The Quarterly Journal of Economics Advance Access published November 18, 2012

GENDER DISCRIMINATION IN JOB ADS:

EVIDENCE FROM CHINA*

Peter Kuhn and Kailing Shen

We study explicit gender discrimination in a population of ads on a Chinese internet job board. Gender-targeted job ads are commonplace, favor women as often as men, and are much less common in jobs requiring higher levels of skill. Employers' relative preferences for female versus male workers, on the other hand, are more strongly related to the preferred age, height and beauty of the worker than to job skill levels. Almost two thirds of the variation in advertised gender preferences occurs within firms, and one third occurs within firm*occupation cells. Overall, these patterns are not well explained by a firm-level animus model, by a glass-ceiling model, nor by models in which broad occupational categories are consistently gendered across firms. Instead, the patterns suggest a model in which firms have idiosyncratic preferences for particular job-gender matches, which are overridden in skilled positions by factors such as thinner labor markets or a greater incentive to search broadly for the most qualified candidate. JEL Codes J16, J63, J71.

---

* Correspondence to Peter Kuhn, epartment of Economics, University of California, Santa Barbara, Santa Barbara, CA USA 93106 fax 805 893 8830, tel 805 893 3666 pjkuhn@econ.ucsb.edu. The authors thank UC Santa Barbara's Institute for Social, Behavioral and Economic Research (ISBER) for financial support. Kuhn thanks UC Berkeley's Center for Labor Economics and the IZA for their generous hospitality during the research on this project. We thank Loren Brandt for providing us a copy of the 2005 Chinese Census microdata. Catherine Weinberger and participants at numerous seminars and conferences provided many helpful comments.

I. INTRODUCTION

A fundamental problem facing empirical research on discrimination is the fact that discrimination is illegal in almost all the markets that are studied. Given market participants' obvious incentives to keep discriminatory actions hidden, researchers have been forced to rely on a variety of indirect methods to infer whether discrimination has occurred. Yet explicitly discriminatory actions, such as posting a job ad which states that a particular race or gender is preferred, were commonplace in the United States as recently as 1960 (Darity and Mason 1998). Indeed, they continue to be both legal and accepted in many if not most of the world's labor markets.

In this paper, we study one of these labor markets –China's—in order to establish a new set of facts about gender discrimination. Using data from a large sample of Internet job ads, we find that explicit gender discrimination is commonplace. For example, during our twenty-week sample period, over one third of the firms that advertised on the job board --which caters to highly educated urban workers seeking private sector jobs-- placed at least one ad stipulating a preferred gender. We also find, perhaps surprisingly, that the share of ads favoring men versus women is roughly equal. Put another way, when it is legal to express gender preferences in job ads, a significant share of employers chooses to do so, and uses gendered ads to solicit women as often as men.

One of the strongest and most robust relationships in our data is a negative *skill-targeting relationship*: as a job's skill requirements rise, the share of ads stipulating a preferred gender declines. This relationship holds whether we measure skill by the job's education requirements, experience requirements, or its advertised wage, and is present in simple comparisons of means as well as regressions that control for detailed job attributes. While we do not find a robust relationship between skill levels and the *direction* of firms' gender preferences, we do find that the direction of firms' preferences is highly correlated with firms' requests for other ascriptive worker characteristics, specifically age, height and beauty. For example, firms' requests for young, tall and attractive workers are strongly associated with explicit requests for women, while requests for older workers correlate with requests for men. Finally, we find that close to two thirds of the variation in firms' advertised gender preferences takes place within firms, and one third is within firm*occupation ("job") cells. While occupation is correlated with firms' gender preferences, the gendering of occupations is not very consistent across firms.

Putting these patterns together, we argue that they are not consistent with a model where advertised gender discrimination is driven primarily by firm-level preferences for one gender over another, nor with 'glass ceiling' models where firms' primary motivation is to restrict women's access to highly-skilled or well-paid jobs. Models in which broad occupational categories are consistently gendered across firms also do a poor job of explaining why some firms request (say) more women than others. Instead, Chinese firms mostly use targeted job ads to divide their pool of less skilled positions into 'male' versus 'female' jobs, while largely abandoning these distinctions at higher skill levels. To account for this pattern, we offer a simple model of hiring from two pools of workers. In the model, firms have job-specific preferences for a worker's type, in which factors such as beauty-related customer discrimination and social perceptions of gender-appropriate work can play important roles. These preferences must, however, be balanced against a desire to fill jobs with the best candidate. When this desire becomes more important (for example, when the pool of potential applicants is small or the job's skill level is high), firms broaden their search to include the less-preferred pool, thus abstaining from type-targeted job ads.

II. RELATED LITERATURE

Legally, and according to most dictionary definitions, "discrimination" refers to taking an action, such as paying a different wage or choosing to hire a person, based not on that person's individual merit but on his/her membership in a particular group.[1] As already noted, when such acts are illegal, economists generally resort to indirect methods, such as searching for unexplained wage gaps, to infer whether they have occurred. Indeed, a large literature has used such indirect methods to study both whether (and where) employer discrimination occurs, and what motivates it (e.g. statistical versus taste motives). Recent reviews are available in Bertrand (2010), Fryer (2010), and Charles and Guryan (2011).

In contrast, to our knowledge, previous studies of explicit discrimination in labor markets are rare. Darity and Mason (1998) reproduce examples of discriminatory ads from U.S. newspapers in 1960, but do not conduct any statistical analysis. Goldin (1990, 2006) examines data from a Department of Labor Women's Bureau survey of employers concerning their 1939 employment policies for office workers, and Lawler and Bae (1998) study gender preferences in a sample of 902 ads placed in an English-language newspaper in Thailand.

[1] For example, the dictionary.com definition is "to make a distinction in favor of or against a person or thing on the basis of the group, class, or category to which the person or thing belongs rather than according to actual merit".

Aside from the above, the discrimination studies that are probably most closely related to ours are correspondence studies such as Bertrand and Mullainathan (2004), who study employers' reactions to identical resumes with randomly assigned race or gender. Also related are full audit studies such as Neumark (1996), which carry this procedure one step further and send matched, trained actors to interview for jobs in response to the callbacks. That said, it is important to note that our ad-based approach measures different aspects of employer discrimination than are measured by the audit/correspondence approach. For example, because ads are formulated before resumes arrive, ad-based measures do not condition on the information that appears in a worker's resume. This means we could conceivably see ad-based discrimination even in situations where audit-based studies show no discrimination: If an employer expects to receive lower quality resumes from the 'disfavored' group, firms might choose to engage in discriminatory job advertising even they would not discriminate between identical resumes from the two groups. At the same time, however, note that targeting an ad excludes not just an individual but an entire group of applicants from consideration; one might expect (and we show theoretically) that such actions will be reserved for cases when firms have relatively extreme preferences for one group over another. Thus, when firms' gender preferences are mild, we might not see gender-targeted ads, but could still observe discrimination in an audit study.

A final difference between ad-based and more familiar measures is that advertised discrimination necessarily involves a conscious decision by the employer to invite only one group to apply. In contrast, audit and other studies are designed to detect both the conscious choices and unconscious biases of employers. In sum, we should expect firms to post discriminatory ads when they *consciously* expect *large* differentials in productivity (or desirability) between two or more groups.[2] These differentials include expected differences in the quality of resumes that would arrive in response to the ad. Taken together, these conditions differ from the conditions in which we would see discrimination in other types of studies, such as audit-based ones.[3]

[2] It is perhaps worth noting that the mere presence of *either* ad-based or audit-based discrimination sheds no direct light on whether the employer's motives for discrimination are taste-based or purely statistical. Distinguishing these sources of discrimination requires additional evidence, such as its distribution across firm and job types, studied later in this paper.

[3] In addition to these conceptual differences, an attractive feature of an ad-based approach (where it is feasible) is its low cost. In contrast to audit and correspondence studies, which typically restrict their attention to a small number of narrowly-defined occupations and/or cities, our approach allows us to paint a statistical portrait of gender discrimination across an entire labor market, spanning a wide array of occupations, industries, provinces, and firm types.

III. DATA

Our data is the universe of unique job advertisements posted on Zhaopin.com, the third largest online job board in China, during four observation periods: May 19 2008 - June 22 2008, January 19 2009 - February 22 2009, May 18 2009 - June 21 2009, and January 18 2010 - February 21 2010.[4] Procedures for downloading the data and defining variables are discussed in the Appendix. Descriptive statistics are provided in Table I. All told, we study a total of 1,051,706 job ads, 41 percent of which require a bachelor's degree or more, 46 percent of which require some post-secondary education, and 13 percent of which require a high school degree or less. Overall, just over one in ten ads on the site expressed a gender preference; this was evenly split between preferences for men and women.[5] About 80 percent of ads required at least a year of experience. Explicit age requirements (which, like gender requirements, also violate some U.S. antidiscrimination laws), appear in 24 percent of all ads. While the mean requested age is quite young (30.6), ads specify a minimum age almost as often as they specify a maximum age.

Of potential relevance to the design of equilibrium job search models (e.g. Rogerson et al. 2005), only about 16 percent of Chinese job ads contain information that could reasonably be construed as a posted wage.[6] Among those ads, the mean wage was 4,279 RMB per month.[7] Another piece of information that is often provided in Zhaopin job ads is the number of positions advertised: this is unspecified about half the time; when specified the modal number is one and the mean is 1.7. Just under 8 percent of the ads explicitly requested that the applicant be physically attractive "*xingxiang*", and 2.6 percent stipulated a minimum height.[8]

Because job ads on Zhaopin are linked to the firm's page on the site, and this page contains some standardized information, we can follow these links to gather measures of firm size and ownership. The average ad was from a firm with 1,565 employees, and an overwhelming share (92.6 percent) of the ads were placed by private-sector firms (both foreign-

[4] Note that firms frequently re-post the same ad; our sample includes each unique ad only once. Our choice of Zhaopin is largely for the technical reason that its site structure allowed us to easily and accurately identify such renewals.
[5] Our measure includes all intensities of preference, though the most typical employer statements were either "female [male] preferred" and "female [male] only".
[6] Zhaopin.com prompts employers to enter a minimum and maximum wage. In most cases these fields are blank; in other cases, employers select both the lowest possible wage (1,000 RMB/month) *and* the highest (50,000 RMB/month) from the drop-down menu for a particular job. After eliminating all ads with a gap of more than 20,000 RMB between the minimum and maximum wage, along with other uninformative answers, we are left with 16 percent of all ads. The number reported is the midpoint between the advertised minimum and maximum.
[7] At the 2009 exchange rate of 6.83RMB/dollar, this corresponds to $627 US, which is high even by urban Chinese standards. Over three quarters of urban workers in China's eight highest-income provinces earned under 1500 RMB, or $220 US per month. See Table A.1.
[8] Loosely, *xianxiang* translates as having a pleasing image, form or figure.

and Chinese-owned). The remaining employers in our sample are State-Owned Enterprises (SOEs).[9] Fully 36 percent of all ads were placed by firms with some foreign ownership. (Our "foreign-owned" category includes Foreign Direct Investment (FDI), joint ventures, and a small number of representative offices). Finally, we note that Zhaopin.com categorizes ads into 39 occupation and 46 industry categories.[10] Especially for the jobs on this site (which are disproportionately skilled, service jobs) these categories provide more detail than is available in social surveys or in available Census microdata. Together with the site's indicators for the province where the job is located, these industry and occupation categories allow us to implement detailed controls for the type of job on offer.

To assess how our sample of jobs ads compares to the overall population of jobs in the Chinese labor market, Appendix Table A.1 compares the characteristics of Zhaopin ads to a representative sample of employed urban workers in China's eight highest-income provinces, taken from the 2005 Census.[11] These eight provinces --Beijing, Shanghai, Guangdong, Jiangsu, Shandong, Tianjing, Zhejiang and Liaoning—coincidentally have the eight highest GDPs per capita and are also the eight provinces with the most Zhaopin ads. The main result is that, compared to employed persons in the same set of provinces, Zhaopin ads are aimed at workers who are much younger, better educated, better paid, and who work in the private sector. For example, while about 28 percent of the workforce in these areas is over 40 years old, only three percent of Zhaopin ads stipulate a desired age over 40.[12] While 77 percent of workers in the high-income provinces had high school education or less in 2005, only 11.4 percent of Zhaopin ads request this level of education, reflecting the job site's focus on skilled occupations and on a highly-educated younger generation of workers. State-owned enterprises employed about 27 percent of workers, but only account for 7 percent of Zhaopin ads. More than three quarters of workers in these provinces earned under 1500 yuan per month, while only 14.5 percent of the

[9] We dropped a very small number of ads from non-profit organizations and for government jobs. In China, recruiting for government jobs takes place through a separate testing and recruitment system.

[10] The most common occupation is sales, at 18 percent of the ads, with IT second at 11 percent. The top five industries were construction, consulting, IT, marketing, and trade, in order.

[11] We expect the mix of Zhaopin ads to differ from a representative sample of jobs for several reasons. First, relative to a sample of occupied jobs, any sample of vacancies will overrepresent entry-level jobs, as well as jobs in expanding and high-turnover occupations and industries. In addition, the vacancies on an internet job board likely require a significantly higher skill level than the median vacancy. Finally, relative to a flow sample of new vacancies, long vacancy spells will be overrepresented in our data; this would affect our estimates if there is parameter heterogeneity in the determinants of ad content that is correlated with vacancy durations.

[12] The working population in these cities is very young by U.S. standards for at least two reasons. First, the Census data include migrants from rural areas and from other cities, who account for 32.5% of the employed urban population and have an average age of 29.8. Second, China has a number of policies that strongly encourage early retirement.

wages posted on Zhaopin ads fall below this level.[13] In sum, even within the high-income provinces that are disproportionately served by Zhaopin, it is clear that Zhaopin.com –while hosting a wide variety of jobs and imposing no restrictions on who can use the site--disproportionately caters to a very special slice of the Chinese labor market: young, highly-educated workers looking for well-paid jobs in the private sector.[14]

To shed additional light on the relation between Zhaopin ads and the overall Chinese workforce, Appendix Table A.1 also compares the gender mix of the workforce across various categories with the gender mix of Zhaopin ads. Overall, the key message of these comparisons is that conditional on stating a gender preference, Zhaopin employers are more likely to request women in the occupations, industries, age categories, education levels, and wage levels where women are also more likely to work. Combined with our anecdotal impressions of Chinese hiring practices, this pattern suggests that voluntary affirmative action --i.e. a practice of deliberately seeking to increase hiring of the minority gender in a given occupation or job—is not a widespread motivation for the gendered job ads in our data.

IV. SOME FACTS

A first, overall look at the types of jobs for which firms say they want male or female employees is provided in Table II. This table reports three fractions --the share of job ads requesting women, those requesting men, and those not indicating a gender preference—for samples of job ads differentiated by indicators of skill demands, by other requested ascriptive characteristics, and by firm characteristics.

Panel A of the Table focuses on the correlation between job skill requirements and advertised gender preferences. It shows that the propensity to gender-target a job ad is strongly related to the job's skill level: only about 6 percent of jobs requiring a university education specify a preferred gender, while 23 percent of jobs requiring high school or less are explicitly gendered in this way. Strikingly, the share of ads requesting men, and the share requesting women, both decline precipitously with education requirements. For the subsample of ads that

[13] In terms of industry and occupation, the IT/ communication and R&D/consulting industries are highly overrepresented on Zhaopin, while manufacturing but also health/education/welfare are underrepresented. Professional and technical workers are highly overrepresented on Zhaopin, while production and construction workers are underrepresented. See the Appendix for additional details.

[14] iResearch Inc, (2010) estimates that 11.2 percent of Chinese firms covered by the State Administration for Industry and Commerce posted ads on at least one Internet Job Board in 2010. Anecdotal evidence from firms posting on these sites suggests that their main alternative to using an Internet job board is on-campus recruiting. At lower skill levels, both newspaper ads and local government employment offices can also play important roles in recruiting. Note, however, that none of these recruiting methods are mutually exclusive; many ads appearing on sites like Zhaopin may also be posted in these other venues.

post a wage, exactly the same pattern is visible: the share of ads requesting men and the share requesting women both decline precipitously with the advertised wage. For example, 27 percent of jobs paying under 1500 yuan/month specify a preferred gender, compared with 7 percent of jobs offering 8000 yuan or more. Together, these patterns are inconsistent with simple glass-ceiling models where the primary purpose of advertised restrictions is to keep women out of highly-skilled or well-paid jobs.[15] Finally, Panel A shows that the amount of gender targeting also falls with a job's experience requirements, though here the decline occurs only among ads requesting women. The lack of a decline for men is related to a tendency for firms to prefer men when they are seeking to hire older workers, documented in Panel B below. Together, we refer to the decline in gender-targeting of job ads with all three measures of job skill requirements in our data as the *skill-targeting relationship*.

Panel B of Table II shows the relationship between advertised gender discrimination and firms' advertised preferences for three other ascriptive worker characteristics: age, beauty and height. One immediate conclusion is that advertised preferences for all these other attributes are complementary with those for gender, in the sense that these ascriptive screens tend to be used together. Thus, only 5.6 percent of ads with no age restrictions are explicitly gendered, compared with 28 percent of ads stipulating a specific age range. 9.1 percent of ads that do not request beauty are gendered, compared with 27.7 percent of ads that do request beauty. And 9.3 percent of ads that have no height requirement are gendered, compared with 56.2 percent of ads that do list a minimum height. In sum, gender preferences are much more likely to appear in job ads when the firm also requests workers of a specific age, height, or workers who are good looking.

A second clear pattern in panel B occurs in the subsample of ads that specifies both a maximum and minimum age, so an "ideal" age (the midpoint of this range) can be calculated. Here, the patterns are starkly opposite for men and women: as the desired age rises, firms are more likely to look for men, and are less likely to be seeking women. Furthermore, these effects are large in magnitude: the share of ads requesting men triples as the desired age rises from under 25 to 35+, and the share of ads requesting women increases eight-fold (from 4 to 33 percent) as the desired age falls from 35+ to under 25. This suggests that, in contrast to skill, which has relatively weak and inconsistent effects on the direction of firms' gender preferences, age plays an important role in firms' preferences for men relative to women for a job.

[15] That said, note that the share of *targeted* ads directed at women does fall with education, experience and the wage. Unlike our overall result on the incidence of targeting, this general 'tilt' towards men as skill requirements rise does not survive controls for other observable ad characteristics. See Table VII.

Finally, panel C of Table II shows that there is no strong relationship between a firm's size and its propensity to gender target its job ads. Thus, to the extent that having many employees is associated with having formal personnel and HRM policies, our data show no obvious incompatibilities between having such policies and engaging in explicit hiring discrimination. That said, foreign-owned firms behave differently from Chinese-owned firms, engaging in less advertised gender discrimination than either private Chinese-owned firms or State-Owned Enterprises. While the differences between SOEs and private Chinese firms are not dramatic, SOEs are somewhat less likely than privately-owned Chinese firms to request women in the hiring process.

Table III explores the striking association between firms' requests for gender, beauty, height and age in more detail by looking at some intersections between these characteristics. Specifically, recall that 5.0 percent of ads explicitly request women (Table I), and that this rises to 23.9 percent if we look only at ads requesting physically attractive applicants (row 1 of Table II). According to Table III, this number rises to 55.9 percent among ads requesting applicants who are tall and good looking, and to 87.1 percent if the applicant is also required to be under 25 years of age. These dramatic changes in conditional probabilities illustrate the high predictive power of demands for other ascriptive characteristics in 'explaining' firms' requests for women. At the same time, it is worth noting that --since only 7.7 percent of ads request beauty and 2.6 percent of ads request height-- this particular constellation of employer desires can only account for only a minority of firms' advertised gender requests. Thus, a more general theory is still needed to understand the wider patterns of gender discrimination in our data.

Yet another perspective on the frequency of discriminatory job ads in China emerges when we organize our data by firms rather than ads. Overall, 73,642 distinct firms placed ads on Zhaopin during our sampling period; thus the typical firm placed 1,051,706 / 73,642 = 14.3 ads. Characteristics of these firms' hiring policies are summarized in Table IV. According to Table IV, 20 percent of the firms in our data placed at least one ad that specifically invited men to apply; for women this number was 25.8 percent. For obvious reasons, these shares rise with the number of ads the firm placed on Zhaopin during our sample period. Thus, for example, among firms that placed more than 50 ads, over 70 percent expressed a gender preference at least once, and 39 percent placed both male-only and female-only ads during our sample period. This suggests, perhaps surprisingly, that a substantial share of the variation in firms' advertised gender preferences in our data might occur within, rather than between firms.

Additional detail on the role of firms, occupations, and their intersection ("job cells") in accounting for patterns of explicit hiring discrimination is provided in Table V, which presents a decomposition of the variance in our three main indicators of discrimination. Our approach follows Groshen's (1991) decomposition of wages into occupation- and firm-specific components. Accordingly, row 4 of Table V reports the R^2 from a regression of advertised discrimination on a full set of firm and occupation fixed effects. Row 1 calculates the minimum contribution of occupation effects to that regression as the reduction in R^2 when the occupation effects are removed from it. Row 2 performs the analogous exercise for firm effects, while row 3 equals row 4 – (row 1 + row 2). In sum, row 4 reports the total variance that can be explained by occupation and firm effects, and rows 1-3 in order partition this variance into components that can be unequivocally assigned to occupation effects, to firm effects, and a portion that cannot be unequivocally assigned to either of these two factors.

Next, row 6 of Table V presents the R^2 from a regression of advertised discrimination on a full set of occupation*firm interactions. Thus, row 6 shows the total variance that can be explained by these "job cells", and row 7 (which is just one minus row 6) is the within-job-cell variance. Finally, row 5 is the difference in R^2 between the regression with occupation*firm interactions (row 6) and the regression with occupation and firm effects (row 4). It is a measure of the extent to which firms gender their occupations differently; for example it would equal zero in column 1 if the tendency for sales jobs to be 'male' was the same in all firms, and similarly for all other occupations.

The decomposition results in Table V are highly consistent across our three indicators of hiring discrimination. For all three measures, the 39 occupation categories in our data, while highly statistically significant as a group, explain only a small share (3 percent or less) of the variance in explicit discrimination. Thus, a simple occupational segregation model, in which differences in occupation mix explain why some firms explicitly search for women and others search for men, is not very powerful.[16] Firm effects consistently explain between 28 and 32 percent of the variation in advertised discrimination, but the total explained by firms and occupations together is between 33.4 and 37.3 percent.[17] Thus, while firm effects are important, almost two thirds of the variation in advertised discrimination occurs within firms. For that

[16] Somewhat more formally, note that according to rows 1 and 4, removing the occupation effects from the column 1 regressions reduces the R^2 very little, from .334 to .334-.017 = .311. Thus, only .017/.334 = 5.1 percent of the between-firm variation in gender preferences can be explained by the fact that firms hire different mixes of occupations.

[17] Blau's (1977) classic study of office workers also finds considerable gender segregation across firms, and shows that this segregation accounts for a substantial share of the gender wage gap.

reason, as already noted, a simple firm-level animus model (like a simple glass-ceiling model) cannot account well for the patterns in our data.

Recalling that Row 5 of the Table measures the extent to which different firms gender occupations differently, we see that this tendency explains a significant share –about one third--of the total variance in advertised discrimination, or about one half of the explained (between-cell) variation. Thus, it appears that occupations are not very consistently gendered across firms. To the extent that our occupation categories measure the type of work performed in a job, this suggests that the broad nature of the work performed may not be a powerful predictor of whether firms seek men or women for a job.[18]

Finally, according to row 6, a full set of occupation*firm interactions explains about two thirds of the variation in advertised discrimination.[19] While these factors are quite powerful, it follows that one third of the variance in discrimination occurs within job cells, i.e. because the same firm, at different times during our 20-week observation window, sometimes (say) requests men for sales jobs and sometimes does not.[20] One interpretation of this pattern is, of course, that individual firms' occupation-specific gender preferences are highly fluid over time. Alternatively, firms might consistently gender jobs at a finer level than our occupation indicators capture, while the mix of detailed job titles they are trying to fill changes from month to month.[21] The latter interpretation is more consistent with Bielby and Baron's (1984) striking evidence of highly detailed gender segregation in U.S. firms.

In sum, Tables I through V have illustrated five key features of advertised gender discrimination in China. First, employers on this website (which caters to highly-qualified young Chinese workers) use gender as an *ex ante* hiring screen quite frequently. Second, explicitly gendered job ads favor women as often as men. Third, job ads are much less likely to be gendered when the job requires more education and/or experience, and when the job offers a high wage. Fourth, gender, as an *ex ante* hiring screen, tends to be used in conjunction with

[18] This inconsistent gendering of occupations, combined with firms' frequent reliance on gendered job ads, is however consistent with a scenario in which firms strive to maintain gender homogeneity in detailed job categories, perhaps for reasons related to worker identities (Goldin 2006, Akerlof and Kranton 2008).

[19] We replicated Table V's decomposition with alternative restrictions on minimum job cell sizes. If we restrict attention to job cells containing at least five ads, the share of variance that is within cells rises to about 50 percent. It remains stable at that level for much larger minimum cell sizes (at least up to 25).

[20] To illustrate the nature of within-job-cell variance in advertised gender preferences, consider for example the 3415 firms that placed 10 or more ads for salespeople. Of these firms, 467 (or 13.7 percent) placed at least one ad requesting a male salesperson and one nongendered ad for sales personnel. 620 (18.2 percent) placed at least one female ad and one nongendered ad, and 167 (or 4.9 percent) placed both a 'male' and a 'female' ad. About the same number (164 firms or 4.8 percent) placed ads of all three types. For obvious reasons, these numbers get larger as we focus on firms that place more ads for salespeople.

[21] A related possibility is that ads directly reflect taste-based discrimination by hiring agents or co-workers, but that hiring is done by more than one agent and for more than one set of co-workers within a firm*occupation cell.

requirements for other ascriptive characteristics. In particular, firms that are looking for older workers are much more likely to be requesting male applicants. But if a firm is looking for someone who is tall, good-looking and under 25, it will almost certainly also request a woman in the ad.

Fifth, almost two thirds of the variation in advertised gender discrimination in our data takes place within firms, and one third occurs within firm*occupation cells. In other words, it is commonplace to see the same firm request (say) men for sales jobs in one ad, then either to request women or not to stipulate a gender preference in another ad for a sales job placed at a different time. This suggests either that firms' gender preferences for particular types of work are highly fluid, or that firms assign genders to jobs at finer levels of detail than our data can measure, which we think is more likely.

V. A MODEL OF DISCRIMINATION IN JOB ADS

As we have noted, a negative skill-targeting relationship is a prominent and robust feature of our data. In this Section we present a simple model of job advertising that is consistent with this pattern, and that helps us think more formally about the conditions under which advertised discrimination is likely to occur. To maximize transparency and intuition, we adopt a nonsequential search framework and take a partial equilibrium approach.[22]

V.A THE MODEL FOR A SINGLE JOB AD

Consider a firm soliciting applications for a single vacant position; applications can come from two groups, labeled *M* and *F*; *M* and *F* also denote the number of applications that would arrive from each group, if invited.[23] Let the net value to the firm of an individual applicant, *j*, in a job with 'standard' skill requirements be

(1) $U_j = v^G + \varepsilon_j \;, \; G \in (M, F)\,,$

where the ε_j represent independent draws from a distribution with *cdf* $F(\varepsilon_j)$. A worker's net value in a job with skill requirement θ is assumed to be θU_j. Importantly, we think of the gender difference in baseline net value, v^M - v^F, as including not only between-group differences in expected revenue productivity and wage costs, but also differences in employer tastes; thus the

[22] While there has been an active theoretical literature on equilibrium labor market search (see Rogerson et al. 2005 for a review), as far as we know none of this literature has studied firms' optimal choice of *how broadly* to search. (i.e. which types of workers, whether differentiated by education, experience, age, or sex, to invite to apply).

[23] In this Section, we assume for simplicity that the number of applicants of a given type that arrive is the same, regardless of whether the ad is targeted to that type or not. We extend the analysis to allow workers to direct their search in response to targeted ads in Section VII.C.

model allows for both taste-based and statistical discrimination.[24] Bearing this in mind we shall often use employers' "preferences towards men" as shorthand for $v^M - v^F$. The only other element of the model is a cost of processing the job applications that arrive; by incurring this cost the firm learns the applicant's idiosyncratic suitability for the job, ε_j. In the basic model described in this section, processing costs are c per application.[25]

Our question in this environment is a simple one: Assuming that the firm hires the worker it likes best (i.e. the one with the highest U_j) from its applicant pool, which groups (M, F, or both) should it invite to apply? Formally, the firm chooses D^M and D^F to maximize:

(2) $$\Pi \equiv Emax(U_j;\, D^M, D^F) \;-\; D^M cM \;-\; D^F cF,$$

where D^M (D^F) is a (0,1) indicator for inviting men (women) to apply, and $Emax(U_j;\, D^M, D^F)$ gives the expected value of the maximum U_j drawn from the sample of applicants defined by D^M and D^F. Since we show below that the firm will never invite only the dispreferred (lower v) group to apply, this boils down to a choice between searching narrowly (in a small pool with a higher expected value), or searching broadly (in a group with a lower overall mean, but more candidates to choose from).

Simple, closed-form solutions to this problem are available when $F(\varepsilon_j)$ has a type-I extreme value distribution, with $F(\varepsilon_j) = exp(-exp(-\varepsilon_j/\beta))$. It follows that $Var(\varepsilon_j) = \beta^2\pi^2/6$, and $E(\varepsilon_j) = \beta\gamma$, where γ is Euler's constant ($\approx$.577). In that case, we can prove:

Lemma 1.

(a) The expected value of the highest U_j in a sample of size G drawn from a single group, M or F, in a job with standard skill requirements, is:

(3) $$U^G{*} = \mu^G + \beta\log(G), \quad G \in (M, F),$$

where $\mu^G \equiv v^G + \beta\gamma$ is the expected net value of a single applicant from group G.

(b) The expected value of the highest U_j drawn from the combined sample of all applicants in a job with standard skill requirements, is:

(4) $$U^C{*} = \beta\log\left[\delta\exp\left(\frac{\mu^M}{\beta}\right) + (1-\delta)\exp\left(\frac{\mu^F}{\beta}\right)\right] + \beta\log C$$

[24] Distinguishing between taste- and statistical motives for discrimination requires evidence on patterns of discrimination across different types of jobs and firms; we present such evidence later in this paper. Here, we simply note that our model's main predictions for when firms will choose to engage in explicit, *advertised* discrimination hold regardless of whether the firm's underlying motive is taste- or productivity-based.

[25] None of the main results change when screening costs are an increasing, nonlinear function of the number of applicants. We consider an alternative screening technology in Section VII.C.

where $C=M+F$ and $\delta = M/(M+F)$.

Proof: See on-line Appendix 1**.**

Lemma 1 shows that the expected value of the best candidate in both the separate (*M* and *F*) pools and the combined (*C*) pools is linear in the log of the pool's size, with a slope parameter (β) that is proportional to the standard deviation of applicant quality. Thus, additional applicants are more valuable when match quality is hard to predict. To keep the notation simple in the remaining analysis, we focus on the case of an equal number of applications from each group, i.e. δ=.5.[26] In that case, (4) can be written:

(5) $$U^{C}* = \mu^{M} + \beta \log\left[1 + (\exp\left(\frac{\mu^{F} - \mu^{M}}{\beta}\right)\right] + \beta \log N ,$$

where N=.5C is the (common) number of applications expected from each group. Now, defining $z \equiv (\mu^M - \mu^F)/\beta$ as the standardized gap in expected net value between the groups,[27] the firm's optimal recruiting policy is described by:

Proposition 1. The firm's optimal recruiting policy is to:

Solicit men only if $z > z^*$,

Solicit women only if $z < -z^*$

Post no advertised restrictions if $-z^* \leq z \leq z^*$

where:

$z^* = -\ln[\exp(cN/\theta\beta) - 1] > 0$ if $cN/\theta\beta \in [0, \ln(2)]$ ("high frictions" case), and

$z^* = 0$ if $cN/\theta\beta > \ln(2)$ ("low frictions" case).

Proof: See on-line Appendix 1**.**

Proposition 1 shows that the factors influencing a firm's optimal recruiting strategy fall naturally into two categories: Unsurprisingly, factors that raise the index *z*, which is the expected productivity, cost and taste advantage of men for the job, raise the 'likelihood' that firms will invite only men to apply, and reduce the likelihood that firms will invite only women.[28] As shorthand, we will refer to these factors as *relative preferences*, or *preferences towards men*. The remaining parameters in the model, *c*, *N*, θ, and β, operate only on the

[26] Results for unequal numbers are presented in an earlier version of this paper (Kuhn and Shen, 2009); nothing of importance changes.

[27] Note that $\beta = \sigma_\varepsilon \sqrt{6}/\pi$, where σ_ε is the standard deviation of net value.

[28] We formalize this likelihood later in this Section**,** where we posit a distribution of expected net values across jobs.

thresholds, z^* and $-z^*$. We refer to these factors as *search-related*. According to Proposition 1, the effects of these two types of factors on the decision to gender-target ads vary between two regions of the parameter space. When $cN/\theta\beta > ln(2)$, the firm's optimal policy is to invite only men when z is positive, and to only invite women when z is negative. Search-related factors have no effect (other than to determine whether one is in this regime). In other words, all jobs are sex-segregated in the sense that no firm searches broadly. We refer to this as the "low frictions" case because it occurs when applications are plentiful; in this case our model essentially specializes to Becker's (1957) model of discrimination in competitive labor markets.[29]

In contrast, untargeted job ads are sometimes the optimal policy in the high-frictions case, i.e. when ($cN/\theta\beta \leq ln(2)$). Here, the firm's optimal policy is to invite women only when z is low, men only when z is high, and to accept applications from both groups for intermediate values of z. Of greater interest, the model has clear predictions for how these thresholds depend on parameter values. Specifically, the four search-related factors (c, N, θ, and β) either move the thresholds $-z^*$ and z^* closer together, making it more likely that firms will engage in gender restrictions of either type, or farther apart, with the opposite effect. In more detail, the model predicts that:

1. Increases in β (i.e. greater idiosyncratic variance of applicant productivity) raise the likelihood that firms will search broadly, inviting both groups to apply.

2. Increases in θ (the job's skill level) also raise the likelihood that firms will search broadly.

3. Increases in c (per-application processing costs) raise the likelihood that firms will search narrowly, inviting only their preferred group for that job to apply.

4. Increases in N (the expected number of applicants) also raise the likelihood that firms will search narrowly.

The intuition for these results is that higher variance in applicant quality raises the option value of searching from a larger pool (i.e. it raises the chance the best candidate will come from

[29] To see this, consider a labor market with many employers, indexed by i, who can hire either men or women. All employers in this market face the same wages, w^M and w^F, but firms' relative tastes for hiring men ($t^M - t^F$) and possibly the expected gender productivity gap, $q^M - q^F$, can vary across firms. Thus, a firm's baseline gender gap in total net value, $v^M - v^F$, can be decomposed into the following components: $v_i^M - v_i^F = (q_i^M - q_i^F) + (t_i^M - t_i^F) - (w^M - w^F)$. If we assume (as Becker does) that men and women are equally productive in these jobs, then in Becker's frictionless world and in our 'low frictions' case, firms where $(t_i^M - t_i^F) > (w^M - w^F)$ will hire only men, and firms where $(t_i^M - t_i^F) > (w^M - w^F)$ will hire only women. The market wage differential ensures that enough firms of either type will exist to employ all the men and women.

the group with the lower expected value), and higher skill levels raise the marginal value to the firm of identifying the best candidate. Thus, both these factors are predicted to reduce the incidence of gender-targeting. In contrast, increases in c and N directly raise the total cost of inviting the 'disfavored' group to apply, thereby reducing the incidence of gender-targeting.

V.B THE MODEL FOR A POPULATION OF JOB ADS

In order to generate predictions for how the share of ads in a population that invite (say) men to apply varies with ads' observable characteristics (including skill levels), we need to specify how those observable characteristics map into the model's parameters, and to introduce a source of unobserved heterogeneity across identical job ads. To that end, we now restrict attention to the high-frictions case, where both targeted and untargeted ads can be optimal. We index ads by i (recall that applicants are indexed by j) and suppose that the employer's net relative valuation of men in the position described in ad i is given by

(6) $$z_i = \mathbf{x_i} \mathbf{b} + v_i$$

where $\mathbf{x_i}$ includes all the observable determinants of firms' preferences towards men (and away from women) for that job, plus a constant term. According to Proposition 1, an ad will then be targeted at men if $v_i > z_i^* - \mathbf{x_i} \mathbf{b}$, targeted at women if $v_i < - z_i^* - \mathbf{x_i} \mathbf{b}$, and will not contain any gender restrictions otherwise. Suppose further that v_i is independently and normally distributed across job ads with mean zero and variance σ_v^2. The likelihood of observing each of the three ad types can then be written:

$$\text{Prob(restrict ad to women)} \equiv P^F = \Phi\left(\frac{-z_i^* - \mathbf{x_i b}}{\sigma_v}\right)$$

(7) $$\text{Prob(no gender restrictions)} \equiv P^C = \Phi\left(\frac{z_i^* - \mathbf{x_i b}}{\sigma_v}\right) - \Phi\left(\frac{-z_i^* - \mathbf{x_i b}}{\sigma_v}\right)$$

$$\text{Prob(restrict ad to men)} \equiv P^M = 1 - \Phi\left(\frac{z_i^* - \mathbf{x_i b}}{\sigma_v}\right),$$

where Φ is the standard normal cdf. If z_i^* and σ_v are constant across observations, (7) describes an ordered probit model.

An important feature of our model, however, is that the ad's observed characteristics – including indicators of its skill requirements as well as observable correlates of application processing costs, expected numbers of applications, and idiosyncratic worker quality—are expected to act on the two thresholds. Specifically, high skill requirements are predicted to move the two thresholds apart by equal amounts; to incorporate this effect, let:

(8) $z_i^* = \exp(-\mathbf{x_i d})$,

which implicitly assumes that any variable that might affect firms' relative valuation of men versus women (z_i) can potentially affect z_i^* as well.

Taken together, (7) and (8) comprise a simple model that can be estimated via maximum likelihood. This model allows all observable characteristics of a job or firm, including the job's skill requirements, to affect the firm's relative preference for men versus women in that job (z), and at the same time to affect the gap between the firm's two thresholds, i.e. level of z^*. Importantly, the effects of any given observable on z^* are identified even if we believe that characteristic also affects a firm's 'tastes towards men' (z).[30] This allows us to test the model's prediction concerning skill requirements without needing to assume, for example, that firms' tastes towards men are independent of the job's skill level.[31]

Unfortunately, while we have estimated the model in (7) and (8) by maximum likelihood with a rich set of covariates, the model is computationally intractable in the presence of large numbers of fixed effects.[32] This is a serious limitation since a key goal of the following Section is to see whether the model's prediction for the effects of skill requirements survives detailed controls for job characteristics that might affect both z and z^*. Fortunately, on-line Appendix 2 shows that under conditions that are approximately satisfied in our data, both **b** and **d** can be estimated (up to a factor of proportionality) by separate OLS regressions. Specifically, **d** can be estimated by an OLS regression of $P^M + P^F$ on **x**, where P^G is a dichotomous indicator for whether the ad states a preference for gender G. In other words, we just regress a dummy for the presence of gender-targeting on the covariates. Under the same conditions, **b** is identified by an OLS regression of $P^M - P^F$ (an outcome which takes the values -1, 0 and 1) on the same covariates.

To sum up, we have developed a model that predicts that increases in a job's skill level, θ, should reduce employers' propensity to gender-target their ads for that job ($-z^*$).[33] Empirically, our vector of observables, **x**, contains three measures of job skill levels (education,

[30] Intuitively, this is because we have data on two distinct outcomes (P^M and P^F). Essentially, **d** is identified by the effects of **x** on their sum ($P^M + P^F$) while **b** is identified by **x**'s effects on their difference ($P^M - P^F$).

[31] It is of course also possible that σ_v varies across ads; for example, we could have $\sigma_v = \exp(\mathbf{x_i g})$. When P^M and P^F are both less than .5 --as is the case at the mean of our data—it is easy to show that an increase in σ_v has the same qualitative effects as an increase in $-z^*$: P^M and P^F both increase. This makes it difficult to separately identify **b** and **g**. Accordingly, we treat σ_v as fixed in this Section, while acknowledging that the estimated effects of parameters on $-z^*$ could also represent effects of the same parameters on σ_v. The consequences of allowing σ_v to depend on the covariates are explored in Section VII.

[32] Results and Stata do-files are available from the authors on request.

[33] Note that the likelihood of gender targeting increases as the lower threshold, $-z^*$, rises towards zero. Accordingly, we use $-z^*$ as shorthand for the propensity to gender-target the job ad at various points in the paper.

experience and the offered wage), and our goal is to estimate their effects on the propensity to gender-target ($-z^*$), controlling for the other elements of **x** *and* allowing for all the components of **x** –including skill levels—to also affect employers' relative preferences for men, z. We have shown that the effects of skill levels on $-z^*$ and z are separately identified, and that under specific conditions we can estimate these effects very simply. All we need to do is run two OLS regressions, one for a (0,1) gender-targeting indicator as a function of **x**, the other for a three-valued indicator (-1, 0 or 1) of the direction of gender-targeting on **x**. We present these estimates in the following Section.

VI. REGRESSION ESTIMATES

This Section presents regression estimates of the determinants of gender-targeting in job ads ($-z^*$), and of firms' preferences in the direction of men (z), using the previous Section's framework. The goals are (a) to see whether the model's prediction of a negative skill-targeting relationship survives detailed controls for other ad characteristics, and (b) to estimate the effects of all the covariates on firms' underlying preferences towards men. Both of these goals are pursued within a framework that allows all the observables to affect both z and $-z^*$.

Accordingly, Table VI presents linear probability model estimates where the dependent variable, $P^M + P^F$ equals one if the ad is gender-targeted (regardless of direction) and zero otherwise. Moving across columns (1) - (3) from left to right, we add increasingly detailed fixed effects, first for occupation, industry and province separately (116 categories), next for these three interacted (22,581 categories), and finally for a full set of occupation*firm interactions (258,751 fixed effects in total).[34] Of special interest are the estimated effects of the indicators of job skill requirements –education and experience—on the propensity to gender-target ads. As predicted by the theory, all the effects are negative. The estimated effects are highly stable across specifications, highly statistically significant, and large in magnitude: for example, jobs requiring university education are 8 to 10 percentage points less likely to be gender targeted than jobs requiring high school or less.[35] Jobs requiring 3-5 years of experience are about 3 percentage points less likely to be gender-targeted than jobs with no explicit experience requirement.

[34] All regressions also control for log firm size, the number of positions advertised, period effects, and whether the job is part-time.

[35] To put this in context, recall from Table II that the mean share of gender-targeted ads for jobs requiring high school or less is 11.3 + 12.0 = 23.3 percent.

It is perhaps worth noting how the estimates of skill requirements in column 3 are identified: essentially, we are comparing two or more ads for the same occupation (say, sales), posted on Zhaopin by the same firm, at two different times during our sampling period, requesting different levels of education or experience. The data show that when a given firm is trying to fill, say, two sales jobs requiring different levels of education, it is significantly less choosy about applicant's gender when filling the position requiring more education. Clearly, the striking correlations between skill and gender-targeting observed in Table II are not an artifact of differences in the mix of occupations and firms that require high versus low educational qualifications.

Column 5 of Table VI probes the association between a job's skill requirements and firms' propensity to gender-discriminate one step further, by asking if jobs with higher posted wages discriminate more or less. Since, as mentioned, only a minority of the ads on the site posted a meaningful wage, the sample size for this exercise is much smaller (172,887 versus 1,051,706 in the full sample); to shed light on the effects of this sample restriction column 4 estimates column 3's model on the subsample with valid wages. According to column 5, controlling for a job's education and experience requirements, jobs that post higher wages are also less likely to be gender-targeted. Education and experience continue to reduce the incidence of gender-targeting, but the education effect is considerably weaker than in column 3.[36]

In the specifications without firm fixed effects, Table VI also shows the effects of firm ownership on the propensity to gender-target in the presence of detailed controls for occupation, industry, province and their interactions. While State-Owned Enterprises engage in about the same amount of gender targeting as private-sector Chinese companies, foreign-owned firms are much less likely to gender-target their job ads than both these groups. This may reflect cultural differences, or the extraterritorial effects of antidiscrimination laws in these firms' home countries.

Shifting our attention from $-z^*$ (the propensity to gender-target) to z (firms' relative preferences for men), Table VII presents results of OLS regressions where the dependent variable is P^M - P^F. The specifications are identical to Table VI. In sharp contrast to Table VI, a job's education requirements have mostly weak and insignificant effects on whether employers prefer men over women for it. The only exception is when controls for offered wages are introduced in column 5: here, men seem to be dispreferred for high-education jobs when the

[36] According to column 4, most of this change in the estimated education effect is not a direct consequence of controlling for wages; instead the effect of education on gender-targeting appears to be smaller in the subsample of jobs with posted wages.

offered wage is held constant. At the same time, however, Chinese employers' preference for men relative to women increases strongly with a job's experience requirements; this tendency is robust to controls for offered wages. Finally, the offered wage has a positive but statistically insignificant effect on firms' preferences towards men.

In sum, Table VII shows that –with the exception of experience-- the effects of a job's skill demands on the direction of firms' gender preferences are typically weaker, and inconsistent across regression specifications and skill measures than the effects of skill demands on gender targeting *per se*. Further, as we show later in this Section, the strong effect of experience is probably not a skill effect, but an artifact of how firms' gender preferences vary with the worker's age. Finally we note that --relative to private-sector Chinese-owned firms-- foreign-owned firms' preferences lean towards women, while SOEs lean towards men. Given that SOEs typically face less product-market competition than privately owned firms, this finding is consistent with an employer taste-based motivation for at least this aspect of advertised discrimination against women.[37]

For additional clues regarding when firms gender-target their ads, and when they prefer men to women, Figure 1 shows occupation fixed effects from regressions identical to column 1 in Tables VI and VII with one exception: education was removed from the list of controls, in order to illustrate its effects in the Figure. Occupations in the Figure are divided into two groups, based on our *a priori* impression of whether they involve a significant amount of customer contact. The six customer-contact occupations, indicated by triangles, are sales, customer service, hospitality/tourism/entertainment ("tourism"), editing/media/film/news ("media"), retail, and "healthcare/beauty/fitness" ("health"). Symbol sizes are proportional to the precision of the estimated fixed effect, and a regression line (estimated with these weights) and 95% confidence band are shown.

Part (a) of Figure 1 shows the estimated fixed effects on the propensity to gender-target ($P^M + P^F$) for the 39 occupations in our data. As predicted by our model, ads for the least-skilled occupational group (labor and domestic service) are almost 30 percentage points more likely to stipulate a preferred gender than in the reference occupation (accounting). The two most positive outliers are administration and tourism. Part (b) of Figure 1 shows the estimated fixed

[37] See Black and Strachan (2001) and Black and Brainerd (2004) for other evidence of the effects of product market competition on gender discrimination. Our findings regarding gender and SOEs are consistent with Zhang, Han, Liu and Zhao (2008), who find that the share of the unadjusted gender wage gap that is not accounted for by observable productivity-related characteristics in China is smaller in market-oriented activities than state-owned ones. For other recent studies of gender differentials in China, see Gustafsson and Li (2000) and Liu, Meng and Zhang (2000).

effects on firms' relative preferences toward men (P^M - P^F). Here, part (a)'s strong negative association is replaced by essentially a zero overall relation with education. Large positive outliers are manual labor, technical occupations, and communications/logistics; large negative outliers (indicating a preference towards women that cannot be accounted for by observable features of the firm, industry, or ad) are tourism, retail, health occupations, and administration. The first three of these are customer-contact occupations; the fourth refers mostly to secretarial jobs. We suspect that customer discrimination plays a role in the former cases, and that managers' tastes might account for the latter. Additional evidence supporting this interpretation is provided in the next Section.

VII. ROBUSTNESS CHECKS, ALTERNATIVE EXPLANATIONS, AND EXTENSIONS

VII.A GENDER AND THE DEMAND FOR OTHER ASCRIPTIVE CHARACTERISTICS

Table III has already suggested a strong interaction between employers' preferences for gender and for other ascriptive worker characteristics, namely workers' age, height and physical attractiveness. In this subsection we explore these interactions further in two main ways. First, we estimate regressions for determinants of firms' advertised requests for other ascriptive worker characteristics –age, beauty, and height— that are specified identically to Tables VI and VII. These regressions show us whether advertised preferences for these other characteristics –which we think of as jointly determined with advertised gender preferences—respond to observables in similar ways to advertised gender preferences. Second, we add controls for these ascriptive characteristics to the regressions for gender discrimination, in order to glean additional descriptive information about how these screens interact.

Table VIII reports OLS regression results with the same specifications as columns (3) and (4) of Tables VI and VII. Because age restrictions, like gender restrictions, are 'bilateral' in our data, columns (1)-(4) of Table VIII adopt a similar strategy of distinguishing firms' propensity to restrict their recruitment to specific worker ages (-z*) from their preferences in the direction of older workers (z). The former effects are estimated by regressions in which the dependent variable equals one if the ad specifies an age range. The latter uses the midpoint of that range, when it is specified, as the dependent variable. Since ads for beauty and height only go in one direction --we encountered no ads requesting unattactive or short people-- we simply present linear probability models for the presence of these restrictions; z and -z* cannot be separately identified here and it is important to take note of this in interpreting the estimates.

Table VIII clearly shows that, like gender-targeting, the incidence of age-targeting declines strongly with jobs' education requirements. For example, even in the most saturated specification (column 2), jobs requiring a university degree are 4.9 percentage points less likely to specify an age range than jobs requiring high school or less. On the other hand, neither experience requirements nor the offered wage have a robust effect on the incidence of age-targeting. Firms' relative valuation of older workers, however, rises with experience requirements and the offered wage; this is perhaps not surprising since experience and wages tend to grow with age. Firms' tendencies to request beauty and height decline with education requirements, experience requirements, and the offered wage, but as noted these estimates cannot distinguish whether this is due to a reduction in the value placed on those attributes (z) or a reduction in the tendency to use them as screens ($-z^*$).[38]

In which occupations do firms look for old, young, beautiful, or tall workers? Online Appendix 4 presents occupation fixed effects for the outcomes in Table VIII, in the same format as Figure 1. It shows that some of the occupations with the highest revealed preferences for female workers in Figure 1 –specifically tourism, retail, and administration—also have high unexplained propensities to request youth, beauty and height. Again, this suggests that customer and supervisor tastes for interacting with attractive women play a significant role in firm's explicit requests for female workers, at least in several key occupations.

In Table IX, we add controls for age, height and beauty requirements to Table VI and VII's regressions for gender discrimination, to see more directly how gender interacts with *ex ante* screening on these other dimensions. Columns 1 and 2 add these controls to the gender targeting regressions in columns 3 and 4 of Table VI; columns 3 and 4 do the same for the "preferences towards men" (P^M - P^F) regressions in Table VII. A first thing to note from Table IX is that the negative skill-targeting relationship for gender discrimination survives the inclusion of controls for screens on other ascriptive characteristics.[39] Thus, the skill-targeting effect is not just a consequence of the tendency for firms to seek tall, good-looking or young women in less-skilled positions. Second, as was noted in Table II's descriptive statistics, screens for all of these ascriptive characteristics are complements with gender screening: gender screens

[38] In specifications without firm fixed effects (not shown; identical to column 2 in Tables VI and VII), we find that foreign-owned firms are much less likely to age-target their ads than Chinese-owned firms and SOEs. Both SOEs and foreign-owned firms are less likely to request beauty, height and a specific age range than private-sector Chinese-owned firms.

[39] In the subsample with posted wages, the education coefficient is still negative but no longer statistically significant in the presence of a wage control.

are more likely to be used when the ad specifies an age range, and when it contains beauty and height requirements.

Third, the positive effect of experience requirements on firms' preferences towards men noted in Table VII falls in magnitude and loses most of its statistical significance when controls for the preferred worker age are added to the regression: firms' tendencies to prefer men when they are looking for experienced workers seem to stem more from a correlation between firms' age and gender preferences than from an experience effect *per se*. Finally, even within firm*occupation cells and controlling for all three measures of job skill requirements (education, experience and the wage), firms tend to strongly favor women over men when they are looking for young, tall and good looking workers. In sum, firms' preferences for a specific package of ascriptive characteristics including youth, beauty and height play an important role in explaining when firms seek to hire women, even within firm*occupation cells. At the same time, the gender-targeting relationship that permeates our data is clearly more than an artifact of firms' 'search for female beauty' in less-skilled jobs.

VII.B ALTERNATIVE EXPLANATIONS OF THE SKILL-TARGETING RELATIONSHIP

One of the main stylized facts documented in this paper is a *negative skill-targeting relationship*: ads for jobs requiring more skill are less likely to stipulate a preferred gender for the employee In Section V we offered a simple model that might account for this relationship; in that model, higher skill demands (θ) reduce the incidence of gender-targeting by increasing the marginal value of identifying the best candidate. Still, it remains possible that the negative skill-targeting relationship in our data is explained by other factors which covary with skill but are not held constant in our statistical analysis. Indeed, some of these factors are suggested by our model itself. The goal of this section is to assess the likely impact of these alternative explanations for the skill-targeting relationship, beginning with four factors -- c, N, σ_ε, and σ_v— that emerge directly from the model.

First, there is some evidence that the cost of assessing the suitability of an individual applicant (c) rises with a job's skill level.[40] But this cannot help explain why gender-targeting falls with skill, since according to Proposition 1 increases in c should make advertised discrimination *more* likely. Intuitively, targeting reduces the number of applications that arrive, which is more useful to the firm when applications are costly to process. On the other hand, it

[40] See Table I in Barron and Bishop (1985). In their employer survey, the total person-hours spent by company personnel recruiting, screening, and interviewing applicants to hire one individual ranged from 7.08 for blue collar workers to 16.99 for managerial personnel.

also seems likely that ads for skilled jobs, on average, attract fewer applicants (N) than ads for less skilled jobs: skilled labor markets may be thinner because skilled workers are more specialized.[41] Since our model predicts that smaller applicant pools lead to a reduction in gender-targeting, variations in market thickness with job skill levels might indeed help account for the negative skill-targeting relationship. A third, and related possibility is that there is, on average, more idiosyncratic variation in the qualifications of applicants to skilled positions than among applicants to unskilled positions (i.e. σ_ε, or β is higher). While this explanation may be hard to assess empirically, it also qualifies as a possible additional source of the skill-targeting relationship in our data.[42]

The fourth and final possible explanation suggested by our model is that the idiosyncratic variance across jobs in the relative ability of men and women to perform them, σ_v, could be lower in a sample of skilled jobs than a sample of unskilled jobs. To see this, refer to equation 7, and suppose now that (like z and z^*) σ_v can also depend on the covariates, for example according to $\sigma_v = \exp(\mathbf{x_i g})$. Suppose also that a minority of ads are targeted at either gender, i.e. $P^M < .5$ and $P^F < .5$ (therefore $-z_i^* - \mathbf{x_i b} < 0$ and $-z_i^* + \mathbf{x_i b} > 0$) at some initial set of parameter and data values, which is clearly the case at the mean of our data. In this case, a small decrease in σ_v has the same qualitative effect as an increase in z_i^*: reducing the share of ads targeted at men and the share targeted at women. Thus it is possible that what appears to be a direct 'price' effect of higher skill requirements is instead a consequence of the fact that the dispersion, across jobs, of jobs' 'gender-suitability' falls with skill levels.

Empirically, why might σ_v might fall with skill? One possibility relates to the fact that skilled jobs do not typically require manual labor. Thus, if there are larger gender gaps in humans' abilities to perform physical tasks than mental tasks, there will be greater cross-job variance in the gender-suitability of unskilled than skilled jobs. While this seems plausible, we note that we can test for this possibility directly, by excluding all occupations that are likely to involve any physical labor from the analysis. Since such jobs are a small fraction of the Zhaopin sample, this has essentially no effect on our main results, including the negative skill-targeting

[41] Baron and Bishop's evidence on this point is not clear, however. The mean number of applicants per job in their survey was essentially same for blue collar versus managerial jobs (7.98 versus 8.08). The highest number of applicants per job was 10.82, for clerical jobs. While we do not have information on the number of applicants per job in our data, we do note that, because Zhaopin.com tends to serve a skilled workforce, its online markets for highly skilled workers –measured by the number of ads—are actually thicker than for less-skilled workers.

[42] A direct test of this explanation would require a scalar measure of the variance of all worker qualifications that are visible to the employer in the hiring process, which could be challenging to construct.

relationship.[43] We conclude that gender differences in the ability to perform physical tasks do not explain the skill-targeting relationship in our data.

Recalling that the job-specific relative net value of men in our model, v_i , includes tastes as well as productivities, a second reason why σ_v might fall with job skill requirements could be that hiring agents and/or the applicant's prospective co-workers simply 'care less' about the applicant's gender at higher skill levels.[44] While this idea has some appeal --for example, educators sometime argue that a key goal of education is to make people comfortable with a wider range of ideas, situations and people— it is not easy to reconcile with at least one pattern in our data, namely the negative effect of experience on gender-targeting. Specifically, if jobs requiring more experience tend to have older hiring agents or co-workers, this 'tastes' hypothesis would require these older persons to have less strict notions of what is a 'proper' gender role for a job than young people. This seems unlikely, given the decline in gender-stereotyping across cohorts in most societies. Also, the fact that the negative skill-targeting relationship holds for all three measures of skill in our data, and for advertised discrimination on several margins (gender, age, height and beauty) strikes us as more supportive of a direct effect of skill than a taste-based explanation. Still, we acknowledge that this pattern of tastes might also contribute to the negative skill-targeting relationship in our data.

A final piece of evidence on the possibility that lower levels of σ_v in skilled jobs explain the skill-targeting relationship is based on the fact that with the right variation in **x**, the effects of observables (including our three skill indicators) on z, z^* and on σ_v are *all* separately identified by a simple sign test, even in the absence of exclusion restrictions. To see this, return to equation (7) and suppose that, for some subsample of our data, more than half the ads are targeted at one of the genders; for example $P^M < .5$ and $P^F > .5$. Now, according to (7) a small decrease in σ_v should still reduce P^M, but should *raise* P^F. More to the point, in a model with a constant σ_v, the effects of any given covariate on P^M and P^F should have the same sign regardless of whether P^M or P^F exceeds .5. If, instead, the effects of a covariate on, say, P^F, switch sign depending on whether P^F is above or below .5, that covariate must have an effect on σ_v.[45]

[43] The only occupations in our data that seem likely to involve any physical labor are construction, manufacturing and "manual labor"; together they constitute less than 11 percent of our sample. In contrast, sales, IT, marketing , accounting, and administration together account for almost half of the ads.

[44] Note that this hypothesis requires a specific pattern of tastes; it is not sufficient, for example, for more-educated hiring agents (or co-workers) to be less biased against women. Instead, what is needed is that as we move down the skill ladder, *some* agents' tastes need to become more intense in favor of women, while other agents' tastes become more intense in favor of men.

[45] Note that this property holds not just for a normal distribution of v, but for any distribution with a median of zero.

To explore this distinction, Table X presents some simple tabulations for two subsets of ads: one in which men are highly favored, the other in which women are highly favored. "Highly male" jobs are defined as in 'technical' occupations (the occupation with the strongest preference towards men in Figure 1b), where the maximum requested age is 25 years or higher. Highly female jobs are defined as jobs that request beauty and stipulate a maximum age under 25. When the job's education requirement was high school or less, more than half of the ads for these types of jobs explicitly requested men and women respectively, which allows us to ask what happens to firms' advertised preferences as we increase the desired skill level, starting from a situation where more than half the ads are targeted to men (women). Given the above discussion, if the effects of higher skill requirements in these highly gendered jobs on P^M and P^F qualitatively mirror their effects in the sample as a whole (which was to reduce both P^M and P^F), this is consistent with our baseline model, where skill requirements operate only through z and z^*. If, instead, the sign patterns are different in these highly-gendered jobs, it appears that skill requirements must also affect σ_v.

The results of Table X clearly favor the "price effect" hypothesis over the preference-heterogeneity explanation of the skill-targeting relationship: as skill requirements in these jobs rise from high school to university, the share of highly male jobs that request men falls from 64 to 51 percent, and the share of highly female jobs that request women declines from 82 to 52 percent. Both these differences are highly statistically significant. While these highly gendered jobs are clearly special, they provide an additional piece of evidence suggesting that a lower cross-job variance of gender-suitability in skilled versus unskilled jobs is not the main source of the negative skill-targeting relationship in our data.

Having considered the possible confounding effects of all the other main parameters in our theoretical model (c, N, σ_ε, and σ_v), we conclude this subsection by considering a final alternative explanation that is outside the model: What if highly skilled jobs are not gender-targeted simply because it is already common knowledge that those jobs are reserved for men? We think this is unlikely given the high degree of overlap between the male and female wage distributions in our data. For example, according to the Census statistics in Table A.1, while women are certainly under-represented among high-wage workers, 25 percent of workers earning over 8000 yuan/month [the top 0.3 percent of the wage distribution] are female. And 31.5 percent of the gendered Zhaopin job ads offering more than 8000 yuan/month request women. In short, women are sufficiently scattered across China's job distribution that gendered ads are likely to be informative about what firms want in most jobs, even at high skill levels.

In sum, our discussion in this Section suggests that at least two unmeasured factors that could be associated with high skill requirements --thinner labor markets and higher idiosyncratic variance of applicant quality-- might (in addition to the direct effect of skill itself) help explain the negative skill-targeting relationship in our data. Two factors that almost certainly do not play a role are skill-related differences in the cost of assessing applicant quality, and a lesser importance of physical strength in skilled jobs. Evidence on a final factor –the notion that hiring agents and co-workers for skilled workers might simply 'care less' whether the successful candidate is male or female—is more mixed, but in our reading the balance of the evidence points away from this story. That said, additional research to distinguish these various mechanisms would be of considerable interest.

VII.C EXTENSIONS

This subsection explores two extensions to our analysis—one theoretical and the other empirical. The theoretical extension adds a type of directed search by workers to our model.[46] Specifically, note that our baseline model assumes that the number of (say) male applications received by a firm that specifically invites men to apply is the same as the number of male applications generated by an ungendered ad (both equal N). But what if φN men respond to an ad directed specifically at men, where $\varphi \in [1,2]$ measures the supply response of a group when an ad is targeted at them? Workers might respond favorably to targeted ads for a number of reasons; for example they might believe they have a better chance of being hired if they are the type of worker the firm is looking for. Or, workers might respond positively to the knowledge that they will be working in an all-male or all-female environment. In other words, this extension allows taste-based discrimination on the part of prospective *applicants* to affect firms' decisions on whether to target their job ads.[47]

In on-line Appendix 3, we solve this expanded model for the optimal z^*, which now depends on c, N, θ, β and φ. As expected, a greater expected supply response from targeting (i.e. a higher φ) raises the likelihood of gender targeting. In addition, the expanded model's qualitative predictions for the effects of c, N, θ, β are all the same as our baseline model; this includes the key prediction that higher skill requirements should reduce the incidence of targeting. Aside from greater generality, the expanded model has two attractive features which

[46] Models of directed worker search are quite common (e.g. Moen 1997), and sometimes incorporate discrimination by firms (Lang, Manove and Dickens 2005). To our knowledge, however, our paper is one of the first attempts to formalize the idea of directed *recruiting* by firms.

[47] As a result, it also opens up another possible explanation for the negative skill-targeting relationship: Firms may gender-target less in skilled jobs if the supply elasticity of skilled *applicants* to gender-targeted ads is lower.

enhance its realism. One of these is the fact that, in contrast to the baseline model, positive application processing costs are no longer essential to the explanation of targeted ads: such ads can be optimal even when processing costs are zero because they attract more applicants of the type the firm likes best. While available evidence suggests that such costs are both real and important, it is an open question whether they are important enough to explain all the ad targeting that occurs.

The second attractive feature of the extended model is that its predictions are less sensitive to small changes in either (a) the level of stigma attached to posting a targeted ad, or (b) what might be called the firm's *pre-screening* technology. To see this, imagine that firms face some stigma if they post a gendered job application, and that it is relatively easy for firms to 'pre-screen' resumes on the basis of easily-observed demographic indicators and simply discard applications belonging to the group with the lower expected value. (Pre-screening only reveals basic demographics and is a distinct process from actually assessing the applicant's ε, which still costs c per applicant). Then even for fairly small levels of social stigma, we should observe no targeted ads in the baseline model: firms will substitute internal pre-screening for ads that ask workers to screen themselves. This is not the case, however, in the expanded model: if targeted ads induce a large enough supply response of the worker types firms prefer, firms may continue to target even in the face of substantial stigma associated with posting a targeted ad. While all the available evidence suggests there is no stigma attached to posting a gendered job at in China, this reduced sensitivity to small changes in stigma is still an attractive feature of the expanded model.

The other extension described in this subsection is an empirical exercise. Specifically, we ask what our Zhaopin results –which are based on a highly specialized segment of the Chinese labor market—imply for the likely incidence of explicit gender discrimination in the Chinese labor market as a whole. To this end, Table XI returns to the high-income provinces sample introduced in Table A.1, and uses regressions similar to those in Tables VI and VII to predict the share of jobs targeted at men, and the share targeted at women. We generate predictions for both the Zhaopin sample, and a sample with the average characteristics of the employed population in those provinces. Since there are essentially no public sector ads nor any ads requesting workers over 50 on Zhaopin, the estimates apply to private sector workers under 50 only.

Row 1 of Table XI uses only the three education categories (high school or less, some college, and university degree), which are available for all observations in both data sets, to predict gender-targeting. Overall, 8.8 percent of Zhaopin ads in the high-income provinces were gender-targeted (a little less than the 10.5 percent in the full Zhaopin sample), with 3.8 percent preferring women and 5.0 percent preferring men. According to Row 1, these shares would rise to 8.8 and 10.3 percent respectively for a sample of job ads that had the same education distribution as the entire employed population in those provinces. Thus, the rate of gender targeting would more than double, from 8.8 to 19.1 percent of all ads. Row 2, which adds controls for industry, occupation, and firm type (private versus SOE) yields broadly similar results, with gender-targeting now rising to 20.8 percent.

Rows 3 and 4 of Table XI alternatively add controls for the 5 wage quintiles or the 3 age groups under 50 in Table II. Recall that offered wages and desired ages are each only present in a minority of Zhaopin ads; thus these two rows are estimated on considerably smaller samples than Rows 1 and 2. Still, the results are suggestive: a population of ads that reflects the education and wage distribution of the employed population in China's high-income provinces would be 21.7 percent gender-targeted, while a population reflecting the education and age distribution would be 40.5 percent gender-targeted. In sum, our regression estimates, combined with information on the mix of jobs in a broader sample of urban Chinese jobs, suggest that explictly gender-targeted job ads may be much more common in the broader Chinese labor market than in the Zhaopin sample we study here.[48]

VIII. CONCLUSIONS

In a legal environment where firms are allowed to engage in explicit gender discrimination when advertising their jobs, when will they choose to do so? Our data show that firms in such an environment will use the option to discriminate much more often when hiring for positions requiring lower levels of skill, whether skill is measured by education requirements, experience requirements, or the offered wage. In this paper we propose a simple model of employer search from two populations that might help explain this pattern. The intuition is straightforward: as skill requirements rise, it becomes increasingly important for firms to identify the best individual candidate for the job. Our model also suggests that other factors which co-vary with skill --such as thinner labor markets or greater idiosyncratic variance in applicant productivity at high skill levels-- may also help account for this skill-targeting effect.

[48] To assess the plausibility of these predictions, we compared them with aggregate quarterly vacancy data from the public employment services of a sample of large cities, provided by China's Ministry of Human Resources and Social Security (MOHRSS) and available online from CEIC Data (2012). In the last quarter of 2011, 36.5 percent of these vacancies were described as 'male', 32.0% were female, and the remaining 31.5%)were not gendered. This large share of gendered job ads in a more representative sample is perhaps not surprising since Table XI's predictions can only control for a small share of the observable ad characteristics.

Another possibility is that gender simply 'matters less' to hiring agents, to incumbent employees or to prospective job applicants in skilled jobs. Sorting out these channels seems a useful avenue for further research.

Regardless of the exact mechanisms behind the skill-targeting effect, our results suggest that skill-upgrading may have a powerful, negative effect on the amount of explicit discrimination that is observed in labor markets. In a sense, this result complements Becker's (1957) notion that enhanced product market competition should also reduce discrimination. Interestingly, however, our model has a prediction for the effects of *labor* market competition that may be surprising: it is not thick but *thin* labor markets (such as those for specialized, skilled workers) that are predicted to reduce firms' reliance on coarse demographic screens. According to the model, discriminatory firms will be more likely to broaden their recruiting efforts to 'disfavored' groups when the expected number of applicants is low.

Abstracting from firms' decisions whether to gender-target their job ads, we have also provided a number of results on the direction of firms' underlying gender preferences. In contrast to its effects on the frequency of gender-targeting, we find that a job's skill level does not have a robust effect on firms' preferences for men relative to women. That said, the direction of firms' gender preferences is strongly related to firms' visions for the appropriate age, height and beauty of the applicant. Firms tend to prefer women when they are looking for young, tall and attractive workers, especially in a small number of customer contact occupations and in the market for secretarial services. On the other hand, firms tend to prefer men when they are searching for older workers; it would be interesting to know more about why this occurs.

Finally, we find that close to two thirds of the variation in estimated gender preferences takes place within firms, and about a third occurs within occupation*firm cells. While inconsistent with a number of simple models, this pronounced within-firm heterogeneity is consistent with a models where firms (or co-workers) value gender homogeneity in detailed job categories *per se*, and with models where beliefs about gender-appropriate roles are more widely held, but apply to much finer occupational categories than we can measure in our data. Additional research to distinguish among these scenarios would be of great interest as well.

University of California, Santa Barbara; National Bureau of Economic Research, and Institute for the Study of Labor (IZA)

Xiamen University and Institute for the Study of Labor (IZA)

Appendix

As noted, our sample consists of all job ads which appeared on Zhaopin.com between during four observation periods: May 19 2008 - June 22 2008, January 19 2009 - February 22 2009, May 18 2009 - June 21 2009, and January 18 2010 - February 21 2010. To collect the data, we developed a web crawler that automatically searches for job ads that were posted on Zhaopin on a given day. On each day of our sampling period, the program started at 11:30 pm for consistency. On the first day of data collection, all ads that were listed as posted on that day were kept. On subsequent days, all ads posted that day are compared with the master list of previously-posted jobs; since many jobs are just renewals that are re-posted we do not download these re-posted jobs but maintain a count of the number of renewals that occurred during this time period.[49] As a result, our data have information on every job that was posted or renewed during this time period, linked to information about the firm posting the job. All of our analysis is restricted to the sample of jobs for which we have matching firm information. The matching rate varied somewhat across specifications but was about 80%.

Age, gender and other job requirements were extracted from each job's html file. For example, in the case of gender, we look for "nue"(female) and "nan"(male) characters in the job description section of the file. We then constructed a match table summarizing about 1468 ways for a job ad to mention "nue"(female) and "nan"(male). After that, we use a program and this match table to derive the gender discrimination variable automatically. We consider our table quite exhaustive. In addition, we visually checked all the ads that mentioned gender in a way that did not match these tables. Only about 100 ads fell into this category. For age variables, we search for "sui" (year of age); our approach could therefore miss jobs that ask for age only using numbers "25-35". Therefore, our indicators should be interpreted as very explicit requirements for gender, age and other characteristics.

Occupation and industry categories are those supplied by Zhaopin.com (firms choose from a list on the website). Note that Zhaopin's occupational categories are not mutually exclusive: firms are allowed to check up to three categories. Since the share of vacancies corresponding to each occupation is generally unspecified when multiple occupations are listed, we restricted our sample to ads for a single occupation; this reduces our sample by about 20 percent. Our results were, however, very similar when we included all ads and classified them according to the first occupation listed, or when we allocated ads fractionally, and equally, across

[49] Re-postting an existing ad entails only a small marginal financial cost, but does require action on the employer's part.

all the occupations they listed. Firms can also list multiple industries; we resolved this by simply choosing the industry that was listed first, and similar robustness checks showed that this has little effect as well. Finally, the analysis sample for the current paper also excludes the approximately 20 percent of ads that did not specify what education level was required (inspection of these ads showed that these were not necessarily unskilled jobs).

As discussed in the paper, there are a number of reasons to expect our Zhaopin sample to be unrepresentative of a broader sample of vacancies or occupied jobs.[50] To assess the representativeness of our sample of Zhaopin ads relative to a population of occupied jobs, columns 1 and 2 of Table A.1 compare the observable characteristics of Zhaopin ads to the employed urban Census population in China's eight highest-income provinces.[51] Together, these eight provinces account for 78 percent of all the Zhaopin ads and for 41 percent of national employment in 2009. They have an employment-weighted annual GDP per capita of RMB 46,930, compared with 31,919 for the nation as a whole.[52] A key distinction from Table I is that, for characteristics such as gender, age, and wage which are often unspecified in job ads, the Zhaopin means in Table A.1 refer to the subset of ads where firms actually specified a preference. For example, according to column 2, of the job ads in high-income provinces that expressed a gender preference, 56.5 percent requested men. The goal is simply to give a rough picture of the types of workers that firms are seeking in the Zhaopin data, and how these compare to urban working people in these eight provinces.

As noted in the paper, compared to employed persons in the same set of provinces, Zhaopin ads serve young, well-educated, well-paid private-sector workers. And while the Zhaopin industry and occupation categories do not correspond neatly to the available Census categories, a few conclusions about industry and occupation mix can be drawn. First, relative to the general working population, the IT/ communication and R&D/consulting industries are highly overrepresented on Zhaopin, together accounting for over 33 percent of Zhaopin ads compared with under 4 percent of total employment. Three industry categories – construction/transporation, trade/hospitality/entertainment, and finance/insurance/real estate— are about equally represented in Zhaopin compared to total employment, while the remaining industries are underrepresented in Zhaopin ads. Notably, even though manufacturing is

[50] One of these reasons is the length-biased sampling associated with constructing a sample of unfilled vacancies. To assess the impact of this biase, we replicated our analysis for a subsample of ads from near the end of each data collection period that consists, almost certainly, of newly posted ads. The results were very similar.

[51] 2005 Census data are from the 1% National Population Sample Survey, conducted by China's National Bureau of Statistics. The microdata were kindly supplied by Loren Brandt of the University of Toronto. Some details on methodology are available at: http://www.stats.gov.cn/eNgliSH/newsandcomingevents/t20060322_402312182.htm

[52] GDP and employment data are from the National Bureau of Statistics.

underrepresented on Zhaopin, almost one quarter of Zhaopin ads are from manufacturing firms. Less detailed matches are possible on occupation, but professional and technical workers are clearly overrerpresented on Zhaopin with 60 percent of the ads, compared with 17 percent of employment.

Columns 3 and 4 of Table A.1 compare the gender mix of Zhaopin ads to the gender mix of employment; in the same spirit as columns 1 and 2, gender mix in the Zhaopin data is represented by the share of gendered ads that are female. According to columns 3 and 4, for workers under 30, the share of gendered ads requesting women closely mirrors women's share of the workforce, at a little over 50 percent. Further, both the share of women in the workforce and the share of gendered Zhaopin ads that are female decline with the worker's age.[53] Turning to education, columns 3 and 4 of Table A.1 show that the female workforce in these cities is a little less likely to have a university degree than the male workforce, and this difference is mirrored, somewhat more strongly, in the Zhaopin job ads. There is also a broad correspondence between industries that tend to employ women (trade/hospitality/entertainment and health/education/ welfare) and industries where job ads are targeted at women. The same is true for occupation, with sales and service occupations being highly female in both Zhaopin and the Census and production/construction being highly male. In both Zhaopin and the Census, the private sector is more 'female' than state-owned enterprises; finally women are underrepresented in high-wage jobs in both Census employment and in Zhaopin ads. Overall, Zhaopin ads tend to request female workers in the types of jobs where, according to the Census, women are already employed.

[53] The considerably more precipitous decline in the Zhaopin ads seems to reflect an interaction between firms' underlying age and gender preferences, which we document in the paper.

Table A.1: Descriptive Statistics, Zhaopin.com Ads versus 2005 Census Employed Population, High-Income Provinces

	Share in Category		Share Female within Category	
	(1)	(2)	(3)	(4)
	Census	Zhaopin	Census	Zhaopin
Gender				
Male	.548	.565		
Age				
30 or below	.404	.521	.511	.555
31-40	.311	.450	.466	.226
41-50	.214	.029	.407	.113
51-60	.071	.001	.197	
Education				
High school or below	.766	.114	.455	.464
College	.135	.431	.465	.508
University	.100	.455	.414	.283
Industry:				
Primary, Manufacturing and Utility	.453	.267	.465	.381
Construction and Transportation	.118	.135	.202	.300
IT and Communication	.016	.185	.405	.494
Trade, Hospitality and Entertainment	.165	.175	.558	.590
Finance, Insurance and Real Estate	.063	.052	.454	.506
R&D and Consulting	.023	.153	.376	.389
Health, Education and Welfare	.102	.033	.617	.687
Public Sector	.060	.000	.318	
Occupation:				
Senior Management	.025	.021	.265	.425
Professional and Technical	.170	.600	.562	.435
Sales and Service	.230	.261	.556	.583
Production and Construction	.443	.119	.399	.121
Public Servants	.132	.000	.342	
Firm ownership:				
Private Sector	.589	.930	.482	.444
SOEs and collectives	.271	.070	.384	.325
Public Administration	.140	.000	.456	
Wage distribution:				
1500 or below	.778	.145	.479	.503
1501-3000	.176	.164	.369	.559
3001-4000	.021	.214	.356	.509
4001-8000	.022	.244	.296	.407
8001 or above	.003	.126	.251	.315

Note: Zhaopin.com distributions refer to ads that stated a preference for the attribute (e.g. age, gender, wage) only. Sample is restricted to urban workers in the eight provinces with the most Zhaopin ads, which are also top 8 provinces in GDP per capita: Beijing, Shanghai, Guangdong, Jiangsu, Shandong, Tianjing, Zhejiang and Liaoning. The share female in Zhaopin ads for workers over 50 is not reported due to the extremely low sample size.

References

Akerlof, George A, and Rachel E. Kranton, "Identity, Supervision, and Work Groups," *American Economic Review*, 98 (2008), 212-17.

Barron, John M., and John Bishop, "Extensive Search, Intensive Search, and Hiring Costs: New Evidence on Employer Hiring Activity," *Economic Inquiry*, 23 (1985) 363-382.

Becker, Gary S., *The Economics of Discrimination*. 2nd ed. (Chicago: University of Chicago Press, [1957] 1971).

Bertrand, Marianne and Sendhil Mullainathan, "Are Emily And Greg More Employable Than Lakisha And Jamal? A Field Experiment On Labor Market Discrimination," *American Economic Review*, 94 (2004), 991-1013.

Bertrand, Marianne, "New Perspectives on Gender," in Orley Ashenfelter and David Card, eds. *Handbook of Labor Economics*, volume 4B. (Amsterdam: North-Holland , 2010) 1543-1590.

Bielby, W. T., and J. N. Baron, "A woman's place is with other women: Sex segregation within organizations,", in B. F. Reskin, ed. *Sex segregation in the workplace: Trends, explanations, remedies*. (Washington, D. C.: National Academy Press, 1984).

Black, Sandra E. and Elizabeth Brainerd, "Importing Equality? The Impact of Globalization on Gender Discrimination," *Industrial and Labor Relations Review*, 57 (2004), 540-59.

Black, S. and P. Strachan, "The Division of Spoils: Rent-Sharing and Discrimination in a Regulated Industry,". *American Economic Review* 91 (2001), 814-831.

Blau, Francine, *Equal Pay in the Office*, Lexington, MA: D. C. Heath, 1977.

CEIC Data, *China Premium Database* (http://www.ceicdata.com/China.html) (2012)

Charles, Kerwin Kofi, and Jonathan Guryan, "Studying Discrimination: Fundamental Challenges and Recent Progress," *Annual Review of Economics*, 3, (2011), 479-511.

Darity, William A. Jr., and Patrick L. Mason, "Evidence on Discrimination in Employment: Codes of Color, Codes of Gender," *Journal of Economic Perspectives* 12 (1998),63-90.

Fryer, Roland G. Jr., "Racial Inequality in the 21st Century: The Declining Significance of Discrimination," in *Handbook of Labor Economics* volume 4B, Orley Ashenfelter and David Card, eds. (Amsterdam: North-Holland, 2010), 855-971.

Goldin, Claudia, *Understanding the Gender Gap: An Economic History of American Women*, (New York: Oxford University Press, 1990).

Goldin, Claudia, "The Rising (and then Declining) Significance of Gender", in *The declining significance of gender?* Francine D. Blau, Mary C. Brinton, and David B. Grusky, eds. (New York: Russell Sage Foundation, 2006).

Gustafsson, Bjorn and Shi Li, "Economic transformation and the gender earnings gap in urban China," *Journal of Population Economics* 13 (2000), 305-329.

Groshen, Erica, "Sources of Intra-industry Wage Dispersion: How Much Do Employers Matter?," *Quarterly Journal of Economics* 106 (1991), 869-84.

iResearch Inc., *iResearch China Online Recruitment Research Report 2010-2011* (http://www.iresearch.com.cn/Report/1583.html), (2011).

Kuhn, Peter and Kailing Shen, "Employers' Preferences for Gender, Age, Height and Beauty: Direct Evidence," NBER working paper no. 15564, 2009

Lang¸ Kevin, Michael Manove and William T. Dickens, "Racial Discrimination in Labor Markets with Posted Wage Offers," *American Economic Review*, 95 (2005), 1327-40.

Lawler, John J. and Johngseok Bae, " Overt Employment Discrimination by Multinational Firms: Cultural and Economic Influences in a Developing Country *Industrial Relations* 37 (1998), 126-152.

Liu, Pak-Wai, Xin Meng and Junsen Zhang, "Sectoral gender wage differentials and discrimination in the transitional Chinese economy," *Journal of Populatio.n Economics* 13 (2000), 331-352.

Moen, Espen R., "Competitive Search Equilibrium," *Journal of Political Economy*, 105 (1997), 385-411.

Ministry of Human Resources and Social Security of the Peoples' Republic of China, *Urban Labor Market Survey* (in Chinese). (http://www.mohrss.gov.cn/page.do?pa=402880202405002801240882b84702d7&guid=1394a37cf8934ba0b8d31a557008f1c4&og=8a81f0842d0d556d012d11155d8c0043) (2012)

Neumark, David, "Sex Discrimination in Restaurant Hiring: An Audit Study," *Quarterly Journal of Economics*, 111 (1996), 915-42.

Rogerson, Richard, Robert Shimer and Randall Wright, "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature*, 43 (2005), 959-988.

Zhang, Junsen, Jun Han, Pak-Wai Liu and Yaohui Zhao, "Trends in the Gender Earnings Gap in Urban China, 1988-2004," *Industrial and Labor Relations Review*, 61 (2008)

TABLE I: Sample means, Zhaopin Job Ads

A. AD CHARACTERISTICS	
Gender requirements	
No gender preference	.895
Prefer male?	.055
Prefer female?	.050
Education requirements	
High school or less	.129
Some postsecondary	.457
University degree	.414
Experience requirements	
None or less than one year	.205
1-3 years	.399
3-5 years	.237
More than 5 years	.158
Age requirements	
No age restrictions	.757
Ad specifies a minimum age	.169
Ad specifies a maximum age	.202
Mean age requested[a]	30.59
Wages	
Wage not specified	.836
Mean Wage, when advertised[b]	4,279
Number of positions advertised	
Unspecified	.481
Mean number, when specified	1.692
Other job characteristics:	
Job is Part Time	.009
Job requires beauty (*xingxiang*)	.077
Job has a height requirement	.026
B. FIRM CHARACTERISTICS	
Firm size (mean number of workers)	1,565
Firm ownership type:[c]	
Private, Domestic	.564
Foreign	.362
State-Owned Enterprise	.074
Number of Ads	1,051,706

Data refer to ads downloaded from Zhaopin.com during four observation periods: May 19 2008 - June 22 2008, January 19 2009 - February 22 2009, May 18 2009 - June 21 2009, and January 18 2010 - February 21 2010.

[a] Midpoint of the maximum and minimum age when both are specified

[b] Wages are measured in RMB per month. Zhaopin prompts firms to specify a minimum and maximum wage. Ads are categorized as not specifying a wage when the (a) either the maximum or minimum is blank, or (b) the maximum and minimum are more than 20,000 RMB/month apart .

[c] "Private, Domestic" includes privately held companies, publicly-traded companies and reformed State-Owned Enterprises. "Foreign" includes Foreign Direct Investment, joint ventures, plus a small number of representative offices. Firm characteristics are reported by firms on each firm's Zhaopin page.

Table II: Share of Job Ads Expressing a Gender Preference, by Ad Characteristics

	Share of Job Ads		
	Requesting Women	With no Gender Preference	Requesting Men?
A. JOB SKILL INDICATORS:			
Education Requirements			
High school or less	.113	.766	.120
Some college	.059	.892	.049
University	.021	.938	.042
Experience requirements			
None or less than one year	.087	.860	.053
1-3 years	.060	.889	.051
3-5 years	.025	.920	.055
More than 5 years	.015	.917	.068
Wages			
Wage not specified	.046	.900	.054
Wage is specified	.072	.870	.058
Wage, if specified:[a]			
under 1500	.167	.734	.099
1500-2999	.114	.808	.078
3000-3999	.053	.899	.048
4000-7999	.034	.918	.048
8000+	.038	.929	.033
B. OTHER ASCRIPTIVE JOB REQUIREMENTS:			
Age requirements			
No age restrictions	.029	.944	.027
Ad specifies a minimum age	.116	.746	.138
Ad specifies a maximum age	.124	.724	.151
Maximum and minimum specified	.131	.721	.149
Mean age, when specified:[b]			
Under 25	.332	.607	.060
25-29	.176	.686	.138
30-34	.083	.756	.161
35+	.041	.771	.188
Job requires beauty (*xingxiang*)?			
No	.034	.909	.056
Yes	.239	.723	.038
Job has a height requirement?			
No	.040	.907	.053
Yes	.419	.438	.142
C. FIRM CHARACTERISTICS			
Firm size (number of workers)			
Under 25	.056	.902	.043
25-99	.059	.888	.053
100-999	.047	.895	.058
1000+	.040	.905	.055
Firm ownership type:[c]			
Private, Domestic	.061	.874	.065
Foreign	.035	.928	.037
State-Owned Enterprise	.042	.893	.065

Notes to Table II:

Data refer to ads downloaded from Zhaopin.com during four observation periods: May 19 2008 - June 22 2008, January 19 2009 - February 22 2009, May 18 2009 - June 21 2009, and January 18 2010 - February 21 2010.
[a] Wages are measured in RMB per month. Zhaopin prompts firms to specify a minimum and maximum wage. Ads are categorized as not specifying a wage when the (a) either the maximum or minimum is blank, or (b) the maximum and minimum are more than 20,000 RMB/month apart.
[b] Midpoint of the maximum and minimum age when both are specified.
[c] "Private, Domestic" includes privately held companies, publicly-traded companies and reformed State-Owned Enterprises. "Foreign" includes Foreign Direct Investment, joint ventures, plus a small number of representative offices. Firm characteristics are reported by firms on each firm's Zhaopin page.

Table III: Share of Job Ads Expressing a Gender Preference, for selected subsamples of ads.

	Share of Job Ads:		
Ad characteristic:	Requesting women	With no Gender Preference	Requesting men
Ad requests beauty	.239	.723	.038
Ad requests both beauty and height	.559	.383	.059
Ad requests beauty, height and maximum age under 25	.871	.103	.026

Table IV: Advertised Gender Preferences, by Firm

	Share of Firms Specifying a Preference For:				
Total Number of Ads Placed by the Firm:	Any Gender	Men	Women	Both Genders	**N**
1	.165	.056	.109	.000	12,688
2-10	.326	.155	.224	.053	40,932
11-50	.548	.345	.392	.189	16,244
51 and over	.707	.535	.559	.387	3,778
All Firms	.367	.200	.258	.091	73,642

Table V: Variance Decomposition

	Dependent Variable		
	(1)	**(2)**	**(3)**
Share of variance explained by:	**Ad Requests Men (P^M)**	**Ad Requests Women (P^F)**	**Ad is Gender-Targeted ($P^M + P^F$)**
1. Occupation	.017	.031	.017
2. Firm	.291	.279	.322
3. Joint Occupation and Firm	.026	.031	.033
4. Total, Occupation plus Firm	.334	.341	.373
5. Occupation*Firm Interactions	.304	.327	.294
6. Total between Job Cells	.639	.668	.667
7. Within Job Cells	.361	.332	.333
8. TOTAL	1.000	1.000	1.000

The variance decomposition follows Groshen (1991). Row 4 is the R^2 in a regression of the dependent variable on a full set of occupation and firm fixed effects. Row 1 refers to the (minimum) contribution of occupation fixed effects to that regression; row 2 to the minimum contribution of firm fixed effects to that regression. Row 3 is the explained variance in that regression that cannot be unambiguously attributed to occupation or firm effects. Row 6 is the R^2 in a regression of the dependent variable on a full set of occupation*firm interactions ('job cells'). Row 5 is the difference between rows 4 and 6, and measures the extent to which different firms 'gender' their occupations differently. Row 7 equals one minus row 6, and measures the the residual variance within job cells, i.e. the extent to which the same firm advertising for a particular occupation sometimes requests men and sometimes does not.

Table VI: Effects of Jobs' Skill Demands on the Probability an Ad is Gender-Targeted ($P^M + P^F$), OLS Estimates:

`	(1)	(2)	(3)	(4)	(5)
Education Requirement:					
Some Postsecondary	-.0744**	-.0681**	-.0599**	-.0291**	-.0209*
	(.0050)	(.0048)	(.0049)	(.0092)	(.0088)
University	-.1006**	-.0946**	-.0806**	-.0372**	-.0202
	(.0057)	(.0056)	(.0057)	(.0113)	(.0106)
Experience Requirement:					
1-3 years	-.0156**	-.0177**	-.0219**	-.0290**	-.0255**
	(.0023)	(.0024)	(.0025)	(.0059)	(.0058)
3-5 years	-.0323**	-.0323**	-.0324**	-.0516**	-.0348**
	(.0027)	(.0027)	(.0032)	(.0070)	(.0069)
More than 5 years	-.0285**	-.0288**	-.0343**	-.0546**	-.0250*
	(.0035)	(.0033)	(.0038)	(.0092)	(.0100)
Log (offered wage)					-.0403**
					(.0058)
Firm Ownership Type:					
Foreign Ownership	-.0314**	-.0304**			
	(.0028)	(.0027)			
State-owned Enterprise	-.0065	-.0032			
	(.0036)	(.0031)			
Fixed Effects (number of groups)	Occ, Ind, Province (116)	Occ*Ind* Province (22,581)	Occ*Firm, Province (258,751)	Occ*Firm, Province (63,333)	Occ*Firm, Province (63,333)
N	1,051,706	1,051,706	1,051,706	172,887	172,887
R^2	.078	.155	.669	.760	.761

Dependent variable in all columns equals one if the ad explicitly requests either men or women, and zero otherwise. All regressions also control for log firm size, the number of positions advertised, period effects, and whether the job is part-time. Standard errors are clustered at the occupation * province level.

Table VII: Effects of Jobs' Skill Demands on Employers' Preferences Towards Men,(P^M - P^F), OLS Estimates:

	(1)	(2)	(3)	(4)	(5)
Education Requirement:					
Some Postsecondary	-.0057	-.0032	-.0127	-.0372**	-.0403**
	(.0067)	(.0065)	(.0083)	(.0141)	(.0136)
University	.0037	.0091	.0000	-.0285	-.0350*
	(.0078)	(.0076)	(.0090)	(.0173)	(.0162)
Experience Requirement:					
1-3 years	.0172**	.0128**	.0132**	.0071	.0058
	(.0027)	(.0027)	(.0029)	(.0072)	(.0075)
3-5 years	.0432**	.0398**	.0387**	.0451**	.0386**
	(.0045)	(.0047)	(.0057)	(.0100)	(.0113)
More than 5 years	.0621**	.0575**	.0489**	.0649**	.0536**
	(.0056)	(.0057)	(.0068)	(.0171)	(.0206)
Log (offered wage)					.0154
					(.0089)
Firm Ownership Type:					
Foreign Ownership	-.0101**	-.0080**			
	(.0022)	(.0021)			
State-owned Enterprise	.0127**	.0144**			
	(.0030)	(.0029)			
Fixed Effects (number of groups)	Occ, Ind, Province (116)	Occ*Ind* Province (22,581)	Occ*Firm, Province (258,751)	Occ*Firm, Province (63,333)	Occ*Firm, Province (63,333)
N	1,051,706	1,051,706	1,051,706	172,887	172,887
R^2	.064	.143	.641	.716	.716

Dependent variable in all columns equals minus one if the ad requests women, one if it requests men, and zero otherwise. All regressions also control for log firm size, the number of positions advertised, period effects, and whether the job is part-time. Standard errors are clustered at the occupation * province level.

Table VIII: Effects of Selected Covariates on Preferences for Other Ascriptive Characteristics:

	Ad is Age-Targeted?[a]		Mean Age Requested, given age-targeted		Ad Requests Beauty?		Ad has a Height Requirement?	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Education Requirement:								
Some Postsecondary	-.0273**	-.0250	.9761**	.0228	-.0297**	-.0112	-.0342**	-.0203**
	(.0045)	(.0143)	(.1768)	(.2714)	(.0046)	(.0103)	(.0037)	(.0059)
University	-.0389**	-.0494**	1.8288**	-.2707	-.0443**	-.0157	-.0390**	-.0253**
	(.0049)	(.0136)	(.2332)	(.3893)	(.0053)	(.0117)	(.0044)	(.0081)
Experience Requirement:								
1-3 years	-.0071**	-.0045	.3579**	.2705	-.0160**	-.0187**	-.0130**	-.0216**
	(.0021)	(.0068)	(.1305)	(.2022)	(.0027)	(.0059)	(.0022)	(.0056)
3-5 years	.0022	-.0068	2.6247**	2.3585**	-.0400**	-.0416**	-.0222**	-.0346**
	(.0027)	(.0097)	(.2032)	(.3547)	(.0055)	(.0080)	(.0037)	(.0091)
More than 5 years	.0153**	.0094	5.0780**	4.3621**	-.0471**	-.0503**	-.0250**	-.0348*
	(.0034)	(.0157)	(.2326)	(.4078)	(.0058)	(.0101)	(.0043)	(.0143)
Log (Wage)		.0083		2.2116**		-.0185**		-.0099*
		(.0069)		(.1944)		(.0056)		(.0041)
Fixed Effects (number of groups)	Occ*Firm, Province (258,751)	Occ*Firm, Province (63,333)	Occ*Firm, Province (50,201)	Occ*Firm, Province (11,797)	Occ*Firm, Province (258,751)	Occ*Firm, Province (63,333)	Occ*Firm, Province (258,751)	Occ*Firm, Province (63,333)
N	1,051,706	172,887	134,768	27,909	1,051,706	172,887	1,051,706	172,887
R^2	.670	.758	.885	.929	.659	.754	.668	.762

** p<.01, * p<.05. OLS estimates. Regressions also control for a the number of vacancies advertised, a dummy for part-time jobs and period fixed effects. Standard errors are clustered at the occupation*province level. Specifications are identical to columns (3) and (4) of Tables VI and VII.

Table IX: Effects of Jobs' Skill Demands on the Probability an Ad is Gender-Targetted ($P^M + P^F$) and on Preferences towards Men (P^M - P^F), OLS Estimates with controls for Other Ascriptive Characteristics

	Ad is Gender-Targeted ($P^M + P^F$)		Preferences Towards Men, ($P^M - P^F$)	
	(1)	(2)	(3)	(4)
JOB SKILL REQUIREMENTS:				
Education Requirement:				
Some Postsecondary	-.0451**	-.0125	-.0274**	-.0473**
	(.0043)	(.0089)	(.0085)	(.0135)
University	-.0619**	-.0072	-.0184*	-.0427**
	(.0051)	(.0106)	(.0093)	(.0164)
Experience Requirement:				
1-3 years	-.0165**	-.0188**	.0076**	-.0031
	(.0021)	(.0054)	(.0025)	(.0062)
3-5 years	-.0235**	-.0236**	.0246**	.0203*
	(.0023)	(.0066)	(.0038)	(.0084)
More than 5 years	-.0250**	-.0156	.0287**	.0297
	(.0029)	(.0093)	(.0045)	(.0159)
Log (offered wage)		-.0384**		.0053
		(.0056)		(.0089)
OTHER ASCRIPTIVE CHARACTERISTICS:				
Ad specifies age range?	.1954**	.0863	-.3792**	-.3013**
	(.0288)	(.0586)	(.0365)	(.0713)
Mean age, when specified (years)	-.0024**	.0013	.0134**	.0107**
	(.0009)	(.0019)	(.0011)	(.0025)
Ad requests beauty?	.0661**	.0705**	-.1141**	-.1390**
	(.0052)	(.0096)	(.0098)	(.0176)
Ad specifies minimum height?	.2545**	.2292**	-.2114**	-.2448**
	(.0161)	(.0245)	(.0183)	(.0469)
Fixed Effects (number of groups)	Occ*Firm, Province (258,751)	Occ*Firm, Province (63,333)	Occ*Firm, Province (258,751)	Occ*Firm, Province (63,333)
N	1,051,706	172,887	1,051,706	172,887
R^2	.684	.772	.652	.727

** p<.01, * p<.05. OLS estimates. Regressions also control for a the number of vacancies advertised, a dummy for part-time jobs and period fixed effects. Standard errors are clustered at the occupation*province level.

Table X: Effects of Education Requirements on Advertised Preferences in Highly Male and Highly Female Jobs

Education Requirement	Share of Ads requesting men in "Highly Male" Jobs:	Share of Ads requesting women in "Highly Female" Jobs:
High school or less	.640	.815
Some college	.534**	.637**
University	.509*	.517**

Highly Male Jobs are Technical Workers, maximum age 25 or over
Highly Female Jobs are Jobs that request beauty, maximum age under 25

**, * refer to differences from row one, significant at 1 and 5 percent respectively.

Table XI: Predicted Share of Gendered Job Ads for All Employed Workers, High-Income Province Sample

		Share of Job Ads:		
Predictions based on:	**Sample**	**Requesting Women**	**With no Gender Preference**	**Requesting Men**
1. Education	**Zhaopin**	.038	.912	.050
	Census	.088	.809	.103
2. Education, Industry, Occupation, Firm Type	**Zhaopin**	.038	.912	.050
	Census	.076	.792	.132
3. Education, Industry, Occupation, Firm Type, Wage	**Zhaopin**	.056	.889	.056
	Census	.090	.783	.127
4. Education, Industry, Occupation, Firm Type, Age	**Zhaopin**	.104	.752	.144
	Census	.125	.595	.280

Zhaopin rows show the predicted share of job ads in each category from a regression based on Zhaopin data using only the regressors listed. The Census rows use that Zhaopin regression to predict the category shares for a sample with mean characteristics of workers in the 2005 Census. Both Zhaopin and Census samples are restricted to urban workers in the eight provinces with the most Zhaopin ads, which are also top 8 provinces in GDP per capita: Beijing, Shanghai, Guangdong, Jiangsu, Shandong, Tianjing, Zhejiang and Liaoning

Figure I:
Occupation Fixed Effects for Gender, by Mean Education Requirements

a) Dependent Variable: Tendency to Gender-Target Ads ($P^M + P^F$)

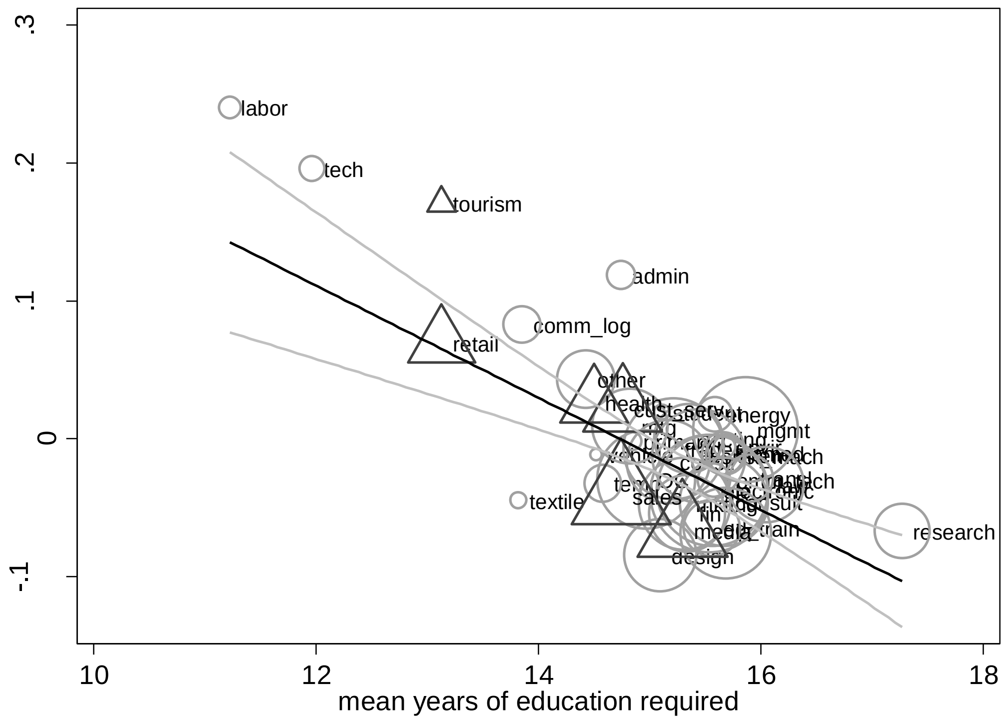

b) Dependent Variable: Employer's Relative Preference towards Men (P^M - P^F)

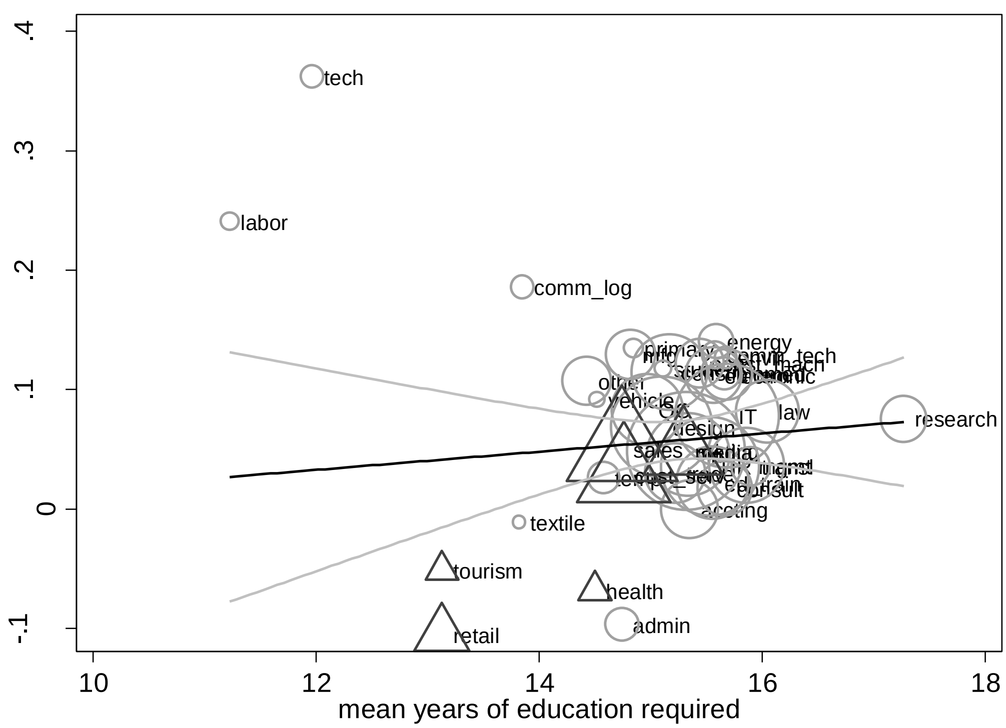

Notes to Figure 1:

Triangles represent customer-contact occupations, circles all other occupations.
Symbol sizes are proportional to the precision of the estimated fixed effect.
Dark line shows a linear regression (estimated with these weights).
Lighter lines show the 95% confidence band.