

NBER WORKING PAPER SERIES

HETEROGENEITY AND UNEMPLOYMENT DYNAMICS

Hie Joo Ahn
James D. Hamilton

Working Paper 22451
<http://www.nber.org/papers/w22451>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2016

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, nor of any other person associated with the Federal Reserve System, nor of the National Bureau of Economic Research. We thank Stephanie Aaronson, Shigeru Fujita and Ryan Michaels for helpful comments on an earlier draft of this paper.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Hie Joo Ahn and James D. Hamilton. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Heterogeneity and Unemployment Dynamics
Hie Joo Ahn and James D. Hamilton
NBER Working Paper No. 22451
July 2016
JEL No. E24

ABSTRACT

This paper develops new estimates of flows into and out of unemployment that allow for unobserved heterogeneity across workers as well as direct effects of unemployment duration on unemployment-exit probabilities. Unlike any previous paper in this literature, we develop a complete dynamic statistical model that allows us to measure the contribution of different shocks to the short-run, medium-run, and long-run variance of unemployment as well as to specific historical episodes. We find that changes in the inflows of newly unemployed are the key driver of economic recessions and identify an increase in permanent job loss as the most important factor.

Hie Joo Ahn
Board of Governors
of the Federal Reserve System
20th & C Street, NW.
Washington, D.C. 20551
HieJoo.Ahn@frb.gov

James D. Hamilton
Department of Economics, 0508
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0508
and NBER
jhamilton@ucsd.edu

Introduction

What accounts for the sharp spike in the unemployment rate during recessions? The answer traditionally given by macroeconomists was that falling product demand leads firms to lay off workers, with these job separations a key driver of economic downturns. That view has been challenged by Hall (2005) and Shimer (2012), among others, who argued that cyclical fluctuations in the unemployment rate are instead primarily driven by declines in the job-finding rates for unemployed workers. By contrast, Yashiv (2007), Elsby, Michaels and Solon (2009), Fujita and Ramey (2009), and Fujita (2011) concluded that flows into the unemployment pool are as important as or more important than the job-finding rates as cyclical drivers of the unemployment rate.

This debate has become particularly important for understanding the Great Recession and its aftermath. In June 2011—two years into the recovery—the unemployment rate still stood at 9.1%, higher than the peak in any postwar recession other than 1982. Even more troubling, the average duration of those unemployed at that time was 40 weeks, about twice the highest value reached in any month over 1947-2005. Of those workers who had been unemployed for less than one month in June 2011, only 57% were still unemployed the next month. By contrast, of those who had been unemployed for more than 6 months as of June 2011, 93% were still unemployed the following month.¹

This fact that the long-term unemployed find jobs or leave the labor force more slowly than others is a strikingly consistent feature in the postwar data, and could be fundamental for understanding the respective contributions of unemployment inflows and outflows during recessions. For example, workers who lose their jobs due to involuntary permanent separation may have a more difficult time finding new jobs than people who quit voluntarily (Bednarzik, 1983; Fujita and Moscarini, 2013). If more of the separations during a recession are involuntary, it could show up as what other researchers have interpreted as a fall in the job-finding rate and increase in the duration of unemployment even if the key driver of the recession was the increase in involuntary separations.

¹The values for f_t^1 and $f_t^{7,+}$ were calculated from

$$f_t^1 = \frac{U_t^1 - U_{t+1}^1}{U_t^1}, \quad f_t^{7,+} = \frac{U_t^{7,+} - (U_{t+1}^{7,+} - U_{t+1}^7)}{U_t^{7,+}}$$

for U_t^n the number unemployed with duration n months at t . The reported series are seasonally adjusted with X-12-ARIMA.

The phenomenon that unemployment exit rates fall with the duration of unemployment has been widely studied, with explanations falling into two broad categories. One possibility is that the experience of being unemployed for a longer period of time directly changes the employment probability for a fixed individual. Following van den Berg and van Ours (1996) we will refer to this possibility as "genuine duration dependence". Individuals lose human capital the longer they are unemployed (Acemoglu, 1995; Ljungqvist and Sargent, 1998), employers may statistically discriminate against those who have been unemployed for longer (Eriksson and Rooth, 2014; Kroft, Lange, and Notowidigdo, 2013)², and individuals may search less the longer they have been unemployed (Faberman and Kudlyak, 2014). We will refer to such negative genuine duration dependence, that is, a condition where a longer period spent in unemployment directly reduces the probability of finding a job, as "unemployment scarring." Another possibility is positive genuine duration dependence. For example, the longer somebody has been unemployed, the more willing that person may be to accept a low-paying job or simply to drop out of the labor force. Meyer (1990) and Katz and Meyer (1990a,b) argued that such effects may become important as unemployment benefits become exhausted. We will refer to the possibility that the probability of exiting unemployment increases as a consequence of a longer duration of unemployment as "motivational" effects.

A quite different explanation for the differences in unemployment exit probabilities across the different duration categories is that there are important differences across job-seekers from the very beginning, arising for example from differences in the reason the individuals left their previous job or in differences in ex ante abilities or motivation across workers. The longer an individual is observed to have been unemployed, the greater the chance that the individual is a member of a group whose unemployment exit probabilities were low to begin with. That such cross-sectional heterogeneity might be important for the question studied by Hall and Shimer was recognized as far back as Darby, Haltiwanger, and Plant (1986), who argued that heterogeneity accounted for falling job-finding rates during recessions in a manner consistent with the traditional macroeconomic interpretation of recessions. A number of researchers have tried to investigate this hypothesis by looking at differences across job seekers in observable characteristics such as demographics, education, industry, occupation, geographical region, and reason for unemployment. Baker (1992),

²Jarosch and Pilossoph (2015) demonstrated that the quantitative magnitude of statistical discrimination found in these studies could in fact be consistent with the claim that cross-section heterogeneity is the primary explanation for the observed tendency of unemployment-continuation probabilities to rise with duration of unemployment.

Shimer (2012), and Kroft, Lange, Notowidigdo, and Katz (forthcoming) found that such variables contributed little to variation over time in long-term unemployment rates, while Aaronson, Mazumder and Schechter (2010), Bachmann and Sinning (2012), Barnichon and Figura (2015), Hall (2014), and Hall and Schulhofer-Wohl (2015) documented important differences across observable characteristics. Elsby, Michaels and Solon (2009) found that incorporating observable heterogeneity reduced the imputed role of cyclical variation in unemployment exit rates.

However, no two individuals with the same coarse observable characteristics are in fact identical. It seems undeniable that a given pool of unemployed individuals that conditions on any set of observed characteristics is likely to become increasingly represented by those with lower ex ante exit probabilities the longer the period of time for which the individuals have been unemployed. Most of the above studies assume that conditional on observable characteristics, unemployed individuals are identical in terms of their transition probabilities into and out of unemployment. The result is that the imputed exit probabilities are determined solely from the most recent labor force statistics as if every month was a new steady state of the economy, not taking into account the fact that each individual has a unique history of unemployment. This approach misses a key feature of economic recessions and unemployment dynamics. Once one acknowledges heterogeneity across workers, the pool of those looking for work at a given point in time— and therefore the exit rates for individuals in that group— depends on the specific history of conditions whereby those individuals came to be unemployed. This means that more information than the current month’s labor force statistics is necessary to account for the different histories of unemployed individuals and thus to credibly analyze the contributions of the inflows and outflows.

A large literature has explored methods to distinguish genuine duration dependence from unobserved cross-sectional heterogeneity.³ A common resolution has been to assume that there is no variation over time in unobserved heterogeneity, in which case identification can be achieved by observing repeated spells of unemployment for a given individual (Honoré,1993). We will demonstrate in Section 1 that an alternative approach is to assume limited variation over time in genuine duration dependence and base inference on the observed panel of aggregate unemployment by duration categories. This approach allows us to study the potential consequences of cyclical variation

³See for example Elbers and Ridder (1982), Heckman and Singer (1984a,b,c), Ridder (1990), Honoré (1993), and van den Berg (2001).

in the matches between unobserved worker characteristics and available jobs.

Van den Berg and van Ours (1996) used a related idea, assuming proportional hazards along with time-invariant genuine duration dependence. In this paper we generalize the approach to let genuine duration dependence change over time with eligibility for unemployment insurance. Most importantly, we will allow for a key role for cyclical variation in worker heterogeneity, something that has never previously been done in either the large microeconomic or macroeconomic literatures on unemployment duration.

Our approach is most similar to that in Hornstein (2012), who used dynamic accounting identities to interpret aggregate panel dynamics in a similar way to that in our paper. However, Hornstein’s model was unidentified— in terms of the discussion of identification in Section 1, his model has 5 unknowns and only 4 equations. As a result, his specification did not allow him to calculate the likelihood function for the observed data or forecasts of unemployment or duration. By contrast, our model generates values for all these along with the optimal statistical inference about the various shocks driving the observed dynamics of unemployment.

Ours is the only paper in the large macro literature on the ins and outs of unemployment that offers a full dynamic description of changes in aggregate unemployment by duration categories. In doing so we resolve a key shortcoming in much of the previous literature. Most previous studies used correlations between unemployment and the steady-state unemployment rate predicted by either inflows or outflows to draw conclusions about how much of the variation in unemployment is due to each factor. However, the unemployment rate is highly serially correlated and possibly nonstationary. What do we even mean by its variance, and how do we distinguish between the contribution to this variance of short-term versus long-term influences? Previous studies often addressed these issues by using some kind of detrending procedures. By contrast, our paper develops a complete statistical model with nonstationary driving processes, which as a by-product generates a forecast of unemployment at any horizon in the future. Since the forecast error at any specified horizon has a stationary distribution and well defined mean squared error whether or not the underlying process is nonstationary, as in den Haan (2000) we can calculate the fraction of the variance in unanticipated changes in unemployment over any horizon that is attributable to the various shocks in the model. This allows us to measure the dynamic contributions of different factors to unemployment and allows us to make very clear statements about the importance for

short-run, medium-run, and long-run dynamics as well as over specific historical episodes. This is one of the key innovations of our approach and is entirely new to this literature.

In Section 1 we introduce the data that we will use in this analysis based on the number of job-seekers each month who report they have been looking for work at various search durations. We describe the accounting identities that will later be used in our full dynamic model and use average values of observable variables over the sample to explain the intuition for how such duration data can be used to separately identify cross-sectional heterogeneity and genuine duration dependence. We also use these calculations to illustrate why cross-sectional heterogeneity appears to be more important than genuine duration dependence in terms of explaining the broad features of these data.

In Section 2 we extend this framework into a full dynamic model in which we postulate the existence of two types of workers at any given date. Type H workers have a higher ex ante probability of exiting unemployment than type L workers, and all workers are also subject to potential scarring or motivational effects. Our model postulates that the number of newly unemployed individuals of either type, as well as the probability for each type of exiting the pool of unemployed at each date, evolve over time according to unobserved random walks. We show how one can calculate the likelihood function for the observed unemployment data and an inference about each of the state variables at every date in the sample using an extended Kalman filter.

Empirical results are reported in Section 3. Broken down in terms of inflows versus outflows, we find that variation over time in the inflows of the newly unemployed are more important than outflows from unemployment in accounting for errors in predicting aggregate unemployment at all horizons. Broken down in terms of types of workers, inflow and outflow probabilities for type L workers are more important than those for type H workers, and account for 90% of the uncertainty in predicting unemployment 2 years ahead. In recessions since 1990, shocks to the inflows of type L workers were the most important cause of rising unemployment during the recession. We find a non-monotonic contribution of genuine duration dependence, with scarring effects dominating up to 1 year but motivational effects apparent for those unemployed longer than a year.

We offer interpretations of our findings in Section 4 by relating our estimated series to those available from other sources. A key difference between type L and type H workers is the circumstances under which they left their previous job. Our imputed series for newly unemployed

type L workers has very similar cyclical dynamics to separate measures of the number of new job-seekers who were involuntarily separated from their previous job for a reason other than what was described as a temporary layoff. Notwithstanding, our estimates imply that type L workers are a strict subset of those involuntarily separated, but also represent subsets of re-entrants to the labor force and other individuals. We conclude that, consistent with the traditional interpretation of business cycles, the key reason that unemployment spikes during recessions is a change in the circumstances under which individuals lose their jobs.

In Section 5 we investigate the robustness of our approach to various alternative specifications, including alternative methods to account for the change in the CPS questionnaire in 1994, allowing for correlation between the innovations of the underlying structural shocks in our model, and the possible effects of time aggregation. While such factors could produce changes in some of the details of our inference, our overall conclusions (summarized in Section 6) appear to be quite robust.

1 Observable implications of unobserved heterogeneity

The purpose of this section is to use steady-state calculations to show how unobserved heterogeneity and genuine duration dependence can be separately identified and to provide the intuition behind some of the results that will be found in Section 3 using our full dynamic model.

The Bureau of Labor Statistics reports for each month t the number of Americans who have been unemployed for less than 5 weeks. Our baseline model is specified at the monthly frequency, leading us to use the notation U_t^1 for the above BLS-reported magnitude, indicating these individuals have been unemployed for 1 month or less as of month t . BLS also reports the number who have been unemployed for between 5 and 14 weeks (or 2-3 months, denoted $U_t^{2.3}$), 15-26 weeks ($U_t^{4.6}$) and longer than 26 weeks ($U_t^{7.+}$). One reason the BLS reports the data in terms of these aggregates is to try to minimize the role of measurement error by averaging within broad groups, an approach that we will also follow in our paper.⁴ Although our theoretical calculations will keep track of durations by individual months, our statistical analysis is all based on the implications for observable broad aggregates. Notwithstanding, when reporting on long-term unemployment, many

⁴In January 2011 the BLS changed the maximum allowable unemployment duration response from 2 years to 5 years. Although this affected the BLS's own estimate of average duration of unemployment, it did not change the total numbers unemployed by the duration categories we use. This is another reason to favor our approach, which relies only on aggregated data.

BLS publications⁵ further break down the $U_t^{7,+}$ category into those unemployed with duration 7-12 months ($U_t^{7,12}$) and those with duration longer than 1 year ($U_t^{13,+}$). Since long-term unemployment is also a major interest in our investigation, we have used the raw CPS micro data from which the usual publicly reported aggregates are constructed to create these last two series for our study.⁶

The five series used in our analysis are graphed in Figure 1, with average values over the full sample reported in the first row of Table 1. Our purpose in this paper is to explore what variation in these duration-specific counts across time can tell us about unemployment dynamics. Our focus will be on the following question— of those individuals who are newly unemployed at time t , what fraction will still be unemployed at time $t + k$? We presume that the answer to this question depends not just on aggregate economic conditions over the interval $(t, t + k)$ but also on the particular characteristics of those individuals. Let w_{it} denote the number of people of type i who are newly unemployed at time t , where we interpret

$$U_t^1 = \sum_{i=1}^I w_{it}. \quad (1)$$

We define $P_{it}(k)$ as the fraction of individuals of type i who were unemployed for one month or less as of date $t - k$ and are still unemployed and looking for work at t . Thus the total number of individuals who have been unemployed for exactly $k + 1$ months at time t is given by

$$U_t^{k+1} = \sum_{i=1}^I w_{i,t-k} P_{it}(k). \quad (2)$$

Heckman and Singer (1984b, Theorem 3.1) demonstrated that when observed duration data are discrete-valued, a discrete number (as opposed to a continuum) of types is fully general. Ham and Rea (1987), Van den Berg and van Ours (1996), and Van den Berg and van der Klaauw (2001) tested for the number of types I and found $I = 2$ provided the best description of the data sets they analyzed, albeit with different formulations of unobserved heterogeneity than we have proposed in (2). Since we aggregate durations into 5 discrete categories, it is only feasible to allow for $I = 2$ types using our data and method. We will designate the two groups as type H and type L in anticipation of the normalization that type L workers have a lower probability of

⁵See for example Bureau of Labor Statistics (2011) and Ilg and Theodossiou (2012).

⁶See Appendix A for further details of data construction.

exiting unemployment. We intend to demonstrate that this offers a simple but compelling way of interpreting the observed data.

We first examine what we could infer about unobserved types based only on the historical average values $\bar{U}^1, \bar{U}^{2.3}, \bar{U}^{4.6}, \bar{U}^{7.12}$, and $\bar{U}^{13.+}$, and then will consider what additional information can be learned from variation over time in these five variables.

1.1 Inference using historical average values alone.

Suppose for purposes of this section only that the number of newly unemployed individuals of each type remained constant over time at values w_L and w_H , respectively, and also that the probabilities that individuals of each type remain unemployed in any given month are constants p_L and p_H . Then (2) would simplify to

$$U^{k+1} = w_L p_L^k + w_H p_H^k. \quad (3)$$

This equation describes the average number of individuals who have been unemployed for $k + 1$ months as the sum of two different functions of k , with each of the two functions being fully characterized by two parameters (w_i and p_i). The solid red curve in Panel A of Figure 2 plots the first function ($w_L p_L^k$), while the dotted blue curve plots the sum. Given observed values of the sum U^n for any four different values of n , we could estimate the four parameters (w_L, w_H, p_L, p_H) to exactly match those four observed numbers, as in Panel A of Figure 2.

As noted above, we regard aggregate measures like $U_t^{4.6}$ as more reliable than a specific estimate such as U_t^5 that could be constructed from CPS micro data. But exactly the same kind of parameter fitting can be done using aggregates like $\bar{U}^{4.6}$. For example,

$$\bar{U}^{2.3} = \bar{U}^2 + \bar{U}^3 = (w_L p_L + w_H p_H) + (w_L p_L^2 + w_H p_H^2).$$

The 4 observed values ($\bar{U}^1, \bar{U}^{2.3}, \bar{U}^{4.6}, \bar{U}^{7.12}$) are sufficient to calculate the four unknowns (w_L, w_H, p_L, p_H) as we illustrate in row 2 of Table 1.⁷ These estimates imply that type H individuals comprise a very high fraction, 78.8%, of the initial pool of unemployed U^1 . But the unemployment-continuation

⁷Specifically, the four functions are obtained from equations (7)-(10) below for the special case when the left-hand variables represent historical averages and on the right-hand side we set $w_{it} = w_i$, $P_{it}(k) = p_i^k$, and $r_i^x = 0$.

probability for type H individuals ($p_H = 0.36$) is much lower than for type L ($p_L = 0.85$). Because the type H are likely to find jobs relatively quickly, there are very few type H individuals included in U^n for durations n beyond 4 months, as seen in Panel A of Figure 2. The key feature of the observed data (represented by the black dots in Figure 2) that gives rise to this conclusion is the fact that the numbers initially drop off very quickly (as most of the type H workers find jobs), but after that much more slowly (as the remaining type L workers continue searching).

Although we did not use the fifth data point, $\bar{U}^{13,+}$, in estimating these parameters, the framework generates a prediction for what that observation would be.⁸ This is reported in the last entry of row 2 of Table 1 to be 614,000 which is quite close to the observed value of 636,000. The feature of the data that produced this result is that the observed numbers fall off at close to a constant exponential rate once we get beyond 4 months, as the simple mixture model would predict.

Alternatively, we could equally well describe the observed averages using a model in which there is only genuine duration dependence (GDD). Suppose that an individual who has been unemployed for τ months has a probability $p(\tau)$ of still being unemployed the following month. We can always write this in the form

$$p(\tau) = \exp(-\exp(d_\tau))$$

for d_τ an arbitrary function of τ , with double-exponentiation a convenient device for ensuring that probabilities are always positive. For example, we could fit the 5 observations in the first row of Table 1 perfectly if we used a 4-parameter representation for d_τ such as⁹

$$d_\tau = \delta_0 + \delta_1\tau + \delta_2\tau^2 + \delta_3\tau^3. \quad (4)$$

A large number of empirical studies have assumed Weibull durations, essentially corresponding to $\delta_2 = \delta_3 = 0$. The values for δ_j that would exactly fit the historical averages are reported in row 3 of Table 1 and the implied function $p(\tau)$ is plotted in panel A of Figure 3. Note that in contrast to the popular Weibull assumption and most theoretical models, the nature of GDD necessary to fit the observed data would have to be nonmonotonic.

⁸Following Hornstein (2012) we truncate all calculations at 48 months in equation (11). Most of the models considered in this paper imply essentially zero probability of an unemployment spell exceeding 4 years in duration.

⁹Specifically, we calculate $U^{k+1} = wp(1)p(2)\cdots p(k)$ and find the values of $w, \delta_0, \delta_1, \delta_2, \delta_3$ to match the observed values in row 1 of Table 1.

If we were willing to restrict the functional form of GDD to the Weibull case, we could also interpret the historical averages as resulting from a combination of unobserved heterogeneity and GDD. Suppose that the probability that an individual of type i who has been unemployed for τ months will still be unemployed the following month is given by

$$p_i(\tau) = \exp\{-\exp[x_i + d_\tau]\} \quad (5)$$

with implied unemployment counts

$$U^{k+1} = \sum_{i=L,H} w_i p_i(1) p_i(2) \cdots p_i(k). \quad (6)$$

The value of x_i for $i = H, L$ reflects cross-sectional heterogeneity in unemployment-continuation probabilities and d_τ captures genuine duration dependence. As noted by Katz and Meyer (1990, p. 992), this double-exponential functional form is a convenient way to implement a proportional hazards specification so as to guarantee a positive hazard¹⁰, a feature that will be very helpful for the generalization in the following section in which we will allow for variation of x_{it} over time. Suppose we were willing to model GDD using a one-parameter function, say $d_\tau = \delta(\tau - 1)$, where a negative value for the scalar δ could represent unemployment scarring. Then we could find a value for the five parameters $w_L, w_H, x_L, x_H, \delta$ so as to fit the 5 time-series averages $\bar{U}^1, \bar{U}^{2.3}, \bar{U}^{4.6}, \bar{U}^{7.12}$, and $\bar{U}^{13.+}$ exactly. These values are reported in row 4 of Table 1. The implied value for δ is close to zero, and the other parameters are close to those for the pure cross-sectional heterogeneity specification of row 2. Thus for this particular parametric example, we would conclude that cross-sectional heterogeneity is much more important than genuine duration dependence in accounting for why observed unemployment-continuation probabilities rise with duration of unemployment. The feature of the data that gave rise to this conclusion is that the 4-parameter pure heterogeneity model gives a very good prediction of all five observations.

¹⁰Consider an individual i who has been unemployed for τ months as of the beginning of month t and let the hazard within month t be $\lambda_{i,t,\tau} = \exp(x_{it}) \exp(d_\tau)$ where the exponentiation is a device to guarantee that the hazard is positive for any x_{it} and d_τ . The meaning of the hazard is that if we divide month t into n subintervals, the probability that individual i exits unemployment in the interval $(s, s + 1/n)$ is $\lambda_{i,t,\tau}/n + o(1/n)$ from which the probability that the individual is still unemployed at the beginning of month $t + 1$ is

$$\lim_{n \rightarrow \infty} [1 - \lambda_{i,t,\tau}/n + o(1/n)]^n = \exp(-\lambda_{i,t,\tau}) = \exp[-\exp(x_{it}) \exp(d_\tau)].$$

1.2 Inference using changes over time.

Next consider what we can discover from using time-series variation in the observed aggregates. Suppose we repeat the above exercises using data only since the Great Recession. Row 6 of Table 1 and Panel B of Figure 2 show the results if we tried to explain these more recent averages entirely in terms of unobserved heterogeneity. The implied value for the unemployment-continuation probability for type L individuals, $p_L = 0.89$, is only slightly higher than the value 0.85 fit to the full historical sample. The reason is that the function \bar{U}^n drops off after $n = 4$ months at only a slightly slower rate than it did historically. However, we would infer that the inflow of new type L individuals, $w_L = 1,065$ is much higher than the historical average value of 679, in order to account for the fact that \bar{U}^n is now dropping off after 4 months from a much higher base. We again find that the 4-parameter model does a reasonable job of anticipating the fifth unused data point.

If we instead tried to explain the recent averages purely in terms of GDD, we would use the parameter values from row 7 of Table 1. These again could fit the data perfectly, albeit with continuation probabilities for which individuals become permanently stuck in unemployment after 2 years¹¹ and a function with odd oscillations (see panel B of Figure 3). Although it is mathematically possible to describe the data with this equation, it would be difficult to motivate a theory of why GDD should have changed shape in this way. It requires for example a steeper initial slope to the curve in panel B of Figure 3 when economic conditions worsened, corresponding to the claim that the scarring associated with unemployment is more severe during a recession. But this is directly contradicted by the experimental finding of Kroft, Lange, and Notowidigdo (2013) that potential employers pay less attention to applicants' duration of unemployment when the labor market is weaker. We will produce additional evidence in Section 4 below on predictability of changes in unemployment that would also be very hard to interpret based on any theory of cyclically changing GDD.

These concerns notwithstanding, would it be possible to allow for both an unrestricted non-monotonic functional form for GDD as well as unobserved heterogeneity? The answer is definitely yes once we take account of changes over time. Suppose for example we were to pool the observations from the first row of Table 1 (the full-sample averages) together with those in row 5 (behavior

¹¹In this case the calculations are of course fundamentally influenced by our convention of truncating unemployment spells at 48 months.

since the Great Recession), giving us a total of 10 observations. If we took the view that the unobserved heterogeneity parameters may have changed over the cycle but that the GDD function d_τ in (5) is time-invariant, we would then be able to generalize d_τ to be a function of τ determined by two parameters, say δ_1 and δ_2 , and use the ten observations to infer ten unknowns (values of w_H, w_L, x_H, x_L for the two subsamples along with the parameters δ_1 and δ_2). Generalizing a little further, if we use observations across 4 different subsamples we could infer values of w_H, w_L, x_H, x_L for each subsample along with a completely unrestricted nonmonotonic GDD function as in (4). In fact, if we were able to use all five observations on $U_t^1, U_t^{2.3}, U_t^{4.6}, U_t^{7.12}, U_t^{13.+}$ for every date t , we could even allow for some modest variation over time in the GDD function $d_{\tau t}$, and indeed such a specification will be included in the general results to be reported below.

Although our approach to identification based on restricting time variation in GDD has been used in some studies of micro labor data such as van den Berg and van Ours (1996), a more common assumption in that literature has been to assume that unobserved heterogeneity is time invariant, with identification dependent on observation of repeated spells of unemployment by the same individual (see for example Alvarez, Borovičková, and Shimer, 2015). Our paper differs from any previous study in either the micro or macro labor literature in focusing on aggregate cyclical variation in unobserved heterogeneity. Documenting its importance for unemployment fluctuations and examining the causes behind it is one of the key original contributions of our paper.

The actual method that we will use for this analysis is in fact far superior to the simple steady-state calculations that we have employed in this section. It takes time for new inflows to start to matter for longer-term unemployment, and underlying labor market conditions are constantly changing. Our dynamic model fully takes into account the implications of the past history for current observed values. We will demonstrate in the next section that if we assume that the values of $w_{Ht}, w_{Lt}, x_{Ht}, x_{Lt}$ evolve gradually over time, we can use observations of $U_t^1, U_t^{2.3}, U_t^{4.6}, U_t^{7.12}, U_t^{13.+}$ in a nonlinear state-space model to form an inference about the changing values of $w_{Ht}, w_{Lt}, x_{Ht}, x_{Lt}$. This approach takes as a starting point the guess that this period's values for these 4 magnitudes are the same values they were last month, and uses differences between the observed values for $U_t^1, U_t^{2.3}, U_t^{4.6}, U_t^{7.12}, U_t^{13.+}$ relative to what the model would have predicted to infer how the 4 magnitudes likely changed from the previous month. Such a procedure can also allow for transient measurement error in each of the 5 observed variables, with an optimal

inference about the unobserved $(w_{Ht}, w_{Lt}, x_{Ht}, x_{Lt})$ feasible from the assumption that any change in observed variables that persists for more than 1 month should be attributed to changes in the underlying $(w_{Ht}, w_{Lt}, x_{Ht}, x_{Lt})$ rather than measurement error.

We have used the time-invariant steady-state calculations in this section primarily to explain the intuition for where the identification is coming from. Nevertheless, it turns out that the key conclusions of the above steady-state calculations— that the majority of newly unemployed individuals can be described as type H who find jobs quickly, that dynamic sorting based on unobserved heterogeneity appears to be much more important than genuine duration dependence in explaining why a longer-term unemployed individual is less likely to exit unemployment, and that the key driver of economic recessions is an increased inflow of newly unemployed type L individuals— will also turn out to characterize what we will find as we now turn to a richer dynamic model.

2 Dynamic formulation

We now consider a state-space model where the dynamic behavior of the observed vector $y_t = (U_t^1, U_t^{2.3}, U_t^{4.6}, U_t^{7.12}, U_t^{13.+})'$ is determined as a nonlinear function of latent dynamic variables— the inflows and outflow probabilities for unemployed individuals with unobserved heterogeneity. Due to the nonlinear nature of the resulting model, we draw inference on the latent variables using the extended Kalman filter.

2.1 State-space representation

We assume smooth variation over time for the latent variables of interest, $w_{Ht}, w_{Lt}, x_{Ht}, x_{Lt}$, with each assumed to follow an unobserved random walk, e.g.,

$$w_{Ht} = w_{H,t-1} + \varepsilon_{Ht}^w.$$

A random walk is by far the most common assumption in dynamic latent variable models as it has proven to be a flexible and parsimonious way to adapt inference to a variety of sources of changing conditions or possible structural breaks.¹² It has the intuitive appeal that we enter period t with the expectation that conditions are similar to those in the previous period, but are prepared

¹²See for example Baumeister and Peersman (2013).

to change that inference on the basis of observed changes in the variables. It is an unambiguous improvement over the steady-state calculations described in the previous section (and invoked in the majority of previous studies in this literature). Note moreover that the random walk specification includes the steady-state formulation as a special case when the variance of ε_{Ht}^w is zero. We have also experimented with a model in which we assume AR(1) dynamics for the latent variables with autoregressive coefficients estimated by maximum likelihood. We found the coefficient estimates to be very close to unity and the resulting inference very similar to those reported for our baseline random walk specification.

Another key detail of our approach is that we allow for the possibility that unemployment counts are all contaminated by error. The durations in CPS are in part self reported and respondents make a variety of errors. We assume that each element of y_t has an associated measurement error $r_t = (r_t^1, r_t^{2.3}, r_t^{4.6}, r_t^{7.12}, r_t^{13.+})'$. Our identification assumption is that the measurement error is white noise, meaning that the inference is only adjusted for changes in the observed variables that prove to be persistent. The observation equations can then be written as follows¹³,

$$U_t^1 = \sum_{i=H,L} w_{it} + r_t^1 \quad (7)$$

$$U_t^{2.3} = \sum_{i=H,L} [w_{i,t-1}P_{it}(1) + w_{i,t-2}P_{it}(2)] + r_t^{2.3} \quad (8)$$

$$U_t^{4.6} = \sum_{i=H,L} \sum_{k=3}^5 [w_{i,t-k}P_{it}(k)] + r_t^{4.6} \quad (9)$$

$$U_t^{7.12} = \sum_{i=H,L} \sum_{k=6}^{11} [w_{i,t-k}P_{it}(k)] + r_t^{7.12} \quad (10)$$

$$U_t^{13.+} = \sum_{i=H,L} \sum_{k=12}^{47} [w_{i,t-k}P_{it}(k)] + r_t^{13.+} \quad (11)$$

where

$$P_{it}(j) = p_{i,t-j+1}(1)p_{i,t-j+2}(2)\dots p_{it}(j). \quad (12)$$

We assume that for type i workers who have already been unemployed for τ months as of time $t - 1$, the fraction who will still be unemployed at t is given by

¹³ As in the steady-state example in Section 1, we consider 4 years to be the maximum unemployment duration considered.

$$p_{it}(\tau) = \exp[-\exp(x_{it} + d_{t\tau})] \quad \text{for } \tau = 1, 2, 3, \dots \quad (13)$$

where $d_{t\tau}$ determines the nature of genuine duration dependence experienced by an unemployed individual with duration of unemployment τ months and x_{it} is a time-varying magnitude influencing the unemployment exit probability for all workers of type i regardless of their duration. Like the inflows w_{LT} and w_{Ht} , we assume that the parameters x_{Lt} and x_{Ht} governing outflow probabilities also follow a random walk. If we assume that the genuine-duration-dependence effects as summarized by $d_{t\tau}$ are time invariant, it is possible to estimate a different value for the parameter d_τ for each τ . We investigated a number of different specifications for d_τ and found the best fit using linear splines at $\tau = 6$ and $\tau = 12$ which we use for the baseline analysis:

$$d_\tau = \begin{cases} \delta_1(\tau - 1) & \text{for } \tau < 6 \\ \delta_1[(6 - 1) - 1] + \delta_2[\tau - (6 - 1)] & \text{for } 6 \leq \tau < 12 \\ \delta_1[(6 - 1) - 1] + \delta_2[(12 - 1) - (6 - 1)] + \delta_3[\tau - (12 - 1)] & \text{for } 12 \leq \tau. \end{cases} \quad (14)$$

Positive δ_j for $j = 1, 2, 3$ imply motivational effects while negative values imply unemployment scarring over the relevant duration ranges.

We can arrive at the likelihood function for the observed data $\{y_1, \dots, y_T\}$ by assuming that the measurement errors are independent Normal,¹⁴ where $R_1, R_{2.3}, R_{4.6}, R_{7.12}$ and $R_{13.+}$ are the standard deviations of $r_t^1, r_t^{2.3}, r_t^{4.6}, r_t^{7.12}$ and $r_t^{13.+}$ respectively:

$$r_t \sim N(0, R)$$

¹⁴The Normality assumption of measurement errors has often been adopted in the literature of unemployment hazards; see for example Abbring, van den Berg and van Ours (2001) and van den Berg and van der Klaauw (2001). Moreover, the identical Kalman filter equations that emerge from an assumption of Normality can also be motivated using a least-squares criterion; see for example Hamilton (1994a, Chapter 13).

$$\underbrace{R}_{5 \times 5} = \begin{bmatrix} R_1^2 & 0 & 0 & 0 & 0 \\ 0 & R_{2.3}^2 & 0 & 0 & 0 \\ 0 & 0 & R_{4.6}^2 & 0 & 0 \\ 0 & 0 & 0 & R_{7.12}^2 & 0 \\ 0 & 0 & 0 & 0 & R_{13,+}^2 \end{bmatrix}.$$

Let ξ_t be the vector $(w_{Lt}, w_{Ht}, x_{Lt}, x_{Ht})'$ and $\varepsilon_t = (\varepsilon_{Lt}^w, \varepsilon_{Ht}^w, \varepsilon_{Lt}^x, \varepsilon_{Ht}^x)'$. Our assumption that the latent factors evolve as random walks would be written as

$$\underbrace{\xi_t}_{4 \times 1} = \xi_{t-1} + \underbrace{\varepsilon_t}_{4 \times 1} \quad (15)$$

$$\underbrace{\varepsilon_t}_{4 \times 1} \sim N(\underbrace{0}_{4 \times 1}, \underbrace{\Sigma}_{4 \times 4})$$

$$\underbrace{\Sigma}_{4 \times 4} = \begin{bmatrix} (\sigma_L^w)^2 & 0 & 0 & 0 \\ 0 & (\sigma_H^w)^2 & 0 & 0 \\ 0 & 0 & (\sigma_L^x)^2 & 0 \\ 0 & 0 & 0 & (\sigma_H^x)^2 \end{bmatrix}.$$

In Section 5 we will also report results for a specification in which the shocks are allowed to be contemporaneously correlated.

Since the measurement equations (7)-(11) are a function of $\{\xi_t, \xi_{t-1}, \dots, \xi_{t-47}\}$, the state equation should describe the joint distribution of ξ_t 's from $t - 47$ to t , where I and 0 denote a (4×4) identity and zero matrix, respectively:

$$\underbrace{\begin{bmatrix} \xi_t \\ \xi_{t-1} \\ \xi_{t-2} \\ \vdots \\ \xi_{t-46} \\ \xi_{t-47} \end{bmatrix}}_{192 \times 1} = \underbrace{\begin{bmatrix} \underbrace{I}_{4 \times 4} & \underbrace{0}_{4 \times 4} & 0 & 0 & \dots & 0 & 0 & 0 \\ I & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ 0 & I & 0 & 0 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \dots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & I & 0 & 0 \\ 0 & 0 & 0 & 0 & \dots & 0 & I & 0 \end{bmatrix}}_{192 \times 192} \underbrace{\begin{bmatrix} \xi_{t-1} \\ \xi_{t-2} \\ \xi_{t-3} \\ \vdots \\ \xi_{t-47} \\ \xi_{t-48} \end{bmatrix}}_{192 \times 1} + \underbrace{\begin{bmatrix} \underbrace{\varepsilon_t}_{4 \times 1} \\ \underbrace{0}_{4 \times 1} \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}}_{192 \times 1}. \quad (16)$$

2.2 Estimation

Our system takes the form of a nonlinear state space model in which the state transition equation is given by (16) and observation equation by (7)-(11) where $P_{it}(j)$ is given by (12) and $p_{it}(\tau)$ by (13). Our baseline model has 12 parameters to estimate, namely the diagonal terms in the variance matrices Σ and R and the parameters governing genuine duration dependence, δ_1 , δ_2 and δ_3 . Because the observation equation is nonlinear in x_{it} , the extended Kalman filter can be used to form the likelihood function for the observed data $\{y_1, \dots, y_T\}$ and form an inference about the unobserved latent variables $\{\xi_1, \dots, \xi_T\}$, as detailed in Appendix B. Inference about historical values for ξ_t provided below correspond to full-sample smoothed inferences, denoted $\hat{\xi}_{t|T}$.

3 Results for the baseline specification

We estimated parameters for the above nonlinear state-space model using seasonally adjusted monthly data on $y_t = (U_t^1, U_t^{2.3}, U_t^{4.6}, U_t^{7.12}, U_t^{13.+})'$ for $t = \text{January 1976 through December 2013}$. Figure 4 plots smoothed estimates for $p_{it}(1)$, the probability that a newly unemployed worker of type i at $t - 1$ will still be unemployed at t . These average 0.35 for type H individuals and 0.82 for type L individuals, close to the average calculations of 0.36 and 0.85, respectively, that we arrived at in row 2 of Table 1 when we were explaining the intuition behind our identification strategy based on steady-state calculations. The probabilities of type H individuals remaining unemployed rise during the early recessions but are less cyclical in the last two recessions. By contrast, the continuation probabilities for type L individuals rise in all recessions. The gap between the two probabilities increased significantly over the last 20 years.

Figure 5 plots inflows of individuals of each type into the pool of newly unemployed. Type H workers constitute 77% on average of the newly unemployed, again close to the value of 79% expected on the basis of the simple steady-state calculations in row 2 of Table 1. Inflows of both types increase during recessions. New inflows of type H workers declined immediately at the end of every recession, but inflows of type L workers continued to rise after the recessions of 1990-91 and 2001 and were still at above-average levels 3 years after the end of the Great Recession. This changing behavior of type L workers' inflows appears to be another important characteristic of jobless recoveries. The Great Recession is unique in that the inflows of type L workers as well as

the continuation probabilities reached higher levels than any earlier dates in our data set.

The combined implications of these cyclical patterns are summarized in Figure 6. Before the Great Recession, the share in total unemployment of type L workers fluctuated between 30% and 50%, falling during expansions and rising during and after recessions. But during the Great Recession, the share of type L workers skyrocketed to over 80%. The usual recovery pattern of a falling share of type L workers has been very slow in the aftermath of the Great Recession.

While the inflows of type H workers show a downward trend since the 1980's, those of type L workers exhibit an upward trend. This difference in the low frequency movements of the two series provides a new perspective on the secular decrease in the inflows to unemployment and the secular rise in the average duration of unemployment. Abraham and Shimer (2001) and Aaronson, Mazumder and Schechter (2010) showed that the substantial rise in average duration of unemployment between mid-1980 and mid-2000 can be explained by the CPS redesign, the aging of the population and the increased labor force attachment of women. Bleakley, Ferris and Fuhrer (1999) concluded that the downward trend in inflows can be explained by reduced churning during this period. Figure 5 shows that the downward trend in the inflows is mainly driven by type H workers. The increased share of type L inflows contributed to the rise in the average duration of unemployment since the 1980's. This suggests that unobserved heterogeneity is important in accounting for low frequency dynamics in the labor market as well as those for business cycle frequencies.

Table 2 provides parameter estimates for our baseline model. We find a value for δ_1 , the parameter that governs genuine duration dependence for unemployment durations less than 6 months, that is near zero and statistically insignificant. The estimate of δ_2 (applying to individuals unemployed for more than 5 months and less than 1 year) is statistically significant and negative. The negative sign is consistent with the scarring hypothesis— the longer someone from either group has been unemployed, provided the duration has been 11 months or less, the more likely it is that person will be unemployed next month. On the other hand, we find a statistically significant positive value for δ_3 (unemployment lasting for a year and over). Once someone has been unemployed for more than a year, it becomes more likely as more months accumulate that they will either find a job or exit the labor force in any given month, consistent with what we have labeled motivational effects. This non-monotonic behavior of genuine duration dependence is displayed graphically in Panel A of Figure 7.

As seen in Panel B of Figure 7, our estimates of genuine duration dependence imply relatively modest changes in continuation probabilities for type L workers for most horizons. And while the implications for long-horizon continuation probabilities for type H workers may appear more significant, they are empirically irrelevant, since the probability that type H workers would be unemployed for more than 12 months is so remote. To gauge the overall significance of genuine duration dependence, we calculated the unemployment level predicted by our model for each date t in the sample if the values of δ_1 , δ_2 , and δ_3 were all set to zero, and found it would only be 2.3% lower on average than the value predicted by our baseline model. Thus although the values of δ_2 and δ_3 are statistically significant, they play a relatively minor role compared to ex ante heterogeneity in accounting for differences in continuation probabilities by duration of unemployment.

3.1 Variance decomposition

Many previous studies have tried to summarize the importance of different factors in determining unemployment by looking at correlations between the observed unemployment rate and the steady-state unemployment rate predicted by each factor of interest alone; see for example Fujita and Ramey (2009) and Shimer (2012). One major benefit of our framework is that it delivers a much cleaner answer to this question in the form of variance decompositions, a familiar method in linear VARs for measuring how much each shock contributes to the mean squared error (MSE) of an s -period-ahead forecast of a magnitude of interest.¹⁵

Our model can be used to account for the difference between the unemployment realization at time $t + s$ and a forecast based on values of the state vector only through date t in terms of the sequence of shocks between t and $t + s$, denoted $\varepsilon_{t+1}, \varepsilon_{t+2}, \dots, \varepsilon_{t+s}$. It is convenient to work with a linear approximation to that decomposition, which we show in Appendix C takes the form

$$y_{t+s} - \hat{y}_{t+s|t} \simeq \sum_{j=1}^s [\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s})] \varepsilon_{t+j} \quad (17)$$

for $\Psi_{s,j}(\cdot)$ a known (5×4) -valued function of $\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s}$. The mean squared error matrix associated with an s -period-ahead forecast of y_{t+s} is then

¹⁵See for example Hamilton (1994a, Section 11.5).

$$\begin{aligned}
E(y_{t+s} - \hat{y}_{t+s|t})(y_{t+s} - \hat{y}_{t+s|t})' &= \sum_{j=1}^s [\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s})] \Sigma [\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s})]' \quad (18) \\
&= \sum_{j=1}^s \sum_{m=1}^4 \Sigma_m [\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s}) e_m] [\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s}) e_m]'
\end{aligned}$$

for e_m column m of the (4×4) identity matrix and Σ_m the row m , column m element of Σ . Thus the contribution of innovations of type L worker's inflows (the first element of $\varepsilon_t = (\varepsilon_{Lt}^w, \varepsilon_{Ht}^w, \varepsilon_{Lt}^x, \varepsilon_{Ht}^x)'$) to the MSE of the s -period-ahead linear forecast error of total unemployment, $\iota_5' y_t$, is given by

$$\iota_5' \sum_{j=1}^s \Sigma_1 [\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s}) e_1] [\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s}) e_1]' \iota_5 \quad (19)$$

where ι_5 denotes a (5×1) vector of ones. Note that as in the constant-parameter linear case, the sum of the contributions of the 4 different structural shocks would be equal to the MSE of an s -period-ahead linear forecast of unemployment in the absence of measurement error. However, in our case the linearization is taken around time-varying values of $\{\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s}\}$. We can evaluate equation (19) at the smoothed inferences $\{\hat{\xi}_{t|T}, \hat{\xi}_{t-1|T}, \dots, \hat{\xi}_{t-47+s|T}\}$ and then take the average value across all dates t in the sample. This gives us an estimate of the contribution of the type L worker's inflows to unemployment fluctuations over a horizon of s months:

$$q_{s,1} = T^{-1} \sum_{t=1}^T \iota_5' \sum_{j=1}^s \Sigma_1 [\Psi_{s,j}(\hat{\xi}_{t|T}, \hat{\xi}_{t-1|T}, \dots, \hat{\xi}_{t-47+s|T}) e_1] [\Psi_{s,j}(\hat{\xi}_{t|T}, \hat{\xi}_{t-1|T}, \dots, \hat{\xi}_{t-47+s|T}) e_1]' \iota_5.$$

Consequently $q_{s,1} / \sum_{m=1}^4 q_{s,m}$ would be the ratio of the first factor's contribution to unemployment volatility at horizon s .

Figure 8 shows the contribution of each factor to the mean squared error in predicting overall unemployment as a function of the forecasting horizon. If one is trying to forecast unemployment one month ahead, uncertainty about future inflows of type H and type L workers are equally important. However, the farther one is looking into the future, the more important becomes uncertainty about what is going to happen to type L workers. If one is trying to predict one or two years into the future, the single most important source of uncertainty is inflows of new type L workers, followed by uncertainty about their outflows. Much of the MSE associated with a

2-year-ahead forecast of unemployment comes from not knowing when the next recession will begin or the current recession will end. For this reason, the MSE associated with 2-year-ahead forecasts is closely related to what some researchers refer to as the "business cycle frequency" in a spectral decomposition. If we are interested in the key factors that change as the economy moves into and out of recessions, inflows and outflows for type L workers are most important. We will provide additional evidence on this point in Section 3.2.

The last panel of Figure 8 breaks these contributions separately into inflows and outflows. Both inflows and outflows are important. However, the uncertainty about future inflows is more important in accounting for the error we would make in predicting total unemployment, accounting for more than 60% of the MSE for any forecasting horizon.

3.2 Historical decomposition

A separate question of interest is how much of the realized variation over some historical episode came from particular structural shocks. As in (17) our model implies an estimate of the contribution of shocks to a particular observed episode, namely

$$y_{t+s} - \hat{y}_{t+s|t} \simeq \sum_{j=1}^s [\Psi_{s,j}(\hat{\xi}_{t|T}, \hat{\xi}_{t-1|T}, \dots, \hat{\xi}_{t-47+s|T})] \hat{\varepsilon}_{t+j|T} \quad (20)$$

where $\hat{\varepsilon}_{t+s|T} = \hat{\xi}_{t+s|T} - \hat{\xi}_{t+s-1|T}$. From this equation, we can estimate for example the contribution of $\varepsilon_{L,t+1}^w, \varepsilon_{L,t+2}^w, \dots, \varepsilon_{L,t+s}^w$ (the shocks to w_L between $t+1$ and $t+s$) to the deviation of the level of unemployment at $t+s$ from the value predicted on the basis of initial conditions at t :

$$\iota_5' \sum_{j=1}^s [\Psi_{s,j}(\hat{\xi}_{t|T}, \hat{\xi}_{t-1|T}, \dots, \hat{\xi}_{t-47+s|T})] e_1 \hat{\varepsilon}_{t+j|T} e_1.$$

Figure 9 shows the contribution of each component to the realized unemployment rate in the last five recessions. In each panel, the solid line (labeled U_{base}) gives the change in the unemployment rate relative to the value at the start of the episode that would have been predicted on the basis of initial conditions. Typically an increase in the inflow of type L workers (whose contribution to total unemployment is indicated by the starred red curves) is the most important reason that unemployment rises during a recession. A continuing increase of these inflows even after the recession was over was an important factor in the jobless recoveries from the 1990 and 2001 recessions.

During the first 8 months of the Great Recession, changes in inflows and outflows of type L individuals were of equal importance in accounting for rising unemployment. But our model concludes that new inflows of type L individuals were by far the most important factor contributing to rising unemployment after July of 2008.

3.3 Features of the data that account for the conclusions

What features of the data lead us to the conclusions in Figures 8 and 9? Type L and type H individuals are not directly observable. Nevertheless, recall for example from Panel A of Figure 2 that our parameter estimates imply that most of the people who have been unemployed for longer than 4 months are likely to be type L individuals. We can thus directly observe an approximation to the unemployment-exit probabilities of type L individuals at any given date simply by looking at the average unemployment-exit probability of those who have been unemployed for 4 months or longer:

$$f_t^{4,+} = \frac{U_t^{4,+} - (U_{t+1}^{4,+} - U_{t+1}^4)}{U_t^{4,+}}. \quad (21)$$

The behavior of this series during the Great Recession is indicated by the blue line with circles in Figure 10. It fell during the first half of the recession but then stabilized, suggesting that ongoing deterioration in the unemployment-exit probabilities of type L workers was not the main factor contributing to rising unemployment during the second half of the recession.

On the other hand, any individual who had been unemployed for exactly 4 months in any given month t was most likely a newly unemployed type L individual at $t - 3$. The red starred line in Figure 10 plots U_t^4 around the Great Recession. This continued to increase long after $f_t^{4,+}$ had stabilized, suggesting that new inflows of type L individuals were the key factor contributing to rising unemployment in the second half of the Great Recession, consistent with the inference from our model in Panel E of Figure 9.

We can summarize the quantitative importance of these observations with the following simple calculations. Suppose that the unemployment-exit probabilities of the long-term unemployed had remained fixed at their value in 2008:M7, namely at $\bar{f}^{4,+} = 0.12$. If we apply this fixed rate to the observed new inflows into this category as measured by U_{t+1}^4 , the number unemployed for exactly 4 months, we would then predict a value for $U_{t+1}^{4,+}$, the number unemployed for 4 months or longer,

according to

$$\hat{U}_{t+1}^{4,+} = U_{t+1}^4 + \hat{U}_t^{4,+}(1 - \bar{f}^{4,+}).$$

If the number of unemployed for 1-3 months had also remained fixed at its value in 2008:M7 ($\bar{U}^{1,3} = 5.687$ million), we would then arrive at a predicted value for total unemployment of $\hat{U}_{t+1} = \hat{U}_{t+1}^{4,+} + \bar{U}^{1,3}$. This series is plotted in Panel A of Figure 11 along with the actual value for total unemployment U_{t+1} . These calculations demonstrate the basis in the observed data for concluding that much of the increase in unemployment during the second half of the Great Recession can be attributed to new inflows of type L individuals alone rather than to any deterioration in the unemployment-exit probability.

We can also use these calculations to see why our analysis reaches a different conclusion from Shimer (2012), who focused on the unemployment-exit probability itself. The aggregate probability is defined as

$$f_t = \frac{U_t - (U_{t+1} - U_{t+1}^1)}{U_t},$$

which is plotted as the solid line in Panel B of Figure 11. We can interpret this as a weighted average of the exit probabilities of those with duration 1-3 months and those with 4 months or longer,

$$f_t = \frac{U_t^{1,3} f_t^{1,3} + U_t^{4,+} f_t^{4,+}}{U_t^{1,3} + U_t^{4,+}},$$

which we use as the definition of $f_t^{1,3}$. We can then calculate what this magnitude would have been predicted to be if $U_t^{1,3}$, $f_t^{1,3}$, and $f_t^{4,+}$ had all remained frozen at their 2008:M7 levels, with the only thing that changed subsequently being the imputed new inflows of type L individuals:

$$\hat{f}_t = \frac{\bar{U}^{1,3} \bar{f}^{1,3} + \hat{U}_t^{4,+} \bar{f}^{4,+}}{\bar{U}^{1,3} + \hat{U}_t^{4,+}}.$$

This series is plotted as the dotted line in Panel B of Figure 11, and shows that much of the observed change in the unemployment-exit probability can be explained by increased inflows of type L individuals alone. It is in sharp contrast to Figure 9 in Shimer (2012), whose graphs purported to show that changes in the composition of the unemployed explain virtually none of the observed changes in exit probabilities. The reason is that his analysis assumed that everyone within a given group is homogeneous and did not make use of their differing individual unemployment

histories. Shimer thus overlooked the single most important driving variable in unemployment dynamics.

Before leaving this issue we should comment on an unresolved controversy in the literature about how to measure outflows from unemployment. Our measure (21) follows van den Berg and van Ours (1996), van den Berg and van der Klaauw (2001), Elsby, Michaels and Solon (2009), Shimer (2012), and Elsby, Hobijn and Şahin (2013) in deriving flow estimates from the observed change in the number of unemployed by duration. An alternative approach, employed by Fujita and Ramey (2009) and Elsby, Hobijn and Şahin (2010), is to look at only those individuals for whom there is a matched observation of unemployment in month t and a status of employment or out of the labor force in month $t + 1$. In the absence of measurement error, the two estimates should be the same, but in practice they turn out to be quite different. In particular, Elsby, Hobijn, and Şahin (2010, Figure 15) used matched flows to calculate a series for a monthly outflow rate from long-term unemployment that remains above 25% throughout the Great Recession, whereas the monthly outflow rate in our Figure 10 falls below 10%. One reason for the discrepancy is misclassification. For example, an individual who goes from long-term unemployed to out of the labor force to back to long-term unemployed in three successive months counts as a successful "graduate" from long-term unemployment using matched flows but is contributing to the stubborn persistence of long-term unemployment when using the stock data. A follow-up paper to Elsby, Hobijn and Şahin (2010) by Elsby et al. (2011) documented that more than half of the newly unemployed individuals reported their duration of unemployment to be 5 weeks or longer. Another important reason is that individuals for whom two consecutive observations are available differ in important ways from those for whom some observations are missing. Abowd and Zellner (1985) and Frazis et al. (2005) acknowledged that these measurement errors are more likely to bias the matched flow data than the stock data and suggested methods to correct the bias.

Since our goal is to understand how the reported stock of long-term unemployed came to be so high and why it falls so slowly, we feel that our approach, which is consistent with the observed stock data by construction, is preferable. In any case, we emphasize that the inflow and outflow rates in Figure 10 were not used in any way in producing our Figures 4-9, which were instead derived solely from the raw data plotted in Figure 1. We report the calculations in Figures 10 and 11 only to provide additional intuition and support for why our findings in Figures 4-9 came out

the way they did.

4 Who are the type L workers?

We noted that many of the individuals that our model designates as type L can be effectively identified ex post by the fact that most of those who have been unemployed longer than 4 months are likely in this group. In this section we discuss the relation between these individuals and various observable characteristics.

4.1 The importance of permanent involuntary separations and recalls

The BLS data include observable characteristics such as age, gender, education, occupation, industry, and reason for unemployment. The consensus of previous studies is that the last category holds the most promise for predicting unemployment duration, though it can only account for a small part of the observed cross-sectional dispersion. Darby, Haltiwanger and Plant (1986) argued that counter-cyclicality in the average unemployment duration mainly comes from the increased inflow of prime-age workers suffering permanent job loss who are likely to have low job-finding probabilities. Bednarek (1983) also noted that permanently separated workers are more likely to experience a long duration of unemployment, while Fujita and Moscarini (2013) showed that the unemployed who are likely to experience long-term unemployment spells tend to be those who are not recalled to work by their previous employers. Shimer (2012) found that the most important potential source of heterogeneity across different workers is differences in the reasons the individuals became unemployed, though he argued that this made only a small empirical contribution to observed cyclical fluctuations in unemployment and job-finding probabilities. Kroft et al. (forthcoming) concluded that observable characteristics could account for almost none of the rise in long-term unemployment during the Great Recession.

Panel A of Figure 12 breaks down people looking for work in terms of the reason they came to be unemployed. Dark bars describe the share of people who have been looking for work for less than one month and white bars the share of those who have been looking for more than 6 months. Permanent job losers and job losers on temporary layoff each account for about one fifth of new entrants into the pool of unemployed. By contrast, those on temporary layoff account for less than

3% of the unemployed with duration longer than 6 months, while around half of the long-term unemployed are accounted for by permanent job losers. This means that the unemployment exit probabilities of permanent job losers are much lower than those of job losers on temporary layoff.

Panel B of Figure 12 plots the inflows to unemployment by reason. Both the inflows of permanent job losers and those on temporary layoff exhibit counter-cyclicalities. They rise as the recession begins and fall as the recession ends. In Panel C of Figure 12 we compare our estimate of the number of newly unemployed type L workers to the number of those newly unemployed who gave permanent separations from their previous job as the reason¹⁶. The two series were arrived at using different data and different methodologies but exhibit remarkably similar dynamics. By contrast, our series for newly unemployed type L workers does not look much like any of the other series in Panel B. Panel D compares the total number of those unemployed who gave permanent separation as the reason to our estimate of the total number of unemployed type L workers, for which the correspondence is even more striking.

In March 2009 there were 1.38 million newly unemployed individuals who reported permanent separation as their reason for unemployment, 454,000 more than in March 2008. In March 2009 there were 3.47 million newly unemployed individuals altogether, 642,000 more than the previous year. This means that $454/642 = 71\%$ of the increase in U_t^1 between 2008:M3 and 2009:M3 was due to permanent separations.¹⁷ There is no question that permanent separations account for much of the increase in newly unemployed type L individuals that we identified in Figure 5 as occurring during this period.

We also repeated calculations like those in Panel A of Figure 11 using only those new inflows into U_t^4 who gave permanent separation as the reason. If we assumed that the unemployment-exit probabilities for this group as well as the number of unemployed in all other groups had remained fixed at their values of 2008:M7, we could account for an increase in total unemployment between 2008:M7 and 2009:M12 of 2.94 million individuals, almost half of the observed total increase of 6.37 million, as a result of inflows of type L individuals who became unemployed as a result of

¹⁶Permanent separations include permanent job losers and persons who completed temporary jobs. The separate series, permanent job losers and persons who completed temporary jobs, are publicly available from 1994, but their sum (permanent separations) is available back to 1976.

¹⁷We seasonally adjusted the number for newly unemployed individuals who reported permanent separation as their reason for unemployment using X-12-ARIMA. We also did the same calculation with publicly available seasonally unadjusted numbers and found that 81% of the increase in U_t^1 between 2008:M3 and 2009:M3 was due to permanent separations.

permanent separations.

4.2 Inference using data that condition on reason for unemployment

To obtain further evidence on the role of observed and unobserved worker characteristics, Ahn (2014) fit models like the one developed here to subsets of workers sorted based on observable characteristics. She replaced our observation vector y_t based on aggregate unemployment numbers with $y_{jt} = (U_{jt}^1, U_{jt}^{2.3}, U_{jt}^{4.6}, U_{jt}^{7.12}, U_{jt}^{13.+})'$ where $U_{jt}^{2.3}$ for example denotes the number of workers with observed characteristic j who have been unemployed for 2-3 months, the idea being that within the group j there are new inflows (w_{jHt} and w_{jLt}) and outflows (p_{jHt} and p_{jLt}) of two unobserved types of workers. Of particular interest for the present discussion are the results when j corresponds to one of the 5 reasons for why the individual was looking for work. Panel A of Figure 13 displays Ahn's estimated values for new inflows of type L workers for each of the categories as well as the sum $\sum_{j=1}^5 \hat{w}_{jLt|T}$. Our series $\hat{w}_{Lt|T}$ inferred from aggregate data is also plotted again for comparison. The sum of micro estimates is very similar to our aggregate estimates, and the individual micro components reveal clearly that those we have described as type L workers primarily represent a subset of people who were either permanently separated from their previous job or are looking again for work after a period of having been out of the labor force.

Ahn (2014) also calculated the models' inferences about the total number of type L individuals in any given observable category j who were unemployed in month t . These are plotted in Panel B of Figure 13. Here the correspondence between the aggregate inference and the sum of the micro estimates is even more compelling, as is the conclusion that type L unemployed workers represent primarily a subset of those permanently separated from their old jobs or re-entering the labor force.

Ahn (2014) found that permanent job losers who are type L account for around 50% of the aggregate type L unemployment and drive most of its counter-cyclicalities. The second most important group is type L re-entrants to the labor force. Considering that permanent job losers are likely to leave the labor force and re-enter to the labor force, there is a high chance that the type L re-entrants used to be permanent job losers before leaving the labor force. In addition, type L people are found disproportionately more among permanent job losers than they are among in other categories. The type L individuals account for one third of the newly unemployed permanent job losers, whereas they only comprise less than one fifth of the inflows in other categories.

Again it is useful to corroborate these conclusions with model-free direct evidence. Our goal is to examine the factors that account statistically for fluctuations in $U_t^{4,+}$, the seasonally adjusted count of individuals who have been unemployed for 4 months or longer. We are interested in the extent to which this can be predicted from the number of newly unemployed individuals with observed characteristic j . We also consider the role of outflows as measured by $F_t = U_{t-1} - U_t^{2,+}$. We summarize the usefulness of different variables for predicting long-term unemployment by estimating 12th-order vector autoregressions of the form

$$x_t = c + \Phi_1 x_{t-1} + \Phi_2 x_{t-2} + \cdots + \Phi_{12} x_{t-12} + \varepsilon_t \quad (22)$$

where x_t is an $(n \times 1)$ vector consisting of $U_t^{4,+}$ along with other variables, Φ_m are $(n \times n)$ matrices, and each row of the system is estimated by OLS.

We first consider a 3-variable system consisting of long-term unemployment along with gross outflows and inflows: $x_t = (U_t^{4,+}, F_t, U_t^1)'$. Key results are summarized in Table 3. Both inflows and outflows are statistically significant predictors of long-term unemployment; an F -test of the hypothesis that the (1,2) elements of $\{\Phi_1, \dots, \Phi_{12}\}$ are all zero rejects with a p -value of 10^{-10} , while the hypothesis that the (1,3) elements are all zero rejects with $p < 10^{-7}$. Of particular interest is a variance decomposition of the VAR, which calculates how much of a 24-month-ahead forecast error $x_{t+24} - \hat{x}_{t+24|t}$ is accounted for by innovations of each of the three variables. Typically in such decompositions the variance in any individual variable x_{it} is mostly accounted for by its own innovations ε_{it} . To try to minimize further any imputed role to innovations in inflows we order inflows U_t^1 last in the Cholesky factorization, meaning that any contemporaneous correlations among the three shocks is imputed to the first two rather than the third. We nevertheless find that inflows account for 34% of the two-year-ahead variance in long-term unemployment. By contrast only 30% can be attributed to outflows.

We next ask whether the composition of inflows has additional explanatory power by looking at a 4-variable VAR in which new inflows of permanently separated workers, $U_{PS,t}^1$, are added to the system. We find that permanently separated workers have significant predictive power even when aggregate inflows U_t^1 are already included in the regression (see Table 3, row 2, column 6). Indeed, when ordered third in the 4-variable VAR, new inflows of permanently separated workers

can account for 41% of the two-year-ahead variance of long-term unemployment and inflows of permanently separated and other workers together contribute 49%.¹⁸

Similar results for predicting longer term unemployment, $U_t^{7,+}$, as reported in row 3. And if we add in new claims for unemployment insurance (denoted S_t), which may be a more reliable measure of new inflows of involuntarily separated workers than estimates based on the CPS, the combined contribution of inflows (U_t^1 , $U_{PS,t}^1$, and S_t) is 63%.

Inflows are also quantitatively very important if we measure variables in terms of fractions rather than aggregate counts. Let $f_t = F_t/U_{t-1}$ denote the unemployment exit probability, $u_t^{4,+} = U_t^{4,+}/U_t$ long-term unemployment as a share of total, $u_{PS,t}^1 = U_{PS,t}^1/U_t^1$ permanent separations as a share of new unemployment, and $s_t = S_t/U_t^1$ new claims for unemployment insurance as a share of new unemployment. In a VAR ordered as $x_t = (u_t^{4,+}, f_t, u_{PS,t}^1, s_t)'$, inflows (as measured by the last two variables) account for 47% of the 24-month-ahead error in forecasting $u_t^{4,+}$ and 45% of the error in forecasting f_t , compared to 35% and 50%, respectively, accounted for by innovations in outflow probabilities f_t .

4.3 Understanding cyclical variation in heterogeneity

The vast majority of newly unemployed individuals will exit unemployment relatively quickly. Even among those who are newly unemployed as a result of a permanent separation, more than half would be designated within our framework as type H . In fact, within the "permanently separated" category, many workers do end up being recalled to their old positions (Fujita and Moscarini, 2013), and such individuals are likely to be included in our type H designation. This is why a much more important predictor of an individual's outcome is how long that individual has been unemployed rather than any observable characteristic. And this is also the key reason why many researchers, whose frameworks assume that all individuals with the same observed characteristic should have the same unemployment-exit probabilities, cannot account for the features that we find in the data.

Our approach also differs radically from the applied micro literature on this topic in that we have put cyclical variation in unobserved heterogeneity front and center of the analysis. Why

¹⁸Calculating the separate contributions of $U_{PS,t}^1$ and U_t^1 is quite sensitive to which is ordered third, since there is a significant contemporaneous correlation between the innovations in $U_{PS,t}^1$ and U_t^1 . However, it is certainly not the case that the former is simply proxying for the latter. When both are included in the regression as in the third row of Table 3, both are statistically significant, with permanent separations producing an F -test with p -value of 0.03 and total new unemployed producing $p < 10^{-5}$.

does unobserved heterogeneity vary cyclically? In normal times there is a tremendous amount of churning in the labor market, with millions of workers entering and exiting the unemployment pool every month even as the overall unemployment rate remains low— see for example, Davis, Faberman and Haltiwanger (2006). Lazear and Spletzer (2012) showed using micro data from JOLTS that churning is procyclical, with quits accounting for the major part of it. However, our measure of type H inflows often rises during recessions. It is clear that in addition to normal churning arising from those who quit their job voluntarily, unemployment due to temporary layoffs is another important part of what we have characterized as type H unemployment. Temporary layoffs rise during recessions, but insofar as many of these individuals often return to their old jobs relatively quickly, our procedure is assigning most of those on temporary layoff to type H rather than type L .

Finally, we emphasize that whether an individual is type L or type H can vary with economic circumstances. An unemployed carpenter who would have little trouble finding a job in normal times may spend a substantial period unemployed during a housing bust. Indeed, the fact that we have identified permanent involuntary separations as a key driver of new inflows of type L individuals is most naturally interpreted as exactly this kind of phenomenon.¹⁹

5 Robustness checks

Here we examine how our conclusions would change under a number of alternative specifications, including changes in the unemployment measures used, alternative specifications of genuine duration dependence, possible correlations among the shocks, and reformulation of the model in terms of weekly rather than a monthly frequency. Further details for all of these alternative specifications are reported in the online appendix.

¹⁹Some readers have asked us whether this view should imply that an individual could change type at some point during a spell of unemployment, for example, when carpenter skills come back in demand. In our framework, the defining characteristic of a type i is the average unemployment continuation probabilities for the group, p_{it} . These probabilities are indeed constantly changing with cyclical conditions under our baseline model. See also Ravenna and Walsh's (2012) DSGE in which the fraction of unemployed type L individuals increases during recessions and is an important propagating factor.

5.1 Accounting for the structural break in the CPS

As noted in Appendix A, a redesign in the CPS in 1994 introduced a structural break with which any user of these data has to deal. Our baseline estimates adjusted the unemployment duration data using differences between rotation groups 1 and 5 and groups 2-4 and 6-8 in the CPS micro data. Here we summarize how our results would change if we were to instead use the adjustment employed by Hornstein (2012).

Table 4 summarizes the implications of alternative specifications for what we see as the most important conclusions that emerge from our baseline analysis. The table breaks down the MSE of a forecast of the overall level of unemployment at 3-month, 1-year, and 2-year forecast horizons into the fraction of the forecast error that is attributable to various shocks. Column 1 gives the numbers implied by our baseline specification and highlights our key conclusion that inflows account for more than half the variance at all horizons. Inflows of type L workers are most important but the outflows of type L workers and the inflows of type H workers are also crucial at a 3-month horizon. At a 1- or 2-year horizon, shocks to inflow and outflow probabilities for type L workers are the most important factors. The table also reports asymptotic standard errors for each of these magnitudes.²⁰

Column 2 of Table 4 reports the analogous variance decompositions when we instead use Hornstein’s data adjustment as described in Appendix A. This produces very little change in these numbers. In column 3 we use only data subsequent to the redesign in 1994 making no adjustment to the reported BLS figures. This reduces the estimated contribution of inflows of type L workers at shorter horizons, but preserves our main finding that for business-cycle frequencies, changes for type L workers account for most of the fluctuations in unemployment, with changes in type L inflows accounting for about half the variance of unemployment at the 2-year horizon. We obtained similar results using the full data set from 1976-2013 with no adjustments for the 1994 redesign (column 4). We also found that the non-monotonic pattern in the genuine duration dependence is preserved regardless of data adjustment methods.

²⁰Standard errors were calculated as follows. For each model, we generated 500 draws for the k -dimensional parameter vector (where k is reported in the first row of the table) from a $N(\hat{\theta}, \hat{V})$ distribution where $\hat{\theta}$ is the MLE and \hat{V} is the $(k \times k)$ variance matrix from inversed hessian of the likelihood function. For each draw of $\theta^{(\ell)}$ we calculated the values implied by that $\theta^{(\ell)}$ and then calculated the standard error of that inference across the draws $\theta^{(1)}, \dots, \theta^{(500)}$.

Note that although we report the log likelihood and Schwarz’s (1978) Bayesian criterion in rows 2 and 3 of Table 4, the values for columns 2-4 are not comparable with the others due to a different definition of the observable data vector y_t .

5.2 Alternative specifications for genuine duration dependence

Our baseline specification assumed that a single parameter δ_1 described genuine duration dependence for any worker unemployed for less than 6 months. We also estimated a model in which each of the observed duration categories (2-3 months, 4-6 months, 7-12 months, and greater than 12 months) was characterized by a different genuine duration parameter, replacing (14) with

$$d_\tau = \begin{cases} \delta_1^A(\tau - 1) & \text{for } \tau < 3 \\ \delta_1^A(3 - 2) + \delta_1^B(\tau - 2) & \text{for } 3 \leq \tau < 6 \\ \delta_1^A(3 - 2) + \delta_1^B(5 - 2) + \delta_2(\tau - 5) & \text{for } 6 \leq \tau < 12 \\ \delta_1^A(3 - 2) + \delta_1^B(5 - 2) + \delta_2(11 - 5) + \delta_3(\tau - 11) & \text{for } 12 \leq \tau. \end{cases}$$

Adding this additional parameter δ_1^B results in only a trivial improvement in the likelihood function and virtually no change in any of the variance decompositions, as seen in comparing columns 1 and 5 of Table 4.

We also estimated a model in which genuine duration dependence varies over time, allowing the coefficients δ_j that characterize genuine duration dependence to take on different values when the national unemployment rate is above 6.5%, times when the labor market is in slack and it is likely that many job losers automatically became eligible for extended UI benefits.²¹ We re-estimated our state space model with $d_{t\tau}$ in equation (13) given by d_τ^0 for dates t for which $u_t \leq 6.5$ and d_τ^E if $u_t > 6.5$ where

$$d_\tau^j = \begin{cases} \delta_1^j(\tau - 1) & \text{for } \tau < 6 \\ \delta_1^j[(6 - 1) - 1] + \delta_2^j[\tau - (6 - 1)] & \text{for } 6 \leq \tau < 12 \\ \delta_1^j[(6 - 1) - 1] + \delta_2^j[(12 - 1) - (6 - 1)] + \delta_3^j[\tau - (12 - 1)] & \text{for } 12 \leq \tau. \end{cases}$$

for $j = 0$ or E .

²¹Vishwanath (1989) and Blanchard and Diamond (1994) developed theoretical models in which genuine duration dependence could be linked to market tightness. See Whittaker and Isaacs (2014) for a detailed discussion of the conditions that can trigger extended unemployment benefits.

Adding 3 new parameters $(\delta_1^E, \delta_2^E, \delta_3^E)$ to the model results in an increase in the log likelihood of 58.2, leading to a rejection (p -value < 0.001) of the null hypothesis that the values of δ_j are constant over time in favor of the alternative that they vary over time depending on eligibility for unemployment benefits. However, this does not change any of our core conclusions, as seen in column 6 of Table 4.

5.3 Allowing for correlated shocks

Our baseline specification assumed that the shocks to w_{Lt} , w_{Ht} , p_{Lt} and p_{Ht} were mutually uncorrelated. It is possible to generalize this in a parsimonious way by allowing a factor structure to the innovations, $\varepsilon_t = \lambda F_t + u_t$, where $F_t \sim N(0, 1)$, λ is a (4×1) vector of factor loadings, and u_t is a (4×1) vector of mutually uncorrelated idiosyncratic components with variance matrix $E(u_t u_t') = Q$:

$$E(\varepsilon_t \varepsilon_t') = \lambda \lambda' + Q$$

$$Q = \begin{bmatrix} (q_H^w)^2 & 0 & 0 & 0 \\ 0 & (q_L^w)^2 & 0 & 0 \\ 0 & 0 & (q_H^x)^2 & 0 \\ 0 & 0 & 0 & (q_L^x)^2 \end{bmatrix}.$$

In this case the variance decomposition (18) becomes

$$\begin{aligned} E(y_{t+s} - \hat{y}_{t+s|t})(y_{t+s} - \hat{y}_{t+s|t})' &= \sum_{j=1}^s [\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s})](\lambda \lambda' + Q)[\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s})]' \\ &= \sum_{j=1}^s [\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s})]\lambda \lambda' [\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s})]' \\ &\quad + \sum_{j=1}^s \sum_{m=1}^4 Q_m [\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s})e_m][\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s})e_m]' \end{aligned}$$

for Q_m the row m , column m element of Q . Because the factor F_t has an effect on all four components, it is not possible to impute the term involving $\lambda \lambda'$ to any one of the four shocks individually. However, we can calculate the portion of the MSE that is attributable to this aggregate factor along with those of each of the individual idiosyncratic shocks in u_t . This is reported in

column 7 of Table 4, and variance decompositions are plotted in Figure 14. The aggregate factor by itself accounts for 57% of the MSE of a 3-month-ahead forecast of unemployment, and inflows and outflows of type H workers account for another 19%. The aggregate factor is strongly correlated with outflows of type L workers. If we isolate the idiosyncratic component of each shock that is uncorrelated with the other three, shocks to inflows of type L workers account for only a quarter of the 3-month-ahead forecast error and almost 1/3 of the 2-year-ahead forecast error. There is essentially no role for the idiosyncratic component of outflows of type L workers, since changes in these outflows are so highly correlated with the other three shocks. This suggests that the probability of exiting unemployment of type L workers is closely related to an aggregate shock and that the compositional change of unemployment can be interpreted as an aggregate phenomenon that is core to the dynamics of economic recessions.

5.4 Time aggregation

Focusing on monthly transition probabilities understates flows into and out of unemployment since someone who loses their job in week 1 of a month but finds a new job in week 2 would never be counted as having been unemployed. Shimer (2012) argued that this time-aggregation bias would result in underestimating the importance of outflows in accounting for cyclical variation in unemployment, and Fujita and Ramey (2009), Shimer (2012) and Hornstein (2012) all formulated their models in continuous time.

On the other hand, Elsby, Michaels and Solon (2009) questioned the theoretical suitability of a continuous-time conception of unemployment dynamics, asking if it makes any sense to count a worker who loses a job at 5:00 p.m. one day and starts a new job at 9:00 a.m. the next as if they had been unemployed at all. We agree, and think that defining the central object of interest to be the fraction of those newly unemployed in month t who are still unemployed in month $t + k$, as in our baseline model, is the most useful way to pose questions about unemployment dynamics. Nevertheless, and following Kaitz (1970), Perry (1972), Sider (1985), Baker (1992), and Elsby, Michaels and Solon (2009) we also estimated a version of our model formulated in terms of weekly frequencies as an additional check for robustness.

We can do so relatively easily if we make a few simplifying assumptions. We view each month t as consisting of 4 equally-spaced weeks and assume that in each of these weeks there is an

inflow of w_{it} workers of type i , each of whom has a probability $p_{it}(0) = \exp[-\exp(x_{it})]$ of exiting unemployment the following week. This means that for those type i individuals who were newly unemployed during the first week of month t , $w_{it}[p_{it}(0)]^3$ are still unemployed as of the end of the month. Thus for the model interpreted in terms of weekly transitions, equation (7) would be replaced by

$$U_t^1 = \sum_{i=H,L} \{w_{it} + w_{it}[p_{it}(0)] + w_{it}[p_{it}(0)]^2 + w_{it}[p_{it}(0)]^3\} + r_t^1.$$

Likewise (8) becomes

$$U_t^{2,3} = \sum_{i=H,L} \sum_{s=1}^4 \{w_{i,t-1}[p_{i,t-1}(1)]^{8-s} + w_{i,t-2}[p_{i,t-2}(2)]^{12-s}\} + r_t^{2,3}$$

for $p_{it}(\tau)$ given by (13)-(14) for $\tau = 1, 2$. Note that although this formulation is conceptualized in terms of weekly inflow and outflows w_i and p_i , the observed data y_t are the same monthly series used in our other formulations, and the number of parameters is the same as for our baseline formulation.

As seen in column 8 of Table 4, the weekly formulation implies a slightly smaller role for inflows than our baseline model. This is to be expected, as allowing for shorter employment spells by construction imputes some people who exit unemployment by obtaining new jobs but then lose them again before the month is over. This may be one reason that Hornstein (2012), who used a model related to ours but in which employment spells could be infinitesimally short, found a smaller role for inflows than we do. One benefit of our formulation is that we can calculate the likelihood function associated with any of the alternative specifications. We find that the weekly model in column 8 has a slightly worse fit to the data than the baseline monthly model in column 1.

6 Conclusion

People who have been unemployed for longer periods than others have dramatically different probabilities of exiting unemployment, and these relative probabilities change significantly over the business cycle. Even when one conditions on observable characteristics, unobserved differences across people and the circumstances under which they came to be unemployed are crucial for

understanding these features of the data.

We have shown how the time series of unemployment levels by different duration categories can be used to infer inflows and outflows from unemployment for workers characterized by unobserved heterogeneity. In contrast to other methods, our approach uses the full history of unemployment data to summarize inflows and outflows from unemployment and allows us to make formal statistical statements about how much of the variance of unemployment is attributable to different factors as well as identify the particular changes that characterized individual historical episodes.

In normal times, around three quarters of those who are newly unemployed find jobs quickly. But in contrast to the conclusions of Hall (2005) and Shimer (2012), we find that more than half the variance in unemployment comes from shocks to the number of newly unemployed. A key feature of economic recessions is newly unemployed individuals who have significantly lower job-finding probabilities. Our inferred values for the size of this group exhibit remarkably similar dynamics to separate measures of the number of people who permanently lose their jobs. We conclude that recessions are characterized by a change in the circumstances under which people become unemployed that accounts for the greater difficulty in finding new jobs during a recession.

References

- Aaronson, Daniel, Bhashkar Mazumder, and Shani Schechter (2010). "What is behind the rise in long-term unemployment," *Economic Perspectives*, QII:28-51.
- Abbring, Jaap H., Gerard J. van den Berg, and Jan C. van Ours (2001). "Business Cycles and Compositional Variation in U.S. Unemployment," *Journal of Business & Economic Statistics*, 19(4): pp. 436-448.
- Abowd, John M., and Arnold Zellner (1985). "Estimating Gross Labor-Force Flows," *Journal of Business and Economic Statistics*, 3(3): 254-283.
- Abraham, Katharine G. and Robert Shimer (2002). "Changes in Unemployment Duration and Labor Force Attachment," in Alan Krueger and Robert Solow, eds., *The Roaring Nineties*, pp. 367-420. Russell Sage Foundation.
- Acemoglu, Daron (1995), "Public Policy in a Model of Long-term Unemployment," *Economica* 62:161-178.
- Ahn, Hie Joo (2014), "The Role of Observed and Unobserved Heterogeneity in the Duration of Unemployment Spells," Working paper, University of California, San Diego.
- Alvarez, Fernando, Katarína Borovičková, and Robert Shimer (2015). "Decomposing Duration Dependence in a Stopping Time Model," working paper, University of Chicago.
- Bachmann, Ronald, and Mathias Sinning (2012). "Decomposing the Ins and Outs of Cyclical Unemployment." Working paper, Australian National University.
- Baker, Michael (1992). "Unemployment Duration: Compositional Effects and Cyclical Variability," *American Economic Review*, 82(1):313-21.
- Barnichon, Regis and Andrew Figura (2015). "Labor Market Heterogeneity and the Aggregate Matching Function," *American Economic Journal: Macroeconomics*, 7(4): 222-49.
- Baumeister, Christiane and Gert Peersman (2013). "Time-Varying Effects of Oil Supply Shocks on the U.S. Economy," *American Economic Journal: Macroeconomics*, 5(4): 1-29.
- Bednarzik, Robert (1983). "Layoffs and Permanent Job Losses: Workers' Traits and Cyclical Patterns," *Monthly Labor Review*, September 1983, pp.3-12, Bureau of Labor Statistics.
- Blanchard, Olivier J. and Peter Diamond (1994). "Ranking, Unemployment Duration, and Wages," *The Review of Economic Studies*, 61(3):417-434,

Bleakley, Hoyt, Ann E. Ferris, and Jeffrey C. Fuhrer (1999). "New Data on Worker Flows During Business Cycle," *New England Economic Review*, July/August:49–76, 1999.

Bureau of Labor Statistics (2011). "How Long before the Unemployed Find Jobs or Quit Looking?", *Issues in Labor Statistics* 2011(1):1-6.

Darby, Michael R., John C. Haltiwanger, and Mark W. Plant (1986). "The Ins and Outs of Unemployment: The Ins Win," NBER Working Paper # 1997, National Bureau of Economic Research, Inc.

Davis, Steven J., R. Jason Faberman, and John Haltiwanger (2006). "The Flow Approach to Labor Markets: New Data Sources and Micro-Macro Links," *Journal of Economic Perspectives*, 20(3):3–26.

Den Haan, Wouter J. (2000). "The Comovement between Output and Prices," *Journal of Monetary Economics*, 46:3-30.

Elbers, Chris, and Geert Ridder (1982), "True and Spurious Duration Dependence: The Identifiability of the Proportional Hazard Model," *Review of Economic Studies*, 49(3): 403-409.

Elsby, Michael W. L., Ryan Michaels, and Gary Solon (2009). "The Ins and Outs of Cyclical Unemployment," *American Economic Journal: Macroeconomics*, 1(1): 84-110.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin (2010). "The Labor Market in the Great Recession," *Brookings Papers on Economic Activity*, Spring 2010: 1-56.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin (2013). "Unemployment Dynamics in the OECD," *Review of Economics and Statistics*, 95(2): 530-548.

Elsby, Michael W. L., Bart Hobijn, Ayşegül Şahin, and Robert G. Valletta (2011). "The Labor Market in the Great Recession: An Update to September 2011," *Brookings Papers on Economic Activity*, Fall 2011: 353-371.

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

Faberman, R. Jason and Marianna Kudlyak (2014). "The Intensity of Job Search and Search Duration," Working Paper 14-12, Federal Reserve Bank of Richmond.

Frazis, Harley J., Edwin L. Robison, Thomas D. Evans, and Martha A. Duff (2005). "Estimating Gross Flows Consistent with Stocks in the CPS," *Monthly Labor Review* September 2005: 3-9.

Fujita, Shigeru (2011). "Dynamics of Worker Flows and Vacancies: Evidence from the Sign Restriction Approach," *Journal of Applied Econometrics*, 26: 89-121.

Fujita, Shigeru and Garey Ramey (2009). "The Cyclicalities of Separation and Job Finding Rates," *International Economic Review*, 50(2):415-430.

Fujita, Shigeru and Giuseppe Moscarini (2013). "Recall and Unemployment," NBER Working Papers, w19640, National Bureau of Economic Research, Inc.

Hall, Robert E. (2005). "Job Loss, Job Finding and Unemployment in the U.S. Economy over the Past Fifty Years," in Mark Gertler and Kenneth Rogoff, eds., *NBER Macroeconomics Annual* 2005, Volume 20, National Bureau of Economic Research, Inc.

Hall, Robert E. (2014). "Quantifying the Lasting Harm to the U.S. Economy from the Financial Crisis," in Jonathan Parker and Michael Woodford, eds., *NBER Macroeconomics Annual* 2014, Volume 29, National Bureau of Economic Research, Inc.

Hall, Robert E. and Sam Schulhofer-Wohl (2015). "Measuring Matching Efficiency with Heterogeneous Jobseekers," NBER working papers, w20939, National Bureau of Economic Research, Inc.

Ham, John C., and Samuel A. Rea, Jr. (1987). "Unemployment Insurance and Male Unemployment Duration in Canada," *Journal of Labor Economics*, 5(3): 325-353.

Hamilton, James D. (1994a). *Time Series Analysis*. Princeton: Princeton University Press.

Hamilton, James D. (1994b). "State-Space Models," in Robert F. Engle and Daniel L. McFadden, eds., *Handbook of Econometrics*, Volume IV, Chapter 50, pp. 3039-3080. Amsterdam: Elsevier.

Heckman, James and B. Singer (1984a). "The Identifiability of the Proportional Hazard Model," *Review of Economic Studies*, 51(2):231-41.

Heckman, James and B. Singer (1984b). "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data," *Econometrica*, 52(2): 271-320.

Heckman, James, and B. Singer (1984c), "Econometric duration analysis", *Journal of Econometrics*, 24: 63-132.

Honoré, Bo E. (1993). "Identification Results for Duration Models with Multiple Spells," *Review of Economic Studies*, 60(1): 241-246.

Hornstein, Andreas (2012). "Accounting for Unemployment: the Long and Short of it," FRB

Richmond Working Paper, No.12-07, Federal Reserve Bank of Richmond.

Ilg, Randy E., and Eleni Theodossiou (2012). "Job Search of the Unemployed by Duration of Unemployment," *Monthly Labor Review*, March 2012, pp.41-49, Bureau of Labor Statistics.

Jarosch, Gregor, and Laura Pilossoph (2015). "Statistical Discrimination and Duration Dependence in the Job Finding Rate," working paper, Princeton University.

Kaitz, Hyman B. (1970). "Analyzing the Length of Spells of Unemployment," *Monthly Labor Review*, November 1970, pp.11-20, Bureau of Labor Statistics.

Katz, Lawrence F. and Bruce D. Meyer (1990a). "The Impact of the Potential Duration of Unemployment Benefits on the Duration of Unemployment," *Journal of Public Economics*, 41:45-72.

Katz, Lawrence F. and Bruce D. Meyer (1990b). "Unemployment Insurance, Recall Expectations, and Unemployment Outcomes," *The Quarterly Journal of Economics*, 105(4):973-1002.

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

Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo, and Lawrence F. Katz (forthcoming). "Long-Term Unemployment and the Great Recession: The Role of Composition, Duration Dependence, and Non-Participation," *Journal of Labor Economics*.

Lazear, Edward P. and James R. Spletzer (2012). "Hiring, Churn and the Business Cycle", *American Economic Review*, 102(3):575-79.

Ljungqvist, Lars, and Thomas J. Sargent (1998). "The European Unemployment Dilemma," *Journal of Political Economy* 106(3):514-550.

Meyer, Bruce D.(1990). "Unemployment Insurance and Unemployment Spells", *Econometrica*, 58(4):757-82.

Perry, George L. (1972). "Unemployment Flows in the U.S. Labor Market," *Brookings Papers on Economic Activity*, 3(2):245-292.

Polivka, Anne E. and Stephen M. Miller (1998). "The CPS after the Redesign: Refocusing the Economic Lens," in John Haltiwanger, Marilyn E. Manser, and Robert Topel, eds., *Labor Statistics Measurement Issues*, pp.249-86.

Ravenna, Federico, and Carl E. Walsh (2012). "Screening and Labor Market Flows in a Model

with Heterogeneous Workers," *Journal of Money, Credit and Banking*, 44(S2): 31-71.

Ridder, Geert (1990). "The Non-Parametric Identification of Generalized Accelerated Failure-Time Models", *Review of Economic Studies*, 57:167-182.

Schwarz, Gideon (1978). "Estimating the Dimension of a Model," *Annals of Statistics* 6:461-464.

Shimer, Robert (2012). "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, 15(2):127-148.

Sider, Hal (1985). "Unemployment Duration and Incidence: 1968-1982," *American Economic Review*, 75(3):461-72.

Van den Berg, Gerald J. (2001). "Duration Models: Specification, Identification, and Multiple Durations," in James J. Heckman and Edward Leamer, eds., *Handbook of Econometrics*, Volume V (North-Holland, Amsterdam).

Van den Berg, Gerald J., and Bas van der Klaauw (2001). "Combining Micro and Macro Unemployment Duration Data," *Journal of Econometrics* 102: 271-309.

Van den Berg, Gerald J., and Jan C. van Ours (1996). "Unemployment Dynamics and Duration Dependence," *Journal of Labor Economics*, 14(1):100-125.

Vishwanath, Tara. (1989). "Job Search, Stigma Effect, and Escape Rate from Unemployment," *Journal of Labor Economics*, 7(4):487-502.

Vroman, Wayne (2009). "Unemployment insurance recipients and nonrecipients in the CPS," *Monthly Labor Review*, October 2009, pp.44-53, Bureau of Labor Statistics.

White, Halbert (1982). "Maximum Likelihood Estimation of Misspecified Models," *Econometrica*, 50:1-25.

Whittaker, Julie M. and Katelin P. Isaacs (2014). "Unemployment Insurance: Programs and Benefits," *Congressional Research Service Report for Congress*, 7-5700.

Yashiv, Eran (2007). "US Labor Market Dynamics Revisited," *The Scandinavian Journal of Economics*, 109(4): 779-806.

Table 1. Actual and predicted values for unemployment on average and since 2007 using different steady-state representations

	Parameter values				Actual or predicted values				
					U^1	$U^{2.3}$	$U^{4.6}$	$U^{7.12}$	$U^{13.+}$
(1)	1976:M1-2013:M12				Observed values				
					3,210	2,303	1,238	1,050	636
(2)	w_H	w_L	p_H	p_L	Fitted (and predicted) values				
	2,531	679	0.360	0.848	3,210	2,303	1,238	1,050	(614)
(3)	w	δ_0	δ_1	δ_2	δ_3	Fitted values			
	3,210	0.0899	-0.3490	0.0110	1.757e-4	3,210	2,303	1,238	1,050 636
(4)	w_H	w_L	$p_H(1)$	$p_L(1)$	δ	Fitted values			
	2,528	683	0.360	0.846	-0.003	3,210	2,303	1,238	1,050 636
(5)	2007:M12-2013:M12				Observed values				
					3,339	2,787	2,131	2,426	1,902
(6)	w_H	w_L	p_H	p_L	Fitted (and predicted) values				
	2,274	1,065	0.329	0.890	3,339	2,787	2,131	2,426	(2,358)
(7)	w	δ_0	δ_1	δ_2	δ_3	Fitted values			
	3,339	0.2382	-0.6644	0.0547	-1.283e-3	3,339	2,787	2,131	2,426 1,902
(8)	w_H	w_L	$p_H(1)$	$p_L(1)$	δ	Fitted values			
	2,307	1,033	0.334	0.900	0.017	3,339	2,787	2,131	2,426 1,902

Notes to Table 1. Table reports average values of U_t^x in thousands of workers over the entire sample and since 2007 along with predicted values from simple steady-state calculations.

Parameters were chosen to fit exactly the values in that row appearing in normal face, while the model's predictions for other numbers are indicated by parentheses.

Table 2. Parameter estimates for the baseline model

σ_L^w	0.0434*** (0.0041)	R_1	0.0981*** (0.0058)	δ_1	0.0053 (0.0138)
σ_H^w	0.0456*** (0.0059)	$R_{2,3}$	0.0759*** (0.0043)	δ_2	-0.0647*** (0.0242)
σ_L^x	0.0446*** (0.0049)	$R_{4,6}$	0.0775*** (0.0068)	δ_3	0.0724*** (0.0250)
σ_H^x	0.0209*** (0.0028)	$R_{7,12}$	0.0597*** (0.0051)		
		R_{13+}	0.0366*** (0.0026)		
No. of Obs.		456			
Log-Likelihood		2,401.6			

Notes to Table 2. White (1982) quasi-maximum-likelihood standard errors in parentheses.

Table 3. Variance decomposition and test of null hypothesis that composition of inflows does not matter in alternative unrestricted vector autoregressions

	Dependent variable (1)	Other variables in VAR (2)	Variance decomposition			F -test (p -value) (6)
			Own (3)	Outflows (4)	Inflows (5)	
(1)	$U_t^{4.+}$	F_t, U_t^1	36%	30%	34%	—
(2)	$U_t^{4.+}$	$F_t, U_{PS,t}^1, U_t^1$	26%	25%	49%	$F(12, 389) = 1.97$ ($p = 0.03$)
(3)	$U_t^{7.+}$	$F_t, U_{PS,t}^1, U_t^1$	24%	30%	46%	$F(12, 389) = 1.79$ ($p = 0.05$)
(4)	$U_t^{4.+}$	$F_t, U_{PS,t}^1, S_t, U_t^1$	24%	14%	63%	$F(24, 377) = 1.78$ ($p = 0.01$)
(5)	$u_t^{4.+}$	$f_t, u_{PS,t}^1, s_t$	18%	35%	47%	$F(24, 389) = 1.11$ ($p = 0.33$)


Notes to Table 3. All results based on a 12-lag VAR estimated 1977:M7-2013:M12 including the variables indicated in columns 1 and 2 with Cholesky ordering from left to right. Variance decompositions refer to contributions to the 24-month-ahead mean-squared error for the variable indicated in column 1.

Table 4. Comparison of variance decomposition across different models


	Shocks	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
# of Param.		12	12	12	12	13	15	16	12
Log-L.		2401.6	2311.0	1149.8	2428.7	2402.0	2459.8	2413.2	2399.3
SIC		-4729.7	-4584.5	-2233.8	-4783.9	-4724.4	-4827.8	-4728.4	-4725.1
3 month	F	- -	- -	- -	- -	- -	- -	0.571 (0.071)	- -
	w_L	0.434 (0.047)	0.423 (0.049)	0.283 (0.063)	0.126 (0.031)	0.429 (0.048)	0.371 (0.045)	0.240 (0.039)	0.396 (0.046)
	w_H	0.223 (0.044)	0.248 (0.047)	0.209 (0.062)	0.396 (0.054)	0.226 (0.044)	0.278 (0.049)	0.128 (0.030)	0.147 (0.038)
	p_L	0.242 (0.039)	0.258 (0.040)	0.307 (0.064)	0.186 (0.039)	0.243 (0.040)	0.211 (0.037)	0.000 (0.051)	0.225 (0.036)
	p_H	0.101 (0.024)	0.071 (0.022)	0.201 (0.051)	0.292 (0.050)	0.103 (0.025)	0.140 (0.031)	0.061 (0.012)	0.232 (0.046)
	Inflows	0.657	0.671	0.492	0.522	0.655	0.649	0.368	0.543
	L group	0.676	0.681	0.590	0.312	0.672	0.582	0.240	0.621
1 year	F	- -	- -	- -	- -	- -	- -	0.615 (0.076)	- -
	w_L	0.545 (0.049)	0.517 (0.050)	0.386 (0.067)	0.292 (0.046)	0.545 (0.052)	0.533 (0.050)	0.311 (0.051)	0.496 (0.050)
	w_H	0.083 (0.020)	0.092 (0.022)	0.094 (0.031)	0.223 (0.042)	0.083 (0.020)	0.108 (0.024)	0.049 (0.013)	0.054 (0.016)
	p_L	0.333 (0.048)	0.365 (0.049)	0.413 (0.066)	0.297 (0.052)	0.333 (0.051)	0.303 (0.049)	0.000 (0.071)	0.354 (0.049)
	p_H	0.039 (0.010)	0.026 (0.009)	0.107 (0.028)	0.188 (0.040)	0.039 (0.011)	0.056 (0.014)	0.025 (0.005)	0.096 (0.024)
	Inflows	0.628	0.609	0.480	0.515	0.628	0.641	0.360	0.550
	L group	0.878	0.882	0.799	0.589	0.878	0.836	0.311	0.850
2 year	F	- -	- -	- -	- -	- -	- -	0.616 (0.078)	- -
	w_L	0.570 (0.050)	0.527 (0.052)	0.434 (0.068)	0.410 (0.049)	0.571 (0.054)	0.533 (0.052)	0.331 (0.055)	0.520 (0.051)
	w_H	0.057 (0.014)	0.062 (0.015)	0.061 (0.021)	0.152 (0.032)	0.057 (0.014)	0.074 (0.017)	0.035 (0.009)	0.038 (0.012)
	p_L	0.346 (0.049)	0.394 (0.051)	0.435 (0.067)	0.306 (0.053)	0.344 (0.053)	0.334 (0.052)	0.000 (0.075)	0.375 (0.050)
	p_H	0.027 (0.007)	0.018 (0.006)	0.070 (0.019)	0.133 (0.031)	0.028 (0.008)	0.039 (0.010)	0.018 (0.004)	0.068 (0.018)
	Inflows	0.627	0.589	0.495	0.562	0.628	0.607	0.366	0.558
	L group	0.916	0.921	0.869	0.716	0.915	0.867	0.331	0.895

Notes to Table 4. (1) Baseline model, (2) alternative data, (3) post 94 data, (4) unadjusted data, (5) unrestricted time-invariant genuine duration dependence (GDD), (6) time-varying GDD, (7) correlated shocks, (8) weekly frequency. SIC calculated as minus twice the log likelihood plus number of parameters

times log of sample size ($T = 456$). Likelihood and SIC for columns 2-4 are not comparable with the others because the data on y_t are different. F denotes aggregate factor. Standard errors in parentheses.




Panel A: U^1 , $U^{2,3}$ and $U^{4,6}$




Panel B: $U^{7,12}$ and $U^{13,+}$

Figure 1. Number of unemployed individuals (in thousands) by duration of time they have already been unemployed as of the indicated date

Notes to Figure 1. Panel A plots the number unemployed for 1 month, 2-3 months, and 4-6 months, while Panel B reports those unemployed 7-12 months and more than 12 months.




Panel A




Panel B

Figure 2. Predicted and actual numbers of unemployed as a function of duration based on constant-parameter cross-sectional heterogeneity specification

Notes to Figure 2. Horizontal axis shows duration of unemployment in months and vertical axis shows number of unemployed for that duration in thousands of individuals. Dots denote imputed values for $\bar{U}^1, \bar{U}^3, \bar{U}^5, \bar{U}^{9.5}$, and \bar{U}^{15} based on equation (6) with w_L , w_H , x_L , x_H , and δ chosen to fit the observed values of $\bar{U}^1, \bar{U}^{2.3}, \bar{U}^{4.6}, \bar{U}^{7.12}$, and $\bar{U}^{13.+}$ exactly. Curves denote predicted values from 4-parameter pure cross-sectional heterogeneity model fit to 1976-2013 historical average values for $\bar{U}^1, \bar{U}^{2.3}, \bar{U}^{4.6}$, and $\bar{U}^{7.12}$ (panel A) and for values since 2007:M12 (panel B).



Panel A



Panel B

Figure 3. Unemployment continuation probabilities as a function of duration based on constant-parameter pure genuine duration dependence specification

Notes to Figure 3. Horizontal axis shows duration of unemployment in months and vertical axis shows probability that individual is still unemployed the following month. Curves denote predicted values from the 5-parameter pure GDD model (4) fit to 1976-2013 historical average values for \bar{U}^1 , $\bar{U}^{2.3}$, $\bar{U}^{4.6}$, $\bar{U}^{7.12}$ and $\bar{U}^{13.+}$ (panel A) and for values since 2007:M12 (panel B). The GDD model exactly fits the dots plotted in Figure 2 for each case.




Figure 4. Probability that a newly unemployed worker of each type will still be unemployed the following month

Notes to Figure 4. The series plotted are $\hat{p}_{it|T}(1)$ for $i = L, H$.




Figure 5. Number of newly unemployed workers of each type

Notes to Figure 5. The series plotted are $\hat{w}_{it|T}$ for $i = L, H$.




Figure 6. Share of total unemployment accounted for by each type of worker






Figure 7. Estimates of genuine duration dependence


Notes to Figure 7. Panel A plots d_τ as a function of τ (months spent in unemployment). Panel B plots average unemployment-continuation probabilities of type H and type L workers as a function of duration of unemployment.



Panel A



Panel B



Panel C

Figure 8. Fraction of variance of error in forecasting total unemployment at different horizons attributable to separate factors

Notes to Figure 8. Horizontal axis indicates the number of months ahead s for which the forecast is formed. Panel A plots the contribution of each of the factors $\{w_{Ht}, w_{Lt}, x_{Ht}, x_{Lt}\}$ separately, Panel B shows combined contributions of $\{w_{Ht}, x_{Ht}\}$ and $\{w_{Lt}, x_{Lt}\}$, and Panel C shows combined contributions of $\{w_{Ht}, w_{Lt}\}$ and $\{x_{Ht}, x_{Lt}\}$.

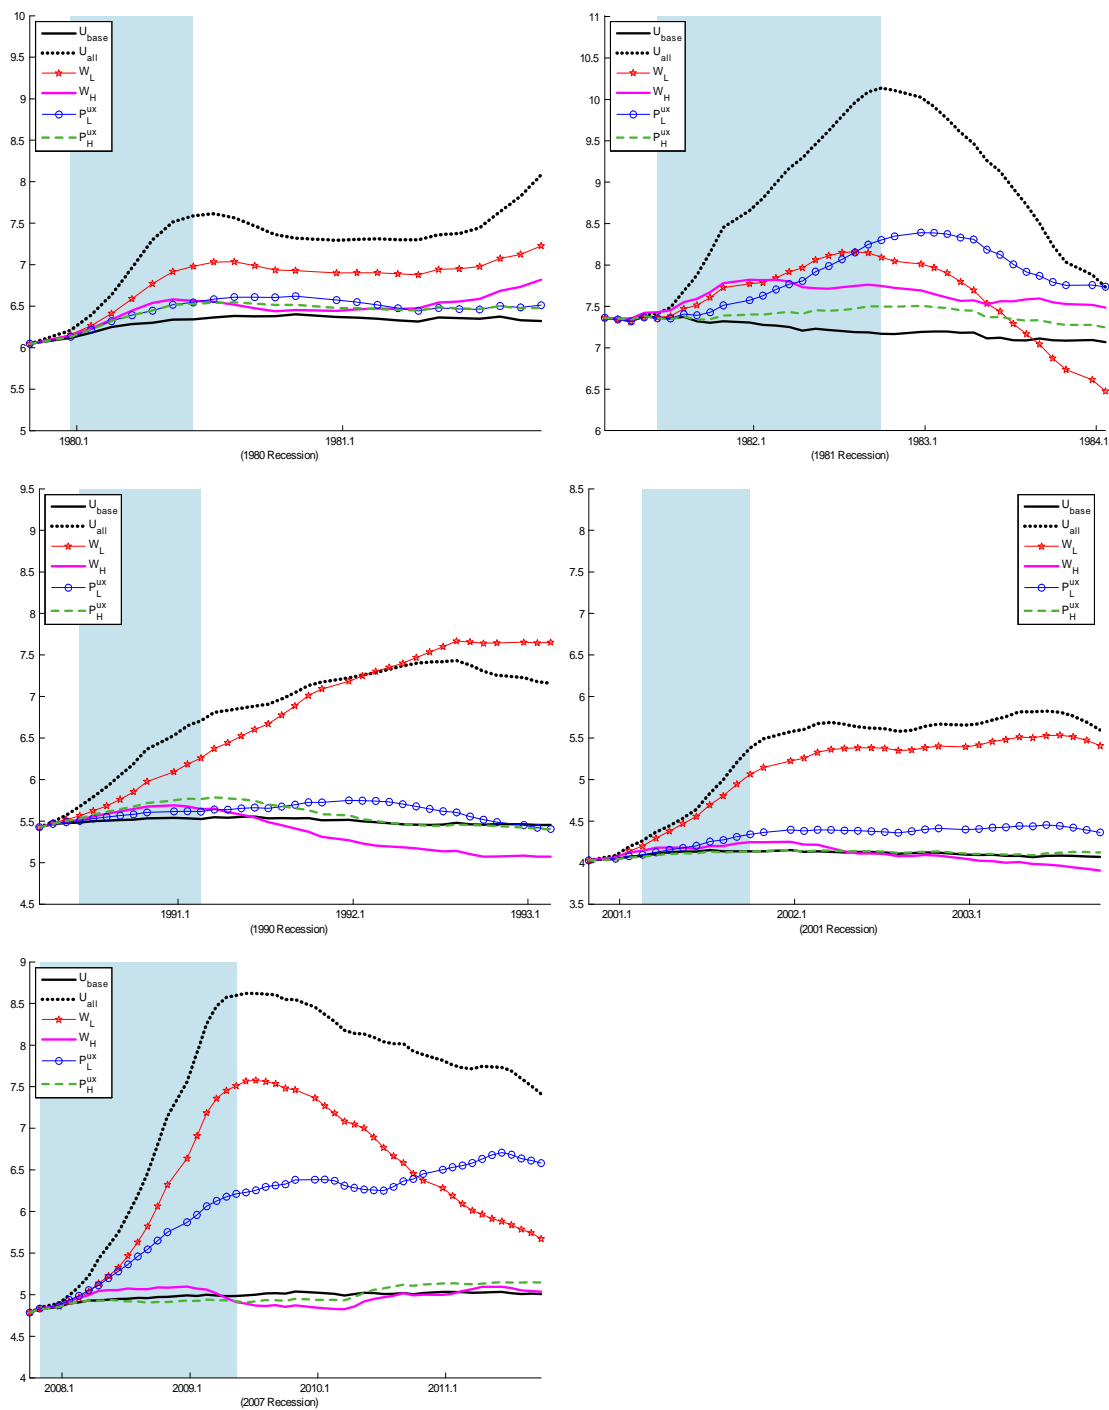


Figure 9. Historical decompositions of five U.S. recessions

Notes to Figure 9. The shaded areas denote NBER recessions.






Figure 10. Number unemployed for 4 months normalized at a value of 1.0 (starred line, left axis) and exit probability of those unemployed for 4 months and over (solid line, right axis) during and after the Great Recession

Notes to Figure 10. The shaded area denotes the Great Recession. Horizontal solid line denotes value of 2008:M7.



Panel A: Total number unemployed (U_t)



Panel B: Exit probability (f_t)

Figure 11. Realized and predicted total number unemployed and unemployment exit probabilities, October 2007 to May 2012

Notes to Figure 11. The shaded area denotes the Great Recession. Units for Panel A are in thousands workers. Predicted fixed exit probability from $U_t^{4,+}$ at 2008:M7 value.

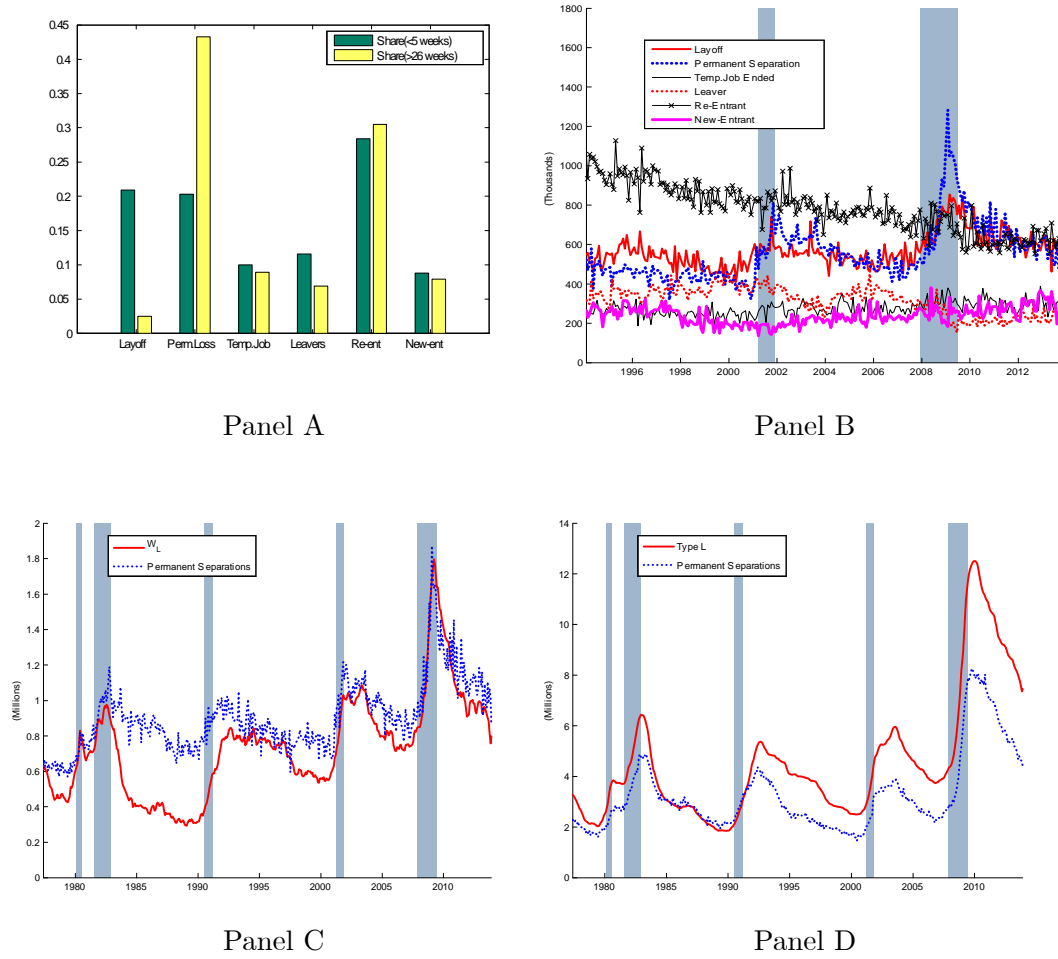




Figure 12. Breakdown of unemployment by reason for unemployment and duration

Notes to Figure 12. Panel A shows 1994-2013 average shares of unemployment by reason. Panel B plots newly unemployed individuals by reason for unemployment. Panel C shows newly unemployed type L individuals and newly unemployed individuals who gave permanent job loss or end of a temporary job as the reason. Panel D shows total numbers of unemployed type L workers compared to total numbers of unemployed who gave permanent job loss or end of temporary job as the reason.




Panel A




Panel B

Figure 13. Inflows and total numbers of type L workers by reason of unemployment


Notes to Figure 13. Panel A shows the number of type L individuals who are newly unemployed by reason of unemployment along with the sum across reasons (thick fuchsia) and inference based on uncategorized aggregate data (dashed black). Panel B shows the number of type L workers who have been unemployed for any duration by reason of unemployment along with the sum across reasons (thick fuchsia) and inference based on uncategorized aggregate data (dashed black). Source: Ahn (2014).



Panel A



Panel B



Panel C

Figure 14. Fraction of variance of error in forecasting total unemployment at different horizons attributable to separate factors in the model with correlated errors

Notes to Figure 14. Horizontal axis indicates the number of months ahead s for which the forecast is formed. Panel A shows the contribution of the aggregate factor F_t along with the idiosyncratic components of $\{w_{Ht}, w_{Lt}, x_{Ht}, x_{Lt}\}$ separately. Panel B shows the combined contributions of idiosyncratic components of $\{w_{Ht}, x_{Ht}\}$ and $\{w_{Lt}, x_{Lt}\}$ along with aggregate factor F_t . Panel C shows the combined contributions of idiosyncratic components of $\{w_{Ht}, w_{Lt}\}$ and $\{x_{Ht}, x_{Lt}\}$ along with aggregate factor F_t .

Appendix

A. Measurement issues and seasonal adjustment

The seasonally adjusted numbers of people unemployed for less than 5 weeks, for between 5 and 14 weeks, 15-26 weeks and for longer than 26 weeks are published by the Bureau of Labor Statistics. To further break down the number unemployed for longer than 26 weeks into those with duration between 27 and 52 weeks and with longer than 52 weeks, we used seasonally unadjusted CPS microdata publicly available at the NBER website (http://www.nber.org/data/cps_basic.html). Since the CPS is a probability sample, each individual is assigned a unique weight that is used to produce the aggregate data. From the CPS microdata, we obtain the number of unemployed whose duration of unemployment is between 27 and 52 weeks and the number longer than 52 weeks. We seasonally adjust the two series using X-12-ARIMA,²² and calculated the ratio of those unemployed 27-52 weeks to the sum. We then multiplied this ratio by the published BLS seasonally adjusted number for individuals who had been unemployed for longer than 26 weeks to obtain our series $U_t^{7.12}$.²³

An important issue in using these data is the redesign of the CPS in 1994. Before 1994, individuals were always asked how long they had been unemployed. After the redesign, if an individual is reported as unemployed during two consecutive months, then her duration is recorded automatically as the sum of her duration last month and the number of weeks between the two months' survey reference periods. Note that if an individual was unemployed during each of the two weeks surveyed, but worked at a job in between, that individual would likely self-report a duration of unemployment to be less than 5 weeks before the redesign, but the duration would be imputed to be a number greater than 5 weeks after the redesign.

As suggested by Elsby, Michaels and Solon (2009) and Shimer (2012) we can get an idea of the size of this effect by making use of the staggered CPS sample design. A given address is sampled for 4 months (called the first through fourth rotations, respectively), not sampled for the next 8

²² An earlier version of this paper dealt with seasonality by taking 12-month moving averages and arrived at similar overall results to those presented in this version. As a further check on the approach used here, we compared the published BLS seasonally adjusted number for those unemployed with duration between 15 and 26 weeks to an X-12-ARIMA-adjusted estimate constructed from the CPS microdata, and found the series to be quite close.

²³ This adjustment is necessary because the published number for unemployed with duration longer than 26 weeks is different from that directly computed from the CPS microdata, although the difference is subtle. The difference arises because the BLS imputes the numbers unemployed with different durations to various factors, e.g., correction of missing observations.

months, and then sampled again for another 4 months (the fifth through eighth rotations). After the 1994 redesign, the durations for unemployed individuals in rotations 2-4 and 6-8 are imputed, whereas those in rotations 1 and 5 are self-reported, just as they were before 1994. For those in rotation groups 1 and 5, we can calculate the fraction of individuals who are newly unemployed and compare this with the total fraction of newly unemployed individuals across all rotations. The ratio of these two numbers is reported in Panel A of Figure A1, and averaged 1.15 over the period 1994-2007 as reported in the second row of Table A1. For comparison, the ratio averaged 1.01 over the period 1989-1993, as seen in the first row. This calculation suggests that if we want to compare the value of U_t^1 as calculated under the redesign to the self-reported numbers available before 1994, we should multiply the former by 1.15. This is similar to the adjustment factors of 1.10 used by Hornstein (2012), 1.154 by Elsby, Michaels and Solon (2009), 1.106 by Shimer (2012), and 1.205 by Polivka and Miller (1998).

For our study, unlike most previous researchers, we also need to specify which categories the underreported newly unemployed are coming from. Figure A1 reports the observed ratios of rotation 1 and 5 shares to the total for the various duration groups, with average values summarized in Table A1. One interesting feature is that under the redesign, the fraction of those with 7-12 month duration from rotations 1 and 5 is very similar to that for other rotations, whereas the fraction of those with 13 or more months is much lower.²⁴ Based on the values in Table A1, we should scale up the estimated values for U_t^1 and scale down the estimated values of $U_t^{2,3}$ and $U_t^{13,+}$ relative to the pre-1994 numbers. The values for $U_t^{4,6}$ and $U_t^{7,12}$ seem not to have been affected much by the redesign. Our preferred adjustment for data subsequent to the 1994 redesign is to multiply U_t^1 by 1.15, $U_t^{2,3}$ by 0.87, $U_t^{13,+}$ by 0.77, and leave $U_t^{4,6}$ and $U_t^{7,12}$ as is. We then multiplied all of our adjusted duration figures by the ratio of total BLS reported unemployment to the sum of our adjusted series in order to match the BLS aggregate exactly.

Hornstein (2012) adopted an alternative adjustment, assuming that all of the imputed newly unemployed came from the $U_t^{2,3}$ category. He chose to multiply U_t^1 by 1.10 and subtract the added workers solely from the $U_t^{2,3}$ category. As a robustness check we also report results using

²⁴One possible explanation is digit preference— an individual is much more likely to report having been unemployed for 12 months than 13 or 14 months. When someone in rotation 5 reports they have been unemployed for 12 months, BLS simply counts them as such, and if they are still unemployed the following month, BLS imputes to them a duration of 13 months. The imputed number of people 13 months and higher is significantly bigger than the self-reported numbers, just as the imputed number of people with 2-3 months appears to be higher than self-reported.

Hornstein’s proposed adjustment in Section 5.1, as well as results using no adjustments at all.

An alternative might be to use the ratios for each t in Figure A1 rather than to use the averages from Table A1. However, as Shimer (2012) and Elsby, Michaels and Solon (2009) mentioned, such an adjustment would be based on only about one quarter of the sample and thus multiplies the sampling variance of the estimate by about four, which implies that noise from the correction procedure could be misleading in understanding the unemployment dynamics.

Table A1. Average ratio of each duration group’s share in the first/fifth rotation group to that in total unemployment

	U^1	$U^{2,3}$	$U^{4,6}$	$U^{7,12}$	$U^{13,+}$
1989-1993	1.01	1.01	0.96	1.02	0.97
1994-2007	1.15	0.87	0.95	1.05	0.77

B. Estimation Algorithm

The system (16) and (7)-(11) can be written as

$$x_t = Fx_{t-1} + v_t$$

$$y_t = h(x_t) + r_t$$

for $x_t = (\xi'_t, \xi'_{t-1}, \dots, \xi'_{t-47})'$, $E(v_t v'_t) = Q$, and $E(r_t r'_t) = R$. The function $h(\cdot)$ as well as elements of the variance matrices R and Q depend on the parameter vector $\theta = (\delta_1, \delta_2, \delta_3, R_1, R_{2,3}, R_{4,6}, R_{7,12}, R_{13+}, \sigma_L^w, \sigma_H^w, \sigma_L^x, \sigma_H^x)'$. The extended Kalman filter (e.g., Hamilton, 1994b) can be viewed as an iterative algorithm to calculate a forecast $\hat{x}_{t+1|t}$ of the state vector conditioned on knowledge of θ and observation of $Y_t = (y'_t, y'_{t-1}, \dots, y'_1)'$ with $P_{t+1|t}$ the MSE of this forecast. With these we can approximate the distribution of y_t conditioned on Y_{t-1} as $N(h(\hat{x}_{t|t-1}), H'_t P_{t|t-1} H_t + R)$ for $H_t = \partial h(x_t) / \partial x'_t|_{x_t = \hat{x}_{t|t-1}}$ from which the likelihood function associated with that θ can be calculated and maximized numerically. The forecast of the state vector can be updated using

$$\hat{x}_{t+1|t} = F\hat{x}_{t|t-1} + FK_t(y_t - h(\hat{x}_{t|t-1}))$$

$$K_t = P_{t|t-1} H_t (H'_t P_{t|t-1} H_t + R)^{-1}$$

$$P_{t+1|t} = F(P_{t|t-1} - K_t H_t' P_{t|t-1}) F' + Q.$$

A similar recursion can be used to form an inference about x_t using the full sample of available data, $\hat{x}_{t|T} = E(x_t|y_T, \dots, y_1)$ and these smoothed inferences are what are reported in any graphs in this paper; see our online appendix for further details.

Prior to the starting date January 1976 for our sample, BLS aggregates are available but not the micro data that we used to construct $U_t^{13,+}$. For the initial value for the extended Kalman filter, we calculated the values that would be implied if pre-sample values had been realizations from an initial steady state, estimating the (4×1) vector $\bar{\xi}_0$ from the average values for $\bar{U}^1, \bar{U}^{2,3}, \bar{U}^{4,6}$, and $\bar{U}^{7,+}$ over February 1972 - January 1976 using the method described in Section 1.1. Our initial guess was then $\hat{x}_{1|0} = \iota_{48} \otimes \bar{\xi}_0$ where ι_{48} denotes a (48×1) vector of ones. Diagonal elements of $P_{1|0}$ determine how much the presample values of ξ_j are allowed to differ from this initial guess $\hat{\xi}_{j|0}$. For this we set $E(\xi_j - \hat{\xi}_{j|0})(\xi_j - \hat{\xi}_{j|0})' = c_0 I_4 + (1 - j)c_1 I_4$ with $c_0 = 10$ and $c_1 = 0.1$. The value for c_0 is quite large relative to the range of $\xi_{t|T}$ over the complete observed sample, ensuring that the particular value we specified for $\hat{x}_{1|0}$ has little influence. For $k < j$ we specify the covariance²⁵ $E(\xi_j - \bar{\xi}_0)(\xi_k - \bar{\xi}_0)' = E(\xi_j - \bar{\xi}_0)(\xi_j - \bar{\xi}_0)'$. The small value for c_1 forces presample ξ_j to be close to ξ_k when j is close to k , again consistent with the observed month-to-month variation in $\hat{\xi}_{t|T}$.

C. Derivation of linearized variance and historical decompositions

The state equation $\xi_{t+1} = \xi_t + \varepsilon_{t+1}$ implies

$$\begin{aligned} \xi_{t+s} &= \xi_t + \varepsilon_{t+1} + \varepsilon_{t+2} + \varepsilon_{t+3} + \dots + \varepsilon_{t+s} \\ &= \xi_t + u_{t+s}. \end{aligned}$$

Letting $y_t = (U_t^1, U_t^{2,3}, U_t^{4,6}, U_t^{7,12}, U_t^{13,+})'$ denote the (5×1) vector of observations for date t , our model implies that in the absence of measurement error y_t would equal $h(\xi_t, \xi_{t-1}, \xi_{t-2}, \dots, \xi_{t-47})$

²⁵In other words,

$$P_{1|0} = \begin{bmatrix} c_0 I_4 & c_0 I_4 & c_0 I_4 & \dots & c_0 I_4 \\ c_0 I_4 & c_0 I_4 + c_1 I_4 & c_0 I_4 + c_1 I_4 & \dots & c_0 I_4 + c_1 I_4 \\ c_0 I_4 & c_0 I_4 + c_1 I_4 & c_0 I_4 + 2c_1 I_4 & \dots & c_0 I_4 + 2c_1 I_4 \\ \vdots & \vdots & \vdots & \dots & \vdots \\ c_0 I_4 & c_0 I_4 + c_1 I_4 & c_0 I_4 + 2c_1 I_4 & \dots & c_0 I_4 + 47c_1 I_4 \end{bmatrix}.$$

where $h(\cdot)$ is a known nonlinear function. Hence

$$y_{t+s} = h(u_{t+s} + \xi_t, u_{t+s-1} + \xi_t, \dots, u_{t+1} + \xi_t, \xi_t, \xi_{t-1}, \dots, \xi_{t-47+s}).$$

We can take a first-order Taylor expansion of this function around $u_{t+j} = 0$ for $j = 1, 2, \dots, s$,

$$y_{t+s} \simeq h(\xi_t, \dots, \xi_t, \xi_t, \xi_{t-1}, \dots, \xi_{t-47+s}) + \sum_{j=1}^s [H_j(\xi_t, \xi_t, \dots, \xi_t, \xi_t, \xi_{t-1}, \dots, \xi_{t-47+s})] u_{t+s+1-j}$$

for $H_j(\cdot)$ the (5×4) matrix associated with the derivative of $h(\cdot)$ with respect to its j th argument.

Using the definition of u_{t+j} , this can be rewritten as

$$y_{t+s} \simeq c_s(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s}) + \sum_{j=1}^s [\Psi_{s,j}(\xi_t, \xi_{t-1}, \dots, \xi_{t-47+s})] \varepsilon_{t+j}$$

from which (17) follows immediately.

Similarly, for purposes of a historical decomposition note that the smoothed inferences satisfy

$$\hat{\xi}_{t+s|T} = \hat{\xi}_{t|T} + \hat{\varepsilon}_{t+1|T} + \hat{\varepsilon}_{t+2|T} + \hat{\varepsilon}_{t+3|T} + \dots + \hat{\varepsilon}_{t+s|T}$$

where $\hat{\varepsilon}_{t+s|T} = \hat{\xi}_{t+s|T} - \hat{\xi}_{t+s-1|T}$. For any date $t + s$ we then have the following model-inferred value for the number of people unemployed:


$$\iota_5' h(\hat{\xi}_{t+s|T}, \hat{\xi}_{t+s-1|T}, \hat{\xi}_{t+s-2|T}, \dots, \hat{\xi}_{t+s-47|T}).$$

For an episode starting at some date t , we can then calculate


$$\iota_5' h(\hat{\xi}_{t|T}, \hat{\xi}_{t|T}, \hat{\xi}_{t|T}, \dots, \hat{\xi}_{t|T}, \hat{\xi}_{t-1|T}, \dots, \hat{\xi}_{t+s-47|T}).$$

This represents the path that unemployment would have been expected to follow between t and $t + s$ as a result of initial conditions at time t if there were no new shocks between t and $t + s$. Given this path for unemployment that is implied by initial conditions, we can then isolate the contribution of each separate shock between t and $t + s$. Using the linearization in equation (17) allows us to


represent the realized deviation from this path in terms of the contribution of individual historical shocks as in (20).




Panel A: U^1




Panel B: $U^{2.3}$



Panel C: $U^{4.6}$



Panel D: $U^{7.12}$



Panel E: $U^{13.+}$

Figure A1. Ratio of each duration group's share in the first and fifth rotation groups to that in all rotation groups