Job Search Behavior among the Employed and Non-Employed*

R. Jason Faberman
Federal Reserve Bank of Chicago

Andreas I. Mueller
Columbia University, NBER and IZA

Ayşegül Şahin
Federal Reserve Bank of New York

Giorgio Topa
Federal Reserve Bank of New York and IZA †

May 2018

Abstract

We develop and implement a unique new survey that focuses on the job search behavior of individuals regardless of their labor force status. We use our survey to study the relationship between search effort and outcomes for employed and non-employed workers. We find that on-the-job search is pervasive: 23 percent of the employed engage in some form of job search over a month. We also find that the employed's search activity exhibits a strong negative relationship with their current wage. Furthermore, we document that the employed fare better than the non-employed in job search: they exert lower effort but receive more offers per unit of search effort and are more likely to receive unsolicited job offers. The employed also receive job offers with substantially higher wages. Differences in observable characteristics of the worker, job, and search process explain only about half of the wage offer differential. Finally, we set up an on-the-job search framework with endogenous search effort, unobserved heterogeneity and censored wage offers. We parameterize the model carefully and match it to a number of key moments from our survey. The model delivers a realistic degree of wage dispersion while implying a reasonable flow value of unemployment, and its implied relationship between search intensity and the wage are remarkably similar to the patterns in our data. Our model also implies that after accounting for unobserved heterogeneity and censoring, the employed enjoy a 16 log point wage offer premium over the unemployed.

Keywords: job search, unemployment, on-the-job search, search effort, wage dispersion
JEL Classifications: E24, J29, J60

*We thank Christine Braun, Fatih Karahan, Rasmus Lentz, Chris Moser, Emi Nakamura, Richard Rogerson, Robert Valletta, and Thijs van Rens, in addition to participants from several conferences and seminars, for useful comments. We also thank Luis Armona, René Chalom, Rebecca Friedman, Thomas Haasl, Max Livingston, Sean Mihaljevich, and Rachel Schuh for their excellent research assistance. The views expressed here are our own and do not necessarily reflect those of the Federal Reserve Banks of New York or Chicago, or those of the Federal Reserve System.

†Addresses: Faberman: Federal Reserve Bank of Chicago, 230 South LaSalle Street, Chicago, IL 60604 (email: jfaberman@chi.frb.org); Mueller: Columbia Business School, Columbia University, Uris Hall, 3022 Broadway, New York, NY 10027 (email: amueller@columbia.edu); Şahin: Federal Reserve Bank of New York, 33 Liberty Street, New York, NY 10045 (email: aysegul.sahin@ny.frb.org); Topa: Federal Reserve Bank of New York, 33 Liberty Street, New York, NY 10045 (email: giorgio.topa@ny.frb.org).

1 Introduction

Job-to-job transitions are an important feature of the U.S. labor market. Fallick and Fleischman (2004) show that, at the end of the 1990s, the monthly job-to-job transition rate was 2.6 percent, which accounted for around 40 percent of all monthly hiring.[1] Despite the critical importance of on-the-job search for understanding labor market dynamics and the central role it has in search theories of the labor market, evidence on the extent and nature of the on-the-job search remains scant.[2] Often, theory must make unguided assumptions about the effort, efficiency, and outcomes of those searching while employed and how they differ from those searching while unemployed. In short, we know little about whether search behavior on-the-job differs in an economically meaningful way from the search behavior of the unemployed.

In this paper, we help fill this void. We document the search behavior and outcomes of employed and non-employed individuals and assess the importance of our findings for search-theoretic models of the labor market. To this end, we design and implement a unique new survey that focuses on job search behavior and outcomes for *all* individuals, regardless of their labor force status. We administer our survey as a supplement to the Survey of Consumer Expectations and have fielded it annually each October since 2013. The survey asks an expansive list of questions on the employment status and current job search, if any, of all respondents, including questions on an individual's search effort and search methods, the incidence of informal recruiting methods, and demographic information.

Our findings provide the most comprehensive evidence to date on the nature of on-the-job

[1]More recently, the job-to-job transition rate has declined (Davis and Haltiwanger, 2014), but so has the transition rate from unemployment to employment, so job-to-job transitions remain an important fraction of aggregate hiring in the U.S. labor market. Furthermore, job-to-job transitions are an important driver of aggregate wage and productivity growth (Faberman and Justiniano, 2015; Moscarini and Postel-Vinay, 2017; Karahan et al., 2017).

[2]Notable exceptions are earlier work by Kahn (1982), Holzer (1987) and Blau and Robins (1990). As we detail in the empirical section, these studies are based on older, discontinued surveys. Researchers have used the more recent American Time Use Survey (ATUS) extensively to analyze the job search behavior of the unemployed (see Krueger and Mueller, 2010, Aguiar, Hurst, and Karabourbanis, 2013, and Mukoyama, Patterson, and Şahin, 2018), but the ATUS is not well-suited to analyze the job search effort of the employed as we discuss in the empirical section. Finally, Kahn (2012) uses labor force data for a sample of 11 European countries and finds that the incidence of job search is higher among those on temporary contracts compared to those on permanent contracts.

search for the U.S. While we uncover various new facts, five key findings stand out. First, on-the-job search is pervasive; 23 percent of the employed report looking for work during our survey months. Second, the intensity of on-the-job search declines strongly with a worker's *residualized* current hourly wage. Third, on-the-job search is more efficient than search by the unemployed: the employed receive a similar number of offers, and a disproportionate number of unsolicited offers, despite exerting a fraction of the search effort of the unemployed. Fourth, the employed appear to sample from a higher-quality job offer distribution than the non-employed. Unconditionally, the wages offered to the employed are 40 log points (49 percent) higher than the wages offered to the non-employed. Accounting for observable worker and job characteristics only reduces the wage offer differential to 25 log points (28 percent). Fifth, despite the poorer quality of their job offers, the non-employed are about one-and-a-half times more likely to accept an offer.

The finding that employed workers face better wage offers suggests that factors that are unique to employment status are important determinants of the hiring process. An obvious concern about this interpretation, however, is that unobserved differences in productivity between employed and non-employed job seekers may be the reason for what appears to be a *wage offer premium*. Those with higher unobserved skills are more likely to be employed and earn higher wages, so a wage offer premium is a natural consequence of this selection effect. Consequently, an individual's prior work history provides a useful proxy for unobserved heterogeneity that may be correlated with one's current labor force status. We take this interpretation seriously and use our survey data to construct individual labor force histories over the previous five years. Controlling for labor force history reduces the *residual* wage offer gap from 25 log points (28 percent) to 20 log points (22 percent). A worker's prior wage can provide additional information regarding unobserved heterogeneity since those with better unobserved skills are more likely to have had higher-paying jobs in the past. Therefore, we additionally use the *prior* wages of workers as a proxy for unobserved heterogeneity. Including prior wages as a control, however, widens the wage gap, an issue we analyze further in our theoretical analysis. Lastly, employed and non-employed workers may have systematic differences in their job search process or in their preferences for non-wage job benefits, yet our estimates of the wage offer premium are little changed when we control for the source of a job offer and the type of non-wage benefits offered. When accounting for all of these factors in our regression analysis, we find that they account for, at most, about half of the wage offer gap between the employed and non-employed.

We next study the quantitative implications of our empirical findings for search theoretic models of the labor market and further assess the importance of unobserved heterogeneity for explaining the wage offer premium of the employed. We set up an on-the-job search framework with additional features that the data support. These include endogenous search effort and exogenous differences in search efficiency and search costs between the employed and unemployed. We parameterize the model carefully by matching it to a number of key moments from our survey data. Our model fits the data well, matching the job-acceptance and thus the job-to-job transition rate of the employed exactly despite not targeting this moment explicitly in our parameterization. It is also able to replicate the negative relationship between search effort and wages observed in the data remarkably well. Given the prominence of this implication in most models with endogenous search effort and on-the-job search, we interpret this success as support for such models as a realistic characterization of the search process.

We also use our model to quantitatively investigate three potential explanations for the wage offer premium for the employed. First, we allow for unobserved *ex ante* worker heterogeneity and the potential selection of low-productivity workers towards unemployment. We discipline the selection effect in our model with the differential in prior wages observed in our data. Second, we allow for censoring of the wage offer distribution to account for our finding that the employed are more selective in the pre-offer stage of the search process. We calibrate the degree of censoring using specific questions on *unrealized offers* in our survey. Finally, we allow for exogenous differences in the wage offer distributions of the employed and unemployed that are independent of censoring and unobserved heterogeneity.

The model estimates suggest that low-productivity workers are negatively selected into unemployment, but the negative selection only accounts for 28 percent (7 log points) of the 25 log point residual wage offer differential after controlling for observable worker and job characteristics.[3] We also provide evidence that controlling for the fraction of time employed over the last five years (in addition to the standard observables) reduces the unexplained wage offer premium only to a limited extent, but in similar magnitudes in the data and the model. This suggests that our model—which exploits prior wages rather than employment histories as a source of variation for the purposes of identification—captures the unobserved heterogeneity that is relevant for the empirically-observed wage offer premium. We find that censoring can account for only about 8 percent (2 log points) of the residual wage offer differential. The fact that we match the job-to-job transition rate exactly mitigates concerns that we underestimate the contribution of censoring since a greater degree of censoring would lead to counterfactually high job acceptance and job-to-job transition rates. The remaining wage offer premium is substantial and amounts to nearly two thirds of the observed differential (16 log points).

An open question remains for why wage offers are of better quality for the employed. We do not model the micro-foundations for the remaining wage offer differential, but its reduced form is consistent with several potential explanations. These include auction models where incumbent and outside employers compete for workers and bid up the wages of the employed (Postel-Vinay and Robin, 2002, and Cahuc, Postel-Vinay and Robin, 2006). It is also consistent with employer discrimination against the unemployed, either to take advantage of their lower reservation wages (as in Carrillo-Tudela, 2009) or through a stigma effect (as in Gibbons and Katz, 1991). Finally, it is consistent with the employed having access to networks that provide better job offers, at least to an extent that is not captured by our controls for the job search process in the regression analysis.

Finally, we assess the implications of our model for the *frictional wage dispersion puzzle*. Our model implies a reasonable flow value of unemployment of 0.74, passing a key test advocated

[3]Our model estimates are consistent with Mueller (2017), who finds in the Current Population Survey that unemployment risk is 36 percent higher for workers below the median residual wage (i.e., after controlling for observables) compared to those above the median residual wage.

by Hornstein, Krusell, and Violante (2011). These authors show that standard search models (such as in McCall, 1970; Mortensen, 1977; Pissarides, 1985) exhibit a tension between observed worker flows, wage dispersion, and the flow value of unemployment. The tension arises because high dispersion in potential wage offers generates a large option value of waiting for a better offer, so only a low (often negative) flow value of unemployment can rationalize the observed flow rate of the unemployed into employment. In our model, the option value of unemployment is limited since employed searchers are relatively more efficient in generating offers and face a better wage offer distribution. Therefore, leaving unemployment for a job does not require unemployment to have a low flow value. Our empirical findings thus constitute an important piece in the resolution of the frictional wage dispersion puzzle.

In sum, our paper makes two key contributions: First, we design and implement a unique new survey and use it to document several new stylized facts on the search process of the employed relative to the unemployed. Our findings speak to margins that are at the heart of on-the-job search models but have been mostly unobservable given the data available thus far. We find that on-the-job search dominates search while unemployed along several key dimensions. Second, we examine these findings through the lens of a job ladder model and find strong empirical support for this class of models. The model implies that the employed enjoy a substantial wage offer premium even after accounting for unobserved heterogeneity and censoring and it provides an intuitive resolution of the frictional wage dispersion puzzle.

The next section describes our survey. Section 3 presents our evidence concerning on-the-job search behavior and job search outcomes by labor force status. Section 4 presents a model of on-the-job search with endogenous search effort and discusses its quantitative implications. Section 5 concludes and offers some additional thoughts on the mechanisms underlying the observed wage offer differential between the employed and non-employed job seekers.

2 Survey Design and Data

Our data come from a supplement to the Survey of Consumer Expectations (SCE), administered by the Federal Reserve Bank of New York. The SCE is a monthly, nationally-representative survey of roughly 1,300 individuals that asks respondents about their expectations about various aspects of the economy. We designed the supplement ourselves and first administered it in October 2013. We have administered it annually since then, and present results for a sample that pools the 2013-15 data together. Our supplement asks a broad range of questions on employment status, job search behavior, and job search outcomes. Demographic data are also available for respondents through the monthly portion of the SCE survey.

The survey asks a variety of questions that are tailored to an individual's employment status and job search behavior. For the employed, including the self-employed, the survey asks questions about their wages, hours, benefits, and the type of work that they do, including questions on the characteristics of their workplace. For the non-employed, the survey asks a range of detailed questions on their most recent employment spell and their reasons for non-employment. The survey also asks questions related to the type of non-employment, including those related to retirement, school enrollment status, and any temporary layoff. We also ask individuals about their prior work history. This includes detailed information about the preceding job of the currently employed.

Regardless of employment status, the survey asks all individuals if they have searched for work within the last four weeks, and if they had not searched, whether or not they would accept a job if one was offered to them. Among the employed, the survey distinguishes between those searching for new work and those searching for a job in addition to their current one. For individuals who have searched or would at least be willing to accept a new job if offered, the survey asks a series of questions relating to their job search (if any), including the reasons for their decision to (not) search. It then asks an exhaustive set of questions on the types of effort exerted when seeking new work (e.g., updating resumes, searching online, contacting employers directly). It also asks about the number of job applications completed within the last four weeks and the number of employer contacts and job offers received. It also probes further to see how those contacts and offers came about, i.e., whether they were the result of traditional search methods or whether they came about through a referral or an unsolicited employer contact. For those who received an offer, including any offers within the last six months, the survey asks about a range of characteristics of the job offer, including the wage offered, the expected hours, its benefits, as well as the type of work to be done and the characteristics of the employer. It also asks what led, or may lead, the respondent to accept or reject the offer, and asks a range of questions about whether there was any bargaining with either the current or future employer. Since only a fraction of respondents in our sample report a job offer in the months leading up to the survey, we ask those who are currently employed a range of additional, retrospective questions about the search process that led to their current job.

Many of the survey questions follow a format similar to the Current Population Survey (CPS), though there are notable differences. The survey identifies the labor force status of respondents at several different points in their employment history: at the time of the survey, at the time of their hiring (if currently employed), and at the time of their job offer (if they received an offer within the last six months). We also impute a labor force status for individuals four weeks prior to the survey. Our ability to identify labor force status at these different points allows us to deal with time aggregation and related issues when comparing the search and job-finding behavior of the employed and non-employed.

We define a respondent's labor force status at the time of the survey in a manner similar to the CPS, but because we ask about search effort more broadly than the CPS, we can generate two measures of unemployment. The Bureau of Labor Statistics (BLS) definition classifies someone as unemployed if they "do not have a job, have actively looked for work in the prior four weeks, and are currently available for work." Those on temporary layoff are also included regardless of search effort or availability. We employ the same definition, but due to the skip logic of the CPS survey design, there are some non-employed in the CPS who are never asked whether they searched for work. These are primarily retired individuals who state that they do not want a job (and are therefore assumed to be unavailable for work). Our survey, however, captures search effort regardless of whether a non-employed individual states that they want work. We define a respondent's labor force status at the time of the survey using the broader "job search" definition of unemployment since we aim to capture overall search activity and its effectiveness within the aggregate economy. The difference is that the "job search" definition includes non-employed individuals who did not state that they want to work but actively searched and are available, while the stricter "BLS definition" includes only those who additionally state that they want work. Given the well-known observation that even workers who state that they are not available for work and they have not searched for work transition from nonparticipation to employment in the CPS, we view our broader "job search" definition of unemployment as a reasonable one.[4]

We identify individuals as either employed or non-employed at the time of their hiring or receipt of a job offer. The survey allows for some greater disaggregation of these labor force statuses and we obtain results similar to those in our main analyses when using the more detailed definitions.[5] We also impute a labor force status four weeks prior to the survey for individuals using a range of their responses on employment status, job tenure, non-employment duration, job offer incidence and timing, etc. We detail our imputation methodology in Appendix B. Having a labor force status for individuals one month prior to the survey is useful for when we apply our empirical findings to the model because the model characterizes a job seeker's search behavior using their labor force status prior to exerting search effort or receiving any job offers.[6]

[4]In Appendix A, we replicate our estimates on search effort and outcomes using the definition consistent with the CPS survey design. The results in the appendix show that those included in the broader "job search" definition represent about 12 percent of those considered out of the labor force under the BLS definition.

[5]Specifically, the labor force status at the time of hiring distinguishes between those who quit from a previous job and those who lost their job immediately prior to starting the current job. The majority of the employed quit from their current job, so the results for this group are very similar to those reported in our analysis. The labor force status at the time of job offer distinguishes between those who were employed either full-time or part-time at the time of the offer. Most individuals were employed full-time, and consequently their results are similar to what we report in our analysis. The vast majority of the non-employed under both definitions report actively searching.

[6]We evaluate the performance of our measure of labor force status along several dimensions in the appendix. We also merged the SCE labor market module to the SCE monthly survey and used the labor market status from the most recent survey available for a given individual, in either September or August of the same year. The results using prior labor force status from the SCE monthly are very similar. See Table B2 in the appendix.

Table 1: Summary Statistics, SCE Labor Supplement vs. Current Population Survey

Labor Force Status	**SCE Labor Supplement** *Job Search Definition*	*BLS Definition*	**Current Population Survey**
Employment-population ratio	0.761	0.761	0.743
	(0.008)	(0.008)	(0.001)
Unemployment rate	8.0	5.3	5.0
	(0.5)	(0.5)	(0.1)
Labor force participation rate	82.8	80.5	78.2
	(0.7)	(0.7)	(0.1)
Demographics			
Percent male	48.7		51.4
	(0.9)		(0.1)
Percent white, non-Hispanic	72.6		63.6
	(0.8)		(0.1)
Percent married	65.5		51.6
	(0.9)		(0.1)
Percent with college degree	33.0		34.1
	(0.9)		(0.1)
Percent aged 18-39	35.2		38.9
	(0.9)		(0.1)
Percent aged 40-59	49.6		49.5
	(0.9)		(0.1)
Percent aged 60+	15.2		11.7
	(0.7)		(0.1)

Note: Estimates come from authors' tabulations from the SCE Labor Supplement and the Current Population Survey (CPS) for data pooled across October 2013, 2014, and 2015. Both samples are for heads of household ages 18 to 64. Job search definition of unemployment includes all non-employed who actively searched and are available for work, regardless of reporting whether they want work. Standard errors are in parentheses.

Our analysis uses a sample from the SCE of individuals aged 18 to 64 pooled across the 2013-15 surveys. This provides just under 2,900 observations. Individuals are only in the SCE for, at most, one year, so our sample is a panel of repeated cross sections rather than longitudinal. By design, the SCE only includes heads of household. The survey does not ask the self-employed about job search, so the self-employed are generally excluded by construction throughout our job search analysis. Table 1 presents basic summary statistics for our analysis sample and a comparable sample using the same months of data from the CPS. The demographic statistics across the two surveys are roughly similar, though the shares that are married and white are both higher in the SCE sample. The employment-to-population ratio, which is unaffected by the differing unemployment definitions, is somewhat higher in the SCE as well. Under the BLS definition, the unemployment rate in the SCE is slightly higher than the CPS rate, but not statistically different. Including the additional job seekers in the "job search" definition increases the unemployment rate considerably, from 5.3 percent to 8.0 percent, suggesting that the BLS definition of unemployment misses some search activity in the economy.

In addition to our main sample, we also focus on two subsamples of the data. The first is the subsample of the currently employed (excluding the currently self-employed). After removing respondents with missing data, this sample includes 1,763 respondents. We use this subsample to examine the job search behavior that led to their hiring at their current jobs. The second is a subsample of all individuals who received a job offer within the last six months. By construction, some of these offers will reflect the respondent's current job, which we identify through a separate question in the survey. After removing offers with only partial data, the sample has 654 observations. We use this sample to examine a range of job offer characteristics, including the offer wage distribution, as well as the characteristics of accepted job offers. Note that we first asked respondents whether they received any offer in the last month, and only if not, did we ask about offers received in the last 6 months. Thus the data allow us to determine the monthly offer rate.

3 Evidence

We now turn to our empirical analysis on job search behavior. We can summarize our main findings as follows: (i) on-the-job search is pervasive among the employed; (ii) the intensity of on-the-job search for the employed declines with their current wage; (iii) employed job seekers search less than the unemployed, but they receive just as many offers, implying that their search is more effective per unit of effort; (iv) the employed receive better offers with higher wages and benefits, even after controlling for their observable characteristics; (v) despite receiving higher-quality offers, they are less likely to accept them.

3.1 Extensive and Intensive Margins of Job Search

We begin with evidence on the basic characteristics of individual job search effort. It is useful to analyze the extensive and intensive margins of job search separately since the distribution of total search effort along both dimensions is informative for thinking about the efficiency of job search.[7] Table 2 reports the incidence of job search by labor force status, which we interpret as the *extensive margin* of job search. By definition, all unemployed, save for those on temporary layoff, search. Since we employ a search-based definition of unemployment, only a minimal amount of those out of the labor force engage in search.[8] Among the employed, roughly 20 percent can be classified as searchers regardless of the criteria we employ to define job search. Over 23 percent of the employed looked for work in the last four weeks, with 20 percent applying to at least one job and a similar amount searching at least once in the last seven days. Around 20 percent of those searching on the job report looking for only part-time jobs. Just over 9 percent of the employed (and nearly 40 percent of employed job seekers) report only looking for an additional job, with no intention of leaving their current job.

Since empirical evidence on the incidence of on-the-job search is scarce and mostly comes from outdated surveys, it is hard to provide a good comparison of the pervasiveness of on-the-job search with other studies. Time use surveys that rely on time-diary data, such as the American Time Use Survey (ATUS), are likely to underestimate search activity since they are based on time use the day prior to the survey.[9] The ATUS reports that, on average, only around 0.6 percent of employed actively searched in the 2013-2015 period. The corresponding fraction is 16.5 percent for the unemployed, revealing the difficulty of comparing daily diary-based measures with traditional surveys.[10] Some older studies that relied on survey data found pervasive job search activity among the employed. For example, according to Black (1980), around 14 percent of

[7] We borrow this distinction from the well-established literature on labor supply. Mukoyama, Patterson, and Şahin (2018) also apply this distinction to the search effort of the unemployed.

[8] By both the "job search" definition and the BLS definition of unemployment, no one outside of the labor force is available for work.

[9] See a detailed discussion in Appendix A.2. In particular, Table A4 in the appendix provides a comprehensive comparison of the SCE and ATUS measures of job search activity.

[10] See also Mueller (2010) for similar evidence for an earlier period.

Table 2: Basic Job Search Statistics by Labor Force Status

	Employed	**Unemployed**	**Out of Labor Force**
Percent that actively searched for work	23.1	99.4	2.1
	(0.9)	(0.6)	(0.7)
Percent that actively searched and are	14.1	99.4	0.0
available for work	(0.7)	(0.6)	(0.0)
Percent reporting no active search or	6.1	0.3	6.0
availability, but would take job if offered	(0.5)	(0.4)	(1.1)
Percent applying to at least one vacancy in last	19.8	93.0	1.8
four weeks	(0.8)	(2.0)	(0.6)
Percent with positive time spent searching in	20.5	85.9	2.6
last seven days	(0.8)	(2.7)	(0.8)
Percent only searching for an	9.2	—	—
additional job	(0.6)		
Percent only seeking part-time work,	20.5	22.6	—
conditional on active search	(1.8)	(3.3)	
Percent only seeking similar work (to most	27.4	5.3	—
recent job), conditional on active search	(2.1)	(1.8)	
N	2,302	163	432

Notes: Estimates come from authors' tabulations from the October 2013-15 waves of the SCE Labor Supplement, for all individuals aged 18-64, by labor force status. Standard errors are in parentheses.

white workers and 10 percent of black workers reported on-the-job search in the 1972 interview of the Panel Study of Income Dynamics (PSID). Similarly, Blau and Robins (1990) report that employed search spells represent about 10 percent of all employment spells in the Employment Opportunity Pilot Projects (EOPP) in 1980. Unfortunately, the main source of labor market statistics for the U.S., the CPS, does not ask questions about job search to employed individuals, but its recent Computer and Internet Use Supplements asked all respondents, regardless of their labor force status, whether they used the internet to search for a job in the past *six months*. Around 28 percent of the employed reported using the internet for job search in the last six months in the 2015 survey. We also asked a question about whether an individual searched in the last *twelve months*. Around 45 percent employed reported searching in the last twelve months using any active search method, including online job search. Given that we designed our survey

Table 3: Intensive Margin: Search Effort by Labor Force Status

	Employed			Unemployed	Out of Labor Force
	Looking for Work	Not Looking	All		
Labor Force Status at Time of Survey					
Hours spent searching, last 7 days	4.35	0.05	1.18	8.46	0.07
	(0.30)	(0.01)	(0.09)	(0.75)	(0.04)
Mean applications sent, last 4 weeks	4.63	0	1.22	8.14	0.08
	(0.49)	(—)	(0.13)	(1.24)	(0.06)
N	508	1,520	2,028	163	432
Labor Force Status in Prior Month					
Mean applications sent			1.19	10.49	0.47
			(0.13)	(1.75)	(0.10)
Mean applications sent, ignoring applications to additional jobs			0.95	10.48	0.47
			(0.13)	(1.75)	(0.10)
N			2,053	117	453

Notes: Estimates come from authors' tabulations from the October 2013-15 waves of the SCE Labor Supplement, for all individuals aged 18-64, excluding the self-employed, by detailed labor force status. The top panel reports results by labor force status at the time of the survey, while the bottom panel reports the results by labor force status in the prior month. See the appendix for how prior month's labor force status is determined. Standard errors are in parentheses.

to cast a wide net to identify any "search activity," we find our estimates for the extensive margin of search reasonable.

Table 3 reports the amount of effort spent on the job search process, the *intensive margin* of job search. We categorize the employed by whether or not they actively look for work.[11] This distinction emphasizes the stark differences in search activity among the employed. We find that the unemployed send substantially more job applications and dedicate more hours to search than the other groups. They put in roughly twice as much effort as the employed that actively look for work. On average, unemployed workers spent around 8.5 hours *per week* on job search and sent 10.5 applications in the *last four weeks*. These findings are remarkably similar to the statistics reported by Barron and Gilley (1981), who use a special survey of the unemployed in the CPS from May 1976. They find that the typical unemployed individual contacted over three employers *per week* and spent approximately eight and two-thirds hours *per week* to make such contacts.

[11] The estimates exclude the self-employed.

3.2 The Distribution of Search Effort and its Relationship to Current Wages

Figure 1 shows the distributions of search time within the last *seven days* and the number of applications sent within the last *four weeks* for employed and unemployed job seekers conditional on searching in the last four weeks.[12] Around half of the employed and about one-third of the unemployed apply to either one or two jobs. About 10 percent of employed job seekers sent more than 10 applications, while just over 25 percent of the unemployed sent more than 10 applications. The right panel of the figure shows the distribution of search time. The differences between the employed and unemployed are more pronounced when we consider the distribution of search time. Around 40 percent of the employed report searching for one hour or less within the last seven days, but around 80 percent of the unemployed searched for longer than one hour. Moreover, searching for longer than 10 hours a week is relatively more common among the unemployed than the employed (around 35 percent vs. 15 percent). Interestingly, even among those people who reported searching in the last four weeks, 25 percent of the employed and 15 percent of the unemployed did not search at all within the last seven days. This observation highlights the intermittent nature of search effort and reinforces our view that ATUS, which is based on a time dairy reported at the daily frequency, greatly understates the extensive margin of job search.[13] Given that the employed are likely to search during work hours (which they would report in the ATUS as their main activity during that time), the bias is likely to be more pronounced for the employed.

We also examine the determinants of search effort while on the job. Dissatisfaction with pay and benefits appears to be the most important reason for search, with 63 percent of employed searchers indicating it as a reason for search.[14] This is consistent with the notion that workers move to more productive, better paid jobs through job-to-job transitions.

[12] Recall from Table 2 that around 23 percent of the employed report actively searching. The remainder is excluded from the analysis to provide a more relevant comparison of distributions.

[13] In Appendix Figure C.2, we report these distributions over a finer grid.

[14] Other important reasons include dissatisfaction with job duties (47 percent) and poor utilization of one's skills or experience (34 percent). Only 12 percent of the employed reported that they searched because they had been given advance notice or otherwise expected to lose their job.

Figure 1: Distribution of Number of Applications Sent in the Last Four Weeks (left panel) and Search Time in Hours in the Last Seven Days (right panel) by Labor Force Status

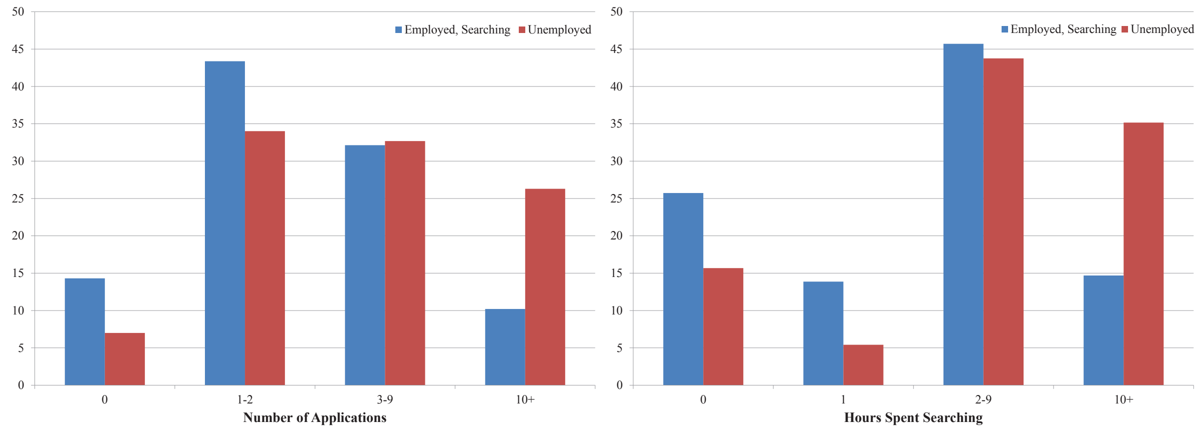

Note: Figure reports the histograms of the number of applications sent in the last four weeks (left panel) and the hours of time spent searching for work in the last seven days (right panel). Estimates are for all individuals, excluding the self-employed, who reported actively searching for work in the 2013-15 labor supplements of the SCE.

A key implication of related job ladder models is that workers near the bottom of the job ladder search harder for a better job while those near the top of the job ladder do not search as hard since their chances of obtaining an offer better than their current job are small. Our data allow for a direct test of this implication. Measurement error and unobserved worker heterogeneity in wages make it difficult to assess the exact position of a worker on the job ladder, but a worker's wage relative to her peers with the same observable characteristics such as age, gender, education, and occupation should still provide a useful proxy for her position on the ladder. Therefore, we estimate a linear regression model of the relationship between a worker's search behavior and her current wage controlling for observable worker characteristics. Our estimates are in Table 4 and show that workers with lower wages in their current job are more likely to engage in search regardless of the definition of search activity that we use. In addition, the overall intensity of search activity, measured by the total number of applications in the last four weeks or the total hours of searching the last seven days, is higher for workers with lower wages.[15]

We also explore potential non-linearities in the search-wage relationship. Figure 2 shows the

[15]In results available on request, we report the effect of various observables on incidence and intensity of search. Females, more educated workers, and workers who identify as black and Hispanic search harder.

Table 4: The Relationship between Search Effort and the Current Wage

	Incidence of Search		Search Effort	
	Active Search	Applied	Applications	Search Time
log current real wage	-0.086^{***}	-0.076^{***}	-0.550^{**}	-0.549^{***}
	(0.016)	(0.015)	(0.220)	(0.135)
Dependent variable mean	0.264	0.226	1.220	1.185
R^2	0.103	0.115	0.050	0.147
N	2,020	2,020	2,020	2,020

Notes: The table reports the estimated relationship from an OLS regression between the dependent variables listed in each column and the (log) real current wage. "Active Search" equals one if an individual actively looked for work in the last four weeks. "Applied" equals one if an individual applied to at least one job in the last four weeks. "Applications" refers to the number of applications sent in the last four weeks. "Search Time" refers to the number of hours spent looking for work in the last seven days. Regressions are sample weighted and control for gender, age, age squared, four education dummies, four race dummies, a homeownership dummy, marital status, marital status×male, the number of children aged 5 and younger, and fixed effects for state and year. Standard errors are in parentheses. *** represents significance at the 1 percent level. ** represents significance at the 5 percent level.

estimates from a locally weighted regression (LOWESS) between the different measures of search effort and the residualized current wage, i.e., the wage conditional on the controls from Table 4. The figure highlights the negative wage-search effort relationship in Table 4 for both the total number of applications and hours of search, and illustrates the quantitatively large decline in search effort from low to high residual wages.[16] These plots strongly confirm the implications of models with endogeneous search effort such as Christensen et al. (2005). Their analysis implied that workers search harder when earning a relatively lower wage, but they lacked direct evidence on job search effort. Our analysis provides this direct evidence. We will return to this issue in the model section below.

3.3 Search Outcomes by Labor Force Status

We have shown that there is considerable job search activity among the employed. We now move on to show how search effort translates into employer contacts, job offers, and new job matches for the employed and unemployed. The fact that our data contain exhaustive information on both search effort and search outcomes at different stages of the process puts us in a unique

[16] Figures C2 and C3 in the Appendix show that similar patterns hold for measures of the incidence of search as well as when we do not control for observable characteristics in the wage.

Figure 2: Job Search Effort by the Current Wage

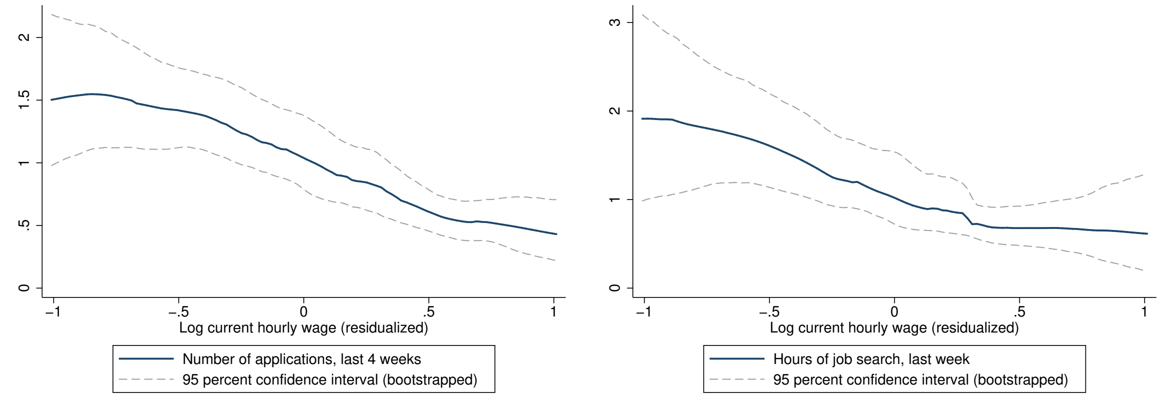

Note: Figure reports the LOWESS estimates (with smoothing parameter 0.8) of the relationship between the measures of search effort listed on each vertical axis and the (log) real current wage of the employed, residualized after controlling for observable worker characteristics (see Table 4 for the list of specific variables). Dashed lines represent 95 percent confidence intervals. The confidence intervals are based on a bootstrap with 500 replications. The estimates use all employed individuals, excluding the self-employed, age 18-64 from the October 2013-15 waves of the SCE Labor Supplement.

position to assess the relative effectiveness of employed versus unemployed search.[17]

The top panel of Table 5 reports search outcomes by labor force status at the time of the survey and shows that those who are employed and looking for work receive the greatest number of employer contacts, interviews, and offers despite the fact that their search effort is about half that of the unemployed. They also receive the most unsolicited employer contacts. These are employer contacts that did not result from a job seeker's search efforts. Overall, those searching on the job receive about 48 percent more contacts and 14 percent more job offers than the unemployed, again, despite exerting about half as much search effort. Those who are employed but not looking for work receive about one-quarter as many contacts and offers as the unemployed despite exerting no search effort. They receive about one-fifth of the offers of those searching on the job.

A potential concern with our estimates is that the outcomes are based on retrospective questions and the respondents' current labor force status may not reflect their labor force status at

[17]This is a long-standing question in the labor-search literature and has important implications for on-the-job search models as we discuss in Section 4. Earlier empirical contributions include Holzer (1987) and Blau and Robins (1990).

Table 5: Search Outcomes by Labor Force Status

	Employed				**Out of**
	Looking for Work	**Not Looking**	**All**	**Unemployed**	**Labor Force**
Labor Force Status at Time of Survey					
Mean contacts received	1.874	0.337	0.742	1.261	0.118
	(0.281)	(0.038)	(0.079)	(0.232)	(0.033)
Mean unsolicited contacts	0.783	0.298	0.426	0.459	0.099
	(0.124)	(0.032)	(0.040)	(0.154)	(0.030)
Mean job interviews (2014-15)	0.460	0.005	0.115	0.354	0.022
	(0.045)	(0.002)	(0.012)	(0.107)	(0.018)
Mean offers	0.425	0.086	0.175	0.373	0.079
	(0.039)	(0.011)	(0.014)	(0.078)	(0.026)
Mean unsolicited offers	0.047	0.046	0.046	0.043	0.053
	(0.01)	(0.009)	(0.007)	(0.016)	(0.023)
Fraction with at least one offer	0.299	0.057	0.118	0.220	0.041
	(0.02)	(0.007)	(0.007)	(0.033)	(0.010)
Fraction with at least one unsolicited	0.041	0.028	0.031	0.043	0.026
offer	(0.009)	(0.005)	(0.004)	(0.016)	(0.008)
Fraction with at least one offer,	0.345	0.086	0.155	0.237	0.059
including unrealized offers	(0.021)	(0.007)	(0.008)	(0.033)	(0.011)
N	508	1,520	2,028	163	432
Labor Force Status in Prior Month					
Fraction with at least one offer			0.105	0.339	0.074
			(0.007)	(0.044)	(0.012)
Fraction with at least one unsolicited			0.030	0.031	0.036
offer			(0.004)	(0.016)	(0.009)
Fraction with at least one offer,			0.144	0.349	0.085
including unrealized offers			(0.008)	(0.044)	(0.013)
Labor Force Status in Prior Month, Ignoring Search Outcomes for Additional Jobs					
Fraction with at least one offer			0.091	0.339	0.074
			(0.006)	(0.044)	(0.012)
Fraction with at least one unsolicited			0.029	0.031	0.036
offer			(0.004)	(0.016)	(0.009)
Fraction with at least one offer,			0.132	0.349	0.085
including unrealized offers			(0.007)	(0.044)	(0.013)
N			2,053	117	453

Note: Estimates come from authors' tabulations from the October 2013-15 waves of the SCE Labor Supplement, for all individuals aged 18-64, excluding the self-employed, by labor force status. The top panel reports results by labor force status at the time of the survey, while the middle and bottom panels report the results by labor force status in the prior month. See the appendix for how prior month's labor force status is determined. Standard errors are in parentheses.

the time of the outcome. Non-random job acceptances by those unemployed at the time of a job offer can create a selection issue.[18] Since the focus is job search behavior by labor force status, we address this issue by constructing a measure of labor force status for the prior month using a wide range of survey questions from the SCE labor supplement.[19] In the middle panel of Table 5, we report offer outcomes by prior labor force status. The results show that the fraction with at least one offer over the last four weeks decreases slightly for the employed, from 11.8 percent to 10.5 percent, but increases substantially for the unemployed, from 22.0 percent to 33.9 percent, when considering labor force status in the "prior month" instead of "at time of survey." This is in line with the expectation that some individuals who are unemployed at the time of the job offer started working by the time of the survey. Another concern is that a substantial fraction of employed workers are only seeking an additional job. Job-to-job transitions, as measured in the CPS and most other household surveys, only capture changes in an individual's main job. Nearly all models of labor market search only consider this type of job-to-job transition as well. In the bottom panel of Table 5, we report offer outcomes ignoring the offers of those who reported only looking for additional work. The fraction of the employed receiving at least one offer falls to 9.1 percent in this case. We use this estimate of the offer rate in our model calibration below.

It is possible that some individuals simply do not pursue offers that they are likely to reject. In this case, the job offers we observe in the data would be *censored*. Most importantly, this type of censoring could be correlated with employment status. To address this issue, our survey asks respondents whether a potential employer indicated that they would be willing to make an offer but the respondent indicated that he or she was not interested. We label these offers as *unrealized* rejected offers as respondents rejected these offers even before a formal offer was made. We indeed find that these unrealized offers are more common for the employed. Among those who did not

[18]This selection issue is similar to the time-aggregation issue that plagues calculations of the separation rate using CPS data.

[19]In Appendix B, we detail our methodology for determining labor force status in the prior month and evaluate our measure along several dimensions. Unfortunately, the survey does not allow us to directly identify whether a respondent was searching on-the-job in the prior month. We also consider an alternative measure of labor force status in the prior month from the monthly SCE panel data and show that we obtain very similar results.

Table 6: Acceptance Decisions by Labor Force Status in Previous Month

	Employed	**Unemployed**	**Out of Labor Force**
Percent of best offers accepted	31.6	53.2	19.7
	(3.3)	(8.3)	(6.9)
Percent of all offers accepted	27.0	48.3	17.5
	(3.0)	(7.9)	(6.4)
Percent of best offers accepted,	30.0	53.2	19.7
ignoring offers for an additional job	(3.5)	(8.3)	(6.9)
Percent of all offers accepted,	26.4	48.3	17.5
ignoring offers for an additional job	(3.3)	(7.9)	(6.4)
N	196	37	34

Note: Estimates come from authors' tabulations from the October 2013-15 waves of the SCE Labor Supplement, for all individuals aged 18-64, excluding the self-employed, by labor force status in the prior month. See the appendix for how prior month's labor force status is determined. Standard errors are in parentheses. The additional work distinction only applies to the employed.

report a formal offer over the last four weeks, about 4 percent of the employed indicated that they rejected such an unrealized offer, compared to only 1 percent of the unemployed. The bottom panel of Table 5 reports the fraction of individuals who received at least one offer, including these unrealized offers. Accounting for unrealized offers raises the fraction receiving a job offer to 13.2 percent for the employed and to 34.9 percent for the unemployed.

Table 6 reports the acceptance rate for offers received within the last four weeks by labor force status in the prior month. The results show that the unemployed are much more likely to accept a given offer, with 48 to 53 percent of their offers accepted, depending on whether we include all offers or just the best offer in the denominator of the acceptance rate, compared to the employed, who accept 26 to 32 percent of their offers, depending on the measure used.[20] In our calibration, we focus on the acceptance rates of the best offers, excluding offers for additional work, which is 30.0 percent for the employed and 53.2 percent for the unemployed.

Finally, we present the distribution of search effort and search outcomes across the different labor force categories. Examining these distributions provides another way of assessing the

[20]The survey only asks respondents if they accepted their best offer. The acceptance rate including all offers assumes that none of the remaining offers were accepted.

Table 7: Distribution of Search Effort and Outcomes by Labor Force Status

	Employed			**Unemployed**	**Out of Labor Force**
	Looking for Work	**Not Looking**	**All**		
Pct. of population	19.4	54.2	73.6	7.3	19.1
Job Search over Last Four Weeks					
Pct. of total applications	59.5	0.0	59.5	39.5	1.0
Pct. of contacts received	55.0	27.6	82.6	14.0	3.4
Pct. of unsolicited contacts	41.5	44.2	85.7	9.2	5.1
Pct. of interviews (2014-15 only)	72.7	2.4	75.1	21.0	3.9
Pct. of offers received	48.1	27.1	75.2	16.0	8.8
Pct. of unsolicited offers received	19.1	52.8	71.9	6.7	21.5

Note: Estimates come from authors' tabulations from the October 2013-15 waves of the SCE Labor Supplement, for all individuals aged 18-64, excluding the self-employed, by labor force status at the time of the survey.

relative efficiency of employed and unemployed job seekers. Table 7 reports the distribution of respondents, job applications, and job search outcomes by labor force status. The unemployed make up just over 7 percent of our sample, but account for nearly 40 percent of all job applications sent. At the same time, they only receive 16 percent of all offers made. In stark contrast, the employed who report not looking for work send no applications by construction but account for around 28 percent of all employer contacts and receive over 27 percent of all job offers. This is due, in part, to the fact that they also account for 44 percent of all unsolicited employer contacts and 53 percent of all unsolicited offers. Those actively searching on the job account for another 48 percent of all job offers. Thus, the job search behavior of the unemployed can be characterized by high effort, but relatively low returns in terms of employer contacts and job offers. The employed, on the other hand, do fairly well regardless of whether they are actually looking for work. Though the unemployed are seemingly less effective in their job search efforts, they are also more likely to accept the offers that they do receive.

3.4 Characteristics of Job Offers and Accepted Jobs

The employed are more effective at generating job offers, but our evidence thus far is silent on whether the employed receive *better* offers than the unemployed. We now examine how the job offers themselves, including all offers and the subset of those that are accepted, differ by employment status. Our survey asks individuals about any offers they received in the last four weeks. For those who received no offer within the last four weeks, it probes further to elicit information on any offers received within the last six months. The survey also elicits the respondent's labor force status at the time of the job offer. It asks a variety of questions about the characteristics of the job offer, including information about the search and bargaining process. It also asks if the offer was accepted (and if it represents their current job).

Table 8 presents the characteristics of best job offers received within the last six months by labor force status (employed vs. non-employed) at the time of the job offer.[21] First, note that over 70 percent of job offers in our sample go to those who were employed at the time of the offer. The results consistently show that the employed receive much better job offers than the non-employed. Unconditionally, the employed receive wage offers that are about 40 log points (49 percent) higher than the wage offers of the non-employed.[22] Even after conditioning on the observable characteristics of the worker and the job offer, the employed enjoy wage offers that are 25 log points (28 percent) higher than the wage offers of the non-employed.[23] The left panel of Figure 3 shows that, even after accounting for these controls, the distribution of wage offers for the employed stochastically dominates the distribution of wage offers for the non-employed.

The remainder of Table 8 shows that job offers received by the employed are superior on other margins as well. Their hours are only 8 log points higher (which is not statistically significant), but they are 20 percentage points more likely to include at least some benefits such as retirement

[21]Starting in 2014, we added a question to the survey that identifies those who searched prior to the receipt of the job offer. Most of the non-employed report actively searching, and in unreported results, we find that the residual wage offer differential that we document is even larger if we restrict the non-employed to those who were searching prior to the job offer.

[22]The offer wage, as well as all other wages in our analysis, refers to the real hourly wage. Respondents report their nominal earnings as an hourly wage, or as a measure of weekly or annual earnings. In the latter cases, we measure the wage as earnings per hour, based on the reported usual hours worked. We convert all wages used into real terms using the Consumer Price Index (CPI).

[23]Our conditional estimates of the offered wage and the subsequent accepted wage control for worker and job characteristics, as well as state and year fixed effects. Our worker controls include sex, age, age squared, marital status, marital status$\times$sex, education, race, homeowner status, and number of household children. Our firm and job controls are the two-digit occupation of the job and the size of the offering firm. We report estimates of the other job offer characteristics that control for observable characteristics in the appendix.

Figure 3: Distribution of All Wage Offers (left panel) and Accepted Wage Offers (right panel)

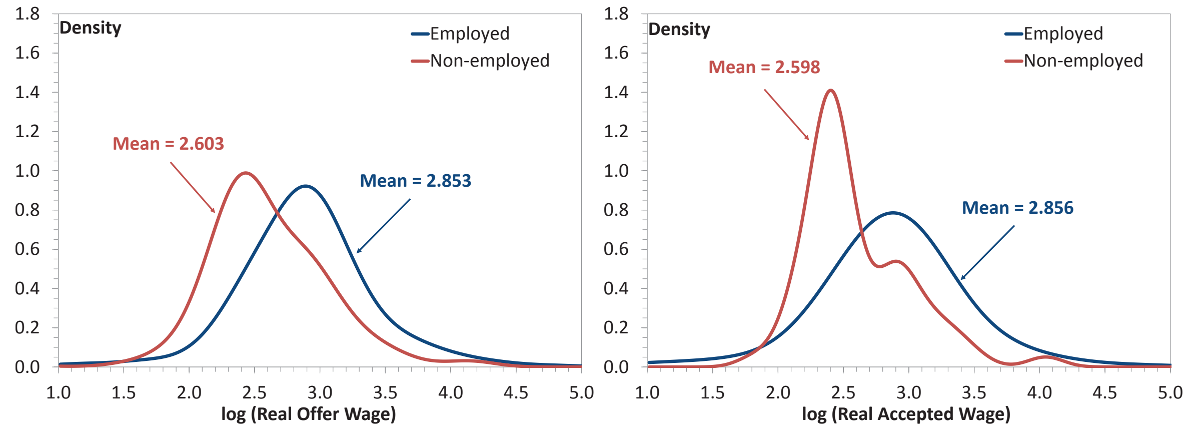

Note: Figures report kernel density estimates of residual log(real wage offer) by labor force status after controlling for observable worker and job characteristics. Estimates are for all (best) job offers received within the last six months by individuals in the October 2013-15 waves of the SCE Labor Supplement.

pay or health insurance. The employed are nearly twice as likely to have received their offer through an unsolicited contact. The employed and non-employed are roughly equally likely to have had a "good idea" of what the job paid prior to receiving the offer. Potentially contributing to the differences in offer wages between the two groups, the employed are significantly more likely to bargain over their offers, with 39 percent of their offers involving some bargaining, compared to 24 percent for the non-employed.[24] Counter-offers by the current employer, defined as anything from matching the outside offer to offering a promotion, pay raise, or some added job benefit, occurred for about 14 percent of the employed who received an offer from an outside firm.

Despite their relatively poor job offers, the non-employed are about one-and-a-half times more likely than the employed to accept a job offer, with 55 percent of offers accepted by the non-employed versus 35 percent by the employed. The acceptance rates are very close to those we obtain using the prior month's labor force status. Table 8 also suggests that a primary reason that the non-employed are more likely to accept their relatively poor job offers is a perceived lack of alternative options. About 27 percent of the non-employed cite a lack of other alternatives as the main reason for accepting an offer, while only 7 percent of the employed cite that as their

[24]These estimates are consistent with Hall and Krueger (2012), who find that around a third of all workers engaged in some bargaining over their pay with their current employer.

primary reason. The right panel of Figure 3 shows that, even after controlling for observed worker and job characteristics, the accepted wage distribution of the employed stochastically dominates the accepted wage distribution of the non-employed.

We can also examine job search retrospectively for those employed at the time of the survey interview by asking them how they came about their current jobs. The advantage of this approach is that we are able to examine their starting wages and previous earnings as a function of their labor force status. This provides us with additional guidance for our model in Section 4.

Table 9 presents the characteristics of the current and previous job by labor force status at the time of hire. We focus on the comparison of the non-employed to those who move directly from employment to their current job. At the time of the survey interview, those hired from non-employment are paid lower wages, have fewer work hours, and are much less likely to have any benefits than those hired while already employed. They are also somewhat more likely to be looking for new work at the time of the survey. Estimates reported in the middle of Table 9 show that most of the wage differences between those hired from employment and those hired from non-employment stem from wage differences at their time of hiring. The real starting wage of those hired from non-employment is 28 log points (32 percent) lower than the real starting wage of those hired from employment, on average. Conditioning on the observable characteristics of the worker and the job reduces the wage difference by nearly half, to about 14 log points (17 percent).[25] Despite the large differences in the wage and hours of the current job across the two labor force categories, the differences in their previous jobs' wages are small and statistically insignificant. This is true for both the unconditional real wage and the wage that controls for observable worker and job characteristics. Note that the smaller difference in the premium in starting wages compared to the difference in offered wages is likely due to a selection issue: poor job offers are less present in the cross-section of current jobs, as individuals accepting these jobs

[25] Our conditional estimates of the starting wage and previous wage use the same worker characteristics as our conditional estimates of the offer wage. The starting wage uses the same firm and job controls, and additionally controls for two-digit industry. The previous wage only includes the two-digit occupation as a job or firm control. We report estimates of the other characteristics of the current job that control for observable characteristics in the appendix.

Table 8: Characteristics of Best Job Offer by Labor Force Status at Time of Offer

	Employed at Offer	**Non-Employed at Offer**	**Difference, E - NE**
Percent of job offers	70.5	29.5	
Offer Wage Estimates			
log real offer wage,	2.893	2.496	0.397
unconditional	(0.041)	(0.057)	(0.075)
Controlling for observable	2.853	2.603	0.250
characteristics	(0.034)	(0.036)	(0.075)
Additional Job Offer Characteristics			
log offer usual hours	3.409	3.333	0.076
	(0.032)	(0.040)	(0.056)
Pct. of offers with no benefits	40.5	60.6	-20.1
	(2.2)	(3.8)	(4.2)
Pct. of offers through an unsolicited	26.2	14.4	11.8
contact	(2.0)	(2.7)	(3.6)
Pct. of respondents with at least a good	54.6	58.8	-4.2
idea of pay	(2.3)	(3.8)	(4.2)
Pct. of offers with some counter-offer	14.2	—	—
given	(1.6)		
Pct. of offers that involved bargaining	38.8	24.4	14.4
	(2.2)	(3.4)	(3.8)
Pct. of (best) job offers accepted	34.6	54.7	-20.1
	(2.2)	(3.9)	(4.1)
Pct. of offers accepted as only option,	6.9	27.3	-20.4
conditional on acceptance	(1.9)	(4.8)	(5.2)
N	489	165	

Note: Estimates come from authors' tabulations from the October 2013-15 waves of the SCE Labor Supplement, for individuals aged 18-64, excluding the self-employed, with at least one job offer in the last six months. Observable characteristics controlled for in the conditional wage estimates include fixed effects for survey year and state as well as a vector of demographic controls: sex, age, age squared, four education categories, four race categories, a dummy for homeownership, the number of children under age 6 in the household, marital status, and marital status×sex. They also include the two-digit SOC occupation of the job and six categories of the firm size of the potential employer. Standard errors are in parentheses.

Table 9: Characteristics of Current and Previous Job, by Labor Force Status at Time of Hire

	Hired from Employment	**Hired from Non-Employment**	**Difference, E - NE**
Share of Employment	69.1	30.9	
Characteristics of Current Job			
log real current wage	3.126	2.841	0.285
	(0.018)	(0.029)	(0.033)
log usual hours	3.679	3.546	0.132
	(0.01)	(0.019)	(0.020)
Median tenure (mos.)	58.0	41.0	17.0
	(2.8)	(3.6)	(4.6)
Pct. with no benefits	16.0	30.9	-16.8
	(1.0)	(2.0)	(2.1)
Percent actively searched	25.4	32.1	-6.6
for work, last four weeks	(1.2)	(2.0)	(2.3)
Starting Wage Estimates			
log real starting wage,	2.938	2.663	0.275
unconditional	(0.018)	(0.028)	(0.033)
Controlling for observable	2.896	2.758	0.139
characteristics	(0.014)	(0.018)	(0.027)
Previous Wage Estimates			
log real previous wage,	2.859	2.821	0.038
unconditional	(0.024)	(0.037)	(0.045)
Controlling for observable	2.834	2.888	-0.054
characteristics	(0.019)	(0.031)	(0.041)
N	1,238	525	

Note: Estimates come from authors' tabulations from October 2013-15 waves of the SCE Labor Supplement, restricted to currently employed individuals aged 18-64, excluding the self-employed, with a reported labor force status at the time of hire and reported current, starting, and previous-job wages and hours. Observable characteristics controlled for in the conditional wage estimates include fixed effects for survey year and state as well as a vector of demographic controls: sex, age, age squared, four education categories, four race categories, a dummy for homeownership, the number of children under age 6 in the household, marital status, and marital status×sex. They also include the two-digit SOC occupation of the current job, as well as the two-digit NAICS industry and six categories of firm size for the current employer. Standard errors are in parentheses.

are more likely to quit and to move to better-paying jobs. If the non-employed get worse offers than the employed, this explains why the gap is smaller among starting wages compared to job offers.[26]

3.5 Accounting for Differences in Job Offers

We can dig deeper into the wage offer difference between the employed and non-employed using responses to a rich set of questions from our survey. Table 8 shows that observable worker and job characteristics explain 37 percent of the wage offer difference. The remaining difference may arise simply because we cannot control for differences that are observed by employers but are unobserved in our data. For example, workers may differ in soft skills such as communication or time-management skills. Those with better soft skills would be more likely to be employed and earn a higher wage. This creates a selection effect that naturally generates a wage gap between the job offers received by the employed versus the unemployed. At the same time, an individual's prior work history provides a useful proxy for this unobserved heterogeneity because it reflects repeated labor market outcomes that were determined at least partly by their unobserved skills. Our survey has detailed questions that allow us to control for an individual's labor force history over the previous five years. As Table 10 shows, controlling for the fraction of the last five years that an individual was employed reduces the wage offer gap from 0.250 to 0.222. When we additionally control for share of the last five years spent unemployed and the share spent as a student, the difference goes down further to 0.205, implying that labor force history can account for an additional 11 percent of the wage offer gap.

While observable worker and job characteristics and labor force history explain half of the wage offer differential, there remains a notable difference between the wage offers of the employed and non-employed workers. Prior wages of workers could also provide additional information regarding workers' unobserved skills. We add the *prior* wage of workers (i.e., the wage of the

[26]Figure C4 in the appendix shows the distribution of starting wages relative to workers' prior wages with and without controlling for observables. Even after conditioning out our controls, those who transition directly from employment receive a 8 log point increase in their wage, on average, while those who were non-employed receive a 13 log point decrease in their wage, on average.

previous job for the employed and the wage of the most recent job for the non-employed) as an additional control in Table 10. We find that the prior wage does not close the gap. On the contrary, the gap widens, an issue that we discuss below in the model section.

Finally, the remaining gap may arise because of differences in the job search process between the employed and non-employed. It is possible for employed workers to have better access to more rewarding job search channels (see, for example, Arbex, O'Dey, and Wiczer, 2016). For example, our empirical analysis shows that the employed are more likely to have received an offer through an unsolicited contact than the non-employed. If these informal offers represent higher-quality jobs, then the higher incidence of unsolicited offers should also contribute to the wage offer gap. Alternatively, non-employed workers may be more likely to pursue jobs with lower wages but better non-wage benefits. In the last row of Table 10, we control for how a job offer came about using dummies for whether the offer was the result of a direct contact by the worker, whether an intermediary (such as an employment agency) was involved, whether it was the result of a referral, or whether the offer was unsolicited. We also control for the (log) hours of the job offer and the incidence of any benefits (categorized into health, retirement, or other benefits). These controls result in little change in our estimate of wage differentials.[27] Thus, while controlling for observable worker and job characteristics, prior labor force history, and the source of the job offer reduces the offered wage gap by about half, a substantial gap between the wages offered to the employed and non-employed remains.

In the next section, we set up a partial equilibrium model of on-the-job search and confront the model with our key empirical findings. Specifically, the model addresses the following facts: (i) on-the-job search is pervasive among the employed; (ii) the intensity of on-the-job search declines with the job seeker's current wage; (iii) employed job seekers search less than the unemployed, but they receive just as many offers; (iv) the employed receive better offers with higher wages and benefits, even after controlling for observable characteristics; (v) despite receiving higher-quality offers, the employed are less likely to accept them.

[27] If we re-estimate the last row of Table 10 excluding the prior wage, we obtain a wage offer gap of 0.193.

Table 10: Offer Wage Gap Estimates, Additional Controls.

Offer Wage Gap Estimates	**E-NE**
log real offer wage, unconditional	0.397
	(0.075)
log real offer wage, controlling for observables	0.250
	(0.075)
log real offer wage, controlling for observable characteristics	0.222
and employment history	(0.087)
Controlling for observable characteristics	0.205
and labor force history	(0.087)
Controlling for observable characteristics, labor force history,	0.288
and prior wage	(0.081)
Controlling for observable characteristics, labor force history,	0.267
prior wage, hours, benefits, and how offer came about	(0.082)

Note: Estimates come from authors' tabulations from the October 2013-15 waves of the SCE Labor Supplement, for individuals aged 18-64, excluding the self-employed, with at least one job offer in the last six months. See note to Table 8 for the included observable characteristics. Employment history controls for the fraction of the prior five years spent employed. Labor force history additionally controls for the fraction of the last five years spent unemployed or in school. The (log) real prior wage is the wage of the previous job for the employed and the most recent job for the non-employed. Additional job controls include (log) hours and dummies for the incidence of health, retirement, or other benefits. Controls for how the job offer came about include dummies for whether it was through a direct employer contact, an intermediary, a referral, or an unsolicited contact. Standard errors are in parentheses.

4 Job Search On and Off the Job: A Theoretical Framework

In this section, we set up a partial equilibrium model of on-the-job search and parameterize it carefully by matching a number of key moments in our data. The theoretical exercise has two main goals. The first is to gauge the relative contribution of various sources for the employed-unemployed wage offer gap observed in the data, and in particular to determine the role of unobserved heterogeneity and censored wage offers. While our empirical estimates attribute more than half of the wage offer gap to labor force status, it is possible that this may be partly due to unobserved differences between the employed and unemployed or due to employed job seekers directing their search towards more acceptable wage offers. The second goal is to assess additional micro and macro implications of our findings.

Our model draws upon the model with endogenous search effort of Christensen et al. (2005), but with four key extensions. First, our model allows for differences in search costs and search efficiency between the employed and unemployed. Second, our model allows for exogenous differences in job offer distributions between the employed and unemployed. Third, wage offers in our model may be censored to account for the possibility that employed job seekers are more selective in the pre-offer stage of the search process. Finally, we allow for *ex ante* heterogeneity in worker productivity. The latter three extensions allow the model to generate a differential in wage offers between the employed and unemployed, which is not present in Christensen et al.'s model.

4.1 Framework

Time is discrete and an individual receives a job offer with probability $\lambda_i(s) = \alpha_i + \beta_i s$, where $s \in [0, \frac{1-\alpha_i}{\beta_i}]$ is the endogenously-chosen level of search effort.[28] The constant α_i reflects the possibility of unsolicited offers, which occur absent of any search effort, and β_i reflects search efficiency. The subscript $i \in \{e, u\}$ captures differences in unsolicited offer arrival rates and search efficiency by employment status. These will generally lead to differential job-offer arrival rates as well. Search effort has an increasing, convex cost that may also vary by employment status, $c_i(s)$, with $c_i', c_i'' > 0$ and $c_i(0) = c_i'(0) = 0$. Existing matches end exogenously at a rate δ, and the discount rate is r.

We introduce *ex ante* heterogeneity in worker productivity, x, into our model. This reflects the heterogeneity that remains unobserved in the data and potentially contributes to the observed wage offer gap between the employed and unemployed. The *ex ante* heterogeneity also affects a worker's cost of search effort and her job separation rate.[29] Under this specification, w reflects the piece rate wage per unit of worker productivity.

[28] Bagger and Lentz (2017) make the same assumption on the functional form of search technology.

[29] We could also extend the model to allow worker heterogeneity to affect the offer arrival rate, but our evidence in Appendix Table C.2 shows that controlling for observable characteristics, including the previous employment history of the worker, do little to affect the likelihood of receiving a job offer by labor force status.

Given this setup, the Bellman equation for the employed is

$$\begin{aligned} W(x,w) &= \max_{\bar{s}_e \geq s \geq 0} \Bigg\{ wx - c_e(x,s) + \frac{1-\delta(x)}{1+r} \Bigg[W(x,w) + \\ & \qquad \lambda_e(s) \int_w^{\bar{w}} [W(x,y) - W(x,w)] dF_e(y|x) \Bigg] + \frac{\delta(x)U(x)}{1+r} \Bigg\}, \end{aligned} \tag{1}$$

where $\bar{s}_e = \frac{1-\alpha_e}{\beta_e}$. The first term on the right-hand side reflects the wage net of search costs. The second term on the right-hand side reflects the continuation value of the job, accounting for the potential separation to either a new job or unemployment. The last term reflects the expected value of a separation to unemployment. As Christensen et al. (2005) show, the value of employment is increasing in the wage. Consequently, optimal search effort will vary with the wage. We assume that search costs are proportional to unobserved productivity, $c_e(x,s) = xc_e(s)$.[30] While this assumption is not essential for our results here, it is consistent with the finding that job-finding rates differ little by skill group (Mincer, 1991, Elsby, Hobijn and Şahin, 2010) or prior wages (Mueller, 2017), and the evidence in Appendix Tables C1 and C2, which shows that observable characteristics have a very limited, if any, effect on job search effort and outcomes by labor force status. The first order condition for an employed individual's search effort, $s_e(x,w)$, is

$$xc_e'(s_e(x,w)) \leq \beta_e \frac{1-\delta(x)}{1+r} \int_w^{\bar{w}} [W(x,y) - W(x,w)] dF_e(y|x), \tag{2}$$

which holds with equality if the optimal search effort is below $\bar{s}$. Since the cost of search effort is increasing and convex, search effort will decline with the wage. Note that, since $s_e(x,w)$ declines with w, search effort declines as workers move up the job ladder.

The Bellman equation for the unemployed is similar in structure. While unemployed, individuals of type x receive a flow utility of unemployment, $b(x)$. Consequently, an unemployed job

[30] See, for example, Cahuc, Postel-Vinay and Robin (2006), and Hall and Mueller (2017), for similar assumptions.

seeker solves

$$U(x) = \max_{\bar{s}_u \geq s \geq 0, R} \left\{ b(x) - c_u(x,s) + \frac{1}{1+r} \left[U(x) + \lambda_u(s) \int_R^{\bar{w}} [W(x,y) - U(x)] dF_u(y|x) \right] \right\}, \tag{3}$$

where $\bar{s}_u = \frac{1-\alpha_u}{\beta_u}$. The first term on the right-hand side reflects the flow value net of search costs and the second term on the right-hand side reflects the continuation value of unemployment, accounting for the probability of finding a job. Similar to our assumption for the employed, we posit that the flow value of unemployment (net of search costs) is proportional to unobserved productivity, i.e. $b(x) - c_u(x,s) = x(b - c_u(s))$.[31] The offered piece-wage rate per unit of worker-level productivity is assumed to be identically distributed across worker types x, i.e., $dF_i(w|x) = dF_i(w)$.

The unemployed of type x will have an optimal reservation wage, $R(x)$, that solves $W(x, R(x)) = U(x)$ and represents the wage at which the unemployed are just indifferent between a job that pays $R(x)$ and unemployment. It is useful to illustrate the first order condition for the unemployed worker's search effort decision:

$$xc_u'(s_u(x)) \leq \frac{\beta_u}{1+r} \int_R^{\bar{w}} [W(x,y) - U(x)] dF_u(y|x). \tag{4}$$

The first order condition differs from (2) in three ways: (i) the search cost function, (ii) the search efficiency parameter, and (iii) the discounted expected gain from accepting a job, which is determined by the reservation wage R and the shape of the wage offer distribution for the unemployed.

Finally, we address our finding that the employed appear to disproportionately reject offers before they are made relative to the unemployed. We model these unrealized offers by assuming that job seekers observe the terms of the offer prior to receiving the formal offer with probability χ_i, and do not pursue the offer further (i.e., reject) if the wage is below the reservation wage.[32] As

[31]This assumption is also consistent with the evidence noted earlier. For example, if we assumed that $b(x) = b$ and $c_i(x,s) = c_i(s)$, then high-x workers would search much more intensively and find many more jobs.

[32]See the appendix of Hall and Mueller (2017) who make the same assumption.

documented in the empirical section, employed workers appear to reject a non-negligible fraction of offers before they are even made. For a worker with productivity x and reservation wage R, one can thus write

$$\tilde{\lambda}_i(x, R) = \lambda_i(s_i(x, R))(\chi_i(1 - F_i(R)) + 1 - \chi_i), \tag{5}$$

$$\tilde{A}_i(x, R) = \frac{1 - F_i(R)}{\chi_i(1 - F_i(R)) + 1 - \chi_i}, \tag{6}$$

where $\lambda_i(s_i(x, R))$ is the probability of receiving an offer, including unrealized offers, $\tilde{\lambda}_i(x, R)$ is the probability of receiving a formal offer, $\tilde{A}_i(x, R)$ is the observed acceptance rate and $1 - F_i(R)$ is the likelihood a potential offer is above the reservation wage threshold.

4.2 Parameterization and Targeted Moments

We calibrate several parameters to match moments from other data sources. We set the frequency of the model to be monthly, with the monthly discount rate matching an annual interest rate of 4 percent. The average monthly job separation rate is set to be 0.015, which matches the average employment-to-unemployment flow rate in the CPS and in the monthly SCE in recent years. We assume that the wage offer distribution is log normal, normalize the mean of the log offered wages to zero, and calibrate the standard deviation of the wage offer distribution to be 0.24 as in Hall and Mueller (2017). This estimate is close to other estimates of frictional wage dispersion, see, e.g., Low, Meghir and Pistaferri (2010) and Tjaden and Wellschmied (2014). We choose this estimate of wage dispersion over one derived from the SCE data because of the relatively small sample of wage offers that we observe in the SCE data. We assume a search cost function of the form $c(s) = ks^{1+(1/\gamma)}$ and follow Christensen et al. (2005) by setting the elasticity parameter γ to a value of 1.19.

We carefully choose the ten remaining parameters of the model by exactly matching them to ten key moments derived from our own survey, as shown in Tables 11 and 12. Recall that we assume an offer arrival probability of the form $\lambda_i(s) = \alpha_i + \beta_i s$ and thus β_i is a key determinant of the returns to search. We normalize the search effort of the unemployed to unity and calculate

Table 11: Targeted Moments

	Data	Model
Search effort of unemployed (normalized)	1	1
Search effort of employed	0.091	0.091
Unsolicited offer rate of unemployed	0.029	0.029
Unsolicited offer rate of employed	0.031	0.031
Formal offer rate of unemployed	0.339	0.339
Formal offer rate of employed	0.091	0.091
Unrealized offer rate of unemployed	0.010	0.010
Unrealized offer rate of employed	0.041	0.041
Acceptance rate of unemployed	0.532	0.532
Residual offered wage differential (E - U)	0.25	0.25
St. dev. of log residual offered wages	0.67	0.67
Residual prior wage differential (E - U)	-0.10	-0.10

the search effort of the employed relative to the normalization. We set the α_i's to match the unsolicited offer rates of the employed and unemployed. We set the β_i's to match the observed job offer arrival rates for all offers (including unsolicited offers) for the employed and unemployed in the data given their relative search effort.[33] The model predicts that the employed have a search efficiency parameter of 0.93, while the unemployed have a search efficiency parameter of 0.32, implying that the employed are almost three times more efficient per unit of search effort.

We assume that the cost function differs for the employed and unemployed by the scaling parameter κ_i and set the κ_i's to match the normalized search efforts of the employed and unemployed. The high level of search efficiency and low level of search effort among the employed implies a very high cost of search for the employed in our calibration. We set the parameters χ_i to match the unrealized offer rates for both the employed and unemployed. The model implies that the censoring of the wage offer distribution due to unrealized, rejected offers is negligible for the unemployed but substantial for the employed, as indicated by the values for χ_i in Table 12. About 6 percent of offers are unrealized for the unemployed but 39 percent of offers are unrealized for the employed.

We set b to match the average acceptance rate of the unemployed. This allows our model, by

[33]We set the offer arrival rates equal to the probability of receiving at least one offer over the course of the last four weeks, ignoring offers for additional jobs, i.e., jobs where a worker does not leave her current employer. We measure search effort as the average number of applications sent over the last four weeks, again ignoring search for additional work.

assumption, to match their job-finding rate, as both the acceptance rate and the offer rate of the unemployed are targets in the calibration. A key test, then, is whether the different models can match the average acceptance rate of the *employed*.[34]

Finally, we parameterize the extent of heterogeneity and negative selection among the unemployed in our model as follows. We assume that there are two types of workers, a low-x and a high-x type. We parameterize each worker type using $x_{low} = -\sigma_x$ and $x_{high} = \sigma_x$, where σ_x is chosen to match the standard deviation of our residual wage offer estimates that control for observed worker and job characteristics. We choose to target the residualized moments since our goal is to quantify the role of *unobserved* heterogeneity. To be more precise, we assume that our observed offered wage, $\tilde{y}$, satisfies

$$log(\tilde{y}) = log(y) + log(x) + \varepsilon_y, \tag{7}$$

where $log(y) \sim N(\mu_i, \sigma_y)$, $log(x) \sim N(0, \sigma_x)$, and $\varepsilon_y \sim N(0, \sigma_{\varepsilon_y})$ are independently distributed, and thus,

$$\sigma_{\tilde{y}} = \sqrt{\sigma_y^2 + \sigma_x^2 + \sigma_{\varepsilon_y}^2}. \tag{8}$$

We assume a moderate degree of measurement error equal to 13 percent of the unconditional variance in offered wages, consistent with Bound and Krueger (1991). We assume that the actual offered wage distribution has a standard deviation of 0.24, consistent with the findings of Hall and Mueller (2017). Given our calibration of both $\sigma_{\varepsilon_y}^2$ and σ_y^2, we get an estimate for $\sigma_x = \sqrt{\sigma_{\tilde{y}}^2 - \sigma_y^2 - \sigma_{\varepsilon_y}^2}$.

One interpretation of the wage offer gap between the employed and unemployed is selection of the unemployment pool towards the low-x types. In our model, negative selection arises because we allow for differences in separation rates across types. This is consistent with the well-known fact that differences in unemployment rates across skill groups are driven by separations

[34] Another option would be to assume that it is equal to a specific value as in Shimer (2005) or Hall and Milgrom (2008), but we prefer our approach of inferring it directly from our data because there is little consensus of what the appropriate level of b is, except that it should not be too low.

and not job finding. We parameterize the difference between $\delta(x_{low})$ and $\delta(x_{high})$ by matching the difference in prior wages between the employed and unemployed that cannot be explained by observable characteristics. We use an estimate of this difference that is comparable to the difference in prior wages reported in Table 9, but it is based on a sample that is restricted to job offers received within the last six months to make sure that our results are not driven by sample composition effects. The estimated log difference in prior wages within this sample is -0.10, i.e. prior wages are about 10 log points higher among the unemployed when compared to the employed.

Intuitively, the larger the difference between $\delta(x_{low})$ and $\delta(x_{high})$, the more negative selection there is among the unemployed and thus the lower the average prior wage of the unemployed. Note that, in the model without negative selection where $\delta(x_{low}) = \delta(x_{high})$, prior wages are higher for the unemployed by about 17 log points. This is because the prior wages of the employed tend to be from jobs further down on the wage ladder, while the prior wages of the unemployed are from jobs prior to a separation and thus further up on the wage ladder. Negative selection among the unemployed through differences in separation rates allows our model to match the difference in prior wages between the employed and unemployed, by reducing the difference in prior wages from 17 to 10 log points.[35]

The extent of negative selection in our model is also consistent with existing estimates in the literature. In our model, the unemployment rate of the low-x type is 8.9 percent and the unemployment rate of the the high-x type is 6.5 percent, implying that the low-x types make up 57.7 percent of the unemployed. As a point of comparison, Mueller (2017) finds in the CPS data that unemployment risk is 36 percent higher for workers below the median residual wage (after controlling for observables) compared to those above the median residual wage, implying that the share of low-residual wage workers among the unemployed is 57.6 percent.

[35]Our model also suggests that including prior wages in a linear regression model to control for unobserved heterogeneity is not necessarily a good idea. Prior wages differ between the employed and unemployed because of differences in unobserved heterogeneity, but also because of differences in their position on the job ladder. Our model, which accounts for both, directly targets the differences in prior wages for the employed and unemployed.

Table 12: Calibrated Parameter Values

(κ_u, κ_e)	(α_u, α_e)	(β_u, β_e)	(χ_u, χ_e)	b	$\mu_{y,e} - \mu_{y,u}$	$x_{high} - x_{low}$	$\delta(x_{low}) - \delta(x_{high})$
(0.17, 3.30)	(0.03,0.05)	(0.32,0.93)	(0.06,0.39)	1.35	0.16	1.08	0.0021

Before moving on to the results, let us note that Table 11 shows that in the SCE data the offer rate of the employed is around 9 percent and the acceptance rate is 30 percent, implying a target of 2.7 percent for the job-to-job transition rate. This rate is above the average employer-to-employer transition rate in the CPS, which averages around 1.9 percent for our sample period.[36] However, there are also job-to-job transitions that do not involve employer changes. For the time period that we consider, around 1 percent of employed workers in the CPS reported that the usual activities and duties of their current job changed since the prior month while another 0.9 percent reported a change in job description. Adding all of these transitions together implies a job-to-job transition rate of 3.8 percent. Our target for the job-to-job transition rate in the SCE (2.7 percent) lies between its strict and broad definitions in the CPS. If we think of the job-ladder model as a change in job description and salary, it is natural to take into account internal moves when calibrating the model.[37]

4.3 Quantitative Implications

We have shown that our model fits the data well, and that it matches additional empirical moments that are not explicitly targeted. Our model also has various micro and macro implications that are important for understanding on-the-job search and how it differs from search by the unemployed. In this section, we focus on three important areas: the implications for the search-wage gradient; the implications for labor force history and unobserved heterogeneity; and the implications for wage dispersion and the flow value of unemployment.

[36] The calculation of the E-to-E transition rate is based on the following question: "Last month, it was reported that you worked for (employer's name). Do you still work for (employer's name) (at your main job)?"

[37] A longstanding literature in personnel economics documents the importance of internal hires. For example, Lazear and Oyer (2004), using employer-employee data from Sweden show that firms fill a substantial fraction of jobs internally and that the rate of internal hiring relative to external hiring increases monotonically at higher levels. At lower levels, around 40 percent of jobs are filled internally while at the higher levels, almost 90 percent of hires come from within the firm. This evidence points to the prominence of within-firm job ladders.

4.3.1 Implications for the Search Effort-Wage Gradient

A key implication of on-the-job search models with endogenous search effort is that employed workers will reduce their search effort as they climb the job ladder. The evidence in Figure 2 supports this important implication, as there is a clear negative relationship between different measures of search effort and the current wage of the worker. In Figure 4, we compare the relationship implied by our model to its counterpart from our survey data. The model produces a remarkably good fit to the data.[38] If anything, the model underpredicts the negative gradient of search effort by current wage. If search costs were less convex (i.e., setting γ higher than 1.19), it would improve on the fit and lead to a steeper decline in search effort, as workers at the bottom of the job ladder would search more intensely. Overall, we take this as clear evidence in support of models with endogenous search effort.

Figure 4: On-the-Job Search Effort by Current Wages (Model vs. Data)

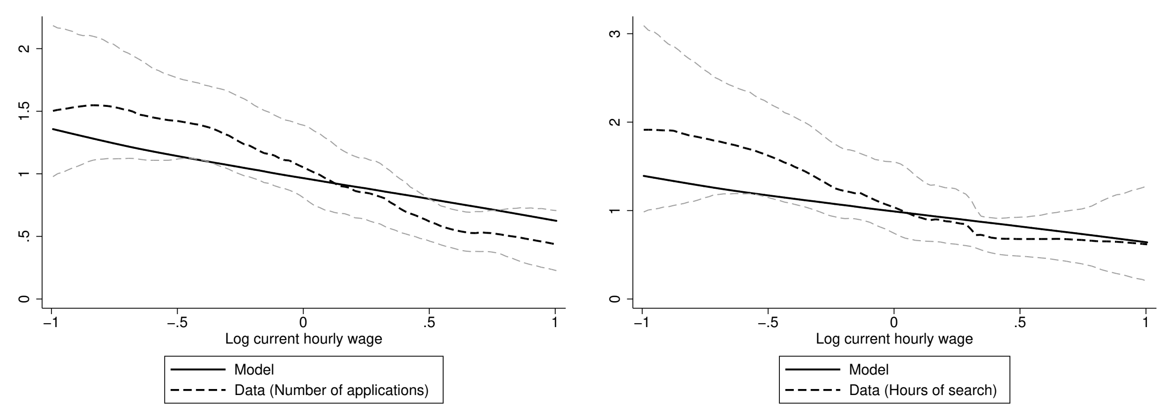

Note: The dashed lines show the 95 percent confidence interval (bootstrapped with 500 replications). For the purposes of the model simulation for these figures, the model was extended from 2 to 11 types.

The endogenous search margin of our model is also key to the model's success in matching the acceptance rate of employed job seekers. As Table 14 shows, our model matches the offer acceptance rate for employed workers in our survey perfectly even though it is not a targeted moment. Workers at the bottom of the wage ladder are more likely to accept an offer and, due

[38] Note that, in the model without unobserved ex-ante worker heterogeneity and measurement error, search effort would decline more steeply with the current wage. Thus, accounting for both of these forces, is important when analyzing the fit of the search-wage gradient.

to endogenously higher search effort, are also more likely to receive an offer than those at the top of the ladder. In Table 14, we show that the offer acceptance rate of the employed drops to 17.6 percent in a basic version of the model with exogenous job search, well below the empirical estimate of 30.0 percent. Without an operative search decision, workers are equally likely to receive offers regardless of their position on the job ladder and more offers are rejected since those who are higher on the wage ladder reject most offers.

4.3.2 Accounting for the *Wage Offer Premium* and Unobserved Heterogeneity

In Subsection 3.5, we examined how much of the wage offer premium could be accounted for by a rich set of controls. In this subsection, we complement our empirical analysis using our model to quantitatively assess possible determinants of the observed wage offer differential between the employed and non-employed. We decompose the differential into three parts: the part due to unobserved heterogeneity, the part due to censoring of the offer distribution, and a residual, unexplained component.

Panel I of Table 13 shows that the model attributes about 7 log points out of the 25 log point residual wage offer differential to unobserved heterogeneity. The contribution of censoring is relatively small, explaining only 2 log points of the wage offer gap. Note that we parameterized the degree of censoring using specific questions on *unrealized offers* in our survey. The fact that we match the job acceptance rate addresses any concern that we underestimate the contribution of censoring since a greater degree of censoring would lead to a counterfactually high job acceptance rate among the employed.[39]

In summary, our calibration suggests that the remaining 16 log points, which is 64 percent of the 25 log point wage offer gap, is accounted for by the residual, unexplained component. This result compares well with our empirical findings, where our imperfect proxies for unobserved heterogeneity suggest a residual wage offer gap of about 20 log points.

Additionally, our model predicts a correlation between *ex ante* unobserved heterogeneity and

[39] See also the robustness section for an alternative calibration of the censoring parameters.

Table 13: Wage Offer Differentials in the Data and the Model

I. Decomposition of offer wage differential	Data	Model Simulation
Wage offer differential	0.250	0.25
- due to worker-heterogeneity	—	0.07
- due to censoring	—	0.02
- due to exogenous differential	—	0.16
II. Decomposition of offer wage differential	**Data**	**Model Regression**
Wage offer differential	0.250	0.25
Controlling for prior employment history	0.222	0.21
Additionally controlling for prior labor force history	0.205	—

Note: Empirical estimates come from authors' tabulations from the October 2013-15 waves of the SCE Labor Supplement, for individuals aged 18-64, excluding the self-employed, with at least one job offer in the last six months. See notes to Tables 8 and 10 for details on controls included in each specification. Standard errors are in parentheses. See text for details of model simulation results.

labor force histories that provides some helpful guidance in addressing unobserved heterogeneity in empirical work. The model implies that even controlling for their *current* labor force status, low-x workers are more likely to be unemployed and less likely to have worked in the past. This observation makes the work history of a worker a useful proxy for unobserved heterogeneity. Our survey allows us to control for workers' labor force history over the last five years. As Panel II of Table 13 shows, controlling for the fraction of the last five years that the worker was employed reduces the wage offer gap from 0.25 to 0.22. When we additionally control for the share of time spent unemployed and the share of time spent as a student, the difference goes down to 0.21. We can connect our quantitative analysis to our data by implementing the same regression analysis on model-generated data from a simulation and compare it to these results. Specifically, we simulate our model and then use the simulated data to replicate the regressions from the empirical analysis. When we run the same regression on the model-generated data, we find that the difference goes down from 0.25 to 0.21, very similar to the data, verifying the usefulness of labor force history as a control for unobserved heterogeneity. The congruence of the regression results in the data and the model is particularly noteworthy because the calibration of our model exploits variation in prior wages rather than employment histories. It is, of course, possible to go one step further and ask how powerful a proxy work history is in controlling for unobserved heterogeneity since our model provides us with an estimate for the role of unobserved heterogeneity. If we shut down the differences in x, the wage offer differential goes down from 0.25 to 0.18 suggesting that controlling for work history accounts for more than half of the contribution of unobserved heterogeneity.

4.3.3 Implications for Wage Dispersion and the Flow Value of Unemployment

Our model also performs demonstrably well in matching the amount of wage dispersion observed in the data, as we show in Table 14. Hornstein, Krusell, and Violante (2011) argue that a standard model of frictional search and matching in the labor market can only account for a tiny fraction of the wage dispersion observed in the data. They find that extending the model to include on-the-job search can increase the ratio of the mean wage to the minimum wage observed in the data (the mean-min ratio), to as high as 1.4, but not nearly as high as the 1.7 ratio they observe in the data. Our benchmark model generates a mean-min ratio of 1.68 after controlling for worker heterogeneity, x. The ability of our model to generate wage dispersion that is consistent with the data is particularly notable because it does so while yielding a reasonable value for the flow utility of unemployment, 0.74, and matching the acceptance rate of the employed in our data. Hornstein, Krusell, and Violante (2011) advocate that any search model that aims to fit transition rates and wage dispersion should have its performance evaluated using its implied flow value of unemployment. Intuitively, our model does well in this respect because both the higher search efficiency and superior job offer distribution of the employed limit the option value of unemployment—i.e. the unemployed perceive little value in waiting for a better offer if they can continue to sample more plentiful and better offers while employed. These results are remarkably robust to a number of alternative calibrations, as we discuss in the robustness subsection.[40]

In Table 14, we consider variations of the model that remove features introduced in our benchmark model. We show that these models cannot match the acceptance rate of the employed or

[40] It is not clear what the empirical target for the dispersion in offered wages should be. Hornstein, Krusell, and Violante (2011) find mean-min ratios of 1.48 to above 2, depending on how narrowly they define a given labor market. Our measure of dispersion is taken from Hall and Mueller (2017) who estimate the standard deviation of wage offers for a given worker at 0.24. This estimate is at the upper end of other papers who identify the dispersion of wage offers from wage changes associated with job switches for a given individual (e.g., Tjaden and Wellschmied, 2014). Interestingly, the mean-min ratio in our model with endogenous search effort exactly matches the lower end of the mean-min wage ratios reported in Hornstein, Krusell, and Violante.

Table 14: Implications for the Wage Dispersion Puzzle

	Data	**Benchmark model:** Endogenous search effort + worker het. + diff. offer dist.	**Restricted model versions:** (1) Endogenous search effort + worker heterogeneity	(2) Endogenous search effort	(3) Exogenous offers
Mean acceptance rate of employed	0.300	0.300	0.215	0.213	0.176
Mean search cost of unemployed		0.17	0.30	0.31	—
Mean search cost of employed		0.06	0.04	0.04	—
$b/E(w)$		0.81	0.69	0.69	0.33
$b/E(w)$ (net of search costs)		0.74	0.50	0.50	0.33
Mean-min ratio (conditional on x)		1.68	1.48	1.48	1.43

generate empirically-plausible values of wage dispersion with a reasonable flow value of unemployment. In particular, we consider three model variations with each sequentially removing one feature from our benchmark model: (1) a model with *ex ante* worker heterogeneity where search is endogenous and search efficiency varies by employment status but wage offers do not; (2) a model *without* worker heterogeneity where search is endogenous and search efficiency varies by employment status but wage offers do not; (3) a basic model *without* worker heterogeneity where search is *exogenous* and search efficiency varies by employment status but wage offers do not. In all three cases, the acceptance rate of the employed never rises above 22 percent, well below the 30 percent observed in the data. Furthermore, the flow value of unemployment remains between 0.33 and 0.50, while the mean-min wage ratio never rises above 1.50. Table D1 in Appendix D reports the full simulation results for these three cases.

4.4 Robustness

We examine the results of five robustness exercises in order to assess our model's quantitative implications. First, we assess an alternative calibration strategy for the censoring of job offers that relies only on matching the acceptance rate of the employed and unemployed. Second, we assess the model for alternative parameterizations of the dispersion in wage offers. Third, we pin down the curvature of the search cost function (the γ parameter) by matching moments of the distribution of search effort. Fourth, we probe the robustness of our findings to alternative assumptions about the unemployed's job offer rate that is more in line with the calibration in Hornstein et al. (2011). Finally, we use the same model, but target raw moments instead of residualized moments in terms of overall wage dispersion and the employed-unemployed wage offer gap. We implement all robustness exercises on our benchmark model and a version of the model where wage offers do not vary by employment status (model 1 in Table 14) to examine whether any of these exercises can further account for any of the observed wage offer differential.

Our first robustness exercise focuses on an alternative way of calibrating the degree of censoring. Recall that one interpretation of our findings is that the employed are more selective in their job search and apply only to jobs that they think would dominate their current job. This implies that workers further up the wage ladder should be more selective. We incorporate this channel into our benchmark model by accounting for censoring of the wage offer distribution by the employed searchers. Our benchmark model suggests that this only accounts for 8 percent (2 log points) of the observed wage offer gap. Our approach, however, only addresses the rejection of unrealized offers. It does not account for censoring due to employers passing on extending job offers to those they feel are out of their reach. We address this possibility by calibrating a version of the model where the offer distributions are assumed to be the same for the employed and unemployed (model 1 in Table 14), but choose the censoring parameter, χ_e, to match the acceptance rate of the employed, which constitutes a natural upper bound on the extent of censoring in employed offers. As we show in Table D3 of the appendix, this version implies a substantial incidence of censoring of the wage offers of the employed, with χ_e of 0.72, and thus endogenously generates a wage offer differential of 0.08. However, this calibration generates only half of the wage offer differential observed in the data (0.12 in the model vs. 0.25 in the data) suggesting that even if we attribute the high acceptance rate of the employed to censoring, the model still falls short of explaining the wage offer differential without differences in wage offers by employment status. At the same time, it is worth pointing out that this version generates a reasonable flow value of unemployment of 0.70 and a mean-min ratio of 1.58. The reason is that this version of the model implies a very high underlying contact rate for the employed, and thus limits the option value of unemployment to a similar extent as our benchmark model.

Our other robustness checks are explained in detail in the appendix and summarized here briefly. If we calibrate the variance of wage offers to be twice as high (i.e., $\sigma_y = 0.34$), the implied wage dispersion, as measured by the mean-min ratio, is naturally larger under this calibration. The implied flow value remains reasonably high, however, at 0.52. If we set $\gamma = 2.5$, we are able to roughly match the observed search effort at the 95th percentile of the distribution. These versions of the model generate more search at the bottom of the wage ladder and less at the top, and further improve the fit of the search-wage relationship (see Figure D1 in the Appendix). In our fourth robustness check, we set the offer rate of the unemployed to 0.60, from 0.339, which implies the unemployment-to-employment transition rate of 0.32 targeted by Hornstein, Krusell, and Violante (2011). The results indicate a flow value of unemployment of 0.58 with a mean-min ratio of 1.69. Thus, our resolution of the puzzle holds. Finally, when we target the raw, unconditional moments in our data, we find that censoring accounts for 2 log points and unobserved heterogeneity accounts for 19 log points of the unconditional 40 log point wage offer differential. This leaves an unaccounted-for wage differential of about 19 log points, which is very close to the 16 log points in our benchmark exercise.

In summary, our model's implications on the wage offer premium, the search-wage gradient, and the wage dispersion puzzle are robust to various alternative specifications. In particular, while some variations attribute a more important role to worker heterogeneity, at least half of the wage offer differential remains unexplained, leaving room for alternative explanations that we discuss in the concluding remarks. As for the search-wage gradient, our baseline calibration does a good job in capturing the relationship in the data and it can further be improved upon, as we briefly discussed in this subsection. The alternative specifications that we discuss above also produce reasonable flow values of unemployment along with substantial wage dispersion.[41]

[41]While there is substantial disagreement in the literature on the exact value of the flow value, ranging from 0.4 (Shimer, 2005) to near unity (Hagedorn and Manovskii, 2008), it is difficult to rationalize a value below 0.4.

5 Concluding Remarks

In this paper, we document new facts about the search effort and search outcomes of the employed and non-employed. We find that search among the employed is pervasive. Nearly one-quarter report searching for work in the previous month. We also find that on-the-job search effort is negatively related to one's current wage. Compared to the non-employed, the employed are more efficient at search. They receive a similar number of offers despite exerting a fraction of the effort. A sizable fraction of job offers go to employed workers not even looking for work, underscoring the importance of unsolicited employer contacts in the job search process. The employed also tend to receive better job offers than the non-employed. Their offered wages are 40 log points (49 percent) higher unconditionally and 25 log points (28 percent) higher after controlling for observable characteristics of the worker and job. Despite their relatively poorer job offers, the non-employed are more likely to accept them.

We apply our findings to an on-the-job search framework with endogenous search effort and unobserved *ex ante* worker heterogeneity. We incorporate our key findings on unsolicited job offers, censoring of the wage offer distribution, and heterogeneity in search efficiency by labor force status. Our model exhibits a remarkably good fit to the data and is able to replicate the job-to-job transition rate and negative relationship between on-the-job search effort and wages well. We find that unobserved heterogeneity explains only 28 percent of the offered wage differential between the employed and unemployed. After accounting for all factors, our model implies that a 16 log point unexplained wage offer gap between the employed and unemployed remains. Finally, our model can replicate the wage dispersion observed in the data with a reasonable flow value of unemployment. Our empirical findings thus constitute an important piece in the resolution of the *frictional wage dispersion puzzle* in the search and matching literature.

An open question remains as to why wage offers appear to be of better quality for the employed compared to the non-employed even after controlling for unobserved heterogeneity and the censoring of job offers. One can think of several possible reasons. Some are more familiar to the literature that has studied wages, job loss, and unemployment, while others are more specific to the job search process. While we leave a thorough evaluation of different mechanisms to future research, we briefly discuss these potential factors here.

One possibility is that human capital depreciates during periods of non-employment. In this case, the employed and non-employed may have a similar wage (and potentially similar skill levels) when they separate from their previous job, but the skills of the non-employed depreciate, leading them to have lower-quality job offers, on average. In this case, accounting for the work history of the employed and non-employed would reduce the gap. We show that controlling for the previous five-year work history (specifically, the fraction of the prior five years spent employed, unemployed, etc.) only reduces the wage gap from 0.250 to 0.205. This suggests that human capital depreciation can explain only a small fraction of the wage offer gap. As discussed, questions about five-year work histories do a good job at capturing unobserved differences in worker productivity.

The presence of bargaining and counter-offers are another way the search process can affect the wage offer gap. We find that 39 percent of offers to the employed, while only 24 percent of offers to the non-employed, involved some bargaining between the individual and the potential employer. In general, a greater propensity to bargain with the potential employer should increase the reported wage offer, all else equal. Moreover, 14 percent of the employed with an outside offer received some form of counter-offer from their current employer. While the latter estimate falls clearly short of the rate of counter-offers in models such as Postel-Vinay and Robin (2002) and Cahuc, Postel-Vinay, and Robin (2006), it is possible that the threat of such counter-offers raises the mean wage offers for the employed even when no such offer occurs in equilibrium.

Finally, our evidence is also consistent with an implicit penalty for job seeking while unemployed, in line with the results of Kroft, Lange, and Notowidigdo (2013). Such a penalty could be due to either explicit discrimination or a perceived signal about a worker's unobserved productivity that employers infer from their employment status (as in Gibbons and Katz, 1991). It is also possible that it is the result of firms discriminating against the unemployed to take advantage of their lower reservation wages (as in Carrillo-Tudela, 2009). Regardless of its source, any such penalty on job search while unemployed would imply that the employed and unemployed draw from differing wage offer distributions.

References

[1] Aguiar, Mark A., Erik Hurst, and Loukas Karabarbounis, 2013. "The Life-Cycle Profile of Time Spent on Job Search," American Economic Review 103(3): 111-116.

[2] Arbex, Marcelo, Dennis O'Dea, and David Wiczer, 2016. "Network Search: Climbing the Job Ladder Faster, Federal Reserve Bank of St. Louis Working Paper No. 2016-009A.

[3] Bagger, Jesper and Rasmus Lentz, 2017. "An Empirical Model of Wage Dispersion with Sorting," forthcoming, Review of Economic Studies.

[4] Barron, John M. and Otis W. Gilley, 1981. "Job Search and Vacancy Contacts: Note," American Economic Review 71(4), 747-752.

[5] Bound, John and Alan B. Krueger, 1991. "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?," Journal of Labor Economics 9(1): 1-24.

[6] Black, Matthew, 1980. "Pecuniary Implications of On-The-Job Search and Quit Activity," Review of Economics and Statistics 62(2): 222-229.

[7] Blau, David M. and Philip K. Robins, 1990. "Job Search Outcomes for the Employed and Unemployed," Journal of Political Economy 98(3): 637-655.

[8] Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin, 2006. "Wage Bargaining with On-the-job Search: Theory and Evidence," Econometrica 74(2): 323-364.

[9] Carrillo-Tudela, Carlos, 2009. "An Equilibrium Search Model When Firms Observe Workers' Employment Status," International Economic Review 50(2): 485-506.

[10] Christensen, Bent Jesper, Rasmus Lentz, Dale T. Mortensen, George R. Neumann, and Axel Werwatz, 2005. "On-the-Job Search and the Wage Distribution," Journal of Labor Economics 23(1): 31-58.

[11] Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger, 2013. "The Establishment-Level Behavior of Vacancies and Hiring," Quarterly Journal of Economics 128(2): 581-622.

[12] Davis, Steven J. and John C. Haltiwanger, 2014. "Labor Market Fluidity and Economic Performance," NBER Working Paper No. 20479, National Bureau of Economic Research.

[13] Elsby, Michael W.L., Bart Hobijn, and Ayşegül Şahin, 2010. "The Labor Market in the Great Recession," Brookings Papers on Economic Activity 41(1): 1-48.

[14] Faberman, R. Jason and Alejandro Justiniano, 2015. "Job switching and wage growth," Chicago Fed Letter No. 337.

[15] Fallick, B., and C. A. Fleischman, 2004. "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows," Finance and Economics Discussion Series 2004-34. Board of Governors of the Federal Reserve System.

[16] Gibbons, R. and Lawrence F. Katz, 1991. "Layoffs and Lemons," Journal of Labor Economics 9(4): 351-380.

[17] Hagedorn, Marcus and Iourii Manovskii, 2008. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited, American Economic Review 98(4): 1692-1706.

[18] Hall, Robert E. and Paul R. Milgrom, 2008. "The Limited Influence of Unemployment on the Wage Bargain," American Economic Review 98(4): 1653-74.

[19] Hall, Robert E. and Alan B. Krueger, 2012. "Evidence on the Incidence of Wage Posting, Wage Bargaining, and On-the-Job Search," American Economic Journal 4(4): 56-67.

[20] Hall, Robert E. and Andreas I. Mueller, 2017. "Wage Dispersion and Search Behavior: The Importance of Non-Wage Job Values," forthcoming in the Journal of Political Economy.

[21] Holzer, Harry J., 1987. "Informal Job Search and Black Youth Unemployment," American Economic Review 77(3): pages 446-452

[22] Hornstein, Andreas, Per Krusell, and G. L. Violante, 2011. "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," American Economic Review 101(7): 2873-98.

[23] Kahn, Lawrence M. and Stuart A. Low, 1982. "The Relative Effects of Employed and Unemployed Job Search, " The Review of Economics and Statistics 64(2): 234-241.

[24] Kahn, Lawrence M., 2012. "Temporary jobs and job search effort in Europe," Labour Economics 19(1): 113-128.

[25] Karahan, Fatih, Ryan Michaels, Benjamin Pugsley, Ayşegül Şahin, and Rachel Schuh, 2017. "Do Job-to-Job Transitions Drive Wage Fluctuations over the Business Cycle?" American Economic Review 107(5): 353-57.

[26] Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo, 2013. "Duration Dependence and Labor Market Conditions: Theory and Evidence from a Field Experiment," Quarterly Journal of Economics 128(3): 1123-1167.

[27] Krueger, Alan B. and Andreas I. Mueller, 2010. "Job Search and Unemployment Insurance: New Evidence from Time Use Data," Journal of Public Economics 94(3-4): 298-307.

[28] Lazear, Edward P. and Paul Oyer, 2004. "Internal and External Labor Markets: A Personnel Economics Approach," Labour Economics 11(5): 527-554.

[29] Lise, Jeremy, 2013. "On-the-Job Search and Precautionary Savings," Review of Economic Studies 80(3): 1086-1113.

[30] Low, Hamish, Costas Meghir and Luigi Pistaferri, 2010. "Wage Risk and Employment Risk over the Life Cycle," American Economic Review 100(4): 1432-1467.

[31] McCall, John Joseph, 1970. “Economics of Information and Job Search,” Quarterly Journal of Economics 84(1): 113-126.

[32] Mincer, Jacob, 1991. “Education and Unemployment,” NBER Working Paper No. 3838, National Bureau of Economic Research.

[33] Mortensen, Dale T., 1977. “Unemployment Insurance and Job Search Decisions,” Industrial and Labor Relations Review 30(4): 505-517.

[34] Moscarini, Giuseppe, and Fabien Postel-Vinay, 2017. “The Relative Power of Employment-to-Employment Reallocation and Unemployment Exits in Predicting Wage Growth,” American Economic Review 107(5): 364-68.

[35] Mueller, Andreas I., 2010. “On-the-Job Search and Wage Dispersion: New Evidence from Time Use Data,” Economics Letters 109(2): 124-127.

[36] Mueller, Andreas I., 2017. “Separations, Sorting and Cyclical Unemployment,” American Economic Review 107(7): 2081-2107.

[37] Mukoyama, Toshihiko, Christina Patterson and Ayşegül Şahin, 2018. “Job Search Behavior over the Business Cycle,” the American Economic Journal: Macroeconomics 10(1) 190-215.

[38] Pissarides, Christopher A., 1985. “Short-run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages,” American Economic Review 75(4): 676-690.

[39] Postel-Vinay, Fabien and Jean-Marc Robin, 2002. “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” Econometrica 70(6): 2295-2350.

[40] Shimer, Robert J., 2005. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” American Economic Review 95(1): 25-49.

[41] Tjaden, Volker and Felix Wellschmied, 2014. “Quantifying the Contribution of Search to Wage Inequality,” American Economic Journal: Macroeconomics 6(1): 134-161.

APPENDIX — FOR ONLINE PUBLICATION

Table of Contents

A Comparison of SCE Labor Survey to External Data

A.1 Results using CPS Definition of Unemployment

Our survey allows for a broader measure of job search among the non-employed than what is possible using the Current Population Survey (CPS). In the CPS, the unemployed are the non-employed who either were on temporary layoff or had actively searched within the last four weeks and were available for work. We use the same definition. The difference in scope between the two surveys is that the CPS does not follow up with certain non-employed individuals (predominantly the retired and disabled) who report that they either do not want work or cannot work, to ask if they had searched.

Our survey suggests that many of these individuals actively search and are available for work. Table A1 shows that they represent just over 12 percent of those counted as out of the labor force under the CPS definition. A similar fraction of those out of the labor force sent at least one application in the prior four weeks or spent some time searching in the previous seven days. As Table 1 in the main text shows, including these individuals in our job search measure increases the measured unemployment rate by 2.7 percentage points. Further analysis (not reported here) suggests that the majority of the difference is due to retired individuals seeking only part-time work. In fact, about half of all individuals actively searching from out of the labor force are only seeking part-time work. Conditional on actively searching, just under 10 percent are looking for work similar to their most recent job.

Tables A2 and A3 replicate our job search analysis using the CPS scope and definition of unemployment using the SCE data. This counts those non-employed who do not explicitly report wanting work as out of the labor force, regardless of whether they later report that they actively searched and are available. The tables correspond to Tables 3 and 5 in the main text, though we only report the replicated results by labor force status at the time of the survey since the difference in the unemployment definition only matters for this period. We determine labor force status in the previous month using responses from a variety of other survey questions that do not directly correspond to the CPS definition. Note also that the results for the employed (regardless of whether they actively searched for work) are the same under both definitions.

The tables show that moving from our job search measure to the CPS measure of unemployment has only a minor effect on our estimates for the intensive margin of search effort and search outcomes for the unemployed. The CPS definition implies a somewhat higher level of search effort. The number of applications rises by 19 percent, to 9.65 per month, and time spent searching rises by 30 percent, to 10.96 hours per week. The number of employer contacts, job interviews, and job offers received all rise somewhat as well, though the differences between the estimates in Table A3 and Table 5 in the main text are not statistically significant. The job search definition and the CPS definition of employment also imply similar ratios of employer contacts per application and mean job offers per application.

Finally, note that the definitions used here have no bearing on our model calibration since it uses job search effort and outcome estimates based on labor force status in the prior month.

Table A1: Basic Job Search Statistics by Labor Force Status, CPS Measure of Unemployment

	Employed	**Unemployed**	**Out of Labor Force**
Percent that actively searched for work	23.1	99.1	14.0
	(0.9)	(0.9)	(1.6)
Percent that actively searched and available	14.1	99.1	12.2
for work	(0.7)	(0.9)	(1.5)
Percent reporting no active search or	6.1	0.4	5.3
availability, but would take job if offered	(0.5)	(0.6)	(1.0)
Percent applying to at least one vacancy in last	19.8	96.2	12.2
four weeks	(0.8)	(1.8)	(1.5)
Percent with positive time spent searching in	20.5	93.5	11.1
last seven days	(0.8)	(2.4)	(1.4)
Percent only searching for an	9.2	—	—
additional job	(0.6)		
Percent only seeking part-time work,	20.5	8.4	49.8
conditional on active search	(1.8)	(2.7)	(6.3)
Percent only seeking similar work (to most	27.4	6.4	9.5
recent job), conditional on active search	(2.1)	(2.4)	(4.2)
No. of Observations	2,302	110	485

Note: Estimates come from authors' tabulations from the 2013-15 panel of the SCE labor survey, for all individuals aged 18-64, by labor force status using the CPS definition of unemployment. Standard errors are in parentheses.

A.2 Search Effort Estimates in the SCE and ATUS

Next, we compare our estimates of the time spent searching for work to comparable estimates from the time diaries of the American Time Use Survey (ATUS). We use the ATUS because existing measures of search effort are rare, particularly when one wants to measure on-the-job

Table A2: Intensive Margin: Search Effort by Labor Force Status, CPS Measure of Unemployment

	Employed			Unemployed	Out of Labor Force
	Looking for Work	Not Looking	All		
Labor Force Status at Time of Survey					
Hours spent searching, last 7 days	4.35	0.05	1.18	10.96	0.55
	(0.30)	(0.01)	(0.09)	(0.97)	(0.12)
Mean applications sent, last 4 weeks	4.63	0.00	1.22	9.65	0.74
	(0.49)	(—)	(0.13)	(1.65)	(0.22)

Note: Estimates come from authors' tabulations from the SCE survey, for all individuals aged 18-64, excluding the self-employed, by detailed labor force status using the CPS definition of unemployment. Standard errors are in parentheses.

Table A3: Search Outcomes by Labor Force Status, CPS Measure of Unemployment

	Employed			Unemployed	Out of Labor Force
	Looking for Work	Not Looking	All		
Labor Force Status at Time of Survey					
Mean contacts received	1.874	0.337	0.742	1.592	0.186
	(0.281)	(0.038)	(0.079)	(0.338)	(0.036)
Mean unsolicited contacts	0.783	0.298	0.426	0.640	0.103
	(0.124)	(0.032)	(0.040)	(0.231)	(0.028)
Mean job interviews (2014-15)	0.460	0.005	0.115	0.550	0.037
	(0.045)	(0.002)	(0.012)	(0.164)	(0.020)
Mean offers	0.425	0.086	0.175	0.389	0.112
	(0.039)	(0.011)	(0.014)	(0.106)	(0.026)
Mean unsolicited offers	0.047	0.046	0.046	0.055	0.050
	(0.010)	(0.009)	(0.007)	(0.022)	(0.020)
Fraction with at least one offer	0.299	0.057	0.118	0.216	0.064
	(0.020)	(0.007)	(0.007)	(0.039)	(0.011)
Fraction with at least one unsolicited	0.041	0.028	0.031	0.055	0.026
offer	(0.009)	(0.005)	(0.004)	(0.022)	(0.007)
Fraction with at least one unsolicited	0.345	0.086	0.155	0.236	0.081
offer, including unrealized offers	(0.021)	(0.007)	(0.008)	(0.041)	(0.012)
N	508	1,520	2,028	110	485

Note: Estimates come from authors' tabulations from the SCE survey, for all individuals aged 18-64, excluding the self-employed, by detailed labor force status using the CPS definition of unemployment. Standard errors are in parentheses. Job interview data are only available for 2014 and 2015.

Table A4: Time Spent Searching for Work, ATUS and SCE Labor Survey Data

	Employed	Unemployed	Out of Labor Force
American Time Use Survey			
Percent reporting time spent searching for work, prior day	0.6	16.5	0.9
Average minutes spent searching, prior day, all respondents	0.8	26.7	1.4
Average minutes spent searching, prior day, conditional on positive search time	145.3	161.8	166.6
N	18,460	1,045	6851
SCE Labor Survey			
Percent reporting time spent searching for work, last seven days	20.5 (0.8)	93.5 (2.4)	11.1 (1.4)
Average minutes spent searching, last seven days, all respondents	62.0 (4.6)	657.6 (57.9)	33.1 (7.3)
Average minutes spent searching, last seven days, conditional on positive search time	299.0 (18.6)	684.9 (60.8)	290.7 (56.3)
N	2,302	110	485

Note: Estimates come from authors' tabulations from the 2013-15 waves of the American Time Use Survey (top panel) and the SCE labor survey (bottom panel), for all individuals aged 18-64, by labor force status. The SCE estimates use the BLS definition of unemployment for determining labor force status. Standard errors are in parentheses. We do not report the standard errors for the ATUS as they were very small.

search.

Table A4 reports our results. We focus on individuals age 18-64 in both surveys and report the average time spent searching for work by labor force status (employed, unemployed, or out of the labor force).[42] We use the CPS definition of unemployment to maintain consistency across the surveys. There are differences between the frequencies over which each survey measures search time, for which we cannot control. The ATUS measures search time for a single day using a detailed time diary, while the SCE asks respondents the number of hours they spent searching for work over the previous seven days.

We find notable differences between the estimates from the SCE and the estimates from the

[42] See also Mueller (2010) for similar statistics for an earlier period.

ATUS, with the SCE estimates implying much more time spent searching for work overall, but less time spent searching (in terms of a daily average) for the subset of individuals who reported positive search time. This is likely due to measurement differences between the surveys and the nature of search. Aside from the obvious frequency differences between the ATUS and SCE data, the ATUS also likely understates the time spent searching for work because individuals only report their primary activity in the time diary. Thus, if an individual is literally searching while on the job, it will likely show up as work time rather than search time. The daily data also likely understate search effort because, as the comparison of the two surveys suggests, search effort is discrete and intermittent. Only 0.6 percent of the employed report any time spent searching on a given day in the ATUS, but 20.5 percent of the employed reported searching within the last seven days in the SCE. Even among the unemployed, who are defined as actively looking for work, the ATUS estimates suggest that only 16.5 percent looked for work on the previous day, while the SCE estimates suggest that 93.5 percent searched within the last seven days. If we condition our estimates on those who reported positive search time, and compare the daily (average) time spent searching between the ATUS and SCE data, we find that the SCE estimates suggest lower search effort per day, on average. Again, this is evidence that search is likely intermittent, with individuals searching for some time on several days during a week than a set amount each day. Consequently, when we couple our evidence from the SCE data with the previous research on on-the-job search mentioned in the main text (e.g., Black, 1980; Blau and Robins, 1990), we feel that our estimates of search effort provide a more comprehensive and reliable measure than the ATUS.

B Measuring Labor Force Status in the Previous Month

B.1 Prior Month's Labor Force Status Based on the SCE Labor Supplement

This appendix details our methodology for determining labor force status in the prior month and evaluates our measure along several comparable dimensions. We derive a labor force status for individuals four weeks prior to their survey interview using a range of survey responses from the SCE labor supplement. We use this measure in our model calibration because it treats the search effort and offer arrivals reported in the survey as subsequent outcomes based on this initial labor force status.

To determine labor force status in the prior month, we first check to see if an individual received an offer in the four weeks prior to the survey. If so, we assign the labor force status at the time they received their job offer (employed or non-employed). This approach assumes that labor force status did not change between the time they received their job offer and four weeks prior to the labor supplement survey. We feel that this is a reasonable assumption given the relatively short time interval. In the 2014 and 2015 waves of the survey, we have additional information on whether an individual was actively searching at the time they received their offer. If so, we count them as unemployed, and if not, we count them as out of the labor force. For those in the 2013 wave, we have to make some modest assumptions to determine whether someone was unemployed or out of the labor force. If an individual who was non-employed at the time of the job offer was employed at the time of the 2013 labor survey, we assume that they were actively searching and count them as unemployed. If they were unemployed at the time of the survey and have been searching for over four weeks, we also count them as unemployed. Otherwise, we count them as out of the labor force in the previous month.

For the remaining individuals (who are the vast majority of respondents), we determine their prior month's labor force status starting with their labor force status at the time of the survey. If an individual did not receive a job offer in the last four weeks but was employed at the time of the survey, we determine their prior month's labor force status as follows: if they report that their current job tenure is at least one month, or if they report tenure of less than a month but with less than two weeks between jobs, we count them as employed in the previous month. Otherwise, we assume that these individuals were actively searching for work and count them as unemployed in the previous month.

If an individual was unemployed at the time of the survey and did not receive an offer in the last four weeks, we count them as employed in the previous month if they were on temporary layoff for less than one month or if their current non-employment spell was one month or less. We count them as unemployed in the prior month if they were on temporary layoff for more than one month or if they report actively searching for work for more than one month. Otherwise, we count them as out of the labor force.

Finally, if an individual was out of the labor force at the time of the survey and did not receive an offer in the last four weeks, we count them as employed if their current non-employment spell is one month or less. We count them as unemployed if they report actively searching for more than one month and they are not currently disabled. Otherwise, we count them as out of the labor force.

Evaluation of this approach suggests that our methodology produces a sensible measure of the prior month's labor force status along several dimensions. First, our estimates imply an employment-to-population ratio of 0.750, an unemployment rate of 6.3 percent, and a labor force participation rate of 80.0 percent. All are roughly comparable to the CPS estimates and the SCE labor supplement estimates (using the CPS definition of unemployment) in Table 1 of the main text. The unemployment rate in the previous month is likely somewhat higher because of our assumption that all those who became employed during the month but did not report a job offer were unemployed when hired.

Second, Table B1 reports labor force transition rates for two data sources. The first source is the SCE labor supplement, which estimates the transition rates using our measure of labor force status in the prior month and labor force status at the time of survey (using the CPS definition described in Appendix A). The second source is the monthly CPS, which measures the transition rates between September and October of each year (2013-15) for individuals in the survey during both months. The transition rates for the SCE are generally very comparable to the transition rates for the CPS. The job-separation rates into unemployment and out of the labor force are nearly identical. The SCE labor supplement has a slightly lower job-finding rate for the unemployed and a notably lower job-finding rate for those out of the labor force. Transitions between unemployment and being out of the labor force are roughly comparable between the two surveys.

B.2 Prior Month's Labor Force Status Based on Monthly SCE Data

Another way to test the validity of our estimates of labor force status in the prior month is to compare it to results based on labor force status for individuals in the regular, monthly SCE for the previous month. The monthly SCE data's measure of labor force status in the previous month

Table B1: Monlthly Labor Market Transition Rates by Labor Force Status

(a) SCE Labor Supplement

	Transition Probability to		
Labor Force Status in Prior Month	**Employment**	**Unemployment**	**Out of the LF**
Employed	0.965	0.012	0.023
Unemployed	0.190	0.550	0.260
Out of the Labor Force	0.014	0.051	0.935

(b) Current Population Survey

	Transition Probability to		
Labor Force Status in September	**Employment**	**Unemployment**	**Out of the LF**
Employed	0.961	0.012	0.027
Unemployed	0.235	0.540	0.225
Out of the Labor Force	0.065	0.040	0.895

Notes: The top panel reports the labor force transition rates using the SCE labor supplements from October 2013-15. It uses the methodology described in the appendix to determine the previous months' labor force status and uses the CPS definition of unemployment for labor force status at the time of the survey. The bottom panel reports the labor force transition rates from the CPS using data matched across September and October of 2013-15.

is generally not consistent with the timing of the SCE labor supplement because individuals may respond to the labor supplement anywhere from a few days to nearly two months after their most recent monthly SCE interview. To deal with this, we assign a prior month's labor force status to individuals in the labor supplement based on the timing between the supplement and their September SCE interview. If the gap between interviews is 22 days or more, we use their September labor force status. If the gap is 21 days or less, or if the September data are missing, we use their August labor force status. We adjust all estimates of search outcomes so that they can be interpreted as monthly rates.

Table B2 replicates the bottom panels of Tables 3 and 5 of the main text using the prior month's labor force status measure derived from the monthly SCE data. The table shows that the estimates are very similar to those estimated using our prior month's labor force status measure derived from the labor supplement. Some minor exceptions exist for the unemployed. For example, application and offer rates are somewhat lower using the monthly SCE measure. Otherwise, the two measures produce nearly identical estimates of search effort and search outcomes.

Table B2: Search Outcomes by Prior Month's Labor Force Status, based on Monthly SCE

	Employed	**Unemployed**	**Out of Labor Force**
Labor Force Status in August/September, Monthly SCE			
Search Effort			
Mean applications sent	1.12	7.81	0.93
	(0.13)	(1.44)	(0.24)
Mean applications sent, ignoring applications to additional jobs	0.89	7.52	0.90
	(0.12)	(1.44)	(0.24)
Search Outcomes			
Fraction with at least one offer	0.110	0.280	0.082
	(0.007)	(0.043)	(0.013)
Fraction with at least one unsolicited offer	0.032	0.007	0.037
	(0.004)	(0.008)	(0.009)
Fraction with at least one offer, including unrealized offers	0.144	0.304	0.099
	(0.008)	(0.044)	(0.014)
Search Outcomes, Ignoring Offers for Additional Jobs			
Fraction with at least one offer	0.095	0.280	0.082
	(0.007)	(0.043)	(0.013)
Fraction with at least one unsolicited offer	0.031	0.007	0.037
	(0.004)	(0.008)	(0.009)
Fraction with at least one offer, including unrealized offers	0.132	0.304	0.099
	(0.008)	(0.044)	(0.014)
N	1,879	108	451

Note: Estimates come from authors' tabulations from the 2013-15 panel of the SCE labor survey, using respondents' labor force status reported in either the August or September waves of the monthly SCE survey, for all individuals aged 18-64, excluding the self-employed. Standard errors are in parentheses.

C Additional Empirical Results

C.1 Results Conditional on Observable Characteristics

In the main analysis, we explore how much wage differentials between the employed and non-employed change when we add controls for observable worker and job characteristics to the offer wage, the wage at the time of hiring, and the previous wage of the currently employed. This subsection examines how much of a gap exists for other job characteristics, and how much differences in search effort and search outcomes by labor force status persist, after controlling for observable worker and job characteristics for these estimates as well.

Table C1: Search Effort by Labor Force Status, Conditional on Observable Worker and Job Characteristics

	Employed				**Out of**
	Looking for Work	**Not Looking**	**All**	**Unemployed**	**Labor Force**
Labor Force Status at Time of Survey					
Hours spent searching, last 7 days	4.33	0.07	1.17	8.56	0.07
	(0.26)	(0.03)	(0.08)	(0.73)	(0.07)
Mean applications sent, last 4 weeks	4.70	-0.05	1.19	8.06	0.28
	(0.47)	(0.04)	(0.13)	(1.24)	(0.08)
N	502	1,503	2,005	160	419
Labor Force Status in Prior Month					
Mean applications sent			1.16	10.52	0.62
			(0.13)	(1.72)	(0.11)
Mean applications sent, ignoring applications to additional jobs			0.91	10.61	0.59
			(0.13)	(1.72)	(0.11)
N			2,018	116	450

Notes: Estimates come from authors' tabulations from the October 2013-15 waves of the SCE Labor Supplement, for all individuals aged 18-64, excluding the self-employed, by detailed labor force status. The top panel reports results by labor force status at the time of the survey, while the bottom panel reports the results by labor force status in the prior month. Standard errors are in parentheses. See appendix text for set of observable worker characteristics used as controls. Observable job characteristics include two-digit SOC occupation, two-digit NAICS industry, and six categories of firm size. Controls also include fixed effects for survey year and state.

We begin by examining the differences in search effort and search outcomes by labor force status after controlling for observables. Throughout the exercise, our worker controls include sex, age, age squared, four education categories, four race categories, a dummy for homeownership, the number of children under age six in the household, marital status, and marital status interacted with sex. The job controls include the two-digit SOC occupation of the job, six categories of the

Table C2: Search Outcomes by Labor Force Status, Conditional on Observable Worker and Job Characteristics

	Employed			Unemployed	Out of Labor Force
	Looking for Work	Not Looking	All		
Labor Force Status at Time of Survey					
Mean contacts received	1.840	0.345	0.732	0.935	0.286
	(0.262)	(0.042)	(0.075)	(0.231)	(0.053)
Mean unsolicited contacts	0.799	0.256	0.397	0.526	0.205
	(0.121)	(0.031)	(0.039)	(0.148)	(0.035)
Mean job interviews (2014-15)	0.432	-0.040	0.105	0.399	0.055
	(0.042)	(0.007)	(0.012)	(0.105)	(0.018)
Mean offers	0.440	0.092	0.183	0.308	0.056
	(0.038)	(0.011)	(0.013)	(0.072)	(0.027)
Mean unsolicited offers	0.058	0.045	0.048	0.045	0.049
	(0.010)	(0.009)	(0.007)	(0.017)	(0.023)
Fraction with at least one offer	0.304	0.059	0.123	0.188	0.027
	(0.019)	(0.006)	(0.007)	(0.031)	(0.010)
Fraction with at least one unsolicited	0.046	0.028	0.033	0.043	0.022
offer	(0.009)	(0.004)	(0.004)	(0.016)	(0.008)
Fraction with at least one offer,	0.350	0.091	0.158	0.209	0.048
including unrealized offers	(0.023)	(0.009)	(0.01)	(0.032)	(0.012)
N	502	1,503	2,005	160	419
Labor Force Status in Prior Month					
Fraction with at least one offer			0.107	0.315	0.069
			(0.007)	(0.041)	(0.012)
Fraction with at least one unsolicited			0.031	0.039	0.032
offer			(0.004)	(0.016)	(0.008)
Fraction with at least one offer,			0.145	0.328	0.083
including unrealized offers			(0.008)	(0.042)	(0.013)
Labor Force Status in Prior Month, Ignoring Search Outcomes for Additional Jobs					
Fraction with at least one offer			0.090	0.322	0.073
			(0.006)	(0.041)	(0.012)
Fraction with at least one unsolicited			0.029	0.041	0.034
offer			(0.004)	(0.016)	(0.008)
Fraction with at least one offer,			0.131	0.335	0.087
including unrealized offers			(0.007)	(0.042)	(0.013)
N			2,018	116	450

Notes: Estimates come from tabulations from the October 2013-15 waves of the SCE Labor Supplement, for all individuals aged 18-64, excluding the self-employed, with at least one job offer in the last six months. Standard errors are in parentheses. See appendix text for set of observable worker characteristics used as controls. Observable job characteristics include two-digit SOC occupation, two-digit NAICS industry, and six categories of firm size. Controls also include fixed effects for survey year and state.

job's firm size, and, when available, the two-digit NAICS industry of the firm. State and year fixed effects are included throughout as well.

Tables C1, C2, and C3 correspond to Tables 3, 5, and 6 in the main text. In general, controlling for observable characteristics does little to alter the original results in the main text. Table C1 shows that search effort is practically unchanged regardless of the measure used or timing of the measurement of labor force status. The search effort of the employed relative to the unemployed, ignoring search for additional work (our preferred measure of relative effort for our model calibration) falls slightly from 0.91 to 0.86.

Tables C2 and C3 show that search outcomes and acceptance rates also change little after controlling for observables. If we ignore the effects of censoring of the wage offer distribution, we can infer the relative search efficiency of the employed to the unemployed directly from the data as $\lambda_i(s) = \alpha_i + \beta_i s$. Using the unconditional estimates from Table 5 in the main text suggests that relative efficiency measured this way is about 2.2 (compared to nearly three in the calibrated model). If we were to instead use the estimates from Table C2, the estimates suggest that the relative efficiency under this method is slightly higher, at 2.5.

Table C4 reports the characteristics of the best job offer for the employed and non-employed after controlling for observable characteristics. Controlling for observables leads to only modest reductions in the observed gaps in job offer characteristics between the employed and non-employed.

Finally, Table C5 reports the characteristics of the current job for those hired from either employment or non-employment after controlling for observable characteristics. Controlling for observables in this case leads to somewhat larger reductions in the observed gaps in job characteristics, but again, the gaps remain statistically significant and quantitatively similar. The one exception is median tenure, whose gap shrinks considerably and becomes statistically insignificant.

C.2 Detailed Distributions of Search Effort

Figure C1 replicates the results for the distribution of search effort from Figure 1 using finer bins for each search effort measure. The greater detail shows that our main result from Figure 1 holds: the search effort of the employed is weighted more towards lower-levels of effort relative to the unemployed.

Table C3: Acceptance Decisions by Labor Force Status in Previous Month, Conditional on Observable Worker and Job Characteristics

	Employed	**Unemployed**	**Out of Labor Force**
Percent of best offers accepted	0.344	0.463	0.321
	(0.024)	(0.060)	(0.051)
Percent of all offers accepted	0.289	0.419	0.324
	(0.023)	(0.060)	(0.047)
Percent of best offers accepted,	0.324	0.493	0.332
ignoring offers for an additional job	(0.024)	(0.061)	(0.050)
Percent of all offers accepted,	0.405	0.508	0.503
ignoring offers for an additional job	(0.022)	(0.055)	(0.041)
N	187	35	28

Notes: Estimates come from tabulations from the October 2013-15 waves of the SCE Labor Supplement, for all individuals aged 18-64, excluding the self-employed, with at least one job offer in the last six months. Standard errors are in parentheses. See appendix text for set of observable worker characteristics used as controls. Observable job offer characteristics include two-digit SOC occupation, and six categories of firm size. Controls also include fixed effects for survey year and state.

Figure C1: Distribution of Number of Applications Sent in the last Four Weeks (left panel) and Search Time in Hours in the Last Seven Days (right panel) by Labor Force Status

Notes: Figure reports the detailed histograms of the number of applications sent in the last four weeks (top panel) and the hours of time spent searching for work in the last seven days (bottom panel). Estimates are for all individuals, excluding the self-employed, in the 2013-15 labor supplements of the SCE.

C.3 Search Effort-Wage Gradient with Different Measures of Job Search

Figures C2 and C3 replicate our estimates of the search effort-wage gradient with alternative measures of search. Figure C2 replicates our results using the residualized real current wage. The top two panels are identical to the results from Figure 2 from the main text. The bottom two panels report the estimates for the incidence of search using measures of whether an individual sent any applications in the last four weeks or spent any time looking for work in the last seven

Table C4: Characteristics of Best Job Offer by Labor Force Status, Conditional on Observable Worker and Job Characteristics

	Employed at Offer	**Non-Employed at Offer**	**Difference, E - NE**
log offer usual hours	3.401	3.350	0.050
	(0.025)	(0.033)	(0.057)
Pct. of offers with no benefits	41.3	57.7	-16.4
	(1.7)	(2.9)	(4.1)
Pct. of offers through an unsolicited	25.9	15.5	10.4
contact	(1.8)	(2.4)	(4.1)
Pct. of respondents with at least a 'good	55.6	57.4	-1.8
idea' of pay	(1.9)	(3.1)	(4.7)
Pct. of offers with some counter-	13.7		
offer given	(1.4)		
Pct. of offers that involved	38.4	23.9	14.5
bargaining	(1.9)	(2.9)	(4.6)
Pct. of job offers accepted	36.1	52.2	-16.1
	(1.8)	(2.8)	(5.1)
Pct. of offers accepted as	6.8	29.2	-22.4
only option	(1.6)	(3.0)	(3.4)
N	488	164	

Notes: Estimates come from tabulations from the October 2013-15 waves of the SCE Labor Supplement, for all individuals aged 18-64, excluding the self-employed, with at least one job offer in the last six months. Standard errors are in parentheses. See appendix text for set of observable worker characteristics used as controls. Observable job offer characteristics include two-digit SOC occupation and six categories of firm size. Controls also include fixed effects for survey year and state.

days. Figure C3 replicates all four exercises using the unconditional rather than the residualized real current wage. All results in both figures show a strong negative relationship between search effort and the current wage.

C.4 Differentials between the Starting and Previous Wage

Figure C4 illustrates the wage differences between those hired from employment and those hired from non-employment for their full wage distributions. It plots the (log) differences in the real starting wage, relative to the real previous wage, for each group, after controlling for observable worker and job characteristics. The relative wage distribution of those hired from employment stochastically dominates the distribution of those hired from non-employment. The figure also shows, however, that a sizable fraction of hires move directly to a lower-wage job and a sizable fraction receive a higher wage after a spell of non-employment. Nevertheless, the left panel of

Table C5: Selected Characteristics of Current Job, by Labor Force Status at Time of Hire, Conditional on Observable Worker and Job Characteristics

	Hired from, Employment	**Hired from Non-Employment**	**Difference, E - NE**
log real current wage	3.080	2.956	0.124
	(0.013)	(0.018)	(0.026)
log offer usual hours	3.666	3.573	0.093
	(0.009)	(0.017)	(0.019)
Median tenure (mos.)	71.4	76.0	-4.6
	(2.3)	(3.4)	(4.1)
Pct. with no benefits	17.8	26.9	-9.1
	(0.8)	(1.6)	(1.8)
Pct. actively searched for work,	26.9	28.6	-1.7
last four weeks	(1.2)	(1.8)	(2.4)
N	1,230	523	

Notes: Estimates come from tabulations from the October 2013-15 waves of the SCE Labor Supplement, for all individuals aged 18-64. Sample is all currently employed, excluding the self-employed with a reported labor force status at the time of hire and reported current, starting, and previous-job wages and hours. Standard errors are in parentheses. See appendix text for set of observable worker characteristics used as controls. Observable job characteristics include two-digit SOC occupation, two-digit NAICS industry, and six categories of firm size. Controls also include fixed effects for survey year and state.

the figure shows that, after conditioning out our controls, those who transition directly from employment receive an 8 log point increase in their wage, on average, while those who were non-employed receive a 13 log point decrease in their wage, on average. The right panel shows that, without any controls for observable worker and job characteristics, the average (log) wage increase for those hired from employment rises from 8 to 9 log points and the average (log) wage decrease for those hired from non-employment changes from 13 to 15 log points.

Figure C2: Job Search Effort by the Current Wage (Residualized)

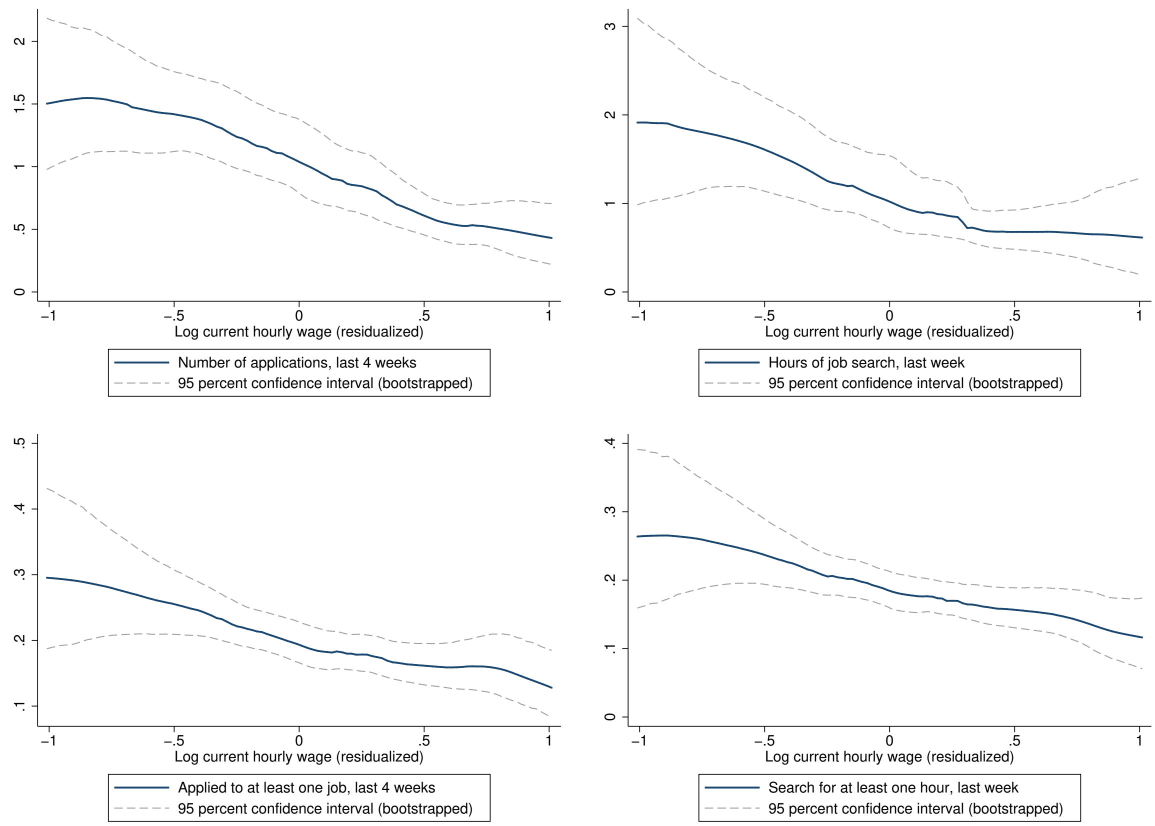

Note: Figure reports the LOWESS estimates (with smoothing parameter 0.8) of the relationship between the measures of search effort listed on each vertical axis and the (log) real current wage of the employed, residualized after controlling for observable worker characteristics (see Table 4 for the list of specific variables). The confidence intervals are based on a bootstrap with 500 replications. The estimates use all employed individuals, excluding the self-employed, age 18-64 from the October 2013-15 waves of the SCE Labor Supplement. Dashed lines represent 95 percent confidence intervals.

Figure C3: Job Search Effort by the Current Wage (Raw)

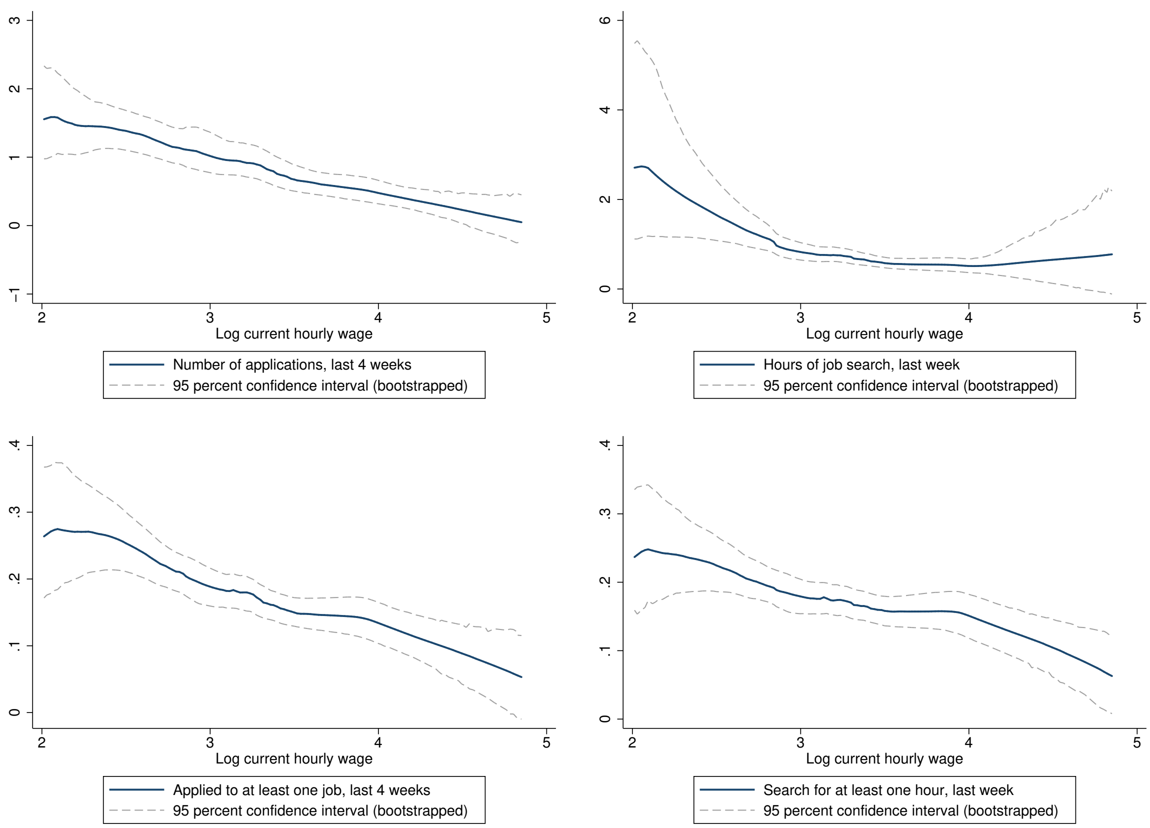

Note: Figure reports the LOWESS estimates (with smoothing parameter 0.8) of the relationship between the measures of search effort listed on each vertical axis and the (log) real current wage of the employed. The confidence intervals are based on a bootstrap with 500 replications. The estimates use all employed individuals, excluding the self-employed, age 18-64 from the October 2013-15 waves of the SCE Labor Supplement. Dashed lines represent 95 percent confidence intervals.

Figure C4: Distribution of Starting Wages Relative to Previous Wage among the Currently Employed Conditional on Observables (left panel) and Without Controls (right panel)

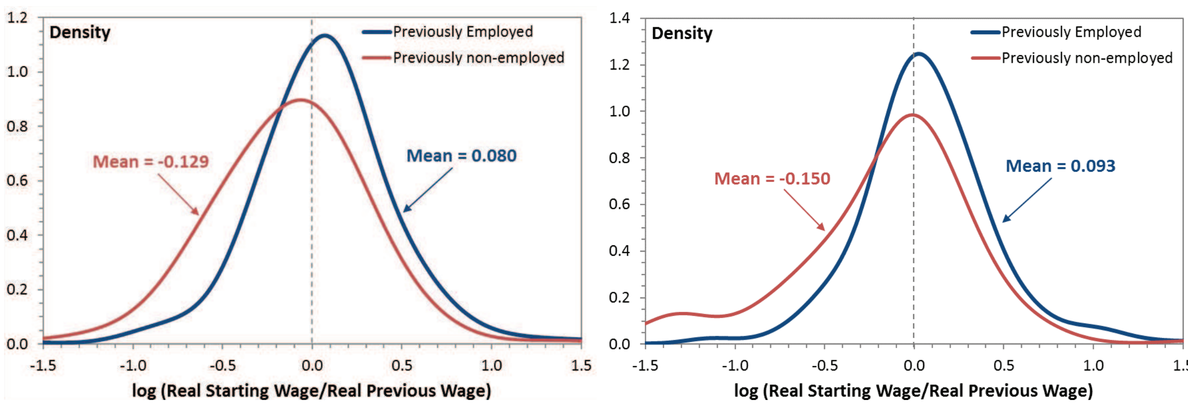

Notes: Figure reports kernel density estimates of the residual of log(real starting wage/real previous wage), where the previous wage refers to final wage of the prior job and the starting wage is for the current job. Estimates are for the sample of the currently employed (excluding self-employed) in the 2013-15 labor supplements of the SCE.

D Robustness of Model Results

D.1 Results for Alternative Models

We calibrate three simpler versions of our benchmark model to evaluate and better understand the implications of our empirical findings for the model's key properties. The first version is a standard model of on-the-job search with exogenous search effort, where we set search effort to unity, ignore the role of unsolicited job offers, and assume that the employed and unemployed draw from the same wage offer distribution. The second version extends the basic model to include endogenous search effort and unsolicited job offers but maintains the assumption of homogeneous workers and an identical wage offer distribution for the employed and unemployed. The third version of the model allows for *ex ante* worker heterogeneity in productivity. Finally our benchmark model from the main text allows for *ex ante* heterogeneity in worker productivity as well as differential wage offer distributions between the employed and unemployed. Table D1 summarizes the calibration and results. The alternative versions of our benchmark model all fail to match the offer acceptance rate of the employed. They also produce less wage dispersion than what is observed in the data and imply a lower flow value of unemployment.

D.2 Robustness Results

Table D2 reports the results of our robustness exercises detailed in the main text. Table D3 replicates our robustness exercises for the version of the model where the employed and unemployed draw from the same wage offer distribution (corresponding to the model results in column (3) of Table D1). Within each table, column (1) repeats the results from our benchmark model for comparison while columns (2)-(6) report the results of our five robustness exercises described in the main text. Column (2) reports the results of our alternative calibration for the censoring of wage offers; column (3) reports the results of increasing the dispersion of wage offers; column (4) reports the results of our alternative calibration of the curvature of the search cost function; column (5) reports the results of our alternative calibration of the job-finding rate; and column (6) reports the results of our use of the unconditional data moments in the calibration.

We focus on the results in Table D2. Robustness results for the remaining two versions of our model are available upon request.

We examine the results of five robustness exercises in order to assess our our model's quantitative implications. We implement the robustness exercises on our baseline model and a version of the model where wage offers do not vary by employment status (model 1 in Table 14) to examine

Table D1: Calibrated Parameter Values and Model Simulation Results

	Data	Benchmark model: Endogenous search effort + worker het. + diff. offer dist.	Restricted model versions: (1) Endogenous search effort + worker heterogeneity	(2) Endogenous search effort	(3) Exogenous offers
κ_u		0.17	0.30	0.31	—
κ_e		3.30	2.32	2.27	—
α_u		0.03	0.03	0.03	—
α_e		0.05	0.05	0.05	—
β_u		0.32	0.32	0.32	0.35
β_e		0.93	0.95	0.94	0.13
χ_u		0.06	0.06	0.06	0.06
χ_e		0.39	0.37	0.36	0.35
b		1.35	1.01	1.02	0.47
$\mu_{y,e} - \mu_{y,u}$		0.16	0.00	0.00	0.00
$x_{max} - x_{min}$		1.08	1.08	—	—
$\delta(x_{min}) - \delta(x_{max})$		0.0021	0.0009	—	—
Targeted moments (means)					
Search effort of unemployed	1	1	1	1	—
Search effort of employed	0.091	0.091	0.091	0.091	—
Unsolicited offer rate of unemployed	0.029	0.029	0.029	0.029	—
Unsolicited offer rate of employed	0.031	0.031	0.031	0.031	—
Formal offer rate of unemployed	0.339	0.339	0.339	0.339	0.339
Formal offer rate of employed	0.091	0.091	0.091	0.091	0.091
Unrealized offer rate of unemployed	0.010	0.010	0.010	0.010	0.010
Unrealized offer rate of employed	0.041	0.041	0.041	0.041	0.041
Acceptance rate of unemployed	0.532	0.532	0.532	0.532	0.532
Residual offered wage differential (E - U)	0.25	0.25	0.04	0.02	0.01
St. dev. of log residual offered wages	0.67	0.67	0.67	0.24	0.24
Residual prior wage differential (E - U)	-0.10	-0.10	-0.10	-0.11	-0.07
Additional moments					
Mean acceptance rate of employed	0.30	0.30	0.215	0.213	0.176
Decomposition of offer wage differential	0.25	0.25	0.04	0.02	0.01
- due to worker-heterogeneity		0.07	0.02	0.00	0.00
- due to censoring		0.02	0.02	0.02	0.01
- due to exogenous differential		0.16	0.00	0.00	0.00
Mean search cost of unemployed		0.17	0.30	0.31	—
Mean search cost of employed		0.06	0.04	0.04	—
$b/E(w)$		0.81	0.69	0.69	0.33
$b/E(w)$ (net of search costs)		0.74	0.50	0.50	0.33
Mean-min ratio (conditional on x)		1.68	1.48	1.48	1.43

Figure D1: On-the-Job Search Effort by Current Wages (Model with $\gamma = 2.5$ vs. Data)

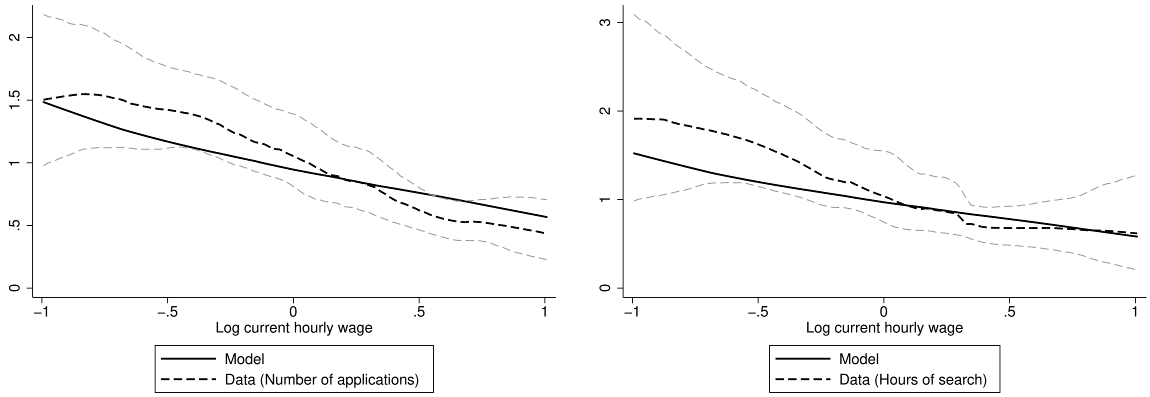

Note: The dashed lines show the 95 percent confidence interval (bootstrapped with 500 replications). For the purposes of the model simulation for these figures, the model was extended from 2 to 11 types.

whether any of these exercises can further account for any of the observed wage offer differential.

1. Alternative calibration strategy for censoring of offers: See main text for a description of the exercise and results.

2. Variance of wage offers: If we calibrate the variance of wage offers to be twice as high (i.e., $\sigma_y = 0.34$), the implied flow value for unemployment is 0.52 and the implied wage dispersion, as measured by the mean-min ratio, is naturally larger under this calibration of the model, increasing from 1.68 to 1.94, as column (2) in Table D3 shows. This version of the model underpredicts the acceptance rate of the employed (0.26 vs 0.30 in the data) and attributes a higher fraction of the wage offer differential to heterogeneity (10 log points vs 7 in baseline). The unexplained wage offer differential is lower than in the benchmark model (12 log points vs. 16 log points).

3. Alternative calibration strategy for the curvature of the search cost function: We also assess to what extent our model matches the observed distribution of job search effort among the employed. We focus on search effort at the bottom of the wage ladder, since it determines how quickly a worker transitions up the ladder should she choose to accept a low-paying job. Our baseline model implies a job search effort of 0.27 (relative to the search effort of the unemployed) at the 95th percentile of the stationary effort distribution, which is somewhat below the 0.38 value observed in the data. The shape of the search effort distribution for the employed is driven by the parameter γ, which determines the convexity of the search cost function. If we set $\gamma = 2.5$, we are able to roughly match the observed search effort at the 95th percentile of the distribution. These versions of the model generate more search at the bottom of the wage ladder and less at the top, consistent with our empirical findings. Implied flow values for unemployment and mean-min ratios are little changed. Finally, this version of the model performs even better in matching the observed relationship between search intensity and the current wage, as shown in Figure D1. [43]

4. Alternative calibration of the job-finding rate: The offer and acceptance rates in our calibration imply a monthly unemployment-to-employment transition rate of 18 percent. This is lower than the transition rate used in Hornstein, Krusell, and Violante (2011), which they target directly. This is partly driven by their sample period (1994-2007) and by their use of total unemployment outflows (both to employment and out of the labor force) when calculating the transition rate. To see how this affects our calibration, we increase the offer rate of the unemployed to 0.60, from 0.339, which implies the unemployment-to-employment transition rate of 0.32 targeted by Hornstein, Krusell, and Violante (2011). The results indicate a flow value of unemployment of 0.58 with a mean-min ratio of 1.69. Thus, our resolution of the puzzle holds.

5. Targeting raw moments instead of residualized differences: Since our model does not feature observed heterogeneity such as education and gender, our baseline calibration targets data moments that control for observable worker and job characteristics. As a robustness check, we also consider a calibration where we target the raw, unconditional moments in our data. In this case, the empirical wage offer gap between the employed and unemployed is 0.40. While the importance of heterogeneity obviously rises relative to our baseline in this case, nearly half of the wage offer differential remains unexplained by either worker heterogeneity or censoring of the wage offer distribution. In this case, we find that censoring accounts for 2 log points and unobserved heterogeneity accounts for 19 log points of the unconditional 40 log point wage offer differential. This leaves about half of the differential (19 log points out of 40) unaccounted for, compared to nearly two-thirds unaccounted for (16 log points out of 25) in our benchmark exercise.

[43] Another robustness check we performed was to target the acceptance rate of all offers as opposed to the acceptance rate of best offers, which ignores the implicit rejection of offers received along with the best offer. While the acceptance rate of all offers is lower for the employed, so is the acceptance rate for the unemployed, and thus the model does no better in matching the acceptance rate of the employed. We believe that targeting the acceptance rate of best offers as our baseline is consistent with using the fraction receiving at least one offer in a given month as our measure of the job offer probability.

Table D2: Alternative Calibrations and Robustness Checks for the Model with Worker-heterogeneity and Different Wage Offer Distributions

	Data	Benchmark	Robustness				
		(1)	(2)	(3)	(4)	(5)	(6)
Calibrated parameter values							
κ_u		0.17	0.17	0.35	0.18	0.30	0.15
κ_e		3.30	3.25	4.95	1.66	3.28	3.49
α_u		0.03	0.03	0.03	0.03	0.03	0.03
α_e		0.05	0.05	0.05	0.05	0.05	0.05
β_u		0.32	0.31	0.32	0.32	0.58	0.32
β_e		0.93	0.92	0.94	0.93	0.93	0.93
χ_u		0.06	0.00	0.06	0.06	0.03	0.06
χ_e		0.39	0.39	0.38	0.39	0.39	0.40
b		1.35	1.34	1.30	1.41	1.23	1.41
$\mu_{y,e} - \mu_{y,u}$		0.16	0.15	0.12	0.14	0.16	0.19
$x_{max} - x_{min}$		1.08	1.08	0.96	1.08	1.08	1.57
$\delta(x_{min}) - \delta(x_{max})$		0.0021	0.0023	0.0045	0.0018	0.0029	0.0042
Targeted moments (means)							
Search effort of unemployed	1	1	1	1	1	1	1
Search effort of employed	0.091	0.091	0.091	0.091	0.091	0.091	0.091
Unsolicited offer rate of unemployed	0.029	0.029	0.029	0.029	0.029	0.029	0.029
Unsolicited offer rate of employed	0.031	0.031	0.031	0.031	0.031	0.031	0.031
Offer rate of unemployed	0.339	0.339	0.339	0.339	0.339	0.600	0.339
Offer rate of employed	0.091	0.091	0.091	0.091	0.091	0.091	0.091
Unrealized offer rate of unemployed	0.010	0.010	0.000	0.010	0.010	0.010	0.010
Unrealized offer rate of employed	0.041	0.041	0.040	0.041	0.041	0.041	0.041
Acceptance rate of unemployed	0.532	0.532	0.532	0.532	0.532	0.532	0.532
Residual offered wage differential (E - U)	0.25	0.25	0.25	0.25	0.25	0.25	0.40
St. dev. of log residual offered wages	0.67	0.67	0.67	0.67	0.67	0.67	0.88
Residual prior wage differential (E - U)	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	0.02
Additional moments							
Mean acceptance rate of employed	0.300	0.300	0.300	0.262	0.310	0.301	0.314
Decomposition of offer wage differential	0.25	0.25	0.25	0.25	0.25	0.25	0.40
- due to worker-heterogeneity		0.07	0.07	0.10	0.08	0.07	0.19
- due to censoring		0.02	0.03	0.03	0.03	0.03	0.02
- due to exogenous differential		0.16	0.15	0.12	0.14	0.16	0.19
Distribution of Employed Job Search Effort:							
95th Percentile	0.38	0.27	0.27	0.25	0.40	0.27	0.27
90th Percentile	0.19	0.22	0.22	0.21	0.28	0.22	0.22
75th Percentile	0.00	0.14	0.14	0.14	0.12	0.13	0.14
50th Percentile	0.00	0.07	0.07	0.07	0.03	0.07	0.06
Unemployment rate of low-x type		0.089	0.089	0.095	0.093	0.052	0.100
Unemployment rate of high-x type		0.065	0.065	0.058	0.061	0.038	0.053
Mean search cost of unemployed		0.17	0.17	0.36	0.19	0.31	0.15
Mean search cost of employed		0.06	0.06	0.09	0.09	0.06	0.07
$b/E(w)$		0.81	0.81	0.68	0.85	0.74	0.82
$b/E(w)$ (net of search costs)		0.74	0.74	0.52	0.78	0.58	0.77
MeanMin-ratio (conditional on x)		1.68	1.69	1.94	1.69	1.69	1.74

Note: (2) Targeting the acceptance rate of employed with χ_e; (3) $\sigma_y = 0.34$; (4) $\gamma = 2.5$; (5) Alternative calibration of the job-finding rate; (6) Targetinng raw moments (not residualized).

Table D3: Alternative Calibrations and Robustness Checks for the Model with Worker Heterogeneity

	Data	Benchmark	Robustness				
		(1)	(2)	(3)	(4)	(5)	(6)
Calibrated parameter values							
κ_u		0.30	0.19	0.49	0.35	0.56	0.30
κ_e		2.32	2.45	3.88	1.20	2.31	2.30
α_u		0.03	0.03	0.03	0.03	0.03	0.03
α_e		0.05	0.09	0.05	0.05	0.05	0.05
β_u		0.32	0.31	0.32	0.32	0.58	0.32
β_e		0.95	1.76	0.94	0.94	0.94	0.95
χ_u		0.06	0.00	0.06	0.06	0.03	0.06
χ_e		0.37	0.72	0.36	0.37	0.37	0.36
b		1.01	1.24	0.99	1.15	0.79	1.01
$\mu_{y,e} - \mu_{y,u}$		0.00	0.00	0.00	0.00	0.00	0.00
$x_{max} - x_{min}$		1.08	1.08	0.96	1.08	1.08	1.57
$\delta(x_{min}) - \delta(x_{max})$		0.0009	0.0019	0.0038	0.0012	0.0012	0.0047
Targeted moments (means)							
Search effort of unemployed	1	1	1	1	1	1	1
Search effort of employed	0.091	0.091	0.091	0.091	0.091	0.091	0.091
Unsolicited offer rate of unemployed	0.029	0.029	0.029	0.029	0.029	0.029	0.029
Unsolicited offer rate of employed	0.031	0.031	0.031	0.031	0.031	0.031	0.031
Offer rate of unemployed	0.339	0.339	0.339	0.339	0.339	0.600	0.339
Offer rate of employed	0.091	0.091	0.091	0.091	0.091	0.091	0.091
Unrealized offer rate of unemployed	0.010	0.010	0.000	0.010	0.010	0.010	0.010
Unrealized offer rate of employed	0.041	0.041	0.163	0.041	0.041	0.041	0.041
Acceptance rate of unemployed	0.532	0.532	0.532	0.532	0.532	0.532	0.532
Residual offered wage differential (E - U)	0.25	0.04	0.12	0.09	0.05	0.04	0.15
St. dev. of log residual offered wages	0.67	0.67	0.67	0.67	0.67	0.67	0.88
Residual prior wage differential (E - U)	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	0.02
Additional moments							
Mean acceptance rate of employed	0.300	0.215	0.300	0.215	0.230	0.217	0.214
Decomposition of offer wage differential:	0.25	0.04	0.12	0.09	0.05	0.04	0.15
- due to worker-heterogeneity		0.02	0.05	0.06	0.03	0.02	0.13
- due to censoring		0.02	0.08	0.03	0.02	0.02	0.02
- due to exogenous differential		0.00	0.00	0.00	0.00	0.00	0.00
Distribution of Employed Job Search Effort:							
95th Percentile	0.38	0.26	0.27	0.25	0.39	0.26	0.26
90th Percentile	0.19	0.21	0.22	0.21	0.28	0.21	0.22
75th Percentile	0.00	0.14	0.14	0.14	0.12	0.14	0.14
50th Percentile	0.00	0.07	0.06	0.07	0.03	0.07	0.07
Unemployment rate of low-x type		0.080	0.085	0.088	0.082	0.047	0.091
Unemployment rate of high-x type		0.074	0.069	0.066	0.071	0.043	0.062
Mean search cost of unemployed		0.30	0.19	0.49	0.35	0.56	0.31
Mean search cost of employed		0.04	0.05	0.07	0.06	0.04	0.04
$b/E(w)$		0.69	0.80	0.58	0.77	0.54	0.69
$b/E(w)$ (net of search costs)		0.50	0.70	0.30	0.56	0.16	0.49
Mean-min-ratio (conditional on x)		1.48	1.58	1.75	1.50	1.49	1.48

Note: (2) Targeting the acceptance rate of employed with χ_e; (3) $\sigma_y = 0.34$; (4) $\gamma = 2.5$; (5) Alternative calibration of the job-finding rate; (6) Targeting raw moments (not residualized).