

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

December 30, 2015

JEL codes: J16, J63, J71

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

When permitted by law, employers sometimes state the preferred age and sex of their employees in job ads. We study this practice using data from one Mexican and three Chinese job boards, showing that it is widely used to request both genders and is especially prevalent in jobs with low skill requirements. Other patterns identified in Kuhn and Shen's (KS, 2013) analysis of a single Chinese job board are also confirmed. We also document a new stylized fact we call the *age twist* in gender profiling: firms' explicit gender requests shift dramatically away from women and towards men when firms are seeking older (as opposed to younger) workers. While some of this twist can be attributed to employers' age-dependent requests for (female) beauty and (male) leadership, the timing of the shift also suggests that young women's movement into childbearing also plays a role.

^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 Wang Yanan Institute for Studies in Economics (WISE), Xiamen University; Research School of Economics, 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 participants at UCSB, Berkeley, Princeton, Columbia, Arizona, Yale, IZA, Cornell, and Binghamton Universities. Participants in the Bergen/Stavanger workshop on Labor Markets, Families and Children; the Trans-Pacific Labor Seminar (TPLS) in Sydney; and discrimination conferences at the Bern University of Applied Sciences and the University of Naples also contributed helpful feedback. 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. Most of the world's workers, however, work in labor markets where employers' sex and age preferences are common features of job ads. Given the widespread interest in gender and age differentials in labor market outcomes around the world, it is perhaps surprising that this practice, which we refer to as age- and gender profiling of job ads, has been subject to very little quantitative analysis.

Recently, however, Kuhn and Shen (KS, 2013) studied patterns of gender profiling in a large sample of job ads on Zhaopin.com, a Chinese Internet job board that serves a highly skilled clientele. Their study documented four main facts about gender profiling in that dataset. First, about one in ten ads (10.5 percent) explicitly requested a worker of a particular gender. Second, 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. Third, 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. Fourth, 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 derive a simple screening cost model can account 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 exclude a set of potential applicants on the basis of an ascriptive demographic attribute without examining those applicants' credentials.

This paper has two main goals. The first is to see whether KS's main findings extend to three new data sets, two from China and one from Mexico. Importantly, the two new Chinese data sets serve workers with a much more representative level of skills than KS's data. In

particular, these datasets represent the entire private-sector labor market of a medium-sized Chinese city and contain a much larger share of unskilled jobs. The Mexican data lets us see whether KS's main results extend to a nation with different economic conditions, labor market institutions and culture. Our second goal is to extend KS's analysis of gender profiling to age profiling, and to the interaction between age and gender profiling in job ads.

Our main findings are fourfold. First, gender profiling is much more common in the three new datasets (at 32, 38 and 72 percent of job ads) than in Zhaopin (10.5%). Age profiling is also much more common in the new datasets, at 43, 72 and 77 percent, versus 13% in Zhaopin. Second, the *negative skill-targeting relationship* is confirmed, for both age- and gender-profiling: jobs with higher skill requirements are less likely to request employees of a specific age or sex. This relationship holds for three different measures of job skill (required education, required experience and the level of the advertised wage), and is present both in the raw data and within cells defined by the interaction of firm and occupation fixed effects. It also holds across data sets: the job boards serving the least-skilled workers have the highest rates of age- and gender-profiling.

Third, the job-specificity of employers' gender profiling choices is confirmed in all three of the new datasets, and the symmetry of employers' requests for men and women is strikingly confirmed in two of them. In the least-skilled of the new job boards, however, requests for men outnumber requests for women by about 30 percent. Finally, we establish a new stylized fact concerning the interaction between age and gender profiling in all four of our data sets.¹ This relationship, which we call the *age twist* in employers' gender requests, is a strong shift in the direction of gender profiling, away from women and towards men, as the desired age of the worker rises from 18 to 45. While some of this twist can be attributed to employers' age-dependent requests for (female) beauty and (male) leadership, most of it occurs within occupation*firm cells; it is present even in the sample of ads that do not explicitly request beauty or leadership and in occupations where such job requirements seem unlikely. The timing of the twist in China coincides quite closely to the ages at which women typically become mothers in that country.

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. Section 7 concludes.

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. Data

The two new 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 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 paper's website.⁴ The most important restrictions affecting all datasets involve dropping observations with missing education requirements and

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.

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 <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.

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.

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.

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.

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.

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 both a minimum age and maximum age (our main summary indicator for whether the ad is age targeted) decline monotonically as we move from the least-skilled job board (XMZYJS) to the most skilled (Zhaopin).⁷ 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 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.⁸ 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.⁹ 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

⁷ This definition allows us to define the requested age as the midpoint of those requirements for all age-targeted ads. All the main results are unchanged if we define age targeting as expressing any sort of age preference.

⁸ Computrabajo has job boards in 20 Spanish-speaking countries. We picked Mexico because it had the most ads.

⁹ Identical ads appearing on multiple dates are included only once. See Delgado Helleseeter (2013) for more details.

10.5 percent).¹⁰ 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.¹¹ 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.¹² Another unique feature of the 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 suggestive way with age and gender profiling Mexico.¹³

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. 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

¹⁰ 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.

¹¹ 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) were the most common.

¹² 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 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.

¹³ 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.

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.

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, 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, Appendix Figures A1 and A2 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 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 Figure 1 are primarily a consequence of the type of work that is done: perhaps some types of work are highly gendered, and others not, and the latter just happen to be more skilled. Columns 2 and 6 add firm fixed effects to this specification: perhaps the skill-targeting relationship results mostly from a pattern where the firms that abstain from age- and gender-targeting (such as, for example, foreign-owned firms) disproportionately happen to employ skilled workers for reasons unrelated to skill *per se*.

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.

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

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 are 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.

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.

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 2 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

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. Neumark, Burn and Button (2015) find considerably less age discrimination in hiring against older women than older men in a large resume correspondence study.

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.

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 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

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.

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.²² We explore possible causes for these age- or experience-dependent gender preferences in Section 6.

5. Understanding the Negative Skill-Targeting Relationship

As we have noted, one interpretation of the robust negative skill-targeting effect in all four of our data sets is as a direct consequence of job skill requirements: 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 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).²³

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 seems unlikely that they can account for the negative skill-targeting relationship that is present in all our data sources.²⁴ 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?

²² Because they allow employers to specify up to five of its 36 occupations for each ad, the XMRC data let us test whether this result is robust to even finer, firm-specific occupation controls. Specifically, instead of using only the first occupation listed (which yields a roughly similar number of categories for XMRC as are available in the other datasets), we added a fixed effect for every one of the 1,701 combinations of these five categories observed in the data, interacted with a firm fixed effect (thus allowing each firm to 'gender' its detailed occupations in its own way). The estimated age twist was almost identical.

²³ A fourth possibility suggested by KS's model is that the idiosyncratic variation in applicant quality (σ_e) 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.

²⁴ 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.

A final possibility suggested by the KS model is that σ_v , the variance *across jobs* in women's (or older workers') comparative advantage, might be greater in less-skilled than more-skilled jobs. In other words, perhaps age and 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 might be more age- and gender-specialized than other jobs.²⁵ 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, 'care less' about the age and gender of people they interact with. In their paper, KS devise a simple test between the variance (σ_v)-based and skill (θ)-based explanations of the negative skill-targeting relationship, based on the fact 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).²⁶

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. While we cannot definitively rule out this explanation with the data at our disposal, 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.²⁷

Summing up, application processing costs and employer stigma do not appear to be likely explanations of the negative skill-targeting relationship, and available evidence suggests that the phenomenon is not caused by a tendency for men and women to be more 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. This leaves either a direct effect of skill requirements on the value of identifying the best job candidate, or a higher level of tightness (fewer applicants per job) to skilled than unskilled jobs. While we suspect that both of these factors play a role, we note that any tendency for labor market tightness to increase with skill would need to be quite universal to explain our results, which are present for both age- and gender-targeting in four different datasets.

²⁵ 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.

²⁶ Details are provided on the paper's website. See Table S-1 of the supplementary tables.

²⁷ 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 declines at the same rate, then firms are probably not substituting internal gender- and age-based filtering of applications for profiling as skill levels rise. Unfortunately, we do not have data on which employees are hired in any of our four datasets.

6. Understanding the Age Twist in Firms' Gender Preferences

In this Section we present six facts that shed some light on possible causes of the age twist in advertised gender requests. We begin by examining the role of gendered employer preferences for beauty and supervision, which are requested quite frequently in Chinese and Mexican job ads, then turn to the possible role of employer preferences for marital status and parenthood, which are rarely mentioned in job ads.

a. Beauty, Supervision and the Age Twist

As noted, it is relatively common for job ads in China and Mexico to request that applicants be physically attractive. If these requests –which could reflect employer, customer and/or co-worker preferences-- are strongly correlated with employers' age and sex requests, the demand for beauty might account for some of the age twist.²⁸ Similarly, if employers (or workers) prefer to have male supervisors, an increasing supervisory component of jobs with age could also contribute to the age twist.²⁹ To explore these ideas, we first ask whether employers' beauty and supervision requests covary with desired age and sex in a way that could account for the age twist. Our focus is on the Zhaopin and Computrabajo datasets, which are the only ones where we have information on employers' requests for beauty and supervision.

Fact 1: Employers' requests for beauty and supervisory duties are strongly associated with their age and gender requests.

This is illustrated in panel (a) of Figure 3 for the Zhaopin data. There, *close to half (45.8 percent) of job ads aimed at women under the age of 30 explicitly request that the applicant be physically attractive.* This number falls to 25.5 percent among women over 30, 9.8 percent for men under 30, and to 4.8 percent for men over 30. Similar but less dramatic trends are present for beauty and photo requests in the Computrabajo data (panels (b) and (c)). Together, these patterns suggest that employers' requests for attractive young women could account for a substantial share of the age twist. Panel (d) of Figure 3 looks at whether the ad says the position entails supervising other workers. As one might expect, supervisory duties increase strongly with age in both 'male' and 'female' job ads. The increase is however much stronger among jobs requesting men, suggesting that preferences to have men in supervisory positions could also account for some of the age twist in our data.

Fact 2: The age twist is positive in almost all occupations, though its magnitude varies considerably. It is especially high in administration, customer service and management occupations.

²⁸ Hamermesh and Biddle (1994), Biddle and Hamermesh (1998) and Mobius and Rosenblat (2006). study the wage and hiring premia associated with beauty in developed-country labor markets.

²⁹ 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.

Additional information on possible causes of the age twist might come from knowing which occupations have the largest twists. To that end, Figures 4-7 show the size of the age twist by occupation in all four of our data sets.³⁰ In all those figures, the age twist is defined as the difference in employers' 'preferences towards men' ($P^M - P^F$) between ads requesting workers under versus over 30 years of age.

Figures 4-7 show that the age twist is positive in almost all occupations in all datasets, including many occupations with no obvious role for beauty (security staff, software, construction, accounting) or supervision (manual workers, logistics workers, manufacturing labor). That said, the age twist does vary substantially in magnitude across occupations. In particular, three occupations stand out as having consistently large twists: administration (which includes receptionists, executive assistants and other office assistants), customer service and management all fall into the top five occupations in three or more of the datasets.³¹ Figures 4-7 also flag two large occupation groups—sales manufacturing workers—both of which have about average-sized twists. Straightforward shift-share calculations show that these five occupations account for more than half of the age twist in all four datasets.

Fact 3: Occupations with high beauty demands and a large supervision component have bigger age twists. Together these two factors can account for 37 percent of the cross-occupation variance in the twist.

What features distinguish the high-twist occupations from the others in Figures 4-7? Inspired by the patterns in Figure 3, 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, i.e. the numbers displayed in Figures 6 and 7. 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*. In the Computrabajo data, the same regression also shows a positive effect, though it is only significant at the ten percent level (column 2).³² The Computrabajo data also allow us to add the occupation's supervision demand to this regression; the results in column 3 show that occupations with frequent leadership duties tend

³⁰ To keep the figures manageable and more comparable across datasets, and to increase the precision of the within-occupation twist measures used later in this Section, Figures 8-11 use occupation groups that are more aggregated than the ones used in our regression analyses. Specifically, compared to the regressions, the number of categories is reduced from 58 to 19, 36 to 23, 39 to 26 and 14 to 13 in XMZYJS, XMRC, Zhaopin and Computrabajo respectively

³¹ Management has a low rank in XMRC, however. Occupations with titles that suggest a customer contact component are indicated by asterisks (*), but do not show consistently above-average age twists.

³² To economize on degrees of freedom, Table 10 uses an aggregate indicator of beauty demand—whether the ad requested beauty or a photo. If these are entered separately, both have positive coefficients but the photo indicator has the stronger and more significant effect.

to ‘turn away’ from seeking female applicants as the worker’s age rises ($p < .10$). Adding this supervision control increases the size and significance of the beauty effect as well, which becomes significant at the 5 percent level. Quantitatively, the R^2 s in Table 11 indicate that gendered preferences for beauty and supervision can account for about 37 percent of the cross-occupational variation in the age twist.

Fact 4: Requests for Beauty and Supervision account for relatively little (8-15 percent) of the age twist across individual job ads.

To what extent can the age pattern of requests for beauty and supervision in Figure 3 explain the covariation between age and the direction of gender requests across ads in our data? Table 11 explores this question by adding controls for ad-level beauty and supervision requests to the column 3 regressions in Tables 8 and 9. While beauty and supervision demands are highly predictive of gender requests, controlling for them does not attenuate the estimated age coefficients very much. Specifically, these controls can account for $(.2647 - .2259) / .2647 = 14.7\%$ of the age twist in the Zhaopin data, and $(.2480 - .2289) / .2480 = 7.7\%$ of the age twist in the Computrabajo data. A similar message emerges from columns (3) and (6) of Table 11, which exclude all ads requesting beauty (or a photo or supervision in the Computrabajo data) from the estimation sample. While the estimated age twist attenuates somewhat further, it remains large in magnitude and highly statistically significant. These facts, together with Figure 4-7’s evidence that the twist is observed even in occupations where it seems unlikely that beauty or supervision play substantial roles, lead us to consider other possible explanations of the age twist.

b. Marriage, Children and the Age Twist

A sizable literature on the male marriage wage premium suggests that firms value marriage among men (Korenman and Neumark 1991, Ahituv and Lerman 2007, Loughran et al. 2009, Petersen et al 2011). And 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). Thus, the fact that both genders are entering marriage and parenthood between age 18 and 45 could also contribute to the age twist.

Fact 5: Explicit employer preferences for workers’ marital status are rare, but are highly correlated with desired gender.

Figure 8 summarizes employers’ marital status requests (only available in the Computrabajo data) by requested gender. While rare overall (comprising only 3.4 percent of all Computrabajo ads), ads requesting a specific marital status exhibit a strong gender pattern. Specifically, among ads requesting men, requests for married applicants outweigh requests for single applicants by a ratio of about 6 to 1. In ads requesting women this pattern is reversed,

with requests for single applicants outweighing those for married applicants by more than four to one. Together these patterns suggest that the movement of men and women into marriage (and parenthood) during the early career could also account for the age twist, especially if firms have preferences for these employee characteristics that are not always expressed in ads.³³ The only additional evidence we can offer on this possibility relates to the precise timing of the age twist.

Fact 6: In China, the age twist in advertised gender preferences is highly concentrated between the ages of 26 and 30. This corresponds quite closely to the ages at which most Chinese women are having their first (and only) child.

The three dashed lines in Figures 9 and 10 show the share of gendered job ads that request women in each of our three Chinese datasets in two-year age bins. For comparability, each dataset has been reweighted to represent the education distribution of working-age urban *hukou* Chinese women in the 2005 one-percent Census microdata file. After reweighting, the structure of the age twist is quite comparable across our three Chinese datasets (especially the two more-skilled ones, XMRC and XMZYJS), and is steepest between age 26 and 30. Indeed, most of the age twist in these two datasets is concentrated into those four years, with about 70 percent of gendered job ads for 26-year-olds requesting women compared to about 35 percent at age 30. This surprisingly concentrated timing of the Chinese age twist offers some hope that we might be able to isolate its proximate causes by looking at the timing of other events in Chinese women's lives.³⁴

To that end, Figure 9 compares the timing of this rapid shift in the mix of job ads to the ages at which urban Chinese women move into marriage in the 2005 Census sample described above. In Figure 9, the blue (rightmost) solid line shows the share of urban Chinese women who are single as a function of age in two-year bins. Interestingly, urban Chinese women's movement into marriage is also highly concentrated into a narrow age range, with the fraction single falling from 80 to about 30 percent between the ages of 22 and 26. This range, however, is earlier than the period when the mix of gendered job ads shifts most strongly away from women (26-30). Thus, the timing of the age twist does not line up well with a simple employer-aversion-to-married-women hypothesis: Chinese employers are still disproportionately requesting female employees (at around 70 percent of ads in XMRC and Zhaopin) at age 26,

³³ Explicit requests for workers with or without children do occur in the Computrabajo data, but are even rarer than marital status requests.

³⁴ The timing of both the age twist and women's movement into marriage and childbirth are much more gradual in Mexico than in China (see Appendix Figure A4). While these two trends are thus consistent with each other, they do not allow for the same test we can do with the Chinese data: asking whether two quite well-defined 'events' happen at the same times in a worker's life.

when a substantial majority of women are already married.³⁵ This timing problem becomes even more severe if employers avoid not just married women but women *who will become married shortly after being hired*: leading the ‘share single’ curve by two years (the leftmost, brown curve) shows the share of women who, if hired today, will be still be single two years from now. To avoid hiring such women, firms will want to start shifting their ads away from women starting even earlier—at age 20.

Figure 10, on the other hand, compares the age twist in urban China to the ages at which urban Chinese women tend to have their first (and usually only) child. Here, the curves match up better: the composition of gendered job ads shifts rapidly away from requesting women during exactly the ages when Chinese women are becoming mothers. While hardly conclusive, this suggests that a broad-based decline in job ads for women who—based on their ages—are having their first child has the potential to account for the age twist in our data.³⁶

7. Discussion

Our analysis of age- and gender-based profiling in job ads shows that this is a widely used practice in Chinese and Mexican labor markets. Consistent with the evidence in Kuhn and Shen (2013), we find that profiling is much more common at low skill levels, whether skill is measured by education requirements, experience requirements or offered wages. As noted, this pattern is consistent with a simple model of hiring from two pools in the presence of application processing costs. The pattern also suggests that skill upgrading in the Chinese and Mexican labor markets may be a potent force in reducing the prevalence of labor market profiling in those economies.

We have also identified a new empirical regularity: a dramatic ‘twist’ in the distribution of employers’ advertised gender requests away from women and towards men as workers age. Like the skill-targeting relationship, this phenomenon is robust to estimation methods and statistical controls, and is present in four data sets drawn from different countries serving very different skill levels. We show that employers are much more likely to explicitly request physically attractive applicants when they are seeking to hire young women than any other demographic group, suggesting that conscious demands for (young, female) beauty may account for some of this age twist; a similar argument applies to preferences for older male supervisors. That said, explicit, gendered employer preferences for beauty and supervision can only account for about 37 percent of the cross-occupation variance in the age twist, and no more than 15 percent of its cross-ad variation.

³⁵ Of course, an employer preference for married men could also explain the age twist. However, when we produced the equivalent to Figure 9 for men, the same timing problem is apparent..

³⁶ Notably, this lack of interest does not appear to be confined to women who are currently caring for a young child: Given the timing of first births and China’s one-child policy, Appendix Figure A3] shows that fewer than 20 percent of 34-year-old urban Chinese women still have a child under the age of 6. If care for pre-school children explains the age twist, we should see an increase in employer interest in hiring women after about age 30, which is not the case.

These facts, plus the fact that the twist remains strong even when all job ads that request beauty or supervision are excluded from our sample, suggest that some additional factor also plays a role. In that regard, we provide some suggestive evidence that the rapid movement of young workers into marriage and parenthood during their 20s may also matter. While explicit requests for marital status are rare, they strongly indicate that employers ‘want men to be married and women to be single’. Also, in China where women’s shift into motherhood is highly concentrated into a narrow age window, the age twist is highly concentrated into those same ages.

While the age twist in advertised vacancies is large in magnitude, a number of *caveats* should be kept in mind when assessing its consequences for more familiar labor market outcomes like wages and employment. One such *caveat* is the fact that even in our less-skilled data sets, a substantial share of age-targeted jobs is open to both genders. This is illustrated in Table A2, which combines nongendered and gender-matched jobs to show that despite the age twist, 42 to 82 percent of job ads targeted at workers over 35 are still open (at least ‘officially’) to women. Similarly, between 45 and 71 percent of the jobs targeted at workers under 25 are at least *open to men*. Thus, workers in relatively low demand according to explicit gender-based requests (young men and older women) are by no means completely excluded from job ads targeted at their age group.


Second, not all ads are age targeted. In particular, the share of ads that are age-targeted—on which our analysis of the age twist is based—is 77.1 percent in XMZYJS, 72.4 percent in Computrabajo, 42.9 percent in XMRC, and 12.8 percent in Zhaopin. To the extent that employers are indeed willing to consider applicants of any age in non-age-targeted ads, this also raises the number of vacancies open to groups that are rarely receive explicit invitations to apply, thus limiting the impact of the profiling patterns documented here.

Bearing these caveats in mind, however, the logic of equilibrium search models suggests profiling practices that exclude specific demographic groups from a substantial share of job vacancies should lead to a variety of behavioral changes among workers. These include applying to more jobs, searching (and perhaps remaining unemployed) longer, setting a lower reservation wage, and choosing to engage in less on-the-job-search for better jobs. Thus, the age twist identified in our data could help explain a number of well known gender gaps in labor market outcomes, including the fact that wages grow considerably less rapidly with age for women than men in virtually all economies where data is available. Understanding the consequences of job market profiling for more familiar labor market outcomes like these is an important question for further research.

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.
- Delgado Helleseeter, 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.
- 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.
- Hamermesh, D. S., & Biddle, J. E. (1994). Beauty and the labor market. *American Economic Review*, 84(5), 1174-1194.
- 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).
- 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
- 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>
- Neumark, David, Ian Burn and Patrick Button. "Is It Harder for Older Workers to Find Jobs? New and Improved Evidence from a Field Experiment" NBER working paper no. 21669, October 2015.
- 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>

**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

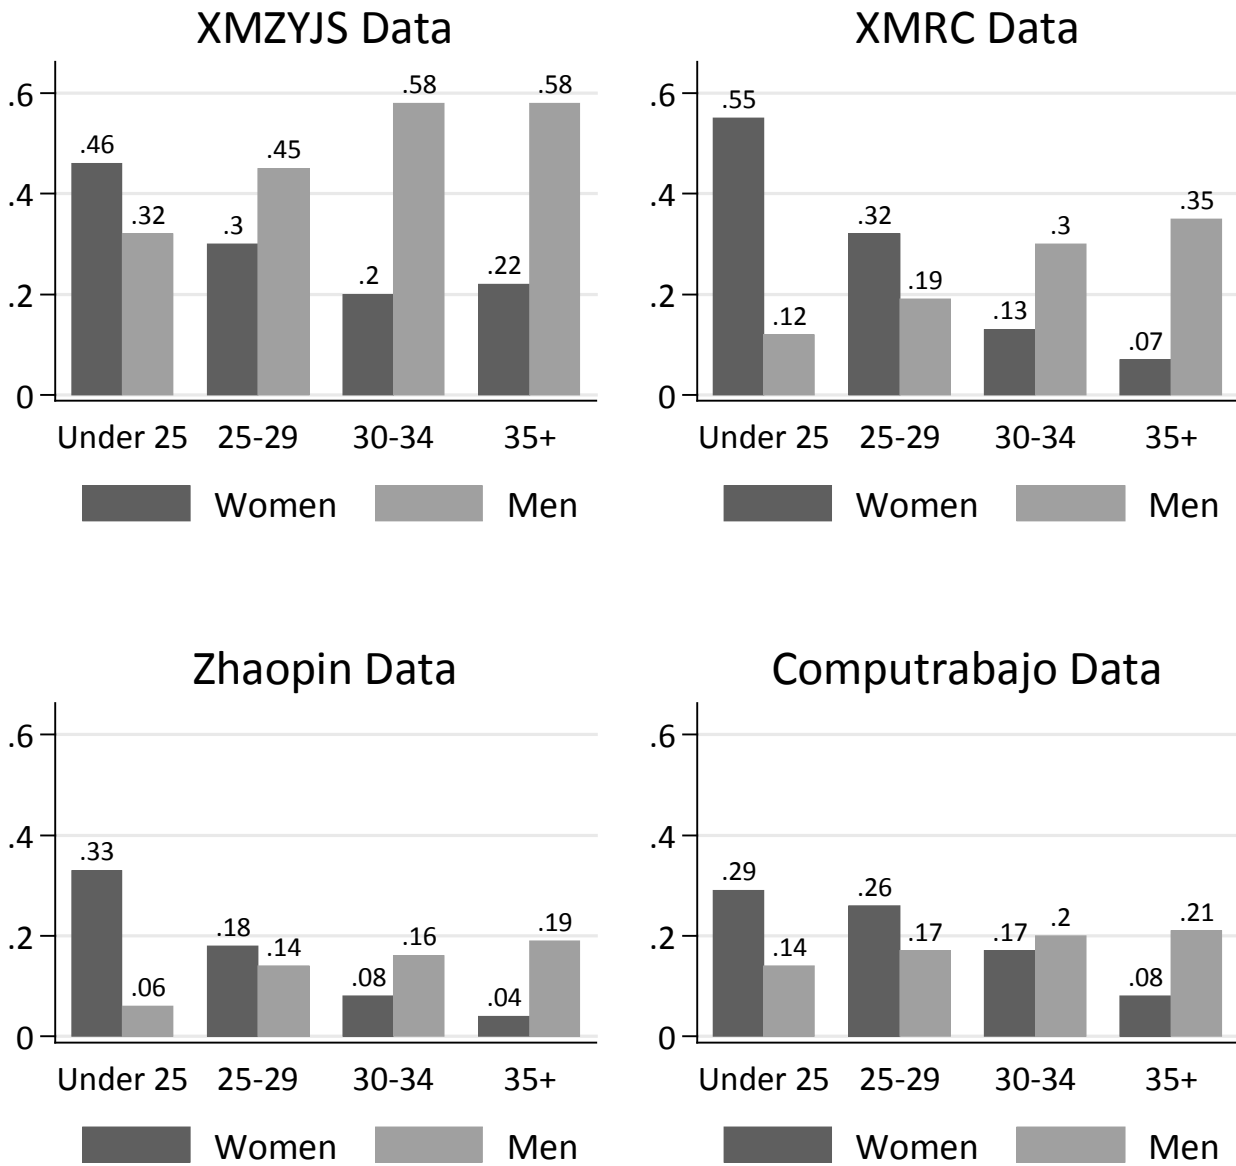


Figure 3: Share of Ads Requesting Beauty, Photo, or Supervisory Duties, by Requested Age and Sex

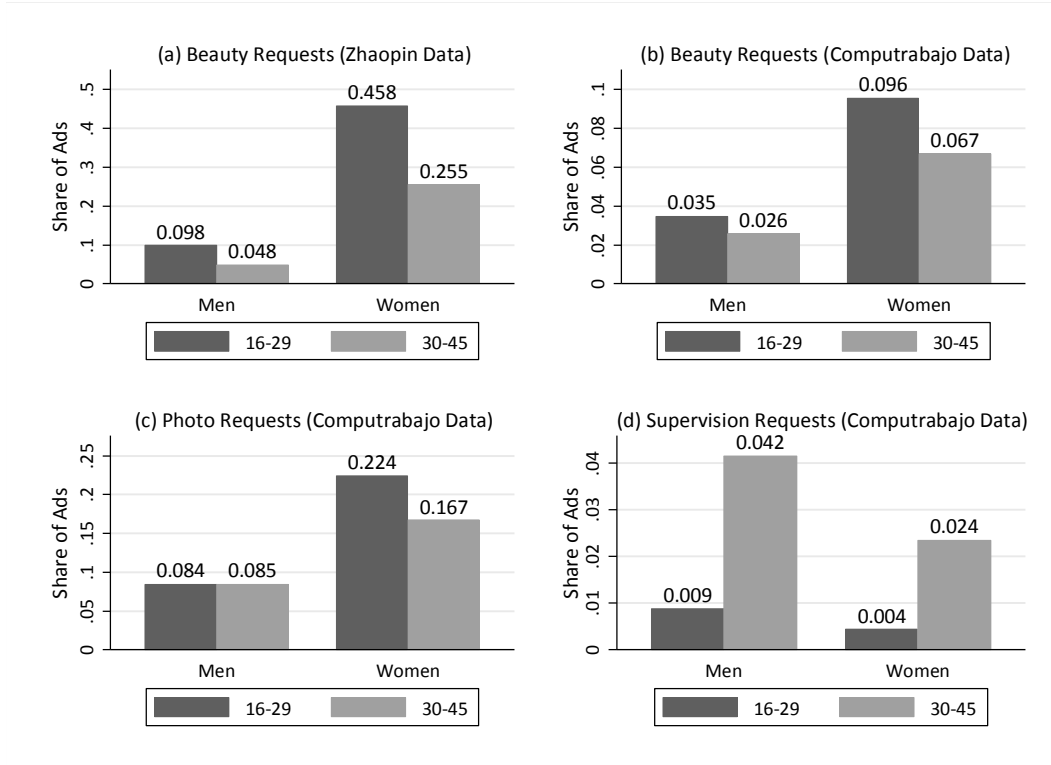


Figure 4: Age Twist by Occupation: XMZYJS Data

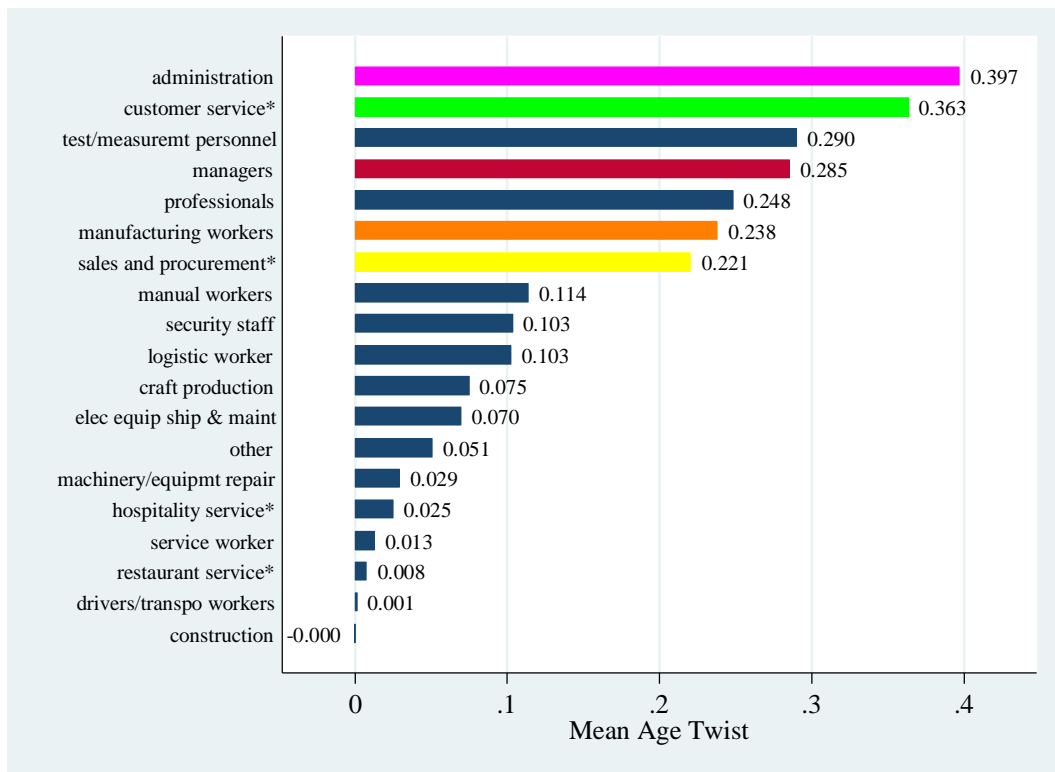


Figure 5: Age Twist by Occupation: XMRC Data

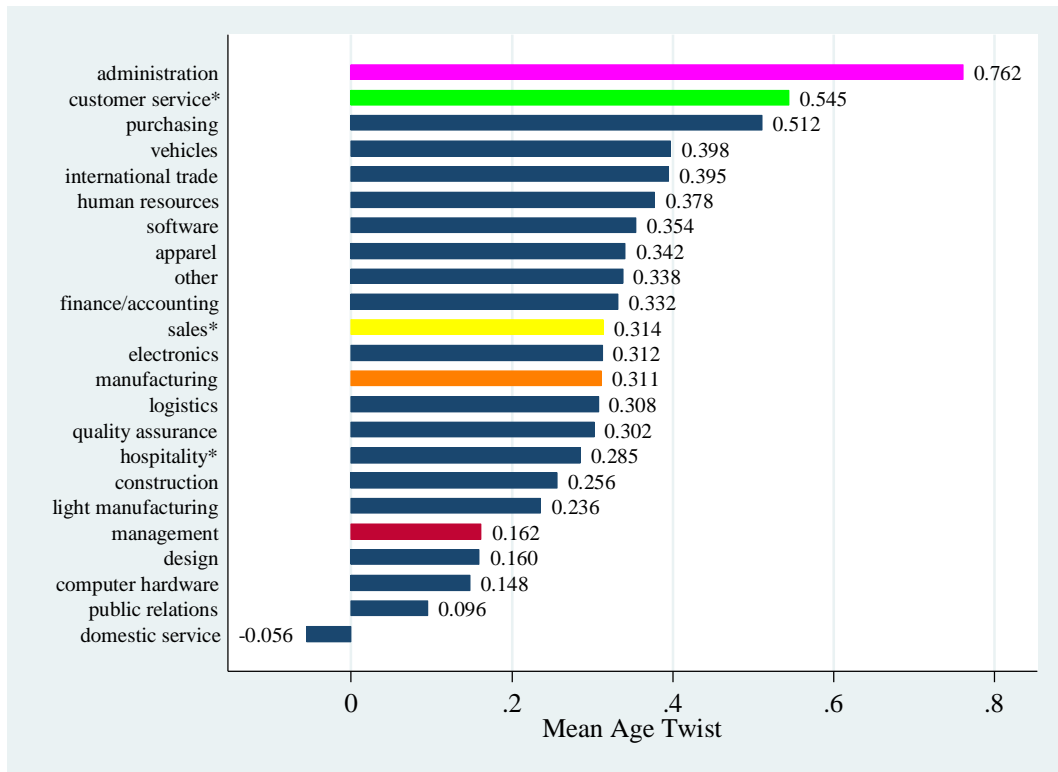


Figure 6: Age Twist by Occupation—Zhaopin Data

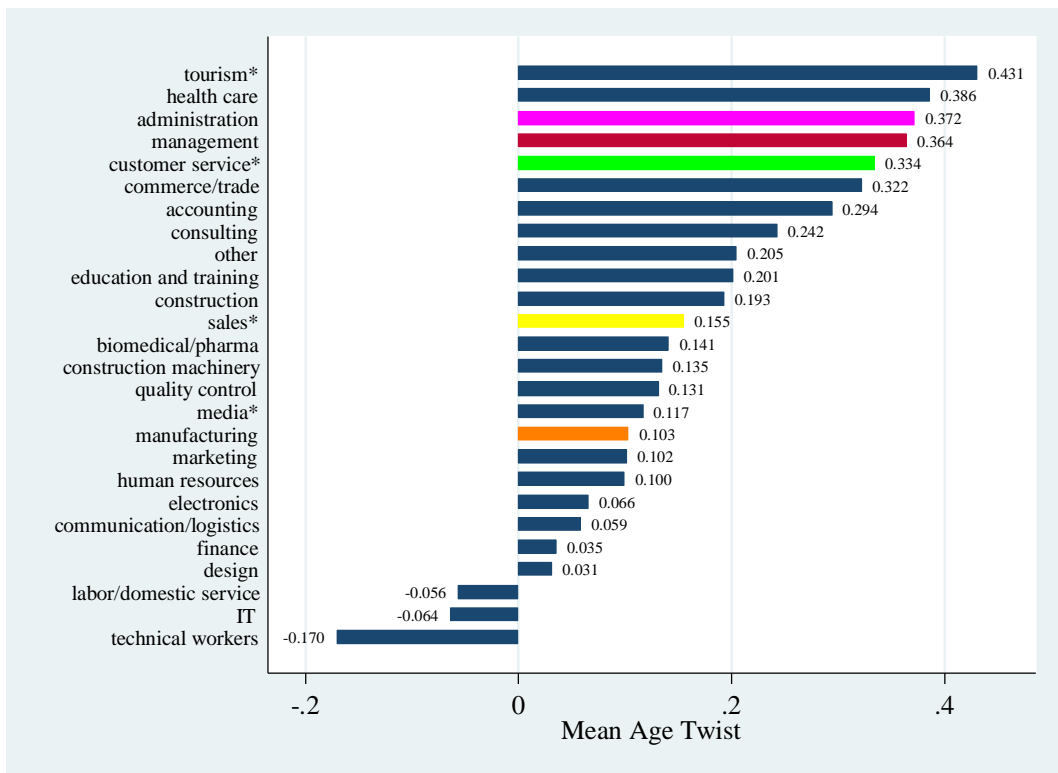
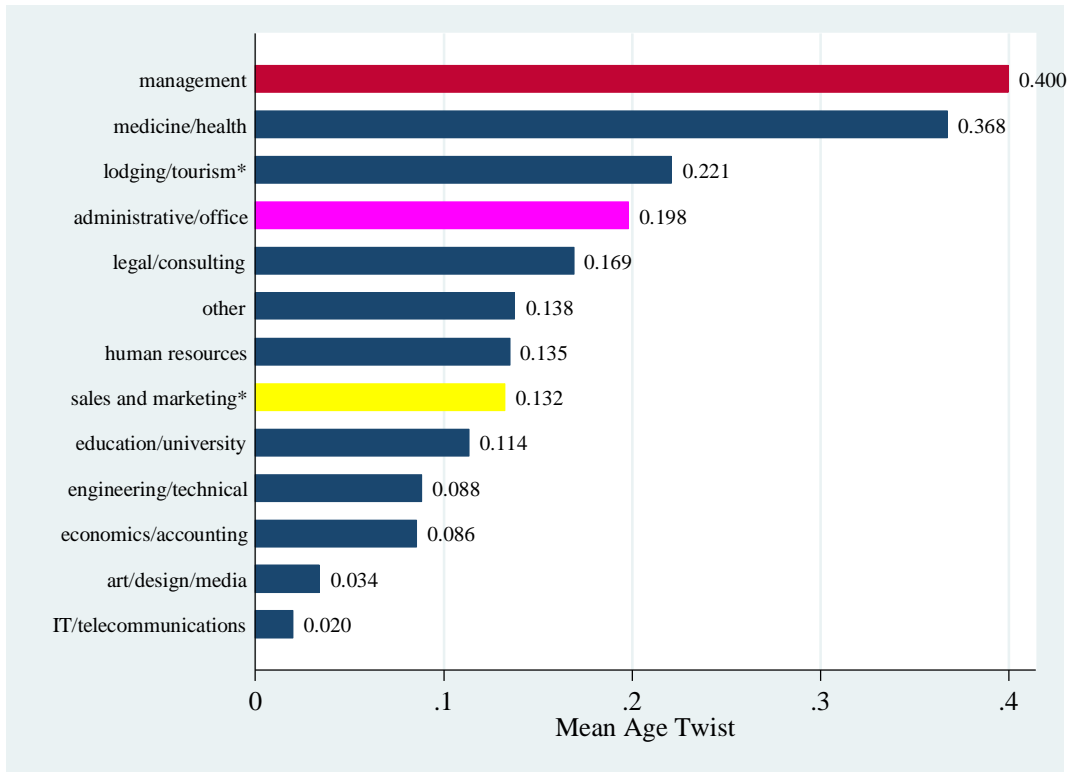


Figure 7: Age Twist by Occupation—Computrabajo Data



Notes to Figures 4-7: the age twist is the difference in employers ‘preferences towards men’ ($P^M - P^F$) between ads requesting workers under versus over 30 years of age. Asterisks (*) indicate that the occupation’s title suggests a customer contact component.

Figure 8: Share of Ads Requesting Single and Married Applicants, by Requested Age and Sex (Computrabajo Data)






Figure 9: Female Share of Gendered Job Ads versus Share of Women who are *Single*, China



Dashed Lines show the share of gendered job ads that request women in each job board. Solid lines show the share of urban Chinese women who are single at that age, or who will still be single two years later. Marital status information is for working=age urban-hukou women in the 2005 Census 1% microdata sample.

Figure 10: Female Share of Gendered Job Ads versus Share of Women who are *Childless*, China



Dashed lines show the share of gendered job ads that request women in each job board. Solid lines show the share of urban Chinese women who are childless at that age, or who will still be childless two years later.

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
Minimum and maximum age specified	0.771	0.428	0.128	0.724
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*** (0.0164)	-0.0142 (0.0149)	-0.0030 (0.0225)	-0.0119 (0.0128)	-0.0603*** (0.0095)	-0.0257*** (0.0058)	-0.0104 (0.0084)	-0.0173*** (0.0040)
College	-0.1896*** (0.0230)	-0.1175*** (0.0187)	-0.1015*** (0.0260)	-0.1254*** (0.0241)	-0.1079*** (0.0256)	-0.0622*** (0.0125)	-0.0657*** (0.0164)	-0.0842*** (0.0228)
Log (offered wage)				0.0956 (0.1155)				0.0741 (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*** (0.0096)	-0.1050*** (0.0073)	-0.0780*** (0.0117)	-0.0708*** (0.0226)	-0.0309*** (0.0083)	-0.0000 (0.0073)	-0.0083 (0.0130)	0.0298 (0.0312)
College	-0.1803*** (0.0195)	-0.1444*** (0.0145)	-0.1230*** (0.0188)	-0.1321** (0.0611)	-0.1068*** (0.0188)	-0.0148 (0.0149)	-0.0231 (0.0230)	0.0207 (0.0670)
Experience required (years)	-0.0263*** (0.0060)	-0.0220*** (0.0073)	-0.0161 (0.0108)	0.0014 (0.0194)	-0.0040 (0.0027)	-0.0030 (0.0027)	-0.0015 (0.0050)	0.0022 (0.0132)
Log (offered wage)				-0.1622*** (0.0533)				-0.0843** (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*** (0.0050)	-0.0687*** (0.0039)	-0.0610*** (0.0049)	-0.0216** (0.0088)	-0.0666*** (0.0045)	-0.0429*** (0.0033)	-0.0266*** (0.0045)	-0.0243* (0.0143)
University	-0.1020*** (0.0056)	-0.0929*** (0.0047)	-0.0824*** (0.0058)	-0.0213** (0.0106)	-0.1136*** (0.0049)	-0.0591*** (0.0035)	-0.0378*** (0.0048)	-0.0486*** (0.0137)
Experience requirement								
Years required	-0.0024*** (0.0004)	-0.0018*** (0.0003)	-0.0017*** (0.0004)	-0.0002 (0.0012)	0.0037*** (0.0004)	0.0034*** (0.0003)	0.0028*** (0.0003)	0.0021 (0.0015)
Log (offered wage)				-0.0419*** (0.0058)				0.0065 (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
<i>N</i>	1,051,038	1,051,038	1,051,038	172,790	1,051,038	1,051,038	1,051,038	172,790
<i>Adjusted R</i> 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*** (0.0117)	0.0113 (0.00780)	0.0023 (0.0079)	-0.0142 (0.0149)	0.0087 (0.0097)	0.0093 (0.0067)	0.0056 (0.0063)	0.0003 (0.0117)
University	-0.0874*** (0.0180)	-0.1109*** (0.0123)	-0.1113*** (0.0133)	-0.0875*** (0.0177)	-0.0445*** (0.0117)	-0.0093 (0.0063)	-0.0101 (0.0067)	0.0301** (0.0130)
Experience requirement								
Years required	0.0077* (0.0041)	-0.0007 (0.0029)	0.0009 (0.0030)	0.0099* (0.0050)	-0.0165*** (0.0034)	-0.0135*** (0.0022)	-0.0131*** (0.0024)	-0.0120** (0.0057)
Log (offered wage)				-0.0283*** (0.0102)				-0.0492*** (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
<i>Adjusted R²</i>	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*** (0.0203)	0.1337*** (0.0178)	0.1035*** (0.0268)	0.0885*** (0.0253)
Age 30-34	0.2304*** (0.0336)	0.2243*** (0.0355)	0.1795*** (0.0408)	0.1550*** (0.0387)
Age 35+	0.2362*** (0.0624)	0.2440*** (0.0748)	0.2206** (0.0942)	0.1975** (0.0893)
Job Requires High//Tech School	0.0381 (0.0265)	-0.0066 (0.0269)	0.0081 (0.0336)	-0.0089 (0.0327)
Job Requires College	-0.0091 (0.0480)	-0.0620 (0.0466)	-0.0312 (0.0693)	-0.0756 (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
<i>N</i>	108,875	108,875	108,875	108,875
<i>Adjusted R</i> 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*** (0.0248)	0.2805*** (0.0313)	0.2457*** (0.0622)	0.2299** (0.0905)
Age 30-34	0.4724*** (0.0440)	0.5398*** (0.0524)	0.5064*** (0.0941)	0.4580*** (0.1206)
Age 35+	0.5533*** (0.0466)	0.6435*** (0.0516)	0.6362*** (0.1025)	0.5754*** (0.1480)
Job requires some postsecondary education	-0.0992** (0.0419)	-0.1282*** (0.0399)	-0.1275** (0.0611)	-0.1474 (0.0917)
Job requires university	-0.0098 (0.0602)	-0.1055* (0.0563)	-0.1512 (0.1069)	-0.1577 (0.2260)
Experience required (years)	0.0218*** (0.0040)	0.0222*** (0.0057)	0.0249** (0.0094)	0.0216 (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
<i>N</i>	17,021	17,021	17,021	8,549
<i>Adjusted R</i> 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*** (0.0226)	0.1627*** (0.0220)	0.1080*** (0.0377)	0.0250 (0.1026)
Age 30-34	0.2374*** (0.0251)	0.2644*** (0.0280)	0.1959*** (0.0443)	0.1144 (0.1286)
Age 35+	0.2718*** (0.0275)	0.3056*** (0.0265)	0.2647*** (0.0410)	0.1568 (0.1167)
Job requires some postsecondary education	-0.0284** (0.0139)	-0.0940*** (0.0164)	-0.0869*** (0.0222)	-0.1075*** (0.0385)
Job requires university	-0.0039 (0.0172)	-0.0825*** (0.0182)	-0.0621** (0.0257)	-0.0559 (0.0482)
Experience required (years)	0.0116*** (0.0015)	0.0119*** (0.0015)	0.0078*** (0.0023)	0.0176** (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
<i>N</i>	134,100	134,100	134,100	27,812
<i>Adjusted R</i> 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*** (0.0136)	0.0548*** (0.0150)	0.0634*** (0.0164)	0.0868** (0.0345)
Age 30-34	0.1303*** (0.0147)	0.1358*** (0.0175)	0.1420*** (0.0180)	0.1556*** (0.0326)
Age 35+	0.2267*** (0.0161)	0.2507*** (0.0186)	0.2480*** (0.0187)	0.2782*** (0.0336)
Job requires some postsecondary education	-0.0542** (0.0240)	-0.0483** (0.0223)	-0.0506** (0.0237)	-0.0601*** (0.0196)
Job requires university	-0.0685*** (0.0199)	-0.0627*** (0.0226)	-0.0643*** (0.0246)	-0.1165*** (0.0357)
Experience required (years)	0.0312*** (0.0064)	0.0276*** (0.0052)	0.0300*** (0.0057)	0.0403*** (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
<i>N</i>	65,516	65,516	65,516	18,943
<i>Adjusted R</i> ²	0.123	0.182	0.236	0.261

Notes: see Table 5

TABLE 10
Effects of Mean Beauty 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
Mean beauty demand	0.8061*** (0.2034)	1.2154* (0.6511)	1.8529** (0.6408)
Mean supervision demand			3.7688* (1.7707)
<i>N</i>	26	13	13
Adjusted R^2	0.3702	0.1716	0.3728

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The age twist is the difference in preference towards men ($P^M - P^F$) between ads requesting workers 30 and older and ads requesting younger workers in the occupation. Regressions are weighted by the number of ads in the occupation. Computrabajo beauty indicator=1 if the ad requests beauty or a photo.

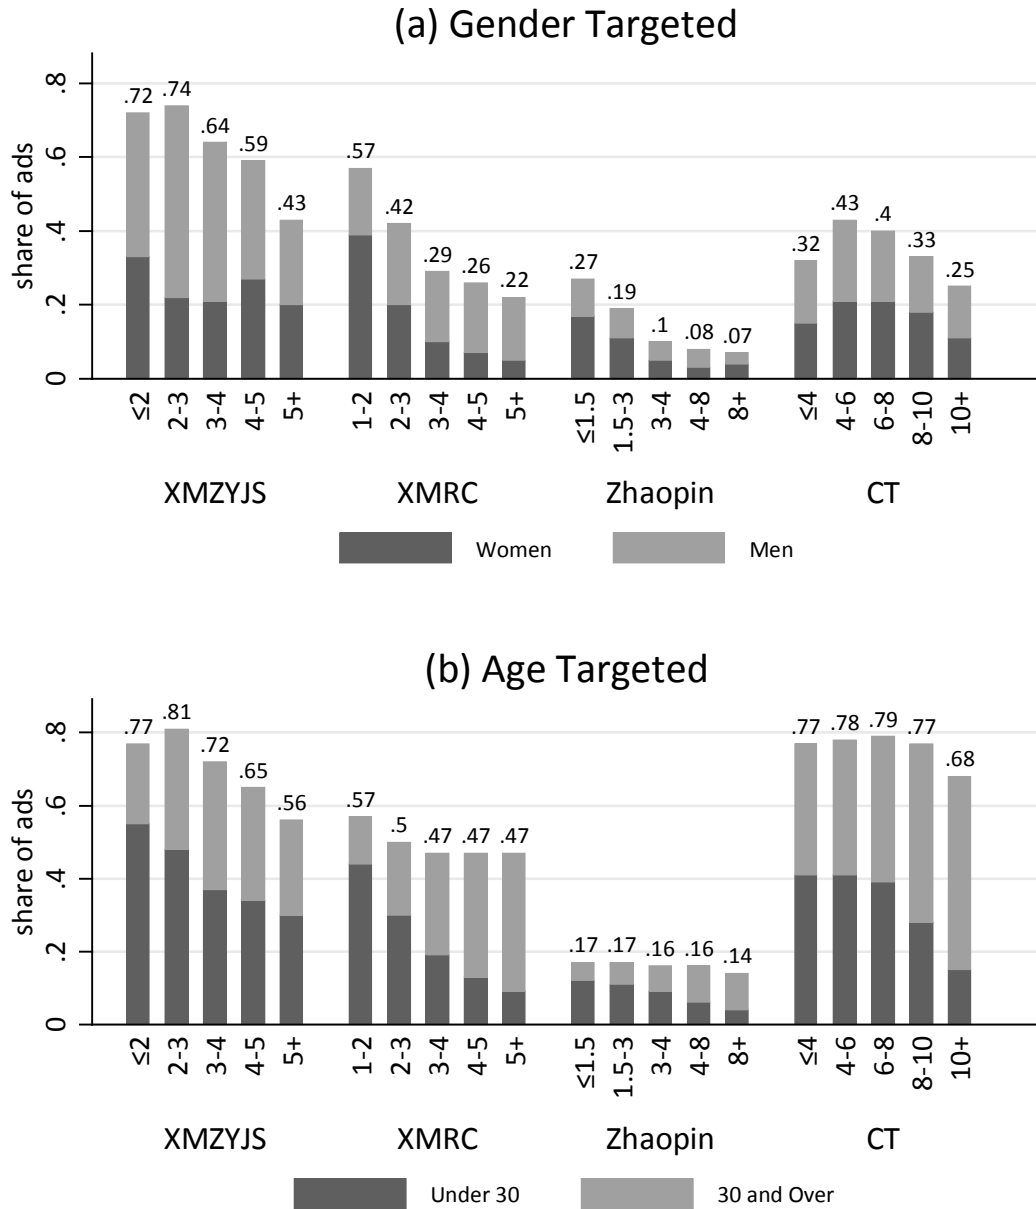
TABLE 11
Effects of Desired Age on the Direction of Firms' Gender Preferences ($P^M - P^F$), Alternative Specifications

	ZHAOPIN DATA			COMPUTRABAJO DATA		
	(1)	(2)	(3)	(4)	(5)	(6)
Age Requested						
25-29	0.1080*** (0.0377)	0.0936** (0.0384)	0.0634 (0.0520)	0.0634*** (0.0164)	0.0576*** (0.0173)	0.0441* (0.0229)
30-34	0.1959*** (0.0443)	0.1677*** (0.0468)	0.1149* (0.0659)	0.1420*** (0.0180)	0.1268*** (0.0176)	0.1045*** (0.0209)
35+	0.2647*** (0.0410)	0.2259*** (0.0424)	0.1815*** (0.0589)	0.2480*** (0.0187)	0.2289*** (0.0180)	0.1993*** (0.0200)
Beauty/supervision controls		Yes			Yes	
Restricted sample			Yes			Yes
Fixed effects	occ*firm, province	occ*firm, province	occ*firm, province	occ*firm, state	occ*firm, state	occ*firm, state
Number of groups	50,048	50,048	42,725	5,774	5,774	5,269
<i>N</i>	134,100	134,100	113,865	65,516	65,516	54,246
R^2	0.702	0.706	0.715	0.236	0.253	0.240

Restricted sample drops all ads requesting beauty from the Zhaopin data, and all ads requesting beauty, a photo or supervision in the Computrabajo data.


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



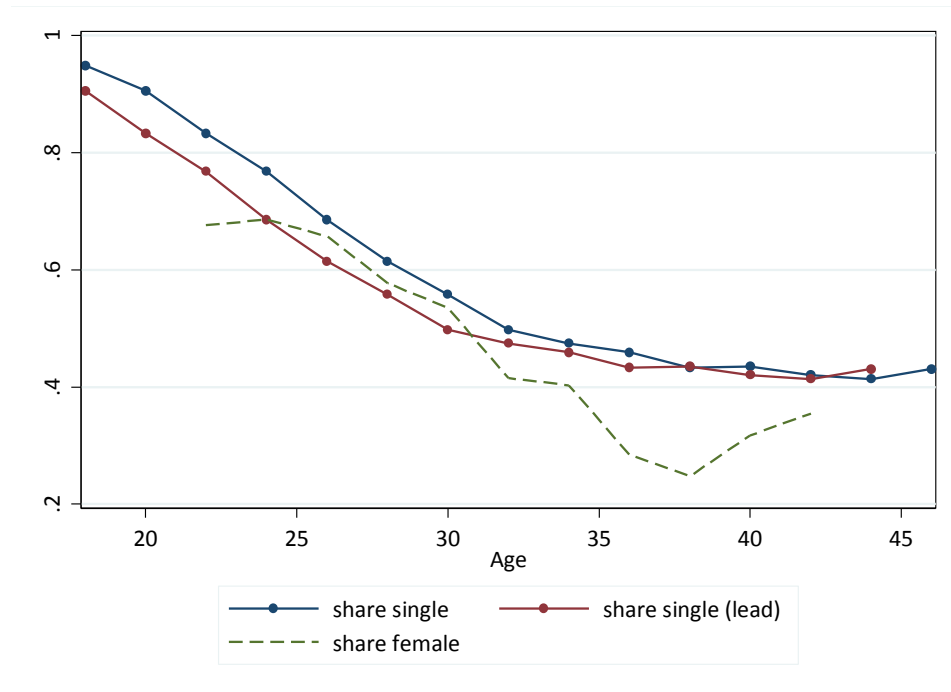
Notes: Experience bins are labeled in years. CT refers to the Computrabajo data.

Figure A3: Female Share of Gendered Job Ads versus Share of Women who have (or will have) a child under age 6, China



Dashed lines show the share of gendered job ads that request women in each job board. Solid lines show the share of urban Chinese women who have a child under age six that year, or within the following two years. See Figures 9 and 10 for sample information.

Figure A4: Female Share of Gendered Job Ads versus Share of Women who are Single, Mexico



Dashed lines show the share of gendered job ads that request women in the Computrabajo dataset. Solid lines show the share of urban Mexican women who are single at that age, or who will still be single two years later. Marital status data are from the 2010 Census microdata file. Sample restricted to urban residents (defined as areas with more than 100,000 inhabitants). Job board data have **not yet been reweighted** so their education distribution matches the above sample's.

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
Share of Age-Targeted Job Ads that are open to Men and Women

	Requested Worker Age			
	Under 25	25-29	30-34	35+
	(1)	(2)	(3)	(4)
A. Ads Open to Women				
XMZYJS	.68	.55	.42	.42
XMRC	.88	.81	.70	.65
Zhaopin	.94	.86	.84	.81
Computrabajo	.86	.83	.80	.79
B. Ads Open to Men				
XMZYJS	.54	.70	.80	.78
XMRC	.45	.68	.87	.93
Zhaopin	.67	.82	.92	.96
Computrabajo	.71	.74	.83	.92

“Open to Men (Women)” means the ad either requests men (women) or does not specify a desired gender.