

Duration Dependence and Composition in Unemployment Spells

James D. Eubanks and [David Wiczer](#)

This article reviews the evidence for duration dependence in job-finding rates and its implications for the unemployment duration distribution. The authors document duration dependence and show that it exists within nearly every demographic subgroup. Then, they examine the implications of duration dependence on unemployment duration, emphasizing that a uniform job-finding rate that does not incorporate duration dependence understates unemployment duration. Finally, they explore a composition-based approach to duration dependence, where they solve for the distribution of pre-existing heterogeneity that is consistent with observed duration dependence. The authors look at how this distribution varies cyclically and, in particular, during the Great Recession. The largest changes occur at the low finding-rate tail of this distribution. (JEL J14, E24)

Federal Reserve Bank of St. Louis *Review*, Fourth Quarter 2016, 98(4), pp. 263-76.
<http://dx.doi.org/10.20955/r.2016.263-276>

Labor markets have improved considerably since the Great Recession. However, one indicator of labor market health, the mean duration of unemployment, remains stubbornly high. The mean duration of unemployment measures the number of weeks the average unemployed worker has been without a job. While the unemployment rate has nearly returned to pre-recession levels, the mean duration of unemployment is still more than 1.6 times higher than it was on the eve of the Great Recession. But the mean duration of unemployment does not fully capture how long most unemployed workers should expect to remain unemployed. Empirical evidence has consistently shown that those unemployed for a longer time—a year, for example—are less likely to find a job in a given month than those unemployed for a shorter time—say, a few weeks. As a result, those who do not find a job in the first few months of unemployment are likely to remain unemployed for a long time, and these long-term unemployed workers increase the mean duration of unemployment beyond what most unemployed workers will experience. This relationship between the length of time a worker has been unemployed and the likelihood of finding a job is known as duration dependence.

James D. Eubanks is a senior research associate and David Wiczer is an economist at the Federal Reserve Bank of St. Louis. The authors thank Carlos Carrillo Tudela, Limor Golan, Christopher J. Neely, and Stephen Williamson.

© 2016, Federal Reserve Bank of St. Louis. The views expressed in this article are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

There are two principal explanations for duration dependence. It could be that some workers are always able to find jobs more quickly than others, perhaps because they have more valuable skills or are more efficient at searching for jobs. These workers are unlikely to remain unemployed for long. After a few months, most will find a job, and the workers who remain unemployed are those who find jobs more slowly. Thus, the job-finding rate will be higher at shorter durations of unemployment and lower at longer durations of unemployment. This outcome is referred to as the “composition” effect. It can also be thought of as “population-level” duration dependence: Each individual’s job-finding rate remains fixed, and duration dependence is visible only at the population level.

It could alternatively be that the duration of unemployment has a direct effect on the job-finding rate. For instance, workers’ skills may deteriorate during unemployment or employers may be reluctant to hire workers unemployed for a long time. In the literature, this is often referred to as “true” duration dependence. To contrast this idea with the composition effect, true duration dependence can be thought of as “individual-level” duration dependence: Each individual’s job-finding rate decreases as unemployment duration increases.¹

Either population-level or individual-level duration dependence could explain the relationship between unemployment duration and the job-finding rate, but in this article we focus on composition, or population-level, duration dependence. Given the single spells of unemployment we observe in the data (from the Bureau of Labor Statistics [BLS] Current Population Survey [CPS]), it would be difficult if not impossible to distinguish between the two different effects. A singular focus on composition allows us to isolate some interesting trends and to explore the implications if this channel were acting alone.

To examine the composition effect, we solve for the distribution of job-finding rates that matches the patterns we observe in the data. When we consider the full sample, two points stand out. First, there is always a subset of workers who find jobs very quickly. Second, there are substantial differences in the rates at which the unemployed find jobs: For example, workers with job-finding rates in the 75th percentile find jobs much more quickly than workers with job-finding rates in the 25th percentile.

When we solve for the distribution of job-finding rates over time, interesting cyclical patterns emerge. The dispersion—the difference between the fastest job-finding rates and the slowest job-finding rates—is countercyclical. This pattern was especially pronounced during the Great Recession. Although many unemployed workers were still able to find new jobs relatively quickly during the recession, the workers with the slowest job-finding rates found jobs at even slower rates than before, which drove the large increase in mean unemployment duration. Furthermore, the job-finding rate of the slowest job finders has been slow to recover to its pre-recession level, resulting in a still-high mean unemployment duration even as the labor market has recovered.

Duration dependence is a consistent feature of the U.S. labor market, both in the aggregate and within demographic and occupational subgroups. That is, even among the subgroups of unemployed workers with the highest average job-finding rates—young workers with in-demand skills, for example—the probability of finding a job still declines as unemployment duration increases. We present data on duration dependence for several subgroups and show

that, even after controlling for age, gender, and occupation, the negative relationship between unemployment duration and the likelihood of finding a job still holds.

A vast literature has studied duration dependence. Most relevant to our work, Clark and Summers (1979) first used linked CPS data to show that duration dependence is a pronounced feature of U.S. data. Much of the literature has focused on how to distinguish between population-level (composition) and individual-level (“true”) duration dependence. Several influential papers in the 1980s (see, for example, Elbers and Ridder, 1982; Heckman and Singer, 1984; and Heckman, 1991) established the difficulty of making this distinction. Generally speaking, one must either impose a structure by making assumptions about the initial distribution of job-finding rates and how each individual’s job-finding rate will decline with unemployment duration, or one must observe multiple spells of unemployment and assume that each worker’s job-finding rate is fixed. Recent notable examples of these two approaches are Hornstein (2012), who uses a model’s structure to quantify the sources of duration dependence, and Alvarez, Borovičková, and Shimer (2016), who use Austrian data with multiple unemployment spells for an individual. Concurrent work by Ahn and Hamilton (2016) use an identification strategy similar to ours to find a discrete type-distribution of ex ante heterogeneity, but they use a different estimation method with some added flexibility. Other papers have tried to identify the source of duration dependence by generating it from other, observable mechanisms. For instance, in Wiczer (2015), occupational reallocation creates finding-rate heterogeneity that generates population-level duration dependence, while Jarosch and Pilossoph (2015) use the response rate from audit studies, which submitted résumés that were equivalent except for the unemployment duration of the candidates.

The article proceeds as follows. First, we describe the data and setup and then present evidence for why duration dependence matters in the aggregate. Then, we examine duration dependence within various demographic subgroups. Finally, we examine the job-finding rate distribution consistent with the patterns observed in the data.

DATA SETUP

Throughout this article, our principal source of data is the CPS. This is a monthly survey that follows workers for four months, breaks for eight months, and then follows them again for four months. As such, we can calculate many flow rates by looking at the change in employment status over two consecutive months. Connecting respondents across months is not always perfect, and by definition one-quarter of the samples in the fourth and eighth months cannot be linked. We use the linkages constructed by Flood et al. (2015) that connect individuals across survey months. We begin our analysis in 1994 because, before the 1994 CPS redesign, many of the statistics of interest were measured differently and could potentially distort our estimates.

The job-finding rate conditional on unemployment duration is featured prominently in our analysis. But, given the relatively short panel, we cannot construct an accurate measure of unemployment duration directly from the data. Instead, we use self-reported unemployment duration. This strategy is not without risk, because self-reported duration may not be

consistent with the employment history that we can observe. In particular, we can detect observably short unemployment stints that are reported as long unemployment stints. In about 17 percent of the cases in which we observe the entire unemployment stint, reported unemployment duration is longer than the duration we observe. We drop these cases, although they may have a benign explanation: Employment status is reported in a particular week of the month (generally the second) and a worker may hold a temporary, intermediate job within a longer unemployment stint and not consider that job to be a true transition that restarts the unemployment duration.

As we will see later, “recall” unemployment is potentially a fairly important form of unemployment transition. Fujita and Moscarini (2013) highlighted the prevalence of these periods of unemployment in which the worker returns to the same employer. They are often very short and a conceptually different form of unemployment because the worker expects to be recalled by the same employer in the majority of cases. Using CPS data, we cannot identify recalls. However, evidence from Fujita and Moscarini (2013), who use Survey of Income and Program Participation data, suggests that many of those with very short unemployment stints, and ex ante fast job finders, are experiencing recall unemployment.

To compute monthly transition rates into and out of unemployment, we must correct for well-known time-aggregation bias. Employment status for the month is observed during one particular week, but there may be transitions that occur in intervening weeks that we cannot observe. To correct for this bias, we follow Elsby, Michaels, and Solon (2009), whose method is similar to the differential equations that Shimer (2012) solves. With these calculations, we can create a time series of separation rates $\{s_i\}$ and finding rates $\{f_i^d\}$ at durations d .

MEAN DURATION AND FINDING RATE HETEROGENEITY

Economists often look at the average job-finding rate and the average job-separation rate—the rate at which employed workers become unemployed—to understand the cyclical properties of the labor market. Although these averages can explain fluctuations in the unemployment rate, they do not provide a good indication of how long an unemployed worker should expect to remain unemployed. If all unemployed workers found jobs at the average rate, unemployment duration would be exponentially distributed. But in reality, many workers are very slow to find jobs and their job-finding rates decrease the longer they remain unemployed. The result is a distribution with more long-term unemployed than an exponential distribution would predict. Thus, the mean unemployment duration is higher than what one might predict from the average job-finding and separation rates. To properly interpret the mean unemployment duration, then, one must understand that unemployed workers find jobs at different rates and, hence, there is duration dependence.

To visualize this, consider unemployment duration using only the average job-finding and separation rates. We could compute the mean unemployment duration of workers in the pool of unemployed workers, which is the headline duration statistic published by the BLS.² In this case, the mean unemployment duration conditional on the history of monthly job-finding and separation rates $\{f_j, s_j\}_{j=t_0}^t$ is

Figure 1
Mean Unemployment Duration



NOTE: The uniform finding rate is the mean unemployment duration predicted by the uniform, average job-finding and separation rates in equation (1). Gray bars indicate recessions as determined by the National Bureau of Economic Research (NBER).

SOURCE: BLS CPS and authors' calculations.


$$(1) \quad E \text{ dur}_t | \{f_j, s_j\}_{j=t_0}^t = \frac{\sum_{\tau=t_0}^t (t-\tau) s_\tau \prod_{j=\tau}^t (1-f_j)}{\sum_{\tau=t_0}^t s_\tau \prod_{j=\tau}^t (1-f_j)},$$

where τ indexes the cohort when the worker entered unemployment and t_0 is sufficiently long before t . The resulting unemployment duration is considerably shorter than the duration actually observe in the data. Figure 1 shows the disparity between the data and the unemployment duration predicted by equation (1). This disparity is caused by the different rates at which unemployed workers find jobs: The expected maximum unemployment duration implied by the average finding rate is less than the expected maximum unemployment duration when some workers find jobs at a slower rate than average.³

To get an idea of how different job-finding rates are across workers, Figure 2 presents the ratio of the mean unemployment duration to the median unemployment duration. For reference, we add the red dashed line at $\frac{1}{\log 2}$, the mean-to-median ratio if the distribution

were exponential, which would be the case if all workers found a job at the same rate and the rate of job separations was constant.

If the distribution of unemployment duration were symmetric, the mean and median would be the same. But, as we've discussed, some workers find jobs at much slower rates than others, and these rates decrease the longer the workers remain unemployed. Thus, the unem-

Figure 2**The Mean-to-Median Ratio of Unemployment Duration**

NOTE: A uniform finding rate with constant separations would predict a mean-to-median ratio of $\frac{1}{\log 2}$, the red dashed line. Gray bars indicate recessions as determined by the NBER.

SOURCE: BLS CPS and authors' calculations.


employment duration distribution is not symmetric, but rather is skewed toward longer durations, resulting in a mean duration that is considerably longer than the median.

The skew of the unemployment duration distribution fluctuates over the business cycle. In the beginning of a recession, there is an influx of unemployed workers. Many of these newly unemployed workers will quickly find new jobs, and their short unemployment durations will reduce the skew and the mean unemployment duration. As the recession ends, fewer workers will become unemployed, and the unemployment duration will again skew toward longer durations.

In the Great Recession and its aftermath, the skew of the unemployment duration steadily increased toward longer durations as both the separation rate fell and the average job-finding rate remained low. In later sections, we explore these low finding rates and how they have contributed to the skewness of the distribution.

DURATION DEPENDENCE WITHIN GROUPS

In this section, we estimate how the likelihood of finding a job decreases with unemployment duration. Even after controlling for differences in job-finding rates based on observable characteristics, workers who have been unemployed for a longer period of time still find jobs more slowly than those who have been unemployed for a shorter period of time.

Figure 3**Unemployment Duration Dependence**

NOTE: The figure displays the negative duration dependence observed in the full sample. The finding rate falls as the duration of unemployment increases.


SOURCE: BLS CPS and authors' calculations.

We estimate a simple linear probability model (LPM) of a job-finding indicator on a flexible form for the baseline hazard, or probability of finding a job, and then a quadratic in age and dummies for gender, two-digit Standard Occupational Classification (SOC)-coded occupation, year, and month.^{4.5} We allow the baseline hazard to be a quadratic in both logs and levels, and higher-order terms in either do not change the profile. The LPM regression is

$$\mathbb{I}_{U \rightarrow E} = \beta_1 \log d + \beta_2 (\log d)^2 + \beta_3 d + \beta_4 d^2 + BX,$$

where X is the vector of characteristics listed above. Figure 3 plots this baseline hazard implied by $\beta_1 - \beta_4$, the log job-finding rate as a function of unemployment duration. We present the predicted finding rate as a log transformation because, as suggested by Elsby, Michaels, and Solon (2009), this is more informative than levels when considering flows. We normalize the first period, duration less than a month, to accentuate the change compared with longer durations.

Notice the very rapid decline in the job-finding rate in the first month of unemployment. It continues to fall as duration increases but levels off considerably after approximately three months. We explore this shape in more detail in the next section, but it suggests that some workers find jobs very quickly in the beginning of their unemployment spell and that the remaining workers search considerably longer.

Figure 4**Unemployment Duration Dependence by Demographic Subgroups**

SOURCE: BLS CPS and authors' calculations.

This same pattern holds when looking within demographic subgroups. In Figure 4 we repeat the same exercise. However, instead of controlling for age and gender, we calculate the results for each subgroup separately. We estimate the baseline hazard with occupation, year, and month dummies for the ≤ 35 -year-old and >35 -year-old age groups. Then we pool across ages and estimate the same statistics separately for males and females. For comparison, we also estimate the full sample, pooling all demographic groups but keeping the occupation, year, and month dummies. We normalize each subsample to the finding rate of the full sample in the first month.

The most salient feature of Figure 4 is that the duration dependence of every demographic group has a pattern similar to that of the full sample. As expected, the young workers find jobs far more quickly than the rest of the sample. However, the difference is small within the first month but increases at higher unemployment durations. However, even young workers have a very pronounced and continued decline in the job-finding rate as unemployment duration increases. Females also find jobs more quickly at first, but again show nearly an equivalent decline in job-finding rates as unemployment duration increases.

The discrepancy between the job-finding rate profile in Figure 3, where we controlled for demographic observables, and that for the full sample in Figure 4, where we did not, shows the role of observable differences in explaining duration dependence. The curves are close but distinct: Not controlling for demographics results in a difference of about 0.3 log points in the finding rate after three months. However, the bulk of duration dependence cannot be

accounted for by observable differences. Instead, it is driven by some unobservable differences that occur across groups. We explore the source of these differences in more detail in the next section.

COMPOSITION BEHIND DURATION DEPENDENCE

We now explore the residual differences in job-finding rates, after controlling for observable characteristics, that would deliver the duration dependence seen in Figure 3. As discussed earlier, the decline in the job-finding rate as unemployment duration increases could stem from a direct effect of unemployment duration on the job-finding rate. We call this individual-level duration independence, or “true” duration dependence. The pattern could also be the result of pre-existing differences in job-finding rates: As unemployment duration increases, the fastest job finders find jobs, so that only the slowest job finders remain unemployed for longer durations. We call this population-level duration dependence, or the composition effect. We explore this effect in this section, assuming for purposes of clarity and simplicity that there is no individual-level duration dependence.

To formalize the composition effect, suppose there are several types of workers, j , who each find a job at a different rate, $f(j)$. The job-finding rate of each type is distributed as $G(j)$. Then, the job-finding rate at duration d , f^d , is the average across these groups. The average is affected by the change in the composition of the population of unemployed workers: Some will have already found a job before d , which are those more likely to be the fast-finding types.


For any d , we can observe f^d , so the distribution $G(s)$ must solve⁶

$$(2) \quad f^d = \frac{\int f(j)(1-f(j))^{d-1} dG(j)}{\int (1-f(j))^{d-1} dG(j)} \quad \forall d \in \{0, 1, \dots\}.$$

Because d takes integer values, we can only solve for G as a discrete distribution, where the number of points is less than or equal to the number of months we estimate. In practice, the number of points is also limited because the finding rate does not change much at high durations. We plot this distribution in Figure 5.

Notice that the finding rate is 1 for a non-negligible group of the population. These people enter unemployment certain to exit within the month. This observation may be evidence of recall unemployment, whereby workers enter unemployment and then quickly return to the same positions they left, which was shown to be quite prevalent by Fujita and Moscarini (2013). These very fast job finders are also consistent with a model of stock-flow matching, as in Ebrahimi and Shimer (2010) and Hawkins and Carrillo-Tudela (2016), whereby workers who first enter unemployment have a large stock from which to search but afterward must wait for the flow. If these patterns in the data were truly generated by stock-flow matching, then our initial assumption that there is no individual-level heterogeneity would be misguided. Indeed, even without accepting a particular explanation, the large mass with a finding rate of 1 can also be seen as heuristic evidence against the composition effect as the sole source of duration dependence.

Figure 5
The Initial Distribution of Job-Finding Rates



NOTE: The figure displays the distribution of job-finding rates consistent with the job-finding rate schedule in Figure 3 and equation (2).


SOURCE: BLS CPS and authors' calculations.

At the other end, there is a large group of workers with very low ex ante finding rates. These are the workers who will remain unemployed for a long period and increase the mean unemployment duration well beyond the median. Recall that we have already conditioned on observable differences in the job-finding rate, so these slow job finders do not belong to any particular group that has a lower than average rate, such as older men. This slow job-finding is due to a residual, unobservable characteristic.

While equation (2) was a steady-state relationship, we can also solve for the dynamics of this distribution. This approach allows us to study the distribution of job-finding rates of the workers entering unemployment in any given period. At any time t , we can define the distribution of types $G_t(j)$ who are entering unemployment as

$$(3) \quad f_t^d = \frac{\int f(j)(1-f(j))^{d-1} dG_{t-d}(j)}{\int (1-f(j))^{d-1} dG_{t-d}(j)}.$$

The timing in equation (3) is such that at duration d in period t , the cohort entered unemployment in period $t-d$, so the distribution of their job-finding rates is G_{t-d} . To have a large enough sample to estimate the finding rates for each duration d at time t , we use a rolling window with observations six months before and after period t .

Figure 6**Log Average and Standard Deviation of the Job-Finding Rate**


NOTE: The figure displays the log of the average job-finding rate for cohorts entering unemployment and the standard deviation of log job-finding rates, where the job-finding rate distribution is defined in equation (3). Gray bars indicate recessions as determined by the NBER.

SOURCE: BLS CPS and authors' calculations.

Figure 6 plots the average and the standard deviation of the job-finding rate distribution for cohorts entering unemployment. Notice that the log of the average job-finding rate is fairly steady and changes little during the 2001 recession. In the Great Recession, though, it dives and then recovers very slowly. The decline slightly predates that recession because we are associating the date with the entering cohort, so the job-finding rates of those who became unemployed a few months before the Great Recession will be lower if they are still in the unemployment pool when the recession begins.

The standard deviation of the log job-finding rate is almost perfectly countercyclical. As the average falls, the deviation rises. This occurs because so much of the movement is among the workers with the slowest job-finding rates. The fastest job finders are a fairly constant fraction of the unemployed population, and they almost always find jobs at the same, near-100 percent rate. When the job-finding rates decline among the slowest job finders, the average job-finding rate decreases and the standard deviation increases. For the same reason, the skewness of the distribution also moves almost in lockstep with the standard deviation.

Figure 7 shows that the Great Recession (2007-09) affected workers with the slowest job-finding rates the most. In fact, almost all of the change in the job-finding rates during the Great Recession comes from these workers. These cohorts entering in the Great Recession had the lowest average job-finding rate of any two-year period in our sample, as might be

Figure 7**Distribution of the Job-Finding Rate for the Great Depression and Its Aftermath**

NOTE: The full sample is the job-finding rate since 1994.

SOURCE: BLS CPS and authors' calculations.

expected. However, the cohorts in subsequent periods also had very low job-finding rates. Our data end with the 2013 cohort. Although their average job-finding rate is slightly higher than that for the Great Recession cohort, the finding rate among the slowest job finders is almost exactly the same.

CONCLUSION

Duration dependence, the decrease in the finding rate of the pool of unemployed workers as unemployment duration increases, implies that some unemployed workers will remain unemployed long term. This finding has implications for aggregate duration statistics: The mean of unemployment duration is significantly higher than the median because these workers pull the mean outward. In the absence of duration dependence, unemployed workers would find jobs at a uniform rate. The mean duration predicted by a uniform rate, however, is consistently lower than the observed mean. Thus, to interpret mean unemployment duration, it is crucial to understand how duration dependence shapes the distribution.

The empirical evidence for duration dependence is remarkably consistent. Duration dependence is present across time and countries and within demographic subgroups. In this article, we use CPS data to document duration dependence within age and gender groups. We also show that, even after controlling for observable characteristics, significant residual

variation in job-finding rates remains. This finding reinforces the conclusions of numerous other researchers who find that duration dependence is the norm.

Duration dependence could result from initial differences in finding rates, which we call population-level duration dependence, or the “composition” effect, or from a direct effect of unemployment duration on job-finding rates, which we call individual-level or “true” duration dependence. Either or both sources could explain the duration dependence we observe in the CPS. To distinguish between the two would require a longer panel or structural assumptions. In this article, we focus solely on the composition effect and explore the implications for unemployment duration if this channel were acting alone.

We first solve for the distribution of job-finding rates that is consistent with observed duration dependence and then examine features of this distribution. We find that there are substantial differences in the rate at which unemployed workers find jobs and that a non-negligible subset of workers always returns to employment very quickly.

We also find interesting cyclical patterns. During the Great Recession, the job-finding rate among the fastest job finders remained relatively constant, but the job-finding rate among the slowest job finders declined, resulting in a decrease in the average job-finding rate. For the same reasons, the dispersion of the finding rate—the difference between the fastest and slowest job finders—increased during the Great Recession. ■

NOTES

- ¹ A notable alternative to these explanations is “stock-flow matching” as in Hawkins and Carrillo-Tudela (2016) or Ebrahimi and Shimer (2010), which also generates duration dependence. Though, this alternative may also be viewed as an explanation for individual-level duration dependence.
- ² Notice that this is a slightly different statistic than another reasonable measure of duration: the expected time before a worker finds a job. The BLS reports average duration among the pool of unemployed, which is more applicable even if the data are treated as a repeated cross section, and we follow them.
- ³ The intuition boils down to Jensen’s inequality. In the data, not everyone finds a job at the same rate, f_i . If there is a unit measure of types, indexed by i , then the average finding rate is $f_i = \int f_i^i di$ and the expected maximum duration before finding a job is $\int \frac{1}{f_i^i} di$. But, if we compute expected maximum duration using the average finding rate, we get $\frac{1}{f_i} \leq \int \frac{1}{f_i^i} di$ because of Jensen’s inequality.
- ⁴ We also estimated a more complicated probit or proportional hazards model. The results change very little, as the predicted probabilities are always fairly close to the middle of [0, 1]. We chose the LPM for the sake of simplicity.
- ⁵ See Wiczer (2015) for the list of occupation categories and a description of their construction mapping Census codes to time-consistent ones.
- ⁶ We have been working with monthly job-finding rates, but the continuous time definition may also be intuitive. If \tilde{f} denotes the continuous time job-finding rate, then the equivalent to equation (2) is $\tilde{f}^d = \frac{\int \tilde{f}(j) e^{-\tilde{f}(j)d} dG(j)}{\int e^{-\tilde{f}(j)d} dG(j)}$.

REFERENCES

- Ahn, Hie Joo and Hamilton, John D. "Heterogeneity and Unemployment Dynamics." Working paper, University of California, Davis, September 2016.
- Alvarez, Fernando; Borovičková, Katarína and Shimer, Robert. "Decomposing Duration Dependence in a Stopping Time Model." NBER Working Paper No. w22188, National Bureau of Economic Research, April 2016.
- Clark, Kim B. and Summers, Lawrence H. "Labor Market Dynamics and Unemployment: A Reconsideration." *Brookings Papers on Economic Activity*, 1979, 1(79), pp. 13-60; <http://dx.doi.org/10.2307/2534304>.
- Ebrahimi, Ehsan and Shimer, Robert. "Stock-Flow Matching." *Journal of Economic Theory*, July 2010, 145(4), pp. 1325-53; <http://dx.doi.org/10.1016/j.jet.2010.02.012>.
- Elbers, Chris and Ridder, Geert. "True and Spurious Duration Dependence: The Identifiability of the Proportional Hazard Model." *Review of Economic Studies*, July 1982, 49(3), pp. 403-09; <http://dx.doi.org/10.2307/2297364>.
- Elsby, Michael W.L.; Michaels, Ryan and Solon, Gary. "The Ins and Outs of Cyclical Unemployment." *American Economic Journal: Macroeconomics*, January 2009, 1(1), pp. 84-110; <http://dx.doi.org/10.1257/mac.1.1.84>.
- Flood, Sarah; King, Miriam; Ruggles, Steven and Warren, J. Robert. *Integrated Public Use Microdata Series, Current Population Survey: Version 4.0*. (Machine-readable database.) Minneapolis: University of Minnesota, 2015.
- Fujita, Shigeru and Moscarini, Giuseppe. "Recall and Unemployment." NBER Working Paper No. 19640, National Bureau of Economic Research, 2013.
- Hawkins, William and Carrillo-Tudela, Carlos. "Endogenous Matching Efficiency: A Stock-Flow Approach." Working paper, 2016.
- Heckman, James J. "Identifying the Hand of Past: Distinguishing State Dependence from Heterogeneity." *American Economic Review*, May 1991, 81(2), pp. 75-79; [http://dx.doi.org/10.1016/0304-4076\(84\)90075-7](http://dx.doi.org/10.1016/0304-4076(84)90075-7).
- Heckman, James J. and Singer, Burton. "Econometric Duration Analysis." *Journal of Econometrics*, January-February 1984, 24(1-2), pp. 63-132.
- Hornstein, Andreas. "Accounting for Unemployment: The Long and Short of It." Working Paper No. 12-07, Federal Reserve Bank of Richmond, November 2012; https://www.richmondfed.org/-/media/richmondfedorg/publications/research/working_papers/2012/pdf/wp12-07.pdf.
- Jarosch, Gregor and Pilossoph, Laura. "Statistical Discrimination and Duration Dependence in the Job Finding Rate." Working paper, Federal Reserve Bank of New York, April 2015; <https://www.kansascityfed.org/~media/files/publicat/events/research/cps/pilossoph.pdf>.
- Shimer, Robert. "Reassessing the Ins and Outs of Unemployment." *Review of Economic Dynamics*, April 2012, 15(2), pp. 127-48; <http://dx.doi.org/10.1016/j.red.2012.02.001>.
- Wiczer, David. "Long-Term Unemployment: Attached and Mismatched?" Working Paper No. 2015-42, Federal Reserve Bank of St. Louis, 2015; <https://research.stlouisfed.org/wp/2015/2015-042.pdf>.