WORKING PAPER #619
PRINCETON UNIVERSITY
INDUSTRIAL RELATIONS SECTION
May 2018
SOLE REVISION
Version: June 15, 2018

Whom Do Employers Want? The Role of Recent Employment and Unemployment Status and Age*

Henry S. Farber
Princeton University

Chris M. Herbst
Arizona State University

Dan Silverman
Arizona State University

Till von Wachter
University of California Los Angeles

Abstract

We use a résumé audit study to better understand the role of employment and unemployment histories in affecting callbacks to job applications. We focus on how the effect of career history varies by age, partly in an attempt to reconcile disparate findings in prior studies. While we cannot reconcile earlier findings on the effect of unemployment duration, the findings solidify an emerging consensus on the role of age and employment on callback. First, among applicants across a broad age range, we find that applicants with 52 weeks of unemployment have a lower callback rate than do applicants with shorter unemployment spells. However, regardless of an applicant's age, there is no relationship between spell length and callback among applicants with shorter spells. Second, we find a hump-shaped relationship between age and callback, with both younger and older applicants having a lower probability of callback relative to prime-aged applicants. Finally, we find that those applicants who are employed at the time of application have a lower callback rate than do unemployed applicants, regardless of whether the interim job is of lower or comparable quality relative to the applied-for job. This may reflect a perception among employers that it is harder or more expensive to attract an applicant who is currently employed.

---

*Farber: farber@princeton.edu; Herbst: Chris.Herbst@asu.edu; Silverman: dsilver3@asu.edu; von Wachter: tvwachter@econ.ucla.edu. We thank Jeong Joo Ahn, Nicole Gandre, Zoram Kaul, Elizabeth Kuttner, Tanairi Ochoa-Martinez, and Jason Shumberger for their first-rate work in fielding the study and preparing the data. Financial support for this project was provided by the Industrial Relations Section at Princeton University. We thank Fabian Lange and participants in a workshop at Brigham Young University for helpful comments. Earlier work on this project formed the basis of Farber's Presidential Address to the Society of Labor Economists, Seattle WA, May 6, 2016.

1 Introduction

The Great Recession led to a dramatic rise both in the unemployment rate and in durations of unemployment. Over the subsequent seven plus years of economic expansion, the unemployment rate has declined sharply but the long-term unemployment rate and the share of long-term unemployment among the unemployed remain at unusually high levels.[1] Moreover, there was a sharp decline in the labor force participation rate and in the employment-population ratio during the Great Recession, with little recovery since.[2] Along with the difficulty the unemployed have in finding a new job, the quality of job matches formed in recessions tends to be lower (e.g., Bowlus 1995, Farber 1999), and job losers often need to pass through a long series of job moves towards higher-paying jobs (e.g., Stevens 1997, von Wachter, Manchester, and Song 2009). Finally, there is mounting evidence that this job search and matching process is particularly challenging for those just entering the labor market (e.g., Kahn 2010, Schwandt and von Wachter 2017) and for older workers who are unlucky enough to lose a job in a recession (e.g., Chan and Stevens 2001, Farber 2017).

The elevated long-term unemployment rate, the substantial fraction of unemployment that is long-term, and the protracted job instability and low job quality that job losers experience reinforce the importance of understanding the determinants of successful job search. In particular, given the job losses associated with a major recession, it is important to understand the roles played by specific features of employment history in determining the employment process.

To better understand the determinants of job search outcomes, our work uses an audit study approach. The approach introduces random variation in the characteristics of fictitious résumés submitted in response to posted job advertisements and examines the effects of these characteristics on callback rates.[3] Recent audit studies have focused on unemployment du-

[1] The long-term unemployment rate has fallen substantially since its peak of 4.2 percent in 2010, but it remains (at about 1.1 percent in 2017) above its rate in earlier strong labor markets. Perhaps more striking, the long-term share of unemployment, which peaked at 44 percent in 2011, remains very high by historical standards at 24 percent in 2017. The long-term share of unemployment was 18 percent in 2007, prior to the Great Recession. These statistics are based on the authors' calculations using data from the Current Population Survey.

[2] The labor force participation rate fell from 66 percent in 2007 to 64 percent in 2012 and fell further to 63 percent in 2017. Over the same period, the employment-population ratio fell from 63 percent in 2007 to 59 percent in 2011 before recovering slightly to 60 percent in 2017. These statistics are based on the authors' calculations using data from the Current Population Survey.

[3] An important early example of an audit study in labor economics is Bertrand and Mullainathan (2004), who study the effect of African-American sounding names on callback rates. Jarosch and Pilossoph (2016)

ration, the role of age, and, to a lesser degree, on the presence and quality of a current (often interim) job. We provide an integrated assessment of the effect of unemployment duration, age, and interim job status on callbacks by randomly varying applicant characteristics in these dimensions on résumés submitted in response to a large number of job postings.

In doing so, the current study seeks to reconcile contrasting results on unemployment duration in the existing literature (reviewed below) and to integrate results on age and recent employment history. Although our earlier work (Farber, et al., 2016; 2017) is methodologically similar to Kroft, et al., (2013), they find a generally negative relationship between the callback rate and unemployment duration while we found no such relationship. In Farber et al. (2017), we attempt a detailed reconciliation with Kroft, et al. (2013), focusing on potential differences between the two analyses with respect to the measure of callback, type of job, applicant education, applicant age, and the particular cities included. Age was the the primary dimension that we could not rule out as the source of the different findings because of a lack of overlap in age between the two studies.[4] Hence, an important goal for this study is to return to the field with a design that varies unemployment durations and includes applicants covering a broad range of ages.

While not developed formally here, a simple Bayesian updating model of employer learning about worker quality from résumé characteristics is a useful framework for interpreting the experimental design and results.[5] A clear implication of such a model is that résumé characteristics that are perceived as negative (positive) will result in lower perceived applicant quality and lower (higher) callback rates. For example, if a longer unemployment spell is a negative signal of worker quality, then callback rates will be lower for applicants with longer unemployment spells. Another implication is that employers, when evaluating older applicants, will place less weight on the duration of a recent unemployment spell because there is additional information available about their labor market performance through previous experience. In other words, unemployment duration will have a stronger relationship with the callback rate for younger workers than for older workers.

In pursuit of an integrated assessment of several obstacles to successful job search after job loss, we fielded an audit study in 2017 varying treatments of unemployment duration,

---

present a clear analysis illustrating the limitations of résumé audit studies in measuring variation in *hiring* that highlights the distinction between a callback and the eventual hiring decision.

[4]Kroft, et al. (2013) analysis focused on young applicants, aged 21-33, while Farber, et al., (2016, 2017) focused on older applicants, aged 35-58.

[5] A more detailed presentation of this model is available in Farber, et al. (2017).

age, and the current employment status of the applicant. We kept all other aspects of the résumés as similar as possible to those in our earlier work. We visited the same eight cities with female applicants of similar education (college graduates) applying for the same type of jobs (white collar office jobs such as administrative or executive assistants, receptionists, secretaries, and the like). Of course we could not hold fixed the state of the labor market. The labor markets in our cities are uniformly stronger in 2017 than they were when we fielded our earlier experiment (2012-2014).

We obtain several key findings. First, our results for the currently employed are clear. Average callback rates are significantly lower for those applicants with an interim job, regardless of its quality than for those who are unemployed. The unemployed had an average callback rate of 12.6 percent, while the rate was 10.8 percent and 10.2 percent for those holding a high and low quality interim job, respectively. Interestingly, this negative effect of holding an interim job obtains only when applying to a relatively high-skilled position. Callbacks for clerical, data entry, office assistant, and other lower-skilled positions are similar for employed and unemployed applicants. Those without a job receive callbacks at a higher rate only for administrative or executive assistant, office manager, or other relatively high skilled positions, which also tend to have higher education requirements. These results suggest that the path back to career work may be more challenging for many workers than is typically understood because, perhaps surprisingly, on-the-job search can be significantly less fruitful when employed than when searching from unemployment.

Second, there are substantial effects of age, but these cut two ways; both the youngest (age 22-23) and the oldest (age 60-61) applicants have significantly lower callback rates than do prime-age applicants. Hence, we confirm the negative effects for older workers reported in other audit studies. We show that the negative effects for younger workers is partly accounted for by the fact that very young job applicants often have less relevant work experience than recommended or required for the job. Any remaining deficit for younger workers dovetails with a longstanding literature on the adverse consequences of early employment shocks.

Third, we find a limited effect of unemployment duration on the likelihood of callbacks. There are no differences in callback rates among applicants with unemployment durations of 24 weeks or less, but applicants who have been unemployed for a year have a lower callback rate. The callback rate was 11.6 percent for applicants with 24 weeks of unemployment, while applicants with 52 weeks of unemployment had a callback rate of 9.1 percent. This finding adds further complexity to the contrasting findings on unemployment duration dependence from prior audit studies.

Finally, we find little evidence of differences in the effect of unemployment duration by age. Hence, there is little reason to think that differences in the ages of the applicants can explain the differences in results across previous studies. We find average callback rates are, for all age categories, at their lowest after a year of unemployment.

One broad conclusion is that, while résumé audit studies may be useful in shedding light on the process of job search and hiring, not enough is known about how, why, and when employers respond to résumé characteristics when making their callback decisions. There is consistent evidence that some features of an applicant's résumé (e.g., age or current employment status) elicit similar responses from employers when studied in different contexts, but it is unclear why this is not true for unemployment duration. These uncertainties suggest caution in interpreting the results from audit studies and potential limits to their external validity, especially in the case of the relationship between callbacks and unemployment duration.

1.1 Related Audit Study Literature

An increasing number of résumé audit studies have been implemented to assess the impact of résumé characteristics on the likelihood of callback. One especially relevant subset of the literature focuses on the effect of the duration of an applicant's unemployment spell on callback rates. These studies have obtained mixed results. Some find a negative effect of unemployment duration on the probability of callback concentrated in the first year of the unemployment spell (Kroft et al., 2013; Ghayad, 2014). In contrast, Nunley et al. (2017) and Farber et al. (2016; 2017) find no effect of unemployment duration on callback rates in U.S. markets. Focusing on Sweden, Eriksson and Rooth (2014) find no effect for highly educated workers and a negative effect for spells lasting *more* than 9 months for less educated workers. Nüß (2018) finds similar results for secondary school graduates in Germany, with lower callback rates beginning at 10 months of unemployment. One difficulty in comparing and assessing the different outcomes of these studies is that they focus on workers of different ages and levels of education. Kroft et al., (2013) and Ghayad (2014) focus on younger workers for a range of education groups, Nunley et al. (2017) focus on younger college graduates, and Farber et al. (2016; 2017) focus on older college graduates.

Motivated by the difficulties unemployed older workers face in the labor market, another group of audit studies addresses how applicant age affects callback rates. The results generally show a negative relationship between age and callbacks for older and middle-aged workers. Lahey (2008) finds large negative effects of age on callbacks for women seeking entry-level positions in the U.S. Neumark et al. (2015; 2016), in a large scale audit study, find significantly lower callback rates for older women but weaker evidence for men in several occupations. Bendick, et al. (1997) find evidence that both older men and older women had lower callback rates than their younger counterparts. Farber, et.al. (2016; 2017) find that older females have lower callback rates than middle-aged females. So far, no study has compared callback rates by age for younger vs. middle aged vs. older workers. This open question, which we address here, takes on greater importance given the increasing evidence that younger workers are particularly hurt by adverse labor market conditions as well (e.g., Kahn 2010, Oreopoulos, et al. 2012, Schwandt and von Wachter 2017).

Lastly, a smaller number of studies assess the effect of holding a job at the time of application on the probability of a callback. Farber, et al. (2016; 2017) find that workers who are currently employed in an interim job at a lower level than the one for which they are applying have a lower callback rate than those who are unemployed. This finding is similar to that in Nunley et al. (2017), who found that recent college graduates in the U.S. had substantially lower callbacks if they were currently employed in jobs that did not require a college education and were not suited to the job for which they were applying. A related but distinct finding of Kroft, et al. (2013) is that individuals who are currently employed in a job *similar* to the one for which they are applying are significantly less likely to receive a callback. These findings are consistent with the idea that employers believe employed job seekers are more difficult to hire, perhaps because they are likely to have a higher reservation wage and be more expensive to recruit. An open question that we address here is whether the job type at which workers are employed matters, or whether employment, per se, lowers callback rates.

2 Research Design

The design of our study closely follows, and then extends, the design in Farber et al. (2017). As with prior correspondence audit studies, the design randomly assigned several job-seeker characteristics to résumés submitted on behalf of fictitious individuals in response to real job advertisements. We then recorded whether each résumé received a positive response. The setting for the experiment consists of two large on-line job boards in the U.S. We used these websites to search for jobs in 8 cities: Charlotte, Chicago, Dallas, Omaha, Pittsburgh, Portland (ME), Sacramento, and Tampa. These cities were originally selected by Farber et al. (2017) to allow for treatment effect differences by local-area unemployment rates. Therefore,

Table 1: Characteristics of Posted Jobs

Job Category	Skill	Percent	Education Category	Percent
Receptionist	Low	22.2	High School Degree	28.3
Office Assistant	Low	6.6	Associate's Degree	8.6
Clerical / Data Entry	Low	4.3	Bachelor's Degree	11.1
Other – Low Skill	Low	1.0	Not Available	52.0
Administrative Assistant	High	52.6		
Executive Assistant	High	6.4		
Office Manager	High	4.5		
Other – High Skill	High	2.3		

Note: Job Category is a classification based on the posted job title. Skill classification is subjective. Education Category indicates the highest education level recommended or required for posted position.

half of the cities studied had relatively low unemployment rates in 2012 (Dallas, Omaha, Pittsburgh, and Portland) while the other half had relatively high unemployment rates in 2012 (Charlotte, Chicago, Sacramento, and Tampa). Unemployment rates fell substantially between 2012 and 2017, but the general ordering of cities as low- and high-unemployment remained largely unchanged.[6]

To reduce idiosyncratic variation in interview requests, we made design choices aimed at ensuring that the résumés and job openings were well-matched. Applications were limited to white collar office positions, including administrative or executive assistants, receptionists, secretaries, and office associates. Our procedure for finding positions to which résumés would be submitted entailed searching for particular keywords in the position description.[7] Importantly, there were more than 1,200 different job titles associated with the jobs openings to which we applied. We then did a textual analysis of the job titles associated with the selected positions in our sample and subjectively assigned the job titles to the 8 categories listed in the left panel of Table 1. Most jobs fall in the 2 categories Administrative Assistant and Receptionist. We further collapse these 8 job categories into 2 groups: 1) "low skill" (Receptionist, Office Assistant, Clerical / Data Entry, and Other – Low), comprising 34.1 percent of postings, and 2) "high skill" (Administrative Assistant, Executive Assistant, Office Manager, and Other – High), comprising 65.9 percent of postings. We will utilize this low-

[6] Given the general strengthening of labor markets over the past 7 years, the distinction between the low- and high-unemployment cities is less salient than it was at the time we designed our previous study, and we do not evaluate the low-high unemployment distinction in the present analysis.

[7] These terms included Administrative Assistant (90.5 percent), Receptionist (6 percent), Office Assistant (2.3 percent), and Executive Assistant (1 percent).

high skill distinction in our analysis.

Given that these jobs are disproportionately held by women, all applicants are assigned female names. Each applicant has a four-year bachelor degree from a non-elite public university or college whose admission rate exceeds 65 percent. The distribution of required or recommended educational attainment in the job postings is listed in the right column of Table 1. Fully 52 percent of postings do not list an education level. A high school degree is specified in 28.3 percent of postings, an associate's degree in 8.6 percent, and a bachelor's degree in 11.1 percent. The posted education level is correlated with the skill level of the job category as we define it.[8] Among postings specifying a bachelor's degree, 84.4 percent are high skill. Among postings specifying an associate's degree, 79.4 percent are high skill. The fraction high-skill falls to 62.6 percent among postings specifying a high school degree. The fraction high-skill is 61.5 percent among postings with no specified education level, suggesting that these postings are more likely to be open to applicants with less education. Given that all our applicants have a four-year bachelor's degree, they may be over-qualified for some of the low-skilled job postings. For this reason, we present some analyses making a distinction between the postings we classify as low-skill and those we classify as high skill.

Many of our applicants have substantial work histories, which are perforce longer for older applicants. The work histories generally consist of three to six white collar jobs. Prior to the current unemployment spell, these work histories had no employment gaps longer than a month in the previous five years. Age and birth year are not listed on the résumés but are instead inferred from information on year of college completion and work experience.

The fictitious résumés were modeled on real résumés posted publicly on a website that was not used for the purposes of data collection for this study. These "source" résumés were used by individuals to apply for positions similar to the target occupations we study but were in cities that were not included in the study. We migrated the characteristics of these source résumés to our fictitious résumés, modified suitably so that, city-by-city, the residential addresses, employers, and post-secondary education institutions are geographically appropriate. We did not migrate the names listed on the source résumés, but instead selected common, cohort-specific names from lists maintained by the Social Security Administration. These résumé names are racially and ethnically neutral in that they are not held disproportionately by a particular racial or ethnic group.[9]

[8] Education level was not used to assign skill category.

[9] The names used in this study were Susan Taylor, Lauren Daniels, Donna Ramsey, Rose Peters, Angela Nelson, Heather Adams, Jennifer Smith, Shannon Robinson, Janice Evans, Linda Carter, Mary Wilcox, and

Fieldwork began in March 2017 and ended in August 2017. A group of research assistants regularly searched both websites for job advertisements that had been posted within the past 7 days and that were located within a 25-mile radius of a given city. The 7-day criterion increased the odds that our résumés were evaluated by employers, while the 25-mile radius ensured that the pool of openings included a sufficiently large number of jobs. We sent 4 résumés in response to each job advertisement. Each résumé was accompanied by a cover letter that provided a broad summary of the applicant's credentials listed on the résumé.[10] On one of the job websites, résumés were submitted directly to employers via email addresses provided by the website host. Résumés on the other website were forwarded to employers through the job board's on-line submission portal. To minimize suspicion that the four résumés were "connected," we let several hours elapse between each submission. The research assistants collected information about each job advertisement and firm, including the job title, name and location of the firm, contact information, industry code, and any experience or education requirements. Altogether our analysis data-set includes 8,488 résumés submitted in response to 2,122 job advertisements.[11]

The essence of the experiment was to manipulate three variables: length of contemporaneous unemployment spell, presence and nature of contemporaneous "interim jobs," and age of applicant. Regarding the first source experimental variation, the work histories were designed such that a base set of résumés indicated that the applicant was just entering unemployment at the time of résumé submission, while other résumés indicated a specific number of weeks of unemployment from the set {4, 12, 24, 52}. The start of the unemployment spell was indicated on the résumé by the end date of the applicant's most recently held job. Therefore, those in the base group listed on her résumé that her most recent job had ended in the month immediately preceding the month in which the application was submitted. The unemployment spells were drawn with replacement from the set {0, 4, 12, 24, 52} and were intended to be assigned with equal probability. There was, however, a glitch in the implementation of the sampling during May, June, and July 2017, and the probabilities were

Joan Mason.

[10] The working paper version of this study (Farber, et al., 2018) contains an appendix that includes a sample of three cover letters and their associated résumés.

[11] The complete data set contains 8,899 résumés submitted in response to 2,304 job postings. A small number of job postings (182) inadvertently received less than four submissions and we delete all the submissions to these 182 postings from our analysis. None of our results is affected by omitting the 411 completed submissions to these 182 postings.

not equal for applications during that period.[12]

The second source of experimental variation assigned applicants to one of three interim-job categories: a low-skilled interim job, a high-skilled interim job, and no interim job. The low-skilled interim job was shown to begin the month in which a given résumé was submitted (following whatever length unemployment spell was assigned to that résumé), and included employment at a chain restaurant, a big box retail store, or a grocery store. These jobs generally involved serving food, stocking shelves, or assisting customers at a register or on a retail floor. Thus this job experience was designed to appear quite different from the career work outlined on the rest of the résumé and is generally of lower skill and responsibility than the posted jobs. The high-skilled interim jobs also began in the month of résumé submission (following whatever length unemployment spell was assigned to that résumé). However, this category included jobs in industries and occupations that matched applicants' broader work history. While the responsibilities on these interim jobs were somewhat downgraded relative to the responsibilities listed for the most recently held job, these were more comparable to the posted jobs and particularly to the lower-skill posted jobs as described above. The third group of résumés did not list an interim job so that these applicants appeared to be currently unemployed along with some unemployment duration as described above. These three treatments were assigned with replacement and with equal probability.

The third source of experimental variation assigned each applicant an age from the set {22-23, 27-28, 33-34, 42-43, 51-52, 60-61}. Although age was not stated on the résumés, it could be inferred by employers from the college graduation dates as well as the start (and end) dates of the jobs listed throughout the employment history. At the beginning of every month we adjusted the employment dates accordingly so that our fictitious applicant did not "age" throughout the study. Age categories were assigned without replacement within a job posting and with equal probability.

The experiment was designed so each dimension of variation (unemployment duration, interim job, and age) was assigned independently. This facilitates the examination of interactions across these dimensions. For example, there might be a differential impact of unemployment duration or interim job by age.

To enable communication between employers and our fictitious applicants, we created a

[12] During these three months, most of the resumes that should have had a 24 week unemployment duration actually had a zero unemployment duration. The observed fractions with each unemployment treatment ranged, over the entire period, from 0.128 with 24 weeks duration to 0.279 with 0 weeks duration. The other three durations (4, 12, and 52 weeks) had fractions very close to 0.20.

telephone number for each city-name combination in the study, for a total of 96 numbers (8 cities * 12 names). An email address was also established for our applicants according to the same rules. The research assistants regularly monitored these accounts and recorded whether résumés received a message from an employer. In particular, they coded three categories of responses: interview requests, a request for more information, and rejections. An interview request is defined as the receipt of either an explicit interview request or an invitation to discuss the résumé and/or position in more detail. A request for more information was coded if the applicant was asked to fill out an application, answer clarifying questions, provide additional information, or interview for a different position. A rejection was coded if the applicant received a message from the employer saying that the résumé was not selected to move to the next round of the hiring/interview process. Résumés that received no correspondence from an employer were ultimately coded as a rejection. Potential employers who did respond were promptly informed that the applicant was no longer available.

3 Results

We begin by analyzing the effect of unemployment duration on the likelihood of a callback. In an attempt to reconcile our earlier findings (Farber, et al, 2016, 2017) with findings of Kroft, et al. (2013), we then investigate how the relationship between unemployment duration and the likelihood of callback varies with age. Next, we consider how the presence of an interim job affects the likelihood of callback. As part of this analysis, we consider the extent to which holding an interim job of different types (low skill vs. high skill) has differential effects on callback rates and how these effects are related to the skill level of the applied-for job (also low skill vs. high skill).

We provide three measures of what constitutes a callback: 1) any callback, 2) a callback asking for more information, and 3) a callback requesting an interview. Overall, the any-callback rate is 11.2 percent, the information-callback rate is 3.0 percent, and the interview-callback rate is 8.2 percent. One plausibility check that our resumes work as intended is that the any-callback rate was significantly (p-value < 0.0005) higher in our low-unemployment cities (at 12.6 percent) than in our high-unemployment cities (at 10.1 percent). The same relationship holds considering information callbacks and interview callbacks separately. In what follows, we present some first-order results that consider different types of callbacks separately, but we largely emphasize any-callbacks.

Table 2: Average Callback Rate, by Unemployment Duration

Weeks U	N Apps	(1) Callback (Any)	(2) Callback (Information)	(3) Callback (Interview)
0	2370	0.110	0.034	0.076
		(0.007)	(0.005)	(0.006)
4	1768	0.118	0.036	0.081
		(0.008)	(0.005)	(0.007)
12	1632	0.123	0.026	0.097
		(0.009)	(0.004)	(0.008)
24	1083	0.119	0.029	0.090
		(0.011)	(0.006)	(0.010)
52	1635	0.091	0.021	0.070
		(0.008)	(0.004)	(0.007)
All	8488	0.112	0.030	0.082
		(0.005)	(0.003)	(0.004)
Pearson χ^2		10.5	8.7	10.4
p-value		0.033	0.068	0.034

Note: Numbers parentheses are standard errors clustered by job id. The Pearson χ^2 and associated p-value refer to a test of independence between unemployment duration and the measure of callback. Sample consists of 8,488 applications for 2,122 job postings.

3.1 Duration of Unemployment

A primary focus of this study is to examine the effect of unemployment duration on the likelihood of receiving callback to a job application. Table 2 contains mean callback rates for each of the three callback measures and each of the five values for unemployment duration. The table also contains the results of Pearson χ^2 tests of the independence of callbacks and unemployment duration. Independence can be rejected for any-callback and for each category of callback, but there is not a progressive relationship between callbacks and unemployment duration. The callback rate is lower for those with the highest unemployment duration (52 weeks) than for those with shorter unemployment spells, but there does not seem to be a distinction in callback rates for those with unemployment spells between 0 and 24 weeks. Pearson χ^2 tests of independence of each measure callback and unemployment duration omitting the 52-week category fails to reject independence (all p-values > 0.4). Hence, the overall rejection of independence is entirely driven by the lower callback rates for applicants with 52 weeks of unemployment.

We will revisit the relationship between the callback rate and unemployment duration in more detail later with particular focus on how the relationship varies with age. But the

Table 3: Average Callback Rate, by Age of Applicant

Age	N Apps	(1) Callback (Any)	(2) Callback (Information)	(3) Callback (Interview)
22-23	1304	0.094	0.028	0.067
		(0.009)	(0.005)	(0.007)
27-28	1572	0.111	0.029	0.083
		(0.008)	(0.004)	(0.007)
33-34	1460	0.129	0.034	0.095
		(0.009)	(0.005)	(0.008)
42-43	1368	0.126	0.037	0.089
		(0.009)	(0.005)	(0.008)
51-52	1439	0.110	0.031	0.080
		(0.009)	(0.005)	(0.008)
60-61	1345	0.097	0.022	0.075
		(0.008)	(0.004)	(0.008)
All	8488	0.112	0.030	0.082
		(0.005)	(0.003)	(0.004)
Pearson χ^2		14.1	6.6	9.0
p-value		0.015	0.252	0.110

Note: Numbers parentheses are standard errors clustered by job id. The Pearson χ^2 and associated p-value refer to a test of independence between age category and the measure of callback. Sample consists of 8,488 applications for 2,122 job postings.

simple comparison of means suggests that, while there are lower callback rates for very long unemployment spells (52 weeks), there is no systematic relationship between callback rates and unemployment duration for spells of one-half year or less.

3.2 Age

Table 3 contains mean callback rates for each of the three callback measures and each of the six age groups. The table also contains the results of Pearson χ^2 tests of independence of callbacks and age category. While independence can be rejected for any-callbacks, there is not a monotone relationship between callbacks and age. It is clear that the callback rate is lower for the youngest (22-23 year old) and oldest (60-61 year old) applicants. The highest callback rate was received by 33-34 and 42-43 year old applicants. This inverted U-shape suggests that employers are most interested in hiring prime-age applicants for these jobs and are less interested in hiring the youngest and oldest applicants.[13]

Interestingly, while existing audit studies on age (including our earlier work) have mainly focused on older applicants, our point estimates indicate that the youngest applicants have significantly lower callback rates relative to prime-age applicants. It may be that employers are less interested in hiring the youngest workers because they have less experience. As noted by Neumark, et al. (2015), when varying age, prior work experience must vary as well. Younger workers simply cannot have as much experience as older workers. To the extent that employers value prior work experience, the lower callback rates for the youngest applicants could be due to their having insufficient experience.

We investigate this directly using information in the posted job descriptions on recommended/required experience for the posted jobs. Of the 2,122 job postings that make up our analysis sample, 820 postings (38.6 percent) did not mention work experience, and 164 postings (7.7 percent) explicitly stated that no experience was required. Among the remaining 53.6 percent of postings that required at least some experience, 570 (26.9 percent of all postings) recommended/required one year or less experience, 491 (23.1 percent of all postings) recommended/required more than one year to 4 years experience, and 77 (3.6 percent of all postings) recommended/required more than 4 years experience (maximum 10). All of our applicants held a BA degree, implying that our youngest applicants (22-23 years old) could not have had more than a year or so of work experience at the time of application and, as a result, may be deemed unqualified for jobs requiring more experience.

Table 4 contains mean callback rates by age category and the posted experience requirement. One obvious pattern is that callback rates for applicants age 22-43 are generally higher for applications to postings with no explicit mention of experience (first row of the table). More generally, the table provides support for a strong interaction between age and posting experience recommendations/requirements in the relationship of callback rates with age. It is only for the youngest two age groups (22-23 and 27-28, estimates in columns 1 and 2 of Table 4) that there is a statistically significant relationship between posted experience and the callback rate. These two age groups show very high callback rates to postings either with no statement of recommended/required experience or with an explicit statement of no experience required. It appears that employers posting jobs where experience is not important prefer younger applicants. Callback rates fall off sharply for younger applicants as recommended/required experience increases. There are not significant differences in callback

[13] The same general pattern hold for information-callbacks and interview-callbacks, but the differences are not statistically significant.

Table 4: Average Callback Rate, by Recommended/Required Experience and Age

Rec/Req Experience	(1) Age 22-23	(2) Age 27-28	(3) Age 33-34	(4) Age 42-43	(5) Age 51-52	(6) Age 60-61
Exp N/A	0.117	0.139	0.145	0.138	0.108	0.104
	(0.015)	(0.015)	(0.016)	(0.016)	(0.014)	(0.014)
Exp = 0	0.176	0.165	0.116	0.065	0.093	0.091
	(0.040)	(0.035)	(0.036)	(0.024)	(0.031)	(0.034)
0 <Exp<= 1	0.075	0.080	0.102	0.115	0.093	0.064
	(0.016)	(0.014)	(0.016)	(0.018)	(0.016)	(0.013)
1 <Exp<= 4	0.056	0.096	0.140	0.144	0.133	0.122
	(0.013)	(0.016)	(0.020)	(0.021)	(0.020)	(0.020)
Exp> 4	0.064	0.036	0.102	0.089	0.155	0.113
	(0.046)	(0.025)	(0.042)	(0.042)	(0.052)	(0.044)
All Exp	0.094	0.111	0.129	0.126	0.110	0.097
	(0.009)	(0.008)	(0.009)	(0.009)	(0.009)	(0.008)
N Apps	1304	1572	1460	1368	1439	1345
Pearson χ^2	18.7	16.6	4.8	4.4	6.3	7.3
p-value	0.001	0.002	0.307	0.179	0.358	0.121

Note: Standard errors clustered by job id are in parentheses. The Pearson χ^2 and associated p-value refer to a test, within age category, of independence between experience category and any-callback.

rates by recommended/required experience for applicants 33 and older (columns 3-6 of Table 4).

3.3 Does the Effect of Unemployment Duration vary with Age?

Recall that Kroft, et al. (2013) found a strong negative relationship between the likelihood of receiving a callback and unemployment duration among younger workers that was concentrated early in the unemployment spell. In contrast, Farber, et al. (2016, 2017) found no such relationship when considering older workers. One possibility is that unemployment could be more salient (get a higher weight in the employer's inference process) for young workers due to their relative lack of experience. Our current analysis (Table 2), using a full spectrum of ages, suggests no relationship between the callback rate and unemployment duration for spells of unemployment up to 24 weeks but a lower likelihood of callback for very long unemployment spells (52 weeks). In this subsection, we examine the relationship between unemployment duration and the likelihood of callback separately by age group.

Our experimental design included six age groups. We showed in the previous subsection that there is virtually no difference in callback rates among the middle four age groups (27-28, 33-34, 42-43, and 51-52). While we do not show the detailed results here, we verified

Table 5: Average Callback Rate, by Unemployment Duration and Age

Weeks U	(1) All Age	(2) Age 22-23	(3) Age 27-52	(4) Age 60-61
0 Weeks	0.110	0.065	0.123	0.096
	(0.007)	(0.013)	(0.009)	(0.016)
4 Weeks	0.118	0.121	0.119	0.110
	(0.008)	(0.019)	(0.010)	(0.018)
12 Weeks	0.123	0.108	0.131	0.097
	(0.009)	(0.021)	(0.010)	(0.020)
24 Weeks	0.119	0.144	0.117	0.107
	(0.011)	(0.027)	(0.013)	(0.022)
52 Weeks	0.091	0.056	0.101	0.075
	(0.008)	(0.016)	(0.010)	(0.016)
F-Statistic	2.89	3.80	1.39	0.62
p-value	0.021	0.005	0.235	0.065
N Applications	8488	1304	5839	1345

Note: Numbers parentheses are robust standard errors clustered by job id. The F-statistic and associated p-value refer to a test of equality of callback rates by unemployment duration within age category.

that there are no significant patterns or differences in the relationship between callback rates and unemployment duration across these four age groups. To simplify the exposition, we proceed showing separately the relationships across 3 age groups (22-23, 27-52, and 60-61). We can think of these as young, prime-aged, and older applicants.

Table 5 contains the mean callback rates by unemployment duration both overall and for each of the three age categories. The first column of the table repeats the means reported in column 1 of Table 2 and shows the lower callback rate for applicants with 52 weeks unemployment. The callback rates in column 2 of the table are for the youngest applicants (22-23 years of age), and they show an inverted U-shape. Callback rates for young applicants with zero unemployment duration and with 52 weeks of unemployment are significantly and substantially lower than for those with 4-24 weeks of unemployment. The callback rates in column 4 of the table are for the oldest applicants (60-61 years of age), and they show a significantly lower callback rate for applicants with 52 weeks unemployment. There are no significant distinctions in callback rates by weeks of unemployment for older applicants with 24 weeks or less unemployment. There are no significant differences in callback rates by unemployment duration for the prime-aged applicants (column 3 of the table).

This analysis indicates that the difference in estimated duration dependence between Farber, et al. (2016, 2017) and Kroft et al. (2013) is likely not due to the fact that the two studies focused on different age groups. While we find a significantly lower callback rate at 52 weeks unemployment for the younger applicants in our study, there is not a general declining pattern of callback rates at lower durations of unemployment. In contrast, Kroft, et al. (2013) found a consistent negative relationship between the callback rate and duration for durations from 3 to 10 months. Unfortunately, our design does not allow us to examine callback rates to applications with callback rates intermediate between 24 and 52 weeks, and the contrast in results remains a puzzle.

3.4 Interim Jobs

Our experimental treatment of interim-job status included three values: 1) no interim job, 2) a lower quality interim job, and 3) a higher quality interim job. We do not have a clear expectation regarding the effect of holding an interim job of either type at the time of application on the likelihood of a callback. One question is whether an interim job is a good signal or a bad signal of worker quality. On the one hand, holding an interim job of any kind could be a positive signal of the applicant's drive or ambition and result in a higher likelihood of callback. On the other hand, a lower quality interim job could indicate that the applicant is lower quality and result in a lower likelihood of callback. If this channel is important we would expect that holding a lower quality interim job would reduce the likelihood of callback relative to otherwise-equivalent holders of a higher quality interim job.[14]

Another possibility is that employers find it easier to attract applicants (get them to accept a job offer) if the applicant is not employed. The fact that an applicant is employed might signal to an employer that the applicant's reservation wage is higher. As such, employers may prefer applicants who are not employed (no interim job). Alternatively, employers may be particularly concerned about pursuing applicants who are working in a higher-quality interim job but less concerned about pursuing applicants with no job or a lower-quality interim job.

Table 6 contains mean callback rates for each of the three callback measures and three interim job treatments. The table also contains the results of a Pearson χ^2 test of independence of callbacks and interim job status. Independence can be rejected for any-callbacks and interview-callbacks, and the pattern is clear. The highest callback rates are for those applicants who hold no interim job. There is no statistically significant difference between

[14] Note that, if employers update substantially based on an interim job, a natural inference is that employers weight current employment more heavily than past experience (all of which is in relatively high quality jobs) when forming an expectation of worker quality.

Table 6: Average Callback Rate, by Interim Job Status

Interim Job	N Apps	(1) Callback (Any)	(2) Callback (Information)	(3) Callback (Interview)
None	2786	0.126	0.033	0.092
		(0.007)	(0.004)	(0.006)
Lower Quality	2861	0.102	0.027	0.075
		(0.006)	(0.004)	(0.006)
Higher Quality	2841	0.108	0.030	0.078
		(0.007)	(0.004)	(0.006)
All	8488	0.112	0.030	0.082
		(0.005)	(0.003)	(0.004)
Pearson χ^2		8.8	2.3	6.4
p-value		0.012	0.323	0.042

Note: Numbers parentheses are standard errors clustered by job id. The Pearson χ^2 and associated p-value refer to a test of independence between interim job status and the measure of callback. Sample consists of 8,488 applications for 2,122 job postings.

callback rates for lower- and higher-quality interim job holders.

As we noted earlier, we classified the job posting to which we submitted résumés as "low-skill" and "high-skill." Given that our applicants' interim jobs also were "low" or "high" quality, it may be that the effect on the callback rate of holding an interim job of a given type could differ with the type of job for which the application is made.[15]

Table 7 contains mean any-callback rates by interim job status and skill level of the posted position. There is no significant relationship between callbacks and interim job status for the applications to the lower quality positions (mostly receptionists, office assistants, and clerical/data entry positions). However, among applications for higher quality positions (mostly administrative and executive assistants and office managers), there is a strong pattern. Holding an interim job of either lower or higher quality results in a significantly lower callback rate relative to those applicants who are unemployed at the time of application.

Overall, the analysis of interim-job status supports the view that employers, particularly those hiring into relatively high-skilled positions, prefer to hire those applicants who are

[15] That there is real information in our categorization of job postings as low-skill or high-skill can be confirmed using information we collected on any stated education requirements or preferences in the job description. Among those postings that did list an education level, the high-skill postings were more likely to list a bachelor's or associate's degree (48.5 percent of the high-skill listings with a stated education level) relative to the low-skill postings (25.7 percent of the low-skill listings with a stated education level). The working paper version of this study (Farber, et al., 2018) contains a detailed breakdown of required/recommended education level by skill level of posting.

Table 7: Average Callback Rate, by Interim Job Status and Position Quality

Position	N Apps	(1) All	(2) No Interim Job	(3) Low Quality Interim Job	(4) High Quality Interim Job	(5) χ^2 [p-value]
Lower Skill	2872	0.101	0.109	0.105	0.090	2.0
		(0.008)	(0.011)	(0.011)	(0.010)	[0.369]
Higher Skill	5616	0.117	0.134	0.100	0.117	10.6
		(0.006)	(0.010)	(0.008)	(0.009)	[0.005]
All	8488	0.112	0.126	0.102	0.108	8.8
		(0.005)	(0.007)	(0.006)	(0.007)	[0.012]

Note: Numbers parentheses are standard errors clustered by job id. The Pearson χ^2 in column (5) and associated [p-value] refer to a test of independence between interim job status and callback for each quality-of-position group. Sample consists of 8,488 applications for 2,122 job postings.

not employed. As we noted earlier, this is consistent with findings in the literature by Kroft, et al. (2013), Farber, et al. (2016, 2017), and Nunley, et al. (2017), and it suggests that employers seeking to fill at least the type of jobs we consider here prefer to hire the unemployed, perhaps because of a perception that they will be easier/cheaper to recruit.[16] The evidence in Table 7, that the preference for the unemployed exists only for applications to the more highly skilled positions, is weakly consistent with this argument since costs of recruitment are likely higher for higher quality jobs.

3.5 Multivariate Analysis

We now turn to a multivariate logit analysis that models the probability of a callback as a function of unemployment duration, age, and interim job status as well as posting-specific characteristics including skill level and education preference/requirement. This analysis first uses both within- and between-posting variation in application characteristics. We choose the logit model for several reasons. In principle, it should provide a better approximation of the functional form for binary choice probabilities with a relative low incidence.[17] Given the canonical sample design of recent audit studies that provide random variation *within* postings, a particular advantage of the logit model is that it provides a consistent approach

[16] While the findings in the literature pertain to a broader groups of job applicants, a particular note of caution regarding the external validity of this interpretation is warranted. Employers for other more skilled jobs may prefer to hire workers who are employed and so be willing to go through the difficulty and bear the expense of attracting the employed.

[17] We have reproduced the logit findings with linear probability and probit models, and the results are not affected by the choice of functional form.

that allows us to obtain estimates that rely on within-posting variation via the Chamberlain (1980) fixed-effects logit model. Finally, the logit model allows us to contrast the fixed-effect estimator with a random-effects logit estimator, which is our preferred specification.

The random-effects model accounts for the fact that job postings are randomly drawn from the underlying population and may differ in their mean callback rate. This model is appropriate (yields consistent estimates) where the baseline variation across job postings in their callback rates is uncorrelated with the observed applicant characteristics of interest. Given our approach in sending résumés to job listings with key characteristics varying randomly, we would not expect the job-specific callback rate to be correlated with résumé characteristics so that estimates derived using the random-effects model should be consistent. More generally, since the treatments were assigned independently to résumés there is no reason to expect that the multivariate analysis in general, and the conditional logit in particular, will affect our main results.

Table 8 contains the main results of our multivariate analysis of the probability of any-callback. We report our findings in terms of odds ratios, which for small probabilities are approximately the ratio of probabilities of callback given a particular treatment vs. a base condition.[18] An odds ratio greater than one implies that the associated factor increases the likelihood of a callback relative to the base group. Similarly, an odds ratio less than one implies that the associated factor decreases the likelihood of a callback relative to the base group. The base group in the table consists of applicants aged 27-52 with zero weeks unemployment and no interim job.[19]

The first column of Table 8 contains results for the basic logit model with controls for weeks of unemployment, age, and interim job status along with city fixed effects (not shown). These results confirm our earlier univariate findings. Applicants with 52 weeks of unemployment have a significantly lower, by more than 20 percent, probability of a callback (odds ratio significantly less than one) than applicants with less unemployment. There are no significant differences in callback rates among the unemployment durations shorter than 52 weeks. Both the youngest (22-23 years old) and oldest (60-61 years old) applicants have

[18] Let $p(1) \equiv Pr\{Callback = 1|X, D = 1\}$ and $p(0) \equiv Pr\{Callback = 1|X, D = 0\}$, where D represents one of our right hand side dummy variables, and X represents the remaining variables in the model. Then the odds ratio R is defined as $R \equiv \frac{p(1)/(1-p(1))}{p(0)/(1-p(0))} = exp\{\beta_D\}$, where β_D is the coefficient on D. Where the probabilities involved are small, the odds ratio is approximately the ratio of probabilities $\left(\frac{p(1)}{p(0)}\right)$.

[19] Given the earlier results on age, we simplified the estimation by combining the 27-28, 33-34, 42-43, and 51-52 year old applicants. There were no distinctions in callback rates among these groups. See Table 3.

Table 8: Logit Model Estimates of Any-Callback Probability : Odds Ratios
Logit, Random-Effects Logit, and Fixed-Effect Logit Models

Variable	(1) Logit	(2) Logit	(3) RE Logit	(4) RE Logit	(5) FE Logit
Applicant Characteristics:					
4 Weeks U	1.046	1.045	0.985	0.982	0.946
	(0.103)	(0.103)	(0.138)	(0.137)	(0.141)
12 Weeks U	1.094	1.077	1.168	1.154	1.151
	(0.111)	(0.109)	(0.165)	(0.163)	(0.172)
24 Weeks U	0.994	0.978	0.974	0.963	0.963
	(0.124)	(0.122)	(0.164)	(0.162)	(0.175)
52 Weeks U	0.779	0.772	0.625	0.621	0.610
	(0.085)	(0.084)	(0.094)	(0.094)	(0.096)
Age 22-23	0.762	0.757	0.580	0.580	0.563
	(0.077)	(0.077)	(0.080)	(0.080)	(0.079)
Age 60-61	0.792	0.790	0.749	0.748	0.786
	(0.078)	(0.078)	(0.101)	(0.101)	(0.109)
Low Interim Job	0.793	0.797	0.715	0.717	0.727
	(0.068)	(0.069)	(0.084)	(0.084)	(0.090)
High Interim Job	0.839	0.843	0.745	0.747	0.740
	(0.073)	(0.074)	(0.088)	(0.088)	(0.094)
Posting Characteristics:					
High Skill Post	—	1.291	—	1.345	—
		(0.150)		(0.230)	
Associate	—	1.031	—	1.076	—
		(0.206)		(0.322)	
Bachelor	—	1.228	—	1.359	—
		(0.210)		(0.369)	
Educ N/A	—	0.908	—	0.848	—
		(0.117)		(0.164)	
$0<\text{Exp}\leq 1$	—	0.750	—	0.679	—
		(0.100)		(0.138)	
Exp >1	—	0.972	—	1.026	—
		(0.126)		(0.207)	
$\hat{\rho}$			0.629	0.625	
			(0.023)	(0.023)	
Log L	-2841.9	-2824.9	-2465.8	-2459.1	-713.2
Sample Size	8488	8488	8488	8488	1936

Note: Numbers in parentheses are asymptotic standard errors (clustered at the job-posting level in columns 1 and 2). The base group consists of applicants with zero weeks unemployment, aged 27-52, and no interim job. The base group in columns 2 and 4 additionally is characterized as a low-skill posting with high school education recommended/required and zero experience recommended/required or no statement regarding experience. All models in columns 1-4 include city fixed-effects.

callback rates about 75-80 percent as large as prime age applicants.[20] Finally, holders of a lower-quality interim job have a 20 percent lower probability of a callback while holders of a higher-quality interim job have about a 16 percent lower probability of callback.

The logit model whose results are presented in column 2 of the table additionally includes controls for characteristics of the job posting. These include indicators for the skill level of the posted job, the stated educational recommendation/requirement, and the stated experience recommendation/requirement. Unlike the characteristics of the applicants, these are characteristics of the job posting and, thus, have no within-posting variation. As expected, none of the estimated odds ratios for our treatment variables (unemployment, age, interim job) are affected in any meaningful way by inclusion of the job posting characteristics.

It is worth noting that the lower callback rate for the youngest applicants is unaffected by controlling for the stated experience recommendation/requirement (column 2). However, this specification constrains the effect of posted experience to be the same across applicant age groups. We relax this constraint later in order to examine whether the lower callback rate for very young applicants is due to the fact, noted above, that these applicants cannot have much prior experience.

Columns 3 and 4 of Table 8 contain estimates of the random-effects logit versions of the specifications in the first two columns of the table. The substantial estimated within-posting cross-observation correlation in the error ($\hat{\rho} = 0.63$, s.e. $= 0.023$) along with the substantial improvement in the log-likelihood when accounting for the random effect strongly implies that there are important posting-specific factors affecting callback rates. Given that the estimated correlation of the unobserved posting-specific factors does not decline when the observed posting-specific factors are included (column 4 of the table), it is clear that there are important unobserved differences across postings that go beyond skill level and the stated required/recommended education and experience levels.

The random-effects estimates of the parameters of interest are very similar to the standard logit estimates and show the same patterns. This is to be expected given that the treatments were assigned randomly to job postings. One small contrast is that the random-effects estimates suggest that the youngest applicants are at an even larger disadvantage, with 22-23 year old applicants having a callback rate about 40 percent lower than prime age workers. Similarly, the interim job penalty in callback rates is larger in the random-effects model.

[20] As we noted earlier (see Table 4 and supporting discussion), the lower callback rate for the youngest applicants is due to the fact that these applicants have very low callback rates to jobs that require prior experience. We explore this further below.

Finally, column 5 of Table 8 contains the estimates of the Chamberlain fixed-effects logit model. The sample is limited to the 1,936 applications to the 484 job posting where there was variation in the outcome (1 to 3 callbacks). The 6,404 applications to the 1,601 postings that received no callbacks and the 148 applications to the 37 postings that received four callbacks are not included. As expected, given random assignment, the fixed-effect estimates of the odds ratios are virtually identical to the random-effects estimates in column 3.

Overall, the results in Table 8 confirm our main findings. First, there is no difference in callback rates by weeks of unemployment for those unemployed 24 weeks or less. However, applicants who were unemployed for 52 weeks did have a lower callback rate. Second, the youngest (22-23) and oldest applicants (60-61) have lower callback rates than prime-age applicants. Finally, holding an interim job, regardless of quality, at the time of application results in lower callback rates.

Next, we estimate the random-effects logit model of the probability of any-callback allowing the effects of both unemployment duration and employment status at the time of application to vary with age. We also take this opportunity to allow age differences in callback rates to vary with the stated experience recommendations/requirements. This model is a less constrained version of the random-effects logit model in column 4 of Table 8, and its estimates are presented in Table 9. The interpretation of this table requires a bit of explanation. All estimates in the table are from the same logistic regression model. The first column of Table 9 contains the odds ratios for the indicated category relative to a baseline consisting of applicants aged 27-52 with zero weeks unemployment and no interim job who have applied for a low-skill posting with a high school education recommended/required and zero experience recommend/required or no statement regarding experience. The estimates show that, at baseline relative to prime age applicants (27-52), the youngest applicants (age 22-23) have a marginally significantly lower likelihood of callback, and postings that recommend/require a 4-year college degree have marginally significantly higher callback rates.

The estimates in columns 2-4 of Table 9 show the odds ratios separately by age category for weeks of unemployment, interim job status, and the recommended/required experience level relative to applicants at the baseline with zero weeks unemployment, no interim job, and relative to postings with zero experience recommended/required or no statement regarding experience.[21] These estimates allow us to examine whether the patterns we found earlier in the relationships between the likelihood of callback and these treatments varies by age.

[21] Note that the odds ratios in columns 2-4 are derived from underlying coefficient estmates that capture the difference in the marginal effect for the interaction relative to the baseline group in column 1.

Table 9: Estimates of Any-Callback Probability, with Treatments interacted with Age
Random-Effects Logit Model: Odds Ratios

Variable	(1) Baseline	Variable	(2) Age 22-23	(3) Age 27-52	(4) Age 60-61
Age 22-23	0.622	4 Weeks U	1.786	0.951	0.688
	(0.238)		(0.671)	(0.158)	(0.247)
Age 60-61	0.823	12 Weeks U	1.657	1.223	0.570
	(0.297)		(0.678)	(0.201)	(0.228)
High Skill Post	1.356	24 Weeks U	2.738	0.855	0.676
	(0.234)		(1.174)	(0.173)	(0.275)
Associate	1.081	52 Weeks U	0.435	0.696	0.444
	(0.327)		(0.205)	(0.122)	(0.177)
Bachelor	1.354	Low Interim Job	0.592	0.683	0.979
	(0.371)		(0.200)	(0.095)	(0.296)
Educ N/A	0.844	High Interim Job	0.651	0.736	0.821
	(0.164)		(0.209)	(0.103)	(0.264)
		$0 < \text{YR} \leq 1$ Exp Req	0.547	0.732	0.565
			(0.197)	(0.158)	(0.210)
		> 1 YR Exp Req	0.355	1.074	1.757
			(0.139)	(0.230)	(0.587)
$\hat{\rho}$	0.630	Log L = -2433.0			
	(0.023)	$N = 8488$ Apps			

Note: Numbers in parentheses are asymptotic standard errors. The estimates in column 1 are odds ratios relative to a baseline group consisting of applicants aged 27-52 with zero weeks unemployment and no interim job who have applied for a low-skill posting with a high school education recommended/required and zero experience or no statement of experience recommended/required. The estimates in columns 2-4 are odds ratios by age relative to the baseline group with zero weeks unemployment and no interim job. City fixed effects are included in the model.

A likelihood-ratio test, based on a comparison of the estimates of the relevant constrained and unconstrained models, of the hypothesis that the effect of unemployment duration on callback does *not* vary by age is marginally rejected at conventional levels (p-value = 0.048). The estimates reported in Table 2 and Table 8 showed that applicants with 52 weeks unemployment had lower callback rates than applicants with shorter unemployment spells but that there were no differences in callback rates across the shorter spells. The estimates in columns 2-4 of Table 9 allow us to examine further whether this pattern varies by age of applicant. These are the multivariate analog of the simple tabulations in Table 5, and they show a similar pattern. While the distinction is sharpest for the youngest applicants, those applicants in every age group with 52 weeks unemployment have significantly lower callback rates than workers with shorter spells of unemployment. There are no significant distinctions by duration of unemployment among applicants in any age group with less than 52 weeks of unemployment. Hence, while our data does not reject that there are slight variations in the effect of unemployment duration by age, there is no evidence here that variation in the relationship between unemployment duration and callback rates by age can account for the differences in results found in previous studies.

The analogous likelihood-ratio test of the hypothesis that the effect of interim job status on callback does *not* vary by age is not rejected (p-value = 0.81). The estimates reported in Table 6 and Table 8 showed a negative relationship between the likelihood of callback and holding an interim job of either lower or higher quality. The estimates in columns 2-4 of Table 9 allow us to examine further whether this pattern varies by age of applicant. These estimates confirm the negative relationship we reported earlier for the youngest and prime-age applicants. However, there appears to be no relationship between the likelihood of callback and interim job status for the oldest applicants.

The multivariate analyses presented in Table 8 constrained the effect of applicant age to be constant across postings with different experience recommendations/requirements. This is not consistent with our finding (Table 4) that the youngest applicants had high callback rates to jobs with little or no experience recommended/required and low callback rates to jobs with more experience recommended/required. Table 9 contains the estimates of the random-effects logit model without this constraint. Not surprisingly, a likelihood-ratio test of the hypothesis that the effect of the stated experience recommendation/requirement on callbacks does *not* vary by age is easily rejected at conventional levels (p-value = 0.004). The estimates clearly show that odds of a callback for younger workers are sharply lower as the experience recommendation/requirement increases. This is not surprising, and it accounts for at least part of the deficit in callback rates for the youngest applicants that we found in Table 8.

4 Final Remarks

The rise and stubborn persistence of long-term unemployment after the Great Recession has motivated a series of audit studies to investigate the determinants of successful job applications. Those studies have generated a variety of results, some of which indicate that age, duration of unemployment, and recent employment history might interact in important ways to influence job application outcomes. To explore those interactions, we fielded a new audit study varying treatments of unemployment duration, age, and the current employment status of the applicant.

The results indicate duration dependence in callbacks only after very long spells of unemployment. This contrasts our own previous findings (Farber et al. 2016, 2017) of no duration dependence even after very long spells. Our new finding is partially consistent with the finding of Kroft et al. (2013) that callback rates are lower for applicants with long unemployment spells (one year or more). However, in contrast to our results, they also find that the callback rate declines steadily as unemployment spell length increases from 3 to 10 months. While our experimental design does not allow us identify effects between 6 and 12 months, we find no decline in callback rates as unemployment spell lengths increase 3 to 6 months. The results reported here are closest to those in Eriksson and Rooth (2014), who found duration dependence of callback rates only after nine months of unemployment.

A simple model of employer learning suggests that older applicants may be at least partially immunized against any stigmatizing effects of long unemployment spells due to having a more substantial work history relative to younger applicants. This immunizing effect of age has the potential to help explain the differing results among prior audit studies. While we do find significant effects of age on callback rates, there is little evidence of an important interaction between age and unemployment duration in the effect on the likelihood of callback. Thus, there is little reason to think that differences in the age of applicants can explain the different results in previous studies.

Our findings do, however, offer insight into a variety of prior results regarding callback rates for the currently employed. Echoing elements of our prior study, Kroft et al. (2013), and Nunley et al. (2017), we find that average callback rates are significantly higher for the unemployed than for those holding an interim job. Importantly we show here that this negative effect of employment holds regardless of the quality of the interim job, but obtains only when applying to a relatively high-skilled position. These results, which indicate some unappreciated challenges of on-the-job search, merit further investigation.

More generally, the variation we observe in some of the qualitative conclusions of audit studies with very similar designs suggests caution in interpreting their results as being generally applicable to other settings. It appears that seemingly small differences in markets, occupations, résumé designs, and other factors may cause substantial differences in how callback rates vary in important dimensions. In this way, inference from audit studies would seem to require more than the usual attention to replication and assessment of external validity.

References

Bendick, Marc, Jr., Charles W. Jackson, and J. Horacio Romero. "Employment Discrimination Against Older Workers: An Experimental Study of Hiring Practices," *Journal of Aging & Social Policy*, 8 (1997), pp. 25-46.

Bertrand, Marianne and Sendhil Mullainathan. "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," *American Economic Review*, 94 (2004), 991-1013.

Bowlus, Audra J. "Matching workers and jobs: Cyclical fluctuations in match quality." *Journal of Labor Economics* 13 (2)(1995) pp. 335-350

Chamberlain, Gary. "Analysis of Covariance with Qualitative Data," *Review of Economic Studies*, 47(1) (1980), 225-238.

Chan, Sewin, and Ann Huff Stevens. "Job Loss and Employment Patterns of Older Workers." *Journal of Labor Economics* 19, no. 2 (2001): 484-521.

Eriksson, Stefan and Dan-Olof Rooth. "Do employers use unemployment as a sorting criterion when hiring? Evidence from a field experiment," *American Economic Review*, 104 (2014), 1014-1039.

Farber, Henry S. "Alternative and Part-Time Employment Arrangements as a Response to Job Loss," *Journal of Labor Economics* 17:S4 (1999), S142-S169

Farber, Henry S. "Employment, Hours, and Earnings Consequences of Job Loss: U.S. Evidence from the Displaced Workers Survey," *Journal of Labor Economics*, 35(S1) (July 2017), pp. S235-S272.

Farber, Henry S., Dan Silverman, and Till von Wachter. "Determinants of Callbacks to Job Applications: An Audit Study," *American Economic Review*, 106(5) May 2016, pp. 314-318.

Farber, Henry S., Dan Silverman, and Till von Wachter. "Factors Determining Callbacks to Job Applications by the Unemployed: An Audit Study," *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 2017.

Farber, Henry S., Chris M. Herbst, Dan Silverman, and Till von Wachter. "Employment, Unemployment, and Age as Factors in Callbacks to Job Applications," Working Paper No. 619, Industrial Relations Section, Princeton University, May 2018. (`http://arks.princeton.edu/ark:/88435/dsp01mp48sg464`)

Ghayad, Rand. “The Jobless Trap.” mimeo, Northeastern University, 2014.

Jarosch, Gregor and Laura Pilossoph. “Statistical Discrimination and Duration Dependence in the Job Finding Rate,” mimeo, Stanford University, September 2016.

Kahn, Lisa B. “The long-term labor market consequences of graduating from college in a bad economy,” *Labour Economics* 17(2) (2010), pp. 303-316.

Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo. “Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment,” *Quarterly Journal of Economics* 129(2) (August 2013) pp. 1123-1167.

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

Neumark, David, Ian Burn, and Patrick Button. ”Is it Harder for Older Workers to Find Jobs: New and Improved Evidence from a Field Study, Working Paper 21669, National Bureau of Economic Research, October 2015. *Journal of Political Economy*, in press.

Neumark, David, Ian Burn, and Patrick Button. “Experimental Age Discrimination Evidence and the Heckman Critique,” *American Economic Review* 106(5) (May 2016), 303-308.

Nüß, Patrick. “Duration Dependence as an Unemployment Stigma: Evidence from a Field Experiment in Germany,” undated manuscript, Macroeconomic Policy Institute (IMK), Christian Albrechts-University, Kiel Germany (2018).

Nunley, John M., Adam Pugh, Nicholas Romero, and Richard Alan Seals, Jr., “Unemployment, Underemployment, and Employment Opportunities: Results from a Correspondence Audit Study of the Labor Market for College Graduates.” *Industrial and Labor Relations Review*, 70(3) (2017), pp. 642-669.

Oreopoulos, Phil, Andrew Heitz, and Till von Wachter. “The Short- and Long-Term Career Effects of Graduating in a Recession: Hysteresis and Heterogeneity in the Market for College Graduates,” *American Economic Journal: Applied Economics*, 4(1) (January 2012), pp. 1-29.

Schwandt, Hannes and Till von Wachter. “Unlucky Cohorts: The Long-Term Effects of Entering the Labor Market in a Recession on Earnings and Health,” Working Paper, April 2017.

Song, Jae and Till von Wachter. "Long-Term Nonemployment and Job Displacement." In: Evaluating Labor Market Dynamics, A Symposium Sponsored by the Federal Reserve Board of Kansas City, Jackson Hole, Wyo., August 2014.

Stevens, Ann Huff. "Persistent Effects of Job Displacement: The Importance of Multiple Job Losses," *Journal of Labor Economics* 15 (1, Part 1) (1997), pp. 165-188.

Von Wachter, Till, Jae Song, and Joyce Manchester. "Long-Term Earnings Losses Due to Mass Layoffs during the 1982 Recession: An Analysis Using US Administrative Data from 1974 to 2004." Working Paper, 2009.