# Shifts in the Beveridge Curve, Job Matching, and Labor Market Dynamics

*Hoyt Bleakley and Jeffrey C. Fuhrer*

*Senior Research Assistant, and Vice President and Economist, respectively, Federal Reserve Bank of Boston.* *The authors thank Lynn Browne, Richard Kopcke, and Scott Schuh for helpful comments, Joe Ritter for access to the gross worker flows data, and Dan Feenberg at the National Bureau of Economic Research for help in obtaining MORG data.*

In the 1940s, when William Beveridge first identified the negative relationship between unemployment and job vacancies that now bears his name, his goal was to determine how far the economy was from "full employment." His definition of full employment—that vacancies slightly exceed unemployment—is quite distant from the later postwar experience. Nevertheless, the Beveridge curve—the scatter plot of unemployment rates versus vacancy rates—is still often used to summarize the state of the labor market. A stylized Beveridge curve is displayed in Figure 1. As the figure shows, the position on the curve can indicate where the economy is in the business cycle: Recessions, for example, are generally times of high unemployment and few job vacancies, corresponding to points on the lower right branch of the curve. In addition, the location of the Beveridge curve relative to the origin has been used to indicate the overall level of labor market activity, sometimes interpreted as the intensity of "reallocation"—the movement of workers from one job to the next, often from one sector to another, in the economy.

However, the Beveridge curve is not a structural economic relationship. That is, workers and firms do not consciously decide to make unemployment negatively related to vacancies. Instead, the decisions of workers and firms about accumulating skills, hiring workers, and setting wages result indirectly in the patterns of unemployment and vacancies that we call the Beveridge curve. Changes in these underlying behaviors will influence both the particular combination of unemployment and vacancies at any time on a fixed Beveridge curve, and the overall location of the curve—close in or further out from the origin. Thus, in order to understand the labor market implications of movements in the Beveridge curve, we must first understand the labor market activities that give rise to the Beveridge curve. This article will begin with simple observations about changes in the location of the Beveridge curve and attempt to explain these changes as the result of shifts in underlying economic processes.

Figure 1

*Stylized Beveridge Curve*

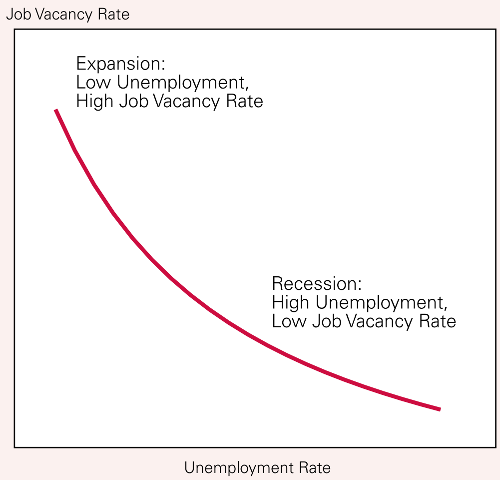


Figure 2

*Gross and Net Employment Flows*

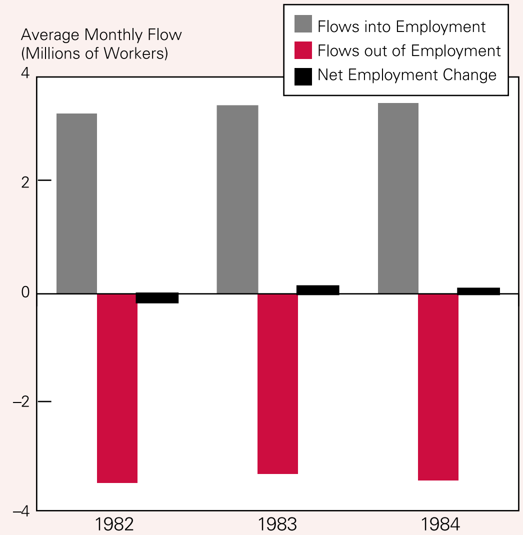


Source: U.S. Bureau of Labor Statistics; Abowd and Zellner (1985).

One key to unraveling the movements of the Beveridge curve will be analysis of both the *stock* of unemployed workers and vacant jobs that make up the points on the Beveridge curve and the *flow* of workers into and out of unemployment. A wealth of recent research has grown from the discovery of the surprisingly large gross flows of jobs and workers—firms' gross creation and destruction of jobs (see Davis, Haltiwanger, and Schuh 1996)—and workers' transitions among the states of employment, unemployment, and not being in the labor force (see, for example, Blanchard and Diamond 1989). These authors and many others suggest that much may be lost in summarizing labor market activity by *net* activity—the net employment flows, for example, published in the U.S. Bureau of Labor Statistics' monthly Household and Establishment Survey releases—and that key features that characterize labor market equilibrium are better understood through the gross flows. Figure 2 displays the relative sizes of the net and gross flows into and out of employment. Although 1982 was a time of underutilized resources and net job loss (about 2 million jobs were lost on net in that year), it was still a year of substantial gross flows both into and out of employment. Moreover, in 1984, a year of significant net employment increase, the gross flows were again large in both directions. Indeed, during the period 1968 to 1984, gross worker flows in and out of employment averaged about 20 times higher than net employment changes.

> *The Beveridge curve—the scatter plot of unemployment rates versus vacancy rates—still is often used to summarize the state of the labor market.*

One key labor market feature that is masked by net flows is the ongoing process of job matching, the process by which workers searching for jobs find a "match" with an employer who has a job vacant. The actual process of searching for work and for workers is no doubt extremely complex. However, a simple approximation to the process at the aggregate level posits that some fraction of the stock of unemployed workers and of the stock of vacant jobs combine each period to make job matches (new hires). This yields a gross flow of newly employed workers that is quite large relative to the net change in the stock of unemployed workers.

The job-matching process is an important component of the labor market. Even in the healthiest of economic climates, large stocks of unemployed workers and vacant jobs wait to be matched. The efficiency with which workers are matched to available jobs will be a key determinant of the duration of unemployment and of the ability of the economy to fully utilize all available resources. Thus, a model of the job-matching process should be a central component of any adequate description of the long-run normal state or "equilibrium" in labor markets.

*A model of the job-matching process should be a central component of any adequate description of the "equilibrium" in labor markets.*

This conversion of the stocks of unemployed workers and vacancies into a flow of new hires is surely one of the mechanisms undergirding the Beveridge curve. However, one cannot look at the plot of the Beveridge curve and unambiguously identify times of high and low job-matching efficiency. That inference can be drawn only when data on the flow of new hires are added to the picture, in essence measuring the "output" that the unemployment and vacancy "inputs" are producing. One goal of this study will be to unravel the extent to which changes in the Beveridge curve over time may be attributed to changes in the process of job matching.

The understanding of these aspects of labor markets is of direct relevance to monetary policy. Because the Federal Reserve is concerned about maintaining stable employment and prices, it needs to know at what level employment should be stabilized. Presumably, the Fed's goal with regard to employment is to remain as close to full employment as possible, given its inflation goals. Thus, the Fed must have the most accurate possible description of full-employment equilibrium in the labor markets.

This article will begin with a brief discussion of the Beveridge curve and an examination of its behavior over the past 30 years. It proceeds to an analysis of the movements of the Beveridge curve. In so doing, the article will discuss in some detail some of the issues surrounding the job-matching process and attempt to estimate the extent to which changes in the job-matching function are responsible for changes in the position of the Beveridge curve. The article will also address other potential sources of shifts in the Beveridge curve, including shifts in the age and gender composition of the labor force and changes in the amount of "churning" in the labor market.

## I. The Beveridge Curve

Figure 3 displays an empirical Beveridge curve—a plot of the civilian unemployment rate (total unemployed as a percentage of the labor force, the horizontal axis) versus the job vacancy rate (imputed vacancies as a percentage of the labor force, the vertical axis), monthly from 1960 to 1996. A number of interesting features are immediately apparent.

First, during most of the past 37 years, unemployment and vacancies have been negatively related. Higher unemployment has been associated with lower vacancies, and lower unemployment with higher vacancies. This negative association is routinely attributed to a common-sense cyclical pattern in the labor market: When the labor market is "tight" and demand for labor is high, most workers who wish to work have found employment, so the unemployment rate is low. Correspondingly, the vacancy rate is relatively high, as employers experience some difficulty finding qualified workers to fill job openings. During a "slack" labor market with weak labor demand, few employers are posting job offerings, so vacancies are low and unemployment is high.

In addition to this common-sense correlation between unemployment and vacancies, the Beveridge curve in Figure 3 provides evidence of two rather prominent shifts. During the 1970s and early 1980s, the curve shifted outward to higher levels of both vacancies and unemployment. The normal cyclical progressions around the curve occurred as before, but in all cases both unemployment and vacancies were higher. However, more recently the Beveridge curve appears to have shifted *back* to approximately its 1960s locus. A clear transition period in the late 1980s (the cluster of boxes outlined in red) has been followed by pairs of vacancy and unemployment rates at least as

Figure 3

*Beveridge Curve with Estimated Shift Lines*

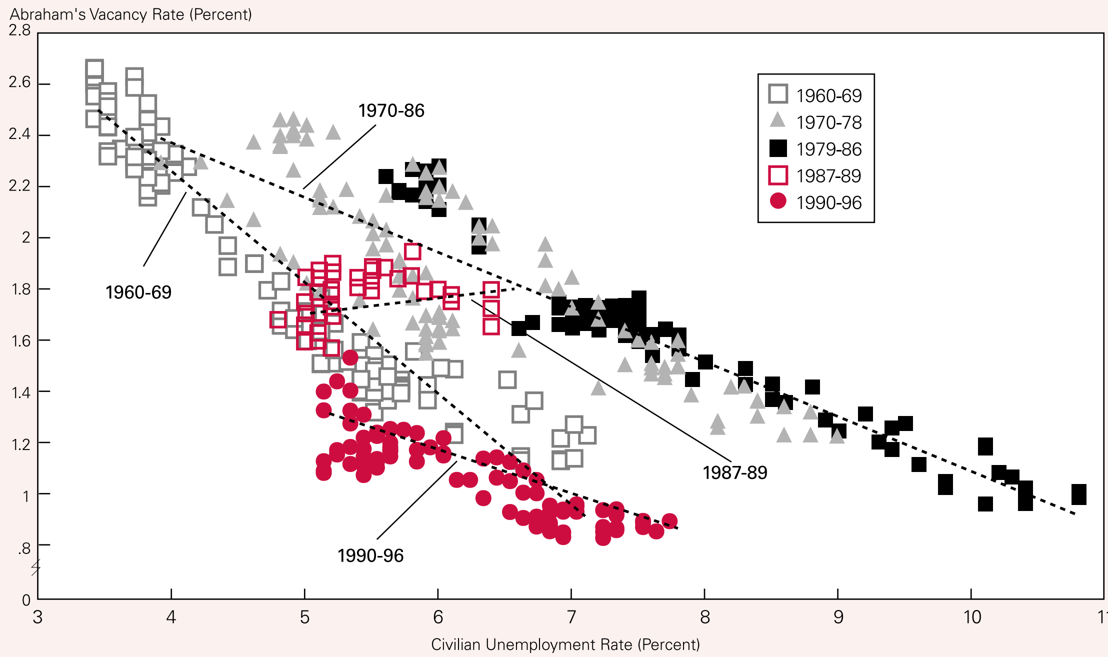


Source: U.S. Bureau of Labor Statistics; Abraham (1987); The Conference Board; and authors' calculations.

low as those of the 1960s. Regression tests of the statistical significance of each of these shifts confirm what is apparent to the eye. Interestingly, while the 1990s look much more like the 1960s than like either of the two succeeding decades, one can handily reject the hypothesis that the 1990s Beveridge curve has the same slope and intercept as the 1960s curve. The shallower slope of the 1990s curve is significantly different from that of the 1960s; it implies that a 1 percentage point decrease in the unemployment rate today is associated with a 0.2 percentage point rise in vacancies, as compared with a 0.4 percentage point rise in vacancies in the 1960s. Nonetheless, the recent location of the Beveridge curve is much closer to its location in the 1960s than to that in the 1970s. Indeed the unemployment and vacancy rates from 1995 and 1996 suggest that the Beveridge curve is moving even further inward—to territory not explored since the 1950s.

The central question of this article is: Why did the Beveridge curve shift? Three candidate explanations for the shift will be considered: a decrease in the degree of churning in the labor market, a slowdown in the growth of the labor force, and an increase in the efficiency of the worker-job matching process. The next section outlines a simple model of the labor market. This is followed by an examination of empirical evidence bearing on each explanation, an attempt to quantify the contribution of each explanation to the observed Beveridge curve shifts, and discussion of some underlying demographic changes that are consistent with these shifts—notably the entrance and absorption into the labor force of the baby boom generation.

## *II. A Simple Model Underlying the Beveridge Curve*

A simple model of unemployment, vacancies, and labor market gross flows will be used as a framework to analyze the candidate explanations for the Bever-

Figure 4

*A Simple Model of Unemployment and Vacancies Underlying the Beveridge Curve*

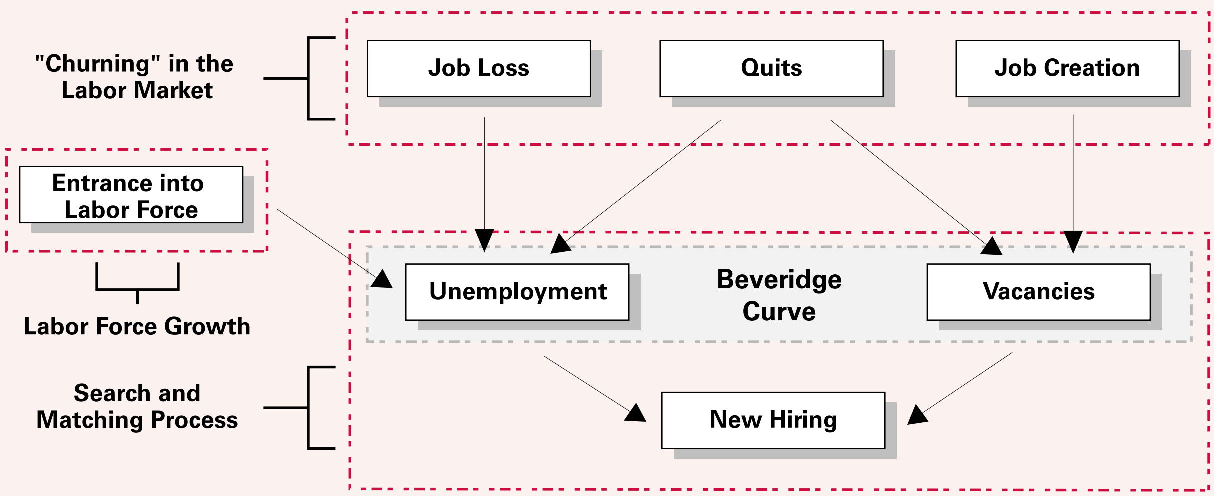


idge shifts. It is important to note that the Beveridge curve is not a fundamental economic relationship in the model. Instead, the flows into and out of unemployment and vacancies, together with the job-matching process, determine the outcomes for unemployment and vacancies that are summarized in the Beveridge curve. In this stylized model, the equilibrium or settling point for the labor market is characterized by a balance of inflows and outflows, not by particular values of the stocks of unemployed workers and vacant jobs.[1]

Figure 4 presents a schematic of the simple model. The Beveridge curve is represented by the grey box outlined at the center of Figure 4; it is determined by the flows of workers and jobs into and out of unemployment and vacancies, as indicated by the arrows. As the figure shows, unemployment arises as the result of flows of job losers, job leavers ("quits" or voluntary job separations), and flows into the labor force. Vacancies arise from the expansion of firms ("job creation") and from quits. The outflows from the Beveridge variables are new hires: Workers leave unemployment upon finding a job and jobs are no longer vacant once a worker is hired. The flows fall into three broad categories: labor market reallocation or "churning," labor force growth, and the search and matching process.

This simple framework is not intended as a complete description of the labor market. Instead, the framework embodies a number of simplifications that allow one to focus only on those features of the labor markets that are believed to help explain shifts in the Beveridge curve. A large number of new entrants to the labor force become employed without a spell of unemployment—college graduates, for example. These flows are ignored in this simple model. The framework also abstracts from movement out of the labor force, the hiring of people who are never observably unemployed, and the closing of vacancies before they are filled.

Job loss, quits, and job creation are related to the overall pace of reallocation or "churning" in the economy. Reallocation occurs even when the economy is stable, as some firms expand and others contract for firm- or industry-specific reasons. The pace of reallocation increases during times of economic upheaval—most often during recessions or rapid expansions—as the health of the economy drives more firms to con-

[1] This is essentially the same as the model presented by Blanchard and Diamond (1989). Interested readers are referred to their paper for further detail.

tract or expand significantly, leading to greater flows of workers and jobs. Thus, changes in the pace of reallocation imply potentially large movements in the *gross* flows in the labor market—flows into or out of employment. More churning also implies lower average job tenure, higher turnover, and more time spent moving among firms (and perhaps sectors) of the economy. An increase in churning means that each month more workers flow into unemployment and more new vacancies are posted. Such an increase would shift the Beveridge curve outward.

*Job loss, quits, and job creation are related to the overall pace of reallocation or "churning" in the economy, which occurs even when the economy is stable as some firms expand and others contract for firm- or industry-specific reasons.*

The growth in the labor force is another important component of the model. As new workers enter the labor market, they join the ranks of the unemployed searching for work. Holding other things fixed, higher levels of labor force growth mean more unemployment, since more workers are searching for jobs at any particular time. In the short run, vacancies may not adjust fully to an increase in labor force growth; in the long run, vacancies will likely increase roughly in line with unemployment. Thus, an increase in the rate of labor force growth would shift the Beveridge curve out and to the right.

Changes in the rate of labor force growth are due primarily to shifts in the demographic composition of the working-age population. For example, a major source of new entrants to the labor force is people finishing school, which is largely a function of the birth rate about 20 years before. Changes in the labor force participation of women and in immigration have also caused changes in the growth rate of the labor force.

Finally, consider the process whereby workers (firms) search for jobs (workers). As Figure 4 shows, this process determines the outflows from unemployment (and vacancies) into new hiring. In reality, this "job-matching" process is far more complicated than depicted here, as it requires finding a match between the skills, location, and industry of workers and firms. This study abstracts from these complications and concentrates simply on how productive this process is: literally, on how many new matches between firms and workers can be fashioned for a given stock of unemployed workers and vacant jobs. The overall efficiency with which workers and vacancies are matched (and new hires generated) determines the outflow rates from the Beveridge variables. A more efficient matching process creates higher outflows from unemployment and vacancies, and thus lower levels of both over time. Thus, greater job-matching efficiency would bring about an inward shift of the Beveridge curve.

In theory, then, shifts in the Beveridge curve can be attributed to changes in each of the three underlying flow components described above. First, increases or decreases in job-matching efficiency would be associated with shifts inward or outward in the Beveridge curve. The reasons behind the changes in efficiency could be many, including improvements in communications that make geographically dispersed job openings known to more potential applicants, and improved mobility of workers that broadens the menu of available jobs for a job-seeker. Second, shifts in the flows into unemployment from outside the labor force would alter the number of unemployed, for a given matching efficiency. For example, a decrease in flows into the labor force would shift the Beveridge curve down and inward. Finally, a decrease in labor market churning—the number of workers searching for jobs and firms searching for workers—would shift the

*Changes in the rate of labor force growth are due primarily to shifts in the demographic composition of the working-age population: people finishing school, changes in the labor force participation of women, and immigration, for example.*

Beveridge curve inward, independent of the state of the job-matching process or of labor force participation.

## III. Evidence for the Explanations of Beveridge Curve Shifts

What do the data show about shifts over time in job matching, labor force growth, and churning? The next few sections will examine evidence bearing on these potential explanations. A model of the job-matching process will uncover shifts in the process that coincide with the shifts in the Beveridge curve. Evidence on changes in churning intensity and labor force growth will be presented to show that these changes also coincide with the shifts in the Beveridge curve. The model outlined above will then be simulated, in order to determine whether the combined effect of these shifts can explain the direction and magnitude of the Beveridge curve shifts observed since 1979.

### The Job Matching Function

Following Blanchard and Diamond (1989), this study employs the useful macroeconomic fiction of the matching function, a highly stylized model of the process by which workers find new jobs. The matching function is roughly analogous to the standard production function in economics: Inputs are combined to produce a flow of output. The inputs to the standard production function are the existing stocks of capital and labor; they are combined to produce a flow of finished goods. In the matching function, the inputs are the existing stocks of unemployed workers and vacancies; the output is a flow of new hires or job matches. Symbolically, the matching function may be written as

$$H_t = A\ f(U_t, V_t)$$

where H stands for new hires, the gross flows into employment in any period; A is the overall productivity factor; U stands for unemployment, or more generally the stock of workers available to fill vacancies; and V is the number of vacant jobs.

A number of data and definitional issues are involved in the estimation of the matching function. The lack of vacancy data is a problem for studying the matching process, as is choosing the most appropriate definition of the pool of workers who are available for a match. For the second issue, careful thought must be given to the treatment of (primarily unionized) workers who have been laid off and are available for recall (and in most cases are subsequently recalled). Moreover, some caution is warranted by the difficulty in measuring workers who make a job match without a spell of unemployment. (See Akerlof, Rose, and Yellen (1986) for a detailed treatment of this last problem.) This study deviates somewhat from the decisions made by Blanchard and Diamond. The choices of data for the estimation are discussed below.

*Vacancy rates*. The fundamental shortcoming of the available vacancy data is that they are not a true vacancy series but are instead derived from the Conference Board's help-wanted advertising index. The true number of vacancies is plausibly related to such a measure. However, the relationship has likely changed over time: The preferred mode of making vacancies known changes; the number of jobs per listing varies; and so on. Katharine Abraham (now Commissioner of the Bureau of Labor Statistics) was the first to construct an adjusted help-wanted index that could reasonably proxy for the vacancy rate (1987). She compared her constructed series to the few published series on vacancies and found the correlation during those short periods to be quite good. She also suggested adjustments to the index to capture the changing relationship between help-wanted advertising and vacancies; her methods have been adopted here.

*The pool of workers available for a match*. While it is conceptually clear that one argument to the matching function is the quantity of workers available for a job match in the current period, it is considerably less clear which data most closely correspond to this quantity. If all people of working age are separated into three groups—employed, unemployed, and not in the labor force—then, in theory, some workers in each of these groups may be available for a job match. The most difficult to measure are those who might make a job match while employed; they are discussed below. Those who are unemployed clearly belong in the group available for job match, with the exception of workers on layoff, awaiting recall. In addition, some workers who are not in the labor force "would be willing to work if a job offer were made to them," and thus might also be properly considered as part of the pool of available workers.

*Layoffs and recalls*. The movements from employment to unemployment and back that occur because of a layoff (subject to recall) and subsequent recall should not be considered as part of the job-matching function to be modeled here. Workers waiting for recall are typically not searching for employment, and firms that recall workers do not advertise vacancies to do so. Job losers on layoff are excluded from the stock of unemployed used in this study, and an estimate of the laid-off workers who are recalled to work is excluded from the estimated flow of new hires.[2]

*Employment-to-employment matches*. The problem is that no data are available from the Current Population Survey (CPS) that reflect these employment-to-employment flows. Akerlof, Rose, and Yellen (1986) estimate that 40 percent of all workers who quit jobs move directly to new jobs without a spell of unemployment. Blanchard and Diamond (1989) use 40 percent of the manufacturing quits, scaled to be representative of the entire economy, as an estimate of the flow of these job-to-job matches. However, the manufacturing quits series is not available after 1981. Moreover, the correspondence of manufacturing quits to the rest of the economy is questionable, and the percentage of quitting workers who are immediately reemployed should not be assumed constant. In addition, the pool of employed workers from which these matches are drawn cannot be distinguished from the pool of employed workers who are not seeking a new job. While unemployed workers can generally be presumed to be looking for a job match (they are classified as unemployed because they answer "Yes" to the question "Are you currently looking for work?" on the CPS survey), the same cannot be said of every worker who is currently employed. Nonetheless, some fraction of the vacancies will be filled by the currently employed. However, there is no way to determine which vacancies are filled by employment-to-employment matches.

In the end, job-to-job matches were excluded from this matching function. No credible method of constructing the gross hiring flows that correspond to such matches was available, nor any obvious way of constructing the stock of employed workers from which such matches are made.

*Matching function data.* The data that enter the matching function are defined as follows:

| Variable | Definition |
|---|---|
| Flow of New Hires | Gross flows from unemployment (U) into employment (E), less imputed outflows from U on layoff. Does not include E→E or N→E flows (flows from not in the labor force to employment). |
| Unemployment Rate | Civilian unemployment rate, less job losers on layoff. (Alternate measure includes not in the labor force, would accept job if offered.) |
| Vacancy Rate | Abraham's (1987) adjusted Help-Wanted Index; adjustments extended using linear trend, as in Blanchard and Diamond (1989). |

### *Estimates of the Matching Function and Its Stability over Time*

While the matching function is obviously a simplification of the complex process by which workers find the right job and employers find the appropriate worker, it nevertheless provides a useful summary of the overall efficiency with which such matches are made. To put some meat on the bones of the stylized matching function described above, it will be assumed that unemployment and vacancies combine to make a match (or new hire) according to a Cobb-Douglas function, $H_t = A_t U_t^{\alpha} V_t^{\beta}$. For this function, hires would increase by a factor of $2^{\alpha+\beta}$ when both unemployment *(U)* and vacancies *(V)* double. If $\alpha+\beta=1$, the matching function exhibits constant returns to scale; that is, a doubling of $V$ and $U$ yields a doubling of new hires. Finally, the overall efficiency of the matching function is summarized in $A$; if $A$ doubles, then the number of new hires doubles for given levels of $U$ and $V$. This concept is equivalent to a boost in overall productivity that raises the level of output for given levels of capital and labor.

The availability of data places two restrictions on the sample period. First, the flows of new hires used in this paper are available only beginning in 1979. Second, the CPS was redesigned beginning in 1994, and flows from 1994 forward may not be comparable with those prior to 1994. As a result, the sample for the initial estimate of the matching function is restricted to the monthly sample from January 1979 through May 1993. This sample spans two important shifts in the Beveridge curve, from the furthest-out curve of the late 1970s and early 1980s to the flattened (transition) period of the late 1980s, and on to the current, inward-

[2] The "new hires" measure is the flow of workers from unemployment to employment from the Current Population Survey, minus estimated recalls. To correct for measurement error, the *average* Abowd-Zellner (1985) adjustment from 1968 to 1986 is applied to the U→E flows in our sample period. (This is a departure from their method; however, the adjustments are not available for the full sample.) Recalls are removed from this flow by subtracting away the outflows from the "unemployed, on layoff" state. We compute the outflow rate by subtracting the net change in job losers on layoff from the gross inflows. (That is, $UL_{in} - UL_{out} = \Delta UL \gg UL_{out} = UL_{in} - \Delta UL$.) For the regressions presented, we assume that this represents laid-off workers who were recalled. Not all of this outflow represents recalls; some of these workers may have left the labor force. However, our regression results are not sensitive to assuming that a smaller fraction of these outflows are in fact recalls. Future work will attempt a better estimate of this flow from layoff to employment.

shifted curve of the 1990s. The matching function will have to explain how the falling stock of vacancies led to the same flow of new hires on average, with the same average stock of unemployed.

Table 1 presents estimates of the matching function for various specifications. The first column shows the baseline specification, which does not impose constant returns to scale and allows a time trend to enter. Perhaps reassuringly (especially given the uncertainties inherent in the construction of the vacancy rate series), both unemployment and vacancies enter as economically and statistically significant determinants of the flow of new hires.[3,4] The sum of the coefficients on unemployment and vacancies is 0.86, suggesting modestly decreasing returns to scale. However, as the last row in the next column shows, one cannot reject the restriction of constant returns to scale in the matching function. In addition, the simplest matching function explains 30 percent of the variation over time in the flows into employment, with no discernible evidence of missed serial correlation in the residuals.[5] Thus, the model provides a reasonably complete accounting of the movement over time in new hires. Finally, the time trend enters significantly and positively, suggesting that matching efficiency has increased over the sample. This result is consistent with the interpretation that some of the inward shift in the Beveridge curve arises from improved job-matching efficiency. Tests were conducted to see whether the shifts in efficiency can be interpreted as occurring around the times indicated by the graph of the Beveridge curve. Moreover, these tests attempt to sort out whether this shift is more properly interpreted as an overall increase in efficiency (a shift in $A$) or an increase in the efficiency with which unemployment or vacancies yield a job match (a shift in $\alpha$ or $\beta$).

*While the matching function is a simplification of the complex process by which workers find the right job and employers find the appropriate worker, it provides a useful summary of the overall efficiency with which such matches are made.*

The second and third columns display the matching function estimates, imposing constant returns to scale and allowing for shifts over time in the overall efficiency and the coefficients on unemployment and vacancies in the matching function. The breakpoint, January 1987, was chosen to conform with the observed shifts in the Beveridge curve in Figure 3.

As the third column shows, the significant trend estimate in the first two columns can be well-approximated by a one-time shift in the matching efficiency parameter $A$ starting in 1987, the beginning of the backward shift in the Beveridge curve. Taking the coefficient estimates literally implies that matching efficiency improves by about 13 percent beginning in 1987. That is, for the same stock of unemployed and vacant jobs, the job-matching process produced a 13 percent larger flow of new hires after 1986 than it did before 1987. The ability of the matching function to explain the time-variation in flows into employment decreases only slightly with this single efficiency shift, to an $R^2$ of 0.28.

The fourth column contains estimates of a shift in the coefficients on unemployment and vacancies, preserving the restriction that the matching function exhibit constant returns to scale. Again, a one-time significant increase in the coefficient on unemployment (with a corresponding decrease in that on vacancies) also approximates the overall increase in matching efficiency over the sample. Such a shift implies that, other things equal, unemployed workers are more productive in the matching process and, on the other side, firms with vacant jobs are somehow less

[3] The matching function was also estimated using the raw data from the help-wanted advertising index, without Abraham's adjustments. The results are similar to those presented in Table 1, suggesting that unaccounted-for shifts in the relationship between the index and vacancies are not critical to these results.

[4] These results are similar to line 6 in Blanchard and Diamond's Table 2. However, in the main results they present (lines 1–5), they include E→E flows, imputed using manufacturing quits, on the left-hand side of the regression. Since manufacturing quits and vacancies are both highly procyclical, this explains both their higher $R^2$ as well as the higher coefficient on vacancies.

[5] The Cobb-Douglas matching function imposes a unit elasticity of substitution between unemployment and vacancies. The constant elasticity of substitution (CES) matching function allows the elasticity to vary, but nests the Cobb-Douglas matching function as a special case. For our data and sample, the CES matching function estimates suggest that the Cobb-Douglas restriction on the elasticity of substitution is entirely appropriate; that is, one cannot reject from the viewpoint of the CES the restriction that the elasticity of substitution equals one. Therefore, only Cobb-Douglas results are reported for the remainder of the study.

Table I
*Regression Results from Cobb-Douglas Matching Function (Estimated in Logarithmic Terms)*

| Dependent Variable: Hires from Unemployed | Baseline (1) | Constant Returns to Scale (2) | (3) | (4) | Instrumental Variables (5) | (6) |
|---|---|---|---|---|---|---|
| Inputs: | | | | | | |
| Unemployment-Layoffs, | .56 | .64 | .69 | .65 | .54 | .76 |
| one lag, except for IV | (.09) | (.02) | (.02) | (.02) | (.10) | (.12) |
| Vacancies | .30 | .36 | .31 | .35 | .31 | .38 |
| one lag, except for IV | (.07) | (.02) | (.02) | (.02) | (.08) | (.09) |
| Time trend x 100 | .13 | .14 | | | .13 | |
| | (.02) | (.02) | | | (.02) | |
| Constant | −.64 | −1.85 | −1.30 | −1.24 | −0.60 | −2.52 |
| | (1.32) | (.08) | (.04) | (.03) | (1.47) | (1.68) |
| Shifts in: | | | | | | |
| Intercept, 1987–1994 | | | .13 | | | .14 |
| | | | (.02) | | | (.02) |
| Unemployment Factor Share, 1987–1994 | | | | .09 | | |
| | | | | (.01) | | |
| Vacancy Factor Share, 1987–1994 | | | | −.09 | | |
| | | | | (.01) | | |
| Test Statistics | | | | | | |
| Adjusted $R^2$ | .30 | .30 | .28 | .29 | | |
| Durbin-Watson | 2.14 | 2.12 | 2.07 | 2.11 | 2.22 | 2.17 |
| p for F-test on constant-returns-to-scale restriction | | .41 | .67 | .52 | | |

Note: Standard errors in parentheses.

productive. Because, as in historical experience, there is a much greater stock of unemployment than vacancies, an increase in productivity of *U* (even at the expense of *V*) results in a net increase in overall match productivity. It is not possible to discriminate between increases in overall matching efficiency (increase in *A*) and increases in *U*-specific efficiency (increases in $\alpha$). Nonetheless, these results point to increased efficiency in the matching process.

Finally, the fifth and sixth columns allow unemployment and vacancies to contemporaneously determine the flow of new hires. These specifications attempt to control for the possible simultaneous determination of *U*, *V*, and *H* by instrumenting for current values of *U* and *V* in the regression.[6] The regressions yield similar results to those in the preceding four columns.

Overall, the results in this section show that the efficiency with which unemployed workers are matched with job vacancies has increased significantly in the late 1980s and early 1990s. One cannot distinguish strongly between an increase in the overall efficiency of the job-matching process and an increase in the contribution to a match from the pool of unemployed workers. Nonetheless, the results suggest that part of the inward shift in the Beveridge curve in the late 1980s may be attributed to improved matching efficiency.

*The efficiency with which unemployed workers are matched with job vacancies has increased significantly in the late 1980s and early 1990s.*

[6] The instruments include lags of *U* and *V*, as well as lags of the federal funds rate, the rate of inflation, and the index of industrial production.

Figure 5

*Labor Force Growth*

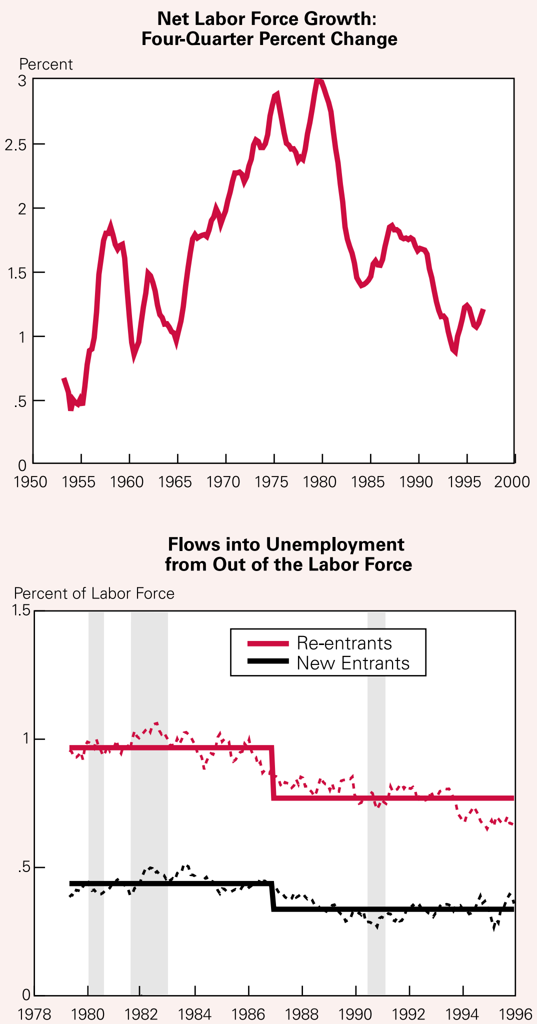


Note: Gray shading shows recessions.
Source: U.S. Bureau of Labor Statistics, *Current Population Survey.*

### *Shifts in Labor Force Growth*

The study now examines the extent to which a slowdown in the rate of labor force growth—specifically, a reduction in the flow of new entrants to the labor market—corresponds to the inward shift in the Beveridge curve. Such a slowdown did indeed take place over the sample period studied, and roughly coincides with the Beveridge shifts. The top panel of Figure 5 shows the net labor force growth rate, which averaged about 2.5 percent per year in the first half of the sample period and about 1.5 percent per year in the second half. Similarly, the bottom panel displays the inflow into unemployment of people who were previously not in the labor force. The solid lines show the average flow rate for the pre- and post-1987 periods; a noticeable and statistically significant drop in the flows of new entrants and reentrants occurs around 1987. This drop can be attributed to two prominent sources: the entrance of the baby bust (the aftermath following the baby boom) into the labor force, and the leveling-off of labor force participation among women. Each of these factors caused a reduction of new entrants to the labor market, implying a lower unemployment rate and a shift of the Beveridge curve to the left and perhaps upward.

These changes likely precipitated only a modest drop in unemployment, however. A rough calculation suggests that they would likely reduce the unemployment rate by less than half a percentage point. When the annual growth rate of the labor force slows from 2.5 percent to 1.5 percent, the monthly inflow rate slows by about 0.08 percentage point. Over the sample period, the average duration of an unemployment spell varied from about two to four months. If this duration applied to new entrants, then the slowdown in labor force growth would reduce the unemployment rate by between 0.16 and 0.32 percentage point.

> *The significant drop in the flows of new entrants and reentrants to the labor force around 1987 can be attributed to the entrance of the baby bust (following the baby boom) and the leveling-off of labor force participation among women.*

Therefore, decreases in the growth rate of new entrants *per se* had only a modest impact on the location of the Beveridge curve. However, the overall effect of these shifts could be greater to the extent that they also imply a change over time in the *composition* of the labor force. For example, as the baby boomers reached their twenties, an age typically associated with greater frequency of unemployment spells (and

Figure 6

*Labor Market Churning*

Flows into and out of Employment as a Percentage of the Working-Age Population

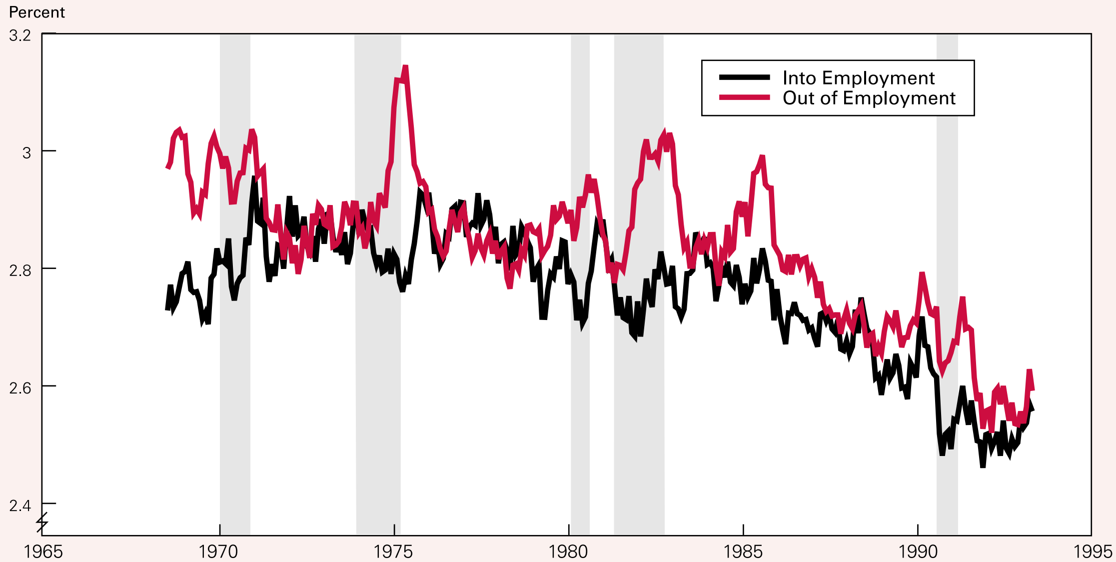


Note: Gray shading shows recessions.
Source: U.S. Bureau of Labor Statistics, *Current Population Survey.*

lower attachment to the labor force), the unemployment rate was likely affected. This possibility will be discussed below.

### *Shifts in Labor Market "Churning"*

A third potential source of the inward shift of the Beveridge curve is a drop in the degree of churning in the labor market. Recall that in the simple model, churning is associated with the overall level of job destruction, job creation, and quits. Figure 6 displays another way of measuring the intensity of churning: the flows into and out of employment. The total flows into employment are represented by the black line, and flows out of employment by the red line. As the figure shows, the amount of churning in the economy has declined noticeably since 1987, the time when a significant inward shift began in the Beveridge curve. Simple regressions confirm the statistical significance of this shift. An exploration of the reasons behind diminished churning is important, but it lies beyond the scope of this paper.

Interestingly, the flow data examined here are *not* consistent with the widespread stories of downsizing and restructuring. It seems more likely that a significant episode of restructuring would be manifested in a significant rise in the flows of job losers relative to the size of the labor force. According to the gross flow data, however, the amount of labor market reallocation has fallen over the past 15 years. This does not disprove the existence of restructuring, but it is difficult to interpret this evidence as supporting the restructuring hypothesis.

Thus, in addition to changes in matching efficiency and the rate of labor force growth that coincide with the Beveridge shift, a significant decrease occurred in the amount of labor market churning.

## *IV. Quantifying the Effects of Shifts in Matching Efficiency and Allocative Intensity*

The previous sections demonstrated that potentially Beveridge-shifting changes in job matching and labor market flows appear to have occurred over the past 20 years. The remaining question is whether these changes were large enough to generate the observed

Figure 7

*Actual and Predicted Shifts in the Beveridge Curve*

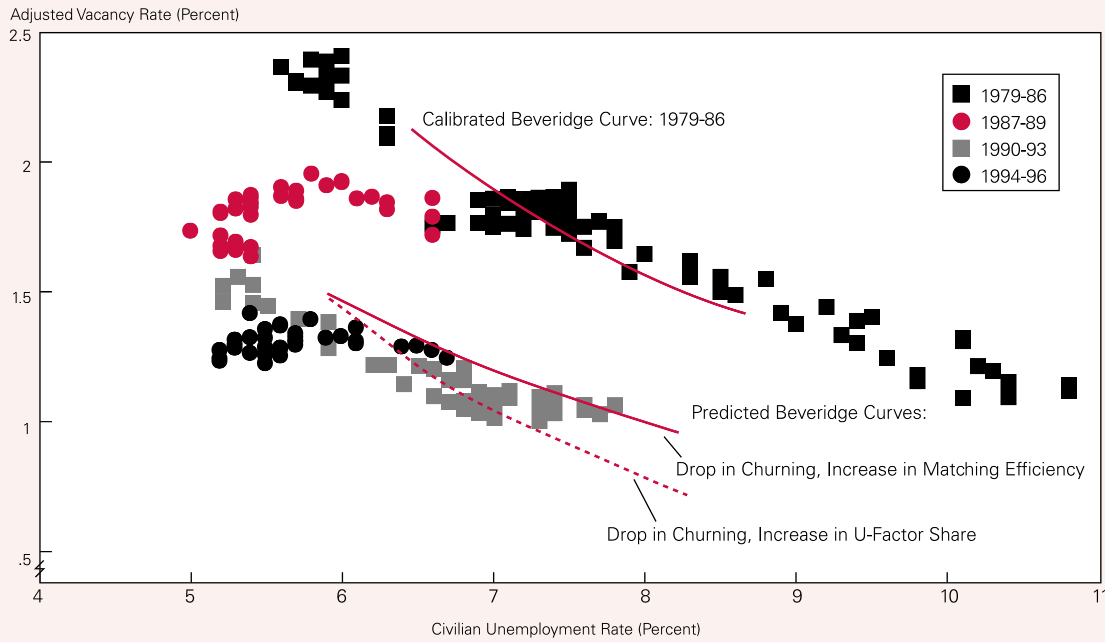


Source: See Figure 3.

shifts in the Beveridge curve. To answer this, the simple model represented in Figure 4 will be used to simulate the impact of the labor-market changes estimated above, and the magnitude and location of the Beveridge shifts implied by the model will be compared with those in Figure 3.

Note that, in the simple model, the general levels of unemployment and vacancies—the overall location of the Beveridge curve relative to the origin—depend on the parameters of the matching function, the rate at which workers quit employment, and the rates at which jobs are created and destroyed. An increase in matching efficiency will decrease the levels of both *U* and *V*; a decrease in the rate of job creation and destruction will also decrease the levels of both *U* and *V*. Because of the symmetry of these processes,[7] changes of either type generate *equal* shifts in the levels of *U* and *V*. The "curve" part of the Beveridge curve is generated by business-cycle variations in the job creation and destruction process, holding fixed the level of job churning and matching efficiency.

The model is calibrated to match the 1979–86 Beveridge curve, and then theoretical shifts in the Beveridge curve caused by changes in the job matching or churning processes are computed.[8] The model is then simulated at different points in the business cycle to trace out the entire curve.

Estimated changes in the matching and churning processes are found to shift the simulated Beveridge curve in a way that matches the observed shift. Figure 7 presents a scatter plot of unemployment and job vacancies along with this study's simulated Beveridge curves. As before, notice (1) the shifted-out Beveridge

[7] A new match causes a one-for-one drop in unemployment and vacancies. Moreover, as defined, an increase in job churning causes equivalent increases in U and V.

[8] The model is calibrated to the pre-1987 Beveridge curve in several ways. First, the inflow rates are selected to imply an equilibrium of roughly 7.5 percent and 1.7 percent for *U* and *V*, respectively. Second, the *U* factor share ($\alpha$) is adjusted downwards to match the slope of the full Beveridge curve. (The matching function is estimated on unemployment excluding layoffs. The slope of the Beveridge curve excluding layoffs matches that predicted using the estimated $\alpha$.) Third, the overall productivity factor, *A*, is set to $e^{-1.3}$, consistent with regression (3).

curve from 1979 to 1986, (2) the transition beginning in 1987, and (3) the shifted-in curve during the 1990s. The simulated curves are superimposed as red lines. The outermost line represents the position of the Beveridge curve during the 1979 to 1986 period. Below and to the left are simulated Beveridge curves resulting from changes in matching and churning. The solid line represents a Beveridge curve shifted by the estimated drop in churning and the increase in matching efficiency. The dashed line represents a shift caused by the churning decrease and the estimated increase in the factor share ($\beta$) for unemployment in matching function. These curves lie near the observed Beveridge curve in the 1990s. Thus, the model's predictions for the shift in the Beveridge curve match up nicely with the observed shift. Since the model is nonlinear, it is difficult to precisely parcel out the part of the shift due to matching and that due to churning. However, simulations suggest that the increase in matching efficiency accounts for slightly more than two-thirds of the shift.

How large is this inward shift? Comparing unemployment and vacancy combinations at roughly comparable points in the business cycle, we find that the unemployment and vacancy rates are about 0.6 percentage point lower in the early 1990s than they were in the early 1980s. Add to this the modest reduction in unemployment (and perhaps vacancies) of approximately 0.2 percentage point due to the slowdown in labor force growth, and essentially all of the inward shift of the Beveridge curve is accounted for.

## *V. Beveridge Curve Shifts and the Baby Boom Generation*

This section will consider whether the baby boom (and subsequent baby bust) affected the Beveridge curve. Labor force growth slowed in the mid 1980s as the baby boom generation completed its entrance into the labor force and the baby bust generation began to enter. At the same time, as the baby boom generation entered their thirties and forties, overall labor force attachment likely increased, and this would have decreased job churning. Moreover, any differences in the job-search behaviors of younger and older workers would have affected the aggregate matching function during this period; unfortunately, no data are currently available that would permit direct exploration of this influence.

It is possible to compute a rough estimate of the baby boom's overall impact on the unemployment rate. For example, Paul Flaim (1990) finds that the baby boom generation put upward pressure on unemployment during the 1960s and 1970s and that this pressure began to ebb in the 1980s. Figure 8 displays unemployment rates by age (the top panels) and the age distribution of the labor force (the bottom panels) for the beginning and ending years of this study's sample, 1979 and 1994.[9] Two features stand out: (1) Younger workers, especially those 25 years of age and under, show higher average rates of unemployment, consistent with a weaker attachment to employment. As the top two panels of Figure 8 show, this pattern is quite stable over the sample. (2) While the relative unemployment rates for age groups remained stable, the age distribution of the labor force did not. The bottom panels of Figure 8 show the movement of the baby boom cohort "hump" from ages 16 to 33 in 1979 to ages 31 to 48 in 1994.

Now consider the effect on unemployment of the aging baby boom cohort. In 1979, the baby boom generation was young and thus more prone to unemployment. In 1994, the same cohort was at an age of lower unemployment. Meanwhile, the baby bust had replaced the boom in the working-age population under 30 years old. Thus, a drop occurred in the size of the age group most prone to unemployment and an increase in an age group less likely to be unemployed, which precipitated a reduction of the overall unemployment rate (holding other factors fixed).

How much of the reduction in the unemployment rate over the period can be attributed to these demographic changes? To answer this, the overall unemployment rate was computed holding each age group's unemployment rate constant and varying only the age distribution of the labor force over time. The evolution of this unemployment rate, driven purely by demographic forces, is displayed in Figure 9.[10] Using 1979 as the base year for the unemployment/age relationship, demographic changes alone

[9] The data used in this section come from the Merged Outgoing Rotation Group files from the Current Population Survey. The top panels of Figure 8 display unemployment rates by age, $(N_U/LF)_{age}$. The bottom panels contain $(N_{age}/LF$, the age distribution of the labor force.

[10] The unemployment rate is defined as $U = N_U/LF$, the number of unemployed as a fraction of the labor force. We can compute the overall unemployment rate as the weighted sum of unemployment rates by age, where the weights are the fraction of the labor force in a given age cohort. That is,

$$U = \sum_{age} (N_U/LF)_{age} \times (N_{age}/LF).$$

We hold $(N_U/LF)_{age}$ constant (using a base year) and allow the age distribution $(N_{age}/LF)$ to vary across time.

Figure 8

*The Baby Boom and the Unemployment Rate*

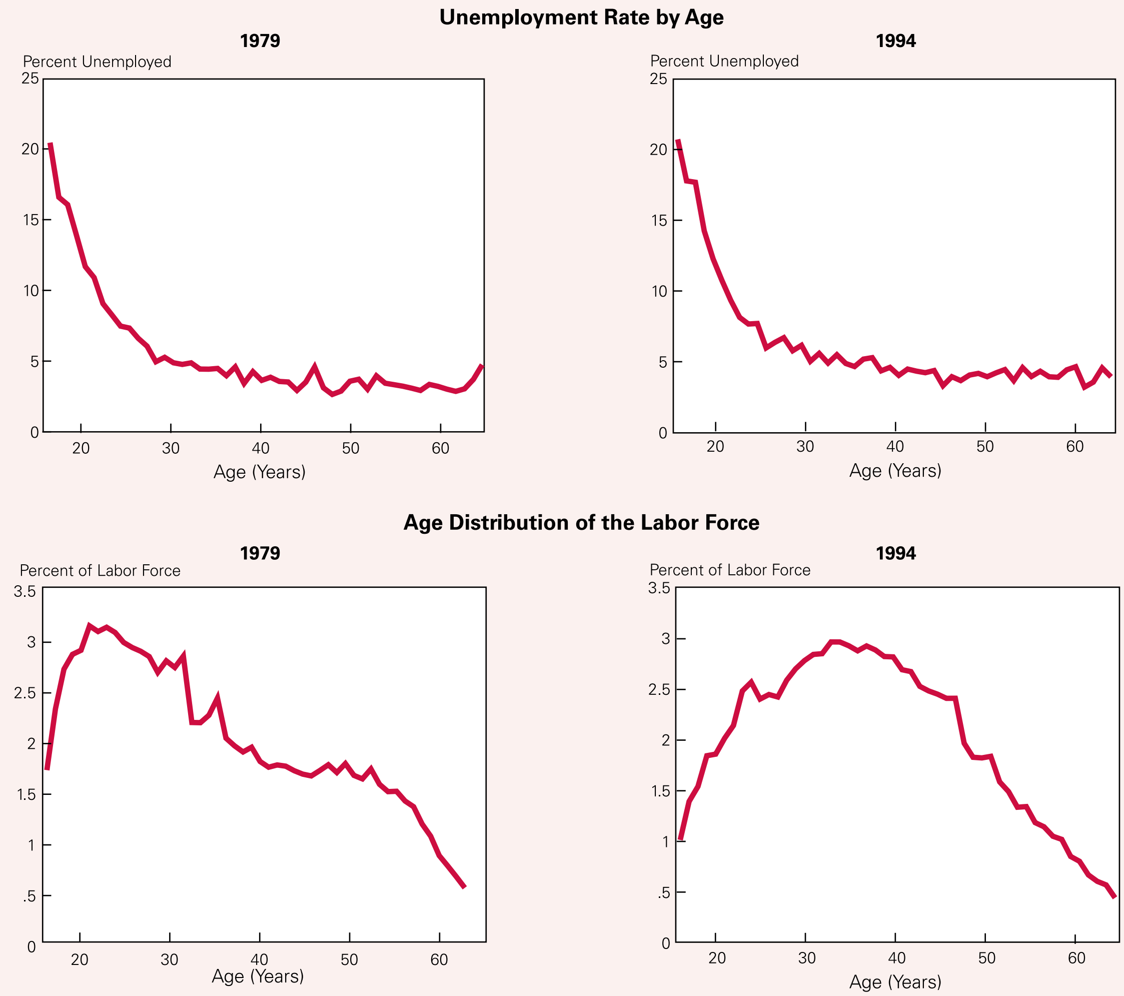


Source: U.S. Bureau of Labor Statistics, *Current Population Survey* and MORG files.

would have moved the unemployment rate from 6 percent in 1979 to 5.25 percent in 1995. Using instead 1995 as the base year, the unemployment rate moves from 6.3 percent to 5.55 percent in the same period. In either case, the overall effect from the aging of the baby boom generation was a drop in the unemployment rate of about 0.75 percentage point between 1979 and 1995. Of course, other pressures moved the unemployment rate up and down during that time period as well, as the figure shows. Nevertheless, according to these simple calculations, demographic changes imply that an unemployment rate of 6 percent in 1979 is comparable to a rate of about 5.25 percent in 1995.

To what extent can the shift in the Beveridge curve be tied to the baby boom? The drop in the unemployment rate implied by demographics is similar in size to the shift predicted by the simple framework developed above. Perhaps demographic change was the driving force behind the drop in job churning, labor force inflows, and increased matching efficiency. However, the timing does not quite work: The demo-

Figure 9

*The Evolution of the Unemployment Rate Due to Demographic Changes*

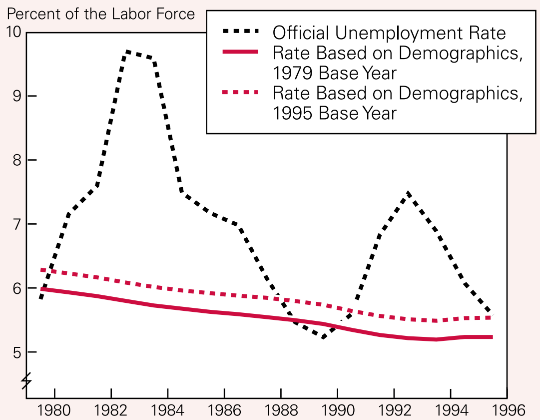


Source: U.S. Bureau of Labor Statistics, *Current Population Survey* and MORG files; authors' calculations.

graphic explanation suggests a gradual change from 1979 to 1995. Instead, a seemingly stable Beveridge relationship has been observed over most of the period, with a few transition years (approximately 1987 to 1989). Further research into the impact of demographic changes on each component of the Beveridge model is needed. Until then, it seems most reasonable to conclude that the baby boom was a contributing factor in the inward shift of the Beveridge curve, while a "true" increase in matching efficiency (that is, one that is independent of demographics) is likely responsible for the remainder.[11]

Finally, it should be noted that the *outward* shift of the Beveridge curve in the 1970s is plausibly linked to *increases* in the unemployment and vacancy rates associated with the entrance of the baby boom generation to the labor force. Whether this increase in unemployment explains most or all of the Beveridge curve shift lies beyond the scope of this paper, largely because of the lack of data for the 1960s and 1970s that would allow us to estimate changes in matching efficiency over this earlier period.

## VI. Conclusion

This study offers several explanations for the apparent shifts in the Beveridge curve during the 1980s and early 1990s. Following the work of Blanchard and Diamond (1989), the Beveridge curve is treated not as a fundamental economic relationship, but as a correlation brought about by an underlying process of labor market flows. In a simple model along these lines, inflows are related to the intensity of labor market churning and to gross labor force growth. Outflows are defined by a simple process of matching unemployed workers with vacant jobs. Shifts in these flow processes are estimated. Most notably, we find a significant increase in efficiency in the aggregate matching process over the period studied. Significant drops in labor force growth and the degree of labor market churning are also observed. These changes combine to produce a predicted shift in the Beveridge curve that closely matches the observed shift. It seems plausible that well-known demographic shifts—the baby boom generation and the increased labor force participation of women—can account for the recent shifts in the Beveridge curve. However, the precise timing and "shape" of these changes do not match the timing of the Beveridge shifts.

*We find a significant increase in efficiency in the aggregate matching process over the period studied. Significant drops in labor force growth and the degree of labor market churning are also observed.*

The "explanations" offered here for Beveridge curve shifts—changes in matching efficiency and the level of churning—are explanations only at a fairly abstract level. Why and how the unemployed may be making job matches more efficiently, and what caused the decreased level of churning are not yet known. These questions remain a subject for further research.

Finally, one glaring omission from the analysis must be noted: namely, any discussion of wages. The wage offered is presumably a key variable in the

[11] Here the effect on unemployment of demographic shifts has been examined. As above, it has been assumed (although no evidence of this has been presented) that vacancies will gradually adjust to important changes in labor force growth and employment attachment, so that the aging of the baby boom plausibly would have a proportionate downward influence on vacancies.

job-matching process at the individual level, as well as in the decision to enter or leave the labor force. Looked at the other way, the process by which job matches occur and the processes that determine the overall level of churning and job entry likely affect the overall level of wages in the economy. For example, Abraham and Medoff (1982) find that instability in the Phillips curve is associated with shifts in the Beveridge curve. The link from job market flows to wages is one that we plan to explore in future work.

## *References*


Abowd, John M. and Arnold Zellner. 1985. "Estimating Gross Labor Force Flows." *Journal of Business and Economic Statistics*, July, pp. 254–83.

Abraham, Katharine G. 1987. "Help-Wanted Advertising, Job Vacancies, and Unemployment." *Brookings Papers on Economic Activity*, no. 1 (June), pp. 207–48.

Abraham, Katharine G. and James Medoff. 1982. "Unemployment, Unsatisfied Demand for Labor, and Compensation Growth, 1956–80." *Brookings Papers on Economic Activity*, no. 1 (June), pp. 49–88.

Akerlof, George, Andrew Rose, and Janet Yellen. 1986. "Job Switching and Job Satisfaction in the U.S. Labor Market." *Brookings Papers on Economic Activity*, no. 2 (Nov.), pp. 495–592.

Blanchard, Olivier Jean and Peter Diamond. 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity*, no. 1 (June), pp. 85–155.

Davis, Steven J., John Haltiwanger, and Scott Schuh. 1996. *Job Creation and Destruction*. Cambridge, MA: The MIT Press.

Flaim, Paul O. 1990. "Population changes, the baby boom, and the unemployment rate." *Monthly Labor Review*, August, pp. 3–10.

Ritter, Joseph. 1993. "Measuring Labor Market Dynamics: Gross Flows of Workers and Jobs." Federal Reserve Bank of St. Louis *Economic Review*, November/December, pp. 39–57.

Tobin, James. 1972. "Inflation and Unemployment." *American Economic Association Papers and Proceedings*, vol. 62 (May), pp. 1–18.

U.S. Department of Labor, Bureau of Labor Statistics. *Current Population Survey*. Merged Outgoing Rotation Group Files. January 1979–December 1995.