Learning and Job Search Dynamics during the Great Recession*

Tristan Potter†
Drexel University

May 26, 2020

Abstract

Krueger and Mueller (2011) document that search effort declined with unemployment duration during the Great Recession. I show that variation in past effort explains this decline. Furthermore, job offers increase subsequent effort. These facts are inconsistent with standard models of search. I introduce a model of sequential search in which workers are uncertain about the offer arrival process and learn through search. Evolving beliefs influence search through two competing channels: the opportunity cost of leisure and the option value of unemployment. Estimation of the model indicates that learning provides a strong account of job search dynamics during the Great Recession.

Keywords: unemployment; sequential search; learning
JEL Classification: J64, E24, D83

*Special thanks to Sanjay Chugh, Ryan Chahrour, and Peter Ireland. Thanks to Susanto Basu, Kit Baum, Simon Gilchrist (discussant), André Kurmann, Andreas Mueller, Tzuo-Hann Law, Fabio Schiantarelli, Shu Lin Wee (discussant), and Matthew Weinberg for insightful comments and conversations. Thanks also to seminar participants at the Federal Reserve Bank of Philadelphia, Wharton, Barcelona GSE Summer Forum (IFU Workshop), Boston College, BLS, Federal Reserve Board, University of Edinburgh, Drexel University, UIUC, University of Mississippi, University of Calgary, CBO, and Middlebury College for valuable input. Any remaining errors are my own.

†School of Economics, LeBow College of Business, Drexel University. Mail: 3220 Market St., 10^{th} floor, Philadelphia, PA 19104. Telephone: 215-895-2540. Email: tristan.l.potter@drexel.edu.

1 Introduction

Why did search effort fall with unemployment duration during the Great Recession? This paper seeks to answer that question by way of three substantive contributions: First, using high-frequency longitudinal data on individuals' search decisions, I show that declining search effort over the unemployment spell—as documented by Krueger and Mueller (2011)—is explained by variation in search effort since job loss. I provide evidence from data on job offers that this reflects job seekers learning about the stochastic process governing the arrival of offers. Second, I propose a model of sequential search to rationalize the empirical results. The model is tractable and sheds light on the mechanisms through which uncertainty and learning influence individuals' search decisions. Finally, I structurally estimate the model and show that learning provides a strong account of the reduced-form empirical results, as well as a number of other dimensions of job search data from the Great Recession. Using only data on realized search decisions and outcomes, I estimate that job losers were excessively optimistic about their job-finding prospects, thus corroborating evidence from existing studies based on subjective probability elicitations about the job-finding process.

The paper begins with an empirical study of job-search dynamics during the Great Recession. The jumping off point is an important study by Krueger and Mueller (2011). They use high-frequency longitudinal data on job search from the Survey of Unemployed Workers in New Jersey (SUWNJ) to document that job seekers' search effort fell over the unemployment spell during the Great Recession. I revisit their analysis and show that the decline in search effort over the unemployment spell is due to variation in individuals' search effort since job loss: When search effort is allowed to depend on both unemployment duration and cumulative search effort since job loss, the former drops out of the model while the latter enters with a highly significant and negative coefficient. To investigate the mechanism driving this result, I turn to data on job offers in the SUWNJ. These data show that search effort jumps discretely after a job offer is received, suggesting that the result is driven by uncertainty and learning about the stochastic process governing the arrival of offers.

These results are important for two principal reasons. First, they suggest that randomness inherent to the search process may induce systematic changes in behavior that impinge on subsequent job-finding prospects. Second, the results challenge a fundamental assumption underlying the canonical theory of sequential job search: that job seekers have complete information about the rate at which job opportunities will arrive during unemployment. In contrast, the findings in this paper suggest that job seekers are uncertain about the availability of work, and that search decisions are driven by learning from experience.

Motivated by this evidence, I introduce a model of sequential search under uncertainty and learning. In the model, at the beginning of each week of unemployment, job seekers choose the amount of time for which a Poisson process with an unobservable arrival rate parameter will run. The amount of time allocated to the Poisson search process, as well as the number of job offers realized during that time, jointly comprise all of the relevant new information made available to job seekers through search. At the end of each week, job seekers update their beliefs to reflect this new information and, if no offer has arrived or if an offer has been rejected, proceed to the next week of unemployment and reoptimize in light of their updated beliefs. Search dynamics during unemployment are thus driven by the dynamic interaction between search decisions and beliefs: Beliefs respond rationally to the outcomes of past search, while search decisions are driven by endogenously evolving beliefs.

This model is the first to integrate learning about the arrival of job offers into a dynamic framework suitable for studying the contours of job search over the spell of unemployment. Its tractability enables transparent characterization of time devoted to job search and the reservation wage at any duration of unemployment in terms of cumulative past search and the stock of job offers received. Indeed, I show that a first-order approximation of the structural model implies the reduced-form regression equation described above. I use the model to decompose the effect of learning on job search into two components: Failing to find work exerts a negative influence on search by reducing the perceived opportunity cost of leisure in the current period, but also stimulates search by reducing the option value of unemployment in future periods. Because the relative strength of these effects varies endogenously as unemployment progresses and job seekers observe the stochastic outcomes of their effort, the model generates rich—and potentially nonmonotonic—search dynamics over the unemployment spell.

Finally, I return to the data and structurally estimate the parameters of the model via a simulated minimum distance procedure. I identify key structural parameters—including those governing the bias in beliefs at the time of job loss—using the reduced-form parameter estimates obtained in the first part of the paper, as well as average search effort, the average arrival rate of job offers, and the average acceptance rate of offers (among those receiving offers) from the SUWNJ data. The estimated model provides a strong account of the data, and in particular, matches the observed effect of cumulative past search and job offers on subsequent search effort. Interestingly, I find that job seekers overestimate their job-finding prospects by roughly 60% at the time of job loss. This result is consistent with evidence from existing studies that also find that job seekers are overly optimistic. However, while existing work typically arrives at this result by exploiting subjective probability elicitations concerning the job-finding process, I identify excessive optimism based on search decisions and outcomes alone.

Understanding the determinants of individuals' search decisions is important. To the extent that search determines the likelihood of finding work, it is inextricably tied to the persistence of income loss associated with unemployment, and thus welfare. Labor-market policies that seek to address unemployment—such as the provision of unemployment insurance benefits, job-search-assistance programs, and employment subsidies—are necessarily predicated on assumptions about the factors that influence individuals' search decisions. This paper argues that the randomness intrinsic to the search process is itself essential to understanding individuals' search decisions and behavior throughout unemployment.

The remainder of the paper is organized as follows. Section 2 discusses related literature. Section 3 studies the empirics of job search during the Great Recession. Section 4 develops the theoretical model and characterizes search and reservation wage dynamics. Section 5 structurally estimates the model. Section 6 concludes. All supplementary material may be found in the online appendix.

2 Related Literature

This paper lies at the intersection of an empirical literature that seeks to understand the determinants of individuals' job-search decisions during unemployment and a theoretical literature that seeks to integrate learning into models of job search.

The paucity of high-frequency longitudinal data on job search has hampered attempts to study the determinants of individuals' job-search decisions during unemployment. Nonetheless, several recent papers have attempted to fill this void. Shimer (2004) uses Current Population Survey (CPS) data to study the determinants of job search in the United States prior to the Great Recession. He measures search effort as the number of reported search methods among actively searching respondents, and finds a hump-shaped profile of job search over the first year of unemployment, peaking at roughly 20 weeks. Mukoyama et al. (2018) update Shimer's analysis by exploiting overlap between the CPS and the American Time Use Survey (ATUS). They construct time-intensity weights for each of the search methods considered in the CPS and use the weights to impute search time for the full CPS sample. Their results corroborate Shimer's findings that search exhibits a hump-shaped profile over the spell of unemployment. Krueger and Mueller (2011) use the SUWNJ—the same data set used in this paper—to show that during the Great Recession, time devoted to job search fell monotonically over the course of unemployment.

This paper complements previous work by developing evidence that, at least during the Great Recession, search decisions were significantly influenced by individuals' experiences while searching for work. Moreover, the theoretical mechanism described herein provides a unified explanation for the fact that job-search effort declined monotonically during the Great Recession, but exhibited a hump-shaped profile in the years prior to the Great Recession: When job seekers' beliefs are sufficiently pessimistic, search declines monotonically throughout unemployment. In contrast, when beliefs are not too pessimistic—as may have been the case prior to the Great Recession—the model implies hump-shaped dynamics qualitatively similar to those documented by Shimer (2004) and Mukoyama et al. (2018).

The paper is also closely related to a literature that studies search in the context of learning. Early examples include Rothschild (1974) and Burdett and Vishwanath (1988), who study search when individuals have incomplete information about the distribution of prices or wages. In this paper, I study search when job seekers have incomplete information about the distribution of offer-arrival rates. Falk et al. (2006a) present evidence from a laboratory experiment that job seekers exhibit substantial uncertainty about their job-finding prospects, and update beliefs based on search outcomes. In a companion paper, Falk et al. (2006b) develop an equilibrium model in which job seekers learn about their linear job-finding probability. Results from the laboratory experiment broadly conform to the message in this paper. However, their theoretical model only accommodates an extensive margin of search, and thus is constrained to focus only on the implications of learning for aggregate labor-market dynamics. This paper, by contrast, is concerned with the contours of the *intensive* margin of search throughout the course of unemployment and with clarifying the underlying mechanisms through which evolving beliefs govern search decisions.

3 Empirics of Job Search during the Great Recession

In this section I study some empirical aspects of job search during the Great Recession using high-frequency longitudinal data on the job search decisions of the unemployed from the SUWNJ.[1]

3.1 Survey description and sample

The SUWNJ is a weekly longitudinal survey of unemployment insurance (UI) benefit recipients in New Jersey beginning in the fall of 2009 and continuing through early 2010. The survey was conducted by the Princeton University Survey Research Center and the data have been made publicly available. The survey covers 6,025 unemployed job seekers for up to 24 weeks for a total of 39,201 weekly interviews. Sampled individuals were asked to participate in a weekly online survey that lasted for a minimum of 12 weeks and, for the long-term unemployed, up to 24 weeks. The weekly survey consisted of questions pertaining to job-search activity, time use, job offers, and consumption. See the online appendix for a more complete description of the survey, and Krueger and Mueller (2011) for a comprehensive description of methodology.

3.2 Evidence from search histories

Using SUWNJ data, Krueger and Mueller (2011) show that job seekers' search effort fell throughout the unemployment spell during the Great Recession. One possible explanation for this observation is that failed job search is discouraging: Repeatedly trying, and failing, to find work creates the impression that suitable work is not available, thus causing job seekers to give up looking. If this is driving the observed decline in search effort during the Great Recession, then search decisions should depend on total time spent searching for work since the time of job loss, not unemployment duration *per se*. I exploit the high-frequency longitudinal nature of the SUWNJ to examine this hypothesis.

[1] All supplementary material may be found in the online appendix.

3.2.1 Empirical strategy

Consider expressing time devoted to job search as a function of unemployment duration, a measure of aggregate labor market slack, and—reflecting the preceding intuition—the total time spent looking for work since job loss:

$$s_{it} = \iota + \kappa d_{it} + \pi \sum_{\tau=1}^{t-1} s_{i\tau} + \gamma u_t + \eta_i + \epsilon_{it}. \tag{1}$$

For individual i in interview week t, s_{it} denotes minutes per day spent on job search, d_{it} denotes unemployment duration, u_t is the (seasonally adjusted) New Jersey unemployment rate, and η_i is a person-specific fixed effect.[2] The coefficients of primary interest are κ and π, which measure the impact of duration and cumulative past search, respectively, on time spent searching for work.

Because of the cohort structure of the data, no individuals in the sample are observed from the beginning of the unemployment spell. This implies that cumulative past search—the variable of interest—is only partially observed, so equation (1) cannot be estimated directly. Accordingly, I take first differences of (1) to clean out all unobservable person-specific terms:

$$\Delta s_{it} = \kappa \Delta d_{it} + \pi s_{it-1} + \gamma \Delta u_t + \Delta \epsilon_{it}. \tag{2}$$

The presence of the lagged-dependent variable on the right-hand side of (2) now gives rise to an endogeneity problem common to dynamic panel models: $\mathbb{E}\left[s_{it-1}\Delta\epsilon_{it}\right] \neq 0$. I address the endogeneity of s_{it-1} by instrumenting with its first lag, s_{it-2}. Under the assumption that ϵ_{it} is serially uncorrelated, $\Delta\epsilon_{it}$ is an MA(1) process, and thus s_{it-2} is a valid instrument for s_{it-1}. The Arellano-Bond test for serial correlation confirms that s_{it-2} is indeed a valid instrument.[3]

I refrain from including further lags of s_{it-1}, because doing so entails considerable loss of data, given that the average individual is observed for fewer than 6 weeks. In the online appendix, I estimate the model using the GMM estimator developed by Arellano and Bond (1991) to exploit additional available moment conditions while mitigating the data loss associated with differencing. Point estimates are consistent with those from the more parsimonious instrumenting strategy discussed above.[4]

3.2.2 Results

Table 1 reports results from two models: (i) a baseline specification that does *not* include as a regressor cumulative past search (Baseline), and (ii) an identical specification augmented with cumulative past search time as described above (Augmented).[5] For each specification, I report results for both the time diary and weekly recall measures of search time.[6]

Two principal results emerge from Table 1. First, the coefficient on cumulative past search is statistically significant and negative for both measures of search effort. Moreover, the augmented model provides a much better fit for the data as measured by the adjusted R^2. Second, when cumulative past search is included as a regressor, unemployment duration ceases to enter the model with a significant negative coefficient. Put

[2] Day-of-week indicators are included in the time diary regressions.

[3] See the online appendix for details.

[4] GMM estimation is implemented via first-differences and forward-orthogonal deviations.

[5] To emphasize the parameters of interest, κ and π, I exclude the estimated coefficient on the unemployment rate from Table 1. Results for the full model are available upon request.

[6] The baseline specification is intended to capture a model similar to that of Krueger and Mueller (2011), relating search effort to unemployment duration. However, their estimation strategy—using a within transformation to purge individual fixed effects—is not feasible in the Augmented model due to the inclusion of cumulative past search, which is partially unobserved. The online appendix reports results from fixed effects estimation of the Baseline model following Krueger and Mueller (2011). Results are similar to those presented in the fourth column of their Table 2.

Table 1: Job search over the unemployment spell

	Time Diary		Weekly Recall	
	Baseline	Augmented	Baseline	Augmented
Duration (κ)	-5.460***	4.528***	-4.585***	6.085*
	(1.504)	(1.604)	(1.507)	(3.296)
Past Search (π)		-0.150***		-0.0906***
		(0.0282)		(0.0248)
Observations	5497	5497	5445	5445
Adjusted R^2	0.066	0.209	0.006	0.086

Robust standard errors in parentheses.

Source: Survey of Unemployed Workers in New Jersey

Notes: Regressions of search effort on (i) unemployment duration (Baseline) and (ii) unemployment duration and cumulative past search (Augmented). All regressions use survey weights. Standard errors are robust and clustered at the individual level. The sample consists of respondents ages 25-54 who have not yet accepted a job offer, are not currently employed, and who do not expect to be recalled by or return to their former employer.

differently, the observed decline in effort over the unemployment spell documented by Krueger and Mueller (2011) can be attributed to variation in past search.

3.2.3 Robustness

I consider several modifications to the model and estimation strategy described above to ensure that the results presented in Table 1 are a robust feature of the data.[7] First, I estimate the model on the exact sample used in Krueger and Mueller (2011).[8] The results are not sensitive to this change. Second, I allow for the possibility that search depends nonlinearly on duration via a cubic polynomial. The results are not sensitive to this change. Third, to address the possibility that search declines over the spell because respondents learn to circumvent the litany of questions on job search by reporting that they have not searched, I restrict attention to individuals reporting strictly positive search effort. The coefficient on cumulative past search for the time diary data is attenuated slightly but remains significant and negative; otherwise, the results do not change.[9] Finally, to address the possibility that the decline in online vacancies between October 2009 and January 2011 in New Jersey is driving the results, I include the Conference Board's Help Wanted OnLine data from New Jersey for the period in question in the analysis.[10] The results are not sensitive to this change.

3.3 Discussion

The results in Table 1 could plausibly be due to job seekers learning about the arrival rate of offers or stock-flow matching. I exploit data on job offers from the SUWNJ to test the learning hypothesis. I then consider a sample of individuals who have been unemployed for a relatively long period to test the stock-flow hypothesis.

[7]Results for all robustness checks are reported in the online appendix.

[8]The sample used throughout the paper is the same as that used in Krueger and Mueller (2011), with two exceptions: I focus on individuals aged 25 to 54 and I restrict attention to individuals who do not expect to return to or be recalled from their last job.

[9]See Davis (2011) in Krueger and Mueller (2011). An alternative approach would be to take differences with respect to interview time (instead of calendar time) and include as a regressor the number of past interviews, which would enable separate identification of duration and interview number. In the present context, this approach suffers from several important limitations, which I discuss at length in the online appendix. When the model is estimated this way, the measured effect of past search remains highly significant and negative, although the interview effects appear to explain most of the decline in effort over the spell.

[10]See Sahin (2011) in Krueger and Mueller (2011).

3.3.1 Learning

The results in Table 1 could be explained by job seekers learning about the process governing the arrival of offers through their idiosyncratic search experiences. If this is the case, then the arrival of a job offer should affect subsequent search decisions. A number of subtle issues potentially impede studying the effect of job offers on search in the framework described above. I exploit the differential timing of job offers and search in the time-diary data to address these issues and examine learning as a potential explanation for the results in Table 1.

Consider augmenting (1) with a variable representing the total number of offers that an individual has received since job loss, analogously to the total search time since job loss:

$$s_{it} = \iota + \kappa d_{it} + \pi \sum_{\tau=1}^{t-1} s_{i\tau} + \phi \sum_{\tau=1}^{t} o_{i\tau} + \gamma u_t + \eta_i + \epsilon_{it} \tag{3}$$

where $o_{i\tau}$ represents the number of offers received by individual i in week τ of unemployment. I focus on the time diary data and therefore index the sum from $\tau = 1$ to t instead of $t-1$.[11] First differencing as before, we obtain

$$\Delta s_{it} = \kappa \Delta d_{it} + \pi s_{it-1} + \phi o_{it} + \gamma \Delta u_t + \Delta \epsilon_{it}. \tag{4}$$

The first column of Table 2 ("Offers") reports results from naïve estimation of (4). The effect of offers is positive and significant at the 5% level: An additional offer is associated with a 36-minute increase in search effort per day. As mentioned above, however, these point estimates are likely contaminated by two biases neglected in the naïve model. I consider these in turn below.

The first possibility is that, if search effort at time t or $t-1$ increases the likelihood of receiving a job offer at time t, then $Cov(o_{it}, \Delta\epsilon_{it}) \neq 0$. The former case is unlikely, as it is unlikely that search effort as measured by the time diary data will affect the likelihood of an offer in the same period, as discussed above. The latter case is more plausible, and will tend to attenuate the estimate of ϕ. To address this concern, I first regress an indicator for whether or not an offer was received in period t (o_{it}) on contemporaneous and past search effort, unemployment duration, individual fixed effects and other controls. I then use the residuals from this regression—the component of the offer purged of the effects of past search—to instrument for o_{it} in (4). The second column of Table 2 ("Offers (residual)") reports the results. Consistent with attenuation resulting from the effect of search on offers, the coefficient on offers increases substantially and is now significant at the 1% level: An additional offer is now associated with a 51-minute increase in search effort per day.[12]

The second possibility is that offer quality may be driving these results: That is, perhaps individuals search more only after receiving particularly high offers that convey positive information about the offer quality distribution, and thus their earnings prospects. If data on individuals' wage expectations were available, one could restrict attention to offers below the expected wage, and thus rule out this possibility. Because such data are not available in the SUWNJ, I proceed by repeating the analysis described in the previous paragraph, restricting attention to individuals whose best offer is below their previously declared *reservation wage*.[13] While imperfect, this approach will successfully rule out offers inducing optimism about earnings prospects and stimulating search so long as search costs are relatively high: As discussed in Burdett

[11] In a given week, job offers are likely to arrive *before* search is recorded in the time diary, whereas it is impossible to disentangle the timing of offers from the timing of search using the weekly recall data. I therefore restrict attention to time diary data. Results using the weekly recall measure are not generally significant.

[12] Because the sample is restricted to individuals who have never accepted an offer, there is no concern that individuals are receiving part-time offers and continuing to search.

[13] Because the reservation wage is potentially affected by the arrival of a job offer, I use its lagged value.

and Vishwanath (1988), in this case the option value of unemployment will be relatively low, implying that no offers above the expected wage will be rejected, so that all individuals who reject offers will also become more pessimistic about their earnings prospects. If such offers affect search, it is unlikely that it is resulting from job seekers learning that better jobs are available; rather, it is consistent with offers affecting search by affecting the perceived availability of jobs in general. The third column of Table 2 ("Offers ($\leq \underline{w}_{t-1}$)") reports results from this final specification. The estimated effect of offers is similar in magnitude to the effect in the previous specification, and remains significant at the 5% level. If offer quality was driving the results, we would expect a significantly smaller coefficient.[14]

Table 2: The effect of job offers

	Offers	Offers (residual)	Offers ($\leq \underline{w}_{t-1}$)
Duration (κ)	3.096	2.748	3.122
	(2.325)	(2.371)	(2.357)
Past Search (π)	-0.0956**	-0.0961**	-0.0947**
	(0.0405)	(0.0403)	(0.0405)
Offers (ϕ)	36.23**	51.03***	44.56*
	(15.11)	(19.36)	(24.41)
Observations	3380	3380	3380
Adjusted R^2	0.153	0.152	0.148

Robust standard errors in parentheses.

Source: Survey of Unemployed Workers in New Jersey

Notes: Regressions of search effort on unemployment duration, cumulative past search, and a job offer indicator. All regressions use survey weights. Standard errors are robust and clustered at the individual level. The sample consists of respondents ages 25-54 who have not yet accepted a job offer, are not currently employed, and who do not expect to be recalled by or return to their former employer. I furthermore restrict the sample to individuals who have been in the survey for at least one month in order to consistently estimate the first stage regression in columns 2 and 3.

Thus, regardless of the specification, job offers appear to have an economically and statistically significant effect on subsequent search effort. These results are robust to inclusion of higher-order polynomials in duration, alternative approaches to controlling for macroeconomic effects, and using various demographic subsamples. Furthermore, the results hold when the total number of offers an individual received in a period is used instead of an indicator for whether or not at least one offer was received.[15] These results suggest an important role for learning about the process governing the arrival of offers in the job search process.

3.3.2 Stock-flow matching

Another possible explanation for the results in Table 1 is the presence of stock-flow matching (Coles and Smith, 1998; Ebrahimy and Shimer, 2006; Coles and Petrongolo, 2008).[16] Specifically, suppose that upon job

[14]Note that this is a particularly conservative approach to purging the effect of learning about the offer quality distribution: By excluding all offers for individuals whose best offer exceeds their reservation wage, I am likely excluding individuals receiving multiple offers some of which are below the reservation wage—individuals whose behavior may in fact also reflect learning about the availability of jobs, not just the offer distribution. Unfortunately, however, this is necessary because data are only available on the wage of the best offer.

[15]Two interesting exceptions to the otherwise robust significance of offers are worth noting: First, inclusion of individuals who expect to be recalled by their former employer reduces the significance of cumulative past search in these regressions (although the signs remain negative and duration still becomes insignificant). Second, inclusion of older individuals—specifically, those 55 and above—reduces the significance of job offers (although the sign remains positive). These results are not surprising in the context of a model of learning: The salience of new information for search decisions is likely very different for individuals expecting to be recalled. Likewise, the information content of offers—and thus the effect of such information on search—is likely to be much smaller for older individuals who have been in the labor market for most of their working lives.

[16]Of course, stock flow matching is unlikely to be able to explain the results in Table 2.

loss, individuals observe a stock of relevant vacancies, and search time is devoted to applying to those jobs. Once that stock has been exhausted, subsequent search is limited by the flow of newly posted vacancies. In this environment, the time devoted to search at the beginning of the unemployment spell corresponds to the rate at which the initial stock is drawn down. Thus, individuals who devote more time to search early in the unemployment spell may more rapidly reduce their search.

As a simple test of whether stock-flow matching is driving the results in Table 1, I restrict the sample to individuals who have been unemployed for over one month. Because they have been unemployed for a relatively long period of time, in a stock-flow model these individuals are more likely to have already drawn down the initial stock of pre-existing vacancies, and thus their current search effort should be dictated by the flow of new vacancies, not their past effort. Results are reported in the online appendix. The measured effect of cumulative past search is nearly identical to that found in Table 1 for both measures of search time, suggesting that stock-flow matching is unlikely to be driving the results.

4 A Model of Sequential Search with Learning

In the canonical theory of sequential search, job offers either arrive every period or arrive stochastically at a known average rate. In this section, I develop a model of search in which job seekers have incomplete information about the availability of work (precisely, the rate at which offers arrive) and learn from their experiences while searching. Specifically, I assume that: (i) Job seekers choose how much time to spend looking for work at the beginning of each period, and (ii) job seekers do not observe the rate at which job offers arrive per unit of time devoted to search. Because the arrival rate is unobserved, job seekers are endowed with a distribution of beliefs that evolves endogenously in response to the arrival of new information. In the model, search time and reservation-wage dynamics over the unemployment spell are driven by the evolution of beliefs, which in turn respond to the idiosyncratic outcomes of search.

In what follows, I present a stylized model of search to illustrate the mechanisms at work. See the online appendix for details of the complete model used for estimation in Section 5.

4.1 Environment

I consider a simple partial equilibrium environment in which search occurs sequentially in the spirit of McCall (1970).

4.1.1 Timing

Unemployment duration is discrete and measured in weeks. Unemployed job seekers maximize the present discounted value of income net of search costs: $E_0 \sum_{t=0}^{\infty} \delta^t (y_t - \eta s_t)$. Search costs may be thought of as monetary costs or the imputed value of forgone leisure.

At the beginning of each week t, job seekers choose to devote fraction s_t of their week to searching for work. While searching, job offers arrive according to a Poisson process with true average rate parameter λ^T.[17] Letting $\tilde{\tau}_t$ denote the stochastic arrival time of the first offer, the *true* probability of a job offer arriving before search ends is given by

$$Pr(\tilde{\tau}_t \leq s_t) \equiv F(s_t; \lambda^T) = 1 - e^{-\lambda^T s_t}. \tag{5}$$

[17] Unemployment duration is discrete, but offers arrive according to a continuous Poisson process within periods. When an offer arrives, I assume that job seekers must stop searching for the remainder of the period to update beliefs and evaluate the offer, so agents never receive more than one offer per week. This assumption could be relaxed by assuming that the number of offers arriving each period follows a Poisson distribution.

If a job offer arrives before search ends ($\tilde{\tau}_t \leq s_t$), the job seeker updates her estimate of λ^T and decides whether to accept the offer, as in a standard McCall (1970)-style search framework. Offers are drawn from a known distribution $\Phi(\omega)$ with density $\phi(\omega)$. If the offer is accepted, the job seeker receives the wage offer for the rest of her life. If the offer is rejected, the job seeker receives flow value of unemployment b and continues searching in the next period. If no offer arrives before search ends ($\tilde{\tau}_t > s_t$), the job seeker receives flow value of unemployment b and updates her estimate of λ^T to reflect the fact that searching for fraction s_t of the week yielded no offers. Figure 1 depicts the timing of the model.

Figure 1: Timing of events

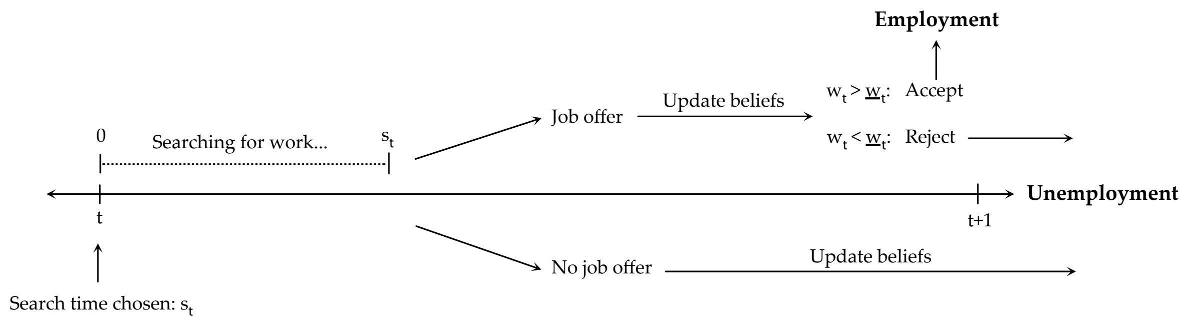

Notes: A typical period of unemployment for a searching worker. Workers choose search effort s_t at the beginning of the period and update beliefs after search ends, but before an offer (if received) is accepted or rejected.

4.1.2 Beliefs

I assume that job seekers do not know the true job offer arrival rate λ^T. Instead, they form beliefs about the value of λ^T, which take the form of a Gamma distribution, parameterized by α_t and β_t. The assumptions that observed arrival times follow a (right-censored) exponential distribution and that beliefs follow a Gamma distribution together imply that beliefs are time-invariant up to parameters α_t and β_t, which affords the model considerable tractability.[18]

The density of beliefs in week t is thus given by

$$\gamma(\lambda; \alpha_t, \beta_t) = \frac{\beta_t^{\alpha_t} \lambda^{\alpha_t - 1} e^{-\beta_t \lambda}}{\Gamma(\alpha_t)}. \tag{6}$$

The mean and variance of the distribution of beliefs in week t are

$$E_t(\tilde{\lambda}) = \frac{\alpha_t}{\beta_t} \qquad Var_t(\tilde{\lambda}) = \frac{\alpha_t}{\beta_t^2}. \tag{7}$$

The parameters of the belief distribution α_t and β_t evolve endogenously over the unemployment spell according to the following laws of motion:

$$\alpha_{t+1} = \begin{cases} \alpha_t + 1 & \text{if} \quad \tau_t \leq s_t \quad \text{(offer)} \\ \alpha_t & \text{if} \quad \tau_t > s_t \quad \text{(no offer)} \end{cases} \tag{8}$$

$$\beta_{t+1} = \begin{cases} \beta_t + \tau_t & \text{if} \quad \tau_t \leq s_t \quad \text{(offer)} \\ \beta_t + s_t & \text{if} \quad \tau_t > s_t \quad \text{(no offer).} \end{cases} \tag{9}$$

[18] See the online appendix for a simple proof of this claim.

Note that α_t counts the number of job offers received since job loss and β_t measures accumulated search time since job loss. The endogeneity of beliefs arises from two sources: (i) the explicit presence of s_t in (8) and (9); and (ii) the fact that whether or not an offer is received implicitly depends on s_t.

Figure 2: Beliefs

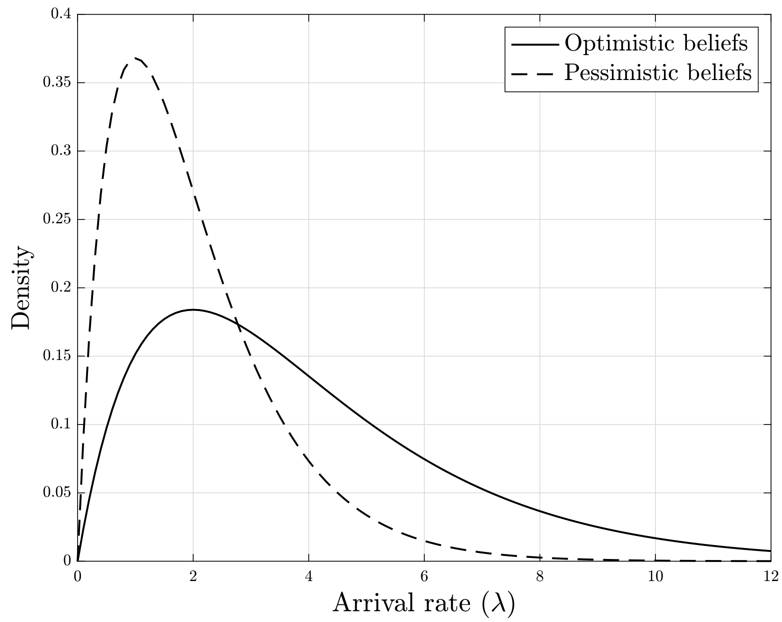

Notes: Two hypothetical belief density functions. The optimistic beliefs depicted correspond to a Gamma density function with shape parameter $\alpha = 2$ and rate parameter $\beta = \frac{1}{2}$; the pessimistic beliefs correspond to $\alpha = 2$ and $\beta = 1$.

Figure 2 depicts two belief distributions associated with different values of α_t and β_t. As more job offers arrive, job seekers become optimistic, and the belief distribution shifts outward. Conversely, as more time is spent searching without receiving an offer, job seekers become pessimistic, and the belief distribution shifts inward.[19] Notice that this specification of beliefs is fairly flexible: Because the belief distribution is fully characterized by two parameters, α and β, which map into the mean and variance of beliefs via (7), any combination of optimism/pessimism (as measured by $\mathbb{E}[\tilde{\lambda}]$) and precision (as measured by $Var(\tilde{\lambda})$) can be accommodated via choice of α and β. This will be important when I turn to structural estimation of individuals' priors at the time of job loss in Section 5.

In keeping with much of the macroeconomic literature on learning, I assume that job seekers optimize within an anticipated-utility framework.[20] This assumption serves to simplify the exposition of the model, and provides a significant reduction in the computational burden associated with estimating the model in Section 5.

[19] Conditional on not receiving an offer, arrival time τ_t is not observed. See the online appendix for discussion of this point as it pertains to conjugacy of the Gamma distribution.

[20] See Kreps (1998).

4.2 Recursive formulation

The value of entering week t unemployed with beliefs characterized by α_t and β_t may be written recursively as

$$V_t^U(\alpha_t,\beta_t) = \max_{s_t}\left\{b - \eta s_t + \delta \int_0^\infty \Big[F(s_t;\lambda)E_t^\omega\left[V_t^O(\omega,\alpha_t,\beta_t)\right] \right. \\ \left. + (1-F(s_t;\lambda))V_{t+1}^U(\alpha_t,\beta_t)]\Big] d\Gamma(\lambda;\alpha_t,\beta_t)\right\} \tag{10}$$

where $V_t^O(\omega,\cdot)$ denotes the value of having offer ω in hand and is given by

$$V_t^O(\omega,\alpha_t,\beta_t) = \max\left\{\frac{\omega}{1-\delta}, V_{t+1}^U(\alpha_t,\beta_t)\right\}. \tag{11}$$

The value of entering week t unemployed is a probability-weighted average of the expected value of receiving a job offer and the value of receiving no offer and remaining unemployed into period $t+1$, less the cost of search. Because λ^T is unobserved, job seekers integrate over possible values of the underlying arrival rate according to the current state of their beliefs, as characterized by α_t and β_t.

4.3 Solution

I solve the model in two stages. First, I characterize behavior at the end of the period for job seekers who have received offers; optimal behavior takes the form of a familiar reservation-wage policy. Second, I determine optimal search time at the start of the period conditional on the reservation-wage policy determined in the first stage.

4.3.1 Reservation wage

Consider first the problem of an unemployed job seeker with a known offer ω in hand. Because the first argument in the max operator in equation (11) is strictly increasing in ω, while the second is constant, the optimal choice between accepting and rejecting the offer may be characterized by a standard reservation-wage policy:

$$V_t^O(\omega,\alpha_t,\beta_t) = \begin{cases} \dfrac{\omega}{1-\delta} & \text{if } \ \omega > \underline{w}_t \\ V_{t+1}^U(\alpha_t,\beta_t) & \text{if } \ \omega \leq \underline{w}_t \end{cases} \tag{12}$$

where the reservation wage is defined by

$$\underline{w}_t = (1-\delta)V_{t+1}^U(\alpha_t,\beta_t). \tag{13}$$

Job seekers choose a threshold wage rate $\underline{w}_t$ such that the present discounted value of accepting an offer $\underline{w}_t$ is equated with the flow value of unemployment b plus the value of remaining unemployed.

4.3.2 Search effort

From 10, the first-order condition for the choice of s_t is given by

$$\eta = \int_0^\infty f(s_t;\lambda)\delta\Big[E_t^\omega\left[V_t^O(\omega,\alpha_t,\beta_t)\right] - V_{t+1}^U(\alpha_t,\beta_t)\Big]\gamma(\lambda;\alpha_t,\beta_t)d\lambda. \tag{14}$$

The expression in brackets is the expected net benefit from receiving an (unknown) offer. Making use of (12) and (13), this term may be written as

$$E_t^{\omega}\left[V_t^O(\omega,\alpha_t,\beta_t)\right]-V_{t+1}^U(\alpha_t,\beta_t)=\frac{1}{1-\delta}\int_{\underline{w}_t}^{\infty}(\omega-\underline{w}_t)\phi(\omega)d\omega. \tag{15}$$

so that the first-order condition reduces to

$$\eta=\int_0^{\infty}f(s_t;\lambda)\left[\frac{\delta}{1-\delta}\int_{\underline{w}_t}^{B}(\omega-\underline{w}_t)\phi(\omega)d\omega\right]\gamma(\lambda;\alpha_t,\beta_t)d\lambda. \tag{16}$$

Job seekers equate the marginal cost of search η with the expected marginal benefit. The expected marginal benefit is the product of the marginal increase in the probability of finding an offer multiplied by the expected net value of an offer, integrated over the unobserved arrival rate λ.

4.3.3 Model dynamics

Using (10) to eliminate the value function from (13) and explicitly integrating over beliefs yields the two key equations that jointly characterize time devoted to job search and the reservation wage:

$$s_t=\beta_t\left[\left(\frac{\delta}{\eta(1-\delta)}\int_{\underline{w}_t}^{B}(\omega-\underline{w}_t)\phi(\omega)d\omega\left(\frac{\alpha_t}{\beta_t}\right)\right)^{\frac{1}{\alpha_t+1}}-1\right] \tag{17}$$

$$\underline{w}_t-b+\eta s_t=\left[1-\left(\frac{\beta_t}{\beta_t+s_t}\right)^{\alpha_t}\right]\left(\frac{\delta}{1-\delta}\int_{\underline{w}_t}^{B}(\omega-\underline{w}_t)\phi(\omega)d\omega\right). \tag{18}$$

Model dynamics are governed by the optimality conditions in (17) and (18), together with the laws of motion for beliefs in (8) and (9).

4.4 Reservation wage

A large literature in empirical labor economics seeks to understand how reservation wages vary with unemployment duration. A robust finding in this literature suggests that reservation wages tend to decline over the course of unemployment (cf. Devine and Kiefer, 1991; Barnes, 1975; Feldstein and Poterba, 1984). Proposition 1 establishes that the model predicts monotonically declining reservation wages as beliefs deteriorate (i.e., as β_t rises).

Proposition 1. *The reservation wage is monotonic in beliefs (β). Deteriorating beliefs reduce the reservation wage.*

Proof. See the online appendix. □

The intuition for this result is straightforward: Reductions in the perceived likelihood of finding work reduce the option value of remaining unemployed—thus making job seekers more willing to accept offers and reducing the reservation wage. Because β_t measures cumulative search effort since job loss, it follows immediately that reservation wages must decline over the unemployment spell in the absence of job offers.

4.5 Search effort

In the model, job seekers learn about the unobserved arrival rate of job offers through their experiences looking for work. The learning process induces changes in the distribution of beliefs through α_t and β_t, which in turn govern search decisions.

4.5.1 The effect of beliefs on search

How exactly does learning affect search decisions? Consider the effect of a deterioration in beliefs—that is, a reduction in β_t—on the optimal choice of search effort. When search ends and no offers have arrived, job seekers update their beliefs to reflect the failure to find work. This updating has two effects on subsequent search effort, as may be seen by differentiating the first-order condition for search effort with respect to β_t:

$$\frac{\partial s_t}{\partial \beta_t} = \underbrace{\frac{\alpha_t s_t - \beta_t}{\beta_t(\alpha_t + 1)}}_{\text{Search productivity effect}} - \underbrace{\left(\frac{\beta_t + s_t}{\alpha_t + 1}\right) \frac{(1 - \Phi(\underline{w}_t))}{\int_{\underline{w}_t}^{B} (\omega - \underline{w}_t)\phi(\omega) d\omega} \frac{\partial \underline{w}_t}{\partial \beta_t}}_{\text{Option value effect}}. \tag{19}$$

On the one hand, a lower perceived probability of finding work means that remaining unemployed is a less attractive option. Just as lower unemployment benefits reduce the option value of remaining unemployed in the standard McCall (1970) model of sequential search, a perceived reduction in the probability of finding work likewise reduces the option value of remaining unemployed in the model described above. I refer to this first effect as the *option value effect*. On the other hand, a lower perceived probability of finding work reduces the opportunity cost of leisure. So long as a hypothetical job offer would increase the expected arrival rate of offers following search, this effect induces a substitution away from time devoted to search.[21] I refer to this second effect as the *search productivity effect*—that is, the effect of beliefs on search holding fixed the option value of unemployment. The combination of these two effects implies that search effort is nonmonotonic in beliefs. Proposition 2 formalizes this.

Proposition 2. *Search effort is nonmonotonic in beliefs (β). Furthermore:*

1. *Deteriorating beliefs exert a positive influence on search through the option value effect;*
2. *Deteriorating beliefs exert a negative influence on search through the search productivity effect if a hypothetical job offer would raise the perceived arrival rate of offers;*
3. *Deteriorating beliefs depress search effort if the search productivity effect is negative and dominates the option value effect.*

Proof. See the online appendix. □

The nonmonotonicity associated with endogenously evolving beliefs implies that search effort can exhibit a hump-shaped profile over the unemployment spell, even in the absence of job offers. Figure 3 provides intuition for why these two effects generate hump-shaped nonmonotonocity.

When job seekers are optimistic (light shading), the perceived increase in the job-finding probability from an extra minute of search, $E^\lambda[\frac{dF}{ds}|\beta]$, is large (the rectangles are tall), while the gain from successful search, $V^E - V^U$, is small (the rectangles are narrow). This means that when a job seeker fails to find work and becomes less optimistic (darker shading), the *reduction* in the marginal benefit of search due to the fall in $E^\lambda[\frac{dF}{ds}|\beta]$ is small because it is scaled by a low value of $V^E - V^U$, whereas the *increase* in the marginal benefit of search due to the rise in $V^E - V^U$ is large because it is scaled by a high value of $E^\lambda[\frac{dF}{ds}|\beta]$. The net effect is thus an increase in the perceived marginal benefit of search (the medium-shade rectangle has greater area than the light-shade rectangle), and thus an increase in search effort. On the other hand, this argument is reversed when job seekers are pessimistic (heavy shading). In this case, the perceived increase in

[21] Interestingly, when a hypothetical offer would instead *reduce* the expected arrival rate of offers $E[\lambda^T|\alpha, \beta]$, an increase in β_t will induce an increase in the perceived marginal effectiveness of search. This happens because, as β_t rises, the perceived marginal effectiveness of search flattens as a function of effort, which implies that effort rises for high values of s_t. As a practical matter, however, this situation only occurs in implausible regions of the parameter space. Estimation of the model in Section 5 confirms this empirically.

Figure 3: Nonmonotonic search dynamics

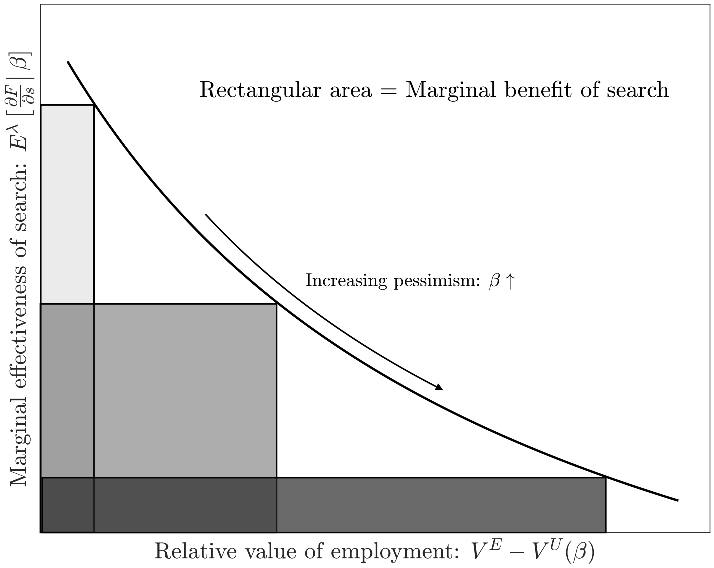

Notes: As searchers become increasingly pessimistic, the perceived marginal benefit of search—corresponding to the area of the pictured rectangles—first increases and then decreases, implying that optimal search effort is nonmonotonic in beliefs.

the job-finding probability from an extra minute of search, $E^{\lambda}[\frac{dF}{ds}|\beta]$, is small (the rectangles are short), while the gain from successful search, $V^E - V^U$, is large (the rectangles are wide). This means that when a job seeker fails to find work and becomes even more pessimistic (darker shading), the *reduction* in the marginal benefit of search due to the fall in $E^{\lambda}[\frac{dF}{ds}|\beta]$ is large because it is scaled by a large value of $V^E - V^U$, whereas the *increase* in the marginal benefit of search due to the rise in $V^E - V^U$ is small because it is scaled by a small value of $E^{\lambda}[\frac{dF}{ds}|\beta]$. The net effect is thus a reduction in the perceived marginal benefit of search (the dark-shade rectangle has less area than the medium-shade rectangle). Note that this argument is just a mechanical application of the "product rule" to differentiating the marginal benefit of job search with respect to β.

Empirically, nonmonotonic search dynamics are a feature of pre-Great Recession data: Both Shimer (2004) and Mukoyama et al. (2018) independently document that search effort appears to exhibit a hump-shaped profile over the first two years of unemployment using CPS data. A credible theory of search should therefore be able to account for this feature of the data. Through the competing search productivity and option value effects described above, the model developed in this paper can do precisely that without relying on other mechanisms. Indeed, when beliefs are sufficiently optimistic at the time of job loss—as one might expect in pre-Great Recession data—the model implies a hump-shaped profile of search effort over the spell. Yet when beliefs are pessimistic, search will decline monotonically.

4.5.2 Relationship to reduced-form results

The model developed in this section was motivated by the observation in Section 3 that total search effort since job loss exerts a negative influence on subsequent search effort. Given that β_t is related to past search effort, equation (19) and Proposition 2 suggest a link between the motivating empirics in Section 3 and the search productivity and option value effects described above. Proposition 3 formalizes this connection.

Proposition 3. *A first-order Taylor expansion of the structural model implies the reduced-form regression estimated in Section 3. Furthermore, the reduced-form coefficient on cumulative past search is negative if the search productivity effect dominates the option value effect.*

Proof. See the online appendix. □

Two things should matter for (Bayesian) job seekers' perceptions of the return to job search effort: (i) how much effort they have previously exerted—that is, the effective amount of time they have spent observing the offer-generating process of job search—and (ii) the observed outcomes of that process. The former is simply past job search, while the latter is past job offers. Proposition (3) is simply a formalization of the preceding intuition afforded by the fact that, via (8) and (9), α_t and β_t admit natural interpretations in terms of observables in the SUWNJ.

The explicit link between the reduced-form analysis in Section 3 and the theoretical model of job search in this section is, so far, only qualitative in nature. What can be said about the model's ability to *quantitatively* account for the reduced-form results in Section 3, and more generally to account for observed search behavior in the data? To answer these questions, Section 5 undertakes a formal quantitative analysis of the structural model.

5 Structural Estimation

This section structurally estimates a version of the model developed in Section 4.[22] There are two principal goals of the exercise: (i) to evaluate the extent to which learning about the arrival rate of job offers is able to account for the reduced-form results in Section 3, as well as observed search behavior more generally; and (ii) to characterize the distribution of beliefs about the job-finding rate in the data. To accomplish these goals, I identify key structural parameters—including those governing beliefs at the time of job loss—using the estimated parameters from the reduced-form analysis in Section 3, in addition to sample averages of search effort, the offer arrival probability, and the offer acceptance probability from the SUWNJ data.

5.1 Model

I enrich the simple model described in Section 4 along three dimensions, all of which are standard in estimated models of job search.[23] Specifically, I allow for on-the-job search, a convex search cost function, and the possibility of offers arriving without search. For a complete description of the model and its solution, see the online appendix. On-the-job search helps resolve a well-known empirical tension in models of sequential search—the fact that plausible transition rates require a counterfactually low degree of wage dispersion—by allowing employed workers to move to higher-wage jobs.[24] Convexity in the search cost function precludes an implausibly large share of high-duration workers from stopping search altogether after failing to receive offers, which would have the effect of counterfactually depressing average search effort in the model. The possibility of offers arriving without search serves to decouple the arrival rate of job offers from search effort to a reasonable extent, and also allows the model to generate plausible sensitivity of effort to cumulative past search (that is, a relatively large negative value of $\hat{\pi}$) without requiring excessive sensitivity of effort to job offers ($\hat{\phi}$). The model is otherwise identical to the model described in Section 4.

[22]See the online appendix for details of the full model.

[23]See, for example, Christensen et al. (2005) and more recently Faberman et al. (2017).

[24]On-the-job search requires that I take a stance on the *perceived* arrival rate of offers while searching during employment relative to unemployment. In the absence of clear empirical or theoretical guidance on this choice, I assume that job seekers believe that time spent searching is equally productive during unemployment and employment. To facilitate solving the model (which has two state variables that evolve stochastically for each simulated individual), I furthermore assume that employment and unemployment are otherwise symmetric.

5.2 Empirical strategy

I use a simulated minimum distance procedure to estimate six structural parameters from the model described in the online appendix:

$$\Theta = [\underbrace{\alpha_0, \beta_0, Bias}_{\text{Beliefs/Job-finding}}, \underbrace{\psi, \Psi, b}_{\text{Physical}}]' \tag{20}$$

where *Bias* is a measure of the distortion in individuals' beliefs about the likelihood of finding work at the time of job loss (defined formally below). The remaining parameters are directly calibrated to the data and estimates from the literature.

Estimation proceeds in three steps. First, I specify the auxiliary model; this is the lens through which I compare the structural model with the data. Next, I estimate the parameters of the auxiliary model—the auxiliary parameters—using the SUWNJ data. Finally, I choose the *structural* parameters Θ so as to minimize the distance between the auxiliary parameters generated by the SUWNJ data and the auxiliary parameters generated by simulating the structural model.

5.2.1 Identification and the auxiliary model

In order to identify the structural parameters Θ, I specify the auxiliary model as three vectors of moments:

$$\Omega = [\Omega_1, \Omega_2, \Omega_3] \tag{21}$$

where

$$\Omega_1 = \hat{\kappa}_{FE} \tag{22}$$

$$\Omega_2 = [\hat{\pi}, \hat{\phi}, \hat{\kappa}] \tag{23}$$

$$\Omega_3 = [\hat{s}, \hat{o}, \hat{a}] \tag{24}$$

Ω_1 contains a single parameter, $\hat{\kappa}_{FE}$, the coefficient on duration in a fixed effects specification in the spirit of Krueger and Mueller (2011) that includes interview effects (see the online appendix for results from this specification and further discussion). Ω_2 contains the three key reduced-form coefficients from Section 3: $\hat{\pi}$ is the measured effect of cumulative past search from the first column of Table 2; $\hat{\phi}$ is the measured effect of job offers from the first column of Table 2; and $\hat{\kappa}$ is the measured effect of duration from the first column of Table 2. The values used in estimation are identical to the values found in Section 3. Ω_3 contains sample averages of minutes per day spent on job search ($\hat{s}$), the probability of receiving an offer ($\hat{o}$), and the probability of accepting an offer conditional on having received an offer ($\hat{a}$).

Letting Ω^e denote the vector of moments obtained from the SUWNJ and $\Omega^m(\Theta)$ denote the moments obtained from simulating the structural model at parameters Θ, we have:

$$\text{SUWNJ:} \quad \Omega^e = [\hat{\kappa}_{FE}, \hat{\pi}, \hat{\phi}, \hat{\kappa}, \hat{s}, \hat{o}, \hat{a}] \tag{25}$$

$$\text{Model:} \quad \Omega^m(\Theta) = [\tilde{\kappa}_{FE}, \tilde{\pi}, \tilde{\phi}, \tilde{\kappa}, \tilde{s}, \tilde{o}, \tilde{a}]. \tag{26}$$

The reduced-form parameters in Ω_1^e and Ω_2^e are taken directly from the estimates in Section 3 and the online appendix, and thus the discussion of the treatment of the data elsewhere in the paper continues to apply. The sample averages in Ω_3^e are simply the predicted values from three linear regression models. The three left-hand side variables are minutes per day spent on job search (from time diary data), an indicator for whether or not an offer was received, and an indicator for whether or not an offer was accepted among individuals who have received offers. In each of the three regressions, the regressors were an indicator for whether or not an individual was eligible for unemployment insurance benefits, an indicator for November 8, 2009 (the date when UI extensions came into effect in New Jersey) and the New Jersey unemployment rate. Predicted values were obtained by evaluating the models at the means of the explanatory variables. The sample is identical to the sample used elsewhere in the paper.

5.2.2 Implementation

Prior to estimation, I fix the weekly discount factor δ to 0.999 and the weekly separation rate ρ to 0.004 following Lentz (2009). I assume that the offer distribution $\Phi(\omega)$ is Lognormal with mean normalized to $m = 1$. I then calibrate the variance ν to match the estimated standard deviation of log job values from Hall and Mueller (2018) of 0.38.[25] Finally, I calibrate the independent probability of an offer (that is, absent search) to $\xi = 0.0184$ to match the probability of receiving an offer among individuals reporting zero search in the SUWNJ.

I assume that newly unemployed job seekers draw unobserved arrival rates λ^T from a Gamma distribution with arbitrary mean (to be estimated) and variance equal to the variance of job seekers' initial beliefs over arrival rates.[26] I directly estimate the bias in individuals' beliefs at the time of job loss, defined as the percentage difference between the *perceived* mean arrival rate (α_0/β_0) and the true population mean arrival rate ($E[\lambda^T]$):[27]

$$Bias \equiv \frac{\alpha_0/\beta_0 - E[\lambda^T]}{E[\lambda^T]}. \tag{27}$$

Following the literature, I assume that the search cost is isoelastic and of the form

$$\eta(s) = \frac{\psi}{\Psi + 1} s^{\Psi+1}, \qquad \psi, \Psi > 0 \tag{28}$$

where ψ and Ψ are parameters to be estimated. Finally, to economize on estimated parameters, I exploit the assumption that, from the standpoint of a job seeker, the only difference between employment and unemployment is the (gross) flow value, which implies that the reservation wage during unemployment is simply equal to the flow value of unemployment: $\underline{w}_t = b$.[28] This immediately implies that, for given values of m and ν (see above), b is directly identified by $\tilde{a} \equiv Pr(\text{accept}|\text{offer}) = 0.74$ and the fact that the offer distribution is Lognormal. That is,

$$b = \Phi^{-1}(1 - \hat{a}) = 0.73. \tag{29}$$

Structural parameters Θ are chosen to minimize the distance between the empirical auxiliary parameters Ω^e and the model-generated auxiliary parameters $\Omega^m(\Theta)$. Formally, the estimator is

$$\hat{\Theta} = \underset{\Theta}{\text{argmin}} \left[\Omega^m(\Theta) - \Omega^e\right] W \left[\Omega^m(\Theta) - \Omega^e\right]', \tag{30}$$

where W is a diagonal weighting matrix containing the inverses of the empirical variances along the main diagonal.[29]

[25] This value is on the high end of estimates found in the literature. Accordingly, in the online appendix, I consider a more standard value of ν chosen to match the estimated standard deviation of log *wages* in Hall and Mueller (2018). This value is also similar to the estimate from Low et al. (2010). Results are not significantly affected by this change.

[26] Beliefs are thus restricted to be consistent with the underlying population distribution of arrival rates up to a mean bias term. I furthermore restrict the variance of the distribution of arrival rates to ensure plausible hazard rates in the first week of unemployment.

[27] Note that the true population mean arrival rate, $E[\lambda^T]$ can be backed out from estimates of α_0, β_0 and $Bias$.

[28] The principal reason for considering the symmetric case is that it dramatically reduces the computational burden associated with estimating the model resulting from endogenously (and stochastically) evolving beliefs. However, the implication that reservation wages do not change over the unemployment spell is broadly in line with the findings of Krueger and Mueller (2016), suggesting that this may not be an unreasonable baseline.

[29] Results are not affected by use of alternative weighting matrices. See the online appendix for results from an estimation

5.3 Results

I begin by considering parameter estimates and model fit.

5.3.1 Parameter estimates

Table 3 reports estimates of Θ. Bootstrapped standard errors are reported in parentheses. Several of the parameter estimates in Table 3 warrant discussion.

Table 3: Parameter estimates

Parameter	Concept	Estimate
Beliefs		
α_0	Initial beliefs (shape)	0.39 (0.09)
β_0	Initial beliefs (rate)	0.27 (0.22)
$Bias$	$\left(\frac{\alpha_0/\beta_0 - E[\lambda^T]}{E[\lambda^T]}\right) \cdot 100$	57% (23%)
Physical		
ψ	Search cost (level)	7.91 (1.12)
Ψ	Search cost (curvature)	0.26 (0.03)
b	Flow value of unempl.	0.73

Source: Survey of Unemployed Workers in New Jersey.

Notes: Point estimates of structural parameters.

First, I estimate that beliefs at the time of job loss are significantly positively biased: The expected arrival rate of offers during search exceeds the average arrival rate by 57%. In more practical terms, given the optimal choice of search effort conditional on these beliefs and the fact that there is a small probability of receiving an offer even if search fails, these results imply a bias of 61% in the perceived probability of receiving a job offer. The result that job seekers overestimate their job-finding prospects is consistent with recent findings from research based on other data sources, as well as direct evidence based on elicited subjective job-finding probabilities from the SUWNJ itself (see, for example, Spinnewijn (2015) and Mueller et al. (2020)). This consistency is particularly noteworthy given that I do not make use of any data on subjective job-finding probabilities in the course of estimation, and suggests that excessive optimism can potentially be identified based on search behavior alone, even in the absence of directly reported perceptions.

Second, as described above, data on the average probability of accepting conditional on receiving an offer directly identifies the *gross* flow value of unemployment, $b = 0.73$. However, computation of the *net* flow value of unemployment requires taking account of search costs. Netting these search costs out, over the entire sample the model implies a net flow value of unemployment of $b - \eta(\bar{s}_{it}) = 0.59$. This value lies squarely within the range found in the literature. Because, on average, search effort declines with unemployment duration in the model, the net flow value of unemployment rises with unemployment duration. However, this is offset by changes in the perceived option value of unemployment, which falls strongly over the spell. Relatedly, I estimate search parameters that suggest a relatively low degree of convexity in the search cost function relative to, e.g., Christensen et al. (2005), who find that a quadratic search cost function provides a reasonably good fit for the data.

Finally, I estimate relatively small values for α_0 and β_0, which together suggest that (i) the distribution of beliefs over the true arrival rate of offers has high variance, and (ii) that there is a large amount of

using an identity weighting matrix. Results from weighting using the inverse of the full empirical covariance matrix are available on request and likewise have a negligible effect on results.

heterogeneity in job-finding rates in the population. Interestingly, the estimated values of α_0 and β_0 also imply that optimal search effort initially *increases* as beliefs deteriorate in the absence of job offers. This is precisely the option value effect of deteriorating beliefs described in Section 4. However, after several weeks of repeatedly failing to find work, search effort begins to fall rapidly, reflecting a dominant search productivity effect from Section 4. This latter effect occurs quickly enough among individuals who fail to receive offers that it vastly outweighs the effect of temporarily increasing effort in the reduced-form regression coefficients used as auxiliary parameters. To see how these two effects vary over time and are affected by job offers, Figure 4 depicts the relationship between search effort and cumulative past effort at the estimated parameter values for (i) an individual who receives no job offers, (ii) an individual who enters unemployment with one job offer, and (iii) an individual who enters unemployment with two job offers.

Figure 4: Search effort and job offers

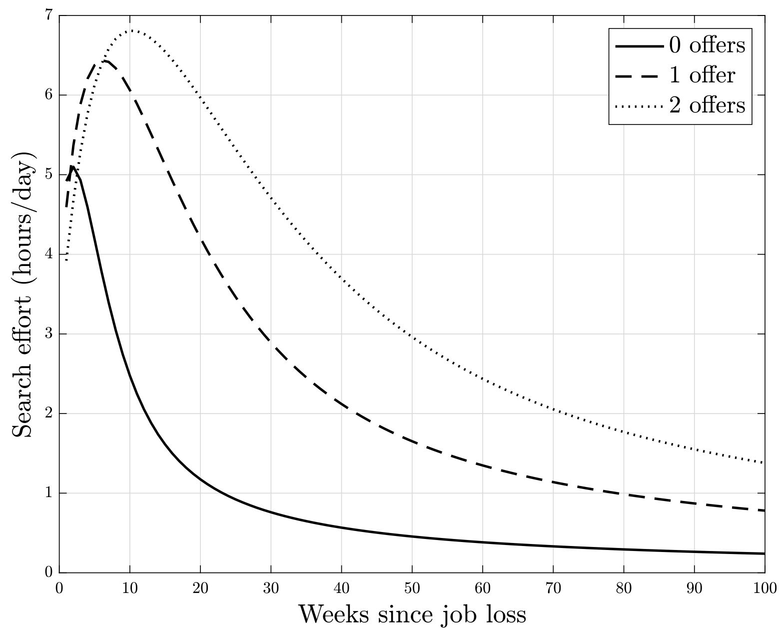

Notes: Search dynamics differ significantly depending on the number of offers workers have received, and thus their beliefs about the rate at which offers will arrive in the future. Good luck early in the unemployment spell stimulates persistently elevated effort.

5.3.2 Model fit

Table 4 reports estimates of auxiliary parameters from the SUWNJ data and from the model simulated at the parameter values in Table 3. For reference, the table also indicates where the empirical auxiliary parameter estimates may be found for more detail on their computation.

The model fits the data well. First, the model matches average time spent on job search ($\hat{s}$), the average probability that an individual receives a job offer ($\hat{o}$), and the average probability that an individual accepts conditional on receiving an offer ($\hat{a}$).[30] Because these moments are precisely estimated in the data, the estimation procedure ascribes high weight to matching them. Furthermore, the model closely matches the negative coefficient on duration ($\hat{\kappa}_{FE}$) in the baseline fixed effects regression in the online appendix, indicating that the model generates an empirically plausible degree of true negative duration dependence in search effort.

[30] This final point is not surprising given that b is calibrated to directly match the probability of acceptance. The values in the model and data are not identical due to the necessity of using a finite number of simulations.

Table 4: Auxiliary Parameters/Moments

Aux. Parameter	Concept	Reference	SUWNJ	Model
FE Baseline				
$\hat{\kappa}_{FE}$	Coefficient: Duration (FE)	Online appendix	-3.639	-3.574
Section 3				
$\hat{\pi}$	Coefficient: Past search	Table 2	-0.096	-0.080
$\hat{\phi}$	Coefficient: Job offers	Table 2	36.232	37.880
$\hat{\kappa}$	Coefficient: Duration	Table 2	3.096	1.391
Averages				
$\hat{s}$	Average: Search (minutes/day)		65.589	65.657
$\hat{o}$	Average: Offer probability		0.0256	0.0241
$\hat{a}$	Average: Acceptance probability		0.7415	0.7435

Source: Survey of Unemployed Workers in New Jersey.

Notes: Auxiliary moments from (i) SUWNJ data and (ii) estimated model.

Next, regarding the key parameters from Table 2 in Section 3, the model closely matches the coefficient on past search ($\hat{\pi}$) as well as the coefficient on job offers ($\hat{\phi}$). The former is slightly underestimated, while the latter is slightly overestimated, but the discrepancies are both small. These results represent strong evidence in support of the central conjecture of the paper—that individuals' idiosyncratic past search experiences play an important role in governing their subsequent behavior. On the other hand, while the model gets the right sign on the coefficient on duration ($\hat{\kappa}$) from Table 2—that is, inclusion of cumulative past search makes the effect of duration positive, as in the data—it underestimates the coefficient's magnitude relative to the SUWNJ. Given that $\hat{\kappa}$ is not statistically distinguishable from zero in the SUWNJ, this is not a particularly surprising result. One possible explanation for the discrepancy between the model and the data along this dimension is the observation that, for most people, search likely requires a fixed time outlay associated with, e.g., updating resumes, scanning help-wanted ads, etc. that, in isolation, would not yield any job offers, but is necessary to begin the offer-generating part of the search process. Consistent with this observation, introducing a fixed time outlay in the model raises $\hat{\kappa}$ while leaving intact $\hat{\kappa}_{FE}$, $\hat{\pi}$ and $\hat{\phi}$.[31]

5.3.3 Counterfactual analysis: Optimism and job search

As discussed above, a key result from estimation of the model is that, consistent with existing literature that exploits subjective probability elicitations, job seekers appear to significantly overestimate their job-finding prospects. How does this excessive optimism affect the trajectory of search effort over the unemployment spell?

To answer this question, I consider two versions of the model: (i) a version evaluated at the baseline estimation results (with a 57% bias in initial beliefs), and (ii) a version in which beliefs at the time of job loss are instead assumed to reflect the true average arrival rate of job offers in the population—that is, a version in which there is no bias. In the latter case, I hold other parameters fixed at their estimated values. Figure 5 depicts average search effort over the first year of unemployment in the two cases.

Interestingly, search effort is uniformly (and significantly) *lower* in the version of the model in which

[31] To understand why a fixed time outlay helps, note that if we neglect other right-hand side variables in (3), we have $\hat{\kappa} = \Delta \bar{s}_{it} - \hat{\pi}\bar{s}_{it-1}$. In the simplest case, adding a fixed amount of search time s^F yields $\hat{\kappa}^F = \Delta \bar{s}_{it} - \hat{\pi}(\bar{s}_{it-1} + s^F) > \hat{\kappa}$, where $\hat{\kappa}^F > \hat{\kappa}$ because $\hat{\pi} < 0$. Indeed, if the fixed time outlay is not very costly relative to "productive" search, a model with a simple linear search cost function and an estimated fixed time outlay provides an even stronger fit for the data than the model described in this section by alleviating the underestimation of $\hat{\kappa}$. Because a fixed time outlay is a non-standard feature in models of search, the preferred specification is the one presented here. Results for this alternative model are available upon request.

Figure 5: Optimism and job search

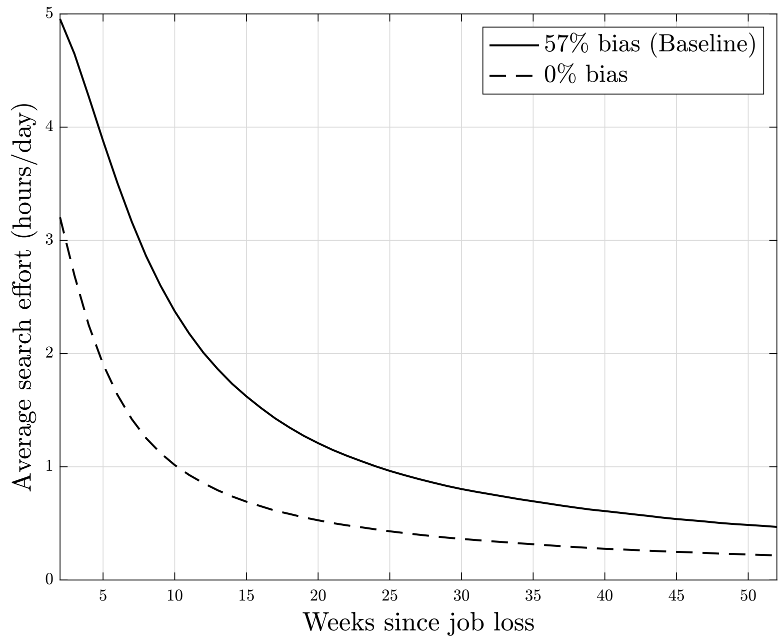

Notes: Counterfactually eliminating the positive bias in beliefs at the time of job loss depresses average search effort by reducing the perceived opportunity cost of leisure.

individuals are not overly optimistic about their job-finding prospects, despite the fact that the underlying economic environment is unchanged. That optimism raises search effort is not obvious a priori: Spinnewijn (2015) argues that job seekers who overestimate how quickly they will find work search too little as a result. However, in the present context, the search productivity effect of Section 4 appears to be operative and dominant, inducing a positive relationship between search effort and beliefs.[32]

6 Conclusion

The essential point of this paper is that job search appears to be an inherently path-dependent process: Bad luck early in the unemployment spell can depress subsequent search effort through its effect on a person's perceptions of their job-finding prospects, thus leading to progressive labor force withdrawal even in the absence of fundamental differences in motivation, search incentives, etc. Unfortunately, without high-frequency longitudinal data such as the SUWNJ, it is difficult to detect such path dependencies. Future research should therefore focus on making such data available in order to better understand what drives individuals' job search decisions and unemployment more broadly.

[32] Of course, one should not draw normative conclusions about the desirability of higher search effort without a more complete model that incorporates, e.g., general equilibrium effects through which increases in aggregate search effort may result in job searchers crowding each other out, or stimulate vacancy posting among firms. These are interesting questions, but beyond the scope of the present analysis.

Appendices

A Data and Robustness

A.1 Sample selection

The Survey of Unemployed Workers in New Jersey (SUWNJ) was conducted by the Princeton University Survey Research Center starting in the fall of 2009 and lasting for up to 24 weeks. A stratified random sampling procedure was used to select participants from the universe of individuals receiving unemployment-insurance (UI) benefits in New Jersey as of September 28, 2009. The original data were stratified by unemployment-duration intervals interacted with the availability of an e-mail address, over-sampling the long-term unemployed and those with e-mail addresses on file. To account for the considerable nonresponse rates, sample weights were created from the underlying administrative records. Because these records contained comprehensive demographic information for the universe from which the sample was drawn, nonresponse weights could be created by comparing the demographic characteristics of respondents and the underlying population of UI-benefit recipients. For a comprehensive description of the survey methodology, the reader is referred to Krueger and Mueller (2011).

Empirical results throughout the paper correspond to a subset of the respondents from the SUWNJ. Unless otherwise noted, the sample includes all prime-age individuals (ages 25-54) who, at the time of the interview, (i) had not accepted a job offer, (ii) did not work for pay in the current week, and (iii) did not expect to be recalled or return to a previous job. The single exception to this is Table A.1, in which I use the sample considered by Krueger and Mueller (2011), who adopt a broader definition of prime age (20-65) and include individuals expecting to be recalled or to return to a previous job.

I likewise follow Krueger and Mueller (2011) in defining time spent on job search. In particular, I only include observations for which at least 14 out of 16 episodes from the time diary were completed and for which respondents indicated at least four different activities over the course of the day. When two activities are reported, each activity is assumed to take 30 minutes. Finally, I trim observations in which time spent on job search exceeds 80 hours per week, and in which time spent on a search method is missing. For further details on construction of search time variables in both papers, see p. 55 of Krueger and Mueller (2011).

A.2 Krueger and Mueller (2011): Fixed Effects (FE) and First Diff. (FD)

Table A.1: Job Search over the Unemployment Spell

	Time Diary		Weekly Recall	
	FE	FD	FE	FD
Duration	-3.310***	-3.568***	-2.637***	-2.912***
	(0.192)	(0.463)	(0.264)	(0.602)
Observations	25366	20537	25640	20486
Adjusted R^2	0.103	0.069	0.020	0.003

Robust standard errors in parentheses.

Source: Survey of Unemployed Workers in New Jersey

Notes: Regressions use survey weights. Standard errors are robust and clustered at the individual level. Following Krueger and Mueller (2011), the sample consists of respondents ages 20-65 who have not yet accepted a job offer and are not currently employed. Sample sizes are smaller in the first-differenced specification due to the necessity of dropping observations without an associated lag.

A.3 Serial correlation test

Table A.2 reports the tests for first- and second-order autocorrelation in the first-differenced residuals developed by Arellano and Bond (1991).[33] If the errors ϵ_{it} of Equation (1) in levels are serially uncorrelated we should expect to see no evidence of second-order autocorrelation in the differenced residuals.[34]

Table A.2: Tests for serial correlation

	Time Diary		Weekly Recall	
	Statistic	p-value	Statistic	p-value
Arellano-Bond test for AR(1)	$z = -7.18$	0.000	$z = -5.40$	0.000
Arellano-Bond test for AR(2)	$z = 0.63$	0.526	$z = -0.75$	0.452

Source: Survey of Unemployed Workers in New Jersey.

H_0: No serial correlation.

The results in Table A.2 suggest that the disturbances ϵ_{it} are serially uncorrelated, and therefore that s_{it-2} is a valid instrument for s_{it-1}.

A.4 Stock-Flow Matching

Table A.3: Job Search over the Unemployment Spell

	Time Diary		Weekly Recall	
	Baseline	Augmented	Baseline	Augmented
Duration (κ)	-5.297***	4.827***	-4.648***	6.040*
	(1.524)	(1.633)	(1.514)	(3.315)
Past Search (π)		-0.154***		-0.0908***
		(0.0290)		(0.0249)
Observations	5464	5464	5416	5416
Adjusted R^2	0.068	0.212	0.006	0.087

Robust standard errors in parentheses.

Source: Survey of Unemployed Workers in New Jersey

Notes: Regressions use survey weights. Standard errors are robust and clustered at the individual level. The sample consists of respondents ages 25-54 who have not yet accepted a job offer, are not currently employed, and who do not expect to be recalled by or return to their former employer. The smaple is furthermore restricted to individuals who have been unemployed for at least 4 weeks, as described in the body of the text.

[33] The test was developed in the context of a GMM framework, but is nonetheless applicable to the simple 2SLS procedure used in the body of the paper.

[34] First-order autocorrelation in the first-differenced residuals results mechanically from the process of taking first differences.

A.5 Interview Effects in FE specification

Table A.4: Job Search over the Unemployment Spell

	Time Diary		Weekly Recall	
	FE	FE w/ interviews	FE	FE w/ interviews
Duration	-4.831***	-3.639***	-2.841***	-3.505***
	(0.435)	(0.808)	(0.679)	(1.155)
Previous interviews		-1.755*		0.992
		(0.975)		(1.324)
Observations	15731	15731	15916	15916
Adjusted R^2	0.120	0.121	0.021	0.021

Robust standard errors in parentheses.

Source: Survey of Unemployed Workers in New Jersey

Notes: Regressions use survey weights. Standard errors are robust and clustered at the individual level. The sample consists of respondents ages 25-54 who have not yet accepted a job offer, are not currently employed, and who do not expect to be recalled by or return to their former employer.

A.6 Robustness

A.6.1 Sample selection, duration trends and calendar time effects

Table A.5: Krueger and Mueller (2011) sample

	Time Diary		Weekly Recall	
	Baseline	Augmented	Baseline	Augmented
Duration (κ)	-5.646***	3.791**	-5.330***	4.287**
	(1.180)	(1.715)	(0.952)	(1.956)
Past Search (π)		-0.138***		-0.0847***
		(0.0240)		(0.0187)
Observations	9542	9542	9387	9387
Adjusted R^2	0.064	0.196	0.004	0.080

Robust standard errors in parentheses.

Source: Survey of Unemployed Workers in New Jersey

Table A.6: Robustness: Duration trends

	Time Diary		Weekly Recall	
	Baseline	Augmented	Baseline	Augmented
Past Search (π)	-0.149***	-0.149***	-0.0918***	-0.0931***
	(0.0276)	(0.0278)	(0.0253)	(0.0254)
Duration (κ)	4.411**	5.670*	6.490**	0.276
	(1.845)	(2.902)	(3.185)	(3.989)
Log(Duration)	15.23		-26.94	
	(18.96)		(25.88)	
Duration^2		-0.0314		0.0652
		(0.0650)		(0.0883)
Duration^3		0.000293		0.0000145
		(0.000412)		(0.000563)
Observations	5497	5497	5445	5445
Adjusted R^2	0.207	0.207	0.087	0.089

Robust standard errors in parentheses.

Source: Survey of Unemployed Workers in New Jersey

Table A.7: Robustness: Calendar time effects (HWOL data)

	Time Diary		Weekly Recall	
	Baseline	Augmented	Baseline	Augmented
Duration (κ)	-7.150***	3.385*	-6.916***	3.997
	(1.220)	(1.789)	(1.380)	(2.751)
Past Search (π)		-0.154***		-0.0908***
		(0.0268)		(0.0250)
Vacancies	-0.00212***	-0.00188***	-0.000880**	-0.000832**
	(0.000483)	(0.000470)	(0.000391)	(0.000367)
Observations	5497	5497	5445	5445
Adjusted R^2	0.073	0.217	0.006	0.087

Robust standard errors in parentheses.

Source: Survey of Unemployed Workers in New Jersey and Conference Board Help Wanted OnLine

A.6.2 Interview effects

In his discussion of Krueger and Mueller (2011), Steven Davis suggests the possibility that respondents may "learn that affirmative responses to questions about job search trigger additional questions" and so "become more likely to falsely report no job search as they gain familiarity with the questionnaire," thus explaining the observed negative duration dependence in time devoted to job search. Following Davis's discussion of Krueger and Mueller (2011), I make use of two pieces of evidence to evaluate the likelihood that interview effects are responsible for my results.

I first estimate equation (2) restricting attention to individuals who have reported searching each week, and who therefore cannot have learned that reporting no search reduces the interview length. Table A.8, below, reports the results.

Table A.8: Robustness: Intensive margin

	Time Diary		Weekly Recall	
	Baseline	Augmented	Baseline	Augmented
Duration (κ)	-9.503***	3.350	-5.644***	4.628
	(2.899)	(5.060)	(1.559)	(3.024)
Past Search (π)		-0.0997***		-0.0694***
		(0.0381)		(0.0206)
Observations	1546	1546	3996	3996
Adjusted R^2	0.072	0.160	0.007	0.073

Robust standard errors in parentheses.

Source: Survey of Unemployed Workers in New Jersey

The coefficients on cumulative past search (for both measures of search) are attenuated relative to the results in Table 1 of the manuscript, but remain highly significant and negative. Moreover, inclusion of cumulative past search continues to render the effect of duration on search insignificant, whereas in the "Baseline" model the coefficient on duration is highly significant and negative.

Next, I consider including the number of past interviews in the regression analysis. There is an important problem with this approach, unique to my analysis of the effect of cumulative past search, which precludes separate identification of the effect of duration from the effect of interview number. To understand the problem, note that (i) differencing the data is necessary in order to clean out unobservable terms from equation (1); and (ii) because we cannot rule out individuals searching during weeks in which they missed/skipped interviews, differences must be taken with respect to *calendar week*, not interview number.[35]

I discuss the first point above in Section 3. Regarding the second point, differencing with respect to interview number treats search in the week prior to the missed interview as the relevant explanatory variable for the change in search between the week prior to the missed interview and the week following the missed interview—thus implicitly assuming that no search occurs during the week of the missed interview. Because there is no basis for this assumption, differences should be taken with respect to calendar week, which in turn implies that the effects of duration and interview number cannot be separately identified.

If we are to disregard the preceding, as well as the fact that skipped interviews are not randomly assigned, we can, as a strictly mechanical matter, estimate (2) with first differences taken with respect to interview number. Table A.9 reports these results for completeness.

[35]In the parlance of Stata, the time series variable must be chosen to be calendar date (or duration), not interview number. Thus, for an individual who responds to the survey for the first three weeks, skips the fourth, and then responds in the fifth and sixth, differenced and lagged variables in week five will be reported as missing/empty.

Table A.9: Robustness: Interview effects

	Time Diary		Weekly Recall	
	Baseline	Augmented	Baseline	Augmented
Duration (κ)	-1.479	-1.205	-0.599	-0.745
	(1.219)	(1.184)	(1.543)	(1.515)
Past Search (π)		-0.109***		-0.0573***
		(0.0267)		(0.0200)
Observations	10155	10155	10125	10125
Adjusted R^2	0.070	0.166	0.004	0.058

Robust standard errors in parentheses.

Source: Survey of Unemployed Workers in New Jersey

The effect of cumulative past search remains highly significant and negative, consistent with job seekers learning about the arrival rate of offers ("Augmented"). However, the presence of the number of past interviews as a control variable renders duration insignificant ("Baseline").[36] Both of these results should be interpreted with caution: The effect of duration (κ) is identified only in a statistical sense, due to non-random assignment of missed interviews (a point made by Krueger and Mueller (2011) and also noted by Davis in his comments). Furthermore, the effect of past search (π) is based on a systematic mis-measurement of past search, resulting from differencing with respect to interview number—see the discussion above.

A.7 GMM

There are two principal drawbacks to the 2SLS procedure used in the body of the text. First, it neglects the additional moment conditions implied by the exogeneity of s_{it-2}. Second, the process of first-differencing induces potential data loss due to missed interviews. I address both of these concerns using the GMM estimators for dynamic panels developed by Holtz-Eakin et al. (1988) and Arellano and Bond (1991).[37]

A.7.1 Differences

To exploit the additional available moment conditions, I estimate the model using the Difference GMM estimator developed by Arellano and Bond (1991). Table A.10 reports the results and Table A.11 reports the associated tests of instrument validity.

A.7.2 Orthogonal deviations

To circumvent the potential data loss associated with differencing, I also estimate a version of the model in which individual effects are purged by taking forward-orthogonal deviations.[38] Table A.12 reports the results, and Table A.13 reports the associated tests of instrument validity.

A.7.3 Discussion

The results above correspond to the two-step estimators with Windmeijer (2005)-corrected standard errors. To avoid instrument proliferation, which can overfit the model and weaken the Hansen test, I restrict attention

[36] The change in the sample associated with differencing with respect to interview number accounts for roughly 45% of the decline in the magnitude of the effect of duration in the "Baseline" column of Table A.9 relative to Table 1 for time diary data, and roughly 35% for the weekly recall data. The balance is accounted for by interview effects.

[37] Specifically, I focus on Difference GMM and Orthogonal Deviations GMM. A System GMM approach is ruled out because for most individuals, the stock variable of interest is itself partially unobserved.

[38] The forward-orthogonal deviation of y_{it} is defined as $y^{\perp}_{it+1} \equiv c_{it}\left[y_{it} - \frac{1}{T_{it}}\sum_{s>t} y_{is}\right]$ where $c_{it} \equiv \sqrt{T_{it}/(T_{it}+1)}$.

Table A.10: Robustness: Two-step GMM (Differences)

	Time Diary		Weekly Recall	
	Baseline	Augmented	Baseline	Augmented
Duration (κ)	-7.297***	-1.131	-3.836***	0.943
	(1.281)	(3.446)	(1.211)	(3.128)
Past Search (π)		-0.0808**		-0.0476**
		(0.0354)		(0.0206)
Observations	8752	6487	8728	6427

Robust standard errors in parentheses.
Source: Survey of Unemployed Workers in New Jersey

Table A.11: Tests of serial correlation and over-identifying restrictions (differences)

	Time Diary		Weekly Recall	
	Statistic	p-value	Statistic	p-value
Arellano-Bond test for AR(1)	$z = -8.66$	0.000	$z = -6.97$	0.000
Arellano-Bond test for AR(2)	$z = -0.38$	0.706	$z = -0.94$	0.346
Sargan test of over-ID restrictions	$\chi^2_8 = 94.67$	0.000	$\chi^2_8 = 48.20$	0.000
Hansen test of over-ID restrictions	$\chi^2_8 = 18.12$	0.020	$\chi^2_8 = 18.07$	0.021

Source: Survey of Unemployed Workers in New Jersey.
H_0 (AB): No serial correlation; H_0 (Sargan/Hansen): Instruments are jointly exogenous.

Table A.12: Robustness: Two-step GMM (Orthog. deviations)

	Time Diary		Weekly Recall	
	Baseline	Augmented	Baseline	Augmented
Duration (κ)	-5.971***	0.235	-4.495***	-0.332
	(0.425)	(1.025)	(0.524)	(1.327)
Past Search (π)		-0.0971***		-0.0428***
		(0.0163)		(0.0132)
Observations	13078	9366	13186	9371

Robust standard errors in parentheses.
Source: Survey of Unemployed Workers in New Jersey

Table A.13: Tests of serial correlation and over-identifying restrictions (orthogonal deviations)

	Time Diary		Weekly Recall	
	Statistic	p-value	Statistic	p-value
Arellano-Bond test for AR(1)	$z = -9.28$	0.000	$z = -7.23$	0.000
Arellano-Bond test for AR(2)	$z = -0.30$	0.761	$z = -1.23$	0.220
Sargan test of over-ID restrictions	$\chi^2_8 = 115.07$	0.000	$\chi^2_8 = 67.06$	0.000
Hansen test of over-ID restrictions	$\chi^2_8 = 15.30$	0.053	$\chi^2_8 = 21.80$	0.005

Source: Survey of Unemployed Workers in New Jersey.
H_0 (AB): No serial correlation; H_0 (Sargan/Hansen): Instruments are jointly exogenous.

to a “collapsed” instrument matrix.

Focusing first on parameter estimates in Tables A.10 and A.12, for both differencing and orthogonal deviations, the estimated coefficients on cumulative past search are highly significant and negative. Furthermore, in both cases, the coefficient on duration is attenuated dramatically by the presence of cumulative past search and becomes insignificant, consistent with the results in the body of the text and the claim that cumulative past search accounts for much of the decline in search over the unemployment spell.

Turning next to the tests of serial correlation and over-identifying restrictions in Tables A.11 and A.13, there is no significant evidence of serial correlation in the differenced errors. This suggests that the second lag and beyond of the dependent variable are valid instruments. The Hansen and Sargan tests, however, reject the null of joint validity for both measures of search time. While there is clearly some disrepancy in these results, Arellano and Bond (1991) use simulated panel data from an AR(1) model to demonstrate that their test for serial correlation has greater power than Hansen-Sargan tests to detect invalidity of lagged instruments due to serial correlation. Thus, to the extent that their analysis is applicable here, there is at least some reason to prefer the Arellano-Bond test of second-order serial correlation when assessing the validity of the instruments.

B Model

B.1 Posterior distribution of beliefs

This section demonstrates that the Gamma distribution is the conjugate prior for the right-censored exponential distribution, and derives the laws of motion for the parameters of the belief distribution with Bayesian updating. Consider an individual who has been unemployed for n weeks. For each week $t = 1, ..., n$ of the unemployment spell, the individual allocates s_t units of time for job search. Define $K \equiv \{t : \tau_t \leq s_t\}$ as the set of weeks in which an offer (below the reservation wage) arrives before search ends. Furthermore, let $n^s \equiv \#K$ and $n^f \equiv n - n^s$. For weeks $t \in K$, individuals observe the exact arrival time $\tau_t \leq s_t$. For the remaining weeks $t \notin K$, individuals only observe that $\tau_t > s_t$.

Because offers arrive according to a Poisson process with unobserved rate parameter λ, arrival times are distributed according to a right-censored exponential distribution with distribution function F and density f. The corresponding likelihood function for λ is thus given by

$$\begin{aligned}
\mathscr{L}(\lambda) &= \mathscr{L}\left(\lambda | \{\tau_t\}_{t\in K}; \{s_t\}_{t\notin K}\right) \\
&= \prod_{t\in K} f(\tau_t|\lambda) \prod_{t\notin K} (1 - F(s_t|\lambda)) \\
&= \prod_{t\in K} \lambda e^{-\lambda\tau_t} \prod_{t\notin K} e^{-\lambda s_t} \\
&= \lambda^{n^s} e^{-\lambda(\sum_{t\in K}\tau_t + \sum_{t\notin K} s_t)} \\
&= \lambda^{n^s} e^{-\lambda(n^s\bar{\tau} + n^f\bar{s})} \qquad (31)
\end{aligned}$$

where $\bar{\tau} \equiv \frac{1}{n^s}\sum_{t\in K}\tau_t$ and $\bar{s} \equiv \frac{1}{n^f}\sum_{t\notin K} s_t$.

Suppose now that prior beliefs over λ follow a Gamma distribution with hyperparameters α_0 and β_0, distribution function $G(\lambda|\alpha_0, \beta_0)$, and density $g(\lambda|\alpha_0, \beta_0)$. Applying Bayes' rule and using the expression for the likelihood function above, the posterior distribution of beliefs over λ is given by

$$\begin{aligned}
p(\lambda) &= \frac{\mathscr{L}(\lambda) g(\lambda|\alpha_0, \beta_0)}{\int \mathscr{L}(\lambda') g(\lambda'|\alpha_0, \beta_0) d\lambda'} \\
&= \frac{\lambda^{n^s} e^{-\lambda(n^s\bar{\tau}+n^f\bar{s})} \beta_0^{\alpha_0} \lambda^{\alpha_0-1} e^{-\lambda\beta_0}/\Gamma(\alpha_0)}{\int (\lambda')^{n^s} e^{-\lambda'(n^s\bar{\tau}+n^f\bar{s})} \beta_0^{\alpha_0} (\lambda')^{\alpha_0-1} e^{-\lambda'\beta_0}/\Gamma(\alpha_0) d\lambda'} \\
&= \frac{e^{-\lambda(\beta_0+n^s\bar{\tau}+n^f\bar{s})} \lambda^{\alpha_0+n^s-1}}{\int e^{-\lambda'(\beta_0+n^s\bar{\tau}+n^f\bar{s})} (\lambda')^{\alpha_0+n^s-1} d\lambda'} \\
&= \frac{e^{-\lambda(\beta_0+n^s\bar{\tau}+n^f\bar{s})} \lambda^{\alpha_0+n^s-1} (\beta_0 + n^s\bar{\tau} + n^f\bar{s})^{\alpha_0+n^s}}{\int e^{-\lambda'(\beta_0+n^s\bar{\tau}+n^f\bar{s})} (\lambda')^{\alpha_0+n^s-1} d\lambda' (\beta_0 + n^s\bar{\tau} + n^f\bar{s})^{\alpha_0+n^s}} \qquad (32)
\end{aligned}$$

Defining $x' \equiv \lambda'(\beta_0 + n^s\bar{\tau} + n^f\bar{s})$, we can rewrite the denominator of (32) in terms of x' as follows

$$\begin{aligned}
&\int e^{-x'} \left(\frac{x'}{\beta_0 + n^s\bar{\tau} + n^f\bar{s}}\right)^{\alpha_0+n^s-1} dx' (\beta_0 + n^s\bar{\tau} + n^f\bar{s})^{\alpha_0+n^s-1} \\
&= \int e^{-x'} \left(x'\right)^{\alpha_0+n^s-1} dx' \\
&= \Gamma(\alpha_0 + n^s). \qquad (33)
\end{aligned}$$

Substituting (33) into (32) and defining $\alpha \equiv \alpha_0 + n^s$ and $\beta \equiv \beta_0 + n^s\bar{\tau} + n^f\bar{s}$, (32) reduces to

$$\begin{aligned} p(\lambda) &= \frac{\lambda^{\alpha-1}e^{-\lambda\beta}\beta^{\alpha}}{\Gamma(\alpha)} \\ &= g(\lambda|\alpha,\beta). \end{aligned} \tag{34}$$

Thus, as claimed in the text, the Gamma distribution with prior hyperparameters α_0 and β_0 is the conjugate prior for the right-censored exponential distribution. Moreover, the posterior hyperparameters α and β, which govern the evolution of beliefs in the model, are defined recursively as

$$\alpha = \alpha_0 + n^s \tag{35}$$

$$\beta = \beta_0 + \sum_{t\in K}\tau_t + \sum_{t\notin K} s_t. \tag{36}$$

Intuitively, the posterior hyperparameters net of their initial values measure the total number of job offers received and the total past time spent looking for work, respectively.

B.2 Proofs of Propositions

For ease of notation, let $\tilde{F}(s_t;\alpha_t,\beta_t)$ denote the perceived probability of receiving a job offer and let $\tilde{f}(s_t;\alpha_t,\beta_t)$ denote the derivative of $\tilde{F}$ with respect to s_t:

$$\begin{aligned} \tilde{F}(s_t;\alpha_t,\beta_t) &\equiv \int_0^{\infty} F(s_t;\lambda)d\Gamma(\lambda;\alpha_t,\beta_t) = 1 - \left(\frac{\beta_t}{\beta_t + s_t}\right)^{\alpha_t} \\ \tilde{f}(s_t;\alpha_t,\beta_t) &\equiv \frac{\partial \tilde{F}}{\partial s_t} = \frac{\alpha_t\beta_t^{\alpha_t}}{(\beta_t + s_t)^{\alpha_t+1}}. \end{aligned}$$

B.2.1 Proposition 1

Proof. Differentiating (13) with respect to β_t and using the envelope theorem, we have

$$\begin{aligned} \frac{\partial \underline{w}_t}{\partial \beta_t} &= (1-\delta)\frac{\partial V_t^U}{\partial \beta_t} \\ &= (1-\delta)\frac{\frac{\partial \tilde{F}}{\partial \beta_t}\frac{\delta}{1-\delta}\int_{\underline{w}_t}^{B}(\omega - \underline{w}_t)\phi(\omega)d\omega}{1-\delta(1-\tilde{F}(s_t;\alpha_t,\beta_t)(1-\Phi(\underline{w}_t)))} \\ &= -(1-\delta)\frac{\frac{s_t}{\beta_t}\tilde{f}(s_t;\alpha_t,\beta_t)\frac{\delta}{1-\delta}\int_{\underline{w}_t}^{B}(\omega - \underline{w}_t)\phi(\omega)d\omega}{1-\delta(1-\tilde{F}(s_t;\alpha_t,\beta_t)(1-\Phi(\underline{w}_t)))} \\ &= -(1-\delta)\frac{\eta s_t/\beta_t}{1-\delta(1-\tilde{F}(s_t;\alpha_t,\beta_t)(1-\Phi(\underline{w}_t)))} \\ &< 0. \end{aligned}$$

□

B.2.2 Proposition 2

Proof.

1. Inspection of equation (19) indicates that the option value effect is positive if and only if $\frac{\partial \underline{w}_t}{\partial \beta_t} < 0$, which is the result proved in Proposition 1.

2. From the expression for the search productivity effect in (19), we can write

$$\frac{\alpha_t s_t - \beta_t}{\beta_t(\alpha_t + 1)} < 0 \iff \frac{\alpha_t}{\beta_t} < \frac{\alpha_t + 1}{\beta_t + s_t}$$
$$\iff E^\lambda \left[\lambda^T\right] < E^\lambda \left[\lambda^T | \text{offer}\right].$$

3. Follows immediately from (19).

$\square$

B.2.3 Proposition 3

Proof. Taking a first-order Taylor expansion of the model around initial beliefs, we obtain

$$s_t \approx s(\alpha_0, \beta_0) + \frac{ds_t}{d\beta_t}\Big|_{\beta_t = \beta_0} (\beta_t - \beta_0) + \frac{ds_t}{d\alpha_t}\Big|_{\alpha_t = \alpha_0} (\alpha_t - \alpha_0).$$

To ease exposition, but without loss of generality, assume that individuals finish searching each day regardless of whether or not an offer arrives. Then the law of motion for β_t is given by $\beta_t = \beta_0 + \sum_{\tau=0}^{t-1} s_\tau$ and the law of motion for α_t is given by $\alpha_t = \alpha_0 + \sum_{\tau=0}^{t-1} \mathbb{1}(\text{Offer in period } \tau) = \alpha_0 + \sum_{\tau=0}^{t-1} o_\tau$. Substituting these laws of motion into the preceding and letting $\pi = \frac{ds_t}{d\beta_t}\big|_{\beta_t=\beta_0}$, $\phi = \frac{ds_t}{d\alpha_t}\big|_{\alpha_t=\alpha_0}$ and $\iota + \eta_i = s(\alpha_0, \beta_0)$, we have

$$s_t \approx \iota + \pi \sum_{\tau=0}^{t-1} s_\tau + \phi \sum_{\tau=0}^{t-1} o_\tau + \eta_i$$

which is just equation (3) from Section 3 with $\kappa = \gamma = 0$ and reindexed so that time begins at $t = 0$ instead of $t = 1$.[39] Note that this implies that π is simply the sum of the search productivity and option value effects, and will therefore be negative so long as the search productivity effect dominates the option value effect. $\square$

[39] There is also a slight difference between the preceding and (3) in the maximum index on the summation of offers. This is simply an artifact of the timing in the time diary data in the SUWNJ that was exploited to identify the effect of offers and is not of any economic significance.

C Estimation Details

C.1 Model

As described in the body of the text, I enrich the simple model described in Section 4 with on-the-job search, a convex search cost function, and the possibility of offers arriving without search.

Formally, let ρ denote the exogenous separation rate, let $\eta_i(s_t)$ denote the total cost of search effort s_t for $i \in \{e, u\}$, and let $\xi_i \geq 0$ denote the exogenous probability that an individual receives an offer independent of search, also for $i \in \{e, u\}$. To ease notation, let $\tilde{F}(s_t) = \tilde{F}(s_t; \alpha_t, \beta_t) \equiv E^\lambda[F(s_t)|\alpha_t, \beta_t]$ denote the perceived probability of receiving an offer given $\xi_i = 0$, search effort s_t, and beliefs α_t and β_t, and let $\tilde{f}(s_t)$ denote the derivative of $\tilde{F}$ with respect to s_t. Because the problem is perceived to be stationary under anticipated utility, we can neglect time subscripts when solving the model.

Let $V^E(w)$ denote the value of being employed at wage w, and let V^U denote the value of unemployment. Then we have

$$\begin{aligned} V^E(w) = \max_{s_e} \Big\{ & w - \eta_e(s_e) \\ & + \delta(1-\rho)\Big[[\tilde{F}(s_e) + (1-\tilde{F}(s_e))\xi_e] \int_0^B \max\{V^E(x), V^E(w)\}\phi(x)dx \\ & + [(1-\tilde{F}(s_e))(1-\xi_e)]V^E(w)\Big] \\ & + \delta\rho V^U \Big\}. \end{aligned} \tag{37}$$

Defining $\tilde{\tilde{F}}_e(s_e) \equiv \tilde{F}(s_e) + (1-\tilde{F}(s_e))\xi_e$ as the composite job-finding probability, we can write

$$\begin{aligned} V^E(w) = \max_{s_e} \Big\{ w - \eta_e(s_e) + \delta(1-\rho)\Big[V^E(w) + \tilde{\tilde{F}}_e(s_e) \int_w^B \Big(V^E(x) - V^E(w)\Big)\phi(x)dx\Big] \\ + \delta\rho V^U \Big\} \end{aligned} \tag{38}$$

which follows from the fact that the solution to the functional equation (37), $V^E(\cdot)$, is an increasing function, so the worker accepts any wage that exceeds her current wage. Letting $\tilde{\tilde{f}}_e(s_e) \equiv \frac{d\tilde{\tilde{F}}_e}{ds_e}$, the corresponding first-order condition for search effort is given by

$$\eta_e'(s_e^*) = \delta(1-\rho)\tilde{\tilde{f}}_e(s_e^*(w)) \int_w^B \Big(V^E(x) - V^E(w)\Big)\phi(x)dx. \tag{39}$$

As expected, search effort is decreasing in w. Using the envelope theorem, we can write

$$\begin{aligned} V_1^E(w) &= 1 + \delta(1-\rho)\left[V_1^E(w) - \tilde{\tilde{F}}_e(s_e^*(w))V_1^E(w)(1-\Phi(w))\right] \\ &= 1 + \delta(1-\rho)V_1^E(w)\left[1 - \tilde{\tilde{F}}_e(s_e^*(w))(1-\Phi(w))\right]. \end{aligned}$$

Solving for $V_1^E(w)$ gives

$$V_1^E(w) = \frac{1}{1 - \delta(1-\rho)\left[1 - \tilde{\tilde{F}}_e(s_e^*(w))(1-\Phi(w))\right]} > 0. \tag{40}$$

Expanding the right-hand side of (39) via integration by parts and using (40), we can write

$$\begin{aligned}\eta_e'(s_e^*(w)) &= \delta(1-\rho)\tilde{\tilde{f}}_e(s_e^*(w))\int_w^B V_1^E(x)(1-\Phi(x))dx \\ &= \delta(1-\rho)\tilde{\tilde{f}}_e(s_e^*(w))\int_w^B \frac{1-\Phi(x)}{1-\delta(1-\rho)\left[1-\tilde{\tilde{F}}_e(s_e^*(x))(1-\Phi(x))\right]}dx. \end{aligned} \tag{41}$$

The value of being unemployed, in turn, is given by[40]

$$V^U = \max_{s_u}\left\{b-\eta_u(s_u)+\delta\Big[V^U+\tilde{\tilde{F}}_u(s_u)(1-\rho)\int_{\underline{w}}^B \Big(V^E(x)-V^U\Big)\phi(x)dx\Big]\right\} \tag{42}$$

where the second line again follows from the fact that $V^E(\cdot)$ is increasing, and the worker's reservation wage $\underline{w}$ solves

$$V^E(\underline{w}) = V^U. \tag{43}$$

Using this fact, we can evaluate $V^E(w)$ at $\underline{w}$ and obtain

$$V^E(\underline{w}) = \frac{1}{1-\delta}\left[\underline{w}-\eta_e(s_e^*(\underline{w}))+\delta(1-\rho)\tilde{\tilde{F}}_e(s_e^*(\underline{w}))\int_{\underline{w}}^B \Big(V^E(x)-V^E(\underline{w})\Big)\phi(x)dx\right]. \tag{44}$$

Similarly,

$$V^U = \frac{1}{1-\delta}\left[b-\eta_u(s_u^*)+\delta(1-\rho)\tilde{\tilde{F}}_u(s_u^*)\int_{\underline{w}}^B \Big(V^E(x)-V^E(\underline{w})\Big)\phi(x)dx\right]. \tag{45}$$

Recalling that $V^U = V^E(\underline{w})$, we can equate (44) and (45) and rearrange to obtain

$$\underline{w}-b = \Big[\eta_e(s_e^*(\underline{w}))-\eta_u(s_u^*)\Big]-\Big[\tilde{\tilde{F}}_e(s_e^*(\underline{w}))-\tilde{\tilde{F}}_u(s_u^*)\Big]\delta(1-\rho)\int_{\underline{w}}^B \Big(V^E(x)-V^U\Big)\phi(x)dx. \tag{46}$$

Using the method described above to simplify the integral, we arrive at

$$\begin{aligned}\underline{w} = b &+ \Big(\eta_e(s_e^*(\underline{w}))-\eta_u(s_u^*)\Big) \\ &- \Big[\tilde{\tilde{F}}_e(s_e^*(\underline{w}))-\tilde{\tilde{F}}_u(s_u^*)\Big]\delta(1-\rho)\int_{\underline{w}}^B \frac{1-\Phi(x)}{1-\delta(1-\rho)\left[1-\tilde{\tilde{F}}_e(s_e^*(x))(1-\Phi(x))\right]}dx.\end{aligned} \tag{47}$$

The solution to the general model is characterized by (47), together with the first-order conditions for search effort among the employed and unemployed above.

A well-known feature of models of on-the-job search is that, when the search environments of the employed and unemployed are symmetric, $\underline{w} = b$. Inspection of (47) indicates that this result obtains in the current context when $\eta_e(s) = \eta_u(s)$, $\xi_e = \xi_u$, and job seekers perceive that the return to search is the same during employment and unemployment. Because of the large degree of ex post heterogeneity in the current model resulting from endogenously evolving beliefs α_t and β_t, maintaining this assumption of symmetry dramatically reduces the computational burden associated with estimating the model, and is therefore the case I focus on in the body of the text.

[40] I assume that newly arrived offers are destroyed at the same rate ρ as offers accruing to currently employed workers so that the problems of employed and unemployed workers are symmetric in this regard.

C.2 Numerical solution and simulation

The model described in Appendix C.1 cannot be solved analytically. I therefore numerically compute search effort on a 5-by-500 grid of values for α_t and β_t. I compute the policy functions as linear interpolations in β_t for each of the five possible values of α_t.

To simulate the model, it is necessary to generate two shock matrices. The first is a 500,000-by-100 matrix of exponential offer-arrival times. The second is a 500,000-by-100 matrix of lognormal wage draws. Because I estimate parameters that govern λ^T (and in some cases ν, although not the case considered in the text), I need to hold constant the underlying stochastic process in the course of estimation. Accordingly, prior to estimation, I generate three 500,000-by-100 matrices of uniformly distributed shocks. These are held fixed throughout the course of estimation. For each value of λ^T considered by the minimization routine, I compute the associated exponential arrival-time shocks by way of an inverse transform sampling procedure using the first matrix of uniform shocks. For each value of ν considered by the minimization routine (when this parameter is being estimated), I compute the lognormal wage shocks by way of a standard Box-Muller transform of the two remaining matrices of uniform shocks. This ensures that the surface of the objective function is stable across iterations, but dependent on λ^T and ν. I simulate 500,000 individuals each for up to 100 weeks of unemployment. Job seekers who accept offers are dropped from the sample, as in the SUWNJ. The sample is sufficiently large to permit replication of the cohort structure of the SUWNJ. Remaining details of the estimation methodology are discussed in Section 5.

C.3 Sensitivity analysis

In order to ensure that the results obtained in Section 5 are not an artifact of particular choices regarding calibration or weighting, I consider two alternative approaches to estimation. First, I re-estimate the model with a lower value for the variance of the offer distribution. In the baseline estimation, I calibrate ν to match the standard deviation of log job values of 0.38 estimated in Hall and Mueller (2018). Because this value is somewhat high relative to existing literature, in the first alternative specification I instead calibrate ν to match the standard deviation of log *wages* estimated in Hall and Mueller (2018) of 0.24. This value is also similar to the estimate from Low et al. (2010). The fourth column in Table C.1 reports the resulting parameter estimates and the fifth column in Table C.2 reports the implied simulated moments. Second, I re-estimate the model using an identity weighting matrix, using percent deviations between empirical and simulated moments instead of level differences so that units do not arbitrarily affect weighting. The fifth column in Table C.1 reports the resulting parameter estimates and the sixth column in Table C.2 reports the implied simulated moments.

Table C.1: Sensitivity analysis: Parameter estimates

Parameter	Concept	Baseline	$\sigma = 0.24$	$W = I$
Beliefs				
α_0	Initial beliefs (shape)	0.39 (0.10)	0.33	0.14
β_0	Initial beliefs (rate)	0.27 (0.17)	0.23	0.06
Bias	$\left(\frac{\alpha_0/\beta_0 - E[\lambda^T]}{E[\lambda^T]}\right) \cdot 100$	57% (23%)	48%	51%
Physical				
ψ	Search cost (level)	7.91 (1.19)	4.47	4.49
Ψ	Search cost (curvature)	0.26 (0.03)	0.28	0.42
b	Flow value of unempl.	0.73	0.83	0.73

Source: Survey of Unemployed Workers in New Jersey.

The main qualitative results obtained in Section 5 are not significantly affected by reducing the offer variance or using an identity weighting matrix. Furthermore, both alternative estimation procedures provide a strong fit for the data. In particular, all three estimations find a large positive bias in beliefs of around 50% and low values for α_0 and β_0, consistent with high variance in beliefs and dispersion in offer arrival rates. Both of the alternative specifications identify a lower value of ψ, although for different reasons: In the case of the low offer variance specification, this results from the fact that search is less valuable when the offer distribution is less spread out, which implies that, all else equal, a lower marginal cost of search is required to generate a given value of average search effort. In the case of the identity weighting matrix, this is likely due to the fact that the estimation procedure places relatively less weight on matching average search effort $\hat{s}$ (which is very precisely estimated in the data), and relatively more weight on matching $\hat{\kappa}$ (which is imprecisely estimated in the data). This is reflected in the fact that the identity-weighted estimation (i) overestimates average search effort in Table C.2 and (ii) finds a higher value for $\hat{\kappa}$, as would be expected from a model with higher average search effort for reasons discussed in the body of the text. Finally, the marginally higher value of b found in the model with lower offer variance is a mechanical result of directly identifying b via the offer distribution in (29).

Table C.2: Sensitivity analysis: Auxiliary parameters/moments

Aux. Parameter	Concept	SUWNJ	Baseline	$\sigma = 0.24$	$W = I$
FE Baseline					
$\hat{\kappa}_{FE}$	Coefficient: Duration (FE)	-3.639	-3.574	-3.579	-3.797
Section 3					
$\hat{\pi}$	Coefficient: Past search	-0.096	-0.080	-0.803	-0.096
$\hat{\phi}$	Coefficient: Job offers	36.232	37.880	36.348	34.524
$\hat{\kappa}$	Coefficient: Duration	3.096	1.391	1.425	2.783
Averages					
$\hat{s}$	Average: Search (minutes/day)	65.589	65.657	65.590	70.185
$\hat{o}$	Average: Offer probability	0.0256	0.0241	0.0237	0.0209
$\hat{a}$	Average: Acceptance probability	0.7415	0.7435	0.7438	0.7438

Source: Survey of Unemployed Workers in New Jersey.

C.4 Intuition for identification

Because of the highly nonlinear nature of the mapping from the structural parameters to the simulated moments, the parameter estimates in Table 3 are, together, simultaneously identified by the seven moments in Table 4 (the exception is b, which is directly identified by $\hat{a}$ via (29)). It is nonetheless instructive to reflect on the key ways in which the auxiliary moments vary with the estimated parameters. To this end, Figure 6 illustrates how each moment in Table 4 varies with at least one parameter in Table 3.

In the first panel of the first row, we see that average search effort $\hat{s}$ falls sharply as ψ rises, reflecting the predictable effect of an increase in the marginal cost of search on individuals' search decisions. In the second panel, in turn, we see that the average probability that an individual receives a job offer, $\hat{o}$, falls as $Bias$ rises. This reflects the fact that, for given beliefs α_0 and β_0, a greater positive bias in those beliefs requires a lower arrival rate of offers during search. The third panel reflects the direct relationship between the gross flow value of unemployment, b, and the probability that an individual accepts conditional on receiving an offer, $\hat{a}$, embodied in (29).

Turning next to the moments from Section 3, in the first panel of the second row, we see that the magnitude of the coefficient on duration from the simple fixed effects regression, $\hat{\kappa}_{FE}$, falls as β_0 rises. In a

Figure 6: Identification

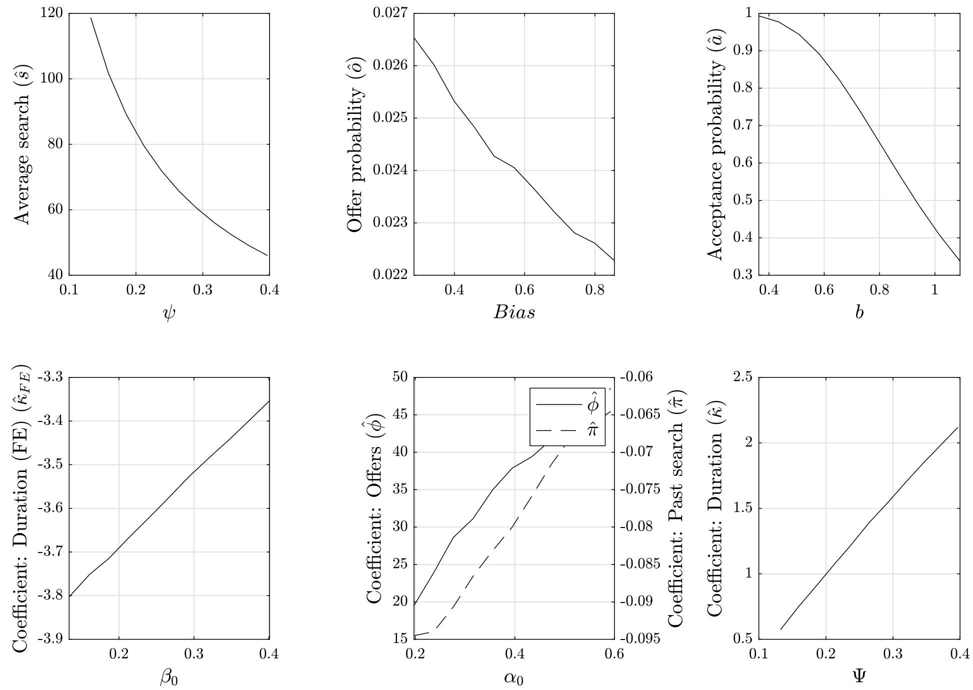

Notes: Panels depicts how moments used in the auxiliary model vary with the estimated structural parameters.

global sense, the data unambiguously want a low value of β_0 in order to generate a high variance in the belief distribution and correspondingly strong learning effects, which in turn allow the model to account for the main results in Section 3. Locally, however—that is, conditional on a relatively low value of β_0—changes in β_0 have a relatively small effect on simulated moments through search decisions (because accumulating search experience quickly comes to dominate β_t), and instead primarily operate through the latent heterogeneity in true arrival rates:[41] For a given value of α_0, a higher value of β_0 implies less heterogeneity in arrival rates, fewer individuals receiving offers early in the spell, and so a lower proportion of slowly decaying search profiles associated with early job offers (see Figure 4), implying less average negative duration dependence as reflected in a value of $\hat{\kappa}_{FE}$ attenuated towards zero.

In the second panel, we see that $\hat{\pi}$ decreases in magnitude, while $\hat{\phi}$ increases sharply, as α_0 rises. That the magnitude of $\hat{\pi}$ falls as α_0 rises reflects the effect of more optimistic beliefs depressing the strong negative influence of past search on subsequent effort. The channel through which $\hat{\phi}$ increases with α_0 is more nuanced: On the one hand, this reflects the fact that when α_0 is higher, individuals enter unemployment with more optimistic beliefs, which in turn delays the rate at which search eventually decays. Higher average search implies that more offers come from search as opposed to arriving stochastically independent of search, which increases the average sensitivity of search to job offers, $\hat{\phi}$. Of course, other parameters that raise average search effort induce a similar effect. $\hat{\phi}$ is disproportionately sensitive to α_0 because higher values of α_0 induce a particularly persistent elevation in search effort among individuals who receive offers, implying that the effect of the first (and subsequent) job offers on search effort increases sharply with α_0.

Finally, the third panel in the second row illustrates that $\hat{\kappa}$ is strongly driven by the degree of curvature in the search cost function. This reflects that a greater degree of curvature in the search cost function precludes complete withdrawal (i.e. zero search) among the long-term unemployed, which in turn allows $\hat{\kappa}$

[41] Indeed, this is why the standard errors in Table 3 are relatively large for β_0 compared with the other parameters.

to increase towards the value found in the data. As discussed in Section 5, the model struggles to match the empirical estimate of $\hat{\kappa}$ because doing so requires relatively high average search effort. See Section 5 for discussion of a simple modification to the model that can generate values of $\hat{\kappa}$ more in line with the data.

References

Arellano, M. and Bond, S. (1991). Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *Review of Economic Studies*, 58.

Barnes, W. F. (1975). Job search models, the duration of unemployment, and the asking wage: Some empirical evidence. *Journal of Human Resources*, 10.

Burdett, K. and Vishwanath, T. (1988). Declining reservation wages and learning. *Review of Economic Studies*, 55.

Christensen, B., Lentz, R., Mortensen, D., Neumann, G., and Werwatz, A. (2005). On-the-job search and the wage distribution. *Journal of Labor Economics*, 23.

Coles, M. and Petrongolo, B. (2008). A test between stock-flow matching and the random matching function approach. *International Economic Review*, 49.

Coles, M. and Smith, E. (1998). Marketplaces and matching. *International Economic Review*, 39.

Devine, T. and Kiefer, N. (1991). *Empirical Labor Economics: The Search Approach.* Oxford University Press.

Ebrahimy, E. and Shimer, R. (2006). Stock-flow matching. Mimeo, University of Chicago.

Faberman, J., Mueller, A., Sahin, A., and Topa, G. (2017). Job search behavior among the employed and non-employed. *NBER Working Paper.*

Falk, A., Huffman, D., and Sunde, U. (2006a). Do i have what it takes? equilibrium search with type uncertainty and non-participation. *IZA Discussion Paper No. 2531.*

Falk, A., Huffman, D., and Sunde, U. (2006b). Self confidence and search. *IZA Discussion Paper No. 2525.*

Feldstein, M. and Poterba, J. (1984). Unemployment insurance and reservation wages. *Journal of Public Economics*, 23.

Hall, R. and Mueller, A. (2018). Wage dispersion and search behavior: The importance of nonwage job values. *Journal of Political Economy*, 126.

Holtz-Eakin, D., Newey, W., and Rosen, H. (1988). Estimating vector autoregressions with panel data. *Econometrica*, 56.

Kreps, D. (1998). Anticipated utility and dynamic choice. In *Frontiers of Research in Economic Theory.* Cambridge University Press.

Krueger, A. and Mueller, A. (2011). Job search, emotional well-being and job finding in a period of mass unemployment: Evidence from high-frequency longitudinal data. *Brookings Papers on Economic Activity*, 42.

Krueger, A. and Mueller, A. (2016). A contribution to the empirics of reservation wages. *American Economic Journal: Economic Policy*, 8.

Lentz, R. (2009). Optimal unemployment insurance in an estimated job search model with savings. *Review of Economic Dynamics*, 12.

Low, H., Meghir, C., and Pistaferri, L. (2010). Wage risk and employment risk over the life cycle. *American Economic Review*, 100.

McCall, J. (1970). Economics of information and job search. *Quarterly Journal of Economics*, 84.

Mueller, A., Spinnewijn, J., and Topa, G. (2020). Job seekers' perceptions and employment prospects: Heterogeneity, duration dependence and bias. *Working Paper*.

Mukoyama, T., Patterson, C., and Sahin, A. (2018). Job search behavior over the business cycle. *American Economic Journal: Macroeconomics*, 10.

Rothschild, M. (1974). Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82.

Shimer, R. (2004). Search intensity. *Unpublished manuscript*.

Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. *Journal of the European Economic Association*, 13.

Windmeijer, F. (2005). A finite sample correction for the variance of linear efficient two-step gmm estimators. *Journal of Econometrics*, 126.