NBER WORKING PAPER SERIES

GENDER DISCRIMINATION IN JOB ADS: THEORY AND EVIDENCE

Peter J. Kuhn
Kailing Shen

Working Paper 17453
http://www.nber.org/papers/w17453

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2011

The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research. 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. Catherine Weinberger, and seminar participants at Xiamen University, Tsinghua University, Renmin University China, ZEW (Centre for European Economic Research), University of Aarhus, the Chinese University of Hong Kong, the Trans-Pacific Labor Seminar, the all-UC labor economics conference, UC Denver, University of Heidelberg, IZA, Tilburg University, University of Amsterdam, Maastricht University and UCSB supplied helpful comments. Additional materials are available at: http://www.econ.ucsb.edu/~pjkuhn/Data/GenderDiscrim/GenderDiscrimIndex.htm

Gender Discrimination in Job Ads: Theory and Evidence
Peter J. Kuhn and Kailing Shen
NBER Working Paper No. 17453
September 2011
JEL No. J16,J63,J71

ABSTRACT

We study firms' advertised gender preferences in a population of ads on a Chinese internet job board, and interpret these patterns using a simple employer search model. The model allows us to distinguish firms' underlying gender preferences from firms' propensities to restrict their search to their preferred gender. The model also predicts that higher job skill requirements should reduce the tendency to gender-target a job ad; this is strongly confirmed in our data, and suggests that rising skill demands may be a potent deterrent to explicit discrimination of the type we document here. We also find that firms' underlying gender preferences are highly job-specific, with many firms requesting men for some jobs and women for others, and with one third of the variation in gender preferences within firm*occupation cells.

Peter J. Kuhn
Department of Economics
University of California, Santa Barbara
2127 North Hall
Santa Barbara, CA 93106
and NBER
pjkuhn@econ.ucsb.edu

Kailing Shen
Wang Yanan Institute for Studies in Economics
(WISE)
Xiamen University
China 361005
kailing.shen@gmail.com

1. Introduction

Not that long ago, it was commonplace for U.S. employers to explicitly discriminate on the basis of age, race and sex in their job advertisements.[1] While these practices are now illegal in most developed countries, the vast majority of the world's population still works in labor markets where such ads are permitted. In this paper, we use this opportunity to study gender discrimination in a large sample of job advertisements taken from an internet job board in China, with a view to understanding the practice and shedding light on the underlying sources of gender differentials in labor markets. To guide our analysis, we develop a simple model of employer search from two populations in the presence of application processing costs.

Applied to the context of gender, our model unsurprisingly predicts that increases in a firm's preferences for men in a job should raise the share of ads for that job that explicitly request male applicants, and reduce the share that request women. Of greater interest, the model also predicts that a set of four 'search-related' factors act in a different way: Increases in application processing costs or in the expected number of applications should raise the likelihood that firms will impose gender restrictions against both men *and* women, while increases in the idiosyncratic variance of applicant productivity and in the job's skill level have the opposite effect (reducing the frequency of discrimination in both directions). We test the latter prediction in our data, and also use the model to interpret the cross-sectional relationship between job and firm characteristics on the one hand, and advertised discrimination on the other.

We find, first of all, that explicit, advertised gender discrimination is commonplace in our data. In particular, over one third of the firms that advertised on the board during our twenty-week sample period placed at least one ad stipulating a preferred gender. Thus, when it is legal to express such preferences, a large share of employers does so. Second, as predicted by our model, we find strong evidence that higher skill requirements –whether measured by the job's education level, the required level of experience, or the offered wage—are associated with lower levels of advertised gender discrimination, whether in favor of men or women. While previous authors (Becker 1957; Black and Strachan 2001) have argued that product market competition may be a

[1] See Darity and Mason (1998) for examples of such ads in the U.S.

potent deterrent to discrimination, our results suggest that rising skill demands may play a similar role, both in explaining the cross-sectional incidence of discrimination across jobs, and in understanding long-term trends and international differences in the extent to which explicit discrimination is accepted and practiced.

Using our model, we are also able to estimate the effects of a wide variety of firm and job characteristics on firms' underlying gender preferences. For example, large firms, state-owned firms, and firms seeking workers with high levels of experience tend to prefer men. Foreign-owned firms, and firms hiring for certain customer-contact occupations tend to prefer women. At the same time, we detect no association between a job's offered wage and firms' underlying relative preferences for men versus women in that job. Thus, *all* of the negative association between a job's offered wage and the probability that women are invited to apply in our data is attributed to search-related factors that vary with job skill levels.

Finally, we find that firms' gender preferences are highly job-specific. For example, over a quarter of the firms that advertised a gender preference requested men for some positions and women for others. More specifically, while firm fixed effects explain about a third of the variation in gender preferences in our data, neither firm-wide preferences towards one gender, nor inter-firm differentials in occupation mix, are very powerful explanations of the pattern of advertised gender preferences. Instead, ads for the same occupation are often targetted at men in some firms and at women in others, and one third of the variation in firms' gender preferences is within firm*occupation cells. Overall, rather than, for example, restricting women's access to highly skilled jobs, Chinese firms mostly use targetted job ads to divide their pool of less skilled positions into 'male' versus 'female' positions.

2. Related Literature

To our knowledge, Darity and Mason (1998) are the only economists who have examined discrimination in U.S. job ads; they reproduce examples of ads from 1960 newspapers, 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. In a sample of several hundred firms, she finds that a majority reserved some positions for men, and others for women. Lawler and

Bae (1998) is the only article we know of that studies discrimination in a recent sample of job ads. Their focus is on the impact of a multinational firm's home country culture on its stated gender preferences in a sample of 902 ads placed in an English-language newspaper in Thailand.[2] Finally, we are aware of two papers --Barron, Bishop and Dunkelberg (1985) and Van Ours and Ridder (1991)— that study employers' advertised hiring restrictions (in their case, minimum education requirements) empirically using microdata on job ads. However, unlike us, both of these papers treat advertised requirements as exogenous vacancy characteristics, rather than as a choice variable for the employer.

Given the lack of previous research on advertised discrimination, it is perhaps useful to contrast ad-based indicators of discrimination with more well-known measures, such as those generated by audit-type studies (e.g. Bertrand and Mullainathan 2004). Bertrand and Mullainathan submitted resumes to U.S. employers *in response to* newspaper job ads, with the respondent's apparent race randomly assigned via race-specific first names. They found large differentials in callback rates for identical black and white resumes. Full audit studies, such as Neumark (1996), carry this procedure one step further and send matched, trained actors to interview for jobs in response to the callbacks.

Ad-based indicators of discrimination differ from audit-type indicators in two ways that have opposite implications for how much discrimination we should expect to observe. First, because ads are formulated before resumes arrive, ad-based measures do not condition on the information that appears in a worker's resume.[3] For this reason, if the employer expects to receive lower quality resumes from the 'disfavored' group, firms might choose to engage in discriminatory job advertising even when audit-type studies would show little or no discrimination. Second, however, according to our model, discrimination in job ads should only be observed when an employer's preferences against (or prior beliefs about) a particular group are intense enough to exceed a strictly

[2] Banerjee et al. (2009) study caste and other preferences in a small sample of marriage ads in India.
[3] Indeed, the fact that no resumes are submitted (and no actors need to be trained) can be seen as an advantage of the ad-based approach: their comparability, realism and representativeness is not at issue. Another key advantage is cost: our data is collected costlessly by a web crawler; thus our sample consists of over one million observations, compared to about 5,000 in Bertrand-Mullainathan and much smaller numbers in most audit studies. The larger sample size, in turn, lets us go beyond the study of a few specific occupations to paint a broad statistical portrait of explicit gender-targetting in this labor market, across occupations, industries, and firm types.

positive threshold. This is because targetting an ad causes a discrete drop in the number of applications received, by discouraging an entire group from applying. This feature of ad-based discrimination suggests, in contrast, that we should see it less frequently than in audit studies where even small employer preferences are in principle detectable.

A third and final difference between ad-based and 'traditional' measures is that advertised discrimination necessarily involves a conscious decision by the employer to invite only one group to apply. In contrast, audit studies should detect both the conscious choices and unconscious biases of employers.[4] In sum, we should expect firms to post discriminatory ads when they *consciously* expect *large* differentials in both the observed (in a resume or interview) *and* unobserved (productivity and taste-related) characteristics of two groups; these conditions differ from the conditions in which we would be more likely to observe discrimination using more familiar methods, such as audit studies or wage regressions.

Finally, we note that neither ad-based nor audit-based measures of discrimination on their own can distinguish whether discrimination is statistical (i.e. based on actual or perceived between-group differences in average productivity that are unobserved by the investigator) or taste-based. Distinguishing these sources of discrimination requires additional evidence, for example on the amount of customer contact in the job, or the amount of product market competition faced by the employer. We examine such evidence here, but the fact that our discrimination measure is derived from job ads does not in itself confer any advantages or disadvantages in distinguishing among these motivations for discriminating.

Theoretically, our model draws on Becker's (1957) seminal model of labor market equilibrium in the presence of discrimination, and on its recent elaboration by Charles and Guryan (2008). We generalize both these models by introducing search frictions and application processing costs: as in Altonji and Pierret (2001), firms in our model can use gender or other demographics as coarse screens for applicant ability; in our case we formally model firms' decisions regarding whether or not to use such screens. Becker's model –where essentially all jobs are segregated in equilibrium-- is a

[4] See Crosby et al. (1980) for a review of studies of unconscious bias.

special case of ours where the idiosyncratic variation in worker quality approaches zero, or where application processing costs approach infinity.[5]

3. A Simple Model of Discrimination in Job Ads

Consider a firm soliciting applications for a single vacant position; applications can come from two distinct groups, labelled *M* and *F*. 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, the gender difference in baseline net value, v^M - v^F, includes not only between-group differences in revenue productivity in the job, but also differences in employer tastes, and in expected wage costs between the groups. We explore these differential sources of expected net value theoretically at the end of the current section, and empirically in Sections 5 and 6. For now we simply work with v^M - v^F as a summary index of all these factors for a particular job, and use employers' "preferences towards men" or "tastes for men" as a shorthand for v^M - v^F.

Simple, closed-form solutions for the model are available when we assume a convenient functional form for $F(\varepsilon_j)$. Specifically, let $F(\varepsilon_j)$ be type-I extreme value, 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 (≈.577). The firm is assumed to choose the individual worker with the highest total value, U_j, from its pool of applicants. The only nontrivial choice facing the firm is which groups (*M*, *F*, or both) to invite to apply; this choice takes into account that it costs the firm a constant amount, *c*, to process each application that arrives, thereby learning its ε_j.[6]

[5] Our model also relates to a literature on optimal search and recruiting that has been inspired by Mortensen and Pissaridies (1994), among others. In most of these models, employers' strategy space is highly restricted: for example in many cases firms' only choices are whether to enter the market and what wage to post (e.g. Moen, 1997; Menzio 2007). Thus, while a number of authors have studied discrimination in an equilibrium search context (e.g. Rosen 1997; Lang, Manove and Dickens 2005), to our knowledge the theoretical search literature has not yet considered the possible optimality of advertised hiring restrictions, which invite some (but not all) worker types to apply.
[6] Note that application processing costs are an essential component of our model: after all, in the absence of processing costs, firms could costlessly duplicate the effects of any advertised job requirement by soliciting applications from everyone, then just discarding the applications from the groups that are not wanted. Application processing costs differ from the usual costs assumed in search models, such as a fixed

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 value of U_j drawn from the sample of applicants defined by D^M and D^F. M and F denote the numbers of applications that would arrive from the two groups, if invited.

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

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:

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

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

Proof: See Appendix 1.

To keep the notation simple in what follows, we shall focus on the case of an equal number of applications from each group, i.e. $\delta=.5$[7] In that case, (4) can be written:

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

where $N=.5C$ is the (common) number of applications expected from each group.

We next define $z \equiv (\mu^M - \mu^F)/\beta$ as the standardized gap in expected net value between the groups.[8] Overall, the firm's optimal recruiting policy is then described by:

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

Solicit men only if $z > z*$,

cost of opening a vacancy and the opportunity cost of keeping the job vacant; they will also be affected by a less familiar set of factors, including the firm's information-processing technology and the complexity of the methods used to evaluate job candidates.

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

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

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

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

where:

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

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

When $cN/\theta\beta \le ln(2)$, 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. We refer to this as the "high variance" case because it corresponds to high levels of variance in idiosyncratic worker quality, i.e. high β; however it also corresponds to higher levels of skill, lower levels of hiring costs, and smaller expected applicant pools. When $cN/\theta\beta > ln(2)$ (the "low variance" case), firms always restrict their ads: to men only when z>0 and to women otherwise.

Overall, Proposition 1 shows that the factors influencing a firm's optimal recruiting strategy fall naturally into two categories: Factors that raise the index z, which is the expected productivity, cost and taste advantage of men for this job, raise the 'likelihood' that firms will invite only men to apply, and reduce the likelihood that firms will invite only women.[9] As a shorthand, we will refer to these factors as *relative preferences*. The remaining parameters in the model, c, N, θ, and β, operate only on the thresholds, z^* and $-z^*$. We refer to these factors as *search-related*.

To understand how the search-related factors work, consider first the high-variance case, where there exist values of z for which it is optimal not to gender-target the job ad. Here, the 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. Specifically, the model predicts that increases in β (i.e. greater idiosyncratic variance of applicant productivity), and increases in θ (the job's skill level) reduce the likelihood that firms will advertise gender restrictions *in either direction*. The intuition is that higher variance raises the option value of searching from a larger pool (i.e. raises the chance the best candidate will come from the group with the lower mean), and higher skill levels raise the importance to the firm of identifying the best candidate. In contrast, increases in c and N move $-z^*$ and z^*

[9] We formalize this likelihood in Section 5, where we posit a distribution of expected net values across jobs that are offered by employers, and estimate a maximum likelihood model of the incidence of advertised gender restrictions across jobs.

closer together, making advertised gender restrictions more likely. This is because both *c* and *N* directly raise the cost of doubling the applicant pool from *N* to 2*N* to include both men and women.

Turning now to the low-variance case, in this region of the parameter space firms will always direct their ads only to the group they prefer. To see how our model specializes to Becker's in this case, consider a labor market with many employers, indexed by *i*, who can hire either men or women for the same task. 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:

(6) $$v_i^M - v_i^F = (q_i^M - q_i^F) + (t_i^M - t_i^F) - (w^M - w^F)$$

If, for simplicity, 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 variance' 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. Our high-variance case thus generalizes Becker's analysis to a case where search frictions cause some job ads *not* to be targetted, as is in fact the case for the majority of job ads in China.

Since gender restrictions occur in only about one in ten ads in our data, we assume in our empirical work that the 'high variance' case applies to all employers (i.e. that there exists *some* expected gender value differential, perhaps infinitesimal, for which every firm would prefer not to issue a discriminatory ad). In our empirical work, we will use this model to interpret patterns of discriminatory advertising. We also test a key prediction of the model: that the incidence of both 'male-only' and 'female-only' ads should decline as the job's skill level rises.

4. Data and Descriptive Statistics

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.[10] Procedures for downloading the data and defining variables are discussed in Appendix 2. Like all samples of job ads, ads on Zhaopin.com will not be representative of all jobs in China; 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. Since we might expect firms in expanding and high-turnover industries to be less selective about their employees' gender, and since we show below that advertised discrimination is less common in skilled occupations, it seems likely that our data underestimates the overall extent of gender-targetting job advertising in China.

Another sampling issue arises from the fact that some of the ads in our data are for multiple vacancies. While we treat the ad as the unit of analysis thoughout this paper, we address this issue by controlling for the number of vacancies the ad advertises.[11] A related issue is that ads for multiple vacancies can specify up to three occupations. Since the share of vacancies corresponding to the individual occupations is generally unknown in these cases, we restrict our sample to ads for a single occupation; this reduces our sample by about 20 percent.[12] Finally, we note that, by construction, our data comprise the stock of unfilled ads rather the flow of new vacancies. This implies that long vacancy spells are overrepresented in our data, which would affect our estimates if there is parameter heterogeneity in the determinants of ad content that is correlated with vacancy durations.[13] To address this concern, we replicated our estimates 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.

Descriptive statistics for our data are provided in Table 1. All told, we study a total of 1,057,538 job ads. Our principal measure of a job's skill level is the education requirement listed by the employer in the ad; by this metric a typical job on Zhaopin is quite highly skilled: about 87 percent of ads (1 - 136595/1057538) require at least some

[10] Note that firms frequently re-post the same ad; our sample treats all such renewals as the same ad. Our choice of Zhaopin is largely for the technical reason that its site structure allowed us to easily and accurately identify such renewals.
[11] We also estimated some specifications where we weight each ad by the number of vacancies it represents; the results are unchanged.
[12] Our results were 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 all the occupations they listed.
[13] See Bergeron et al. (2008) for a recent discussion of the effects of length-biased sampling.

post-secondary education.[14] Eighty-one percent require some experience, with the modal experience requirement between one and three years. A little more than one in ten ads expressed a gender preference; this was evenly split between men and women.[15] As suggested by our model, however, employers' propensity to express a gender preference is strongly, and negatively related to the job's educational requirements: almost a quarter of jobs requiring high school education or less were explicitly 'gendered'; this fraction falls to a little over six percent for jobs requiring a university education.

About half of the ads in our sample specified the number of vacancies that were available. Half of these, in turn, were for a single position. That said, a significant share of the ads were for large numbers of job openings. More than half of the ads were placed by privately owned, Chinese firms (this includes privately held companies, publicly-traded companies and former State-Owned Enterprises where a majority of shares are still owned by the state). Another 36 percent of ads were from employers with some foreign connection. Most of these were Foreign Direct Investment (FDI) and joint ventures, though a small number of representative offices are also included. A further seven percent of ads were for jobs in State-Owned Enterprises. We also observed some ads from non-profit employers (e.g. in education or health care) or (local, provincial or central) public service; while these are a very small share of the total, our large sample size allows us to study them as well.

A complementary picture of the frequency of discriminatory job ads in China emerges when we organize our data by firms rather than ads. Overall, 74,202 distinct firms placed ads on Zhaopin during our sampling period; thus the average number of ads per firm was 1,057,538/74,202 = 14.25. Characteristics of these 74,202 firms' hiring policies are summarized in Table 2. According to Table 2, 19.9 percent of the firms in our data placed at least one ad that invited only 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, and with the number of distinct occupations for which it advertised on Zhaopin during our sample period. Thus, for example, among firms that placed more

[14] By far the most common occupation (of the 40 categories used by Zhaopin on its site) is sales, at about 22 percent of the ads, with IT second at about 13 percent. The top five industries were construction, consulting, IT, marketing, and trade. 24 and 17 percent of the ads were for jobs in Beijing and Shanghai respectively, but all Chinese provinces are represented in our data.

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

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.

In sum, Tables 1 and 2 illustrate three key features of advertised gender discrimination in China. First, a large share of employers engage in hiring practices that would be considered an illegal form of discrimination in the United States. Although we do not address the larger question of whether Chinese or American laws are more appropriate, we do discuss some consequences of these legislative differences in the conclusion to this paper. Second, advertised gender preferences vary significantly across jobs within the same firm. Finally, advertised gender preferences in favor of both men and women decline markedly as the job's advertised education requirement rises. We explore all of these, and many other features of discrimination in job ads in more detail in the following section.

5. Empirical Determinants of Gender Discrimination in Job Ads

To bring Section 2's model of advertised discrimination to the data, let the employer's net relative valuation of men in the position described in ad i be given by

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

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 targetted at men if $v_i > z_i^* - \mathbf{x_i}\,\mathbf{b}$, targetted 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 cdf $F(v_i)$. The likelihood of observing each of the three possible ad types can then be written:

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

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

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

If $z_i^* = z^*$ (a constant for all observations) and F is $N(0, \sigma_v^2)$, (8) describes an ordered probit model, which can be used to estimate the parameter vector **b** up to a constant of proportionality (σ_v). An important feature of our theoretical model, however, is that a subset of the ad's observed characteristics –in particular the indicators of its skill requirements, but also any observable correlates of application processing costs, expected numbers of applications, and idiosyncratic worker quality—are expected to act on the two thresholds in this three-outcome ordered probit in a somewhat unusual way. Specifically, high skill requirements are predicted to move the two thresholds apart by equal amounts; this is reflected in (8) by the fact that the thresholds equal $-z^*$ and z^* respectively.

In order to model the dependence of these thresholds on ads' observable characteristics, we therefore assume in addition that:

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

which implicitly assumes that any variable that might affect firms' relative valuation of men versus women (z_i) can at least potentially affect z_i^* as well.[16]

Taken together, (8) and (9) plus a distributional assumption for F comprise a simple model of our data 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. the length of the interval of z's in which the firm chooses not to target its ads. The latter is summarized by the scalar z^*. Importantly, the effects of any given observable on z^* are identified even if we believe that observable also affects a firm's 'tastes towards men' (z).[17] 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. It also allows us to identify the effects of observable covariates on the firm's relative preferences for men while adjusting for differences between ads in search-related factors that would otherwise contaminate our estimates by affecting the firm's overall preferences to search narrowly versus broadly.

Finally, we note that (9) is formulated such that a positive d_k coefficient implies that covariate x_k *reduces* z^*, thus making discrimination of both types more common. This allows us to interpret the **b** parameters, loosely, as the effects of the covariates on firms' relative preferences towards men (and away from women), and the **d** parameters

[16] Note also that, according to our theoretical model, $z_i^* = -\ln[\exp(Q_i)-1]$, where $Q_i = c_i N_i / \theta_i \beta_i$; therefore (by re-arranging) the estimated ratio Q_i can be recovered from our parameter estimates for each ad via the expression $Q_i = \log\{\exp[-\exp(-\mathbf{x_i d})]+1\}$.

[17] Intuitively, as we show below, this is because we have data on two distinct outcomes (P^M and P^F); **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 ($P^M - P^F$). A critical assumption for identification, however, is that the variance of v_i not depend on the parameters, $\mathbf{x_i}$. We discuss how the interpretation of our results changes when σ_v is not constant in Section 6(c) below.

as firm's tendencies to gender-target their ads in general (i.e. as the effects of the covariates on *minus* z^*).

Maximum likelihood estimates of the parameter vectors **b** and **d** in (8) and (9) using the standard normal distribution for F are presented in columns (1) and (2) of Table 3. In addition to the variables shown, these estimates also control for the number of vacancies advertised, part-time jobs, and period fixed effects. Net of all the controls, column 1 shows that firms' underlying preferences for men (z) vary in a nonmonotonic way with the job's education requirements, at first falling, then rising. In contrast, as predicted by our model, the estimated effects of the propensity to discriminate in general are negative and monotonic, even while controlling for the observable ad characteristics in the Table 3 regressions.

According to Table 3, experience requirements have a similar qualitative effect on z^* as do education requirements: higher experience requirements are associated with broader search strategies and less total discrimination. To the extent that experience requirements also indicate a higher overall skill level for the job, this also supports our model's predictions. Interestingly, however, experience requirements also have a highly significant, positive association with firms' preferences towards men. In other words, we find that Chinese employers act as if men were better suited than women to jobs requiring high levels of experience.

Table 3 also shows that employers' relative preferences towards men, z, are higher in bigger firms, state-owned enterprises, and in government jobs, and lower in firms with some foreign ownership (all relative to Chinese private sector firms). Turning to firms' overall propensity to to engage in advertised discrimination, this is lower in larger firms and in state-owned enterprises, and especially in foreign-owned firms. The opposite is the case for government jobs in China. We comment more on possible explanations for these patterns later in this section.

Our goal in the remainder of this section is to assess whether the above patterns, especially those involving skill requirements, are robust to a much more detailed set of controls for job characteristics. Such characteristics include occupation and industry fixed effects that would capture unobserved correlates of both skill requirements and employers' gender preferences. Introducing large numbers of fixed effects is however not computationally tractable in maximum likelihood models such as the one in columns

(1) and (2) of Table 3. For that reason we next develop a linear approximation to our model and illustrate its properties in the remainder of Table 3.

To that end, Appendix 3 shows that, up to a factor of proportionality and under certain conditions, the parameter d_k is equal to x_k's average marginal effect on the probability that a firm discriminates *in some direction* in its ad. This marginal effect, in turn, can be estimated by an OLS regression of $P^M + P^F$ on **x**, where P^K is an indicator for whether the ad states a preference for gender K. Under the same conditions, b_k is identified by an OLS regression of P^M - P^F on the covariates. Further, the above conditions (specifically, that P^M and P^F have the same mean) are approximately satisfied in our data.

The remaining columns of Table 3 examine the accuracy of this approximation in our data, by comparing the marginal effects estimated by ML versus OLS approaches. Columns (3) and (4) compute the estimated marginal effects $\partial(P^M + P^F)/ \partial x_k$ and $\partial(P^M - P^F)/ \partial x_k$ from our ML estimates of **b** and **d** for an ad with mean characteristics ($\mathbf{x} = \bar{\mathbf{x}}$). Columns (5) and (6) compute the same marginal effects from a linear probability model regression. Clearly, these estimated marginal effects (and their standard errors) are very similar. In the remainder of the paper we will therefore estimate the determinants of $z*$ and z using OLS regressions of $P^M + P^F$ and $P^M - P^F$ on **x** respectively.

Table 4 presents results of OLS regressions where the dependent variable is $P^M + P^F$, where P^M (P^F) is a (0,1) indicator variable for whether the ad is targetted at men (women).[18] As argued, up to a constant, coefficients in these regressions are estimates of each variable's effects on $-z*$, with positive coefficients moving the thresholds $-z*$ and $z*$ closer together, thus raising the frequency of advertised discrimination in both directions. Column 1 simply adds province fixed effects to the regression reported in column (8) of Table 3; the results change very little. Moving across the columns, we then add occupation and industry fixed effects(column 2), fixed effects for occupation* industry interactions (column 3), for occupation*industry*province interactions (column 4); for individual firms and occupations (column 5), and finally in column 6 for a full set of 260,228 firm*occupation interactions. As already noted, a key motivation for adding these fixed effects is to ascertain whether firms' tendency to refrain from gender

[18] Coefficients of the 'underlying' regressions for P^M and P^F separately can be computed by simple addition or subtraction of the coefficients in Tables 4 and 5. Tables (including standard errors) are also available from the authors.

discrimination when jobs require a higher skill level is associated with the skill requirement itself, rather than other features of the type of work that is required, the industry, or even the specific firm doing the hiring.

While the magnitudes of the estimated effects vary somewhat across specifications, Table 4 clearly shows that the tendency to refrain from gender discrimination as skill requirements rise is a highly robust feature of our data. Even when we compare the same firm hiring for the same occupation at different times (i.e in the column 6 specification), that firm is less likely to stipulate a gender preference when the job's required skill level is high than when it is lower. The effects also remain large in magnitude: raising education requirements from high school or less to some postsecondary reduces the probability that a firm will engage in some form of gender discrimination by 6 percentage points, relative to a mean of about 10 percentage points.

As in columns (3) and (5) of Table 3, all the columns of Table 4 show that, like education requirements, higher experience requirements strongly reduce gender-targetting in job ads (though the effect is not monotonic at the highest experience levels in all specifications). Once again, these estimated effects are consistent with our predictions regarding skill requirements and are large in magnitude: raising experience requirements from less than one year to 3-5 years reduces the probability that a firm will engage in some form of gender discrimination by about 3 percentage points in all specifications.

While it is not possible to estimate the effects of firm characteristics in the presence of firm fixed effects, columns (1)-(4) of Table 4 show that the estimated effects of firm size and type on advertised discrimination are highly robust to specification. For example, larger firms are less likely to restrict their job ads to a particular gender, but the effect is small. Foreign ownership reduces the chance that firms will engage in some form of advertised gender discrimination by about 3 percentage points in all specifications. Since this effect is robust to industry, occupation and other controls, it seems likely to reflect a pure effect of those international connections, such as international differences in corporate culture, or the influence of laws in the home country on a firm's operations in China. The only other employer type that differs significantly from the base category (Chinese private sector firms) is the public sector, which is much more likely to express a gender preference. These large effects however also come with large standard errors, since our sample of public sector ads is very small.

Table 5 presents results of OLS regressions where the dependent variable is P^M - P^F. As argued above, up to a scale factor these coefficients give the effect of observable characteristics on an employer's net relative valuation of men (relative to women) for the job on offer (*z*). Importantly, under the assumptions of our model, these estimates are purged of the effects of observables on the propensity to search broadly (without restrictions) or narrowly that would contaminate other estimates, especially if the same observables affect search-related factors (such as the number of expected applications and application processing costs) as affect employers' relative valuations.

In sharp contrast to Table 5, a job's education requirements have mostly weak and insignificant effects on whether employers prefer men versus women to fill it. Thus, the strong effects of education on the tendency to discriminate is almost entirely attributable to its effects on thresholds, which were predicted by our model. In an important sense, then, Chinese employers are not using discriminatory job ads to 'keep women out of' jobs requiring high levels of education. The same, however cannot be said about experience: employers are much more likely to prefer men to women in jobs with higher experience requirements than in other jobs.[19] Large employers also appear to have greater preferences in the direction of men.

Table 5 also shows a negative effect of foreign ownership on firms' preferences towards men, though this effect (of less than a percentage point) is much smaller than their negative effect on the propensity to discriminate in general. SOEs and governments, on the other hand, seem to have a stronger preference in favor of men than other firms (the government effect very large but with a correspondingly large standard error). Given that these two employer types face less product-market competition than others in our sample, this finding is consistent with the prediction of taste-based discrimination models, that employers will be more able to indulge discriminatory tastes when competition is low.[20]

---

[19] It may be interesting to note that these effects would be masked in a simple regression of P^M on experience requirements: those coefficients are essentially zero in the column 6 specification because the effect of experience on preferences is totally masked by the negative effect of experience on thresholds.
[20] 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, SOEs and the public sector are consistent with Zhang, Han, Liu and Zhao (2008), who find that the share of the unadjusted

It may also be of interest to examine the estimated occupation fixed effects in Tables 4 and 5 for clues regarding *when* firms want men versus women in a job. These fixed effects, shown in Figure 1, are derived from regressions identical to column 2 in Tables 4 and 5 with one exception: education was removed from the list of controls, in order to illustrate its effects in the Figure itself. Occupations in the Figure are divided into two groups, based on our *a priori* impression of whether they are likely to involve a considerable 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 inverse of the variance of the estimated fixed effect, and a regression line (estimated with these weights) and the 95% confidence band is shown.

Part (a) of Figure 1 shows the estimated fixed effects on $-z^*$ for the 40 occupations in our data from the Table 4 regressions. As predicted by our model (and as was apparent from the regressions), 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 (sales). Aside from public service, which is based on a very small sample, the two most positive outliers in the tendency to gender-target ads are administration and tourism occupations.

Part (b) of Figure 1 shows the estimated fixed effects from the Table 5 regressions for firms' relative preferences toward men (z). Here, the strong negative association in part (a) is replaced by essentially a zero overall relation to education. Large positive outliers in the employers' preferences towards men are occupations involving manual labor, technical occupations, and communications/logistics; large negative outliers (indicating a preference for 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

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).

secretarial jobs. We suspect that customer discrimination plays a role in the former cases, and that managers' tastes might account for the latter.[21]

Next, we note that the R^2s in Tables 4 and 5 provide additional information on the structure of the cross-sectional variation in employers' gender preferences (z) and search-related factors ($-z^*$). Notably, less than 10 percent of the variation in either of these factors is explainable by a full set of occupation fixed effects, or even by a full set of over 1600 occupation*industry interactions. Firm fixed effects add a considerable amount of explanatory power, raising R^2 to .28 and .37 respectively. We note, however, that removing occupation effects from the model with firm effects (not shown) reduces R^2 imperceptibly: thus, *almost none of the between-firm differences in gender preferences can be explained by the fact that firms hire different mixes of occupations*. On the other hand, R^2 for both z and $-z^*$ rise to about two thirds with a full set of occupation*firm interactions. This provides further support for the notion that firms' gender preferences are job-specific, in the sense that they vary across occupations within firms, *and* that different firms tend to have different gender preferences for the same occupations. Finally, we note that the remaining one third of the variation in employers' gender preferences (and of search-related factors as well) is within firm*occupation cells. This suggests that firms engage in gender-job typing, or even gender segregation, at the level of highly detailed job descriptions, as documented, for example by Bielby and Baron (1984).

Finally, as outlined in Appendix 4, we calculated the share of total (observed and unobserved) variation in firms' advertised preferences for men that was attributable to variation in z versus $-z^*$. We did the same for women. In both cases, for all regression specifications reported in Tables 4 and 5 we found that about half of both the observed and unobserved variation was attributable to z versus $-z^*$, with a small covariance term. In other words, about half the cross-sectional variation in advertised discrimination in favor of, say, men, is not associated with differences in firms' net preferences for men over women as employees (whether related to tastes, expected productivity differences, *or* wage differences). Instead this variation is associated with search-related factors that make some firms reluctant to restrict their applicant pools and others happy to do so.

[21] In results available from the authors, we show that this same set of four occupations also exhibited a highly disproportionate propensity to request that applicants be tall and physically attractive.

This underscores the importance of accounting for search-related factors when attempting to infer employers' gender preferences from data on discriminatory job ads.

6. Understanding the Effects of Education Requirements on Advertised Discrimination: Alternative Explanations and Robustness Tests

Our main empirical result in this paper is that advertised discrimination against both men and women declines with a job's education and experience requirements. The explanation we have offered for this result, based on our simple model of job advertising, is that higher skill requirements (θ in our model) directly raise the marginal return to identifying the best candidate in the pool of potential applicants. In this section we consider some other possible explanations of our main result, and marshall any available evidence to assess those explanations.

a) Education and the Tastes of Hiring Agents: Evidence on Other Types of Discrimination

One class of alternative explanations argues that the association between education and discrimination in our data is not driven by differences in the price of discrimination across skill levels, but by differences in preferences for discrimination between persons in charge of hiring skilled versus unskilled workers. Perhaps more-educated people are in charge of hiring for jobs requiring more education, and more-educated people simply have less discriminatory tastes? Of course, as our theoretical model makes clear, for this alternative explanation to work, it is not sufficient, for example, for more-educated hiring agents to be less biased against women. What is needed, instead, is a particular sort of covariance between tastes and education: highly-educated agents should not care which gender does any particular job, but as we move we move down the skill ladder, *some* hiring agents' tastes need to become more intense in favor of women, while other hiring agents' tastes become more intense in favor of men.[22] This is of course possible, but if hiring agents favor their own gender, then to be

[22] Formally, there are at least two ways to conceptualize this pattern of tastes. The most direct would be to hypothesize that more-educated people are more likely to simply feel it is ethically wrong to have or express a gender preference. While one could incorporate such preferences into our model, this seems rather close to assuming the result one wishes to explain. More subtly, it could be the case that the variance *across jobs* in the extent to which hiring agents think men versus women will perform better (specifically the variance of v_i in equation 7) is lower in a sample of skilled jobs than in a sample of unskilled jobs.

consistent with our data the ratio of male to female hiring agents would need to remain relatively constant as we move down the skill ladder, which is not the case in most organizations.

Turning to evidence that bears on the above explanation, we note first that if higher education requirements operate by changing the marginal financial return to identifying the best candidate, higher education requirements should affect not only firms' decisions to discriminate on the basis of gender, but their tendency to use other demographic characteristics as screens for applicant ability as well. To explore this question, Table 6 presents descriptive statistics on advertised requirements in our data for three additional ascriptive demographic characteristics: age, beauty and height.[23] While of course it is also possible that hiring agents' tastes for workers' ages, appearance and tastes also covary with education in the same way as their gender preferences, our confidence in the parsimonious "price of discrimination" hypothesis is increased if all these other types of advertised discrimination diminish with skill levels.

According to Table 6, slightly less than one in four ads on Zhaopin.com expressed an age preference; perhaps surprisingly, minimum age requirements were almost as common as maxima. Unlike gender and age requirements, which can go in both directions (excluding either men or women, ruling out both the too-old and the too-young), we could find no ads in our sample of over a million that requested short or unattractive applicants. All together, 7.7 percent of ads requested that the applicant be physically attractive ("*xingxiang*"). Finally, at 2.6 percent of job ads, height requirements are the least common of the ascriptive job requirements on which we have data. Turning to our main hypothesis, Table 6, like Table 1, shows a strong negative association between a job's skill requirements and its tendency to have age, beauty and height requirements. Among jobs requiring high school education or less, 40, 15 and 9 percent had an age, beauty or height requirement; for jobs requiring university education these numbers fall to 17, 4 and 1 percent respectively.

Since this is formally equivalent to a situation where men's and women's *actual* relative productivity varies less across jobs as skill levels rise, we address it in more detail when we discuss the effects of variation in σ_v and σ_ε across job skill levels below.

[23] We also searched for evidence of ethnic or racial discrimination in our sample of ads. Only 56 ads in our sample explicitly requested that the applicant be Han (China's dominant ethnic group). Ads requesting minority ethnicities were actually more common.

Table 7 confirms the robustness of the above patterns by showing linear probability model regression results with the same specifications as columns (4) and (6) of Tables 4 and 5. Because age restrictions, like gender restrictions are 'bilateral' in our data, columns (1)-(4) adopt a similar strategy of separately identifying firms' preferences in the direction of older workers (*z*) (or alternatively, their 'ideal' employee age for that job) from their propensity to impose age restrictions in general ($-z^*$). The latter effects come from regressions in which the dependent variable equals one if the ad had either a minimum or minimum age (or both). The regression underlying the former uses the difference between two dummies (for the presence of a minimum and maximum age respectively) as the dependent variable. Since ads for beauty and height only go in one direction, 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 7 clearly shows that, even in the presence of highly detailed controls, firms' propensity to age-target ads (in any way) declines sharply with jobs' education requirements: in the most saturated regression specification (column 4) jobs requiring some postsecondary education (a university degree) are 8 (11) percentage points less likely to specify an age requirement of any kind. Firms' relative valuation of older workers also seems to rise with education requirements, but the effect is much smaller in magnitude and not consistently statistically significant. For beauty and height it is not possible to separate z and z^* factors, but the strong negative effect of job skill requirements on firms' propensities to specify these requirements is again consistent with a strong negative effect on $-z^*$ combined with weaker (or zero) effects on z.

Turning to the other covariates in Table 7, consistent with our findings for gender we also find that firms with some foreign ownership are much less (8 percentage points) likely to advertise any age restrictions for their jobs, and less likely to request either beauty or a minimum height. We also find that non-profit employers and state-owned enterprises have strong preferences for younger workers, while foreign-owned firms have a slightly greater preference for older workers than Chinese private-sector firms.

b) Education and the Tastes of Hiring Agents: Effects of Other Indicators of Skill Requirements

A second way to distinguish explanations based on the tastes of hiring agents from a price-of-discrimination-based explanation of the education effects in our data is based on the following intuition: If education reduces the amount of advertised discrimination because it raises the monetary gains from identifying the best candidate for the job, then other factors that raise these monetary stakes for employers should have the same effect—reducing advertised discrimination against both men and women. Indeed, we have already shown this for experience requirements: higher experience requirements also reduce advertised discrimination in both directions. This pattern is probably harder to explain using tastes than the education patterns: if jobs requiring more experience tend to have more experienced (and presumably older) hiring agents, a parallel tastes story would require these older agents 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 successive cohorts in most societies.[24]

Of course, in addition to education and experience requirements, another obvious indicator of a job's skill level (and thus of the financial consequences of not hiring the best person for the job) is the job's offered wage level. So far, we have presented no data on advertised wages in this paper, mostly because only about 16.5 percent of the ads in our sample contain usable information about offered wages.[25] Thus, any analysis of wage effects must be conditional on the decision to advertise the wage, a decision we know little about.[26] With this limitation in mind, Table 8 presents coefficients on a job's advertised wage when it is added to the regressions for gender in columns (4) and (6) of

[24] To test this "age and tastes hypothesis" more directly we restricted our sample to ads that specified both a minimum and maximum age, then divided the resulting sample into three groups of ads of roughly equal size: those where the midpoint of the advertised age range was under 28, between 28 and 33, and over 33. The share of job ads with *any* gender requirement was .35, .27, and .22 for the three groups respectively. If tastes of hiring agents are to explain this pattern, it must therefore be the case that agents hiring older workers (who are presumably older themselves) have less-strict notions of proper gender types for jobs.

[25] On Zhaopin.com, firms typically indicate wages by entering a minimum and maximum amount into fields on a form. Overall, only 31.5% of ads contained any information in these fields whatsoever; an additional 14 percent selected both the lowest minimum and highest maximum allowed (1000 and 50,000 yuan/month respectively) allowed on the web form. An additional one percent specified a wage range in excess of 20,000 yuan. All together, this leaves us with only about 16.5 percent of ads that had reasonably informative wage data. Among these ads, the mean midpoint of the advertised wage range was 4279 yuan/month.

[26] See Brencic (2011) for a recent empirical study of this question, which uses partially-directed search models such as Menzio's (2007) as a source of hypotheses and analytical framework.

Tables 4 and 5, and to the equivalent regressions for age, beauty and height in Table 7. Taken together, the results once again strongly support the notion that any factor that raises the marginal financial rewards to identifying the best candidate reduces firms' reliance on coarse demographic screens, including gender, in the hiring process. Columns (1)-(3) also indicate that firms' net preferences towards men for a job are essentially unrelated to the job's overall wage level. Notably, this finding is not consistent with "glass ceiling" motivations for gender discrimination. Simply put, 'keeping women out of high-paid jobs' is not a good overall description of firms' motivations for gender discrimination in this data.

c) Other Model Parameters affecting the Cost of Discrimination.

So far in this paper, we have focused the test of our model on its implications for the effects of the job's skill demands (θ) on advertised discrimination. While we believe this question has considerable intrinsic interest, our main justification for this focus is simply the availability of data: we have multiple estimates of skill demands (education, experience and wage), but we do not have good measures of the other fundamental parameters of our simple model, namely c, M, σ_ε, and σ_v. Since all of these factors can co-vary with a job's skill level, differences in these unmeasured factors across job skill levels could also explain our main empirical result. In the remainder of this section, we discuss each of these potential confounds in turn.

First, it seems plausible that the cost of assessing the suitability of an individual applicant (c) rises with a job's skill level, and available evidence supports this notion.[27] Could this explain why there is less advertised gender discrimination for highly skilled jobs? Absolutely not, according to our model. Instead, as Proposition 1 indicates, increases in c should make advertised discrimination *more* likely: evaluating a large group of applicants is now more costly, so the firm is grateful for any screen might raise the average quality of applicants while reducing the size of the applicant pool. Put a different way, unless one believes that applicant screening costs *fall* with job skill levels, variation in screening costs with skill cannot explain our paper's main result.

[27] See for example Table 1 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.

On the other hand, it is also possible that ads for skilled jobs, on average, attract fewer applicants (M) than ads for less skilled jobs: in some sense, skilled labor markets may be thinner because skilled workers are more specialized. In contrast to the evidence on c above, however, the available evidence on this is not very clear.[28] A related possibility is that there is, on average, more idiosyncratic variation in the true, underlying qualifications of applicants to skilled positions than among applicants to unskilled positions (i.e. σ_{ε},, or β is higher). Because it requires a scalar measure of the variance of all worker qualifications that are visible to the employer in the hiring process, this is a difficult hypothesis to test directly. Importantly, however, even if thinner markets and greater idiosyncracy of applicants play a significant part in explaining our main result, our results still suggest that skill upgrading in a labor market will reduce overt, advertised discrimination, as long as idiosyncracy and thinness are in a sense intrinsic to skilled labor markets. Only the mechanism via which skill operates is different.

A final possibility is that the idiosyncratic variance *across jobs* in the relative ability of men and women to perform them, σ_v, is *lower* in a sample of skilled jobs than a sample of unskilled jobs.[29] One sense in which this might be true is if unskilled jobs involve more manual labor, and if there is are larger gender gaps in humans' abilities to perform physical tasks than mental tasks. For example, men might have an advantage in strength-intensive occupations and women in jobs requiring fine motor skills, while the genders are equally able at non-physical tasks. (Or put even more succinctly, perhaps men's and women's minds are more similar than their bodies). While this seems plausible, we note that the share of jobs in our sample that are likely to involve any physical labor is quite small.[30] Furthermore, excluding all of these occupations from our estimation sample has essentially no effect on the our main results for education and experience in Table 4.

---

[28] The mean number of applicants per job in Barron and Bishop's 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 (see note 30).

[29] Note that formally, this explanation is identical to the argument that hiring agents' tastes are less variable across jobs at high skill levels (see footnote 22), except that the heterogeneity across jobs now refers to the gender gap in actual productivity: both hypotheses correspond to a lower level of σ_v in a sample of jobs requiring more skill.

[30] 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.

In sum, in our assessment the evidence in this section largely supports the notion that increases in the financial benefits associated with identifying the ablest job candidate play an important role in explaining why advertised gender discrimination is lower among jobs requiring higher levels of skill. While we cannot (and do not wish to) rule out the possibility that other factors associated with higher skill levels, such as thinner labor markets or greater heterogeneity in applicant productivity, also contribute to this pattern, we note that as long as these additional factors are in some sense intrinsic features of skilled labor markets, our analysis continues to suggest that rising skill requirements in a labor market will reduce the incidence of explicit discrimination – whether on the basis of gender, age, height or other characteristics such as beauty-- in job ads.

7. 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. This pattern holds both for discrimination against women and for discrimination against men. We see this potent role of skill demands as a deterrent to discrimination as complementary to the role of product market competition as emphasized by Becker (1957), and suspect that rising skill demands may play an important role in explaining why nations tend to abandon explicit discrimination as they develop economically.

What underlying economic processes might explain this robust effect of skill demands? In this paper, we present a simple model of employer search from two populations. According to the model, regardless of whether firms prefer men or women for the job they are advertising, firms should be less likely to restrict their search to their preferred gender as the job's skill requirements rise. The intuition is straightforward: as skill requirements rise, it becomes increasingly important for firms to identify the best individual candidate for the job. While other factors, such as thinner labor markets or a different cross-sectional distribution of tastes among hiring managers for skilled positions, could also account for the same pattern in the data, we present a number of pieces of evidence which suggest that the above simple and direct mechanism likely plays a central role.

Another useful feature of our model is that it allows us to separately identify the effects of observable ad and firm characteristics on 'search related' factors –which reduce advertised discrimination against both men and women--, and employers' net preferences towards men (and away from women) for a particular job. Since search-related factors account for almost half of the cross-job variation in advertised discrimination, it is important to account for them when attempting to make inferences about firms' gender preferences. Using the framework of our model, we find that employers' underlying preferences for men versus women in a job are unrelated to the job's education and wage level, but increase with the job's required level of experience. Large firms, and state-owned enterprises also tend to disproportionately prefer men over women, while foreign owned firms seem to prefer women.

Finally, an important feature of our data is that employers' estimated gender preferences vary widely across jobs *within* firms. Indeed, the broad statistical portrait that emerges from our data is one where gender-targetted ads are primarily used to allocate some of each firm's unskilled jobs to men, and others to women. Further, since these occupational gender patterns appear to differ across firms, their explanations may need to consider factors beyond the nature of the job itself. One such possibility is the existence of productivity gains (or wage savings) associated with gender *homogeneity* in a detailed job title (see for example Goldin 2006). This suggests that future research into the causes of gender differentials, especially in emerging-country labor markets such as China's, might do well to focus on the causes and consequences of within-firm gender segregation.

Should China and other developing nations follow the U.S.'s example and ban the mention of sex, age, and perhaps other ascriptive characteristics in job ads? While such normative questions are beyond the scope of this paper, we note on the one hand that allowing explicit discrimination in job ads does make it easier for employers to engage in taste-based discrimination. On the other hand, we note that a much higher share of U.S. than Chinese workers is employed in highly skilled jobs, where both our model and data indicate that advertised discrimination is less useful to employers. Also, our results suggest that three ongoing developments –upskilling of the Chinese labor force, expansion of the share of Chinese firms that face meaningful product market competition, and an increasing presence in China of firms with some foreign ownership—may have powerful effects in reducing the incidence of discriminatory job ads even in the absence of policy interventions.

References

Altonji, J. and C. Pierret. "Employer Learning and Statistical Discrimination" *Quarterly Journal of Economics*, Vol. 116, No. 1. (Feb., 2001), pp. 313-350

Angrist, Joshua D., and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton: Princeton University Press, 2009.

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

Banerjee, Abhijit, Esther Duflo, Maitreesh Ghatak, and Jeanne Lafortune. "Marry for What: Caste and Mate Selection in Modern India" NBER working paper no. 14958, May 2009.

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

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.

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

Bergeron, Pierre-Jerome, Masoud Asgharian, and David B. Wolfson. "Covariate Bias Induced by Length-Biased Sampling of Failure Times" *Journal of the American Statistical Association* Vol. 103 (482). p 737-42. (June 2008).

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(4) (September 2004): 991-1013.

Bielby, W. T., and J. N. Baron. (1984) "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.

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.

Brencic, Vera. "Wage Posting". Unpublished paper, University of Alberta, 2011.

Charles, Kerwin Kofi and Jonathan Guryan. 2008 "Prejudice and Wages: An Empirical Assessment of Becker's The Economics of Discrimination" *Journal of Political Economy*, 116(5): 773-809

Crosby, F., Bromley, S., & Saxe, L. (1980). "Recent unobtrusive studies of Black and White discrimination and prejudice: A literature review" *Psychological Bulletin*, 87, 546-563.

Darity, William A. Jr., and Patrick L. Mason. "Evidence on Discrimination in Employment: Codes of Color, Codes of Gender". *Journal of Economic Perspectives* 12(2) (Spring 1998):63-90.

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?* Eds. Francine D. Blau, Mary C. Brinton, and David B. Grusky. 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* (2000) 13: 305-329

Kuhn, Peter and Kailing Shen. “Employers’ Preferences for Gender, Age, Height and Beauty: Direct Evidence”. NBER working paper no. 15564, December 2009

Lang¸ Kevin, Michael Manove and William T. Dickens. “Racial Discrimination in Labor Markets with Posted Wage Offers” *American Economic Review*. Vol. 95 (4). p 1327-40. September 2005.

Lawler, John J. and Johngseok Bae. 1998. “ Overt Employment Discrimination by Multinational Firms: Cultural and Economic Influences in a Developing Country *Industrial Relations* 37(2): 126-152.

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

Menzio, Guido. “A Theory of Partially Directed Search” *Journal of Political Economy*. Vol. 115 (5). p 748-69. October 2007.

Moen, Espen R. “Competitive Search Equilibrium” *Journal of Political Economy*. Vol. 105 (2). p 385-411. April 1997.

Mortensen, Dale, and Christopher Pissaridies. **“**Job Creation and Job Destruction in the Theory of Unemployment.” *Review of Economic Studies*. July 1994; 61(3): 397-415

Neumark, David. “Sex Discrimination in Restaurant Hiring: An Audit Study” *Quarterly Journal of Economics* 111(3) (August 1996): 915-42.

Van Ours, Jan, and Geert Ridder. “Job requirements and the recruitment of new employees” *Economics Letters* 36 (1991): 213-218.

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*, Vol. 61, No. 2 (January 2008).

Appendix 1: Proofs

Proof of Lemma 1:

We begin by normalizing the net value of an applicant, defining $u_j \equiv U_j/\beta = v^G/\beta + e_j$, where $e_j = \varepsilon_j/\beta$ follows a "standard" extreme value distribution with $\mathrm{Var}(e_j) = \pi^2/6$ and $\mathrm{E}(e_j) = \gamma$, j indexes applicants, and $G \in (M, F)$ indexes groups. This normalization does not affect the firm's optimal selection of a worker –the draw of e_j that maximizes u corresponds to the draw of ε_j that maximizes U-- and the maximized value of U can be calculated as βu^*, where u^* is the maximized value of u. Further, this normalization expresses the problem in a standard multinomial logit format, which allows us to draw on some results from that literature.

Among these, it is well known that the expected value of the maximum of $v^G/\beta + e_j$ when e_j is independently drawn N times from a "standard" extreme value distribution is $v^G/\beta + \gamma + \log(N)$.[31] Multiplying through by β the expected maximum of U when the firm samples from either the M or F pool separately is

(A1) $$U^{G}* = v^G + \beta\gamma + \beta\log(G),$$

which proves part (a) of the Lemma.

Turning back to standardized net values, the expected value of the highest u_j in the "combined" sample, u^C*, equals $u^M*q^M + u^F*(1- q^M)$, where u^G* is the expected value of the best *overall* worker given the best overall worker is from group G, and q^M is the probability that the best overall worker turns out to be an M. Again using results from the MNL literature, we know that $u^M* = v^M/\beta + \gamma - \log(p^M)$, where:

$$p^M = \frac{\exp(v^M/\beta)}{M\exp(v^M/\beta) + F\exp(v^F/\beta)}$$

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

$$p^F = \frac{\exp(v^F/\beta)}{M\exp(v^F/\beta) + F\exp(v^F/\beta)}.$$

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

$$q^M = \frac{M\exp(v^M/\beta)}{M\exp(v^M/\beta) + F\exp(v^F/\beta)}.$$

Note that, as the variance of individual productivity (β) falls towards zero, the probability that the best overall worker will be from the group with the higher net value (v) approaches one; conversely as β approaches infinity, q^M approaches the share of M's in the population, i.e. $M/(M+F)$.

[31] See Arcidiacono and Miller (2008, p. 8) for a general proof.

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

$$u^{C}* = \gamma + log\,[\,Mexp(v^{M}/\beta) + Fexp(v^{F}/\beta)\,]\,.$$

Letting $\delta = M/(M+F) \equiv M/C$ be the fraction of M's in the combined pool, this becomes:

$$u^{C}* = \gamma + log\,[\,\delta\, exp(v^{M}/\beta\,) + (1\text{-}\delta\,)exp(v^{F}/\beta\,)\,] + \log C\,.$$

Multiplying through by β, the corresponding maximized unstandardized value is therefore:

$$U^{C}* = \gamma\,\beta + \beta\, log\,[\,\delta\, exp(v^{M}/\beta\,) + (1\text{-}\delta\,)exp(v^{F}/\beta\,)\,] + \beta \log C\,.$$

Expressing this in terms of the means of the unstandardized distributions, $\mu^{G} \equiv v^{G} + \beta\gamma$, yields after some algebra:

(A2) $$U^{C}* = \beta \log\left[\delta \exp\left(\frac{\mu^{M}}{\beta}\right) + (1-\delta)\exp\left(\frac{\mu^{F}}{\beta}\right)\right] + \beta \log C\,,$$

which proves part (b).

Proof of Propositions 1 and 2:

Consider first the difference in the firm's objective function ("profits") between recruiting strategy M (invite men only) and C (invite all). Expected profits from inviting only men to apply are (from (2) and (3)):

$$\Pi^{M} = U^{M}* - cN = \mu^{M} + \beta\log(N) - cN.$$

Expected profits from a combined strategy are (from (2) and (5):

$$\Pi^{C} = U^{C}* - 2cN = \mu^{M} + \beta \log\left[1 + (\exp\left(\frac{\mu^{F} - \mu^{M}}{\beta}\right)\right] + \beta \log N - 2cN.$$

Subtracting and simplifying yields:

(A3) $$\Pi^{M} - \Pi^{C} = -\beta \log[1 + \exp(-z)] + cN \equiv R^{M}(z, \beta) + cN$$

where $N = .5C$ is the (equal) number of applications expected from each group, and $z = \dfrac{\mu^{M} - \mu^{F}}{\beta}$ is the standardized expected net value advantage of group M.

By symmetry,

(A4) $\Pi^F - \Pi^C = -\beta\log[1+\exp(z)]+cN = R^M(-z, \beta) + cN \equiv R^F(z, \beta) + cN$

Inspection of (A3) and (A4) shows that $R^M(z, \beta)$ is always negative and increasing in z. $R^M(-z, \beta)$ is thus also always negative and decreasing in z. Further, these two functions intersect when $z = 0$, at the level $-\beta\log 2$, as shown in Figure A1 below.

When $-cN > -\beta\log 2$, as is the case when cN=.4 in Figure A1, it follows from (A4) and (A5) that firms prefer to advertise only to men when $z > z^*$, to advertise only to women when $z<-z^*$, and not to restrict their ads when $-z^*<z<z^*$. When $-cN < -\beta\log 2$, as is the case when cN=1 in Figure A1, the combined search strategy is never preferred to *both* of the restricted strategies, so the firm restricts to men when z>0, and to women when z<0.

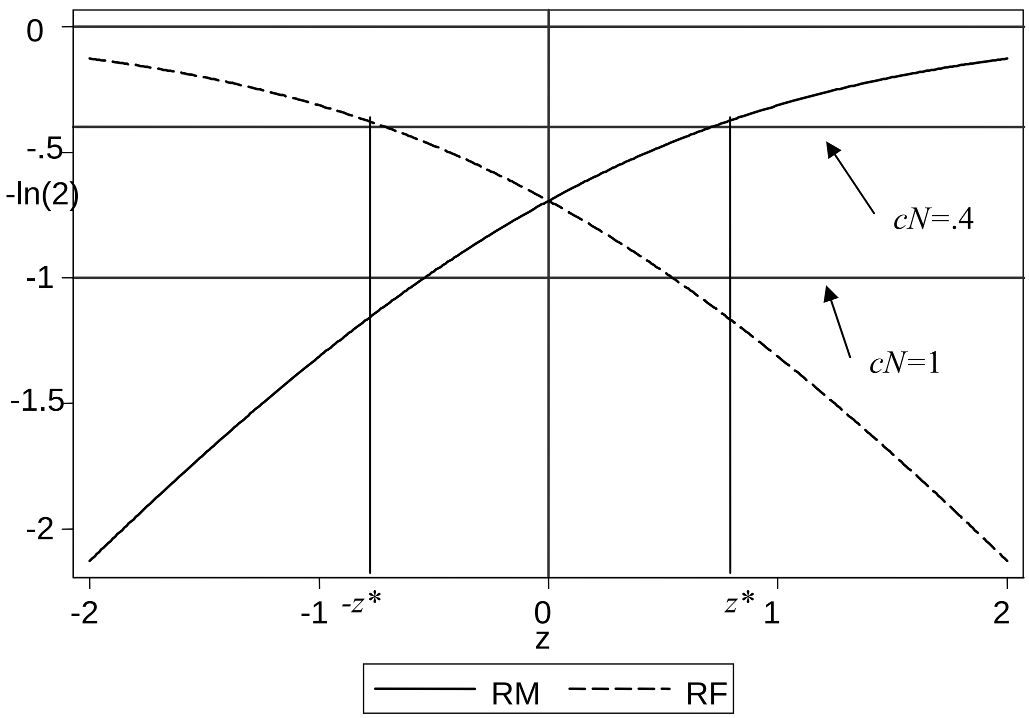

Figure A1: Gains and Costs to Restricted Searches, β=1

The critical values, z^* and $-z^*$, can be found by setting A3 or A4 equal to zero, yielding:

(A5) $z^* = -\log[\exp(cN/\beta)-1]$

This completes the proof for a job with standard skill requirements. To generalize the analysis to any job, recall that by assumption skill requirements are represented by a constant, θ, that multiplies all realizations of U_j. Thus, changes in θ multiply both the means (μ^M and μ^F) and standard deviation $(\sigma = \beta\pi/\sqrt{6})$ of the applicants' net values (U_j) by the same constant, leaving the standardized expected gender gap in net values for a particular job, $z = \dfrac{\mu^M - \mu^F}{\beta}$ unchanged. Therefore, the only change in firms' decision rules involves the threshold, z^*, which becomes, for a job with skill requirement θ : $z^* = -\log[\exp(cN/\theta\beta) - 1]$, as claimed in Proposition 2.[32]

Appendix 2: Data

As noted, our overall 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. 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%.

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); our approach could therefore miss jobs that ask for age only using numbers "25-35". Therefore, the variables that we use here should be interpreted as having 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 when submitting their ad). Note that our occupation and industry dummy variables are not mutually exclusive, as firms are allowed to check multiple categories. As noted in the paper, we handle this in the results reported here by restricting attention to ads for a single occupation. 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).

---

[32] Since for given $\mu^M - \mu^F$, z depends by definition on β, Proposition 1 and 2's predictions for the effects of the idiosyncratic variance parameter, β, must be interpreted as conditional on a given z, i.e. conditional on a given *standardized* gap in expected values. In consequence the effects of β are, in fact, identical to the effects of rescaling workers' value by θ. However, it is easy to show that the effects of an increase in β conditional on a given *absolute* gap in expected values, i.e. a given $\mu^M - \mu^F$, is qualitatively the same but stronger in magnitude. To see this, note that for given $\mu^M - \mu^F$, an increase in β 'shrinks' all the z's towards the origin. This effect reinforces the effects of β on the thresholds (z^*'s), thus reinforcing the decline in the 'probability' that firms will adopt gender restrictions of either kind.

Appendix 3: A Linear Approximation to the Model

Consider first the effects of an infinitesimal increase in covariate x_k on the probability that an ad is targetted at men (P^M) or women (P^F) respectively in our model. Differentiating the corresponding terms in equation (8), these marginal effects are given by:

(A6)
$$\partial P^M / \partial x_k = \left(d_k \exp(-\mathbf{xd}) + b_k\right) f(v_2{*})$$
$$\partial P^F / \partial x_k = \left(d_k \exp(-\mathbf{xd}) - b_k\right) f(v_1{*}),$$

where $v_1{*}$ ($v_2{*}$) is the value of v below (above) which firms advertise a preference for men (women).

Combining these to yield the marginal effect of x_k on the probability that a firm engages in *some* type of gender discrimination:

(A7) $$\partial(P^M + P^F)/\, \partial x_k = d_k \exp(-\mathbf{xd})\left(f(v_1{*}) + f(v_2{*})\right) + b_k\left(f(v_2{*}) - f(v_1{*})\right).$$

Now suppose for a moment that $f(v_i)$ is uniform with density $m/2$. In this case, (A7) reduces to

(A8) $\partial(P^M + P^F)/\, \partial x_k = d_k m \exp(-\mathbf{xd})$, while a parallel argument for (P^M - P^F) yields:

(A9) $\partial(P^M - P^F)/\, \partial x_k = b_k m.$

In other words, if $f(v_i)$ is uniform*,* then up to a factor of proportionality (equal to $m\exp(-\mathbf{xd})$ for an individual observation or $mE(\exp(-\mathbf{xd}))$ for the average marginal effect in the estimation sample), the parameter d_k is equal to x_k's marginal effect on the probability that a firm discriminates *in some direction* in its ad. This is true regardless of x_k's effects on the firm's preference towards men, z, because any such effects subtract out of (A7). Similarly, b_k is identified by the marginal effect of x_k on the difference between the probabilities, P^M -P^F. More generally, when $f(v_i)$ is not uniform, the average marginal effects become:

(A10) $$E[\partial(P^M + P^F)/\, \partial x_k] = d_k E\left\{\exp(-\mathbf{xd})[f(v_1^*) + f(v_2^*)]\right\} + b_k E\left\{f(v_2^*) - f(v_1^*)\right\},$$

(A11) $$E[\partial(P^M - P^F)/\, \partial x_k] = d_k E\left\{\exp(-\mathbf{xd})[f(v_2^*) - f(v_1^*)]\right\} + b_k E\left\{f(v_2^*) + f(v_1^*)\right\},$$

Now, if the share of ads targetted at men versus women is similar ($P^M \approx P^F$), then for any *symmetric* $f(v_i)$, the expected difference in densities $E\left\{f(v_2^*) - f(v_1^*)\right\}$ will be approximately zero. This, combined with the fact that the covariate vector **xd** only moves $v_1{*}$ and $v_2{*}$ further apart or closer together, means that both the second term in (A10) and the first in (A11) will be close to zero. Thus, (A8) and (A9) remain approximately true for the average marginal effects in the estimation sample, with the appropriate expectations replacing the constants $m\exp(-\mathbf{xd})$ and m respectively.

Finally, we appeal to the fact that linear probability models typically estimate the average marginal effects well, even in limited dependent variable contexts. (According to Angrist and Pischke (2009), the correspondence is exact in the case of a single, dummy regressor (pp. 96-98), and approximate in more general applications (pp. 104-107). In sum, according to our argument, the OLS coefficient in a regression of $P^M + P^F$ on x_k will be approximately equal to gd_k, and the OLS coefficient in a regression of P^M - P^F on x_k will be approximately equal to hb_k, where $g =$ $E\{\exp(-\mathbf{xd})[f(v_1^*)+f(v_2^*)]\}$ and $h = E[f(v_1^*)+f(v_2^*)]$.

Appendix 4: Variance Decomposition

In this exercise, we decompose the cross-sectional variation in advertised preferences for men (P^M) and in advertised preferences for women (P^F) in a linear probability model into following two components, plus the covariance between them:

i) Variation in firms' ***relative preference*** for men versus women in the job: productivity differences, tastes and wage differences. (associated with parameter vector **b** below)

ii) Variation across jobs and firms in ***search-related factors***: expected number of applicants, application processing costs, and variance in match quality? (associated with parameter vector **d** below)

Specifically, we suppose that:

Prob (request male) = $P_i^M = (\mathbf{b} + \mathbf{d})\mathbf{x}_i + e_i^M$

Prob (request female) = $P_i^F = (\mathbf{-b} + \mathbf{d})\mathbf{x}_i + e_i^F$, where i indexes ads.

Thus, the parameter b_k gives the effect of ad characteristic x_k on firms' relative preference towards men (z), and d_k gives its effect on firms' propensity to restrict in general ($-z^*$). (To simplify the presentation, note that this notation departs slightly from the body of the paper, where **b** and **d** determine the underlying indexes, not the probability itself.)

Then, if we estimate separate linear probability models for P_i^M and P_i^F :

$P_i^M = \boldsymbol{\beta}^{\mathbf{M}}\mathbf{x_i} + e_i^M$

$P_i^F = \boldsymbol{\beta}^{\mathbf{F}}\mathbf{x_i} + e_i^F$, the estimated coefficients will exactly identify **b** and **d**, because

(A12) $\boldsymbol{\beta}^{\mathbf{M}} = \mathbf{b} + \mathbf{d},$ and

(A13) $\boldsymbol{\beta}^{\mathbf{F}} = \mathbf{-b} + \mathbf{d}$.

After solving (A12) and (A13), for **b** and **d**, columns (1)-(6) of Table A1 then decompose the explained variation in P_i^M and P_i^F using:

$\text{Var}(\boldsymbol{\beta}^{\mathbf{M}}\mathbf{x}) = \text{Var}(\mathbf{bx}) + \text{Var}(\mathbf{dx}) + 2\text{Cov}(\mathbf{bx},\mathbf{dx})$

$\text{Var}(\boldsymbol{\beta}^{\mathbf{F}}\mathbf{x}) = \text{Var}(\mathbf{bx}) + \text{Var}(\mathbf{dx}) - 2\text{Cov}(\mathbf{bx},\mathbf{dx})$.

The regression specifications underlying columns (1)-(6) are identical to those in Tables 4 and 5. Column 7 decomposes the unexplained variance in the same manner, using the residuals e_i^M and e_i^F, from the column 6 regressions.

Table A1: Variance Decomposition for Determinants of Gender Preferences

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Explained Share of Variance with the following Fixed Effects:						**Unexplained Variance in Specification (6)**
	Province	Occupation, Province	Occ*Ind, Province	Occ*Ind* Province	Firm, Occ	Firm*Occ	
Variance in Advertised Preferences for Men:							
1. Total	.026	.061	.082	.149	.324	.639	.361
2. Related to tastes and productivities (z)	.007	.032	.042	.073	.143	.324	.182
3. Related to frictions ($-z^*$)	.020	.035	.043	.070	.166	.303	.150
4. Related to covariance between these	.000	-.007	-.004	.006	.016	.013	.030
Variance in Advertised Preferences for Women:							
5. Total	.030	.081	.098	.150	.319	.671	.329
6. Related to tastes and productivities (z)	.008	.035	.046	.080	.156	.354	.198
7. Related to frictions ($-z^*$)	.021	.039	.047	.077	.181	.331	.164
8. Related to covariance between these	.000	.007	.004	-.006	-.017	-.014	-.032

With the exception of column 1, rows 2 and 3 are approximately equal. So are rows 6 and 7. We conclude that variation in z and z^* contribute about equally to the observed variance in advertised preferences for men. The same is true for women.

Table 1: Descriptive Statistics

	Advertised Education Requirement			
A. AD CHARACTERISTICS	High School	Some Postsecondary	University	Combined
Gender requirement				
No gender preference	.766	.892	.938	.895
Prefer male?	.120	.049	.042	.055
Prefer female?	.114	.059	.021	.050
Advertised Experience requirement				
Unspecified	.343	.204	.136	.194
1 Year or less	.030	.013	.005	.012
1 to 3 years	.467	.482	.287	.399
3 to 5 years	.120	.209	.305	.237
5 to 10 years	.037	.085	.232	.139
10 Years or above	.004	.007	.037	.019
Job is Part Time	.015	.008	.009	.009
Number of positions advertised:				
Unspecified	.546	.483	.456	.480
1	.159	.260	.358	.287
2	.089	.104	.094	.098
3-5	.102	.094	.067	.084
6-15	.070	.046	.021	.038
16-50	.028	.013	.005	.011
51+	.006	.001	.000	.001
B. FIRM CHARACTERISTICS				
Firm size (number of workers):				
1-19	.067	.071	.081	.075
20-99	.333	.336	.316	.327
100-499	.332	.325	.335	.330
500-999	.108	.097	.095	.098
1,000-9,999	.124	.134	.138	.135
10,000 +	.035	.036	.036	.036
Firm ownership type:*				
Private, Domestic	.672	.606	.477	.561
Foreign	.259	.326	.430	.360
NonProfit	.003	.003	.005	.004
State-Owned Enterprise	.066	.065	.087	.074
Government	.000	.000	.001	.001
Number of Ads	136,595	482,697	438,246	1,057,538

* “Private, Domestic” includes privately held companies, publicly-traded companies and reformed State-Owned Enterprises where a majority of shares are still owned by the state. “Foreign” includes Foreign Direct Investment, joint ventures, plus a small number of representative offices.

Table 2: 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	12,834
2-10	.326	.155	.224	.053	41,233
11-50	.549	.345	.392	.188	16,343
51 and over	.708	.536	.560	.388	3,792
All Firms	.367	.199	.258	.091	74,202

Table 3: Maximum Likelihood and OLS estimates of structural parameters (b and d) and associated marginal effects.

	Maximum Likelihood Estimates				*OLS Estimates*	
	b	***d***	$\partial(P^M+P^F)/\partial x$	$\partial(P^M-P^F)/\partial x$	$\partial(P^M+P^F)/\partial x$	$\partial(P^M-P^F)/\partial x$
	(1)	(2)	(3)	(4)	(5)	(6)
Education Requirement:						
Some Postsecondary	-.1090**	-.2899**	-.0962**	-.0300**	-.1215**	-.0267**
	(.0041)	(.0028)	(.0009)	(.0008)	(.0009)	(.0010)
University	.0173**	-.4289**	-.1390**	-.0101**	-.1616**	-.0109**
	(.0046)	(.0030)	(.0010)	(.0009)	(.0010)	(.0011)
Experience Requirement:						
1-3 years	.0948**	-.0516**	-.0150**	.0166**	-.0202**	.0273**
	(.0039)	(.0024)	(.0008)	(.0008)	(.0008)	(.0009)
3-5 years	.3084**	-.1047**	-.0283**	.0560**	-.0304**	.0656**
	(.0048)	(.0028)	(.0009)	(.0009)	(.0009)	(.0010)
More than 5 years	.4664**	-.1022**	-.0246**	.0865**	-.0157**	.0878**
	(.0060)	(.0033)	(.0010)	(.0011)	(.0010)	(.0011)
Log (Firm Size)	.0213**	-.0040**	-.0009**	.0040**	-.0013**	.0046**
	(.0008)	(.0005)	(.0002)	(.0002)	(.0002)	(.0002)
Firm Ownership Type:						
Foreign Ownership	-.0359**	-.1266**	-.0418**	-.0109**	-.0380**	-.0101**
	(.0035)	(.0019)	(.0006)	(.0007)	(.0006)	(.0007)
Non-profit	-.0192	-.0182	-.0063	-.0043	-.0097	-.0038
Organization	(.0248)	(.0143)	(.0046)	(.0048)	(.0048)	(.0051)
State-owned	.0777**	-.0241**	-.0064**	.0142**	-.0081**	.0147**
Enterprise	(.0059)	(.0035)	(.0011)	(.0011)	(.0012)	(.0012)
Government	.2534**	.1665**	.0588**	.0540**	.0743**	.0821**
	(.0513)	(.0338)	(.0107)	(.0096)	(.0110)	(.0118)

** p<0.01, * p<0.05.

Regressions also control for the number of vacancies advertised, a dummy for part-time jobs and period fixed effects in both equations. The omitted firm type is for-profit firms with no government or foreign connection. Sample size = 1,057,538 ads

Table 4: Effects of Selected Covariates on Employers' Tendency to Gender-Target Ads (-*z), Linear Probability Model Estimates:**

	(1)	(2)	(3)	(4)	(5)	(6)
Education Requirement:						
Some Postsecondary	-.1204**	-.0743**	-.0688**	-.0681**	-.0923**	-.0602**
	(.0073)	(.0050)	(.0048)	(.0048)	(.0044)	(.0049)
University	-.1585**	-.1005**	-.0947**	-.0944**	-.1220**	-.0808**
	(.0092)	(.0057)	(.0055)	(.0056)	(.0059)	(.0058)
Experience Requirement:						
1-3 years	-.0203**	-.0158**	-.0176**	-.0179**	-.0250**	-.0220**
	(.0030)	(.0023)	(.0023)	(.0024)	(.0024)	(.0025)
3-5 years	-.0302**	-.0324**	-.0327**	-.0324**	-.0394**	-.0325**
	(.0041)	(.0027)	(.0027)	(.0027)	(.0044)	(.0032)
More than 5 years	-.0154**	-.0290**	-.0277**	-.0292**	-.0359**	-.0345**
	(.0048)	(.0035)	(.0034)	(.0033)	(.0049)	(.0038)
Log (Firm Size)	-.0017*	-.0028**	-.0034**	-.0034**		
	(.0007)	(.0006)	(.0006)	(.0005)		
Firm Ownership Type:						
Foreign Ownership	-.0309**	-.0313**	-.0308**	-.0305**		
	(.0025)	(.0027)	(.0026)	(.0027)		
Non-profit	-.0070	-.0059	-.0072	-.0013		
Organization	(.0104)	(.0099)	(.0097)	(.0099)		
State-owned	-.0054	-.0061	-.0049	-.0034		
Enterprise	(.0042)	(.0036)	(.0032)	(.0032)		
Government	.0699**	.0680**	.0742**	.0711*		
	(.0272)	(.0230)	(.0287)	(.0313)		
Fixed Effects (number of groups)	Province (31)	Occ, Ind, Province (115)	Occ*Ind, Province (1,638)	Occ*Ind* Province (22,909)	Occ, Firm, Province (74,232)	Occ*Firm, Province (260,322)
R^2	.043	.078	.096	.156	.366	.669

** p<0.01, * p<0.05. 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. The omitted firm type is for-profit firms with no government or foreign connection. Sample size = 1,057,538 ads.

Table 5: Effects of Selected Covariates on Employers' Relative Valuation of Men (*z*), Linear Probability Model Estimates:

	(1)	(2)	(3)	(4)	(5)	(6)
Education Requirement:						
Some Postsecondary	-.0256*	-.0052	-.0053	-.0029	-.0443**	-.0125
	(.0112)	(.0067)	(.0064)	(.0064)	(.0094)	(.0082)
University	-.0071	.0044	.0057	.0097	-.0323**	.0005
	(.0125)	(.0077)	(.0075)	(.0076)	(.0103)	(.0089)
Experience Requirement:						
1-3 years	.0275**	.0170**	.0138**	.0127**	.0258**	.0131**
	(.0036)	(.0027)	(.0026)	(.0027)	(.0037)	(.0030)
3-5 years	.0657**	.0430**	.0407**	.0397**	.0624**	.0387**
	(.0070)	(.0045)	(.0045)	(.0047)	(.0077)	(.0057)
More than 5 years	.0869**	.0616**	.0599**	.0573**	.0787**	.0486**
	(.0082)	(.0056)	(.0056)	(.0057)	(.0089)	(.0068)
Log (Firm Size)	.0042**	.0046**	.0045**	.0047**		
	(.0007)	(.0006)	(.0007)	(.0007)		
Firm Ownership Type:						
Foreign Ownership	-.0072**	-.0101**	-.0088**	-.0080**		
	(.0022)	(.0022)	(.0021)	(.0021)		
Non-profit	-.0028	.0018	-.0010	.0059		
Organization	(.0092)	(.0100)	(.0092)	(.0094)		
State-owned	.0160**	.0129**	.0144**	.0141**		
Enterprise	(.0035)	(.0030)	(.0030)	(.0029)		
Government	.0831	.1087*	.0908	.0666		
	(.0417)	(.0502)	(.0469)	(.0329)		
Fixed Effects (number of groups)	Province (31)	Occ, Ind, Province (115)	Occ*Ind, Province (1,638)	Occ*Ind* Province (22,909)	Occ, Firm, Province (74,232)	Occ*Firm, Province (260,322)
R^2	.014	.064	.084	.144	.282	.640

** p<0.01, * p<0.05. 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. The omitted firm type is for-profit firms with no government or foreign connection. Sample size = 1,057,538 ads.

Table 6: Advertised Preferences for Ascriptive Characteristics other than Gender

	Advertised Education Requirement			
	High School	Some Postsecondary	University	Combined
Ad has no age restrictions	.604	.733	.831	.757
Ad has a minimum age requirement	.287	.192	.108	.169
Ad has a aximum age requirement	.348	.218	.139	.202
Job requires beauty (“xingxiang”)	.146	.091	.041	.077
Job has a height requirement	.093	.022	.009	.026

Table 7: Effects of Selected Covariates on Preferences for Other Ascriptive Characteristics:

	Relative Valuation of Older Workers (Any Minimum Age? – Any Maximum Age?)		Propensity to Age-Target Ads (Any Minimum Age? + Any Maximum Age?)		Request Beauty?		Height Requirement?	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Education Requirement:								
Some Postsecondary	.0123**	.0219**	-.1468**	-.0784**	-.0248**	-.0298**	-.0336**	-.0342**
	(.0045)	(.0038)	(.0078)	(.0089)	(.0042)	(.0046)	(.0039)	(.0037)
University	-.0051	.0319**	-.2627**	-.1108**	-.0420**	-.0442**	-.0367**	-.0389**
	(.0046)	(.0042)	(.0101)	(.0100)	(.0043)	(.0053)	(.0039)	(.0044)
Experience Requirement:								
1-3 years	-.0056**	-.0047*	-.0341**	-.0089*	-.0157**	-.0162**	-.0126**	-.0132**
	(.0021)	(.0023)	(.0056)	(.0044)	(.0020)	(.0027)	(.0016)	(.0022)
3-5 years	.0120**	.0108**	-.0001	.0140**	-.0424**	-.0400**	-.0200**	-.0224**
	(.0029)	(.0031)	(.0070)	(.0052)	(.0041)	(.0055)	(.0024)	(.0038)
More than 5 years	.0240**	.0243**	.0535**	.0584**	-.0516**	-.0473**	-.0202**	-.0252**
	(.0036)	(.0036)	(.0093)	(.0073)	(.0045)	(.0059)	(.0025)	(.0043)
Log (Firm Size)	.0000		.0136**		-.0026**		.0005	
	(.0008)		(.0021)		(.0006)		(.0003)	
Firm Ownership Type:								
Foreign Ownership	.0140**		-.0844**		-.0072**		-.0033**	
	(.0018)		(.0068)		(.0023)		(.0010)	
Non-profit	-.0784**		.0255		.0089		.0078	
Organization	(.0112)		(.0201)		(.0118)		(.0061)	
State-owned	-.0382**		-.0285**		-.0113**		-.0052**	
Enterprise	(.0037)		(.0070)		(.0028)		(.0014)	
Government	-.0583		.0295		.0075		.0345	
	(.0371)		(.0843)		(.0569)		(.0325)	
Fixed Effects (number of groups)	Occ*Ind* Province (22,909)	Occ*Firm, Province (260,322)	Occ*Ind* Province (22,909)	Occ*Firm, Province (260,322)	Occ*Ind* Province (22,909)	Occ*Firm, Province (260,322)	Occ*Ind* Province (22,909)	Occ*Firm, Province (260,322)
R^2	.090	.650	.155	.692	.170	.659	.217	.668

** p<0.01, * p<0.05. 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. The omitted firm type is for-profit firms with no government or foreign connection. Sample size = 1,057,369 ads.

Table 8: Effects of Advertised Wages on Gender and Other Preferences

	Dependent Variable:					
	(1)	(2)	(3)	(4)	(5)	(6)
Fixed Effects Included in Regression:	Relative Valuation of Men	Tendency to Gender-Target Ads	Relative Valuation of Old Workers	Tendency to Age-Target Ads	Request Beauty?	Height Requirement?
1.Occupation*Industry* Province	.0043 (.0053)	-.0497** (.0040)	.0256** (.0059)	.0191 (.0106)	-.0188** (.0042)	-.0026 (.0027)
2. Occupation*firm, Province	.0150 (.0089)	-.0405** (.0058)	.0098 (.0061)	.0136 (.0132)	-.0180** (.0055)	-.0097* (.0041)

Notes: ** $p<0.01$, * $p<0.05$. OLS estimates. Sample size = 173,945 ads.

All regressions also control for education and experience requirements, the number of vacancies advertised, a dummy for part-time jobs, and period fixed effects. Row 1 regressions also control for firm size and ownership. Standard errors are clustered at the occupation*province level. All wages are measured in logs.

Figure 1: Occupation Fixed Effects, by Mean Education Requirements

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

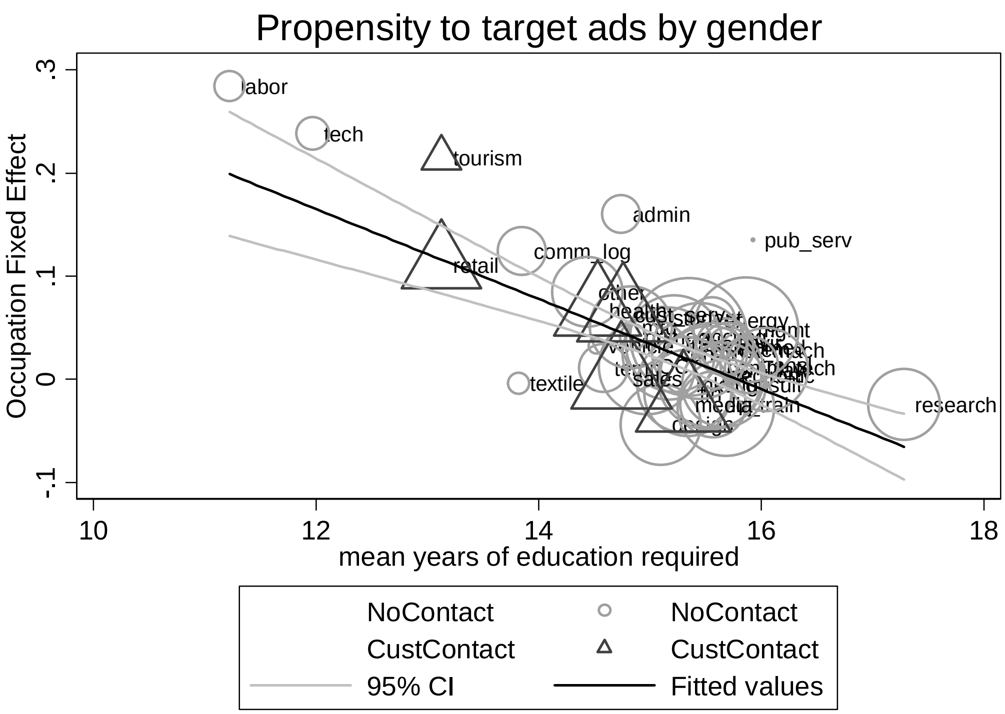

b) Dependent Variable: Employer's Relative Preference towards Men (P^M - P^F)