

NBER WORKING PAPER SERIES

THE AGE TWIST IN EMPLOYERS' GENDER REQUESTS:
EVIDENCE FROM FOUR JOB BOARDS

Miguel Delgado Hellesester
Peter Kuhn
Kailing Shen

Working Paper 22187
<http://www.nber.org/papers/w22187>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

April 2016, Revised November 2017

Previously circulated as "Age and Gender Profiling in the Chinese and Mexican Labor Markets: Evidence from Four Job Boards." This research is supported by the National 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 Shuo Zhang for excellent research assistance. All authors contributed equally to the paper. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Miguel Delgado Hellesester, Peter Kuhn, and Kailing Shen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Age Twist in Employers' Gender Requests: Evidence from Four Job Boards
Miguel Delgado Helleseeter, Peter Kuhn, and Kailing Shen
NBER Working Paper No. 22187
April 2016, Revised November 2017
JEL No. J16,J63,J71

ABSTRACT

When permitted by law, employers sometimes state the preferred age and gender of their employees in job ads. We study the interaction of advertised requests for age and gender on one Mexican and three Chinese job boards, showing that firms' explicit gender requests shift dramatically away from women and towards men when firms are seeking older (as opposed to younger) workers. This 'age twist' in advertised gender preferences occurs in all four of our datasets and survives controls for occupation, firm, and job title fixed effects. Together, observed characteristics of job ads (including the job title) can account for 65 percent of the twist; within this 'explained' component, just three factors: employers' requests for older men in managerial positions, and for young women in customer contact and helping positions, account for more than half. The latter requests are frequently accompanied by explicit requests for physically attractive candidates. Based on its timing, the remaining portion of the twist, which occurs within job titles, appears to be connected to a differential effect of parenthood on firms' relative requests for men versus women.

Miguel Delgado Helleseeter
Economics Program
California State University Channel Islands
One University Drive
Camarillo, CA 93012 USA
805-437-8545
miguel.delgadohelleseeter@csuci.edu

Kailing Shen
Research School of Economics
ANU College of Business & Economics
HW Arndt Building (25a)
The Australian National University
Canberra ACT 0200
Australia
kailing.shen@gmail.com

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

“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, are employed in labor markets where employers’ age and gender 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 in job ads, has been subject to very little quantitative analysis.

Recently, 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. On that job board, about one in ten ads explicitly requested a worker of a particular gender. In addition, KS documented three main stylized facts. First, employers’ advertised gender preferences were *symmetric* overall, in the sense that a roughly equal number of ads requested men and women. Second, advertised gender preferences were highly *job-specific*: firm plus occupation fixed effects explained less than half of the variance in stated preferences, and a large share of firms requested men for some jobs and women for others. 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. KS developed a simple screening cost model to account for these observed patterns.

While KS brought an entirely new type of data to bear on the study of gender differentials in labor markets, their findings are based on a single job board that serves a highly skilled clientele, so their results are not directly informative about the levels and patterns of advertised gender preferences in other parts of the Chinese labor market, or about other countries where profiling is permitted. And while KS presented some parallel evidence on age profiling, they did not study the interaction between employers' age and gender preferences.

Motivated by these gaps, this paper studies the patterns identified by KS, and the interaction between firms' age and gender requests in three new data sets, two from China and one from Mexico. Importantly, the two new Chinese job boards 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 large Chinese city and contain a much larger share of unskilled jobs. In order to see whether KS's main results are unique to China, we also collected data from another middle-income country that permits gendered job ads—Mexico.¹ While China and Mexico have very similar GDP per capita and both have weak anti-discrimination laws, historical differences that could generate distinct patterns of explicit demand for female labor include China's planned economy --which encouraged very high female labor force participation-- and Mexico's Catholic traditions encouraging traditional family structures. More recently, China's one-child policy is an important institutional difference, contributing to China's much lower fertility and higher female labor force participation rate than Mexico's during our

¹ To identify an appropriate non-Chinese dataset to study, we started by examining the country-specific sites of Computrabajo.com, which is the leading job board in Latin America. We then chose Mexico because it was the largest, had a similar GDP to China's, and (like many other boards) also hosted gender-specific job ads. Because Computrabajo is a national site with a heavy focus on Mexico's large cities, it provides a natural comparison to Zhaopin.com.

sample period. Patterns that are shared by the two countries despite these differences are likely driven by a more universal set of causes, which we attempt to isolate in the paper.

Our main findings are threefold. 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, at 43, 72 and 77 percent, versus 13% in Zhaopin. These patterns are consistent with the negative skill-targeting relationship: age- and gender profiling are much more common on job boards catering to less-skilled workers. Second, the three main patterns identified by KS in Zhaopin.com —symmetry, job-specificity and the negative skill-targeting relationship— are confirmed within each of the new datasets as well. Confirmation of the negative skill-targeting relationship provides additional support for the idea that ongoing increases in the demand for skilled labor will reduce the incidence of explicit age- and gender-targeting of job ads in developing economies.

Finally, we establish a new stylized fact concerning the interaction between age and gender profiling.² This relationship, which we call the *age twist* in employers' gender requests, is a strong shift in the *direction* of employers' gender requests, away from women and towards men, as the desired age of the worker rises from 18 to 45. This age twist survives controls for occupation, firm, and firm-by-occupation fixed effects, and —like KS's results discussed above— is present in all four datasets, including Computrabajo. To help us understand this twist, we introduce information derived from job titles. While job titles have not been widely used in economic research, they are more detailed (and more predictive of wages and application

² 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 job boards.

decisions) than six-digit SOC codes (Marinescu and Wolthoff, 2015). We find that observed aspects of job ads (including the titles) can account for 65 percent of the age twist; of that explained portion, just three factors: employers' explicit requests for older men in managerial roles, and for young women in a handful of customer contact and helping roles (like front desk, customer service and clerk positions) explain more than half. These requests for young women, in turn, are highly correlated with explicit requests for physical attractiveness. For example, in the Zhaopin data, almost half (45.8 percent) of job ads aimed at women under the age of 30 explicitly request that the applicant be physically attractive, compared to 9.8 percent for men under 30.³

We also find that the age twist remains substantial in magnitude and highly statistically significant even *within* the most detailed job titles available to us. Based on its timing, this portion of the twist appears to be connected to an effect of parenthood on firms' relative requests for men versus women. Section 2 of the paper describes all four of the data sets at our disposal. Section 3 describes our replication of KS's main results in our three new datasets. Section 4 documents and attempts to understand the causes of the age twist in advertised gender preferences, and Section 5 concludes.

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.⁴ In part because Xiamen was one

³ Similar but less patterns are present in Computrabajo, the only other data set where we have information on employers' explicit beauty requests.

⁴ 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

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 the local government's 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 it operates purely as a job-posting service: workers cannot post resumes or apply to jobs on 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. 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 nationally recognized as the dominant job board for skilled workers in Xiamen, possibly due to

region. In contrast, the 2012 national populations of Norway, Sweden and Denmark are 5.0, 5.6 and 9.6 million respectively.

its close links with the local government and social security bureau. Details on how we constructed an analysis sample for XMZYJS, XMRC, and our other two datasets are available in Appendix 1.⁵ The most important restrictions affecting all datasets involve dropping job ads with missing education requirements and missing occupation information, and all jobs requesting workers over the age of 45.⁶ In all four datasets we also drop a very small number of ads for public-sector jobs. Recruiting for these jobs takes place via other mechanisms.

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 who on average are considerably more skilled than even those on XMRC.

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.⁷ 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. XMRC, in turn, hosts considerably less-skilled ads on average than

⁵ All the paper's appendices are available at:
<https://sites.google.com/site/pjkuhn83/data/age-and-gender-profiling>

⁶ Throughout the paper, an ad's 'requested age' refers to the midpoint of the requested minimum and maximum age. Ads with a requested age above 45 are extremely rare in all our data sets. As we discuss later in the paper, this is likely because job turnover, and the tendency to recruit workers via formal job ads, both decline rapidly with the worker's age.

⁷ These were placed by 8,839; 6,680; and 73,636 firms; thus we observe an average of 16.0, 5.9 and 14.3 ads per firm.

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 27.6 to 30.5 years.

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.⁸ Similarly, the share of ads that specify both a

⁸ Following KS, we include all expressed degrees of preference in our indicators of age and gender targeting. In all four job boards, however, the most common way to express these preferences is either to fill in a dedicated field or simply to write: "Age: 20-24" and/or "Sex: male" in the job description. 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.

minimum age and a maximum age (our main summary indicator for whether the ad is age targeted) declines monotonically as we move from the least-skilled job board (XMZYJS) to the most skilled (Zhaopin). Taken together, these patterns 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). This may be due to the large number of less-skilled, manual jobs on that board.

To provide some context for our ad samples in Table 1, Table 2 provides comparable statistics for private-sector employees in Xiamen and in urban China, respectively. For comparability with our job board samples, these data (taken from the 2005 Chinese Census) exclude workers over the age of 45.⁹ Table 2 shows, first of all, that—at least when we focus on workers under 45—Xiamen's labor forces is quite representative of the rest of urban China in its gender, education, age, experience and wage composition. Turning to comparisons between our job board data and the urban labor force, we note first that the balanced gender mix of urban

⁹ 'Urban China' in Table 2 and throughout this paper refers to China's largest cities—specifically the four municipalities directly under the jurisdiction of the central government (Beijing, Shanghai, Tianjin and Chongqing) plus the 15 subprovincial cities. As we argue below, our job boards are broadly representative of these cities, which are more prosperous and modern than the rest of urban China. Together, they have a population of 228 million people, which is 18 percent of the country. Labor markets in rural areas and smaller cities may operate differently from the markets studied in this paper.

Chinese workers (54 percent male) mirrors the balanced mix of job ads on XMRC, but is more female than XMZYJS's ad mix. Eighty-six percent of Xiamen's workers have high school education or less, while the share of ads requesting such workers is 93.7 percent on XMZYJS, 49.3 percent on XMRC and 12.9 percent on Zhaopin. Thus, while XMZYJS serves a labor market that is slightly less skilled than the Chinese average, the other two boards, as already noted, serve a highly skilled slice of the Chinese labor market.¹⁰

Once we focus on workers under 45, there is not much difference between the age of the workers requested on these three job boards and the urban Chinese average either: XMZYJS and XMRC request workers who are one to two years younger than the typical Xiamen employee, and Zhaopin's mean age request of 30.5 is very similar to the national average. Compared to mean monthly wages of about 2,200 RMB in Xiamen and urban China, mean advertised wages on XMZYJS, XMRC and Zhaopin are 1,810, 2,556 and 4,279 RMB respectively, reflecting the job boards' focus on different skill levels relative to the national average.¹¹

Our Mexican data is a sample of job ads posted on Computrabajo.com.mx. 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 website that disproportionately serves highly skilled workers. The

¹⁰ Since there was considerable educational upgrading in China between the 2005 Census and our ad sampling dates, these comparisons probably understate XMZYJS's focus on unskilled workers relative to the city's labor force.

¹¹ Comparing the occupation mix of our job boards with the urban labor force is difficult because of the idiosyncratic job classification systems used by each board; mapping these into Census categories is a fairly subjective exercise. With these cautions in mind, Table A1.3 attempts to compare the broad occupation distributions on the boards to the urban labor force as defined in Table 2. Broadly speaking, jobs in production, construction and manufacturing are under-represented on our job boards (especially on Zhaopin), while professional and technical jobs are over-represented on all our boards. Essentially, our job board sample over-represents highly skilled occupations, which are less likely to be gender-profiled than others.

construction of our Computrabajo sample was very similar to our Chinese datasets, and is also described in Appendix 1. As reported in Table 1, our Computrabajo analysis sample contains 90,487 ads.¹² As in Zhaopin, a large share of ads on Computrabajo (42 percent) require a university education, though jobs requiring high school or less are much more common on Computrabajo than on Zhaopin. Consistent with this skill difference, Computrabajo jobs require less experience on average 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). As on Zhaopin and XMRC, Computrabajo's gender-targeted ads are also equally split between men and women.

Table 2 also presents comparable statistics for private-sector employees in urban Mexico, from the 2010 Mexican Census. As expected, the 42 percent of Computrabajo ads that target university-educated workers is much higher than the 21 percent of Mexico's urban labor force that is university-educated. Urban Mexico's labor force is more male than China's (at 64 percent male), but like our Chinese boards the average age requested on Computrabajo is slightly younger than the average urban Mexican worker. Similarly, reflecting Computrabajo's focus on skilled workers, the offered wage, at 7640 pesos per month (approximately 580 U.S. dollars based on exchange rates for that period), is somewhat higher than the average wage of employed, urban workers (6793 pesos/month). Both are well above the average wage in the country as a whole.

Overall, once we focus on the labor market for workers under the age of 45, the characteristics of the workers requested in our four job boards are broadly similar to the

¹² These were placed by 1,944 firms for an average of 46.5 ads per firm.

characteristics of employed, private-sector workers in urban China and Mexico. In cases where the samples differ, they do so in ways that we would expect, based on the labor market segments targeted by each board.

3. Replicating KS: Job-Specificity and the Negative Skill-Targeting Relationship

Table 1 has already confirmed the rough symmetry of employers' requests for men and women in our three new datasets, and shows that the negative skill-targeting relationship holds across datasets: job boards serving skilled workers are much less likely to gender- and age-target their job ads than other job boards. In this Section we summarize evidence showing that the negative skill-targeting relationship also holds within each of the new datasets and is robust to regression controls. We also show that—as in Zhaopin—employers' gender requests in all the new datasets are highly job-specific in the sense that they vary more within firms than across firms. In all cases, the details of the analysis are provided in the Appendix.

3.1 The Job-Specificity of Employers' Gender Requests

Since firms typically place multiple ads on Zhaopin, KS were able to ask what share of firms post at least one ad requesting a man *and* at least one requesting a woman. Among firms who placed more than 50 Zhaopin job ads, 38.6 percent placed at least one ad that favors men and one ad that favors women. In XMZYJS, XMRC, and Computrabajo respectively, these shares were 93.4, 96.2, and 89.2 percent respectively (see Appendix 2 for details). These patterns support the idea that gender-typing of jobs *within* firms plays a major role in the age twist.

A second way to assess the relative role of firms and jobs in accounting for patterns of gender requests is to decompose the variance of advertised gender preferences across ads into components that are associated with occupations, with firms, with differences in the way firms assign genders to occupations, and within-job cell (firm*occupation cell) variation. The

occupation controls used in this analysis are the categories that each of the job boards uses to classify ads on their site, which number 58 in XMZYJS, 36 in XMRC, 39 in Zhaopin and 14 in Computrabajo respectively. Here, Appendix 2 shows that less than 40 percent of the variance in advertised gender preferences can be accounted for by occupation and firm fixed effects in all four of our datasets. In addition, all datasets exhibit a large amount of within-job-cell (firm*occupation) variation in advertised gender preferences, ranging from about a third of the total in Zhaopin and XMRC to over 70 percent in Computrabajo.

Taken together, the above evidence suggests that models based on *firm-level* tastes for men relative to women are likely to be of little help in understanding patterns in gender profiling. A more appropriate class of models would view firms as having a portfolio of jobs, some of which they view as best suited to men, and others to women. The importance of the *type of work done* rather than the identity of the firm doing the hiring for gender profiling is dramatically confirmed in Section 4 of the current paper, where we introduce detailed job descriptions derived from job titles in addition to firm fixed effects.

3.2 The Negative Skill-Targeting Relationship

As noted, KS found that Zhaopin jobs requiring higher skill levels were less likely to be gender-targeted, and less likely to be age- targeted. Using education requirements as our main skill measure, Figure 1 of this paper confirms that this relationship holds for both age- and gender-targeting within all four of our datasets. 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 tendency for employers to prefer one gender in skilled jobs and the other in unskilled jobs.

KS also showed that the skill-targeting relationship was robust to different measures of skill (education, experience and/or wages), and to regression controls for occupation effects, firm effects and their interaction. This is confirmed for our three new datasets in Appendix 3: The relationship persists when we use required experience or the offered wage as our measure of skill requirements, and in regression specifications that are as similar as possible to KS's in the three new datasets, including occupation and occupation*firm controls.

3.3 What Explains the Negative Skill-Targeting Relationship?

As we have noted, one interpretation of the skill-targeting relationship, suggested by KS's model, 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 predict that firms should search more broadly as jobs' skill demands rise. KS, however, also identify other possible causes of the skill-targeting relationship, including skill-based variation in stigma against employers posting targeted ads, in application processing costs, in the expected number of applicants per position, and in the idiosyncratic variance *across jobs* in their relative suitability for men versus women (σ_v in the model). In the latter case, the argument is that while men and women may be quite different in their suitability for unskilled jobs (with systematic differences in fine manual dexterity versus gross motor strength, for example), such gender differences are less important in skilled, or mental tasks.

In Appendix 4, we replicate a series of tests designed by KS to distinguish among the above explanations, and conclude, as KS do, that uncontrolled variation in stigma, in application processing costs, and σ_v are not likely explanations, but that a decline in the number of applications per job (i.e. increasing labor market tightness) with skill could help account for the

negative skill-targeting relationship (provided that this tightness pattern applies in all four of our datasets and in the presence of all our regression controls).

Summing up, evidence from our three new datasets shows that the negative skill targeting relationship first identified by KS extends to the use of both gender and age targeting, in four different samples of job ads serving workers of very different skill levels in two emerging economies. The relationship is robust to controls for the type of work done (firm*occupation fixed effects) and holds for three different measures of skill requirements: education, experience, and the level of the posted wage. Additional tests show that the relationship is most likely a direct result of skill requirements (or of systematically tighter labor markets for skilled jobs in all four of our datasets) than of other unmeasured factors. A potential implication is that the rapid, ongoing skill upgrading in many emerging economies may itself lead to a voluntary reduction in employers' use of coarse demographic screens like age and gender to pre-judge the suitability of potential job applicants.

4. The Age Twist in Employers' Gender Preferences

4.1 Documenting the Age Twist

The age twist in employers' gender preferences in all four of our datasets is shown in Figure 2. While the share of ads making a gender request (the sum of the male and female bars) does not vary much with the worker's age, in all cases the *direction* of firms' advertised gender preferences shifts strongly away from women and towards men as the desired age of the worker 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 (versus a man) 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. 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.¹³

What explains this striking pattern? As a first step in understanding it, we ask to what extent these correlations could be associated with differences in the types of skills employers demand from workers of different ages, or by the possibility that men and women may work at different firms as they age. To that end, we run regressions of the 'direction' of employers' gender preferences on various features of the job ad, including occupation indicators. Following KS, our measure of the direction of employers' gender preferences is given by $(P^M - P^F)$, where P^M (P^F) is an indicator for whether the ad requests a man (woman). Thus, $(P^M - P^F)$ 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.¹⁴

¹³ Similar patterns have been observed for callbacks in U.S. correspondence studies. For example, Lahey (2008) found that younger women were 40 percent more likely to receive callbacks than older women. Neumark, Burn and Button (2015) find considerably more age discrimination in hiring against older women than older men in a large resume correspondence study.

¹⁴ 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, if we model the probability of preferring men conditional on having a

To interpret the effect of occupation and other controls on the estimated age twist, let $\beta^U > 0$ (the unadjusted age twist) be the coefficient on an age indicator in a regression without the above controls:

$$P^M_i - P^F_i = \alpha + \beta^U A_i + \varepsilon_i, \quad (1)$$

and let β^A (the regression-adjusted twist) be the coefficient from a regression that controls for whether the job is, say, in occupation 1, X^1 :

$$P^M_i - P^F_i = \alpha + \beta^A A_i + \gamma^1 X^1_i + \varepsilon_i. \quad (2)$$

Using the omitted-variable bias formula, controlling for occupation 1 helps account for the age twist ($\beta^A < \beta^U$) iff:

$$\gamma^1 D^1 > 0, \quad (3)$$

where D^1 is estimated from the auxiliary regression:

$$X^1_i = a + D^1 A_i + e_i. \quad (4)$$

Thus there are two types of occupations that can help explain the age twist: occupations for which firms disproportionately request men ($\gamma^1 > 0$), and which request older workers more frequently than younger workers ($D^1 > 0$) (henceforth ‘growing male’ occupations), and occupations for which firms disproportionately request women ($\gamma^1 < 0$), and which are less frequently requested as workers age ($D^1 < 0$) (henceforth ‘shrinking female’ occupations). Examples might be management and low-level administrative positions, respectively.

Estimates of the above types of regressions for each of our datasets are presented in Tables 3-6. Aside from the age controls, the specifications are identical to the regressions used to investigate the skill-targeting relationship in Appendix 3. Column (1) of each table has no

gender preference, or if we estimate separate regression models for P^M_i and P^F_i , relative to not stating a gender preference.

controls, thus reproducing the results in Figure 2. Column 2 adds controls for the job's education requirement, firm characteristics (size and ownership type), the number of job openings advertised and occupation fixed effects.¹⁵ Column 3 replaces the firm characteristics by a full set of firm fixed effects, and column four interacts the firm and occupation fixed effects.

Overall, there are two main messages from Tables 3-6. First, occupation and firm controls account for a substantial share (ranging from 32 to 66 percent depending on the data set and which age category is compared to ages 18-24) in XMZYJS and Zhaopin, but for much more modest shares in the remaining datasets. Indeed, the attenuating effect of these variables is essentially absent in Computrabajo, which has the coarsest occupation categories. Second, the estimated twist remains large in magnitude and highly statistically significant in all four datasets, even in the presence of column 4's occupation*firm controls. Essentially, these column 4 regressions compare two job ads for the same occupation in the same firm (say, sales at Dell Computer) requesting workers of different ages. They show that when an older worker is requested, the firm is more likely to request a man relative to a woman compared to when a young worker is requested.

4.2 Understanding the Age Twist: Job Titles and the Role of Gendered Tasks

The attenuation of the estimated age twist when occupation controls are introduced in Tables 3-5 suggests that more detailed age-related differences in the type of work employers offer to men and women of different ages might account for more of the twist. For example, suppose that employers specifically prefer young women over other age-sex groups for

¹⁵ In three of the four datasets, geographical controls (for the location of the job) are also included: districts of the Xiamen metro area in XMZYJS, Chinese provinces in Zhaopin, and Mexican states in Computrabajo. Zhaopin (the only dataset that has a separate industry variable) adds industry effects here also.

customer-contact jobs. Or suppose that employers prefer their truck drivers (who are overwhelmingly male) to be older because young men are seen as unreliable. A large or influential set of jobs like this —where preferences or expected productivity depend on age-sex interactions— could explain the age twist.

To explore this idea further, we are able to access much more detailed information about job duties —the job title— in one of our datasets: XMRC. Job titles are the main heading in every job ad. They provide a brief description of the job, and can run up to 9 words in XMRC. For example, here is a random sample of ten job titles on the XMRC website: administration, front desk administration assistant, project engineer, quality control, shift leader, customer service maintenance specialist, ME product engineer Xiamen, experienced two-dimension designer, customer service engineer, and front desk clerk. Job titles appear to provide considerably more relevant information about the type of work than is available even in the most granular standardized occupational classification systems. For example, Marinescu and Wolthoff (2015) found that job titles on Careerbuilder.com were much more predictive of advertised wages than 6-digit SOC codes, and were essential controls for identifying the effect of advertised wages on the number and quality of applications an ad received.

Job titles are highly predictive of employers' gender requests. This is illustrated in Table 7, which focuses on just the first word of the title and lists the 25 'most female' words, the 25 'most male' words and the 25 words that are least likely to be explicitly gendered.¹⁶ Some first words which occur very frequently in the data (like front desk, clerk, and teller) are strongly associated with explicit requests for female employees, while others (security, electronic, and

¹⁶ Table 7 restricts its attention to first words that occur in at least 30 ads.

driver) are strongly associated with requests for men.¹⁷ This suggests that job titles might help us understand not only employers' gender preferences, but how they interact with their age preferences.

To create a sample of XMRC job titles for analysis, we translated and cleaned all job titles that appeared more than once in the XMRC dataset. This gave us a dataset of 9,884 job ads (relative to 17,021 in Table 4) which appears to be representative: running the Table 4 regressions on this subsample yielded almost identical results. To quantify the ability of these new job title variables to account for the age twist, it is helpful to have a scalar measure of the magnitude of the twist (as opposed to the four age categories used in Figure 2 and Tables 3-6). Accordingly, in Table 8 we measure the job's requested age by a single dummy variable for whether the requested age of the worker is 30 or higher.¹⁸

Columns 1 and 2 of Table 8 reproduce the specification in columns 1 and 2 of Table 4 for our sample of translated titles with a scalar age twist indicator. They show, once again, that adding broad occupation controls attenuates the estimated twist somewhat: from 44.2 to 37.2 percentage points. Columns 3-6 successively add fixed effects for the first $k = 1, \dots, 4$ words of the raw job title, while column 7 has fixed effects for every title, regardless of length. Column 8 adds a full set of firm fixed effects to the regression. According to the Table, controlling for just the first word of the job title attenuates the estimated age twist much more ($.372 - .257 = 11.5$ percentage points) than a full set of occupation controls ($.442 - .372 = 7.0$ percentage points).

¹⁷ "Electronic" as a first title word was followed by a wide variety of words, including design(er), engineer, technician and apprentice. The two most common complete job titles starting with "electronic" are electronic machine repair, and electronic service maintenance.

¹⁸ About 42 percent of age-targeted job ads request workers in this age range (30 to 45 years old).

Adding the second and third words attenuates the estimated twist further (to 17.0 percentage points), after which neither additional words nor a full set of firm fixed effects changes the age coefficient substantially. In Table 8's most saturated specification (column 8), the estimated age twist equals .155. Thus, in total, a differing mix of detailed types of work and (to a lesser extent) a different mix of firms hiring old versus young workers, together can account for $.442 \times .155 / .422 = 65$ percent of the age twist.

What specific types of jobs account for employers' declining explicit interest in female workers, and increasing requests for male workers as workers age? To answer this question, we manually inspected large numbers of titles and noted that their most frequently-occurring words fell into four main categories: words denoting the job's *rank or skill level*; words suggesting frequent *contact* with customers, clients or business associates; miscellaneous words describing the worker, job or firm (but not the job duties), and (most commonly) words describing the type of work done ('*duties*'). We then coded these words as follows.

Rank and skill level variables

Job title words describe a position's rank or skill level in three main ways. The first indicates supervision of others; we categorized such *management* words (or word groups, like division vice manager) into ten detailed ranks ranging from division manager to supervisor. We also created a more aggregated three-fold categorization (higher management, managers and chiefs, and leaders and supervisors) and an aggregate *Managers* indicator for whether any managerial or supervisory duties were mentioned.

A second set of frequently-occurring words denoting rank or skill indicated a subordinate rank in the sense that the job entails helping or assisting someone else. These were assistant,

clerk, secretary, and helper. We created dummy variables for the occurrence of each of these words, plus an aggregate *Assistants* indicator for the presence of any one of them.

A third group of words suggested a high level of expertise in a certain area without implying supervision or authority over others. We created dummies for each of these words: specialist, consultant, elite, master, professional, technician, senior and mentor, plus an aggregate *Experts* indicator for the presence of one or more of them.

Contact variables

Motivated by models of customer discrimination, we also coded the following frequently-occurring words suggesting that the job requires frequent contact with customers, clients or business associates: front desk, customer service, teller, cashier, sales and representative. We created dummies for the occurrence of each of these words, plus an aggregate “*Contacters*” indicator for the presence of any of them.

Miscellaneous worker and job attributes

We created dummies for each of the following words indicating desired aspects of the applicant: experienced, new graduate, reserve cadre, and English; as a group we refer to these as *worker words*.¹⁹ Similarly, we created a set of *job words* dummies for each of the following descriptors of the job or firm that do not refer specifically to the tasks to be performed: shift, team, indoor, outdoor, part time, full time, domestic, international, and urgent (which indicates the firm wants to fill the position quickly).

¹⁹ “Reserve cadre” in private sector firms refers to entry level jobs that involve rotating between different departments in preparation for promotion to a managerial position. Only 57 of the job ads in the Table 9 regressions (or 0.58 percent of the total) are for reserve cadre positions.

Duties

By far the most common job title words describe the type of work done, as in typical occupational classification systems. To capture these *duty words*, after encoding all the *managers, helpers, contact, worker* and *job* words listed above and removing frequent nuisance or connector words (specifically in, of, the, for, to, on, be, by, with, one, as, or), we created dummy variables for all 515 remaining title words. A list of the 100 most common of these words is provided in Appendix 5. The top ten duty words, in order, are administration, engineer, HR, designer, business, warehousing, accounting, division, QC, and business person.

We then run a series of regressions for $P^M - P^F$ (the explicit ‘maleness’ of the job) starting with the variables used in Table 4, then adding increasingly detailed controls for the above aspects of the job. For reference, columns 1 and 2 of Table 9 reproduce columns 1 and 2 of Table 8, which do not rely on the job title controls. Column 3 adds just four variables derived from the job title—the aggregate *Managers, Assistants, Contacters* and *Experts* indicators. Of these, two—*Assistants* and *Contacters*—have very large and statistically significant effects, indicating that jobs requesting helping or contact duties are much more likely to request women relative to men than other jobs (corresponding to shifts of 48 and 41 percentage points respectively). Requests for *Experts* have a precise zero partial correlation with firm’s gender requests. Finally, in this specification jobs that request managers (of all ranks) are actually slightly more likely to request women relative to men (a 3 percentage point effect). Together, adding just these four indicator variables to the regression attenuates the estimated age twist from 37.2 to 27.6 percentage points, a substantially greater reduction than is obtained from the full set of occupation controls in column 2.

Column 4 replaces these four broad indicators —*Managers, Assistants, Contacters* and *Experts*— by indicators for their individual constituent words (only the coefficients for the managers, assistants and contacters words are reported).²⁰ Column 4 also adds the miscellaneous *worker words* and *job words* extracted from the title. This has essentially no effect on the estimated age twist, but indicates that employers’ strong preference for female *Assistants* is driven by three individual job words: assistant, clerk and secretary. The strong preference for female contacters comes from all of the following words: front desk, customer service, teller, cashier and sales. (‘Representative’ is (insignificantly) male). Interestingly, it is primarily the lower-ranked managerial positions —especially supervisors— where women are disproportionately requested. Finally, it is worth noting that the *job words* indicators (not reported in the table) for *part time* work, *outdoor* work and for vacancies that must be filled *urgently* are all correlated with gender requests in directions one might expect— outdoor work (part-time work) predicts explicit requests for men (women), and employers are more likely to invite women to apply when the job must be filled urgently. These words, however, do not account for a quantitatively meaningful share of the age twist, primarily because all three occur relatively rarely in our sample of job titles.

Column (5) adds controls for all the 515 remaining words in the job title, which overwhelmingly describe the specific tasks to be performed.²¹ These ‘duty’ words further attenuate the estimated twist to 18.9 percentage points. Importantly, these detailed task controls

²⁰ Only two of the eight ‘experts’ words had a statistically significant effect on the direction of employers’ gender requests: “specialists” were disproportionately female, while “technicians” were disproportionately male. Controlling for these variables does not account for a quantitatively meaningful share of the age twist, however.

²¹ Of these, 374 were retained in the regression after eliminating collinear words.

change the sign of the estimated manager coefficients relative to columns 3 and 4. While overall employers disproportionately request women for managerial positions, column 5 indicates that this is because women are requested for tasks that employ low-ranked leaders (especially supervisors). Once we look within detailed tasks, employers actually express a substantial explicit preference for men in leadership roles, especially at higher levels. Finally, column (6) replaces the firm size and ownership variables by a full set of firm fixed effects, with little change to the estimated coefficients and very little additional effect on the estimated twist.

Overall, comparing columns (1) and (6) of Table 9, adding all the controls now available in the XMRC data reduces the estimated age twist from .442 without any controls to .181. These controls thus account for $.442 - .181 = .261$, or 59 percent of the overall age twist in employers' gender requests on the XMRC website. The difference between this 59 percent and the 65 explained percentage in Table 8 is due to details of job titles that are not captured by our theoretically motivated, dimension-reducing coding of the title words.

4.3 Decomposing the Age Twist: Which Job Characteristics Matter Most?

While successively adding covariates to a regression sheds some light on which ones are most important in accounting for an empirical regularity, the estimated importance of different covariates is not invariant to the order in which they are added to the regression. Thus, while Table 9 suggests that employers' demand for young, female *Assistants* and *Contacters* plays an important role in the age twist, in this Section we use an approach recently extended by Gelbach (2016) to uniquely decompose the unadjusted age twist in our data (.442) into components associated with various covariate groups. In addition to order invariance, Gelbach's approach also has the advantage of yielding consistent estimates of economically and econometrically meaningful population parameters that are based on the omitted variable bias formula, and of

producing consistent estimates of the standard errors of the covariate groups' contributions. Such estimates are often absent in previous decomposition exercises.

In our context, as suggested by equations (1)-(4) the ability of a job characteristic, X , to account for the age twist depends both on its effects on employers' gender preferences (γ) and on its correlation with employers' age requests (D). Generalizing equation (2) to the case of multiple regressors,

$$P^M_i - P^F_i = \alpha + \beta^A A_i + \gamma^1 \mathbf{x}^1_i + \gamma^2 \mathbf{x}^2_i + \varepsilon_i, \quad (5)$$

where \mathbf{x}^1_i and \mathbf{x}^2_i are vectors of variables in two different categories (for example, contact variables and top-100 words), the omitted-variable bias formula allows us to partition the explained portion of the age twist into components related to each of the two variable groups:

$$\text{'Explained' gap} \equiv \beta^U - \beta^A = \gamma^1 \mathbf{d}^1 + \gamma^2 \mathbf{d}^2 \quad (6)$$

where each element of \mathbf{d}^j (D^{kj}) is estimated from a different auxiliary regression:

$$X^{kj}_i = a + D^{kj} A_i + e_i. \quad (7)$$

Figure 3 uses equation (6) to decompose the explained twist of .261 in Table 9 into the contributions of eleven distinct variable groups, showing 95 percent confidence bands.²² Of these, the largest single contributor (.076) is the occurrence of a 'helping' word in the job title. Five other categories of variables also contribute substantial amounts, specifically the detailed 'duty' words from the job title (.067), the presence of a 'contact' word in the job title (.048), an indication in the title that the job contains some managerial duties (.027), differences in occupation mix between jobs requesting older versus younger workers (.025), and firm fixed effects (.015). The remaining five variable categories: education requested, number of vacancies

²² Contributions and their standard errors were computed using Gelbach's Stata code b1x2.do.

advertised, expertise words, miscellaneous worker words, and miscellaneous job words, contribute negligible amounts. Importantly, the large role of job titles relative to the quite muted role of a full set of firm fixed effects underscores KS's *job-specificity* result: explicit gender preferences are very much about the *type of work* that is done and much less about which firm is doing the hiring.²³

Figures 4-6 provide additional detail on the contributions of individual variables within the five quantitatively important variable categories in Figure 3. According to Figure 4, just three occupations — construction, manufacturing, and administration ($.0091 + .0034 + .0123 = .0248$) — can account for the entire occupation mix effect (though some smaller countervailing effects exist). Of these, construction and manufacturing are ‘growing male’ occupations (occupations that constitute a larger share of explicitly gendered job ads for older workers than younger workers), while administration is an influential female occupation that shrinks in size with age.

Figure 5(a) shows that growing requests (with worker age) for higher management positions (which are disproportionately male) account for the majority of the contribution of the *Managers* variables to explaining the twist. Figure 5(b) shows that 62 percent ($.047/.076$) of the *Assistants* effect comes from explicit requests for female clerks, with smaller contributions from assistants and secretaries (Figure 6). According to Figure 5(c), the largest single contributor to the *Contacters* effect on the twist comes from explicit requests for young women in “front desk” positions, with customer service and teller titles also contributing substantial shares.

²³ This large role of job titles relative to firm fixed effects is not a consequence of the large number of job titles included in the model. Indeed column 8 of Table 9 controls for 515 job duty words compared to 3,633 firm IDs; despite this the duty words account for much more of the twist.

Finally, turning to the job ‘duty’ words, Figure 6 shows that ‘growing male’ title words contribute considerably more to the age twist than ‘shrinking female’ words (though both are counteracted by substantial contributions from shrinking male and growing female words). To provide some context for these effects, the ‘growing male’ job title words making the ten largest contributions to the age twist are driver (.0293), electronic (.0135), security (.0126), engineer (.0052), chef (.0046), high (.0031), plumber (.0029), molding (.0026), fitter (.0023), equipment (.0022), and machine (.0020). Together, these words suggest that employers’ disproportionate requests for older versus younger men in a variety of semi-skilled occupations also play a quantitatively significant role in the twist.²⁴

Summarizing all the above information, Figure 7 apportions the explained age twist to its main sources. Most strikingly, 53% of the explained age twist is attributable to just three sets of factors: employers’ preferences for older men in managerial positions (9%); employers’ preferences for young women in four customer contact jobs: front desk, customer service, teller and cashier (17%); and employers’ preferences for young women in three ‘helping’ positions: clerk, assistant and secretary (27%). That said, a disproportionate number of requests for older men in the following occupations (and job titles) also play a role: construction and manufacturing occupations; and driver, electronic and security as job title words.

Importantly, the results in Table 9 and Figures 3-7 are robust to all the different ways we have explored to encode the rich information contained in job titles. For example, very little

²⁴ Employers’ explicit requests for older men in some of these occupations may be related to a perception that men become more reliable workers with age. We discuss this possibility further in Section 4.5.

changes when we control for much smaller numbers (than the full 515) of job title words.²⁵ The intuition is simply that infrequently-appearing job duties can only have a limited explanatory effect on a bounded outcome variable; thus the main results are driven by specific title words that are already listed in Figure 7. Appendix 5 shows that the results are also robust to replacing dummies for occurrence of individual ‘duty’ words by fixed effects for 880 mutually exclusive combinations of all those words.²⁶ Results are also very similar when we model requests for men (P^M) and requests for women (P^F) as separate dependent variables instead of as a difference. Here, we find that the share of ads explicitly requesting men rises with age and the share of ads explicitly requesting women falls with age. Both of these effects attenuate as controls are added, in much the same way that their difference attenuates. Indeed, the estimated effects of age on ($P^M - P^F$) from this procedure are arithmetically identical to those reported in Tables 8 and 9.

4.4 Exploring the *Contacters* and *Assistants* Effects: The Role of Beauty

The previous section estimated that employers’ preferences for young women in contact and helping positions can account for a substantial share of the age twist (about 47 percent of the explained component, and $.47 * 59 = 28$ percent of the total). But why do employers prefer young women in these roles? Motivated by literatures on beauty and customer discrimination in labor markets, here we ask whether a desire for physical attractiveness plays a role.²⁷

²⁵ Appendix Figure A5.1 replicates Figure 3 using only the 50 most common job duty words. Even here, the *Assistants* category is the largest single contributor, and the 50 most common duty words still contribute much more to the twist than the complete set of firm fixed effects.

²⁶ After dropping collinear categories, 647 of these 880 categories were retained in the regression.

²⁷ 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

Unfortunately, our analysis is limited by the fact that we do not have data on explicit beauty requests in the XMRC data.²⁸ Still we can draw some inferences using data on employers' explicit requests for beauty in the Zhaopin and Computrabajo datasets. As documented in Table 1, employers request attractive workers in 4.0 percent of ads on the Computrabajo site, and in 7.7 percent on Zhaopin. These beauty indicators were derived by searching the text of the detailed job descriptions for phrases indicating a preference for a visually attractive candidate, then compiling a list of the most common forms of such requests. Appendix 1 lists the ten most common requests in both datasets.²⁹ In the Computrabajo data we have access to a second indicator of firms' interests in the candidate's attractiveness: 10.9 percent of Computrabajo ads instructed applicants to include a photo with their application.³⁰

Using the above data we are able to establish the following facts concerning the role of beauty in the age twist.

markets. Kapoor and Magesan (2016) document customer discrimination in the form of productivity-enhancing effects of facial beauty among restaurant workers.

²⁸ When they are made, requests for beauty generally appear in the detailed job description portion of the ad (not the job title). We have this information in the two datasets we scraped (Zhaopin and Computrabajo), but this information was not provided to us by XMRC and XMZYJS, who gave us data extracts directly.

²⁹ 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.

³⁰ 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 Zhaopin's data division chief, Jian Hao) already include photos. Accordingly, we do not use the photo request information in the Zhaopin data.

Fact 1. Employers are more likely to request beauty, or to request a photograph, when they are requesting young women than any other demographic group.

For the Zhaopin data, Fact 1 is illustrated in panel (a) of Figure 8. 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 confirm that employers' requests for beauty covary strongly with their age and gender requests, and do so in a direction that could help account for the age twist.

Fact 2. In the datasets where we have information on explicit employer requests for beauty, those requests account for 7 to 15 percent of the age twist.

To estimate the extent to which employers' *explicit* beauty requests can account for the estimated age twist, we added a control for such requests to the gender preference regressions in Tables 5 and 6. The results, reported in Appendix 5, show that beauty (and photo) requests strongly predict explicit requests for women, and that controlling for beauty attenuates the estimated age twist by 15.2 percent in Zhaopin and by 7.5 percent in Computrabajo.

A final piece of evidence concerning the role of beauty in the age twist comes from *which* occupations have the largest age twists.

Fact 3: Looking across all four of our datasets, the age twist is especially high in administration, customer service and management occupations. The age twist is also significantly higher in occupations with higher-than-average beauty demands.

Appendix 5 shows the size of the age twist in each occupation in all four of our data sets. While the twist is positive in almost all occupations in all datasets, three occupations stand out as having consistently high twists: administration, customer service, and management.³¹ The disproportionate role of these occupations is consistent with the substantial role of contact, helping and managerial job title words in the XMRC data, because these title words appear very frequently in those occupations.³² Appendix 5 also reports the results of cross-occupation regressions in the Zhaopin and Computrabajo datasets. In both datasets, occupations with a higher demand for beauty exhibit significantly larger age twists ($p < .01$ in Zhaopin and $p < .10$ in Computrabajo).

In sum, data on employers' explicit requests for physically attractive applicants that is available in two of our four datasets suggests that requests for young, female beauty, especially in certain key occupations like administration and customer service, account for seven to 15 percent of the age twist. It is worth noting, however, that according to KS's model firms will not always explicitly request worker characteristics they believe are correlated with productivity: requests for *any* characteristic—including beauty—should only be made when employers' priors about its productivity effects are sufficiently extreme or precise. If that is the case, the true contribution of beauty preferences to the age twist, especially in contacting and helping jobs, may be considerably greater than is indicated by explicit beauty requests only.

³¹ The one exception is that management has a low twist in XMRC.

³² The five most frequent first job title words in XMRC's administration occupation are administration, front desk, HR, clerk and general; by far the most common first title word in the customer service occupation is customer service. Thus, these two occupations capture the dominant 'helping' word and the two most important 'contact' words in accounting for the age twist (see Figure 7).

4.5 Exploring the Within-Title Twist: Marriage, Parenthood and Labor Force Attachment

In section 4.3, we established that 65 percent of the age twist can be ‘explained’ by job titles and other ad characteristics. What factors might account for the remaining 35 percent, which occurs within detailed job titles? Having accounted for job characteristics, some promising candidates are worker characteristics like marriage, parenthood, and any related changes in labor force attachment, turnover and work intensity. Marriage and parenthood change substantially between the ages of 18 and 45 and might have differential effects on men’s and women’s attractiveness to employers.³³ To explore these ideas, this Section begins by establishing two key facts about employers’ expressed marital status preferences, and about the timing of the age twist in China and Mexico. These facts suggest a connection between the twist and parenthood. We conclude by asking whether this apparent connection might be linked to changes in men’s and women’s relative labor force attachment, turnover, and work intensity.

Fact 4: Explicit employer preferences for workers’ marital status are rare in China and Mexico. When they occur, however, their direction is highly correlated with explicit gender requests: employers ‘want’ men to be married, and women to be single.

³³ For example, a sizable literature on wages suggests that employers 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).

Information on employers' requests for marital status are only available in our Computrabajo data, where only 3.4 percent of ads explicitly indicate a preferred marital status.³⁴ That said, the pattern of these requests strongly suggests that employers prefer married men and single women. Specifically, among ads requesting men, Figure 9 shows that requests for married applicants outweigh requests for single applicants by a ratio of about six to one. In ads requesting women this pattern is reversed, with requests for single applicants outweighing those for married applicants by more than four to one.

Fact 4 suggests that the movement of men and women into marriage during the early career has the potential to account for some of the age twist. Quantitatively, however, adding controls for marital status requests into the Table 6 regressions has a negligible effect on the estimated twist, since these requests are so rare. As in the case of beauty requests, however, this may understate the true contribution of employers' marital status preferences to the twist if only the most extreme marital status preferences are actually written into job ads.

A second clue regarding the role of employers' marital status and parenthood preferences looks at the precise timing of the age twist, and takes advantage of the fact that Chinese women's entry into marriage and motherhood is highly concentrated into a specific age range.

Fact 5: In urban 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 urban Chinese women are having their first (and only) child.

³⁴ Requests for marital status and children generally occur in the detailed job description section of the ad, which we only have in the Zhaopin and Computrabajo datasets. Such requests appear to be almost nonexistent in Zhaopin. Explicit requests for workers with or without children also occur in the Computrabajo data, but are even rarer than marital status requests.

The three dashed lines in Figure 10 (a) and (b) 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 Chinese women in the 2005 one-percent Census microdata file. After reweighting, the timing of the age twist is quite comparable across our three Chinese datasets (especially the two more-skilled ones, XMRC and XMZYJS). Remarkably, most of the age twist is concentrated in the four years between age 26 and 30, with about 70 percent of gendered job ads for 26-year-olds requesting women compared to about 35 percent at age 30.³⁵

Panel (a) of Figure 10 compares the timing of this rapid shift in the gender mix of job ads to the ages at which urban Chinese women move into marriage in the 2005 Census sample described above. In panel (a), the blue (rightmost) solid line shows the share of urban Chinese women who are single in two-year age bins. Interestingly, urban Chinese women's movement into marriage is also highly concentrated into a narrow age range, with the fraction single falling from over 90 to under 30 percent between the ages of 20 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 marriage bar (Goldin 1991) or employer-aversion-to-married-women hypothesis: Chinese employers are still disproportionately requesting female employees (at around 70 percent of gendered ads in XMRC and Zhaopin) at

³⁵ The timing of both the age twist and women's movement into motherhood are much more gradual in Mexico than in China (see Appendix 5.5). While these two trends are 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 typical worker's life.

age 26, when a substantial majority of women are already married.³⁶ This timing problem becomes even more severe if we expect employers to avoid not just married women but women *who are likely to marry 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 still be single two years from now. To ensure they hire such women, firms would need to start shifting their ads away from women starting even earlier—at age 20.

Panel (b) of 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 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 reluctance to hire mothers relative to fathers may also contribute to the age twist. Notably, however, the share of ads requesting women does not recover at ages when children are in school: Given the timing of first births and China’s one-child policy, Appendix Figure A5.8 shows that the share of urban Chinese women *without* a child under the age of 6 begins to increase rapidly after age 34. If care for *pre-school-age* children explains the age twist, we should see increased employer interest in hiring women at these ages, which is not the case.

Is the age twist in employers’ gender requests between ages 26 and 30 in urban China accompanied by similarly pronounced changes in labor force activity at those ages? This question is explored in Figure 11 and Appendix 6.³⁷ Perhaps surprisingly, Figure 11 shows that

³⁶ Of course, an employer preference for married men could also explain the age twist. However, when we produced the equivalent to Figure 10 for men, the same timing problem was apparent.

³⁷ Parallel labor force activity data for Mexico is also provided in Appendix 6. Although, as noted, the timing of Mexico’s age twist is much more gradual than China’s, the main patterns are consistent in the sense that there are no pronounced changes in the gender gap in turnover,

while men's and women's employment rates rise during their 20s as many are completing school, both genders' employment rates are remarkably stable between the ages of 26 and 40. Men's employment rate holds steady at about 90 percent, and women's at 70 percent. Thus, the age twist in employers' advertised requests for female employees does not coincide with an exodus of mothers from the labor force.³⁸

If not labor force withdrawal, could the within-job-title age twist be connected to changes in other aspects of labor force activity, such as turnover and work intensity? Our best available measure of turnover in the Chinese Census is the share of employed workers with a contract of two years duration or less (suggesting a lower level of past and prospective job attachment than a longer contract), shown in Figure A6.2. While women are somewhat more likely to have such short-term contracts, and while employed men and women both transition out of such short-term contracts at a substantial rate as they age, the gender gap in the prevalence of such contracts is remarkably constant over the work life.³⁹

private sector share, and work hours during the period when the twist occurs. The one exception is for labor force participation rates: Women's participation is much lower in Mexico than China, and the gender gap in labor force participation does increase with age (Figure A6.7). That said, the change in women's relative participation seems small relative to the pronounced age twist in the gender mix of job ads.

³⁸ Figure A6.1 also rules out the possibility that women are leaving (or men entering) the private sector (to which all our job ad datasets are confined) around the time the twist is occurring: Chinese working women are actually more likely to work in the private sector than men, but the gender gap remains highly stable over the life cycle.

³⁹ It is also worth noting that the expected effects of turnover on the share of job ads requesting workers of a given gender are far from clear. On the one hand, employers may wish to avoid recruiting workers they expect to turn over quickly. On the other, workers who turn over rapidly need to be replaced more often.

If men become more committed to their jobs when they become fathers, and women less, we should expect to see sizable shifts in work hours around the age of parenthood. Figure A6.3 shows that this is not the case either. While employed, private-sector men report working 1-2 more hours per week than women do, both men and women work long hours —between 49 and 53 hours per week—, and the gender-hours gap is remarkably small and constant between the ages of 26 and 42. Related, Figure A6.4 shows that part-time work is extremely rare among both men and women in urban China, and A6.5 shows that large and relatively constant shares of both men and women work more than 50 hours per week.

Having ruled out large changes in labor force attachment and work hours, what other factors might account for the within-title age twist around the age that Chinese workers are getting married? One possibility is an effect of motherhood and fatherhood *per se* on workers' productivity. For example, men could be becoming more reliable workers as they become fathers, or simply as they enter their 30s. Women could devote less energy to their jobs at the same ages. Indeed, even though employed women's work hours do not change around the age of motherhood, time-use data show a large (nine hours per week) increase in Chinese women's unpaid work time around age 25 (Dong and An 2013, Table 4), which could affect employers' propensities to request female job applicants. Alternatively, customer and co-worker tastes could play a role. *Even within job titles*, co-workers and customers may expect to interact with a certain type of person (attractive young women selling clothes, older men supervising younger co-workers), and employers may be reluctant to violate those expectations. Such factors would also help explain the fact that employers' very low propensity to request women on all our job boards does not recover long after their children have reached school age.

We conclude this section by discussing one potential explanation of the age twist we cannot address directly. This explanation results from the fact that —like all vacancy-based data— our data only capture recruiting that occurs via the formal process of posting a job ad. Thus, while we are confident that our two Xiamen-based datasets provide a reasonably complete picture of the formal vacancies in a large Chinese city, our data cannot rule out the effects of a gender-specific change in the way firms recruit workers as they age. In this regard, it seems likely that the share of workers who are recruited via formal methods, i.e. by posting ads and soliciting applicants, declines as workers age: As they gain labor market experience, workers are likely to acquire more connections and to develop a reputation, and to substitute towards less formal job search methods that exploit these advantages. To explain the age twist, however, women’s use of informal job search networks would need to rise *more rapidly* with age than men’s. While this seems unlikely —if anything, we might expect women to acquire fewer ‘old boy’ connections as they age— we are not aware of any data that can conclusively rule it out.⁴⁰

5. Discussion

According to Dictionary.com, discrimination is “...making a distinction in favor of or against, a person or thing based on the group, class, or category to which that person or thing belongs rather than on individual merit.” According to that definition, the age- and gender-targeted job ads studied in this paper are an explicitly discriminatory practice. In this paper, we show that this practice is widely used in both China and Mexico, and establish some robust empirical facts concerning how it is used. For example, age- and gender-profiling is much more

⁴⁰ The U.S. Current Population Survey collects data on search methods, but only for unemployed workers (who we would expect to be highly selected into formal search). The NLSY97 collects detailed search method data for both employed and unemployed workers (see for example Kuhn and Mansour 2014), but only covers young workers.

common in less-skilled jobs—a pattern that is consistent with Kuhn and Shen’s screening-costs model of recruiting. This pattern also suggests that ongoing skill upgrading may reduce age- and gender-based profiling in emerging-economy labor markets.

Another robust feature of age- and gender profiling identified in this paper is a pronounced ‘twist’ in employers’ advertised gender requests away from women and towards men as workers age. According to our estimates, 65 percent of this age twist is associated with observed features of job ads, including the detailed descriptions of job ranks and duties contained in job titles. In particular, age-related shifts in the task mix of job ads, away from female-intensive tasks (like helping activities and customer contact) and towards male-intensive tasks (like higher management and driving) play a prominent role. Evidence on explicit employer requests for beauty suggests that the above ‘contacting’ and ‘helping’ effects reflect, at least in part, a preference for *attractive* young women. For example, in the Zhaopin data almost half (45.8 percent) of job ads aimed at women under the age of 30 explicitly request that the applicant be physically attractive, compared to 9.8 percent for men under 30. In Computrabajo, 22 percent of ads requesting young women request a photo, compared to 8 percent for young men.

The complement to the above finding is that 35 percent of the age twist occurs *within job titles*; this within-title twist is also quantitatively and statistically significant. While we cannot precisely decompose this unexplained or ‘pure’ age effect into portions associated with specific measured characteristics, two pieces of evidence suggest a connection to parenthood. First, in Mexico —where explicit marital status requests are sometimes made in job ads— employers who make marital status requests overwhelmingly request married men and single women. Second, in China the timing of the age twist in job ads coincides with the very narrow range of ages when women have their first child. Importantly, and somewhat surprisingly, Chinese

women do not exhibit marked changes in labor force attachment, turnover, or hours of work at these ages, however. This suggests that other factors may be at work, including men's increasing relative effort levels and reliability in market work, and customer- and co-worker tastes that vary with age *within* job titles.

Beyond documenting the widespread use of age- and gender-profiling in job ads and establishing key facts about how this practice is used, we believe that our paper has two broader implications for the study of gender differentials in labor markets. First, we document a powerful role of detailed job duties—as measured by job titles—in explaining employers' explicit gender preferences. In Table 7, for example, 85 percent of ads for front desk personnel explicitly request women, and 88 percent of ads for security personnel explicitly request men. This extreme, job-based heterogeneity in employers' gender preferences has implications for how gender preferences in the workplace should be conceptualized, and for how audit studies of gender preferences should be conducted and interpreted. Audit studies, which are widely viewed as a gold standard for measuring discrimination, often only pick a small share of relatively un-specialized job types to study.⁴¹ Thus, audit studies may fail to capture the high degree of job-specificity in employers' gender preferences: These preferences are decidedly *not* about liking or disliking women per se, as a pure Beckerian model might suggest. Gender preferences in the workplace are much more about strong and widespread preconceptions, shared by employers,

⁴¹ In addition to cost, a key reason for doing this is the practical difficulty of constructing plausible resumes for a large variety of jobs, many of which are highly specialized. Thus, among well-known studies, both Bertrand and Mullainathan (2004) and Kroft et al. (2013) restrict their attention to four occupations: sales, administrative support, clerical, and customer service. Oreopoulos (2011) considers more occupations, but restricts his attention to jobs that required three to seven years of experience and an undergraduate degree, and to applications that possess those qualifications. Lahey (2008) sends resumes only to entry-level jobs, defined as jobs requiring no more than one year of education plus experience.

workers and customers, about *what type of work is appropriate* for men and women of different ages. Audit studies will not reveal the structure of these detailed interactions unless they submit applications to a large and representative sample of jobs, including ‘gender-inappropriate’ ones. Thus, one advantage of job ad content studies like ours—which work with the universe of ads on a job board—is to reveal the strong interaction of employers’ preferences with the detailed job requirements that are stated in job titles.

Second, a key feature of the gender wage gap is the fact that it widens as workers age in almost all labor markets.⁴² While this gap in wage growth rates is frequently attributed to gender gaps in career interruptions and human capital investments, our paper suggests that another factor might also contribute to this phenomenon: the share of jobs that are actually ‘open’ to female applicants declines with age. As we have argued, possible causes for this decline include changing relative productivities of men and women, and customer and co-worker preferences for leadership and beauty that vary with the worker’s age. This paper suggests that the potential role of these factors in explaining women’s lower rate of age-related wage growth with age in most economies is worthy of additional exploration.⁴³

⁴² Figures A6.6 and A6.13 confirm that the gender (log) wage gap widens considerably with age in China and Mexico as well.

⁴³ While the age twist in the mix of gendered job ads is dramatic and pervasive, it is important to bear in mind that the twist’s effect on factors like age-wage profiles is tempered by the presence of nongendered jobs that are available to workers at different ages, especially at higher skill levels. That said, the extent to which non-gendered jobs really play such an ‘absorbing’ role depends on employers’ hiring practices in those jobs, as well as in the gendered jobs. Studying hiring (or at least callback) patterns in labor markets with gendered job ads is another promising direction for future research.

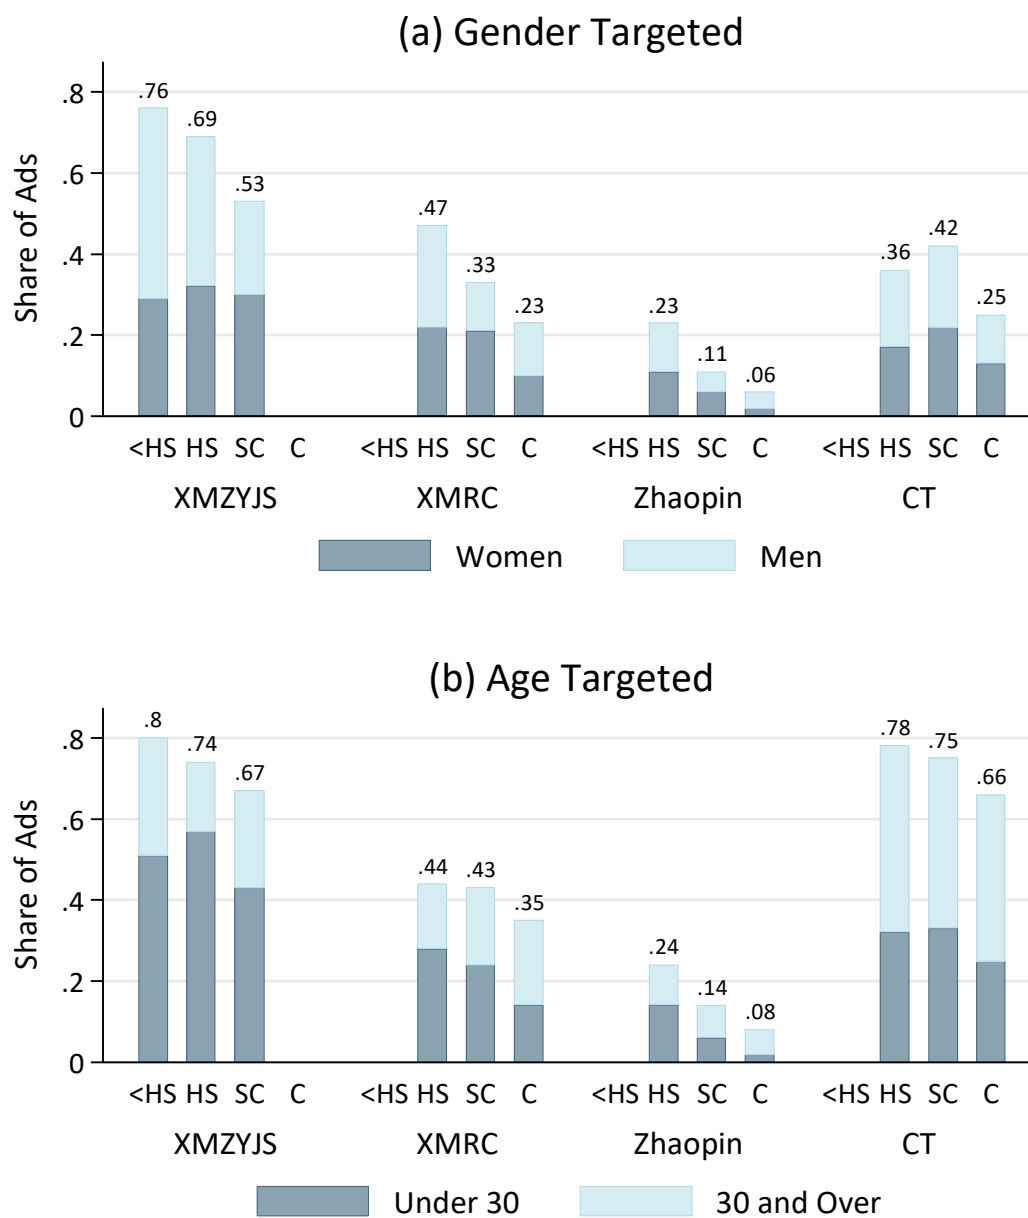
REFERENCES

- Ahituv, Avner and Robert I. Lerman (2007). "How Do Marital Status, Work Effort, and Wage Rates Interact?" *Demography* 44(3): 623-647 DOI: 10.1353/dem.2007.0021
- Altonji, Joseph G., and Charles R. Pierret. 2001. "Employer learning and statistical discrimination." *Quarterly Journal of Economics*, 116(1): 313-350.
- Biddle, Jeff E. and Daniel S. Hamermesh (1998). "Beauty, Productivity, and Discrimination: Lawyers' Looks and Lucre," *Journal of Labor Economics* 16(1): 172-201.
- Brencic, Vera (2010). "Wage Posting: Evidence from Job Ads," *Canadian Journal of Economics* 45(4): 1529-1559.
- Dong, Xiao-yuan and Xinli An. "Gender Patterns and Value of Unpaid Care Work: Findings From China's First Large-Scale Time Use Survey" *Review of Income and Wealth* (61) 3, September 2015: 540-560.
- Fernandez-Kranz, Daniel, Aitor Lacuesta, and Nuria Rodriguez-Planas (2013). "The Motherhood Earnings Dip Evidence from Administrative Records," *Journal of Human Resources* 48(1): 169-197.
- Gelbach, Jonah B. 2016. "When Do Covariates Matter? And Which Ones, and How Much?" *Journal of Labor Economics* 34(2.1): 509-530.
- Goldin, Claudia. 1991. "Marriage Bars: Discrimination Against Married Women Workers, 1920's to 1950's", in Rosovsky, H., D. Landes, P. Higgonet (eds.) *Favorites of Fortune: Technology, Growth and Economic Development Since the Industrial Revolution*. Cambridge, MA: Harvard University Press.
- Hamermesh, Daniel S. and Jeff E. Biddle (1994). "Beauty and the Labor Market," *American Economic Review* 84(5): 1174-1194.
- Kahn, Lisa B., and Fabian Lange. 2014. "Employer Learning, Productivity, and the Earnings Distribution: Evidence from Performance Measures." *The Review of Economic Studies*, 81(4): 1575-1613.
- Kapoor, Sasha and Arvind N. Magesan. "Having it Easy: Consumer Discrimination and Specialization in the Workplace" unpublished paper, University of Rotterdam, 2016.
- Korenman, Sanders and David Neumark (1991). "Does Marriage Really Make Men More Productive?" *Journal of Human Resources* 26(2): 282-307.
- Korenman, Sanders, and David 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* 128(1): 287-336.
- Kuhn, Peter and Hani Mansour. (2014) "Is Internet Job Search Still Ineffective?" *Economic Journal* 124 (581): 1213-1233.
- Lahey, Joanna N. (2008). "Age, Women, and Hiring: An Experimental Study," *Journal of Human Resources* 43: 30-56.

- Loughran, David S. and Julie Zissimopoulos (2009). "Why Wait? The Effect of Marriage and Childbearing on the Wages of Men and Women," *Journal of Human Resources* 44(2): 326-349.
- Mobius, Markus M. and Rosenblat, Tanya 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 (2015). "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.

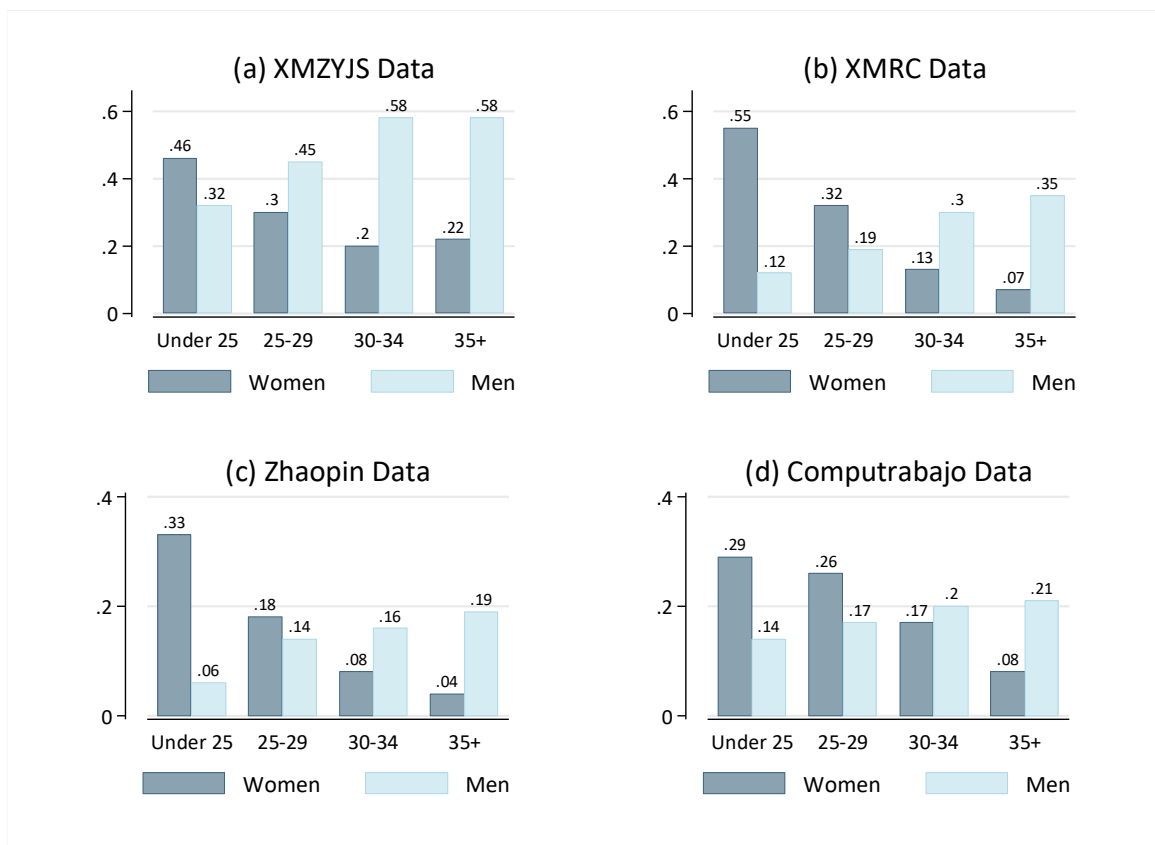
FIGURES

**Figure 1: The Negative Skill-Targeting Relationship—
Share of Ads That Are Age and Gender Targeted, by Desired Education**

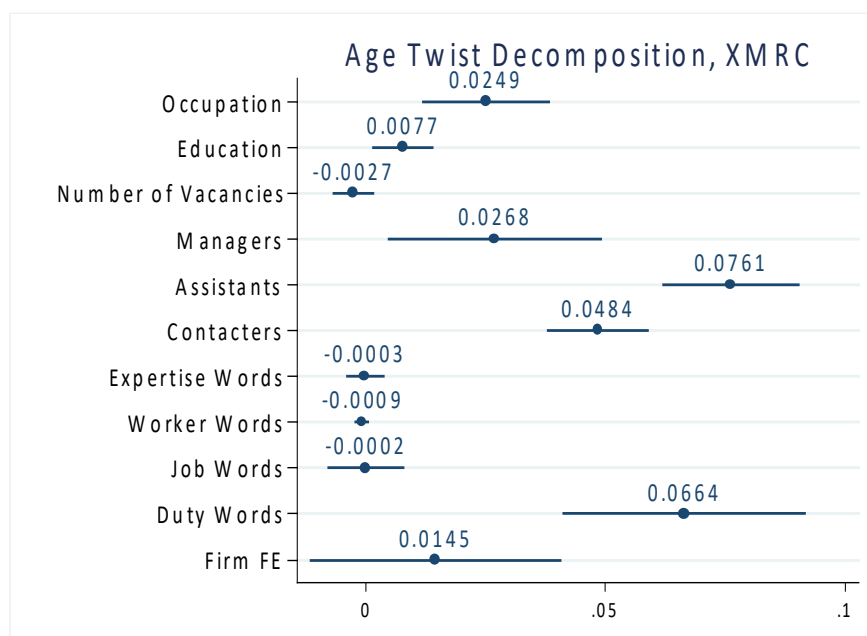


Notes: <HS, HS, SC, and C denote less than high school, high school, some college, and college or more respectively. CT refers to the Computrabajo data.

**Figure 2: The Age Twist—
Share of Ads Requesting Women and Men, by Desired Age**



Notes to Figure 2: The share of job ads in each of the four age groups (under 25, 25-29, 30-34 and 35+) are .30, .39, .26 and .06 in XMZYJS; .16, .41, .30 and .12 in XMRC; .12, .30, .33 and .25 in Zhaopin; and .07, .33, .41 and .19 in Computrabajo.

Figure 3: ‘Explained’ Components of the Age Twist, XMRC

Contributions are measured in percentage points, and are computed using the omitted variable bias formula from columns 1 and 6 in Table 9. The total explained age twist equals .0261. Lines indicate 95% confidence intervals.


Figure 4: Sources of the Occupation Mix component of the Age Twist, XMRC.

Figure 4 decomposes the occupation mix component in Figure 3 (.025 percentage points in total) into four types of occupations. Individual occupations making the ten (three) largest *absolute* contributions are:

-growing male occupations ($\gamma > 0, d > 0$): **construction (.0091)**, **manufacturing (.0034)**, logistics (.0025), management (.0010)

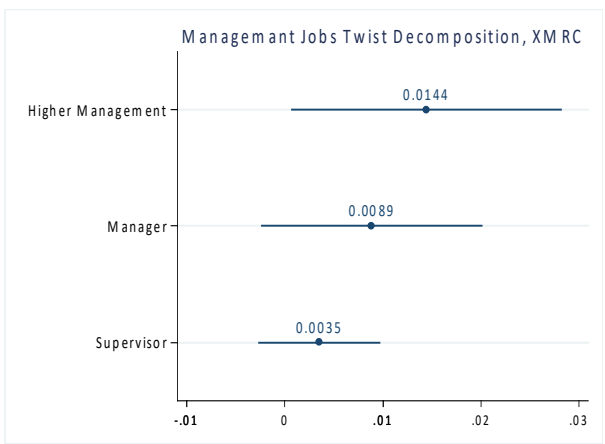
-shrinking male occupations ($\gamma > 0, d < 0$): software (-.0023)

-growing female occupations ($\gamma < 0, d > 0$): home service (-.0022), finance/accounting (-.0014)


-shrinking female occupations ($\gamma < 0, d < 0$): **administration (.0123)**, sales (.0028), international trade (.0020)

Figure 5: Sources of the Managers, Assistants and Contacters Components of the Age Twist

a) Managers:



b) Assistants:



c) Contacters:

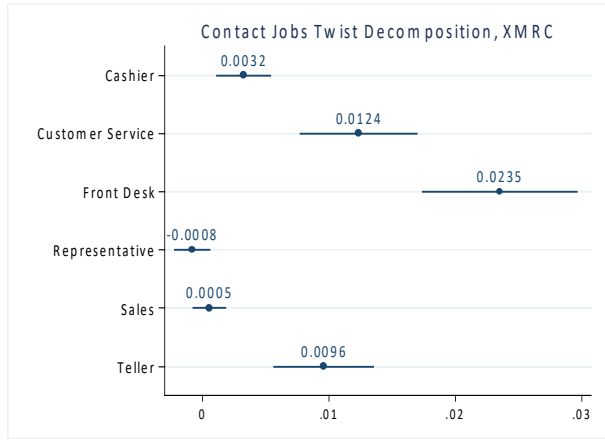


Figure 5 allocates the contributions of the three components (.027, .076, and .048 percentage points respectively) to the more detailed sub-components shown.

Figure 6: Sources of the ‘Duty Words’ Component of the Age Twist, XMRC

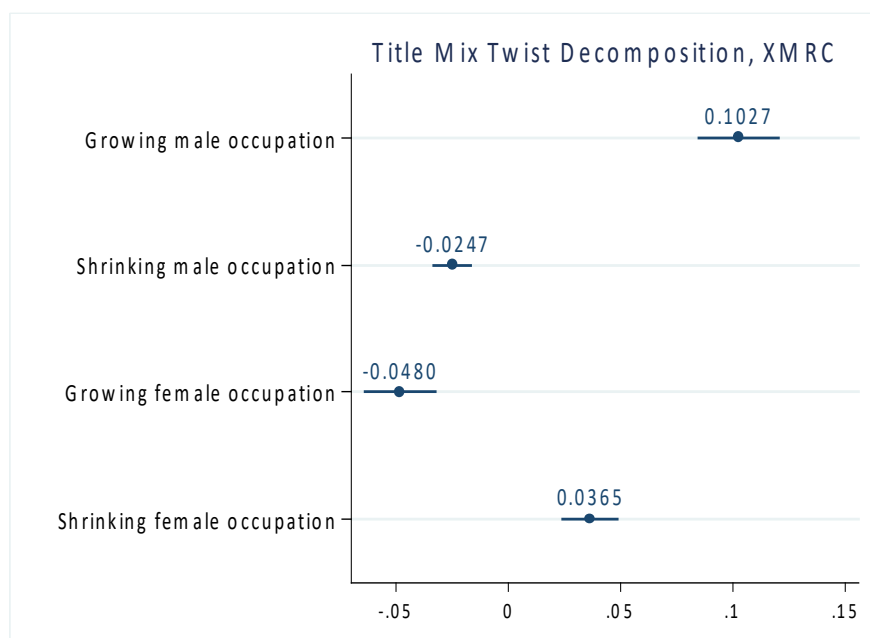


Figure 6 decomposes the ‘detailed title words’ component in Figure 3 (.067 percentage points in total) into four types of words. Individual words making the twenty (**five**) largest *absolute* contributions are:

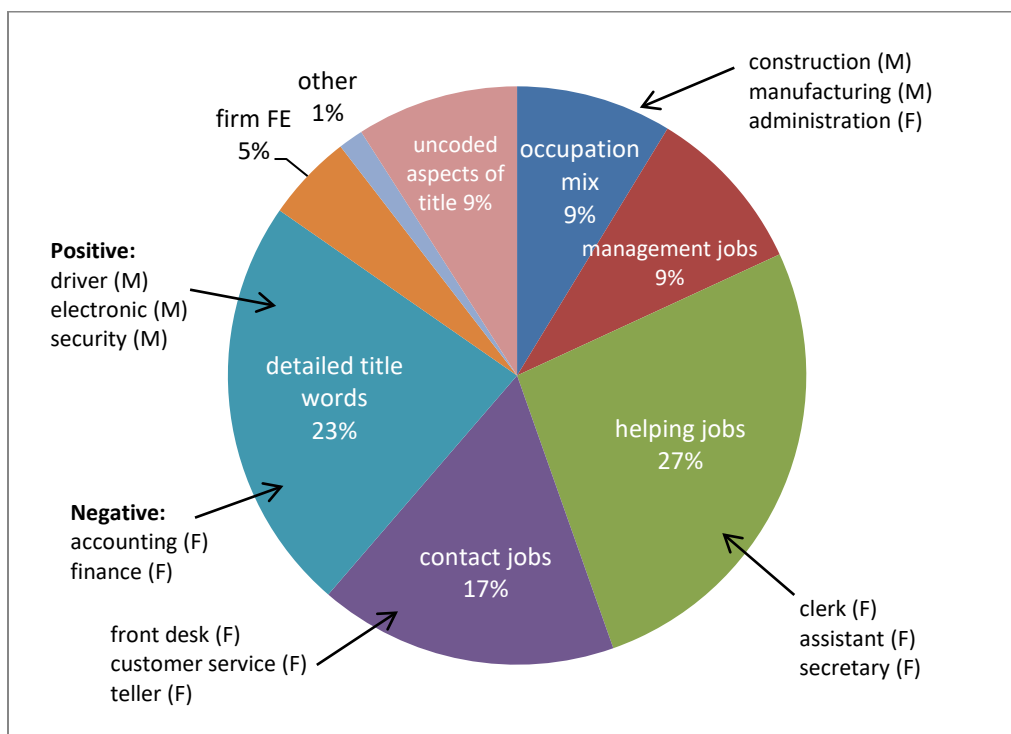
growing male words ($\gamma > 0, D > 0$): **driver (.0293)**, **electronic (.0135)**, **security (.0126)**, engineer (.0052), chef (.0046), high (.0031) plumber (.0029), molding (.0026), fitter (.0023), equipment (.0022), machine (.0020)

shrinking male words ($\gamma > 0, D < 0$): network (-.0026), warehousing (-.0025),

growing female words ($\gamma < 0, D > 0$): **accounting (-.0109)**, **finance (-.0060)**, cleaning (-.0050), division (-.0044) voltage (-.0041)

shrinking female words ($\gamma < 0, D < 0$): statistician (.0049), cuisine(.0021).

Figure 7: Sources of the Explained Portion of the Age Twist, XMRC




The explained portion of the age twist (comparing columns 1 and 8 of Table 8) is $44.2 - 15.5 = 28.7$ percentage points, or 65 percent of the unadjusted twist. Figure 7 summarizes the main factors contributing to this explained twist. “Uncoded aspects” refers to the difference between fixed effects for every unique title (Table 8) and Table 9, which codes ranks and individual words in the titles.

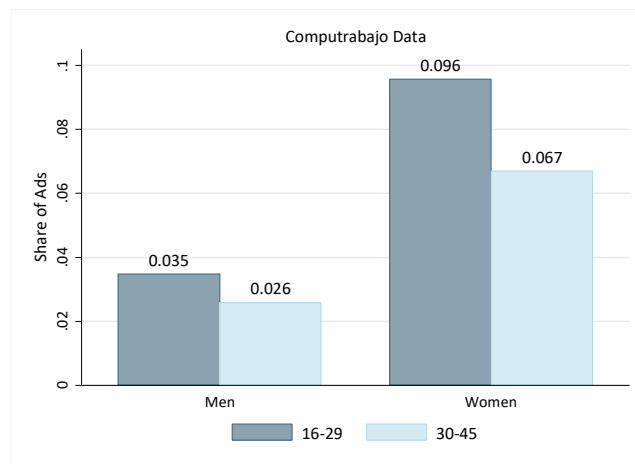
Arrows show the top three absolute occupation contributions, the top five absolute title-words contributions, and the most important helping and contact job words. F and M denote ‘female’ and ‘male’ occupations or title words respectively ($\gamma > 0$). Title words make *positive* contributions to explaining the twist when $\gamma D > 0$. Thus ‘growing’ (with age) ($D > 0$) male title words like ‘driver’ and ‘electronic’ help account for the twist, while growing female title words like accounting and finance actually contribute negatively to the twist.

Figure 8: Beauty requests by age and gender

a) Share of ads requesting beauty, Zhaopin data



b) Share of ads requesting beauty, Computrabajo data:



c) Share of ads requesting a photo, Computrabajo data

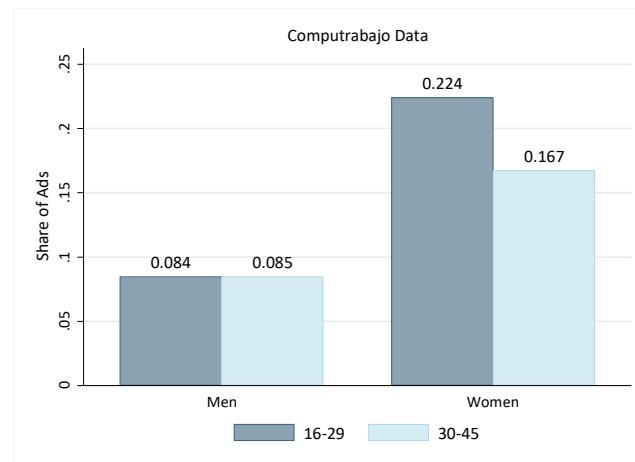


Figure 9: Share of Ads Requesting Single and Married Applicants, by Requested Age and Gender (Computrabajo Data)






Figure 10: Timing of Marriage and Motherhood relative to the Age Twist, China

a) Female Share of Gendered Job Ads versus Share of Women who are *Single*




Solid lines show the share of urban Chinese women who are single at that age, or who will still be single two years later.

b) Female Share of Gendered Job Ads versus Share of Women who are *Childless*



Solid lines show the share of urban Chinese women who are childless, or who will still be childless two years later.

Dashed lines in both parts of the Figure show the share of gendered job ads that request women in each job board. Marital status, education and fertility information is for urban women in the 2005 Census 1% microdata sample.

Figure 11: Employed Share of the Population, China

Lines show the share of urban Chinese residents who are employed.

Source: Authors' calculations from the Chinese 2005 1% Census.

TABLES

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		0.373	0.457	0.151
University	0.064	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	1,810	2,556	4,279	7,640
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
Beauty Indicators				
Ad requests beauty?	n/a	n/a	0.077	0.040
Photo required with application	n/a	n/a	n/a	0.109
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.

TABLE 2: Sample Means: Employed, Private Sector Population in Xiamen, Urban China, and Urban Mexico

	(1)	(2)	(3)
Ad Characteristics	Xiamen	Urban China	Urban Mexico
Gender			
Male	53%	54%	64%
Female	47%	46%	36%
Education			
Junior middle school or less	61%	61%	49%
High school or Tech school	25%	24%	29%
Some postsecondary	8%	9%	21%
University	5%	6%	21%
Potential Experience			
under one year	1%	1%	5%
1-3 years	12%	8%	6%
4-5 years	10%	9%	7%
More than 5 years	77%	81%	81%
Age (years)	29.5	30.5	31.4
Monthly wage (RMB or Pesos)	2,242	2,203	6,793
Number of observations	1,057	86,950	465,039

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.

Potential experience equals age minus years of schooling minus 6.

Chinese data are from the 2005 Census, 1% sample, and refer to persons aged 16-45, currently living in urban regions, who are currently employed in the private sector (i.e. excluding SOEs, government and collectives). "Urban China" comprises the four municipalities directly under the jurisdiction of the central government (Beijing, Shanghai, Tianjin and Chongqing) plus the 15 subprovincial cities: Changchun, Chengdu, Dalian, Guangzhou, Hangzhou, Harbin, Jinan, Nanjing, Ningbo, Qingdao, Shenyang, Shenzhen, Wuhan, Xiamen, and Xi'an. Chinese wages have been adjusted for per capita GDP growth between 2005 and 2010 using IMF GDP statistics.

Mexican data are from the 2010 Census microdata file, and refer to persons aged 16-45 who are currently employed in the private sector (i.e. excluding legislative bodies, federal, state, and municipal administration, international organizations, public schools, public hospitals, central banking, and other public sector services). "Urban Mexico" is defined as areas ("localidades") with more than 100,000 inhabitants.

TABLE 3: Effects of Desired Age on Firms' Gender Preferences ($P^M - P^F$), XMZYJS DATA

	(1)	(2)	(3)	(4)
Age 25-29	0.2903*** (0.0170)	0.1117*** (0.0147)	0.1337*** (0.0166)	0.1035*** (0.0207)
Age 30-34	0.5201*** (0.0228)	0.2304*** (0.0192)	0.2243*** (0.0214)	0.1795*** (0.0268)
Age 35+	0.5019*** (0.0259)	0.2362*** (0.0232)	0.2440*** (0.0312)	0.2206*** (0.0385)
Job Requires High//Tech School		0.0381*** (0.0118)	-0.0066 (0.0132)	0.0081 (0.0169)
Job Requires College		-0.0091 (0.0218)	-0.0620** (0.0308)	-0.0312 (0.0378)
Fixed effects	None	district, occupation	occ, firm	occ*firm
number of groups	0	63	7,590	22,173
N	108,875	108,875	108,875	108,875
Adjusted R2	0.055	0.211	0.288	0.335

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

Column 2 controls for log firm size and firm ownership type. Columns 2-4 control for number of positions advertised.

TABLE 4: Effects of Desired Age on Firms' Gender Preferences ($P^M - P^F$), XMRC DATA

	(1)	(2)	(3)	(4)
Age 25-29	0.2935*** (0.0171)	0.2293*** (0.0161)	0.2861*** (0.0224)	0.2521*** (0.0443)
Age 30-34	0.5905*** (0.0182)	0.4890*** (0.0175)	0.5613*** (0.0248)	0.5296*** (0.0502)
Age 35+	0.7099*** (0.0216)	0.5915*** (0.0216)	0.6865*** (0.0325)	0.6820*** (0.0618)
Job requires some postsecondary education		-0.0935*** (0.0118)	-0.1219*** (0.0169)	-0.1210*** (0.0334)
Job requires university		0.0077 (0.0184)	-0.0875*** (0.0271)	-0.1312** (0.0536)
Fixed effects	none	occupation	occ, firm	occ*firm
number of groups	0	36	4,412	9,739
N	17,021	17,021	17,021	17,021
Adjusted R2	0.105	0.264	0.512	0.767

*** p<0.01, ** p<0.05, * p<0.1. Standard errors (in parentheses) are clustered by firm. Column 2 controls for log firm size and firm ownership type. Columns 2-4 control for number of positions advertised.

TABLE 5: Effects of Desired Age on Firms' Gender Preferences ($P^M - P^F$), ZHAOPIN DATA

	(1)	(2)	(3)	(4)
Age 25-29	0.2331*** (0.0363)	0.1584*** (0.0186)	0.1660*** (0.0184)	0.1100*** (0.0326)
Age 30-34	0.3491*** (0.0283)	0.2564*** (0.0193)	0.2794*** (0.0200)	0.2062*** (0.0385)
Age 35+	0.4186*** (0.0285)	0.3098*** (0.0207)	0.3389*** (0.0212)	0.2861*** (0.0479)
Job requires some postsecondary education		-0.0216* (0.0129)	-0.0895*** (0.0141)	-0.0825*** (0.0175)
Job requires university		0.0171 (0.0148)	-0.0688*** (0.0174)	-0.0506** (0.0233)
Fixed effects	none	occ, ind, province	occ, province, firm	occ*firm, province
Number of groups	0	116	24,428	50,048
N	134,100	134,100	134,100	134,100
Adjusted R2	0.059	0.186	0.454	0.701

*** p<0.01, ** p<0.05, * p<0.1. Standard errors (in parentheses) are clustered by firm. Column 2 controls for log firm size and firm ownership type. Columns 2-4 control for number of positions advertised.

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

	(1)	(2)	(3)	(4)
Age 25-29	0.0611** (0.0287)	0.0494* (0.0261)	0.0596** (0.0265)	0.0686** (0.0280)
Age 30-34	0.1812*** (0.0195)	0.1433*** (0.0203)	0.1470*** (0.0198)	0.1541*** (0.0212)
Age 35+	0.2772*** (0.0234)	0.2459*** (0.0234)	0.2685*** (0.0216)	0.2671*** (0.0208)
Job requires some postsecondary education		-0.0473*** (0.0174)	-0.0433** (0.0179)	-0.0448** (0.0192)
Job requires university		-0.0557*** (0.0154)	-0.0526*** (0.0155)	-0.0532*** (0.0173)
Fixed effects	none	occ*state	occ*state, firm	occ*firm, state
Number of groups	0	425	2,054	5,774
N	65,516	65,516	65,516	65,516
Adjusted R2	0.020	0.127	0.207	0.302

*** p<0.01, ** p<0.05, * p<0.1. Standard errors (in parentheses) are clustered by firm. Column 2 controls for log firm size and firm ownership type.

TABLE 7: Most-female, most-male and least-gendered Job Titles, XMRC

Most-female job titles			Most-male job titles			Least-gendered job titles			
Job title	Share female	<i>N</i>	Job title	Share male	<i>N</i>	Job title	Share not gendered	<i>N</i>	
1	front_desk	0.849	305	goods	0.879	33	designer	0.933	30
2	clerk	0.815	211	security	0.877	138	advertising	0.844	32
3	teller	0.772	189	machine	0.875	40	website	0.806	31
4	secretary	0.763	38	equipment	0.844	32	two_dimension	0.806	103
5	cashier	0.744	39	electronic	0.816	125	recruiting	0.778	36
6	document	0.683	41	driver	0.755	220	business_person	0.775	160
7	statistician	0.655	58	mold	0.667	75	planning	0.763	76
8	assistant	0.641	39	injection	0.654	78	product	0.759	54
9	office	0.639	72	milling	0.645	31	senior	0.755	53
10	administration	0.578	431	technician	0.580	50	marketing	0.752	250
11	customer_service	0.576	165	mechanical	0.576	33	investment_attraction	0.744	39
12	merchandiser	0.545	33	engineer	0.529	34	vice	0.738	42
13	accounting	0.498	267	service	0.511	47	training	0.690	58
14	manager	0.474	76	engineering	0.467	30	customer	0.646	48
15	finance	0.459	233	warehouse	0.459	418	telephone	0.645	31
16	general	0.451	273	chef	0.433	30	design	0.627	67
17	HR	0.444	417	electrical	0.412	34	custom	0.618	34
18	shop	0.423	52	construction	0.395	43	project	0.608	51
19	international_trade	0.421	242	material	0.361	36	QC	0.591	340
20	business	0.401	367	technical	0.360	50	cleaning	0.590	39
21	cleaning	0.385	39	manufacture	0.310	323	internet	0.589	112
22	procurement	0.355	211	project	0.294	51	electrical	0.588	34
23	telephone	0.355	31	internet	0.268	112	chief	0.577	78
24	operating	0.354	79	production	0.261	69	international_trade	0.545	242
25	customer	0.333	48	QC	0.212	340	shop	0.538	52

Job titles refer to the first word of the title. Share female (male) is the share of job ads with that title explicitly requesting female (male) applicants. Share not gendered is the share not requesting a specific gender. *N* = number of ads. Sample restricted to titles with at least 30 ads.

Table 8: Regressions for $P^M - P^F$ using all the Table 4 controls, plus fixed effects for the first k words of the raw job title, XMRC

	(1) No controls	(2) Table 4 Controls (col. 2)	(3) Col. 2 plus first title word	(4) First Two Words	(5) First Three Words	(6) First Four Words	(7) Entire Title	(8) Entire Title plus Firm FE
Age30plus	0.4420*** (0.0148)	0.3723*** (0.0138)	0.2574*** (0.0285)	0.1978*** (0.0214)	0.1703*** (0.0204)	0.1650*** (0.0206)	0.1653*** (0.0207)	0.1548*** (0.0197)
Some Post-Secondary		-0.0928*** (0.0150)	0.0557** (0.0256)	0.0655*** (0.0240)	0.0603*** (0.0230)	0.0612*** (0.0233)	0.0611*** (0.0234)	0.0526** (0.0209)
University		0.0350 (0.0244)	0.1750*** (0.0470)	0.1904*** (0.0340)	0.1776*** (0.0305)	0.1754*** (0.0313)	0.1746*** (0.0314)	0.1329*** (0.0354)
Observations	9,884	9,884	9,884	9,884	9,884	9,884	9,884	9,884
R-squared	0.0891	0.2678	0.4466	0.5429	0.5763	0.5803	0.5804	0.7853

*** p<0.01, ** p<0.05, * p<0.1. Standard errors (in parentheses) are clustered by occupation.
High-dimensional two-way fixed effects in column 8 are computed using Stata's felsdreg.do command.

Notes: Columns 3-6 successively add fixed effects for the first $k = 1, \dots, 4$ words of the job title. Columns 7 and 8 have fixed effects for every title, regardless of length.

TABLE 9: Effects of Desired Age on Firms' Gender Preferences ($P^M - P^F$) controlling for words used in job titles, XMRC

	(1)	(2)	(3)	(4)	(5)	(6)
Age30plus	0.4420*** (0.0148)	0.3723*** (0.0138)	0.2758*** (0.0143)	0.2727*** (0.0143)	0.1888*** (0.0138)	0.1811*** (0.0228)
Some Post-Secondary		-0.0928*** (0.0150)	-0.1014*** (0.0141)	-0.0920*** (0.0145)	0.0443*** (0.0151)	0.0272 (0.0254)
University		0.0350 (0.0244)	0.0341 (0.0239)	0.0423* (0.0241)	0.1592*** (0.0247)	0.1102*** (0.0413)
Managers			-0.0339** (0.0165)			
Assistants			-0.4751*** (0.0214)			
Contacters			-0.4146*** (0.0210)			
Experts			-0.0024 (0.0200)			
Higher Management				-0.0012 (0.0318)	0.1472** (0.0719)	0.2265** (0.1083)
Manager				-0.0453** (0.0184)	0.0922*** (0.0237)	0.0571 (0.0368)
Supervisor				-0.1771*** (0.0392)	0.0855** (0.0422)	0.0731 (0.0647)
Assistant				-0.4300*** (0.0265)	-0.3148*** (0.0293)	-0.3329*** (0.0450)
Clerk				-0.5750*** (0.0249)	-0.4919*** (0.0261)	-0.5094*** (0.0452)
Secretary				-0.5724*** (0.0528)	-0.5331*** (0.0546)	-0.6560*** (0.0891)
Helper				0.4207 (0.4218)	1.8063*** (0.2790)	2.2749*** (0.2953)
Front Desk				-0.3831*** (0.0285)	-0.3520*** (0.0296)	-0.4156*** (0.0507)
Customer Service				-0.3848*** (0.0396)	-0.3628*** (0.0431)	-0.4750*** (0.0713)
Teller				-0.4494*** (0.0382)	-0.5798*** (0.0457)	-0.5576*** (0.0748)
Cashier				-0.5554*** (0.0762)	-0.5346*** (0.0777)	-0.4461*** (0.1267)
Sales				-0.2949*** (0.1008)	-0.1635 (0.1271)	-0.2291 (0.2186)
Representative				0.0954 (0.0627)	0.1711*** (0.0624)	0.1279 (0.0980)
Basic Ad Characteristics		Yes	Yes	Yes	Yes	Yes
Occupation Fixed Effects		Yes	Yes	Yes	Yes	Yes
All Job Title Duty Words					Yes	Yes
Firm Fixed Effects						Yes
Observations	9,884	9,884	9,884	9,884	9,884	9,884
R-squared	0.0891	0.2678	0.3334	0.3441	0.4868	0.7150

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

Notes:

Basic ad controls are education requested, number of vacancies, firm size and location.

Occupation fixed effects are for 36 job categories used by XMRC.

Columns 4-6 also control for the miscellaneous *worker words* and *job words* extracted from the title.

Columns 5 and 6 include fixed effects for all 515 remaining title words, describing duties. Of these, 374 were retained in the regression after eliminating collinear words.