Employers' Preferences for Gender, Age, Height and Beauty: Direct Evidence

Peter Kuhn[a]
University of California, Santa Barbara

Kailing Shen[b]
Xiamen University

November 18, 2009

We study firms' advertised preferences for gender, age, height and beauty in a sample of ads from a Chinese internet job board, and interpret these patterns using a simple employer search model. We find that these characteristics are widely and highly valued by Chinese employers, though employers' valuations can be highly specific to occupations and jobs. Consistent with our model, advertised preferences for gender, age, height and beauty all become less prevalent as job skill requirements rise. Cross-sectional patterns suggest some role for customer discrimination, product market competition, and corporate culture. Using the recent collapse of China's labor market as a natural experiment, we find that firms' advertised education and experience requirements respond to changing labor market conditions in the direction predicted by our model, while firms' advertised preferences for age, gender, height and beauty do not.

[a](corresponding author) Department of Economics, University of California, Santa Barbara, Santa Barbara, CA USA 93106, NBER and IZA, pjkuhn@econ.ucsb.edu.; [b]Wang Yanan Institute for Studies in Economics (WISE), Xiamen University, China 361005 and IZA, klshen@xmu.edu.cn
The authors thank UC Santa Barbara's Institute for Social, Behavioral and Economic Research (ISBER) for financial support. Kuhn thanks UC Berkeley's Center for Labor Economics for their generous hospitality during the research on this project. Catherine Weinberger, and seminar participants at Xiamen University, the Chinese University of Hong Kong, the Trans-Pacific Labor Seminar, the all-UC labor economics conference, UC Denver, and UCSB supplied helpful comments.

1. Introduction

Over the past half century, economists have devoted considerable effort to measuring the strength and source of employers' preferences for certain demographic attributes of their employees, including race, gender, age, and physical attractiveness. A well known difficulty with measuring these preferences is the fact that it is either illegal, or highly controversial, for employers to publicly express them in many countries.[1] These data limitations have led researchers to pursue a variety of indirect strategies, including the study of wage regression residuals (Oaxaca 1973), studies of the evolution of returns to cognitive ability over workers' careers (Altonji and Pierret 2001), and studies that use aggregate data on discriminatory attitudes (Charles and Guryan 2008). Alternatively, some investigators have turned to resource-intensive audit studies (e.g. Yinger 1986, Neumark 1996), which typically focus on very specific jobs and industries.

The starting point for this paper is the observation that this fundamental data problem in the economics of discrimination does not apply to the world's largest labor market: China's.[2] This allows us to measure employers' preferences more directly, by studying posted requirements for a worker's gender, age, height and physical appearance in a large sample of ads on an internet job board.[3] Interpreting these advertising decisions in light of a simple employer search model, we draw a number of inferences about the magnitude, pattern and likely sources of employers' preferences for these attributes. We find, first of all, that employer preferences for these attributes are widely advertised in China: 34 percent of firms that advertised on the board during our five-month sample period placed at least one ad stipulating a preferred gender; 47, 29 and 10 percent respectively expressed a preference for age, physical attractiveness, and height respectively. Ninety percent of firms who placed 50 or more ads expressed a preference for at least one of these characteristics. We also find a high degree of job-specificity in firms' 'discriminatory' ads: it is commonplace for firms to require, say, beauty for some jobs, but not for others, or to solicit men for some jobs and women for others. Since, even for conservative parameter values, our model suggests that discriminatory ads should only be used when employers' group-based preferences are particularly intense, we conclude that the above characteristics are widely *and* highly valued by Chinese employers.

[1] While gender discrimination is prohibited in the US, age discrimination is barred by the 1967 Age Discrimination in Employment Act only for workers over 40 in firms with at least 20 workers. A handful of U.S. jurisdictions, including Michigan and California, prohibit discrimination on the basis of height, weight and/or looks (Hamermesh 2010); obesity is partially covered by the Americans with Disabilities Act.

[2] The only Chinese anti-discrimination law of which we are aware is contained in the August 2007 *Law of Employment Promotion*, which states: "When recruiting, employers, except for jobs or positions claimed to be unsuitable for females by the country, should not refuse to employ women based on gender or raise the employment standard for women. When recruiting, the employers should not set rules restricting females' marriage or fertility." (chapter 3, Rules 25-28). No sanctions appear to be prescribed for violations.

[3] To our knowledge, the only other article to have studied explicit discrimination in job ads is Lawler and Bae (1998). Their focus is on the impact of a multinational firm's home country culture on its overt gender discrimination in a sample of 902 ads placed in an English-language newspaper in Thailand. Darity and Mason (1998) reproduce examples of discriminatory U.S. job ads from before the Civil Rights era, but do not conduct a statistical analysis. Banerjee et al. (2009) study caste and other preferences in a small sample of marriage ads in India.

Consistent with our model, we also find that Chinese firms are more likely to express preferences for each of these "US-prohibited" attributes when education and experience requirements for the job are low. This finding suggests that our internet-based sample of ads for highly skilled jobs may understate employers' preferences for "US-prohibited' characteristics in China as a whole. Cross-sectional patterns in our data also suggest some role in job ads for customer discrimination (especially for beauty and height among women); for product market competition (state-owned enterprises discriminate more); and for corporate culture (foreign-owned firms discriminate much less). Finally, using the recent collapse of China's labor market --which occurred at different rates in different occupations and provinces-- as a natural experiment, we find that firms' posted education and experience requirements respond to changing labor market conditions in the direction predicted by our search model, while firms' advertised preferences for age, gender, height and beauty do not.

2. Related Literature

a) Discrimination

As noted, economists have studied gaps in labor market outcomes related to gender (e.g. Bertrand et al. 2009), age (e.g. Lazear 1979), beauty (e.g. Hamermesh and Biddle 1994) and height (e.g. Case and Paxson 2008) for some time. Possible explanations for such gaps include the tastes of employers, co-workers and/or customers (Becker 1964, Mobius and Rosenblat 2006, Charles and Guryan 2008), and differences in the amounts and types of skills held by these groups.[4] To the extent that the above characteristics are used as a signal of employee skills that are not measured at the time of hire, such employer behavior would constitute what is commonly called statistical discrimination.[5]

As also noted, attempts to measure the amount of taste-based and/or statistical discrimination face significant obstacles in labor markets where such practices are frowned upon or illegal. Not only is it difficult to distinguish taste-based from statistical discrimination; it is also difficult to differentiate both types of discrimination from a scenario where firms do not use these criteria at all, but some groups outperform others in a group-blind hiring process. These difficulties have led some investigators to conduct audit studies, in which matched pairs of interviewees (actors) of different race or gender respond to ads and attend interviews (see for example Yinger 1986, Neumark 1996). Well-known concerns with audit studies include imperfect matching (and training) of actors, the inability to double-blind (auditors may wish to please the researcher, and do not care about actually landing the job), expense (which limits both sample size and the

---

[4] See for example Niederle and Vesterlund 2007 and Black and Spitz-Oener 2007 regarding gender differences in skill mix and personality; Cerella 1991, Papalia 2002 and Charness and Villeval 2009 for age differences; and Case and Paxson 2008 regarding height, health and cognitive ability.

[5] Also, we note two additional possible explanations for group-based hiring: (1) monopsonistic wage discrimination by employers (Oaxaca and Ransom, 2008), and (2) –specifically for the case of age— deferred-compensation contracts that might make firms unwilling to hire older workers at the same wage offered to new junior employees in the same position (e.g. Lazear 1979). Whether the latter truly constitutes 'discrimination' is open to debate.

scope of studies, which are typically focused on narrowly-defined occupations), and the need to deceive employers, which raises issues of informed consent.

In the study that is probably most closely related to ours, Bertrand and Mullainathan (BM) (2004) essentially conduct just the first part of an audit study, which occurs before the actual interview. BM submit resumes to employers in response to newspaper job ads, with the respondent's apparent race randomly assigned via race-specific first names. They then study the callback rates for these resumes. Advantages of this approach include the fact that resumes can be 'perfectly' matched, the absence of actors and interviews (which greatly limits the scope for experimenter demand effects and reduces expenses), and less employer deception than audit studies. Remaining concerns with BM's methodology include the question of whether these resumes are seen as realistic by employers (given the level of racial segregation in most U.S. cities, is it plausible that a black applicant attended the same school as his matched white applicant?), the fact that race is only inferred (persons with distinctively-black names may not be representative of all black applicants), the fact that callbacks do not imply a job offer, and the absence of data on the population of 'real' resumes received by the firm for the position, either from the ad that was responded to or via other recruitment channels.

A key difference between our study and BM's is our focus not on employer responses to identical resumes, but on employer preferences that are advertised before any applications arrive. One obvious advantage of this approach is that no resumes need to be constructed or submitted; thus their comparability, realism and representativeness are not at issue. This also makes our approach much cheaper; our data is collected automatically by a web crawler, thus our sample consists of over 500,000 observations, compared to about 5,000 in BM and much smaller numbers in most audit studies. This sample size, in turn, lets us study not just a few specific occupations but the population of all jobs on offer on this website. We can thus compare discrimination across occupations, industries, firm types, time periods, labor market conditions, and worker characteristics (gender, age, height and beauty), and study some ways in which these conditions and characteristics interact.[6]

This said, it is important to note that our approach measures a different aspect of employer discrimination than BM's. One difference is that discrimination in our data is expressed *ex ante*, i.e. before the employer learns the contents of any applicant's resume. Thus, we measure employer preferences (or estimates/perceptions of expected applicant quality) that are not conditional on the characteristics observed in a resume. To the extent that the disfavored group has, on average, lower quality observables, our measure should therefore indicate higher levels of discrimination than conditional-on-observables indicators like BM's. On the other hand, while BM's measure can, at least in principle, detect small employer preferences for one group over another, we observe discrimination only when the employer does not want to see *any* applications from the disfavored group.

[6] Our approach also shares some limitations with BM and with audit studies. For example, like them, we do not know how many (or what types of) applications these firms received, either in response to the ads we observe, or through other channels. Thus, for example, it is possible, though we think rather unlikely, that firms with affirmative-action type motivations request (say) female applicants in Zhaopin ads because of a shortage of female applications from other recruitment channels for the same job.

As our model shows, employers' preferences might need to exceed quite a high threshold for this to optimal; for this reason we should see discrimination *less* frequently using our method than using BM- and audit-type measures. Importantly, our model provides testable predictions about the conditions under which employers will wish to impose such 'blanket' restrictions, which differ from the type of discriminatory behavior that is typically measured in U.S. studies.

Our study also contributes to a literature on the effects of market competition on discrimination, which dates back to Becker (1957). Interestingly, both Becker's original model, and recent empirical studies (for example Black and Strachan 2001 and Black and Brainerd 2004) focus on the effects of competition in *product* and *capital* markets on the demand for discrimination. Empirical study of the effects of labor market conditions, beginning with Freeman's (1973) classic analysis of black wages and unemployment, to our knowledge, has been confined to the study of time-series patterns in black versus white unemployment rates in the U.S.[7] We contribute to this literature, first, by providing a simple search model within which to interpret the effects of labor market conditions on firms' demand for discrimination, and by providing a new empirical approach, based on job ad content.

b) Employer Search

The main outcome of interest in the current paper is a firm's decision, when posting a job advertisement, to specify that the applicant should be of a particular gender, age, or have some other characteristic. Aside from possible differences in the firm's motivation, this is conceptually the same as a practice that is universally accepted in job advertising: specifying that the applicant should, for example, have a high school degree, or at least two years of experience. Indeed, our theoretical and empirical analysis treats all advertised job requirements isomorphically. Since advertised job requirements have received surprisingly little attention in the literature on optimal employer search and recruiting, our paper contributes to that literature as well.

Recent treatments of employer search or recruitment take two main approaches. The first, inspired among others by Mortensen and Pissaridies (1994), models search equilibria in which both workers and firms optimize. In these models, employers' strategy space is typically highly restricted; for example in many models the only choices firms make are whether to enter the market and what wage to post; sometimes firms can choose to advertise wages or working conditions (e.g. Moen, 1997; Menzio 2007). Further, in most of these models, workers are identical.[8] Thus, to our knowledge, the equilibrium search literature has not yet considered the possible optimality of advertised hiring restrictions of the type documented in this paper, which invite some (but not all) worker types to apply.

[7] For a recent update to Freeman's analysis, see Couch and Fairlie (2005).
[8] For an excellent recent exception, see Teulings and Gautier (2004). Lang, Manove and Dickens (2005) develop a posted-wage search model with two worker types; they show that small employer preferences for one group can generate large equilibrium differentials in outcomes. Rosen's (1997) model has two worker types and assumes workers are better informed about match quality than firms; discriminatory equilibria can arise even when firms have no group-based preferences.

A second branch of the employer search literature looks more closely at the dynamics of employee recruitment using data on individual vacancies (e.g. Burdett and Cunningham 1998). A key question here is whether vacancy durations are better described by a sequential employer search model (where firms set a reservation level of worker quality, then simply wait until an application exceeding that quality arrives), or a nonsequential one (where a large batch of applications arrives shortly after a vacancy is advertised, and vacancy durations largely consist of the time taken by the firm to select the best candidate from this pool).[9] Van Ours and Ridder's (1992, 1993) results tend to favor the nonsequential model. Van Ommeren and Russo (2009) make the subtler point that a nonsequential model best characterizes the type of employer search (i.e. publicly advertised vacancies, as opposed to less formal methods) that is studied in this paper. For this reason, we use a simple nonsequential model here.

In contrast to the general equilibrium literature, two papers in the above 'micro' literature --Barron, Bishop and Dunkelberg (1985) and Van Ours and Ridder (1991)— have studied employers' advertised job requirements empirically. However, unlike us, both of these papers treat advertised requirements as exogenous vacancy characteristics (that might affect vacancy durations), rather than as a choice variable for the employer. Thus, to our knowledge, ours is the first paper to write down a theoretical model of employers' decisions regarding what job requirements to advertise (whether 'discriminatory' or not), and the first to study the determinants of this decision empirically.

3. A Model

Consider a firm inviting applications for a vacant position; applications can come from two distinct groups labeled A and B, where A and B also represent the number of applications that would be received from each group. Under what conditions will the firm prefer to invite only one of these two groups to apply? To address this question, let the value to the firm of an individual applicant, j, be given by $U_j = v^A + \varepsilon_j$, and $U_j = v^B + \varepsilon_j$ for groups A and B respectively, where the ε_j represent independent draws from the same *cdf*, $F(\varepsilon_j)$.[10] For the job in question, we assume that members of group A are on average more highly desired, i.e. that $v^A > v^B$. The firm is assumed to choose the worker with the highest total value, U_j, from its pool of applicants. The cost of processing an application (thus learning its ε_j) is c per application.[11] All of our analysis assumes what

[9] See Morgan (1983) for a model that combines elements of both approaches.

[10] Importantly, in this formulation, the worker's total value to the firm, U_j, includes both his/her actual productivity and any discriminatory tastes or inaccurate perceptions the employer might hold. It should also be interpreted as net of wages paid; thus for example v^A will be higher if the As can be paid a lower wage than the Bs for the job that is advertised. This noted, for brevity we often refer to U_j (and related parameters such as v^A and v^B) simply as workers' "productivities" in this section.

[11] Note that application processing costs (as distinct from the vacancy costs more typically modelled, such as a fixed cost of opening a vacancy and the opportunity cost of keeping the job vacant) 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. Note also that application processing costs will, in general, be affected by different factors (including, for example, the firm's information-processing

we call "free random disposal" of applications; this means that a firm can avoid processing costs on any random sample of applications by discarding them.[12]

We begin by denoting the expected maximum value of U_j in a pool of n applicants with 'baseline" productivity $v \in \{v^A, v^B\}$ by $H(v, n)$; under quite general conditions $H_1 > 0$, $H_2 > 0$, and $H_{22} < 0$.[13] If the firm chooses a search strategy to maximize expected profits net of search costs, $H(v,n) - cn$, then it simply selects the highest level of profits from the following three cases:[14]

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

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

Because it is difficult to make useful comparisons between the G and H functions for an arbitrary underlying *cdf* of match qualities $F(\varepsilon_j)$, (and because we wish to parameterize certain features --in particular the variance-- of this distribution and calibrate the model numerically), we assume that ε_j follows a Type-1 extreme-value distribution, which yields a simple closed-form expression for the expected value of the maximum in any given number of draws. In particular, we assume that $F(\varepsilon_j) = exp(-exp(-\varepsilon_j/\beta))$; thus $\mathrm{Var}(\varepsilon_j) = \beta^2\pi^2/6$, and $\mathrm{E}(\varepsilon_j) = \beta\gamma$, where γ is Euler's constant (≈.577). Under these assumptions, Appendix 1 shows that:

$$H(v^J,A) = v^J + \beta\gamma + log(J) \equiv \mu^J + \beta log(J), \text{ and} \tag{1}$$

$$G[v^A,A;\ v^B,B] = \mu^C + log(M), \tag{2}$$

where $\mu^J \equiv v^J + \beta\gamma,\ J \in \{A, B\}$, is the expected productivity of a type-J worker, $M=A+B$, and:

$$\mu^C = \beta \log\left[\delta \exp\left(\mu^A/\beta\right) + (1-\delta)\exp\left(\mu^B/\beta\right)\right]. \tag{3}$$

technology and the complexity of the methods used to evaluate job candidates) than the vacancy costs that are typically included in search models.

[12] More precisely, when the firm disposes of a subset of applications in a mixed pool of As and Bs, we assume that share of the two groups in the deleted samples equals their share in the population. Thus, while disposal is random with respect to the ε_j's, there is no randomness in the shares of the two types in the sample of retained applications. In the absence of free random disposal, there are conditions (specifically when applications are plentiful overall and the less-productive group (B) is smaller in number), where the firm's optimal policy is to invite only group B to apply, simply to save on processing costs. We do not think this is very likely.

[13] See for example Stigler (1961), p. 215.

[14] Our analysis assumes that applicants comply with firms' announced preferences by not applying where they are not 'wanted'. This is incentive compatible for applicants if there is any cost to submitting an application.

In (3), $\delta = A/M$ is the fraction of group A (the more 'desired' group) in the combined pool.

As noted, a comparative static of interest to us is the effect of expanding the supply of applications on firms' preferences among strategies A through C. Supply is parameterized by M, the number of applicants in the combined pool; since we are considering applications for a single vacancy, M is effectively the unemployment/ vacancy (u/v) ratio facing a single firm. We wish to characterize how the firms' preferences among strategies change as we expand M starting from its smallest possible value, keeping the shares of the As and Bs constant, i.e. holding $A=\delta M$ and $B=(1-\delta)M$. The predicted effects of M are summarized in:[15]

Proposition 1: If $\mu^A - \mu^C + \beta log(\delta) + c(1-\delta)/\delta > 0$, firms will invite only the favored group (A) to apply at all levels of application supply, M. Otherwise, there exists a critical value of M, denoted by $\tilde{M}$, above which firms invite only the As to apply, and below which firms invite applications from all.

Proof: In Appendix 1.

Proposition 2: $\tilde{M}$ falls (making Strategy A more 'likely') when:

-the baseline productivity difference between the groups, $v^A - v^B$, rises;
-application processing costs, c, rise;
-the standard deviation of unobserved worker productivities, $\sigma = \beta\pi/\sqrt{6}$, falls.

Proof: In Appendix 1.[16]

Next, we define a neutral increase in a job's skill level as one that multiplies the productivity, U_j, of every worker by the same factor, $\theta > 1$. Appendix 1 also shows that:

Corollary: A neutral increase in a job's skill level, θ, raises $\tilde{M}$, making Strategy A less likely.

The intuition is that, as skill demands rise, workers' idiosyncratic ability 'matters more' to job performance (see for example Gibbons and Waldman 1999); thus the option

---

[15] To focus on nontrivial cases, all of our theoretical results assume that $M^* > 1/\delta$, where M^* is the (common) number of applications that maximizes the firm's profits under all three strategies (A, B, and C). This simply states that applications are scarce at the minimum level of market supply when there is only one applicant of the "preferred" type available, i.e. when A=1.

[16] We have shown that Proposition 1 is robust to various changes in assumptions. For example, it must hold in the following sense for any processing cost function that satisfies the second order condition for an optimal number of applications N^* holds: If N^* exists, then any increase in labor supply (M) beyond N^* increases the likelihood that firms will advertise hiring restrictions. Also, it is reasonable to ask what happens if we introduce a distinction between two aspects of application processing: (a) simply screening applications for the presence of the preferred or dispreferred demographic attribute, and (b) determining an applicant's match quality, ε. Even if (a) can be done very cheaply, we argue that firms will still choose to advertise their preferred worker attributes, since advertising the restriction has the same effect as screening and costs even less (typically nothing for internet job ads).

value of searching within the disfavored group rises. In sum, our model predicts that firms' decisions to target their job ads to a particular demographic group should become more frequent (1) the greater the number of applications the firm expects to receive (higher M); (2) the greater the firm's preference for one demographic group over another (v^A - v^B); (3) the higher are application processing costs (c); (4) the lower is the dispersion of unobserved worker productivity (β); and (5) the lower is the skill level of the job (θ).

We conclude our theoretical discussion with an analysis of expected magnitudes. Essentially, our model says that firms should advertise that they are not interested in receiving applications for a job from a specific group (say women, or applicants with less than high school education) when the firm's *ex ante* assessment of the chance that the best overall candidate will come from that group is so low that the expected benefits of examining that group's applications are outweighed by the costs of processing them. If marginal processing costs are low, and if the variance of unobserved productivity is not too low, we should therefore expect to see advertised hiring restrictions relatively rarely, even in the presence of strong preferences for the favored group, i.e. high $v^A - v^B$.

To illustrate this property of our model more formally, define $[\mu^A - \mu^B]^*$ as the critical level of $\mu^A - \mu^B$ above which a firm prefers to target its ad at group A only; $[\mu^A - \mu^B]^*$ is implicitly defined by equation (A4). Normalize the expected net productivity (over his entire tenure with the firm) of a type A worker at μ^A =1, and let δ =.5. Next, note that equation (A4) depends not on c and M individually, but only on their product, cM, i.e. on the total costs of processing all applications that arrive when no restrictions are imposed. To our knowledge, the only study that quantifies these costs is Barron and Bishop (1985). Specifically, their Table 1 reports that the *total* man-hours spent by company personnel recruiting, screening, and interviewing applicants to hire one individual was between 10.60 and 16.99 hours for occupations typical of ads on Zhaopin (professional/technical, managerial, clerical and sales).[17] Assuming that the expected tenure of a newly hired worker is one year[18] and assigning the same value of time to the hirers and hirees, this implies a ratio of total hiring costs to expected employee value of $[15/(40 \cdot 4.33 \cdot 12)] = .0072$ for the preferred group. Finally, the remaining parameter we require is β, or equivalently, σ—the standard deviation of idiosyncratic worker net productivity. As a starting point, suppose that the 95th percentile of the type-A net productivity distribution is twice the mean; this implies that σ=.75.

The top left entry of Table 1 indicates that, in the base case scenario described above, the less favored group (say, women) would need to have an annual expected net productivity of *minus* 1.97 (compared to men's expected net productivity of one), or an expected productivity disadvantage of at least 297 percent, for it to make sense for firms

[17] A handful of other studies provide data on the number of applicants per position, M on its own. For example, Barron and Bishop 1985, and Barron, Bishop and Dunkelberg 1985 report means of about 6 and 9 respectively for the U.S. in the early 1980s; van Ours and Ridder (1992) report a mean of 11 in the Netherlands in 1986/7). While it seems likely that the number of applicants to internet job ads might greatly exceed these values, we also note that the internet also greatly reduces the marginal cost of processing ads, c. Thus, it is not all clear that our estimate of cM – the total cost of processing all ads that arrive—is very different for our Chinese internet sample than in these pre-internet studies.

[18] This corresponds roughly with the mean duration of all uncensored jobs (13.7 months) in Farber's (1994) study of inter-firm worker mobility.

to exclude women from the application process.[19] Referring back to Proposition 2, this very high level of $[\mu^A - \mu^B]^*$ would only increase further if total application processing costs (cM) were lower (or equivalently if a typical worker stayed at the firm for more than a year), or if σ was higher than these base-case values. Thus, we only explore the sensitivity of our results to parameter changes in the other direction in Table 1.

Row 1 of Table 1 shows, first, that if instead of taking a total of 15 employee hours to recruit for a position that will be occupied for one year, it took about a week (.02 ≈ 1/52), the massive critical productivity disadvantage falls from 297% to a still-massive 237%. The remaining entries in Row 1 imagine that recruiting costs for a one-year position amount to one month and two months respectively. Even in the latter case (column 4), it still makes sense to exclude women from consideration for the job only when they are 84% less valuable than men on average, or worse. Column 1 examines the consequences of reducing the idiosyncratic variation in worker net productivity, σ, while holding application processing costs, cM, at their base-case value. When σ falls to .375, the 95th percentile male applicant is only 50 percent more valuable to the firm than the mean male applicant. Here, firms will wish to exclude women if their mean net productivity is 128 percent below men's or worse. This number falls to 59% when σ = .20, which corresponds to a ratio of the 95th percentile to the mean of 1.27.

The remainder of Table 1 shows that hiring restrictions against the less-favored group can be in the firm's interest for relatively low (say, under 20%) values of the expected net productivity differential between the two groups only if (a) it is highly costly for firms to process applications *and* (b) employee performance is highly predictable, with the 'best' employees little better than the average.[20] Overall, we conclude from Table 1 that, for what seem to be realistic costs of employee selection and realistic levels of employee heterogeneity, employers would need to expect very large productivity differentials between the groups --or to have very strong tastes for one group over another-- in order for it to make sense to exclude one group from the application process. On *a priori* grounds, we thus expect such restrictions to be relatively rare.

4. Data and Descriptive Analysis

Our data is the universe of unique job advertisements posted between May 16 and July 29, 2008 and between Dec 17, 2008 and Feb 28, 2009, on Zhaopin.com, the third-

---

[19] At this point, it may be worth recalling that between-group differences in μ in our model include both employer tastes and employer expectations of actual productivity. Further, they are net of wages paid. This does not affect the interpretation of our results if firms post wages, i.e. if they pay group-independent wages within the job that is advertised. If, on the other hand, firms can pay lower wages to the disfavored (B) group than the A's within the same job, then 297% actually understates the firm's tastes for (or perceived productivity advantage of) the favored group.

[20] We also examined the effects of changes in the relative numbers of the two groups, δ, on $[\mu^A - \mu^B]^*$. Varying δ from .05 to .95 did not change our baseline results materially. This is because a small expected number of (say) female applicants, (a) reduces the expected costs of processing women's applications, but (b) reduces the likelihood that this (small) additional pool will yield the best candidate. These two tendencies offset each other.

largest online job board in China.[21] Information on the advertising firm's characteristics was added to the data by following links to the firms' websites that were provided in the ads. Procedures for downloading the data and defining variables are discussed in Appendix 2. Clearly, ads on Zhaopin.com will not be representative of all jobs in China; like all samples of job ads they will overrepresent jobs in expanding and high-turnover occupations and industries. In addition, the jobs on Zhaopin.com likely require a significantly higher skill level than the median job in China. Since, according to our own data, discrimination according to gender, age, height and beauty is more common in less-skilled occupations, our data likely underestimates its extent in China.

Another sampling issue arises from the fact that some of the ads in our data are for multiple vacancies.[22] While most of our analysis treats the job ad as the unit of analysis (using the number of vacancies it advertises as a control variable of interest) we report some specifications which weight the sample by the number of vacancies. 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 typically unknown in these cases, we restrict our sample to ads for a single occupation; this reduces our sample by about 20 percent.[23] Finally, we note that, by construction, our data is a 'stock' sample of "ads in progress" rather than a flow sample of newly-posted (or just-filled) ads. This implies that long vacancy spells are overrepresented in our data, which would affect our estimates if there is parameter heterogeneity that is correlated with vacancy durations. To address this concern, we replicated our estimates for a subsample of ads that consists, almost certainly, of newly posted ads.[24]

Descriptive statistics for our data are provided in Table 2. All told, we study a total of 633,664 job ads; as expected these tend to be for relatively highly-skilled jobs: about 70 percent of ads require at least some post-secondary education.[25] Seventy percent require some experience, with the modal experience requirement between one and three years.[26] One in ten ads expressed a gender preference; this was evenly split

[21] Firms frequently re-post the same ad; as detailed below our sample treats such 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 renewals.
[22] The most extreme case was a single ad for 8000 vacancies at a newly-opened chemical plant; our results are not materially affected by excluding this observation.
[23] 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. Importantly, we do not exclude ads for multiple occupations when we calculate our indicators of labor market conditions (for example the number of competing ads in an occupation/region/period cell), which are the key regressors of interest in Section 5 of the paper.
[24] Appendix 2 describes the creation of this "inflow sample". See Bergeron et al. (2009) for a recent discussion of the effects of length-biased sampling.
[25] By far the most common occupation (of the 38 categories used by Zhaopin on its site) is sales, at about 20 percent of the ads, with computer-related occupations second at about 9 percent. The top five industries were consulting, IT service, construction, software and internet/e-commerce. 26 and 18 percent of the ads were for jobs in Beijing and Shanghai respectively, but all Chinese provinces are represented in our data.
[26] Additional information about skill requirements, as well as about how firms respond to labor supply shocks, might of course be gleaned from the wage levels attached to jobs advertised on Zhaopin. Unfortunately, only about one eighth of the ads in our sample contained any usable information about wages. Since this small a sample is likely to be highly unrepresentative, we decided not to include wages in our analysis of ad contents.

between men and women.[27] About one in four ads expressed an age preference; perhaps surprisingly, minimum age requirements were almost as common as maxima. Indeed, 12 percent of ads specified both a minimum and a maximum age, and fewer than one percent of advertised minima were at age 16—the lowest observed in our data. That said, firms seem overall to be seeking younger workers: the mean minimum and maximum ages were about 25 and 35 respectively, and fewer than 5 percent of the maximum ages were over 45.

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), none of the 633,664 ads in our sample requested short or unattractive applicants. All together, 8 percent of ads requested that the applicant be physically attractive ("*xingxiang*"). Given the relative rarity of gender restrictions, a large majority of the requests for beauty in our sample were in ads that did not express a gender preference. Among ads that did express a gender preference, however, beauty was much more valued among women than men: fully one out of every three ads targeted at female applicants required the applicant to be physically attractive, compared to 6.7 percent for ads targeted at men. This contrasts starkly with some U.S. studies (e.g. Hamermesh and Biddle 1994), which did not find that the labor market returns to beauty were highly 'gendered'.

At two percent of job ads, height requirements are the least common of the 'US-prohibited" job requirements we study here. That noted, a fascinating aspect of these height restrictions is that, when they exist, they are *always* gender-specific. In some sense, this is true by definition in ads that are gender-targeted. More surprising is the fact that the majority of height requirements --which appear in ads without gender restrictions-- explicitly list two, gender-specific height requirements, with the female about 8 cm (3.15 inches) shorter. This suggests that these height requirements are not associated with specific job tasks (e.g. operating a particular machine), but either used as signals of health or cognitive ability (as argued by Case and Paxson 2008), or valued in themselves, perhaps by a firm's customers.[28]

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 35 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 eight 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 federal) public service; while these are a very small share of the total, our large sample size allows us to study them as well.

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

[28] 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.

One advantage of working with job board data is the fact that our data is a census of the ads posted there. Thus, to the extent that ads on Zhaopin provide a representative picture of all vacancies that are relevant to workers who search on line, we can use our data to characterize the labor market environment each ad inhabits. To that end, the final panel in Table 2 shows our baseline indicators of the labor market conditions facing each ad, computed at the level of the occupation/province/period cell.[29] Note that all three of our empirical indicators actually measure excess labor *demand*, or labor market tightness (i.e. v/u); thus they are inversely related to M in our theoretical model. The first of these (*MRR*, or the mean renewal rate) is the mean number of times an ad in that cell was renewed. "Renewing" an ad means simply re-posting the same ad on Zhaopin on a day after it was first posted; in consequence it will be listed as "posted today" on the site. Since it is essentially free for employers to renew ads, employers do so very frequently and workers pay little attention to ads posted more than a few days ago. Thus, the number of renewals is probably the best available indicator of a vacancy's duration in our data.[30] On average, an ad in our data was renewed about 13 times; this declined by 27 percent (from about 15 to 11) between our two observation windows, suggesting a considerable decline in excess demand for labor. This is supported by our two other indicators of labor market tightness, namely the total number of competing ads (*NCA*) in an occupation*province cell, or the total number of competing vacancies (*NCV* = *NCA* times the mean number of positions per ad in the cell), which declined between our two periods from 4499 to 2867 and from 19185 to 12275 respectively, which is 36 % by either measure.

Interestingly, means of the US-prohibited job requirements in the first part of Table 2 change very little between periods despite this large decline in competition for workers. Advertised education requirements fell, but both of these trends were almost certainly affected by the nature of the demand shock that hit Chinese labor markets in this period, which disproportionately affected highly-skilled occupations and the prosperous eastern provinces.[31] A complete analysis of the effects of labor demand shocks would net out such economy-wide factors, as we do in Section 5 of the paper.

A different picture of the frequency of 'discriminatory' job ads in China emerges when we organize our data by firms rather than ads. Overall, 63,507 distinct firms placed ads on Zhaopin during our sampling period. The median (mean) (maximum) number of ads placed per firm was 4 (9.97) (6,316); the median (mean) (maximum) number of

---

[29] With 38 occupations, 31 provinces, and 2 periods, there are 2356 possible cells. Later in the paper, we also provide results for labor market indicators defined at the occupation*province*industry*period level.

[30] In principle, it would be possible (if we resampled ads from Zhaopin) to compute the date when a firm's ad actually disappeared from Zhaopin.com, but conversations with Zhaopin officials suggest that this would yield a considerably poorer measure of the vacancy's duration, since many firms simply allow the vacancy to remain on the site after it is filled. In fact, this practice is encouraged by Zhaopin, because it raises the job ad count Zhaopin can advertise to prospective applicants.

[31] For example, the correlation between an occupation's mean educational requirement and the between-period change our three measures of the supply of applications (MRR and the log of *NCA* and *NCV* respectively) is -.27, -.32 and -.30 across our 38 occupations. Across China's 31 provinces, the correlations between these three labor supply measures and the province's initial log GDP per capita were -.60, -.74 and -.63 respectively. Thus, the declines in excess labor demand were largest in skilled occupations and high-income provinces.

distinct occupations advertised for over our entire sampling period (out of a possible 38) was 2 (3.42) (36). Characteristics of these 63,507 firms' hiring policies are summarized in Table 3. According to Table 3, 18.7 percent of the firms in our data placed at least one ad that invited only men to apply; for women this number was 23.5 percent. 29 percent of firms requested beauty in at least one of their Zhaopin ads, and 9.7 percent advertised a height requirement. Sixty-one percent placed an ad specifying at least one of the four 'US-prohibited' attributes in our data (gender, age, height and appearance). 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 than 50 ads, about half directed at least one of their ads specifically at women; half did the same for men, and 36 percent placed both male-only and female-only ads during our sample period. Ninety percent of these firms requested a US-prohibited attribute in at least one ad.

Together, the statistics in Table 3 thus illustrate two additional features of advertised discrimination in China. First, a large share of employers engage in what some Western observers would call explicit discrimination. Second, firms' discriminatory preferences appear to be strongly tied to specific jobs and occupations within the organization. This suggests that firms' relative preferences for groups ($v^A - v^B$ in our model) likely vary not just in size but in sign across jobs within firms.

Table 4 concludes our descriptive analysis of firms' advertised job requirements by presenting a correlation matrix among all of them. The pattern that emerges is clear. First, aside from the mechanical negative correlation (at the ad level) between a preference for men or women, all entries but one in rows 1-6 –which refer to the "U.S. prohibited" requirements-- are positive. The same is true of the correlation between the two "U.S. allowed" requirements (education and experience) in the bottom right of the table. Finally, consider the 2×6 matrix of correlations between these two sets of characteristics in the bottom left of the table. All but one of these 12 coefficients are negative and statistically significant (the one positive correlation is between experience and a preference for men). Thus, Chinese employers disproportionately advertise for US-prohibited characteristics in jobs where US-allowed job requirements are low. A particularly striking feature of the above pattern is that it holds for gender and age preferences in either direction: When skill level is measured by education and experience requirements, preferences for men and for women, and both minimum and maximum age requirements, are all more common in less-skilled than more-skilled job advertisements. Perhaps surprisingly, it follows that prohibiting 'discriminatory' ads in China would disproportionately impact employers of less-skilled labor.[32]

[32] In additional analysis available from the authors, we consider the likely impact of banning 'discriminatory' advertising in China by asking how accurately a jobseeker who observes all features of an ad on Zhaopin other than gender, age, beauty and height would be able to predict those requirements. Since all of these R^2s are low, the impact of prohibiting a characteristic on the 'information' available to jobseekers is roughly proportional to those characteristics' unconditional variance. This means that (a) considerable information would be lost from a ban, and (b) more information would be lost in ads for unskilled jobs, since 'discriminatory' ads are more common for those jobs.

5. Labor Market Conditions and the Demand for Discrimination

Our goal in this section is to estimate the causal effects of labor market tightness, as measured by *MRR*, *LCA* and *LCV*, on firms' advertised hiring standards, including 'discriminatory' standards, where *MRR*, *LCV* and *LCA* represent conditions in the ad's occupation*province or occupation*industry*province cell.[33] Since we include both period effects, and fixed effects for the above cells, only the differential within-cell temporal variation in market conditions between our two periods is used to identify the effects of labor market conditions on ad content.

A natural concern with this exercise is that, in our theoretical model, *M* represents labor market conditions that are exogenous to the firm which is placing an ad, while *MRR*, *LCA* and *LCV* in our data –which measure excess demand in the local labor market-- may also be correlated with unobserved determinants of the advertising firm's demand for labor. We approach this issue in several ways. First, while it is likely that periods in which our indicators of excess demand in the labor market, --*MRR*, *LCA* and *LCV*[34]—are high, will also be periods in which the firms posting ads in our sample experience positive product demand shocks, we note that this is a concern only if those unobserved product demand shocks affect the firm's relative demand for different types of labor. Unobserved demand shocks, if they impact the firm's desire to hire men and women, or attractive versus unattractive workers equally within an occupation, are not a concern.

Second, if unobserved industry demand shocks do affect the firm's relative demand for labor of different types, theory offers some guidance regarding this effect's direction: If good product market conditions operate primarily by raising the price of the firm's output (thus raising both μ^A and μ^B equiproportionately), then $\mu^A - \mu^B$ will tend to be higher when our indicators of excess labor demand are high. Thus, if anything, our estimates of the effects of *MRR, LCA* and *LCV* on the firm's demand for its 'preferred' type of labor (whether this is beautiful workers, or well-educated workers) will be biased in the direction of finding a positive effect. We will interpret our results with this in mind. Finally, later in this section we isolate the effects of labor market conditions for the advertised job that are driven by changes in demand outside the advertising firm's industry. These changes should be less correlated with changes in product demand at the firm placing the ad.

Table 5 presents the results of our baseline specification, using the mean ad renewal rate in an ad's occupation*province*period cell as our indicator of labor market conditions. Each column presents results for a different dependent variable; all regressions include fixed effects for occupation*province and for period. Except for experience, where --as predicted-- a high mean renewal rate reduces firms' advertised demands for experience, *MRR* has no statistically significant effect on any other ad characteristics, including the US-prohibited characteristics and the education

[33] Our regressions also control for a direct indicator of the likely number of applications per position (specifically the number of vacancies the firm is trying to fill with the ad). Although the variation in this indicator is less clearly exogenous to the advertising firm, its estimated effects are also consistent with our model.

[34] *LCA* and *LCV* denote the logs of *NCA* and *NCV* respectively.

requirement. It is worth noting that Table 5's estimated effects for the six US-prohibited characteristics are not only statistically indistinguishable from zero, but small in magnitude.[35] Finally, we note that the *p*-value for a chi-squared test of the hypothesis that all the US-prohibited attributes (columns 1-6) are unaffected by the mean renewal rate is .6515. For the hypothesis that neither of the US-allowed features of job ads (education and experience) were affected by *MRR/10*, this *p*-value is .0008.[36] This insensitivity of advertised 'discriminatory' employer preferences, and sensitivity of nondiscriminatory preferences, is mirrored in a wide variety of robustness checks below.

Before turning to those robustness checks, it is interesting to consider the coefficients on some of the cross-sectional 'control' variables in Table 5. Like *MRR*, the number of positions a firm is hoping to fill with the current ad is another indicator of our theoretical construct *M*, because a higher number of positions to be filled implies a correspondingly lower number of applications per vacancy. Thus we would expect firms hiring for a large number of positions to reduce their hiring standards, thus imposing fewer discriminatory restrictions and requiring lower levels of education and experience.[37] Table 5 indicates that this is indeed the case for education and experience, but not for the US-prohibited ad characteristics: Mirroring what we find for our 'market conditions variables (*MRR*, *LCA* and *LCV*), firms become choosier on education and experience requirements, but not on their preferences for attributes like age, height and gender, when workers are easier to find.

According to Becker's (1971) model of taste-based discrimination, firms operating in competitive product and capital markets should be less able to discriminate than firms in product markets that are sheltered from competition or takeover. At the same time, a number of authors, including Lawler and Bae (1998) have argued that corporate culture affects a firm's tendency to discriminate. For both these reasons it is interesting to look at the correlation of firm ownership types with advertised discrimination in our data. Consistent with the corporate culture hypothesis, firms with some foreign ownership require considerably higher levels of "US-allowed" qualifications (.85 years of education and half a year of experience) and are less likely than the reference category (Chinese private-sector firms) to express preferences for all US-prohibited attributes. The size, robustness and statistical significance of this effect on the use of US-prohibited attributes dwarfs any effect of local labor market conditions estimated in this paper. Like foreign-owned firms, Chinese State-Owned Enterprises require high levels of education and experience. Consistent with Becker's hypothesis (and with an employer taste for men and youth) however, they are more likely than

---

[35] For example, recall from Table 2 that the mean decline in *MRR/10* was about .4 (corresponding to a 27 percent decline in vacancy durations), and that the share of ads requesting men in our first sampling period was .053. Applying the coefficient and standard error for *MRR/10* in Table 5 indicates that with 95 percent certainty, this 27 percent decline in vacancies neither reduced firms' demand for men below .043, nor raised it above .060.

[36] These *p*-values adjust for clustering at the occupation-province level, as well as correlation of errors across the 10 equations in Table 5. See Stata's *suest* command.

[37] A competing explanation of such a pattern is that heterogeneity in the jobs to be filled might rise with the number of positions on offer, making a specific requirement less appropriate. This is at least partly addressed by the fact that we restrict our sample to ads for a single occupation only.

Chinese private-sector firms to prefer men, to disprefer women, and to specify a maximum age.[38]

A final set of cross-sectional 'control' variables of some interest are the fixed effects related to occupation. To recover aggregate occupation effects, we estimated a simpler version of the regressions in Table 5.[39] Figures 1-6 plot the resulting occupation fixed effects against the occupation's mean education requirement, for our six US-prohibited requirements. Occupations 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" ("beautyfit"). 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.

Figures 1-6 illustrate once again that advertisements with 'discriminatory' employer preferences are disproportionally found in occupations requiring low education levels: this is strikingly true for all such preferences. Inspection of the figures for outliers yields no obvious results for the age restrictions (except possibly a disproportionate advertising of age restrictions for higher management positions), but the following occupations seem to stand out as positive outliers in firms' demand for women, beauty and height: retail, tourism, beauty/fitness, "administration" and perhaps customer service. Of these, only administration was an occupation we did not classify as involving customer contact. Subsequent inquiries into the nature of jobs in the administrative category revealed that secretarial services are a major component.[40] We also note that tourism and retail seem to be negative outliers in their demand for men. In sum, an examination of cross-sectional occupational patterns suggests some support for the notion that customers, plus persons who hire secretaries, prefer to interact with tall, attractive women.[41]

[38] On the other hand, SOEs are less likely to request beauty and height and to specify a minimum age. We note that the Chinese public sector also disproportionately prefers young men. Our findings regarding gender, SOEs and the public sector are consistent with Meng (1998); Gustafsson and Li (2000)); Liu, Meng and Zhang 2000; and Zhang, Han, Liu and Zhao (2008), who find that the share of the unadjusted gender wage gap that is not accounted for by observable productivity-related characteristics in China is smaller in market-oriented activities than state-owned ones.

[39] Specifically, these regressions control for the number of positions advertised, firm size, firm ownership, part-time jobs, fixed effects for period 2 and for number of positions not stated, plus a full set of province and industry fixed effects. Measures of labor market conditions are not included because they do not have a clear interpretation in the absence of occupation*province fixed effects. Skill indicators, such as education, are not included so we can illustrate their effects in Figures 1-6.

[40] Confirming a strong employer preference for women in these positions, Bertrand and Mullainathan (2004) were forced to stop sending male resumes in response to administrative job ads due to extremely low callback rates.

[41] To test this hypothesis somewhat more formally, we re-estimated the regressions underlying Figs 1-6, adding controls for education and experience requirements, and replacing the occupation fixed effects by a single dummy variable for customer contact occupations. Customer-contact occupations were 2.0 percentage points less likely to request male applicants, 1.0 percentage points more likely to request female applicants, 3.0 and 3.4 percentage points more likely to specify a minimum and maximum age respectively, and 4.6 and 1.5 percentage points more likely to request beauty and height. All of these coefficients had *p*-values of below .001. Recalling that the sample means of beauty and height are 8 and 2.4 percent respectively, the effects are clearly large in magnitude as well.

Turning back to the effects of labor market conditions on firms' demand for discrimination, Table 6 explores the effects of alternative indicators of those conditions. Row 1 reproduces Table 5's results for the mean renewal rate (*MRR*) at the occupation*province*period level; Row 3 and 5 present coefficient estimates from regressions where *MRR* is replaced by the log of the number of competing ads in the cell, and the log of the competing vacancies (*LCA* and *LCV* respectively). Estimated effects of market conditions on firms' demands for education and experience are similar to row 1 (though note that these are in different units), and become statistically significant for both education and experience in the *LCV* specification. Rows 2, 4 and 6 of Table 6 change the level at which the above labor market variables are calculated to the occupation*industry*province*period. The advantage to doing this is a decline in measurement error, if workers in a given occupation have industry-specific skills, for example if computer and furniture salespeople are not perfect substitutes. Such a decline in measurement error should reduce the standard errors on *MRR*, *LCA* and *LCV*, and make these coefficients more negative. The disadvantage is that it may increase the expected negative correlation between the hiring firm's own demand for labor and *MRR*, *LCA* and *LCV* (at least if firms raise their relative demands for their preferred worker types when product demand is strong). This should make the estimated coefficients on *MRR*, *LCA* and *LCV* less negative.

If anything, defining labor markets at the occupation*province*industry level tends to reduce the standard errors of the coefficients on US-allowed ad characteristics, increasing their statistical significance. Now, high vacancy rates (intense competition for workers) lead to statistically significant declines in education requirements for two of the measures, and significant declines in experience requirements for all three measures. In all specifications, the data decisively reject the null hypothesis that both of these ad characteristics are unaffected by labor market conditions. Together, these results suggest that measurement error, rather than unobserved relative demand shocks, is the greater concern with our baseline estimates. For the US-prohibited ad characteristics, on the other hand, the estimated effects of labor market conditions are mostly insignificant, and inconsistent in sign when they are significant, for all the labor measures in Table 6. For the most part, tests of cross-equation restrictions do however reject the hypothesis that all these coefficients are zero, though somewhat less decisively than for the US-allowed coefficients.

Table 7 combines information from occupations and industries to compare the effects of two competing labor market indicators in the same regression. One refers to conditions in an ad's own occupation*industry*region*period cell; this is the indicator used in rows 2, 4 and 6 of Table 6. The other summarizes the conditions pertaining to that same occupation in other industries in the same province (and period). As noted, the idea is that, if other industries employing, say, software programmers, are expanding, this will reduce the availability of software programmers to one's own industry, without necessarily being correlated with the state of product demand at the advertising firm. In Table 7, the coefficients on conditions in the ad's own occupation and industry are similar to those in Table 6: US-allowed features of job ads respond to labor market conditions in the direction predicted by our simple search model, while US-prohibited characteristics do not. Effects of conditions in other industries for the same occupation on both types of ad characteristics are mostly insignificant, with the exception of experience requirements, which are statistically significant in two of three cases. Still, this effect on experience requirements provides additional support for the interpretation of our estimated coefficients on *MRR*, *LCA* and *LCV* as responses of ad-posting firms to supply shocks in the labor market where they are hiring.

Table 8 conducts three additional robustness checks, all of them now measuring labor market conditions at the occupation*province *industry*period level, or finer. The first asks whether our results change when we treat vacancies, rather than ads (which sometimes announce multiple vacancies) as the unit of analysis. To address this issue, Panel A of the Table restricts the sample to those ads which reported the number of vacancies available (including one), and replicates the regression results in rows 2, 4 and 6 of Table 6 weighting each observation by the number of vacancies it represents. Panel B of Table 8 addresses the concern that our results may be biased because, by construction, the estimation sample consists disproportionately of job ads that took a long time to fill. To address this concern, we constructed a subsample of ads that, almost surely, were first posted on Zhaopin after the start of the observation period in which it appears. The procedure for defining this sample is described in Appendix 2. We then re-estimated rows 2, 4 and 6 of Table 6 on this subsample. Finally, Panel C pushes our analysis of the effects of market conditions one step further by replacing the occupation province*industry fixed effects in rows 2, 4 and 6 of Table 6 by occupation*province* firm effects. These estimates remove any heterogeneity arising from a changing mix of firms that were hiring on Zhaopin in our two time periods: identification now comes from comparing the ads issued by the same firm, advertising for the same occupation in the same province, before and after China's recent labor market collapse.

Taken together, the results of these robustness checks are broadly similar, with mostly insignificant and inconsistently-signed effects of labor market conditions on the incidence of US-prohibited restrictions, and fairly consistent evidence that advertised education and experience requirements are less restrictive in tight labor markets. This is confirmed by tests of cross-equation restrictions: the null hypothesis that both of the US-allowed ad characteristics are unresponsive to labor market conditions is rejected in all specifications at the .005 level or better; the null hypothesis that all the US-prohibited characteristics are unresponsive is not always rejected at conventional significance levels. Finally, we note that the results in Table 8 (as well as previous tables) are most robust when experience requirements are the dependent variable: in every specification, firms reduce the amount of experience they require in their posted job advertisements in tight labor markets. This seems sensible, since experience might be the easiest of the hiring criteria to adjust in tight labor markets, compared for example to education requirements, gender mix, and standards for employees' physical appearance.

Our final set of robustness checks is summarized in Table 9. These address the concern that our finding of an insignificant effect of labor market conditions on advertised 'discriminatory' restrictions may be driven by the fact that some of those restrictions are quite rare (in the case of height, comprising only two percent of ads). Thus there may be little effect of market conditions on the demand for these characteristics, simply because the vast majority of firms do not care about them (or care in different directions, as in the case of gender). Columns 1-6 of Table 9 address this question by restricting the estimation sample for each of the six US-prohibited characteristics to occupations whose demands for the characteristic under study were in the top quartile in period 1 of our sample. Thus, for example, column 1 asks whether those occupations that were most inclined to explicitly seek male labor in the summer of 2008 were likely to intensify their requests for male labor when workers became much easier to find six months later. Again, while there is the odd significantly negative coefficient, there are some significant positive coefficients as well, and most coefficients are statistically insignificant. The test statistics in column 9 fail to reject the null hypothesis that all six coefficients are zero three of six times at the one-percent level.

Columns 7 and 8 address this question in a slightly different way for the one US-prohibited ad characteristic, age, that we can measure in a (partially) continuous way. Here, we restrict our sample to the approximately one quarter of ads that specified an age preference, and regress the level of the minimum and maximum age (in years) on labor market conditions. If firms that care about age become less 'choosy' regarding ages when workers are hard to find, we would expect high levels of *MRR*, *LCA* and *LCV* to reduce the minimum age, and to raise the maximum, thus widening the range of acceptable ages. This is not what we see: Looking at the occupation*province*industry based results (which arguably delineate labor markets more accurately), we find that the both the minimum and maximum ages tend fall by about the same amount when relatively few applications are available. While this might be interpreted as firms shifting their recruiting towards (less-valued) younger workers in tight labor markets (and is consistent with the trend towards lower experience requirements in those situations), this hypothesis is hard to reconcile with the leftward movement (and indeed existence) of the maximum age limit.

In sum, this section finds that firms reduce their advertised education and experience requirements in tight local labor markets. This effect is most robust for the case of experience, which is arguably the easiest of these job requirements to adjust in the short run. In contrast, we do not detect any robust or consistent effect of labor market conditions on firms' advertised requirements for 'US-prohibited' employee attributes, specifically gender, age, beauty and height. Indeed, compared to the effects of other factors, such as foreign ownership of the firm, or attributes of the occupation such as the amount of customer contact, effects of labor market conditions on the demand for US-prohibited characteristics are miniscule. While we can do little more than speculate about the possible reasons for this difference, one possibility might involve nonconvexities associated with 'appropriate' employee types for a job: If, for example, all incumbent employees in a job are male, the costs of hiring a single female for that job may greatly exceed the expected net productivity difference between an individual male and female applicant; among other possibilities, incumbents may harass (or customers may shun) workers of a different type (Lindbeck and Snower, 1988). This meshes with the suggestion of our search model, that productivity differentials would need to be quite large to justify excluding one group from the application process completely.[42]

[42] It also echoes China's own employment law (see footnote 2) which explicitly exempts jobs that are 'claimed to be unsuitable for females by the country' from antidiscrimination legislation, and with our own evidence that firms' demands for gender and other US-prohibited characteristics vary considerably across jobs within the firm.

6. Summary

We have studied patterns of firms' advertised preferences for an employee's gender, age, physical appearance and height in a sample of internet job ads, and have interpreted these patterns using a simple employer search model. We find that job advertisements specifying these requirements are commonplace. Clearly, a large fraction of Chinese employers find these characteristics to be useful screens, either because employers value the characteristics intrinsically or because the characteristics signal other desired attributes that are not easily measured at the time of hire. Further, our model suggests that firms will only advertise these preferences when they are intense.

Consistent with our model, firms become less likely to request 'US-prohibited' employee attributes in their job ads as job skill levels rise. We find that discriminatory restrictions are highly job-specific, with the same firm often imposing them in some jobs and not in others, or in different directions (e.g. against men or against women) in different jobs. Our model also predicts that firms should relax both their 'discriminatory' and nondiscriminatory advertised job requirements in tight labor markets. Using data from China's recent labor market collapse, we find that only the latter are sensitive to local, occupational labor market conditions. This lack of a response to labor market conditions is particularly striking when compared to other estimated effects, for example the large and robust differences between foreign- and domestic-owned firms, which may capture, for example, the effects of corporate culture.

Finally, in addition to offering a new approach to the empirical study of labor market discrimination, our paper contributes to the literature on employer search and recruitment. We develop, to our knowledge, the first theoretical model of employers' choice of advertised hiring criteria, and conduct the first empirical study of the determinants of these criteria, whether 'discriminatory' like gender, or not (like education). The model also provides a natural framework within which to study an old question in labor economics –the effects of labor market tightness on hiring discrimination.

We conclude with a brief discussion of the main 'fact' that makes this paper possible: why is it acceptable for employers to specify, say, a preferred employee gender or age in some countries but not in others? While we cannot provide a definitive answer, we offer two observations. One is our own result that employer preferences for age, gender, beauty and height all decline with a job's skill level. Thus, the benefits to employers of expressing such preferences may be considerably lower in a highly-skilled labor market like the U.S. than in a developing country like China. Second, we note that, in contrast to the 'US-allowed' characteristics (education and experience), the 'US-prohibited' characteristics are all *ascriptive*, in the sense there is little or nothing an applicant can do to alter his or her gender, age, beauty, or height.[43] The notion that job opportunities should not be allowed to depend on ascriptive (as opposed to earned) characteristics may be a stronger feature of American political culture than others.

[43] That said, see Lee's (2009) anecdotal account of efforts taken by Chinese job applicants to increase their beauty and their height.

References

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

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. 23 (3). p 363-82. July 1985.

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.

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, Lawrence Katz, and Claudia Goldin. "Dynamics of the Gender Gap for Young Professionals in the Corporate and Financial Sectors". NBER Working Paper No. 14681, January 2009.

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.

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

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

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

Burdett, Kenneth and Elizabeth J. Cunningham. "Toward a Theory of Vacancies" *Journal of Labor Economics*. July 1998; 16(3): 445-78

Case, A. and C Paxson, "Stature and Status: Height, Ability and Labor Market Outcomes" *Journal of Political Economy* 116(3) (June 2008): 499- 532.

Cerella, J. "Age effects may be global, not local: comment" *Journal of Experimental Psychology: General* (1991) (120): 215-23.

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

Charness, Gary, and Marie-Claire Villeval "Cooperation and Competition in Intergenerational Experiments in the Field and the Laboratory" *American Economic Review*. Vol. 99 (3). p 956-78. June 2009.

Couch, Kenneth A., and Robert Fairlie. "Last Hired, First Fired? Black-White Unemployment and the Business Cycle". unpublished paper, University of Connecticut, November 2005.

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.

Farber, H. "The Analysis of Inter-Firm Worker Mobility". *Journal of Labor Economics* 12 (1994): 554-593.

Freeman, Richard B. "Changes in the Labor Market for Black Americans, 1948-72". *Brookings Papers on Economic Activity* 1, (1973): 67-131.

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

Goldin, C. and C. Rouse. "Orchestrating Impartiality: The Impact of 'Blind' Auditions on Female Musicians" *American Economic Review* September 2000; 90(4): 715-41

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

Hamermesh, Daniel S. *Beauty Pays.* Princeton: Princeton University Press, forthcoming 2010.

Hamermesh, Daniel S. and Jeff E. Biddle. "Beauty and the Labor Market" *The American Economic Review*, Vol. 84, No. 5 (Dec., 1994), pp. 1174-1194

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. " Overt Employment Discrimination by Multinational Firms: Cultural and Economic Influences in a Developing Country *Industrial Relations*, 37, 2, April 1998 (p 126-152)

Lazear, Edward P. "Why Is There Mandatory Retirement?" *Journal of Political Economy*. Vol. 87 (6). p 1261-84. December 1979.

Lee, Don. "Plastic Surgery Just to Get a Job". *Los Angeles Times*, March 30, 2009. page B1.

Lindbeck, Assar, and Dennis J. Snower. "Cooperation, Harassment, and Involuntary Unemployment: An Insider-Outsider Approach" *American Economic Review*, Vol. 78, No. 1 (Mar., 1988), pp. 167-188

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

Meng, Xin and Junsen Zhang. "The Two-Tier Labor Market in Urban China Occupational Segregation and Wage Differentials between Urban Residents and Rural Migrants in Shanghai" *Journal of Comparative Economics* 29, 485–504 (2001)

Meng, Xin "Male–female wage determination and gender wage discrimination in China's rural industrial sector" *Labour Economics* 5 (1998): 67–89.

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

Mobius, Markus M., and Tanya S. Rosenblat. "Why Beauty Matters" *American Economic Review*. Vol. 96 (1). p 222-35. March 2006.

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.

Niederle, Muriel, and Lise Vesterlund. "Do Women Shy Away from Competition? Do Men Compete Too Much?" *Quarterly Journal of Economics*, Aug2007, Vol. 122 Issue 3, p1067-1101

Oaxaca, R. "Male-Female Wage Differentials in Urban Labour Markets." *International Economic Review* 14 (1973): 693-709

Oaxaca, Ronald, and Michael Ransom. " New Market Power Models and Sex Differences in Pay", Princeton Industrial Relations Section working paper no. 540, December 2008.

Papalia, Diane et al., *Adult Development and Aging* (2nd edition) Boston: McGraw Hill, 2002.

Rosen, Asa. "An equilibrium search-matching model of discrimination." *European Economic Review* 41 (1997) 1589-1613

Stigler, George J. "The Economics of Information". *Journal of Political Economy* 69(3) (June 1961): 213-225.

Teulings, Coen N. and Pieter A. Gautier. "The Right Man for the Job" *The Review of Economic Studies*, Vol. 71, No. 2 (Apr., 2004), pp. 553-580

Van Ommeren, Jos and Giovanni Russo. "Firm Recruitment Behaviour: Sequential or Non-Sequential Search?" IZA Discussion Paper No. 4008, February 2009

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

Van Ours, Jan, and Geert Ridder. "Vacancies and the Recruitment of New Employees". *Journal of Labor Economics* 10(2) (April 1992): 138-155.

Van Ours, Jan, and Geert Ridder. "Vacancy Durations: Search or Selection?" *Oxford Bulletin of Economics and Statistics* 55(2) (1993).

Yinger, John. "Measuring Racial Discrimination with Fair Housing Audits: Caught in the Act." *American Economic Review*. Vol. 76 (5). p 881-93. December 1986.

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).

Table 1: Simulation Results: Critical Values of Group B's mean productivity disadvantage (%) for alternative parameter values:

	(1)	(2)	(3)	(4)
	Application processing costs (cM) as a share of Total Employee Net Value:			
Standard Deviation of Idio-syncratic Employee Value (σ):	**.0072 (base)**	**.020**	**.083**	**.167**
.750 (base)	297%	237%	153%	84%
.550	205%	161%	98%	66%
.375	128%	98%	55%	32%
.200	59%	42%	19%	5%

Notes: Numbers in the table represent the percentage gap between the mean productivities of the two groups ($100*(\mu^A - \mu^B)/\mu^A$) at which a firm is indifferent between excluding group B and not. σ=.750 implies that an applicant of a given type at the 95th percentile of the productivity distribution is twice as productive as the mean for that type. The remaining values of σ imply ratios of the 95th percentile to the mean of 1.74, 1.5, and 1.27 respectively.

Table 2: Descriptive Statistics

	Period 1	Period 2	Combined
A. ADVERTISED JOB REQUIREMENTS			
Gender requirement			
No gender preference	0.898	0.896	0.897
Prefer male?	0.053	0.052†	0.053
Prefer female?	0.049	0.052	0.050
Age requirement			
No age restrictions	0.762	0.755	0.759
Minimum age requirement?	0.158	0.174	0.165
Maximum age requirement?	0.193	0.200	0.196
Both minimum and maximum age requirement	0.113	0.129	0.120
Minimum age, when required	24.59	24.59	24.59
Maximum age, when required	35.57	35.71	35.63
Age range, when maximum and minimum listed	11.51	11.81	11.65
Beauty requirement			
Job requires beauty (“xingxiang”)	0.076	0.085†	0.080
Requires beauty given preference for male	0.052	0.063	0.056
Requires beauty given preference for female	0.329	0.338	0.333
Requires beauty given no gender preference	0.064	0.071	0.067
Height requirement			
Any height requirement?	0.022	0.027	0.024
Two gender-specific height requirements	0.013	0.016	0.014
A single height requirement and preference for men	0.003	0.004	0.003
A single height requirement and preference for women	0.006	0.008	0.007
One gender-independent height requirement	0.000	0.000	0.000
Minimum male height, when required (centimeters)	171.1	170.7	170.9
Minimum female height, when required (centimeters)	163.1	162.7	162.9
Advertised Education requirement			
Unspecified	0.208	0.221	0.214
Grade 9	0.006	0.011	0.008
High school	0.090	0.103	0.096
Post secondary	0.370	0.372	0.371
University	0.316	0.281	0.301
Master's degree	0.009	0.011	0.010
PhD	0.001	0.001	0.001
Advertised Experience requirement			
Unspecified	0.302	0.305	0.304
No experience	0.003	0.004	0.003
1 Year or less	0.012	0.011	0.011
1 to 3 years	0.358	0.367	0.362
3 to 5 years	0.201	0.191	0.197
5 to 10 years	0.109	0.106	0.108
10 Years or above	0.014	0.016	0.014

….Table 2, Continued	Period 1	Period 2	Combined
B. OTHER AD CHARACTERISTICS			
Job is Part Time	0.011	0.020	0.015
Length of Job description:			
Number of lines	7.08	6.90	7.01
Number of characters	451.8	453.3†	452.4
Number of positions advertised:			
Unspecified	0.502	0.515	0.508
1	0.259	0.244	0.253
2	0.100	0.094	0.098
3-5	0.085	0.085†	0.085
6-15	0.039	0.044	0.041
16-50	0.012	0.015	0.013
51+	0.002	0.003	0.002
C. FIRM CHARACTERISTICS			
Firm size (number of workers):			
1-19	0.091	0.102	0.096
20-99	0.359	0.374	0.365
100-499	0.307	0.311	0.309
500-999	0.095	0.081	0.089
1,000-9,999	0.118	0.107	0.113
10,000 +	0.030	0.025	0.028
Firm ownership type: **			
Private, Domestic	0.549	0.592	0.567
Foreign	0.370	0.311	0.345
NonProfit	0.003	0.004	0.003
State-Owned Enterprise	0.078	0.092	0.084
Government	0.001	0.001	0.001
D. LABOR MARKET MEASURES:			
Mean refreshments per ad in occupation*province cell (MRR)	14.87	10.90	13.19
Number of ads in occupation*province cell (NCA)	4499	2867	3809
Number of vacancies in occupation*province cell (NCV)	19185	12275	16262
N	365615	268049	633664

Notes: Sample sizes vary slightly across rows; reported numbers are for gender restrictions. Period 1 is May 16/08 - July 29/08; Period 2 is Dec 17/08 - Feb 28/09. † indicates the difference between periods is *not* statistically significant at the 5% level.
* In inches, these are 5 feet 7 inches for men, and 5 feet 4 inches for women, respectively.
** "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 3: Prevalence of "Discriminatory" Job Ads, by Firm.

	Share of Firms Specifying a Preference For:											
	Gender				Age							
	Any Gender	Men	Women	Both Genders	Any Age Preference	Minimum Age	Maximum Age	Both Min and Max	Beauty	Height	Any US-Prohibited Trait	**N**
A. Total Number of Ads Placed by the Firm:												
1	.1643	.0673	.0970	.0000	.2698	.1698	.2345	.1345	.1046	.0267	.3887	8650
2-10	.3037	.1496	.2070	.0528	.4217	.3119	.3727	.2630	.2492	.0734	.5876	40873
11-50	.5376	.3509	.3895	.2029	.7056	.5946	.6549	.5439	.5108	.1992	.8257	12407
51 and over	.6607	.5149	.5079	.3621	.8326	.7749	.7996	.7419	.6741	.2930	.8998	1577
All Firms	.3393	.1868	.2352	.0826	.4667	.3593	.4196	.3122	.2912	.0971	.6148	63507
B. Total Number of Occupations Sought by the Firm:												
1	.1680	.0716	.1048	.0085	.2937	.1994	.2558	.1616	.1324	.0424	.4109	18086
2-10	.3905	.2126	.2735	.0956	.5184	.4033	.4669	.3519	.3383	.1082	.6845	43217
11-20	.7365	.6167	.5456	.4258	.8678	.8000	.8320	.7642	.6606	.3217	.9189	2095
21 and over	.8349	.7890	.7064	.6606	.9633	.9541	.9174	.9083	.8532	.4587	.9817	109
All Firms:	.3393	.1868	.2352	.0826	.4667	.3593	.4196	.3122	.2912	.0971	.6148	63507

Table 4: Pairwise Correlations Among Advertised Job Requirements

	Male	**Female**	**Minimum Age?**	**Maximum Age?**	**Beauty**	**Height**	**Education (years)**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**
1. Prefer male?							
2. Prefer female?	-.0452						
3. Any minimum age?	.1512	.1201					
4. Any maximum age?	.2065	.1640	.5950				
5. Requires beauty?	-.0191	.2147	.1133	.1352			
6. Any height requirement?	.0467	.1570	.1102	.1409	.3282		
7. Years of education required	-.0958	-.1248	-.1335	-.1437	-.1215	-.1689	
8. Years of experience required	.0185	-.0906	-.0206	-.0267	-.1087	-.0731	.2528

Note: All *p*-values are less than .001.

Table 5: Effects of Mean Ad Renewal Rate and Firm Characteristics on Job Ad Content

	DEPENDENT VARIABLE:							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Prefer male?	Prefer female?	Minimum Age?	Maximum Age?	Request Beauty?	Height Requirement?	Education (years)	Experience (years)
Mean Renewal Rate/10	-.00337	.00339	-.00349	-.00394	.00339	-.00256	.11269	-.22332**
	(.00327)	(.00319)	(.00974)	(.01000)	(.00502)	(.00327)	(.09239)	(.06253)
Number of Vacancies	-.00065**	-.00007	.00236*	.00136	.00056	.00070*	-.04133*	-.24196**
Advertised/10	(.00022)	(.00011)	(.00111)	(.00071)	(.00034)	(.00034)	(.01922)	(.03783)
Log Firm Size	.00178**	-.00375**	.01646**	.01236**	-.00224**	.00144**	.19397**	-.01586
	(.00037)	(.00053)	(.00152)	(.00145)	(.00060)	(.00040)	(.00941)	(.00928)
Firm Ownership Type:								
Foreign Ownership	-.01790**	-.01602**	-.03854**	-.05961**	-.01911**	-.00871**	.84959**	.55510**
	(.00141)	(.00185)	(.00372)	(.00387)	(.00257)	(.00143)	(.04210)	(.02782)
Non-profit	.01161	-.01963**	-.01275	.06241**	.00218	.00573	.67614**	.31217**
organization	(.00744)	(.00679)	(.01505)	(.01749)	(.01024)	(.00504)	(.17786)	(.10580)
State-owned	.00820**	-.00977**	-.01700**	.01234*	-.01698**	-.00469**	.22354**	.17736**
enterprise	(.00243)	(.00251)	(.00482)	(.00488)	(.00306)	(.00132)	(.04970)	(.02778)
Government	.06937**	-.02959	-.05110	.14807**	.07850	.06160	1.22066	.28165
	(.02377)	(.02069)	(.03803)	(.05058)	(.06795)	(.04745)	(.65169)	(.25142)
Period 2	-.00149	.00267	.01355**	.00216	.00404	.00254	.03932	.01135
	(.00171)	(.00183)	(.00381)	(.00413)	(.00216)	(.00174)	(.04380)	(.02326)
Sample Size	633,617	633,617	633,617	633,617	633,617	633,617	633,617	441,239
R^2	.058	.063	.050	.066	.097	.103	.190	.129

** p<0.01, * p<0.05. OLS estimates. Regressions also control for a full set of occupation*province fixed effects, plus a dummy for part-time jobs and for number of vacancies not specified. Standard errors are clustered at the occupation*province level. The omitted firm type is for-profit firms with no government or foreign connection.

Table 6: Effects of Alternative Measures of Labor Market Conditions (*M*) on Job Ad Content

	(1) Male	(2) Female	(3) Minimum Age?	(4) Maximum Age?	(5) Beauty	(6) Height	(7) Education	(8) Experience	(9) Pr (1)-(6)=0	(10) Pr (7)-(8)=0
MEAN RENEWAL RATE/10:										
In occupation/province cell	-.00336	.00340	-.00350	-.00394	.00339	-.00255	.11256	-.22011**	.6515	0.0008
	(.00327)	(.00318)	(.00973)	(.00999)	(.00501)	(.00327)	(.09230)	(.06216)		
In occupation/province /industry cell	-.00095	.00494*	.00786*	.00135	.00386	-.00050	-.01165	-.10688**	.0427	0.0000
	(.00153)	(.00225)	(.00396)	(.00426)	(.00257)	(.00115)	(.03933)	(.02387)		
LOG (COMPETING ADS) :										
In occupation/province Cell	.00094	.01391**	-.01554*	-.00254	.00597	.01016**	-.10504	-.30228**	.0000	.0000
	(.00348)	(.00316)	(.00669)	(.00704)	(.00522)	(.00353)	(.09261)	(.05111)		
In occupation/province /industry cell	-.00369*	.00104	-.00310	-.01008**	-.00228	.00272	-.13903**	-.08192**	.0004	.0000
	(.00159)	(.00168)	(.00331)	(.00340)	(.00239)	(.00145)	(.03263)	(.02245)		
LOG (COMPETING VACANCIES) :										
In occupation/province cell	-.00038	.01043**	-.00611	.00752*	.00320	.00505**	-.11309**	-.09302**	.0000	.0000
	(.00164)	(.00208)	(.00320)	(.00371)	(.00231)	(.00187)	(.03949)	(.02248)		
In occupation/province /industry cell	-.00080	.00059	.00139	-.00101	-.00035	.00401**	-.13113**	-.06002**	.0006	.0006
	(.00114)	(.00134)	(.00228)	(.00238)	(.00170)	(.00108)	(.02221)	(.01451)		

** p<0.01, * p<0.05. Regressions also control for a full set of occupation*province or occupation*province*industry fixed effects, firm type, firm size, plus dummies for part-time jobs, period 2, and for number of vacancies not specified. Standard errors and *p*-values are clustered at the occupation*province or occupation*province*industry level and adjusted for cross-equation correlations.

Table 7: Distinguishing the effects of ads in Own Industry versus other Industries

	(1) Male	(2) Female	(3) Minimum Age?	(4) Maximum Age?	(5) Beauty	(6) Height	(7) Education	(8) Experience	(9) Pr (1)-(6)=0	(10) Pr (7)-(8)=0
MEAN RENEWAL RATE:										
In occupation/province /industry	-.00118	.00480*	.00805*	.00106	.00398	-.00076	-.01017	-.10579**	.0289	.0001
	(.00154)	(.00227)	(.00400)	(.00430)	(.00260)	(.00115)	(.03975)	(.02409)		
In occupation/province, *other* industries	-.00008	-.00009	-.00035	.00010	-.00005	-.00003	.00285	-.00202*	.2000	.0157
	(.00011)	(.00008)	(.00019)	(.00018)	(.00009)	(.00006)	(.00156)	(.00097)		
LOG (COMPETING ADS) :										
In occupation/province /industry	-.00374*	.00104	-.00371	-.01003**	-.00239	.00278	-.14054**	-.08603**	.0005	.0000
	(.00160)	(.00168)	(.00334)	(.00344)	(.00241)	(.00146)	(.03298)	(.02266)		
In occupation/province, *other* industries	-.00002	-.00001	-.00045*	.00005	-.00008	.00006	-.00100	-.00279**	.0923	.0241
	(.00011)	(.00008)	(.00021)	(.00018)	(.00009)	(.00006)	(.00153)	(.00103)		
LOG (COMPETING VACANCIES) :										
In occupation/province /industry	-.00094	.00050	.00078	-.00170	-.00053	.00372**	-.13011**	-.05890**	.0011	.0000
	(.00112)	(.00137)	(.00234)	(.00233)	(.00174)	(.00103)	(.02307)	(.01482)		
In occupation/province, *other* industries	.00000	.00000	-.00009*	.00000	.00002	.00000	-.00027	-.00021	.5412	.6805
	(.00003)	(.00003)	(.00005)	(.00005)	(.00003)	(.00002)	(.00040)	(.00032)		

** p<0.01, * p<0.05. Regressions also control for a full set of occupation*province*industry fixed effects, plus a dummy for part-time jobs and for number of vacancies not specified. Standard errors are clustered at the occupation*province*industry level. "Own" MRR, and "Other industry" LCA and LCV are divided by 10; "Other industry" MRR is divided by 100.

Table 8: Additional Robustness Checks

	(1) Male	(2) Female	(3) Minimum Age?	(4) Maximum Age?	(5) Beauty	(6) Height	(7) Education	(8) Experience	(9) Pr (1)-(6)=0	(10) Pr (7)-(8)=0
A. Weight by Number Of Vacancies:										
Mean Renewal Rate/10	-.00847	.01302	.03158	.04004*	.03824*	.02053	.25537	-.19572**	.0176	.0000
	(.00469)	(.00864)	(.02022)	(.01887)	(.01646)	(.01364)	(.18998)	(.04481)		
Log number of Competing ads	.00566	.01029	-.01886	-.00769	.00025	.01292	.07271	-.15610**	.2337	.0024
	(.00453)	(.00907)	(.01617)	(.01610)	(.01278)	(.01127)	(.12606)	(.04734)		
Log number of Competing Vacancies	-.01310**	-.00004	.02662	.01048	.01295	.02573*	-.76032**	-.17683**	.0003	.0003
	(.00433)	(.00617)	(.01911)	(.01625)	(.01419)	(.01125)	(.10092)	(.02853)		
B. Inflow Sample only:										
Mean Renewal Rate/10	-.00163	.00474	.00697	.00239	-.00004	-.00250	.09129	-.09186*	.1297	.0029
	(.00250)	(.00276)	(.00562)	(.00611)	(.00395)	(.00174)	(.05217)	(.03612)		
Log number of Competing ads	-.00385	-.00276	-.00148	-.00979*	-.00818**	.00093	-.14707**	-.04395	.0061	.0028
	(.00238)	(.00223)	(.00440)	(.00464)	(.00312)	(.00164)	(.04542)	(.02960)		
Log number of Competing Vacancies	-.00041	-.00145	-.00064	-.00241	-.00427	.00321**	-.11528**	-.05866**	.0042	.0042
	(.00174)	(.00162)	(.00319)	(.00310)	(.00223)	(.00118)	(.03008)	(.01973)		
C. Firm Fixed Effects:										
Mean Renewal Rate/10	-.00031	.00258	.00273	.00330	.00737	-.00458*	-.02040	-.22739**	.0697	.0001
	(.00316)	(.00246)	(.00555)	(.00590)	(.00431)	(.00203)	(.06011)	(.05285)		
Log number of Competing ads	-.00235	.00654**	-.02570**	-.01711*	-.01255*	-.00170	.02268	-.32471**	.0018	.0000
	(.00339)	(.00250)	(.00726)	(.00715)	(.00556)	(.00232)	(.04675)	(.05484)		
Log number of Competing Vacancies	.00070	.00433*	-.00985*	-.00283	-.00936**	-.00007	.00047	-.16533**	.0025	.0025
	(.00237)	(.00171)	(.00492)	(.00502)	(.00344)	(.00132)	(.03223)	(.03205)		

** p<0.01, * p<0.05. All regressions also control for firm type, firm size, plus dummies for part-time jobs, period 2, and for number of vacancies not specified. Panels A and B also control for a full set of occupation*province*industry fixed effects; Panel C for a full set of occupation*province*firm effects. Standard errors are clustered at the occupation*province*industry or occupation*province*firm level.

Table 9: Sample Restricted to Occupations with High Period -1 Preferences for Each 'Prohibited' Attribute

DEPENDENT VARIABLE:

	Sample: Occupations in Top Quartile of First-Period Preferences for the Characteristic						Sample: Ads with an Age Preference			
	(1) Male	(2) Female	(3) Minimum Age?	(4) Maximum Age?	(5) Beauty	(6) Height	(7) Min Age (years)	(8) Max Age (years)	(9) Pr (1)-(6)=0	(10) Pr (7)-(8)=0
MEAN RENEWAL RATE/10:										
In occupation/province cell	-.00118	.00779	-.00781	-.02088	.01410	-.00449	-.03943	-.24112	.1145	.5657
	(.00883)	(.00639)	(.01815)	(.01849)	(.01240)	(.00752)	(.11734)	(.22737)		
In occupation/province /industry cell	-.00174	.01054*	.01608	-.00144	.00792	.00002	-.01945	-.09450	.0185	.5744
	(.00408)	(.00516)	(.00887)	(.00965)	(.00611)	(.00309)	(.06882)	(.09029)		
LOG (COMPETING ADS) :										
In occupation/province Cell	-.00175	.01962*	-.02546	.01553	.00057	.02426**	-.22364	-.98697**	.0007	.0001
	(.00840)	(.00919)	(.01543)	(.01563)	(.01396)	(.00897)	(.18257)	(.23504)		
In occupation/province /industry cell	-.00789*	.00059	.00275	.00311	-.01124	.00376	-.14863*	-.17312	.0347	.0443
	(.00402)	(.00428)	(.00731)	(.00713)	(.00645)	(.00408)	(.06297)	(.09672)		
LOG (COMPETING VACANCIES) :										
In occupation/province cell	-.00154	.01820**	-.00574	.01768*	.00544	.01159**	-.21167**	-.58775**	.0000	.0000
	(.00330)	(.00528)	(.00501)	(.00689)	(.00536)	(.00333)	(.06897)	(.09600)		
In occupation/province /industry cell	-.00254	-.00079	.00396	.01066*	-.00395	.00923**	-.24349**	-.26864**	.0000	.0000
	(.00288)	(.00242)	(.00447)	(.00459)	(.00365)	(.00254)	(.04532)	(.06007)		

** p<0.01, * p<0.05. Regressions also control for a full set of occupation*province or occupation*province*industry fixed effects, firm type, firm size, plus dummies for part-time jobs, period 2, and for number of vacancies not specified. Standard errors and p-values are clustered at the occupation*province or occupation*province*industry level. Occupations are weighted by total number of ads in computing quartiles.

Figure 1: Occupation Fixed Effects for US-Prohibited Characteristics, By Education Requirement and Customer Contact

a)

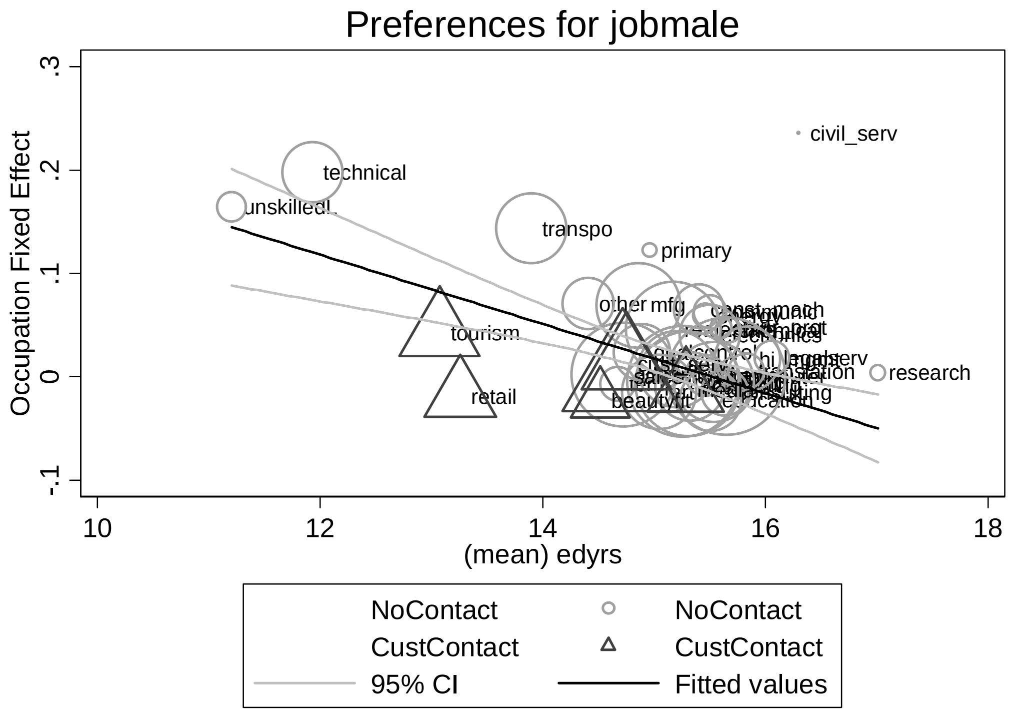

b)

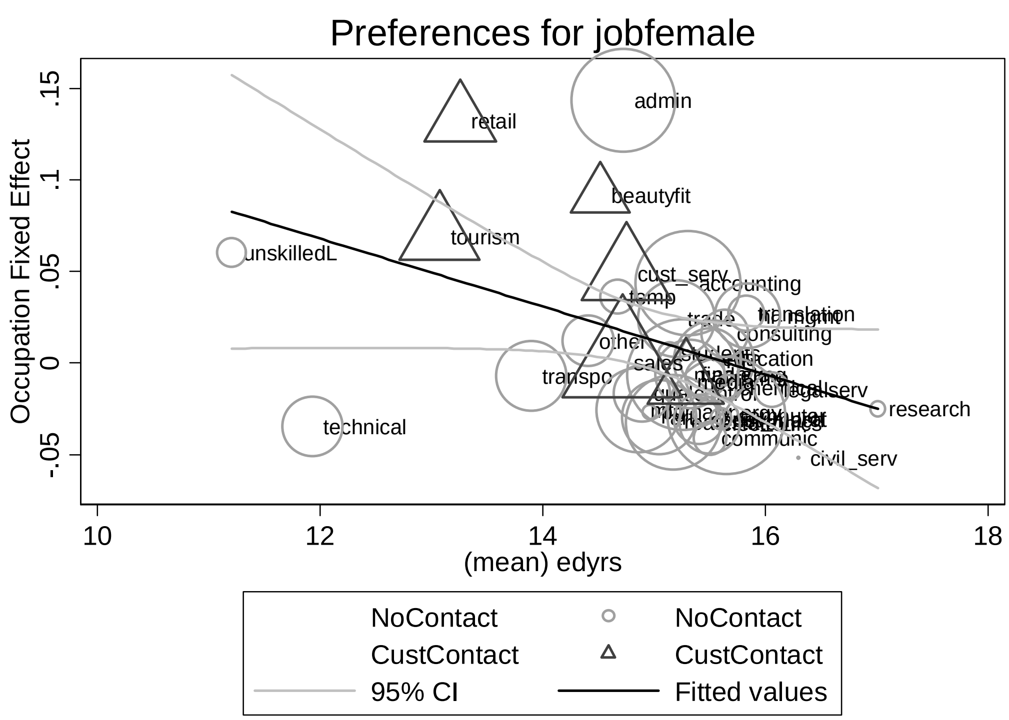

Note: Symbol size is proportional to the inverse of the variance of the estimated fixed effect.

c)

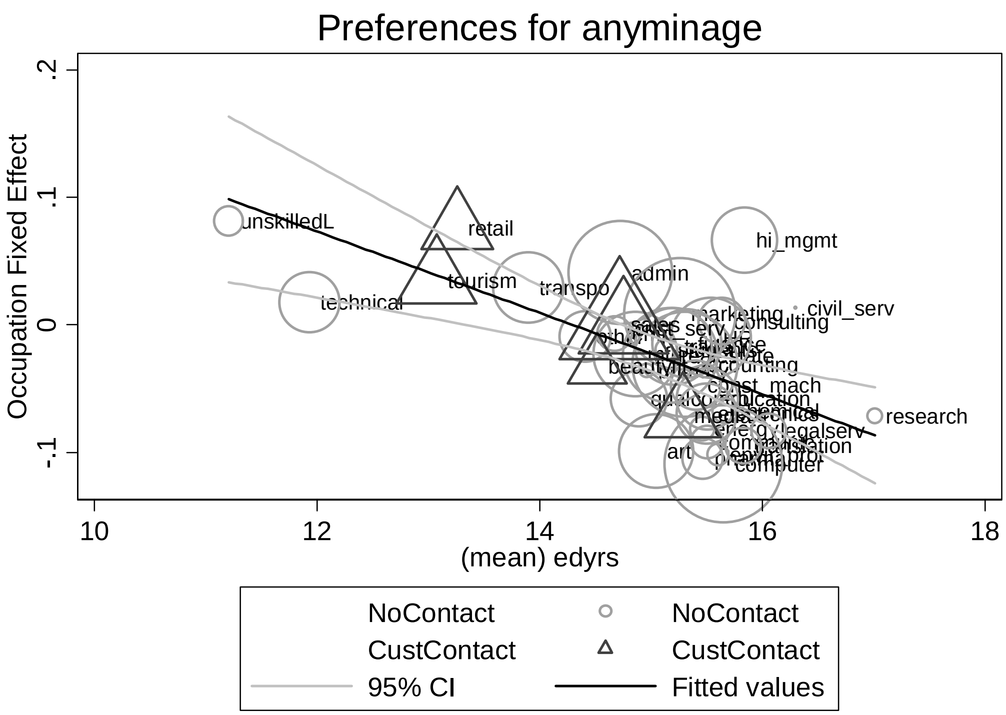
Preferences for anyminage
Occupation Fixed Effect
.2
.1
0
-.1
10
12
14
16
18
(mean) edyrs
unskilledL
retail
hi_mgmt
tourism
admin
technical
transpo
civil_serv
research
art
computer
NoContact
CustContact
95% CI
Fitted values

d)

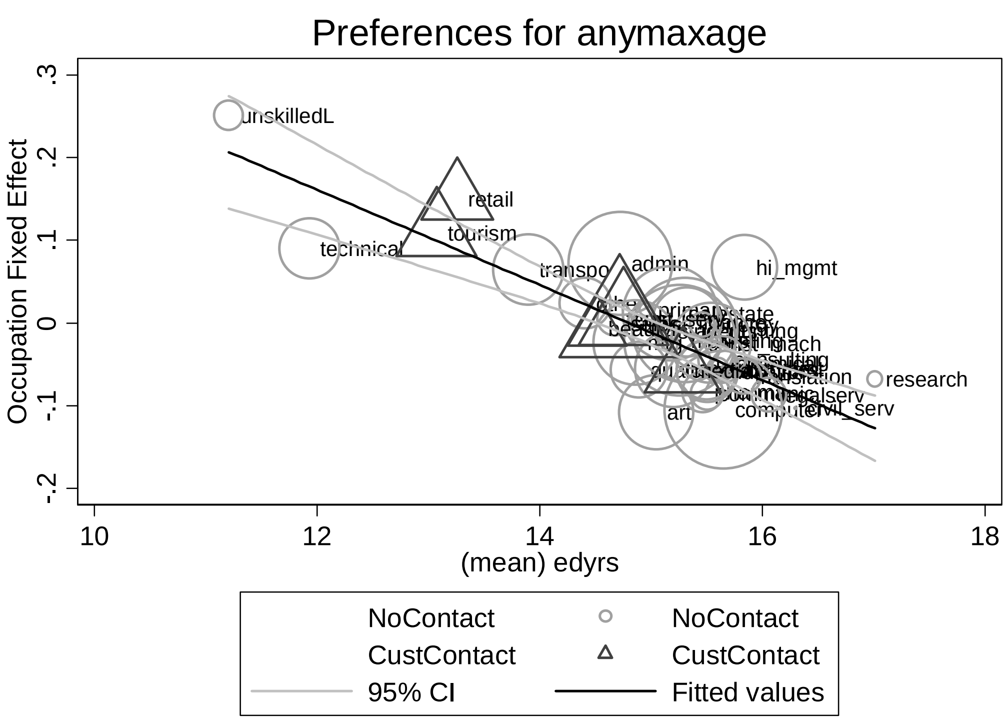
Preferences for anymaxage
Occupation Fixed Effect
.3
.2
.1
0
-.1
-.2
10
12
14
16
18
(mean) edyrs
unskilledL
retail
tourism
technical
transpo
admin
hi_mgmt
research
art
NoContact
CustContact
95% CI
Fitted values

e)

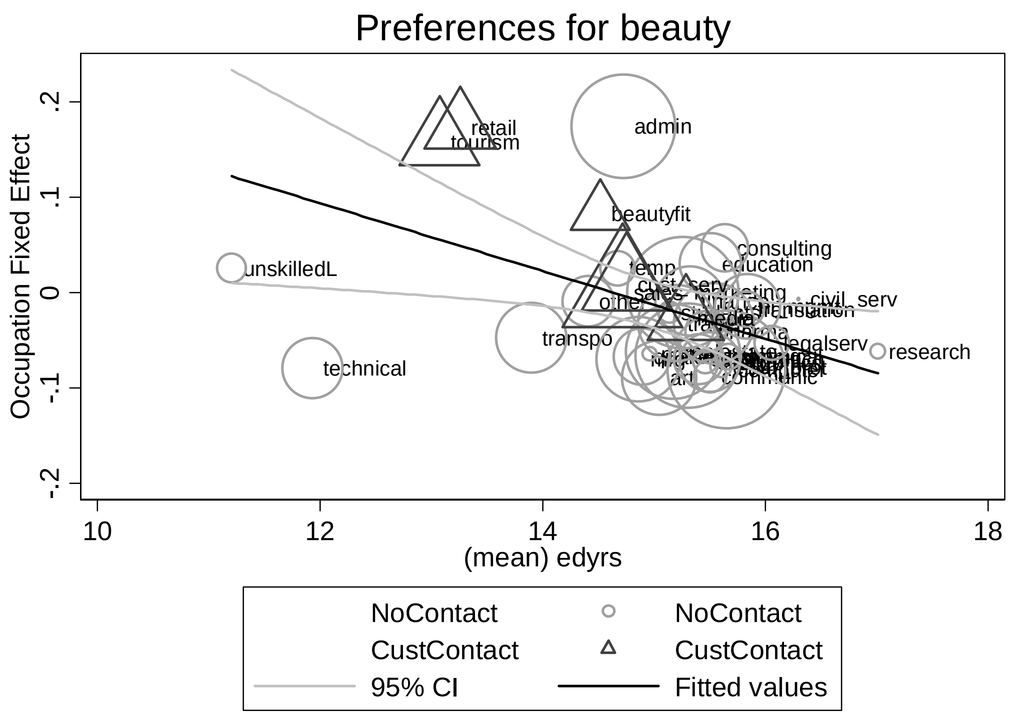
Preferences for beauty
Occupation Fixed Effect
(mean) edyrs
NoContact
CustContact
95% CI
Fitted values
retail
tourism
admin
beautyfit
consulting
education
unskilledL
civil_serv
transpo
legalserv
research
technical

f)

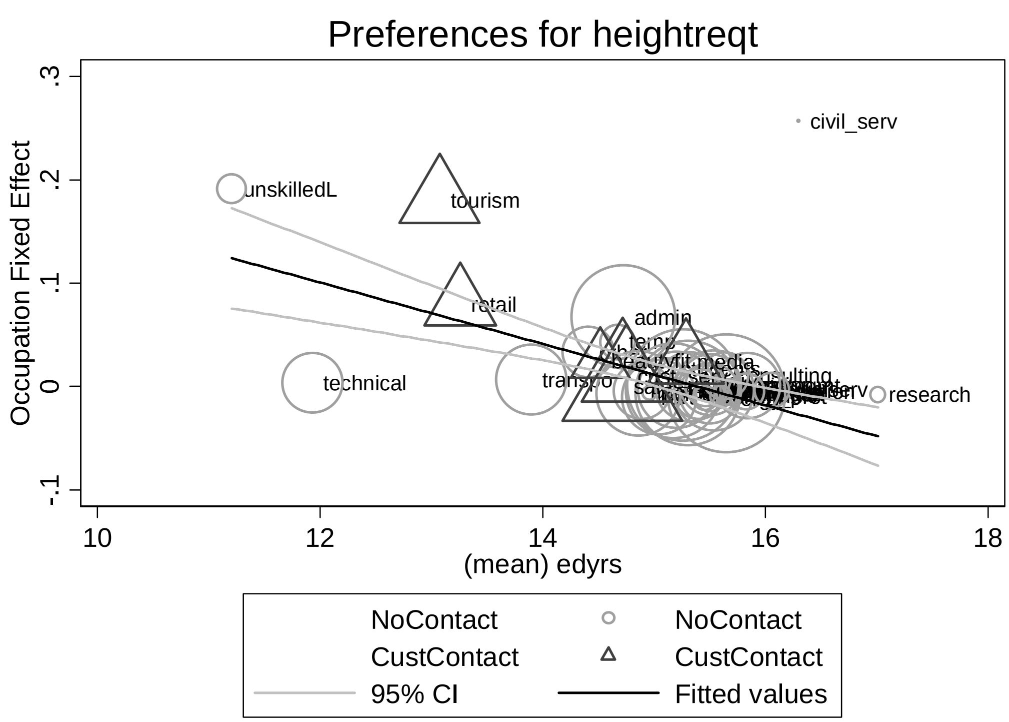
Preferences for heightreqt
Occupation Fixed Effect
(mean) edyrs
NoContact
CustContact
95% CI
Fitted values
civil_serv
unskilledL
tourism
retail
admin
technical
transpo
research

Appendix 1: Proofs

Expected Value of the Maximum:

We begin by normalizing the value of the worker, defining $u_j \equiv U_j/\beta = v^J/\beta + e_j$, where $e_j = \varepsilon_j/\beta$ follows a "standard" extreme value distribution with $\mathrm{Var}(e_j) = \pi^2/6$ and $E(e_j) = \gamma$. This 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 casts expresses the problem in a standard multinomial logit format and allows us to draw on some results from that literature.

Among these, it is well known that the expected value of the maximum of $v^J/\beta + e_j$ when e_j is independently drawn J times from a "standard" extreme value distribution is $v^J/\beta + \gamma + \log(n)$.[44] Multiplying through by β the expected maximum of U when the firm samples from either the A or B pool separately thus equals $v^J + \beta\gamma + \beta\log(J)$, as claimed for these two strategies.

In general, the expected value of the highest u_j in the "combined", C sample, $u^{C*} = u^{A*} q^A + u^{B*}(1- q^A)$, where u^{J*} is the expected productivity of the best overall worker given the best overall worker is of type J, and q^A is the probability that the best overall worker turns out to be drawn from pool A. Using results from the MNL literature, we know that $u^{A*} = v^A/\beta + \gamma - log(p^A)$, where:

$$p^A = \frac{\exp(v^A/\beta)}{A\exp(v^A/\beta) + B\exp(v^B/\beta)}$$

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

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

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

$$q^A = \frac{A\exp(v^A/\beta)}{A\exp(v^A/\beta) + B\exp(v^A/\beta)} = Aq^A.$$

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

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

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

$$u^{C}* \;=\; \gamma + \;\; log\;[\; Aexp(v^{A}/\,\beta) + Bexp(v^{B}/\,\beta)\;]\;.$$

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

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

The corresponding maximized unstandardized value is therefore:

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

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

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

as claimed. ■

Proof of Proposition 1

We recall first that the firm's choices between the three recruitment strategies yield the following levels of expected profits:

Strategy A: Invite A's only:	$E(\pi) = \mu^{A} + \beta log(A\,) - cA$
Strategy B: Invite B's only:	$E(\pi) = \mu^{B} + \beta log(B\,) - cB$
Strategy C: Invite all comers:	$E(\pi) = \mu^{C} + \beta log(M) - cM$

where none of the μ's depend on the number of applicants in either pool. Thus, the profit-maximizing number of applicants if the firm could choose the number of applications it receives --i.e. the n that maximizes $\beta log(n\,) - cn$-- therefore equals $\beta/c \equiv N^*$, irrespective of whether the firm chooses strategy A, B or C. Thus the optimal sample size rises with the variance of unobserved worker qualifications, β, and falls with application processing costs.

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

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

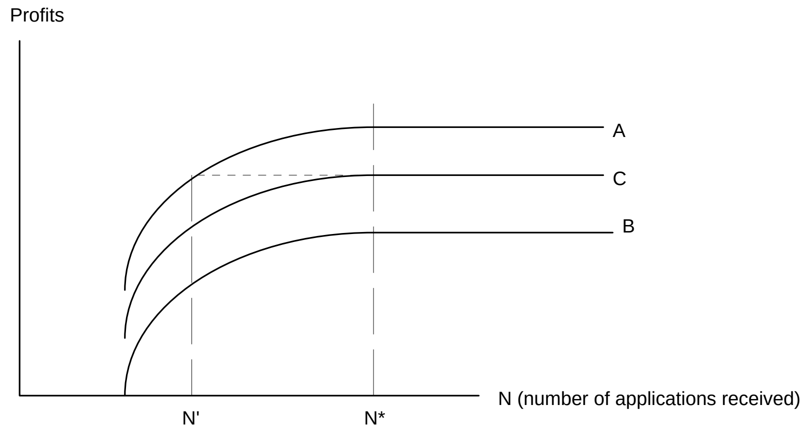

Now imagine that the firm is faced with fixed numbers of applicants under each of its possible recruiting strategies: A, B, and $M=A+B$. The fact that the three above functions are nondecreasing immediately implies that the firm will never choose recruiting strategy B for any level of labor supply M: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

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

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

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

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

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

which is increasing in M.

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

$$\Pi^A(\delta M) - \Pi^C(N^*) = \mu^A + \beta log(\delta M) - c\delta M - \Pi^C(N^*) \qquad (A2)$$

Because $\Pi^C(N^*)$ is independent of M, and because $\Pi^A(\delta M) = \mu^A + \beta log(\delta M) - c\delta M$ is increasing in M when $\delta M < N^*$, ($A2$) is also increasing in M.

Region 3 ($N^* < \delta M < M$). The difference in profits between strategies A and C is given by $\Pi^A(N^*) - \Pi^C(N^*) = \mu^A - \mu^C$, which is independent of M.

Combining our results for Regions 1-3, and verifying that profits do not fall over the boundaries between these regions (this is straightforward), completes the proof that the profit differential is nondecreasing in *M*. We now make use of this result.

Recall that the smallest possible value of *M*, M_{min}, equals $1/\delta$ (corresponding to $A=\delta M = 1$), and that Assumption 1 guarantees that M_{min} occupies Region 1. Substituting these values into (1), the profit differential between strategies A and C is $\mu^A - \mu^C + \beta \log(\delta) + c(1-\delta)/\delta$; firms will prefer to hire only *A*s when the supply of applications is at a minimum when:

$$\mu^A - \mu^C + \beta \log(\delta) + c(1-\delta)/\delta > 0 \qquad (A3)$$

Further, since the profit advantage of strategy A is nondecreasing in *M*, (A3) implies that firms will choose strategy A at all other levels of *M* as well. This proves the first part of the Proposition.

When *(A3)* is violated, $\Pi^A - \Pi^C < 0$ and firms prefer strategy C when $M=M_{min}$.
But we know that $\Pi^A - \Pi^C > 0$ in Region 3. Thus, because $\Pi^A - \Pi^C$ is monotonically increasing throughout Regions 1 and 2, there must exist a critical value of M, $\tilde{M}$, below which firms prefer strategy A and above which they prefer strategy C. This proves the remainder of the Proposition. ■

Proof of Proposition 2

Depending on parameter values, $\tilde{M}$ can fall either into Region 1 or 2. If it falls into Region 1, then it is implicitly defined by the expression:

$\Pi^A(\delta M) - \Pi^C(M) = \mu^A - \mu^C + \beta\log(\delta) + (1-\delta)cM = 0,$ i.e.

$$\mu^A - \beta \log\left[\delta \exp\left(\frac{\mu^A}{\beta}\right) + (1-\delta)\exp\left(\frac{\mu^B}{\beta}\right)\right] + \beta \log \delta + (1-\delta)cM = 0$$

which can in turn be rewritten:

$$-\beta \log\left[\delta + (1-\delta)\exp\left(\frac{\mu^B - \mu^A}{\beta}\right)\right] + \beta \log \delta + (1-\delta)cM = 0. \qquad (A4)$$

To understand the comparative statics of optimal recruiting strategies, it is helpful to consider the interpretation of the three terms in (A4). The first term equals $\mu^A - \mu^C$; it is positive reflecting the direct advantage of Strategy A: the typical applicant is more productive when hiring is restricted to the As. The second term is negative; this cost of Strategy A reflects the loss of option value from sampling fewer applicants. This cost is larger in absolute value the smaller the share of the As in the population and the greater the variance in unobserved worker productivity. The final ter m is positive, reflecting the savings in application processing costs under Strategy A, relative to C. (It literally equals the total cost of processing all the B candidates' applications.)

Now, since the LHS of (A4) is strictly increasing in M, any parameter change that raises the LHS will reduce the value of $\tilde{M}$. It follows immediately that increases in c reduce $\tilde{M}$, and (given that the first term is positive) increases in μ^A - μ^B do the same. To see the effect of the variance parameter, β, it is instructive to initially consider some limiting cases. When the variance (β) approaches zero, both the first and second terms approach zero, leaving only the third, positive term; thus Strategy A is preferred regardless of any other parameter values. This makes sense: when worker productivity is perfectly predicted by the worker's type there is no point in interviewing any Bs because there is no possibility any B will be better than the first A who is interviewed.

Conversely, as the variance (β) approaches infinity, term 1 approaches unity, while term 2 falls without limit towards minus infinity. Now, the option value of Combined strategy C must therefore dominate, and A is dispreferred. Finally, to see that, in general, an increase in β acts to reduce the LHS of (A4), thus raising $\tilde{M}$ and making Strategy C more "likely", differentiate (A4) with respect to β to obtain:

$$\frac{\partial LHS}{\partial \beta} = -\log D \; + \; \left(\frac{\mu^B - \mu^A}{\beta}\right)(1-\delta)\exp\left(\frac{\mu^B - \mu^A}{\beta}\right) \; + \; \log\delta \; < \; 0 \qquad \text{(A5)}$$

Since the middle term in (A5) is negative, the sign follows from the fact that

$$D \equiv \left[\delta + (1-\delta)\exp\left(\frac{\mu^B - \mu^A}{\beta}\right)\right] > \delta .$$

Finally we turn to the case where $\tilde{M}$ falls into Region 2. Now $\tilde{M}$ is defined as the value of M that sets ($A2$) equal to zero, i.e. that satisfies:

$$\Pi^A(\delta M) - \Pi^C(N^*) \; = \; \mu^A + \beta \log(\delta M) - c\delta M - \mu^C - \beta \log(N^*) + cN^* = 0 \qquad \text{(A6)}$$

The fact that we are in region 2 implies that the LHS of (A6), like (A4), is monotonically increasing in M. Thus any parameter change that raises the LHS will reduce the value of $\tilde{M}$. To see how the preceding proofs are adapted to Region 2, consider the case of a small increase in c. This reduces $\Pi^A(\delta M)$ by δM, and reduces $\Pi^C(N^*)$ by (approximately, ignoring integer issues in applying an envelope result to $\Pi^C(N^*)$) by $N^* > \delta M$, thus raising the LHS, and reducing $\tilde{M}$. The comparative statics for μ^A - μ^B and β proceed analogously. ■

Proof of Corollary:

We consider the case of Region 1 only for brevity, and thus focus on equation (A4). By definition, the term $(\mu^B - \mu^A)/\beta$ in A4 is unaffected by a neutral increase in skill requirements. Thus, the only term affected by an increase in θ is β, which is proportional to θ. Holding $(\mu^B - \mu^A)/\beta$ constant, (A5) thus becomes:

$$\frac{\partial LHS}{\partial \beta} = -\log D \; + \; \log\delta \; < \; 0 .$$

Appendix 2: Data

As noted, our overall sample consists of all job ads which appeared on Zhaopin.com between May 16 and July 29, 2008, and between Dec 17, 2008 and Feb 28, 2009. At the end of each day, our program automatically searches for job ads that were posted on Zhaopin that day. The program starts at 11:30pm sharp each day for consistency. On the first day of data collection, all ads that were posted that day were kept. On subsequent days, all ads posted that day are compared with the master list of previously-posted jobs; since many such jobs are just renewals that are re-posted (employers can re-post and existing ad; this entails a small marginal financial cost but does require action on the employer's part), we do not download these refreshed jobs but maintain a count of the number of renewals that occur during this time period. A similar procedure was applied to the list of firms. As a result, our data have information on every job that was posted or renewed during this time period, linked to information about the firm posting the job. All of our regression analysis is restricted to the sample of jobs for which we have matching firm information. The matching rate varies somewhat across specifications but was about 80.2%

Age, gender and other job requirements were extracted from each job's html file. For example, in the case of gender, we look for "nue"(female) and "nan"(male) characters in the job description section of the file. We then constructed a match table summarizing about 1468 ways for a job ad to mention "nue"(female) and "nan"(male). After that, we use a program and this match table to derive the gender discrimination variable automatically. We consider our table quite exhaustive. In addition, we also visually check all the job ads that mentioned gender in a way that did not match these tables. Only about 100 jobs out of our entire sample fell into this category. For age variables, we search for "sui" (year of age); 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. (This is the case both when a single ad is for multiple vacancies and when it is not). Finally, our data on job ads was merged with a number of province-level characteristics, taken from 2000 Census and 2001 National Census of Basic Units of China accessed on November 2 2008 through http://www.acmr.com.cn's Support System for China Statistics Application.

To construct our "inflow sample" of job ads, we first examined the empirical distribution of dates that an ad first appears on the job site during our sampling period. As Figure A1 (which refers to our first sampling window, in the summer of 2008) shows, this distribution has a large spike on the first day we collected ads, then declines rapidly, reflecting the fact that most jobs "posted today" after our first day of data collection were in fact just repostings or renewals of jobs that had been posted earlier. After about a month, however, the empirical distribution of new jobs (that we have not seen before on the site) becomes quite constant. This suggests the sample of ads newly appearing on the site after that time are essentially all new; we thus define our "inflow" sample as all ads which appear in our data for the first time after June 1, 2008. For our second observation window, this date was January 1, 2009.

Figure A1: Flows of ads by date first observed

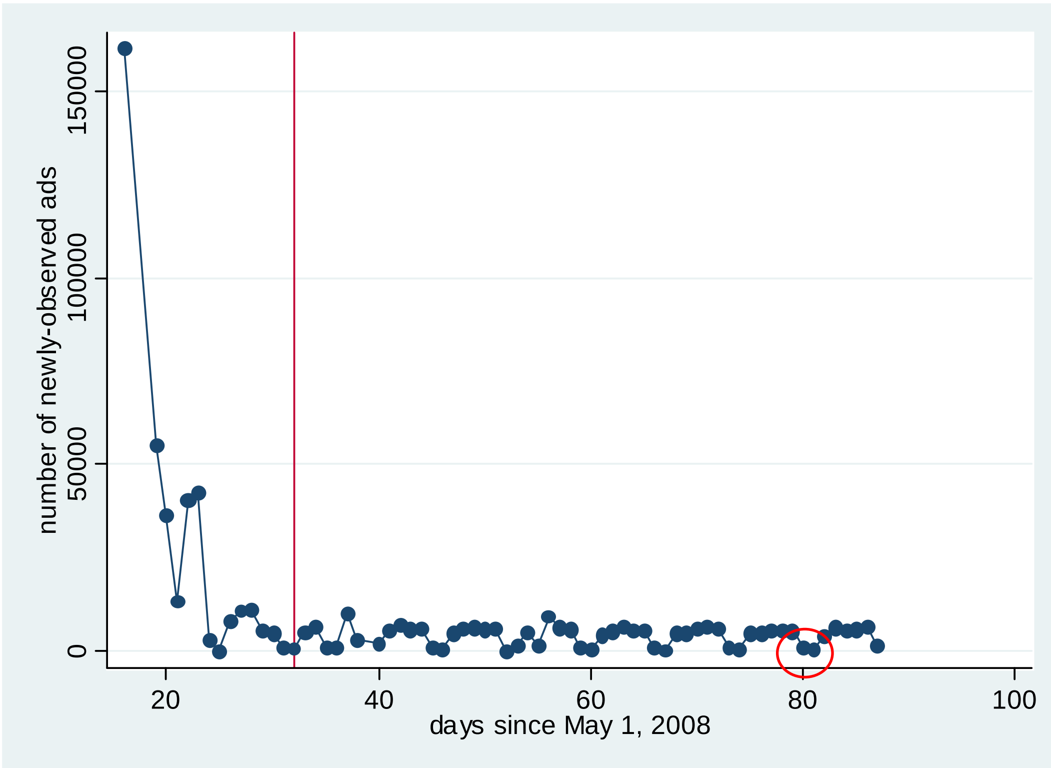

Notes:

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