Job History, Non-Cognitive Skills, and Employability

Alain Cohn[*] Michel André Maréchal[†] Frédéric Schneider[‡]
Roberto A. Weber[†]

May 2, 2017

Abstract

We use laboratory and field experiments and a survey experiment with Human Resources professionals to investigate whether frequent job changes provide a signal of poor non-cognitive skills and whether firms therefore use employment histories to discriminate against employees who switch jobs frequently. Across all three studies, we find consistent evidence that workers who change jobs less frequently have, or are perceived to have, better non-cognitive skills, and are thus more employable. These findings highlight the importance of job history as a signal of non-cognitive skills in labor markets, and point to a cost of frequent job changes for workers.

[*] University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637
[†] Department of Economics, University of Zürich, Blümlisalpstrasse 10, 8006 Zürich.
[‡] School of Management, Yale University, 165 Whitney Ave, New Haven, CT 06511

1. Introduction

While traditional accounts of human capital mainly emphasize the importance of cognitive or physical skills (e.g., Becker 1964), more recent research highlights the relevance of social and behavioral skills for the productivity of workers (Bowles, Gintis, and Osborne 2001; Heckman, Stixrud, and Urzua 2006) and argues that the labor market increasingly rewards such skills (Deming 2015). These skills—often grouped broadly under the label "non-cognitive skills"—involve, for example, a worker's reliability, trustworthiness, self-control, loyalty, and ability to work well with others (e.g., Heckman and Rubinstein 2001; Dohmen et al. 2009; Lindqvist and Vestman 2011).

An important open question is how information regarding non-cognitive skills is conveyed in labor markets. One piece of observable and typically verifiable information in most job applications is employment history—what positions an applicant has previously held, at which firms, and for how long.[1] Could this information provide a signal of a prospective employee's non-cognitive skills?[2]

In this paper, we propose that employers will often view frequent job changes as reflective of a worker's poor non-cognitive skills. In turn, employers will, *ceteris paribus*, find workers who change jobs frequently less desirable, particularly in contexts where non-cognitive skills are important.[3] Our conjecture thus ascribes a potentially powerful role to employment history—a widely available type of information in labor markets—in providing a signal of desirable labor market qualities.

[1]Referrals by existing employees (Rees 1966; Pallais and Sands 2016; Burks et al. 2015) and social networks (Granovetter 1974; Gërxhani, Brandts, and Schram 2013) may also be channels through which employers can obtain information about prospective workers' abilities, including their non-cognitive skills. Once a candidate is hired, the employer can gather further information by directly observing workplace behavior (Bartling, Fehr, and Schmidt 2012).

[2]Publicly observable histories also form the basis of an extensive literature on screening and signaling in labor markets (Spence 1973; Arrow 1973; Stiglitz 1975; Waldman 1984). This literature has typically focused on educational attainment as a signal of human capital, i.e., cognitive abilities that may facilitate learning and performing work-related tasks (Tyler, Murnane, and Willett 2000; Bedard 2001).

[3]The popular business press often recognizes that frequent job changes can be associated with perceptions of "disloyalty, fickleness, and unreliability" (Trikha 2012; Suster 2010). Others have noted that workers are heterogeneous in their propensity to remain with specific employers, and that this corresponds to stable individual characteristics (Ghiselli 1974; Blumen, Kogan, and McCarthy 1955).

Why should applicants' job histories convey information about their non-cognitive skills? Most employment relationships require a worker to follow directions from supervisors, get along well with others and exhibit reliability and self-control. Hence, employees who do these things are often more valuable to employers and less likely to quit jobs due to personal conflicts. On the other hand, workers with poor non-cognitive skills are more likely to experience workplace conflicts and either leave or be terminated.

To provide an initial test of our hypothesis, we analyzed data from the National Longitudinal Survey of Youth 1997 (NLSY97) for relationships between individual characteristics related to non-cognitive skills and the number of jobs individuals have held in their careers.[4] Controlling for various individual-specific covariates, we find that people who are more likely to break rules, have been arrested by the police, and drink at work switch jobs significantly more often. Moreover, the personality trait conscientiousness (i.e., the tendency to act dutifully and self-disciplined) is negatively associated with the number of job changes.[5]

The NLSY97 data further suggest that changing jobs more frequently is associated with an increased likelihood of being unemployed. Figure 1 shows coefficient estimates from regressions of current unemployment status on the number of previous jobs for different sub-populations and a number of controls. For example, an individual with eight instead of three different jobs since the age of 20 (25th vs. 75th percentile) has more than double the probability of being currently unemployed (2.2 vs. 4.6%).[6] Hence, the perils of frequent job switching for employment outcomes seem to be a broad phenomenon that applies across many demographic groups, occupations, and industries.

While this correlational analysis is suggestive of our hypothesized relationship, potentially unobservable variables do not allow for a causal interpretation of the association between non-cognitive skills, job changes, and employability. Therefore, in what fol-

[4] The NLSY97 is a large, nationally representative panel of young Americans, covering a wide range of jobs and industries in the US labor market.

[5] Details of this analysis are provided in the Online Appendix.

[6] We also find a negative relationship between the number of previous jobs and current income from wages and salaries. See Online Appendix for further details of this analysis.

Figure 1: Relationship between number of previous jobs and current unemployment status

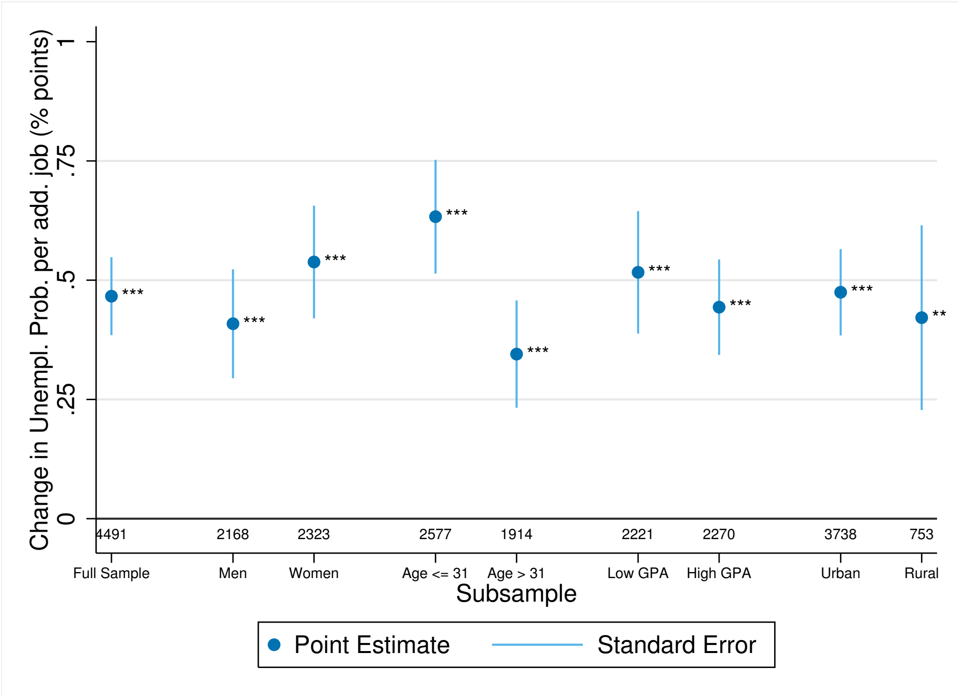

Coefficients from OLS regressions of current unemployment status (as of the last interview wave in 2013/14) on the number of past jobs (since the age of 20), controlling for past unemployment, highest academic degree, high-school GPA, age, gender, ethnicity, geographical region, urban/rural area, and month in which the interview was conducted. Each dot represents a separate regression, corresponding to the sub-population indicated on the horizontal axis. The numbers at the bottom indicate sample size. The stars next to the dots indicate significance with *** $p < 0.01$ and ** $p < 0.05$.

lows, we provide causal evidence of our hypothesized relationship using complementary laboratory, field, and survey experiments. We do not claim to provide a comprehensive interpretation of tenure-employment relationships, but rather propose one particular mechanism through which employment history can impact subsequent labor market outcomes.[7] Specifically, we test two hypotheses regarding the role of employment history as

[7]There may be several reasons for either a positive or negative relationship between job mobility and employability. For example, workers who switch employers more often may accumulate a larger stock of general human capital—that is, skills and knowledge that are useful across jobs, firms, and industries (Mincer 1958; Becker 1962). Moreover, the reasons behind job changes are undoubtedly important for subsequent labor market outcomes (Jovanovic 1979; Topel and Ward 1992), and job mobility may have differential impacts at different points in a worker's career (Bartel 1980; Mincer and Jovanovic 1982; Farber 1999).

a signal of non-cognitive skills. Our first hypothesis is that frequent job changes provide a signal of poor non-cognitive skills. In other words, workers who change jobs frequently will tend to be lower on desirable attributes like reliability, self-control, and ability to work well with others. Our second hypothesis is that, as a consequence of the relationship between job changes and non-cognitive skills, employers will prefer workers with fewer employment changes. We find support for both hypotheses, in the lab as well as in the field. While our field experiment provides the most compelling evidence of the economic significance of our findings, the lab and survey experiments deliver the clearest insights into the precise mechanism driving the relationship between job changes and employment outcomes.

The lab experiment allows us to isolate non-cognitive skills from other possible channels through which a relationship between past and future employment might occur. For example, it eliminates heterogeneity in workers' task-related skills, effort cost, experience, and recruitment costs as confounding factors. Firms value workers to the extent that they reliably comply with requests for high effort. Since workers with a greater tendency to provide voluntary effort are more valuable to firms, firms can benefit from using informative signals regarding non-cognitive skills, and should thus favor contracting with more reliable and cooperative workers.

Our laboratory results show that, first, workers who switch employers less frequently tend to be those who are more reliable and cooperative. Second, following an exogenous unemployment shock that requires all workers to find a new employer, job histories facilitate the signaling of these positive traits—workers with fewer job changes receive more job offers and earn a greater income than those who have switched jobs more often. Finally, by turning off the ability of firms to observe work histories we show that this information is crucial in firms' attempts to identify reliable workers. Hence, the results demonstrate that frequent job changes can serve as a signal of poor non-cognitive skills and influence employability.

We then test whether the phenomenon we identify in the laboratory is also relevant for real labor markets. We report a field experiment that studies whether frequent job changes make prospective employees less desirable to firms. Specifically, we sent resumes to several open positions for administrative and clerical jobs and then measured whether the applicants receive invitations to job interviews. The resumes varied, by random assignment, the candidates' job history.[8] For every open position, we sent two applications: one with four shorter periods of tenure at different firms, and one with a single period of tenure with the same total length. We counterbalanced other aspects of the resumes. In two waves of data collection, we observe significantly higher callback rates for the applicant with fewer job changes. That is, workers who change jobs more frequently are less desirable in the field study, just as they are in our laboratory study. This result is robust to variations in the economic environment, industry, and job characteristics. Moreover, the size of the effect we observe in the field experiment is substantial—the difference in callback rates for applicants with one versus four prior employers is similar in magnitude to the difference between applicants with one versus eight months of unemployment (Kroft, Lange, and Notowidigdo 2013) and white versus black candidates (Bertrand and Mullainathan 2004).

Finally, we complement the field experiment with a survey study to obtain more information on what inferences prospective employers make when receiving the resumes in the field study. Specifically, we approached professionals with experience in human resources (HR) management to survey their impressions of the resumes used in the field study. We find that HR professionals attribute poorer non-cognitive skills to a resume with more frequent job changes—specifically, worse evaluations for the characteristics "reliable," "team oriented," and "perseverant." Moreover, these perceived differences in non-cognitive skills, as opposed to task-related skills and experience, largely explain the HR professionals' greater stated willingness to invite applicants with fewer job changes

[8]Many studies have used this method to test for other aspects of labor market discrimination, such as race, gender or unemployment duration (Riach and Rich 2002; Bertrand and Mullainathan 2004; Carlsson and Rooth 2007; Oberholzer-Gee 2008; Kroft, Lange, and Notowidigdo 2013; Eriksson and Rooth 2014; Deming et al. 2016; Bartoš et al. 2016).

for an interview. Thus, the survey experiment provides evidence confirming that the resumes in the field study create different perceptions of candidates' non-cognitive skills, and that these perceptions are important drivers of callbacks.

Our evidence that employers discriminate against frequent job changes may have implications that go beyond the value of work history as a signal of non-cognitive skills. For instance, workers may be unwilling to undertake job changes out of fear of the negative impact on future prospective employers' perception of their attributes. This inertia or friction in job mobility may create inefficient matching between employees and employers. Labor market frictions are a key feature of modern search theory in macroeconomics because they provide possible explanations for the existence of unemployment and wage inequality (e.g., Petrongolo and Pissarides 2001; Rogerson, Shimer, and Wright 2005). Previous work has focused primarily on structural factors for why workers may refuse job offers and wait for more attractive ones, such as how quickly they can sell their houses (Head and Lloyd-Ellis 2012). Our paper adds to this literature by proposing a mechanism for labor market frictions that arises endogenously, through employers' preference for workers with better non-cognitive skills and the limited information available to employers on this characteristic.

Our study is also related to a large empirical literature studying the relationship between job mobility and wage growth. While some of these studies find that mobility and wage growth are positively related (Topel and Ward 1992; Becker and Hills 1983; Bartel 1980), others find a negative relationship (Light and McGarry 1998; Mincer and Jovanovic 1982; Borjas 1981). Our paper contributes to this literature by examining the impact of variations in job mobility that have limited relationships with task-related skills or ability. We provide one possible mechanism—the signaling of non-cognitive skills—through which prior mobility may affect future outcomes, though our focus is on employability rather than wages.[9]

[9] A separate strand of literature explores how job tenure with a particular firm relates to wage profiles (Dustmann and Meghir 2005; Altonji, Smith, and Vidangos 2013; Bagger et al. 2014). This is distinct from our study because we focus on job tenure solely for its signaling purposes when changing jobs

The remainder of this paper is structured as follows. The next section presents the design and results of the laboratory experiment. Section 3 reports the field experiment and the related survey study of HR professionals. Finally, in Section 4 we provide a broad interpretation of the combined results and conclude.

2. Laboratory Experiment

Our laboratory experiment uses a setting in which a worker's value to firms is determined by her non-cognitive skills (i.e., reliability and cooperativeness) and where task-related skills and experience are held constant. Specifically, we employ a widely used experimental labor market paradigm in which incomplete contracts create incentives for inefficient shirking by workers. Workers are valuable to firms if they voluntarily provide high effort. To study whether employers use employment histories as a signal of relevant non-cognitive skills, we exogenously manipulate whether they have access to workers' job histories.

2.1. Experimental Design

Each experimental labor market consists of 17 participants, of which seven are randomly assigned the role of a firm; the remaining ten participants are assigned the role of a worker. Each participant is identifiable through a permanent ID number. The experiment lasts 30 periods. In any given period, each firm can hire at most one worker, and each worker can work for at most one firm. Because labor supply exceeds labor demand, some workers are unemployed in a given period.

Every period has two stages: a hiring stage and a work stage. In the hiring stage, firms can post two types of wage offers: *i)* public wage offers, which any worker can accept, and *ii)* private wage offers, which target specific workers. Each offer contains a binding wage, $w \in \{1, 2, \ldots, 100\}$, and a desired effort level, $\hat{e} \in \{1, 2, \ldots, 10\}$. A worker can accept any public offer or any private offer directed at her. A private offer is thus a clear

between firms.

indication that a firm prefers a particular worker. At the end of the hiring stage, up to seven firms and workers are matched in an employment relationship for that period.

The second stage is the work stage, in which employed workers select an effort level to provide. Workers choose "effort" by selecting a number, $e \in \{1, 2, \ldots, 10\}$, which implies monetary costs according to an effort cost schedule, $c(e)$ (see Table 1). The worker's payoff from employment is equal to the wage minus the effort costs: $\pi_{\text{worker}} = w - c(e)$. [10] Because workers simply choose a number, we eliminate any task-specific skill and experience differences between workers. Nevertheless, the employer cares about the worker's effort choice: the firm earns 10 ECU per unit of worker effort e, but also has to pay the wage w: $\pi_{\text{firm}} = 10e - w$.[11]

Table 1: Workers' effort cost

e	1	2	3	4	5	6	7	8	9	10
$c(e)$	0	1	2	4	6	8	10	12	15	18

While aggregate payoffs are maximized if workers provide maximum effort, the worker's monetary incentive—in the absence of repeated-game incentives—is to shirk and provide minimal effort. Effort in this context is thus a one-dimensional proxy for the voluntary provision of costly, but productive effort at work—i.e., a measure of an employee's cooperativeness, reliability, and diligence.

To study the role of work histories as a signal of these non-cognitive skills, we vary experimentally whether workers' employment histories are available to firms. In the "History" condition, each firm sees a table on the computer screen listing all ten workers in the labor market, sorted by their ID number. The table indicates, for all previous periods, either the ID of the firm that hired the worker in that period or whether that worker was unemployed.[12] This provides prospective employers with a simple version of

[10] This is a standard approach in experimental labor markets. Brüggen and Strobel (2007) show that such numerical effort choices produce similar behavior as real effort decisions.

[11] Unemployed workers receive $\pi_{\text{unempl}} = 5$; firms without a worker receive a payoff of zero in that period. All payoffs are denoted in "Experimental Currency Units" (ECU), converted into Swiss Francs at a rate of 20 ECU = 1 CHF ($\approx$ 1.06 USD) at the end of the session.

[12] If the worker was unemployed in a particular period, the cell is filled with a dash. Importantly, the

the employment histories typically contained in job applications, including job changes and spells of unemployment. By contrast, the work history table is absent in the "No History" condition.[13]

Our two hypotheses are that work histories provide a signal of non-cognitive skills and that firms use this signal when deciding which workers to employ. We expect that workers who provide higher voluntary effort will tend to be those who remain longer with the same employer. In addition, when employment histories are available, we expect that firms will use this information to make targeted offers to workers with fewer prior job changes. By contrast, if the number of previous employers is not diagnostic of future effort choices or firms do not appreciate the signaling value of previous changes, then we should not observe that the availability of employment histories affects labor market outcomes.

To investigate whether firms use employment histories to screen for high-effort workers, we implement an exogenous layoff shock that forces all firms to seek a new worker. From period 17 onwards, we remove both the option for firms to make private offers to the worker they had hired in period 16, and the option for workers to see or accept public offers from the firm they had worked for in period 16. This change is permanent, meaning that no market participant is allowed to interact with their partner from period 16 in any of the remaining periods. This shock introduces an exogenous layoff, which requires all workers to search for new employment opportunities.[14] This design feature allows us to investigate which workers firms find desirable in a context where all workers are simultaneously—and for exogenous reasons—searching for new employment.

---

table does not show workers' effort or wages, only the firm for which they worked (see Online Appendix). Workers could see a similar table that listed the firms by their ID number and listed which workers worked for a particular firm across periods.

[13]Note, however, that in both conditions firms have private information about the workers they had previously employed.

[14]Participants did not know that this shock would happen in period 17. They were informed that this restriction would come into effect at some point "between period 10 and period 20." We did this to rule out that firms would strategically separate from long-term employees in period 16 just to be able to re-hire them in period 17.

Procedures

We conducted the study between December 2012 and May 2013, and additional sessions in June 2015, at the Laboratory for Behavioral and Experimental Economics at the University of Zurich. Each session was randomly assigned to one of the two treatment conditions. All interactions between participants took place via the z-Tree computer interface (Fischbacher 2007). Computer stations were separated by partition walls, ensuring anonymity of the participants. The participants received detailed written instructions and then had to complete a comprehension check to make sure that they understood the rules of the experiment (See Online Appendix). We read instructions aloud to establish common knowledge. [15] We recruited a total of 561 participants using the software h-root (Bock, Baetge, and Nicklisch 2014). Of these, 272 (16 markets) were in the No History and 289 (17 markets) in the History condition.

2.2. Results

Are work histories an informative signal of voluntary effort provision?

Figure 2 shows the relationship between workers' effort and their employment history during the first 16 periods of the experiment. In the History condition, workers who had a single employer throughout periods 1 to 16 provided an average effort of 9.2, close to the maximum of 10. Average effort decreases with the number of pre-shock employers to a level of 4.7 for workers with six different pre-shock employers ($p = 0.040$; Mann-Whitney-U test, henceforth denoted as MWU).[16] Similarly, workers in the No History condition with one employer also exerted higher effort on average than those who changed jobs more frequently (9.2 for one employer vs. 5.9 for six employers; $p < 0.001$, MWU). Hence, regardless of whether work histories are available, workers who act more cooperatively

[15]Sessions lasted slightly under two hours, and participants earned an average of 51 Swiss Francs (about 54 US dollars).

[16]Since observations are not independent within markets we use a cluster-robust version of the MWU test (see Datta and Satten 2005).

and reliably are also those with fewer changes.

Figure 2: Voluntary Effort and Number of Employers

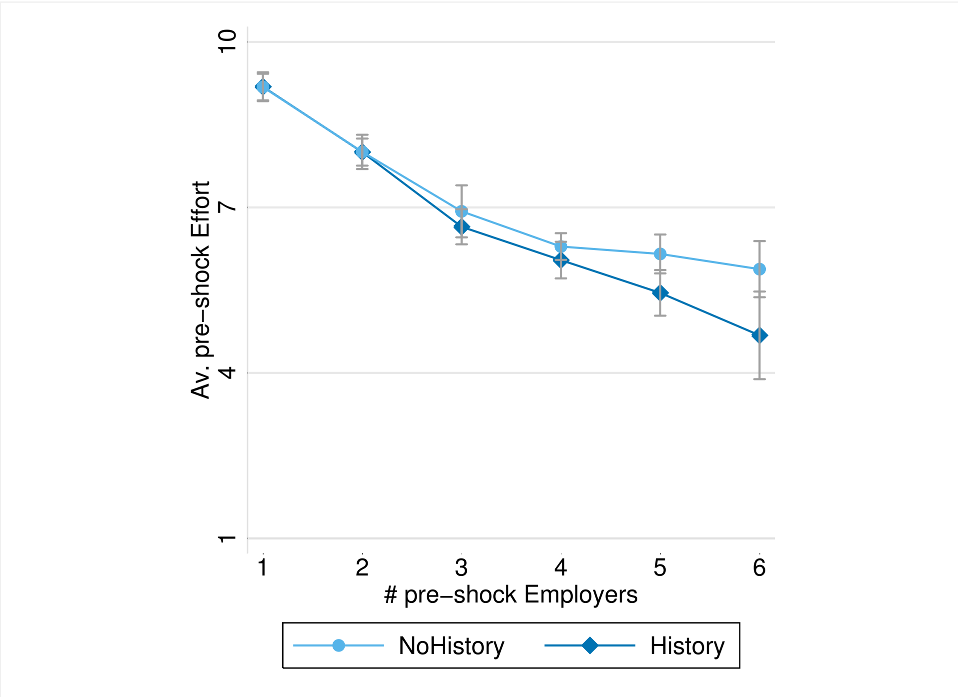

Average effort a worker exerted in periods 1 to 16 in relation to the number of different employers that the worker had during that phase. Unit of observation: worker. Error bars calculated using 1000 bootstrap pseudo-samples, accounting for clustered standard errors at the labor market level. There is a negative relationship between effort exerted in periods 1 to 16 and the number of employers workers had during that phase.

We confirm this non-parametric result with a regression analysis (see Appendix for details). We find that an increase of average pre-shock effort by one unit is associated with a decrease in the number of pre-shock employers by 0.36 in the History condition and 0.34 in the No History condition (both $p < 0.001$, t-tests). Further, we cannot reject the null hypothesis that the coefficient is the same in both conditions (0.787, t-test).

Result 1 (Employment history and effort)

Frequent job changes are indicative of lower effort provision. This relationship holds for workers in the History and the No History condition.

Do firms prefer workers with stable employment?

As shown above, employment histories provide valuable information about workers' reliability in providing voluntary effort. Do firms take this information into account when making job offers?

Figure 3 indicates that firms indeed use workers' employment histories to screen for high-effort workers. In period 17, workers with one pre-shock employer receive 84% more private offers in the History compared to the No History condition ($p = 0.007$, MWU). And while the number of offers drops sharply in the History condition with the number of pre-shock employers, we observe no such trend in the No History condition.

The regression analysis in Table 2 estimates the relationship between the frequency of job changes and employability while controlling for prior unemployment spells. Specifically, we estimate a regression model using the following equation:

$$y_i = \alpha + \beta_1(N_i - 1) + \beta_2 U_i + \varepsilon_{im}. \tag{1}$$

We regress the number of private offers received by an employee in period 17, y_i, on the number of pre-shock employers minus 1, i.e., $N_i - 1$. Thus, the constant reflects the number of private offers obtained by a worker with one pre-shock employer.[17] We additionally control for the number of periods unemployed before the shock, U_i, which is also observable for prospective employers in the History condition. Column 1 shows that, controlling for unemployment spells, each additional employer before the shock significantly reduces the number of private offers in period 17 by 0.219 in the History condition ($p = 0.005$, t-test). By contrast, column 2 shows that in the No History condition—where information about job changes is private information—the coefficient of the number of additional employers is close to zero and statistically insignificant ($p = 0.936$, t-test). In column 3, we pool the observations from both treatments and include a dummy for the History treatment, as well as the corresponding interaction terms.

[17] No worker was unemployed in all 16 pre-shock periods.

Figure 3: Private Offers in Period 17

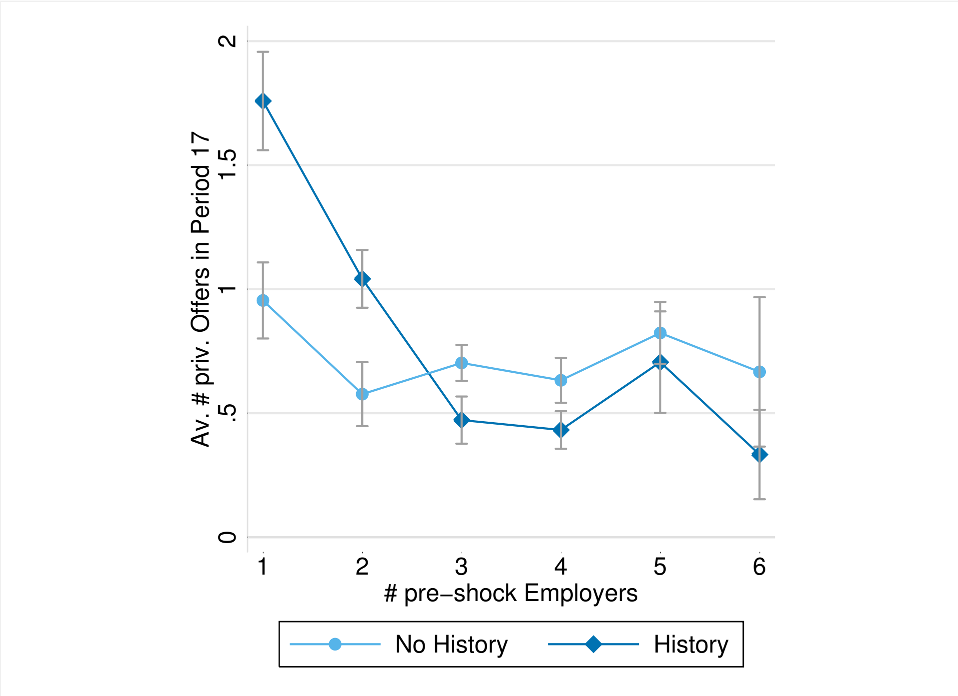

Number of private employment offers that a worker receives from firms at the beginning of period 17 (directly after the shock) as a function of the number of different employers that the worker had before the shock (periods 1 to 16). The unit of observation is a worker. Error bars calculated using 1000 bootstrap pseudo-samples, accounting for clustered standard errors at the labor market level. In the No History condition, where firms rely solely on own information from their previous employment relations with workers, there is almost no effect of number of previous employers. In the History condition, where firms can observe all workers' employment histories before the shock, there is a pronounced negative effect of the number of previous employers.

The results confirm that the coefficient for the number of additional employers differs significantly between the History and the No History condition ($p = 0.019$, t-test).[18]

Our experimental design allows us to follow workers for the remaining 14 periods after the unemployment shock (i.e., periods 17 to 30). The "lifetime" loss in earnings for

[18]We also find that every additional period of unemployment reduces the number of private offers by 0.074 in the History condition ($p = 0.009$, t-test, see column 1) and 0.050 in the No History condition ($p = 0.023$, t-test, see column 2). One interpretation for the negative relationship between unemployment and employability in the No History condition is that firms do not make private offers in period 17 to workers they had "fired" before the shock. Although firms cannot make private offers to workers they had employed in Period 16, they can re-hire workers they had employed in earlier periods.

Table 2: Regression analysis of private offers in Period 17

Condition	(1) History	(2) No History	(3) Pooled
# Employers	-0.219***	-0.005	-0.005
	(0.068)	(0.056)	(0.055)
# Periods Unemployed	-0.074***	-0.050**	-0.050**
	(0.025)	(0.020)	(0.020)
History			0.673***
			(0.239)
History × # Employers			-0.215**
			(0.087)
History × # Periods Unemployed			-0.024
			(0.031)
Constant	1.639***	0.965***	0.965***
	(0.193)	(0.147)	(0.145)
R^2	0.254	0.047	0.185
N	170	160	330
Clusters	17	16	33

OLS regressions, standard errors in parentheses, adjusted for clustering at the labor market level, using White sandwich estimators. Unit of observation: worker.
Significance levels: * p<0.1, ** p<0.05, *** p<0.01.
Dependent variable: number of private offers to worker after the shock (period 17).
Independent variables: Constant: the baseline is a worker in the (No) History condition who was continuously employed by the same firm for all 16 periods before the shock. "History:" dummy for History treatment condition; "# Employers:" number of additional pre-shock employers; "# Periods Unempl.:" number of pre-shock periods the worker was unemployed.

workers with unsteady pre-shock job histories in the History condition is sizable: workers with five or six pre-shock employers earned, on average, 261 ECU. By contrast, workers with one or two pre-shock employers made, on average, 428 ECU after the layoff shock, or about 64% more ($p < 0.001$, MWU). In the No History condition, the difference in earnings between these two groups of employees is much smaller (312 ECU vs. 360 ECU, $p < 0.001$, MWU).

Result 2 (Employment history and job outcomes)
When employment histories are available, workers with fewer previous employers receive more private job offers and achieve higher profits. These relationships are much weaker when employment histories are not available.

We have seen that, when firms can screen for workers with stable job histories, workers with frequent employment changes incur significant losses; but this screening also has broader labor market implications. For instance, as Figure 4 shows, the availability of employment histories influences the duration of employment and unemployment spells. On average, employment relations last longer when employment histories are available (2.4 vs. 3.2 periods, $p = 0.002$, MWU, see also column 4 in Table 5), and workers also remain unemployed longer (1.9 vs. 2.3 periods, $p = 0.001$, MWU). Further analysis reveals that this is driven by both the demand and supply side of the labor market.[19]

Result 3 (Labor market frictions)
Employment relationships and unemployment spells last longer, and the number of employer changes decreases when job histories are publicly observable.

3. Field Experiment

Our laboratory experiment suggests that frequent job changes provide a negative signal of prospective employees' non-cognitive skills. The lab study also shows that employers use this information, if available, to determine which workers to seek out for employment. We now examine whether these results hold in an actual labor market with real firms. To this end, we conducted a field experiment using the correspondence method, a well-established method that has been used to study racial and gender discrimination in hiring (e.g., Riach and Rich 2002; Bertrand and Mullainathan 2004; Carlsson and Rooth 2007; Pager, Western, and Bonikowski 2009; Oreopoulos 2011). Specifically, we applied to a large number of job openings using fictitious applicants and then measured whether the prospective employers called back the applicants for a job interview. We exogenously

[19] Firms are more likely to make job offers to their current workers in the History than in the No History condition ($p < 0.001$, t-test), and workers are more likely to accept offers from their current employers ($p = 0.028$, t-test). The p-values come from OLS regressions of the respective variables on a treatment dummy with cluster-robust standard errors at the labor market level. We also find that observable job histories make private reputation portable: firms in the History condition are more willing to make offers to workers they had not employed previously. In period 17, only 38% of employed workers had previously worked for the same firm in the History condition, while in the No History condition firms hired workers that they had employed before in 68% of cases. OLS regressions confirm that this difference is statistically significant ($p = 0.006$, t-test).

Figure 4: Hysteresis with Job Market History

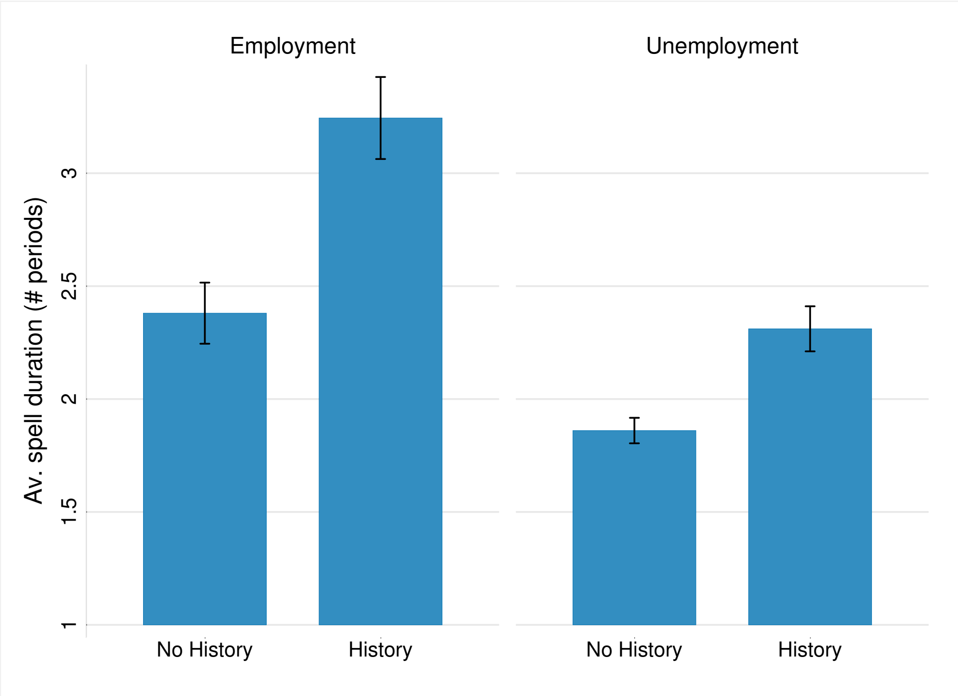

Average duration of employment and unemployment spells, in number of periods, across conditions. Unit of observation: labor market. Left Panel: the average employment spell is about 36% longer in treatment History relative to No History. Right Panel: the availability of employment histories increases unemployment spells by 24%.

varied the number of previous jobs that our candidates held to examine whether this would influence the probability of a callback. We also ran a complementary survey experiment with HR professionals to assess the role of non-cognitive skills in explaining why frequent job changes can harm an applicant's employment prospects. We conducted the field experiment in two waves. The first wave took place between May and June 2012, and the second wave one year later, from April to June 2013. We sent out a total of 1,680 applications for positions in the German-speaking part of Switzerland (680 in the first wave and 1,000 in the second wave). We employed the same strategy as most correspondence studies and focused on commercial jobs (i.e., administrative and clerical work) to have enough job vacancies with similar skill requirements. Our sample of postings covers a broad spectrum of commercial jobs, including jobs in customer services, sales support, and management assistance.[20] Over two waves of data collection, we surveyed all administrative and clerical job ads posted on four large job search websites. To obtain reasonably high callback rates, we focused on job postings that were no older than ten days and that offered a position in the broader area of Zurich or adjacent cantons (i.e., close to the applicants' home address).

The resumes

For each wave, we created four identities for the fictitious job applicants, two male and two female. We used names from a list of the most common first and family names in Switzerland to avoid ethnic discrimination, and employed photos from students who gave us their permission to use them for the study. To track responses, we gave each identity a unique home address, email address, and cell phone number.[21] We took great care to make the resumes look authentic and appealing. To this end, we consulted Human Resources professionals and used templates from the Swiss Association of Commercial Employees and related websites.

We sent out two applications for every open position. The pairs of applications described virtually identical applicants in all observable characteristics, except for the frequency of job changes. Both candidates were 26 years old and well-qualified, as they had a diploma in commercial studies with high grades. They both had eight years of work experience in exactly the same job functions. To differentiate the two resumes, we described the functions using different terms (e.g., human resources vs. personnel management) and changed the order in which the functions appeared on the resumes. Both applicants were currently employed when we sent out the applications. We further gave both a set of complementary qualities that employers typically desire for commercial workers, such as

[20] Administrative jobs constitute about 11% of Switzerland's total workforce (Swiss Federal Statistical Office 2008).

[21] Incoming calls were automatically redirected to a voice mail box. Email addresses had different providers to minimize suspicion. To collect responses by postal mail we used real postal addresses and tagged the letter boxes with the corresponding names. However, only about 2% of the employers contacted the applicants via postal mail.

relevant computer and language skills. To minimize suspicion, we used a different formatting and layout for the two resumes. We counterbalanced the two formatting schemes across treatments.

Treatments

For each identity, we implemented a version of the resume with continuous employment at a single firm ("One Employer" condition) and a version with comparable experience but four different firms ("Four Employers" condition). The Four Employers resume signals that the applicant had moved rather frequently from one employer to the next. After a degree in commercial education, the candidate made horizontal moves between four firms every twenty to twenty-four months.[22] In contrast, the applicant with the One Employer resume had spent his or her entire post-education career at the same company. Both resumes exhibited a total of eight years of work experience in exactly the same departments (i.e., administration, accounting, human resources, customer service, and purchase).

For each job ad, we sent a Four Employers and a One Employer resume. We randomized which of the two applicants was assigned the Four and One Employer resume, respectively, and then submitted both resumes, in randomized order, a couple of hours apart. Both had the same gender, which was determined at random unless an employer explicitly asked for candidates of a particular gender.

In the first wave, the Four Employers resumes had short gaps between jobs. Although short breaks between jobs are the norm, they could potentially affect callback rates because employers may consider them as unemployment spells that could signal low productivity (Oberholzer-Gee 2008; Kroft, Lange, and Notowidigdo 2013; Eriksson and Rooth 2014). We therefore removed the gaps from the Four Employers' work history in the second wave.

[22] The companies were chosen from a list of employers that offer commercial positions from a vocational counseling website.

Measuring callbacks

We recorded all incoming responses within seven weeks after submitting the applications; however, most employers contacted the applicants within two weeks. Because we are interested in whether the employers exhibit a preference for one candidate, we define a callback as an explicit request for an interview or a message stating that one of the applicants is shortlisted for interview.[23] Two research assistants who were blind to the experimental conditions coded the responses according to these pre-defined rules. To minimize the inconvenience caused to the employers, we declined interview invitations within 24 hours.[24]

3.1. Results of the Field Experiment

In total, we sent 1,680 applications to 840 job vacancies in a broad range of industries (see Table 6 in the Appendix). Most ads were for jobs in private limited liability companies (87.7%), followed by public employers or NGOs (8.8%), and organizations of other legal forms (3.5%, e.g., single proprietors or cooperatives). 75.4% looked for full-time employees (i.e., at least four days a week).[25] Because we used a within-subjects design—i.e., each prospective employer received a One Employer and a Four Employer resume—our sample of firms is, by construction, balanced across treatments.

We obtained callbacks for 17.1% of the applications; 57.9% of the applications were immediately rejected, 14.6% remained unanswered, and 10.4% received requests that more documents would be needed (without receiving an interview invitation or being short-listed).[26] The average response time was 8.3 days. Most responses came in by

[23]Like other correspondence studies (e.g., Bertrand and Mullainathan 2004; Eriksson and Rooth 2014; Kroft, Lange, and Notowidigdo 2013), we do not observe whether an applicant eventually gets the job, but simply whether a prospective employer contacts the candidate for a job interview. It is reasonable to expect that an invitation for an interview reflects an employer's hiring preference and that differences in interview rates translate into differences in hiring rates.

[24]The studies were approved by the Human Subjects Committee of the Faculty of Economics, Business Administration, and Information Technology of the University of Zurich.

[25]The sample includes job openings placed by employment agencies (16.2%); the results do not change if we exclude these observations from the analysis.

[26]Our results are qualitatively the same if we treat requests for additional documents as callbacks (see

email (85%), followed by phone call (13%), and postal mail (2%).

The results from the first wave show that the Four Employers profile led to a substantially lower callback rate (see Panel A in Figure 5). While the Four Employers resume had a callback rate of 16.8%, the rate was 23.2% for the One Employer condition, i.e., 40% higher. The treatment difference is statistically significant ($p = 0.003$, McNemar test).[27]

Figure 5: Job changes and employability

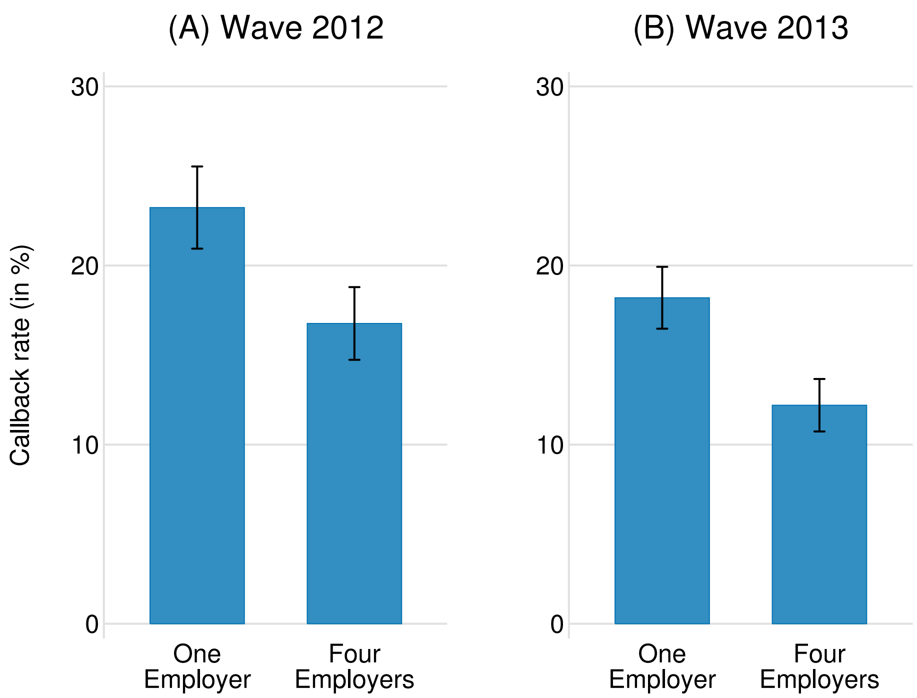

Error bars indicate standard error of the mean. Panel A displays average callback rates by treatment for the 2012 wave. Panel B shows the results for the 2013 wave where the Four Employers resume did not contain any employment gaps between job changes.

For the applicant resumes of the 2013 wave, we removed all gaps between job changes. Panel B in Figure 5 shows that the results replicate when the Four Employers resume had no gaps between jobs. The treatment effect in the 2013 wave is similar to the one

Table 8 in the Appendix).

[27]We use the non-parametric McNemar test for paired observations which compares how often one profile is preferred over the other (see Siegel and Castellan 1988).

in the 2012 wave: the callback rate is almost 50% higher in the One Employer (18.2%) than in the Four Employers treatment (12.2% , $p = 0.001$, McNemar test).[28]

Our treatment effect across both waves is sizable compared to other correspondence studies. For example, the standardized effect size (i.e., Cohen's d) in our study is similar in magnitude to the difference in callback rates that Kroft, Lange, and Notowidigdo (2013) found between applicants with one and eight months of unemployment. It is also similar to the difference between white- and black-sounding names reported in Bertrand and Mullainathan (2004).[29]

A regression analysis corroborates the preceding non-parametric results. Specifically, we estimate the following linear probability model:

$$y_{ij} = \alpha + \beta_1 * N_{ij} + \beta_2 * \mathbf{X}_{ij} + \beta_2 * \mathbf{Z}_j + \epsilon_{ij}. \tag{2}$$

The dependent variable y_{ij}, indicating whether applicant i received a callback for a vacancy j, is regressed on a dummy variable, N_{ij}, for the Four Employers treatment. We control for month of the year, gender of the applicant, gender of the recruiting manager (i.e., contact person), and gender match between the two. We also add dummies for employment agencies and part-time positions, as well as the firms' industry and legal form. Finally, we include the (log) driving distance to the workplace and monthly local labor

[28]Overall, the callback rate in both treatments was lower in 2013 than in 2012 ($p < 0.001$, MWU). One possible reason is that the applicants faced tougher labor market conditions in 2013. Monthly regional labor market statistics (SECO 2013) show that the average number of candidates per job increased from 8.8 to 10.4, and that the local unemployment rate rose from 2.7 to 2.8 between the first and the second wave. An occupation-specific but less direct indicator of labor market conditions is the average response time in our field experiment, which we can use as a proxy for the number of applications the HR recruiters had to assess at that time. In line with the aggregate labor market data we find a significant increase in average response time from 7.7 work days in 2012 to 8.7 work days in 2013 ($p = 0.025$, MWU). As shown in the regression analysis, the effect of multiple previous employers is neither more nor less pronounced when workers compete more for jobs.

[29]Kroft, Lange, and Notowidigdo (2013) found that callback rates dropped from roughly 7% to 4%. With a standard deviation of 0.212, this corresponds to a standardized mean effect (i.e., Cohen's d) of 0.142. Bertrand and Mullainathan (2004) found callback rates of 9.7% and 6.4% for white and black-sounding names, respectively. The standard deviation in callbacks was 0.272, implying a Cohen's d of 0.121. In our study, pooling both waves, we found a 20.2% callback rate for the One Employer condition and a 14.1% callback rate for the Four Employer condition. The standard deviation in callbacks was 0.377, resulting in a Cohen's d of 0.164.

market conditions (i.e., the number of applicants per open position and the employment rate at the cantonal level). The vectors Z_i and X_{ij} represent the control variables measured at the vacancy level and those that vary within vacancies, respectively. We allow for idiosyncratic variation with the error term, ϵ_{ij}. We report OLS estimates and correct the standard errors for clustering at the vacancy level. The results remain qualitatively the same if we use a Probit model instead.

Table 3: Regression analysis of job changes and employability

Dependent variable	Callback = 1					
	(1)	(2)	(3)	(4)	(5)	(6)
Four Employers	-0.062***	-0.060***	-0.061***	-0.061***	-0.060***	-0.074***
	(0.014)	(0.018)	(0.018)	(0.018)	(0.018)	(0.019)
Four Emp. X 2012		-0.005	-0.005	-0.005	-0.005	0.002
		(0.028)	(0.028)	(0.028)	(0.028)	(0.029)
Wave 2012		0.050*	0.042	0.032	0.039	0.043
		(0.029)	(0.029)	(0.030)	(0.031)	(0.031)
Industry experience						0.063**
						(0.031)
Constant	0.202***	0.182***	0.510***	0.268	0.154	0.142
	(0.014)	(0.017)	(0.155)	(0.192)	(0.246)	(0.246)
Additional controls?						
Month				Yes	Yes	Yes
Gender/gend. match			Yes		Yes	Yes
Firm/job character.			Yes		Yes	Yes
Driving distance			Yes	Yes	Yes	Yes
Labor market				Yes	Yes	Yes
Observations	1680	1680	1680	1680	1680	1680
F	20.328	8.271	5.383	5.488	5.110	4.642
Prob> F	0.000	0.000	0.000	0.000	0.000	0.000

OLS regressions, cluster-robust standard errors at the job ad level.
Dependent variable: Dummy indicating a callback.
Independent variables: "Four Employers:" dummy for Four Employers resume; "Wave 2012:" dummy for the first wave of the study (in 2012); "Industry experience:" dummy whether the applicant had previous work experience in the corresponding industry; "Month:" dummies for month when the application was submitted; "Gender/gender match:" dummies for the gender of the applicant and recruiting manager, and the corresponding interaction term between the two; "Firm/job characteristics:" dummies for industry, legal form, employment agency, and part-time jobs; "Driving distance:" log of the distance between home and work address (in meters and assuming traveling by car using Google Maps); "Labor market:" monthly local unemployment rate and the number of applicants per open position (statistics from State Secretariat for Economic Affairs (SECO)).
Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

Table 3 presents regression results for various combinations of control variables. Column 1 shows the results without control variables. We find a significant 6.2% point reduction in the callback rate in the Four Employers relative to the One Employer treatment ($p < 0.001$, t-test). In Column 2 we test whether the treatment effect is significantly different between the two waves by including a dummy for the 2012 wave and its interaction with treatment Four Employers. The interaction effect is small and statistically insignificant ($p = 0.867$, t-test), suggesting that the treatment effect is stable across waves. Columns 3 through 5 indicate that the Four Employers effect is robust in magnitude and significance if we control for a variety of background variables.[30]

Result 4 (Multiple employers and employability in the field)
Applicants with more frequent job changes are significantly less likely to receive callbacks for interviews. The effect is similar regardless of whether resumes include short employment gaps between jobs.

Although we find a consistently negative effect of the Four Employers resume, we want to emphasize that our results do not imply that more frequent job changes will *always* reduce employability. A higher frequency of changes can, in principle, also signal desirable qualities, such as that a worker gained more transferable human capital due to more diverse work experiences (Mincer 1958; Becker 1962). For example, the probability that our applicants had previous work experience in the industry of the prospective employer was naturally higher for the worker with more frequent job changes (50 vs. 32.6%, $p < 0.001$, χ^2-test). To explore the extent to which more diverse industry experience had a compensating positive effect on employability we additionally included a dummy variable "Industry experience" in our regression model. This variable takes a value of one if the applicant has ever worked in the industry of the prospective employer and zero otherwise. Column 6 of Table 3 shows that industry experience significantly increases the probability of a callback by 6.3% (p = 0.040, t-test). Hence, job changes can increase the callback rate through a higher likelihood of relevant work experience. However, this means that we

[30]We additionally examined possible sources of heterogeneity in the treatment effect, including the applicants' gender, job vacancies placed by employment agencies, full-time positions, driving distances to the work place, as well as monthly regional labor market conditions. However, none of the interactions reaches statistical significance at conventional levels. We also find that the point estimates of the treatment effect are negative for all but one of the eight applicant identities. All of these additional tests are available from the authors upon request.

may underestimate the pure effect of job changes on employability. Indeed, we find that the magnitude of the coefficient for the Four Employers treatment is approximately 23% greater in Column 6 than in Column 5, where we do not control for industry experience. In other words, had the One and Four Employers candidates had similar levels of industry experience, employers would have discriminated even more strongly against the Four Employers candidate.

3.2. Survey Experiment

Although our correspondence study shows large effects of frequent job changes on a candidate's chances to be invited for a job interview, , it does not tell us why the Four Employer candidates were less desirable. To get a sense of how hiring managers evaluate the two sets of applicants, we complement the correspondence study with a survey experiment with Human Resources (HR) professionals in which we obtain their perceptions of the candidates from the field experiment.

In the survey experiment, we were interested whether those who make hiring decisions perceive the applicants with four previous employers as having lower non-cognitive skills than those with one previous employer. To answer this question, we went to a job fair for graduating students in Zurich in April 2014 to recruit HR professionals. At this event, mostly large companies from diverse industries (e.g., engineering, electronics, telecommunication, and consulting) presented themselves to job seekers.[31] Each company had its own booth, at which company representatives, including recruiters, were available for questions, such as what kind of employees the firm seeks or how the application process works. We approached each booth and asked whether the most experienced HR representative would be available for a short survey. A total of 83 HR professionals completed the questionnaire.

Since we expected a smaller sample size than in our field experiment, we selected two male candidates from our pool of applicants used in the field experiment. Each survey

[31] See Table 9 in the Appendix for descriptive statistics of our survey sample.

participant was shown a Four Employers and a One Employer resume, side by side. We randomized which of the two candidates would be the one with the greater number of job changes and counterbalanced the order (i.e., left or right) in which the candidates were presented.

In the survey, participants rated both candidates on ten characteristics using 7-point Likert scales, ranging from 1 "does not apply at all" to 7 "applies fully."[32] The characteristics can be broadly divided into task-related skills and experience (captured by the items "skilled," "experienced in commerce," and "multi-talented"), and non-cognitive skills ("able to work in teams," "willing to adapt," "perseverant," "honest," "reliable," "self-directed," and "goal-oriented"). We further asked participants how likely they would be to call back a candidate for an interview had the applicant applied at their firm, on a scale from 1 "very unlikely" to 7 "very likely."

The survey responses allow us to examine which qualities HR professionals associate with the different resumes from the field experiment, and which of these qualities are likely responsible for the difference in callback rates we find in the field experiment.

3.3. Results of the Survey Experiment

To distinguish between task-related skills and experience (we use the term "task-related skills" for both, task-related skills and experience, henceforth) from non-cognitive skills, we created an index for each of the two dimensions by averaging the ratings for a respondent's perception of the individual qualities within a particular dimension. Table 11 in the Appendix presents the results for each individual characteristic. Figure 6 reveals that, relative to the candidates with fewer previous employers, the Four Employers candidates score 0.40 points lower on non-cognitive skills ($p < 0.001$, Wilcoxon signed rank test, henceforth denoted as WSR). By contrast, the difference in task-related skills between the two candidate profiles is small (0.03 points) and statistically insignificant ($p = 0.651$, WSR).

[32]The Online Appendix provides a copy of the survey, translated from German to English.

Figure 6: HR professionals' perceptions of One and Four Employers candidates

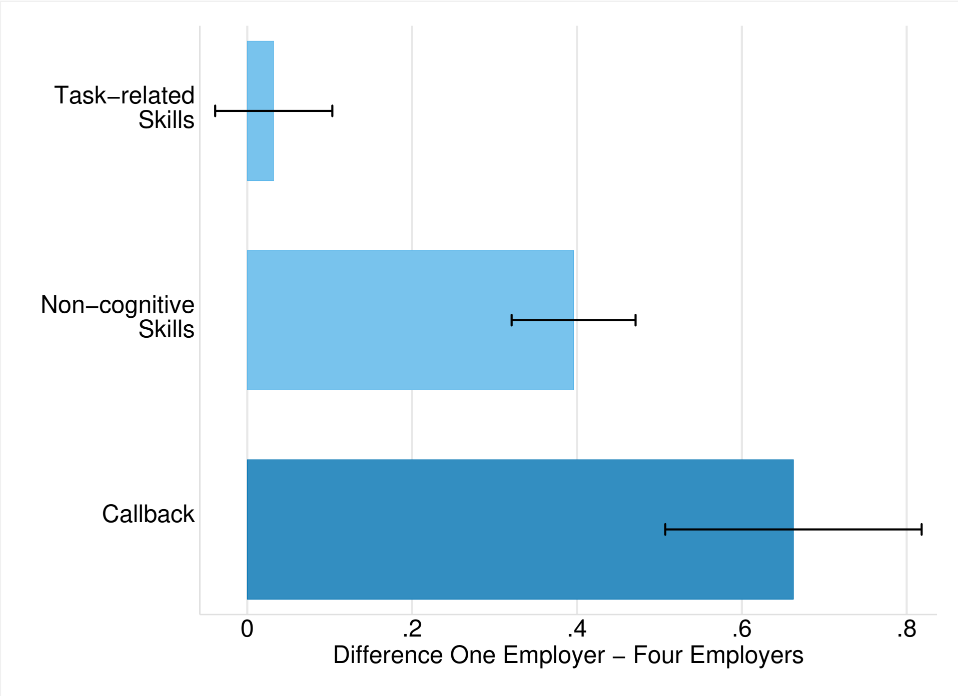

Average difference in ratings, based on a 7-point Likert scale, between the One Employer and Four Employers candidate. Error bars indicate standard error of the mean.

Moreover, the HR professionals indicate that they would be more likely to call back the One Employer than the Four Employers candidate for a job interview ($p < 0.001$, WSR). We thus replicate that employers are more likely to invite those candidates for an interview who change jobs less frequently, confirming the key result from our field experiment with a separate sample of HR professionals.

Result 5 (HR professionals' perceptions of applicants)

HR professionals perceive applicants with more frequent job changes to have lower non-cognitive skills than those with fewer changes. In contrast, applicants with varying numbers of job changes are perceived similarly in terms of task-related skills and experience. Moreover, we replicate our previous result that HR professionals prefer candidates with fewer changes.

To assess the extent to which the perceived difference in applicants' qualities can ac-

Table 4: Regression analysis of HR professionals' perceptions of the applicants

	(1)	(2)	(3)	(4)
Four Employers	-0.663***	-0.282**	-0.643***	-0.292**
	(0.156)	(0.123)	(0.142)	(0.124)
Non-cognitive skills		0.962***		0.932***
		(0.185)		(0.205)
Task-related skills			0.617***	0.043
			(0.188)	(0.219)
Constant	5.518***	0.626	2.237**	0.547
	(0.129)	(0.978)	(1.032)	(1.160)
adj. R^2	0.053	0.275	0.152	0.271
N	166	166	166	166
% explained	—	57.5	3.0	55.9

OLS regressions, cluster-robust standard errors in parentheses, clustered at the recruiter level. Unit of observation: recruiter-resume (2 resumes per recruiter).
Significance levels: * p<0.1, ** p<0.05, *** p<0.01.
Dependent variable: callback rating for a resume (7-point Likert scale).
Independent variables: "One Employer:" dummy variable for resume with only one employer; "Task-related skills" unweighted average of ratings on "skilled," "experienced in commerce," and "multi-talented;" "Non-cognitive skills:" unweighted average of ratings on "able to work in teams," "willing to adapt," "perseverant," "honest," "reliable," "self-directed," and "goal-oriented."
% explained: result of Oaxaca-Blinder decomposition of Four Employers effect: how much of the 0.663 point treatment difference in invitation ratings is explained by the difference in the respective regressors?

count for the difference in the callback rates, we estimate the following regression model:

$$y_{ij} = \alpha + \beta_1 * N_{ij} + \beta_2 * A_{ij} + \beta_2 * S_{ij} + \epsilon_{ij}. \tag{3}$$

The dependent variable y_{ij} is the stated likelihood of a callback (between 1 and 7) a recruiter j assigns to a candidate i. N_{ij} is a dummy variable indicating the Four Employers treatment. We also include the applicants' score for non-cognitive skills (A_{ij}) and task-related skills and experience (S_{ij}). We estimate the model using OLS and correct the standard errors to account for dependence in the error term ϵ_{ij} at the recruiter level. Column 1 in Table 4 reports the unconditional effect of the Four Employers treatment: callback likelihood ratings are, on average, 0.66 points lower in the Four Employers than in the One Employer treatment ($p < 0.001$, t-test). In column 2, we add the non-cognitive skills score and find that the coefficient is close to one and highly significant ($p < 0.001$, t-test). That is, an increase in the perceived level of non-cognitive skills by one point increases the callback likelihood rating by roughly one point. Crucially, the Four Employers treatment effect shrinks from -0.663 to -0.282, which corresponds to a reduction of 57.5%. This decrease suggests that over half of the treatment effect can be explained by the fact that recruiters rate the Four Employers candidate lower on non-cognitive skills than the One Employer candidate.[33]

By contrast, although the task-related skills score is positively associated with the callback likelihood, it does not explain much of the treatment effect (see column 3). While an increase in task-related skills by one point increases the callback likelihood rating by about 0.6 points ($p = 0.002$, t-test), the Four Employers coefficient decreases by only 3%. Hence, perceptions of task-related skills and experience are predictive of the callback likelihood, but they do not help explain why the Four Employers candidate is less likely to be invited for a job interview than the One Employer applicant.

Finally, column 4 includes both scores simultaneously as regressors. The coefficients of both the Four Employers treatment and non-cognitive skills score remain virtually unchanged compared to column 2. By contrast, the coefficient of task-related skills is close to zero and statistically insignificant ($p = 0.843$, t-test). Hence, perceived non-cognitive skills are more strongly related to callback likelihood ratings than perceived task-related skills.

Result 6 (Explanatory power of non-cognitive skills)
Recruiters report they are less likely to call back applicants with more frequent job changes, and they largely do so because they perceive them to have poorer non-cognitive skills relative to those with fewer changes. Perceptions of task-related skills and experience do not explain the difference in the callback likelihood rating.

4. Conclusion

This paper puts forth a novel interpretation of the relationship between job changes and employability. We argue that job changes can provide a signal of a worker's non-cognitive

[33] This result is equivalent to a pooled Blinder-Oaxaca decomposition (Blinder 1973; Oaxaca 1973).

skills, such as cooperativeness, reliability, and ability to work well with others. Our motivating hypothesis is that workers who are less cooperative, reliable, team-oriented, and generally more difficult to get along with will often be, holding all else equal, the ones who change jobs more frequently. Therefore, we expect that prospective employers will use employment history as a signal of non-cognitive skills and discriminate against employees who change jobs more frequently. An inspection of the National Labor Survey of Youth 1997 provides some correlational support for these predictions.

Building on these suggestive relationships, we combine lab, field, and survey experiments to test our hypotheses directly. In the laboratory study, we find that employment history provides a signal of non-cognitive skills. Workers who switch employment less frequently are more likely to fulfill employers' effort requests. Firms recognize this and exhibit a preference for hiring workers with fewer job changes when this information is available. In the field experiment we sent out pairs of resumes for open job positions—one resume in which the applicant changed jobs frequently and another in which the applicant remained with a single employer. As in the laboratory, we find that employers exhibit a preference for candidates with fewer job changes: frequent changes result in substantially lower callback rates. To verify that the differential demand for the candidates from the field experiment is due to firms' perceptions of the candidates' non-cognitive skills, we conducted a survey of HR professionals. The results confirm that a primary inference recruiters make from the resumes is that workers who switch jobs more frequently have poorer non-cognitive skills—particularly regarding reliability, perseverance, and ability to work in teams. Moreover, this perception accounts for a large part of recruiters' stated preferences for the applicant with fewer prior job changes.

Hence, two central results emerge from our studies. First, firms prefer workers who change jobs less frequently, at least in the contexts we study. Second, changing jobs less frequently is positively correlated and perceived to be correlated with measures of non-cognitive skills.

Several additional observations arise from our studies. In the laboratory, when firms can observe work histories, those workers with fewer job changes earn considerably more following a shock in which everyone has to search for new employment. We also find greater history dependence in labor market outcomes when employment histories are available: workers tend to stay either employed or unemployed for longer periods. In combination, these findings suggest that concerns about appearing to have poor non-cognitive skills may, in some cases, create labor market inefficiencies. Perhaps most importantly, they suggest a possible friction in labor market mobility—workers may fear changing jobs due to the impact on their perceived non-cognitive skills.

In our studies, frequent job changes hurt employability. However, there may be contexts in which frequent changes convey desirable qualities, such as varied experience, larger professional networks, and greater ambition. Any of these aspects may mitigate or entirely counteract the effects we observe in our studies. Indeed, we find that industry experience—which is more likely for an applicant with more frequent job changes—increases the likelihood of a favorable response from a prospective employer in our field experiment. Thus, even in our data, there are ways in which employment changes can be beneficial. Furthermore, our analysis focuses primarily on horizontal job changes; in contrast, frequent vertical moves may be much less likely to be perceived negatively because they lead to more challenging and better-paid positions. Hence, we acknowledge that there may be contexts in which the total effect of the frequency of job changes is positive. Our argument is that where non-cognitive skills and reliability are important relative to concerns like those above, the market will interpret frequent job changes as a negative signal of this quality, so that workers' labor market prospects may be harmed by changing jobs frequently. We leave for future work to identify the critical boundaries on our findings.

References

Altonji, Joseph G., Anthony A. Smith, and Ivan Vidangos. 2013. "Modeling Earnings Dynamics." *Econometrica* 81 (4): 1395–1454.

Arrow, Kenneth J. 1973. "Higher Education As a Filter." *Journal of Public Economics* 2 (3): 193–216.

Bagger, Jesper, Fran Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin. 2014. "Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics." *American Economic Review* 104 (6): 1551–96.

Bartel, Ann P. 1980. "Earnings Growth on the Job and between Jobs." *Economic Inquiry* 18 (1): 123–137.

Bartling, Björn, Ernst Fehr, and Klaus M Schmidt. 2012. "Screening, Competition, and Job Design: Economic Origins of Good Jobs." *American Economic Review* 102 (2): 834–864.

Bartoš, Vojtěch, Michal Bauer, Julie Chytilová, and Filip Matějka. 2016. "Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition." *The American Economic Review* 106 (6): 1437–1475.

Becker, Brian E., and Stephen M. Hills. 1983. "The Long-Run Effects of Job Changes and Unemployment among Male Teenagers." *Journal of Human Resources* 18 (2): 197–212.

Becker, Gary S. 1962. "Investment in Human Capital: A Theoretical Analysis." *Journal of Political Economy* 70 (5): 9–49.

———. 1964. *Human capital; a theoretical and empirical analysis, with special reference to education.* New York: Columbia University Press.

Bedard, Kelly. 2001. "Human capital versus signaling models: university access and high school dropouts." *Journal of Political Economy* 109 (4): 749–775.

Bertrand, Marianne, and Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review* 94 (4): 991–1013.

Blinder, Alan S. 1973. "Wage discrimination: reduced form and structural estimates." *Journal of Human Resources* 8 (4): 436–455.

Blumen, Isadore, Marvin Kogan, and Philip J. McCarthy. 1955. The Industrial Mobility of Labor as a Probability Process. Ithaca, NY: Cornell University Press.

Bock, Olaf, Ingmar Baetge, and Andreas Nicklisch. 2014. "hroot: Hamburg Registration and Organization Online Tool." *European Economic Review* 71:117–120.

Borjas, George J. 1981. "Job Mobility and Earnings over the Life Cycle." *Industrial and Labor Relations Review* 34 (3): 365–376.

Bowles, Samuel, Herbert Gintis, and Melissa Osborne. 2001. "The Determinants of Earnings: A Behavioral Approach." *Journal of Economic Literature* 39 (4): 1137–1176.

Brüggen, Alexander, and Martin Strobel. 2007. "Real Effort Versus Chosen Effort in Experiments." *Economics Letters* 96 (2): 232–236.

Burks, Stephen, Bo Cowgill, Mitchell Hoffman, and Michael Housman. 2015. "The Value of Hiring through Employee Referrals." *Quarterly Journal of Economics* 130 (2): 805–839.

Carlsson, Magnus, and Dan-Olof Rooth. 2007. "Evidence of Ethnic Discrimination in the Swedish Labor Market Using Experimental Data." *Labour Economics* 14 (4): 716–729.

Datta, Somnath, and Glen A Satten. 2005. "Rank-sum Tests for Clustered Data." *Journal of the American Statistical Association* 100 (471): 908–915.

Deming, David J. 2015. *The Growing Importance of Social Skills in the Labor Market.* Technical report, Working Paper 21473. National Bureau of Economic Research.

Deming, David J, Noam Yuchtman, Amira Abulafi, Claudia Goldin, and Lawrence F Katz. 2016. "The Value of Postsecondary Credentials in the Labor Market: An Experimental Study." *American Economic Review* 106 (3): 778–806.

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde. 2009. "Homo Reciprocans: Survey Evidence on Behavioural Outcomes." *Economic Journal* 119 (536): 592–612.

Dustmann, Christian, and Costas Meghir. 2005. "Wages, Experience and Seniority." *Review of Economic Studies* 72 (1): 77–108.

Eriksson, Stefan, and Dan-Olof Rooth. 2014. "Do Employers use Unemployment as a Sorting Criterion when Hiring? Evidence from a Field Experiment." *American Economic Review* 104 (3): 1014–1039.

Farber, Henry S. 1999. "Mobility and Stability: The Dynamics of Job Change in Labor Markets." In *Handbook of Labor Economics,* edited by Orley. Ashenfelter and David E. Card, vol. 3, Part B, 2439–2483. Elsevier.

Fischbacher, Urs. 2007. "z-Tree: Zurich Toolbox for Ready-made Economic Experiments." *Experimental Economics* 10 (2): 171–178.

Gërxhani, Klarita, Jordi Brandts, and Arthur Schram. 2013. "The Emergence of Employer Information Networks in an Experimental Labor Market." *Social Networks* 35 (4): 541–560.

Ghiselli, Edwin E. 1974. "Some Perspectives for Industrial Psychology." *American Psychologist* 29 (2): 80–87.

Granovetter, Mark. 1974. *Getting a job: A study of contacts and careers.* Cambridge, MA: Harvard University Press.

Head, Allen, and Huw Lloyd-Ellis. 2012. "Housing Liquidity, Mobility, and the Labour Market." *Review of Economic Studies* 79 (4): 1559–1589.

Heckman, James J, and Yona Rubinstein. 2001. "The Importance of Noncognitive Skills: Lessons from the GED Testing Program." *American Economic Review* 9 (2): 145–149.

Heckman, James J, Jora Stixrud, and Sergio Urzua. 2006. "The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior." *Journal of Labor Economics* 24 (3): 411–482.

Jovanovic, Boyan. 1979. "Firm-specific Capital and Turnover." *Journal of Political Economy* 87 (6): 1246–1260.

Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo. 2013. "Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment." *Quarterly Journal of Economics* 1123 (3): 1123–1167.

Light, Audrey, and Kathleen McGarry. 1998. "Job Change Patterns and the Wages of Young Men." *Review of Economics and Statistics* 80 (2): 276–286.

Lindqvist, Erik, and Roine Vestman. 2011. "The Labor Market Returns to Cognitive and Noncognitive Ability: Evidence from the Swedish Enlistment." *American Economic Journal: Applied Economics* 3 (1): 101–128.

Mincer, Jacob. 1958. "Investment in Human Capital and Personal Income Distribution." *Journal of Political Economy* 66 (4): 281–302.

Mincer, Jacob, and Boyan Jovanovic. 1982. "Labor Mobility and Wages." In *Studies in Labor Markets,* edited by Sherwin Rosen, 21–64. Chicago: University of Chicago Press for NBER.

Oaxaca, Ronald. 1973. "Male-female Wage Differentials in Urban Labor Markets." *International Economic Review* 14 (3): 693–709.

Oberholzer-Gee, Felix. 2008. "Nonemployment Stigma as Rational Herding: A Field Experiment." *Journal of Economic Behavior & Organization* 65 (1): 30–40.

Oreopoulos, Philip. 2011. "Why Do Skilled Immigrants Struggle in the Labor Market? A Field Experiment with Thirteen Thousand Résumés." *American Economic Journal: Economic Policy* 3 (4): 148–171.

Pager, Devah, Bruce Western, and Bart Bonikowski. 2009. "Discrimination in a Low-wage Labor Market a Field Experiment." *American Sociological Review* 74 (5): 777–799.

Pallais, Amanda, and Emily Glassberg Sands. 2016. "Why the Referential Treatment? Evidence from Field Experiments on Referrals." *Journal of Political Economy* 124 (6): 1793–1828.

Petrongolo, Barbara, and Christopher A. Pissarides. 2001. "Looking into the Black Box: A Survey of the Matching Function." *Journal of Economic Literature* 39 (2): 390–431.

Rees, Albert. 1966. "Information Networks in Labor Markets." *American Economic Review* 56 (1/2): 559–566.

Riach, Peter A, and Judith Rich. 2002. "Field Experiments of Discrimination in the Market Place." *Economic Journal* 112 (483): F480–F518.

Rogerson, Richard, Robert Shimer, and Randall Wright. 2005. "Search-theoretic Models of the Labor Market: A Survey." *Journal of Economic Literature* 43 (4): 959–988.

SECO, Swiss State Secretariat for Economic Affairs. 2013. "Monthly Labor Market Statistics." `https://www.amstat.ch/`.

Siegel, Sidney, and N. Castellan John. 1988. Nonparametric Statistics for the Behavioral Sciences. New York: McGraw-Hill.

Spence, Michael. 1973. "Job Market Signaling." *Quarterly Journal of Economics* 87 (3): 355–374.

Stiglitz, Joseph E. 1975. "The Theory of "Screening," Education, and the Distribution of Income." *American Economic Review* 65 (3): 283–300.

Suster, Mark. 2010. *Never Hire Job Hoppers. They Make Terrible Employees.* `http://www.businessinsider.com/mark-suster-never-hire-job-hoppers-never-they-make-terrible-employees-2010-4`. Blog.

Swiss Federal Statistical Office. 2008. *Swiss Labour Force Survey.* `http://www.bfs.admin.ch/bfs/portal/en/index/themen/03/22/publ.html?publicationID=5625`.

Topel, Robert H., and Micheal P. Ward. 1992. "Job Mobility and the Careers of Young Men." *Quarterly Journal of Economics* 107 (2): 439–479.

Trikha, Ritika. 2012. *How a Gen-y Job Seeker Can Justify Their Job Hopping.* `http://money.usnews.com/money/blogs/outside-voices-careers/2012/10/23/how-a-geny-job-seeker-can-justify-their-job-hopping`. Blog.

Tyler, John H, Richard J Murnane, and John B Willett. 2000. "Estimating the Labor Market Signaling Value of the GED." *Quarterly Journal of Economics* 115 (2): 431–468.

Waldman, Michael. 1984. "Job Assignments, Signalling, and Efficiency." *RAND Journal of Economics* 15 (2): 255–267.

Appendices

A. Laboratory Experiment: Appendix

We examined the relationship between voluntary work effort and the frequency of job changes using Ordinary Least Squares (OLS) regressions. Our analysis is based on the following linear regression model:

$$N_i = \alpha + \beta(e_i - 1) + \varepsilon_{im}. \tag{4}$$

The dependent variable, N_i, is the number of employers a worker i had in the 16 periods before the turnover shock and e_i is the worker's effort level in periods 1 to 16. We use $(e_i - 1)$ in the regression model so that the constant, α, can be interpreted as the number of employers of a worker who provided the minimum effort of 1 before the shock.[34] We allow the error terms, ε_{im}, to be correlated within each labor market.

Table 5: Regression analysis of number of employers

Condition	(1) History	(2) No History	(3) Pooled
Avg. Effort Periods 1-16	-0.357***	-0.342***	-0.342***
	(0.036)	(0.043)	(0.042)
History			-0.276
			(0.405)
History X Avg. Effort 1-16			-0.015
			(0.055)
Constant	5.051***	5.327***	5.327***
	(0.295)	(0.286)	(0.281)
adj. R^2	0.337	0.243	0.303
N	170	160	330

OLS regressions, standard errors in parentheses, adjusted for clustering at the session level using White sandwich estimators. Unit of observation: workers.
Significance levels: * p<0.1, ** p<0.05, *** p<0.01.
Dependent variable: number of different employers before the shock (periods 1 to 16).
Independent variables: Constant: average number of pre-shock employers for a worker in the No History condition (History condition for column 1) who provided minimum effort; "Effort Periods 1-16:" effort provided by the worker in periods 1 to 16 (subtracting 1 $(e_i - 1)$ to facilitate interpretation of the constant); "History:" dummy for History treatment condition; "History × Avg. Effort 1-16:" interaction between History dummy and pre-shock effort.

Column 1 in Table 5 reports the regression results for the History treatment. The constant of about 5 indicates that a worker who provided the minimum effort before the shock had, on average, five different employers (out of a maximum of 7) during that time. Increasing first-period effort by one unit is associated with a reduction of the number of pre-shock employers by about 0.36 ($p < 0.001$, t-test). We observe a similar pattern in the No History condition, as shown in column 2. Providing minimum effort

[34]Every participant had at least one employer before the shock. We obtain similar results if we use, instead, first-period effort as an explanatory variable or if we control for the number of periods unemployed before the shock.

results in 5.3 pre-shock employers, and increasing effort by one unit reduces the number of pre-shock employers by about 0.34 ($p < 0.001$, t-test). In column 3, we pool the data from both treatments and additionally include a dummy for the History treatment as well as its interaction with the number of employers. This allows us to test whether the relationship between effort and number of pre-shock employers is stronger in the History condition. Yet, both the coefficient of the History dummy and the interaction term are insignificant ($p = 0.501$ and 0.787, t-tests), confirming that the relationship between effort provision and job history is similar in both conditions. Together, these findings support our prediction that workers who change jobs frequently are less reliable and cooperative.[35]

B. Field Experiment: Appendix

Table 6: Descriptive statistics

	Mean	Sd
May	0.673	0.470
June	0.199	0.399
Industry: cars	0.026	0.160
Industry: bank	0.019	0.137
Industry: chemical	0.023	0.149
Industry: service and admin	0.235	0.424
Industry: trade	0.115	0.320
Industry: tourism	0.007	0.084
Industry: construction/housing	0.086	0.280
Industry: logistics	0.031	0.173
Industry: communication	0.036	0.186
Industry: electro/metal industry	0.151	0.358
Industry: food industry	0.014	0.119
Industry: legal	0.036	0.186
Industry: public administration	0.031	0.173
Industry: insurance	0.012	0.109

	Mean	Sd
Industry: travel agency	0.005	0.069
Industry: health service	0.023	0.149
Industry: hospital	0.031	0.173
Industry: transport	0.007	0.084
Industry: fiduciary	0.096	0.295
Industry: other	0.017	0.128
Legal: public or ngo	0.088	0.284
Legal: LLC	0.877	0.328
Legal: other	0.035	0.183
Employment agency	0.170	0.376
Part-time job	0.175	0.380
Avg. ln(driving distance)	9.704	1.322
Male HR person	0.321	0.467
Male applicant	0.487	0.500
Applicants per vacancy	9.709	5.013
Local unemployment rate	2.781	0.406

[35]Note that this relationship alone does not tell us the reasons behind job changes—that is, whether a worker left the employer for a better offer elsewhere or whether the current employer did not make another offer to the worker. Our data indicate that job changes tend to be driven by employers. Specifically, in 86% of the cases in which workers changed jobs, they did not receive a private offer from their old employer. On the other hand, 91% of private offers from a worker's previous employer are accepted.

Table 8: Regression analysis: alternative callback definition

Dependent variable	Callback = 1					
	(1)	(2)	(3)	(4)	(5)	(6)
Four Employers	-0.033**	-0.032*	-0.033*	-0.034**	-0.033*	-0.056***
	(0.014)	(0.017)	(0.017)	(0.017)	(0.017)	(0.019)
Four Emp. X 2012		-0.003	-0.003	-0.003	-0.003	0.007
		(0.028)	(0.029)	(0.029)	(0.029)	(0.029)
Wave 2012		0.068**	0.064**	0.020	0.043	0.049
		(0.032)	(0.032)	(0.035)	(0.035)	(0.035)
Industry experience						0.107***
						(0.037)
Constant	0.292***	0.264***	0.781***	0.835***	0.566*	0.546*
	(0.016)	(0.020)	(0.179)	(0.254)	(0.304)	(0.302)
Additional controls?						
Month				Yes	Yes	Yes
Gender/gend. match			Yes		Yes	Yes
Firm/job character.			Yes		Yes	Yes
Driving distance			Yes	Yes	Yes	Yes
Labor market				Yes	Yes	Yes
Observations	1680	1680	1680	1680	1680	1680
F	5.881	3.829	6.993	5.286	7.146	6.212
Prob> F	0.016	0.010	0.000	0.000	0.000	0.000

This table shows OLS coefficient estimates (standard errors adjusted for clustering at the job advertisement level are reported in parentheses).
Dependent variable: dummy indicating a callback.
Independent variables: "Four Employers:" dummy for Four Employers profile; "Wave 2012:" dummy for first wave of study (in 2012); "Industry experience:" dummy whether applicant has had previous work experience in the corresponding industry; "Month:" dummies for month when application was sent; "Gender/gender match:" dummies for gender of applicant and HR person, and corresponding interaction term; "Firm/job characteristics:" industry dummies, legal form dummies, employment agency dummy, and part-time job dummy; "Driving distance:" log of distance in meters by car (using Google Maps); "Labor market:" monthly local unemployment rate and number of applicants per open position (statistics from State Secretariat for Economic Affairs (SECO)).
Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

C. Survey Experiment: Appendix

Table 9: Descriptive statistics of the participants.

Variable	**Mean**	**Median**
Firm size (employees)	24'892	1'300
Staff at booth	3.5	3
# resumes/month	54.5	30
Years HR experience	6.7	5
% female	59	—
Age (10-year bracket)	—	25–35
Sample size	83	

Industry	**#**
Plant Engineering/-Construction	14
Electrical Ind./Electronics	12
IT / Telecom	10
Consulting	12
Mechanical Engineering	8
Chemical Ind./Pharma	5
Medical Technology	3
Financial Services/Banking	3
Optomechanics	2
Consumer Goods	2
Other	12
Total	83

Table 11: Difference in ratings of the 10 different characteristics (One Employer Rating rating minus Four Employers rating), mean and p-value of paired t-test. $N = 83$.

Characteristic	Mean Diff.	p-value (corr.)
perseverant	1.24	<0.001***
reliable	0.77	<0.001***
teamwork	0.40	<0.001***
honest	0.27	0.199
skilled	0.19	0.229
willing to adapt	0.34	0.299
goal-oriented	-0.17	0.989
self-directed	-0.07	1.000
multi-talented	-0.05	1.000
experienced	-0.05	1.000

* p<0.10, ** p<0.05, *** p<0.01, Holm-Bonferroni correction.