepsilon_j + \mu_{ij} \qquad \text{(A17.1)}$$

In words, we start with the specification in Table 1 of the paper, then interact the ad ban's effects by job type (*F*, *N* and *M* jobs) with our new job characteristics variables (g_j). Following column 4 of Table 1, the control vector, X_{jt}, includes quartics in the job's age and calendar time, and job ad fixed effects are included in the regression.

The results are shown in Table A17.1. They indicate, first of all, that the estimated effects of the ban in 'standard' jobs (5-day weekly work schedule, regular day shifts, and none of the seven benefits) are very similar to our Table 1 estimates: Female applicant shares fell (dramatically) in *F* jobs, rose slightly in *N* jobs, and rose (substantially but less dramatically) in *M* jobs. The latter increase in women's representation, however, varied with the job's work schedule: It was 5.6 percentage points in jobs with five working days per week; about half that large in 6-day jobs, and less than one percentage point in jobs with 6.5 to 7 working days per week. A similar pattern applies to the *decline* in women's representation in *F* jobs after the ban: It was 12.7 percentage points in jobs with five working days per week and more than two percentage points greater in jobs with 5.5 or 6 working days per week. Thus, long work schedules *attenuated* women's entry into *M* jobs after they were opened up to women, and *intensified* women's exit from *F* jobs after they were opened up to men. A similar

[36] The job characteristics studied in this Section are all constructed from dedicated fields in a job ad's profile. Additional information about job characteristics is contained in the open text job description of the job ad, which could be exploited using natural language processing methods.

[37] 5.5 days a week most likely means alternating between 2-day weekends and 1-day weekends.

type of heterogeneity is observed across jobs with different daily schedules, though here the significant coefficients are largely confined to the *M* jobs: In *M* jobs with a regular (single shift) daily schedule, women's applicant share rose by 4.7 percentage points. This increase was three percentage points smaller in jobs with 2 or 3 shifts, and four percentage points smaller in jobs with irregular shifts.

Finally, part (c) of Table A17.1 shows effect heterogeneity with respect to advertised benefits. We see no heterogeneity with respect to the five most commonly advertised benefits, which are typically offered together. All five of these benefits are mandatory in China; enforcement of these requirements is also strong and compliance is high. Thus, advertising these benefits appears to convey little information to workers. There is, however, some heterogeneity with respect to the other two benefits -- housing provident funds and commercial health insurance -- in applicants' responses to the ban in *F* jobs: Women's withdrawal from *F* jobs is mitigated when jobs advertise a housing provident fund, and accentuated when jobs offer commercial health insurance.

Some key institutional details regarding these two benefits can help us make sense of these findings. While it is mandatory for Chinese employers to provide a housing provident fund, enforcement and compliance (in contrast to the five 'standard' benefits) are reputedly weak. Thus, mentioning a housing fund in a job ad publicly signals an employers' commitment to provide this benefit (rather than, for example, requiring prospective employees to request it during the recruitment process). Thus, reluctance to 'ask' (Roussille, 2021) could account for women's attraction to job ads mentioning a housing fund. Commercial health insurance, on the other hand, is not government-mandated. Instead it is widely considered to be a 'luxury' benefit used to attract high ability candidates to competitive positions. Thus, women's tendency to apply to less-competitive positions (Flory, Leibbrandt and List 2015) could account for women's relative avoidance of jobs mentioning commercial health insurance.

Together, the heterogeneity analyses in this Appendix increase our confidence in our main estimates, since they show that women's responses to the ban varied in plausible ways across jobs with different advertised work schedules, and with different advertised benefits.

Table A17.1 Treatment Effect Heterogeneity in the Ban's Effects on the Female Share of Applications, by Advertised Working Time Arrangements and Advertised Benefits

	Share of Job Ads, by job's gender preferences			Effects of the Ban on the Female Share of Applications		
	F Job	*N* Job	*M* Job	*F* Job	*N* Job	*M* Job
	(1)	(2)	(3)	(4)	(5)	(6)
			a. Working days per week			
a. Post effect in omitted category (5 days)	42.83%	47.31%	29.76%	–0.1272***	0.0071***	0.0563***
Interactions with working days category:						
working days not stated	19.19%	24.97%	34.68%	0.0041	–0.0055***	–0.0266***
≤ 4.5 days	0.20%	0.17%	0.20%	–0.0069	–0.0050	–0.0533***
5.5 days	15.41%	12.02%	7.73%	–0.0218***	0.0017	–0.0084
6 days	22.06%	15.18%	26.26%	–0.0264***	–0.0063**	–0.0267***
6.5 to 7 days	0.32%	0.35%	1.38%	0.0214	0.0044	–0.0467***
			b. Daily working schedules			
b. Post effect in omitted category (regular day)	52.92%	44.67%	41.29%	–0.1367***	0.0050***	0.0472***
Interactions with schedule category:						
schedule not stated	44.02%	52.22%	50.28%	0.0072	0.0007	–0.0096***
2 shifts	2.22%	1.80%	5.36%	–0.0352	–0.0028	–0.0308***
3 shifts	0.24%	0.32%	1.04%	0.0067	–0.0188*	–0.0347***
4 shifts	0.00%	0.05%	0.15%	0.0000	–0.0078	–0.0080
regular night	0.21%	0.18%	0.22%	–0.0034	0.0028	–0.0039
irregular	0.38%	0.78%	1.66%	0.0344	0.0045	–0.0408***
			c. Advertised benefits			
Main *Post* effect (no benefits offered)				–0.1342***	0.0058***	0.0391***
Interactions with individual benefits:						
Pension	25.97%	20.99%	27.98%	–0.0524	0.0193*	0.0014
Health insurance	26.15%	21.10%	27.96%	0.0282	–0.0037	–0.0097
Unemployment insurance	25.25%	20.40%	27.33%	0.0399*	–0.0016	0.0147
Injury insurance	25.47%	20.53%	27.64%	–0.0144	–0.0059	–0.0200
Fertility insurance	25.17%	20.25%	26.85%	–0.0027	–0.0116	0.0139
Housing provident fund	7.33%	9.39%	10.95%	0.0237***	0.0034	0.0066
Commercial health insurance	5.67%	5.66%	9.58%	–0.0251***	–0.0038	0.0011

Notes:

1. Columns (1), (2) and (3) show the mean of the job characteristic by job type (*F*, *N* and *M*); columns (4), (5) and (6) give the estimated coefficients based on the Table 1 column (4) specification with interaction terms added.
2. All job ads are used here.
3. Observations are ad-week cells; the dependent variable is the share of applications from female applicants. The estimated coefficients are from three regressions, one for each panel. The regressions all have sample size of 1,428,730 and R^2 of 0.647. The regression is weighted by the number of applications, and standard errors are clustered by firm ID.
4. * p<.10, ** p<.05, *** p<.01.
5. The non-wage benefit dummies in panel c are not mutually exclusive.

Appendix 18: Summary of Robustness Tests

In Tables A18.1, A18.2 and A18.3 we collect our estimates of the robustness tests we conducted in Appendices 5, 7, 8, 9, 10, 11, 12 and 15 to provide an overall sense of the messages they convey. Table A18.1 shows that our estimates of the ban's effects on the gender composition of application and call-back pools in Tables 1 and 2 are highly robust to all the specification tests we conducted. Table A18.2 collects our robustness checks of the ban's effects on application flows, quality and yield in previously gendered jobs, in panels c and d of Tables 4 -- 6. For application flows, the estimated coefficients are highly robust, supporting our interpretation of the main mechanisms underlying the ban's integrating effects. For quality and yield, the results support our conclusion of a likely increase in quality and no detectable change in yield.

Finally, Table A18.3 collects our robustness checks of the ban's effects on aggregate application flows, in panel a of Tables 4 -- 6. The robustness tests support our conclusion that if anything, application pools increased in size and quality, while there was no detectable change in application yield. (There is one estimate of a significant decline in yield (from the DiD) analysis, but we believe that this results from the presence of pre-trends for that outcome.)

Table A18.1: Summary of Robustness Checks -- Gender Mix Outcomes

Outcomes	(1)	(2)	(3)	(4)
	Female Share of Applications (Table 1)		Female Share of Call-Backs (Table 2)	
	in *F* jobs	in *M* jobs	in *F* jobs	in *M* jobs
Main Result in the Paper:	–0.1393***	0.0348***	–0.1039***	0.0246***
Robustness Test:				
Appendix 5.1 Functional Form for Time Trends[a]	–0.1390***	0.0350***	–0.1037***	0.0248***
Appendix 5.2 Functional Form for Ad Age Trends[a]	–0.1391***	0.0351***	–0.1039***	0.0247***
Appendix 7 Call-Back Sample Only	–0.1387***	0.0342***	n/a	n/a
Appendix 8.1 Application Reads	n/a	n/a	–0.1327***	0.0317***
Appendix 8.2 Imputed Hires A[b]	n/a	n/a	–0.0925***	0.0259***
Appendix 8.2 Imputed Hires B[b]	n/a	n/a	–0.1229***	0.0363***
Appendices 9.1 & 9.2 Bandwidth (grid search)[c]	All	All	All	All
Appendix 9.3 CCT Optimal Bandwidth[d]	–0.1657***	0.0583***	–0.0902***	0.0144
Appendix 10 Placebo Bans[e]	lowest	highest	lowest	highest
Appendix 11 DiD Approach[f]	–0.1343***	0.0359***	–0.1006***	0.0311***
Appendix 12 Re-weight to national job board[g]	All	All	All	Some

Notes:

1. * $p<.10$, ** $p<.05$, *** $p<.01$.
2. Coefficients show the effect of the ban on the female share in *M* and *F* jobs, relative to *N* jobs from column 4 of Tables 1 and 2, and the corresponding robustness tests.
3. Appendices 6, 13, and 17 are excluded because they conducted heterogeneity analyses, not robustness tests.
4. Appendix 14 is excluded because count data models do not apply to the outcomes in this table (female shares).
5. n/a indicates that the test was not applicable to this outcome.

a: from column 5 of Tables 5.1.1 -- A5.2.2 (separate quartics before and after the ban)
b: from panel b, column 4 of Table A8.2.1 or A8.2.3 (X=Y=7)
c: "All" indicates that all bandwidths (from 3 to 51 weeks) produced very similar estimates to our main estimates (52 weeks)
d: from row 3 Table A9.3.1 (most saturated CCT regression specification)
e: "lowest" ("highest") indicates that the actual treatment effect was much lower (higher) than all the placebo estimates
f: estimates are from column 4 of Tables 11.1.1 and 11.1.2
g: "All" means that all of the re-weighting schemes produced a statistically significant estimate, very similar to the main estimates. "Some" means that some of the treatment effects (on M jobs) were statistically insignificant after re-weighting.

Table A18.2: Summary of Robustness Tests -- Arrivals, Quality and Yield in Previously Gendered Jobs

		(1)	(2)	(3)	(4)
		F to F	**M to F**	**F to M**	M to M
		a. Application Arrivals			
Main Result in the Paper:		–0.0160***	**0.0388*** **	**0.0122*** **	0.0043
Robustness Test:					
Appendix 7	Call-Back Sample Only	–0.0221***	**0.0434*** **	**0.0127*** **	–0.0005
Appendices 9.2	Bandwidth (grid search)[a]	–0.0175***	**0.0365*** **	**0.0112*** **	–0.0048
Appendix 11	DiD Approach[b]	–0.0171***	**0.0513*** **	**0.0189*** **	0.0026
		b. Application Quality			
Main Result in the Paper:		0.0268	**0.1064***	**0.0217**	0.0298**
Robustness Test:					
Appendix 7	Call-Back Sample Only	0.0370*	**0.1224***	**–0.0074**	0.0265
Appendices 9.2	Bandwidth (grid search)[a]	0.0430***	**0.0733***	**–0.0034**	0.0184
Appendix 11	DiD Approach[b]	0.0278**	**0.0559**	**0.0450**	–0.0002
		c. Application Yield			
Main Result in the Paper:		0.0017	**0.0004**	**–0.0101**	–0.0030
Robustness Test:					
Appendix 7	Call-Back Sample Only	n/a	**n/a**	**n/a**	n/a
Appendices 9.2	Bandwidth (grid search)[a]	–0.0038	**–0.0031**	**–0.0271**	–0.0001
Appendix 11	DiD Approach[b]	–0.0122	**0.0134**	**–0.0377**	–0.0068

Notes:

1. **Bold** indicates gender-mismatched applications
2. * *p*<.10, ** *p*<.05, *** *p*<.01.
3. n/a means the test does not apply to this outcome

a: The reported coefficients are for a 55 day window -- the most different from the baseline. Coefficients for shorter windows are very similar.

b: Weekly arrivals coefficients in Table A11.2.1 have been converted to daily arrivals.

Table A18.3: Summary of Robustness Tests -- Aggregate Search Frictions

		(1)	(2)	(3)
Outcomes		Application Arrivals (Table 4.A)	Application Quality (Table 5.A)	Application Yield (Table 6.A)
Main Result in the Paper:		0.0058***	0.0169***	–0.0028
Robustness Tests:				
Appendix 7	Call-back Sample Only	0.0004	0.0219***	n/a
Appendix 9.2	Bandwidth (grid search)	weakly positive[a]	strongly positive[b]	0[c]
Appendix 10	Placebo bans	weakly positive[d]	weakly positive[d]	0[e]
Appendix 11	DiD approach	0.0097***[f]	0.0093**	–0.0055**
Appendix 15	Count data model	0.0101***	n/a	n/a

Notes:

1. * p<.10, ** p<.05, *** p<.01.
2. n/a means the test is not applicable to the outcome in question

a: 8 of the 10 significant estimates are positive; the rest are insignificant
b: All 25 estimates are positive and significant
c: All 25 estimates are insignificant
d: Our positive main result is greater than most but not all placebo estimates,
e: Our insignificant, negative main result is in the middle of the placebo estimates.
f: Table A11.2.1's coefficient for applications per week has been converted to a daily rate.