Explicit Discrimination in Hiring:

Evidence from a Chinese Internet Job Board

Peter Kuhn[a]
University of California, Santa Barbara

Kailing Shen[b]
Xiamen University

February 19, 2009

Preliminary and Incomplete: Please do not Quote

We study patterns of firms' advertised gender preferences in a sample of job ads posted on an internet board in China, and interpret these patterns using a simple nonsequential employer search model. We find that firms are less likely to express a gender preference in *either* direction when the expected number of applicants per vacant position is low, when the firm has some foreign ownership, and when the job demands a high level of education or management experience. At the same time, firms' advertised preferences move towards men and *away* from women as firms grow in size, when they are state-owned, and when the job requires a higher level of experience.

[a](corresponding author) Department of Economics, University of California, Santa Barbara, Santa Barbara, CA USA 93106, NBER and IZA, pjkuhn@econ.ucsb.edu.; [b]Wang Yanan Institute for Studies in Economics (WISE), Xiamen University, China 361005 and IZA, klshen@xmu.edu.cn
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 for their generous hospitality during the research on this project.

1. Introduction

A central difficulty in studying employment discrimination in most developed countries is the fact that discrimination is illegal. As a result, economists are usually forced to rely on indirect measures of discrimination (such as the 'unexplained' effect of gender or race in a wage regression), which are subject to well known problems.[1] Direct evidence of discriminatory actions by employers is notoriously difficult to establish in court, and is largely absent in academic research based on large representative samples.

In this paper we study the prevalence and patterns of explicit gender discrimination in hiring using a large sample of ads from an internet job board in China.[2] In China, legal and social sanctions against specifying a desired gender (and indeed age, height or geographical origin) of prospective employees are largely absent; we therefore observe employer preferences among these worker attributes directly and use observed patterns of these preferences to shed light on several well-known theories of discrimination, including taste-based models, and models based on productivity differences between groups.[3] A key question we address is the role of labor market conditions –in particular the expected scarcity of applicants for a position—in employers' decisions to specify a gender preference in their job ads.

We begin by providing a theoretical context for the study of explicit discrimination in job advertising: why, and under what conditions, might we expect an employer to invite only members of a specific demographic group to apply for a particular vacancy? We then use those theoretical perspectives to interpret the patterns of explicit gender discrimination we see in a large sample of job ads.

2. A Model

In our data, we observe firms explicitly stating that they are seeking workers with certain easily observable attributes, such as age, sex and education, for a particular position. Under what conditions might such restrictions be in the interests of a profit- or utility-maximizing employer, and what is the best way to model employers' decisions to impose these restrictions? Before laying out a specific model, we make two observations.

---

[1] A few economists, e.g. Kuhn (1987) and Antecol and Kuhn (2000) have studied workers' reports of perceived employment discrimination. Charles and Guryan (2008) examine the association between state-level mean survey reports of discriminatory attitudes (interpreted as pertaining to employers and co-workers) and black-white wage differences.

[2] In future work, we will study patterns of employer discrimination along other margins, including age, height and hukao (residence permit) using this data as well.

[3] Of course, if actual productivity differences are employers' primary motivation for preferring one demographic group over another for a particular job, it is not clear whether employers' preferences are properly labelled discriminatory. Throughout this paper we use "discrimination" to refer to any employer preference for one demographic group over another for a particular job, regardless of its source.

First, for firms to wish to advertise a preferred level of an easily-observable worker attribute, it must be the case that it is costly simply to receive and process job applications, even if the only processing that is done is simply to discard applications not meeting these easily observable requirements.[4] After all, in the absence of such costs, firms could costlessly duplicate the effects of any advertised requirement by soliciting applications from all, then just discarding the applications not meeting certain simple criteria. For this reason, processing costs, and their structure, play an important role in any analysis of discrimination in job advertising, ours included.

For simplicity, in most of our analysis below, we assume a constant per application processing cost. Payment of this cost reveals not only the easily observable aspects of the applicant (such as gender, age and education) but also the subtler ones, including the value of the worker-specific match with the job. We also discuss, however, how our main results change when we introduce a distinction between the cost of screening applications for whether they satisfy simple criteria such as age or sex, and the cost of learning the match-specific component. All of our analysis assumes what we call "free random disposal" of applications; this means that a firm can avoid processing costs on any random sample of applications by discarding them.[5]

A second initial observation concerns the type of search model in which we embed our analysis of advertised hiring restrictions.[6] We use a simple nonsequential search model, based on evidence in Van Ours and Ridder (1992, 1993) and Weber (2000), who show that vacancy durations are not well described by a sequential model where firms set a reservation level of worker quality, then wait until an application exceeding that quality arrives. Instead, the vast majority of applications arrive shortly after a vacancy is advertised, and vacancy durations largely consist of the time taken by the firm to select the best candidate from this pool. Thus, to focus on essentials we consider a static model where the application pool from each demographic group that is potentially invited to apply is fixed in size; the firm's two decisions are (a) to determine which demographic groups to invite to apply, then (b) to find the best candidate in the applicant pool.[7]

---

[4] The fact that firms in developed economies routinely list education requirements in job ads (which is of course legal) is *prima facie* evidence of such costs.

[5] If applications arrive at a relatively constant rate, this could be achieved simply by closing off the application process at a particular date. Stopping the process is less effective if, as is commonly the case, most applications arrive *en masse* shortly the job ad is posted.

[6] To our knowledge the only empirical or theoretical models of employers' advertised job requirements are Barron, Bishop and Dunkelberg (1985) and Van Ours and Ridder (1991). Both of these papers treat these requirements as exogenous vacancy characteristics, rather than a choice variable for the employer. Kuhn (1993) models employers' choices of gender-based hiring restrictions, but not in a search context.

[7] That said, similar results to ours can be likely be derived in a sequential search context. For example, Burdett and Cunningham (1998, p.449) argue that an increase in the arrival rate of applicants in a sequential employer search model leads firms to raise the reservation ability level, i.e. to become more selective (see Mortensen 1986, p. 865 for a proof). This is similar in spirit to our result that a larger overall expected applicant pool raises the likelihood that firms will 'tighten' their advertised restrictions on such characteristics as age, sex and education.

We turn now to a formal model. Consider a firm inviting applications for a vacant position; applications can come from two distinct groups labelled *A* and *B,* where *A* and *B* also represent the number of applications that would be received from each group. Let the value to the firm of an individual applicant, *j*, be given by $u_j = v^A + e_j$, and $u_j = v^B + e_j$ for groups *A* and *B* respectively, where the e_j represent independent draws from the same *cdf*, *F(e).*[8] For the position in question, we assume that members of group *A* are on average more highly desired, i.e. that $v^A > v^B$. The firm is assumed to choose the worker with the highest total value, u_j, from its pool of applicants. The cost of evaluating an application (thus learning its e_j) is assumed to be *c* per application; we assume that the firm can avoid this cost (thereby learning nothing about the applicant's e_j) by simply discarding a random subset of applications received.[9] The question we pose is whether, and under what conditions, an expected-profit-maximizing firm will wish to invite only one of the two above groups to apply.[10]

We begin by denoting the expected maximum value of u_j in a pool of *n* applicants with 'baseline" productivity $v \quad \{v^A, v^B\}$ by *H(v, n)*; under quite general conditions $H_1 > 0$, $H_2 > 0$, and $H_{22} < 0$.[11] If the firm chooses a search strategy to maximize expected profits net of search costs, $H(v,n) - cn$, then it simply selects the highest level of profits from the following three cases:[12]

A: Invite *A*'s only:	$E(\ ) = H(v^A,A) - cA$
B: Invite *B*'s only:	$E(\ ) = H(v^B,B) - cB$
C: Combined search— invite all:	$E(\ ) = G[v^A,A;\ v^B,B] - c(A+B),$

where $G[v^A,A;\ v^B,B]$ is the expected value of the best match chosen from a sample composed of *A* applications with baseline value v^A, and *B* applications with baseline value v^B.]

[8] The worker's total 'value', u_j, includes both actual productivity as well as any discriminatory tastes or inaccurate perceptions the employer might hold concerning the two groups. It should also be interpreted as net of wages paid; thus our formulation implicitly abstracts from individual-specific wage bargaining. Our specification does, however, allow for the possibility that the wage attached to a job could differ between

Because it is difficult --except in degenerate cases such as $v^A=v^B$-- to make useful comparisons between the *G* and *H* functions for an arbitrary underlying *cdf* of match qualities *F(e)*, we derive most of our remaining results for the case where e_j follows a Type-1 extreme-value distribution, i.e. where $F(e) = exp(-exp(-e))$. In this case, $H(v^J,J) = v^J + \gamma + log(J),\ J\ \ \{A, B\}$, where γ is Euler's constant (≈.577).[13] Further, Appendix 1 shows that $G[v^A,A; v^B,B] = \gamma + v^C + log(M)$, where $M \equiv A+B$. $v^C \equiv log[\delta\, exp(v^A) + (1-\delta)exp(v^B)]$ thus functions as the'base' productivity in the combined "*C*" sample, and $\delta = A/M$ is the fraction of group *A* (the more 'desired' group) in the combined pool. Thus, a useful feature of the extreme value case is that, in the combined sample, the marginal return to an additional applicant is independent of the applicant's type.

In sum, in the extreme-value case the firm's choices between the three recruitment strategies yield the following levels of expected profits:

Strategy A: Invite *A*'s only:	$E(\ ) = v^A + \gamma + log(A) - cA$
Strategy B: Invite *B*'s only:	$E(\ ) = v^B + \gamma + log(B) - cB$
Strategy C: Invite all comers:	$E(\ ) = v^C + \gamma + log(M) - cM$

The profit-maximizing number of applicants if the firm could choose the number of applications it receives --i.e. the *n* that maximizes $v\ + \gamma + log(n) – cn$-- therefore equals $1/c \equiv N^*$, irrespective of whether the firm chooses strategy A, B or C.

A comparative static of interest to us is the effect of expanding the market (making applications more plentiful overall) on firms' preferences among the three strategies above. Market thickness is parameterized by *M*, the number of applicants in the combined pool; we consider how the firms' preferences among strategies change as we expand *M* starting from its smallest possible value, keeping the shares of the *A*s and *B*s available constant, i.e. holding $A=\delta M$ and $B=(1-\delta)M$. To focus on nontrivial cases, all of our theoretical results are conditional on Assumption 1, which simply states that applications are scarce at the minimum level of market size we consider, i.e. when there is only one applicant of the "preferred" type available; this level is just $M_{min} = 1/\delta$.

Assumption 1: $N^* > M_{min}$.

Proposition 1: When $v^A - v^C – log(1/\delta) + c(1-\delta)/\delta > 0$, firms will invite only the favored group (*A*) to apply at all levels of "market tightness", *M*. Otherwise, there exists a critical value of market thickness, $\tilde{M}$, above which firms invite only the *A*s to apply, and below which firms invite applications from all.

Proof: In Appendix 1.

[13] See Arcidiacono and Miller (2008, p. 8) for a general proof; our case is an application of their results to the multinomial logit (MNL) case.

P**roposition 2:** If the productivity difference between the groups, $v^A - v^B$, rises, or as processing costs, *c*, rise, $\tilde{M}$ falls.

Proof: In Appendix 1.

Proposition 2 implies that advertised hiring restrictions become more 'likely' when processing costs, and between-group productivity differences, are high.

How robust are Propositions 1 and 2 to changes in assumptions? First, while we have assumed constant per-application processing costs, it is straightforward to show that our main results are unchanged as long as economies of scale in processing costs are not too great.[14] Second, it is reasonable to ask what happens if we introduce a distinction between two types of application processing costs: (a) "pre-screening" applications for the presence of the preferred or dispreferred demographic attribute, and (b) determining an applicant's individual match quality, *e*. If (a) can be done very cheaply, will firms still advertise hiring restrictions? Clearly they will, since simply announcing a preference in a job ad has the same effect as pre-screening and costs nothing.[15] Finally, we note that, for the most part our model's results are also unchanged if we drop the assumption of free random disposal of applications. The main difference is that, when applications are extremely plentiful (i.e. when *M* is very large) and when the less-productive group, *B*, is smaller in size (i.e. $\delta > .5$), the firm might choose to invite only the *less*-productive group to apply just to save on processing costs. We do not think this is very likely.

We conclude our theoretical discussion with a comment about expected magnitudes. Essentially, our model says that firms advertise that they are not interested in receiving applications from a specific demographic group when the firm's *ex ante* assessment of the chance that the best overall candidate will come from that group is so low that the expected benefits of examining that group's applications are outweighed by the costs of processing them. If marginal processing costs are low, we should therefore expect to see advertised hiring restrictions relatively rarely— these should be seen only when firms are highly confident of a large gap between the suitabilities of the two groups for the job in question.[16] Put another way, according to our model, the incidence of advertised hiring restrictions should not be interpreted as the share of employers who who prefer one demographic group over another for a particular type of vacancy. Instead, the rate of advertised hiring restrictions gives the share of vacancies in which

---

[14] Further, our results must hold in the following sense for *any* processing cost function that satisfies the second order condition for an optimal number of applications (i.e. as long as N^* exists): If N^* exists, then additional applications beyond N^* only raise costs. Thus any increase in market thickness beyond N^* increases the likelihood that firms will advertise hiring restrictions.

[15] That said, if there are social or legal sanctions to advertising a preference, pre-screening might dominate advertising the firm's preferences.

[16] Possible sources of such large perceived gaps include nonconvex effects of the share of *A*s and *B*s in a job on costs or productivity. For example, if it is illegal to pay different wages to men and women *in the same job*, then adding a single man or woman to a previously completely segregated job category could have very large effects on the wage bill. Alternatively, the first representative of a group in a previously segregated job could be subject to considerable harassment and intimidation, which could have large negative productivity effects.

firms' preferences for a certain group are sufficiently intense, i.e. these preferences must exceed a (positive) threshold, which rises as application processing costs fall.

3. Data

Our data are a sample of job advertisements posted between May 16 and July 12, 2008 on Zhaopin.com, the third-largest online job board in China.[17] Procedures for downloading the sample and defining variables are discussed in Appendix 2. Clearly, ads on Zhaopin.com are not representative of all jobs in China; like all samples of job ads they will overrepresent jobs in expanding and high-turnover occupations and industries. In addition, the jobs on Zhaopin.com likely require a significantly higher skill level than the median job in China (for one thing, Zhaopin is not accessible to workers without internet access or who are illiterate). Since, according to our own data, gender discrimination is more common in less-skilled occupations, our data likely underestimate the extent of such discrimination in China.[18]

Another sampling issue reflects the fact that some of the ads in our sample are for multiple vacancies.[19] While most of our analysis treats the unit of analysis as the job ad (using the number of vacancies it advertises as a control variable of interest) we estimated some specifications which weight the sample by the number of vacancies. Finally, we note that, by construction, our sample is a stock sample of "ads in progress" rather than a flow sample of newly-posted (or just-filled) ads. This implies that long vacancy spells are overrepresented in the data, which would affect our estimates if there is parameter heterogeneity that is correlated with durations (for example if the sensitivity of gender restrictions to occupation is greater in jobs that take a long time to fill). To address this concern, we also replicated our estimates for a subsample of ads that consists, almost certainly, of newly posted ads.[20]

Descriptive Statistics of our sample are provided in Table 1. All told, we have data on almost half a million job ads; these tend to be for relatively highly-skilled jobs in large urban centers. Indeed about 70 percent of ads require at least some post-secondary education, and 13 percent require management experience. 26 and 18 percent of the ads were for jobs in Beijing and Shanhai respectively. By far the most common occupation is sales, at 20 percent of the ads, with computer-related occupations second at about 9 percent. The top five industries were consulting, IT service, construction, software and internet/e-commerce. Only 1.4 percent of the ads are for part time jobs. A typical ad in

---

[17] The two larger sites are www.51job.com, 15% owned by the Japanese firm Recruit, and Chinahr.com (45% owned by Monster.com). (Zhaopin is owned by the Australian firm, Seek). Our choice of Zhaopin is largely due to technical reasons—unlike the other firms at the time we collected the data, Zhaopin provided a stable link to each job and each firm, which allows us to easily check whether an ad has been renewed.

[18] Of course, similar representativeness issues affect research that uses U.S. internet job postings, such as (see, e.g. Brencic and Norris 2008, Kroft and Pope 2008).

[19] The most extreme case was a single ad for 8000 vacancies at a newly-opened chemical plant. Since our sample is so large, our results are not materially affected by excluding this observation or by assigning it a weight of 8000 times an ad for a single vacancy.

[20] Appendix 2 describes the creation of this "inflow sample".

our sample was renewed 15 times, and half of the ads specified the number of vacancies that were available. Half of these, in turn, were for a single position. As already noted, a significant share of the ads, however, were for large numbers of job openings, sometimes in excess of 50. Our sample of job ads covers the entire spectrum of firm sizes. Finally, some comments about firm types are in order. The first three firm categories available to us have some international connection: Foreign Direct Investments (FDI), Representative Offices and Joint Ventures; together these comprise about 36.9 percent of the ads in our sample. Publicly-traded and privately-held Chinese for-profit companies comprise another 38.4% of the sample together; a further 9.9 percent of the ads are from corporations that were once State-Owned Enterprises.[21] SOEs themselves account for 7.3% of the ads, and a miniscule share come from the public sector (local, state and national government) comprises a miniscule share (0.06%); most of the public sector's recruiting is done via other channels.

Finally, 4.6 percent of the ads in our sample expressed a preference for a male applicant; 4.8 percent for a female applicant.[22] This relatively low incidence of advertised gender preference is perhaps not surprising, given the previous section's theoretical discussion.

4. Predictions

The model presented in Section 2 predicts that firms' decisions to target their job ads to a particular demographic group should respond unambiguously to three distinct types of factors: explicit discrimination in job ads should become more frequent the greater the number of applications the firm expects to receive (higher M), the greater the firm's preference for one demographic group over another ($v^A - v^B$) and the higher are application processing costs, c.[23] In this section we discuss how these theoretical constructs are mapped into testable predictions using the data available to us.

a. Expected Market thickness, M:

In our data, we can construct the following indicators of M: (1) the number of positions the firm wishes to fill with the current ad (more positions indicative of a *lower* M per vacancy); (2) the provincial unemployment rate; and (3) the mean number of times an ad was renewed in an ad's occupation/province-specific cell. To derive indicator (3), we first restrict attention to ads that impose no gender restrictions (this eliminates any

[21] The share of publicly traded companies may seem very low to American readers; some of this may be due to a conflation of the "publicly held company" and "corporation" categories in our data. That said, in June 2007 there were only 1477 publicly traded companies in all of China that had shares traded domestically. Both formal and informal minimum requirements for firms to go public are very strict in China.

[22] This includes all "intensities" of preference, though the most typical employer statements were either "female[male] preferred" and "female[male] only".

[23] The other parameter in the model, δ (the share of the favored group in the expected applicant pool) has complex and ambiguous predicted effects.

direct effects of gender restrictions on the speed with which an ad is filled). Then we calculate the mean number of times an ad is renewed in each cell, as an indication of how hard it is to fill a job in that cell.[24]

Finally, we note that the direction of predicted effects of market size, *M*, do not depend on *which* group the firm favors (i.e. the assignment of the labels A and B is arbitrary); since in our data we observe that women are solicited for some jobs and men for others, our model predicts that the likelihood that firms advertise a preference in *either* direction should increase with our measures of market thickness, *M*.

b. Intensity of preference for the favored group ($v^A - v^B$):

One obvious reason why $v^A - v^B$ might vary across jobs is that on average, men and women possess different skill bundles, and that jobs vary in the mix of skills they require.[25] Another well known possible source of employer preferences for one demographic group over another is discriminatory tastes .

s in

t

their job ads. Notably, *product market* competition (unlike labor market competition, *M*) should only affect firms' gender-based hiring restrictions if those restrictions are based on factors unrelated to workers' true productivity; thus we take any estimated effects of product market competition on hiring restrictions as *prima facie* evidence of taste-based discrimination.

A final possible determinant of the expected group productivity differential, v^A - v^B, is the overall skill level required in the job, measured in our data by its advertised education and experience requirements. Of course, if higher overall skill demands raise both v^A and v^B by the same amount, no change in advertised gender preferences is predicted. On the other hand, economists sometimes argue that higher-ranked positions in a firm are jobs where individual-specific skill differences (*e* in our model) "matter more" to the firm (see for example Gibbons and Waldman 1999). In this case, in jobs requiring more education and experience, the variance of *e* is is higher relative to the mean difference between groups, v^A - v^B, making gender restrictions less in firms' interests. Accordingly we estimate the sensitivity of gender restrictions to a job's education and experience requirements.

c. Application processing costs

The only (crude) proxy we have in our data for marginal application processing costs is the size (number of employees) in the hiring firm. If there are economies of scale in processing costs, then large firms should have lower marginal processing costs and have less of a need to place *ex ante* restrictions in their job ads. Of course, if for some reason discriminatory tastes are more common among large employers, this could reverse the above effects.

5. Results

In all of the regression results reported here, an individual observation, *i*, represents a single job ad (if an ad is renewed --which occurs frequently in our data--, we make note of that fact, but do not treat it as a new observation). Similarly, if an ad is for more than one opening, we again make note of that fact, but again we treat the ad as a single observation.[28] The main results we report are for the entire sample of ads we collected (which includes a significant number that were 'in progress' when we first started downloading data), but Appendix 3 also reports results for a 'flow' sample of ads that were placed for the first time during our observation window.

All the regressions we present are linear probability models in which the dependent variable is whether an ad mentioned a preference for an applicant of a specific gender. As should be clear from the previous section, some of the covariates available to us vary at the ad level, *i*, for example the number of positions being advertised. Others

[28] That said, in some robustness checks, we weight ads by the number of positions they represent, to obtain estimates that are representative of the universe of *positions* advertised.

vary at the firm level, *f*, (for example, firm ownership and size); note that our data allow us to identify multiple ads placed by the same firm. Still other covariates vary at the province level, *p* (for example the unemployment rate), and by occupation, industry and occupation/province cells.

Tables 2 and 3 present linear probability estimates of firms' advertised preferences for men and women respectively. Each table reports the results of three different specifications. In column 1, the only regressors are those that vary at the ad or firm level, specifically the number of positions advertised, the education and experience requirements, whether the position is part time, and the firm's size and ownership category. Column 2 adds to this specification two indicators of market thickness, *M*, that vary at higher levels of aggregation, specifically the provincial unemployment rate and the renewal probability of ads in that occupation/province cell. Column 3 adds a full set of industry, occupation and province fixed effects to the specification; these of course absorb all the variation in provincial unemployment rates, but still allow us to estimate the effects of the number of positions advertised and the renewal rate.

Recalling that the omitted category for the 'number of positions' indicator variables is ads that do that do not specify the number of positions, Table 2 strongly supports our model's prediction that firms who are seeking to fill a larger number of vacancies are less likely to express a preference for male applicants. With one exception (the 16-50 openings category), Table 3 shows that the same is true regarding preferences for female applicants. Also as predicted, firms are more likely to specify a preferred gender (whether male or female) when job applicants are expected to be plentiful, i.e. in provinces with high unemployment rates. The same is also true for our other indicator of market thickness, the mean number of renewals per ad in an ad's occupation-province cell: A high renewal rate indicates that other firms' vacancies in that cell tend to stay open a long time, suggesting that qualified applicants are scarce. When this is the case, firms are again less likely to express gender preferences (in either direction) in their on-line job ads. The estimated effect, however, becomes statistically insignificant when we control for a full set of occupation, industry and province fixed effects.

Turning to firm ownership types, the three categories with some foreign involvement (FDI, representative offices and joint ventures) are less likely than the reference category (privately held firms) to express a gender preference (in either direction) in their ads; this effect is highly robust to occupation, industry and province fixed effects. This could either reflect greater exposure of such firms to product-market competition, or a different business culture that disapproves of such restrictions. Compared to the omitted category (privately held firms) publicly-traded corporations are less likely to express a gender preference; this is also true for former SOEs that are now "Corporations" (though not so much for men). Finally, two firm types are consistently more likely than privately-held firms to express a preference for men and *less* likely to express a preference for women: state-owned enterprises and the public sector itself. Since these employers are more likely to invite one group to apply but less likely to invite the other group to apply, according to our model these differences cannot easily be explained by the overall tightness of the labor market conditions they face or in their costs of processing applications. Instead this pattern is consistent only with a stronger preference for male versus female employees, v^A - v^B, or with a greater latitude to indulge those preferences due to a lack of competition.[29]

According to Table 2, large firms are more likely to solicit applications targetted at men only, and less likely to solicit applications targetted at women. These trends are essentially monotonic across all firm sizes (with the exception of the very largest category for men only) and highly robust to controls for occupation, industry, job skill requirements and region. They are not consistent with our hypothesis that, since larger firms have more sophisticated (and automated) human resource management systems (and therefore lower per-application processing costs), they should have less of a need to target their ads to specific demographic groups. Aside from the possibility of greater insulation from competitive forces among the larger firms, it is hard to think of a plausible explanation for this very robust finding.

Finally, we consider the effects of a job's skill requirements on the prevalence of advertised gender preferences. According to Table 2, jobs that require a great deal of experience are more likely than other jobs to explicitly solicit male applicants, and *less* likely to solicit female applicants. In this sense, experience and "maleness" appear to be complements in our data, while experience and "femaleness" are substitutes. Interestingly, this is not true of education requirements or of the need for *management* experience: jobs that request a high level of education or management experience are much less likely to specify a gender preference *in either direction* than lower-skilled jobs. Unlike our results for experience *per se*, these findings are consistent with the view that idiosyncratic individual ability "matters more" in jobs with high skill demands, making gender restrictions a less effective way to screen out unqualified applicants.

One possible concern regarding the preceding analysis is that by construction, the estimation sample consists disproportionately of job ads that took a long time to fill. To address this concern, we constructed a subsample of ads that, almost surely, includes only ads that were *first* posted on Zhaopin after the start of our observation period (i.e. May 16, 2008). The procedure for defining this sample is described in Appendix 2. We then re-estimated Tables 2 and 3 on this subsample. The results are reported in Appendix 3 and are very similar to our full-sample results; the one difference of note is that the mean ad renewal rate now has statistically significant effects in the presence of occupation, industry and province fixed effects. For men, however, its sign is contrary to our theoretical expectations.[30]

---

[29] Since the dramatic SOE reforms of the mid 1990s, the remaining Chinese SOEs have tended to be among the more profitable types of firms in China. Another explanation for SOEs' preferences for men that has been suggested is that SOEs face stricter legal requirements than other firms to pay very generous maternity benefits.

[30] To address a different concern about the representativeness of our sample, we also re-estimated Tables 2 and 3, weighting each ad by the number of positions it advertised. (This generates estimates that are representative of all vacancies, rather than all ads.) The results were once again very similar.

Another possible concern is the effect of unmeasured firm-specific characteristics on some of our results, including the "market thickness" effects that very strongly and consistently suggest that firms are less likely to announce gender preferences in their job ads when workers are hard to find. To address this concern, Table 4 adds a full set of firm fixed effects to our analysis (in addition to the province, industry, and occupation effects already present). This means, of course, that we can only estimate the effects of factors affecting a job ad that vary across ads within firms. That said, for the variables whose impact we can examine, the results from Table 4 are quite consistent with Tables 2 and 3 (though the standard errors are, not surprisingly, higher). For example, looking within firms, we now find that firms are less 'choosy' about an applicant's gender when they are filling a large number of vacancies than a smaller number.

In the presence of firm fixed effects, a higher ad renewal rate in an occupation-province cell (indicating that vacancies are hard to fill) still reduces the incidence of advertised gender requirements (though the effect is statistically insignificant for men). Like Tables 2 and 3, Table 4 again suggests that experience is a complement with "maleness" and a substitute with "femaleness", while jobs requiring high levels of education or management experience are less likely to specify a gender preference in either direction. We conclude that unobserved firm characteristics are not responsible for these highly robust patterns in our data either.

7. Summary

This paper has studied patterns of firms' advertised gender preferences in a sample of internet job ads posted in China, and has attempted to interpret these patterns using a simple nonsequential employer search model. According to the model, increases in market thickness (i.e. in the expected number of applicants per position) should increase firms' propensity to specify a preferred gender (either male *or* female) in their job advertisements. We find considerable support for this prediction using three separate indicators of market thickness: the number of vacancies the firm hopes to fill with the ad, the provincial unemployment rate, and the mean number of times that employers renew job ads in an occupation/province cell.

We also examine the hypothesis that *if* firms have Beckerian discriminatory preferences for males over females, they should be more likely to express these preferences in job ads when product market competition is low. We interpret our findings that firms' advertised preferences move towards men and away from women as firms grow in size and when they are state-owned as consistent with this hypothesis.

In our sample, firms' advertised preferences also move towards men and away from women as the job's required experience level job rises. At the same time, advertised gender preferences for *both* men and women become less frequent when jobs require high levels of education or management experience. We interpret the latter two findings (for education and management experience) as consistent with the notion that gender is a less informative signal in skill-intensive jobs where 'ability matters more'.

In ongoing and future work we hope to improve on the current analysis using improved indicators of market thickness, product market competition, and application processing costs, and by exploiting temporal variation in the available measures of these key theoretical contructs. We also plan to extend our analysis to firms' advertised preferences regarding the applicant's age, height, and *hukao* (residence permit).

References

Antecol, Heather and Peter Kuhn. "Gender as an Impediment to Labor Market Success: Why do Young Women Report Greater Harm?", *Journal of Labor Economic*s, 18(4), (October 2000): 702-28.

Arcidiacono, Peter and Robert A. Miller. "CCP Estimation of Dynamic Discrete Choice Modelswith Unobserved Heterogeneity" Unpublished paper, Carnegie Mellon University, February 20, 2008

Bacolod, Marigee and Bernardo S. Blum. "Two Sides of the Same Coin: U.S. "Residual" Inequality and the Gender Gap" unpublished paper, University of California-Irvine, 2006.

Barron, John M., John Bishop, and William C. Dunkelberg. "Employer Search: The Interviewing and Hiring of New Employees". *Review of Economics and Statistics* 67 (1) (Feb. 1985): 43-52.

Black, Sandra E. and Elizabeth Brainerd. "Importing Equality? The Impact of Globalization on Gender Discrimination" *Industrial and Labor Relations Review*. July 2004; 57(4): 540-59

Black, S. and P. Strachan. "The Division of Spoils: Rent-Sharing and Discrimination in a Regulated Industry". *American Economic Review* 91(4) (September 2001): 814-831.

Black, Sandra E. and Alexandra Spitz-Oener "Explaining Women's Success: Technological Change and the Skill Content of Women's Work" NBER Working paper no. 13116, May 2007

Brencic, Vera and John B. Norris. "Employers' Online Search: An Empirical Analysis" unpublished paper, University of Alberta, 2008.

Brown, Meta, Christopher J. Flinn and Andrew Schotter "Real-Time Search in the Laboratory and the Market". Unpublished paper, New York University, September 2007

Burdett, Kenneth and Elizabeth J. Cunningham. "Toward a Theory of Vacancies" *Journal of Labor Economics*. July 1998; 16(3): 445-78 (show that higher search costs reduce the employer's reservation match productivity).

Charles, Kerwin Kofi and Jonathan Guryan. "Prejudice and Wages: An Empirical Assessment of Becker's The Economics of Discrimination" *Journal of Political Economy,* 2008, vol. 116, no. 5: 773-809

Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg (2006): "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model," American Economic Review, 96, 1043-1068.

Gibbons, R. and M. Waldman, "A Theory of Wage and Promotion Dynamics within Firms". *Quarterly Journal of Economics* 114 (November 1999): 1321-1358.

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

Kroft, Kory and Devin G. Pope. “Does Online Search Crowd Out Traditional Search and Improve Matching Efficiency? Evidence from Craigslis

Table 1: Descriptive Statistics

Total Sample	499,549
Education	
No Restrictions	21.4%
Grade 9	0.5%
High School	8.3%
Post Secondary	38.3%
University	30.5%
Master's degree	1.0%
PhD	0.7%
Require management experience	13.1%
Work Experience requirement	
No Experience required	32.1%
1 Year or Less	1.4%
1 to 3 Years	35.7%
3 to 5 Years	19.1%
5 to 10 Years	10.3%
10 Years or Above	1.4%
Gender Requirement	
Prefer male?	4.6%
Prefer female?	4.8%
Job Location (top 5)	
Beijing	26.3%
Shanghai	18.1%
Guangdong	7.8%
Shandong	7.2%
Jiangsu	5.4%
Occupation[a] (top 5)	
Sales	20.1%
Computer hardware and software/Internet/IT	9.2%
Marketing/PR	7.7%
Administrative/Logistics	5.7%
Accounting/Audit/Statistics	5.4%
Part Time Position?	1.4%
Provincial Unemployment Rate[b]	6.1%
Mean Number of Renewals per Ad	15.0
Number of Positions	
Unspecified	50.3%
1	24.1%
2	9.8%
3-5	9.0%
6-15	4.7%
16-50	1.7%
Over 50	0.4%

Table 1, continued:

Firm Ownership Type:[c]	
Foreign Direct Investment (FDI)	19.8%
Representative Office	1.4%
Joint venture	15.7%
Publicly Held Company	0.7%
Privately Held Company	37.7%
Corporation—former SOE	9.9%
Non-profit organization	0.4%
State-owned (SOE)	7.3%
Public sector	0.1%
Other	7.0%
Industry[a] **(top 5):**	
Consulting(Management/Legal/Accounting)	14.5%
IT Service (System/Database/Maintenence)	12.2%
Construction	11.6%
Software	9.6%
Internet/E-Commerce	9.6%
Firm Size:	
1-19	9.0%
20-99	36.0%
100-499	30.9%
500-999	9.1%
1,000-9,999	11.5%
10,000 +	3.5%

[a] Ads may specify more than one occupation and industry.

[b]This is the ratio of unemployed workers (job losers plus persons who have never had a job but are looking for one) to the *working population*.

[c] "FDI" denotes an enterprise that is solely owned by foreign investors. Representative offices can only be set up by foreign firms in China; they have no legal authority to sign contracts on their own and all their funding must come from outside China. "Publicly Held Companies" refers to firms whose stocks are traded on an exchange. "Corporations" also have multiple share-holders; most are reformed State-Owned-Enterprises (SOE) with the majority of shares still owned by the state. Non-profit organizations provide social services, such as education, healthcare, etc. "Public" denotes local, province or national government employment.

Table 2: Determinants of Firms' Advertised Preferences for Men

	Specification		
	Ad- or firm-level covariates only	**Ad- or firm-level and aggregate covariates**	**Ad- or firm-level covariates plus occ, ind and province fixed effects**
INDICATORS OF MARKET THICKNESS (*M*):			
Number of Positions			
1	.0192(.0008)***	.0168(.0008)***	.0143(.0008)***
2	.0216(.0010)***	.0189(.0010)***	.0154(.0010)***
3-5	.0094(.0011)***	.0079(.0011)***	.0072(.0011)***
6-15	.0077(.0014)***	.0082(.0014)***	.0090(.0014)***
16-50	-.0076(.0023)***	-.0060(.0023)***	-.0036(.0023)*
Over 50	-.0160(.0050)***	-.0125(.0050)***	-.0078(.0050)*
Mean Ad Renewal Rate	---	-.0033(.0001)***	-.0000(.0001)
Provincial Unemployment Rate	---	.0025(.0001)***	---
DETERMINANTS OF (v^A - v^B) and *c*:			
Part Time Position?	-.0337(.0025)***	-.0190(.0025)***	-.0137(.0027)***
Firm Ownership Type:			
FDI	-.0273(.0008)***	-.0246(.0009)***	-.0263(.0009)***
Representative Office	-.0171(.0028)***	-.0154(.0028)***	-.0124(.0028)***
Joint venture	-.0168(.0009)***	-.0132(.0009)***	-.0114(.0009)***
Publicly-Traded Company	-.0207(.0035)***	-.0179(.0035)***	-.0139(.0034)***
"Corporation"	-.0073(.0011)***	-.0108(.0011)***	-.0054(.0011)***
Non-profit	-.0068(.0047)*	-.0013(.0047)	.0117(.0047)***
State-owned enterprise (SOE)	.0075(.0012)***	.0104(.0012)***	.0085(.0012)***
Public sector	.0803(.0117)***	.0811(.0117)***	.0790(.0116)***
Other	.0039(.0012)***	-.0032(.0012)***	-.0029(.0012)**
Firm Size:			
20-99	.0108(.0011)***	.0089(.0011)***	.0062(.0011)***
100-499	.0218(.0011)***	.0199(.0011)***	.0152(.0011)***
500-999	.0197(.0014)***	.0167(.0014)***	.0108(.0014)***
1,000-9,999	.0295(.0013)***	.0283(.0013)***	.0223(.0014)***
10,000 +	.0028(.0019)*	.0011(.0019)	.0060(.0020)***
Management Experience?	-.0069(.0009)***	-.0063(.0009)***	-.0042(.0009)***
Experience Requirement:			
Up to one year	-.0179(.0026)***	-.0151(.0026)***	-.0119(.0026)***
1-3 years	.0009(.0008)	.0009(.0008)	-.0016(.0008)**
3-5 years	.0167(.0009)***	.0149(.0009)***	.0053(.0009)***
5-10 years	.0317(.0012)***	.0286(.0012)***	.0160(.0012)***
> 10 years	.0176(.0026)***	.0152(.0026)***	.0015(.0026)
Education Requirement:			
Grade 9	.1409(.0042)***	.1369(.0041)***	.0885(.0041)***
High school (12 years)	.0435(.0012)***	.0436(.0012)***	.0343(.0012)***
Post-secondary (15 years)	-.0136(.0009)***	-.0094(.0009)***	.0030(.0009)***
University (16 years)	-.0240(.0010)***	-.0177(.0010)***	-.0044(.0010)***
Master's degree (19 years)	-.0260(.0030)***	-.0160(.0030)***	-.0044(.0030)*
PhD (21 years)	-.0469(.0111)***	-.0422(.0110)***	-.0321(.0110)***

Notes: The omitted category for "number of positions" is "number of positions not stated". Education and experience requirements are relative to no stated requirement; omitted firm type and size are privately held firms and fewer than 20 employees respectively. See Table 1 for definitions of firm types. Sample size is 499,549. ***, ** and * denote significance at 1%, 5% and 10% respectively.

Table 3: Determinants of Firms' Advertised Preferences for Women

	Specification		
	Ad- or firm-level covariates only	**Ad- or firm-level and aggregate covariates**	**Ad- or firm-level covariates plus occ, ind and province fixed effects**
INDICATORS OF MARKET THICKNESS (*M*):			
Number of Positions			
1	.0285(.0008)***	.0282(.0008)***	.0191(.0008)***
2	.0067(.0011)***	.0063(.0011)***	.0097(.0010)***
3-5	-.0059(.0011)***	-.0061(.0011)***	-.0012(.0011)
6-15	-.0037(.0015)***	-.0035(.0015)***	.0025(.0014)**
16-50	.0146(.0024)***	.0150(.0024)***	.0209(.0023)***
Over 50	-.0020(.0051)	-.0013(.0051)	-.0160(.0051)***
Mean Ad Renewal Rate	---	-.0006(.0001)***	-.0000(.0001)
Provincial Unemployment Rate	---	.0006(.0001)***	---
DETERMINANTS OF (v^A - v^B) and c:			
Part Time Position?	-.0015(.0026)	.0010(.0026)	.0079(.0027)***
Firm Ownership Type:			
FDI	-.0260(.0009)***	-.0256(.0009)***	-.0221(.0009)***
Representative Office	-.0247(.0028)***	-.0243(.0028)***	-.0254(.0028)***
Joint venture	-.0170(.0009)***	-.0165(.0009)***	-.0135(.0009)***
Publicly-Traded Company	-.0098(.0036)***	-.0093(.0036)***	-.0113(.0035)***
"Corporation"	-.0136(.0011)***	-.0140(.0011)***	-.0097(.0011)***
Non-profit	-.0213(.0048)***	-.0203(.0048)***	-.0018(.0048)
State-owned enterprise (SOE)	-.0122(.0012)***	-.0118(.0012)***	-.0092(.0012)***
Public sector	-.0202(.0120)**	-.0200(.0120)**	-.0267(.0117)**
Other	-.0080(.0012)***	-.0093(.0012)***	-.0074(.0013)***
Firm Size:			
20-99	-.0004(.0011)	-.0007(.0011)	-.0011(.0011)
100-499	-.0058(.0011)***	-.0061(.0012)***	-.0080(.0012)***
500-999	-.0081(.0014)***	-.0085(.0014)***	-.0160(.0014)***
1,000-9,999	-.0175(.0014)***	-.0176(.0014)***	-.0227(.0014)***
10,000 +	-.0341(.0020)***	-.0342(.0020)***	-.0366(.0020)***
Management Experience?	-.0061(.0009)***	-.0059(.0010)***	-.0113(.0009)***
Experience Requirement:			
Up to one year	.0103(.0026)***	.0108(.0026)***	.0048(.0026)**
1-3 years	-.0145(.0008)***	-.0144(.0008)***	-.0097(.0008)***
3-5 years	-.0390(.0010)***	-.0393(.0010)***	-.0302(.0010)***
5-10 years	-.0481(.0012)***	-.0486(.0012)***	-.0396(.0012)***
> 10 years	-.0384(.0026)***	-.0389(.0026)***	-.0309(.0026)***
Education Requirement:			
Grade 9	.0454(.0042)***	.0447(.0042)***	.0365(.0042)***
High school (12 years)	.0525(.0013)***	.0525(.0013)***	.0468(.0012)***
Post-secondary (15 years)	.0100(.0009)***	.0107(.0009)***	.0149(.0009)***
University (16 years)	-.0141(.0010)***	-.0131(.0010)***	.0011(.0010)
Master 's degree (19 years)	-.0316(.0031)***	-.0299(.0031)***	-.0100(.0031)***
PhD (21 years)			

Table 4: Estimates with Firm Fixed Effects

	Preference for men	**Preference for women**
Number of Positions		
1	.0050(.0009)***	.0135(.0009)***
2	.0079(.0011)***	.0070(.0011)***
3-5	.0045(.0012)***	-.0033(.0012)***
6-15	.0079(.0016)***	-.0037(.0016)**
16-50	-.0036(.0026)*	.0043(.0026)*
Over 50	-.0123(.0053)**	-.0045(.0054)
Mean Ad Renewal Rate	-.0002(.0002)	-.0007(.0002)***
Part Time Position?	-.0153(.0030)***	.0038(.0030)
Management Experience?	-.0016(.0010)**	-.0107(.0010)***
Experience Requirement:		
Up to one year	-.0060(.0026)**	.0166(.0027)***
1-3 years	-.0023(.0009)***	-.0158(.0009)***
3-5 years	.0054(.0010)***	-.0371(.0010)***
5-10 years	.0137(.0012)***	-.0413(.0013)***
> 10 years	.0032(.0025)	-.0352(.0026)***
Education Requirement:		
Grade 9	.0641(.0042)***	.0316(.0043)***
High school (12 years)	.0173(.0014)***	.0421(.0014)***
Post-secondary (15 years)	-.0140(.0010)***	.0184(.0011)***
University (16 years)	-.0167(.0012)***	.0028(.0012)***
Master (19 years)	-.0179(.0031)***	-.0056(.0031)**
PhD (21 years)	-.0390(.0112)***	.0012(.0113)

The omitted category for “number of positions” is “number of positions not stated”. Education and experience requirements are relative to no stated requirement. Sample size is 499,549.
***, ** and * refer to significant at the 1%, 5% and 10% level respectively.

Appendix 1: Proofs

Expected Value of the Maximum in a Combined Sample:

In general, the expected value to the firm of its most-preferred worker in a pool of A applicants with productivity $u_j = v^A + e_j$ and B applicants with productivity $u_j = v^B + e_j$ is given by $V^A q^A + V^B (1- q^A)$, where V^J is the expected productivity of the best overall worker given the best overall worker is of type J, and q^A is the probability that the best overall worker turns out to be drawn from pool A. Again, using results from the MNL literature, we know that $V^A = v^A + \gamma - log(p^A)$, where:

$$p^A \quad \frac{\exp(v^A)}{A\exp(v^A) \quad B\exp(v^B)}$$

is the probability that an individual type-A applicant turns out to be the best in the entire, combined pool. Similarly, we have $V^B = v^B + \gamma - log(p^B)$, where:

$$p^B \quad \frac{\exp(v^B)}{A\exp(v^A) \quad B\exp(v^B)}.$$

Finally, the probability that the firm's preferred applicant from this combined pool is drawn from the *A*'s is just:

$$q^A \quad \frac{A\exp(v^A)}{A\exp(v^A) \quad B\exp(v^B)}.$$

Combining all the necessary expressions and simplifying, the expected productivity of the best worker from the combined pool can be written as:

$$\gamma + log(A exp(v^A) + B exp(v^B)).$$

Letting $\delta = A/(A+B) \equiv A/M$ be the fraction of *A*'s in the combined pool and defining v^C as $log[\delta \, exp(v^A) + (1- \delta) exp(v^B)]$, this becomes just:

$$\gamma + v^C + log(M). \qquad \blacksquare$$

Proof of Proposition 1

First, note that under free random disposal, firms will discard any applications in excess of N^* (in the 'combined' strategy C, we assume this leaves the firm with the same share of *A*s in its applicant pool as in the population, δ). Letting N be the number of applications *received* by the firm, it follows that, for all three recruiting strategies, profits will be increasing in N for $N < N^*$, independent of N for $N > N^*$, and (consequently) nondecreasing in N overall, as depicted in Figure A1 below.

Figure A1: Profits as a Function of Applications Received under Alternative Recruiting Strategies

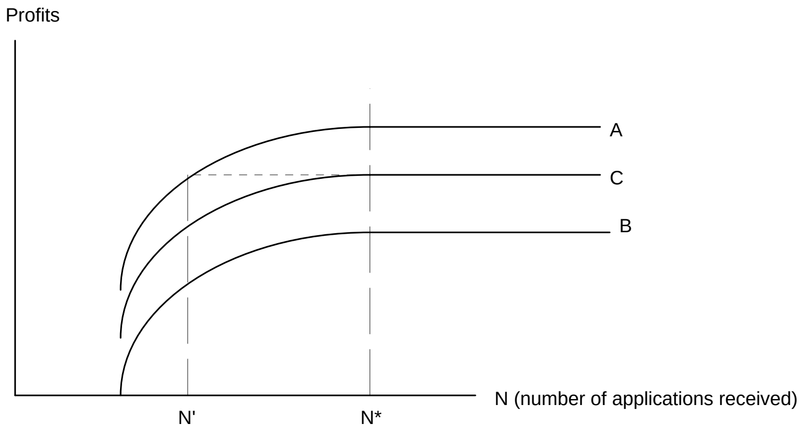

Now imagine that the firm is faced with fixed numbers of applicants under each of its possible recruiting strategies: A, B, and $M=A+B$. The fact that the three above functions are nondecreasing immediately implies that the firm will never choose recruiting strategy B for any level of market thickness M: It is always dominated by the combined strategy C, because (since $A+B= M > B$) firms receive more applications under strategy C, and because profits are higher at any given number of applications.

Thus, the only remaining comparison we need to make is between Strategies A and C. We proceed by first noting that, as M rises from its minimum possible level, it must pass through three regions:

Region 1: (both As and Bs are scarce): $\delta M < M < N^*$
Region 2: (only As are scarce): $\delta M < N^* < M$
Region 3: (neither type is scarce): $N^* < \delta M < M$

Next, we show that the profit differential between strategies A and C is nondecreasing in M. In addition the differential is strictly increasing in M in Regions 1 and 2, and does not vary with M in Region 3. To do this, we first define the functions $\Pi^K(n)$ as $v^K + \gamma + log(n) – cn$, K $\{A, B, C\}$, and their corresponding maxima, $\Pi^K{}^* = \Pi^K(N^*)$, then proceed by region:

Region 1 ($\delta M < M < N^*$). In this case the difference in profits between strategies A and C is given by:

$$\Pi^A(\delta M) - \Pi^C(M) = v^A - v^C + log(\delta) + (1-\delta)cM \,, \qquad (A1)$$

which is increasing in M.

Region 2 ($\delta M < N^* < M$). Now, the difference in profits between strategies A and C is given by:

$$\Pi^A(\delta M) - \Pi^C(N^*) = v^A + \gamma + log(\delta M) - c\delta M - \Pi^C(N^*) \tag{A2}$$

Because $\Pi^C(N^*)$ is independent of M, and because $\Pi^A(\delta M) = v^A + \gamma + log(\delta M) - c\delta M$

$\Pi^A(\delta M) - \Pi^C(N^*) = v^A + \gamma + log(\delta M) - c\delta M - v^C - \gamma - log(N^*) + cN^* = 0$

Noting that $log(\delta M) - c\delta M$ is increasing in δM in Region 2, an increase in $v^A - v^C$ means that M must fall to preserve the equality. Now, note that a small increase in c reduces $\Pi^A(\delta M)$ by δM, and reduces $\Pi^C(N^*$

Appendix 2: Data

As noted, our overall sample consists of all job ads which appeared on Zhaopin.com between May 16 and July 12, 2008 inclusive. At the end of each day, our program automatically searches for job ads that were posted on Zhaopin that day. The program starts at 11:30pm sharp each day for consistency. On the first day of data collection, all ads that were posted that day were kept. On subsequent days, all ads posted that day are compared with the master list of previously-posted jobs; since many such jobs are just renewals that are re-posted (employers can re-post and existing ad; this entails a small marginal financial cost but does require action on the employer's part), we do not download these refreshed jobs but maintain a count of the number of renewals that occur during this time period. A similar procedure was applied to the list of firms. 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 regression analysis is restricted to the sample of jobs for which we have matching firm information. The matching rate varies somewhat across specifications but was about 80.2%

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 also visually check all the job ads that mentioned gender in a way that did not match these tables. Only about 100 jobs out of our entire sample fell into this category.

For age variables, we search for "sui" (year of age); for hukou variables, we search for "hu" (part of the hukou term). Our approach could miss jobs that ask for age only using numbers "25-35", or jobs that vaguely just ask for individuals to come from a certain hometown (which could potentially have hukou implications). Therefore, the variables that we use here should be interpreted as having very explicit requirements for gender, age and hukou.

Occupation and industry categories are those supplied by Zhaopin.com (firms choose from a list on the website when submitting their ad). Note that our occupation and industry dummy variables are not mutually exclusive, as firms are allowed to check multiple categories. (This is the case both when a single ad is for multiple vacancies and when it is not).

Finally, our data on job ads was merged with a number of province-level characteristics, taken from 2000 Census and 2001 National Census of Basic Units of China accessed on November 2 2008 through http://www.acmr.com.cn's Support System for China Statistics Application.

To construct our "inflow sample" of job ads, we first examined the empirical distribution of dates that an ad first appears on the job site during our sampling period. As Figure A1 shows, this distribution has a large spike on the first day we collected ads, then declines rapidly, reflecting the fact that most jobs "posted today" after our first day of data collection were in fact just repostings or renewals of jobs that had been posted earlier. After about a month, however, the empirical distribution of new jobs (that we have not seen before on the site) becomes quite constant. This suggests the sample of ads newly appearing on the site after that time are essentially all new; we thus define our "inflow" sample as all ads which appear in our data for the first time after June 1, 2008. (We also experimented with an "outflow" sample, consisting of vacancies that were last observed before June 30, i.e. almost 6 weeks before the end of data collection. The results were similar.)

Figure A1: Flows of ads by date first observed

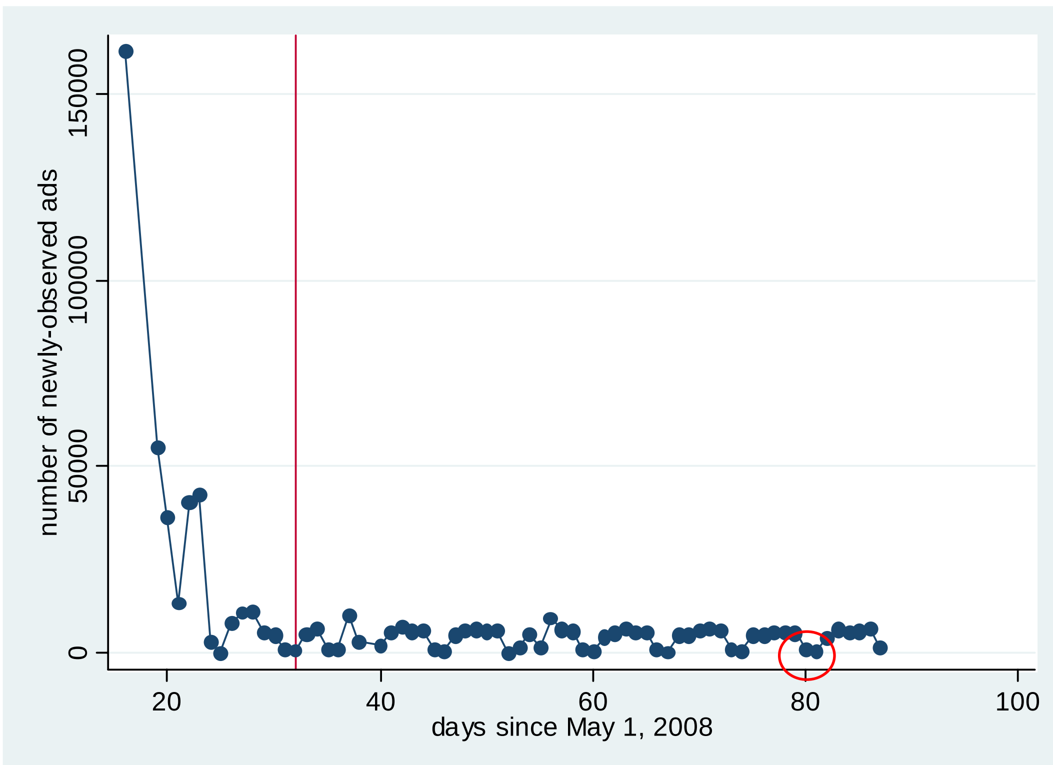

Notes:

-vertical line indicates June 1, the beginning of our "inflow sample"
-circled points show a weekend

Appendix 3: Supplementary Tables

Table A1: Determinants of Firms' Preferences for Men (Entry flow sample only)

	Specification		
	Ad- or firm-level covariates only	**Ad- or firm-level and aggregate covariates**	**Ad- or firm-level covariates plus occ, ind and province**
[redacted]			
State-owned enterprise (SOE)	.0057(.0020)***	.0079(.0020)***	.0068(.0020)***
Public sector	.0834(.0184)***	.0853(.0183)***	

Table A2: Determinants of Firms' Preferences for Women: Entry flow sample only

	Specification		
	Ad- or firm-level covariates only	**Ad- or firm-level and aggregate covariates**	**Ad- or firm-level covariates plus occ, ind and province fixed effects**
INDICATORS OF MARKET THICKNESS (*M*):			
Number of Positions			
1	.0293(.0013)***	.0290(.0013)***	.0191(.0013)***
2	.0074(.0017)***	.0071(.0017)***	.0111(.0017)***
3-5	-.0058(.0018)***	-.0060(.0018)***	.0016(.0017)
6-15	-.0078(.0024)***	-.0074(.0024)***	-.0018(.0024)
16-50	.0152(.0040)***	.0157(.0040)***	.0208(.0039)***
Over 50	-.0134(.0131)	-.0135(.0131)	-.0139(.0128)
Mean Ad Renewal Rate	---	-.0007(.0001)***	-.0004(.0002)**
Provincial Unemployment Rate	---	.0011(.0002)***	---
DETERMINANTS OF (v^A - v^B) AND *c*:			
Part Time Position?	-.0148(.0047)***	-.0127(.0047)***	-.0033(.0049)
Firm Ownership Type:			
FDI	-.0345(.0015)***	-.0343(.0015)***	-.0297(.0015)***
Repres. Office	-.0174(.0051)***	-.0171(.0051)***	-.0131(.0050)***
Joint venture	-.0240(.0017)***	-.0237(.0017)***	-.0182(.0017)***
Publicly-Traded Company	-.0152(.0058)***	-.0144(.0058)***	-.0103(.0057)**
"Corporation"	-.0197(.0016)***	-.0196(.0016)***	-.0161(.0016)***
Non-profit	-.0195(.0087)**	-.0181(.0087)**	.0029(.0086)
State-owned enterprise (SOE)	-.0217(.0021)***	-.0215(.0021)***	-.0188(.0020)***
Public sector	-.0197(.0190)	-.0192(.0190)	-.0293(.0186)*
Other	-.0136(.0019)***	-.0153(.0020)***	-.0097(.0020)***
Firm Size:			
20-99	-.0003(.0017)	-.0008(.0017)	-.0022(.0017)*
100-499	-.0054(.0018)***	-.0058(.0018)***	-.0097(.0018)***
500-999	-.0029(.0023)	-.0036(.0023)*	-.0095(.0024)***
1,000-9,999	.0023(.0022)	.0022(.0022)	-.0044(.0023)**
10,000 +	-.0252(.0035)***	-.0254(.0035)***	-.0240(.0036)***
Management Experience?	-.0139(.0016)***	-.0135(.0016)***	-.0173(.0016)***
Experience Requirement:			
Up to one year	-.0106(.0040)***	-.0095(.0040)***	-.0123(.0039)***
1-3 years	-.0187(.0013)***	-.0186(.0013)***	-.0141(.0013)***
3-5 years	-.0477(.0016)***	-.0480(.0016)***	-.0367(.0016)***
5-10 years	-.0526(.0021)***	-.0532(.0021)***	-.0413(.0021)***
> 10 years	-.0182(.0049)***	-.0186(.0049)***	-.0082(.0049)**
Education Requirement:			
Grade 9	.0466(.0060)***	.0463(.0060)***	.0398(.0059)***
High school (12 years)	.0545(.0019)***	.0547(.0019)***	.0525(.0019)***
Post-secondary (15 years)	.0125(.0014)***	.0134(.0014)***	.0198(.0014)***
University (16 years)	-.0136(.0016)***	-.0126(.0016)***	.0032(.0017)**
Master's degree (19 years)	-.0340(.0061)***	-.0322(.0061)***	-.0106(.0060)**
PhD (21 years)	-.0342(.0228)*	-.0348(.0228)*	-.0214(.0224)

See previous tables for variable definitions. Sample size is 206,203.
***, ** and * refer to significant at the 1%, 5% and 10% level respectively.