

4

The Labor Market and the Phillips Curve

A New Method for Estimating Time Variation in the NAIRU

William T. Dickens

The non-accelerating inflation rate of unemployment (NAIRU) is frequently employed in fiscal and monetary policy deliberations. The U.S. Congressional Budget Office uses estimates of the NAIRU to compute potential GDP, that in turn is used to make budget projections that affect decisions about federal spending and taxation. Central banks consider estimates of the NAIRU to determine the likely course of inflation and what actions they should take to preserve price stability. A problem with the use of the NAIRU in policy formation is that it is thought to change over time (Ball and Mankiw 2002; Cohen, Dickens, and Posen 2001; Stock 2001; Gordon 1997, 1998). But estimates of the NAIRU and its time variation are remarkably imprecise and are far from robust (Staiger, Stock, and Watson 1997, 2001; Stock 2001).

NAIRU estimates are obtained from estimates of the Phillips curve—the relationship between the inflation rate, on the one hand, and the unemployment rate, measures of inflationary expectations, and variables representing supply shocks on the other. Typically, inflationary expectations are proxied with several lags of inflation and the unemployment rate is entered with lags as well. The NAIRU is recovered as the constant in the regression divided by the coefficient on unemployment (or the sum of the coefficient on unemployment and its lags).

The notion that the NAIRU might vary over time goes back at least to Perry (1970), who suggested that changes in the demographic composition of the labor force would change the NAIRU. He adjusted the unemployment rate to account for this. By 1990 several authors, including Gordon (1990) and Abraham (1987), had suggested that the NAIRU was probably lower in the 1960s than in the 1970s and 1980s. This

adjustment was initially accommodated by adding dummy variables or splines for certain periods to the Phillips curve regression. However, when it began to appear that the U.S. NAIRU was coming down in the 1990s, new methods were developed to track its changes. Staiger, Stock, and Watson (1997), Gordon (1997, 1998) and Stock and Watson (1999) applied time-varying coefficient models and structural break models to NAIRU estimation, and typically found evidence that it rose in the late 1960s or early 1970s and declined in the 1990s.¹ However, the timing and the magnitudes of the estimated changes differed markedly depending on the specification used. Furthermore, confidence bounds on the estimated NAIRUs were so large that the estimates had little value for policy.²

This paper presents a new approach to estimating time variation in the NAIRU. A major problem with Phillips curve-based estimates is that the complicated relationship between inflation, its own lags, supply shocks, and unemployment and its lags makes it possible to explain any particular incidence of high or low inflation a number of different ways. This problem is the root cause of both the lack of robust results and the large confidence intervals around NAIRU estimates derived from Phillips curve estimates. This paper explores an alternative source of information about time variation in the NAIRU. To the extent that such changes are due to changes in the efficiency of the labor market, these changes are reflected not just in the relationship between inflation and unemployment, but also in the relationship between unemployment and job vacancies. That relationship is much simpler and consequently much easier to model in a robust fashion. Combined estimates of the Phillips curve and Beveridge curve—the relationship between unemployment and vacancies—yield remarkably consistent estimates of the timing of changes in the NAIRU.

The next section provides a brief introduction to the literature on the Beveridge curve and on how it has shifted over time. It argues that because the Beveridge curve is much simpler and potentially better fitting than the Phillips curve, it provides a better basis for discerning shifts in the efficiency of the labor market's functioning. These shifts appear to be quite large. The second section develops a theory linking shifts in the Beveridge curve to shifts in the NAIRU. The third section presents estimates of the Beveridge curve model developed in the second section. These esti-

mates turn out to be very robust and motivate the model developed in the fourth section.

The fourth section presents estimates of a linearized version of the model using a Kalman filter. The filtered series is essentially a weighted average of the residuals of the Beveridge and Phillips curves that has been scaled to satisfy an identifying constraint—this constraint is that the coefficient on the filtered variable must be the same as minus the coefficient on the unemployment rate in the Phillips curve. As might be expected, given how precisely the Beveridge curve is estimated, the filter puts nearly all the weight on the Beveridge curve residuals. Estimates of a restricted version of the model suggest that the information from the Beveridge curve adds significantly to the explanatory power of the Phillips curve. The Beveridge curve and Phillips curve NAIRUs look fairly similar, a result which supports the theory behind both curves. Confidence intervals that account for both forecast and parametric uncertainty are about 40 percent larger for Phillips curve NAIRU series than for series derived from the combined Beveridge curve-Phillips curve model. While estimates of the magnitude of the fluctuations in the NAIRU based on the joint Beveridge curve-Phillips curve model are still fairly uncertain, there is little uncertainty about the timing of the fluctuations.

1. The Beveridge Curve

The Beveridge curve describes a convex relationship between job vacancies and unemployment. It is named after Lord William Beveridge, reflecting his work defining full employment in terms of the ratio of unemployment to job vacancies (Beveridge 1945, 18–20). Hansen (1970) was the first to propose a formal model to explain the nature and shape of the relationship based on disequilibrium in two labor markets. Blanchard and Diamond (1989) offer an alternative model based on a matching function.

Until recently there was no vacancy data for the United States, but starting with Abraham (1983) the Conference Board's help-wanted index has been used to construct a proxy for the number of vacancies in several studies (Abraham 1983, 1987; Blanchard and Diamond 1989; Bleakley and Fuhrer 1997; Medoff 1983; Valetta 2006). Abraham (1983) argued for several adjustments to the help-wanted index to take account of

changes in the structure of the newspaper industry and changes in the use of advertising by business in response to equal opportunity laws and regulations. However, Zargorsky (1998) provides convincing evidence that, except for an adjustment for scale (due to Konstant and Wingard 1968), up to at least 1994 the help-wanted index tracks vacancies well without any adjustment.

Sometime after 1994 this relationship between job vacancies and unemployment falls apart. By the time the Bureau of Labor Statistics began conducting the Job Openings and Labor Turnover Survey (JOLTS) series in December 2000, measures of the vacancy rate constructed from the help-wanted index were running well below the numbers coming out of the JOLTS series. This divergence could have been anticipated given the explosive growth of the Internet as a way for people to find and apply for jobs. Monster.com, one of the first job-matching services on the Internet, started operating in 1994. At that time it listed only 200 jobs (Hernandez 2008). In late 1998 it was still listing only 50,000 jobs. But by May 1999 Monster.com was listing 204,000 jobs, held more than 1.3 million active resumes, and was recording 7.6 million hits per month (Answers.com, 2008). By July 2002 Monster was receiving over 14 million unique hits per month (Hernandez 2008). Today, sites like Careerbuilder.com and Monster.com are only some of many ways that workers connect to job openings through the Internet. Many companies' web sites advertise employment opportunities, and employment agencies use the Internet to troll for jobs and workers to fill them. There is no doubt that a smaller fraction of jobs are listed in newspaper help-wanted advertisements today than was the case 15 years ago.

In the work that follows, the help-wanted index is used with only a scale adjustment, as suggested by Zargosky (1998). From the above it seems likely that this index remains a reliable measure of job vacancies at least up to the end of 1997, but probably not much beyond that point. Thus the years 1998–2000 are dropped in the work presented here. After 2000 the JOLTS data are used to measure the number of vacancies.

Figure 4.1 presents a plot of the vacancy rate (vacant jobs over labor force) versus the unemployment rate from 1954 to 1997 and from 2001 to 2007. Two things are apparent in the graph. First, the Beveridge curve is by no means a stable trade-off. The same vacancy rate was associated with a much higher rate of unemployment in the 1980s than in the

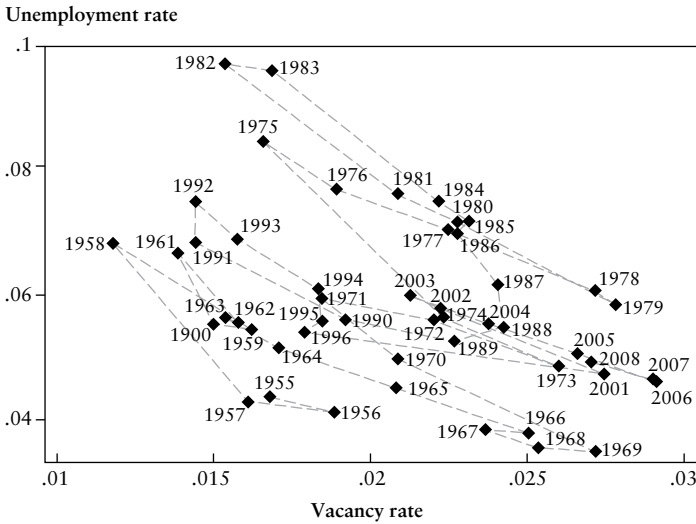


Figure 4.1
 Beveridge Curve 1954–1996, 2001–2007
Source: Author’s computations.

1950s and 1960s. The trade-off in the 1990s and 2000s seems to have improved notably. Starting with Abraham (1987) several authors have offered theories to explain these movements (Cohen, Dickens, and Posen 2001; Valletta 2006).

While the trade-off moves around quite a bit, there do appear to be long periods in which the relationship is relatively stable. From 1958 through 1970 the vacancy and unemployment rates move back and forth in a relatively tight band. There is a similar period from about 1975 through 1986, then another from 1989 to 1997, and then again starting in 2001. The relationship between vacancies and unemployment over these different periods looks remarkably similar. As a result, detecting the magnitude and timing of shifts in the position of the relationship is relatively easy.

If these changes do reflect changes in the efficiency of the functioning of the labor market then these should correspond to large changes in the NAIRU. This possibility provides the point of departure for this paper. What is needed is a theory to guide the measurement of the movements in the vacancy-unemployment relationship and to translate it into movements in the NAIRU.

2. The Model

The model is an extension of Blanchard and Diamond's (1989) continuous-time labor market model modified to yield a NAIRU. Each firm in the economy hires one worker and faces a nominal price for its product at time t , the natural log of which for firm i is denoted

$$(1) \quad p_i(t) = p(t) + z_i(t),$$

where $p(t)$ is the natural log of the aggregate price level³ at time t and $z_i(t)$ is the natural log of the real price entrepreneur i faces. While $p(t)$ changes continuously, the values $z_i(t)$ change in jumps that take place at a rate s . When these change, a new $z_i(t)$ is drawn from a uniform distribution with support on the interval $[a, b]$.

Firms know that the natural log of their real costs for production (including the expected amortized cost of capital) will be w , where ($a < w < 0 < b$), but they do not know the current price level. Thus they do not know the real profits they will be able to make should they choose to produce. Before this information is revealed they must make an irreversible purchase of capital (though they do not have to pay for the capital until it is delivered, and delivery can be delayed till a worker is hired if this is a new firm). Thus, both currently active firms and new firms will decide to produce when faced with a new price if

$$(2) \quad p_i(t) - [p(t) + e(t)] = z_i(t) - e(t) > w,$$

where $e(t)$ is the error in their perception of the log of the current price level common to all entrepreneurs and, thus, the term in brackets is their perception of the log of price level at time t . Active firms (those who currently employ a worker) can continue employing the same worker once the new capital investment is made. New firms must post a vacancy and wait to find a worker before they can begin producing.

If we now assume that $b - a = 1$, then a fraction

$$(3) \quad F = \min(1, \max([1 - w] - e(t), 0))$$

of active firms facing new prices will choose to continue to operate and a fraction $1 - F$ will cease to operate. A fraction F of new firms will choose to post a vacancy while $1 - F$ choose not to post a vacancy and dissolve, as do operating firms that receive a new price and decide not to operate.

It is assumed that the capital cost is sufficiently large relative to the largest possible error in perception that true real prices will never be less than variable costs. Thus, firms that decide to post a vacancy or operate will continue to do so at least until they receive a new price, even if they have underestimated the real price as they are still covering a fraction of the cost of capital.

It is next assumed that unemployed workers are matched with vacant jobs at a rate $M(U, V)$, where U and V are the number of unemployed workers and vacancies respectively. M is assumed to be homogenous degree 1 with $dM/dU > 0$ and $dM/dV > 0$. The labor force contains L workers so that the equations of motion for the vacancy rate and the unemployment rate are given by⁴

$$(4) \quad dV = cgJ^* F - cV(1 - F) - M(U, V)$$

and

$$(5) \quad dU = c(L - U)(1 - F) - M(U, V).$$

New potential firms are created at a rate cgJ^* , where c is the constant rate at which old firms receive new prices, and g and J^* are constants to be defined later. New vacancies are thus created at the rate cgJ^*F (the first term in equation 4). Vacancies disappear when workers are matched to those vacancies (the last term in equation 4) or when a firm with a posted vacancy receives a new price perceived as being too low to be profitable (the second term in equation 4). Workers become unemployed when their firm receives a new price that is perceived to be unprofitable (the first term in equation 5) and leave unemployment when matched with a job (the second term in equation 5).

A permanent increase in F will cause a permanent increase in the number of vacancies and a decline in the number of unemployed, while a decline will have the opposite effects. Following Blanchard and Diamond (1989), this equilibrium locus is defined as the Beveridge curve. The equation that defines it implicitly can be found by setting dV and dU to zero and substituting F out of (4) and (5). Doing this and dividing by the number of workers in the labor force squared, L^2 , yields

$$(6) \quad (1 - u) = [1 + j/(g j^*)] m(u, v) / c,$$

where lowercase letters denote the value of their uppercase counterpart divided by L and $j = v + 1 - u$, or the ratio of jobs to workers.

The long-run equilibrium of the model is defined as the values of v and u that are obtained when $e(t) = 0$. Since the right-hand sides of both (4) and (5) must equal zero in equilibrium, if we set them equal to each other we see that equilibrium also implies $j(1 - F) = g j^* F$. So if we normalize j^* to equal the equilibrium value of j we get that $g = (1 - F)/F$, which in equilibrium equals $w/(1 - w)$. Thus, the term in brackets in equation (6) becomes $1/w$. Note also that in equilibrium if $v + 1 - u = j^*$ then

$$(7) \quad v = (j^* - 1) + u.$$

Together (6) and (7) determine the long-run equilibrium values of v and u —the latter being the NAIRU.

Figure 4.2 plots examples of equation (6) and equation (7) showing how the NAIRU is derived. Equation (6) has the familiar convex shape associated with the Beveridge curve.⁵ Equation (7) is a 45-degree line, the intercept of which is equal to the excess of the vacancy rate over the unemployment rate in equilibrium (i.e., $j^* - 1$).

3. Estimating the Vacancy-Unemployment Relationship

To obtain a Beveridge curve equation that can be estimated, equation (6) must be linearized in logs. Approximating $m(v, u) = A(t) v^b u^{1-b}$, (6) can be rewritten as

$$(6') \quad \ln\left(\frac{1-u}{u}\right) = \ln(A(t)/c) + b \ln(v/u) + \ln\left[1 + \frac{j}{g j^*}\right].$$

Treating the last term as an error term yields the Beveridge equation to be estimated

$$(6'') \quad \ln\left(\frac{1-u_t}{u_t}\right) = A'_t + b \ln(v_t/u_t) + \mu_t,$$

where A'_t is an appropriately scaled time-varying parameter that reflects changes in the efficiency of the matching process. The final term, the log of one plus the ratio of jobs to the number of jobs in equilibrium divided by g , should vary only very slightly compared to the log of the ratio of the vacancy to the unemployment rate.

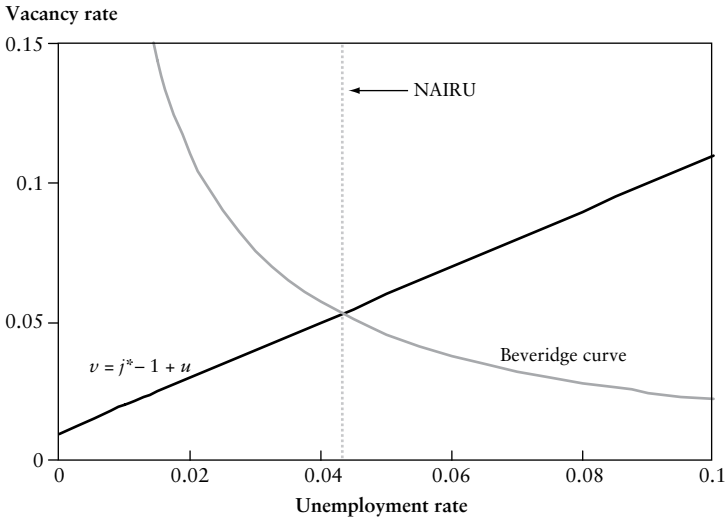


Figure 4.2
 Determination of NAIKU from Beveridge Curve
 Source: Author's computations.

Although equation (6") specifies a single variable linear relationship between $\ln((1 - u)/u)$ and $\ln(v/u)$, it cannot be estimated directly with ordinary least squares. The $A'(t)$ term is time-varying, and from inspection of figure 4.1 there is good reason to believe that the variation in that term would be correlated with the v/u ratio. It would bias the estimates of the coefficient of $\ln(v/u)$ if the time-varying component of $A'(t)$ was treated as a component of the regression error term. Further, the ratio of jobs to the number of jobs in equilibrium will be positively correlated with $\ln(v/u)$, which will tend to bias the estimate of the coefficient of $\ln(v/u)$ downward (though probably only slightly).

From figure 4.1 it appears that the variation in A' is at a much lower frequency than the movement along the Beveridge curve that is reflected in the co-movement of $\ln((1 - u)/u)$ and $\ln(v/u)$. Three different approaches are taken to removing this low frequency variation. First, equation (6") is estimated using only subperiods where the $v - u$ relationship seems stable based on inspection of figure 4.1. Second, the low frequency variation is filtered out of the data and the model estimated only on the filtered data. Finally, both the left- and right-hand sides of (6") are first differenced.

With the low frequency variation in A' removed, the relationship in (6'') should fit well if the approximations used to construct it are good.

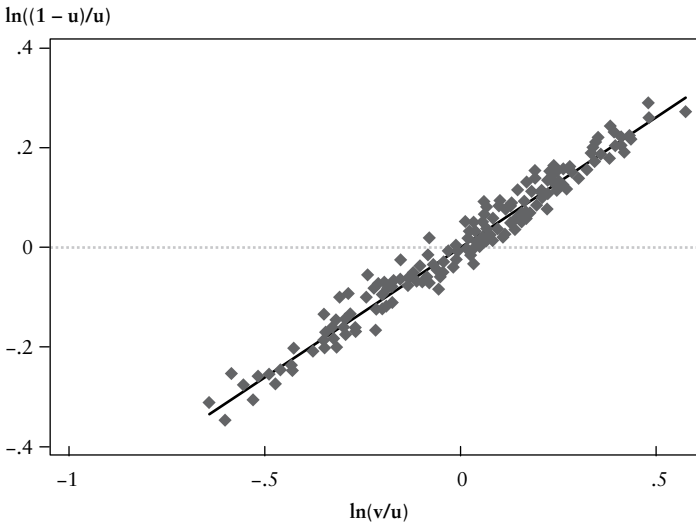
In fact, the relationship between the high frequency variation in the left- and right-hand sides of equation (6'') are remarkably well described by a simple linear relationship as can be seen in figure 4.3. In the bottom panel, the differenced data are plotted against each other. In the top panel data that have been passed through a 25-quarter centered moving average filter are plotted against each other. In this case an unemployment rate that has been age-adjusted, as in Shimer (1999), is used rather than the total unemployment rate. The R^2 s for both regressions are .90 or higher as the observations are tightly packed around a line with a slope that is only slightly larger than .50—the value one would expect if unemployed workers and job vacancies had the same impact on the matching rate.

Nor are these two relationships atypical. Table 4.1 presents 32 different estimates of the coefficient of $\ln(v/u)$ using two different measures of unemployment (age-adjusted unemployment on the top half and total unemployment on the bottom half) and a number of different methods to remove the low-frequency variation. The instrumental variables (IV) estimates are constructed using four lagged values of the log of the vacancy rate.⁶ All of the estimated values of b fall in the interval from .45 to .56 and all are precisely estimated. It is also worth noting that the IV estimates do not vary much from the ordinary least square estimates. There is simply too little error in the relationship for endogeneity of the right-hand-side variable to matter.

4. Joint Estimation of the Phillips and Beveridge Curves

The estimation done in this section proceeds under the assumption that a single unobserved factor moves both the equilibrium ratio of jobs to workers (j^*) and the constant $A'(t)$ in the Beveridge curve, and that the relationship is deterministic. If both are arbitrary functions of that unobservable variable, equation (7) is substituted into equation (6), and j is set equal to j^* (as it is by definition in equilibrium). Equation (6) implicitly defines the NAIRU as a function of the unobservable driver. Inverting that function, linearizing it, and substituting it for the unobservable in

Filtered Model 25-Quarter Moving Average Using Age-Adjusted Unemployment
 $R^2 = 0.97, b = 0.53$



First Difference Model, no constant, Using Total Employment
 $R^2 = 0.902, b = 0.542$

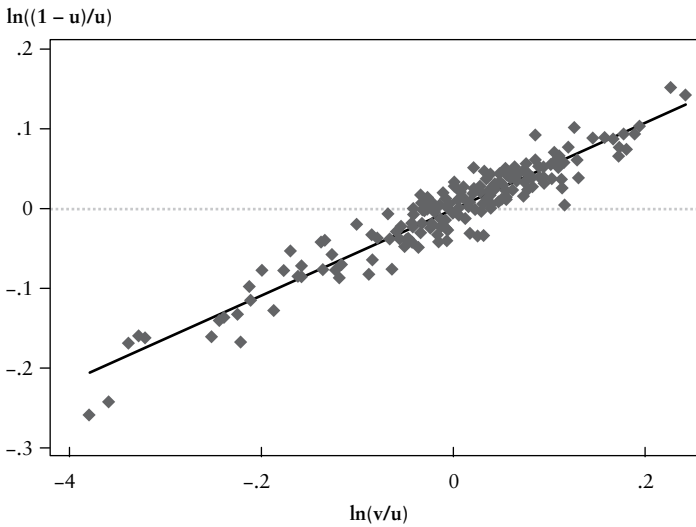


Figure 4.3

Two Examples of Model Fit

Source: Author's computations.

Note: The sample periods used are 1954–1996 and 2001–2007.

Table 4.1
Alternative Estimates of the Vacancy-Unemployment Relationship

	First Difference		Filtered (two-sided MA)				Stable Time Periods			
	With Constant	No Constant	9 qtr. window	17 qtr. window	25 qtr. window	1958:Q3–1968:Q4	1975:Q2–1985:Q4	2001:Q1–2007:Q2	2001:Q1–2007:Q2	
Age-adjusted Unemployment	OLS Coef	0.53	0.54	0.53	0.53	0.52	0.52	0.52	0.47	
	std. error	0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	
IV	Coef	0.56	0.54	0.52	0.52	0.53	0.50	0.50	0.45	
	std. error	0.03	0.03	0.02	0.01	0.01	0.01	0.01	0.01	
Total Unemployment	OLS Coef	0.54	0.54	0.53	0.53	0.50	0.50	0.50	0.49	
	std. error	0.02	0.02	0.02	0.01	0.01	0.01	0.01	0.01	
IV	Coef	0.56	0.55	0.52	0.52	0.50	0.50	0.50	0.49	
	std. error	0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	

Note: The regressions using first differenced data use both the JOLTS and the Help-Wanted Index; The regressions using filtered data use only the Help-Wanted Index data from 1955:Q1–1996:Q4; Stable Time Periods 1958:Q3–1968:Q4 and 2001:Q1–2007:Q2; Help-Wanted Index, 2001:Q1–2007:Q2, JOLTS.

The instrumental variables used are 4 lags of the vacancy rate for the Stable Time Period Regressions, 4 lags of the change in the vacancy rate are used for the first-difference regressions and 4 lags of the filtered vacancy rate for the filtered regressions.
Source: Author's computations.

the determination of $A'(t)$ yields the Beveridge curve equation to be estimated

$$(8) \quad \ln\left(\frac{1-u_t}{u_t}\right) = A'' - au_t^* + b'\ln(v_t / u_t) + \mu_t,$$

where u^* is the NAIRU, which is unobservable.

A standard price-price Phillips curve is also estimated of the form

$$(9) \quad \pi_t = \sum_{i=S_u}^{L_u} c_i(u_t^* - u_{t-i}) + \sum_{i=1}^{L_{inf}} d_i\pi_{t-i} + x_t'h + \varepsilon_t,$$

where π_t is inflation in quarter t , L_u and L_{inf} are the number of lags of unemployment and inflation included respectively, S_u is either 0 or 1 depending on whether contemporaneous unemployment is included in the equation, x_t is a vector of dummy variables capturing supply shocks, and h is a conforming vector of coefficients. The coefficients d_i are constrained to sum to 1 so that u^* can be interpreted as the NAIRU in the absence of any observed supply shocks.

It is assumed that the NAIRU u^* evolves as a random walk with an innovation that is independent of the innovations in equations (8) and (9). To identify the model it is further assumed that $\text{cov}(\mu_t, \varepsilon_t) = 0$ so that the only source of correlation between the unobservables in equations (9) and (10) is the common NAIRU.

The model is estimated using a Kalman filter. The constraint that the NAIRU must have the same coefficient as the unemployment rate in the Phillips curve, and the restriction on the covariances of the error terms, are adequate to completely identify all the model parameters. The approach used here is similar to that taken by Basistha and Startz (2008) to estimating the NAIRU with multiple indicators.

Table 4.2 presents six different estimates of the Beveridge curve-Phillips curve model. The first column presents a specification using data from 1955:Q1 to 1997:Q4 and from 2001:Q1 to 2008:Q3 using the Consumer Price Index (CPI) as the inflation measure, and the civilian unemployment rate as the unemployment measure. Included in the Phillips curve equation are contemporaneous unemployment, twelve lags of the inflation rate, and three lags of the unemployment rate. All the parameters of the model are estimated with a fair degree of precision and

Table 4.2

Kalman Filter Estimates of Beveridge Curve-Phillips Curve System and Ordinary Least Squares Phillips Curve with Kalman NAIRU

Parameter	Specification					
	Beveridge curve error variance unconstrained			Beveridge curve error variance constrained to zero		
	55-97 01-08 12 lags CPI unemp. + 3 lags	60-95 01-08 4 lags GDP consumption deflator age adj. unemp.	55-96 8 lags GDP deflator 2 lags unemp.	55-97 01-08 12 lags CPI unemp. + 3 lags	60-95 8 lags core CPI age-adj. unemploy.+ 2 lags	65-96 1 lag GDP deflator 4 lags unemp.
Phillips curve						
Sum coefficient unemployment	-.73 (.21)	-.30 (.09)	-.40 (.11)	-.73 (.21)	-.51 (.20)	-.16 (.13)
Average** Kalman gain	1.5E-7	3E-8	3E-7	0*	0*	0*
Beveridge Curve						
Constant	4.57 (.38)	4.14 (.21)	4.59 (.42)	4.56 (.37)	4.22 (.27)	3.96 (.43)
NAIRU	-19.9 (6.2)	-13.4 (3.5)	-20.4 (6.9)	-19.8 (6.1)	-14.8 (4.5)	-10.7 (6.6)
Ln(v/u)	.52 (.01)	.50 (.01)	.52 (.01)	.51 (.01)	.50 (.01)	.50 (.01)
Post-2000 dummy	-.20 (.09)	-.09 (.08)		-.19 (.08)		
Average** Kalman gain	-.059	-.080	-.057	-.057	-.080	-.051
s.d. NAIRU innovation	.001 (.0003)	.001 (.0003)	.001 (.0004)	.001 (.0003)	.001 (.0003)	.001 (.0009)
Log Likelihood	1063.3	1029.5	951.4	1060.7	783.6	732.4
NAIRU in Phillips curve						
Constant	.004 (.011)	.008 (.007)	.000 (.008)	.002 (.012)	.014 (.011)	.004 (.007)
NAIRU	.86 (.37)	.66 (.42)	.90 (.43)	.85 (.37)	.55 (.47)	.57 (.99)
Chi-squared test Constant = 0 NAIRU=1	.40 [.82]	1.15 [.56]	.26 [.88]	.41 [.81]	1.26 [.53]	.60 [.74]

Standard errors in parenthesis, significance levels in square brackets.

*Constrained to zero by assumption that Beveridge curve error variance equals zero.

** Average across all time periods.

Source: Author's computations.

the coefficient on $\ln(v/u)$ is very precisely estimated and falls in the same narrow range as the single equation estimates. The unobserved NAIRU variable enters the Beveridge curve equation with a very precisely estimated coefficient.

The help-wanted index is used to estimate the vacancy series during the earlier period and the JOLTS survey during the latter period. Can the same model fit both periods? Originally the model was estimated with a dummy variable for the latter period in the Beveridge curve and interactions between that dummy and the NAIRU and between the dummy and $\ln(v/u)$. Neither of the interactions was statistically significant individually or jointly in any specification tried, so these were dropped from the model with virtually no impact on any other model parameters. Only the dummy variable supplementing the intercept was retained.

The most interesting result in column 1 is the relative magnitude of the Kalman gain for the Phillips curve and Beveridge curve residuals. The Kalman filter model constructs the NAIRU as a weighted average of the residuals of the two equations plus the previous period's estimate of the NAIRU. The Kalman gain is the weight put on each of the two residuals. Given the estimated parameters, the residuals of the Phillips curve play virtually no role in constructing the NAIRU, while those of the Beveridge curve play a major role. The estimated variance of the Beveridge curve innovation (μ_t) is so close to zero⁷ that the model identifies the NAIRU as being nearly exactly proportional to the difference between the left-hand side of equation (8) and the constant plus $b \ln(v/u)$. This result is not unique to the model in the first column. In every specification presented—in fact, in every one of the several dozen specifications tried—the model chose to equate the NAIRU with the Beveridge curve residual nearly exactly. Thus the right three columns of the table present three specifications with the constraint that the Beveridge curve residual is exactly proportional to the NAIRU. The first of these replicates the specification in the first column. Comparing the results in the first and third columns shows what little effect the constraint has on the estimated coefficients.

That the Phillips curve residuals are seen as uninformative with respect to the magnitude of the model NAIRU might indicate that what is being measured is not a NAIRU, but simply the time variation in the inter-

cept of the Beveridge curve. On the other hand, it could be that the time variation in the intercept of the Beveridge curve very precisely measures movements in the NAIRU so that information from Phillips curve residuals is superfluous. We saw in table 4.1 and figure 4.3 just how well the Beveridge curve fits. There is no straightforward test of whether the model's NAIRU matters for explaining inflation, since it must be assumed that the coefficient on the unemployment rate is the same as the coefficient on the NAIRU for identification. However, once the variance of the Beveridge curve innovation is restricted to zero, so that the Beveridge curve residual is assumed to be proportional to the NAIRU, a constant and a separate coefficient for the NAIRU can be added to the Phillips curve. The bottom three lines of table 4.2 show the results of doing this for the six specifications presented there. If the Kalman filter model is correct then the constant term in the Phillips curve should equal zero, and the coefficient on the NAIRU should equal 1. In none of the six specifications can either hypothesis be rejected individually or jointly. Further, in most of the specifications presented the hypothesis that the coefficient on the NAIRU is equal to zero can be rejected at the .10 level or better in a one-tailed test.

When models with every possible combination of inflation measure (CPI, CPI-core, GDP deflator, GDP consumption deflator), inflation lag structure, and unemployment lag structure were estimated, the results were remarkably consistent. There were 159 specifications that could not be rejected as being overly constrained compared to the specification in column one of table 4.2. Of those, there was not one in which the hypotheses that the constant was equal to zero or the coefficient on the NAIRU is equal to one could be rejected individually or jointly at the .10 level (two-tailed test). Yet in 92 of the 159 specifications, the hypothesis that the coefficient on the NAIRU is zero was rejected at the .05 level and at the .10 level in 150 specifications (one-tailed test).

The results aren't quite as good when the model is estimated with age-adjusted unemployment. There are 169 specifications that cannot be rejected when compared to one with 12 lags of inflation, contemporaneous unemployment, and 4 lags of unemployment. Once again, there is not a single specification where the hypothesis that the constant in the Phillips curve equals zero or the hypothesis that the coefficient on the

NAIRU equals 1 can be rejected at the .10 level individually or jointly. However, there are only 43 specifications where the hypothesis that the coefficient on the NAIRU is zero can be rejected at the .10 level and only 9 where it can be rejected at the .05 level. Perhaps once the variation in the NAIRU due to the age structure of the population is taken into account by adjusting the unemployment rate there is little additional information in the Beveridge curve residuals. On the other hand, the Beveridge curve model might be considered misspecified when age-adjusted unemployment is used, since the vacancy rate hasn't been similarly adjusted.

It appears that the low frequency movements in the Beveridge curve probably belong in the Phillips curve as an indicator of variation in the NAIRU—at least if unadjusted unemployment is used. The next question is whether there is any more variation in the NAIRU than the variation due to labor market efficiency reflected in the Beveridge curve residuals. With the constraint that there is no innovation in the Beveridge curve, it is possible to estimate a model with two unobservables so that the NAIRU is the sum of the Beveridge curve residual and a filtered version of the Phillips curve residual. When this model is estimated in any of a wide range of specifications the likelihood is maximized when the variance of the innovation in the Phillips curve NAIRU is zero. Thus the hypothesis that there is nothing more to variation in the NAIRU than that captured by the Beveridge curve residual cannot be rejected.

How does the standard time-varying NAIRU estimated using only the Phillips curve compare to the Beveridge curve based NAIRU? Figure 4.4 presents examples of both. The gray line in figure 4.4 depicts the NAIRU derived from the Beveridge curve model from column 1 of table 4.2. The heavy black line in figure 4.4 depicts a NAIRU estimated using only the residuals of a Phillips curve with the same specification as that used to estimate the combined Beveridge curve-Phillips curve model. The two are similar in many respects. Except for a short period in the late 1960s to early 1970s, and the lack of a bulge in the Phillips curve NAIRU in the mid-1980s, the two track each other fairly closely. When the unemployment rate was high relative to the job vacancy rate it was also high relative to the inflation rate.

If the Beveridge curve and the Phillips curve NAIRU look similar, what is the advantage of the latter? Confidence intervals for both series

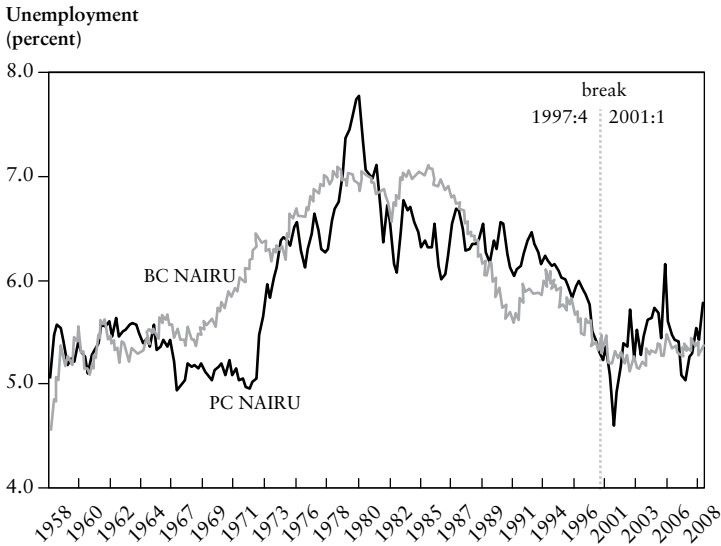


Figure 4.4
Phillips Curve and Beveridge Curve NAIURU
Source: Author's computations.

were constructed for the two models in figure 4.4 that take account of both forecast and parametric uncertainty by computing 10,000 Monte Carlo trials. Despite the Beveridge curve model having several additional parameters, the 90-percent confidence intervals for the Phillips curve NAIURU were about 40 percent larger on average.⁸ Other specifications for the two models yielded similar results—the Beveridge curve-based NAIURU had narrower confidence intervals in every specification tried.

The very precise and similar estimates of the Beveridge curve series across many different specifications, along with the narrow confidence intervals on the computed NAIURU series, suggests that there should be considerably more certainty about the position of the Beveridge curve NAIURU than there is about NAIURUs estimated from the Phillips curve alone. This is somewhat true. Figure 4.5 presents the average value of the Beveridge curve NAIURU estimated across the 159 specifications using total unemployment. The specifications varied the lags of unemployment and inflation and the inflation measure as described above. Also plotted in figure 4.5 are the minimum and maximum values in each quar-

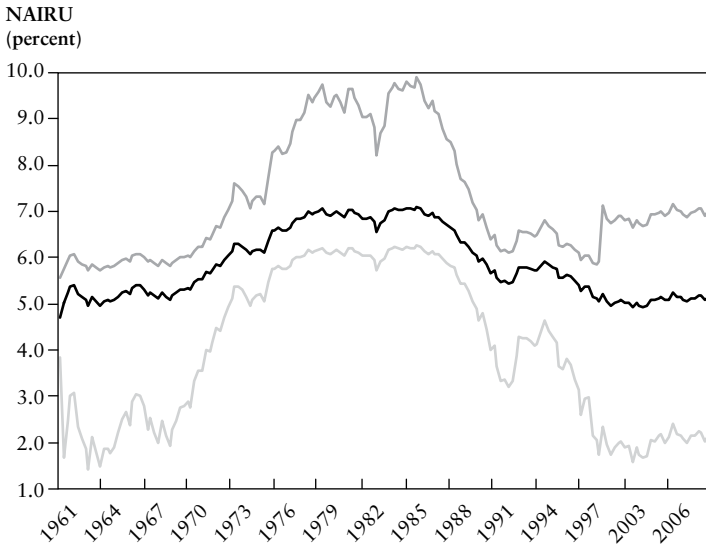


Figure 4.5
NAIRU with Upper and Lower Bounds
Source: Author's computations.

ter of the 90-percent confidence intervals for the NAIURU computed for each of the 159 specifications. The bounds were estimated by simulating parametric and forecast uncertainty with 10,000 Monte Carlo trials each.

Allowing for forecast, specification, and parametric uncertainty, as the bounds in figure 4.5 do, considerable uncertainty about the position of the NAIURU at any given time remains. This is particularly true at the moment because the switch from the help-wanted series to the JOLTS series adds substantially to uncertainty. Still, the results reported here improve on estimates based on only the Phillips curve in at least one dimension—there is little uncertainty about the timing of major changes in the NAIURU. All estimates show a substantial rise in the NAIURU during the 1970s and a decline in the late 1980s and early 1990s. While NAIURU values much above 6 percent can be ruled out during the 1960s and the mid-to-late 1990s, values less than that can be ruled out for the decade starting in 1978. This provides more guidance to policymakers than past estimates.

5. Conclusion

This paper has presented a new method for estimating time variation in the NAIRU using the vacancy-unemployment relationship. A simple theory of this relationship based on a matching model suggests equations that do an uncannily good job of fitting transformed vacancy and unemployment data. When the Beveridge curve model is estimated simultaneously with a Phillips curve, the parameter estimates for both equations are reasonable and the parameters of the Beveridge curve are estimated with particular accuracy. The estimates suggest that the NAIRU is nearly exactly proportional to the residual in the Beveridge curve. When this constraint is imposed it is possible to test whether the Beveridge curve residuals help explain inflation. In the 328 specifications tried, there were none in which the hypothesis that the scaled Beveridge curve residual was the NAIRU could be rejected. The hypothesis that the Beveridge curve residual did not help explain inflation could be rejected at least at the .10 level in nearly all specifications using the total unemployment rate and many where age-adjusted unemployment was used. A model augmented with a time-varying NAIRU estimated as the sum of the scaled Beveridge curve residual plus a filtered version of the Phillips curve error was estimated. The hypothesis that the filtered Phillips curve error did not help forecast inflation could not be rejected and the resulting NAIRU series differed little from the Beveridge curve residual alone. A standard Phillips curve NAIRU resembles the Beveridge curve NAIRU in the timing of its movements, which validates the theory on which both the Beveridge curve and the Phillips curve NAIRU are based. However, estimates of the Beveridge curve NAIRU are more precise.

Despite the very precise estimates of the parameters of the Beveridge curve, when forecast, specification, and parametric uncertainty are taken into account, the data are consistent with a fairly wide range of values for the NAIRU at each point in time. This is particularly true since 2001, when the JOLTS vacancy rate series replaces the help-wanted series. Still, there appears to be considerable information in the Beveridge curve model about movements in the NAIRU. Estimated NAIRU series differ in the magnitude of the fluctuations, but hardly at all in their average value or the timing of the fluctuations. Further, as we get more experience

with the JOLTS vacancy series, uncertainty about the current NAIRU will decline since the increased uncertainty post-2000 is due entirely to uncertainty about the magnitude of the post-2000 dummy variable in the Beveridge curve.

■ *I would like to thank Vania Stavrakeva for her excellent research assistance and Robert Valletta for the use of his age-adjusted unemployment series. I would also like to thank the Brookings Institution and the Russell Sage Foundation for research support. This paper further develops work I did with Jessica Cohen and Adam Posen on Beveridge curve shifts (Cohen, Dickens, and Posen 2001). A conversation with Jim Stock was helpful in thinking through identification issues and Richard Startz pointed out a problem with the model and made several helpful comments about presentation. I would also like to thank Olivier Blanchard for comments on an earlier draft of the paper and for pointing me toward two conference papers he wrote in the late 1980s that took a very similar approach to the one taken here. This paper differs from those in the way the Beveridge curve is estimated, the joint estimation of the Phillips curve and Beveridge curve, and its focus on estimating a time-varying NAIRU.*

Notes

1. In these papers only the constant term, or the NAIRU, was allowed to vary. When Brainard and Perry (2000) estimated Phillips curves allowing all parameters to vary they found that the constant and the coefficient on unemployment were relatively stable (and thus so was the NAIRU). Instead they explained the different behavior of inflation over the decades by variation in the sum of the coefficients on lagged inflation.
2. See, for example, the results in Staiger, Stock, and Watson (1997) or compare figures 3 and 4 in Gordon (1997).
3. Where the aggregate price level is the geometric average of all prices in the economy.
4. In a departure from Blanchard and Diamond (1989), quits and layoffs are both assumed to arise from the breakdown of a match, which is signified by the arrival of a new price for the entrepreneur.
5. The figure assumes a log-linear matching function with constant returns to scale and equal weight on vacancies and unemployment.

6. Using either the level, the filtered value, or first differences corresponding to the treatment of the variables in the specification being estimated.

7. In fact, in all of the dozens of specifications tried the variance of the Beveridge curve residual was estimated to be slightly negative. This is possible because the variance is scaled and then added to the forecast error variance due to the innovation in the NAIRU so that the total forecast error variance remains positive. The slight negative value was often statistically significantly different from zero, a result that suggests that the assumptions about the innovations are not exactly correct. In fact, the small positive auto-correlation (about .25 in specifications where it was inspected) in the estimated changes in the NAIRU could explain this.

8. In general the likelihood function for the Kalman filter model for the Philips curve NAIRU pushed the variance of the innovation of the NAIRU to zero so that value was fixed in order to compute the NAIRU shown in figure 4.4. The variance of the innovation was chosen so that the Beveridge curve and Phillips curve NAIRUs had the same variance.

References

- Abraham, Katharine G. 1983. "Structural/Frictional vs. Deficient Demand Unemployment: Some New Evidence." *American Economic Review* 73(4): 708–724.
- Abraham, Katharine G. 1987. "Help-Wanted Advertising, Job Vacancies, and Unemployment." *Brookings Papers on Economic Activity* 1: 207–248.
- Answers.com. 2008. "Hoover's Profile: Monster Worldwide, Inc." Available at <http://www.answers.com/topic/monster-worldwide-inc?cat=biz-fin>. Accessed June 12.
- Ball, Lawrence, and N. Gregory Mankiw. 2002. "The NAIRU in Theory and Practice." *Journal of Economic Perspectives* 16(4): 115–136
- Basistha, Arabinda, and Richard Startz. 2008. "Measuring the NAIRU with Reduced Uncertainty: A Multiple-Indicator Common-Cycle Approach." *Review of Economics and Statistics* 90(4): 805–811.
- Beveridge, William H. 1945. *Full Employment in a Free Society*. New York: W.W. Norton & Company.
- Blanchard, Olivier Jean, and Peter A. Diamond. 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity* 1: 1–60.
- Bleakley, Hoyt, and Jeffrey C. Fuhrer. 1997. "Shifts in the Beveridge Curve, Job Matching, and Labor Market Dynamics." *New England Economic Review* (September/October): 3–19. Available at <http://www.bos.frb.org/economic/neer/neer1997/neer597a.pdf>.
- Brainard, William C., and George L. Perry. 2000. "Making Policy in a Changing World" Perry, George and James Tobin (eds), In *Economic Events, Ideas, and Policies: The 1960s and After*, ed. George L. Perry and James Tobin, 43–68. Washington, DC: Brookings Institution Press.

Cohen, Jessica, William T. Dickens, and Adam Posen. 2001. "Have the New Human Resource Management Practices Lowered the Sustainable Unemployment Rate?" In *The Roaring Nineties: Can Full Employment be Sustained?*, ed. Alan B. Krueger and Robert M. Solow, 219–259. New York: The Russell Sage Foundation and the Century Foundation Press.

Gordon, Robert J. 1990. "U.S. Inflation, Labor's Share, and the Natural Rate of Unemployment." In *Economics of Wage Determination*, ed. Heinz König, 1–34. New York: Springer Verlag.

Gordon, Robert J. 1997. "The Time-Varying NAIRU and its Implications for Economic Policy." *Journal of Economic Perspectives* 11(1): 11–32.

Gordon, Robert J. 1998. "Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU." *Brookings Papers on Economic Activity* 2: 297–346.

Hansen, Bent. 1970. "Excess Demand, Unemployment, Vacancies, and Wages." *Quarterly Journal of Economics* 84(1): 1–23.

Hernandez, Judy. 2008. "The Monster.com Story." Available at <http://www.careertrainer.com/Request.jsp?lView=ViewArticle&Article=OID%3A55989>. Accessed June 12.

Konstant, Raymond A., and Irvin F. Wingard. 1968. "Analysis and Use of Job Vacancy Statistics: Part I." *Monthly Labor Review* 91(August): 22–31.

Medoff, James L. 1983. "U.S. Labor Markets: Imbalance, Wage Growth and Productivity in the 1970s." *Brookings Papers on Economic Activity* 1: 87–128.

Perry, George L. 1970. "Changing Labor Markets and Inflation." *Brookings Papers on Economic Activity* 1: 411–441.

Shimer, Robert. 1999. "Why is the U.S. Unemployment Rate So Much Lower?" In *NBER Macroeconomics Annual 1998*, ed. Ben S. Bernanke and Julio J. Rotemberg, 11–61. Cambridge, MA: The MIT Press.

Staiger, Douglas, James H. Stock, and Mark W. Watson. 1997. "How Precise are Estimates of the Natural Rate of Unemployment?" In *Reducing Inflation: Motivation and Strategy*, ed. Christina D. Romer and David H. Romer, 195–246. Chicago: University of Chicago Press.

Staiger, Douglas, James H. Stock, and Mark W. Watson. 2001. "Prices, Wages, and the U.S. NAIRU in the 1990s." In *The Roaring Nineties: Can Full Employment be Sustained?*, ed. Alan B. Krueger and Robert M. Solow, 3–60. New York: The Russell Sage Foundation and the Century Foundation Press.

Stock, James H. 2001. "Monetary Policy in a Changing Economy: Indicators, Rules, and the Shift Towards Intangible Output." Paper presented at Bank of Japan Conference on Monetary Policy.

Stock, James H., and Mark W. Watson. 1999. "Forecasting Inflation." *Journal of Monetary Economics* 44(2): 293–335.

Valletta, Robert G. 2006. "Why Has the U.S. Beveridge Curve Shifted Back? New Evidence Using Regional Data." Federal Reserve Bank of San Francisco Working Paper Series No. 2005-25. Available at <http://www.frbsf.org/publications/economics/papers/2005/wp05-25bk.pdf>.

Zagorsky, Jay L. 1998. "Job Vacancies in the United States: 1923 to 1994." *Review of Economics and Statistics* 80(2): 338-345.

Comments on “A New Method for Estimating Time Variation in the NAIRU”

by William T. Dickens

Olivier Blanchard

Dickens’s paper offers a promising strategy to identify shifts in the natural rate of unemployment by looking jointly at the Beveridge curve and the Phillips curve. In my comments, I want to do two things. First, sketch the conceptual framework that allows one to extract information about the natural rate of unemployment from the Beveridge curve and the Phillips curve, and the factors behind its movement. I believe this is the framework that underpins Dickens’s analysis; all I want to do is to make it more transparent. Second, bring the framework to the U.S. data. The answers one gets from this exercise are surprisingly clear-cut and, I think, interesting, as they do not conform to my (and I suspect your) priors. This is, of course, just a first look at the data, but it shows how useful this approach can potentially be for thinking about changes in the natural rate of unemployment.

The theoretical framework can be characterized simply by two equations. The first equation is a relation between the flow into unemployment and the flow out of unemployment and back to employment:

$$(1) \quad s(1-u) = mf(u, v).$$

Equilibrium in the labor market is characterized by the equality of separations and hires. The left side of the equation captures separations, where s is the separation rate and u is unemployment normalized by the labor force. The right side of equation (1) captures hires, and is given by a matching function, whose output is a flow of new hires. The function matches unemployed workers with vacancies, v . The function is increasing in u and v , while m is a scale parameter denoting the efficiency of the matching process. Equation (1) describes a negative relation between

unemployment and job vacancies—the Beveridge curve. The position along the Beveridge curve is related to the state of the business cycle. Recessions are periods when many unemployed workers are pursuing few vacancies. Conversely, in a tight labor market, which is generally associated with high labor demand, more vacancies are searching among fewer unemployed workers.

For our purposes, we are interested not so much in movements along the Beveridge curve, but in the factors that lead to *shifts* in the curve. Given the way equation (1) is written, shifts in the curve arise from two sources. The first is a change in s . A decrease in the separation rate means lower flows through unemployment. This could result from less labor market reallocation and/or more relative flows directly from employment (or entrance into the labor force) to employment. The latter case could be the result, for example, of temporary employment agencies now making it possible for workers to transition from one job to another without becoming unemployed in the process. A lower s , for a given amount of vacancies, implies less unemployment, and thus a shift of the Beveridge curve inward. The other source of shifts in the curve is a change in m , the efficiency of matching. It is reasonable to think that in recent years the technology for matching workers and jobs has become better, so that not as many workers or as many vacancies are needed to generate a certain flow of hires. The advent of Monster.com is an obvious example. Improvements in matching efficiency shift the Beveridge curve inward.

The second equation in the framework determines wages. The relationship says that the wage firms can offer must be equal to the wage implied by bargaining:

$$(2) \quad \bar{w} = w \left(\frac{v}{u}, z \right).$$

The left side of the equation denotes the wage that is consistent with normal profits. The right side of the equation is a wage function that can be derived in the context of a wage bargaining model where the surplus from matching a firm and a worker is shared in some proportion (see, for example, Pissarides 1985). According to this function, the labor market variable that matters in determining the outcome of wage bargaining is the ratio of vacancies to unemployment, v/u . This ratio determines the

bargaining power that each party possesses. A high v/u indicates that workers' bargaining strength is high relative to firms', and this yields a higher wage rate. The wage-determination function is thus increasing in v/u . For given \bar{w} and z , equation (2) describes a positive linear relationship between v and u . The parameter z summarizes the factors that affect bargaining, which may arise from the presence of unions or other institutional features pertaining to the labor market that influence the bargaining power of workers relative to firms.

Equation (2) abstracts from the presence of nominal rigidities. Wage dynamics, however, can be introduced in (2) to yield a Phillips curve specification that relates wage inflation, price inflation, v/u , and z . Movements in z will shift the Phillips curve relationship. Increasing globalization, to the extent that it reduces the bargaining power of workers, will shift the Phillips curve inward in the (u, w) space.

Using equations (1) and (2), it is possible to derive the steady-state equilibrium levels of u and v . In particular, the natural rate of unemployment can be written as:

$$(3) \quad u^* = u^*(s, m, z).$$

The natural rate of unemployment is a function of the separation rate, the efficiency of the matching process, and the factors that affect the bargaining power of workers. From equation (3), it is evident that we cannot obtain an estimate of the natural rate of unemployment by estimating either the Beveridge curve or the Phillips curve alone. Both the Beveridge curve and the Phillips curve are needed to back out s , m , and z , the three factors that affect u^* . This simple framework thus illustrates why the strategy pursued in Dickens's paper to estimate the natural rate of unemployment is potentially a good one: estimation of both the Beveridge curve and the Phillips curve can capture the three factors influencing the natural rate of unemployment.

Having laid the framework for thinking about movements in the natural rate of unemployment in the context of the Beveridge curve and the Phillips curve, I will now bring this framework to the data. The first observation that I would like to make, which I think is not controversial, is that low frequency movements in the unemployment rate strongly suggest a decline in the natural rate of unemployment since the late 1980s,

by a magnitude of at least 1 percentage point. The question is then what accounts for the decline in the natural rate of unemployment. Is it movements in s , m , or z ? As documented in Dickens's paper, the Beveridge curve has shifted noticeably inward after a transition period in the late 1980s. This inward shift in the Beveridge curve points to s and m , and not to z —the Phillips curve shifter—as potentially important factors in lowering the natural rate of unemployment. Is it possible to assess the independent contribution of s and m to the decline in u^* ? The answer to this question is relatively straightforward. We can observe the separation rate s directly in the data. An estimate of m can be obtained from estimating a matching function, as in Blanchard and Diamond (1989, 1991). Performing this exercise, it becomes apparent that a lot of the action is coming from the separation rate, which has declined noticeably since the late 1980s. This is shown in figure 4.6, which plots a time series for the separation rate from employment and for the unemployment rate, using Bureau of Labor Statistics data. The separation rate includes separations from employment to unemployment, from employment to out of the labor force, and from employment to employment.

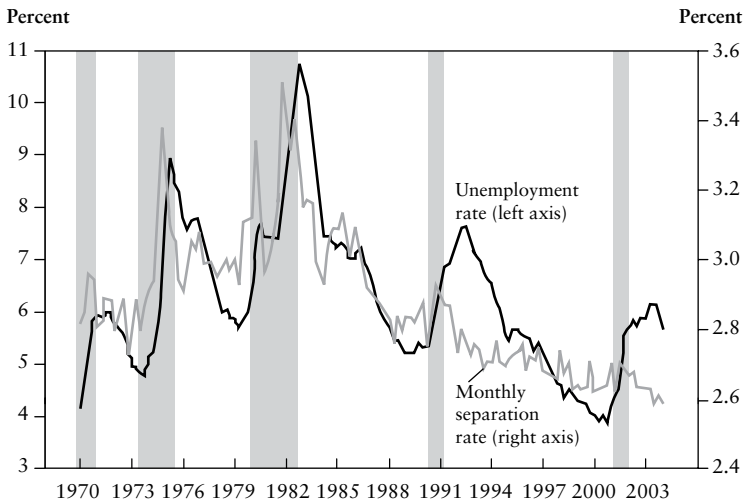


Figure 4.6

The Unemployment Rate and the Separation Rate from Employment, 1970–2004

Source: Author's computations from Bureau of Labor Statistics data.

Note: Shaded areas show NBER-dated recessions.

The sharp decline in the separation rate suggests that s played a prominent role in the decline of the natural rate of unemployment. Note that the picture shows that s has declined from a monthly 3 percent rate to about 2.5 percent—a 15 percent decline. This is roughly the same decline that seems to have occurred for the natural rate of unemployment, which is usually estimated to have declined from 6 percent to about 5 percent, approximately a 15 percent decline. In sum, to a close approximation, it appears that the decline in the natural rate of unemployment can be explained entirely by a decline in the separation rate. Changes in matching efficiency and in the bargaining power of workers—with this last factor being my prior as the most likely explanation for a decline in u^* —do not appear to account for an important part of the story.

The next step, which goes beyond the scope of Dickens's paper but is nonetheless an important issue, is to understand what lies behind the decline in the separation rate. Lower worker flows could be the result of demographics. An older labor force is less inclined to move among jobs. This would result in lower worker flows even in an environment in which job flows are approximately given. But it is possible to look directly at job flows, and here we see that the decline in worker flows is not just due to demographics, as figure 4.7 shows. There has been a decline in job creation and in job destruction; that is, the amount of job reallocation in the U.S. economy has fallen. This decline does not explain entirely the decline in worker flows, but roughly two-thirds of it. Still, the reasons why churning in the labor market has decreased remain to be investigated. More wage flexibility, better inventory control, and more integrated chains of production could have contributed to the decline in job creation and job destruction.

In sum, Dickens's paper outlines a promising approach, and I hope that my comments have highlighted the usefulness of this approach to both estimating the natural rate of unemployment and discriminating among different reasons for movements in the natural rate of unemployment. The implementation of this approach to U.S. data delivers clear leads, in that it downplays decreases in the bargaining power of workers and more efficient firm-worker matching as explanations for a lower natural rate of unemployment. Somewhat surprisingly, the data point to lower worker flows, and in turn to lower job flows as the most important factor affect-

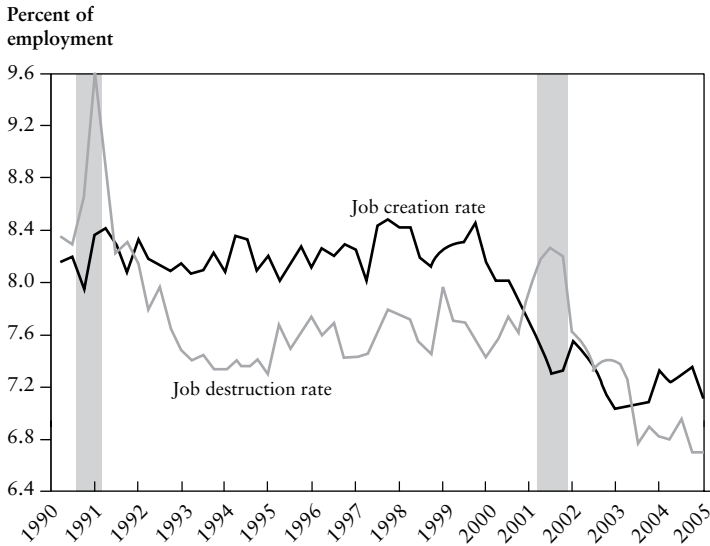


Figure 4.7

Quarterly Job Flows in the Private Sector, 1990–2005

Source: Faberman (2006); tabulated from BLS Business Employment

Note: Dynamics (BED) micro data. Shaded areas show NBER-dated recessions.

ing the decline in u^* . What lies behind the decrease in job flows remains an open issue that we need to address. But at least we know what to look for, and how to interpret it.

References

Blanchard, Olivier Jean, and Peter A. Diamond. 1989. “The Beveridge Curve.” *Brookings Papers on Economic Activity* 1: 1–60.

Blanchard, Olivier Jean, and Peter A. Diamond. 1991. “The Aggregate Matching Function.” Working Paper No. 3175. Cambridge, MA: National Bureau of Economic Research.

Faberman, R. Jason. 2006. “Jobs Flows and the Recent Business Cycle: Not All ‘Recoveries’ Are Created Equal.” BLS Working Paper No. 391. Washington, DC: U.S. Bureau of Labor Statistics. Available at <http://www.bls.gov/ore/pdf/ec060030.pdf>.

Pissarides, Christopher A. 1985. “Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages.” *American Economic Review* 75(4): 676–690.

Comments on “A New Method for Estimating Time Variation in the NAIRU” by William T. Dickens

Christopher A. Pissarides

It is fitting that there should be a paper on the Beveridge curve at this conference. Like the Phillips article, the original paper on the Beveridge curve was also published in 1958, by Dow and Dicks-Mireaux in *Oxford Economic Papers*. The article was mainly about measurement and made the case for a good correlation between vacancies and unemployment, using British historical data. It offered no theory. It launched a literature, known as “UV-analysis,” on the measurement of vacancies and unemployment and on their relation to excess demand in the labor market. It was concerned with the problem of finding how much unemployment can be reduced with Keynesian demand management policy, given the frictions in the labor market, and in this sense it was a precursor of the later critiques of the Phillips curve. Lipsey (1965) brought the Phillips curve and UV-analysis together, in a paper that addressed many of the issues addressed in Bill Dickens’s paper.

Commenting on this paper has become, in Dickens’s words, an attempt to hit a “moving target.” In order to avoid writing a comment that may turn out to be irrelevant I have therefore decided to comment less directly on what Bill says, and focus instead on the problem that he has posed and discuss some thoughts on how to go about modeling it.

Dickens suggests using information derived from the Beveridge curve to calculate changes in the NAIRU. I totally agree with this objective—ever since its inception, the Beveridge curve has been used to classify reasons for changes in unemployment. These exercises were a precursor to his task. Dickens’s question can be rephrased to the question, did unemployment between t and $t + 1$ change because of a change in the NAIRU,

or because of nominal shocks? Or, more generally, did unemployment change because of a real shock or because of a nominal shock?

Distinguishing between changes in the NAIRU and other changes in unemployment requires two equations. One is the Beveridge curve, which is an equilibrium equation that summarizes the speed of structural change and the frictions in the labor market. The other equation is essentially an equation for the demand for labor. In my view, the best way to think about the Beveridge curve is in terms of the flows in and out of unemployment. By definition, the change in unemployment between period t and $t + 1$ is

$$(1) \quad u_{t+1} - u_t = \text{inflows in } t - \text{outflows in } t,$$

with the stocks measured at the beginning of the period. For the flow terms we can write

$$(2) \quad \text{inflows} = \text{new entry} + \text{job separations, and}$$

$$(3) \quad \text{outflows} = \text{exits} + \text{job acceptances.}$$

The Beveridge curve is defined as the combination of unemployment and vacancies that equates the inflows with the outflows. Writing a theory of the Beveridge curve amounts to modeling each one of the four terms in (2) and (3), and tracing the combinations of vacancies and unemployment that maintain the equality between the inflows and the outflows in the absence of shocks.

Perhaps surprisingly at first, but on reflection not so surprisingly, we get a good approximation to the dynamics of unemployment if we treat unemployment as if it were always on the Beveridge curve (Pissarides 1986, Shimer 2007). It might be surprising at first because with the change in unemployment given by the difference between inflows and outflows, and the Beveridge curve defined as the locus of equality between inflows and outflows, how does unemployment change if we are always on the Beveridge curve? The best way to think about this conundrum is in terms of speeds of adjustment and the length of the period. Treat unemployment as the only unknown in the inflows = outflows condition and assume the period is a quarter. If one of the four terms in (2) and (3) changes because of a shock, unemployment changes fast to restore equality between the new inflows and outflows. In other

words, although the labor market is characterized by frictions, given the size of the shocks that we normally observe, the frictions are sufficiently small that unemployment jumps between one flow equilibrium and the next within a quarter.¹

Consider now the shocks that might make unemployment change in the context of the Beveridge curve. The search and matching theory makes the job acceptance flow the key to the entire framework (Pissarides 2000, chapter 1). In its simplest form it assumes constant job separation rates $s(1 - u_t)$, either zero or constant entry and exit rates, and that the rate of job acceptance is given by the aggregate matching function, $m(u_t, v_t)$. The matching function gives the number of new jobs formed as a function of the workers available to take new jobs, and the number of vacant jobs, v_t . Let f_t denote the average rate of job finding, defined by $f_t = m(1, v_t/u_t)$, and assume that entry and exit are zero. The Beveridge curve is

$$(4) \quad u_t = \frac{s}{s + f(v_t / u_t)}.$$

If nominal shocks have any influence on unemployment in this framework, the channel through which they have it is the vacancy rate, v_t . The vacancy rate is given by the second equation of the system, the demand for labor. If, for example, a positive nominal shock that raises inflation increases the demand for labor because of nominal stickiness somewhere in the system, the vacancy rate increases above trend and unemployment falls. The implied negative relation between unemployment and inflation is the essence of the Phillips curve, and the channel that links the change in the demand for labor with unemployment is the vacancy rate and the matching function.

In terms of the Beveridge diagram derived from (4), the fall in unemployment induced by the nominal shock is represented by a movement along the Beveridge curve. If one were to accept the simple framework underlying equation (4) as a complete characterization of the dynamics of unemployment, the vacancy rate is the only channel through which nominal shocks can be transmitted to unemployment. Any other changes in unemployment, for given vacancies, are changes in the NAIRU. These changes are associated with changes in the rate of labor turnover, s , changes in the matching efficiency of the labor market, represented by

shifts in $f(\cdot)$ for given v_t/u_t , and with changes in the rate of entry into and exit from the labor force. For example, demographics shift the NAIRU, potentially by changing all terms in (4), the rate of labor turnover, the matching efficiency of the labor market, and the rate of entry and exit from the labor force. Unemployment insurance shifts the NAIRU by changing the intensity of search, the efficiency of matching, and so on.

In my view, the best way to uncover changes in the NAIRU associated with shifts in the matching efficiency of the labor market is not to estimate the entire Beveridge relation, as Bill has attempted to do, but to estimate the matching function directly (or the job-finding rate). When I did this for Britain in 1986 I found that most of the changes in unemployment were associated with changes in the NAIRU, although changes in the vacancy rate also played a role. This was to be expected, given that when unemployment was trending up between the late 1960s and the early 1980s the vacancy rate was fluctuating around a flat trend. Several estimates of matching functions by other authors can be used to decompose changes in unemployment between changes due to the vacancy rate and changes due to other factors.² The U.S. experience since 2001, when reliable vacancy data became available through JOLTS, is probably unique in that it attributes virtually all changes in unemployment, save for a small error term, to changes in the vacancy rate, a property that has been emphasized in some of Shimer's recent influential work (for instance, see Shimer 2005 and Elbrahimi and Shimer 2008).

Dickens finds something similar in his estimated Beveridge curves. However, this finding does not necessarily imply a constant NAIRU, even in the simple framework of equation (4). There might be causes of changes in the vacancy rate, which keep the Beveridge curve fixed, and which are real and associated with changes in the NAIRU. For example, consider material shocks. If the price of raw materials goes up and real wages are subject to inertia, vacancies might fall dramatically. Unemployment rises through a movement down the Beveridge curve. The Beveridge curve has no obvious reason to shift in this case.

This is why we need to estimate a second equation, preferably simultaneously with the matching function, before we can confidently calculate the NAIRU. The second equation is a demand for labor equation and is derived from a conventional model of the firm with costs of adjustment due to frictions. The difference between investment-type quadratic adjust-

ment costs and matching frictions is that the costs of adjustment with frictions depend on the tightness of the labor market. At high vacancy-to-unemployment ratios these costs and frictions are higher, because there is more competition between firms for the pool of unemployed workers. The implication of this property is that we can write the dynamic demand for labor equation as a vacancy supply equation and estimate it in terms of all the conventional labor demand regressors, including price misperceptions (Pissarides 1986; Yashiv 2000).

Dickens has a second equation in his model but it is not a labor demand equation. His equation is similar to the one that featured in the very first models of the Beveridge curve (Dow and Dicks-Mireaux 1958). It is essentially the 45-degree line through the origin, which defines the locus of equality points between u and v as the equilibrium points. Modern approaches to the Beveridge curve derive the second equation from optimizing models of the firm and show that the slope of the second curve is a function of the model's parameters.

A more important point about the second equation, however, is this: are we justified in focusing on the vacancy rate as the only variable that can transmit nominal shocks to unemployment? In the context of Phillips curve analysis we are asking whether all shocks to the unemployment rate other than those acting through the vacancy rate are shocks to the NAIRU. In the context of Beveridge curve analysis the question is whether the simple framework in (4) is sufficient.

There has been a lot of work on this issue recently, with reference mainly to business cycle fluctuations in unemployment. These high frequency fluctuations are also the ones that Bill studies in his paper. The upshot of the discussion is that business cycle fluctuations in unemployment are driven both by fluctuations in the inflow rate and the outflow rate (see Shimer 2007, Fujita and Ramey 2007, and Petrongolo and Pissarides 2008). Moreover, for cyclical fluctuations one can ignore the movement in and out of the labor force and focus on movements between employment and unemployment. In that context, the consensus is that about two-thirds of fluctuations are due to the outflow rate, for which the matching function approach serves us well, and another third to the inflow rate. The inflow rate in (4) is the parameter s . The recent empirical literature on the ins and outs of unemployment says that s should not be a parameter but a cyclical variable. A complete model of the NAIRU

derived from the Beveridge curve should account for the endogeneity of job separations.

Although good theoretical models of the endogeneity of job separations exist, it is much more difficult to find good empirical or quantitative models of separations.³ I think this is likely to be the main sticking point in the task that Dickens set himself. Because job separations vary and are negatively correlated with job accessions, it is plausible to assume that these are driven by the optimizing decisions of firms and workers in response to shocks. Some of those shocks are nominal, and if there are nominal rigidities of the kind analyzed in Phillips curve models, some changes in the parameter s in (4) are changes associated with nominal shocks, namely, not changes in the NAIRU. But changes in s shift the Beveridge curve. It follows that in a general model of the NAIRU there are changes in unemployment that are not caused by changes in vacancies, and which are not changes in the NAIRU. Therefore, identifying all changes in unemployment that take place for a given vacancy rate as changes in the NAIRU would be a mistake.

A challenge that is facing both search and matching theory and modern Phillips curve analysis is how to explain the fact that on average about one-third of fluctuations in unemployment are due to shocks to job separations (or, at least, to the unemployment inflow rate) and yet for long stretches of time the vacancy-unemployment scatter of points is tightly distributed around a fixed Beveridge curve. As far as I know there is no paper in the literature yet that does that, and so there is no model that can convincingly be used to provide a framework for the estimation of the NAIRU from Beveridge curve analysis. But following the approach that I outlined in this comment, under the assumption that all nonrandom shifts in the Beveridge curve are changes in the NAIRU, is a good first approximation to the data.⁴

Notes

1. In my examination of British and other European data, the only time that the assumption of flow equality in quarterly data was not a good working assumption was the two-year period of the large “Thatcher shock,” 1979–1981. See Pissarides (1986) and Petrongolo and Pissarides (2008). Shimer (2007) does not report any period when this assumption was badly violated for the United States.

2. See Blanchard and Diamond (1989) for early U.S. estimates and Petrongolo and Pissarides (2001) for a survey of several estimates.
3. For the theory see Mortensen and Pissarides (1994) and Caballero and Hammour (1996). For more discussion of the empirics of job separations see Davis and Haltiwanger (1999).
4. See Ebrahimi and Shimer (2008) for a promising attempt at explaining simultaneously the tightly distributed points in Beveridge space and the variance in the separation rate. They focus on the post-2001 data, when there are no shifts in the Beveridge curve. The problem of reconciling periodic shifts with long periods of tightly distributed $u - v$ points remains.

References

- Blanchard, Olivier Jean, and Peter A. Diamond. 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity* 1: 1–60.
- Caballero, Ricardo J., and Mohamad L. Hammour. 1996. "On the Timing and Efficiency of Creative Destruction." *Quarterly Journal of Economics* 111(3): 805–852.
- Davis, Steven J., and John C. Haltiwanger. 1999. "Gross Job Flows." In *Handbook of Labor Economics*, Vol. 3B, ed. Orley Ashenfelter and David E. Card, 2711–2805. Amsterdam: Elsevier.
- Dow, J. C. R. and L. A. Dicks-Mireaux. 1958. "The Excess Demand for Labour: A Study of Conditions in Great Britain, 1946–56." *Oxford Economic Papers*, New Series 10(1): 1–33.
- Ebrahimi, Eshan, and R. Shimer. 2008. "Stock-Flow Matching." Working paper. Available at <http://robert.shimer.googlepages.com/stock-flow-print.pdf>
- Fujita, Shigeru, and Garey Ramey. 2008. "The Cyclicalities of Separation and Job Finding Rates." Forthcoming in *The International Economic Review*. Available at http://weber.ucsd.edu/~gramey/Pubs/FR_IER_1-10-08.pdf.
- Lipsey, Richard G. 1965. "Structural and Deficient-Demand Unemployment Reconsidered." In *Employment Policy and the Labor Market*, ed. Arthur M. Ross, 210–255. Berkeley: University of California Press.
- Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies* 61(3): 397–415.
- Petrongolo, Barbara, and Christopher A. Pissarides. 2001. "Looking into the Black Box: A Survey of the Matching Function." *Journal of Economic Literature* 39(2): 390–431.
- Petrongolo, Barbara, and Christopher A. Pissarides. 2008. "The Ins and Outs of European Unemployment." *American Economic Review* 98(2): 256–262.
- Pissarides, Christopher A. 1986. "Unemployment and Vacancies in Britain." *Economic Policy* 1(3): 500–541.

Pissarides, Christopher A. 2000. *Equilibrium Unemployment Theory*, second edition. Cambridge, MA: The MIT Press.

Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95(1): 25–49.

Shimer, Robert. 2007. "Reassessing the Ins and Outs of Unemployment." Working Paper No. 13421. Cambridge, MA: National Bureau of Economic Research.

Yashiv, Eran. 2000. "The Determinants of Equilibrium Unemployment." *American Economic Review* 90(5): 1297–1322.