Age and Gender-Based Screening in Employee Recruitment:

Evidence from Four Job Boards

Preliminary—please do not quote without permission

Miguel Delgado Helleseter

Peter Kuhn

Kailing Shen

November 6, 2013

When permitted by law and custom, employers sometimes engage in *ex ante* screening of job candidates using demographic criteria like age and gender. We study this practice using four samples of job ads, taken from job boards in China and Mexico. Consistent with Kuhn and Shen's (2013) screening cost model, we find a strong, negative skill-targeting relationship in all four data sets: employers are less likely to express an explicit age or gender preference in a job ad as the job's skill requirements rise. We also find that employers' age and gender preferences interact in a strong and consistent way across all four data sets: men's (revealed) relative desirability as workers rises dramatically between the ages of 18 and 40. While some of this pattern is can be linked to an age-dependent demand for feminine beauty, much of it is hard to explain with standard models of gender differences in labor markets.

This research is supported by National Natural Science Foundation of China through Grant 71203188, "Impacts of *Hukou*, Education and Wage on Job Search and Match: Evidence Based on Online Job Board Microdata". We thank Anne Pisor for interesting discussions on related literature in anthropology, and participants in UCSB's and UC Berkeley's labor lunch seminars for helpful comments.

1. Introduction

Americans, and persons in most developed countries, seem to strongly prefer to be evaluated on their individual merits, not based on their membership in a particular demographic group. Faced with large numbers of applications, however, employers everywhere can find it difficult to evaluate each application they receive on its detailed merits. This creates a demand for simple filters that allow firms to narrow down applicant pools before engaging in more intensive scrutiny. While some filters –such as minimum levels of education or experience—are commonly used throughout the world, others –such as age, race and sex—are explicitly prohibited in many developed countries. As evidenced by a voluminous literature devoted to detecting labor market discrimination in developed nations, this can make it hard for researchers to infer how firms really value workers of different ages, genders and races.

In most of the world's labor markets, however, the above prohibitions do not exist; indeed it is commonplace in many countries for employers to include explicit sex and age preferences in job ads. Job ads from those countries may therefore provide a unique laboratory from which to learn how employers value those worker characteristics. Pursuing this idea, Kuhn and Shen (2013) (henceforth KS) studied firms' use of gender-targeted job ads in China. In their paper, KS present a simple model in which application processing costs sometimes motivate firms to invite only applicants from a particular demographic group; a key prediction of their model is that this tendency –to narrow one's search to a preferred group—should be less frequent in jobs that demand higher skill levels than lower ones. The intuition is simple: limiting one's search to the higher-expected-value group becomes less attractive as it becomes more important to identify the single best candidate for the job.

In their data, KS document a *negative skill-targeting relationship* that is consistent with this simple idea. Specifically, in a cross-section of job ads, the probability that an ad is explicitly targeted at *either* gender declines with three different indicators of the job's skill level: the required levels of education and experience, and the level of the offered wage. This relationship is robust across samples and specifications, and persists even within firm*occupation cells. In other words, the same firm, advertising at two different times for (say) sales personnel, will be less likely to express a gender preference when the sales job requires a higher level of skill.

In addition to shedding light on the conditions under which employers will engage in *ex ante* screening of job candidates using characteristics like age or gender, demographically-targeted job ads also shed light on the conditions under which employers believe that men are more suitable employees than women, or vice versa. In their analysis, KS identified a number of patterns in firm's relative preferences for men versus women, but did not investigate them in detail, in part because of concerns that their sample of job ads was relatively idiosyncratic: an

Internet job board in China (Zhaopin.com), which disproportionately serves a highly skilled labor market.

Motivated by that observation, this paper has three main goals. First, we wish to see whether the negative skill-targeting relationship identified by KS for gender is robust across labor markets. If so, that would provide support for the simple model of directed recruiting proposed in their paper. Second, we extend KS's analysis of gender-targeting to the case of age-targeting; if their screening model is valid it should apply to age-based screening as well. Finally, in this paper we are particularly interested in whether there are any universal patterns in the conditions under which firms prefer women versus men, and older versus younger persons as employees. If such patterns exist, they might shed light on more fundamental determinants of age and gender differentials in labor markets. To accomplish these goals, this paper introduces data from three job boards in addition to KS's Zhaopin.com: two from China, and one from Mexico. Aside from the intrinsic value of replication, the two new Chinese datasets are of interest because they serve workers with a much more representative level of skills. In fact, together they represent the entire private-sector labor market of a medium-sized Chinese city. Our Mexican data is similar to Zhaopin.com in that it is a national board serving a highly skilled population. It allows us to see whether KS's main results extend to a different culture (which, however, shares the feature that it tolerates explicit age and gender-based filtering in job ads).

Our main results are as follows. First, we find strong confirmation of the negative skill-targeting relationship: employers are less likely to express an explicit age *or* gender preference in a job ad as the job's skill requirements rise. This is true both across data sets --i.e. the job boards that serve less-skilled workers exhibit much more age- and gender-targeting-- and within data sets, both in the unadjusted means and in the presence of detailed controls including firm*occupation fixed effects. Second, we identify a new empirical regularity concerning the *direction* of firms' age and gender preferences that is present in all our datasets and robust to the same detailed controls. Specifically, as workers age from around 18 to 40, employers' explicitly advertised gender preferences shift strongly away from women and towards men. We refer to this phenomenon as *age-dependent gender preferences*, and attempt to understand it more fully in the paper.

The remainder of the paper is organized as follows. Section 2 describes all four of the data sets at our disposal. Sections 3 and 4 present our main results, first as simple means, then in a regression framework. Section 5 considers alternative explanations for the negative skill-targeting relationship, while Section 6 explores possible explanations of the age-dependent gender preferences in our data. Section 7 briefly explores the relative role of firms versus jobs in accounting for advertised gender preferences in our four data sets, and Section 8 concludes.

2. Data

The first of the three Chinese data sets used in this paper is a sample of job ads taken from Zhaopin.com, the third-largest Internet job board in China. For comparability, we use the same analysis sample as KS (2013), which is based on all unique ads posted in four five-week observation periods in 2008-2010. Detailed descriptions of this data are available in KS; key for our purposes here is that Zhaopin serves a highly skilled labor market, and operates nationally. Since skill requirements play a central role in our analysis, all of our analysis samples --including the Zhaopin sample-- are restricted to the large majority of job ads in each data set that explicitly specify an education requirement.

Our two new Chinese data sources are job boards serving the city of Xiamen, a southern coastal city about the size of Los Angeles.[1] In part because Xiamen was one of the five economic zones established immediately after China's 1979 economic reforms, it is highly modernized relative to other Chinese cities, with an economy based on electronics, machinery and chemical engineering. The city hosts one of Dell Computer's five global production and service centers. One of our two Xiamen-based job boards, XMRC (http://www.xmrc.com.cn), is a for-profit company sponsored by the local government.[2] Its mandate is to serve the market for skilled workers in the Xiamen metropolitan area. XMRC operates like a typical U.S. job board: both job ads and resumes are posted online, workers can submit applications to specific jobs via the site, and firms can contact individual workers through the site as well.

In contrast, our other Xiamen-based job board, XMZYJS (the Xia-Zhang-Quan 3 city public job board, www.xmzyjs.com), is operated directly by government employees of the local labor bureau. Unlike XMRC it has a history as a brick-and-mortar employment service, and its mandate is to serve the less-skilled portion of the area's labor market. In contrast to XMRC, XMZYJS is purely a job-posting service: workers cannot post resumes or apply to jobs through the site. In fact, while XMZYJS now posts all its job ads online, many of these ads are viewed in XMZYJS's offices by workers who visit in person. This is done both on individual computer terminals and on a large electronic wall display. Applications are made by calling the company that placed the ad, or by coming to a specific window on XMZYJS's premises that has been reserved by the employer at a posted date and time.

[1] Based on the 2010 Chinese and U.S. censuses, Xiamen's metropolitan area (which our data represent) is composed of three major cities --Xiamen, Zhangzhou and Quanzhou—with a combined population of about 16.5 million, compared to 17.8 million for the Los Angeles metropolitan region. In contrast, the 2012 national populations of Norway, Sweden and Denmark (frequently used in empirical labor market studies that are accurately described as 'nationally representative') are 5.0, 5.6 and 9.6 million respectively.

[2] XMRC's direct sponsor is the Xiamen Human Resources Service Center. Like all our data sets, XMRC serves only private sector employers (including State-Owned Enterprises, but excluding government employees). Recruiting for public-sector jobs takes place via a different process.

Together, the XMRC and XMZYJS sites provide a fairly complete picture of the active labor market of Xiamen. By design, XMZYJS aggregates job postings from all local and specialized job boards for less-skilled workers in the metropolitan area, and XMRC is the main job board for skilled workers in the area. While there is potentially some cross-posting of job ads across the two sites, descriptive statistics on the types of jobs on offer suggest the sites do, indeed, serve very different populations. For our analysis, XMRC provided us with all the ads for jobs in Xiamen that received their first application between May 1 and October 30, 2010. Details on how we constructed an analysis sample for this and our other new datasets are provided in Appendix A. The most important restrictions involved dropping observations with missing education requirements and missing occupation information. XMZYJS, on the other hand, provided all the ads for jobs that appeared in calendar 2010; very similar restrictions were used to construct its analysis sample.

Summary statistics for our three Chinese samples are presented in columns 1-3 of Table 1; they are arranged in order of increasing skill of the workforce served. All told, we have 141,284, 39,746, and 1,051,706 ads in our XMZYJS, XMRC and Zhaopin samples respectively. Reflecting the job boards' varying skill levels, education requirements are lowest in XMZYJS, with (60+33=) 93 percent of ads requiring a high-school education or less, compared to 49 percent of ads on XMRC.[3] XMRC, in turn, hosts considerably less-skilled ads on average than Zhaopin.com, where only 13 percent of ads required high school education or less. These differences in skill are mirrored in the advertised wage levels, which are lowest in XMZYJS and highest in Zhaopin. Interestingly, the share of ads that post a wage declines monotonically from 100% in XMZYJS (where employers are required to provide this information), to 100-58 = 44 percent in XMRC and 14 percent in Zhaopin. This pattern is consistent with Brencic's (2012) finding that wage posting is negatively associated with job skill levels. Also consistent with the skill differences across the three Chinese datasets, requested experience levels are much lower in XMRC than in Zhaopin, and are not even a designated field on XMZYJS ads.

In contrast to Zhaopin and XMRC, all XMZYJS ads specify the number of positions that are open, and mean number of positions (8.7) is vastly higher than on the other two Chinese job boards. Again, this almost certainly reflects the fact that XMZYJS is where Xiamen's employers go to recruit production workers and to fill other less-skilled positions. That said, "ideal" job candidate ages (taken as the midpoint of the minimum and maximum requested ages when both are stipulated) are quite similar across all our data sets, ranging from 27.7 to 30.7 years. Indeed, requests for workers under 18 and over 40 are extremely uncommon on all four of these job boards; in consequence all of our results in the paper should be interpreted as

[3] Together, these statistics are roughly consistent with available data that is representative of Xiamen's labor force. In particular, statistics from the 2005 1% National Population Sample Survey indicate that 73.6 percent of workers in Xiamen had a high school degree or less.

applying to workers between 18 and 40 years of age. The relative youth of the workers sought on these job boards could reflect a number of factors, including the fact that entry-level positions are likely to be overrepresented in a sample of vacant jobs.

Finally, turning to age- and gender-targeting in job ads, row 1 of Table 1 shows that over two thirds (100-28=72 percent) of job ads on XMZYJS are gender-targeted, compared to 38 percent on XMRC and 10.5 percent on Zhaopin.com. This pattern is strongly consistent with the negative skill-targeting relationship identified in KS. Interestingly, however, while the share of ads favoring men versus women is roughly equal on both XMRC and Zhaopin, many more ads favor men than women on XMZYJS (42 versus 30 percent of the total). Also consistent with a model in which *ex ante* demographic screens decline in value as skill levels rise, the share of ads that specify a minimum age and the share that specify a maximum age both decline monotonically as we move from the least-skilled job board (XMZYJS) to the most skilled (Zhaopin). In sum, KS's prediction that age- and gender-targeting of job ads should decrease with a job's skill level is strongly confirmed when we compare mean levels of targeting across job boards serving very different skill levels in China. Related, statistics from our least-skilled data set (which is also the most representative of China's entire workforce) suggest that age- and gender targeting of job ads is extremely widespread in China as a whole, with (as already noted) 72 percent of job ads directed at a specific gender and 77 percent stipulating a maximum age.

Our Mexican data is a sample of job ads posted on Computrabajo.com.mx.[4] Of the new data sets explored in this paper, the Computrabajo data are most similar to Zhaopin in the sense that they come from a national online site that disproportionately serves highly skilled workers. Thus, especially in comparison with Zhaopin, they should allow us to ascertain whether the theoretically predicted skill-targeting relationship, as well as patterns of employers' perceived gender-suitability across job and worker types transcend national and cultural boundaries. Computrabajo is also the only one of the four datasets in Table 1 where we are able to study how firms' advertised gender preferences interact with their preferences for a worker's marital status.

To construct an analysis sample from the Computrabajo website, we collected advertisements daily for approximately 18 months between early 2011 and mid-2012 using a web crawler. Both the standardized fields and the open text portions of each ad were parsed to extract variables for the analysis.[5] We use the universe of unique advertisements posted, where unique advertisements are defined as ads that appeared once during this time, plus

[4] Computrabajo.com actually operates job boards in 20 Spanish-speaking countries. We picked Mexico because it was the largest.

[5] See Delgado Helleseter (2013) for additional details on how data were collected, observations were identified, and variables were extracted from each advertisement.

identical advertisements appearing on multiple dates (in which case only the first appearance is included in the sample). Additional sample restrictions are similar to those used in our Chinese datasets, and are described in Appendix A.

Overall, our Computrabajo analysis sample contains 90,561 ads. Consistent with its position as a national job site, a large share of ads (42 percent) requires a university education. Interestingly, this share is essentially identical to the share in Zhaopin, though jobs requiring high school or less are much more common on Computrabajo. Consistent with this skill difference, Computrabajo jobs require less experience than Zhaopin jobs, and are more likely to post a wage. Computrabajo ads are also gender-targeted three times as often as Zhaopin ads (at 32 versus 10.5 percent).[6] As in Zhaopin, gender-targeted ads are equally split between men and women. Of our four datasets, Computrabajo and Zhaopin are the only ones where we have information on employers' requests for physically attractive workers: at 13.9 versus 7.7 percent of ads, this number is considerably higher in Computrabajo than in Zhaopin.[7]

Finally, a unique feature of the Computrabajo data is the availability of information on the employer's preferred marital status. Overall, 3.4 percent of Computrabajo ads explicitly indicate the worker's preferred marital status. Of these requests, 2.1 percent are for married persons and 1.3 for single persons. Despite being relatively uncommon compared to age- and gender-based screens, as we shall see screening on marital status interacts in a strong and interesting way with those screens in Mexico. [8]

3. Results: Means

Tables 2-5 show the means of four variables –the share of job ads requesting women, the share requesting men, the share expressing no gender preference, and the share stipulating a desired age range for the employee—for our four data sets in turn. Turning first to the *Zhaopin* data in Table 2, we see four clear patterns. The first, shown in Panel A, is the *negative skill-targeting relationship*: the share of jobs that are gender-targeted declines strongly with all three indicators of the job's skill level (education, experience, and the wage). Specifically, (100-

---

[6] The average advertised salary in Computrabajo is 7,642 Mexican pesos per month (approximately 580 U.S. dollars on average based on exchange rates for that period), which is well above average salaries in Mexico. It is about 200 U.S. dollars above the average in the XMRC data.

[7] The beauty indicator in our Mexican data indicates the presence of one or more of a number of closely-related keywords, the most common of which is "buena presentation". In Zhaopin it indicates a request for *xingxiang*, which translates roughly as having a pleasing form or image.

[8] Following KS's Zhaopin analysis, we include all expressed degrees of preference in our indicators of employer targeting on gender, age and marital status. In all cases, however, the most common way to express these preferences is similar to how education and other requirements are posted. For example, just as it is typical to write "Education: high school", many ads expressing a gender preference would write "Sex: male", or simply "male". In Mexico it is possible to extract additional information on the employer's gender preferences from gendered job titles, for example whether the ad requests an 'abogado' or an 'abogada'. Since we do not attempt to extract this information, our statistics will underestimate the amount of gender-targeting in Mexican job ads.

76.6 =) 23.4 percent of jobs requiring high school or less specify a preferred gender, compared to only 6.2 percent of jobs requiring a university education. Similarly, the share of jobs that are gendered in this way falls from 26.6 percent for jobs paying under 1500 RMB/month to 7.1 percent for jobs paying over 8000 RMB/month, with a weaker negative relationship for experience levels. In all cases except experience, this decline applies both to the share of jobs targeted at men and the share targeted at women. Confirming the negative skill-targeting relationship, the probability that an ad is age-targeted also falls with skill, though the relationship is much stronger for the case of education than our other two indicators of skill.

Second, panel B of Table 2 shows that –among jobs that specify an age range—the direction of employers' gender preferences is highly age-dependent. Specifically, if a firm is looking for a worker under the age of 25, chances are one in three that it is explicitly searching for a woman. The comparable figure for men is 6.0 percent. On the other hand, if a firm is looking for a worker over the age of 35, the chances that it is explicitly looking for a man are 18.8 percent, compared to 4.1 percent for women. As noted, we refer to this phenomenon as *age-dependent gender preferences*.

Third, Panel B of Table 2 shows that *age and gender screens are complements*, in the sense that employers tend to use these screens together. Thus, for example, (100-72.1 =) 27.9 percent of ads that specify both a maximum and minimum age are gender-targeted, compared to only 5.9 percent of non-age-targeted ads. And finally, the Zhaopin data show that employers value physical attractiveness much more among women than among men: 23.9 percent of ads requesting beauty are directed explicitly at women, compared to 3.8 percent for men. Later in the paper, we explore the role played by these highly gendered preferences for beauty in explaining employers' age-dependent gender preferences.

Having reviewed these patterns in the *Zhaopin* data, we next ask to what extent they appear in the three new datasets now at our disposal. Turning first to the *skill-targeting relationship*, we see first that it is strongly confirmed in the *XMRC* data. While the overall levels of age- and gender-targeting are higher in XMRC than Zhaopin --reflecting its lower mean level of skill--, the first three rows of Table 3 show that the prevalence of gender-targeting declines monotonically with the job's required education level, from 100-53 = 47 percent of jobs requiring high school or less to 24 percent of jobs requiring university or higher education. As in Zhaopin, ads requesting men and ads requesting women both become less common as education levels rise, and at roughly the same rate. Similarly, the share of ads that are gendered falls from 46 percent for jobs requiring one year of experience to 25 percent when more than five years are required, with an even larger decline as offered wage levels rise from

1000 to over 5000 RMB/month.[9] The share of ads that are age-targeted also declines with all three indicators of the job's skill level.

In the *XMZYJS* data, the least-skilled of our four datasets, we have only two measures of skill: education requirements and the offered wage.[10] While an even higher share of jobs is gendered in this data, the share of jobs that is gendered is again much lower at the highest education and wage levels than the lowest.[11] A similar pattern is evident for age-targeting, though as for gender-targeting, the very lowest education and wage categories do not fit neatly into the overall trend. Finally in the *Computrabajo* data, the incidence of age-targeting declines quite smoothly with all three of our skill indicators (education, experience and the wage). There is no association between experience and gender-targeting, though the gender-targeting does appear to fall with education and offered wages.[12] In sum, the negative skill-targeting relationship first observed by KS in the Zhaopin data appears to be confirmed within each of the three new datasets studied in this paper.

Is the phenomenon of *age-dependent gender preferences* also confirmed in our three new datasets? Panel B of Tables 3-5 shows the share of job ads that are targeted at men and women, separately by age in the XMRC, XMZYJS and Computrabajo data respectively. These results for all four datasets are also graphed in Figure 1. While the overall levels of gender-targeting differ substantially across these four datasets, the same pattern of strongly age-dependent gender preferences is visible in all of them: Comparing ads requesting workers under 25 to those requesting workers 35 and older, the share of ads explicitly requesting female workers declines from 54.7 to 6.7 percent in the XMRC data, from 46.4 to 22.0 percent in the XMZYJS data, and from 28.9 to 8.0 percent in the Computrabajo data. The share of ads explicitly requesting men, on the other hand, *increases* from 12.3 to 35.4 percent in the XMRC data, from 32.2 to 58.3 percent in the XMZYJS data, and from 14.0 to 20.9 percent in the Computrabajo data.[13]

Other noteworthy features of the descriptive statistics in Tables 2-5 include the following. First, age and gender screens are complements on all four job boards, in the sense that they tend to be used together. For example in XMRC, (100-49.4 =) 50.6 percent of ads that

[9] In the XMRC data, like the Computrabajo data, we cannot distinguish in any meaningful way between ads that require no experience and those that do not list an experience requirement.

[10] Recall that in contrast to XMRC, all XMZYJS ads post a wage; in addition since our education categories are quite course the wage is probably the best overall skill measure in this data set. Reflecting the low skill levels of these jobs, XMZYJS does not provide a specific field for employers to list experience requirements.

[11] There is some non-monotonicity involving the lowest education and wage categories, but this is relatively minor compared to the overall trend.

[12] There is some nonmonotonicity affecting the lowest education and wage categories; interestingly this disappears in the presence of controls. See Section 4.

[13] Very few ads on any of our job boards list desired ages above 40.

specify a desired age range also specify an explicit gender preference, compared to only 27.5 percent of ads do not specify an age preference. The equivalent numbers in XMZYJS are 76.9 and 14.3 percent. In Computrabajo they are 38.0 and 14.1 percent; in Zhaopin 27.9 and 5.6 percent. Second, like Zhaopin, the Computrabajo data allows us to study firms' advertised preferences for physical attractiveness, and how these preferences interact with firms' gender preferences. As in the Zhaopin data, Table 5 shows that beauty is requested much more frequently in job ads where firms are seeking women than when they are seeking men. Specifically, while jobs that do not request beauty invite men and women to apply with roughly equal frequency, jobs that request beauty request women more than twice as often as men (the rates are 28.5 versus 12.2 percent).

Third, a unique feature of the Computrabajo data is that they include information on firms' advertised preferences for marital status. As for beauty, we find that firms' marital status preferences interact very strongly with their gender preferences. Specifically, advertisements requesting a person who is single request women 60.1 percent of the time while requesting men only 14.0 percent of the time. In contrast, advertisements requesting a married person request women 8.4 percent of the time, but request men 52.9 percent of the time. Thus, according to the Computrabajo data, *employers want women to be single (and good-looking) and men to be married*. In what follows, we'll assess the role of these preferences for looks and marital status in accounting for the pervasive phenomenon of age-dependent gender preferences in our data.

4. Results: Regressions

4.1 Gender Targeting Regressions

Tables 6-9 present identically-specified regressions for each of our four datasets in turn. In columns 1-4 the dependent variable is an indicator for *whether* the ad is gender-targeted (regardless of direction); in columns 5-8 the outcome is whether the ad is age-targeted (regardless of what age is requested). In other words, if P^F is an indicator for whether the ad specifically requests women, and P^M is an indicator for whether it requests men , then the dependent variable in columns 1-4 is just P^F+P^M, which equals either zero or one. In columns 5-8, the outcome equals one if an ad specifies both a maximum and minimum age, and zero otherwise.[14] Four specifications are presented for each of these two dependent variables. The first (in columns 1 and 5) includes occupation fixed effects.[15] The goal is to see whether the

[14] Similar results are obtained when we use less restrictive definitions of age-targeting (e.g. an indicator for whether *any* sort of age preference is specified).

[15] Depending on the dataset, industry and province or state effects may also be present, depending on relevance and availability. (For example, province effects are irrelevant in our two samples from the city of Xiamen, and Zhaopin is the only dataset with an industry variable.

patterns identified in Tables 2-5 and in Figure 1 are primarily a consequence of the type of work that is done: perhaps some types of work are highly gendered, and others not, and the latter just happen to be more skilled. Columns 2 and 6 add firm fixed effects to this specification: perhaps the skill-targeting relationship results mostly from a pattern where the firms that abstain from age-and gender-targeting (such as, for example, foreign-owned firms) disproportionately happen to employ skilled workers for reasons unrelated to skill *per se*.

Finally, the remaining columns (3, 4, 7 and 8) include fixed effects for "job cells", i.e. for the interaction of firms with occupations. Here, the estimates tell us whether the same firm, advertising at two different times for the same occupation (say, sales), is more likely to gender-target its ad when seeking a highly educated salesperson than a salesperson with less education. If the negative skill-targeting relationship persists even within job cells, this suggests that it is more likely driven by a factor that is directly related to skill levels, rather than factors that vary across the different types of jobs men and women occupy within the same firm. Since wage posting is universal in only one of our four datasets, columns 3 and 6 present these regressions without controlling for the offered wage; columns 4 and 8 then add an offered wage control at the cost of a substantial reduction in sample size in some datasets. [16]

For the *Zhaopin* data, Table 6 shows that the probability an ad is gender-targeted is negatively related to all three of our skill measures (education, experience and the wage), across all three specifications. The probability of age targeting is also negatively related to education requirements across all three specifications, though no robust pattern is present for experience and the offered wage. The *XMRC* data in Table 7 show negative and statistically significant effects of education and experience on both gender and age targeting in the presence of occupation fixed effects (columns 1 and 5), and for the most part when occupation and firm effects are entered separately (columns 2 and 6). These become insignificant for age – but not for gender—when we look within job (firm*occupation) cells, perhaps because there is not much variation in education and experience requirements within those cells. Notably, however, even within these detailed cells, jobs that offer higher wages are much less likely to be gender-targeted *or* age-targeted. This is consistent with the notion that most of the variation in skill levels within job cells is associated with different offered wage levels than in our broad measures of education and experience qualifications.

The *XMZYJS* data in Table 8 show a robust, monotonic, quantitatively large, and highly statistically significant negative relationship between a job ad's education requirements and the probability the ad is gender targeted. The same is true for age-targeting, and both the age- and gender-targeting relationships with education are present within firm*occupation cells as well

[16] Columns 1-3 of Tables 6-9 correspond, respectively, to columns 1 ,3 and 5 of KS's Table VI. A parallel correspondence applies to our analysis of firms' gender preferences in Table 10-13.

as less saturated specifications. The XMZYJS data, however, do not show a statistically significant effect of offered wages on age- or gender-targeting in the presence of education controls. Turning to the Computrabajo data in Table 9, the non-monotonic patterns observed in the raw data for the effects of education are also observed in the column 1 regressions. However, as we add additional covariates in columns 2 and 3 the estimated pattern becomes monotonic, consistent with the skill-targeting hypothesis. Jobs requiring university are much less likely to be age or gender-targeted in the presence of occupation*state or job cell fixed effects. These effects persist for gender but not for age when an offered wage control is added. In both cases, however, the offered wage has a strong, negative effect on targeting, suggesting it is the best measure of skill within job cells. Unexpectedly, the lowest experience category (1-3 years) actually has more gender- and age-targeting than the omitted category. This is most likely because (in contrast to XMRC) experience is not a dedicated field in the Computrabajo data. In consequence the omitted experience category contains both jobs that require no experience and jobs that do, but where we failed to detect the requirement.

In sum, Tables 6-9 show that in almost all instances, the negative skill-targeting relationship for both age and gender in the unconditional means of all four of our datasets persists in the presence of detailed controls for the types of work that is being performed and the type of employer posting the ad. Indeed, in most cases the relationship persists even within firm*occupation cells, suggesting that the phenomenon is closely tied to skill levels *per se*, rather than the fact that skilled workers do different types of jobs –perhaps ones where 'gender matters less' for productivity—than less-skilled workers. We explore this 'gender matters less in skilled jobs' hypothesis further in Section 5.

4.2 Regression estimates of age-dependence in firms' gender preferences

In Tables 10-13 we turn our attention from the *presence* of gender targeting to its *direction*: under what conditions do job ads favor men versus women? Following KS, we regress a simple outcome measure equal to (P^M-P^F) on the desired age specified in the ad and other covariates. This outcome variable equals -1 when the job requests women, zero when the ad is not gender-targeted, and 1 when the ad requests men; KS show that under reasonable conditions this approach reveals the determinants of firms' underlying assessments of men's and women's relative suitability for the job being advertised.[17] With the exception that all columns now include the indicators of job's requested age level as our main regressors of interest, the specifications are identical to Tables 6-9. As for our skill-targeting regressions, the goal is to measure to what extent the strong age-dependence of firms' gender preferences

[17] The key conditions are that roughly the same number of ads are targeted at men versus women in the sample as a whole, and that the distribution of men's and women's unobserved relative values across jobs is symmetric. That said, very similar results are obtained if we estimate ordered probit models, or if we simply model the probability of preferring men conditional on stating a gender preference.

stems from firms' tendencies to use men and women for different types of work (as measured by occupation or firm*occupation fixed effects and skill level controls), or is more directly related to the worker's age *per se*.

In the *Zhaopin* data (Table 10), columns (1) and (2) show that firms' preferences tilt strongly towards men as workers age between 18 and 40, even in the presence of occupation*firm fixed effects. Interestingly, however, adding a control for the offered wage in column (3) suggests that –at least at fixed wage levels—this age-dependence is more strongly linked to firms' desired experience level than the worker's age per se. This suggests that, at least in Zhaopin, firms' within-job-cell tendency to prefer men when seeking older workers can be accounted for by a strong tendency to request men (relative to women) *when firms are seeking experienced workers*. As we shall see, this is not the case in our three other datasets, where age *per se* seems to matter in all specifications.

The *XMRC* data in Table 11 show a strong positive effect of desired age on firms' preferences towards men in all specifications; further the magnitude of this effect is highly stable as we add more detailed controls. The effect is also large in magnitude: ads for workers over 35 have a differential probability of hiring men (P^M-P^F) that is .6 units greater than ads for workers under 25. The same is true in both the *XMZYJS* and *Computrabajo* data (Tables 12 and 13), though the estimated magnitudes are lower, reflecting the lower overall prevalence of gender targeting in these job boards.

In sum, Tables 10-13 show that the pattern of strong age-dependence in gender preferences in the means of our four data sources persists in the presence of detailed controls for the types of work that is being performed and the type of employer posting the ad. Indeed, in all but one case the relationship persists even within firm*occupation cells, suggesting that the phenomenon is closely tied to workers' age *per se*, rather than the fact that older workers do different types of jobs within firms. In the one exception –the Zhaopin data—firms' advertised preferences for experience are better predictors of the direction of their gender preferences within job cells than their age preferences. We explore possible causes for these age- or experience-dependent gender preferences in Section 6.

5. Patterns in the Tendency to Age- or Gender-Target: Alternative Interpretations

So far we have argued that the robust negative skill-targeting effect observed in all four of our data sets is a direct consequence of a higher level of skill: Since higher skill levels 'raise the stakes' –making it more important to identify the best job candidate—firms should search more broadly as jobs' skill demands (indexed by θ in KS's model) rise. That said, KS's model also identifies a number of other factors that are predicted to affect the use of demographic screens in the hiring process. To the extent that these additional factors covary in the right direction with a job's skill level, they could also explain the negative skill-targeting relationship we see in the data. In this section we briefly discuss the possible role of these factors, which are application processing costs (*c*), the expected number of applicants per position (*N*), and the unexplained variance across jobs in their relative suitability for men versus women (σ_v).[18]

Turning first to application processing costs (*c*), KS's model predicts that *ex ante* screening should become more common as these costs rise (because discouraging one group from applying saves on application processing costs). Since both intuition and available evidence suggest that application processing costs rise with jobs' skill levels, it is clear that uncontrolled covariation of processing costs with job skill levels cannot account for the negative skill-targeting relationship that is present in all our data sources.[19] On the other hand, KS's model also predicts that *ex ante* screening should become less common as the number of applicants per job (*N*) shrinks. Thus, unmeasured covariation between skill levels and labor market tightness could explain the negative skill-targeting relationship if labor markets for skilled workers are on average tighter. Again, the intuition is simple: why would you rule out an entire group of applicants when applicants on the whole are scarce?

A final possibility raised by KS's model is that σ_v, the variance across jobs in men's and women's relative ability to perform them, might be greater in less-skilled than more-skilled jobs. In other words, perhaps 'gender matters less' for performance in skilled than in unskilled jobs. For example, if men and women are more different physically than mentally, jobs requiring manual labor might be more gender-specialized than other jobs.[20] KS provide evidence against this interpretation by dropping all jobs likely to require any physical labor from the Zhaopin data set; when they do so, the negative skill-targeting relationship is unchanged, both within and across jobs. They also devise a simple test between the σ_v–based and θ –based explanations of the negative skill-targeting relationship, using data on jobs that are highly gendered (i.e. jobs where more than half the ads explicitly requests one of the two genders). When we perform similar tests in the current data, as in KS our results favor a direct effect of skill (θ). Details are provided in Table O-1 of the on-line appendix.[21]

---

[18] A fourth possibility suggested by KS's model is that the idiosyncratic variation in applicant quality (σ_ε) is higher at higher skill levels. In practice this has very similar implications to, and is very hard to distinguish from the direct effects of skill demands (θ). (The distinction hinges on whether there is a bigger difference between the quality of a good and bad lawyer than a good and bad legal assistant, *over and above* the difference that follows from the jobs' skill demands.). Interested readers should consult KS for additional details.

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

[20] Note that effects of this nature would need to occur within firm*occupation cells to explain the regression results in all four data sets at our disposal.

[21] Available at www.xxxx.edu

In sum, we can easily rule out application processing costs that vary by skill as explanations of the robust skill-targeting effect in all our data sources, and available evidence suggests that the phenomenon is not caused by a tendency for men and women to be more fundamentally similar to each other when performing skilled versus unskilled work. We do acknowledge that the skill-targeting phenomenon could be attributable to systematically greater tightness of labor markets at higher skill levels, in addition to the direct effects of skill demands. But since the negative skill-targeting relationship appears in all four of our data sets, any such tendency of labor market to 'tighten up' as skill demands rise would need to be quite universal to explain our results.

6. Patterns in Firms' Relative Preferences for Men: Some Explanations

One possible explanation for the fact that firms disproportionately value youth among female workers is employers' demands for physical attractiveness, which could be strongly focused on young women.[22] Another is a differential effect of marital status on men's and women's (actual or perceived) productivity. If marriage is perceived to raise men's productivity, and to have weaker positive (or even negative) effects on women's productivity, then the fact that both genders are rapidly entering marriage between age 18 and 40 could also explain this sharp reversal in employers' advertised preferences.[23]

To explore these issues, Table 14 presents regressions from Zhaopin and Computrabajo that add controls for whether the job ad requests beauty, as well as for the ad's requested marital status to the regressions in Tables 10 and 13. In both data sets and in all regression specifications, employers' requests for physically attractive workers are highly gendered: requests for beauty strongly increase the chances a firm is looking for a woman, relative to a man. Similarly, the Computrabajo data show that employers strongly 'want men to be married and women to be single'. Neither the estimated beauty nor marital status effects vary appreciably in magnitude as controls are added, including controls for occupation*firm cells and for the level of offered wages. This suggests that firms' preferences for beauty and marital status are not easily accounted for by differences in the type of jobs men and women occupy. Instead, firms' tendencies to prefer women who are attractive and single over other women, and their tendency to prefer married to single men, are just as strong within job cells as they are overall.

---

[22] A number of studies have demonstrated a positive productivity effect of beauty in various jobs, and in many cases disproportionately for women. For example, Landry et. al (2006) document a large positive effect of beauty on women's productivity in a field experiment on charitable giving.

[23] A similar, well known argument also applies to gender differences in the effect of marriage on expected *turnover*. Marriage can also reduce women's relative value to firms if unmarried women are valued in the workplace for sexual or marriage-market reasons.

Comparing the age coefficients in columns 1-3 of Table 14 to their counterparts in Table 10 shows very little effect of adding a beauty control in the Zhaopin data. In the presence of offered wage controls (column 4), however, the insignificant age effects in Table 10 diminish in magnitude (and remain statistically insignificant) when a beauty control is added in Table 14. As in Table 10, this appears to be a consequence of a much stronger, positive effect of desired *experience* on preferences for men when offered wages are held constant. Overall, however, adding controls for whether beauty is requested in the job ad has little effect on the estimated relationships between age and employers' gender preferences in the Zhaopin data. Turning to the Computrabajo data in part B of Table 14, the age coefficients are highly significant in all specifications, and in all cases only slightly smaller in magnitude than in Table 10. Taken together, the results in Table 14 suggest that the age effects on employers' gender preferences observed in the raw data and in the Table 10-13 regressions are not merely a consequence of firms' gendered preferences for beauty or marital status, but might indicate something more fundamental about the effects of age *per se*.

Additional insight into the motivations behind firms' preferences for attractive young women is available from Figures 2-5, which present occupation fixed effects from the regressions in column 1 of Tables 10-13.[24] To illustrate the effects of customer contact on firms' gender preferences, occupations that involve customer contact are denoted by triangles.[25] A common pattern in all four figures is that the customer contact occupations consistently fall below the regression line, and in most cases significantly so. One other occupation that consistently lies significantly below the line is administrative staff. Conversations with job board officials revealed that this occupational category consists mostly of secretaries, receptionists and other office assistants. Together, these patterns suggest that customer and supervisor preferences for interacting with attractive, young (and possibly unmarried) women may also play an important role in advertised age and gender preferences.

To see to what extent the demand for attractive, young (and in some cases unattached) women in customer-contact and administrative jobs can account for the strong age-dependence of firms' gender preferences in our data, we re-estimated Tables 10-13 excluding all ads requesting a customer-contact occupation, all ads in administrative occupations, and all ads requesting beauty or a specific marital status. The results, available from the authors, were strikingly similar to Tables 10-13. Indeed, in two of the data sets (Zhaopin and XMZYJS) the estimated age effects were somewhat larger in magnitude and more consistently statistically

[24] Education controls are removed from those regressions in order to illustrate the effect of education on gender preferences in the Figures.

[25] Our definition of customer contact occupations is based simply on the occupation's name as provided by the job boards.

significant.[26] Results for XMRC were unchanged in almost all respects, while the Computrabajo estimates became slightly smaller in magnitude but unchanged in statistical significance.

Despite the occupation, firm, and job-cell fixed effects in Tables 10-13, it remains possible that the age-dependence in firms' gender preferences we estimate is caused by differences in the types of work firms think is suitable for men and women that are present *within* the occupation and job-cell categories available to us. To address this concern, we were able to construct a much more detailed occupation classification in one of our datasets: XMRC. Specifically, we take advantage of the fact that the XMRC site allows a firm to list up to 5 of its 36 occupation categories for each job ad.[27] While few ads use all five labels, the combinations that are used provide a much finer grid of job types than classifications that assign a single, low-dimensional label to each job. Indeed, this interaction yields 1,556 distinct occupation combinations compared to the 36 used in Table 11. Re-estimating Table 11 using these finer occupation controls yields almost identical age coefficients, even in columns (3) and (4) where we interact these finer occupation labels with firm fixed effects (thus allowing each firm to 'gender' its detailed occupations in its own way, without restriction).[28] Even within these detailed job cells, we find that firms 'want women to be young and men to be older', suggesting that those preferences are not purely an artifact of differences in the type of work firms assign to men versus women.

Despite the regressions in Table 14 that control for advertised marital status preferences, it is of course still possible that age-dependent gender preferences are explained by women's movement into marriage and childbearing between the ages of 18 and 40, if firms do not always screen explicitly on these characteristics *when* they care about them. While it is hard to definitively rule out such an explanation, it does raise the issue of why –if marriage and children are firms' primary concern—firms wouldn't screen on these characteristics directly in regulatory environments where this is permitted. More to the point, if caring for young children is employers' main concern with female employees, it's unclear why firms avoid women over 35 as strongly as our data suggest. Based on 2005 Census data, the mean age at first birth among urban Chinese women is 26.5. Since only about 14 percent of these women have more than one child, the probability of currently having a child under the age of five is

[26] Especially interesting here are the column 4 results, which showed an insignificant age effects in Table 10 (for Zhaopin); these coefficients change little in value when we restrict our attention to non-customer—contact occupations, but become statistically significant at the 5 percent level.

[27] All our XMRC estimates reported so far are based on the first occupation listed in the ad, which yields a similar number of categories as our other three data sets. But almost half of the ads in our XMRC data (44.6 percent) list more than one occupation. While one reason for this could be to describe different positions in ads that refer to more than one vacancy, the propensity to list more than one occupation is essentially the same in ads for a single opening (42.7 percent) as it is overall. Thus we interpret multiple occupation listings primarily as a means of providing more detail about the job's duties.

[28] These results are also available from the authors.

maximized (at 44 percent) at age 29, and falls sharply to 25 percent at 32, 9 percent at 35 and 2 percent at 40. In urban Mexico, where fertility rates are much higher, mean age at first birth is 24, but the probability of currently having a child under the age of five also peaks earlier, (at 51 percent) at 24 as well.[29] The share of women with a child under five then declines monotonically to 38 percent at 30, 29% at 35, and 10 percent at 40. Thus if care of young children explains firms' declining attraction to women as they age, one would expect this decline to reverse after age 29 in China, and even earlier in Mexico.[30] Comparing the 30-34 and 35+ categories, that is not the case in any of our data sets.[31] Thus, any child-care based explanation of the age pattern in employers' gender preferences in our data would need to be based on the care of school-aged children (and in the case of China, in almost all cases the care of a *single* school-aged child), and would seem to require greater maternal time and energy investments to care for school-aged children than for younger children.[32]

Explaining the correlation of firms' gender preferences with firms' desired levels of worker experience (as we observe in the *Zhaopin* data) is also not straightforward. While firms might expect men's turnover rates to fall more rapidly with age than women's, it is much harder to explain why firms' relative preference for men rises with *past* experience, which is what the job ads request. Indeed, if firms are worried about women's future labor force attachment, we would expect past experience to be a stronger signal of permanence for women than men.

Another possible interpretation of the advertised preference patterns in all our data sources is that they represent variation not in firms' preferences (as we have been assuming), but in the relative wage costs of hiring different types of workers. For example, if the gender wage gap was higher among young than older workers, or among single than married workers, firms might prefer to hire women when they are young and single, not because they value those traits in women, but because women are cheaper (relative to men) when they are young and unmarried. By the same logic, firms might prefer attractive women in customer-contact occupations, not because customers value female beauty, but because young, attractive women disproportionately enjoy interacting with customers. This would lead young women to

[29] Mexican fertility statistics come from the 2010 Census. We define 'urban' as living in an area with more than 100,000 inhabitants.

[30] Census data on women's employment rates in both countries also shows no evidence of declining labor force attachment between the ages of 30 and 40: In both countries, women's employment rates are roughly unchanged between 30 and 40. In China, women's employment rate is constant at around 70% in that age range, and is independent of a woman's marital status or whether she has had a child.

[31] There is a slight reversal in the unadjusted means for the XMZYJS data, but this disappears in the presence of controls.

[32] Care of aging parents and parents-in-law could also be a demand on women's time in China, though in most cases this would not be a major factor for women in the age ranges (18-40) studied in this paper.

'crowd' into customer-contact occupations, bidding down their relative wages in those jobs and making them more attractive to firms than other worker types.

Overall, we think that 'cost'- or 'supply'- based explanations of the advertising patterns in our data are highly unlikely. One reason is that uncontrolled variation in relative compensation costs should be less of an issue in our regression specifications that control for the level of the posted wage, and in specifications with firm*occupation cell fixed effects. This is because firms should have less latitude to pay different wages to workers based on their age and sex in these situations. Yet in most cases, adding these controls does not affect any of our main estimates. More importantly, the patterns of gender wage gaps needed to explain the main patterns in advertised gender preferences in our data are at variance with available evidence. Specifically, to explain our result that firms' preferences shift toward men as the candidate's ideal age and experience rise, and when candidates are married, then the gender gap would need to decline with age and experience, and marriage would need to reduce men's wages, relative to women's. But evidence from our own data on offered wages, and well known stylized facts from other countries, strongly indicate that the opposite is the case.[33] Similarly, for variation in men's relative wage costs to explain our findings for beauty, we'd need women to earn a smaller beauty premium than men. This is not the case in our data either.

In sum, viewed in the context of the cross-sectional wage patterns in our data, the advertising patterns in our data point to an employer-demand driven scenario where patterns in firms' advertised requests reflect variation in employer demand for worker attributes like age and sex, not variation in the market costs of those attributes. Indeed, taken together, they suggest a scenario in which firms' preferences for (say) older men both appear directly in job ads *and* drive up older men's relative wages in the labor market. In such a scenario, firms' advertised requests for older men actually understate their preferences for those workers, because older men's higher relative compensation costs attenuate firms' attraction to them.

A final, more speculative explanation of the apparently universal age pattern in firms' explicit gender preferences in our data is suggested by the striking degree to which they correspond to the theory and evidence regarding human mate value in evolutionary psychology (see for example Sugiyama 2005 for a recent review). According to that perspective, women's reproductive value is maximized shortly after menarche, then declines rapidly with age and

---

[33] Online Appendix Table O-2 shows that regressing offered wages on desired age and gender and their interaction (for the sample of ads that contains all three variables) reveals a flatter age-wage profile for women than men in all four? of our job board data sets. Similar regressions show that women also have a flatter experience-wage profile in the three data sets that contain an experience variable. Finally, consistent with the marriage-premium literature (e.g. Dougherty 2006), Table O-3 confirms that men experience a marriage-wage premium and women a marriage-wage penalty in the Computrabajo data.

with parity.[34] Furthermore, because women typically lactated for about two thirds of their reproductive years in the small forager societies where humans' gender preferences likely evolved, women's reproductive capacity was an extremely scarce fitness resource, not just from men's perspective but for the entire community.[35] Accordingly, one might expect men's preferences to be highly attuned to cues of future reproductive value, i.e. youth, health and low parity (see Symons, 1995), and human communities in general to place a high value on this scarce resource. Since this pattern corresponds precisely to firms' preferences for female youth, beauty and single marital status in our data, it raises the possibility that employers' gender preferences may also be associated with evolved human preferences that placed a high value on a historically scarce resource: young, healthy, fertile women.[36]

7. More on the Variation in Gender Preferences: The Role of Firms versus Jobs

While the regressions reported in Tables 6-14 estimated the determinants of firms' gender preferences both within and across firms and jobs, so far we haven't explored the relative importance of firm- versus job- level factors in explaining when firms want to hire men versus women. In this section we briefly report on two exercises that illustrate the striking importance of *job*-level factors in firms' preferences in all our data sets. Together, the results have useful implications for the types of models of firms' preferences that can adequately account for the patterns in our data.

a. Gender-targeting at the firm level

Since the vast majority of employers post multiple ads on any given job board, it may be of interest to ask how frequently a typical *firm* engages in gender targeting. To that end, Table IV in KS showed the incidence of gender-targeting by firm (rather than by ad). One key result was that the share of firms that issued a gender-targeted ad at some point during the 20-week observation period is much higher than the share of *ads* that are gender targeted. Specifically, among firms who placed more than 50 Zhaopin job ads, 70.7 percent issued at least one gender-targeted ad during the 20-week observation window. In XMRC, XMZYJS and Computrabajo respectively, Table O-4 in the online appendix shows that these shares were 100, 99.1, and 98.6 percent respectively. While these higher numbers reflect both a greater overall incidence of gender targeting in the new datasets and the longer observation windows, they clearly illustrate that gender-targeting of job ads is commonplace in all these labor markets.

---

[34] According to Sugiyama, a typical forager woman loses one sixth of her reproductive capacity with each birth.

[35] Based on data from the Yanomamö, Symons (1995) estimated that a typical woman was fertile on just 78 of 8,030 days, or less than one percent of the time between menarche and menopause. According to Sugiyama's review (2005) similar results apply to other forager societies.

[36] For well-known reasons, male mate value is not linked as closely to youth as women's from an evolutionary perspective. Instead one would expect women to evaluate men more on the basis of revealed productive capacity and reliability as a provider, which may increase with age (and might be signaled by past work experience).

The second main result in KS's Table IV was that many firms gender-target in both directions, issuing some ads that request men and others that request women. Specifically, among firms who placed more than 50 Zhaopin job ads, 38.7 percent placed at least one ad that favors men *and* one ad that favors women. In XMRC, XMZYJS and Computrabajo respectively, Table O-4 shows that these shares were 96.2, 92.4 and 89.2 percent respectively. This dramatically illustrates that a large share of the variation in advertised gender preferences in all these markets takes place *within firms*: models based on firm-level tastes for men relative to women are likely to be of little help in understanding these patterns.

b. Variance decomposition

Pursuing the above idea further, Table V in KS decomposed the variance of P^F, P^M and P^F+P^M across ads into components that are associated with occupations, with firms, with differences in the way firms gender their occupations, and within job (firm*occupation) variation. Table O-5 of the online appendix reproduces that table, and replicates it for the XMRC, XMZYJS, and Computrabajo datasets respectively. Methodological details of the decomposition are provided in KS. Common features of all four datasets include the following. First, occupation alone explains very little of the variation in advertised gender preferences (less than 10 percent in all cases and usually much less). Second, in all three Chinese datasets, occupation and firm effects together (but not interacted) explain between 30 and 42 percent of the variance in advertised gender preferences (this number is 17-20 percent in Computrabajo). Thus, in all four datasets the majority of the variance in advertised gender preferences is *within firms*: it is not so much that some firms prefer to hire women and others prefer men, but more that most firms prefer men for one subset of their jobs and women for a different subset.

Third, adding occupation*firm interactions accounts for between about 10 and 42 percent of the variance, suggesting that even within data sets *occupations are not consistently gendered across firms*. Finally, in all datasets there is a large degree of within-job-cell variation in advertised gender preferences, ranging from about a third of the total in Zhaopin and XMRC to over 70 percent in Computrabajo.

Together, the results in this section highlight the substantial heterogeneity in employers' gender preferences within jobs inside the same firm. This phenomenon is hard to reconcile with models based on *firm*-level tastes for men versus women. Because gender-typing of occupations is not very consistent across firms, our results are also inconsistent with a scenario where the same broad occupations are universally viewed as 'women's work' and others as men's. In contrast to these heterogeneous occupation-based gender preferences, firms' age-based gender preferences seem to be highly robust both across and within firms, and across and within occupations.

8. Discussion

Our analysis of employers' use of *ex ante* age- and gender-based hiring screens has yielded two main results. First, we confirm the existence of a robust negative skill-targeting relationship in all four data sets available to us: employers are less likely to use age and gender as screening mechanisms as jobs' skill demands rise. As noted, these results are consistent with a simple model of employer hiring from two pools in the presence of application processing costs, exposited in Kuhn and Shen (2013). Second, we identify common patterns in employers' requests for male versus female employees across different demographic and job categories; these patterns are highly robust to estimation methods, to statistical controls, across data sets drawn from different countries, from different regions of the same country, and serving very different skill levels. Specifically, treating job ad contents as descriptions of the 'ideal' employee from the employer's point of view, the ideal male employee is older, experienced and married. The ideal female employee is young, physically attractive, and single. These patterns are evident in all four data sets whenever the relevant employee characteristic is measured, and are evident even when comparing firms' requests for male and female employees within detailed firm*occupation cells.

What explains these apparently universal patterns in men's and women's revealed relative desirability to employers? While some patterns are consistent with existing models – for example the male marriage-premium literature suggests some reasons why firms might prefer married to single men--, others are harder to understand. For example, while firms might be concerned that women may quit or reduce their effort levels due to the presence of young children in the home, this does not explain firms' strong apparent disinterest in hiring women (relative to men) when seeking workers over the age of 35, compared to workers between 25 and 35. And while employers' preferences for female beauty, especially in customer-contact and administrative occupations, account for some of the shift in firms' gender preferences as workers age, this shift in preferences remains very strong in occupations and in jobs that neither request physically attractive workers nor appear to require customer contact. Further, the 'twist' in firms' gender preferences with employee age is present even within firm*occupation cells, so it does not appear to be related solely to differences in the type of work that firms think men and women are suitable for at different ages.

Overall, our results suggest that further investigation of the strong age-dependence in employers' explicitly-announced gender preferences is worthy of further study. Perhaps the flatter age-earnings profiles among women that are observed in almost all industrialized societies are connected not just with gender differences in *workers'* preferences and in optimal human capital investment decisions, but with more deeply seated human preferences that place a sharply declining social value on women, relative to men, as they age.

References

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

Brencic, Vera. "Wage Posting: Evidence from Job Ads" *Canadian Journal of Economics* 45(4) (Nov 2012): 1529-1559.

Delgado Helleseter, Miguel. "English Skills and Wages in a Non English Speaking Country: Findings from Online Advertisements in Mexico" unpublished paper, University of California, Santa Barbara, February 2013.

Dougherty, Christopher. "The Marriage Earnings Premium as a Distributed Fixed Effect" *Journal of Human Resources* 41.2 (Apr 2006): 433-443.

Kuhn, Peter and Kailing Shen (2013) "Gender Discrimination in Job Ads: Evidence from China". *Quarterly Journal of Economics*, (February): 287–336.

Landry, Craig E., Andreas Lange, John A. List, Michael K. Price and Nicholas G. Rupp. "Toward an Understanding of the Economics of Charity: Evidence from a Field Experiment" *The Quarterly Journal of Economics* (2006) 121 (2): 747-782. doi: 10.1162/qjec.2006.121.2.747

Sugiyama, Lawrence S. 2005. "Physical Attractiveness in Adaptationist Perspective", In: Buss DM. (Ed). *Evolutionary Psychology Handbook*. New York: Wiley. Pp. 292-343.

Symons, D. (1995). "Beauty is in the adaptations of the beholder", In P. R. Abramson & S. D. Pinkerson (Eds.), *Sexual nature, sexual culture* (pp. 80–118). Chicago: University of Chicago Press

Figure 1: Share of ads requesting women and men, by desired age

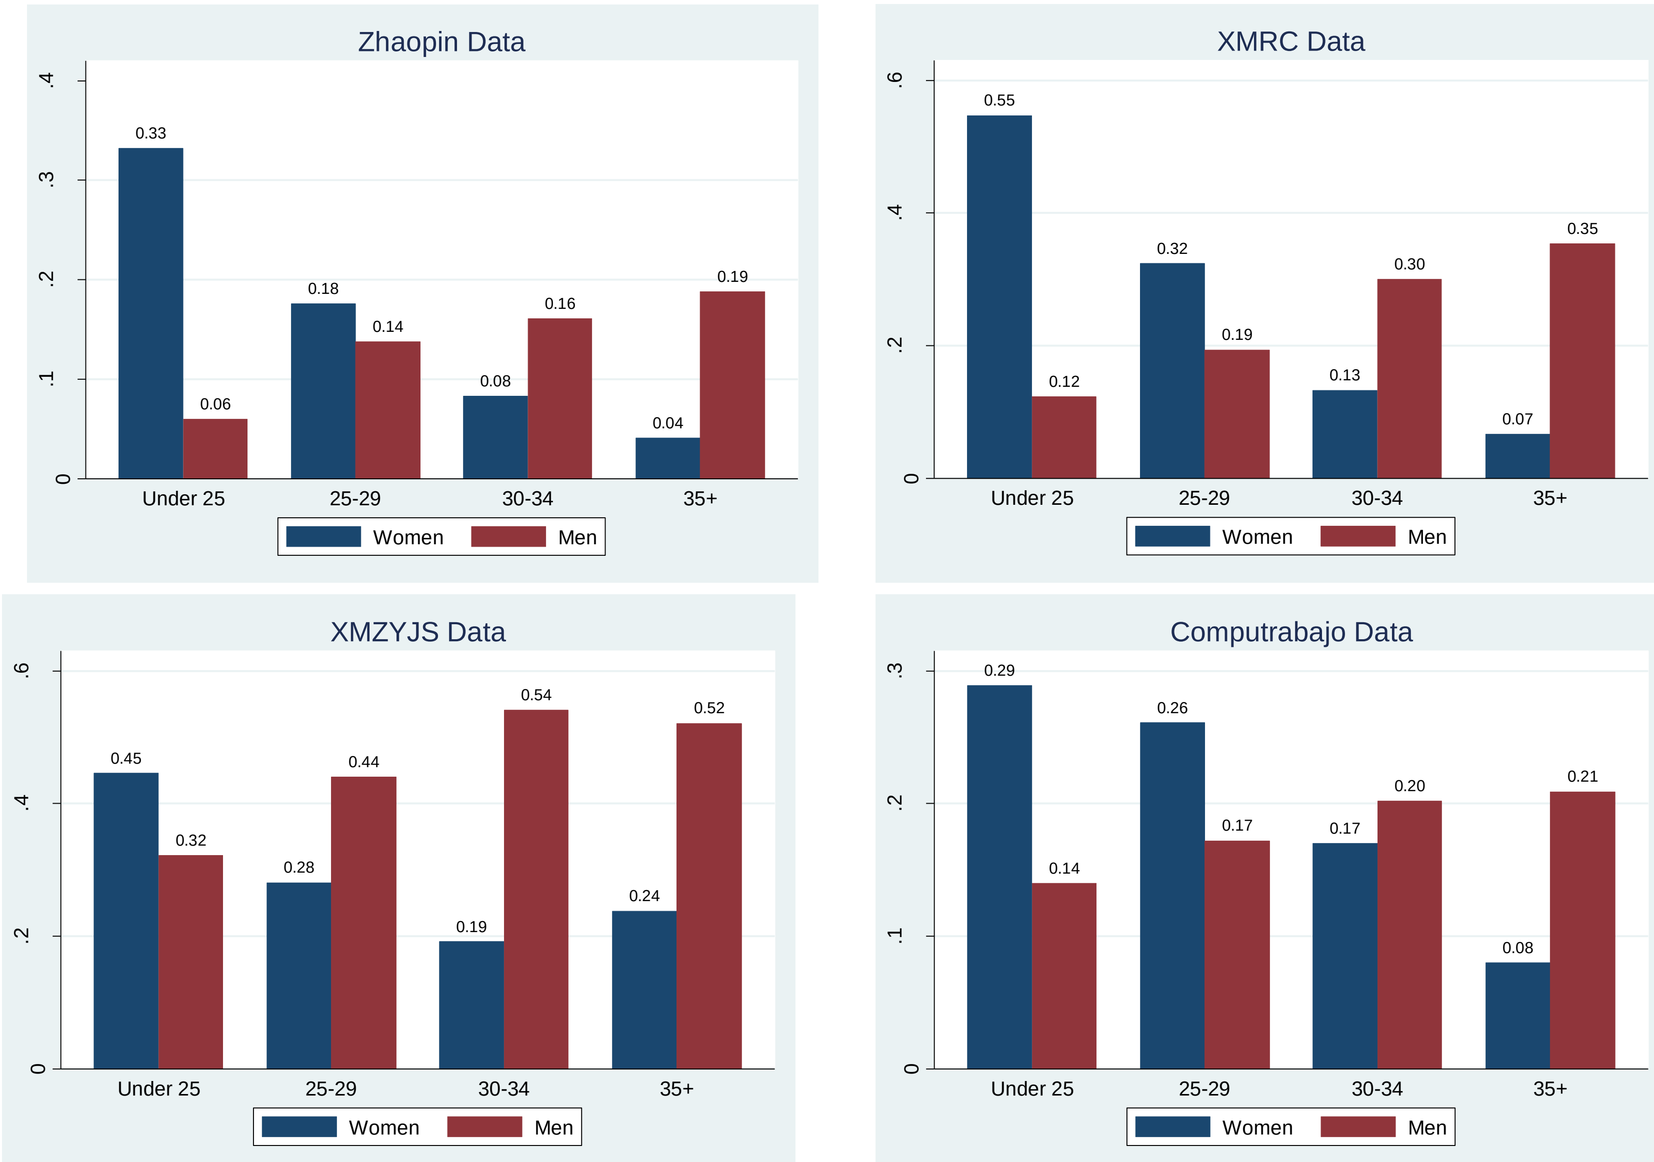

Figure 2: Occupation fixed effects in preferences toward men (P^M - P^F), Zhaopin Data:

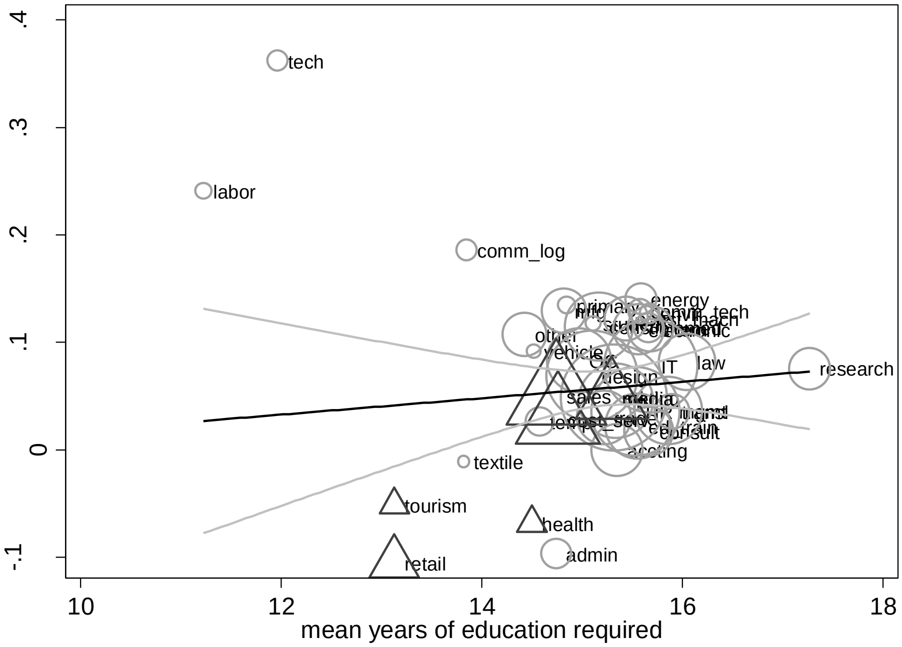

Figure 3: Occupation fixed effects in preferences toward men (P^M - P^F), XMRC Data:

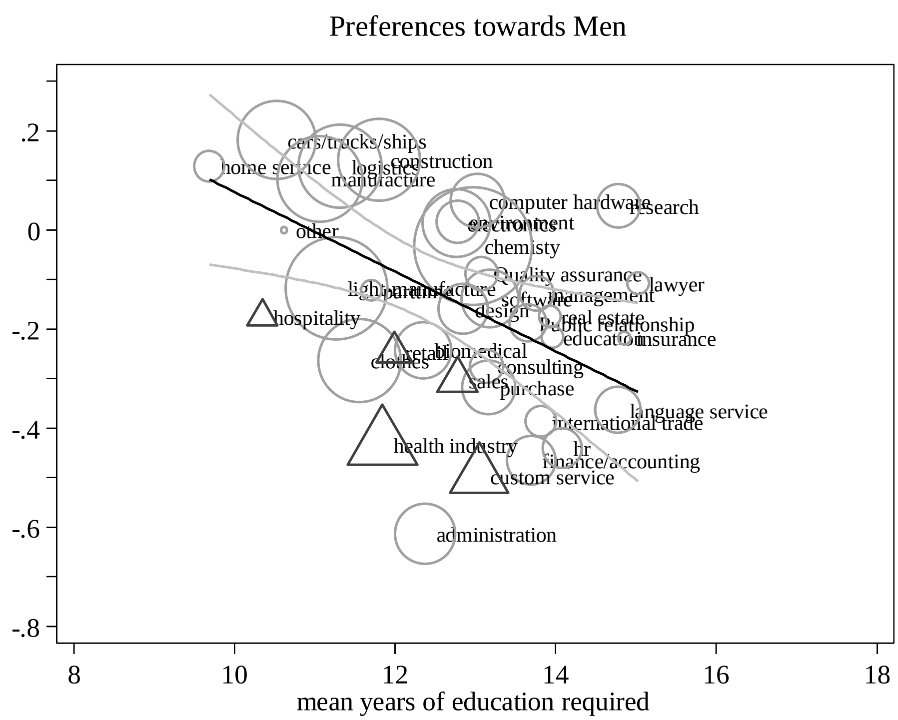

Note: Symbol size in Figures 2-5 is inversely proportional to the standard error of the estimated fixed effect. Customer contact occupations indicated by triangles.

Figure 4: Occupation fixed effects in preferences toward men (P^M - P^F), XMZYJS Data:

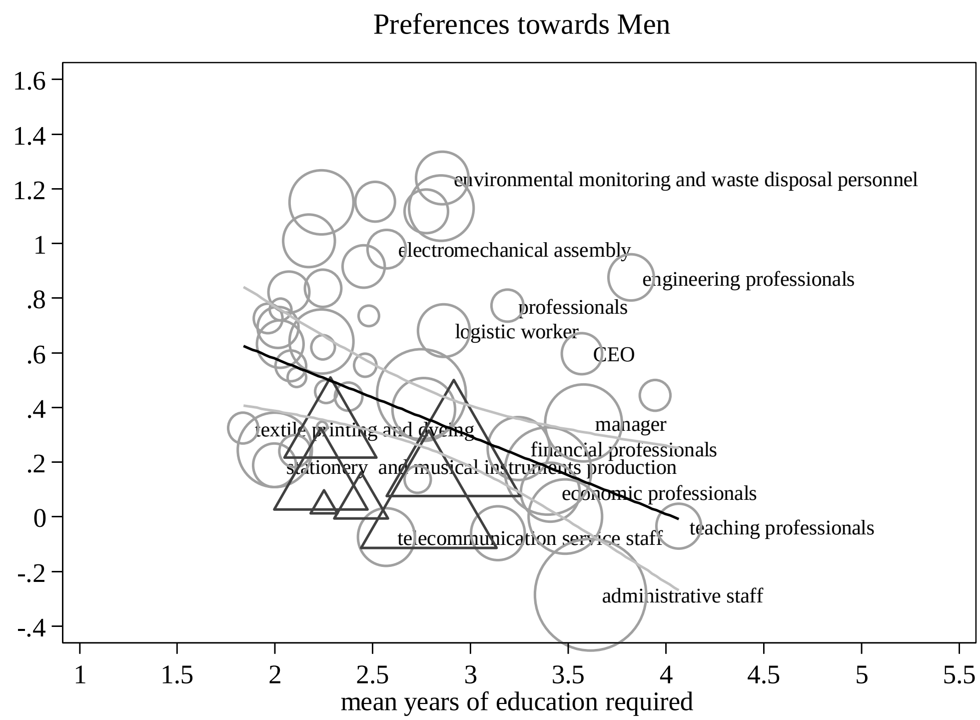

Note: only outliers are labeled. The six customer contact occupations (triangles) are commercial service, procurement staff, restaurant service, hospitality service, medical assistant and service worker.

Figure 5: Occupation fixed effects in preferences toward men (P^M - P^F), Computrabajo Data:

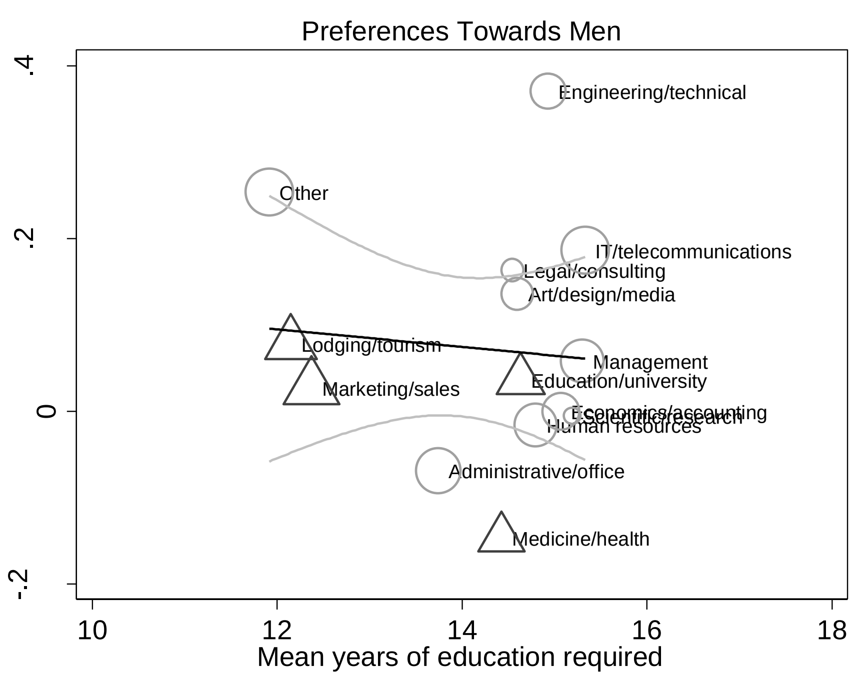

Appendix A: Sample Construction

XMZYJS:

To construct our analysis sample from the population of ads in our frame, we excluded ads with the following characteristics: a minimum requested age below 18 or a maximum requested age over 60, an offered wage above 10000 yuan/ month, missing number of vacancies, jobs located outside the city of Xiamen, required education level is illiterate/other/undergraduate or above, missing community information, jobs that require a *Zhicheng* rank above the assistant level profession rank level, jobs in agriculture, ads with missing education requirement, missing occupation, and ads with missing firm information.

XMRC:

To construct our analysis sample from the population of ads in our frame, we excluded ads with a minimum requested age below 18 or a maximum requested age over 60; ads offering more than 10,000 yuan/month or less than 1000 yuan/month; ads requesting a master's, professional or PhD degree (all of these were rare); ads for more than 10 vacancies (since job descriptions tended to be vague); ads for jobs located outside Xiamen city; ads with missing firm information; ads from firms located outside mainland China; ads with missing occupation information; and ads without a stated education requirement.

Zhaopin:

Sampling procedures for the Zhaopin data are described in KS (2013). They are very similar to the procedures in the XMZYJS and XMRC data.

Computrabajo:

As in all the data sets described in Table 1, our Computrabajo analysis sample restricts attention to ads where the required level of education is explicitly stated. Ads offering a monthly salary below 1,700 Mexican pesos, which is below the minimum wage, or above 100,000 Mexican pesos per month were dropped from the sample.[37]

[37] Many Computrabajo advertisements indicate that the advertised salary either includes or does not include benefit payments, bonuses, commissions, and other extras. For the analysis presented here, ranges were created based on the advertised amounts, ignoring this additional language.

TABLE 1

Sample Means: XMZYJS, XMRC, Zhaopin and Computrabajo Job Ads

	(1)	(2)	(3)	(4)
Ad Characteristics	XMZYJS	XMRC	Zhaopin	Computrabajo
Gender requirements				
No gender preference	0.277	0.616	0.895	0.680
Prefer male?	0.421	0.186	0.055	0.161
Prefer female?	0.303	0.199	0.050	0.159
Education requirements				
Junior middle school or less	0.604	0.493	0.129	0.429
High school or Tech school	0.332			
Some postsecondary	0.064	0.373	0.457	0.151
University		0.134	0.414	0.420
Experience requirements				
none mentioned or under one year	n/a	0.510	0.205	0.809
1-3 years	n/a	0.412	0.399	0.166
4-5 years	n/a	0.061	0.237	0.022
More than 5 years	n/a	0.018	0.158	0.002
Age requirements				
No age restrictions	0.000	0.482	0.757	0.214
Ad specifies a minimum age	1.000	0.505	0.169	0.768
Ad specifies a maximum age	0.771	0.442	0.202	0.743
Mean age, when specified (years)	27.65	28.91	30.59	30.66
Wages				
Wage not specified	0.000	0.584	0.836	0.725
Mean wage, when advertised	1810	2555	4279	7642
Preferred Marital Status				
Single person preferred	n/a	n/a	n/a	0.013
Married person preferred	n/a	n/a	n/a	0.021
Number of positions advertised				
Unspecified	0.000	0.055	0.481	n/a
Mean number, when specified	8.687	1.794	1.692	n/a
Job requires beauty	n/a	n/a	0.077	0.139
Number of ads	141,284	39,746	1,051,706	90,561

Notes: Wages are in RMB/month in XMZYJS, XMRC and Zhaopin; in Mexican pesos/month in Computrabajo.

The average exchange rate is 6.77 RMB per U.S. dollar in the year of 2010, our data period.

The average exchange rate is 13.2 MXN per U.S. dollar in the period of our Computrabajo data.

In all data sets the mean age is the midpoint of the minimum and maximum, conditional on both being specified.

n/a denotes data are not available in that data set

TABLE 2
Share of Job Ads Expressing a Gender or Age Preference, by Ad Characteristics
ZHAOPIN DATA

	Share of job ads			
	(1)	(2)	(3)	(4)
	Requesting women	With no gender preference	Requesting men	Specifying an age range[1]
A. JOB SKILL INDICATORS				
Education Requirements				
High school or less	0.113	0.766	0.120	0.239
Some college	0.059	0.892	0.049	0.143
University	0.021	0.938	0.042	0.077
Experience requirements				
None or less than one year	0.087	0.860	0.053	0.159
1-3 years	0.060	0.889	0.051	0.123
3-5 years	0.025	0.920	0.055	0.112
More than 5 years	0.015	0.917	0.068	0.125
Wages				
Wage not specified	0.046	0.900	0.054	0.122
Wage is specified	0.072	0.870	0.058	0.161
Wage, if specified:				
under 1500	0.167	0.734	0.099	0.178
1500-2999	0.114	0.808	0.078	0.171
3000-3999	0.053	0.899	0.048	0.160
4000-7999	0.034	0.918	0.048	0.155
8000+	0.038	0.929	0.033	0.142
B. ASCRIPTIVE JOB REQUIREMENTS				
Age requirements				
No age restrictions	0.029	0.944	0.027	0.000
Ad specifies a minimum age	0.116	0.746	0.138	0.757
Ad specifies a maximum age	0.124	0.724	0.151	0.635
Maximum and minimum specified	0.131	0.721	0.149	1.000
Mean age, when specified:				
Under 25	0.332	0.607	0.060	1.000
25-29	0.176	0.686	0.138	1.000
30-34	0.083	0.756	0.161	1.000
35+	0.041	0.771	0.188	1.000
Job requires beauty (*xingxiang*)?				
No	0.034	0.909	0.056	0.118
Yes	0.239	0.723	0.038	0.249

TABLE 3
Share of Job Ads Expressing a Gender or Age Preference, by Ad Characteristics
XMRC DATA

	Share of job ads			
	(1)	(2)	(3)	(4)
	Requesting women	With no gender preference	Requesting men	Specifying an age range[1]
A. JOB SKILL INDICATORS				
Education requirements				
High school or less	0.216	0.534	0.250	0.445
Some postsecondary	0.210	0.672	0.118	0.435
University	0.105	0.761	0.134	0.351
Experience requirements				
experience requirement missing	0.203	0.624	0.174	0.329
1 years	0.305	0.536	0.159	0.560
2-3 years	0.170	0.620	0.210	0.525
4-5 years	0.067	0.693	0.241	0.497
More than 5 years	0.037	0.749	0.214	0.516
Wages				
Wage not specified	0.168	0.653	0.180	0.366
Wage is specified	0.242	0.564	0.194	0.517
Wage, if specified				
1000-2000	0.387	0.436	0.177	0.564
2000-3000	0.198	0.587	0.215	0.501
3000-4000	0.096	0.715	0.189	0.466
4000-5000	0.069	0.744	0.187	0.468
more than 5000	0.054	0.777	0.168	0.474
B. ASCRIPTIVE JOB REQUIREMENTS				
Age requirements				
No age restrictions	0.136	0.725	0.139	0.000
Ad specifies a minimum age	0.258	0.514	0.228	0.850
Ad specifies a maximum age	0.270	0.495	0.234	0.969
Maximum and minimum specified	0.272	0.494	0.234	1.000
Mean age, when specified				
Under 25	0.547	0.330	0.123	1.000
25-29	0.324	0.481	0.194	1.000
30-34	0.133	0.567	0.300	1.000
35+	0.067	0.579	0.354	1.000

1. Both a maximum and minimum age are specified.

TABLE 4
Share of Job Ads Expressing a Gender or Age Preference, by Ad Characteristics
XMZYJS DATA

	Share of job ads			
	(1)	(2)	(3)	(4)
	Requesting women	With no gender preference	Requesting men	Specifying an age range[1]
A. JOB SKILL INDICATORS				
Education requirements				
Junior Middle School or less	0.290	0.240	0.470	0.802
High School/Tech School	0.325	0.307	0.368	0.734
College	0.305	0.461	0.234	0.672
Wage				
2000 or less	0.332	0.277	0.392	0.764
2000-3000	0.223	0.258	0.519	0.806
3000-4000	0.206	0.363	0.430	0.719
4000-5000	0.269	0.416	0.316	0.649
more than 5000	0.200	0.574	0.226	0.556
B. ASCRIPTIVE JOB REQUIREMENTS				
Age requirements				
No age restrictions	0.143	0.857	0.000	0.000
Ad specifies a minimum age	0.303	0.276	0.421	0.771
Ad specifies a maximum age	0.317	0.231	0.452	1.000
Maximum and minimum specified	0.317	0.231	0.452	1.000
Mean age, when specified				
Under 25	0.464	0.214	0.322	1.000
25-29	0.299	0.252	0.448	1.000
30-34	0.197	0.227	0.576	1.000
35+	0.220	0.196	0.583	1.000

1. Both a maximum and minimum age are specified.

TABLE 5
Share of Job Ads Expressing a Gender or Age Preference, by Ad Characteristics
COMPUTRABAJO DATA

	Share of job ads			
	(1)	(2)	(3)	(4)
	Requesting women	With no gender preference	Requesting men	Specifying an age range[1]
A. JOB SKILL INDICATORS:				
Education requirements				
High school or less	0.165	0.648	0.187	0.776
Some postsecondary	0.215	0.582	0.203	0.752
University degree	0.133	0.748	0.119	0.662
Experience requirements				
Not stated or less than one year	0.156	0.695	0.148	0.726
1-3 years	0.182	0.643	0.175	0.727
3-5 years	0.117	0.686	0.196	0.663
More than 5 years	0.056	0.696	0.248	0.498
Wages				
Wage not specified	0.152	0.697	0.151	0.710
Wage is specified	0.179	0.635	0.186	0.762
Wage, if specified				
under 4,000	0.154	0.675	0.171	0.772
4,000-5,999	0.208	0.566	0.225	0.786
6,000-7,999	0.212	0.594	0.193	0.786
8,000-9,999	0.177	0.670	0.153	0.778
10,000+	0.113	0.746	0.141	0.677
B. ASCRIPTIVE JOB REQUIREMENTS				
Age requirements				
No age restrictions	0.069	0.859	0.071	0.000
Ad specifies a minimum age	0.185	0.630	0.184	0.943
Ad specifies a maximum age	0.190	0.621	0.189	0.975
Maximum and minimum specified	0.192	0.620	0.189	1.000
Mean age, if max and min specified:				
Under 25	0.289	0.571	0.140	1.000
25-29	0.261	0.567	0.172	1.000
30-34	0.170	0.628	0.202	1.000
35+	0.080	0.711	0.209	1.000
Job requires beauty				
No	0.139	0.694	0.167	0.709
Yes	0.285	0.592	0.122	0.816
Marital status preferences				
Prefer single	0.601	0.259	0.140	0.927
Prefer married	0.084	0.387	0.529	0.900
No marital status preference	0.155	0.692	0.153	0.718

1. Both a maximum and minimum age are specified.

TABLE 6
Effects of Jobs' Skill Demands on the Probability an Ad is Age- or Gender-Targeted
ZHAOPIN DATA

	Prob(Gender Targeted)				Prob(Age Targeted)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Education requirement								
Some postsecondary	-0.0744***	-0.0678***	-0.0599***	-0.0209**	-0.0675***	-0.0435***	-0.0273***	-0.0250*
	(0.0050)	(0.0039)	(0.0049)	(0.0088)	(0.0045)	(0.0033)	(0.0045)	(0.0143)
University	-0.1006***	-0.0912***	-0.0806***	-0.0202*	-0.1152***	-0.0600***	-0.0389***	-0.0494***
	(0.0057)	(0.0048)	(0.0057)	(0.0106)	(0.0050)	(0.0035)	(0.0049)	(0.0136)
Experience requirement								
2-3 years	-0.0156***	-0.0208***	-0.0219***	-0.0255***	-0.0145***	-0.0049***	-0.0071***	-0.0045
	(0.0023)	(0.0019)	(0.0025)	(0.0058)	(0.0032)	(0.0018)	(0.0021)	(0.0068)
4-5 years	-0.0323***	-0.0338***	-0.0324***	-0.0348***	-0.0038	0.0030	0.0022	-0.0068
	(0.0027)	(0.0026)	(0.0032)	(0.0069)	(0.0038)	(0.0022)	(0.0027)	(0.0097)
More than 5 years	-0.0285***	-0.0329***	-0.0343***	-0.0250**	0.0154***	0.0218***	0.0153***	0.0094
	(0.0035)	(0.0030)	(0.0038)	(0.0100)	(0.0048)	(0.0027)	(0.0034)	(0.0157)
Log (offered wage)				-0.0403***				0.0083
				(0.0058)				(0.0069)
Fixed effects	occ, ind, province	occ, province, firm	occ*firm, province	occ*firm, province	occ, ind, province	occ, province, firm	occ*firm, province	occ*firm, province
number of groups	116	73,712	258,782	63,364	116	73,712	258,782	63,364
N	1,051,706	1,051,706	1,051,706	172,887	1,051,706	1,051,706	1,051,706	172,887
R^2	0.078	0.331	0.562	0.623	0.058	0.385	0.562	0.617

*** p<0.01, ** p<0.05, * p<0.1. Standard errors (in parentheses) are clustered by occupation.

Regressions without firm fixed effects include controls for log firm size and firm ownership type.

All regressions control for number of positions advertised.

TABLE 7
Effects of Jobs' Skill Demands on the Probability an Ad is Age- or Gender-Targeted
XMRC DATA

	Prob(Gender Targeted)				Prob(Age Targeted)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Education requirement								
Some postsecondary	-0.1168***	-0.1037***	-0.0761***	-0.0709***	-0.0303***	0.0008	-0.0071	0.0297
	(0.0097)	(0.0075)	(0.0121)	(0.0226)	(0.0083)	(0.0073)	(0.0131)	(0.0314)
University	-0.1822***	-0.1453***	-0.1232***	-0.1342**	-0.1069***	-0.0145	-0.0223	0.0194
	(0.0197)	(0.0145)	(0.0189)	(0.0627)	(0.0184)	(0.0150)	(0.0230)	(0.0675)
Experience requirement								
2-3 years	-0.0680***	-0.0658***	-0.0677***	-0.0259	-0.0208**	-0.0185**	-0.0094	-0.0140
	(0.0122)	(0.0124)	(0.0231)	(0.0336)	(0.0082)	(0.0073)	(0.0143)	(0.0207)
4-5 years	-0.1188***	-0.1013***	-0.0896*	0.0312	-0.0322**	-0.0309*	-0.0316	-0.0208
	(0.0278)	(0.0308)	(0.0504)	(0.0842)	(0.0145)	(0.0164)	(0.0257)	(0.0706)
More than 5 years	-0.1610***	-0.1349***	-0.0925	0.0635	-0.0041	0.0073	0.0226	0.0675
	(0.0322)	(0.0418)	(0.0705)	(0.1617)	(0.0205)	(0.0213)	(0.0394)	(0.1096)
Log (offered wage)				-0.1646***				-0.0793**
				(0.0520)				(0.0367)
Fixed effects	occupation	occupation, firm	occ*firm	occ*firm	occupation	occupation, firm	occ*firm	occ*firm
number of groups	36	6,716	20,625	10,688	36	6,716	20,625	10,688
N	39,746	39,746	39,746	16,550	39,746	39,746	39,746	16,550
R^2	0.102	0.23	0.301	0.337	0.075	0.362	0.425	0.455

*** p<0.01, ** p<0.05, * p<0.1. Standard errors (in parentheses) are clustered by occupation.

Regressions without firm fixed effects include controls for log firm size and firm ownership type.

All regressions control for number of positions advertised.

TABLE 8
Effects of Jobs' Skill Demands on the Probability an Ad is Age- or Gender-Targeted
XMZYJS DATA

	Prob(Gender Targeted)				**Prob(Age Targeted)**			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Education requirement								
High School/Tech School	-0.0595***	-0.0143	-0.0030	-0.0120	-0.0605***	-0.0260***	-0.0105	-0.0174***
	(0.0164)	(0.0149)	(0.0225)	(0.0128)	(0.0096)	(0.0059)	(0.0085)	(0.0040)
College	-0.1897***	-0.1177***	-0.1016***	-0.1255***	-0.1081***	-0.0627***	-0.0659***	-0.0844***
	(0.0229)	(0.0187)	(0.0260)	(0.0241)	(0.0256)	(0.0126)	(0.0166)	(0.0228)
Log (offered wage)				0.0957				0.0741
				(0.1154)				(0.0836)
Fixed effects	occupation	occ, firm	occ*firm	occ*firm	occupation	occ, firm	occ*firm	occ*firm
number of groups	58	8,916	27,669	27,669	58	8,916	27,669	27,669
N	141,284	141,284	141,284	141,284	141,284	141,284	141,284	141,284
R^2	0.063	0.387	0.520	0.522	0.040	0.558	0.639	0.640

*** p<0.01, ** p<0.05, * p<0.1. Standard errors (in parentheses) are clustered by occupation.

Regressions without firm fixed effects include controls for firm ownership type
and for work location within the Xiamen metropolitan area.

All regressions control for number of positions advertised.

TABLE 9
Effects of Jobs' Skill Demands on the Probability an Ad is Age- or Gender-Targeted
COMPUTRABAJO DATA

	Prob(Gender Targeted)				Prob(Age Targeted)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Education requirement								
Some postsecondary	0.0502***	0.00909	0.000211	-0.0155	0.00402	0.00645	0.00287	-0.00145
	(0.0117)	(0.00788)	(0.00800)	(0.0152)	(0.00915)	(0.00642)	(0.00608)	(0.0117)
University	-0.0884***	-0.113***	-0.113***	-0.0875***	-0.0531***	-0.0145**	-0.0148**	0.0286**
	(0.0181)	(0.0124)	(0.0133)	(0.0177)	(0.0126)	(0.00631)	(0.00668)	(0.0130)
Experience requirement								
1-3 years	0.0433***	0.0305***	0.0316***	0.0334***	0.0318***	0.0192***	0.0176***	0.0213***
	(0.00641)	(0.00473)	(0.00505)	(0.00936)	(0.00814)	(0.00502)	(0.00490)	(0.00813)
More than 3 years	0.0266*	0.00328	0.00882	-0.0269	-0.0235	-0.0286***	-0.0339***	-0.0831***
	(0.0143)	(0.0118)	(0.0135)	(0.0288)	(0.0149)	(0.0107)	(0.0124)	(0.0257)
Log (offered wage)				-0.0255**				-0.0505***
				(0.0105)				(0.0124)
Fixed effects	Occ*State	Occ*State, Firm	Occ*Firm, State	Occ*Firm, State	Occ*State	Occ*State, Firm	Occ*Firm, State	Occ*Firm, State
Number of groups	441	2,384	6,978	3,053	441	2,384	6,978	3,053
N	90,561	90,561	90,561	24,896	90,561	90,561	90,561	24,896
R^2	0.072	0.219	0.303	0.348	0.077	0.343	0.407	0.401

*** p<0.01, ** p<0.05, * p<0.1. Standard errors (in parentheses) are clustered by occupation*state.

Regressions without firm fixed effects include controls for log firm size and firm ownership type.

All regressions control for number of positions advertised.

TABLE 10
Effects of Age and Skill Requirements on the Direction of Firms' Gender Preferences (P^M-P^F)
ZHAOPIN DATA

	(1)	(2)	(3)	(4)
Age Requested				
25-29	0.1529***	0.1612***	0.1074***	0.0240
	(0.0227)	(0.0220)	(0.0378)	(0.1021)
30-34	0.2294***	0.2528***	0.1838***	0.1064
	(0.0253)	(0.0276)	(0.0444)	(0.1284)
35+	0.2635***	0.2942***	0.2508***	0.1461
	(0.0274)	(0.0258)	(0.0404)	(0.1162)
Education requirement				
Some postsecondary	-0.0313**	-0.0970***	-0.0920***	-0.1094***
	(0.0139)	(0.0164)	(0.0223)	(0.0380)
University	-0.0098	-0.0881***	-0.0699***	-0.0590
	(0.0173)	(0.0185)	(0.0262)	(0.0475)
Experience requirement				
2-3 years	0.0179**	0.0019	-0.0020	0.0229
	(0.0075)	(0.0078)	(0.0105)	(0.0232)
4-5 years	0.0715***	0.0660***	0.0609***	0.0808**
	(0.0081)	(0.0099)	(0.0157)	(0.0368)
More than 5 years	0.1091***	0.0980***	0.0732***	0.1645***
	(0.0106)	(0.0118)	(0.0163)	(0.0564)
Log (offered wage)				-0.0541
				(0.0462)
Fixed effects	occ, ind, province	occ, province, firm	occ*firm, province	occ*firm, province
Number of groups	116	24,492	50,232	11,828
N	134,768	134,768	134,768	27,909
R^2	0.189	0.455	0.702	0.734

Regressions without firm fixed effects include controls for log firm size and firm ownership type.
Standard errors (in parentheses) are clustered at the occupation level. *** p<0.01, ** p<0.05, * p<0.1.
All regressions control for number of positions advertised.

TABLE 11

Effects of Age and Skill Requirements on the Direction of Firms' Gender Preferences (P^M-P^F)

XMRC DATA

	(1)	(2)	(3)	(4)
Age Requested				
25-29	0.2221***	0.2729***	0.2385***	0.2299**
	(0.0240)	(0.0308)	(0.0620)	(0.0931)
30-34	0.4602***	0.5244***	0.4904***	0.4560***
	(0.0427)	(0.0521)	(0.0952)	(0.1286)
35+	0.5518***	0.6401***	0.6334***	0.5865***
	(0.0465)	(0.0518)	(0.1016)	(0.1482)
Education requirement				
Some postsecondary	-0.1019**	-0.1304***	-0.1295**	-0.1478
	(0.0424)	(0.0400)	(0.0606)	(0.0904)
University	-0.0087	-0.1035*	-0.1502	-0.1564
	(0.0599)	(0.0555)	(0.1050)	(0.2273)
Experience requirement				
2-3 years	0.0765***	0.0957***	0.1117***	0.0752
	(0.0180)	(0.0202)	(0.0302)	(0.0742)
4-5 years	0.1420***	0.1436***	0.1653***	0.0933
	(0.0240)	(0.0363)	(0.0462)	(0.1458)
More than 5 years	0.0889***	0.1005***	0.1338*	0.1318
	(0.0260)	(0.0351)	(0.0698)	(0.1925)
Log (offered wage)				0.1570*
				(0.0878)
Fixed effects	occupation	occ, firm	occ*firm	occ*firm
number of groups	36	4,562	10,570	6,062
N	17,040	17,040	17,040	8,559
R^2	0.264	0.336	0.387	0.410

Regressions without firm fixed effects include controls for log firm size and firm ownership type.

Standard errors (in parentheses) are clustered at the occupation level. *** p<0.01, ** p<0.05, * p<0.1.

All regressions control for number of positions advertised.

TABLE 12

Effects of Age and Skill Requirements on the Direction of Firms' Gender Preferences (P^M-P^F)

XMZYJS DATA

	(1)	(2)	(3)	(4)
Age Requested				
25-29	0.1117***	0.1338***	0.1033***	0.0882***
	(0.0203)	(0.0178)	(0.0267)	(0.0253)
30-34	0.2306***	0.2241***	0.1793***	0.1548***
	(0.0336)	(0.0356)	(0.0407)	(0.0387)
35+	0.2382***	0.2417***	0.2183**	0.1958**
	(0.0616)	(0.0741)	(0.0931)	(0.0884)
Education requirement				
High School/Tech School	0.0379	-0.0064	0.0084	-0.0086
	(0.0265)	(0.0270)	(0.0337)	(0.0328)
College	-0.0094	-0.0620	-0.0310	-0.0754
	(0.0480)	(0.0466)	(0.0693)	(0.0691)
Log (offered wage)				0.2171***
				(0.0472)
Fixed effects	occupation	occ, firm	occ*firm	occ*firm
number of groups	57	7,609	22,212	22,212
N	108,971	108,971	108,971	108,971
R^2	0.211	0.288	0.335	0.338

Regressions without firm fixed effects include controls for firm ownership type.
and for work location within the Xiamen metropolitan area

Standard errors (in parentheses) are clustered at the occupation level. *** p<0.01, ** p<0.05, * p<0.1.

All regressions control for number of positions advertised.

TABLE 13

Effects of Age and Skill Requirements on the Direction of Firms' Gender Preferences (P^M-P^F)

COMPUTRABAJO DATA

	(1)	(2)	(3)	(4)
Age requested				
25-29	0.0465***	0.0579***	0.0667***	0.0899***
	(0.0136)	(0.0148)	(0.0162)	(0.0341)
30-34	0.137***	0.143***	0.150***	0.163***
	(0.0148)	(0.0172)	(0.0176)	(0.0321)
35+	0.237***	0.262***	0.260***	0.294***
	(0.0158)	(0.0183)	(0.0185)	(0.0332)
Education requirement				
Some postsecondary	-0.0513**	-0.0453**	-0.0471**	-0.0585***
	(0.0241)	(0.0223)	(0.0238)	(0.0195)
University	-0.0621***	-0.0567**	-0.0576**	-0.114***
	(0.0195)	(0.0222)	(0.0240)	(0.0355)
Experience requirement				
1-3 years	0.0226**	0.0101	0.0110	0.0287*
	(0.0101)	(0.00923)	(0.00997)	(0.0163)
More than 3 years	0.125***	0.0956***	0.103***	0.127***
	(0.0211)	(0.0209)	(0.0213)	(0.0415)
Log (offered wage)				0.0164
				(0.0259)
Fixed effects	Occ*State	Occ*State, Firm	Occ*Firm, State	Occ*Firm, State
Number of groups	425	2,070	5,779	2,572
N	65,590	65,590	65,590	18,963
R^2	0.128	0.208	0.303	0.361

Standard errors (in parentheses) are clustered at the occupation*state level.

*** p<0.01, ** p<0.05, * p<0.1.

TABLE 14
Effects of Age and Skill Requirements on the Direction of Firms' Gender Preferences (P^M-P^F)
with Beauty and Marital Status Controls

	(1)	(2)	(3)	(4)
A. ZHAOPIN DATA				
Age requested				
25-29	0.1367***	0.1425***	0.0932**	0.0051
	(0.0204)	(0.0209)	(0.0384)	(0.1087)
30-34	0.2003***	0.2188***	0.1563***	0.0657
	(0.0224)	(0.0265)	(0.0467)	(0.1440)
35+	0.2277***	0.2509***	0.2128***	0.0807
	(0.0237)	(0.0239)	(0.0417)	(0.1342)
Ad requests beauty?	-0.1956***	-0.1993***	-0.1819***	-0.2500***
	(0.0126)	(0.0123)	(0.0251)	(0.0891)
Education requirement				
Some postsecondary	-0.0342***	-0.0986***	-0.0940***	-0.0943***
	(0.0131)	(0.0156)	(0.0216)	(0.0322)
University	-0.0129	-0.0903***	-0.0740***	-0.0474
	(0.0166)	(0.0178)	(0.0258)	(0.0480)
Experience requirement				
2-3 years	0.0167**	0.0010	-0.0035	0.0230
	(0.0072)	(0.0078)	(0.0104)	(0.0228)
4-5 years	0.0666***	0.0619***	0.0573***	0.0810**
	(0.0080)	(0.0099)	(0.0158)	(0.0347)
More than 5 years	0.1027***	0.0925***	0.0701***	0.1704***
	(0.0106)	(0.0116)	(0.0163)	(0.0563)
Log (offered wage)				-0.0537
				(0.0447)
Fixed Effects	occ, ind, province	occ, province, firm	occ*firm, province	occ*firm, province

TABLE 14, Continued
Effects of Age and Skill Requirements on the Direction of Firms' Gender Preferences (P^M-P^F)
with Beauty and Marital Status Controls

	(1)	(2)	(3)	(4)
B. COMPUTRABAJO DATA				
Age requested				
25-29	0.0274*	0.0399**	0.0523***	0.0780**
	(0.0144)	(0.0157)	(0.0173)	(0.0323)
30-34	0.101***	0.109***	0.120***	0.136***
	(0.0144)	(0.0168)	(0.0175)	(0.0300)
35+	0.188***	0.215***	0.219***	0.257***
	(0.0149)	(0.0175)	(0.0181)	(0.0328)
Ad requests beauty?	-0.164***	-0.172***	-0.174***	-0.180***
	(0.0149)	(0.0173)	(0.0178)	(0.0257)
Ad requests married?	0.345***	0.316***	0.280***	0.248***
	(0.0189)	(0.0198)	(0.0217)	(0.0538)
Ad requests single?	-0.343***	-0.366***	-0.329***	-0.267***
	(0.0246)	(0.0252)	(0.0262)	(0.0440)
Education requirement				
Some postsecondary	-0.0457*	-0.0398*	-0.0420*	-0.0533***
	(0.0233)	(0.0213)	(0.0228)	(0.0205)
University	-0.0561***	-0.0505**	-0.0515**	-0.113***
	(0.0190)	(0.0218)	(0.0236)	(0.0341)
Experience requirement				
1-3 years	0.0230**	0.0119	0.0121	0.0273*
	(0.0103)	(0.00913)	(0.0101)	(0.0163)
More than 3 years	0.113***	0.0850***	0.0946***	0.129***
	(0.0209)	(0.0209)	(0.0214)	(0.0431)
Log (offered wage)				0.0228
				(0.0241)
Fixed Effects	Occ*State	Occ*State, Firm	Occ*Firm, State	Occ*Firm, State

*** p<0.01, ** p<0.05, * p<0.1. Standard errors (in parentheses) are clustered at the occupation*state level.
Specifications are identical to Tables 10 and 13, with beauty and marital status controls added.