The Pervasive Importance of Tightness in Labor-Market Volatility *

Robert E. Hall
Hoover Institution and Department of Economics,
Stanford University
National Bureau of Economic Research
rehall@stanford.edu; stanford.edu/∼rehall
Sam Schulhofer-Wohl
Federal Reserve Bank of Chicago
Sam@chi.frb.org

April 11, 2017

Abstract

The distinctive contribution of the unemployment model of Diamond, Mortensen, and Pissarides was the creation of an economically coherent concept of labor-market tightness. In a tight market, jobseekers find jobs quickly and employers take longer to fill jobs. The evidence in this paper on both outflow rates from unemployment and on job-filling rates shows that a single index factor describes tightness. Outflow rates from unemployment and the duration of vacancies move in parallel as functions of the single index.

Two flows determine the unemployment rate: inflows—such as job loss—and outflows—such as job finding. Recessions generally involve a bulge of inflows. Tightness controls the outflow rate from unemployment. A non-parametric decomposition of unemployment into components of the inflow and the outflow rate shows that the outflow rate plays a substantially larger role in unemployment volatility than does the inflow rate, though the latter matters, especially at the beginning of a recession. The DMP model's concept of tightness is central to the understanding of fluctuations in unemployment and thus in employment and output.

JEL E24, J64

*Hall's research was supported by the Hoover Institution. The views expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of Chicago or the National Bureau of Economic Research. Complete backup for all of the calculations will be available from stanford.edu/∼rehall

The modern theory of unemployment portrays job-seeking as a transitory state at the individual level. People flow into unemployment for heterogeneous reasons, some the result of personal choice—entry or re-entry to the labor market and quitting an earlier job—some the result of an employer's unilateral decision—loss of a job, with or without a chance at recall to that job—and some joint—the ending of an explicitly temporary job. Job-seekers spend weeks or months in unemployment, with substantial monthly probabilities of flowing out, either by taking a job, or by ending search and leaving the labor force.

Diamond, Mortensen, and Pissarides's model of unemployment considers both the flow into unemployment and the flow out. Because the model delivers a coherent version of the previously murky concept of labor-market tightness, it is common to strip the model down by assuming a constant rate of flow out of jobs and into unemployment so as to focus on the determination of tightness and thus of the job-finding rate. But many authors, from the beginning of the DMP literature, made the inflow to unemployment an endogenous variable.

This paper studies unemployment data with the aim of clarifying the roles of inflows to and outflows from unemployment. It joins a substantial literature on the "ins and outs" of unemployment. It concentrates on heterogeneity in the flows. The first step is to demonstrate that labor-market tightness is powerful and pervasive. Across data on unemployment outflow rates for 18 categories of workers and data on vacancy durations for 16 industries, the paper finds that all 34 measures move together. A single driving force affects all types of workers and all industries.

The paper then turns to the role of fluctuations in tightness in the determination of unemployment. Recessions in general and the last one in particular involve sharp increases in unemployment followed by protracted declines. During the sharp increase phase, there is a substantial increase in the number of jobseekers who lost earlier jobs they believed were permanent. Inflows to unemployment are less persistent that changes in the outflow rate. In a decomposition of the overall movements of unemployment that recognizes the heterogeneity of unemployment by both the events that resulted in unemployment and the duration of a spell to date, fluctuations in outflow rates are substantially more important than fluctuations in inflow rates, averaged over expansions and contractions.

This paper is not about distinguishing true duration effects from heterogeneity among the unemployed. And it is not concerned with the time-aggregation issue that has occupied much of the earlier literature. It is about the relevance of the basic ideas of the DMP search-and-matching model, specifically about its concept of labor-market tightness. In that model, an adverse shock slackens the labor market—outflow rates of jobseekers fall and the recruiting employers fill jobs more quickly. Tightness determines the success of job-seekers and the success of recruiters, in opposite directions. The simplest version of the DMP model, which ignores heterogeneity among job-seekers, has the striking property that a burst of separations has no effect on tightness and only a transitory effect on unemployment. The persistence of unemployment arises almost entirely from the persistence of movements in incentives for job creation. With heterogeneity, the changes in the composition of unemployment that accompany sharp adverse shocks makes some contribution—in concept, potentially a lot—to the determination of the level of unemployment. One of the purposes of the paper is to measure tightness in a way that does not conflate it with the composition effects.

1 Related Literature

Lilien (1982) argued for a large role reallocation in recessions and thus made a case that rising inflows were an important component of rising unemployment, but Abraham and Katz (1986) were influential critics of that view. They pointed out that measures of labor-market tightness moved together across most sectors, while if some sectors were shedding workers and others absorbing them, that correlation would be absent. See Hall (2017) for recent evidence confirming that observation and arguing for a single aggregate driving force for labor-market tightness, namely discount rates.

Shimer (2005) found a preponderant role for variations in outflows, supporting the assumption from Mortensen and Pissarides (1994) that the separation rate could be taken as a constant. As he notes on p. 33, it's important to distinguish outflow rates from unemployment—outflows per unemployed person—from flows of the number of people departing unemployment. The latter is not very cyclical but the rate is highly pro-cyclical. An important property of the DMP class of models, as Shimer, p. 40, emphasized, is that a burst of job loss with no change in job-creation incentives has essentially no effect on market tightness. This point was implicit in Abraham and Katz (1986) much earlier.

den Haan, Ramey and Watson (2000) studied a model focusing on displacement; they assumed no tightness effects—the job-finding rate is a constant in their model. Fujita and Ramey (2009) disagreed with Shimer, finding that fluctuations in inflow rates to unemployment account for almost half of unemployment volatility on a contemporaneous basis and

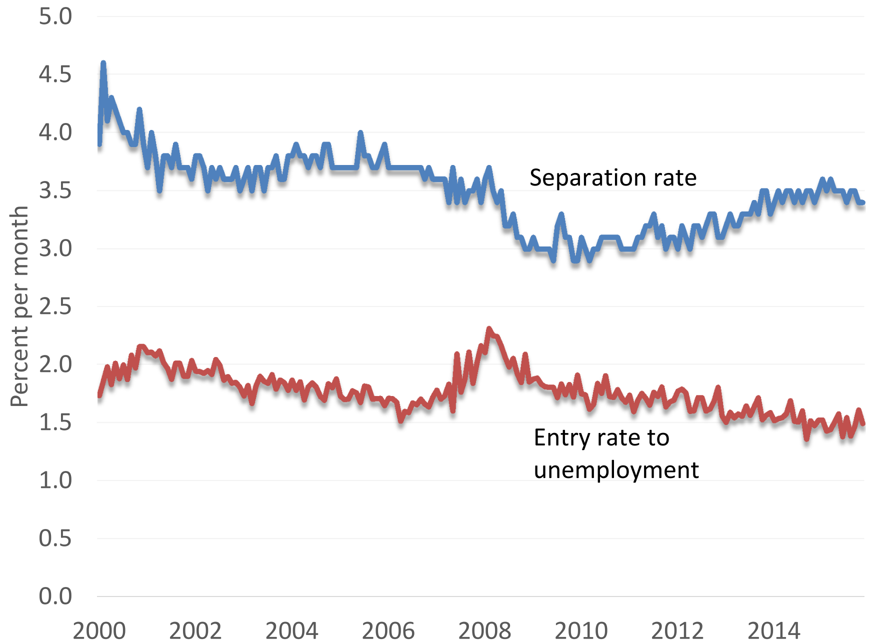

Figure 1: Separation Rate from Jobs and Inflow Rate to Unemployment

even more considering dynamics. Their work focuses on cross-correlations rather than on effects within a model of general or labor-market equilibrium.

Hall (2005) makes the distinction between separations from jobs, which were not elevated as a result of the recession of 2001, and inflow rates to unemployment, which rise in recessions because the fraction of job-losers rises in recessions. Perry (1972) pointed this out. In consequence, there is a big distinction between the cyclatility of separations and of inflow rates to unemployment. It has never been in dispute that the latter were countercyclical. See Figure 1.

Recently, Ahn and Hamilton (2016) concluded that displacements are the main channel of increasing unemployment following adverse shocks. Part of the debate is no more than semantic, because a good part of the slowdown in job-finding arises from changes in the composition of jobseekers, toward those, especially displaced workers, who have lower job-finding rates.

2 Measures of tightness

In the DMP model, tightness can be measured from two perspectives:

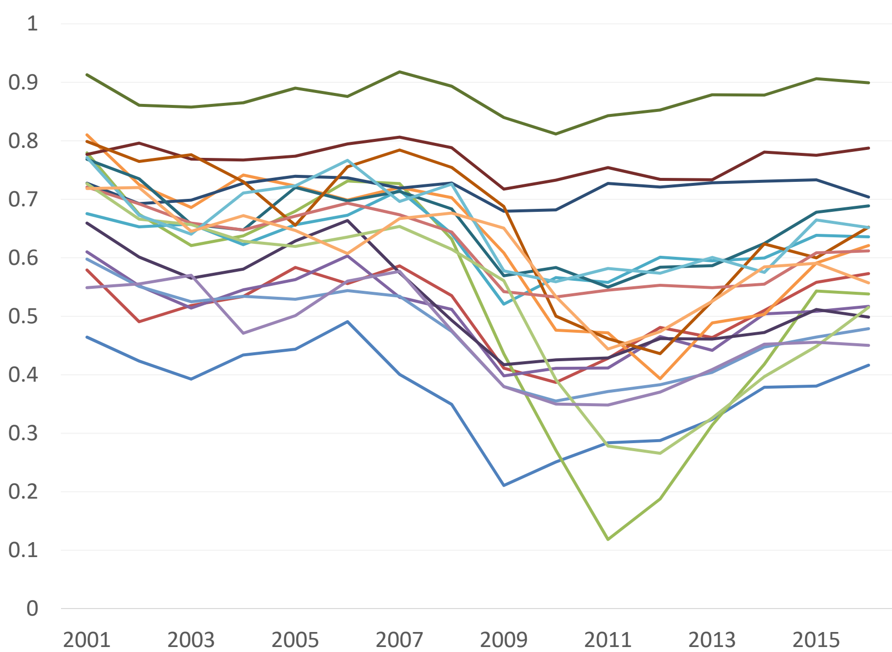

Figure 2: Unemployment Outflow Hazards

Jobseeker perspective: The outflow hazard from unemployment, the fraction of unemployed workers in month $t-1$ in monthly duration category $d-1$ who are not present in month t in category d. There are 6 categories of unemployment and 4 duration categories and thus 18 distinct measures of the unemployment exit rate. Figure 2 shows the raw data for the 18 measures, without attempting to label them.

Employer perspective: The average duration of vacancies, the fraction of vacancies filled in a month (the ratio of vacancies to hires). Figure 3 shows the raw data for 19 industries, including federal and state-local governments.

Both measures rise with tightness. They have entirely different data sources—measurement errors do not impart any common movements to the two. In both cases, the common source of movement is apparent, though the measures differ in their average levels and in the magnitude of the response to the underlying impetus.

Adjusting the measures for level and response to a common driving force reveals a close affinity among the measures. The next two sections of the paper describe a statistical model of the affinity and apply it to the data shown in Figure 2 and Figure 3. To provide a sense of the typical match of the adjusted measure, Figure 4 shows the results, for one of the

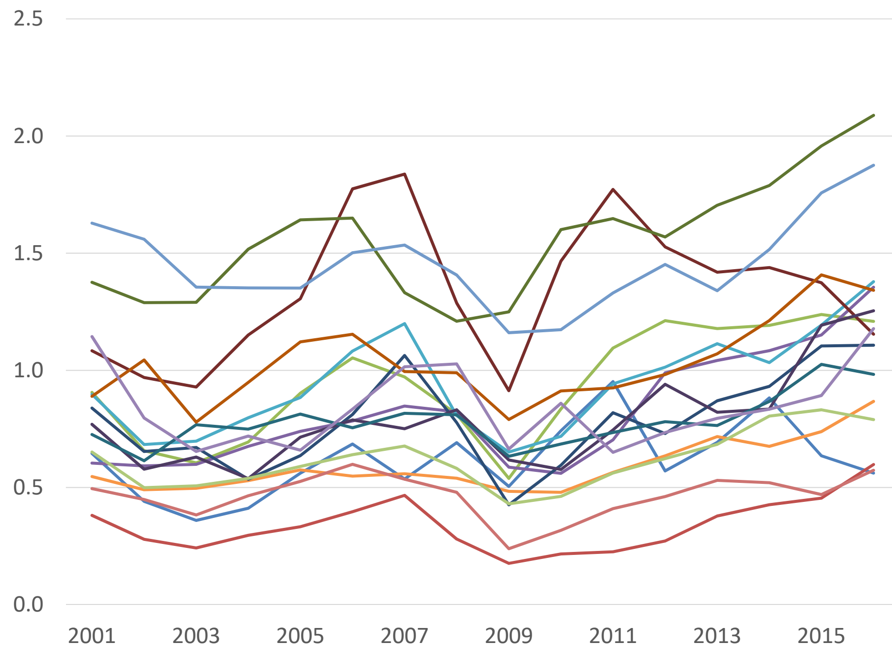

Figure 3: Vacancy Duration by Industry

unemployment exit hazards and one of the vacancy durations. The exit hazard for is for losers of permanent jobs, in the lowest of three unemployment duration categories. The vacancy duration is for manufacturing durables.

The setup in this paper inherits some of the features of Hall and Schulhofer-Wohl (2017). notably the disaggregation of unemployment by the six reasons for entry to unemployment in the Current Population Survey, and the use of the duration of vacancies from JOLTS as the measure of tightness. That paper focuses on the flows from a variety of originating labor-market states into new jobs. Only about a fifth of flows into jobs are out of any of the six unemployment states. Similarly, the outflow hazards from unemployment studied in this paper include large flows out of the labor force as well as flows into jobs. It is a sign of the overwhelmingly pervasive influence of labor-market tightness that the hazards found in the earlier paper are similar to the ones in this paper.

3 Statistical Framework

We consider the following model of the relation between tightness and an underlying unobserved tightness factor, T_t:

$$x_{i,t} = \alpha_i + \gamma_i t + \beta_i T_t. \tag{1}$$

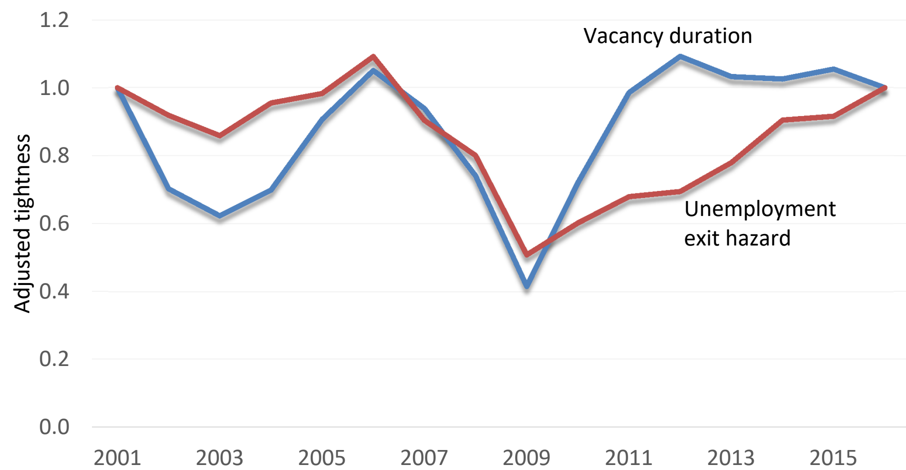

Figure 4: Scaled and Detrended Exit Hazard for Losers of Permanent Jobs and Vacancy Duration In Manufacturing Durables

Here $x_{i,t}$ is tightness measure i in year t, α_i is a measure-specific constant, γ_i is a measure-specific trend, and β_i is a measure-specific response to tightness.

In the classical factor-analysis literature, and in the related vector-autoregression literature, the factor T_t is taken as a serially uncorrelated random variable. With only a single unobserved factor, and the standard assumption that normalizes either the loading coefficient β or the variance of the factor, no issue of identification arises. The intercept and trend coefficient can be estimated by regression. The approach here is different—the aim is to recover T_t as a time fixed effect, common across the measures indexed by i. A substantial issue of identification then arises, as expressed in the following (Δ is the time-difference operator):

Theorem. *Suppose equation* (1) *has a solution* $\{\alpha_i, \gamma_i, \beta_i, T_t\}$. *Then it has a three-dimensional family of solutions indexed by* $\{\lambda, \tilde{T}_1, \tilde{T}_2\}$ *with*

$$\Delta^2 \tilde{T}_t = \frac{\Delta^2 x_{1,t}}{\lambda \beta_1}, t > 2, \tag{2}$$

$$\tilde{\gamma}_i = \Delta x_{i,2} - \lambda \beta_i \Delta \tilde{T}_2, \tag{3}$$

$$\tilde{\alpha}_i = x_{i,1} - \tilde{\gamma}_i - \lambda \beta_i \tilde{T}_1. \tag{4}$$

Any well-chosen restriction on the parameters of dimension at least three will pin down the solution. We use the following:

- Take the initial value T_1 to be the average, across industries, of the observed durations of vacancies in 2001
- Take the terminal value to be the average for the terminal year
- Take the trend growth rates, γ_i, to be zero for all of the vacancy durations—this is intuitive but much stronger than the minimum needed for identification, which could be a single zero restriction

With these identifying assumptions, and the addition of an idiosyncratic random component $\epsilon_{i,t}$ satisfying $\mathbb{E}_{i,t}\ \epsilon_{i,t} = 0$, the parameters can be estimated by nonlinear least squares. Standard errors of the estimates are calculated from the linearization of the model at the estimated values.

4 Data

The Bureau of Labor Statistics' website publishes data from the Current Population Survey on the unemployment count, cross-classified by duration and reason for unemployment. Categories for duration are less than 5 weeks, 5 through 14 weeks, 15 through 25 weeks, and 26 or more weeks. Reasons for unemployment refer to the circumstances that resulted in a person searching actively for work:

- On temporary layoff
- Permanent job losers
- People who completed temporary jobs
- Job leavers
- Re-entrants

- New entrants

From the data on unemployment counts, We calculate approximate continuation rates for unemployment as the number of people remaining unemployed in higher categories as the fraction of those who were counted in the lowest category in the relevant earlier surveys. That category is unemployed less than 5 weeks and is taken as a measure of the inflow to unemployment in this paper. The rules for calculating the approximate denominator for each of the three continuation rates are: in the short-duration category (those counted in this month's survey as unemployed 5 to 14 weeks), the continuation rate is the sum of the inflows one and two months ago; in the medium duration category (those counted in this month's survey as unemployed 15 to 25 weeks), the rate is the inflow 3,4, and 5 months ago; and for the long duration category, unemployed 26 or more weeks, the inflow 6 to 17 months ago.

The unemployment hazard rate is the difference in the continuation rate in one category and the rate in the next lower category, divided by the rate in the lower category. Note that the hazard is a per-category frequency, not a per-month frequency. The choice is purely a normalization—the difference between the two definitions of the hazard is absorbed by the βs.

5 Results

Table 1 shows the estimates of the coefficients β_i reflecting the loading of the unemployment outflow hazard on the unobserved common tightness factor, T_t. The columns refer to the duration categories defined earlier. The estimates are generally around 0.5—an increase in tightness of 0.1 months of vacancy duration corresponds to an increase in the outflow hazard of 5 percentage points. The standard errors are close to 0.3 in all cases, so there is some ambiguity about the size of each one, but none about the strength of the general relation between tightness and the unemployment exit outflow hazards.

Table 2 shows the estimates of the β_is for the 16 private industries distinguished in JOLTS. These estimates are more precise than those for the unemployment exit hazards. Because the unobserved factor T_t is normalized at the beginning and end of the sample at the average of the reported vacancy durations, the loading factors are distributed around one. The goods economy—manufacturing, wholesale trade, and transportation—have coef-

	Short	*Medium*	*Long*
Lost permanent job	0.80	0.68	1.47
	(0.29)	(0.29)	(0.32)
Temp job ended	0.61	0.48	0.66
	(0.29)	(0.29)	(0.29)
Laid off	0.14	0.25	0.30
	(0.28)	(0.28)	(0.28)
Quit	0.65	0.60	0.56
	(0.29)	(0.29)	(0.29)
Re-entrant	0.62	0.49	0.64
	(0.29)	(0.29)	(0.29)
New entrant	0.60	0.58	0.24
	(0.29)	(0.29)	(0.28)

Table 1: Estimated Loading Factors on Tightness, Unemployment Exit Hazard

ficients above one, along with finance and real estate. Entertainment and mining have low coefficients.

Figure 5 shows the estimates of the unobserved tightness factor, interpreted as the duration of vacancies in months of an industry with a loading coefficient of one. The bars mark one standard error above and below the estimates. The factor tracks the business cycle, but with an upward trend. Tightness fell during and after the recession in 2001, rose to a peak in 2007, fell precipitately after the financial crisis, and grew substantially during the recovery.With consideration of the trend, the labor market was considerably slacker in 2009 than in its previous trough, in 2003.

The one-dimensional tightness factor accounts for much of the common movements of the 34 separate measures of tightness from the perspectives of jobseekers and the employers. One way to see this is in terms of the principal components. Table 3 describes the results of extracting the principal components of the tightness measures. It gives the fraction of the variance of the data explained—in the usual sense of this type of analysis—by the 15 principal components, in declining order of the fraction explained. The left column considers both the unemployment outflow hazards and the vacancy durations. The first principal component is essentially a positively weighted average of all of the vacancy durations, with small coefficients, positive and negative, applied to the outflow hazards. In other words, the best single index of the joint movement of all the variables extracts the movement from the

Industry	*Loading coefficient*	*Standard error*
Mining and logging	0.45	(0.20)
Construction	0.83	(0.21)
Durable goods	1.72	(0.23)
Nondurable goods	1.87	(0.24)
Wholesale trade	1.69	(0.23)
Retail trade	0.83	(0.21)
Transportation, warehousing, and utilities	1.49	(0.22)
Information	1.00	(0.21)
Finance and insurance	1.75	(0.23)
Real estate and rental and leasing	1.55	(0.23)
Professional and business services	0.80	(0.21)
Educational services	1.32	(0.22)
Health care and social assistance	1.22	(0.22)
Arts, entertainment, and recreation	0.53	(0.20)
Accommodation and food services	0.94	(0.21)
Other services	0.76	(0.20)

Table 2: Estimated Loading Factors on Tightness, Vacancy Duration

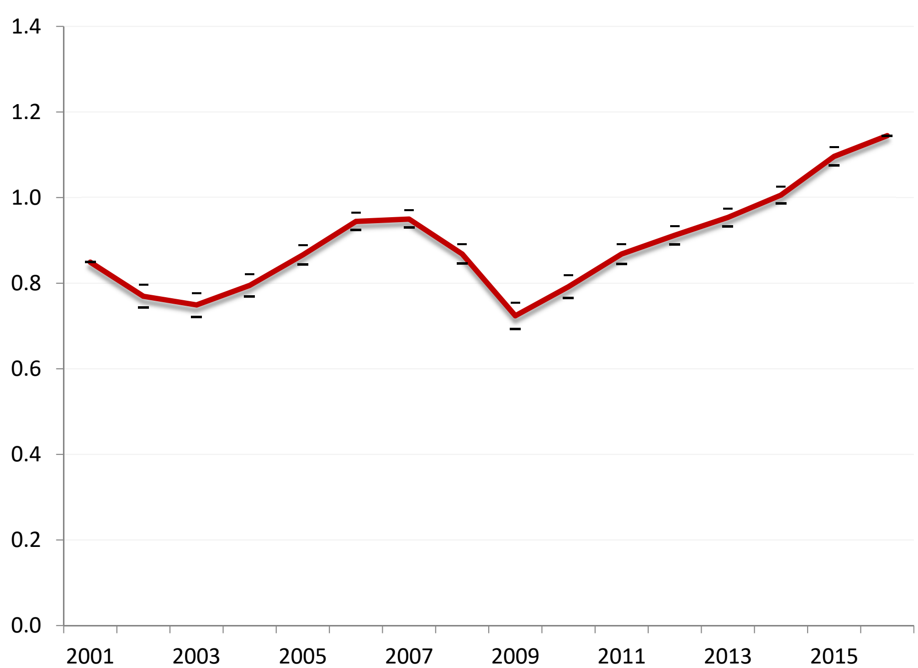

Figure 5: Estimates and Standard Errors of Tightness by Year

Principal component number	*Percent of variance*		
	All 34 measures	*Unemploy-ment exit hazard measures*	*Vacancy duration measures*
1	51.8	87.9	63.8
2	24.6	6.7	17.0
3	11.0	1.5	7.7
4	4.0	1.3	3.2
Sum of higher	8.6	2.6	8.3

Table 3: Principal Components of Tightness Measures

vacancy durations and explains the movements of the hazards using that information. That component explains just over half of the variance of the joint sample. The second principal component is a positively weighted average of the outflow hazards. It explains most of the rest of the joint variance.

The middle column of Table 3 shows that the unemployment outflow hazards have a great deal of common movement—a single component accounts for almost 90 percent of the joint variance. The right columns shows that a single component explains a substantial fraction of the vacancy duration, but there are other movements not associated with the first component that add to almost 40 percent of the variance. In interpreting all of these results, it is important to keep in mind that principal components approach does not separate the role of the idiosyncratic element $\epsilon_{i,t}$ from the principal components. The finding that the first principal component does not explain all of the variance is not cause for rejection of the single-factor model.

The residuals $x_{i,t} - \alpha_i - \gamma_i t - \beta_i T_t$ from the estimated model quantify the disturbances $\epsilon_{i,t}$ in the usual sense of regression residuals. The hypothesis that a single factor dominates the systematic movements of both measures of tightness implies that the true correlation of all pairs of measures, within and across the two types, should be zero. If there is another factor that influences a subset of the measures, and the influence has the same sign for all members of the subset, those the residuals for those measures will be positively correlated. With only 16 observations, the measured correlations have a good deal of sampling variation. Figure 6 shows the distribution of the 561 correlations. There is only a slight tendency for the positive correlations to be more numerous than the negative ones. The obvious candidates for subsets

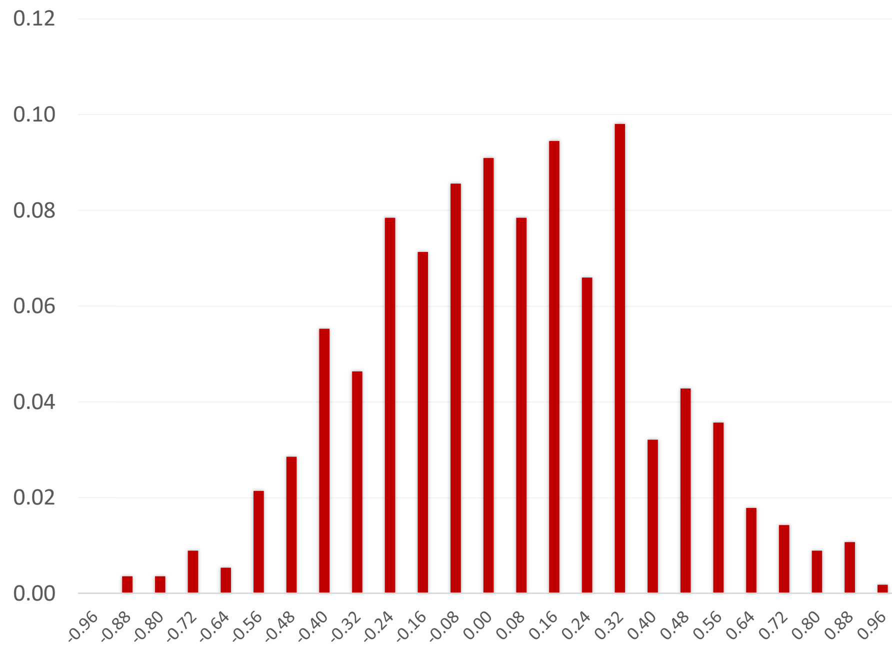

Figure 6: Distribution of Correlations of Residuals

are the outflow hazards from unemployment, one the one hand, and the vacancy durations, on the other. The average correlation among the hazard residuals is 0.271, giving mild support to the idea already mentioned that there are common movements of the hazards not associated with the first principal component. The average correlation among the duration residuals is 0.037, effectively zero, as is the average correlation between a hazard residual and a duration residual, at -0.041.

Another implication of the hypothesis of this paper is that the dispersion of the idiosyncratic residual should be fairly small relative to the dispersion of the corresponding tightness measure. That is, the typical coefficient of multiple correlation, R, should be reasonably large. Figure 7 shows the distribution of R across the 34 equations. The majority of the Rs are above 0.5.

In summary, the hypothesis that 34 different measures of labor-market tightness, 18 from unemployment outflow rates and 16 from durations of vacancies, holds up reasonably well. Most of the measures have significant idiosyncratic variation, but there is little sign of any second factor. To a reasonable approximation some single driving force acts on tightness.

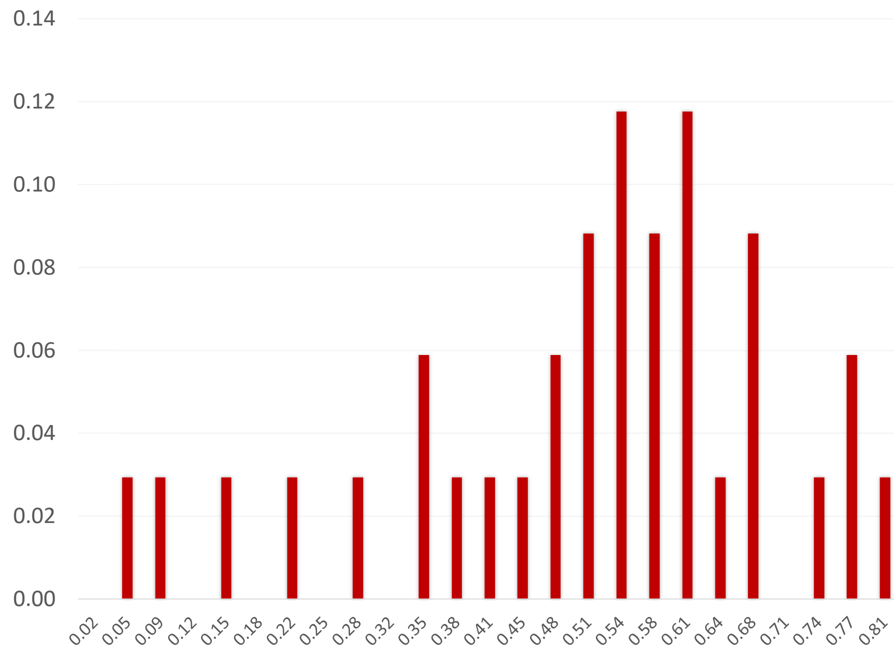

Figure 7: Distribution of Multiple Correlation Coefficient

6 The Role of Tightness in Unemployment Volatility

The second phase of this paper demonstrates that tightness has a major—even dominant—role in the movements of unemployment. An extensive literature on the "ins and outs" of unemployment considers the roles of inflows to and outflows from unemployment. Neither flow has a causal role—both respond, generally in the same direction, to the driving forces of aggregate fluctuations.

A threshold remark is that the inflow is different from the separation flow and the outflow is different from the job-finding flow. Only around 20 percent of people separating from jobs flow into unemployment. The majority take other jobs immediately and another fraction leave the labor force. Similarly, many flowing out of unemployment leave the labor force.

6.1 Displacement

A displaced worker is one who has lost a job previously thought to be permanent. Some definitions restrict attention to those who had at least three years of tenure at the lost job, but the CPS lacks the job history information that would be needed to make that limitation. Those whose current unemployment resulted from loss of a permanent job constitute the closest match to the concept of displacement. The restriction to the reason for the current

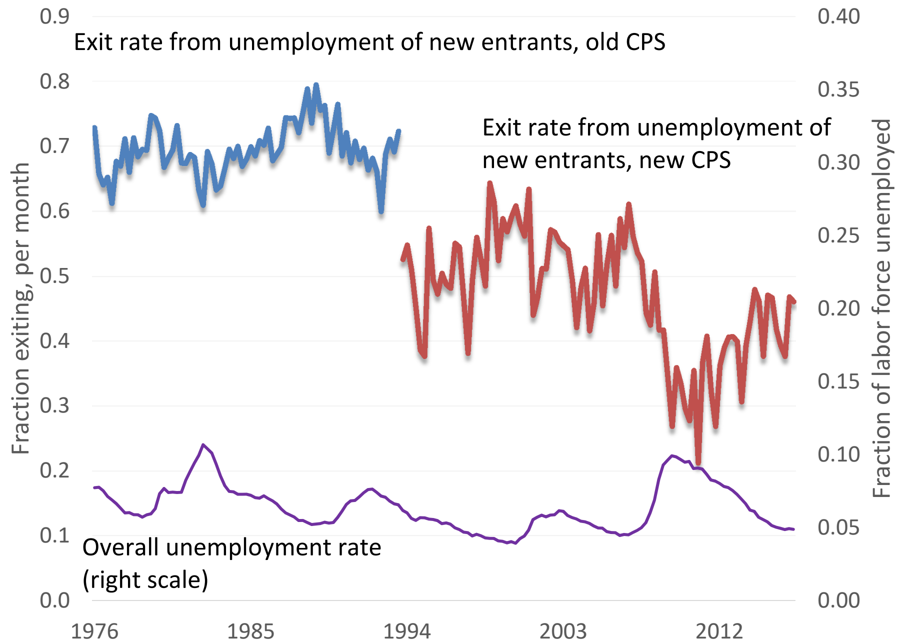

Figure 8: Monthly Exit Rate from Unemployment, Labor-Force Entrants

spell of unemployment results in some blurring of the CPS evidence on displacement. There is growing agreement, supported by the labor-economics literature, that displaced workers tend to have high mobility among the three primary statuses in the labor market: working, looking for work, and spending time out of the labor market. The current status of a job-seeker who experienced displacement in the previous few years may not be a good identifier of the target group because the job-seeker will have rotated through other statuses in the interim, so the displaced workers are scattered among those who held temporary jobs after displacement, those who were laid off from jobs soon after taking them because employers found them unsuitable, those who quit unsatisfactory jobs, and those who left the labor force and have since returned.

There is one category cannot contain displaced workers, namely new entrants to the labor force, comprising mainly individuals who have recently finished school. Job-finding rates of new entrants have roughly the same elasticity with respect to a measure of market tightness as rates for other categories of job-seekers, refuting the extreme hypothesis that the labor market, contrary to the DMP principles, has non-cyclical tightness and that recessions are purely the result of a burst of job loss. Their rate fell in half during the Great Recession. See Figure 8.

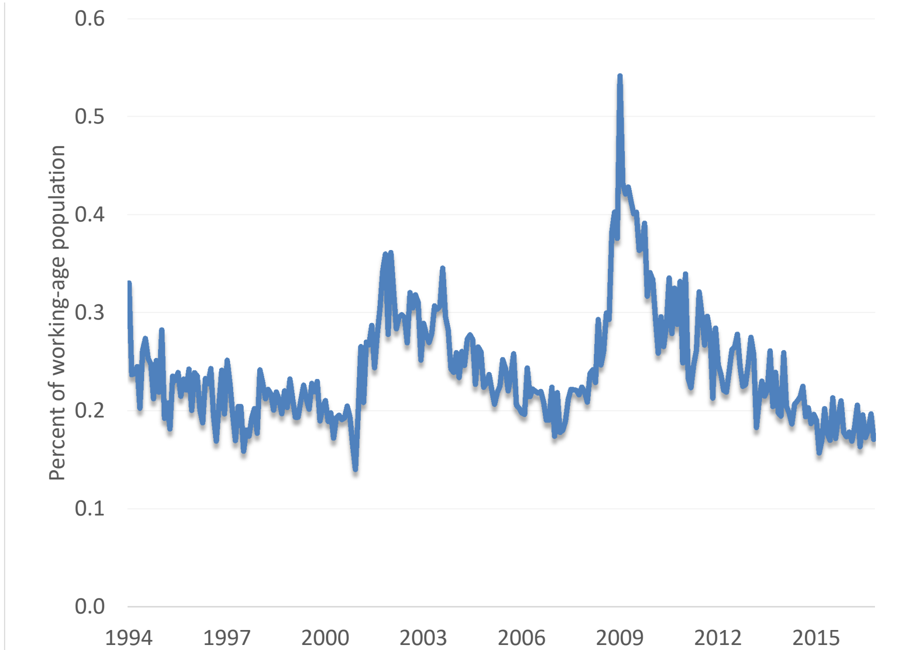

Figure 9: Percent of the Working-Age Population Who Lost a Permanent Job within the Past 4 Weeks, Current Population Survey

The number of individuals in the category *lost permanent job, less than 5 weeks of unemployment* may be the best available measure of displacement. Figure 9 shows the data over the period starting in 1994 when the CPS has included that category.

Many papers have emphasized that much of the volatility of unemployment arises from long-term spells–see Hornstein (2012) in particular. The change in composition by duration arises for two reasons: (1) displacement job loss (measured as loss of permanent jobs in the CPS) becomes a much larger fraction of inflows to unemployment, and (2) outflow rates for all types of unemployed individuals fall, so the spells within each type of unemployment become longer. The switch toward more long-term unemployment does not resolve the question of the relative role of inflows and outflow rates in the overall volatility of unemployment

7 Decomposition with Heterogeneous Inflows to Unemployment

7.1 Framework

The objective is a non-parametric decomposition of unemployment flows into inflows and outflows, based on data on unemployment levels by duration categories and reasons for unemployment. The durations are grouped into categories indexed by j, such as $j = 2$, the category comprising those unemployed 15 to 26 weeks. The decomposition is separate for each reason. Inflows are taken to be the level of unemployment in the shortest duration category, numbered $j = 0$, with less than 5 weeks of unemployment to date. Outflows are measured as the continuation rate $c_{i,j,t}$, the fraction of the inflow in earlier weeks now counted in duration category j who are still unemployed. Note that the continuation rate is the empirical survival rate in unemployment from the beginning of the spell, not the monthly outflow hazard from unemployment studied earlier in the paper. The inflow category has $c_{i,0,t} = 1$. The earlier inflow is denoted $z_{i,j,t}$ and is measured as the number of people in the lowest duration category now counted in duration category j, divided by the working-age population in the current month, t (not the population earlier when they flowed into unemployment). This convention simplifies the math by eliminating the population from any of the equations but taking exact account of its role. We use population as the scaling factor for unemployment rather than the labor force, to avoid influences from unemployment's role in the labor force and from the decline in the labor force that accelerated after the crisis recession. The conclusions would be essentially the same with the conventional definition of the unemployment rate, however.

This framework has some elements in common with Ahn and Hamilton (2016), particularly its focus on continuation rates revealed in unemployment counts within duration groups and its use of unfiltered data. But its approach to heterogeneity is complementary to theirs. They model heterogeneity in terms of a pair of unobserved types that differ in outflow rates. Their main results use aggregate unemployment disaggregated by duration alone, whereas this paper uses data disaggregated by the six categories by reason distinguished in the CPS and by duration. This paper does not tackle the tricky and controversial task of dealing with heterogeneity within each of the six categories. It is clear that distinguishing the six categories takes account of a great deal of heterogeneity.

With this notation, the unemployment/working-age-population ratio for unemployment reason i satisfies

$$u_{i,t} = \sum_j c_{i,j,t} z_{i,j,t} \tag{5}$$

by construction. The linearization around the time-invariant path $\{\bar{z}_{i,j}, \bar{c}_{i,j}\}$, the sample means of the inflows and outflow rates, is

$$u_{i,t} = \sum_j \left[\bar{c}_{i,j}\bar{z}_{i,j} + \bar{c}_{i,j}(z_{i,j,t} - \bar{z}_{i,j}) + \bar{z}_{i,j}(c_{i,j,t} - \bar{c}_{i,j})\right] + \eta_{i,t}. \quad (6)$$

The linearization makes a four-component decomposition of unemployment comprising, for each reason i, (1) a constant; (2) an *inflow component* corresponding to the counterfactual where outflow rates are constant but inflows have their actual values; (3) an *outflow component* corresponding to the reverse counterfactual, where inflow rates are constant but inflows have their actual values; and (4) a small *interaction component*, $\eta_{i,t}$, that arises from nonlinearity. The inflow and outflow components have means zero, but the interaction component has a nonzero mean to the extent that c and z are correlated.

The data contain large seasonal fluctuations. Rather than use seasonally adjusted data, We do all of the calculations in the original monthly data, then present the results as annual averages.

7.2 Results

Figure 10 shows the results for the first three reason categories that have been distinguished in the Current Population Survey since 1994, *loss of permanent job*, *temporary job ended*, and *on layoff*. For each reason category, the graph on the left shows the inflow component (heavy blue line) together with the actual unemployment rate (lighter red line) and the graph on the right shows the outflow component in the same format, along with the interaction component, shown as a thin black line. The vertical axis is in percentage points of the working age (16 plus) population. The sample period contains two recessions, one in 2001, the *tech recession* and the other from the end of 2007 until mid-2009, the *crisis recession*. Before 1994, the CPS did not distinguish loss of a permanent job from the ending of a temp job. The 1994 revision of the CPS also introduced discontinuities in other categories. We have calculated similar results for the unemployment-reason categories in the CPS from 1976 through 1993. The results for that period reveal important changes that occurred prior to 1994, but do not alter the basic conclusion of this paper about the relative roles of displacement and market-tightness in the U.S. labor market over the entire period starting in 1976. For the tech recession, the discussion compares 2002 to 2000, and for the crisis recession, 2010 to 2007. In the tech recession, total unemployment as a fraction of the population rose by 1.2 percentage points, while in the crisis recession, it rose by 3.2 percentage points.

Figure 10(a) and Figure 10(b) show unemployment resulting from the loss of a permanent job—the category that includes displacement. Displacement-related unemployment was a significant contributor to the bulge in overall employment in the tech recession (0.72 percentage points out of a total of 1.17) but a smaller fraction of the much larger increase in overall unemployment in the crisis recession (1.93 out of 3.18). The inflow component rose by 0.45 percentage points in the tech recession and by 0.58 percentage points in the crisis recession. The outflow component, shown in the right-hand graph, is smaller than the inflow component in the tech recession (0.35 points compared to 0.45 points) and much more in the crisis recession (1.00 points compared to 0.58 points). However, the story is complicated by the interaction component of 0.35 points that cannot be split between the components. The interaction component is almost negligible for all other categories and all other episodes within the loss of permanent job category. It arises because of uniquely high coincident increases in both the inflow and outflow components during the crisis recession. Thus displacement-category unemployment rose so high and remained high so long—in the sense used in this paper and the earlier related literature—because of a rise in inflows to that category and an increase in the continuation rate for unemployment in that category. Recall that the continuation rate is measured separately by duration group, so in principle it does not reflect the shift in composition to individuals in higher duration categories with higher continuation rates. But the absence of a breakdown for duration within the existing 6-months plus category somewhat clouds this principle.

Figure 10(c) and Figure 10(d) describe unemployment resulting from the ending of temporary jobs. The graph on the left showed that almost no changes occurred in the inflow component in either recession. As shown on the right, essentially all of the increase in this category arose from lengthening of unemployment spells. The contribution to the unemployment increase in the crisis recession was 0.21 percentage points of the population.

Figure 10(e) and Figure 10(f) continue with the layoff category—people not working but expecting to return to existing jobs. In both recessions, the increase in the layoff inflow component was higher than the increase in the outflow. In the tech recession, the inflow was 0.07 percentage points and the outflow was 0.05, while in the crisis recession, the inflow rose by 0.09 and the outflow by 0.08. In the single year 2009, the inflow component spiked at 0.14 while the outflow was more stable at 0.06. In the recessions between 1976 (when

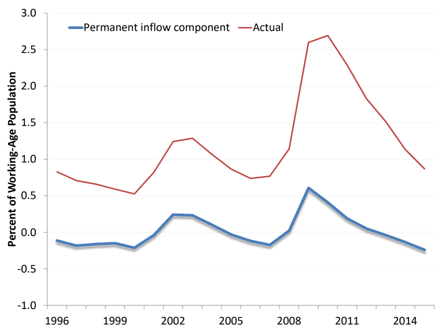

(a) Inflows, loss of perm job

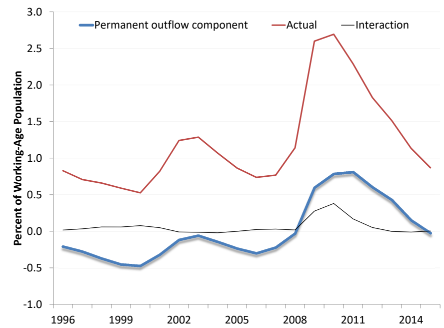

(b) Outflow rates, loss of perm job

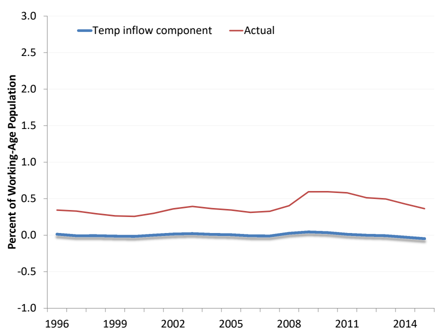

(c) Inflows, temp job ended

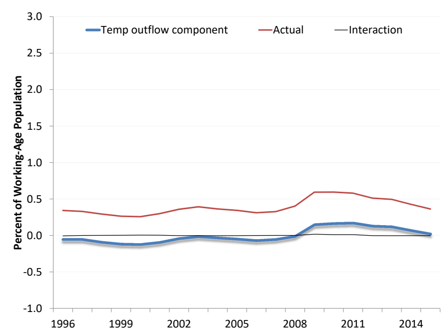

(d) Outflow rates, temp job ended

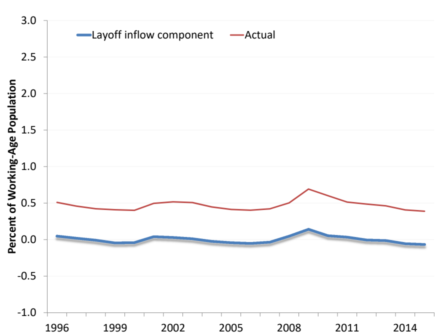

(e) Inflows, layoff

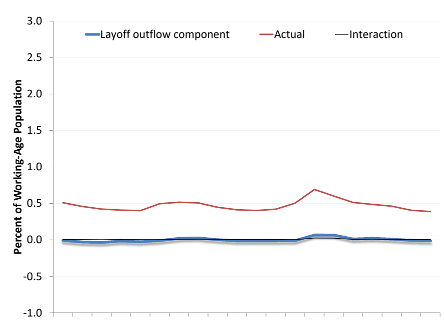

(f) Outflow rates, layoff

Figure 10: Inflow and Outflow Components, Unemployment by Reason of Loss of Permanent Job, End of Temp Job, and Layoff

layoff unemployment was first reported separately in the CPS) and 1994, layoffs were a larger factor in recession unemployment bulges.

Figure 11 shows similar graphs for the remaining three categories of unemployment by reason. Figure 11(a) and Figure 11(b) refers to unemployment following a quit. The unemployment rate from quits, shown by the thin red line at the top, was roughly constant at 0.38 percentage points of the working-age population. The constancy arises from substantial but mostly offsetting cyclical movements of the inflow and outflow components. In the tech recession, the inflow component fell by 0.06 percentage points while the outflow component rose by the same amount. In the crisis recession, the inflow component rose by 0.15 points and the outflow component by 0.24. The slackening of the labor market in the crisis recession considerably raised the continuation rate in unemployment among those whose unemployment spells began with a quit.

Figure 11(c) and Figure 11(d) shows the striking difference between the movement of the components of unemployment in the case of spells that began with re-entry to the labor force. In normal times, this category and the first, loss of permanent job, account for the largest of the six categories, averaging about one percent of the working-age population each. Jointly they account for more than half of the normal overall unemployment rate of 3.5 percent of that population. The inflow component declined steadily, with little cyclical movement—it fell by 0.04 percentage points in the tech recession and by 0.15 points in the crisis recession. Only the latter is visible in the graph, but it is small relative to the large size of the component. The outflow component rose by 0.17 percentage points in the tech recession and by 0.79 points in the crisis recession. Re-entrants behaved qualitatively like quitters—their inflow components fell in recessions and their outflow components rose, but the rise was much bigger for the large category of re-entrants.

Figure 11(e) and Figure 11(f) describe the final component, new entrants to the labor force. In normal times, this category accounts for about 0.3 percentage points of unemployment as a fraction of the population. The tech recession resulted in only small changes in the inflow and outflow components. In the crisis recession, the inflow component rose slightly, by 0.03 points, but the outflow component rose significantly relative to normal, by 0.21 points.

To summarize with respect to the breakdown between the inflow and outflow components, in the tech recession, the inflow component rose more than the outflow component in the lost-permanent-job category (by 0.10 percentage points of the population) and in the layoff

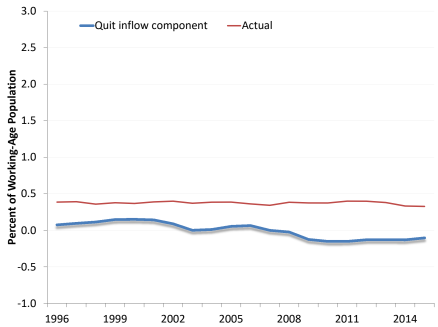

(a) Inflows, quit

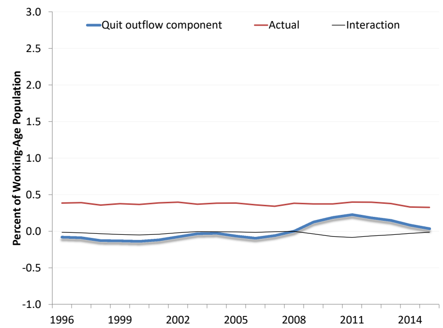

(b) Outflow rates, quit

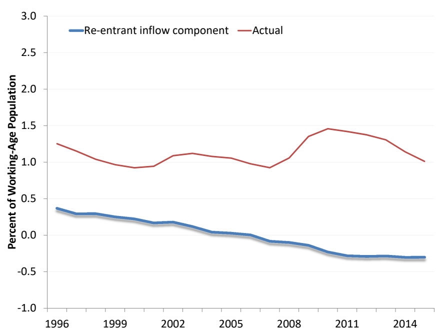

(c) Inflows, re-entrant

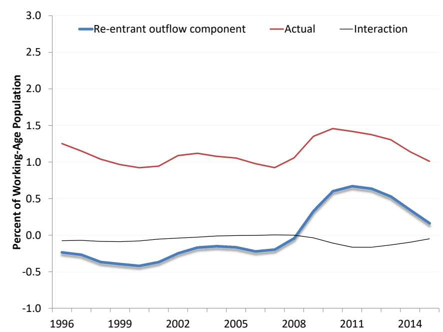

(d) Outflow rates, re-entrant

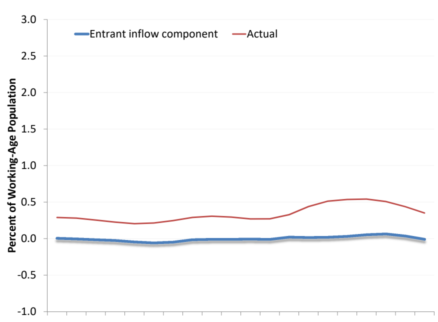

(e) Inflows, new entrant

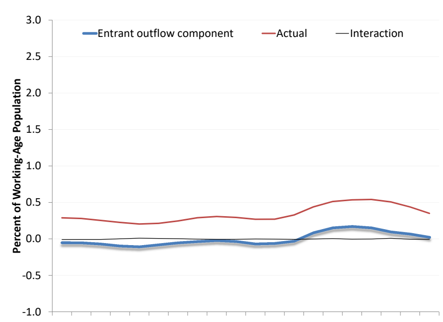

(f) Outflow rates, new entrant

Figure 11: Inflow and Outflow Components, Unemployment by Reason of Quit, Re-entry, and New Entry

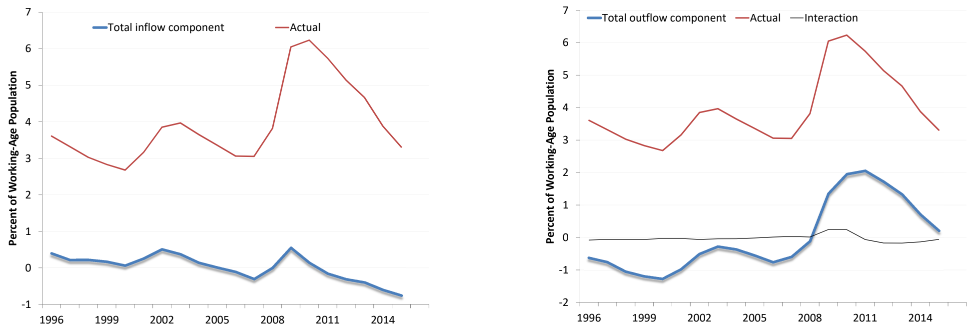

(a) Inflows (b) Outflows

Figure 12: Inflow and Outflow Components for Total Unemployment

category (by 0.02 percentage points). In the other four categories, the outflow component rose more than the inflow component. In the crisis recession, the pattern was the same, except that the increase in the outflow component for loss of permanent job was considerably larger than the increase in the inflow component

7.3 Aggregation

The natural form of aggregation is to add the six categories described above. It is natural because each is stated as the ratio of the category to the entire working age population, so the aggregate is interpreted directly as the overall fraction of the population. Figure 12 shows the sums of the components for the six categories. The inflow component shows pronounced increases in both recessions—0.44 percentage points of the population in both cases. But the outflow component increased by more in both recessions—0.76 percentage points in the tech recession and 2.53 points in the crisis recession. The sum of the interaction terms across the six categories is 0.21, so there is no significant ambiguity about the conclusion that, in the crisis recession, a slacker labor market—in the sense of higher continuation rates in unemployment—were the dominant source of the rise in unemployment. Further, as the figure reveals, the bulge in inflows to unemployment was not nearly as persistent as the bulge in outflow unemployment.

Plainly, the outflow component, corresponding to changes in outflow rates associated with labor-market tightness, are quite dominant in overall volatility. The theory of tightness

	Standard deviations, percentage points of population					
	Percent <5 weeks	*Inflow*	*Outflow*	*Residual*	*Sum*	*Actual*
Layoff	55	5.0	2.7	0.6	8.3	7.5
Permanent	24	22.2	40.9	10.3	73.4	65.6
Temp	33	2.0	9.5	0.6	12.2	10.7
Quit	37	10.7	11.8	2.4	24.8	2.1
Re-entrant	30	22.8	37.4	5.2	65.4	17.1
Entrant	31	3.2	8.6	0.5	12.3	11.3
Total		35.3	109.2	10.7	155.2	107.8

Table 4: Summary of Results

in the DMP model is highly relevant for understanding the movements of unemployment. Variations of inflows matter as well.

Adding the categories together leaves the interaction effect allocated to neither inflow nor outflow rates. Only in the case of loss of permanent jobs does this matter at all. The sum of the interaction effects is smaller than the effect for loss of permanent jobs because the interaction effect for quits is negative.

8 Ahn and Hamilton

Ahn and Hamilton (2016) seek to distinguish genuine duration dependence from unobserved cross-sectional heterogeneity. The challenge to identification under this assumption is the subject of a large literature over the past 40 years–see Alvarez, Borovičková and Shimer (2016) for an extensive discussion and results for Austrian workers with multiple spells. Ahn and Hamilton argue that one can separate the two influences by looking at duration categories, under certain assumptions. They emphasize the importance of not detrending the data and thus considering nonstationary driving forces.

Ahn and Hamilton posit two types of individuals, H (fast exiters) and L (slow). Inflows and outflows evolve as random walks. At every forecast horizon, inflows account for more of the forecast variance than do outflows.They find that type L workers are more important than type H workers, and account for 90 percent of the uncertainty in predicting unemployment 2 years ahead. They also find that heterogeneity is more important than genuine duration-dependence. They summarize their results as "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."

They find that the contribution to unemployment of duration dependence is tiny. The two big factors in recessions are increased inflows to unemployment, especially of the L-type unemployed (Figure 5) and decreased outflow rates of both types (their Figure 4). They go on to study unemployed individuals who lost permanent jobs to confirm their findings. The results in their Figure 12 confirm Hall and Schulhofer-Wohl's (2017) results showing lower job-finding rates among the unemployed who had suffered the permanent loss of a job compared to other categories of the unemployed.

Ahn and Hamilton's finding that changes in inflow to unemployment are the more important of the two flows disagrees substantially with the finding of this paper—that cyclical changes in outflow rates dominate the movements of overall unemployment. Tracking down the sources of this disagreement is under way.

9 Sources of Cyclical Fluctuations in Unemployment

This paper supports the view that cyclical changes in labor-market tightness are the main driving force of unemployment, though it gives cyclical fluctuations in inflows a chance and they do have an important transitory role, as shown in Figure 12(a). Thus the DMP model's basic idea about labor-market tightness seems highly relevant. The idea applies to all driving forces that change the incentives to create jobs. In the earlier DMP literature, the driving force as taken to be productivity. An increase in workers' productivity not fully matched by an increase in wages enhanced the present value J of a newly hired worker. The model maps J into tightness through a simple and intuitive equation,

$$\kappa \cdot \frac{1}{q} = \kappa \cdot T = J. \tag{7}$$

This zero-profit condition says that employers adjust recruiting effort to the point that the marginal vacancy posting just pays for itself in present value. On the left side, κ is the per-period cost of maintaining a vacancy and $1/q$ is the expected number of periods it takes to fill a vacancy, denoted T earlier in the paper. q is the period-period probability of filling a vacancy, which varies negatively with tightness. The job value, J, is the difference between the present value of the worker's contribution to revenue less the present value of the wage.

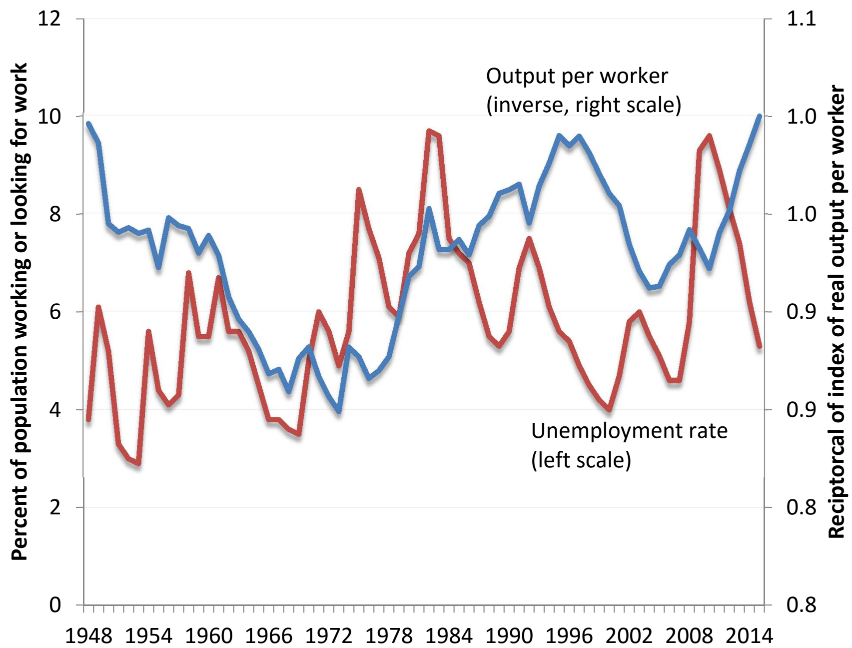

Figure 13: Unemployment and Detrended Output per Worker, 1948 to 2015

When some force causes J to rise, recruiting becomes sufficiently more aggressive to tighten the market by increasing T and thus raising the expected cost of hiring one new worker.

The question about the model that is unresolved today, more than 20 years after the publication of the canon of the model, Mortensen and Pissarides (1994), is: What force depresses the payoff to job creation in recessions? In that paper, and in hundreds of successor papers, the force is a drop in productivity. But unemployment does not track the movements of productivity in U.S. data, as Figure 13 (from Hall (2017)) shows.

A more promising line of thought starts from the observation that the valuation of investments, financial and real, seems to be quite volatile relative to the realizations of their subsequent values. A line of research culminating in Campbell and Shiller (1988), recently surveyed in **?**, focuses on the volatility of the stock market. Occasionally, Investors appear to become apprehensive about holding claims to future payoffs. Modeling in the Campbell-Shiller tradition interprets their actions in terms of an increase in the time-discount rate, so the effect is greater for payoffs in the more distant future. Presumably the higher discounts immediately visible in a crisis in the stock market also apply to the payoffs from real investments in job creation and physical capital. A term sometimes applied is in that setting is "loss of confidence." Figure 14 suggests a fairly close association, especially in recent decades, between unemployment and the stock market.

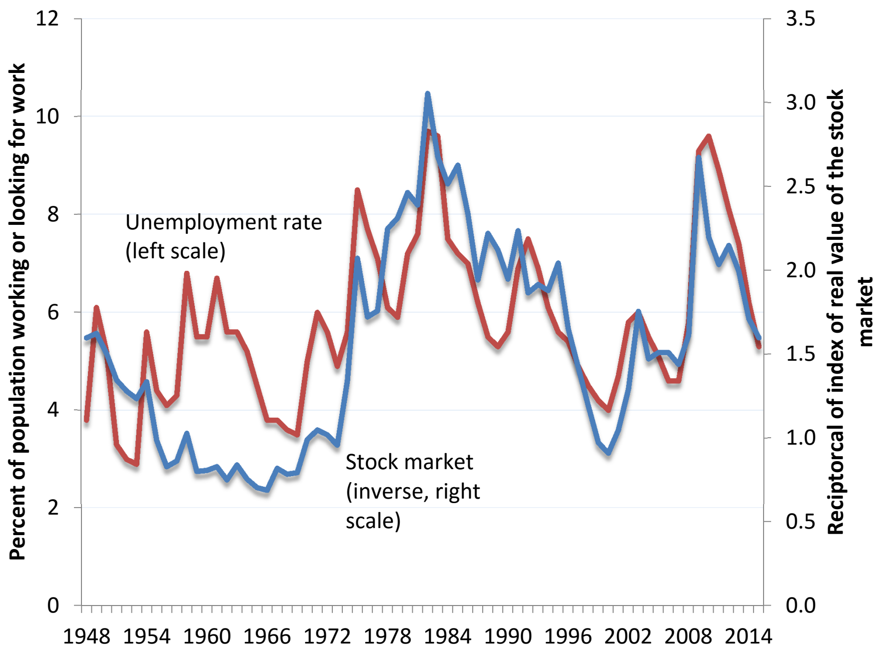

Figure 14: Unemployment and the Real, Detrended Value of the S&P Stock-Market Index, 1948-2015

10 Concluding Remark

The data point strongly in the direction of an aggregate driving force operating mainly through labor-market tightness, quite uniformly across industries and categories of unemployment. One such force is confidence, possibly measured from financial discounts, but other aggregate forces are still in the running.

References

Abraham, Katharine G. and Lawrence F. Katz, "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?," *Journal of Political Economy*, June 1986, *94* (3), 507–522. Part 1.

Ahn, Hie Joo and James D. Hamilton, "Heterogeneity and Unemployment Dynamics," August 2016. Department of Economics, University of California, San Diego.

Alvarez, Fernando, Katarína Borovičková, and Robert Shimer, "Decomposing Duration Dependence in a Stopping Time Model," 2016. University of Chicago.

Campbell, John Y. and Robert J. Shiller, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1988, *1* (3), 195–228.

den Haan, Wouter J., Garey Ramey, and Joel Watson, "Job destruction and the experiences of displaced workers," *Carnegie-Rochester Conference Series on Public Policy*, 2000, *52*, 87 – 128.

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

Hall, Robert E., "Job Loss, Job Finding, and Unemployment in the U.S. Economy over the Past Fifty Years," *NBER Macroeconomics Annual*, 2005, *20* (1), 101–137.

——, "High Discounts and High Unemployment," *American Economic Review*, February 2017, *107* (2), 305–330.

—— and Sam Schulhofer-Wohl, "Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Jobseekers," February 2017. Hoover Institution, Stanford University.

Hornstein, Andreas, "Accounting for Unemployment: The Long and Short of It," Technical Report, Federal Reserve Bank of Richmond 2012. Working Paper 12-07.

Lilien, David M., "Sectoral Shifts and Cyclical Unemployment," *Journal of Political Economy*, 1982, *90* (4), 777–793.

Mortensen, Dale T. and Christopher Pissarides, "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 1994, *61*, 397–415.

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

Shimer, Robert, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 2005, *95* (1), 24–49.