

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@frbchi.org

July 5, 2017

Abstract

A distinctive contribution of the unemployment model of Diamond, Mortensen, and Pissarides is the creation of an economically coherent concept of labor-market tightness. In a tighter market, jobseekers find jobs more quickly and employers take longer to fill jobs. The evidence in this paper on individual job-finding rates and on job-filling rates of employers shows that a single index factor describes tightness. Jobfinding rates and jobfilling rates move in parallel as functions of the single index. We account for heterogeneity in both rates. For the jobfinding rate, we use success rates for jobseekers over one-month and 12-month periods of potential search in data from the Current Population Survey, and we distinguish 16 categories of jobseekers, based on their activities and experiences leading up to the measurement date. For the jobfilling rate, we use the ratio of hires to job openings in the Job Openings and Labor Turnover Survey, distinguishing between 16 industries. We conclude that the DMP model's concept of tightness is central to the understanding of fluctuations in unemployment and thus in employment and output. We show that our index of labor-market tightness is highly correlated with movements of total hours of work and substantially correlated with movements of aggregate output, although noncyclical influences—total factor productivity and labor-force participation—are important sources of output volatility as well.

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 labor-market volatility rests on the central concept of market tightness. In a tighter market, jobseekers find jobs more quickly and employers fill jobs more slowly. Diamond, Mortensen, and Pissarides’s model of unemployment (Mortensen and Pissarides (1994)) provides a rigorous framework for understanding tightness, a previously murky concept. Tightness varies in response to changes in the incentives facing employers to create jobs. We show that tightness is at the heart of fluctuations in unemployment and employment, and that variations in tightness have effects that pervade the entire labor market. These variations have a large role, along with stochastic variations in productivity growth, in the volatility of output. No model of business-cycle fluctuations can claim realism without embodying an account of tightness. We also show that tightness moves in unison across the entire U.S. labor market—a single latent factor characterizes the movements of the jobfinding process for many types of workers and recruiting success in many private industries.

A key object in the DMP model is the matching function, which shows how the labor market maps stocks of jobseekers and recruiting efforts into flows of hires. We develop a matching-function specification that recognizes heterogeneity among types of jobseekers and among employers in different industries. Jobseeking success rates respond positively to tightness and recruiting success rates respond negatively. From data from the Current Population Survey, we measure coefficients relating jobfinding rates to tightness, modeled as a latent statistical factor. From the Job Openings and Labor Turnover Survey, we measure coefficients relating jobfilling rates by industry to the same latent factor. We exploit the statistical principle that, in a body of indicators that each respond to a single latent common factor plus an idiosyncratic random factor, the common latent factor is the first principal component of the indicators.

The paper then turns to the role of fluctuations in tightness in the determination of unemployment, total labor input, and output. Recessions in general and the last one in particular involve sharp increases in unemployment followed by protracted declines. The DMP model treats unemployment as a state variable, responding to tightness, which is a jump variable, through a law of motion that involves lags. Nonetheless, unemployment tracks tightness quite accurately in annual data. Total labor input in the economy is also highly correlated with tightness, even though fluctuations in labor-force participation and variations in hours per worker are not intrinsically connected to tightness but do influence labor input.

And the relation of tightness to output fluctuations is strong, despite the importance of fluctuations in the growth of total factor productivity.

1 Tightness as a Sufficient Statistic for Jobfinding and Jobfilling Rates

1.1 The standard DMP setup

We start from the constant-returns matching function relating the flow of hires H to the stocks of jobseekers N and vacancies V :

$$H = m(N, V). \tag{1}$$

Tightness, called θ , is defined as the ratio of vacancies to jobseekers :

$$\theta = \frac{V}{N}. \tag{2}$$

The jobfinding rate f is:

$$f(\theta) = \frac{H}{N} = m(1, \theta). \tag{3}$$

The jobfilling rate q is:

$$q(\theta) = \frac{H}{V} = \frac{f(\theta)}{\theta}. \tag{4}$$

Theorem 1. *In the standard DMP setup, with fixed matching function, tightness θ is a sufficient statistic for the jobfinding rate $f(\theta)$ and the jobfilling rate $q(\theta)$.*

In other words, two observed variables, the job-finding rate and vacancy-filling rate, are functions of only a single variable, the tightness measure, taken as the vacancy/unemployment rate, θ .

1.2 Using T as the measure of tightness with one type of jobseeker

In this setup, tightness is defined as vacancy duration:

$$T = \frac{V}{H}. \tag{5}$$

Constant returns of the matching function implies

$$1 = m\left(\frac{N}{H}, \frac{V}{H}\right) = m\left(\frac{1}{f(T)}, T\right). \tag{6}$$

This defines the jobfinding rate function $f(T)$ implicitly. The job-filling rate is

$$q(T) = \frac{H}{V} = \frac{1}{T}. \quad (7)$$

Theorem 2. *In the DMP setup with vacancy duration T serving as the measure of tightness, with fixed matching function and one type of jobseeker, tightness T is a sufficient statistic for the jobfinding rate $f(T)$ and the jobfilling rate $q(T)$.*

Again, f and q are functions of a single measure of tightness, T .

1.3 Multiple jobseekers and multiple types of employers

Hall and Schulhofer-Wohl (2017) introduced a matching function that accommodates multiple types of jobseekers. Now we consider a generalization of that matching function that also accommodates multiple industries. Let i index heterogeneous categories of jobseekers with jobfinding rates $f_i(T)$ and let j index heterogeneous employers with jobfilling rates $q_j(T)$. Both depend on a measure of tightness, T , that is not observed directly. The function $f_i(T)$ is continuous and increasing, with $f_i(0) = 0$ and $\lim_{T \rightarrow \infty} f_i(T) = 1$ and the function $q_j(T)$ is continuous and decreasing, with $q_j(0) = 1$ and $\lim_{T \rightarrow \infty} q_j(T) = 0$. The vector \mathbf{N} gives the number of jobseekers of each type and the vector \mathbf{V} gives the number of vacancies for each type of employer.

Definition: A constant-returns matching function with heterogeneous jobseekers and employers is a function $m(\mathbf{N}, \mathbf{V})$ giving the volume of matches as a function of the vectors of jobseekers and vacancies that is first-degree homogeneous and semi-strictly increasing in the vector $\mathbf{X} = [\mathbf{N}, -\mathbf{V}]$.

Theorem 3. *Let*

$$G(\mathbf{X}, T) = \sum_i N_i f_i(T) - \sum_j V_j q_j(T). \quad (8)$$

For any \mathbf{X} , there exists a unique value of tightness, $T(\mathbf{X})$, such that $G(\mathbf{X}, T(\mathbf{X})) = 0$. $m(\mathbf{X}) = \sum_i N_i f_i(T(\mathbf{X}))$ is a constant-returns matching function.

Proof: $G(\mathbf{X}, 0) = -\sum_j V_j < 0$. $\lim_{T \rightarrow \infty} G(\mathbf{X}, T) = \sum_i N_i > 0$, so there is a $T^*(\mathbf{X})$ with $G(\mathbf{X}, T^*(\mathbf{X})) > 0$. $G(\mathbf{X}, T)$ is continuous in T , so there is a $T(\mathbf{X})$ such that $G(\mathbf{X}, T(\mathbf{X})) = 0$ by the intermediate value theorem. m gives the volume of matches because it equals the number of jobs found, which by construction equals the number of jobs filled. $T(\mathbf{X})$ is unique

because $G(\mathbf{X}, T)$ is strictly increasing in T . m is first-degree homogeneous because a scaling of \mathbf{X} leaves $T(\mathbf{X})$ unchanged and raises $\sum_i N_i f_i(T(\mathbf{X}))$ in proportion. \square

Theorem 4. *In the setup with multiple types of jobseekers and multiple types of employers, and with fixed parameters of the jobfinding rate functions $f_i(T)$ and the jobfilling rate functions $q_j(T)$, tightness $T = T(\mathbf{X})$ is a sufficient statistic for $f_i(T)$ and $q_j(T)$.*

We take the local linear approximation,

$$f_i(T) = \alpha_i^f + \beta_i^f T \tag{9}$$

and

$$q_j(T) = \alpha_j^q - \beta_j^q T. \tag{10}$$

In many cases, we stack the parameter vectors as α and β .

Theorem 5. *In the setup with multiple types of jobseekers and multiple types of employers, and with fixed parameters α and β , T is a sufficient statistic for the jobfinding rates, $\alpha_i^f + \beta_i^f T$, and the jobfilling rates, $\alpha_j^q - \beta_j^q T$.*

1.4 Implications

Equilibrium in the DMP class of models occurs at the point of zero profit

$$\kappa \cdot T = J$$

Here J is the capital gain an employer receives for a successful hire. It is the difference between the present value of the new worker's contribution to revenue, P , and the present value of the worker's wage, W :

$$J = P - W. \tag{11}$$

Under the assumption of parameter stability, the sufficient-statistic property of T has the strong implication that only T belongs on the right-hand side of an equation for the jobfinding and vacancy-filling rates—all other variables are excluded. Joint fluctuations in jobfinding and vacancy duration result from changes in the job value J and not from any shift of the matching function. For example, there cannot be seasonal shifts in the jobfinding or jobfilling rates arising from the presence of seasonal shifts of the matching function—any seasonality of jobfinding and jobfilling must result from contemporaneous seasonal movements of J . These

implications rely on the zero-profit condition and thus need not hold for employers to whom a zero-profit condition does not apply. For this reason, we will apply our model only to data on private-sector vacancies and exclude government vacancies.

2 Statistical Framework

We consider the following model of the relation between an indicator (jobfinding or jobfilling rate) and the underlying unobserved tightness factor, T_t :

$$x_{i,t} = \alpha_i + \beta_i T_t + \epsilon_{i,t}. \quad (12)$$

Here $x_{i,t}$ is indicator i in month t , α_i is a measure-specific constant, β_i is the response of measure i to tightness, T_t is the latent tightness index discussed earlier, and $\epsilon_{i,t}$ is an idiosyncratic component, including measurement error, assumed to be uncorrelated over time and across tightness measures. We let i run over both the jobfinding series and the jobfilling series. This statistical model imposes strong restrictions on the data. In particular, to the extent that there are secular trends or predictable seasonality in the tightness data, those trends must enter the data through underlying tightness T_t ; the data series should not be separately detrended or seasonally adjusted before we estimate the common factor T_t . This strong restriction is consistent with our theory of the matching function with multiple categories of jobseekers and employers, under which trends or seasonality in productivity, labor demand, reservation wages, or any other variable affecting jobseekers' or firms' behavior affect job-finding rates or vacancy duration only by inducing trends or seasonality in aggregate tightness T_t . We test this restriction below.

If the variance of $\epsilon_{i,t}$ is the same across measures i and over time, and if $\epsilon_{i,t}$ is uncorrelated across i and t , as we have assumed, T_t is the first principal component of the collection of time series $\{x_{i,t}\}$. The first principal component is defined as the eigenvector corresponding to the largest eigenvalue of the variance-covariance matrix of $\{x_{i,t}\}$. Under the assumption that $\epsilon_{i,t}$ is homoskedastic and uncorrelated across i and t , there are only two eigenvalues, the larger of which is $\sigma_\epsilon^2 + \mathbb{E}(T_t^2)$ and corresponds to the eigenvector $\{T_t\}$. As long as the heteroskedasticity and serial correlation are bounded, Chamberlain and Rothschild (1983) show that the first principal component will still be a consistent estimator of T_t .

There are good reasons to believe that the measurement error has different variance for different tightness measures—perhaps so different that the principal-components estimator

will not work. The job-finding rate comes from an entirely different data source from the vacancy duration, so the variances of the measurement errors associated with the two different data sources may differ substantially. And within each of these two categories of tightness measures, different measures are estimated from samples of different sizes. For example, in an average month, 39 observations contribute to the estimated job-finding rate for unemployed people who have entered the labor force within the past three weeks, while 46,500 observations contribute to the estimated rate of job-to-job transitions. Based on the effect of sample size on the variance of an estimate, we would expect the variance of the measurement error to be more than 1,000 times larger for new entrants to the labor force than for job-to-job transitions. Other differences may also cause the measurement error to vary across tightness measures—for example, probabilities close to 0 or 1 can be estimated more precisely than probabilities close to 0.5.

In our baseline specification, we standardize each series $x_{i,t}$ to have mean 0 and variance 1:

$$\tilde{x}_{i,t} = \frac{x_{i,t} - \mu_i}{\sigma_i}, \quad (13)$$

where μ_i and σ_i are the time-series mean and standard deviation of $x_{i,t}$. We calculate the first principal component of the standardized indicators.

In most cases, standardization should reduce heteroskedasticity. If the response coefficient β_i were the same for all i , heteroskedasticity in $\epsilon_{i,t}$ would be the only source of different standard deviations σ_i for the different series i , and standardization would remove all of the heteroskedasticity in $\epsilon_{i,t}$. In practice, because β_i varies across series, standardization only approximately removes the heteroskedasticity. In the limit as the differences arising from heterogeneity in the β_i dominate the differences arising from the heterogeneity of the variances of $\epsilon_{i,t}$, standardizing makes the heteroskedasticity worse—see Greenaway-McGrevy, Han and Sul (2012). We check how our results differ if without standardizing the data. In the appendix, we also report a robustness check where we normalize each series by an estimate of the variance of the measurement error in that series, an approach that makes more statistical assumptions but is potentially more robust to differences in β_i .

Principal-components analysis provides not only a method for recovering the common latent tightness factor T_t , but also a way to test for whether the data actually satisfy the single-factor model in equation (12), or the standardized version of the equation. Bai and Ng (2002) provide a set of criterion functions for estimating the number of factors in a

collection of time series. Each criterion function is a penalized measure of the residual sum of squares from the factor model; the estimated number of factors is the number that minimizes the criterion function. We use these criterion functions to estimate the number of factors that drive our time series. We do this for the raw data, for the standardized data, and for data that we have detrended and seasonally adjusted by obtaining residuals from regressions of each series on a linear trend and month-of-year dummies. Our theoretical prediction is that one factor (or a very small number of factors) should drive all of the series, even before we remove trends and seasonality, but after we remove heteroskedasticity by standardizing the data. Like Chamberlain and Rothschild’s, Bai and Ng’s method assumes that heteroskedasticity and cross-sectional and time series dependence of the errors $\epsilon_{i,t}$ are bounded but does not require that the errors actually be white noise. Thus, the method is appropriate even if standardization does not entirely remove heteroskedasticity, and even if there is some cross-sectional or time series dependence in the data. This dependence could arise, for example, because the CPS is a rotating panel and because the samples used to calculate the different job-finding rates are not independent.

In our statistical model, equation (12), the sign and overall level of the index of tightness, T_t , are not identified—any candidate could be multiplied by any positive or negative factor constant across i and the corresponding β_i coefficients divided by the same factor. We normalize T_t so that it has the same standard deviation within our sample as does the aggregate ratio of the number of job openings to the number of monthly hires in JOLTS, and so that it is positively correlated with this ratio. These choices are a pure normalization to make the results easy to relate to published data and have no substantive content.

3 Data

For the job-finding rate, we use data from Hall and Schulhofer-Wohl (2017)—see that paper and its online replication files for the details. We create 16-month activity histories from the Current Population Survey for all respondents of working age. Then we calculate the frequencies of employment following an observation month, up to 15 months later, conditional on the activity in the observation month. For example, for people observed in third month of their presence in the survey, who were classified as unemployed on account of loss of a permanent job, we calculate the fraction who were employed eight months later, in the 11th month of their presence in the survey. We refer the length of the time from the conditioning




Figure 1: Jobfinding Rates for 16 Categories of Jobseekers over One-Month Span

month to the month of potential employment as the *span*—eight months, in this example. We calculated the conditional frequencies of employment for all respondents for all possible spans. These are 1, 2, 3, 12, 13, 14, and 15 months—shorter spans treat short-lived and long-lived jobs equally while longer spans focus on the transition rate into longer-lasting jobs. (In what follows, we report results only for 1- and 12-month spans; results with 2- and 3-month spans are similar to those with 1-month spans, and results with 13- to 15-month spans are similar to those with 12-month spans.) We considered 16 job-seeking categories: 13 categories of unemployed workers distinguished by the reason for unemployment and reported duration of unemployment; 2 categories of workers out of the labor force, distinguished by whether they say they want to work; and employed workers, who may make a job-to-job transition. In our earlier paper, we adjusted the job-finding rates for changes in demographics, but the adjustments had essentially no effect, so here we use the original tabulations without adjustment. Figure 1 shows annual averages of the raw data for the 16 measures for 1-month spans, without attempting to label them. Figure 2 shows the same measures for 12-month spans. The rates vary a great deal in average level, but common movement is apparent, especially the plunge in the recession that began at the end of 2007.

Figure 3 shows the data from JOLTS on (the negative of) jobfilling rates, as annual averages. Note that most of the rates exceed one—meaning that the typical volume of hires




Figure 2: Jobfinding Rates for 16 Categories of Jobseekers over 12-Month Span

during a month exceeds the number of jobs open at the end of the month. The data on jobfilling rates should be considered a daily or continuous rate stated at monthly levels. There is a good deal of heterogeneity across industries in levels and in relative volatility, but no doubt that they tend to move together.

Both the jobfinding rates and the jobfilling rates vary with tightness, in opposite directions. 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.

Figure 4 shows all 32 monthly time series we analyze in standardized form—16 series on job-finding rates by category of job seeker, for one-month job-finding spans, and 16 series on jobfilling rates by industry. Figure 5 shows the same series aggregated to annual frequency. The series generally move closely together and most of the differences between them appear to be white noise, a hypothesis we will test formally below.




Figure 3: (Negative of) Jobfilling Rate by Industry from the 16 Industries in JOLTS




Figure 4: Standardized Monthly Data on Jobfinding and Jobfilling Rates, for One-Month Jobfinding Spans




Figure 5: Standardized Annual Data on Jobfinding and Jobfilling Rates, for One-Month Jobfinding Spans

4 Results

Table 1 shows the estimated number of factors driving the collection of 32 time series, for different versions of the data. Bai and Ng provide three criterion functions for estimating the number of common factors. We almost always find the same number of common factors by all three criteria. The different criterion functions apply only to the estimate of the number of factors; the factors themselves are estimated identically, by principal components. The table also shows the percentage of the variance in the data explained by each of the first two factors.

We analyze six versions of the data. The first is the raw data. In the next version, we detrend each series by computing residuals from a linear regression on a time trend. In the third version, we detrend and seasonally adjust each series by computing residuals from a linear regression on a time trend and month-of-year dummies. The remaining three versions are identical to the first three except that we standardize each series to have mean 0 and variance 1, after detrending and seasonally adjusting, if applicable. We consider seven variants of each of these six versions—each variant uses a different number of months as the span to measure job-finding rates. Table 1 shows only the results for standardized data and

	<i>Bai-Ng criterion</i>			<i>Variance explained</i>	
	<i>1</i>	<i>2</i>	<i>3</i>		
<i>One-month span</i>	<i>Number of factors</i>			<i>First factor</i>	<i>Second factor</i>
<i>Adjustment</i>					
None	3	3	3	26	19
Trend	2	2	2	26	17
Trend and seasonals	1	1	1	32	7
<i>12-month span</i>					
<i>Adjustment</i>					
None	2	2	2	22	16
Trend	2	2	2	25	12
Trend and seasonals	1	1	2	28	9

Table 1: Number of Factors According to the Bai-Ng Criteria, and Percentage Explained by the First Two Factors

for 1- and 12-month spans. Results for all 42 estimations are available in the online backup for the paper.

If we do not standardize the data, we almost always estimate that 30 or 31 factors underly the data. In other words, heteroskedasticity is so high that each series is portrayed as having its own driving force. Even so, the series clearly have much in common, with the first principal component explaining about half of the total variance of the data, and the second principal component explaining around 20 percent of the total variance. These findings hold whether we remove trends and monthly seasonals before estimating the factors or not, and regardless of the length of the job-finding span.

Reducing heteroskedasticity by standardizing the data reveals that the data have low dimension. Without detrending or seasonally adjusting the data, we almost always find that two factors drive the 32 time series, though for 1- and 2-month job-finding spans we sometimes find three factors. In the standardized data, separately detrending the series generally does not reduce the estimated number of factors, indicating that, to the extent a trend is relevant in explaining the behavior of the data, it is a common trend, as our model predicts. Further adjusting the data, by removing monthly seasonals, tends to reduce the estimated number of factors to 1 instead of 2. This finding suggests there may be some role for seasonality in explaining the data beyond what the model allows, but it is a limited role, much less than the number of factors we would find if each series had its own



Figure 6: Annual Averages of Values of the Latent Tightness Index Extracted Using One-Month and 12- month Spans, with Detrending and Seasonal Adjustment

seasonal pattern. In the standardized data as well, the decomposition of variance shows that the series have much in common, with the first principal component—our estimate of T —explaining 20 to 32 percent of the total variance. Note that the theory does not call for T to explain all or even a majority of the variance. Each series is potentially measured with error, and these measurement errors contribute to the variance, perhaps substantially. What the theory demands is that, apart from the single common factor T , the series appear to be white noise—which is what we find.

Figure 6 shows our estimates of T for one- and 12-month spans, extracted from the standardized data, with detrended and seasonally adjusted data, smoothed by taking annual averages of the underlying monthly estimates. The values of T are quite similar. They track the business cycle, with a contraction in 2001 and a deeper contraction in 2008 and 2009. These values are normalized to have the same standard deviation as the overall average duration of vacancies reported in JOLTS and are deviations from means.

Table 2 shows the estimates of the coefficients β_i reflecting the relation of the jobfinding rate to the unobserved common tightness factor, T_t , based on the detrended and seasonally adjusted data. The loadings are all positive, as theory leads us to expect, but there are significant differences in the loadings of different series. For unemployed people, the esti-

<i>Industry</i>	<i>Span, months</i>	
	<i>1</i>	<i>12</i>
Out of labor force	0.021	0.048
Want job	0.087	0.179
Working	0.013	
Recently laid off	0.149	0.157
Recently lost permanent job	0.324	0.402
Temp job recently ended	0.163	0.180
Recently quit	0.295	0.229
Recently entered	0.179	0.158
Recently re-entered	0.266	0.261
On layoff for months	0.248	0.150
Lost permanent job months ago	0.254	0.432
Temp job ended months ago	0.207	0.249
Quit months ago	0.250	0.430
Entered months ago	0.196	0.328
Re-entered months ago	0.215	0.242
Long-term unemployed	0.159	0.196

Table 2: Estimated Slopes of the Jobfinding Rate with Respect to Tightness

mates are generally between 0.15 and 0.4—an increase in tightness of 0.1 months of vacancy duration corresponds to an increase in the jobfinding rate of 1.5 to 4 percentage points—but there is considerable heterogeneity within the job-finding rates. In particular, the job-finding rates of people who have quit or lost permanent jobs have the strongest relation to T , while the job-finding rates of people on temporary layoff move less with T .

Table 3 shows the estimates of the β s for the 16 private industries distinguished in JOLTS. Because the unobserved factor T_t is normalized to match the standard deviation of the average ratio of vacancies to hires as published in JOLTS, the loading factors for the negative of industry hires-vacancies ratios are distributed around minus one. Industries in the goods economy—manufacturing, wholesale trade, and transportation—have coefficients above one in absolute value, along with mining, construction, and real estate. Health and financial services have low coefficients in absolute value. The jobfilling rate in the construction industry is about 20 times more responsive to T than the jobfilling rate in finance. Jobfilling rates are generally more responsive to T than are jobfinding rates.

<i>Industry</i>	<i>Span, months</i>	
	<i>1</i>	<i>12</i>
Mining	-2.923	-2.356
Construction	-11.121	-10.896
Durables	-2.726	-2.729
Nondurables	-2.283	-2.192
Wholesale	-2.369	-2.469
Retail	-1.611	-1.583
Transportation	-2.922	-2.953
Information	-1.444	-1.418
Financial	-0.401	-0.554
Real estate	-1.611	-1.481
Professional services	-0.915	-0.980
Education	-1.240	-1.144
Health	-0.501	-0.504
Entertainment	-7.044	-7.327
Accommodation	-2.459	-2.506
Other services	-1.075	-0.667

Table 3: Estimated Slopes of the Industry-Level Jobfilling Rate with Respect to Tightness

References

- Bai, Jushan and Serena Ng, “Determining the Number of Factors in Approximate Factor Models,” *Econometrica*, 2002, *70* (1), 191–221.
- Chamberlain, Gary and Michael Rothschild, “Arbitrage, Factor Structure, and Mean-Variance Analysis on Large Asset Markets,” *Econometrica*, 1983, *51* (5), 1281–1304.
- Greenaway-McGrevy, Ryan, Chirok Han, and Donggyu Sul, “Asymptotic Distribution of Factor Augmented Estimators for Panel Regression,” *Journal of Econometrics*, 2012, *169* (1), 48—53.
- Hall, Robert E. and Sam Schulhofer-Wohl, “Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Jobseekers,” *American Economic Journal:Macroeconomics*, 2017. Forthcoming.
- Mortensen, Dale T. and Christopher Pissarides, “Job Creation and Job Destruction in the Theory of Unemployment,” *Review of Economic Studies*, 1994, *61*, 397–415.