NBER WORKING PAPER SERIES

AGE DISCRIMINATION IN HIRING:
EVIDENCE FROM AGE-BLIND VS. NON-AGE-BLIND HIRING PROCEDURES

David Neumark

Working Paper 26623
http://www.nber.org/papers/w26623

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2020

I am grateful to Ian Burn for helpful comments, research assistance, and collaboration on many projects related to age discrimination. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w26623.ack

Age Discrimination in Hiring: Evidence from Age-Blind vs. Non-Age-Blind Hiring Procedures
David Neumark
NBER Working Paper No. 26623
January 2020
JEL No. J14,J7

ABSTRACT

I study age discrimination in hiring, exploiting a difference between age-revealed and partially age-blind hiring procedures. Under the first hiring procedure, age is revealed simultaneously with other applicant information and job offer rates are much lower for older than for younger job applicants. Under the second hiring procedure, interview selections are based on detailed, age-blind on-line applications, while subsequent interviews are not age-blind. Older applicants are not under-selected for interviews, but after in-person interviews when age is revealed, older applicants still face a much lower job offer rate. This evidence is strongly consistent with age discrimination in hiring.

David Neumark
Department of Economics
University of California, Irvine
3151 Social Science Plaza
Irvine, CA 92697
and NBER
dneumark@uci.edu

“Sometimes, not knowing what you're doing allows you to do things you never knew you could do.”

– Nell Scovell

I. Introduction

Economists studying discrimination in hiring increasingly rely on audit or correspondence studies. Audit studies use actual applicants coached to act alike, and measure discrimination as differences in job offers; correspondence studies create fake applicants (on paper, or electronically), and measure discrimination as differences in “callbacks” for job interviews. In both types of studies, group membership of the applicants is varied randomly, and hence is independent of other applicant characteristics. These field experiments are generally viewed as the most reliable means of inferring labor market discrimination (e.g., Fix and Struyk, 1993), although they are, of course, not infallible.[1] For example, Neumark and Rich (2019) show that a number of field experiments of labor market discrimination appear to overstate discrimination, owing to an econometric problem first identified by Heckman and Siegelman (1993).[2]

In contrast to these field experiment studies, non-experimental studies of labor market discrimination use the often rich sets of control variables available (at least in some data sets) to try to account for potential productivity differences between groups. This approach has been most commonly applied to the study of wage discrimination (see Altonji and Blank, 1999), but it has also been applied to hiring discrimination (e.g., Holzer, 1998). The obvious challenge in using non-experimental data is that we may not adequately control for differences between groups, hence finding spurious evidence of discrimination (or, conceivably, spurious evidence of non-discrimination, depending on the unobservables).

In this paper, I take a different approach, in the context of age discrimination in hiring. I use non-

[1] For a review of this experimental evidence, see Neumark (2018).
[2] Heckman and Siegelman show that when the groups in question have different variances of the unobservables, audit/correspondence study methods can give very misleading results. Neumark (2012) develops a solution to this problem, which is the basis of the analysis in Neumark and Rich (2019).

experimental data – stemming from a lawsuit over age-based hiring discrimination for a restaurant chain that opened locations throughout the country.[3] However, I exploit a difference in hiring procedures that alters when those responsible for hiring become aware of the age of applicants. Because hiring managers cannot act on a worker's age until they become aware of it, I use evidence on age differences in outcomes under the different hiring procedures to identify age discrimination.

Under the first hiring procedure, age is revealed at the same time as all other information about the applicant (because everything is done in-person). Under this non-age-blind hiring procedure, job offer rates for older applicants (aged 40 and over) are substantially lower than for younger applicants (by 68%, on a baseline job offer rate for younger applicants of 14.1%). Under the second hiring procedure, selections for interview are made on-line, and are age-blind. Subsequently, those selected are interviewed. Under this procedure, older applicants are selected for interviews – the age-blind hiring procedure – at equal or higher rates to younger applicants. However, after the interview – where age is revealed in the same way as for the paper, in-store applicants – the job offer rate is much lower for older applicants (by 40%, on a baseline job offer rate among interviewees of 30% for younger applicants). The net result is a lower job offer rate for older applicants (by 46%, on a baseline hiring rate of 10% for younger applicants). However, the age difference in treatment only occurs at the non-age-blind stage of the hiring process.

The comparison of outcomes by age in these two hiring procedures, and in the two stages of the second procedure, provide clean evidence on age discrimination in hiring. With data from only the first procedure (paper applications), the evidence could face the same challenge as other evidence based on standard non-experimental – that the age difference could reflect unobservables associated with age. But this alternative interpretation is contradicted by what happens under the second hiring procedure. Recruiters do not under-select older applicants when have detailed information related to their qualifications and experience, but *do not have* information on applicants' ages; but they do under-select

[3] All results reported in this paper are from reports filed in the case, although I cannot reveal the identities of those involved. The same confidentiality restrictions preclude me from sharing the data.

once age is revealed at the interview stage. I also provide some corroborating evidence that customer discrimination at least partly underlies the evidence of age discrimination.

II. Related Prior Research

The evidence I obtain based on differences in the observability of age to those making hiring decisions parallels some other key papers on discrimination. Goldin and Rouse (2000) is a non-experimental study of how the switch to blind auditions for major orchestras affected the selection of female auditionees. The variation in this study arises from the adoption by orchestras, over time, of blind auditions where the musician plays behind a screen and other steps are taken to ensure that the musician's identity is not known when selection decisions are made. The authors find that the selection of females increased because of blind auditions, suggesting that there was discrimination against women prior to the adoption of blind auditions.[4] Paralleling the argument in the Introduction, Goldin and Rouse cast their study as complementary to audit or correspondence studies of discrimination.

Åslund and Nordströum Skans (2012) also find evidence consistent with anonymization helping applicants who might experience discrimination. They study anonymized job application procedures in non-experimental data in Sweden, and find increased interviews and job offers for women, but for ethnic minorities only increased interviews.

Two recent experimental studies manipulate the information about applicants, with the experimental treatment being to anonymize applicants and see if selection for interviews is affected. Interestingly, these studies sometimes find that anonymization does not increase interviews/hiring of the group against which discrimination might be thought to occur, which might be interpreted as evidence against a finding of discrimination.

Krause et al. (2012) studied economics Ph.D. applicants to a European research institute, randomizing the anonymization of demographic information before the applications went to the hiring

[4] The use of blind auditions is chosen by the orchestra, and the authors are careful to rule out selection on using blind auditions as an explanation of their findings.

committee.[5] The key result is that female applicants were more likely to be invited for interviews in the *non-anonymous* sample, whereas this advantage was erased in the anonymous sample. Taken literally, this means that there was discrimination in favor of female candidates in the non-anonymous setting, which could not occur with anonymous applications. On the other hand, the hiring committee was aware of the experiment, and this could have motivated them to be non-discriminatory in evaluating the non-anonymous applications. There was, however, no such pattern regarding non-Western applications.

A second application of this method, to hiring of minority job candidates in France, found that minorities fared worse under anonymization, getting a smaller share of interviews (Behaghel et al., 2015). (The hiring gap also widened, but not as much as the interview gap.) The authors suggest that this happens because anonymization prevents firms from downweighting, for minority applicants, negative characteristics of job applicants that are associated with minority group membership.[6] Behaghel et al. note that the firms that agreed to participate were similar on most observables, except that they hired more minorities. When they joined the study, then, those that received anonymous applications may have been unable to continue preferential hiring of minorities. Under this interpretation, extending anonymization to all employers would have ambiguous effects, depending on the extent to which non-participating employers engage in discrimination against minorities, and anonymization prevents this.

Although we generally think of experimental evidence as superior, in these two studies, at least, experimenter effects may contaminate the results. Thus, it is possible that non-experimental data on differences in outcomes depending on the information available about applicants could – in some circumstances at least – be more reliable.

Another potential advantage of my evidence relative to experimental studies is that the latter rely almost exclusively on differences in callback rates (based on correspondence studies, rather than audit studies). Some past research has found that evidence on callback rates is predictive of discrimination at

[5] After the hiring process was complete, applicants were asked for permission to use their data in the study; 65 percent agreed.
[6] They find evidence consistent with this interpretation, based on ratings of applicants in both treatments by counselors from the public employment service.

the hiring stage. Studies of ethnic discrimination by the International Labor Organization (ILO), discussed in Riach and Rich (2002), provide estimates of differences at the selection for interview stage and the job offer stage, and find that around 90 percent of the discrimination that is detected occurs at the selection for interview stage. And Neumark (1996) finds similar evidence in an audit study of sex discrimination that also included a callback stage. However, a recent study of discrimination in France against North Africans suggests that correspondence study evidence on callback rates may sometimes fail to detect evidence of discrimination (Cahuc et al., 2019). This conclusion is based on evidence that the observed callback rate is lower only in the private sector, while discriminatory preferences and beliefs are similar in the two sectors – which the authors interpret as similarly discriminatory hiring rates in both sectors.[7] The evidence in the current paper circumvents this issue by looking at actual hiring behavior.

In a combination of experimental and non-experimental methods, Agan and Starr (2018) study the effects of ban-the-box initiatives that change whether employers are aware of criminal backgrounds, to identify the effect of information about workers' criminal backgrounds on black-white differences in callbacks for interviews in a correspondence study. This study finds that restricting information on criminal background seems to have increased statistical discrimination against blacks. The change in information from ban-the-box initiatives parallels, in the present paper, the differences in information available under the alternative hiring procedures.

III. Hiring Procedures

The company used two different hiring procedures, switching from a paper application process completed in stores to an on-line application process during the sample period. This generated variation across restaurants largely based on when they opened (and it is the hiring associated with the restaurant openings that I study).

Under one hiring procedure ("paper applications" completed in stores), applicants for employment submit a paper application to company personnel in person and go through a pre-screening

[7] The authors do not have an independent estimate of discrimination at the hiring stage, but cite other evidence on similar rates of under-representation of North Africans in the private and public sectors.

interview. This pre-screening interview is based on a quite limited set of questions. Based on the pre-screening interview, an applicant is or is not forwarded on in the process. If they do advance in the process, the applicant could then go through a second interview and final interview with a company manager. There is no data on the interview selection step, but rather just information on one distinct outcome – whether a job was ultimately offered.

From the point of view of information about age, the key feature of the paper application hiring procedure is that the age of the applicants is observed early in the process – insofar as the screeners can approximate age based on visual appearance. Because, from the start of the process, the hiring process was not “age-blind,” it is appropriate to treat the application and selection for interview steps as combined.

In the second hiring procedure (“electronic applications”), there are two steps. In the first step, people apply on-line for jobs, proving detailed information on their education, their previous work experience and history, their availability for work, whether they have reliable transportation, and more. They also respond to a battery of about 100 questions that are used to provide a quantitative talent assessment on dimensions that include customer service, engagement, retention, and teamwork. These questions lead to two numerical scores – one based on experience, availability, and flexibility, which by and large is based on the same pre-screening information used in the paper application process, and a second based on the battery of talent assessment questions (both on a 1 to 100 scale). They also lead to a three-category indicator suggesting that the manager can move ahead confidently with the candidate, move ahead with caution, or that the manager should not consider the candidate unless other options have been exhausted; this is based primarily on the talent assessment. The data include both of these metrics.

In the second step, selected applicants are interviewed, in person, and job offer decisions are made after the interview(s). Importantly, no direct information on age is elicited in the first step. In contrast, the second step is not “age-blind” because applicants interview in-person with company managers. I have data on both selection for interviews and whether interviewees were offered a job. Because the first step of the process is age-blind, and the second is not, it is appropriate to study the two outcomes separately (although I also look at overall hiring out of the applicant pool to provide a comparison with the paper applications).

IV. Identifying Discrimination

If I only had the paper applications with their single outcome based on age and other applicant characteristics, I would be limited to testing whether there is a difference in job offer rates between older and younger applicants conditional on controls for these characteristics – an approach usually termed the "residual" approach to testing for labor market discrimination (Oaxaca, 1973). With detailed data on applicant characteristics the residual approach can of course be informative. However, there is always the potential objection that a researcher was unable to control for an omitted variable correlated with age that could explain the age difference in job offer rates.

This is where the comparison with the outcomes from the two-step process for the electronic applications is critical. I can run a similar analysis to that for the paper applications – asking whether there is a lower job offer rate to older applicants, conditional on the control variables (which are more detailed than for the paper applications). But I can also break the process into two steps. In the first step of this process, armed with detailed information on applicant characteristics, but blind to age, I can test whether the company under-selects older applicants for interviews. And I can then separately examine job offers to interviews – decisions that are *not* age blind.

It turns out that the company does not under-select older applicants for interviews. Rather, older applicants experience adverse outcomes (fewer job offers) only at the interview stage, when age is revealed. This evidence from the electronic applications thus substantially undermines any argument about an omitted variable correlated with age that explains lower job offer rates for older applicants. Put differently, the differences in outcomes for older vs. younger applicants in the two stages of the hiring process for electronic applications, in comparison to the difference in outcomes in the single-stage hiring process for paper applications, provides far more rigorous evidence of age discrimination than can be gleaned solely from the analysis of job offers by age for the paper applications controlling for detailed applicant characteristics.

There are two potential arguments that could be made against the claim that the difference in outcomes for older vs. younger applicants in the two stages of the process of using electronic applications, in comparison to what happens with the single-stage process for paper applications, provides rigorous evidence of age discrimination. First, there could be an omitted variable associated with the interview that explains the lower job offer rate for older applicants at the second stage for the electronic applications, which was not revealed or detected in the first stage of the process, but which *was* revealed in the single-stage process for the paper applications. However, given the detailed assessment of qualifications and experience elicited in the first step of the hiring process for the electronic applications, it is hard to imagine what legitimate, job-related difference is revealed in the interview that could justify the age difference in job offer rates that occurs at the second stage, despite the absence of under-selection of older applicants for interviews.

Second, the difference in hiring procedures arises largely across restaurants, so that identification mainly comes from differences in hiring procedures across restaurants rather than a change in hiring procedures for given restaurant locations.[8] Thus, one could potentially be concerned that there are other differences in hiring procedures across restaurants that could potentially explain the differences in findings by age (although it is not clear what these would be). However, the company in question had uniform hiring processes across restaurants. Although hiring decisions were made at the restaurant level, all hiring managers went through the company's manager-in-training program that covered hiring, and a training in opening new restaurants, which covered hiring and the company's interview and selection process. The company used the same paper or electronic application forms across each restaurant, and company documents and testimony pointed to the same general hiring criteria and processes across paper application restaurants and electronic application restaurants (aside from the differences I have described).

That said, this discussion reveals the near-impossibility of definitively ruling out all alternative

[8] As discussed below, only a handful of restaurants have data using both types of hiring procedures.

non-discriminatory explanations of the evidence. The best one can do, it might be argued, is to explain why the non-discriminatory explanations are highly implausible.

V. Empirical Approach

In this section I briefly outline the models I estimate, prior to explaining the data in detail and presenting the results. For the data from both paper and electronic applications, I estimate linear probability models for whether applicants were offered a job (*OJ*).[9] I have an indicator of whether the applicant was aged 40 or over (*Age40P*), which makes the applicant eligible for protection under the Age Discrimination in Employment Act. I have dummy variables for each separate restaurant location (*R*). The location controls are potentially important because the age composition of the workforce may differ across locations, and if hiring rates also differ across locations this could generate a spurious relationship between job offer rates and age. I have data on numerous characteristics of the applicants – with more-detailed control variables for the electronic applications than the paper applications (I denote these *X* without distinguishing between the paper and electronic controls). Among the latter, I also have detailed assessment scores from the first stage of the electronic application process, before age is revealed. In addition, for the hiring process for the electronic applications, I have data on both selection for interviews – decisions made before age is revealed (*SEL*) – and job offers.

Thus, I first estimate models for job offers, using data from either the paper or electronic applications. Letting i index individuals and j restaurant locations, these models take the form:

$$OJ_{ij} = \alpha + \beta_{OJ} \cdot Age40P_{ij} + X_{ij}\gamma + R_j\delta + \varepsilon_{ij} \ . \qquad (1)$$

Evidence that $\beta_{OJ} < 0$ would indicate – based on the residual approach to discrimination – that there is age discrimination in hiring. However, I can obtain much more compelling evidence from the data for the electronic applications.

Using the additional information for the electronic applications, I estimate similar regressions – but for selections for interviews, and for job offers conditional on selection for interviews:

[9] The estimated marginal effects from probit models were nearly identical.

$$SEL_{ij} = \alpha + \beta_{SEL} \cdot Age40P_{ij} + X_{ij}\gamma + R_j\delta + \eta_{ij} \tag{2}$$

and

$$OJ_{ij} = \alpha + \beta_{OJ|SEL} \cdot Age40P_{ij} + X_{ij}\gamma + R_j\delta + \upsilon_{ij} \ \ | \ \{SEL_{ij} = 1\} \ \ . \tag{3}$$

Evidence that $\beta_{OJ|SEL} < 0$ while $\beta_{SEL} = 0$ would provide much more compelling evidence of age discrimination in hiring, because information on age is only available at the job offer stage and not the selection stage.

There is no reason to expect υ in equation (3) to be correlated with *Age40P* because of the selection decision (equation (2)), because the selection decision is age-blind. Indeed, consistent with this, the estimates of β_{SEL} in equation (2) are always near zero and generally statistically insignificant. Finally, I have data on whether the jobs were for "front-of-house" jobs (hosts, bartenders, and servers) or "back-of-house" jobs (line cooks, prep cooks, and stewards). "Front-of-house" jobs entail customer interaction and providing services to guests. I use differences for these two types of jobs (denoted *FOH* and *BOH*) to explore the potential role of customer discrimination. I also estimate all three specifications for front-of-house and back-of-house jobs. Stronger evidence of age discrimination from front-of-house jobs would suggest that customer discrimination plays a role (or at least perceived customer discrimination on the part of the employer).

VI. Data

There are multiple components of the data I use. In this section, I describe each of them in turn and provide more detail.

Hires

One data source I use is the employee roster covering employees by location, at new restaurants that opened during the sample period (2010-2016). The employee roster identifies hires at these locations around the time of the location opening. I define as hired workers those who were hired at, rehired at, or transferred to the location of a new restaurant, starting from the time the first hire, rehire, or transfer was made at a location, and ending a year after the location opening date, hence capturing the hiring that occurred when a location opened. I focus only on hourly employees, excluding managers.[10]

Of course, I am ultimately interested in whether job applicants are offered jobs, since the job offer is the decision the employer makes. One potential issue is if there are many individuals included in the employee roster that I could not match to an application, which would imply that there is missing applicant data on that person. If this happened a lot, we would have to be concerned that the application data are not complete. In addition, I need the applicant data to capture applicant characteristics. Thus, I use the employee roster to assess the completeness of the application data for each restaurant location. I discarded a small number of locations that have somewhat incomplete application data. After doing this, there is a small number of individuals in the employee roster but not the application data, whom I discard.

I use the employee roster for two other purposes. First, as discussed below, I sometimes rely on the roster for data on the ages of applicants. Second, while for most applicants I have information on whether a job offer was made, this information is sometimes missing, and if a person shows up on the employee roster but I do not have information on a job offer, I code them as having been offered a job.[11]

Data on applicants

For each restaurant location, the application data I was provided covers the earliest applications prior to the restaurant opening through one year after the restaurant opened. (This pertains to both the electronic and paper applications.)

<u>Coding age</u>

I obtain information on ages of applicants from different sources depending on the available data. But I am able to verify that these sources are very consistent with each other (when data are available from more than one source). First, dates of birth are available from employees in the employee roster. I always use this information as definitive when I have it. The applicant data do not contain date of birth. Date of birth on most applicants was acquired by submitting information on applicants to a company

---

[10] There are sometimes multiple observations at a location – for example, when an employee was a new hire and then was rehired at the same location. I use the initial hire.

[11] As short-hand, I always refer to the outcome as a “job offer.”

called Accurint, which has a proprietary algorithm to identify peoples' dates of birth based on other information about those people – which in the case of the applicant data included name, address, home phone number and/or cell phone number, and more.[12] If Accurint was able to find the applicant in the data sources to which they match, Accurint provided us with the date of birth of the best match. The dates of birth were either a full date of birth (containing a month, day, and year) or a partial date of birth (missing day or month). Accurint was able to provide at least a partial date of birth for 86.96% of electronic applicants and 81.09% of paper applicants (Appendix Table A1).

To check the date of birth results provided by Accurint, I compared the dates of birth Accurint provided for all employees to the company data. The year matches exactly for 94.17% of cases. (When I looked at exact birth date, rather than only year, the match rate was still 86.99%.) There is a small number of larger discrepancies, but these are very infrequent. Based on these results, I am highly confident in using the date of birth information identified by Accurint. To be clear, however, these checks were done only for the employee data for which I have a second and presumably reliable reading of date of birth.

For the electronic applications, there was also a method to recover an estimate of date of birth for some applicants for whom Accurint did not provide a date of birth or who did not appear on the employee roster. In particular, in the on-line application process, applicants were asked and sometimes report their year of high school graduation. Given that it is very common to graduate high school at or very near age 18, year of high school graduation can convey fairly accurate information on an applicant's age. To check this, I compared age based on year of high school graduation to age from the employee roster or from Accurint (when only the latter was available).[13] The match rate on age in years (calculating age as the year of application minus the year of graduation plus 18) was 99.88%. I therefore used year of high

[12] For the paper applications, to the extent that there are Social Security numbers, these were provided to Accurint. Email addresses were available for electronic applications. Accurint relies on multiple public records data sources, such as Department of Motor Vehicle databases, as well as some private information sources.
[13] There is no equivalent check on the birth dates in the data for paper applications, because year of high school graduation is not reported often in the paper applications.

school graduation when it was reported but age information was missing from the other sources.[14] Appendix Table A2 shows the source of date of birth information for the electronic and paper applications, as well as the percentages for which age could not be assigned.

One other question I can address using the employee roster is whether there is any pattern to the ages of people that Accurint had difficulty coding. I do this by comparing the age distribution for employees for whom Accurint returns a date of birth to the age distribution for employees for whom Accurint does not return a date of birth. As shown in Appendix Figure A1, these distributions look quite similar (albeit with lower frequency for the latter group). That is, there is no apparent tendency for Accurint to fail to return dates of birth for young versus old employees.

I determined the youngest an applicant could be on the day they applied using the dates of birth from Accurint. For the applicants who were matched to a full date of birth, I know their exact age when they applied for a position with the company. For the applicants who were not matched to a full date of birth, I used the information that was provided by Accurint to provide a lower bound on the age of the applicant.[15] If the date of birth for an applicant was missing the month, December was used for the month of birth. If the date of birth for an applicant was missing the day of the month they were born, I use the last day of the month.

<u>Application outcomes and applicant characteristics</u>

For the electronic applications, the applicant database from the company includes information on interviews and job offers, capturing the dates that each applicant was interviewed and extended a job offer (if they were). There are some people who appear on the employee roster for whom job offer information

[14] This does raise the question of whether the reported year of high school graduation was looked at, or used as an indicator of age, by hiring managers when selecting applicants to interview at the first step of the electronic application process, which I have characterized as age-blind. However, deposition testimony about the hiring process indicated that the applicant's date of high school (or college) would not be a factor in selecting who to interview. This is borne out by the evidence that older applicants are not under-selected at this stage of the application process.

[15] When I do not know the exact date of birth, there is a small possibility that an applicant is incorrectly assigned to the under 40 age group (as opposed to the 40 or over group). However, this misclassification error creates, if anything, a bias *against* finding an age difference in job offer rates.

is not recorded, cases that I recode as having received an offer at the same location (since clearly they were hired). This required matching the names and restaurant locations between the applicant data and the employee roster – matches that are not always perfect because of things like different spelling of names or shortening of names (e.g., Tim vs. Timothy).[16]

In this hiring process, managers were instructed to move only one application through the system if an applicant submitted an application for more than one position. Thus, for applicants who submitted more than one application, I keep the application that went farthest in the hiring process, so that I will not miss a job offer if it occurred.[17]

For paper applications, I was provided with the paper applications filled out in person, with identifiers for restaurant location. As noted above, the applicant first fills out a paper application for a pre-screening interview, with a limited number of questions. I was also provided with some booklets capturing the interviews and assessments that could follow the pre-screening interview, which include, most importantly, information on whether a job was offered. The interview assessment results ask open-ended questions about work experiences, covering topics such as handling stress, teamwork, and customer service. In principle the candidate is given a score on each of these, but recording of these scores is very incomplete.

The applications and the interview booklets were coded to make them machine-readable, working with an outside vendor (Bluestar) to ensure very high accuracy, including extensive review of the coding in process. Any data that were difficult to read were indicated as missing or potentially problematic, and

---

[16] I do this using the "reclink" fuzzy matching algorithm in Stata. The algorithm determines the best match for each record in one of the files. It computes a score between 0 and 1 (higher match scores reflect a higher probability that the two records are a match, and 1 indicates a perfect match) and proposes a match if two records have a match score between 0.6 and 1. Matches are based on the name on the application and the name of the employee on the roster. Matches were conditional on the employee and the electronic application originating at the same location (because it seems far less likely that an applicant and employee record with similar names at different locations are the same person). If the match between an employee and an application was scored 0.95 or better, it was coded as a match.
[17] If two applications stopped at the same stage (most common when they were both rejected without being interviewed), I randomly chose one and the other was discarded (using a random number generator), so as to not count a person's application as rejected more than once.

in my analysis I rely only on data that could be accurately coded without guesswork. Moreover, even if there is some coding error (and, like with all datasets, this cannot be ruled out) there is no reason to believe the small number of errors that might have entered the data via this coding are correlated with age in a way that would bias estimated differences in outcomes with respect to age.

I use data from the initial paper application as control variables in my analysis of the paper applications. Those variables I capture and use in my analysis (which include all but a short work history, which cannot be coded consistently), include, among others: meeting minimum age requirements; legal work status; whether the person is a current employee of the company; salary expectations; availability of reliable transportation; record of felony convictions; shift availability; high school diploma; and previous company experience. (See Appendix Table A3.)

I did not use the information in the interview booklets to construct controls for the regression analyses for three reasons. First, there are very few booklets. I was able to match booklets to only about 13% of paper applications. Second, there is a lot of missing information in the booklets. For example, the boxes for interview scores, both the total score and the individual components, are missing for 66% of booklets. Likewise, boxes for the assessment results are missing for 26.2% of booklets. Third, it appears that the booklets for applicants who were given a job offer were provided by the company at a much higher rate than the booklets for applicants who were not given a job offer. The majority of booklets, 51%, are coded as receiving a job offer, with 49% not coded as receiving a job offer.[18] However, I do capture job offers from the booklets, when the information is available.

Given that the interview booklets had job offer information, I had to match the interview booklets to the paper applications, again using names and location. The fuzzy matching algorithm described above was used. I coded an applicant as receiving a job offer in any of the following cases: if there was a date of offer recorded; if the rate of pay was recorded; if the position was accepted; if there was a start date given; or if there was an orientation date given. Job offers were also coded as having been given if the

[18] In contrast, the percentage of paper applicants hired, for the restaurants for which I use paper applications, ranges from 4.60% to 17.07%.

name on the application appeared in the employee roster at the same location, using the same matching procedure as for the electronic applications.

If the same applicant (same first name, middle name, and last name) submitted more than one paper application to the same location, only one paper application was kept. Unlike for the electronic applications, there is no way to select the application that went farthest, so I simply randomly select one of the applications.

For the electronic applications, there is a much more-detailed set of control variables available. These are listed in Appendix Table A3, and include things such as: willingness to be trained in different jobs, to stay late, to work multiple shifts; prior job experience; availability of transportation; past felony conviction; and assessment results based on the on-line applications (indicated in Appendix Table A3 under both “three-category ranking” and “numerical score”).

Data used in study

Table 1 lists, for each restaurant location, information on applications, job offers, and other features of the data, as well as an indication of how I analyze the data from each location (i.e., which type of analysis). To begin, columns (2) and (3) list the number of electronic and paper applications. Most locations have exclusively one kind of application process, although in a few cases there is a mix, because restaurants switched from the paper system to the electronic system during the hiring period for that location’s opening.

Columns (4)-(9) provide additional information on hires and applicants. In particular, they focus on two issues. First, do I have application data for individuals that are actually hired? I answer this by asking whether there are applications for all or almost all people who are listed on the employee roster as a hire during the relevant time period and at a relevant location. As column (5) shows, the percentage (share) of hires not represented among the applications is typically in the 6-30% range, although it is much higher in a few cases, such as Location 7 and Location 15, which have 59.35% and 66.47% hires missing from applicant data, respectively. Overall, there is application data for 82.51% of the individuals hired. This is not necessarily a problem, as I can still garner information on job offers for the applicants, and as long as there is not systematic difference in the ages of those hired for whom there is or is not application data, no bias is introduced.

Second, is the set of applications I have potentially biased, or does it provide a reliable sample of the applicant pool? Columns (7) and (8) point to more problematic issues for a few restaurants. Here, I report the share of applicants that are hired, and the share of paper applicants that are hired. In general, the share of applicants that are hired ranges between around 3% and 17%. However, for three locations – Location 2, Location 3, and Location 5 – this share is much higher, between 41% and 68% (column (8)). Thus, it appears that for these stores I am missing data on a very large share of applicants who were *not* hired. Moreover, for Location 12, although the share of applicants hired overall is not high (8.28%), the share of paper applicants hired is very high (84.92%). This reflects that I have almost no paper applications that did not result in a hire; based on the share of applicants that are hired in other locations (around 2 to 17 percent), I am missing paper applications from many people who were not hired. This type of missing data is more problematic, because it can arise from the company not retaining applications from a large share of applicants who were not offered jobs. Given that the company could have discarded applications from the group it tended not hire, this type of missing data would more plausibly create bias in estimated age differences if they existed in the hypothetical complete data.

Table 1 then shows, in columns (9) and (10), the data I use for each location, if I do, and why. First, all of the restaurants with only electronic applications are included in the analysis of electronic applications. In addition, for restaurants with both electronic applications and paper applications, I include the electronic applications from those restaurants in the analysis of electronic applications. Restaurants for which electronic applications are included in the analysis of electronic applications are indicated in column (10) with either "EA" or "EA and PA."

Second, among the restaurants with only paper applications, those that do not have an inordinately high share of applicants hired (column (9)) are included in the analysis of paper applications. In addition, for restaurants with a sizable number of paper applications, even if there are also electronic applications, I include the paper applications from those restaurants in the analysis of paper applications (with one exception noted below). Restaurants for which all or some applications are included in the analysis of paper applications are indicated in column (10) with either "PA" or "EA and PA."

There are some restaurant locations I will explain in more detail. One location is Location 12, where although the share of applicants hired overall is not high (8.28%), the share of *paper* applicants hired is very high (84.92%). I cannot rely on the paper applications for reasons explained above, but I have no reason to believe the electronic applications are unreliable; hence, as indicated in column (9), only the electronic applications are used. For location 15, although the share of hires missing from the applicant data is very high (66.47%), I can study the electronic applicant data that does not have a high share of hires; hence, column (9) indicates "EA." Finally, for Location 7, although the share of hires missing from applicant data is very high (59.35%), the share of the applicant pool that are hires (column (7)) and the share of paper applicant pool that are hires (column (8)) are very low; hence, I use both types of applications.

VII. Results

Overall job offer rates

I first report the results for overall job offer rates, for both the paper applications and the electronic applications. For the paper applications, this is all I can do, because there is only one step in the hiring process. Afterwards, I delve into the richer and more compelling analyses available for the electronic applications.

For the paper applications, recall that the key feature of the hiring process for these application is that, from the start of the process, the hiring process was not "age-blind." I begin with a specification (column (1)) that just includes the dummy variable of interest and controls for each restaurant location, to capture differences in job offer rates across location that could be correlated with age. I then estimate a specification with richer controls available from the paper applications (column (2)). I use linear probability models, with standard errors clustered as the restaurant location level.

As reported in Table 2, for the analysis of applicants who applied using paper applications, in the specification with restaurant fixed effects only, the estimated difference in the probability that older applicants receive job offers is −0.091, or 9.1 percentage points lower. The estimated difference in the probability of job offers for older applicants is strongly statistically significant. Relative to the job offer rate of 14.1% for younger applicants, the job offer rate to older applicants is lower by 64.6%. When I add the full set of controls (which is fairly limited for the paper applications), the estimated difference in the probability that older applicants receive offers is −0.096, or 9.6 percentage points lower, slightly larger than the estimate without the additional controls. The estimated difference in the probability of job offers for older applicants is strongly statistically significant. Relative to the job offer rate of 14.1% for younger applicants, the job offer rate to older applicants is lower by 67.9%.

The evidence from the paper applications suggests substantial discrimination against older workers in hiring. One could argue, however, that other qualifications of older vs. younger applicants explain the age difference in job offer rates – and indeed in the paper applications I do not have rich information on applicant characteristics. In the data on the electronic applications, I have a considerably richer set of control variables, including metrics from the company's own assessment tool.

For the analysis of overall job offer rates for the electronic applications, I report results for three different sets of controls. First, I estimate the age difference in job offer rates controlling only for restaurant location (with restaurant fixed effects), in column (3). Second, I use the detailed data available to control for applicants' qualifications, schedule availability, and other factors (column (4)); as noted earlier (and in the table notes) there are many more controls than for the paper applications. Finally, there are also direct measures of applicant assessments computed as part of the on-line application process (the three-category rankings and numerical scores). Given that these are intended to be summary measures of the evaluation of on-line applications, and the company chooses to employ them, I also estimate specifications dropping the other controls (except the restaurant fixed effects) and retaining only these evaluation scores (column (5)).

For the overall analysis of applicants who applied using on-line, in the specification with restaurant fixed effects only, the estimated difference in the probability that older applicants receive job offers is −0.019, or 1.9 percentage points lower. This estimated is statistically significant, although smaller than the difference in job offer rates for the paper applications. Relative to the job offer rate of 10.0% for younger applicants, the job offer rate to older applicants is lower by 18.7%.

When I add the full set of controls, the estimated difference in the probability that older applicants receive job offers becomes larger negative – a statistically significant estimate of −0.055, or 5.5 percentage points lower. It is noteworthy that this estimated age difference is larger than the estimate without controls. This is consistent with older applicants being more qualified than younger applicants, a result that comes out more strongly when I present evidence on the two stages of the hiring process for the electronic applications. This same result – a negative differential for older applicants – is confirmed in the last column of Table 2. Here I use the company's metrics for evaluating the electronic applications. The estimated difference in the probability that older applicants receive offers is −0.046, or 4.6 percentage points lower. This estimated is statistically significant, although smaller than the difference in job offer rates for the paper applications, and again much larger than the estimate without individual-level controls in the first column. Relative to the job offer rate of 10.0% for younger applicants, the job offer rate to older applicants is lower by 45.8%.

The estimated differences in overall job offer rates in columns (2), (4), and (5) of Table 2 – the models with controls, for paper and electronic applications – are qualitatively similar. In all cases, the evidence indicates statistically significantly lower job offer rates to older applicants, conditional on their qualifications and characteristics. However, for the electronic applications we have the additional information, implied by the difference in results including and excluding the individual controls, that older applicants are more qualified, which is why the estimated age difference becomes larger when the control variables are added. I next turn to the evidence from the two steps of the hiring process for the electronic applications, which provides additional evidence on this latter point from the first, age-blind step of the hiring process, and provides more compelling evidence that the overall differences in job offer rates reflect age discrimination.

Evidence on the two steps of the hiring process for electronic applications

I now turn to separate analyses of the two steps of the hiring process for the electronic applications: selection for interviews in the age-blind first step; and job offers after the interview in the non-age-blind second step. As discussed above, the differences in results for these two stages can provide more rigorous evidence on age discrimination.

In Table 3, for the model with restaurant fixed effects only, in the first, age-blind step older applicants are 0.058 or 5.8 percentage points more likely to be selected for interview. This difference is statistically significant. In contrast, in the second step, which is not age-blind, older applicants (in this case, among those selected for interviews), are disfavored. The estimated probability that a job offer is made to an older interviewed applicant is lower by 0.081, or 8.1 percentage points, which is strongly statistically significant. Column (2) shows that relative to the selection rate of 30.0% for younger applicants, the selection rate for older applicants is lower by 27.0%.

Thus, the implication from columns (1) and (2) of Table 3 is that based on the on-line application and assessment, older workers are seen as more qualified for the job, on average. Only at the interview stage, when age is revealed, does the job offer rate shift *against* older job applicants. This conclusion is reinforced in the remaining columns that add the alternative sets of control variables.

Whether I include the full controls or just the assessment measures, the estimates in columns (3) and (5) indicate that, in the first, age-blind step of the hiring process in which applicants are selected for interviews, there is no substantive difference in the treatment of older vs. younger applicants. The estimated age gap is small (−0.004 or 0.011, or minus 0.4 to 1.1 percentage points) and statistically insignificant.

In contrast, as we might anticipated given older workers' better qualifications and assessments, in the models with controls – columns (4) and (6) – there is now even stronger evidence pointing to age discrimination in the second, non-age-blind step of the hiring process. In column (4), the estimated probability that a job offer is made to an older interviewed applicant is lower by 0.128, or 12.8 percentage points, which is strongly statistically significant. Relative to the job offer rate of 30.0% for younger applicants, the job offer rate for older applicants is lower by 42.8%. And in column (6), using the company's assessment scores, the corresponding estimates are very similar, with a job offer rate for older applicants lower by a strongly statistically significant 11.9 percentage points, or 39.5%.

The evidence can be summarized as follows. First, older applicants are more qualified in terms of applicant characteristics and evaluations used by the company in their on-line application system. Second, older applicants are not under-selected for interviews at the first, age-blind step of the hiring process for electronic applications, conditional on these characteristics and evaluations. Third, at the second, non-age-blind step of this process when applicants are interviewed, the largely neutral selection of older applicants turns into a far lower job offer rate to older applicants. Fourth, in the one-step process for paper applications, where age is revealed early, there is a similar far lower job-offer rate for older applicants.

This set of results is strongly consistent with age discrimination. The only tenable alternative explanation is that something is revealed at the interviews that marks older applicants as less qualified. But given that in the on-line system the company is relying on its metrics for evaluating candidates based on detailed information, it is difficult to imagine what information related to qualifications for the job could be revealed in the interviews but not the assessments, and hence the far more plausible interpretation is that age, per se, is the driving factor in reducing job offer rates to older applicants. We next turn to some potentially corroborating evidence that the evidence reflects age discrimination.

Customer discrimination?

A standard hypothesis for why employers discriminate against a particular group of workers is that employers believe that their customers have discriminatory tastes that are prejudicial to that group, so that hiring from that group would make the business less profitable (Becker, 1971). Jobs at the restaurant company studied in this paper can be divided into “front-of-house” jobs (hosts, bartenders, and servers) and “back-of-house” jobs (line cooks, prep cooks, and stewards), with the front-of-house jobs being those that entail customer interaction and providing services to guests.

I therefore repeat versions of the analyses reported above, but for front-of-house and back-of-house jobs separately. The evidence indicates that the age gaps in the job offer rate are larger for front-of-house than for back-of-house jobs, which is consistent with customer discrimination as at least a partial source of the evidence reported above. Moreover, this evidence of customer discrimination is in contrast to the alternative hypothesis that something is revealed in the job interviews that provides a non-discriminatory explanation of the lower job offer rate to older applicants, as it is not clear why this unobservable would differ between front-of-house and back-of-house jobs.

The evidence is reported in Table 4. Results are reported for the last set of specifications for the paper and electronic applications (using the full controls for the former and the assessment controls for the latter), but the results are the same using the full controls for the electronic applications. Columns (1) and (2) provide results for the overall job offer rate for the paper applications – the one-step process that is not age-blind. In this case, the estimated differential for older applicants is −0.101 or 10.1 percentage points for front-of-house jobs, compared with −0.068 or 6.8 percentage points for back-of-house jobs. Both estimated differences are statistically significant. The offer rate for older applicants is 75.5% lower for front-of-house jobs, and 52.3% lower for back-of-house jobs, consistent with customer discrimination partly driving the age difference in job offer rates.

Columns (3)-(5) present the evidence for electronic applications for front-of-house jobs. The estimated age difference in the overall hiring rate is −0.048 (statistically significant). At the first, age-blind stage of selection for interviews, there is no evidence of lower selection rates for older applicants; the estimates is positive (0.008), very small, and insignificant. However, there is a large and significant difference in the direction of discrimination against older applicants in the second, non-age-blind job-offer stage, with a statistically significant estimate of −0.123.

For back-of-house jobs, the point estimates reported in columns (6)-(8) are not that different from the front-of-house estimates. The estimated age difference in the overall hiring rate is −0.041 (statistically significant). At the first, age-blind stage of selection for interviews, there is no evidence of lower selection rates for older applicants; the estimates is positive (0.020), small, and insignificant. There is a large and significant difference in the direction of discrimination against older applicants in the second, non-age-blind job-offer stage, with a statistically significant estimate of −0.112.

However, while the point estimates for front-of-house and back-of-house jobs are fairly similar, the relative differences are much larger. For overall job offer rates, the offer rate for older applicants is 55.3% lower for front-of-house jobs, but only 27.7% lower for back-of-house jobs, because the baseline job offer rate is lower for front-of-house jobs. And the relative difference is also more pronounced for job offers conditional on selection for interview, again because the baseline (conditional) job offer rate is lower for front-of-house jobs. The job offer rate for older applicants conditional on selection for interview is 45.7% lower for front-of-house jobs, but only 28% lower for back-of-house jobs. Thus, the evidence is consistent with customer age discrimination playing a role in generating the age discrimination documented in these data.

VIII. Conclusions

In this paper, I study age discrimination in hiring using non-experimental data, exploiting a difference in hiring procedures that alters when those responsible for hiring become aware of the age of applicants. Because hiring managers can of course not act on a worker's age until they become aware of it, I use the difference in hiring procedures to identify age discrimination. Under the first hiring procedure, age is revealed at the same time as all other information about the applicant (because everything is done in-person). Under the second hiring procedure, selections for interview are made on-line, and are age-blind. Subsequently, those selected are interviewed – which is of course not age-blind.

I find evidence strongly consistent with age discrimination. Under either hiring procedure, overall job offer rates to older applicants are much lower than to younger applicants. However, under the hiring procedure with a first age-blind step and a second step that is not age-blind, we learn the following. First, older applicants are more qualified in terms of applicant characteristics and evaluations used by the company in their on-line application system. Second, older applicants are not under-selected for interviews at the first, age-blind step of the hiring process for electronic applications. Third, at the second, non-age-blind step of this process, when applicants are interviewed in person, the neutral selection of older applicants turns into a far lower job offer rate to older applicants. Fourth, in the one-step process for paper applications, where is not age-blind, there is a similar far lower job-offer rate for older applicants. The most plausible interpretation is that age, per se, is the driving factor in reducing job offer rates to older applicants.

There is some corroborating evidence that age discrimination explains the findings. In particular, the age differences in job offer rates that I estimate are larger for front-of-house jobs that involve customer service and interaction than for back-of-house jobs that do not. This evidence suggests the customer discrimination (or at least the company's perception of customers' preferences) at least partly underlies the evidence of age discrimination that I find.

Finally, there may be a more general conclusion we can draw from this evidence regarding hiring procedures that can reduce discrimination. The evidence can be interpreted as showing that hiring procedures that elicit detailed information on job applicants without revealing membership in groups that may experience discrimination can in fact eliminate discriminatory outcomes. It may not be realistic to expect employers to forego interviews completely, relying on hiring procedures that are completely blind to demographic or other characteristics of workers. But for some, perhaps lower-skilled jobs this might be more realistic. (And it might be possible in unusual cases, such as the orchestra auditions studied by Goldin and Rouse (2000).) More generally, it is certainly possible to push hiring procedures blind to age, race, sex, etc., further along in the hiring process – as illustrated by the studies by Krause et al. (2012), Åslund and Nordströum Skans (2012), Behaghel et al. (2015), and Agan and Starr (2018). The evidence I find of discrimination that surfaces at the stage of face-to-face interviews should motivate the development of hiring procedures that downweight the importance of these interviews, or that make interviews more neutral in their evaluations of job candidates.

On the other hand, in some of these studies (Krause et al., 2012; Behaghel et al., 2015), women or minorities fared worse when applications were anonymized. One explanation for this kind of evidence is that anonymization impeded affirmative action efforts to hire more women or minorities. Conversely, the Goldin and Rouse (2000) study may have found that anonymization helped women because orchestras were *not* engaged in affirmative action. I suspect that anonymization helps older applicants in the data studied in the present paper, and would more generally, because employers do not engage in affirmative action with respect to older workers. Indeed, Executive Order 11246 (amended), which is generally viewed as providing the legal basis for affirmative action in employment in the United States, refers to "race, color, religion, sex, sexual orientation, gender identity, or national origin"[19] – but *not* to age.

[19] See https://www.dol.gov/ofccp/regs/statutes/eo11246.htm.

References

Agan, Amanda, and Sonja Starr. 2018. "Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment." *Quarterly Journal of Economics*, Vol. 131, pp. 191-235.

Altonji, Joseph G., and Rebecca M. Blank. 1999. "Race and Gender in the Labor Market." In Ashenfelter and Card, eds., Handbook of Labor Economics, Volume 3, Part C, pp. 3143-259. Amsterdam: Elsevier.

Åslund, Olof, and Oskar Nordströum Skans. 2012. "Do Anonymous Job Application Procedures Level the Playing Field?" *Industrial and Labor Relations Review*, Vol. 65, pp. 82-107.

Becker, Gary S. 1971. The Economics of Discrimination, Second Edition. Chicago: University of Chicago Press.

Behaghel, Luc, Bruno Crépon, and Thomas Le Barbanchon. 2015. "Unintended Effects of Anonymous Resumes." *American Economic Journal: Applied Economics*, Vol. 7, pp. 1-27.

Cahuc, Pierre, Stéphane Carcillo, Andreea Minea, and Marie-Anne Valfort. 2019. "When Correspondence Studies Fail to Detect Hiring Discrimination." IZA Discussion Paper No. 12653.

Fix, Michael, and Raymond Struyk. 1993. Clear and Convincing Evidence: Measurement of Discrimination in America. Washington, DC: The Urban Institute Press.

Goldin, Claudia, and Cecilia Rouse. 2000. "Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians." *American Economic Review*, Vol. 90, pp. 715-41.

Heckman, James, and Peter Siegelman. 1993. "The Urban Institute Audit Studies: Their Methods and Findings." In Fix and Struyk, eds., Clear and Convincing Evidence: Measurement of Discrimination in America, pp. 187-258. Washington, D.C.: The Urban Institute Press.

Holzer, Harry J. 1998. "Why Do Small Establishments Hire Fewer Blacks Than Large Ones?" *Journal of Human Resources*, Vol. 33, pp. 896-914.

Krause, Annabelle, Ulf Rinne, and Klaus F. Zimmermann. 2012. "Anonymous Job Applications of Fresh Ph.D. Economists." *Economics Letters*, Vol. 117, pp. 441-4.

Neumark, David. 2018. "Experimental Research on Labor Market Discrimination." *Journal of Economic Literature*, Vol. 56, pp. 799-866

Neumark, David. 2012. "Detecting Evidence of Discrimination in Audit and Correspondence Studies." *Journal of Human Resources*, Vol. 47, pp. 1128-57.

Neumark, David. 1996. "Sex Discrimination in Restaurant Hiring: An Audit Study." *Quarterly Journal of Economics*, Vol. 111, pp. 915-41.

Neumark, David, and Judith Rich. 2019. "Do Field Experiments on Labor and Housing Markets Overstate Discrimination? A Re-Examination of the Evidence." *ILR Review*, Vol. 72, pp. 223-52.

Oaxaca, Roland. 1973. "Male-Female Wage Differentials in Urban Labor Markets." *International Economic Review*, Vol. 13, pp. 693-709.

Riach, Peter A., and Judith Rich. 2002. "Field Experiments of Discrimination in the Market Place." *The Economic Journal*, Vol. 112, pp. F480-518.

Table 1: Treatment of Restaurants Locations and Paper and Electronic Applications from Restaurants in Analysis

Location	Electronic applicants	Paper applicants	Total hires	Share of hires missing from applicant data	Share paper applications	Share of applicant pool that are hires	Share of paper applicant pool that are hires	Analysis: hiring out of applicants	Comments
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
1	392	1179	184	9.78%	75.05%	10.57%	13.31%	EA and PA	
2	0	329	189	29.10%	100.00%	40.73%	40.73%		High % of paper applicants hired, suggests many paper applications missing
3	0	212	177	18.08%	100.00%	68.40%	68.40%		High % of paper applicants hired, suggests many paper applications missing
4	0	2531	151	5.96%	100.00%	5.61%	5.61%	PA	
5	0	314	177	22.60%	100.00%	43.63%	43.63%		High % of paper applicants hired, suggests many paper applications missing
6	386	1071	190	22.63%	73.51%	10.09%	11.48%	EA and PA	
7	849	891	214	59.35%	51.21%	5.00%	4.60%	EA and PA	
8	157	854	156	19.87%	84.47%	12.36%	13.93%	EA and PA	
9	0	2763	174	13.22%	100.00%	5.47%	5.47%	PA	
10	0	943	194	17.01%	100.00%	17.07%	17.07%	PA	
11	1290	1337	191	28.80%	50.89%	5.18%	5.83%	EA and PA	
12	1564	126	167	16.17%	7.46%	8.28%	84.92%	EA	High % of paper applicants hired, suggests many paper applications missing
13	3376	0	166	6.02%	0.00%	4.62%	N/A	EA	
14	3409	0	154	13.07%	0.00%	3.93%	N/A	EA	
15	1158	8	170	66.47%	0.69%	4.89%	100.00%	EA	High % of paper applicants hired, suggests many paper applications missing EA only, excluding the 8 paper applications
16	4385	6	181	7.73%	0.14%	3.80%	16.67%	EA	EA only, excluding the 6 paper applications
17	2641	0	237	7.59%	0.00%	8.29%	N/A	EA	
18	2345	1	231	12.99%	0.04%	8.57%	100.00%	EA	EA only, excluding the 1 paper applications
19	3327	0	189	12.17%	0.00%	4.99%	N/A	EA	
20	3597	3	204	13.73%	0.08%	4.89%	0.00%	EA	EA only, excluding the 3 paper applications
21	4839	0	265	8.30%	0.00%	5.02%	N/A	EA	
22	1642	0	250	20.00%	0.00%	12.18%	N/A	EA	
23	3240	0	207	13.53%	0.00%	5.52%	N/A	EA	
24	1346	0	183	13.66%	0.00%	11.74%	N/A	EA	
25	3352	0	201	6.47%	0.00%	5.61%	N/A	EA	
26	10387	0	182	7.73%	0.00%	1.62%	N/A	EA	
27	4740	0	216	7.41%	0.00%	4.22%	N/A	EA	
28	5258	0	163	16.56%	0.00%	2.59%	N/A	EA	
29	1064	0	213	28.64%	0.00%	14.29%	N/A	EA	
30	1985	0	181	23.76%	0.00%	6.95%	N/A	EA	
31	5896	0	202	10.89%	0.00%	3.05%	N/A	EA	
32	1382	0	194	17.53%	0.00%	11.58%	N/A	EA	
33	2985	0	198	13.13%	0.00%	5.76%	N/A	EA	
34	1620	10	198	14.65%	0.61%	10.37%	50.00%	EA	EA only, excluding the 10 paper applications
35	1193	0	177	9.60%	0.00%	13.41%	N/A	EA	
All locations	79,805	12,578	6,726	17.49%	13.62%	6.01%	12.00%		

Table 2: Estimated Differences in Probabilities of Job Offers to Applicants Aged 40 and Over, Paper and Electronic Applications, Linear Probability Models

	Paper applications		**Electronic applications**		
	Restaurant controls	**Full controls**	**Restaurant controls**	**Full controls**	**Assessment categories and scores**
	(1)	(2)	(3)	(4)	(5)
Age 40 or over	-.091 (.020)	-.096 (.017)	-.019 (.007)	-.055 (.007)	-.046 (.007)
Baseline job offer rate for those under age 40	14.1%	14.1%	10.0%	10.0%	10.0%
Percent difference relative to baseline rate for under age 40, age 40 or over	-64.6%	-67.9%	∨-18.7%	-55.4%	-45.8%
Sample size	8,485	8,485	47,667	47,667	47,667
% of sample size age 40 and over	22.1%	22.1%	13.4%	13.4%	13.4%

Estimates are from linear probability models. Of the 12,578 total paper applications, only one application per applicant per location is kept. After removing duplicate applications, the number of paper application is 11,377. I also excluded applications from locations not used in the analyses of paper applications. I further restrict the sample to only applications where I have information on date of birth. Of the total 79,805 electronic applications, only one application per applicant is kept per location. This reduces the number of applications to 51,800. From the 51,800 applications, only applicants where I have information on date of birth are used. The first row of table reports the estimated coefficient, and below it the standard error of the coefficient estimate, clustered at the restaurant level. The table also reports the baseline job offer (or selection) rate and the implied percent difference associated with age 40 or over. All models include dummy variables for restaurants and an intercept. For the paper applications, models in the "Full Controls" column also include the following control variables: felony; high school diploma; legal right to work; whether age is over 18; previous company experience; and dummy variables for missing data. For the electronic applications, models in the "Full Controls" column also include the following control variables: assessment; the assessment categories (dummy variables) and scores (1-100); legal right to work; transportation; age over 21 (separately for bartender and non-bartender applicants); felony; educational degree level and type; type and amount of restaurant experience; training information (willingness to trained as trainer); whether seeking full-time or part-time work; shift availability; previous company experience; and dummy variables for missing data. Among the potential control variables, only those missing for fewer than 25% of observations are used in the analysis. See Appendix Table A3 for details of the controls used.

Table 3: Estimated Differences in Probabilities of Selection for Interviews and Job Offers in Two Steps of Hiring Process for Applicants Aged 40 and Over, Electronic Applications, Linear Probability Models

	Restaurant controls		**Full controls**		**Assessment categories and scores**	
	Selection for interview	**Job offers for those selected for interview**	**Selection for interview**	**Job offers for those selected for interview**	**Selection for interview**	**Job offers for those selected for interview**
	(1)	(2)	(3)	(4)	(5)	(6)
Age 40 or over	.058 (.013)	-.081 (.016)	-.004 (.010)	-.128 (.014)	.011 (.011)	-.119 (.014)
Baseline job offer/selection rate for those under age 40	33.4%	30.0%	33.4%	30.0%	33.4%	30.0%
Percent difference relative to baseline rate for under age 40, age 40 or over	17.3%	-27.0%	-1.2%	-42.8%	3.2%	-39.5%
Sample size	47,667	16,322	47,667	16,322	47,667	16,322
% of sample size age 40 and over	13.4%	15.5%	13.4%	15.5%	13.4%	15.5%

Notes to Table 2 (in reference to the electronic applications) apply.

Table 4: Estimated Differences in Probabilities of Job Offers and Selections for Interviews for Applicants Aged 40 and Over, Paper and Electronic Applications, Front-of-House and Back-of-House Jobs, Linear Probability Models, Full Controls

	Paper Applications, Full Controls		**Electronic Applications, Assessment Categories and Score Controls**					
	Overall		**Front-of-house**			**Back-of-house**		
	Front-of-house	**Back-of-house**	**Overall job offers, front-of-house**	**Selection for interview, front-of-house**	**Job offers for those selected for interview, front-of-house**	**Overall job offers, back-of-house**	**Selection for interview, back-of-house**	**Job offers for those selected for interview, back-of-house**
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Age 40 or over	-.101 (.025)	-.068 (.010)	-.048 (.007)	.008 (.014)	-.123 (.017)	-.041 (.009)	.020 (.011)	-.112 (.017)
Baseline job offer/selection rate for those under age 40	13.4%	13.0%	8.7%	32.3%	26.9%	14.9%	37.5%	39.8%
Percent difference relative to baseline rate for under age 40, age 40 or over	-75.5%	-52.3%	-55.3%	2.6%	-45.7%	-27.7%	5.2%	-28.0%
Sample size	6,289	3,069	36,162	36,162	11,894	11,505	11,505	4,428
% of sample size age 40 and over	18.9%	28.4%	10.6%	10.6%	12.3%	22.0%	22.0%	24.0%

Notes to Table 2 apply. All models include dummy variables for restaurants and an intercept.

Appendix Table A1: Success Rate of Accurint Finding Ages of Applicants

	Electronic applications	Paper applications
Total applications	51,800 [100%]	10,388 [100%]
Number of applications Accurint found date of birth	45,043 [86.96%]	8,424 [81.09%]
Number of applications Accurint did not find date of birth	6,757 [13.04%]	1,963 [18.90%]

Note: Accurint was credited with finding an age if they reported the year of birth for an applicant. The first number in each cell is the number of applicants. The share of total applications is reported below.

Appendix Table A2: Source of Age Data by Application Type

a. Electronic applications	Number of applications	Share of applications
Total number of electronic applications	51,800	100.00%
Number of applications where age was found	47,667	92.02%
Age found using employee roster (primary)	4,111	
Age found using Accurint (secondary)	41,508	
Age found using year of HS graduation (tertiary)	2,048	
Number of applications where age was not found	4,133	7.98%
b. Paper applications	Number of applications	Share of applications
Total number of paper applications for restaurants included in the analysis	10,388	100.00%
Number of applications where age was found	8,485	90.56%
Age found using employee roster (primary)	981	
Age found using Accurint (secondary)	7,504	
Number of applications where age was not found	1,903	9.44%

Note: If an application could be matched to the employee, the age was determined using the reported date of birth in the employee roster. If the applicant did not appear in the employee roster, the date of birth reported by Accurint was used to determine the age. For electronic applicants who reported a high school year of graduation and whose age was not identified using the employee roster or Accurint, the year of high school graduation was used to estimate the age. All ages are determined on the date the applicant applied. Applications with no recorded date of application are excluded.

Appendix Table A3: List of Controls in Regression Analyses

Variable	Coding	Control in the paper application analysis	Control in electronic application analysis
Are you over the age of 18?	Yes/No	Yes	No
Are you seeking part time or full time work?	Categorical		Yes
Are you willing to be cross-trained?	Yes/No		Yes
Are you willing to stay late in an emergency?	Yes/No		Yes
Are you willing to work both a lunch and dinner shift in the same day?	Yes/No		Yes
Are you willing to work both part time and full time?	Yes/No		Yes
“Pre-screen” assessment percentile	Numeric		Yes
“Pre-screen” assessment categories	Red, Yellow, Green		Yes
Do you have a high school diploma?	Yes/No	Yes	Yes
Do you have experience as a certified trainer?	Yes/No		Yes
Do you have legal right to work	Yes/No	Yes	Yes
Do you have previous full service restaurant experience as a cook/chef?	Yes/No		Yes
Do you have previous full service restaurant experience as a food-server/bartender?	Yes/No		Yes
Do you have previous full service restaurant experience as a host/service assistant?	Yes/No		Yes
Do you have previous full service restaurant experience in other kitchen functions?	Yes/No		Yes
Do you have previous full service restaurant experience serving wine/beer/alcohol?	Yes/No		Yes
Do you have reliable transportation?	Yes/No		Yes
Have you ever been convicted of a felony?	Yes/No	Yes	Yes
Have you ever been employed at a company restaurant?	Yes/No	Yes	Yes
How many years if experience do you have in full service restaurant?	Categorical		Yes
If applying for server or bartender are you legal to serve alcohol?	Yes/No		Yes
Is your schedule flexible so you can attend training?	Yes/No		Yes
Location	Categorical	Yes	Yes
Position applying for	Categorical	Yes	Yes
Talent assessment score	Numeric		Yes
What is your highest level of education	Categorical		Yes
What is your highest level of restaurant experience?	Categorical		Yes
What shifts can you work Friday?	Categorical		Yes
What shifts can you work Monday?	Categorical		Yes
What shifts can you work Saturday?	Categorical		Yes
What shifts can you work Sunday?	Categorical		Yes
What shifts can you work Thursday?	Categorical		Yes
What shifts can you work Tuesday?	Categorical		Yes
What shifts can you work Wednesday?	Categorical		Yes
Willing to work holidays?	Yes/No		Yes

Appendix Figure A1: Distribution of Employee Age on Employee Roster by Whether Accurint Returned a Date of Birth

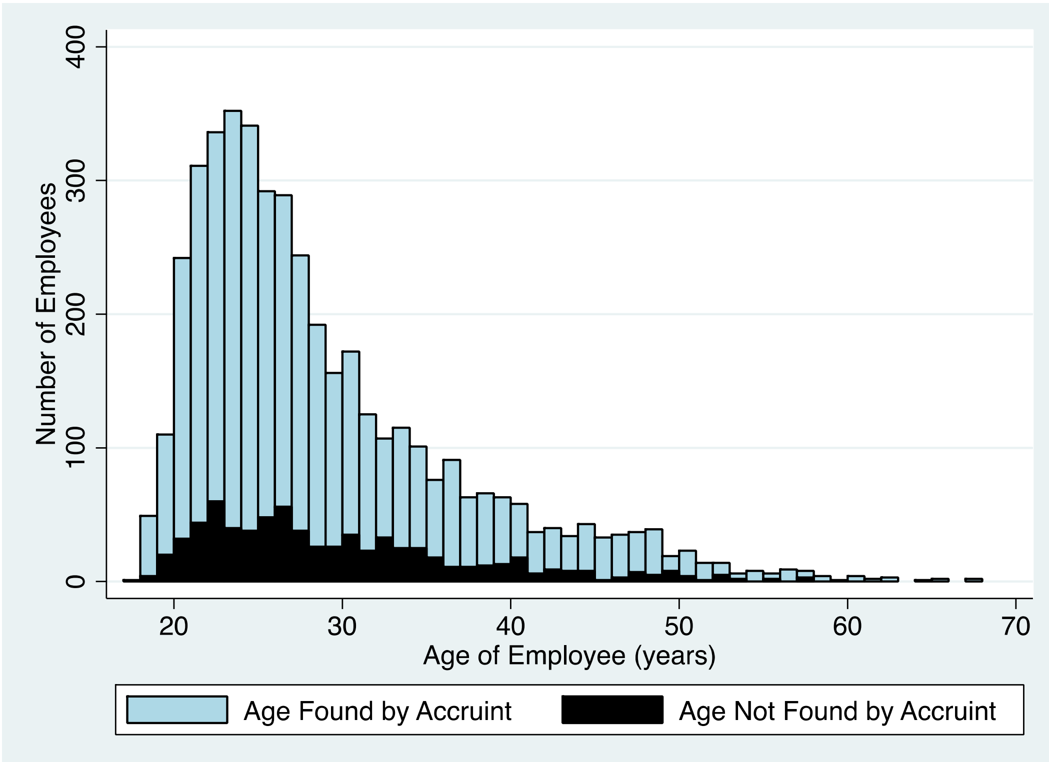

Note: There are 6,726 employees on the employee roster. For the 5,100 employees for whom I am able to match to an application in the sample of applications used in the analyses, I calculate the age of employees rounded to the nearest year.