Age and Gender Profiling in the Chinese and Mexican Labor Markets: Evidence from Four Job Boards

Miguel Delgado Helleseter[a]

Peter Kuhn[b]

Kailing Shen[c]

October 29, 2015

JEL codes: J16, J63, J71

Keywords: gender, discrimination, age, China, Mexico, Internet, beauty, search, screening

When permitted by law, employers sometimes state the preferred age and sex of their employees in job ads. We extend Kuhn and Shen's (2013) study of this practice to three new samples of job ads from China and Mexico, focusing on the interaction between firms' age and gender preferences. In all four data sets, we find that both age and gender profiling become less common as job skill requirements rise (the "negative skill-targeting relationship"). Also, firms' explicit gender requests shift dramatically away from women and towards men when firms are seeking older (as opposed to younger) workers (the "age twist" in gender preferences). Some of this twist can be attributed to employers' age-dependent demands for female beauty and male leadership. Employers also appear to exhibit a marked and widespread lack of explicit interest in women who (based on their age) have likely become mothers, regardless of their children's ages.

[a] California State University, Channel Islands, Miguel.delgadohelleseter@csuci.edu ; [b] Corresponding author: University of California, Santa Barbara, NBER and IZA, peter.kuhn@ucsb.edu; [c]Australian National University and IZA, kailing.shen@anu.edu.au.

This research is supported by National Natural Science Foundation of China grant 71203188, "Impacts of *Hukou*, Education and Wage on Job Search and Match: Evidence Based on Online Job Board Microdata". We thank Anne Pisor and Corinne Low for interesting discussions on related literature, and seminar/conference participants UCSB, UC Berkeley, Princeton, Columbia, Arizona, Yale, IZA, the Bergen/Stavanger workshop on Labor Markets, Families and Children, and the Trans-Pacific Labor Seminar (TPLS) in Sydney for helpful comments. All authors contributed equally to the paper.

"I have heard the wish expressed that one could be a girl and a beautiful girl from the age of thirteen to the age of twenty-two and after that to be a man." Jean de la Bruyère, Les Caractères 1688

1. Introduction

While it is commonplace for job ads to specify desired levels of education, experience and other worker qualifications, explicit employer requests for workers of a particular age, race or gender are prohibited in many developed countries. This can make it hard for researchers to infer how firms value those attributes, leading researchers to adopt a variety of indirect methods to infer firms' preferences. In most of the world's labor markets, however, it is commonplace for employers to include explicit sex and age preferences in job ads. In consequence, these job ads provide a unique laboratory from which to learn how firms value a number of employee characteristics.

Pursuing this idea, Kuhn and Shen (2013; henceforth KS) studied employers' gender-targeting policies in a large sample of job ads on a Chinese Internet job board. Their study highlighted three main findings. First, they found that employers' advertised gender preferences were *symmetric*, in the sense that a roughly equal number of ads requested men and women, both overall and conditioning on skill requirements. Second, advertised gender preferences were highly *job-specific*: neither firm- nor occupation fixed effects explain much of the variance in stated preferences, and a large share of firms request men for some jobs, and women for others, even within the same occupation. Third, KS identified an empirical regularity called the *negative skill-targeting relationship*: firms' propensities to gender-target their ads (in either direction) were strongly and negatively related to the job's skill level.

While the above patterns can be explained by a number of factors, KS presented evidence suggesting that a simple screening cost model accounts for the observed patterns well. In this model, employers choose whether to search narrowly –inviting only a 'preferred' demographic group to apply—or to search broadly, but lean towards broad searches when it becomes more important to identify the single best candidate for the job. While the model makes clear that employers' use of demographically-targeted job ads could be motivated either by employer tastes or by perceived or actual productivity differences between groups, it also clarifies that demographically-restricted searches constitute *prejudice* in the original sense of the word: employers have decided to judge a potential applicant on the basis of an ascriptive demographic attribute without examining that applicant's credentials.

This paper has two main goals. The first is to see whether KS's main findings extend to different data sets and labor markets, and whether they apply to both age- *and* gender-targeting of ads in those markets. We find that they do, using three new datasets, two from China and one from Mexico. Importantly, these new data sets serve workers with a much more representative level of skills than KS's data, which came from a national job board serving a highly skilled clientele. In fact, the two new Chinese datasets represent the entire private-sector labor market of a medium-sized Chinese city. The Mexican data shows that KS's main results extend to a nation with different economic conditions, labor market institutions and culture.

Second, we establish a new empirical regularity concerning the interaction between firms' age and gender preferences in all our data sets.[1] This relationship, which we call the *age twist* in employers' gender preferences, is a strong shift in the *direction* of employers' advertised gender preferences, away from women and towards men, as a worker's age rises from 18 to 45. We explore a number of explanations for this pattern, showing that some of it can be attributed to age-dependent demands for female beauty, for male leadership, and to gender differences in firms' marital status preferences. We also suggest that this dramatic shift in firms' gender preferences with age may play a hitherto-unexplored role in accounting for a well-known stylized fact: that women's age-earnings profiles are flatter than men's. The age-related twist also sheds interesting new light on KS's 'symmetry' result: While this symmetry is surprisingly robust to controls for jobs' skill demands, it breaks down within (desired) age groups. The overall rough balance in firms' requests for women and men in all our datasets is actually the net result of a strong preference for young women and a strong preference for older men.

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. Section 5 considers alternative explanations for the negative skill-targeting relationship, while Section 6 explores possible explanations of the age twist in advertised gender preferences. Section 7 concludes.

2. Data

Two of the three Chinese data sources used in this paper are job boards serving the city of Xiamen, a southern coastal city about the size of Los Angeles.[2] 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. One of these job boards, XMZYJS (the Xia-Zhang-Quan 3 city public job board, www.xmzyjs.com), is operated directly by government employees of the local labor bureau. Like state-operated Job Centers in the U.S., XMZYJS has a history as a brick-and-mortar employment service. XMZYJS's mandate is to serve the less-skilled portion of the area's labor market, and operates purely as 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

[1] KS noted this age twist in their data, but did not study it in any detail. In part, this was due to a concern that the highly-skilled Zhaopin data might not be representative of other Chinese employers.

[2] Based on the 2010 Chinese and U.S. censuses, Xiamen's metropolitan area has 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 are 5.0, 5.6 and 9.6 million respectively.

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. Our other Xiamen-based job board, XMRC (http://www.xmrc.com.cn), is a for-profit, privately-operated company that is sponsored by the local government.[3] 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.

Together, the XMRC and XMZYJS sites provide a fairly complete picture of the active private-sector 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, XMZYJS provided all the ads for jobs that appeared in calendar 2010; details on how we constructed an analysis sample for this and all our datasets is available on the authors' websites.[4] The most important restrictions affecting all datasets involve dropping observations with missing education requirements and missing occupation information, and all jobs requesting workers over the age of 45.[5] XMRC provided us with all the ads for jobs in Xiamen that received their first application between May 1 and October 30, 2010; our XMRC sampling criteria are almost identical to those in XMZYJS.

To facilitate comparison with KS's data, we also present results from *Zhaopin.com*, with sample definitions and regression specifications that are as identical as possible to our XMZYJS and XMRC analyses. Zhaopin.com is the third-largest Internet job board in China; it operates nationally and serves workers that, on average, are considerably more skilled than even those on XMRC. As in KS, this sample is based on all unique ads posted in four five-week observation periods in 2008-2010. In contrast to XMRC and XMZYJS where the data were supplied by the job boards, our Zhaopin data were collected by a web crawler. Details are provided in KS.

[3] Like all our data sets, XMRC serves private sector employers almost exclusively. Recruiting for public-sector jobs, and most recruiting for State-Owned-Enterprises (SOEs) takes place via a different process.

[4] See the paper's website: http://www.econ.ucsb.edu/~pjkuhn/Data/Age&Gender/Age&GenderIndex.htm for all supplementary materials. Like all our datasets, our XMZYJS data is a sample of vacancies that were unfilled at some point during a sampling period. Since XMZYJS's sampling period is an entire calendar year and most vacancies last only a few weeks, this leads to only a minor overrepresentation of long-duration vacancies in the data. Still, to assess the likely effects of duration-bias in sampling, we replicated our main analysis in the Zhaopin data –which has the shortest sampling windows--for an 'inflow' sample of ads that was first posted well after the start of our observation window. There was very little change in the results.

[5] Ads specifically requesting workers over 45 are extremely rare in all our data sets; thus our data are not very informative about firms' hiring preferences in that age range.

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 141,188, 39,727, and 1,051,038 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.[6] 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 = 42 percent in XMRC and 16 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 the XMZYJS site.

In contrast to Zhaopin and XMRC, all XMZYJS ads specify the number of positions that are open, and the mean number of positions (8.7) is much 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 30.5 years. In addition to our restriction to ads requesting workers aged 45 and under, 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.

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. Similarly, 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). Taken together, these patterns in age- and gender-targeting across data sets strongly confirm the negative skill-targeting relationship predicted by KS's model. Related, statistics from XMZYJS (which, among our datasets, is 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. Interestingly, the share of ads favoring men versus women is roughly equal on XMRC, matching

[6] 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.

the symmetry found in Zhaopin. That said, this symmetry breaks down somewhat at the lowest skill levels, with more ads favoring men than women on XMZYJS (42 versus 30 percent of the total).

Our Mexican data is a sample of job ads posted on Computrabajo.com.mx.[7] 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 skill-targeting relationship and the age twist in firms' gender preferences transcend national and cultural boundaries.

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. We use the universe of unique advertisements posted during this period.[8] Additional sample restrictions are similar to those used in our Chinese datasets, and are described on the paper's website. As reported in column 4 of Table 1, our Computrabajo analysis sample contains 90,487 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).[9] Reflecting the overall symmetry in the Zhaopin and XMRC data, Computrabajo's gender-targeted ads are also equally split between men and women.

Of our four datasets, Computrabajo and Zhaopin are the only ones where we have information on employers' explicit requests for physically attractive workers: at 7.7 versus 3.9 percent of ads, this number is considerably higher in Zhaopin than in Computrabajo. These 'beauty' indicators were derived by searching the text of ads for phrases indicating a preference for a visually attractive candidate, then compiling a list of the most common forms of such requests.[10] In the Computrabajo data we also have access to a second indicator of firms' interests in the candidate's attractiveness: 10.9 percent of Computrabajo ads explicitly instructed applicants to include a photo with their application.[11] Another unique feature of the

[7] Computrabajo has job boards in 20 Spanish-speaking countries. We picked Mexico because it had the most ads.
[8] Identical ads appearing on multiple dates are included only once. See Delgado Helleseter (2013) for more details.
[9] The average advertised salary in Computrabajo is 7,640 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.
[10] Appendix Table A1 shows the ten most frequent forms of beauty requests in the Zhaopin and Computrabajo data. In Mexico, requests for a good or excellent "presentación" dominated; In China, requests for good image and temperament (形象气质佳) or regular facial features (五官端正) were the most common.
[11] We detected a photo request in only 1.2 percent of Zhaopin ads, and these are not strongly correlated with age or beauty requests, or with other features of the data. We suspect that this is because a large majority of resumes

Computrabajo data is information on whether the job requires some sort of leadership duties; specifically 1.6 percent of ads say that the job involves supervising others. Finally, 3.4 percent of Computrabajo ads explicitly indicate the firm's preferred marital status for the worker. Of these, 2.1 percent request married applicants and 1.3 percent request single persons. Despite being relatively uncommon compared to age- and gender-based screens, we shall see that supervision demands and marital status screens interact in a strong and informative way with age and gender screens in Mexico.[12]

3. Results: The Negative Skill-Targeting Relationship

Figure 1 shows the share of jobs that request a specific gender, and the share of ads that request a specific worker age in each of our four datasets, broken down by the job's education requirement.[13] While the overall amount of gender targeting, as already noted, is much higher in the job boards serving less skilled workers (like XMZYJS), Figure 1 also shows a strong, negative relationship between education requirements and employers' use of both gender- *and* age-based screens *within* each of the four datasets.[14] By disaggregating gender-targeted ads into those requesting women versus men, Figure 1 also shows that in most cases, explicit requests for *both* women and men decline with job skill requirements. This shows that the skill-targeting relationship is not an artifact of a change in employers' preferences for any one gender with skill. It also confirms a key feature of KS's symmetry result: that a rough equality of requests for men and women persists even within job skill categories. Part (b) of Figure 1 indicates that a similar pattern applies to age targeting.

In addition to a job's education requirement, our data contains two additional indicators of skill requirements: the level of the posted wage (when one is posted) and the requested level of experience. While these skill indicators are not as universally advertised as education requirements in our data, it still seems important to ask whether age- and gender-targeting of job ads is negatively associated with these skill indicators as well. To that end, Figures A1 and A2 in the Appendix replicate Figure 1 using wages and experience requirements as the skill indicator respectively. Once again, a robust decline in both age- and gender-targeting of job

---

on the website (at least 80 percent according to Jian Hao, Zhaopin's data division chief) already include photos. Accordingly, we do not use the photo request information in the Zhaopin data.

[12] 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 simply to write: "Age: 20-24" and/or "Sex: 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 'abogado' or 'abogada'. Since we did not extract this information, our statistics will underestimate the amount of gender-targeting in Mexican job ads.

[13] In Figure 1 and throughout the paper, an ad is defined as age-targeted if it specifies both a maximum and minimum age. (This allows us to define the requested age as the midpoint of those requirements for all age-targeted ads.) Results are unchanged if we define age targeting as expressing any sort of age preference.

[14] The only exception to a strict monotonic decline is for the lowest education category in the Computrabajo data. This exception disappears in the presence of regression controls. See Table 5.

ads is present in all four datasets for these two additional measures of skill requirements. Overall, Figures 1, A1 and A2 show that the negative skill targeting relationship first identified by KS extends to the use of both gender and age screens, in four different samples of job ads serving workers of very different skill levels in two countries.[15]

To see whether the above relationships can be accounted for by observable differences between skilled and unskilled jobs, Tables 2-5 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 dependent variable equals one if the ad specifies both a maximum and minimum age, and zero otherwise. Four specifications are presented for each of these two dependent variables. The first (in columns 1 and 5) includes occupation fixed effects in addition to the covariates shown.[16] The goal is to see whether the patterns identified in Figures 1-3 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 work men and women do 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.[17]

---

[15] The share of ads requesting workers age 30 and over does not always fall with requested experience or the wage, due to the mechanical connection between age and experience and the age-wage profile.

[16] Depending on the dataset, industry and province/state/district 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.

[17] Columns 1,3 and 4 of Tables 2-5 correspond, respectively, to columns 1 ,3 and 5 of KS's Table VI, though the results for Zhaopin differ slightly due to differences in specification and sample.

Starting with the XMZYJS data, Table 2 shows 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 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. The *XMRC* data in Table 3 show negative and statistically significant effects of education on gender targeting in all specifications; education also reduces age targeting in the presence of occupation fixed effects but the relationship is insignificant in other specifications. Experience effects on age and gender-targeting and small and imprecisely estimated, but always negative when significant, and the offered wage has a strong negative effect on both age and gender-targeting. This wage effect 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.

For the *Zhaopin* data, Table 4 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 of our regression specifications. The probability of age targeting is also negatively related to education requirements across all specifications. The anomalous positive effect of the amount of experience required on age-targeting is probably related to the mechanical relationship between a candidate's age and the amount of experience he or she can have. Turning to the Computrabajo data in Table 5, 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. Education effects on age targeting do not show a consistent pattern in the Computrabajo data, but experience has consistently negative effect, as predicted. Offered wages have a strong, negative effect on both gender-and age targeting.

In sum, Tables 2-5 show that in almost all instances, the negative skill-targeting relationship for both age and gender in all four of our datasets persists in the presence of detailed controls for the type 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. Results: The Age Twist in Firms' Gender Preferences

The age twist in firms' gender preferences in all four of our datasets is shown in Figure 2. In all cases, the data show that firms' gender preferences shift strongly away from women and towards men as the desired age of the workers they are seeking rises. Consider for example the XMZYJS data. If a firm on this job board is looking for a worker under the age of 25, the odds that it is seeking a woman are about 1.4 to one (46/32). On the other hand, among ads requesting workers over the age of 35, requests for men outnumber requests for women by more than 2.5 to one (58/22 = 2.64). This pattern is even more dramatic in our other three datasets. For example in XMRC, ads requesting women outnumber ads requesting men by more than four to one when firms are seeking workers under the age of 25, while ads requesting men outnumber ads requesting women by five to one when firms are seeking workers over 35. In Zhaopin, ads run more than five to one in favor of women for workers under 25, and more than four to one in favor of men for workers 35 and over. Finally, in Computrabajo these ratios are 2/1 favoring women among under 25's and 2.5/1 favoring men among over 35's.[18]

In Tables 6-9 we assess whether the age twist documented in Figure 4 survives regression controls, in a parallel fashion to the analysis of gender and age targeting in Tables 2-5. 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 desirability for the job being advertised.[19] With the exception that all regressions now include indicators of job's requested age level as our main regressors of interest, the specifications are identical to Tables 2-5. As for our skill-targeting regressions, the goal is to measure to what extent the age twist in firms' gender preferences 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 XMZYJS data (Table 6), all the specifications show a highly significant twist in firms' gender preferences towards men as workers age between 18 and 45. This is the case even in the presence of occupation*firm fixed effects, and even when a control for the level of the offered wage is added in column (4). The magnitude of the estimated 'twist' is large, and is relatively unaffected by adding any of the control variables. Table 6 also shows that at least 100-33.8=66.2 percent of the variance in firms' gender preferences occurs within occupation*firm cells (this number is 59.0, 26.6 and 73.9 in XMRC, Zhaopin and Computrabajo specifically), illustrating the high level of job specificity of those preferences. It follows that

[18] In a correspondence study of women's job applications in two U.S. cities, Lahey (2008) found that younger women were 40 percent more likely to receive callbacks than older women. We are not aware of any audit or correspondence studies that compare the effects of age for male versus female candidates.

[19] 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.

models with *firm*-level tastes towards one gender are not good candidates to explain our data; indeed it is very common for firms to explicitly request women in some ads and men in others, even within occupational groups, suggesting that gender-typing of jobs within plays an important role in all our datasets.[20]

Like XMZYJS, the XMRC data in Table 7 show a strong positive effect of desired age on firms' preferences towards men in all specifications; 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 requesting men (P^M-P^F) that is .55 higher than ads for workers under 25. In the *Zhaopin* data (Table 8), columns (1) through (3) show that firms' preferences tilt strongly towards men as workers age between 18 and 45, even in the presence of occupation*firm fixed effects. Interestingly, however, adding a control for the offered wage in column (4) increases the standard errors on these effects substantially, making them statistically insignificant.[21] Finally, like the XMZYJS and XMRC data, the *Computrabajo* data in Table 9 show a robust, highly statistically significant, and quantitatively large age twist against women in firms' gender preferences in all specifications.

In sum, Tables 6-9 show that –like the negative skill-targeting relationship-- the age twist in gender preferences in all four of our data sources persists in the presence of detailed controls for the type of work that is performed and the type of employer posting the ad. Indeed, in all but one case (where it is hard to distinguish age and experience effects and the results depend on how the age variable is specified) 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. We explore possible causes for these age- or experience-dependent gender preferences in Section 6.

5. Explaining the Negative Skill-Targeting Relationship

As we have noted, one interpretation of the robust negative skill-targeting effect observed in all four of our data sets is as 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—simple screening models such as KS's predict that firms should search more broadly as jobs' skill demands (indexed by θ) rise. That said, their 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. In this section we briefly

[20] Additional documentation of the role of occupations, firms and jobs in accounting for advertised gender preferences is available on the paper's website. See the "Variance Decompositions" document.

[21] The negative age effect is, however, significant in this specification when a continuous age measure is used. Note that the estimated coefficient on experience rises sharply in magnitude when a wage control is introduced, suggesting that –at least at fixed wages—the age twist is at least partly associated with a tendency for employers to request men when they are seeking *experienced* workers in the Zhaopin data. This is not the case in our other two datasets with experience indicators –XMRC and Computrabajo—where it is always age *per se* that matters.

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).[22]

Turning first to application processing costs (*c*), the screening model predicts that *ex ante* screening should become more common as *c* rises (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.[23] On the other hand, the 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 suggested by the KS model is that σ_v, the variance across jobs in men's and women's (or older versus younger workers') relative ability to perform them, might be greater in less-skilled than more-skilled jobs. In other words, perhaps age or gender 'matter less' for performance in skilled than in unskilled jobs. For example, if men and women are more different physically than mentally, and if age affects the performance of physical tasks more than mental tasks, jobs requiring manual labor –which tend to less education and experience-- might be more age- and gender-specialized than other jobs.[24] A related hypothesis applies this same logic to employers, co-workers' or customers' tastes: perhaps skilled co-workers, as well as employers or customers of skilled workers, just 'care less' about the age and gender of people they interact with. In their paper, KS devise a simple test between the σ_v–based and θ –based explanations of the negative skill-targeting relationship, based on the idea that the models have different predictions for the effects of skill in jobs that are that are highly gendered (i.e. jobs where more than half the ads explicitly requests one of the two genders) than in other jobs. When we perform similar tests in the current data, as in KS our results favor a direct effect of skill (θ).[25]

---

[22] A fourth possibility suggested by KS's model is that the idiosyncratic variation in applicant quality (σ_ε) is higher at higher skill levels. Empirically, this 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.

[23] 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.

[24] Note that effects of this nature would need to occur within firm*occupation cells to explain our results. KS also show that the skill-targeting effect persists when all jobs requiring manual labor are dropped from the sample.

[25] Details are provided on the paper's website. See Table S-1 of the supplementary tables.

Stepping outside KS's baseline model, one final factor that might explain the negative skill-targeting relationship is an increase in the stigma associated with posting an age- or gender-targeted ad with the job's skill level. Specifically, if employers can cheaply filter applications by age and gender after they have been submitted, employers who intend to consider only applications from one age group or gender can easily avoid stigma by accepting applications from everyone, then simply discarding applications from the non-desired group. By this logic, an increase in employer stigma with skill would lead firms to substitute internal for external filtering, thus reducing their use of age- and gender-targeted ads without necessarily changing their hiring decisions. While we cannot definitively rule out this explanation with the data at our disposal, we note that our conversations with job board officials suggested that very little stigma is associated with posting an age- or gender-targeted job ad in China, at any skill level.[26]

In sum, we can 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, nor by a tendency for firms, co-workers or customers to 'care less' about age and gender in more skilled environments. We do acknowledge that --in addition to the direct effect of skill on the value of identifying the best job candidate-- the negative skill-targeting relationship could be attributable to systematically greater tightness of labor markets at higher skill levels, or to a greater stigma associated with posting an age- or gender-targeted ad at high skill levels. But since the negative skill-targeting relationship appears for both age and gender in all four of our data sets, any such tendency of labor market to 'tighten up', or for employer stigma to fall as skill demands rise, would need to be quite universal to explain our results.

6. Explaining the Age-Related Twist in Firms' Gender Preferences

a. Occupation shifts

As discussed in Section 4, one possible explanation of the age twist in employers' gender preferences is based on gender differences in actual or perceived suitability for different types of work. For example, among jobs targeted at younger workers, the mix of tasks demanded by firms might consist disproportionately of tasks that women are considered to be better at than men, while the opposite might be true for older workers. Of course, if tasks are measured by firm*occupation cells, then the regression results in Tables 6-9 suggest that this is not the main explanation of the age twist in firms' gender preferences: adding firm*occupation fixed effects does not substantially reduce the estimated size of the twist.

---

[26] One possible test of the stigma hypothesis would be to see whether gender segregation in actual hiring declines with skill at the same rate as gender targeting in job ads. If it does, then firms are probably not substituting internal for external filtering as skill levels rise. Unfortunately, we do not have data on which employees are hired in any of our four datasets.

To address the possibility that the results in Tables 6-9 are an artifact of the relatively broad occupation categories used there, we were able to construct a much more detailed occupation classification in one of our datasets: XMRC.[27] Specifically, we take advantage of the fact that the XMRC website allows a firm to list up to 5 of its 36 occupation categories for each job ad.[28] 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 1,701 distinct occupation combinations compared to the 36 used in Table 7. Re-estimating Table 7 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). We conclude that the age twist in firms' gender preferences in our data is not a consequence of firms demanding a different mix of tasks from older versus younger workers.

b. Occupational characteristics

The fact that the estimated age twist in firms' gender preferences survives detailed occupation controls does not, of course, imply that the twist is equally strong in all occupations. To see which occupations exhibit the strongest age twist, Figures 3-6 show the means of our "preferences for men" variable ($P^M - P^F$) by occupation, separately for ads requesting workers under 30 versus 30 and older. Thus, for example, Figure 3 shows that among ads for administrative staff under age 30, ads for women outnumber ads for men by over 60 percentage points in the XMZYJS data. In contrast, ads for administrative personnel aged 30 and older favor women by only a 25 percent margin. Occupations in the upper right quadrant of the figures are targeted disproportionately at men regardless of age; occupations in the lower left are targeted disproportionately at women regardless of age. An occupation's (vertical *or* horizontal) distance from the 45-degree line shows the size of its age twist. Occupations above the 45-degree line experience a twist towards men as workers age, while occupations below it experience a shift away from men.

As a group, Figures 3-6 reveal the following. First, confirming our regression results, the age twist in firms' gender preferences is broadly based: in all datasets, a large majority of the points are above the 45-degree line. A striking related fact is that there are no observations in the bottom right quadrant in any dataset: *No* occupations switch from preferring men to women among young workers to preferring women to men among older workers. Related, the bottom left quadrant is relatively empty as well: while many occupations have a strong

[27] The number of occupation categories in Tables 2-9 is roughly comparable across the data sets, at 58, 36, 39 and 14 for XMZYJS, XMRC, Zhaopin and Computrabajo respectively.

[28] All our XMRC estimates reported so far are based on the first occupation listed in the ad. Multiple occupations are, however, listed in 38.2 percent of the ads. Since the propensity to list more than one occupation is essentially the same in ads for a single opening (64.7 percent) as it is overall, we interpret multiple occupation listings primarily as a means of providing more detail about each job's duties.

preference for young women, occupations with strong preferences for older women (for example, values of $P^M - P^F$ below -.2) are essentially nonexistent, in striking contrast to the large number of occupations that strongly prefer men in both age groups. Instead, a rough overall visual impression from all the figures is of two types of occupations. One group of occupations lies fairly close to the 45-degree line in about the rightmost two thirds of the figures. At younger ages, these jobs range from having a small preference for women to strong preferences for men (e.g. real estate and construction in XMRC). When older workers are sought, all of these jobs prefer men, though the size of their twist towards men is roughly similar across occupations. Other occupations in this group include technical workers, engineering, manufacturing workers and quality control, though no single occupation consistently stands out as an outlier.[29]

The second group of occupations is located around the horizontal axis to the left of the origin: Some of these occupations exhibit a considerably stronger shift in employers' preferences with age, in several cases switching from intense preferences for women to gender neutrality when workers under versus over 30 are sought. Most of the customer-contact occupations in our data –denoted by triangles—belong to this group, but probably the most consistent outlier is administration, which shifts from a 70 percentage point preference for women to gender neutrality in the XMRC data, and from a 35 percent female advantage to gender neutrality in Zhaopin.[30] Similar shifts occur in XMZYJS and Computrabajo, though employers exhibit a small female preference even among older workers in these datasets.[31] While we had no idea what types of jobs the "administrative" occupation category represented when we started working with job board data, subsequent conversations with job board officials revealed that these ads were primarily for secretaries, administrative assistants and receptionists.

The tendency for customer-contact and administrative jobs to strongly favor young women over young men, while being largely gender neutral among older workers suggests that some trait that is disproportionately present and highly valued among young women in those types of jobs might help explain why these occupations experience such a strong 'age twist'. One candidate, suggested by a recent ethnographic study of young female jobseekers in China (Hua, 2013) is employers' demand for beauty. A role for beauty is also suggested by the statistics in Figure 7(a), which shows that employers' beauty requests are highly gendered in the two datasets where that information is available. For example, in Zhaopin, ads that do not request beauty are not very likely to request either a man or a woman specifically, though they

---

[29] Occupations that involve supervising others (primarily "management") are denoted by squares in Figures 2-5. Ads for these jobs tend to be gender neutral among young workers, but with some male preference among older workers. This shift is neither large in magnitude nor statistically different from other occupations, however.
[30] Consistent with an employer preference for attractive women in customer service occupations, Landry et. al (2006) find a large positive effect of beauty on women's productivity in a field experiment on charitable giving.
[31] The age twist in administrative jobs is significantly greater than average in all our datasets except Computrabajo.

do request men twice as often as women (at 6 versus 3 percent of ads). Ads that request beauty, on the other hand, are six times as likely to request a woman as a man, at 24 versus 6 percent. Viewed a different way, 37 percent of ads requesting women in the Zhaopin data explicitly request that the applicant be physically attractive, compared to 5 percent in ads requesting men. A similar pattern is present in the Computrabajo data, though the overall level of gender targeting is higher.[32] Further corroboration is provided by the data on photo requests in Figure 7(b): Computrabajo ads that do not request the applicant provide a photo invite men and women to apply with equal frequency, but women are sought more than twice as often as men when applications require a photo.

Another type of explanation for the age twist is a trait that is disproportionately present, or valued, in older men. One candidate, suggested by the literature on gender and leadership, could be a preference to have men in leadership positions.[33] More specifically, if older workers, on average, are expected to do more supervision, and if firms view men as better or more acceptable supervisors, employers' preferences could shift towards men as they are hiring older workers. This idea is also supported in Figure 7(b), which shows that firms request men and women at equal rates in jobs that do not involve supervising other workers. However, if supervisory duties are involved, men are invited to apply more than twice as often as women (at 29 versus 12 percent of ads).

To assess the role of employer demands for beauty and supervision in explaining which occupations had the biggest age twist, Table 10 presents some simple cross-occupation regressions in the Zhaopin and Computrabajo datasets. The dependent variable in these regressions is the size of an occupation's age twist away from women in Figures 5 and 6 (i.e. their distance from the 45 degree line). The single regressor in the Zhaopin data is a continuous variable equal to the share of ads in the occupation that explicitly request an attractive applicant. Column 1 of the table strongly indicates that occupations with a higher demand for beauty exhibit larger shifts in employers' preferences away from women as workers age.[34] In the Computrabajo data, we are able to add measures of the occupation's mean requests for photos and for engaging in supervision to the same regression. When we do so, column 2 shows that occupations where a large share of jobs require supervising others also have a

[32] This strong association between employers' demands for beauty and job ads directed at women contrasts with a number of studies of beauty in the U.S. and U.K. that do not find large gender differences in the effects of beauty on wages (Frieze et al. 1991, Hamermesh and Biddle 1994, Biddle and Hamermesh 1998, Harper 2000, and Mobius and Rosenblatt 2006) or on a candidate's success in interview or resume audit studies (Dipboye et al. 1977, Heilman and Saruwatari 1979, and Mack and Rainey 1990). While part of this difference may reflect the countries we study we also note that none of the U.S. and U.K. beauty studies examine how gender and beauty interact with age in the hiring process.

[33] Reuben et al. (2012) study experiments in which groups performing a real-effort task choose their own leaders. They find that women are selected much less often as leaders than is suggested by their past performance.

[34] This relationship remains significant when the administrative occupation is dropped from the sample; thus it is not driven by firms' desires for attractive young women in administrative jobs alone.

stronger age twist, suggesting that a preference for male leaders accounts for some of the twist. When both beauty and photo requests are entered as separate regressors (column 2) only the latter has a statistically significant association with the age twist, but a combined beauty-photo indicator (column 3) is highly statistically significant.

c. Explicit ad-level employer requests: beauty, marital status and leadership

Another source of information about the motivations for employers' gender preferences is to ask how those preferences covary at the *ad* level with other worker characteristics that are sometimes requested by firms. In this section, we focus on three such characteristics: beauty, leadership and marital status. The reasons why beauty and leadership might account for the age twist have already been discussed. 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 entering marriage between age 18 and 45 could also explain the sharp reversal in employers' advertised gender preferences.[35] This idea is supported by Figure 8, which shows that employers' marital status requests (only available in the Computrabajo data) are highly gendered. Specifically, among ads requesting single applicants, ads requesting women outnumber those requesting men by more than four to one (60 versus 14 percent). Among ads requesting married applicants, men are requested more than six times as frequently as women (53 versus 8 percent).

To assess the ability of employers' explicit, ad-level requests to account for the age twist in our data more formally, Appendix Tables A2 and A3 add an indicator for whether the job ad requests beauty in the Zhaopin and Computrabajo datasets to the regressors in Tables 8 and 9. In the Computrabajo data, we also add indicators for whether the ad explicitly requests a single or married applicant to these regressions, plus indicators for whether the ad requests a photo or supervisory duties. Consistent with Figure 7(a), Tables A2 and A3 show that employers' requests for physically attractive workers are highly gendered in both data sets and in all regression specifications: requests for beauty strongly increase the chances a firm is looking for a woman, relative to a man. In addition, consistent with the male marriage premium and motherhood penalty literatures, the Computrabajo data show that employers strongly 'want men to be married and women to be single'. Finally, consistent with the notion that employers prefer men for leadership roles, ads for supervisory jobs are much more likely to request men than women. Note that neither the estimated beauty, marital status nor supervision effects

[35] A sizable literature on the male marriage wage premium (Korenman and Neumark 1991, Ahituv and Lerman 2007, Loughran et al. 2009, Petersen et al 2011) suggests that firms value marriage among men. While marriage does not appear to have strong effects on women's wages in developed countries, there is widespread evidence of a motherhood penalty (Korenman and Neumark 1992, Fernandez-Kranz et al 2013). We explore the effects of motherhood in the next subsection.

vary appreciably in magnitude as controls for occupations and occupation-firm interactions are added.[36]

In sum, Tables A2 and A3 show that explicit employer requests for beauty, marital status and leadership are highly correlated with firms' gender preferences in our data, not only overall but even within firm*occupation cells. But can these *explicit requests* account for the broadly-based age twist in the data? Clearly not: comparing the age coefficients in columns 1-3 of Tables A2 and A3 to their counterparts in Tables 8 and 9 shows very little effect of adding these new controls. In part, this is because explicit requests for beauty, marital status and supervision are relatively rare, and probably do not capture all the variation in firms' demands for those attributes. Accordingly, we consider other possible causes of the twist below.

d. Unstated employer preferences

A key prediction of KS's screening model is that even when a characteristic is predictive of expected job performance, employers should not necessarily use it as an *ex ante* hiring screen. Instead, employers should only announce their preferences for an attribute in a job ad when those preferences are intense, or when the attribute is a precise predictor of productivity. For this reason, we should expect Table A2 and A3's regression controls for explicit beauty, marital status and leadership requests to give only a lower-bound estimate of their role in explaining the age twist. For example, a preference for men in supervisory roles could account for a broad-based age twist if supervisory responsibilities in most jobs tend to increase with age even when this is not explicitly specified in job ads. A widespread desire to interact with attractive young women could have a similar effect if those beauty preferences are not always advertised. Sex-specific employer preferences for marriage can have a similar effect as a cohort of young workers forms new unions; these preferences are not measured in three of our datasets, and may be only poorly measured in Computrabajo where fewer than four percent of ads express a marital status preference. Possibly most important, parenthood may play a role, as suggested by a large literature documenting a negative relationship between women's labor market outcomes and childbearing across a variety of countries (Korenman and Neumark 1992, Loughran et al. 2009, Fernandez-Kranz et al. 2013). Thus, women's movement into their childbearing years between the ages of 18 and 45 could play a large role in the age twist against them, even if firms rarely express explicit preferences concerning motherhood in job ads.[37]

While we cannot measure firms' demands for workers with different childbearing risks or expected childcare duties directly, one way to assess their relevance to the age twist in employers' advertised gender preferences is to ask whether the age pattern in marriage and

---

[36] The effect of supervisory duties, however, vanishes when we add controls for offered wages; this could reflect a need to pay men more to take supervisory positions than women.

[37] While a few Computrabajo ads do, in fact mention a preference for women without children, there are not enough explicit requests to allow a quantititative analysis.

child care responsibilities matches the pattern in the age twist in the two countries we study. To that end, Figure 9 plots the share of urban men and women who are married as a function of age, using data from the 2005 and 2010 Censuses of China and Mexico respectively. Unsurprisingly, both men and women are moving rapidly into marriage between the ages of 18 and 45 in both countries. Importantly, though, marriage rates are quite stable in both countries after age 30, so marriage cannot easily account for the continued age twist in employers' gender preferences after age 30. That said, the large decline in firms' explicit requests for young women between the ages of 18 and 29 in all four of our datasets could indeed be linked to an unstated employer preference for married men and single women, combined with men's and women's rapid movement into marriage between those ages.

In the same vein, Figure 10 plots the share of Chinese and Mexican women who *currently* have a child under the age of 6 as a function of age. In urban China, this share is maximized (at 44 percent) at age 29. After that, the share 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, the probability of currently having a child under the age of six also peaks at age 29 (46 percent). The share of women with a child under six then declines to 36 percent at 34, 22 percent at 38, and just below 10 percent at 42. Put a different way, Figure 10 shows that the age bin in Figure 2 with the highest share of mothers with preschool-aged children is the 25-29 group. Thus, if care of pre-school-age children explains firms' declining attraction to women as they age, Figure 10 suggests that 30-34 and 35-45-year-old women should be *more* attractive to employers than 25-29-year olds. According to Figure 2, that is not the case in any of our data sets.[38] Thus, any child-care based explanation of the age twist in employers' gender preferences would need to argue that the care of older children requires greater maternal time and energy than caring for pre-school-age children. This seems especially unlikely in China, where only about 14 percent of women have more than one child, and child care by grandparents is widespread.[39]

Related, Figure 11 presents data on the share of urban women who are employed by age, from the same Census sources as Figures 9 and 10. While women's employment rates are lower at most ages in Mexico than in China, women in both countries exhibit *increasing* labor force attachment between the ages of 18 and 25, and constant employment rates (of about 70% in China and 50% in Mexico) between 25 and 45.[40] Neither country shows any sign of a dip in women's employment rates at the ages when most women are having and caring for their

[38] There is a slight reversal in the XMZYJS data, but this disappears in the presence of controls.

[39] Note also that a substantial share of urban Chinese workers leave their children behind in rural areas (Démurger and Xu, 2013). Care of aging parents could be a demand on older women's time in China, though in most cases this would not be a major factor for women in the age ranges (18-45) studied in this paper.

[40] In this age range, urban Chinese women's employment rates also do not differ between married and single women, or between women with and without children, raising additional concerns about the role of family responsibilities in explaining firms' preferences for men among workers in this age group.

youngest children. Even more than the fertility patterns in Figure 10, this pattern raises challenges for an explanation of the age twist in firms' gender preferences that is based on age-related changes in women's labor force commitment.[41]

Finally, to assess the possible role of unstated employer preferences for beauty and leadership, Figures 12 through 14 show the detailed age pattern in advertised requests for beauty, photos and supervisory duties in the Zhaopin and Computrabajo data, and compare these to the detailed age pattern in advertised requests for men and women. Since these figures are informative about unadvertised preferences only to the extent that they are positively correlated with advertised preferences, they are suggestive only. That said, the patterns are striking, especially in the Zhaopin data, where the age pattern in (overall) beauty requests coincides almost exactly with the age pattern in explicit requests for female job applicants.[42] In Computrabajo, explicit demands for beauty fall faster with age than requests for female workers, while requests for photos fall more slowly. All three series, however, level out after age 35, suggesting a possible connection. In the same vein, Figure 14 shows employers' requests for men and their demands for supervisors both beginning to rise at about age 28, then leveling out, also suggesting a possible relationship.

Overall, based on recent age patterns in childbearing and labor force participation in these two economies, employer preferences for workers without child care responsibilities, and age patterns in women's labor force commitment do not seem to be promising explanations of the observed age twist in advertised gender preferences. More promising are unadvertised, gendered employer preferences for their workers' marital status, which could help account for the rapid age twist in firms' explicit gender preferences near the start of the work life (between the ages of 18 and about 29). Unstated preferences to have 'older' men in leadership positions could also account for the rapid increase in advertised requests for men around age 28. Finally, the remarkable coincidence of the age patterns in employers' requests for beauty and their requests for women, especially in the Zhaopin data, suggests that unadvertised age-dependent demands for female beauty may also play a role.

e. Wages and supply-related factors

Another possible interpretation of the advertised preference patterns in all our data sources is that they represent variation not in firms' preferences or productivity assessments (as we have been assuming), but in the relative wage costs of hiring different types of workers.

---

[41] On the surface, the positive correlation between firms' requests for experienced workers and their requests for men in the Zhaopin and Computrabajo data –see column 4 of Tables 8 and 9— seems consistent with the idea that firms see men as less likely to quit. On reflection, however, this is not at all obvious: The ads are more likely to insist on minimum levels of *past* experience for men, but 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.

[42] Recall from Figure 7 that requests for women and beauty are far from perfectly correlated in the microdata: 76 percent of ads requesting women do not request beauty, and 63 percent of jobs requesting women do not request beauty. Nevertheless, the strong coincidence in the age pattern of these requests seems highly suggestive.

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, just 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 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 wage cost- or supply- based explanations of the advertising patterns in our data are highly unlikely. One reason is that the patterns of gender wage gaps needed to explain the age twist are at odds 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 wage 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.[43] 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.

Second, if a disproportionate number of employer requests for young women and older men is an indicator of a higher level of employer demand for those workers, we should see evidence that jobs targeted at those workers are harder to fill. Pursuing this idea, we turn to the Zhaopin data, where we have access to an indicator of vacancy duration. Specifically, we know the number of times each ad was renewed during our observation window. As on many online platforms, workers viewing Zhaopin ads pay closer attention to newly-posted ads than old ones; in consequence it is common for employers to repeatedly re-post the same job ad until it is filled. Indeed an average ad in the Zhaopin data was renewed 6.8 times. Figure 15 shows how this renewal rate varies with the age requested in the ad, separately for jobs requesting men, jobs with no gender preference, and jobs requesting women. 95% confidence intervals for local polynomial regressions with optimal epanechnikov kernels are shown.

Confirming our previous analysis, Figure 15's wide confidence intervals for 'male' and nongendered jobs targeted at workers under 22 or so reflect the small number of ads targeted at those groups. In contrast, the renewal rate for 'female' jobs has very tight confidence bands among the youngest workers where ads are frequent, and wide bands among ads requesting older women. More importantly, Figure 15 indicates that ads for young women in their late

[43] See the Supplementary Tables on the paper's website. Table S-2 reveals a flatter age-wage profile for women than men in all four of our data sets (in ads where age, gender and wage information are present). The same is true for experience-wage profiles in the three data sets that contain an experience variable. Finally, Table S-3 confirms that men experience a marriage-wage premium and women a marriage-wage penalty in the Computrabajo data.

teens and early 20s are apparently much harder to fill than ads for women 25 and older, while ads for men become harder to fill as older candidates are sought. Indeed, the amount of apparent excess demand for very young women in our data is striking: ads for 17- and 18-year old women are renewed almost 15 times on average, compared to an average of 6 or 7 times for women aged 25-35, and only four to six times for men under 20.[44]

In sum, when viewed in the context of the cross-sectional wage patterns and vacancy durations 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 disproportionately in job ads *and* drive up older men's relative wages. 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 should attenuate firms' attraction to them.[45]

f. Summing up: what explains the age twist?

Overall, the analysis in this section establishes the following facts about the age twist in employers' explicit gender preferences. First, subsections (a) and (b) showed that the twist is broadly based: it is present within almost all broad occupations and even within the most detailed firm*occupation cells we can construct. Second, however, the cross-occupation analysis in subsection (b) showed that the shift is stronger in *occupations* where employers explicitly value beauty and where supervisory responsibilities are more prevalent. This suggests that employer demands for attractive young women, and preferences to have men in leadership roles may play some role in explaining the age twist. Third, subsection (c)'s ad-level analysis showed that explicit requests in a job ad for beauty and single marital status are highly correlated with firms' requests for women, while requests for married workers and leadership strongly predict explicit requests for men. That said, explicit requests for beauty, marital status and leadership do not account for a significant share of the age twist in our ad-level analysis,

---

[44] Figure 15's vacancy duration analysis highlights the fact that all our main estimates of firms' age and gender preferences incorporate two distinct sources of firms' excess demands for different worker types: (1) firms may post more *new* ads for certain types of workers, and/or (2) ads for that type of worker are overrepresented in the vacancy stock because they are harder to fill. Since both are indicators of higher employer demand, we do not distinguish them in our main analysis. However, since our sampling windows are long relative to most vacancy durations, and all our main results persist for an inflow sample of vacancies (see footnote 4), we know that the patterns in Tables 6-9 are not primarily due to selection by duration; they mostly reflect a higher inflow of ads for the employers' most desired groups.

[45] In a labor market with roughly equal numbers of male and female candidates but a highly disproportionate share of ads targeted at (say) older men, the excess supply of older women can be absorbed in two ways—longer search durations, or simply by women taking a higher share of the non-gendered jobs. Thus, while shorter vacancy durations for older women are supportive of a demand-driven scenario, they are not essential for markets to clear.

possibly because those requests are rare and only represent especially intense employer preferences for those attributes.

What, then, might account for the broad-based shift in employers' advertised preferences away from women as workers age? Obvious candidates are *un-advertised* employer preferences for beauty, leadership, marital status and parenthood that vary with age and gender. While we cannot definitively rule these explanations in or out, we note that explanations based on women's higher child care responsibilities and lower levels of labor force commitment do not match well with age trends in childbearing and labor force participation in the countries we study. At the same time, the strong predictive power of explicit requests for beauty, marital status and leadership *when they are made*, and the coincidence in the age pattern of employers' requests for these characteristics with their gender requests, suggests that (perhaps optimally) unstated employer preferences for beauty, marital status and leadership may play a larger role in the age twist than our ad-level regressions suggest.

7. 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 *ex ante* screening mechanisms as jobs' skill demands rise. As noted, these results are consistent with a simple model of hiring from two pools in the presence of application processing costs, though they could also be explained by a greater level of stigma for posting gendered and age-targeted ads at higher skill levels. In either case, the strong tendency for firms to abstain from explicit age- and gender discrimination as skill requirements rise suggests that skill upgrading may be a potent force in reducing the prevalence of those employer practices.

Second, we identify a new empirical regularity—a dramatic 'twist' in employers' advertised gender preferences away from women and towards men as workers age. Like the skill-targeting relationship, this phenomenon is highly robust to estimation methods and statistical controls, and is present in four data sets drawn from different countries and serving very different skill levels. This age twist sheds interesting new light on KS's 'symmetry' result, that even within skill groups, employers tend to request men and women with roughly equal frequency. While confirming this symmetry, we show that it is actually the net result of a strong preference for young women and a strong preference for older men.

What explains this apparently universal pattern in the effect of age on men's and women's relative desirability to employers? In the paper we argue that age patterns in childbearing and labor force attachment in the countries we study do not fit the timing of the age twist well. We also show that explicit employer requests for beauty, marital status and leadership are highly correlated with firms' requests for young women and older men, and argue that –since these requests are (optimally) made only when employers' preferences are especially intense—unstated, gendered preferences for beauty, marital status and leadership might account for a substantial share of the age twist in our data.

Given that beauty plays at least some role in the age twist, what might explain its effect? One interpretation that is consistent with our data is the notion that female beauty is an asset that is valued in the labor market, but depreciates rapidly with age. Indeed some observers (including Hakim 2011) have referred to such an asset as *erotic capital*, and argue that it plays an important but underappreciated role in markets for labor, marriage, dating and sex.[46] Data collected by anthropologists suggests that women's physical attractiveness is more important to potential mates than men's in almost all societies, and evidence on reproductive patterns in small forager societies has been argued to support an evolutionary basis for men's attraction to young, attractive, low-parity women.[47] To the extent that men do much of the hiring for the jobs we study, or that customers and co-workers derive utility from interacting with attractive young women, it is not inconceivable that preferences of this type play a role in the 'age twist' we observe. While economists studying age-wage profiles generally attribute women's flatter wage profiles to gender differences in human capital acquisition and specialization in home production, our results suggest that the depreciation of young women's 'erotic capital', as well as cultural norms favoring older men in leadership positions, may play a role as well.

---

[46] The young female jobseekers interviewed by Hua (2013) in Beijing explicitly refer to beauty as 'capital', and refer to their expenditures on plastic surgery as investments in their labor market prospects. Etcoff (1999) provides a wide-ranging discussion of the asset value of female beauty over the life cycle in modern societies, and Low (2014) studies the effects of women's depreciating reproductive capital in U.S. marriage markets.

[47] Ford and Beach (1951) found that women's physical attractiveness received more "explicit consideration" in partner preference than men's in almost two hundred tribal cultures, and Buss (1989) found that men valued physical attractiveness and good looks in a partner more than women did in 34 of the 37 cultures he studied. Evidence from small forager societies suggests that youth, attractiveness and low parity are predictive of women's future reproductive capacity (Sugiyama 2005); given the high scarcity of female reproductive capacity in those societies a male preference for them could have substantial fitness value.

References

Ahituv, Avner; Lerman, Robert I. (2007) "How do marital status, work effort, and wage rates interact?" *Demography* 44(3): 623-647 DOI: 10.1353/dem.2007.0021

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.

Biddle, J. E., & Hamermesh, D. S. (1998). Beauty, productivity, and discrimination: Lawyers' looks and lucre. *Journal of Labor Economics, 16*(1), 172-201.

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

Buss, David M. (1989) "Sex differences in human mate preferences: Evolutionary hypotheses tested in 37 cultures". *Behavioral and Brain Sciences* 12 pp. 1-14.

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.

Démurger, Sylvie and Hui Xu. "Left-Behind Children and Return Decisions of Rural Migrants in China" IZA discussion paper no. 7727 (November 2013).

Dipboye, R. L., Arvey, R. D., & Terpstra, D. E. (1977). Sex and physical attractiveness of raters and applicants as determinants of resumé evaluations. *Journal of Applied Psychology, 62*(3), 288-294. doi:http://dx.doi.org/10.1037/0021-9010.62.3.288

Etcoff, Nancy. *Survival of the Prettiest: The Science of Beauty.* New York, Doubleday, 1999.

Fernandez-Kranz, Daniel; Lacuesta, Aitor; Rodriguez-Planas, Nuria. (2013) "The Motherhood Earnings Dip Evidence from Administrative Records" *Journal of Human Resources* 48(1): 169-197.

Ford, C. S., and Beach, F. A. *Patterns of Sexual Behavior.* New York: Harper and Row, 1951.

Frieze, I. H., Olson, J. E., & Russell, J. (1991). Attractiveness and income for men and women in management. *Journal of Applied Social Psychology, 21*(13), 1039-1057.

Hakim, Catherine. *Erotic Capital: The Power of Attraction in the Boardroom and the Bedroom* New York: Basic Books, 2011

Hamermesh, D. S., & Biddle, J. E. (1994). Beauty and the labor market. *American Economic Review, 84*(5), 1174-1194.

Harper, B. (2000). Beauty, stature and the labour market: A British cohort study. *Oxford Bulletin of Economics and Statistics, 62*, 771-800.

Heilman, M. E., & Saruwatari, L. R. (1979). When beauty is beastly: The effects of appearance and sex on evaluations of job applicants for managerial and nonmanagerial jobs. *Organizational Behavior & Human Performance, 23*(3), 360-372.

Hua, Wen. *Buying Beauty: Cosmetic Surgery in China* Hong Kong: Hong Kong University Press, 2013.

Korenman, Sanders and David Neumark (1991) "Does Marriage Really Make Men More Productive?"Journal of Human Resources 26(2): 282-307

Korenman, S., and D. Neumark (1992) “Marriage, Motherhood, and Wages” *Journal of Human Resources* 27(2): 233-255

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

Lahey, Joanna N. (2008) “Age, Women, and Hiring: An Experimental Study” *Journal of Human Resources* 43:30-56 (December).

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

Loughran, David S. and Zissimopoulos, Julie M. “Why Wait? The Effect of Marriage and Childbearing on the Wages of Men and Women” *Journal of Human Resources* Volume: 44 Issue: 2 Pages: 326-349 Published: SPR 2009

Low, Corinne. "Pricing the Biological Clock: Reproductive Capital on the US Marriage Market" unpublished paper, Columbia University, January 1, 2014

Mack, D., & Rainey, D. (1990). Female applicants' grooming and personnel selection. *Journal of Social Behavior & Personality, 5*(5), 399-407.

Mobius, M. M., & Rosenblat, T. S. (2006). Why beauty matters. *American Economic Review, 96*(1), 222-235. doi:http://dx.doi.org/10.1257/000282806776157515

Petersen, Trond, Andrew M. Penner, and Geir Hogsnes (2011) “The Male Marital Wage Premium: Sorting versus Differential Pay” *Industrial and Labor Relations Review* 64(2): 283-304.

Reuben, E., Rey-Biel, P., Sapienza, P., & Zingales, L. (2012). The emergence of male leadership in competitive environments. *Journal of Economic Behavior and Organization, 83*(1), 111-117. doi:http://dx.doi.org/10.1016/j.jebo.2011.06.016

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: The Negative Skill-Targeting Relationship--
Share of Ads That Are- Age and Gender-Targeted, by Desired Education

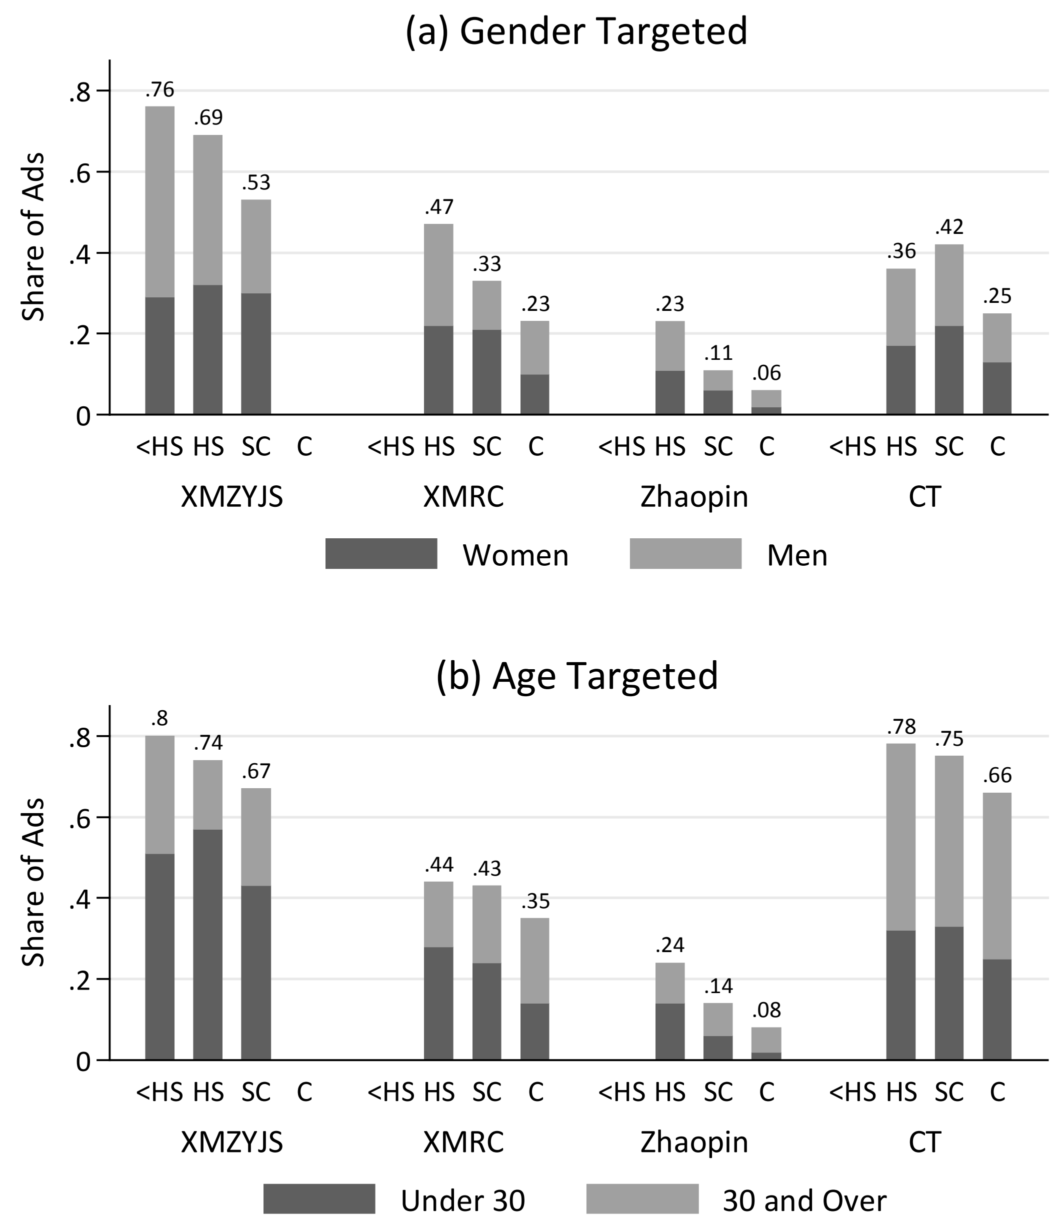

Notes: <HS, HS, SC, and C denote less than high school, high school, some college, and college or more respectively. CT refers to the Computrabajo data.

Figure 2: The Age Twist--

Share of Ads Requesting Women and Men, by Desired Age

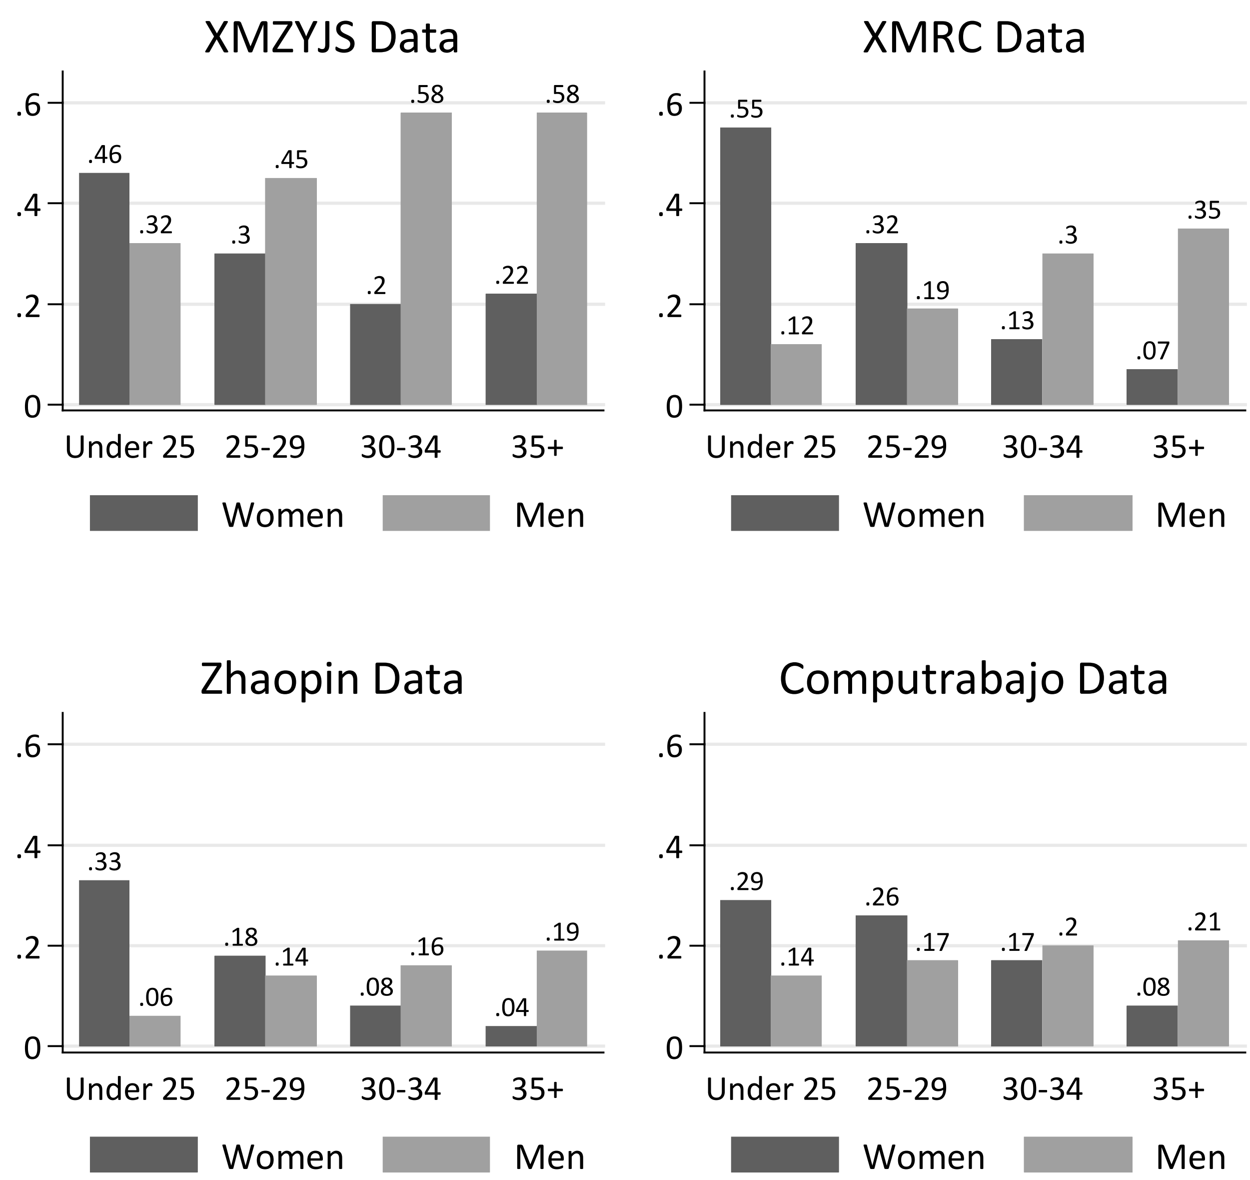

Figure 3

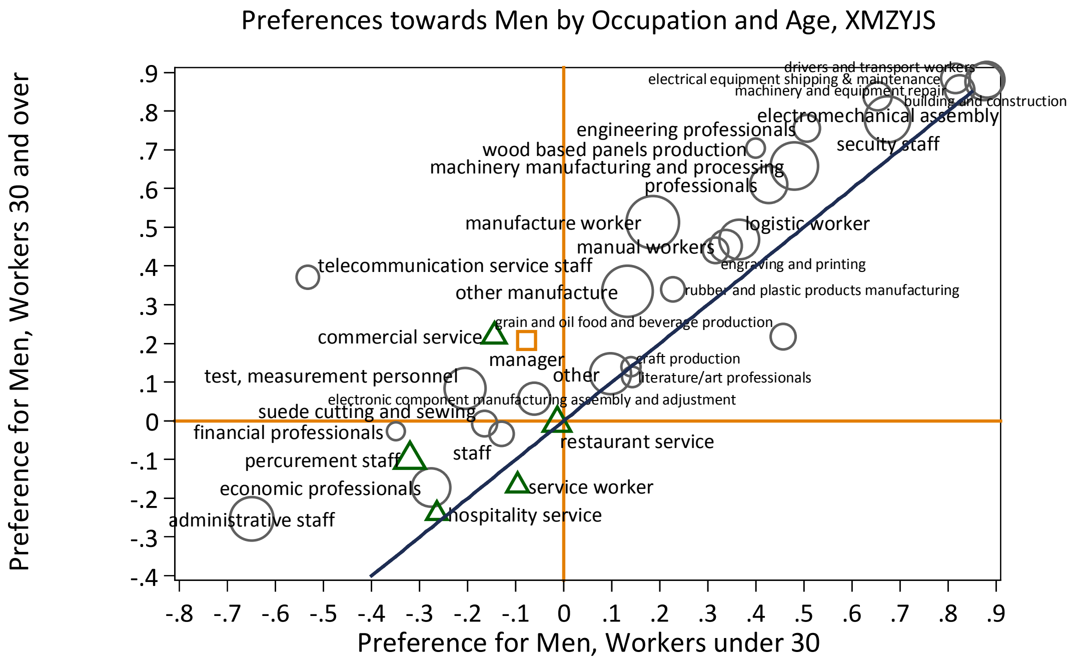

Figure 4

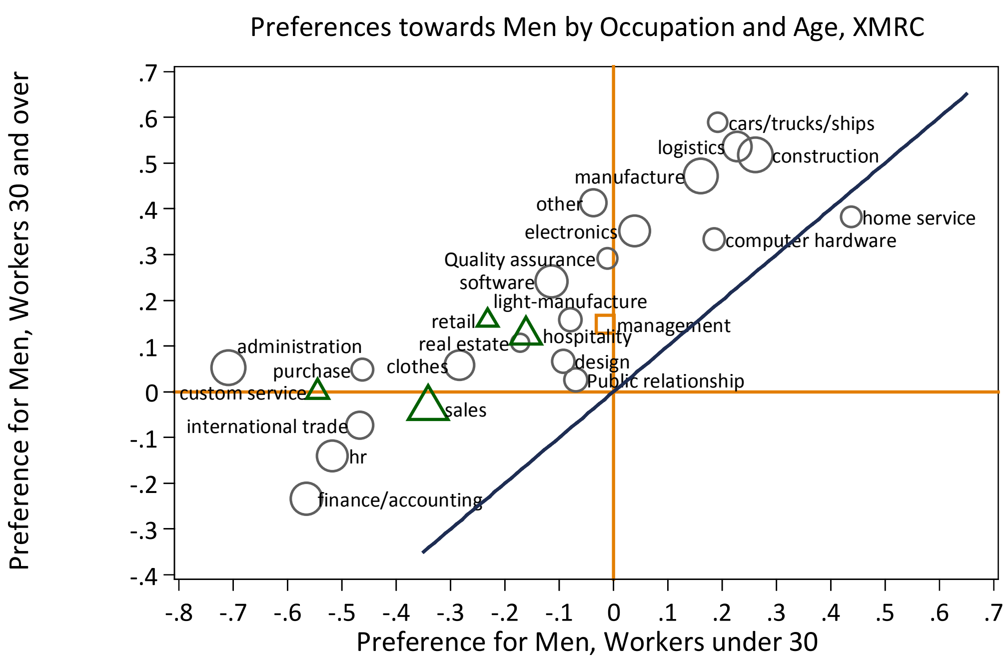

Figure 5

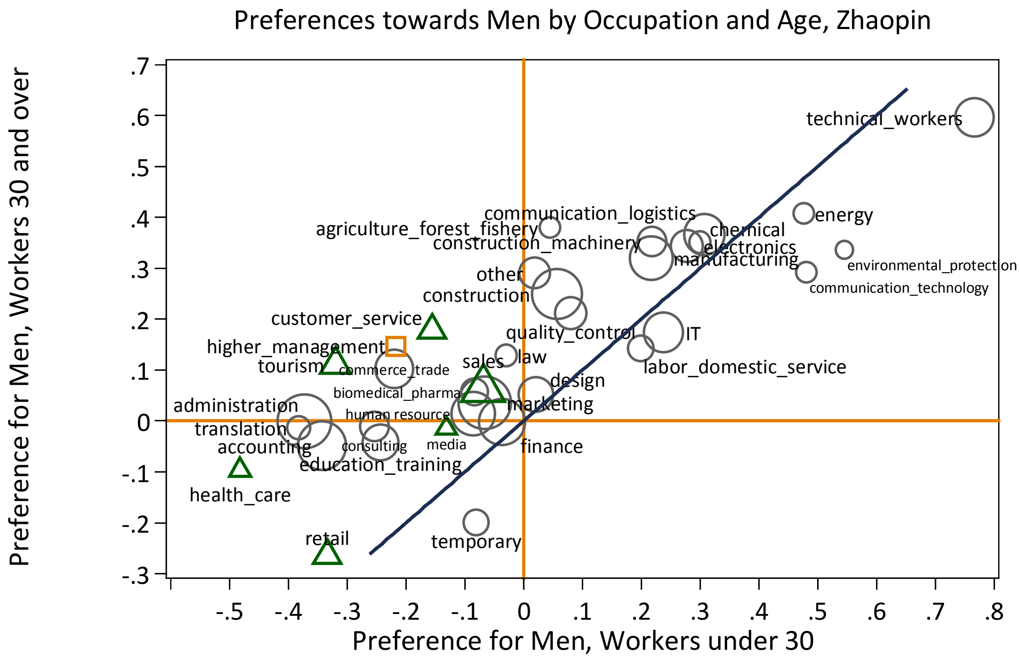

Figure 6

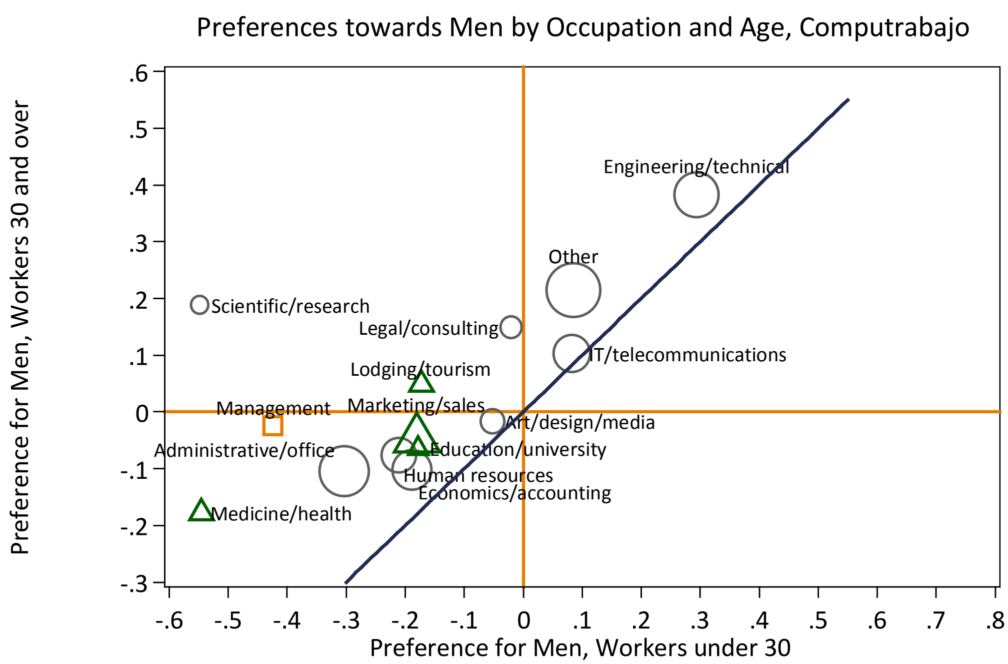

Symbol size in Figures 3-6 is proportional to the square root of the number of ads in the occupation. Customer contact occupations indicated by triangles, managerial ones by squares. Occupations with fewer than 100 observations are excluded.

Figure 7: Share of Ads Requesting Women and Men, by Advertised Beauty, Photo and Leadership Requests

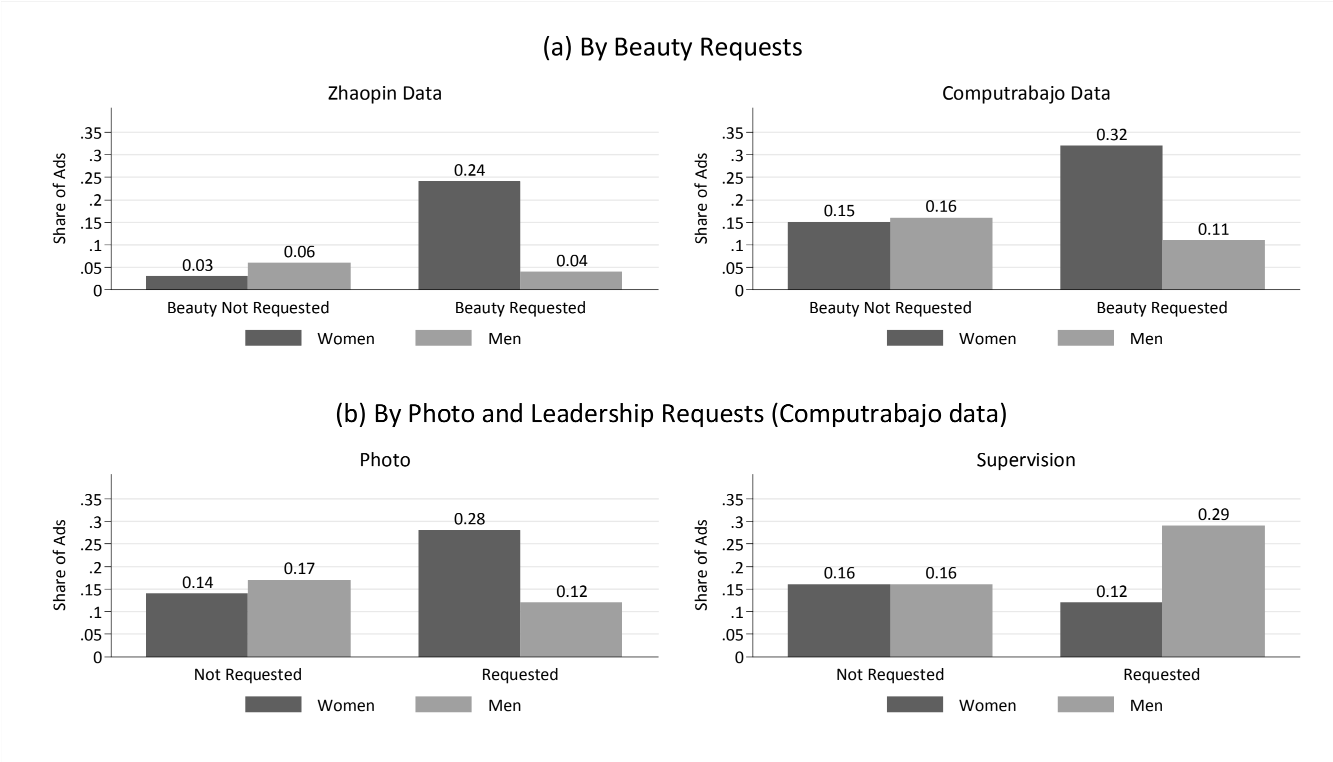

Figure 8: Share of Ads Requesting Women and Men, by Ads' Marital Status Requests, Computrabajo Data

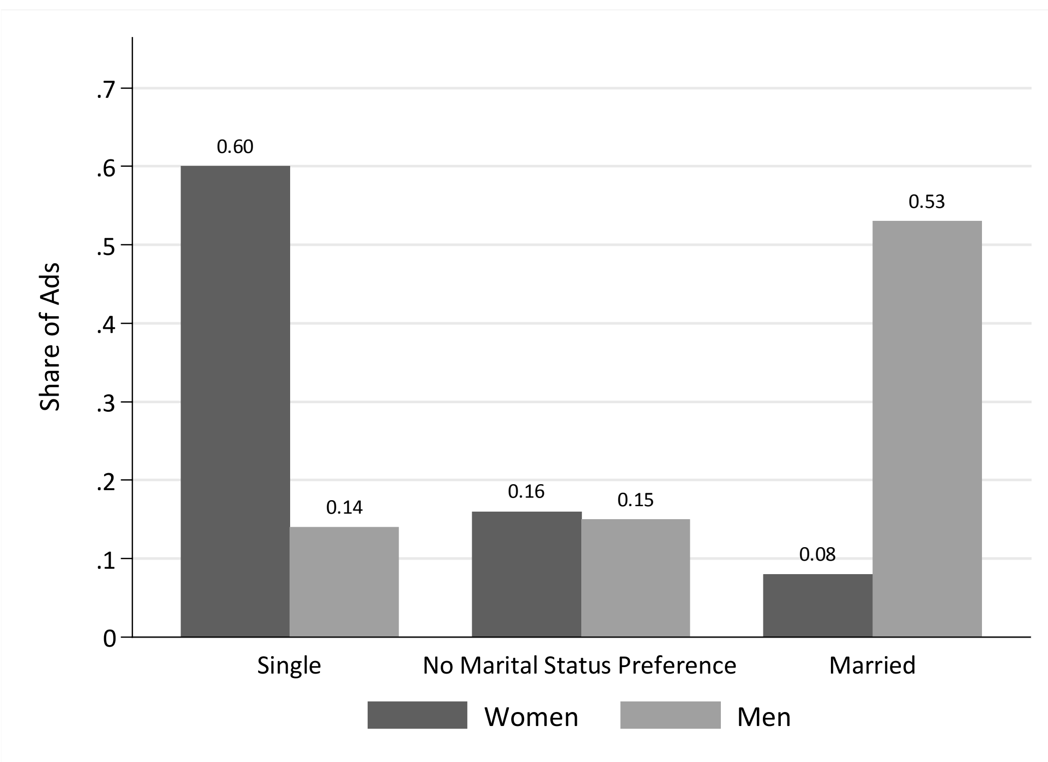

Figure 9: Percent of Men and Women Married or Cohabiting, by Age

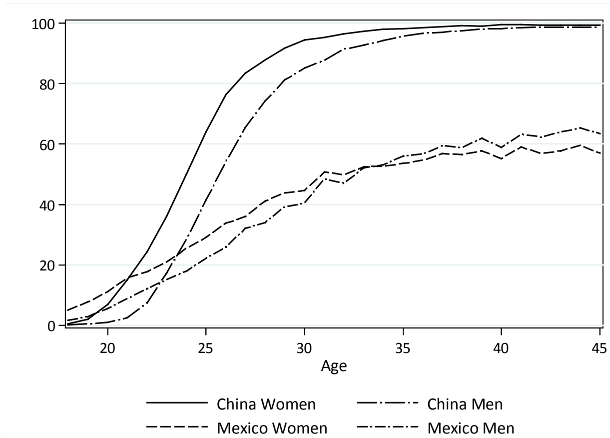

Figure 10: Percent of Women with a Child under Age 6, by Age

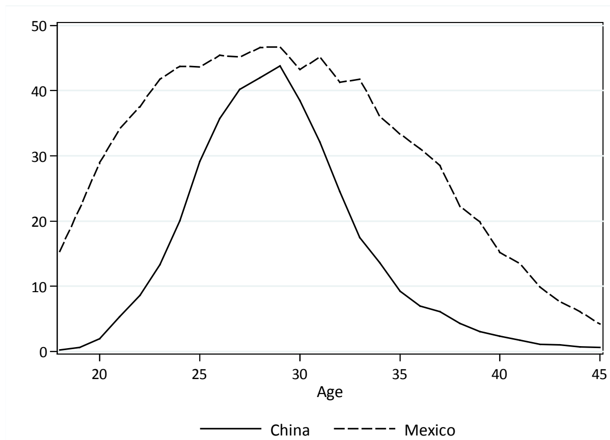

Figure 11: Percent of Women Employed, by Age

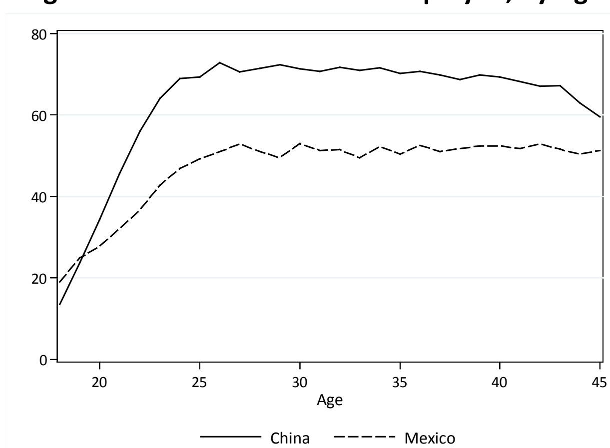

Source for Figures 9-11: Authors' calculations from the Chinese 2005 and Mexican 2010 Census. Samples are restricted to urban residents (defined as areas with more than 100,000 inhabitants in Mexico and as city residents in China).

Figure 12: Share of Ads Requesting Women and Beauty by Age, Zhaopin Data

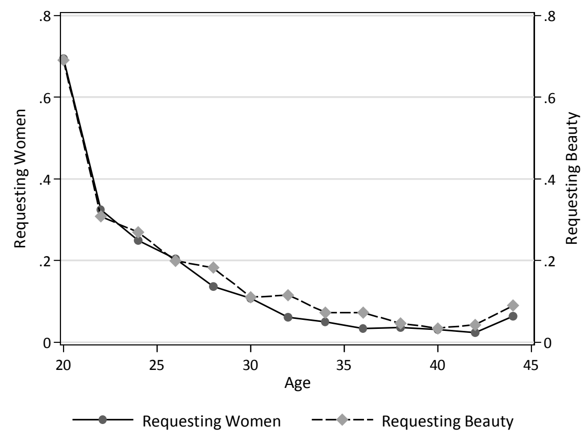

Figure 13: Share of Ads Requesting Women, Beauty and Photo by Age, Computrabajo Data

Figure 14: Share of Ads Requesting Men and Supervision by Age, Computrabajo Data

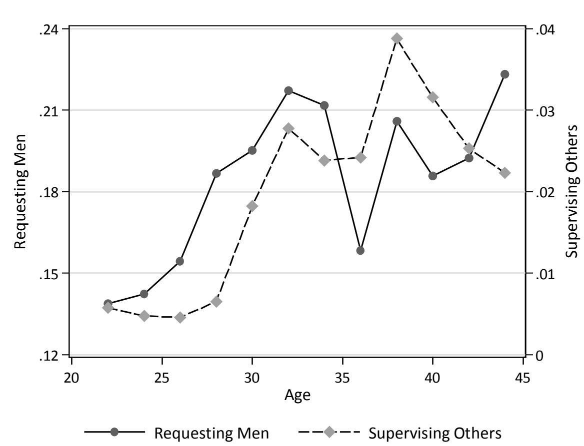

Figure 15: Mean Number of Ad Renewals by Requested Age, Zhaopin Data

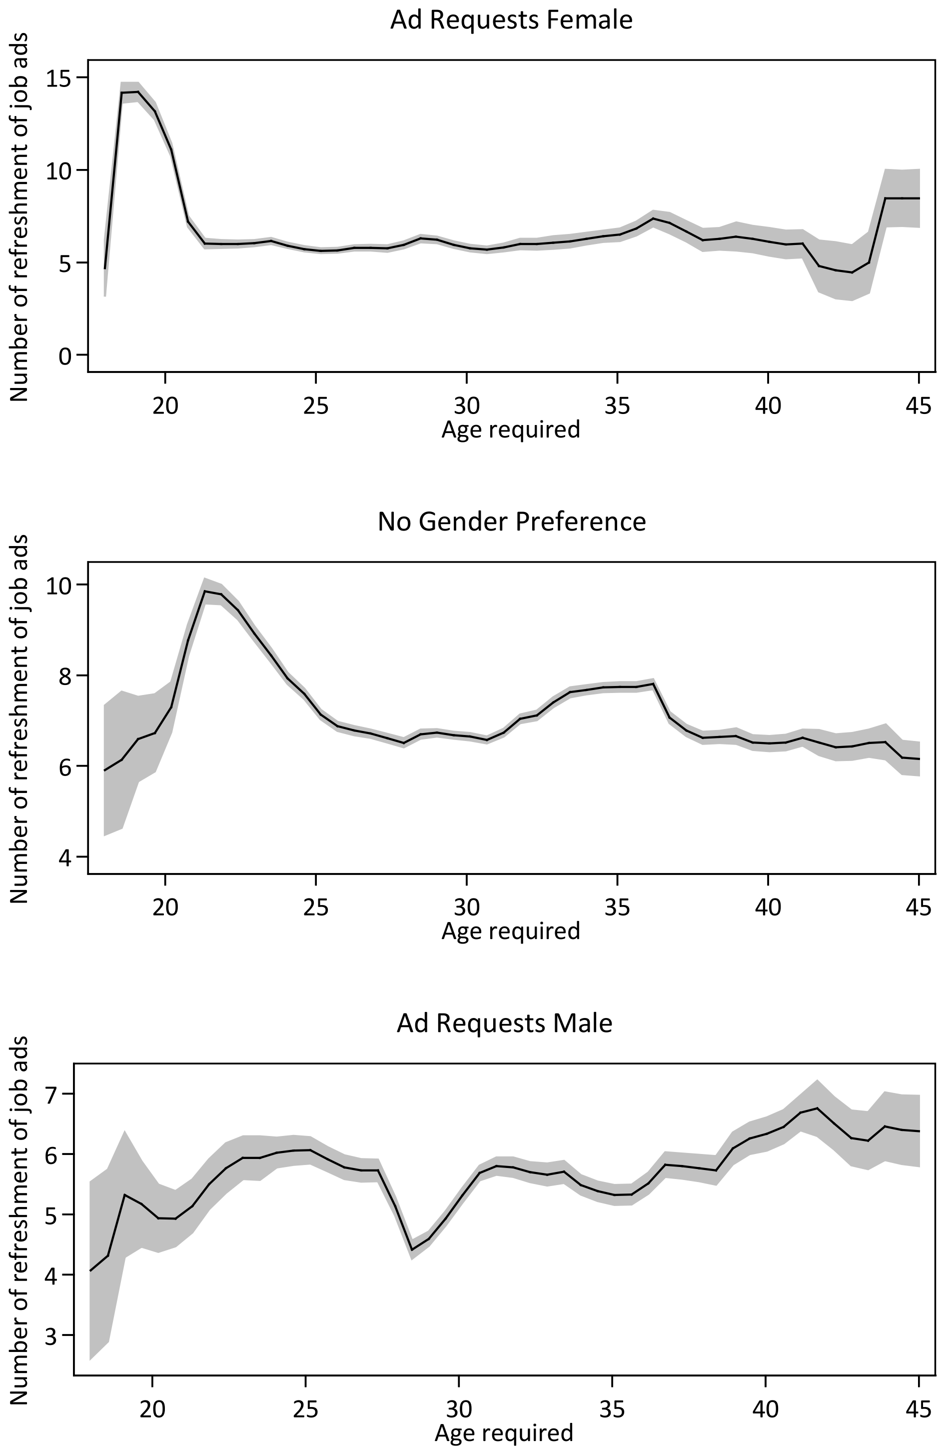

Note: gray areas show the 95% confidence intervals. Solid lines are estimated using an epanechnikov kernel-weighted local polynomial.

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	0.185	0.055	0.160
Prefer female?	0.303	0.199	0.050	0.160
Education requirements				
Junior middle school or less	0.604	0.493	0.129	0.429
High school or Tech school	0.333			
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	0.639
1-3 years	n/a	0.412	0.400	0.336
4-5 years	n/a	0.061	0.237	0.022
More than 5 years	n/a	0.017	0.158	0.002
Age requirements				
No age restrictions	0.000	0.482	0.758	0.214
Ad specifies a minimum age	1.000	0.504	0.169	0.768
Ad specifies a maximum age	0.771	0.442	0.201	0.742
Mean age, when specified (years)	27.63	28.89	30.50	30.64
Wages				
Wage not specified	0.000	0.584	0.836	0.725
Mean wage, when advertised	1810	2556	4279	7640
Number of positions advertised				
Unspecified	0.000	0.055	0.481	n/a
Mean number, when specified	8.691	1.794	3.258	n/a
Ad requests beauty?	n/a	n/a	0.077	0.040
Photo required with application	n/a	n/a	n/a	0.109
Job requires supervising others	n/a	n/a	n/a	0.016
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 ads	141,188	39,727	1,051,038	90,487

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: Effects of Jobs' Skill Requirements 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)
High//Tech School	-0.0595***	-0.0142	-0.0030	-0.0119	-0.0603***	-0.0257***	-0.0104	-0.0173***
	(0.0164)	(0.0149)	(0.0225)	(0.0128)	(0.0095)	(0.0058)	(0.0084)	(0.0040)
College	-0.1896***	-0.1175***	-0.1015***	-0.1254***	-0.1079***	-0.0622***	-0.0657***	-0.0842***
	(0.0230)	(0.0187)	(0.0260)	(0.0241)	(0.0256)	(0.0125)	(0.0164)	(0.0228)
Log (offered wage)				0.0956				0.0741
				(0.1155)				(0.0837)
Fixed effects	district, occupation	occ, firm	occ*firm	occ*firm	district, occupation	occ, firm	occ*firm	occ*firm
number of groups	64	8,897	27,633	27,633	64	8,897	27,633	27,633
N	141,188	141,188	141,188	141,188	141,188	141,188	141,188	141,188
Adjusted R^2	0.063	0.387	0.520	0.522	0.040	0.558	0.639	0.641

*** p<0.01, ** p<0.05, * p<0.1. Standard errors (in parentheses) are clustered by occupation. All regressions control for number of positions advertised.

TABLE 3: Effects of Jobs' Skill Requirements 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)
Some Post-Secondary Education	-0.1175***	-0.1050***	-0.0780***	-0.0708***	-0.0309***	-0.0000	-0.0083	0.0298
	(0.0096)	(0.0073)	(0.0117)	(0.0226)	(0.0083)	(0.0073)	(0.0130)	(0.0312)
College	-0.1803***	-0.1444***	-0.1230***	-0.1321**	-0.1068***	-0.0148	-0.0231	0.0207
	(0.0195)	(0.0145)	(0.0188)	(0.0611)	(0.0188)	(0.0149)	(0.0230)	(0.0670)
Experience required	-0.0263***	-0.0220***	-0.0161	0.0014	-0.0040	-0.0030	-0.0015	0.0022
(years)	(0.0060)	(0.0073)	(0.0108)	(0.0194)	(0.0027)	(0.0027)	(0.0050)	(0.0132)
Log (offered wage)				-0.1622***				-0.0843**
				(0.0533)				(0.0380)
Fixed effects	occupation	occ, firm	occ*firm	occ*firm	occupation	occ, firm	occ*firm	occ*firm
number of groups	36	6,716	20,618	10,682	36	6,716	20,618	10,682
N	39,727	39,727	39,727	16,540	39,727	39,727	39,727	16,540
Adjusted R^2	0.102	0.230	0.301	0.336	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 control for log firm size and firm ownership type.

All regressions control for number of positions advertised and an indicator for missing experience requirement.

TABLE 4: Effects of Jobs' Skill Requirements 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)
Some Post-Secondary education	-0.0752***	-0.0687***	-0.0610***	-0.0216**	-0.0666***	-0.0429***	-0.0266***	-0.0243*
	(0.0050)	(0.0039)	(0.0049)	(0.0088)	(0.0045)	(0.0033)	(0.0045)	(0.0143)
University	-0.1020***	-0.0929***	-0.0824***	-0.0213**	-0.1136***	-0.0591***	-0.0378***	-0.0486***
	(0.0056)	(0.0047)	(0.0058)	(0.0106)	(0.0049)	(0.0035)	(0.0048)	(0.0137)
Experience requirement								
Years required	-0.0024***	-0.0018***	-0.0017***	-0.0002	0.0037***	0.0034***	0.0028***	0.0021
	(0.0004)	(0.0003)	(0.0004)	(0.0012)	(0.0004)	(0.0003)	(0.0003)	(0.0015)
Log (offered wage)				-0.0419***				0.0065
				(0.0058)				(0.0071)
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	73,706	258,685	63,342	116	73,706	258,685	63,342
N	1,051,038	1,051,038	1,051,038	172,790	1,051,038	1,051,038	1,051,038	172,790
Adjusted R^2	0.078	0.331	0.562	0.623	0.058	0.386	0.562	0.619

*** 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 ownership. All regressions control for number of positions advertised and an indicator for missing experience requirement.

TABLE 5: Effects of Jobs' Skill Requirements 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)
Some Post-Secondary education	0.0527***	0.0113	0.0023	-0.0142	0.0087	0.0093	0.0056	0.0003
	(0.0117)	(0.00780)	(0.0079)	(0.0149)	(0.0097)	(0.0067)	(0.0063)	(0.0117)
University	-0.0874***	-0.1109***	-0.1113***	-0.0875***	-0.0445***	-0.0093	-0.0101	0.0301**
	(0.0180)	(0.0123)	(0.0133)	(0.0177)	(0.0117)	(0.0063)	(0.0067)	(0.0130)
Experience requirement								
Years required	0.0077*	-0.0007	0.0009	0.0099*	-0.0165***	-0.0135***	-0.0131***	-0.0120**
	(0.0041)	(0.0029)	(0.0030)	(0.0050)	(0.0034)	(0.0022)	(0.0024)	(0.0057)
Log (offered wage)				-0.0283***				-0.0492***
				(0.0102)				(0.0114)
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	2,383	6,973	3,051	441	2,383	6,973	3,051
N	90,487	90,487	90,487	24,876	90,487	90,487	90,487	24,876
Adjusted R^2	0.066	0.197	0.244	0.256	0.074	0.327	0.358	0.317

*** p<0.01, ** p<0.05, * p<0.1. Standard errors (in parentheses) are clustered by occ*state. Regressions without firm fixed effects include controls for log firm size and firm ownership type.

All regressions include an indicator for missing experience requirement.

TABLE 6: Effects of Desired Age on Firms' Gender Preferences (P^M-P^F), XMZYJS DATA

	(1)	(2)	(3)	(4)
Age 25-29	0.1117***	0.1337***	0.1035***	0.0885***
	(0.0203)	(0.0178)	(0.0268)	(0.0253)
Age 30-34	0.2304***	0.2243***	0.1795***	0.1550***
	(0.0336)	(0.0355)	(0.0408)	(0.0387)
Age 35+	0.2362***	0.2440***	0.2206**	0.1975**
	(0.0624)	(0.0748)	(0.0942)	(0.0893)
Job Requires High//Tech School	0.0381	-0.0066	0.0081	-0.0089
	(0.0265)	(0.0269)	(0.0336)	(0.0327)
Job Requires College	-0.0091	-0.0620	-0.0312	-0.0756
	(0.0480)	(0.0466)	(0.0693)	(0.0691)
Log (offered wage)				0.2170***
				(0.0470)
Fixed effects	district, occupation	occ, firm	occ*firm	occ*firm
number of groups	63	7,590	22,173	22,173
N	108,875	108,875	108,875	108,875
Adjusted R^2	0.211	0.288	0.335	0.338

Notes: See Table 2

TABLE 7: Effects of Desired Age on Firms' Gender Preferences (P^M-P^F), XMRC DATA

	(1)	(2)	(3)	(4)
Age 25-29	0.2270***	0.2805***	0.2457***	0.2299**
	(0.0248)	(0.0313)	(0.0622)	(0.0905)
Age 30-34	0.4724***	0.5398***	0.5064***	0.4580***
	(0.0440)	(0.0524)	(0.0941)	(0.1206)
Age 35+	0.5533***	0.6435***	0.6362***	0.5754***
	(0.0466)	(0.0516)	(0.1025)	(0.1480)
Job requires some postsecondary education	-0.0992**	-0.1282***	-0.1275**	-0.1474
	(0.0419)	(0.0399)	(0.0611)	(0.0917)
Job requires university	-0.0098	-0.1055*	-0.1512	-0.1577
	(0.0602)	(0.0563)	(0.1069)	(0.2260)
Experience required	0.0218***	0.0222***	0.0249**	0.0216
(years)	(0.0040)	(0.0057)	(0.0094)	(0.0286)
Log (offered wage)				0.1622*
				(0.0901)
Fixed effects	occupation	occ, firm	occ*firm	occ*firm
number of groups	36	4,561	10,563	6,056
N	17,021	17,021	17,021	8,549
Adjusted R^2	0.263	0.335	0.386	0.410

Notes: see Table 3.

TABLE 8: Effects of Desired Age on Firms' Gender Preferences (P^M-P^F), ZHAOPIN DATA

	(1)	(2)	(3)	(4)
Age 25-29	0.1529***	0.1627***	0.1080***	0.0250
	(0.0226)	(0.0220)	(0.0377)	(0.1026)
Age 30-34	0.2374***	0.2644***	0.1959***	0.1144
	(0.0251)	(0.0280)	(0.0443)	(0.1286)
Age 35+	0.2718***	0.3056***	0.2647***	0.1568
	(0.0275)	(0.0265)	(0.0410)	(0.1167)
Job requires some postsecondary education	-0.0284**	-0.0940***	-0.0869***	-0.1075***
	(0.0139)	(0.0164)	(0.0222)	(0.0385)
Job requires university	-0.0039	-0.0825***	-0.0621**	-0.0559
	(0.0172)	(0.0182)	(0.0257)	(0.0482)
Experience required	0.0116***	0.0119***	0.0078***	0.0176**
(years)	(0.0015)	(0.0015)	(0.0023)	(0.0073)
Log (offered wage)				-0.0476
				(0.0458)
Fixed effects	occ, ind, province	occ, province, firm	occ*firm, province	occ*firm, province
Number of groups	116	24,428	50,048	11,789
N	134,100	134,100	134,100	27,812
Adjusted R^2	0.188	0.455	0.702	0.734

Notes: see Table 4

TABLE 9: Effects of Desired Age on Firms' Gender Preferences (P^M-P^F), COMPUTRABAJO DATA

	(1)	(2)	(3)	(4)
Age 25-29	0.0441***	0.0548***	0.0634***	0.0868**
	(0.0136)	(0.0150)	(0.0164)	(0.0345)
Age 30-34	0.1303***	0.1358***	0.1420***	0.1556***
	(0.0147)	(0.0175)	(0.0180)	(0.0326)
Age 35+	0.2267***	0.2507***	0.2480***	0.2782***
	(0.0161)	(0.0186)	(0.0187)	(0.0336)
Job requires some postsecondary education	-0.0542**	-0.0483**	-0.0506**	-0.0601***
	(0.0240)	(0.0223)	(0.0237)	(0.0196)
Job requires university	-0.0685***	-0.0627***	-0.0643***	-0.1165***
	(0.0199)	(0.0226)	(0.0246)	(0.0357)
Experience required	0.0312***	0.0276***	0.0300***	0.0403***
(years)	(0.0064)	(0.0052)	(0.0057)	(0.0135)
Log (offered wage)				0.0099
				(0.0266)
Fixed effects	occ*state	occ*state, firm	occ*firm, state	occ*firm, state
Number of groups	425	2,054	5,774	2,570
N	65,516	65,516	65,516	18,943
Adjusted R^2	0.123	0.182	0.236	0.261

Notes: see Table 5

TABLE 10
Effects of Mean Beauty, Photo and Supervision Requests on the Magnitude of an Occupation's Age Twist, Zhaopin and Computrabajo Data

	(1)	(2)	(3)
	Zhaopin Data	**Computrabajo Data**	**Computrabajo Data**
Occupation's mean beauty demand	0.7875***	1.5560	
	(0.2114)	(1.0819)	
Occupation's mean photo demand		2.0534**	
		(0.8149)	
Occupation's mean beauty or photo demand			2.0779***
			(0.6461)
Occupation's mean supervision demand		4.2368*	4.0820**
		(1.9751)	(1.8248)
N	36	14	14
Adjusted R^2	.269	0.381	0.412

*** p<0.01, ** p<0.05, * p<0.1. Occupations with under 100 observations were excluded, and occupations were weighted by the number of ads.

Appendix: For Online Publication

Figure A1: Share of Ads that are Age and Gender Targeted, by the Job's Posted Wage

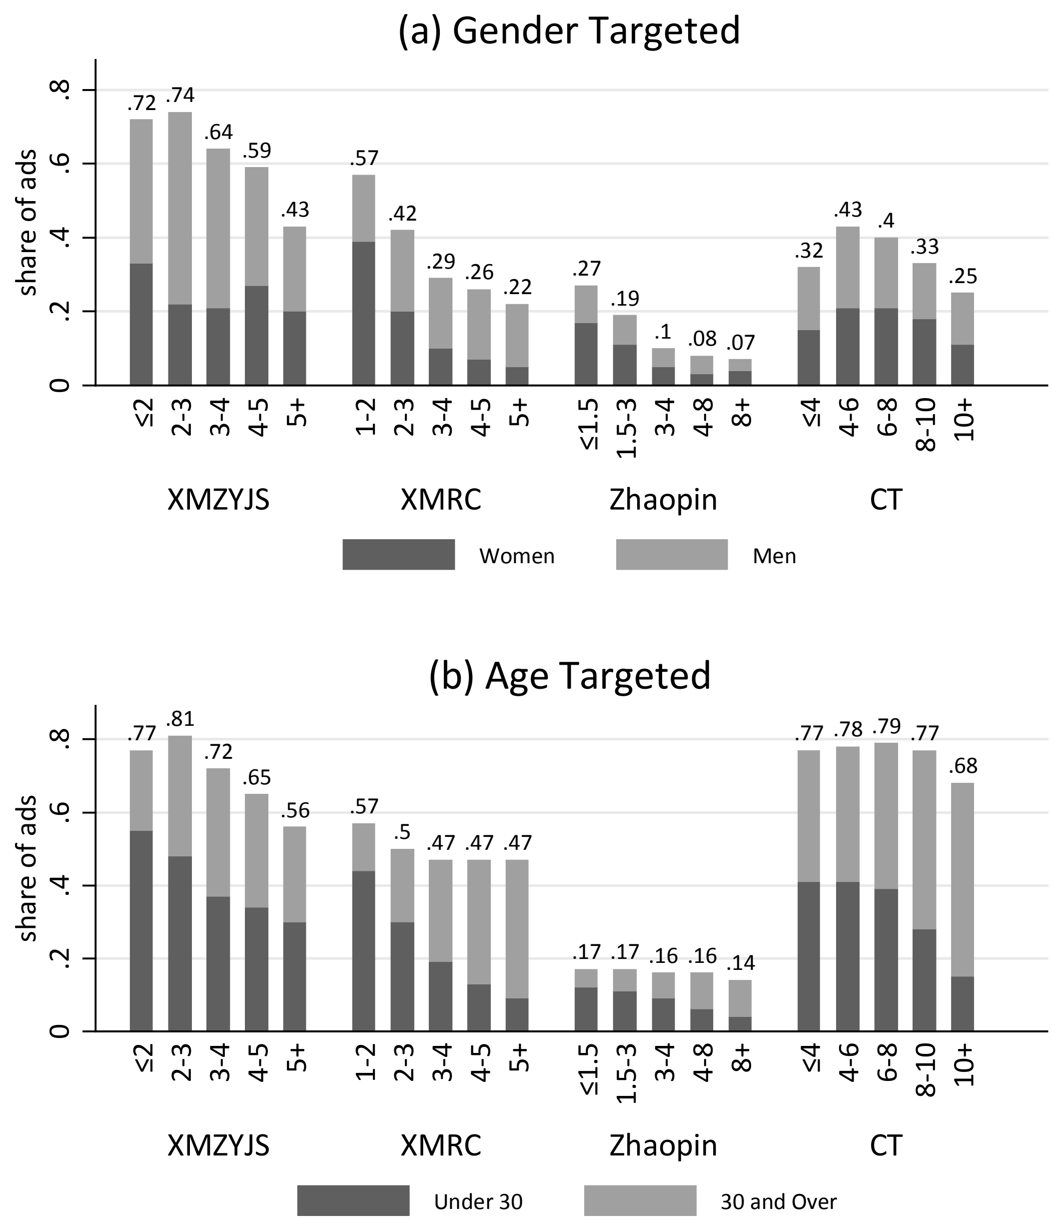

Notes: Wage bins are labeled in thousands of RMB or Mexican pesos per month. CT refers to the Computrabajo data.

Figure A2: Share of Ads that are Age and Gender Targeted, by the Job's Experience Requirement

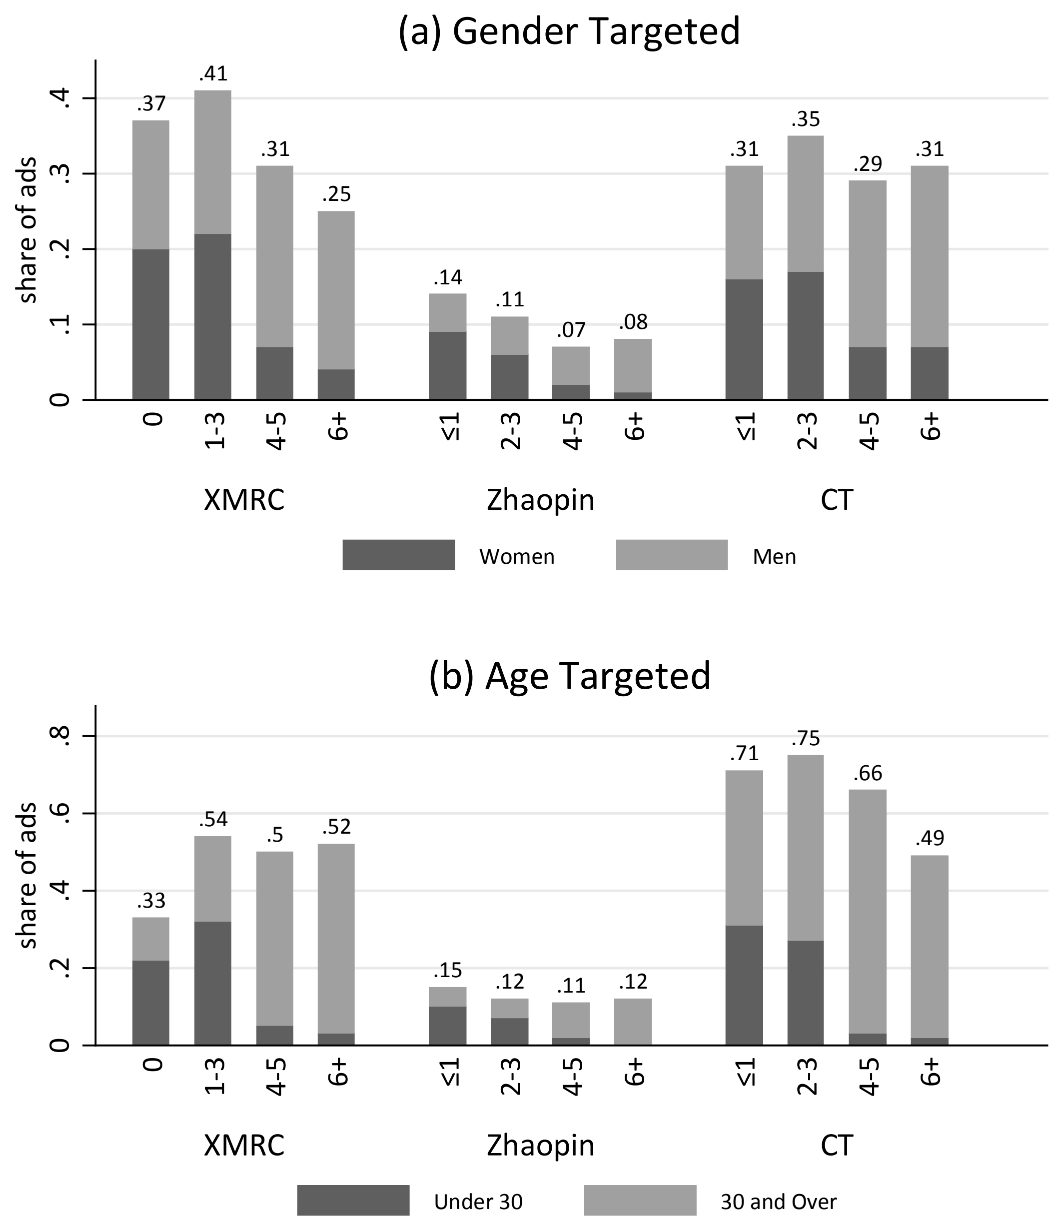

Notes: Experience bins are labeled in years. CT refers to the Computrabajo data.

TABLE A1
Ten Most Frequent Beauty Requests, Zhaopin and Computrabajo Data

A. Zhaopin Data

Rank	Percentage	Chinese text	Translation
1	30.05	形象气质佳	good image and temperament
2	9.54	五官端正	has regular facial features
3	7.35	形象良好	good image
4	5.76	品貌端正	well-shaped figure and decorous/straight appearance
5	4.99	形象好	good image
6	4.56	形象气质	image and temperament
7	4.22	形象佳	good image
8	4.11	形象好，气质佳	good image and temperament
9	3.68	相貌端正	good appearance
10	2.49	形象气质良好	good image and temperament
Others	23.25		
Total	100.00		

B. Computrabajo Data

Rank	Percentage	Spanish text	Translation
1	50.70	Excelente presentación	Great appearance
2	24.26	Buena presentación	Good appearance
3	6.92	Buena presencia	Good appearance (presence)
4	6.53	Excelente presencia	Great appearance (presence)
5	4.17	Buena imagen	Good image
6	2.46	Sin tatuajes	No tattoos
7	1.57	Talla	Size (dress size)
8	0.81	Estatura mínima	Minimum height
9	0.59	Excelente estado de salud	Great health
10	0.36	Condición física	Physical condition (fitness level)
Others	1.62		
Total	100.00		

TABLE A2
Effects of Age and Skill Requirements on the Direction of Firms' Gender Preferences (P^M-P^F) with Beauty Controls, Zhaopin data

	(1)	(2)	(3)	(4)
Age requested				
25-29	0.1366***	0.1436***	0.0936**	0.0059
	(0.0203)	(0.0209)	(0.0384)	(0.1093)
30-34	0.2075***	0.2295***	0.1677***	0.0736
	(0.0222)	(0.0269)	(0.0468)	(0.1448)
35+	0.2353***	0.2615***	0.2259***	0.0911
	(0.0237)	(0.0246)	(0.0424)	(0.1352)
Ad requests beauty?	-0.1963***	-0.2004***	-0.1833***	-0.2491***
	(0.0126)	(0.0124)	(0.0252)	(0.0893)
Education requirement				
Some postsecondary	-0.0316**	-0.0958***	-0.0891***	-0.0924***
	(0.0131)	(0.0155)	(0.0215)	(0.0324)
University	-0.0076	-0.0850***	-0.0665***	-0.0440
	(0.0164)	(0.0175)	(0.0253)	(0.0484)
Experience requirement				
Years required	0.0110***	0.0113***	0.0077***	0.0179**
	(0.0015)	(0.0015)	(0.0023)	(0.0072)
Log (offered wage)				-0.0468
				(0.0444)
Fixed Effects	occ, ind, province	occ, province, firm	occ*firm, province	occ*firm, province
Number of groups	116	24,428	50,048	11,789
N	134,100	134,100	134,100	27,812
Adjusted R^2	0.203	0.464	0.706	0.741

*** p<0.01, ** p<0.05, * p<0.1. Standard errors (in parentheses) are clustered at the occupation level. Specifications are identical to Table 8, with beauty controls added.

.

TABLE A3
Effects of Age and Skill Requirements on the Direction of Firms' Gender Preferences (P^M-P^F) with Beauty, Marital Status and Supervision Controls, Computrabajo Data

	(1)	(2)	(3)	(4)
Age requested				
25-29	0.0250*	0.0372**	0.0488***	0.0703**
	(0.0150)	(0.0161)	(0.0176)	(0.0319)
30-34	0.0920***	0.0998***	0.1095***	0.1227***
	(0.0147)	(0.0170)	(0.0176)	(0.0296)
35+	0.1744***	0.2007***	0.2029***	0.2348***
	(0.0153)	(0.0176)	(0.0180)	(0.0331)
Ad requests beauty?	-0.1637***	-0.1592***	-0.1699***	-0.1831***
	(0.0229)	(0.0233)	(0.0233)	(0.0290)
Ad requests photo?	-0.1466***	-0.1566***	-0.1534***	-0.1565***
	(0.0187)	(0.0196)	(0.0202)	(0.0318)
Ad requests married?	0.3434***	0.3149***	0.2772***	0.2498***
	(0.0189)	(0.0198)	(0.0215)	(0.0535)
Ad requests single?	-0.3420***	-0.3638***	-0.3261***	-0.2602***
	(0.0249)	(0.0255)	(0.0266)	(0.0453)
Job requires supervising others?	0.1054***	0.0847***	0.0780**	-0.0301
	(0.0278)	(0.0273)	(0.0310)	(0.0567)
Education requirement				
Some postsecondary	-0.0504**	-0.0441**	-0.0468**	-0.0543**
	(0.0233)	(0.0216)	(0.0231)	(0.0211)
University	-0.0655***	-0.0584***	-0.0605**	-0.1164***
	(0.0196)	(0.0222)	(0.0244)	(0.0338)
Experience required	0.0291***	0.0258***	0.0293***	0.0424***
(years)	(0.0060)	(0.0051)	(0.00550)	(0.0135)
Log (offered wage)				0.0146
				(0.0247)
Fixed Effects	occ*state	occ*state, firm	occ*firm, state	occ*firm, state
Number of groups	425	2,054	5,774	2,570
N	65,516	65,516	65,516	18,943
Adjusted R^2	0.146	0.203	0.253	0.275

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

Specifications are identical to Table 9, with beauty, photo, supervision, and marital status controls added.