Labor Market Search With Imperfect Information and Learning*

John J. Conlon Laura Pilossoph Matthew Wiswall
Basit Zafar †

August 2018

Abstract

We investigate the role of information frictions in the US labor market using a new nationally representative panel dataset on individuals' labor market expectations and realizations. We find that expectations about future job offers are, on average, highly predictive of actual outcomes. Despite their predictive power, however, deviations of ex post realizations from ex ante expectations are often sizable. The panel aspect of the data allows us to study how individuals update their labor market expectations in response to such shocks. We find a strong response: an individual who receives a job offer one dollar above her expectation subsequently adjusts her expectations upward by $0.47. The updating patterns we document are, on the whole, inconsistent with Bayesian updating. We embed the empirical evidence on expectations and learning into a model of search on- and off- the job with learning, and show that it is far better able to fit the data on reservation wages relative to a model that assumes complete information. The estimated model indicates that workers would have lower employment transition responses to changes in the value of unemployment through higher unemployment benefits than in a complete information model, suggesting that assuming workers have complete information can bias estimates of the predictions of government interventions. We use the framework to gauge the welfare costs of information frictions which arise because individuals make uninformed job acceptance decisions and find that the costs due to information frictions are sizable, but are largely mitigated by the presence of learning.

JEL Codes: D83, D84, J64
Keywords: expectations, learning, labor search, optimism

*The authors thank Fatih Karahan and Gregor Jarosch for valuable feedback, and thank seminar participants at the 2018 CESifo Subjective Expectations Workshop, the 2018 NBER Summer Institute, Cambridge University, Essex, Northwestern University, NY Fed Brown Bag Series, Princeton Labor Working Lunch, Royal Holloway, SF Fed, UChicago Harris, UCL, and University of Michigan for invaluable comments. All errors that remain are ours. The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System as a whole.

†Conlon: Harvard University, Department of Economics; E-mail: johnconlon@g.harvard.edu. Pilossoph: Research and Statistics, Federal Reserve Bank of New York; E-mail: pilossoph@gmail.com. Wiswall: University of Wisconsin-Madison, Department of Economics; E-mail: matt.wiswall@gmail.com. Zafar: Arizona Status University, Department of Economics; E-mail: basitak@gmail.com

1 Introduction

When individuals decide on whether or not to accept a job opportunity, expectations play a central role. In particular, individuals with a job offer in hand weigh whether to accept it, or to reject it and wait for a better opportunity. This decision process takes center stage in the hallmark models of labor market search and unemployment (Pissarides, 1985; Mortensen and Pissarides, 1994; Burdett and Mortensen, 1998), and key to these calculations are workers' beliefs about the likely value of future offers. These setups and nearly all research in this area (with a few notable exceptions, which we turn to below) have in common the default assumption that workers' beliefs are correct: workers know the distribution of wages from which they draw job offers. These "complete information" assumptions are often made, yet rarely tested, because from standard observational data on employment and wage outcomes, it is difficult to separate heterogeneity in expectations from heterogeneity in productivity and received offers.

This paper offers novel empirical evidence on information frictions in the labor market, and uses this data, in combination with labor market models, to quantitatively assess their welfare consequences. The core of our study is a collection of new panel data on workers' expectations (about the number and value of future job offers) along with actual employment outcomes from a representative survey of US household heads. The unique expectations and realizations data come from a series of new questions we designed and added to the Federal Reserve Bank of New York's Survey of Consumer Expectations (SCE), Labor Market Survey. Our data cover the July 2014 to November 2017 waves, and include over 4,300 unique individuals surveyed every 4 months for up to one year, for a total of 8,883 person-wave observations over 11 survey waves. Modeled after the Current Population Survey (CPS), the panel survey includes the standard information on employment and wages, and a collection of new questions we designed on expectations about future job arrivals and wage offers. The panel structure of our data allows us to study not only the distribution of individual expectations at a point in time, but also how each individual's expectations relate to realizations in the following 4-month period (e.g., wage offers), allowing for an analysis of the accuracy of the expectations and how individuals learn from their labor market experiences. In addition, we collect detailed information not only on accepted offers, but also on offers individuals end up rejecting which, to the best of our knowledge, makes our dataset unique.

We begin by showing empirically how beliefs about the labor market relate to actual realizations. We then further exploit the panel aspect of the data to show how beliefs are updated over time in relation to individual labor market experiences. We find that workers' beliefs predict subsequent outcomes. For example, the average respondent's *expected* wage offer (over the next four months) is \$32.3 per hour. This is quite close to the average *actual* received offer wage, \$30.4 per hour, suggesting that these expectations data have useful informational content. How-

ever, "shocks" – the difference between realized and expected salary offers – are often sizable.[1] Although the median shock is close to zero (-$0.6), the 10th percentile is -$14.6/hour (an over-estimation of the realized hourly wage by $14.6), and the 90th percentile is +$8.2/hour. This heterogeneity in accuracy could play an important role in analyzing heterogeneity in labor market dynamics, which we investigate more formally in the context of the model of labor market search and learning we develop below.

In terms of learning, we find that individuals respond to these labor market shocks in a way that is consistent with the presence of information frictions: a positive shock, indicating a higher than expected salary offer, causes individuals to update their beliefs about future wages upward. The slope of the updating curve is about 0.47; that is, every additional unexpected dollar offered increases expected earnings in the next period by about $0.47. We find little evidence that learning is consistent with standard Bayesian updating. First, when splitting the sample by the precision of reported expectations using elicited information on the individual's prior belief, we find that individuals who are more certain about the expected offer update their beliefs by more in response to a shock. Moreover, relative to the Bayesian benchmark that we construct for each respondent (again using data on the self-reported precision of their prior belief), the updating seems to be excessive– the average slope under Bayesian updating would be about 0.16, substantially lower than what we observe in the data. This over-reliance of respondents on recent information is consistent with individuals using a "representative" heuristic (Grether, 1980).

The second part of our paper incorporates the expectations and learning data into an otherwise standard model of on-the-job search, and uses the model to quantify the welfare costs of information frictions. We introduce expectations and learning by allowing workers to have individual-specific beliefs about the mean of the offer distribution they face, which may differ from the true mean (Burdett and Vishwanath, 1988). Job offers arrive in both unemployment and employment and, as job offers arrive, workers update their beliefs about future expected wage offers using the information contained in the offer they receive, if any. Because an individual's beliefs about the offer distribution affect their job acceptance decisions, those with a relatively pessimistic view of their job offer prospects have lower reservation wages and will accept more offers; those with relatively optimistic beliefs have higher reservation wages, and are more "picky" about the jobs they take.

We use our new data on expectations and learning to estimate the learning rule, and identify and estimate the remaining parameters of the model using the probabilities of receiving offers

[1]The fact that realized wage offers differ from expected wage offers is not necessarily evidence of a "shock". For example, random search models would predict that this type of event happens all the time, even when workers have complete information. However, the fact that we see workers updating their beliefs in response to these events, as we describe below, suggests that the realized deviation contains information useful for future predictions. Hence, we refer to this deviation as a "shock" going forward.

on- and off- the job, as well as observed transition rates to and from employment. In particular, we estimate preference parameters (the flow value of unemployment and the cost of making a job to job transition) via Simulated Method of Moments (SMM) and match the unemployment-to-employment and employment-to-employment transition rates we observe, taking into account the heterogeneity in reported beliefs and learning. We allow rich patterns of heterogeneity in model parameters, including by the worker's education, and find that the parameter estimates one would obtain under the complete information assumption substantially differ from the estimates under our model of incomplete information and learning. This is not surprising given the large degree of heterogeneity in expectations which we document.

We also show that counterfactual changes in the flow value of unemployment (e.g. changes in unemployment benefits) have different implications under a complete information benchmark relative to our model of incomplete information and learning. Under the complete information model, changes in the value of unemployment would have a larger effect on transitions out of unemployment than in our baseline model with heterogeneous beliefs and learning.[2] In our model, unemployed individuals have heterogeneous responses to changes in the value of unemployment because of heterogeneity in their beliefs. The different predictions of the two models indicate that assuming incorrectly that workers have complete information can bias estimates of the effects of government interventions, such as increases in unemployment benefits.

To validate our model, we use data on respondents' self-reported reservation wages (data not used in the model estimation) to investigate how our model-predicted reservation wages compare with reported reservation wages in the data. We find that our model does far better in capturing heterogeneity in reported reservation wages than the predictions of a model estimated assuming complete information, as reported reservation wages have a much higher correlation with the reservation wages predicted by our general model (heterogeneous beliefs and learning) than with reservation wages implied by the complete information model.

In our final exercise, we compute welfare in our estimated model and compare it to welfare under two alternative environments: (1) complete information, and (2) heterogeneous beliefs, but no learning. The former informs us about the importance of information frictions more broadly for welfare, while the latter informs us about the importance of learning for welfare. We find that non-college (college) workers in our sample would be, on average, willing to pay \$175 (\$817) per year to have complete information. Given average annual full-time earnings of \$53,444 (\$83,734) for non-college (college) workers, this is not a very high willingness to pay to remove information frictions. However, we find that the average willingness to pay for complete information is substantially larger when workers are unable to learn– \$1,525 (\$,3,174) per year for non-college (college) workers. That is, learning mitigates most of the damage caused by

[2]As we describe later, this is not a general result, as the difference in predictions depends on the distribution of beliefs in our sample.

information frictions. Although subjective beliefs are quite heterogeneous, and some workers' beliefs are biased, the high rate of learning implies that workers' beliefs converge towards the truth quite quickly.

We next assess how the costs of information frictions compare to those of more commonly studied frictions in the labor market. To do so, we ask how the welfare gains from removing information frictions compare to those of reducing search frictions. We find that the increase in utility when moving to complete information is roughly equivalent to raising the arrival probability of offers in non-employment by 5% (15%) for non-college (college) workers. On the other hand, the gains of moving from the no-learning case to the case with complete information are equivalent to raising the arrival probability of offers in non-employment by about 40% (100%) for the non-college (college) workers. Thus, we find welfare costs arising from information frictions are as important as the traditional search frictions which is the focus of the previous literature. We emphasize that our results are arguably a lower bound on the importance of incomplete information, because we assume that workers are not fully informed about only one aspect – the mean of the offer distribution – of the labor market.

Our paper proceeds as follows. In the next section, we review some related literature and clarify our contributions. Section 3 describes the survey and provides a description of the sample. Section 4 looks at the labor market expectations and realizations data from the survey and analyzes how respondents update their expectations. Section 5 develops a model of on-the-job search that allows for heterogeneous beliefs and learning, and Section 6 outlines the estimation of the model. Section 7 presents the model estimates and validates the model using evidence on reported reservation wages. Section 8 quantifies the welfare of information frictions, and Section 9 concludes.

2 Related Literature

Our work is closely related to the literature on job search among the employed and non-employed, but extends the basic setting to allow for deviations from complete information, as in Burdett and Vishwanath (1988).[3] We build on Burdett and Vishwanath (1988) in three important ways. First, we provide new and rich empirical evidence on the nature of beliefs and how they evolve with heterogeneous labor market experiences which inform our model. Standard data sets usually only contain information on accepted offers and actual transitions; there is typically no information on other offers that individuals may have received and rejected, or on respondents' labor market expectations. Second, we extend the Burdett and Vishwanath (1988) theoretical

[3]Our framework studies beliefs and learning about the mean of the wage offer distribution. For work which studies learning about one's own quality, job finding rate, or productivity, see Gonzalez and Shi (2010) and Doppelt (2016). For a model in which both firms and workers learn about their joint match quality, see Jovanovic (1979).

framework to allow for search on-the-job and a learning rule that is not necessarily Bayesian, but instead is informed by the data on belief updating we collect. Finally, we estimate our model and provide calculations for the welfare costs of information frictions.

Central to our analysis is the data on reported expected wage offers as well as rejected offers. Spinnewijn (2015) also uses elicited expectations from a sample of job seekers in Michigan and Maryland surveyed between 1996 and 1998 by Price et al. (1998), and compares their expectations of being employed to realized employment, finding that unemployed workers tend to overestimate the rate at which they will find jobs, leading them to under-save. In another recent paper, Potter (2017) studies Bayesian learning in a search environment, but workers learn instead about the arrival rate of offers. Potter estimates his model using survey data on unemployed workers in New Jersey during the Great Recession. Like these papers, we also find that unemployed respondents over-estimate offer arrival rates, and provide new evidence that the same pattern holds for employed workers. However, we find no evidence of such a bias in the case of expectations regarding wage offers. Furthermore, we also find that the costs of biased beliefs, both from under-estimating and over-estimating future offers, can be significant, but that these costs are mitigated by the ability of workers to learn over time. Faberman, Meuller, Sahin, and Topa (2017), using a different supplemental cross-sectional annual survey to the same datasource that we use (SCE), study the relative efficacy of search between the employed and non-employed using information on search behavior as well as rejected wage offers. We exploit the additional panel structure in our data to study learning and how beliefs are updated over time in response to labor market outcomes.[4]

We also use the data we have on reported reservation wages as a test of our model against a complete information environment. There are a handful of other papers which also have information on reported reservations wages in addition to accepted offers and labor market transitions. Krueger and Mueller (2016) collect rich panel data from unemployed job seekers in New Jersey. They analyze how reported reservation wages evolve over the spell of unemployment and find that reservation wages for the unemployed start out high, and do not adjust downward enough, providing suggestive evidence that workers "persistently misjudge their prospects or anchor their reservation wage." Le Barbanchon, Rathelot, and Roulet (forthcoming) use administrative data in France, where unemployed individuals must declare their reservation wage in order to claim unemployment insurance, to examine the relationship between potential benefit duration and reported reservation wages, for which they find an elasticity of zero. We instead study the relationship between job-finding rates and changes in the value of unemployment/leisure within

[4]Typically longer panels, that collect data on both expectations and realizations, are needed for such an investigation, which are fairly rare in the context of labor market outcomes. One exception is Stephens (2004) who, using the Health and Retirement Study, finds that job loss expectations are in fact predictive of subsequent job losses. The ex post accuracy of expectations has been investigated in many other contexts, including survival (Hurd and McGarry, 2002) and education (Jacob and Wilder, 2011).

our structural model, and contrast the predictions of our learning model to the predictions of the complete information environment. Our approach thus contrasts with the methods put forth by Lancaster and Chesher (1983), who show how to "deduce" the parameters governing this elasticity using reported reservation wages and expectations rather than estimating them within a structural model. Our more structural approach, however, allows us to ask questions regarding welfare under different counterfactual models of learning.

More broadly, our paper is related to a growing literature which collects and uses subjective expectations data to understand decision-making under uncertainty (see Manski, 2004, for an earlier survey of this literature). Recent work in this area has investigated the importance of information gaps in decision-making. For example, Wiswall and Zafar (2015) embed an information experiment (that provides information on returns associated with different college majors) in a survey that elicits subjective expectations data directly from college students. They find that students are quite misinformed about returns to different majors, and use these data to estimate a model of college major choice. We similarly use our data on expectations to study the importance of information gaps, and quantify the information frictions both in terms of welfare losses and their search friction equivalents. Our results are also consistent with other recent work that uses different ways to show that job seekers may be misinformed about some aspect of the job search environment. For example, Belot, Kircher, and Muller (2016) find that providing job seekers with relevant alternative search opportunities based on actual labor market data - ones that would not have otherwise been utilized - increases their interview prospects. In particular, they show that individual search strategies indeed react to the news provided by the experiment, which is what we might expect if, as the authors state, "seekers lack relevant information". Likewise, using high-frequency panel data on individuals' job applications, Kudlyak, Lkhagvasuren, and Sysuyev (2014) show that job seekers learn over the search process.

To summarize, we believe the paper makes three main contributions. First, to our knowledge, it is the first paper to empirically investigate the nature of learning in the labor market. Rather than assuming homogenous prior beliefs and Bayesian learning, we allow for heterogeneous beliefs and non-Bayesian learning, informed by our direct panel data on beliefs. Second, unlike most prior work in this area that focuses on the unemployed, our sample is fairly representative of the population of US job seekers, and includes unemployed and employed workers. Third, we quantify the importance of information frictions by embedding the subjective data and the data-based updating into a job search model that is disciplined by our novel data.

3 Data

3.1 Survey Design and Administration

Our data come from the Survey of Consumer Expectations (SCE) Labor Market Survey. To the original survey, we added two broad sets of questions: (1) an "Experiences" component that collects data on labor market outcomes, such as offers received in the past 4 months, earnings, search behavior, reservation wages, and labor market transitions, and (2) an "Expectations" component that collects data on expectations regarding future job offers, labor market transitions, and earnings.[5] Respondents also provide information on many demographic variables. When appropriate, questions have built-in logical checks (for instance, percent chances of an exhaustive set of events have to sum to 100). Item non-response is extremely rare, and almost never exceeds one percent for any question.

The SCE is fielded by the Federal Reserve Bank of New York. The SCE is an internet-based monthly survey of a rotating panel of approximately 1,300 household heads from across the US. The survey, as its name suggests, elicits expectations about a variety of economic variables, such as inflation and labor market conditions. Respondents participate in the panel for up to twelve months, with a roughly equal number rotating in and out of the panel each month.[6] Respondents are invited to participate in at least one survey each month.

The survey is conducted online by the Demand Institute, a non-profit organization jointly operated by The Conference Board and Nielsen. The sampling frame for the SCE is based on that used for The Conference Board's Consumer Confidence Survey (CCS). Respondents to the CCS, itself based on a representative national sample drawn from mailing addresses, are invited to join the SCE internet panel. The response rate for first-time invitees is around 55%. Respondents receive $15 for completing each survey.[7]

Active panel members who had participated in an SCE monthly survey in the prior three months are eligible to participate in the Labor Market Survey. The structure of the survey – with both forward-looking and retrospective questions – combined with the panel structure makes this data well-suited to study learning in the context of labor markets. Because respondents are in the SCE panel for up to 12 months, they may end up taking between one and three Labor Market Surveys during their tenure on the panel.

[5] See `https://www.newyorkfed.org/microeconomics/sce/labor#/` for details.

[6] In that sense, the panel differs from most online survey platforms where samples are not refreshed regularly. The concern that the sample consists of enthusiastic survey-takers is then less applicable in our case.

[7] See Armantier et al. (2017) for a technical background of the SCE, and visit `www.newyorkfed.org/microeconomics/sce.html` for additional information.

3.2 Descriptive Statistics

Table 1 shows that our dataset consists of 8,883 total observations, covering the period from July 2014 to November 2017. Our sample includes 4,388 unique individuals, of whom 33.6% have taken 1 labor market survey, 32.2% have taken 2 surveys, and the remainder 3 surveys. The first column of Table 1 shows the characteristics of our sample, while the second column shows the characteristics of household heads in the Current Population Survey (CPS) over the period January 2015-September 2017. Our sample aligns well with the demographic characteristics of the United States household heads along most dimensions. For instance, the average age of our respondents is 45.8 years, 51.8% are males, and 76.8% are white; the corresponding numbers among US household heads in the CPS are 44.7 years, 51.6%, and 77.0%. Our sample, however, is significantly more educated: 57.3% of our respondents have at least a Bachelors' degree, while only 35.9% of the US household heads fall into this category. This may partly be a result of differential internet access and computer literacy across education groups in the US population. We return to this point later when we estimate our model of learning on this sample, where we estimate model parameters separately for college and non-college educated respondents.

Turning to variables related to the labor market, although there are differences between our sample and the CPS sample, we see that our sample compares favorably with national-level statistics. Although our sample is more likely to be working full-time (69.4% versus 64.9%), the unemployment rate is similar. [8] Likewise, for both college and non-college workers, the hourly wage conditional on working full-time in our data is similar to that in the CPS. For example, the average hourly wage for college workers in our data is \$40.3 versus \$42.0 in the CPS. In fact, Appendix Figure A1 shows that the distributions of full-time earnings for college and non-college workers are very similar in the two datasets.

Before proceeding, it is best to clarify some key data choices that we make. First, we convert all earnings variables (realizations and expectations) to July 2017 dollars, based on the CPI. Second, we convert all annual earnings variables into hourly earnings, assuming people work 52 weeks a year and 40 hours a week if full-time and 20 hours a week if part-time. We directly ask whether received offers were part-time or full-time, so we adjust offer salaries accordingly. If someone is working part-time, we assume her beliefs (including reservation wage) are about part-time work. If someone is not working part-time (including non-employed), we assume her beliefs are about full-time work. Third, we drop individuals whose wages (for current job, offers, or expectations) are less than \$4.81/hr (\$10,000 a year, full-time). We also drop respondents whose revisions in beliefs (about earnings, salary wages) between surveys, revisions in realizations, or the gap between the realization and the previous period expectation is greater

[8] In addition, Appendix Table A1 shows the labor force transition rates in the SCE and the CPS. While the horizons for the two surveys differ, 4-months (3-months) for the SCE (CPS), the transition rates in the two samples are comparable.

than \$48.1 (\$100,000 a year, full time). These criteria drop less than 10% of the observations. We also winsorize the top 1% of each of these variables within education groups. That is, within education groups, we set everyone whose value is above the 99th percentile to the 99th percentile.

4 New Evidence on Expectations and Labor Market Dynamics

In this section, we explore our novel data on expectations and labor market dynamics. We analyze two key areas of the labor market: the arrival rate and wage distribution of job offers. The panel structure of our data allows us to study not only the distribution of individual expectations at a point in time, but also how each individual's expectations (expected number and wage of future job offers) relate to realizations in the following 4-month period (actual number and wages of received job offers). This in turn allows an analysis of the accuracy of expectations, how beliefs form, and how individuals learn from their experiences. In the later sections, we use these data to estimate a model of labor market search with heterogeneous beliefs and learning.

4.1 Expectations about Job Offers

We begin our analysis with the arrival rate of job offers, the key distinction between classical models of the labor market where workers are always presented with a job at some prevailing wage and frictional-based search models where job offers are not always available. We asked unemployed and out-of-the-labor-force (OLF) workers the following question about expected job offers:

"*What do you think is the percent chance that within the coming four months, you will receive at least one job offer? Remember that a job offer is not necessarily a job you will accept.*"

For currently employed individuals, the question wording was slightly different, and we asked about job offers from "*another employer.*" Answers to this question can range from 0 percent chance (probability 0) to 100 percent chance (probability 1) of receiving a job offer. In the remainder of the paper we express all probabilistic data as probabilities $[0, 1]$ rather than percent chances $[0, 100]$. In addition to questions about the percent chance of receiving any offer, we also asked all respondents about their expected number of job offers over the next 4 months.

Panel A of Table 2 presents summary statistics on expected job offers. The overall average expected probability of receiving at least one job offer in the next 4 months was 0.32. For employed individuals, it was 0.32. For unemployed workers, the average expected rate was higher

at 0.61, and for OLF individuals the average expectation was 0.19. The difference in expectations is consistent with the higher search intensity for unemployed workers versus employed workers, and lowest search intensity for OLF respondents. The heterogeneity in expectations is apparent from the large standard deviations, reported in parentheses.

The second set of statistics in Panel A is for the expected number of offers in the next 4 months. The average number of expected job offers is 0.8, reflecting that a majority of workers do not expect any job offer at all (the median is 0). Again, following the patterns with the expected arrival rate of any job offer, unemployed workers on average expect more job offers (2.0), followed by employed workers (0.8 offers), and OLF individuals (0.6 offers).

Conditional on reporting a non-zero likelihood of receiving at least one offer in the next four months, respondents are also asked about their expectation regarding salary offers. We denote these individual-specific expectations as $E_i(w_{i,t})$ where the i subscript on the expectations symbol represents an individual's belief as reported in the data. The expectation is over the wage offer in period t, which in our data collection is the next 4 months. The earnings beliefs ($E_i(w_{i,t})$) is elicited as follows:

"*Think about the job offers that you may receive within the coming four months. Roughly speaking, what do you think the average annual salary for these offers will be?*"

Because this question is conditional on expecting (with non-zero probability) to receive at least one offer, the sample sizes for these questions are necessarily smaller. We construct hourly wage offers from the salary expectations by dividing the reported belief by 40 hours for full-time workers, and 20 hours for part-time workers.

The final row in Panel A of Table 2 shows that the average expected salary offer is $32.3 per hour. Employed workers expect higher salary offers ($34.0) compared to unemployed workers ($20.1). The large standard deviations indicate that there is substantial heterogeneity in respondents' wage expectations.

Finally, we examine how these labor market expectations are related to key demographics and labor market activities. The top panel of Table 3 shows that many of the relationships are sensible. In each of these regressions, we include survey fixed effects (for month and year of survey) which absorbs any aggregate business cycle effects or idiosyncrasies associated with particular surveys. We see that individuals who report searching for jobs, on average, expect to receive 1.1 more job offers than their counterparts who are not searching (the mean of this variable is 0.84). Those searching also expect to receive higher wage offers. The expected offer, unsurprisingly, is positively correlated with one's current salary, education, and age. Conditional on these other characteristics, being unemployed is not systematically related with expectations regarding number of offers. However, non-employed individuals expect to receive substantially lower offer wages.

4.2 Realized Job Offers

Panel B of Table 2 presents our survey data on labor market realizations. An important feature of our data is that we collect data on job offers received, regardless of whether the job offer was actually accepted. In standard data sets, accepted job offer data is usually the only data available to analyze labor market dynamics.

In our survey, we asked the following question:

"*How many job offers did you receive in the last 4 months? Remember a job offer is not necessarily a job that you accepted.*"

Panel B of Table 2 reports that about 19 percent of all individuals received at least one job offer in the past 4 months, and the average number of job offers received was 0.34. The asterisks in panel B indicate that the mean realizations are statistically significantly different from the mean expectations. Interestingly, employed and unemployed workers in our sample receive job offers at roughly the same rate, despite their different expectations (see above). OLF individuals received the fewest job offers, with only 9 percent receiving any job offer at all.

Survey respondents' received salary offers are quite similar, on average, to their 4 month prior expectation. The average received salary offer was \$30.4 compared to an average expectation of \$32.3 (they, however, differ statistically from each other at the 10% level). The median expected and realized salary offers were quite similar, around \$25 for both. Following the pattern in expectations reported in Panel A, employed individuals reported receiving higher salary offers than unemployed workers (\$30.8 versus \$24.2).[9]

The lower panel of Table 3 shows the correlates of labor market realizations. Although searching for a job leads to a slightly higher number of offers, the returns to search effort are substantially lower than what respondents expected (Panel A). In addition, searching for a job is not systematically related to the received wage offer.

Appendix Table A2 replicates Table 2, but restricts it to the subset of respondents who take at least two labor market surveys, so that we have their data on expectations in one survey and realizations in the subsequent survey. Generally, we see similar patterns in this table. One point worth noting is that our final dataset contains 545 instances where the respondent receives at least one offer and has data on her prior expectations. This subsample forms the basis of our analysis that investigates the accuracy of expectations, and the role of learning.

[9]This is consistent with Faberman et al. (2017), who using a separate module added to the SCE, find that employed individuals fare better than the non-employed in job search.

4.3 Accuracy of Expectations and Shocks

We next compare the accuracy of expectations at the individual respondent level, exploiting the panel structure of our data. Figure 1 plots the expected number of offers on the horizontal axis (0 to 5) and the corresponding average number of actual offers received 4 months later. The increasing height of the bars indicates that a higher expected number of offers is positively correlated with more actual offers. However, individuals tend to expect a much higher number of offers than what they receive.[10] This gap between expectations and realizations appears to be mostly among the individuals who expected a high number of offers: those expecting 5 offers actually received less than 2 offers on average. At the lower end, those who expected 0 offers had fairly accurate expectations, as they actually received less than 0.1 offers on average (i.e., about 90 percent were exactly correct in predicting no offer would be received).

Figure 2 performs a similar exercise as in Figure 1, but compares actual and expected salary offers. For each decile of the expected salary offer distribution, we compute the mean actual offer received. This figure reveals a high positive correlation in expectations and actual salary offers 4 months later. Regressing the actual wage offer onto the expectation (reported in the prior survey) yields an estimate that is not statistically different from 1.

Although expectations and realizations have a high aggregate correlation, there is still considerable individual-level heterogeneity in the accuracy of expectations. We construct the salary offer "shock" for each individual, defined as the realized salary minus expected salary:

$$w_{i,t} - E_i(w_{i,t}).$$

Here $E_i(w_{i,t})$ is i's expectation about period t earnings reported in the prior survey (4 months earlier). Figure 3 shows the distribution of the shocks for our sample. The figure indicates that about half the sample experienced a positive shock (realized salary offer was better than the 4-month prior expectation) and about half the sample experienced a negative shock (realized salary worse than expected). Although the median and average shock was close to 0 (mean being -$1.8 and median being -$0.6), the size of the shocks for some sample respondents was sizable, with the 10 percentile at -$14.6/hour (an over-estimation of the realized hourly wage by $14.6), and the 90th percentile at $8.2/hour (an under-estimation of the realized offer wage by $8.2.).[11]

[10] The fact that expectations about number of offers are systematically higher than actual number of offers does not necessarily mean that expectations are biased. For example, such a pattern may arise if individuals stop search as soon as they receive an offer that is much better than they accepted. If that were the case, the gap between expectations and realizations should be larger for those individuals who accept an offer. We do not find evidence of such a pattern, suggesting that expectations of number of offers are at least partly biased. In addition, as we show below, the fact the respondents systematically update their beliefs about number of offers in response to the gap between the realized and expected number of offers indicates that the deviation has informational content.

[11] It is worth noting that we find little correlation between shocks in number of offers and wage offers. The

We next investigate correlates of the size of the shock in Table 4. The first column shows estimates of univariate regressions, where the absolute log shock is regressed onto each demographic variable, one at a time. We see that older and OLF individuals, on average, have substantially larger absolute shocks. The finding that these individuals are less accurate on average is perhaps not surprising, since earnings dispersion increases with age (Heathcote, Storesletten, and Violante, 2014; Guvenen et al., 2016), and OLF individuals, by definition, are less attached to the labor force and hence less likely to be well-informed. The absolute shock is, on average, smaller for male respondents and those with higher current/past salaries. More educated individuals do not have smaller average absolute shocks.

These relationships are qualitatively the same when we estimate a multivariate regression in the second column of Table 4. The R squared reported in the last row indicates that this rich set of demographic controls explains only a fifth of the variation in the absolute shocks. This heterogeneity in accuracy in expectations could potentially play an important role in analyzing the distribution of labor market dynamics, a topic we return to later.[12]

4.4 Learning

With our panel data, we can study how expectations are updated in response to realized salary offers, and use the data to describe the labor market learning process directly. For each respondent, we construct the *change* in salary expectations between two consecutive 4 month periods:

$$E_i(w_{i,t+1}) - E_i(w_{i,t}),$$

where $E_i(w_{i,t})$ is respondent i's expectation about period t earnings formed in the prior 4 month period, and $E_i(w_{i,t+1})$ is respondent i's expectation about future $t+1$ earnings formed in the current period t.

Figure 4 plots the change in expected salary $E_i(w_{i,t+1}) - E_i(w_{i,t})$ relative to the shock reported by that individual, $w_{i,t} - E_i(w_{i,t})$. The figure shows the binscatter, where the shock is binned into deciles, as well as a line based on a linear regression of the change in expectations onto the shock. The slope of the linear regression line in Figure 4 indicates the direction and degree of expectations updating in response to the new information (i.e., the shock). The upward slope indicates that on average individuals are responding in a logical way to the shock: a positive shock, indicating higher than expected salary offers, causes individuals to update their beliefs

Spearman rank correlation is in fact negative (-0.10; p-value = 0.04). So it is not the case that individuals who are optimistic about the number of offers are also optimistic about the wage offer.

[12]We can likewise investigate the correlates of the shock in the number of offers (results available from the authors). We find that it is those individuals who report actively searching for jobs that make substantially larger errors. We also find that males, on average, have larger errors. However, a rich set of covariates can explain less than 5% of the variation in the shocks.

about future wages upward. The slope of the updating curve, which is flatter than the 45 degree line, indicates that individuals do not fully (one-to-one) update in response to the current shock, and that prior expectations (perhaps gained from years of past labor market experience) still inform expectations.

Table 5 analyzes expectations updating using a simple regression analysis. The dependent variable here is the change in expected earnings, $E_i(w_{i,t+1}) - E_i(w_{i,t})$, and the main right-hand side variable is the shock, $w_{i,t} - E_i(w_{i,t})$:

$$E_i(w_{i,t+1}) - E_i(w_{i,t}) = \zeta_0 + \zeta_1(w_{i,t} - E_i(w_{i,t})) + \varepsilon_{it}. \quad (1)$$

ζ_1 reflects the weight the respondent puts on the shock. Another way to view this updating equation (1) is to rearrange it as:

$$E_i(w_{i,t+1}) = \zeta_0 + (1 - \zeta_1)E_i(w_{i,t}) + \zeta_1 w_{i,t} + \varepsilon_{it}. \quad (2)$$

Equation (2) shows that the "posterior" expectation $E_i(w_{i,t+1})$ is the weighted average of the prior belief $E_i(w_{i,t})$ and the signal (the offer wage one receives), $w_{i,t}$. In a standard Bayesian learning model, the ζ_1 parameter has a particular value that depends on the individual's uncertainty relative to the strength of the signal. Because our data includes beliefs (before and after the signal is received) and the signal value itself (realized wage offers), we can simply estimate the updating equation and weighting parameter ζ_1 without any restriction. This allows us to freely analyze the learning process, "letting the data speak for itself," without assuming that individuals are standard Bayesian learners.

The first column in Table 5 shows that the estimated ζ_1 is around 0.47, i.e., every \$1 in unanticipated salary offer in the current period increases expected earnings in the next period by \$0.47. Mirroring Figure 4, the effect has the sensible positive sign, but is far less than 1, indicating that current shocks do not cause "full" updating to the most recent wage offer. The regression also has a high R^2, indicating that these shocks can explain a significant proportion of the updating behavior observed in the data. Column (2) adds additional demographic variables. We see that the estimate of ζ_1 remains unchanged.[13] The next two columns of the table show slightly higher responsiveness to the shocks for the sub-sample of non-college respondents, relative to college respondents (a ζ_1 estimate of 0.56 versus 0.45), though the difference is not statistically significant at conventional levels (p-value = 0.28). Finally, the last two columns of the table present estimates of the same specification as in the first two columns, except in logs.

[13]One concern is that measurement error in reported expectations could bias the estimate of ζ_1. To investigate this, we estimate equation (2) where we instrument the respondent's expectation in period t with the lagged expectation reported in period $t-1$. This exercise can only be done for the subset of respondents whom we observe in all three waves. We find that the estimate of ζ_1 in this instrumented regression is 0.44, almost identical to the estimate of 0.47 in the baseline regression. This suggests that measurement error in reported expectations is not a serious concern.

We see that the estimates are similar to those in the first two columns.[14]

A signal from the labor market (that is, a received offer) may enable the respondent to learn about her own type and/or may provide information about the labor market (though aggregate time-varying labor market effects are soaked up by the survey fixed effects in the specifications in Table 5). Our data do not allow us to disentangle these two potential channels. Although it may be useful to understand what exactly the workers are learning about, we take the learning process as we observe in the data directly to our search model, without having to take a stand about whether workers are learning about their own type or about the labor market.[15]

We next investigate the heterogeneity in how respondents revise their expectations in response to the shocks. The first two columns of Appendix Table A3 shows that respondents who are older are more responsive to the shock, relative to their younger counterparts (p-value = 0.10). In a stationary labor market with Bayesian learning, individuals with greater experience in the labor market should be less responsive to any new signals. However, if labor market volatility increases with age (i.e., older workers face more uncertainty), this finding can be rationalized.[16] The next two columns show that the response to the shock does not differ by whether the respondent reported doing anything to look for a new job. Columns (5) and (6) of the table show that the response to positive shocks (underestimation of offer wages) is similar to that of negative shocks.

The survey also collected respondents' uncertainty about wage offers over the next four months.[17] Bayesian-consistent updating would predict that more uncertain respondents should be more responsive to the shock. This is investigated in the last two columns of Table A2. We see little evidence of this: in fact, the ζ_1 estimate is larger for respondents with lower uncertainty (though the estimate is not statistically different from that of their more uncertain

[14]We similarly see that expectations about number of offers are systematically revised in response to the shock in the number of offers: an unanticipated additional offer increases expected number of offers in the next period by 0.50.

[15]Given that we have a panel survey, individuals may also be learning about how to answer and take the survey over time. Although we cannot rule this out entirely, we do not find evidence of this. For example, our estimate of ζ_1 is virtually the same whether we use only the first two observations for a given individual in the panel, or the last two.

[16]Among older workers, the median salary of those who get offers is the 48th percentile of the salary distribution, compared to the 51st percentile for those who do not get offers. Among younger workers, the median salary of those who get offers is the 49th percentile, compared to 50th percentile for those who do not get offers. So, selection into who receives offers does not appear to be driving the difference in updating between older and younger workers.

[17]Respondents were asked: *"Think again about the job offers that you may receive within the coming four months. What do you think is the percent chance that the job with the best offer will have an annual salary for the first year of Less than 0.8*X; Between 0.8*X and 0.9*X; Between 0.9*X and X; Between X and 1.1*X; Between 1.1*X and 1.2*X; More than 1.2*X"*, where X is the respondent's expectation about the salary for the best offer they expect to get in the coming four months. Responses to the bins were required to sum to 100. The number of bins receiving non-zero probability should go up as a respondent becomes more uncertain. Here, we classify respondents who assign non-zero mass to above-median number of bins as "High Uncertainty" respondents.

counterparts).

Finally, one may ask how the estimated ζ_1 compares with what it would be if individuals updated in a Bayesian manner. Appendix section B shows how we construct the individual-specific Bayesian-updating parameter. The Bayesian ζ_1, averaged across all respondents, is 0.14, substantially smaller than the estimate of 0.43 that we obtain directly from the data (column 5 of Table 5). In fact, we find that the observed (individual-specific) updating parameter is higher than the Bayesian-predicted updating parameter for 65.3% of our respondents.[18] Thus, respondents seem to put excessive weight on the signals, relative to the Bayesian benchmark that we construct. Previous psychological and experimental economics literature on belief updating has documented several systematic biases in beliefs, including the tendency of individuals to rely too heavily on recent information- the literature refers to this as the "representative" heuristic (Grether, 1980; Kahneman and Tversky, 1982). Our results suggest, as does this earlier literature, that Bayesian updating may not be a good assumption in this setting.

4.5 Reservation Wages

To measure how respondents' expectations affect their willingness to accept future job offers, we also elicited their subjective reservation wages. Our survey included the following question:

"*Suppose someone offered you a job today in a line of work that you would consider. What is the lowest wage or salary you would accept (BEFORE taxes and other deductions) for this job?*"[19]

We then convert respondents' answers into hourly wages.[20] Among employed respondents, the average reservation wage was \$38.8/hr, compared to \$27.9/hr for the non-employed. Intuitively, employed individuals should generally have reservation wages somewhat larger than their current wage to justify paying the cost of switching jobs, and this prediction is borne out in the data. The average difference between employed respondents' reservation wage and their current salary is \$5.3, and this difference is positive for 70% of employed respondents.

Table 6 regresses respondents' reported reservation wage on their expected average salary offer, and (for the employed) their current salary. We see that, even accounting for respondents'

[18]Construction of the Bayesian-updating parameter requires assumptions on the variance of the offer distribution, that are described in the appendix. Conversely, one could rationalize the updating parameter of 0.43 through Bayesian updating by backing out the variance of the offer distribution that would be consistent with the .43 slope. This variance would be about 30-40% of the variance that we arrive at.

[19]This wording is similar to the one first used in the May 1976 Current Population Survey, and since used in many studies including Krueger and Mueller (2016).

[20]Prior to March 2017, respondents could report their reservation wage in annual, monthly, biweekly, weekly, or hourly terms. Starting in March 2017, they could only report their reservation wage in annual terms. As above, we assume part-time workers consider job offers with 20-hour work weeks, while full-time workers and the non-employed consider 40-hour work weeks. We set reservation wages to missing if revision in hourly reservation wage is greater than \$48.1.

current salary, reported reservations wages are significantly related to their expectations about future offers, indicating that individuals who expect higher future wages, have higher reported reservation wages. For the employed, the coefficient on expectations is actually larger than that on current salary, though this difference is only statistically significant for college-educated respondents.

These results point to the potential importance of incorporating subjective expectations when modeling labor market transitions. We return to these data below to test whether our model of job search with learning better predicts respondents' reservation wages than a more traditional model with perfect information.

5 A Model of Labor Search with Heterogeneous Beliefs and Learning

In this section we develop a model of job search augmented to allow for heterogeneous worker beliefs and learning. In the following section, we use our data on labor market expectations to estimate this model and quantify the importance of heterogeneous beliefs and learning for welfare.

5.1 Environment

Our modeling framework is based on a standard model of labor market search, allowing for on-the-job search (Burdett 1978). Time t is discrete. Workers are either employed or unemployed.[21] All individuals discount the future at rate $\beta \in (0, 1)$.

Workers search for jobs in unemployment and on-the-job, and in each period job offers arrive in those states with probability λ_u^* and λ_e^*, respectively. Jobs are destroyed with an exogenous probability δ^*. Each job offer is associated with a wage offer w, where wage offers are distributed according to F^*. We assume that log wage offers are Normally distributed with mean μ^* and variance σ^{*2}. As detailed below, we allow for heterogeneity in the wage offer distributions and arrival rates of jobs by the worker's observable and unobservable characteristics.

In each period, if an unemployed worker receives a job offer, she decides whether to accept the offer and leave unemployment or remain unemployed and enjoy the value of leisure b. If an employed worker receives an outside offer, she decides whether or not to accept the offer and switch firms after paying some moving cost m.[22] Current firms do not compete with the

[21] Although we refer to non-employed as unemployed, our empirical work does not distinguish between unemployed and out of the labor force (OLF) states, but rather combines them into one category. This is motivated by the large number of OLF individuals reporting that they receive job offers that transition directly to employment.

[22] We introduce a moving cost because our data clearly indicate reservation wages for employed workers that differ from their current wage. Without a moving cost, the model would predict a reservation wage that is equal

outside wage offers their workers receive. Throughout, we use $*$ to indicate "true" or "actual" probabilities of receiving wage offers or jobs being destroyed, to distinguish these from the workers' beliefs.

5.2 Beliefs and Learning

We augment this otherwise standard model of on-the-job search with a general model of beliefs and learning. We start with the notion that workers make decisions with possibly limited knowledge about future wage offers.[23] We assume that workers know the dispersion in (log) wage offers given by σ^*, but do not necessarily know the mean of the (log) wage offer distribution μ^*. Define μ_t as the worker's current (date t) *belief* about μ^*; for simplicity, we have dropped the i subscript on the worker's belief. We say that a worker's beliefs are *incorrect* if $\mu_t \neq \mu^*$.[24] Note that the true distribution of wage offers F^* is not indexed by t because we assume that it is not time varying. Worker beliefs at time t, characterized by μ_t and distribution F_t, potentially evolve over time due to learning.

Workers update their beliefs about μ^* by learning from their labor market experiences. Figure 5 summarizes the timing of the model. Workers enter the period with some beliefs summarized by μ_t. Employed workers earn their wage w, while unemployed workers earn their value of unemployment b. Employed workers might separate from their current employer due to job destruction, in which case we assume they cannot search for new employment within the period.[25] Employed workers who have not separated and unemployed workers may receive job offers. Conditional on these outcomes, workers then form beliefs about future wage offers in the next period, where the new belief is μ_{t+1}.

We denote the worker's updating or learning process as the mapping L from a worker's current beliefs μ_t and realized wage offers w_t to next period's updated beliefs μ_{t+1}, $L : \mu_t, w_t \mapsto \mu_{t+1}$. We parametrize this mapping as:

$$\mu_{t+1} = (1 - \zeta)\mu_t + \zeta \ln w_t. \tag{3}$$

Note that (3) is the analog of (2), but in log wages not levels. For the case where no offer is

to the current wage.

[23] We focus here on beliefs and learning about wage offers as in (Burdett and Vishwanath, 1988), and not on learning about arrival rates or job destruction rates, though our data contain information on the latter. This is because individuals arguably have more ability to control the number of offers they receive (through their search effort) than the wage offer, and we do not have rich data on search effort to endogenize it. Moreover, because we model the information frictions pertaining to wage offers which are on average unbiased, our exercise can be seen as offering a lower bound on the importance of information frictions. For a model which features Bayesian learning about the arrival rate of offers, see Potter (2017).

[24] To connect this model to the notation in the reduced form analysis above, we can write that in period t each worker i has belief $\mu_{i,t} = E_i(\ln w_{i,t})$.

[25] This is a simplification. The qualitative results would not change significantly if we instead allowed separated workers to search as well.

received, we assume there is no learning about wage offers.[26] ζ is the "learning" or "updating" parameter, which we identify directly from observed updating in our data, as described below. $\zeta = 0$ implies no learning: $\mu_{t+1} = \mu_t$. $\zeta = 1$ implies immediate updating to the latest wage offer: $\mu_{t+1} = \ln w_t$. In Bayesian learning, ζ is a function of the time-varying precision of the worker's belief entering this period and the underlying (known) dispersion of the wage offer distribution. We do not assume that workers are Bayesian type learners, and allow for non-Bayesian learning that we infer from the data directly.

5.3 Perceived Values of Unemployment and Employment

We next characterize the perceived flow values of unemployment and employment.[27] We characterize each of these values conditional on the initial beliefs of the worker entering period t, given by μ_t.

For a worker with beliefs μ_t, the perceived value of unemployment is

$$U(\mu_t) = b + \beta\lambda_u^* \int_y \max\{W(y, \mu_{t+1}), U(\mu_{t+1})\} dF_t + \beta(1 - \lambda_u^*)U(\mu_t), \tag{4}$$

where b is the flow value of unemployment, and F_t is the worker's belief regarding the wage offer distribution, given her current belief μ_t.[28] There are two possible future events: i) if the worker receives a job offer, she updates her beliefs to μ_{t+1}, and ii) if she does not receive a job offer, her belief remains at μ_t.[29] For a particular wage offer realization y, μ_{t+1} is the updated mean of the (log) wage offer distribution given by the learning process in equation (3).

For a worker with current beliefs μ_t, the value of employment at wage w is

$$\begin{aligned} W(w, \mu_t) = w &+ \beta\delta^* U(\mu_t) \\ &+ \beta(1 - \delta^*)\lambda_e^* \int_y \max\{W(y, \mu_{t+1}) - m, W(w, \mu_{t+1})\} dF_t \\ &+ \beta(1 - \delta^*)(1 - \lambda_e^*)W(w, \mu_t), \end{aligned} \tag{5}$$

where δ^* is the job destruction probability and λ_e^* is the probability of receiving an offer while

[26] In a possible extension to this model, in which there is learning about the arrival rates of jobs, receiving no offer in a period could cause the worker to update their beliefs about the arrival rate. In the data, we do observe that absolute revisions for expectations about wage offer are smaller for individuals who do not receive any job offer, providing some support for our assumption that workers only learn if they receive an offer.

[27] We refer to these as perceived to distinguish them from actual values of employment $W^*(\cdot, \cdot)$ and unemployment $U^*(\cdot)$ which we describe in Section 8. These are perceived because they are based on the workers' beliefs about the offer distribution rather than the actual one.

[28] The variance of the worker's prior is not a state variable because we assume they are certain about their beliefs.

[29] Importantly, the expectation over wage offers is with respect to F_t given belief μ_t, not the updated belief about the wage offer distribution F_{t+1} given updated belief μ_{t+1}, because the worker is forming expectations without having received any offers yet, and hence using her current beliefs when entering the period.

employed. There are three possible future events: i) if the worker's job is destroyed, then she becomes unemployed and there is no learning (her belief in unemployment remains at μ_t), ii) if the job is not destroyed, but the worker does not receive an offer, then she remains at her current job and her beliefs remain at μ_t, iii) if the job is not destroyed and the worker receives an offer, then her beliefs are updated to μ_{t+1} given the learning process (3). If the worker accepts the offer y, she receives $W(y, \mu_{t+1}) - m$, where m is the "moving cost" of changing jobs. μ_{t+1} is the updated belief given the received wage offer y following equation (3).

5.4 Reservation Wages

We next move to providing the reservation wage policy in our model with learning. One of the key features of a model with learning is that the perceived values of unemployment U and employment W (equations (4)-(5)) depend on the wage offers received, as realized offers provide information about the unknown wage offer distribution. In particular, the value of search (either in unemployment or on-the-job) *increases* as the individual receives a wage offer that is higher than the worker's current belief. As Burdett and Vishwanath (1988) point out, models with learning are therefore not necessarily characterized by standard reservation wage policies, as individuals could accept relatively low offers, but reject high wage offers which deliver updated beliefs that make continuing to search more attractive. Burdett and Vishwanath (1988) also provide general conditions for a reservation wage policy to exist, and then prove that in their specific Bayesian model with Normally distributed wage offers that a reservation wage policy does indeed exist.

Our setting is slightly more complicated because we allow for on-the-job search, job switching costs, and job destruction, in addition to heterogeneous beliefs and non-Bayesian learning. We take two approaches to analyzing this issue in our model. First, in Appendix C, we extend the Burdett and Vishwanath (1988) results to provide sufficient conditions for a reservation wage policy to exist in our model. We then numerically test these assumptions and confirm that a reservation wage offer policy exists at our estimated primitive parameters.

We now turn to characterizing this reservation wage policy. For an unemployed worker who enters the period with beliefs μ_t, the reservation wage is defined as the wage offer which makes them indifferent between accepting and rejecting the offer. Crucial to our model is that, in general, this reservation wage takes into account the effect that this offer has on their future beliefs. The reservation wage in unemployment $y_r^u(\mu_t)$ is defined by the following equation:

$$W\Big(y_r^u\big(\mu_t\big), \mu_{t+1}\Big) - U\Big(\mu_{t+1}\Big) = 0, \tag{6}$$

where the updated belief μ_{t+1} depends on the offer y and current beliefs μ_t according to the learning process (3). At offers $y > y_r^u(\mu_t)$, the unemployed worker with current beliefs μ_t will

accept the job and leave unemployment.

The reservation wage for an employed worker depends both on the worker's current belief about the wage offer process μ_t and her current wage w. The reservation wage $y_r^e(\mu_t, w)$ for an employed individual earning wage w is defined by the following equation:

$$W\Big(y_r^e\big(w,\mu_t\big),\mu_{t+1}\Big) - W\Big(w,\mu_{t+1}\Big) = m, \tag{7}$$

where the updated belief μ_{t+1} depends on the offer y and current beliefs μ_t according to the learning process (3). At offers $y > y_r^e(w,\mu_t)$ the worker with initial beliefs μ_t and wage w will switch to a new employer.

5.5 Employment State Transitions

Finally, we can use the model structure to characterize the employment state transitions. Conditional on current beliefs μ_t, the probability an unemployed worker moves to employment is given by:

$$\mathbf{pr}(UE|\mu_t) \equiv \mathbf{Pr}(\text{E in } t+1 \mid \mu_t, \text{ U in } t) = \lambda_u^*(1 - F^*(y_r^u(\mu_t))), \tag{8}$$

where λ_u^* and F^* are the true probabilities of job arrivals and job offers, and $y_r^u(\mu_t)$ is the reservation wage for a worker who begins the period with beliefs μ_t.

Similarly, the probability an employed individual at wage w enters the period with beliefs μ_t transitions to a new job is given by:

$$\mathbf{pr}(EE|\mu_t, w) \equiv \mathbf{Pr}(\text{E in } t+1 \mid w,\mu_t, \text{ E in } t) = (1-\delta^*)\lambda_e^*(1 - F^*(y_r^e(w,\mu_t))), \tag{9}$$

where λ_e^* is the true probability that a job offer arrives in employment, δ^* is the true separation probability, and $y_r^e(w,\mu_t)$ is the reservation wage for an employed worker at wage w and beliefs μ_t.

Finally, the probability an employed individual separates to unemployment is simply

$$\mathbf{pr}(EU|\mu_t, w) \equiv \mathbf{Pr}(\text{U in } t+1 \mid w,\mu_t, \text{ E in } t) = \delta^*. \tag{10}$$

Note that the existence of a reservation wage policy for employed workers, as discussed above, implies that there are no offer realizations that cause a currently employed worker to voluntarily decide to become unemployed. This could have occurred if, for example, a current wage offer caused the worker to update their beliefs positively and the probability that an offer arrives in unemployment was sufficiently larger than that in employment. In this case, once a worker believes that they face a very favorable offer distribution, they might want to move to unemployment to take advantage of a higher arrival rate of offers. Our primitive assumptions on

the reservation wage policy, discussed in Appendix C, rule this out. And, we have confirmed numerically, that at our estimated parameters, the reservation wage policy exists.

5.6 Model Solution

For a given set of parameters, we solve the model using value function iteration. The model solution provides reservation wage functions for unemployed and employed workers following equations (6) and (7), respectively, as well as unemployment-to-employment and employer-to-employer transition probabilities following equations (8) and (9), all of which depend on current beliefs μ_t. For details on the numerical solution to the model, see the Appendix, Section D.

5.7 Learning and Labor Market Search

Before moving to the empirical estimation, we pause to examine how heterogeneous beliefs and learning affect labor market search in this framework. Figure 6 plots the policy functions (reservation wages) defined in equations 6 and 7. Figure 6 also gives us a sense of the importance of optimism or pessimism, or more generally how heterogeneity in beliefs affects behavior. We define "optimists" as individuals whose current belief exceeds the true wage offer parameter ($\mu_t > \mu^*$), and "pessimists" as the converse. If all individuals were pessimists, reservation wages in unemployment would be relatively low and transition rates would be high. Similarly, if all individuals were optimistic, reservation wages would be relatively high and transition rates would be low.

The rate of learning about wage offers also affect reservation wages and thus affect transition probabilities from non-employment to employment, as well as from employer-to-employer. First, with no learning ($\zeta = 0$), individuals who are more optimistic believe they have higher option values from search and thus have higher reservation wages. Therefore, all else equal, an optimistic individual will be less likely to exit non-employment as compared to a pessimistic individual when faced with the same wage offers. This can be seen in the increasing reservation wage function in black for unemployment in the left panel and employment in the right panel.[30]

Introducing learning ($\zeta > 0$) flattens the reservation wage rule everywhere, as current beliefs begin to matter less for outcomes. In the extreme case where beliefs are completely transitory ($\zeta = 1$), reservation wages are independent of current beliefs, both in unemployment and employment. In our model timing, current values depend on current beliefs, but if individuals receive an offer and $\zeta = 1$, they update their beliefs immediately to the new offer- whether they accept the offer or not - and then choose their reservation wage accordingly. Therefore, in the

[30]For employed workers, an important determinant of reservation wages is also the current wage. If an individual is in a job with a current wage far above the mean of the perceived offer distribution, this worker's beliefs will not have a large affect on employment transitions because this individual's reservation wage will be high anyway due to a high current wage.

model with $\zeta = 1$, no matter the initial beliefs, if two individuals receive the same offer they update their beliefs to the same point and will thus have the same reservation wage.

6 Identification and Estimation

In general, separately identifying worker beliefs and learning from other model primitives (such as the value of unemployment) would be difficult without strong parametric assumptions. One such commonly used set of assumptions is to assume all workers have the same beliefs, those beliefs are correct, and the worker therefore never learns (or needs to learn) anything in the labor market. Another set of previous assumptions imposes Bayesian learning on all workers (Burdett and Vishwanath, 1988). As we detail in this Section, we use our unique data on labor market expectations to identify a richer model, one that allows for heterogeneous, possibly incorrect beliefs, and general, non-Bayesian, learning. In addition, our data on job offers received, regardless of whether accepted or not, allows us to directly identify job arrival probabilities and the actual wage offer distribution. As we detail below, we use our rich dataset to form a simple and robust multi-step estimator of the model parameters.

6.1 Identification

We first discuss the identification of the model parameters. One model period is assumed to be four months, as our panel is interviewed once every four months. We set the discount rate exogenously at $\beta = 0.984$ to match an annual discount rate of 5%. The remaining parameters to be identified are the job destruction and arrival probabilities $\delta^\star, \lambda_u^\star, \lambda_e^\star$, the actual wage offer distribution F^* for each individual in our sample, the distribution of worker beliefs about F^*, G, learning behavior L, the flow value of unemployment b, and the employer-to-employer moving cost m. Throughout, we introduce various sources of heterogeneity in these primitives, in particular allowing for different wage offers, transition probabilities, and job destruction probabilities by the worker's education level and other characteristics.

Job Destruction and Arrival Probabilities Using standard data that includes only employment transitions and unemployment durations, job offers are not observed and the offer arrival probability is identified given the particular model structure. Our richer survey data, in which respondents directly report any job offers, regardless of whether they accept them or not, allows us to identify the job offer arrival probabilities robustly, and without any bias due to model mis-specification. We identify the offer arrival probabilities $\lambda_u^\star$ and $\lambda_e^\star$ using the proportion of unemployed and employed respondents receiving a job offer. We identify the job destruction/separation probabilities $\delta^\star$ using the proportion of workers moving from employ-

ment to non-employment. We identify the arrival and separation probabilities separately for groups of workers defined by education (college or non-college), as discussed below.

Actual Wage Offer Distribution As with arrival probabilities, standard labor market data cannot non-parametrically identify the wage offer distribution from accepted wages. The standard approach is to use the model structure, in particular the implied reservation wages, to identify the offer distribution from the endogenously determined observed accepted offers. Instead, we use our data on actual wage *offers*, whether they were accepted or not, to identify the offer distribution.[31]

Importantly, we do not want to assume that all workers in our sample face the same wage offer distribution, as this would imply that any deviation of beliefs from this incorrectly assumed homogeneous wage offer distribution would reflect some degree of "pessimism" or "optimism," rather than actual wage offer heterogeneity. Our approach is to allow for a rich heterogeneous wage offer distribution, as much as our data allows.

The starting point for our empirical model for log wage offers $log(y)$ is to assume that each individual i faces an actual log wage offer distribution given by $\ln y_i \sim N(\mu_i^*, \sigma_i^{*2})$, where μ_i^* and σ_i^* are the mean and standard deviation of individual i's specific log wage offer distribution. We assume that μ_i^* is a function of both observable characteristics X_i for each individual (age, education, gender, and industry), and an unobservable η_i. We assume that η_i has a discrete distribution with two types, $\eta_i \in \{\eta_1, \eta_2\}$. μ_i^* is then:

$$\mu_i^* = \alpha' X_i + \eta_i.$$

We estimate the support of η_i, along with the probability that an individual is a type η_1, π_1, using observed log wage offers in our sample. We allow for the unobserved type support $\{\eta_1, \eta_2\}$ and probability mass parameter π_1 to vary by the worker's education. For details on how we implement this in practice, see Appendix F. The distribution of unobserved types is identified from "residual" heterogeneity in wage offers, conditional on observable covariates X.

Because individuals can now be one of two unobserved types, we amend the transition probabilities in equations (8)-(9) to account for this unobserved heterogeneity and "integrate out" the unobserved type heterogeneity in the standard way. The augmented individual transition probabilities (conditional on covariates X) are thus given by:

[31] One concern could be that the observed empirical offer distribution is truncated from below. Workers may not actively pursue offers from employers with very low wages and hence not report them in the survey. This is certainly plausible and might lead us to under predict the degree of wage dispersion in the data. First, using simulated data from the estimated model, we can show that this is not the case, suggesting that the truncation of low offers is not a serious issue. In addition, in the reported data, we see that in more than half of the cases (54%), individuals receive offers that are below their current wages.

$$\mathbf{pr}(UE|X,\mu_t) = \lambda_u^*\Big[\pi_1\big(1-F^*(X,\eta_1)(y_r^u(\mu_t))\big) + (1-\pi_1)\Big(1-F^*(X,\eta_2)(y_r^u(\mu_t))\Big)\Big], \quad (11)$$

$$\mathbf{pr}(EE|X,\mu_t,w) = (1-\delta^*)\lambda_e^*\left[\pi_1(1-F^*(X,\eta_1)(y_r^e(w,\mu_t))) + (1-\pi_1)(1-F^*(X,\eta_2)(y_r^e(w,\mu_t)))\right]. \quad (12)$$

Beliefs and Learning Rather than assume all workers share a common prior belief, we identify the distribution of beliefs about the mean of the (log) wage offer distribution, $\mu_{i,t}$, by directly using our data on elicited beliefs.[32] We impose no parametric distribution on the $\mu_{i,t}$ distribution in our sample, and use the sample distribution to non-parametrically identify the beliefs distribution.

We identify the learning process using our panel data on belief updating and the wage offers that individuals received. Given prior belief in period t, $\mu_{i,t}$, a wage offer $w_{i,t}$, and an updated belief $\mu_{i,t+1}$, the learning process takes the form:

$$\mu_{i,t+1} = (1-\zeta)\mu_{i,t} + \zeta \ln w_{i,t}. \quad (13)$$

We allow for ζ to vary by the worker's education. Note that the ζ we estimate from data on belief-updating need not be Bayesian; we do not impose any relationship between the updating rule and an individual's level of uncertainty.

Unemployment Value and Moving Costs There are two remaining parameters: the value of non-employment or "leisure" b and the value of the current employer or "moving cost" m. We identify these free parameters from the observed employment transitions (unemployment-to-employment and job-to-job). In particular, we use the model-generated policy rules (reservation wages) - which are a function of an individual's labor market status and current beliefs - to compute the individual's implied reservation wage and thus probability of making a transition to employment or a transition to another employer for the unemployed and employed, respectively. Given the structure of the model, there is a unique and one-to-one mapping between the transition rates and the b and m values for any labor market. We allow for the b and m values to vary by the worker's education.

[32] Because our elicited beliefs data is in levels, we estimate the individual $\mu_{i,t}$ using the reported expectation in levels given the estimated σ^*.

6.2 Estimator

We use the identification arguments constructively to form a two step estimator. All model parameters are estimated by at least education level: non-college (less than a 4-year college degree) and college (college degree or more). As previously noted, estimation of the actual wage offer distribution allows even richer heterogeneity based on age, gender, and industry, in addition to education, and allows for unobserved heterogeneity as described above.

Step 1: Estimate arrival and separation probabilities, actual wage offer distributions, beliefs and the learning rule Arrival and separation probabilities are estimated directly from our sample. The procedure for estimating the individual-specific wage offer parameters μ_i^* and σ_i^* has been outlined earlier. To estimate the learning process ζ, we use the panel data on updated beliefs.

Step 2: Estimate flow value of unemployment (b) and moving costs (m) In the second step of the estimator, we take the Step 1 values and estimate the (b, m) combination which minimizes the distance between the empirical probabilities $\hat{pr}(UE)$ and $\hat{pr}(EE)$ and simulated ones $pr(UE|b, m)$ and $pr(EE|b, m)$. The simulated transition probabilities depend on the Step 1 parameter estimates, including the data on beliefs. As with all of the other model parameters, we estimate the model separately for the college and non-college groups. Our estimator for (b, m) is thus defined as:

$$[\hat{b}, \hat{m}] = \arg\min_{b,m} \Big(\hat{M} - M(b, m)\Big)' \Big(\hat{M} - M(b, m)\Big), \tag{14}$$

where the vector of sample moments (sample transition probabilities) are given by

$$\hat{M} = [\hat{pr}(UE) \;\; \hat{pr}(EE))]',$$

and the vector of model-implied transition probabilities is given by

$$M(b, m) = [pr(UE|b, m) \;\; pr(EE|b, m)]'.$$

The model-implied transition probabilities are given in (11) and (12), which we average over all individuals in the sample to obtain a group-level transition probability. We use simulation to compute the transition probabilities, the details of which are in the Appendix section E.

We compute the estimator using an approximation to the objective function (as a function of the unknown parameters b, m). First, we compute each model for 11 b values and 11 m values (11 x 11 unique values), each of which gives rise to some non-employment to employment and employer-to-employer transition probabilities. We then approximate the transition probabilities

as functions of (b, m) using Chebyshev polynomials, and find the (b, m) combinations that minimize the distance between the observed transition probabilities and the model-generated ones. Our estimator performs very well–the fit is near perfect–as discussed below.

Inference To account for the multiple steps of our estimator and the clustered nature of our panel data, we conduct inference for our model estimates using a panel bootstrap over all steps of the model estimation. For a sample size of N respondents, we re-sample with replacement each respondent $i = 1, \ldots, N$ and take all of their sample data for every wave of the survey they participated in. For each bootstrap sample of size N, we then re-estimate all parameters in sequence, including arrival rates of jobs and job destruction probabilities, actual wage offer distribution, belief distribution and learning process, and finally the b, m values. For $B = 100$ bootstrap re-samples, we compute standard errors and confidence intervals. Our block/panel bootstrap over all model estimation steps accurately reflects the clustered dependence of the data (by individual) and any sample variability introduced in each estimation step.

7 Results

7.1 Parameter Estimates and Sample Fit

Table 7 summarizes the parameter estimates for $\sigma^*, \lambda_e^*, \lambda_u^*, \delta^*, \zeta, \eta_1, \eta_2, \pi_1$ for each education group. As might be expected, the offer probabilities are higher for the college educated relative to the non-college educated, while the separation probability is lower. Additionally, the variance of wage offers is larger for the college educated. The updating parameter, ζ, is lower for the college educated, indicating that they tend to place less weight on new offers received relative to current beliefs as compared to the non-college educated.

The middle panel of Table 7 reports the parameter estimates for the value of leisure and moving costs $(\hat{b}, \hat{m})$ as well as the fit of each moment for each model described above. Given the model is just-identified, we are able to fit the employment transitions for each model almost exactly. While the estimated b levels are lower for the non-college sample, the estimated value of leisure relative to the mean offered wage in our non-college and college samples is .24 and .08, respectively, suggesting that the college educated have a higher opportunity cost of work. Similarly, moving costs as a fraction of the average hourly wage offer for employed workers in our sample is 3.2 and 6.6 for the non-college and college workers, respectively. This is consistent with the notion that more educated workers are in specialized jobs that have higher costs of moving.

Figure 7 plots the distribution of implied errors (truth - belief) in our sample, both for the non-college (left panel) and college (right panel) samples. For both samples, the median error is

slightly positive ($0.33 and $1.27 for the non-college and college samples, respectively), implying that the median worker is pessimistic. However, both samples have a wide range of errors, with a large number of both optimistic and pessimistic individuals in each.

7.2 Model Validation: Reservation Wages

Our estimation strategy employs information on labor market flows and how beliefs are updated given offers that individuals receive. However, we do not make use of any information on reported reservation wages and their relationship with beliefs in our estimation. As discussed above, we have found that in our data, reported beliefs and reservations wages are correlated, consistent with our model predictions. In this sub-section, we use the reservation wage data to validate our model predictions directly. In particular, we ask how well our model does relative to a more standard model of complete information ("CI") in predicting reported reservation wages.

To do so we solve and estimate a complete information model in which individuals know the true mean of the (log) wage offer distribution and there is no learning necessary, following the same procedure as outlined in Section 6.[33] This generates estimates for the value of leisure and moving costs under the complete information model, $(\hat{b}_{CI}, \hat{m}_{CI})$, as well as policy functions (reservation wages) for the unemployed and employed under complete information. In this case, all unemployed workers of the same type have a common reservation wage, and reservation wages in employment depend only on current wages. We then also use our estimated model with subjective beliefs and learning to generate reservation wages implied by the learning model, which instead are a function of each individual's current belief $\mu_{i,t}$ and current wage $w_{i,t}$ if employed. Finally, we compare how these model-predicted reservation wages fare in explaining heterogeneity in self-reported reservation wages in our sample.

For this comparison, we do two things. First, for each of the models, we compute the correlation between model-predicted reservation wages and reported reservation wages, separately for the employed and unemployed. Second, for each of the models, we regress reported reservation wages on reservation wages implied by the model, demographics, and a constant, separately for the employed and unemployed. The R^2 from this regression is what we are interested in. Table 8 presents these statistics for our learning model, which we refer to as the "baseline" model, and the "complete information" model. The first two columns show statistics for the employed. Both the correlation coefficient and the R^2 are higher for the baseline model. For example, the baseline model yields a correlation coefficient of 0.58 for the non-college sample (and an R^2 of 0.38) versus 0.52 (and 0.33) for the complete information model. Note that reservation wages for employed workers depend on current wages in both models, and so the complete information model is already able to pick up a significant amount of reported reservation wage heterogeneity

[33] That is, beliefs about the mean of the offer distribution, $\mu_{i,t}$, are now set equal to their true value μ_i^* for all i and t. See the Appendix G for details regarding the estimation and fit of the perfect information model.

without needing heterogeneity in beliefs. Nonetheless, belief heterogeneity increases the fit of model-generated reservation wages with reported ones for employed workers.

More importantly, when we shift to the last two columns of Table 8, we see that the baseline model does significantly better in picking up the heterogeneity in reported wages for the unemployed. For the non-college and college groups, respectively, the R^2 is about 0.3 and 0.35 under the baseline model, and less than 0.10 for the complete information model for both groups. And, the correlation coefficient under the baseline model is more than five times that of the complete information model. This leads us to conclude that the baseline model – with heterogeneous beliefs and learning – does a much better job of matching the data on reservation wages.[34]

7.3 Policy Analysis: Counterfactual Changes in b

The two panels of Figure 8 plot the UE rate for our baseline model (blue lines) for counterfactual changes in the flow value of unemployment b, for the non-college and college subsamples. In both panels, we also plot the transition rate estimated from our data (the estimation "target") as a horizontal line in red. The intersection of the empirical target with the model-implied transition curves is at our estimated b value reported in Table 7, giving a visual depiction of how this data moment identifies the b value. All other values of b are "counterfactual." In interpreting these counterfactual results, it is important to remember that the flow value in our model represents any unemployment benefits and the individual's taste or distaste for unemployment relative to employment. Changing b counterfactually in these exercises can be interpreted as a change in the government-provided unemployment benefits.

In both panels, we see the expected negative slope: a higher flow value of unemployment b implies a lower rate of leaving unemployment (the blue lines).[35] Increasing b by \$7 over the ranges shown in the panels leads to a decline of 7.5 percentage points in the transition rate into employment for the non-college group; increasing b by \$5 over the ranges shown in the panels leads to a decline of about 2 percentage points in the transition rate into employment for the college group.

One contribution of our framework is to show the importance of incorporating data on prior beliefs and learning relative to the standard model assuming complete information. To quantify

[34] A direct consequence of this analysis is that our model should be able to capture a larger degree of cross-sectional wage dispersion relative to a standard model without heterogeneous beliefs (Hornstein, Krusell, and Violante (2011) so long as beliefs about μ^* are not above it for all individuals). To confirm this, we simulate labor market histories for a large number of workers, assuming that they enter the labor market with beliefs $\mu_{i,0}$ equal to their true offer distribution parameters, μ_i^*. We then calculate the standard deviation of log wages in our baseline estimated model (within different μ^* groups, averaged across groups) as well as the complete information model with the same parameters. When moving from complete information to our model with learning, the mean standard deviation of log wages rises from .25 to .27 for the non-college sample, and from .3 to .33 for the college sample in our baseline model relative to complete information.

[35] The slope need not be linear, and will in general depend on the belief composition of our sample, over which we compute the transition probabilities.

this comparison, we use our estimated complete information model, which assumes that all workers have the same beliefs about the wage offer distribution and those beliefs are correct (equal to the heterogeneous wage offer distribution we estimate): $\mu_{it} = \mu_i^*$ for all i, t. For this counterfactual complete information model, we perform the same exercise as with our baseline model, and compute the implied transition probabilities for various b values. These are plotted in the black lines in the figure. For both subsamples, we see that the complete information model predicts more responsiveness to changes in the value of unemployment.[36] The difference is particularly transparent in the case of the non-college-educated sample. For a change in b from \$1 to \$7, the baseline model predicts a drop of approximately 7.5 percentage points in the UE transition rate, while the complete information model predicts a full percentage point more than that. The different predictions of the two models indicates that assuming incorrectly that workers have complete information can bias estimates of the predictions of government interventions, such as increases in unemployment benefits.

8 The Welfare Costs of Information Frictions

In this section, we quantify the utility/welfare costs of information frictions and translate this quantity into its search friction equivalent. To do so, we conduct the following exercises. We solve our preferred baseline model - the model with heterogeneous beliefs and learning- with the estimated $(\hat{b}, \hat{m})$ from Table 7 and compute the implied employer-to-employer transition probabilities, unemployment-to-employment transition probabilities, and the expected present discounted value of utility. We then repeat this exercise assuming instead that (i) individuals have complete information, where their initial beliefs about the wage offer distribution are correct $\mu_{i,t} = \mu_i^*$ for all i and t and there is no learning, and (ii) individuals have imperfect priors, but do not learn over time ($\zeta = 0$ but $\mu_{i,t} \neq \mu_i^*$). A comparison of the baseline model with (i) informs us about the welfare consequences of imperfect information. A comparison of the baseline model with (ii) is informative about the role of learning in this context.

Welfare Calculations In our model with incorrect beliefs, welfare calculations are not as straight-forward as they would be in standard models. This is because individual behavior is a function of *beliefs* about the distribution of offers, but realized welfare is a function of that behavior and the *actual* distribution of offers. Therefore, to compute welfare, we keep preferences, the learning rule, and the associated reservation wage polices as estimated, but compute welfare using the *actual* distribution of wage offers as follows.

The *actual* present-discounted value of unemployment $U^*(\mu_t)$ and employment $W^*(w, \mu_t)$ are defined as:

[36] Note that the two models also have different moving costs operating in the background.

$$U^*(\mu_t) = b + \beta\lambda_u^* \int_{y_r^u(\mu_t)} W^*(y, \mu_{t+1}) dF^*(y) + \beta\lambda_u^* \int^{y_u^r(\mu_t)} U^*(\mu_{t+1}) dF^*(y)$$
$$+\beta(1-\lambda_u^*)U^*(\mu_t), \tag{15}$$

and

$$W^*(w, \mu_t) = w + \beta\delta^* U^*(\mu_t) + \beta(1-\delta^*)\lambda_e^* \int_{y_r^e(w,\mu_t)} \{W^*(y, \mu_{t+1}) - m\} dF^*(y)$$
$$+\beta(1-\delta^*)\lambda_e^* \int^{y_r^e(w,\mu_t)} W^*(w, \mu_{t+1}) dF^*(y) + \beta(1-\delta^*)(1-\lambda_e^*)W^*(w, \mu_t), \tag{16}$$

respectively, where the learning rule continues to follow equation (3), and $\zeta = 0$ for the model without learning and the model with complete information. These values differ from the perceived values described in equations (4) and (5) because they impose the policy functions from the model with learning we have estimated (that is, the $y_r^u(\mu_t)$ and $y_r^e(w, \mu_t)$ for the baseline case), but use the actual offer distributions that individuals face, rather than the individual's belief about the offer distribution.[37] Therefore, these values capture actual mean present discounted values of income and thus utility for our (linear utility) model with learning. To provide some interpretation of the magnitudes of welfare that follow, note that earning \$1 in hourly wages in our model forever (without the possibility of separation) has a present discounted value (PDV) of utility of $\frac{1}{1-\beta} = \frac{1}{1-0.984} = \$62.5$.

We start with the complete information model, in which perceived and actual expected welfare are identical. The results are reported in the first row of each panel in Table 9, with the top panel corresponding to the non-college sample and the bottom panel corresponding to the college sample. The first two columns report the implied mean transition probabilities in our sample. The last two columns depict the average PDV of utility in annual earnings relative to the CI model and the interquartile range of this mean gain, respectively.[38] By definition, these measures are zero for the first row, which corresponds to the complete information model.

Benchmark 1: Baseline Model To compare our baseline framework to a model with complete information, we then calculate the same objects under our baseline model, the results of which are reported in the second row of each panel in Table 9. For non-college workers, moving from the baseline model to a model with complete information reduces the one-period ahead

[37] For the empirical priors and no learning model (Benchmark 2 below), we impose the policy functions from the model with empirical priors and no learning, but again uses the actual offer distributions that individuals face rather than the individual's belief about the offer distribution, with the additional assumption that beliefs are fixed at their starting values.

[38] In particular, this is the implied values computed in Equations (15) and (16), multiplied by $(1-\beta) * 2000$ to express it in annual dollars.

unemployment-to-employment transition probability by roughly 1.65 percentage points, and the employer-to-employer transition probability by 0.71 percentage points. In general, the direction of these changes is ambiguous, and depends on the distribution of beliefs about μ_i^* in our sample relative to each individual's actual μ_i^*. For example, if most people are pessimistic ($\mu_{i,t} < \mu_i^*$), then moving to a model with complete information should, on average, lower unemployment-to-employment transition probabilities. Moving to complete information increases an individual's beliefs about their outside option of search in unemployment (to the truth), consequently raising reservation wages and lowering job finding rates. Indeed, we estimate that unemployed individuals are on average pessimistic, with the average unemployed worker under-estimating the mean of their offer distribution by \$.031.[39]

More importantly, a non-college worker would be willing to pay \$175.24 more *per year* to have complete information, while a college worker would be willing to pay \$816.94 more per year for complete information. Given that annual earnings for full-time non-college (college) workers in our sample are \$53,444 (\$83,734), this is a rather low willingness to pay for complete information. The direction of this change, unlike the transition probabilities, will always be negative as we move from CI to our baseline model. Under complete information, individuals' reservation wages are chosen so as to maximize expected lifetime utility. Therefore, any deviation from this optimal behavior must lower lifetime utility.[40] Another way to interpret these numbers is that a non-college (college) worker would be willing to make a *one-time* payment of \$10,953 (\$51,059) to have complete information about the mean offer they would receive.

As we will see next, the welfare losses of information frictions are significantly mitigated by the fact that individuals are able to learn from their labor market experiences. If instead individuals never learned and had their beliefs forever fixed at their initial values, the gains from complete information would be substantially larger.

Benchmark 2: Empirical Priors and No Learning To compare our baseline model to the model with empirical priors, but without learning, we assume that individuals have the same initial beliefs $\mu_{i,t}$, but that they do not update their beliefs in response to their labor market experiences ($\zeta = 0$). We then solve for the same objects as above, using the same parameters as were used in the model with learning, the results of which are reported in the third row of each panel in Table 9. Moving from a world with complete information to a world with heterogeneous beliefs without learning significantly lowers expected utility, for both education groups. On average, welfare falls by about \$1,525 dollars per year for the non-college, and by \$3,174 for the college-educated. That is, a non-college (college) worker who cannot learn from their labor market experiences would be willing to make a *one-time* payment of \$95,303

[39] For reference, Figure 7 depicts the distribution of errors in the full sample.

[40] Note that this does not preclude short-run losses from moving to complete information.

($198,368), respectively, to have complete information. The average willingness to pay in this case is nearly 9 (4) times as high as in the case with learning for non-college (college) workers. This suggests that learning mitigates most of the damage caused by information frictions.

Although the mean losses are economically significant, they mask a considerable amount of heterogeneity, as is evident by the large inter-quartile range of welfare losses presented in the last column of Table 9. Figure 9 plots the welfare gains of moving from the model with no learning to the model with complete information as a function of initial errors in beliefs. If beliefs are initially correct and there is no learning, naturally the gains are zero as the two models are equivalent. This can alternatively be seen in Figure 10, which plots reservation wages in the two environments as a function of initial errors. Here, reservation wages are the same in the two models when beliefs are correct. Moving away from correct beliefs in either direction (pessimism or optimism) raises the welfare gains of complete information: the farther beliefs are from the truth, the more costly are the errors associated with the incorrect reservation wage rules. Moreover, for non-trivial mistakes, the gains are quite large, suggesting that individuals would be willing to pay a significant amount of money for the opportunity to learn.

Comparing Search vs. Information Frictions Search frictions are embodied in the job arrival probabilities λ_u and λ_e, with values less than 1 implying that workers need to wait for job offers. To quantify how information frictions compare to search frictions, we look for the rise in offer arrival probabilities λ_u and λ_e that would deliver the same mean increase in expected utility under our baseline model. The remaining rows of each panel in Table 9 show what happens to transition probabilities and expected utility as either λ_u or λ_e increase from their estimated values by different amounts (2, 10, and 50 percent). For these calculations, we report welfare gains in Table 9 relative to our baseline model.

The increase in expected utility when moving to complete information from our baseline model is roughly equivalent to increasing the arrival probability in non-employment by 5% (15%) for non-college (college) workers. On the other hand, the gains of going from the no-learning case to the complete information world are equivalent to raising the arrival probability in non-employment by about 40% (100%) for the non-college (college) workers, indicating that welfare costs arising from information frictions in the labor market are as large as non-trivial changes in search frictions.

9 Conclusion

We collect rich data on individuals' labor market expectations and realizations to understand the importance of learning and information frictions in the labor market. Although expectations have high predictive power, we find that the errors in forecasts are sizable. However, we docu-

ment that individuals learn from their experiences, and revise their expectations in meaningful ways to new information that arrives. We embed these data and learning patterns into a model of search on- and off- the job, where workers are uncertain about the mean of the offer distribution. We find that the model with heterogeneous beliefs and learning fits the data significantly better than one estimated assuming complete information, and that policy inference would be biased if one incorrectly assumed that workers had complete information. Finally, we show that the welfare gains from eliminating information frictions are nontrivial, and would be even larger if individuals did not learn from their labor market experiences. In the absence of learning, welfare losses are comparable to a decline in the arrival rate of offers in non-employment of more than 40 percent.

This is the first paper to empirically investigate the role of learning in the labor market and to quantify the role of information frictions in job search. We believe there are at least two avenues for future research. The first is related to understanding the learning process better. Although we show that workers learn, our data do now allow us to decompose how much of this is because of workers learning about their own type versus learning about the labor market. We also find a rather high rate of learning in the data. Our panel is short, and we are limited in making progress on this front. Given the importance of learning and information frictions that we document, it would be useful to collect similar data over longer periods. Second, our model focuses on learning on only one dimension, the mean of the offer distribution. Workers may be misinformed about other aspects of the labor market too (such as the arrival rates or the dispersion of the offer distribution). In that sense our setup may provide a lower bound on the importance of information frictions. Work that incorporates these additional dimensions (and collects data to do so) would be useful for understanding job search behavior better and our results highlight the need for collecting such data and relaxing the complete information assumption in labor search models.

References

[1] Armantier, O., G. Topa, W. van der Klaauw, and B. Zafar. 2017. "An Overview of the Survey of Consumer Expectations," Economic Policy Review, forthcoming.

[2] Michèle, Belot, Kircher, Philip, and Muller, Paul. 2016. "Providing Advice to Job Seekers at Low Cost: An Experimental Study on Online Advice," Working Paper.

[3] Burdett, Kenneth. 1978 "A Theory of Employee Job Search and Quit Rates." The American Economic Review 68, no. 1: 212-20.

[4] Burdett, K., and Dale Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment," The International Economic Review, 39(2): 257-273.

[5] Burdett, K., and T. Vishwanath. 1988. "Declining Reservation Wages and Learning," The Review of Economic Studies, 55(4): 655-665 .

[6] Cogley, Timothy and Sargent, Thomas J. 2008. "Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making, " International Economic Review, Vol. 49, Issue 1, pp. 185-221.

[7] Doppelt, Ross 2016. "The Hazards of Unemployment: A Macroeconomic Model of Job Search and Resume Dynamics , " Working Paper.

[8] Faberman, J., A. Mueller, A. Sahin, and G. Topa. 2017. "Job Search Behavior among the Employed and Non-Employed," NBER Working Paper No. 23731.

[9] Fernandez-Blanco, Javier and Edgar Preugschat. "On the Effects of Ranking by Unemployment Duration." European Economic Review 104 (May. 2018) : 92-110 .

[10] Grether, David. 1980. "Bayes Rule as a Descriptive Model: The Representativeness Heuristic.," The Quarterly Journal of Economics, 95 (3): 53757.

[11] Gonzalez, F, and S. Shi. 2010. "An Equilibrium Theory of Learning, Search, and Wages," Econometrica, 78(2): 509-537.

[12] Guvenen, F., F. Karahan, S. Ozkan, and J. Song. 2016. "What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?" Working Paper.

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

[14] Hornstein, A., P. Krusell, and G. Violante. 2011. "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," American Economic Review, Vol. 101(7), 2873-2898.

[15] Heathcote, J., K. Storesletten, and G. Violante. 2014. "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," American Economic Review, Vol. 104(7), 2075-2126.

[16] Hurd, Michael and Kathleen McGarry. 2002. "The Predictive Validity of Subjective Probabilities of Survival," Economic Journal, 112(482): 966-985.

[17] Jacob, Brian and Tamara Wilder. 2011. "Educational Expectations and Attainment," in Whither Opportunity? Rising Inequality and the Uncertain Life Chances of Low-Income Children, edited by Greg J. Duncan and Richard J. Murnane. New York, NY: Russell Sage Press.

[18] Jovanovic, Boyan. 1979. "Job Matching and the Theory of Turnover , " Journal of Political Economy, Vol. 87, No. 5, Part 1, pp. 972-990.

[19] Kahneman, Daniel, and Amos Tversky. 1982. "Subjective probability: A judgment of representativeness," in Daniel Kahneman, Paul Slovic and Amos Tversky (eds) Judgment under Uncertainty: Heuristics and Biases. Cambridge: Cambridge University Press.

[20] Krueger, A., and A. Mueller. 2016. "A Contribution to the Empirics of Reservation Wages," American Economic Journal: Economic Policy, V8(1): 142-179.

[21] Kudlyak, Marianna, Damba Lkhagvasuren and Roman Sysuyev. 2014. "Systematic Job Search: New Evidence from Individual Job Application Data." Working Paper

[22] Lancaster, Tony and Chesher, Andrew. 1983. "An Econometric Analysis of Reservation Wages," Econometrica, vol. 51(6): 1661-1676.

[23] Le Barbanchon, Thomas, Roland Rathelot, and Alexandra Roulet. 2018. "Unemployment Insurance and Reservation Wages: Evidence from Administrative Data," Journal of Public Economics, forthcoming.

[24] Manski, Charles. 2004. "Measuring Expectations." Econometrica, 72(5): 1329-1376.

[25] Mortensen, Dale, and Christopher Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment," Review of Economic Studies, V61(3): 397-415.

[26] Pissarides, C. 1985. "Short-run Dynamics of Equilibrium Unemployment, Vacancies, and Real Wages," American Economic Review, V75(4): 676-690.

[27] Potter, T. 2017. "Learning and Job Search Dynamics during the Great Recession," Working Paper.

[28] Price, Richard H., Amiram D. Vinokur, George Howe, and Robert D. Caplan. 2004. "Preventing Depression in Couples facing Job Loss, "

[29] Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," American Economic Review, Vol. 95(1), 25-49.

[30] Spinnewijn, Johannes. 2015. "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs," Journal of the European Economic Association, Vol. 13(1), 130-167.

[31] Stephens, Mel. 2004. "Job Loss Expectations, Realizations, and Household Consumption Behavior," The Review of Economics and Statistics, Vol. 86(1): 253-269.

[32] Wiswall, M., and B. Zafar. 2015. "Determinants of College Major Choice: Identification using an Information Experiment," Review of Economic Studies, 82, 791824.

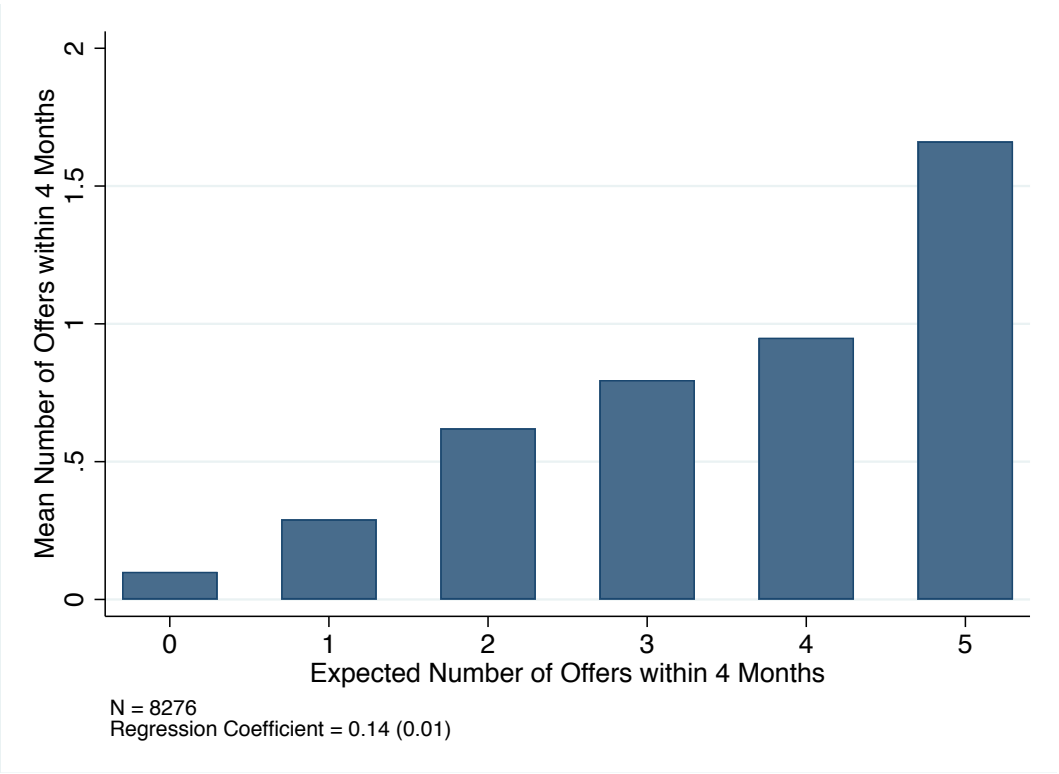


Figure 1: Expectations vs. Realizations: Number of Offers

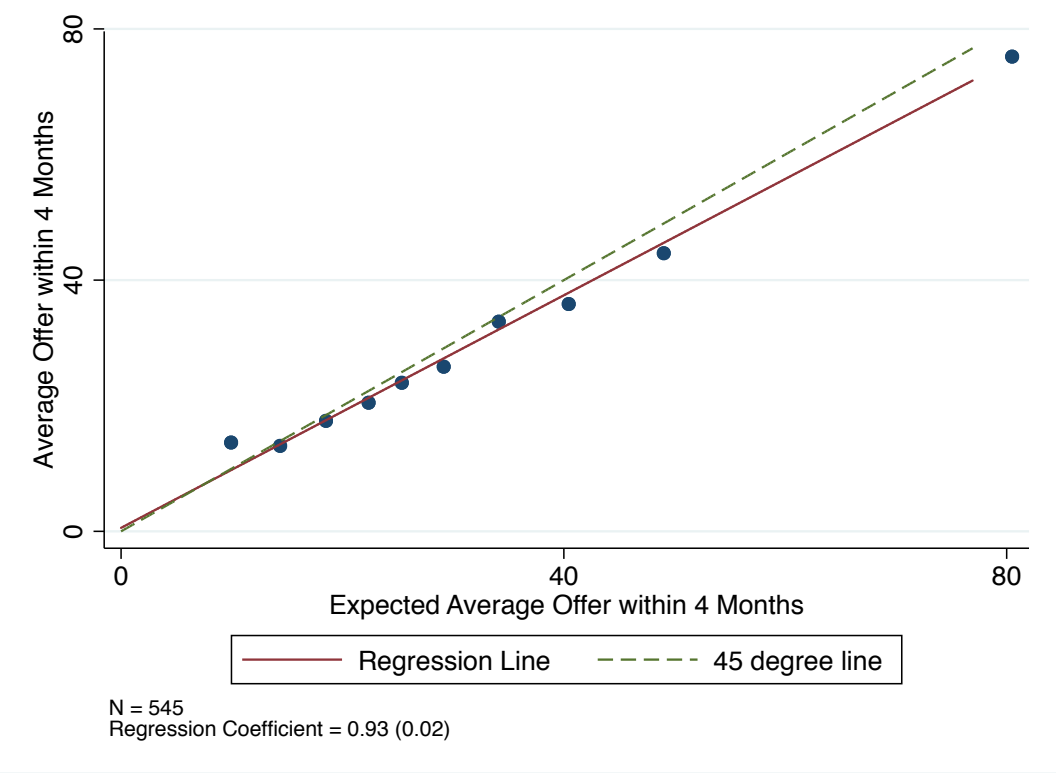


Figure 2: Expectations vs. Realizations: Offer Salary



Figure 3: Distribution of Offer Shocks



Figure 4: Updating Offer Expectations

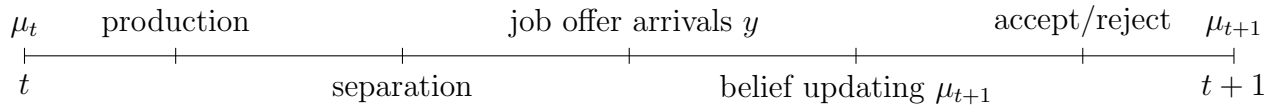
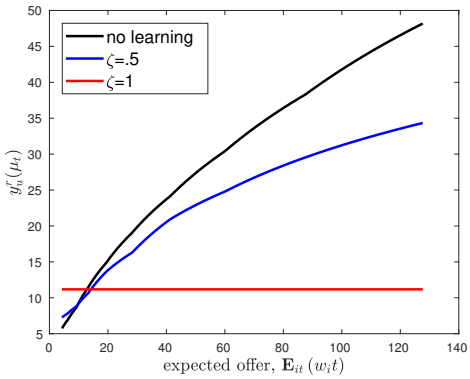
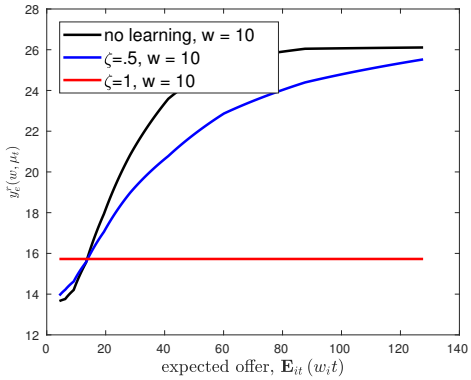


Figure 5: Model Timing



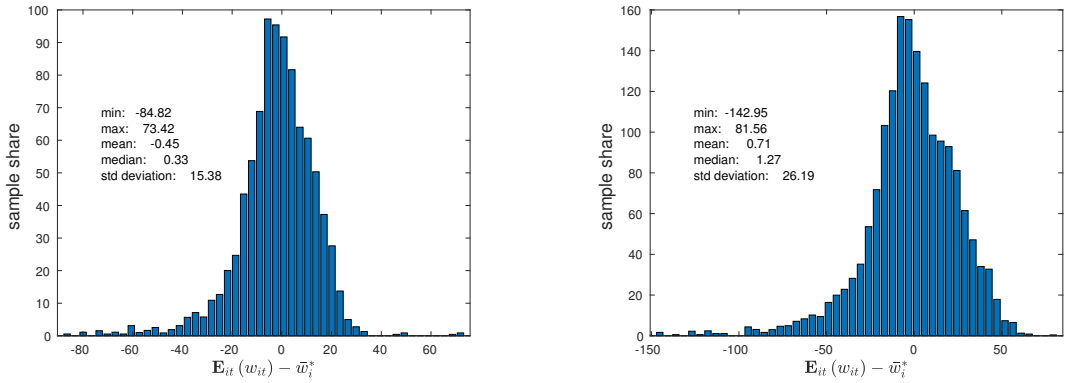
Reservation Wage, Unemployment



Reservation Wage, Employment

Figure 6: The Effect of Beliefs and Learning on Reservation Wages

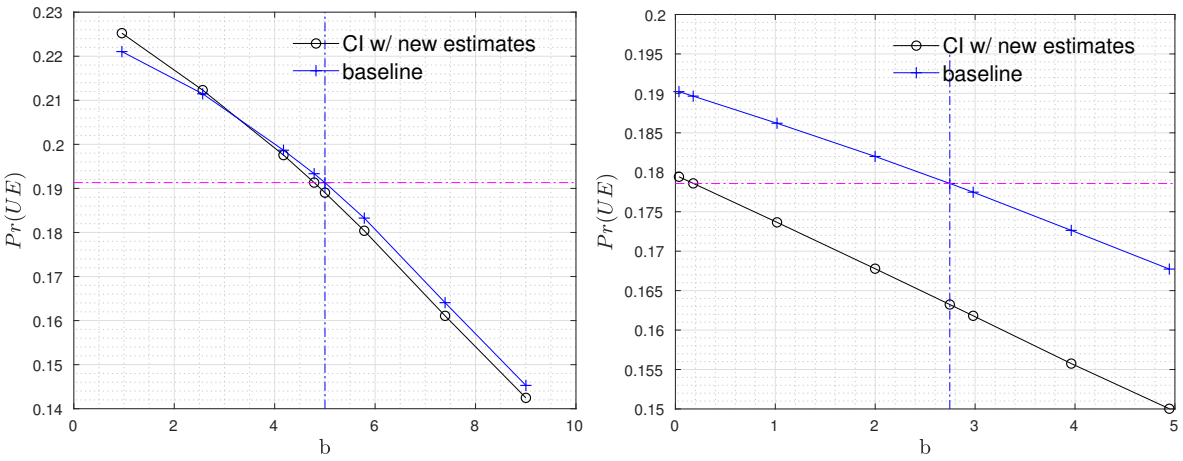
Notes. The left panel plots the reservation wage for unemployed workers for different expected offers for the model with no learning (black line, $\zeta = 0$), for the model with some learning (blue line, $\zeta = .5$), and for the extreme case of updating to current offers (red line, $\zeta = 1$). The right panel plots the reservation wage for the same scenarios for an employed worker at a particular wage level $w = 10$.



Non-college Sample College Sample

Figure 7: Estimated Sample Distribution of Errors

Notes. The left panel plots the distribution of errors for the non-college sample (left panel) and the college sample (right panel). Errors are defined as the deviation in implied means of the log wage offer distribution, $\exp(\mu_i^* + \frac{\sigma^*}{2}) - \exp(\mu_{i,t} + \frac{\sigma^*}{2})$

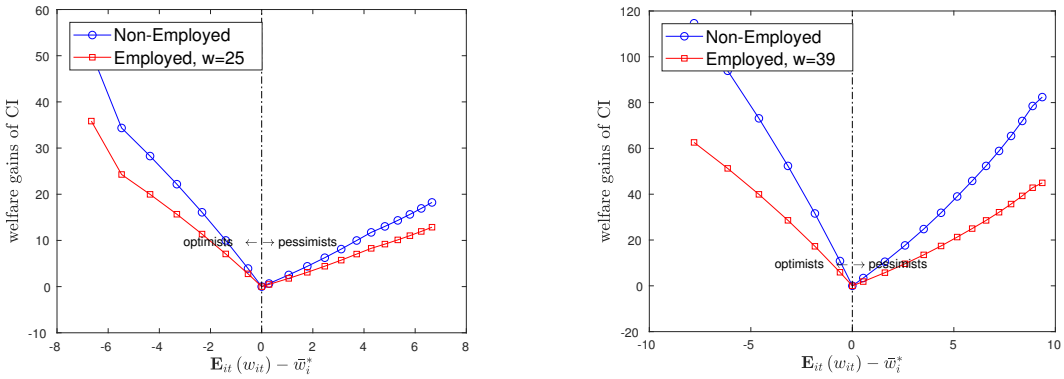


(a) Non-College Sample

(b) College Sample

Figure 8: The Effect of b on Transition Probabilities

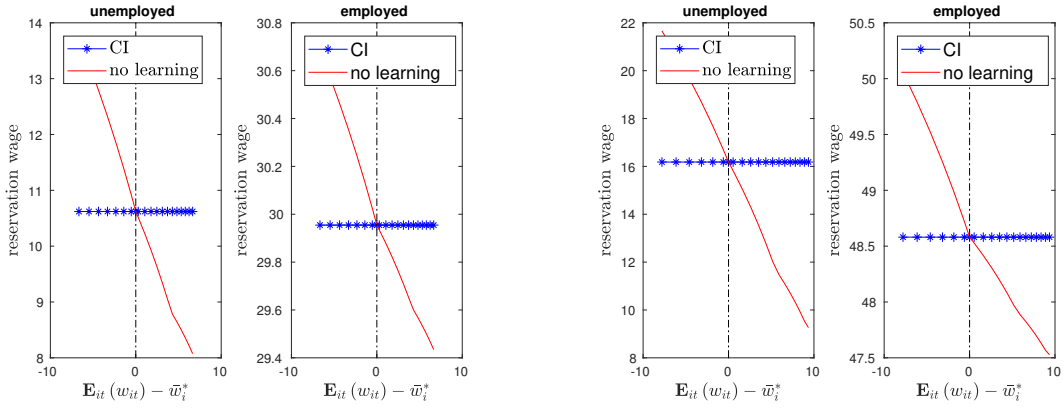
Notes. Panel (a) plots the unemployment to employment transition probability, $Pr(UE)$, for different values of b for both the complete information model (black line) and the model with empirical priors and learning (blue line) for the non-college group. Panel (b) plots the same objects for the college subsample. In all panels, the horizontal dotted line plots the targeted transition probability from the estimation, while the horizontal line corresponds to the estimated b in our baseline model. Note that, while not pictured, each model has a different estimated moving cost as well.



Non-college Sample College Sample

Figure 9: Welfare Gains as a Function of Errors in the Model with no Learning

Notes. The left panel plots the welfare gains for unemployed workers (blue line) and employed workers (red line) for different deviations from a particular μ^* for the non-college sample. The right panel plots the welfare gains for unemployed workers (blue line) and employed workers (red line) for different deviations from a particular μ^* for the college sample. The x-axis is the implied deviations in expected average offers.



Non-college Sample College Sample

Figure 10: Deviations in Reservation Wages as a Function of Errors in the Model with no Learning

Notes. The left panel plots the reservation wage for unemployed workers (left plot) and employed workers (right plot) for different deviations from a particular μ^* for the non-college sample for both the rational expectations model (blue line) and the model with no learning (red line). The right panel plots the reservation wage for unemployed workers (left plot) and employed workers (right plot) for different deviations from a particular μ^* for the college sample for both the rational expectations model (blue line) and the model with no learning (red line). The x-axis is the implied deviations in expected average offers.

Table 1: Demographic Characteristics

		SCE	CPS[a]	p-value[b]
Oberservations		8883	1458056	
Number of unique individuals		4388		
% 1 Survey	Mean	33.61		
% 2 Surveys	Mean	32.18		
Male	Mean	51.80	51.58	0.69
Age	Mean	45.83	44.70	0.00
		(12.12)	(12.30)	
	Min.	22.00	22.00	
	Max.	65.00	65.00	
Bachelor's Degree or higher	Mean	57.31	35.88	0.00
White (Not Hispanic)	Mean	76.82	77.03	0.65
Married/partner[c]	Mean	65.16	58.91	0.00
Number of children	Mean	0.64	0.88	0.00
	Min.	0.00	0.00	
	Max.	20.00	9.00	
Employed	Mean	79.17	74.49	0.00
Unemployed	Mean	3.60	3.24	0.06
OLF	Mean	17.22	22.28	0.00
Unemployment Rate	Mean	4.35	4.16	0.43
Working Full-time[d]	Mean	69.35	64.80	0.00
Working Part-Time	Mean	9.38	8.95	0.17
Full-time Hourly Wage (College)	Mean	40.25	41.95	0.00
Full-time Hourly Wage (Non-college)	Mean	25.69	23.52	0.00

Table reports means, with standard deviations in parentheses.

[a] Statistics are based on the January 2015- September 2017 CPS monthly data.

[b] p-value of the equality of the CPS (in col 2) and SCE (in col 1) means.

[c] The SCE asks "Are you currently married or living as a partner with someone?". CPS respondents are classified as married or living with a partner if they are "Married, spouse present," "Married, spouse absent," or if someone else in their household reported themselves as the head of household's "Partner roommate" or "Unmarried partner."

Table 2: Labor Market Expectations and Realizations - Full Sample

	All	Employed	Unemployed	OLF
Panel A: Expectations				
Received at least 1 offer	0.32	0.32	0.61	0.19
	(0.46)	(0.47)	(0.49)	(0.39)
	[0.00]	[0.00]	[1.00]	[0.00]
	8276	6999	295	982
Number of offers	0.84	0.82	2.04	0.60
	(2.42)	(2.43)	(2.17)	(2.26)
	[0.00]	[0.00]	[2.00]	[0.00]
	6538	5575	228	735
Salary Offer	32.32	34.00	20.15	19.80
	(23.47)	(23.81)	(17.76)	(15.00)
	[24.92]	[27.95]	[14.79]	[14.83]
	5895	5194	262	439
Panel B: Realizations				
Received at least 1 offer	0.19**	0.20***	0.20***	0.09**
	(0.39)	(0.40)	(0.40)	(0.28)
	[0.00]	[0.00]	[0.00]	[0.00]
	7411	6303	257	851
Number of offers	0.34***	0.36***	0.40**	0.17
	(1.00)	(1.01)	(1.12)	(0.81)
	[0.00]	[0.00]	[0.00]	[0.00]
	7411	6303	257	851
Salary Offer	30.43*	30.78*	24.24	28.63
	(25.77)	(25.07)	(31.35)	(32.56)
	[24.04]	[24.44]	[13.72]	[19.41]
	1374	1252	50	72

Table reports means, standard dev in parentheses, and medians in brackets. Sample size in the last row of each variable panel. ***, **, * denote that mean realizations differ from mean expectations at the 1%, 5%, and 10% levels, respectively.

Table 3: Determinants of Job Offer Expectations and Realizations

	Number of Offers	Log Average Offer
Panel A: Expectations		
Unemployed	0.23	-0.27***
	(0.28)	(0.08)
Out of labor force	-0.15*	-0.31***
	(0.08)	(0.04)
Searched for job	1.07***	0.05***
	(0.10)	(0.02)
Searched for job * Non-employed	0.33	-0.08
	(0.35)	(0.09)
Log Current/Most Recent Salary	-0.15***	0.35***
	(0.06)	(0.02)
Male	0.15*	0.12***
	(0.09)	(0.02)
Age	-0.70*	0.11*
	(0.39)	(0.07)
BA or higher	-0.07	0.28***
	(0.07)	(0.02)
Constant	2.53***	-0.81***
	(0.64)	(0.24)
Survey Fixed Effects	Yes	Yes
Mean of Dep. Variable	0.84	3.27
N	6538.00	5895.00
R^2	0.04	0.46
Panel B: Realizations		
Unemployed	-0.27**	0.00
	(0.13)	(0.24)
Out of labor force	-0.16***	-0.02
	(0.03)	(0.09)
Searched for job	0.12***	-0.01
	(0.04)	(0.03)
Searched for job * Non-employed	0.30**	-0.04
	(0.15)	(0.25)
Log Current/Most Recent Salary	-0.02	0.52***
	(0.02)	(0.04)
Male	0.02	0.12***
	(0.03)	(0.03)
Age	-0.41***	0.30**
	(0.12)	(0.12)
BA or higher	-0.02	0.17***
	(0.03)	(0.03)
Constant	0.71***	-2.84***
	(0.18)	(0.40)
Survey Fixed Effects	Yes	Yes
Mean of Dep. Variable	0.34	3.17
N	7411.00	1374.00
R^2	0.01	0.54

Table reports OLS estimates. Robust standard errors (clustered at the individual level) in parentheses: ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 4: Correlates of the Absolute Offer Wage Shock (logs)

	(1)	(2)
Unemployed	-0.031	-0.013
	(0.041)	(0.128)
Out of labor force	0.348***	0.311***
	(0.106)	(0.102)
Searched for job	-0.033	-0.042
	(0.028)	(0.026)
Searched for job * Non-employed		-0.036
		(0.137)
Log Current/Most Recent Salary	-0.133***	-0.126***
	(0.027)	(0.028)
BA or higher	-0.031	0.034
	(0.032)	(0.029)
Age/100	0.338***	0.307***
	(0.121)	(0.113)
Male	-0.074**	-0.037
	(0.029)	(0.027)
Constant		1.523***
		(0.294)
Survey Fixed Effects	Yes	Yes
Mean of Dep. Variable	0.286	0.286
N	545	545
R^2		0.227

OLS estimates presented. Dependent variable is the absolute shock (the gap between the actual offer and the expected offer) in logs. Robust standard errors in parentheses. ***, **, * denote sig. at 1%, 5%, and 10% levels, respectively.

[a] Column (1) shows the estimates from univariate regressions. Each cell in the first column is an estimate from a separate regression. Columns (2) and (3) show estimates from a multivariate regression.

Table 5: Revision of Expected Average Offer, by Education

	(1) All	(2) All	(3) College	(4) Non-College	(5) All	(6) All
Shock (Realization$_t$ - Exp$_{t-1}$)[a]	0.47***	0.47***	0.45***	0.56***		
	(0.06)	(0.06)	(0.07)	(0.12)		
Log Shock (Log Realization$_t$ - Log Exp$_{t-1}$)[a]					0.43***	0.40***
					(0.06)	(0.06)
Survey Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Demographics Included?[b]	No	Yes	Yes	Yes	No	Yes
Mean of dependent variable	-0.13	-0.13	0.05	-0.49	-0.01	-0.01
Observations	493	493	331	162	493	493
R^2	0.32	0.38	0.39	0.56	0.26	0.33

OLS estimates presented. Dependent variable in columns 1-4 is the revision (t minus t-1) expectation of average offer. In columns 5 and 6, the dependent variable is the revision in the log of expectation of average offer. Robust standard errors in parentheses. ***, **, * denote sig. at 1%, 5%, and 10% levels, respectively.

[a] Shock is defined as the realized wage offer in time t minus the expected wage offer reported at t-1.

[b] Demographics include age, household income, numeracy, and dummies for race, gender, parenthood, and education.

Table 6: Determinants of Log Reservation Wages

	(1)	(2)	(3)	(4)
	Employed		Non-Employed	
	College	Non-College	College	Non-College
Log Expected Average Salary Offer	0.53***	0.44***	0.73***	0.64***
	(0.03)	(0.04)	(0.03)	(0.05)
Log Current Salary	0.35***	0.39***		
	(0.03)	(0.04)		
Survey Fixed Effects	Yes	Yes	Yes	Yes
Demographics Included?[a]	Yes	Yes	Yes	Yes
Mean of dependent variable	3.64	3.19	3.14	2.79
Observations	3281	1802	368	312
R^2	0.76	0.66	0.59	0.50

OLS estimates presented. Dependent variable is log reported reservation wage. Robust standard errors (clustered at the individual level) in parentheses. ***, **, * denote sig. at 1%, 5%, and 10% levels, respectively.

[a] Demographics include age, numeracy, and dummies for race, gender, parenthood, and education.

Table 7: Parameter Estimates

Parameter	Description	Non-College	College
β	discount rate	0.984	0.984
δ	separation rate	0.039	0.020
		(0.006)	(0.004)
λ_u	offer arrival rate, in unemployment	0.252	0.292
		(0.042)	(0.039)
λ_e	-, in employment	0.198	0.203
		(0.015)	(0.008)
$\sigma^{\star^2}$	$w \sim logN(\mu_i^\star, \sigma^\star)$	0.317	0.414
		(0.024)	(0.023)
ζ	slope of updating rule	0.534	0.372
		(0.090)	(0.059)
η_1	worker type 1	-0.325	-0.294
		(0.012)	(0.024)
η_2	worker type 2	0.422	0.380
		(0.024)	(0.031)
π_1	probability type 1	0.564	0.564
		(0.003)	(0.003)
b	value of leisure	5.002	2.747
		(2.643)	(1.524)
m	moving cost	65.903	223.330
		(30.181)	(63.337)
$Pr(UE)$	UE transition probability - model	19.130	17.858
$Pr(UE)$	UE transition probability - data	19.130	17.857
$Pr(EE)$	EE transition probability - model	5.449	5.315
$Pr(EE)$	EE transition probability- data	5.450	5.315

The discount rate is set ex ante, while the remaining parameters are estimated from the corresponding moments in our data, following the methods described in Section 6. Bootstrapped standard errors from $B = 100$ bootstraps in parenthesis.

Table 8: Validating the Model: Reservation Wage, By Employment Status

		Employed		Non-Employed	
		Non-College	College	Non-College	College
Baseline	Correlation Coef.	0.577***	0.735***	0.499***	0.547***
	R squared	0.383	0.563	0.292	0.348
Conplete Information	Correlation Coef.	0.523***	0.654***	0.090***	0.106***
	R squared	0.325	0.458	0.077	0.059

The first and third rows show the correlation coefficients between model-predicted reservation wages and self-reported reservation wages in the SCE data, by education and employment. The second and fourth rows show the R^2 for regressions of self-reported reservation wages on model-predicted reservation wages, including controls for demographics as well as industry and time fixed-effects, by education and employment. *** denotes significance at $p < .01$.

a

Table 9: The Welfare Costs of Information Frictions

Non-College

	Pr(UE)	Pr(EE)	Δ Welfare	IQR Δ Welfare
Complete Information	17.48	4.74	0.00	0.00
Baseline	19.13	5.45	-175.24	55.44
Empirical Priors, No Learning	18.56	5.57	-1524.84	1202.60
Bayesian Learning	24.94	6.27	-467.10	638.25
Increase $\lambda_u = 0.25$ in Baseline:				
2 percent	18.45	5.57	104.70	74.75
10 percent	18.04	5.57	499.54	348.84
50 percent	16.67	5.57	1893.40	1610.14
100 percent	15.66	5.57	2009.56	4548.75
Increase $\lambda_e = 0.20$ in Baseline:				
2 percent	18.61	5.56	80.09	88.84
10 percent	18.80	5.52	388.40	439.10
50 percent	19.99	5.35	1718.43	1952.51
100 percent	21.58	5.20	3050.89	3466.39

College

	Pr(UE)	Pr(EE)	Δ Welfare	IQR Δ Welfare
Complete Information	13.97	4.18	0.00	0.00
Baseline	17.86	5.32	-816.94	343.54
Empirical Priors, No Learning	18.12	5.50	-3173.89	2606.80
Bayesian Learning	29.17	6.39	-434.10	1763.05
Increase $\lambda_u = 0.29$ in Baseline:				
2 percent	18.01	5.50	145.83	106.43
10 percent	17.59	5.50	701.43	498.00
50 percent	16.13	5.50	2738.40	2477.22
100 percent	14.98	5.50	3245.63	6451.30
Increase $\lambda_e = 0.20$ in Baseline:				
2 percent	18.16	5.49	233.07	266.06
10 percent	18.32	5.45	1127.93	1322.11
50 percent	19.19	5.26	4469.56	6238.61
100 percent	20.42	5.06	8817.79	10416.73

The first column reports the one-period-ahead UE transition probability for each model, while the second column does the same for EE transition probability. The third column reports the mean difference in present discounted value of income (expected utility) as described in Equations 15 and 16 (converted to an annual equivalent by multiplying by $2000 * (1 - \beta)$) relative to the complete information model in the first three rows, and relative to the baseline model in the rows moving arrival rates. The final column reports the interquartile range of these gains.

A Additional Tables and Figures

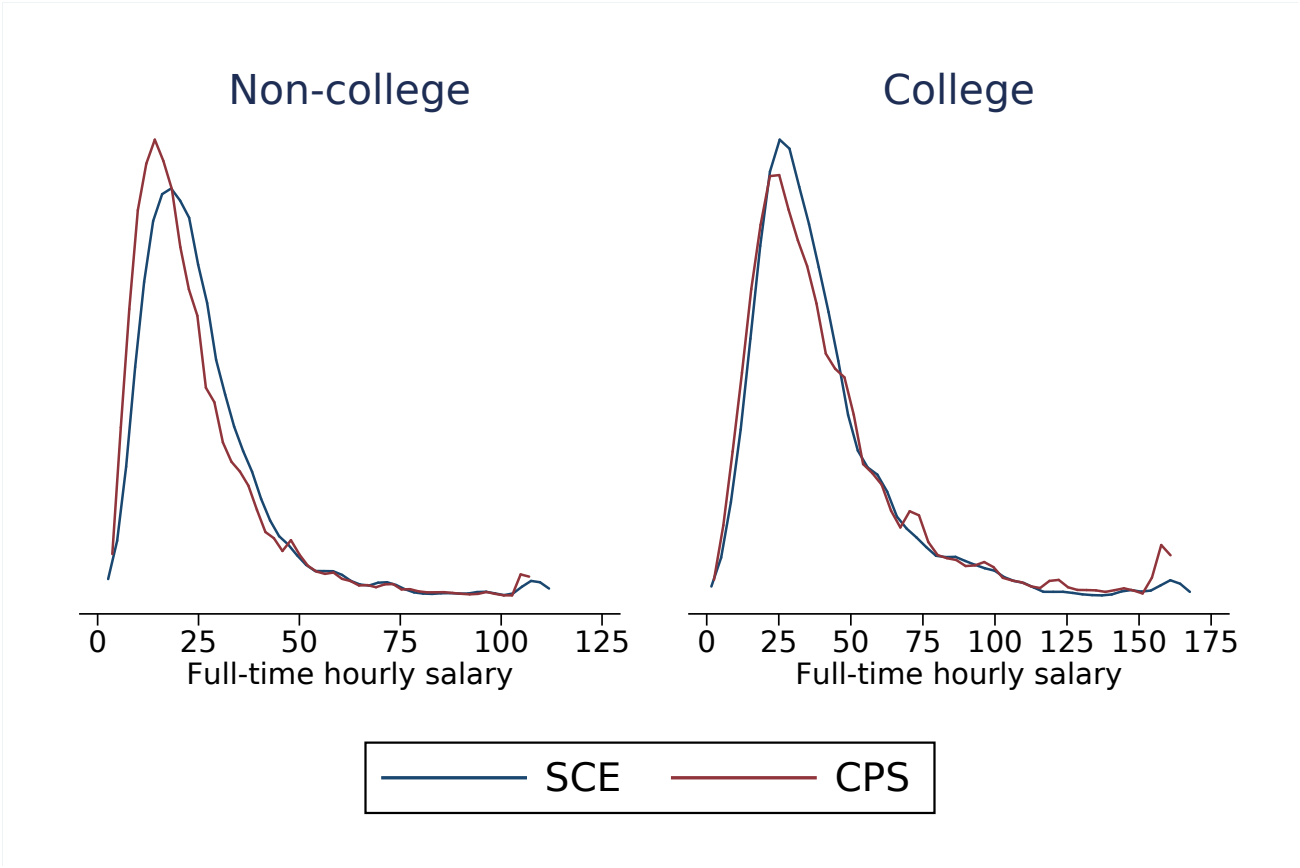


Figure A1: Comparing salary distributions in the CPS and SCE

Table A1: Labor Force Transition Rates

		Same Employer	New Employer	Unemployed	OLF	N
Panel A: SCE 4-month transition rates						
LFS (t=0)	Employed	92.7	4.6	1.5	1.3	3083
	Unemployed		29.1	56.4	14.4	118
	OLF		5.1	2.2	92.7	419
Panel B: CPS 3-month transition rates						
LFS (t=0)	Employed	90.0	4.9	1.8	3.2	369301
	Unemployed		36.5	38.8	24.7	19709
	OLF		8.6	3.5	87.9	106987

Panel A shows the percent of SCE respondents in the labor force status listed in the rows that transition into the labor market status listed in the column headings in the next four months. Panel B shows the percent of CPS respondents in the labor force status listed in the rows that transition into the labor market status listed in the column headings in the next three months.

Table A2: Labor Market Expectations and Realizations by Current Employment Status - Consistent Sample

	All	Employed	Unemployed	OLF
Panel A: Expectations				
Received at least 1 offer	0.32	0.33	0.54	0.17
	(0.47)	(0.47)	(0.50)	(0.38)
	[0.00]	[0.00]	[1.00]	[0.00]
	3670	3134	118	418
Number of offers	0.86	0.85	1.84	0.58
	(2.46)	(2.50)	(2.46)	(1.93)
	[0.00]	[0.00]	[1.00]	[0.00]
	2863	2489	85	289
Salary Offer	32.22	33.88	21.46	23.39
	(21.43)	(21.81)	(17.09)	(12.85)
	[27.18]	[29.11]	[16.83]	[24.26]
	545	468	49	28
Panel B: Realizations				
Received at least 1 offer	0.17***	0.17***	0.43***	0.12
	(0.38)	(0.37)	(0.50)	(0.33)
	[0.00]	[0.00]	[0.00]	[0.00]
	3670	3134	118	418
Number of offers	0.28***	0.28***	0.73***	0.16
	(0.89)	(0.88)	(1.37)	(0.72)
	[0.00]	[0.00]	[0.00]	[0.00]
	2863	2489	85	289
Salary Offer	30.39	31.54	23.00	24.05
	(22.63)	(23.35)	(16.48)	(15.56)
	[24.92]	[25.77]	[15.45]	[19.64]
	545	468	49	28

Table reports means, standard dev in parentheses, and medians in brackets. Sample size in the last row of each variable panel. Sample is restricted to respondents who report both an expectation(Panel A) of the outcome and a realization (Panel B) in the subsequent survey. ***, **, * denote that mean realizations differ from mean expectations at the 1%, 5%, and 10% levels, respectively.

Table A3: Revision of Expected Average Offer, by Demographics

	(1) Age < 45	(2) Age ≥ 45	(3) Did Not Search	(4) Searched[d]	(5) Positive Shock[e]	(6) Negative Shock	(7) Low Uncertainty[f]	(8) High Uncertainty
Shock (Realization$_t$ - Exp$_{t-1}$)[a]	0.39***	0.55***	0.47***	0.43***	0.42***	0.47***	0.49***	0.40***
	(0.07)	(0.08)	(0.09)	(0.09)	(0.16)	(0.09)	(0.08)	(0.10)
Survey Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demographics Included?[b]	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean of dependent variable	0.22	-0.55	-0.34	0.91	3.93	-3.01	-0.65	0.92
Observations	274	219	260	183	205	288	317	175
R^2	0.40	0.49	0.39	0.42	0.23	0.39	0.43	0.43
p-value[c]	0.10		0.65		0.77		0.48	

OLS estimates presented. Dependent var is the revision (t minus t-1) expectation of average offer. Robust standard errors in parentheses. ***, **, * denote sig. at 1%, 5%, and 10% levels, respectively.

[a] Shock is defined as the realized wage offer in time t minus the expected wage offer reported at t-1.

[b] Demographics include age, household income, numeracy, and dummies for race, gender, parenthood, and education.

[c] p-value of the equality of the estimate on the shock variable for each pair of columns (1 and 2; 3 and 4; 5 and 6; 7 and 8).

[d] Searched is the subsample that did anything in the last for months to look for new work.

[e] Positive shock is the subsample which underestimates the offer wage.

[f] Low Uncertainty is the set of respondents who assign non-zero probability to below-median number of bins when reporting the subjective distribution of offer wages expected in the next four months.

Table A4: Complete Information Parameter Estimates

Parameter	Description	Non-College	College
β	discount rate	0.984	0.984
δ	separation rate	0.039	0.020
		(0.006)	(0.004)
λ_u	offer arrival rate, in unemployment	0.252	0.292
		(0.042)	(0.039)
λ_e	-, in employment	0.198	0.203
		(0.015)	(0.008)
$\sigma^{\star^2}$	$w \sim logN(\mu_i^\star, \sigma^\star)$	0.317	0.414
		(0.024)	(0.023)
ζ	slope of updating rule	0.534	0.372
		(0.090)	(0.059)
η_1	worker type 1	-0.325	-0.294
		(0.012)	(0.024)
η_2	worker type 2	0.422	0.380
		(0.017)	(0.031)
π_1	probability type 1	0.564	0.564
		(0.003)	(0.003)
b	value of leisure	4.786	0.178
		(2.595)	(0.137)
m	moving cost	47.754	137.737
		(19.967)	(41.883)
$Pr(UE)$	UE transition probability - model	19.132	17.859
$Pr(UE)$	UE transition probability - data	19.130	17.857
$Pr(EE)$	EE transition probability - model	5.449	5.315
$Pr(EE)$	EE transition probability- data	5.450	5.315

The discount rate is set ex ante, while the remaining parameters are estimated from the corresponding moments in our data, following the methods described in Section 6. Bootstrapped standard errors from $B = 100$ bootstraps in parenthesis.

B Bayesian Updating

To compute the Bayesian updating parameter, we use the following question from SCE labor market survey:

"*Think again about the job offers that you may receive within the coming four months. What do you think is the percent chance that the job with the best offer will have an annual salary for the first year of...*

*Less than .8*X dollars* _____
*Between .8*X and .9*X dollars* _____
*Between .9*X and X dollars* _____
*Between X and 1.1*X dollars* _____
*Between 1.1*X and 1.2*X dollars* _____
*More than 1.2*X dollars* _____"

where X is the respondent's answer to the question:

"*Think about the job offers that you may receive within the coming four months. Roughly speaking, what do you think the annual salary for the best offer will be for the first year? Note the best offer is the offer you would be most likely to accept.* "

For each respondent, we fit a log-normal distribution to the answers to the first question. Let the standard deviation of this estimated distribution be ϕ_i. An individual's Bayesian updating parameter is then

$$\zeta_i = \frac{n_i \phi_i^2}{\sigma^{*2} + n_i \phi_i^2},$$

where σ^* is the variance of the log offer wage distribution, and is calculated as described in Appendix F. n_i is the number of offers the individual i received.

C Proofs

We make the following three assumptions:

Assumption 1. *If* $y > y'$, *then* $W(y, \mu_{t+1}(y)) - W(y', \mu_{t+1}(y')) > U(\mu_{t+1}(y)) - U(\mu_{t+1}(y'))$

Assumption 2. *If* $y > y'$, *then* $W(y, \mu_{t+1}(y)) - m - (W(y', \mu_{t+1}(y')) - m) > W(w, \mu_{t+1}(y)) - W(w, \mu_{t+1}(y'))$ $\forall w$

Assumption 3. $\exists$ $\hat{w}$ *such that* $U(\mu_{t+1}(\hat{w})) > W(y, \mu_{t+1}(\hat{w}))$.

The first two assumptions mimic the second assumption in Burdett and Vishwanath (1988), but account for search on the job. They both imply that receiving a larger offer never increases the expected returns from search in the current state by "too much". The third assumption plays the same role as their assumption that search is always profitable. Informally, an easy way to see why one needs these assumptions is to differentiate the left hand and right hand sides of the reservation wage rules. Reservation wage strategies will only make sense if for every y above the reservation wage, the value of employment exceeds the value of non-employment. These assumptions guarantee that will be the case. More formally,

Proposition 1. *If the above assumptions hold, for any μ_t there exists a reservation wage in non-employment $y_u^r \equiv y_u^r(\mu_t)$ such that*

$$W(y, \mu_{t+1}) \lesseqgtr U(\mu_{t+1}) \quad as \quad y \lesseqgtr y_u^r(\mu_t)$$

and a reservation wage in employment $y_e^r \equiv y_e^r(w, \mu_t)$ for every w such that:

$$W(y, \mu_{t+1}) - m \lesseqgtr W(w, \mu_{t+1}) \quad as \quad y \lesseqgtr y_e^r(w, \mu_t)$$

Proof. Start with the reservation wage in non-employment. Suppose that y_u^r exists and satisfies (6). Then for any $y > y_u^r$ it follows that $W(y, \mu_{t+1}(y)) > U(\mu_{t+1}(y))$ and the offer is accepted. Similarly, for any $y < y_u^r$ it follows that $W(y, \mu_{t+1}(y)) < U(\mu_{t+1}(y))$ and the offer is rejected. To show existence, note that $h(y) = U(\mu_{t+1}(y)) - W(y, \mu_{t+1}(y))$ is strictly decreasing in y since the first assumption states that the derivative of U with respect to y is less than the derivative of W with respect to y. The third assumption implies that $h(\hat{w}) > 0$. By the first assumption, $h(y) < 0$ as $y \to \infty$. Therefore, $\exists\ y_u^r$ such that $W(y_u^r, \mu_{t+1}) = U(\mu_{t+1})$. A similar proof holds for the reservation wage in employment. □

D Computation

The model is solved as follows. Assume there is a belief grid for μ^*, $\mathbf{p}$, with N_p points, and we discretize the wage offer distribution to N_w^l grid points, but we approximate the value functions on $N_w \leq N_w^l$ of those points. We are then looking for a numerical approximation at (N_w, N_p) points for $W(\cdot, \cdot)$ and at (N_p) points for $U(\cdot)$. To solve for the value functions, we proceed as follows:

1. For a good initial guess, solve the Anticipated Utility model (Cogley and Sargeant (2008)) which has the following value functions:

$$U_{AU}(\mu_t) = b + \beta\lambda_u^* \int_y \max\{W_{AU}(y, \mu_t), U_{AU}(\mu_t)\} dF_t(y) + \beta(1 - \lambda_u^*)U(\mu_t) \tag{17}$$

$$W_{AU}(w,\mu_t) = w + \beta\delta^* U_{AU}(\mu_t)$$
$$+\beta(1-\delta^*)\lambda_e^* \int_y \max\{W_{AU}(y,\mu_t) - m, W_{AU}(w,\mu_t)\} dF_t(y)$$
$$+\beta(1-\delta^*)(1-\lambda_e^*) W_{AU}(w,\mu_t) \tag{18}$$

These are almost the same value functions as our model, but they do not incorporate that an individual takes into account the effect that offers have on future beliefs (so everything on the right hand side remains μ_t rather than μ_{t+1}. We use the solution to this model as the starting guess for the values $W(\cdot,\cdot)$ and $U(\cdot)$ for the true model.

2. Use these values and the corresponding reservation wages from the Anticipated Utility model as starting guesses, proceed via value function iteration conditional on a guess for reservation wages. That is, start with some guess for the reservation wages (the first of which comes from the AU model just described). Given these policy rules, use the initial guess of the value functions to evaluate the right-hand-side of Equations (4) and (5). To do so, we need the value of employment for each possible offer, taking into account what effect that offer has on future beliefs. We get this value by interpolating the guess at the offer y and its corresponding belief $\mu_{t+1}(y,p)$ given some current belief p and some model-specific updating rule.

3. Proceed via value function iteration until convergence.

4. Update the reservation wages to be consistent with the new value functions following Equations (6) and (7), and repeat the above steps again.

5. Continue in this way until both wages and values have converged.

6. Once the model has converged, we have approximations for $W(w,p)$, $U(p)$, $y_e^r(w,p)$ and $y_u^r(p)$.

E Simulation

To simulate the model, we proceed as follows. For any set of parameters, we first solve the model following Section D above. Using the policy functions, for each individual in our sample, we can interpolate on the reservation wage functions at each individual's $\mu_{i,t}$ (or $\mu_{i,t}$ and current wage, if employed) to get their model-implied reservation wages. Given these reservation wages and the assumed parameters which govern the offer distribution for each individual, μ_i^* and σ_i^*, we can compute the probability an unemployed individual transitions from unemployment to employment as $\lambda_u^*(1 - F^*(y_u^r(\mu_{i,t})))$, and the probability that an employed individual at some wage w transitions to another employer, $(1-\delta^*)\lambda_e^*(1 - F^*(y_e^r(w,\mu_{i,t})))$.

F Estimating Individual Wage Offer Distributions

To estimate the parameters for individual wage offer distributions μ_i^*, σ^*, we take the following steps separately for each education group. First, we regress log wage offers on observable characteristics:

$$ln(y_i) = \alpha' X_i + \eta_i + \sigma^* \epsilon_i, \quad \epsilon_i \sim N(0,1)$$

to obtain an estimate for α, $\hat{\alpha}$. Given this estimate, we then search for values of $\eta_1, \eta_2, \pi_1, \sigma^*$ to minimize the distance between model-implied percentiles of the wage offer distribution (percentiles $\{5, 10, 15, ..., 95\}$), and their empirical counterparts. Throughout, we apply a normalization to η_2 to ensure that $\mathbf{E}(\eta_i) = 0$. Once we have these estimates, we assume that an individual's (log) wage offer distributed is normally distributed with mean $\mu_i^* = \hat{\alpha} X_i + \eta_1$ with probability π_1 and with mean $\mu_i^* = \hat{\alpha} X_i + \eta_2$ with probability $1 - \pi_1$. In both cases the variance is given by σ^*, which is assumed to vary by education, but not by individual.

G Complete Information Model

In the complete information model, the value of unemployment is given by:

$$U_{CI} = b + \beta \lambda_u^* \int_y \max\{W_{RE}(y), U_{RE}\} dF^*(y) + \beta(1 - \lambda_u^*) U_{RE} \tag{19}$$

while the value of employment at some wage w is given by:

$$\begin{aligned} W_{CI}(w) = w &+ \beta \delta^* U_{RE} \\ &+ \beta(1-\delta^*)\lambda_e^* \int_y \max\{W_{RE}(y) - m, W_{RE}(w)\} dF^*(y) \\ &+ \beta(1-\delta^*)(1-\lambda_e^*) W_{RE}(w) \end{aligned} \tag{20}$$

We estimate this model by solving the same steps as we do for our learning model, except we do not use any of the information on beliefs. Table A4 reports the parameter estimates for the complete information model. Note that, with the exception of b, m all parameters are the same since they are independent of the model specification. Regarding the new estimates for b, m in the complete information framework, moving costs are lower for both groups. This is because in our sample, lower wage workers (the set of workers likely to move to higher paying jobs and thus likely to be affected by moving costs) tend to be pessimistic. All else equal, they would therefore move less under complete information. Therefore, to match the data under complete information, moving costs are lower to induce more movements. Given the lower moving costs,

the value of leisure then adjusts to match the transition rates from non-employment in the data. In general, the direction that the value of leisure will move to under complete information is ambiguous, and depends in part on what happens to the newly estimated moving cost, since moving costs also affect reservation wages in non-employment.