OLIVIER JEAN BLANCHARD
Massachusetts Institute of Technology
PETER DIAMOND
Massachusetts Institute of Technology

The Cyclical Behavior of the Gross Flows of U.S. Workers

THE U.S. LABOR MARKET is characterized by high rates of job creation and job destruction, and by large flows of workers into and out of employment. In a previous paper, we developed a conceptual framework to interpret the dynamic behavior of both the levels, or stocks, of employment, unemployment, and vacancies as well as the flows into and out of these states.[1] At that time, we focused only on the behavior of stocks without examining the behavior of the flows themselves. In this paper, we intend to rectify that by refining and extending the pictures both of job creation and destruction, and of the flows of workers sketched in our earlier paper.

We rely on three data sets. Our primary data source is the Current Population Survey (CPS), which gives monthly gross flows of workers between employment, unemployment, and "not in the labor force." These data cover the period 1968–86 and are disaggregated by age and sex. Whenever available, we use the data as corrected by John Abowd and Arnold Zellner.[2] The other two sets cover manufacturing, and are

We thank the National Science Foundation for financial assistance, and the Bureau of Labor Statistics for data and advice. We thank Shyam Parekh, Hugh Courtney, and Jae Woo Lee for research assistance. We thank John Abowd, Bruno Contini, Steve Davis, John Haltiwanger, Jeff Miron, Kevin Murphy, Larry Summers, and Mark Watson for comments and advice.

1. Blanchard and Diamond (1989a).
2. Abowd and Zellner (1985). We discuss their adjustment methods at a later point in the paper.

collected from firms rather than workers. We use the manufacturing data to check and refine the picture we derive from the CPS flows. The first manufacturing data set, collected by the Bureau of Labor Statistics, records monthly gross flows into and out of manufacturing employment; it is disaggregated by the reason for the worker's move—quits versus layoffs, rehires versus new hires. Unfortunately, collection of these data was discontinued in 1981. The second set was put together by Steve Davis and John Haltiwanger from the Longitudinal Research Datafile, and gives quarterly net changes in employment by establishment; therefore, it allows for disaggregation by the size of the employment change, as well as by the type and size of the firm or the sector.[3] These data are available for the period 1972–85.[4] All three data sets correspond to different definitions of the flows and allow for different types of decomposition. Thus, our overall examination yields a sharp, rich picture of the labor market.

Within the context of the data sets, we focus on two main facts. Fact one concerns *jobs* and the process of job creation and destruction, and confirms the recent findings by Davis and Haltiwanger concerning manufacturing employment.[5] Suppose that cyclical fluctuations in the flows of job creation and destruction were symmetrical, with recessions, for example, leading to increases in job destruction that are equal to decreases in job creation. Assume also that quits, and thus replacements of quits, are procyclical—low in recessions and high in booms.[6] We would then expect cyclical fluctuations in the flow into employment to exceed those in the flow out of employment. That is, in a recession, the

3. Davis and Haltiwanger (1990).

4. With the passage of time, data series have become longer and more abundant. An early investigation, carried out by Perry (1972), had to infer the flow data by recording the stocks of unemployed at different intervals. Also, see two other early studies of the cyclical behavior of CPS flows by Smith, Vanski, and Holt (1974) and the Bureau of Labor Statistics (1977). Further analysis of their cyclical behavior was carried out by Clark and Summers (1978), who extended their earlier work on labor market transitions (1979).

We limit ourselves to U.S. data. A recent study by Burda and Wyplosz (1990) looks at many of the same issues as we do using data from the United States, United Kingdom, Germany, and France. They also refer to European research on gross flows.

5. Davis and Haltiwanger (1990). While "destruction" has become an accepted term, it has some wrong connotations. Destruction here means that the job is not filled again when the worker leaves: either the job disappears forever, or is filled when conditions change. As such, "termination" or "closure" may be more appropriate terms.

Also, Davis and Haltiwanger define "destruction" to include an employment decrease from worker unavailability after a quit. We adjust for this element.

6. This assumption reflects the true behavior of quits.

increase in job destruction would be partly offset by lower quits, and the decrease in job creation would be reinforced by lower replacements for those quits.

Interestingly, we find the opposite to be true for both the CPS and the manufacturing flow data. The amplitude of fluctuations in the flow out of employment is larger than that of the flow into employment. This, in turn, implies a much larger amplitude of the underlying fluctuations in job destruction than of job creation. Reduced employment in recessions results more from high rates of job destruction than low rates of job creation. Similarly, booms are times of low job destruction rather than high job creation. The Davis-Haltiwanger data, which come closest to directly measuring job creation and destruction, confirm these ideas. This finding contrasts with many characterizations of cyclical fluctuations. It rules out a Schumpeterian view of cyclical fluctuations, with booms as times when inventions are implemented yielding high job creation. It also rules out a view of the cycle in which movements in aggregate demand lead to symmetric effects on the rate of job creation and destruction.[7]

The second fact concerns *workers,* and is based on an examination of the CPS flows between employment, unemployment, and "not in the labor force"—or *E, U,* and *N* in what follows. It is well known that in the United States only half of the average flow into *E* comes from *U*, with the other half coming from people classified as not in the labor force. It is also well known, at least since the work of Kim Clark and Lawrence Summers, that the distinction between these two pools is fuzzy, with many workers going back and forth between the two states.[8] Therefore, *U* and *N* are perceived to be very similar states. However, we find sharp differences between the cyclical behavior of the flows between *E* and *U* on the one hand and the flows between *E* and *N* on the other. In particular, we find that the flow from *E* to *U,* which we also call the *EU* flow, increases in a recession while the flow from *E* to *N,* or the *EN* flow, decreases. We also find that the *UE* flow increases in a recession, while the *NE* flow decreases. We refine this aggregate picture by examining flows by sex and age, and find clear cyclical differences

7. In a later section of the paper, we advance several tentative explanations for this result, all based on the idea that recessions are periods of cleaning up, leading to additional job destruction. However, detailed pursuit of an explanation would require looking at different data from those in our paper. Thus, we leave that to future research.

8. Clark and Summers (1979).

among different age groups—in particular between young, mature, and older workers, and between males and females.

This finding leads us to develop a model of the labor market that allows for two types of workers, "primary" and "secondary." These two types are represented in different proportions in the various age and sex groups. They may differ in behavior in several ways. Secondary workers may quit more. They may search less when not employed. Firms, when they have the choice, may lay off secondary workers first, or rank primary over secondary workers when hiring. We show that a simple model—with similar layoff rates but different quit rates, with primary workers going into *U* and secondary workers going into *N*, and with the ranking of primary workers above secondary workers in hiring—can explain the basic characteristics of the flows of workers between *E* and both *U* and *N*. We conclude by discussing extensions and implications of our model. In particular, we speculate about its implications for wages but stop short of pursuing the matter empirically.[9]

We divide the paper into five parts. The first briefly describes the CPS data and the Abowd-Zellner adjustment, and gives the adjusted mean flows between *E, U,* and *N*. The second describes the methodology used throughout the paper to characterize the "cyclical behavior" of the series. Put simply, we characterize the behavior of gross flows when the economy moves along the Beveridge curve. The third section focuses on the flows between employment and "nonemployment," which is the sum of *U* and *N*. This section relates the results obtained from the CPS data to those obtained using manufacturing data and to the Davis-Haltiwanger results. The fourth section examines the flows between *E* and each of the two nonemployment states, *U* and *N*; we look at both aggregate flows and flows disaggregated by age and sex. Finally, we propose a conceptual framework based on the distinction between primary and secondary workers with different labor market behavior.

The Gross Flows: Unadjusted and Adjusted Data

Since 1949, gross flows of workers, by age, sex, and race, have been tabulated monthly from the CPS. Except for brief episodes, however,

9. We base our speculation on recent theoretical work; see Blanchard and Diamond (1990a, 1990b).

those tabulations have not been published because the Bureau of Labor Statistics (BLS) perceives them to be of poor quality. Therefore, we begin by discussing their shortcomings and how they may affect our study.

Unadjusted Gross Flows

BLS has identified two problems with the CPS data: "missing observations" and "classification error."[10] For the period that concerns us, 1968–86, the average size of the CPS interview group is approximately 50,000. Because the survey is conducted using rotating panels of interviewees, only 75 percent are present in any two consecutive months and thus only that fraction can be used to compute gross flows. And, of that fraction, 7.5 percent of those interviewed in one month are not located in the next, and another 7.5 percent cannot be located in the previous month's group of interviewees. Thus, the problem of "missing observations." If the missing observations were random, then one could just look at those individuals for whom observations were available for two consecutive months. However, the evidence suggests otherwise. The other problem of "classification error" arises because the wrong answer is recorded for some individuals and they are improperly classified.

Both problems bias the measured flows and generate additional measurement noise (that is, beyond conventional sampling error). In attenuated form, the problems also affect the stocks published by BLS, which rely on the whole CPS sample. We present an example that simply focuses on the effects of classification error and shows the effects at work.[11]

Suppose that individuals can be in only two states, employment and nonemployment; in other words, no distinction is made between unemployment and not in the labor force. Suppose that the sets of people in each state are constant and equal in number to E and M respectively. Thus, all measured transitions are spurious. Assume that the probability that an employed worker says that he is not employed is χ, and that the

10. For a more detailed description of the CPS survey, see Bureau of Labor Statistics (1982). For a discussion of the problems associated with the reported gross flows, see Abowd and Zellner (1985) and Poterba and Summers (1986).

11. Abowd and Zellner (1985) and Poterba and Summers (1986) address the problems with the CPS data in a more formal manner.

probability that a nonemployed worker says that he is employed is θ. Assume that misclassifications are independent across individuals and time. Let the measured stocks be denoted by E^* and M^*; let the measured flows be denoted by EM, for the flow from E to M, and ME, for the flow going in the other direction. Then, ignoring classification error in the previous month, their expected values will be given by:

$$\mathcal{E}(EM) = \chi E, \text{ and } \sigma^2(EM) = \chi(1 - \chi)E,$$

$$\mathcal{E}(ME) = \theta M, \text{ and } \sigma^2(ME) = \theta(1 - \theta)M,$$

$$\mathcal{E}(E^*) = E + (-\chi E + \theta M), \text{ and } \sigma^2(E^*) = \theta(1 - \theta)M + \chi(1 - \chi)E,$$

$$\mathcal{E}(M^*) = M + (\chi E - \theta M), \text{ and } \sigma^2(M^*) = \theta(1 - \theta)M + \chi(1 - \chi)E.$$

The spurious transitions will lead to upward biases in the two measured flows, but the biases will partly offset each other in the measured stocks. In addition, classification error leads to noise in the measured series; the noise is relatively larger for the flows than for the stocks. Thus, nonrandom missing data and classification error lead to both bias and additional noise in unadjusted gross flows. If either the pattern of missing data or of classification error varies cyclically, the bias also has a cyclical component.

Adjusted Gross Flows

Several studies have proposed corrections for the missing observations and classification error problems.[12] Abowd and Zellner have constructed adjusted series, both aggregated and disaggregated by sex, that we shall use whenever available. Abowd and Zellner make two sets of corrections. They allocate missing data to the unadjusted gross flows using a fixed allocation pattern; as a result, the time series behavior of the implied stocks—E, U, and N—fits the time series of the actual stocks as closely as possible. They then use reinterview survey information to correct for classification error. BLS reinterviews approximately 3,000 households each month. This information yields estimates of the equivalents of χ and θ, which can be used to estimate the bias and correct the original flows.

12. See Abowd and Zellner (1985) and Poterba and Summers (1986).

Abowd and Zellner's adjustments allow for—and thus correct for—cyclical movements in classification error, but not in missing data. More specifically, they estimate classification errors for each quarter by averaging the results of that quarter's reinterview surveys. This very flexible adjustment should remove the effects of cyclical classification error. However, they assume a constant allocation pattern of missing data over the sample. They test this assumption by dividing their sample into two subsamples. While they find the set of parameters to be significantly different across subsamples, they conclude that the difference is of little importance and that most coefficients are quite similar.

The adjusted gross flow series, like the unadjusted ones, exhibit very strong seasonal movements. Because we found that the number of working days has an important effect, and thus is evidence of changing seasonality, we have adjusted all series by the Census Bureau's seasonal adjustment program, X11. We did so with some reluctance.[13]

Basic Characteristics of the Adjusted Gross Flows

In what follows, we use the gross flow series adjusted by Abowd and Zellner and adjusted for seasonality. We denote the three stocks as E, U, and N and the six flows as EU, EN, UE, UN, NE, and NU.

Figure 1 gives the average values of the gross flows and stocks from January 1968 to May 1986. The stock numbers are from the full CPS sample. We give two numbers for N: first, the total number of people not in the labor force; second, and in brackets, the number of people classified as not in the labor force but who "want a job." This pool, which roughly equals the unemployment pool, probably includes many of the people, new entrants excepted, who go into and out of N. The numbers in parentheses are the original unadjusted gross flow numbers from the CPS.

The figure shows that the flows between E and N are as large as the flows between E and U.[14] The fact that U and the "want-a-job" pool in N are roughly equal implies similar hazard rates from either U or N, to E. Hazard rates are the average individual probabilities of moving from

13. Our reluctance stems from the fact that the statistical properties of the program X11 are not well understood.

14. These flow numbers do not include people who join or leave the civilian noninstitutional population during a given month. Those flows average 350,000 for the flow into the population and 200,000 for the flow out, and are small compared to the flows in the figure.

Figure 1. Average Values of Gross Stocks and Flows for Employment, Unemployment, and Not in the Labor Force, January 1968–May 1986[a]

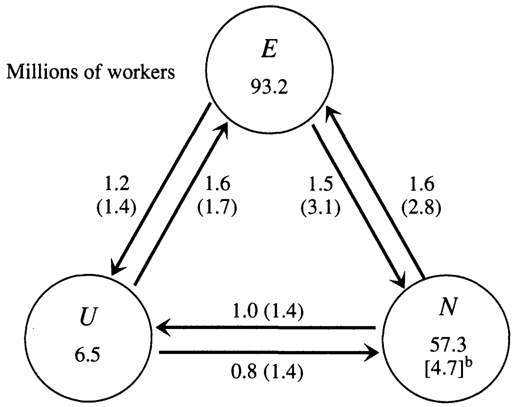

Source: Stock numbers are from the Current Population Survey (CPS). For the flow data, we use the Abowd-Zellner adjusted gross flow series. The original unadjusted numbers from the CPS appear in parentheses. All numbers are in millions.

a. The variables E, U, and N represent employment, unemployment, and not in the labor force respectively.

b. The bracketed stock figure for N equals the number of people who "want a job."

one pool to another. The figure also shows how the Abowd-Zellner adjustment decreases the flows into and out of N, while leaving the mean flows between E and U largely unaffected.

Table 1 documents the importance of the seasonal component for the gross flows. For the table, we regress each gross flow on twelve monthly dummies and compute the standard deviations of the fitted values and of the residuals.[15] The table gives the largest and smallest seasonal coefficients, together with their respective months. Flows are defined by two months of data; we identify a flow by its second month. The seasonal component is most important for the flows between employment and not in the labor force. Seasonality accounts for 85 percent of the variance in the NE flow and 79 percent of the variance in the EN flow. June shows the largest flow from N to E, and September shows the largest flow from E to N, pointing to the importance of the school calendar. The same set of regressions applied to disaggregated age groups indicates that 16–19 year olds account for about three-quarters of the seasonal movements between E and N.

15. This is not the method we use to seasonally adjust the series for figure 1. This is only a convenient descriptive device.

Table 1. Seasonal Component of Adjusted Gross Flows of Workers, January 1968–May 1986

Thousands of workers

Flows	*Monthly gross flows*		*Standard deviations*[a]		
	Largest	*Smallest*	σ	σ_s	σ_{ns}
EU	1,777 (Jan)	1,058 (Mar)	452	187	422
EN	3,488 (Sep)	982 (Mar)	698	647	269
EU + EN	4,798 (Sep)	2,040 (Mar)	887	738	504
UE	1,945 (Jul)	1,123 (Jan)	429	229	372
NE	2,691 (Jun)	885 (Dec)	528	459	269
UE + NE	4,301 (Jun)	2,099 (Dec)	758	598	477
NU	1,854 (Jun)	675 (Dec)	421	270	331
UN	1,077 (Aug)	672 (Jun)	289	101	278

Sources: Abowd and Zellner (1985) and authors' own calculations using the Abowd-Zellner adjusted gross flows.
a. The standard deviation σ is the standard deviation of the gross flow; σ_s is the pseudo standard deviation of the fitted values in a regression of the gross flow on twelve monthly dummies; σ_{ns} is the standard deviation of the residual of the regression.

Finally, figure 2 gives three different series for the gross flow from unemployment to employment; the other flows have similar features. The three series are the Abowd-Zellner adjusted series, before and after seasonal adjustment, and a filtered series, which is obtained as a centered seven-month moving average of the seasonally adjusted series. The shaded areas correspond to recessionary periods as identified by the National Bureau of Economic Research (NBER). The figure makes two points. First, the seasonal component is indeed important. Second, the seasonally adjusted gross flow series still exhibits high frequency movements, or, put another way, appears to have a large noise component atypical of standard economic time series. The third series shows that these high frequency movements are substantially smoothed by even a short filter, and that a clear cyclical pattern then emerges.

This raises the question of what causes the noise component. How much of the noise is due to randomness in the data, to measurement error, to inadequate seasonal adjustment, and to spontaneous, short-lived bursts of mobility? We cannot quite tell, but several pieces of evidence may be relevant here. Thus, before discussing the noise component's potential sources, it is useful to know the results derived from regressing the flows on current and lagged values of the stocks; the results are reported in table 2. On average, the ratio of the standard deviation of the residuals to the mean is 8.1 percent. Moreover, the

Figure 2. Flows out of Unemployment into Employment, January 1968–May 1986

Millions of workers

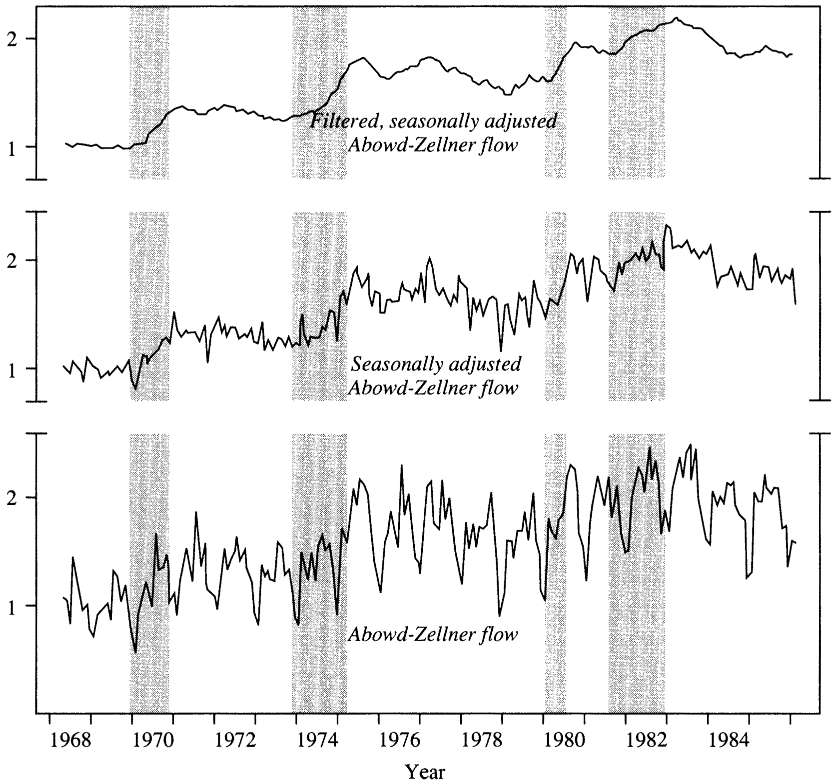

Sources: Abowd and Zellner (1985). We use the Census Bureau's seasonal adjustment program, X11. We filter the series using a centered seven-month moving average. Shaded areas represent recessions.

correlation matrix of residuals is nonnegative, with larger correlations between the two flows connecting pairs of stocks. Furthermore, the two flows that have the largest seasonal components, those between E and N, also display the largest residuals.

While the overall CPS sample is substantial, the bulk of the population remains in E or N from month to month. In other words, a transition is a relatively uncommon event. Thus, the sample size for measuring transitions is small. In a typical month, only 7.5 percent of the usable sample shows any transition at all (see figure 1). Therefore, the expected number of transitions averaged over the six flows is about 750 people, with a standard deviation around 3.7 percent. With this sample size, the monthly variation from randomness is substantial, nearly half the mean size of the residuals, 8.1 percent, that we found when regressing the flows on the stocks.

It is worth considering other potential sources of noise. One possibility is that the Abowd-Zellner adjustment, while it removes the bias, does not remove the noise created by the two measurement problems. If classification errors are uncorrelated across individuals, then we can use the simple example described earlier to approximate the contribution made by the Abowd-Zellner adjustment. With probabilities of 0.2 percent that an employed worker is recorded as unemployed and 1.0 percent that an employed worker is recorded as not in the labor force, the standard deviations of the noise in the *EU* and *EN* flows due to classification error are 1.7 and 3.8 percent of the means respectively.[16] Correction of these errors in the following month would generate spurious return flows, contributing to the apparent sawtooth nature of the noise.

In thinking about classification error, one must decide whether the important errors come from misclassification of flows (which would cause negative correlation among residuals) or spurious generation of flows (which would cause a positive correlation). Positive correlation occurs when the error is corrected the following month, and there is positive serial correlation in the aggregate likelihood of error, for example, over the cycle. It may be useful to remember that a single interviewee reports for the entire household. The presence of positive correlation suggests that the latter source of error may be more important.

The fact that the noise component appears larger for the flows between *E* and *N* than for other flows suggests that inadequate seasonal adjustment also contributes to measurement error. It is clear that the seasonal adjustment program, X11, neither removes white noise seasonal components nor allows for interactions between the business cycle and seasonality. Both factors could be present. As we shall see, the flows from manufacturing data, which are obtained from firms, are much smoother, underlining the importance of measurement error in the CPS flows.

16. The probability coefficients are averages over 1977–82, obtained from Abowd and Zellner (1985, table 6).

The Cyclical Behavior of the Gross Flows

In this section, we develop a methodology to characterize the cyclical behavior of the gross flows. While the details are somewhat involved, the basic idea is a simple one: we identify cyclical disturbances as those that move the economy along the Beveridge curve, or more accurately, a Beveridge loop, and trace the effects of such disturbances on the flows.

One may question why we do not use an even simpler approach: for example, regressing the flows on unemployment and tracing the effects of unemployment on the flows.[17] We have two reasons—one theoretical, and the other empirical—for not following this course. First, we believe that movements in unemployment also come from disturbances other than those to aggregate activity, even if the aggregate activity shocks dominate at business cycle frequencies. Second, even given unemployment, other variables such as employment and vacancies help predict the flows.

Basic Methodology

In our previous Brookings paper, we assumed that the behavior of three stock variables—employment, E, unemployment, U, and vacancies, V—was the result of the dynamic effects of three underlying sources of shocks. We called these sources aggregate activity, reallocation intensity, and labor force disturbances.

We let $\mathbf{X} \equiv [E\ U\ V]'$ be the vector composed of the three stock variables, and $\boldsymbol{\epsilon} \equiv [\epsilon_c\ \epsilon_s\ \epsilon_f]'$ be the vector of the three underlying white noise innovations to aggregate activity, reallocation intensity, and labor force disturbances respectively. Thus, we assumed:

$$\mathbf{X} = \mathbf{A}(\mathbf{L})\boldsymbol{\epsilon}, \tag{1}$$

where $\mathbf{A}(\mathbf{L})$ was an infinite lag polynomial. In the present context, $\mathbf{A}(\mathbf{L})$ could be considered the convolution of two lag polynomials, one giving the dynamic effects of innovations on the underlying disturbances, and the other giving the dynamic effects of disturbances on the stock variables. To identify the effects of the disturbances on the stock

17. Perry (1972) used this approach.

variables, we made assumptions about the contemporaneous effects of those innovations as summarized in the matrix $\mathbf{A}_0$. We return to this issue later.

For our present purpose, we maintain the assumptions of our earlier paper. We also assume that the same three underlying disturbances determine the dynamics of the flows, but that these flows have an additional noise component. Think for the moment of this noise as measurement error present in the flows and not present in the stocks. We return to this assumption below. More formally, let $\mathbf{F}^*$ be the nonobservable vector of the flows free of noise, with $\mathbf{F}^* \equiv [EU^*\ UE^*\ EN^*\ NE^*\ UN^*\ NU^*]$, and let $\mathbf{F}$ be the vector of the corresponding observable flows. Let $\boldsymbol{\phi}$ be a 6×1 vector of noise. We assume:

$$\mathbf{F} = \mathbf{F}^* + \boldsymbol{\phi}, \tag{2}$$

$$\mathbf{F}^* = \mathbf{B(L)}\boldsymbol{\epsilon}, \tag{3}$$

where $\mathbf{B(L)}$ is again an infinite lag polynomial.

Clearly, stock-flow identities imply relations between $\mathbf{F}^*$ and the first two components of $\mathbf{X}$, and thus between $\mathbf{A(L)}$ and $\mathbf{B(L)}$. However, we shall not impose those restrictions in what follows. Finally, we assume that the disturbances, $\boldsymbol{\epsilon}$, and the noise, $\boldsymbol{\phi}$, have zero cross correlation at all leads and lags.

Under these assumptions, we can do two things. We can clean the flows of their noise and we can characterize the cyclical behavior of the flows, defined as the dynamic effects of ϵ_c on the flows. Equations 1–3 imply the existence of a relation between stocks and flows:

$$\mathbf{F} = \mathbf{C(L)X} + \boldsymbol{\phi}, \text{ where } \mathbf{C(L)} = \mathbf{B(L)}[\mathbf{A(L)}]^{-1}. \tag{4}$$

Under the assumption that $\boldsymbol{\epsilon}$ and $\boldsymbol{\phi}$ are uncorrelated, this relation can be estimated by ordinary least squares, and the fitted values from the regressions give us an estimate of $\mathbf{F}^*$. Note that this relation does not depend on the particular identification assumptions for the different shocks.

Furthermore, given specific identification assumptions, we can trace the effects of a shock to ϵ_c on U, E, and V using equation 1. Then, using the auxiliary regressions of equation 4 that characterize the dynamic relation between stocks and flows, we can trace the effects of ϵ_c on the flows.

Assumptions about Disturbances and Noise

There are several assumptions, implicit and explicit in our methodology, that we would like to discuss before proceeding with the estimation. The first assumption is that the stocks and flows are determined by three, and only three, underlying disturbances (plus noise in the case of flows). While there is no reason to believe that three disturbances, rather than two or four perhaps, dominate the dynamics of stocks and flows, we did find in our previous paper that such an assumption allowed for a plausible interpretation of the behavior of the stocks. Moreover, we can indirectly test this assumption. For example, if there were only one or two dominant disturbance sources, only one or two of the three stock variables would be significant. If, however, it is true that no more than three disturbances dominate the dynamics of all variables in the labor market, then variables other than those already included in **X** should be redundant in equation 4 and their coefficients should be insignificant. As summarized in table 2, there is little evidence to support the significance of additional variables.

The second assumption is that the noise in the flows is not present in the stocks. From our construction of stocks and flows, we know that this cannot be literally true. If we think the noise in the flows comes from sampling error, missing data, classification error, or unadjusted seasonality, then this noise will remain in the stocks, although probably to a much smaller extent. Thus, the assumption that no noise exists in the stocks is an approximation. We can also test this idea: if it is approximately correct, the high frequency noise present in **F** should be absent from the fitted values in equation 4. While the fitted values we obtain from the test still show some high frequency movement, there is much less than in the original series.

The third assumption is that the underlying disturbances and the noise are uncorrelated. This idea relates to our earlier discussion of whether the pattern of missing data and classification error has a cyclical component, and if so whether the Abowd-Zellner adjustment removes it. We concluded that the Abowd-Zellner adjustment corrects for cyclical movements with classification error, but does not adjust for missing data.

There is a general, but again informal, test of all three assumptions taken together. Under these assumptions, the impulse responses for the stocks that come directly from estimation using the stocks should be close to those that come from estimation using the flows and the accumulation identities.

Projections of the Flows on the Stocks

Table 2 gives the basic results of our projections of the flows on the stocks. For the period January 1968 to May 1986, we regress each flow on a constant, a time trend, and the current and four lagged values of *E*, *U*, and *V*. We choose lag length based on likelihood ratio tests; the results that follow are unaffected by the choice of longer lag lengths. Own lagged flows are insignificant when entered into the equation. The availability of Abowd-Zellner corrected gross flows determined the sample period. Civilian employment, *E*, and unemployment, *U*, come from the CPS survey. The variable *V* is the vacancy series constructed in our previous paper.

While these regressions have no structural interpretation, the table presents some interesting results. First, for regressions without a lagged dependent variable, the R^2's are high (around 0.9) for all the flows except for those between *E* and *N* (which are around 0.6). Second, Durbin-Watson statistics are mostly consistent with a white noise residual, although the Q-statistics indicate some remaining autocorrelation. Third, all three stock variables are highly significant in most equations, suggesting the presence of at least three underlying disturbances. As discussed earlier, the correlation among the residuals can be used to assess the source of the apparent white noise component. All correlations are small and positive or are zero, with a mean of 0.12; they are slightly larger between flows between the same two states. For example, the correlation between the residuals of the *EU* and *UE* flows is 0.31.

The bottom panel tests for the significance of other variables in the flow equations. The variables are tested one at a time and fall into three groups. In the first group we include industrial production, manufacturing employment, and capacity utilization because one might expect the manufacturing sector to behave differently from the economy as a whole and to help predict the flows given aggregate variables. The second group includes insured unemployment, the number of people unemployed 27 weeks or more, and the number of people who have become unemployed

Table 2. Regression of Labor Force Flows on Labor Force Stocks, January 1968–May 1986[a]

		Summary statistics		*Significance tests*[c]				
Flows	*Time*[b]	R^2	*DW*	*SF(Q)*	*SF(E)*	*SF(U)*	*SF(V)*	σ/μ[d]
EU	0.3	0.95	1.53	**	...	**	**	6.7
EN	6.6	0.56	1.20	**	**	**	...	9.2
UE	−10.0**	0.92	1.70	**	**	**	**	5.8
NE	4.1	0.55	1.73	**	**	**	*	9.3
UN	1.8	0.90	1.40	**	...	**	...	9.0
NU	2.9	0.91	2.02	*	**	**	...	8.6

Sources: Unemployment and employment stocks are from the Current Population Survey; *V* is the vacancy series constructed in Blanchard and Diamond (1989a); the flow data are from the Abowd-Zellner adjusted gross flow series. All variables are seasonally adjusted. One asterisk represents significance at the 5 percent level. Two asterisks represent significance at the 1 percent level.

a. We use ordinary least squares regression of each flow (measured in thousands) on a constant, time, and current and four lagged values of *E*, *U*, and *V*.

b. Time equals coefficient on the time variable.

c. We use SF(Q) to denote the significance level of the $Q(42)$ statistic that the residuals are not white noise. We use SF(X), where $X = E, U, V$, to denote the significance level of the test that the set of coefficients on current and lagged X equals zero.

d. Here we divide the standard deviation by the mean, shown in percentages.

	Significance tests for other variables[e]						
	IP	*EM*	*CAPU*	*U27+*	*Ulayoff*	*Uinsur*	*Wantjob*
EU	...	*	*	**	**	**	...
EN	...	...	*	**	...	**	...
UE	...	...	...	**	...	*	...
NE	...	...	...	...	...	...	...
UN	*	*	*	...	...	**	*
NU	...	...	...	...	...	...	*

	Correlation of residuals					
	EU	*EN*	*UE*	*NE*	*UN*	*NU*
EU						
EN	0.12					
UE	0.31	0.05				
NE	0.11	0.36	0.12			
UN	0.11	0.22	0.00	0.04		
NU	0.00	0.18	0.12	0.01	0.17	...

e. Variables are described in text.

through the loss of their job. This group may detect whether compositional effects are important and could help predict the flows. The third group includes only one variable, the number of people classified as not in the labor force but who want a job. This variable may serve as a proxy for the group of people who move between *N*, and either *U* or *E*.[18] While most of these variables affect some of the flows, only insured and long-term unemployment have significance levels comparable with those of vacancies. Moreover, the significance of these two variables disappears when longer lags are used for unemployment, which suggests that they capture a longer distributed lag effect of unemployment than allowed for

18. The total civilian noninstitutional population is a very smooth series. Thus, when *N* is included, the number of people in the civilian population, but not in the labor force, leads to a high degree of collinearity between time, *E*, *U*, and *N*.

in the initial specification. Overall, these regressions provide little evidence against our assumption that three disturbances dominate labor market dynamics.

Identifying the Cyclical Behavior of the Gross Flows

In our previous paper, we identified aggregate activity disturbances as those that moved unemployment and vacancies in opposite directions—that is, along a Beveridge curve or loop—for some time after the initial shock. Under that identification assumption, we found that aggregate activity disturbances generated large thin loops in the Beveridge space and accounted for most of the movements in unemployment at high and medium frequencies. Thus, we have characterized the cyclical behavior of the gross flows as that part of the movement in gross flows accounted for by the aggregate activity disturbance. It is in this sense that we here characterize cyclical movements as the movements in unemployment along the Beveridge curve.

More precisely, we first estimate a vector autoregressive (VAR) system in E, U, and V. We use the same system as in our previous paper, except for the sample period, for which we use January 1968–May 1986 in order to preserve the symmetry with the flows, and for the fact that the variables are specified in levels rather than logs. Using the identification restrictions above, we transform the VAR in the form of equation 1 and trace the effects of a single one-standard-deviation shock in the aggregate activity variable on E, U, and V. Finally, using the set of regressions in equation 4, we trace the effects of the shock on the six flows among E, U, and N.[19]

In figure 3, we plot the movement in U and V implied by a one-standard-deviation negative shock to aggregate activity. We find it easier—for reasons that are not clear to us—to think of recessions rather than expansions; thus the choice of a negative shock. In the following sections, we examine the movement in flows associated with the movement in U and V. In figure 4, we plot the part of unemployment that is the result of aggregate activity shocks against total unemployment for the entire period. The figure shows the close relation between the two at business cycle frequencies.

19. A technical discussion is presented in Appendix A.

Figure 3. Response of Unemployment and Vacancies to a Negative Shock[a]

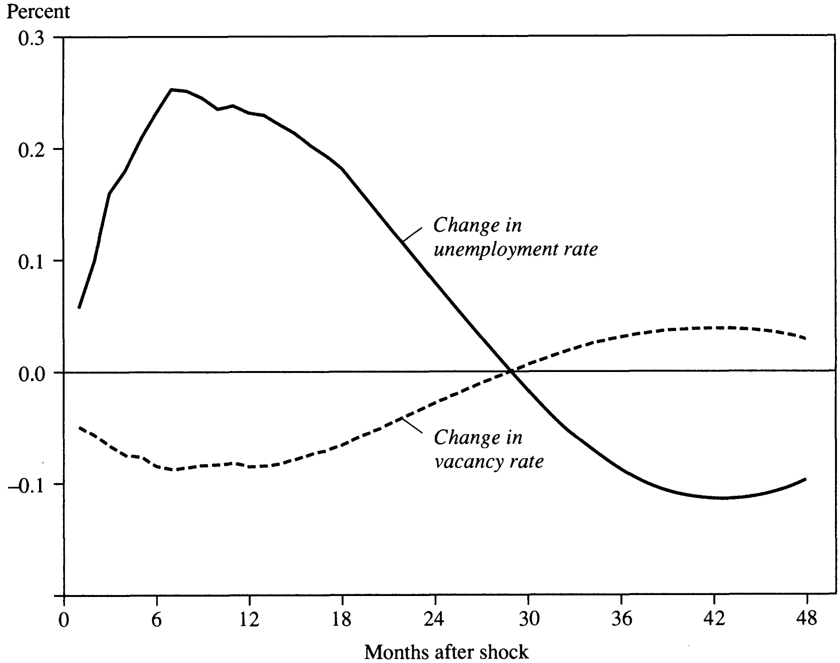

Source: Authors' own calculations.
a. The shock is defined as a one-standard-deviation negative innovation in the aggregate activity disturbance.

Flows into and out of Employment: Job Creation and Job Destruction

In this section, we concentrate on the flows between employment, E, and "nonemployment," M, defined as the sum of those unemployed and those not in the labor force; that is, $U + N = M$. We start with results from the CPS data and show that cyclical movements in employment are associated with larger variations in the flow out of employment than the flow into it. We then compare those results to the ones we obtained using the two manufacturing data sets. We conclude that there is strong evidence that cyclical fluctuations are associated with much larger movements in job destruction than job creation and discuss a number of potential explanations.

Figure 4. Cyclical versus Total Unemployment, January 1968–May 1986

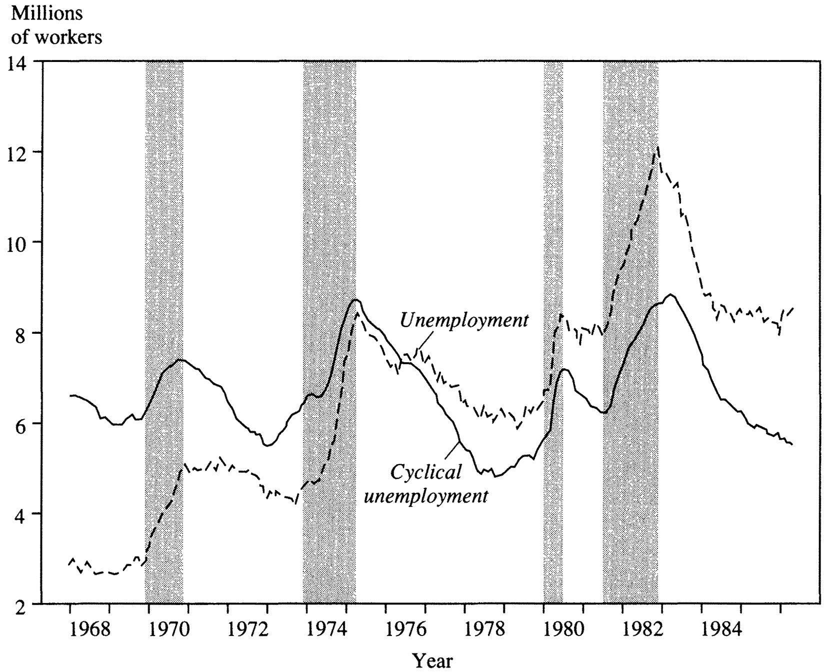

Source: Authors' own calculations. Shaded areas represent recessions.

Flows into and out of Employment from the CPS

The results displayed in figure 5 give the effects of a one-standard-deviation negative innovation in the aggregate activity disturbance on employment and the flows into and out of employment. All numbers are in thousands of workers and refer to differences with the no-innovation case. The numbers alongside the arrows are the cumulated changes in the flows from month one (when the shock takes place) to month *n*. The numbers in parentheses are standard deviations obtained by Monte-Carlo simulation. The numbers in the circles are changes in the stocks: the first is the number implied by cumulation of the estimated response of the flows, the second (in brackets) is the response of employment obtained directly from estimation of equation 1. Comparison of the two numbers serves as a rough check on the reliability of the approach.

Figure 5. Response of Gross Flows into and out of Employment to an Aggregate Activity Shock, over Selected Intervals[a]

Thousands of workers

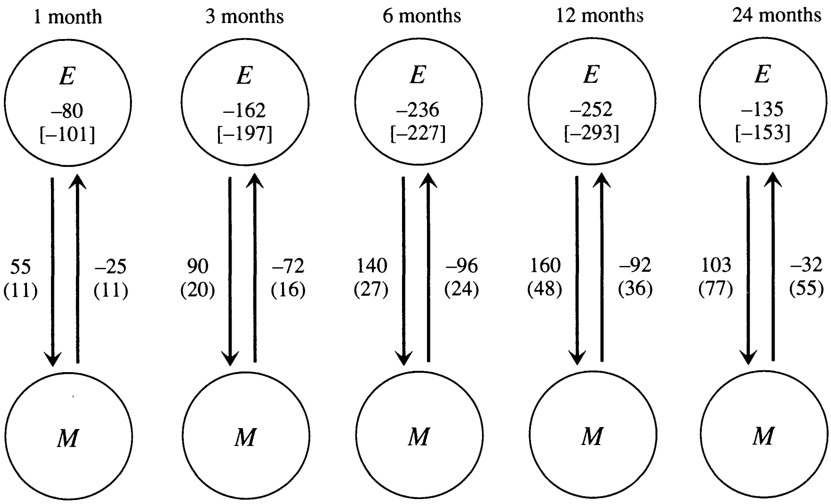

Source: Authors' own calculations using stock numbers from the Current Population Survey, the Abowd-Zellner adjusted gross flow series, and the vacancy series in Blanchard and Diamond (1989a). The variables E and M represent employment and nonemployment respectively.

a. The figure shows the effects on employment stocks and flows of a one-standard-deviation negative innovation in the aggregate activity disturbance. Numbers refer to differences with the no-innovation case: numbers alongside the arrows are the cumulated flow changes from the month of the shock to month n; numbers in circles are the stock changes, and numbers in brackets are stock responses obtained by direct estimation of equation 1. Standard deviations appear in parentheses.

A negative innovation, or a recessionary shock, leads to a decrease in employment, which peaks twelve months later at around 252,000 workers. Employment then slowly returns to normal. The responses of the two employment series—the one directly estimated and the other obtained by accumulation of flow responses—are close, which is good news for our approach and its underlying assumptions. The interesting feature of the figure is the behavior of the two gross flows. The increase in the flow out of employment accounts for more of the contraction in employment than does the decrease in the flow into employment. This remains true throughout the adjustment process. Had we looked at the effect of a positive innovation, and thus at an increase in employment, the linearity of our model would have implied that most of the adjustment was again through variations in the outward flow, rather than the inward flow.

Figure 6. Flows into and out of Employment, January 1968–May 1985

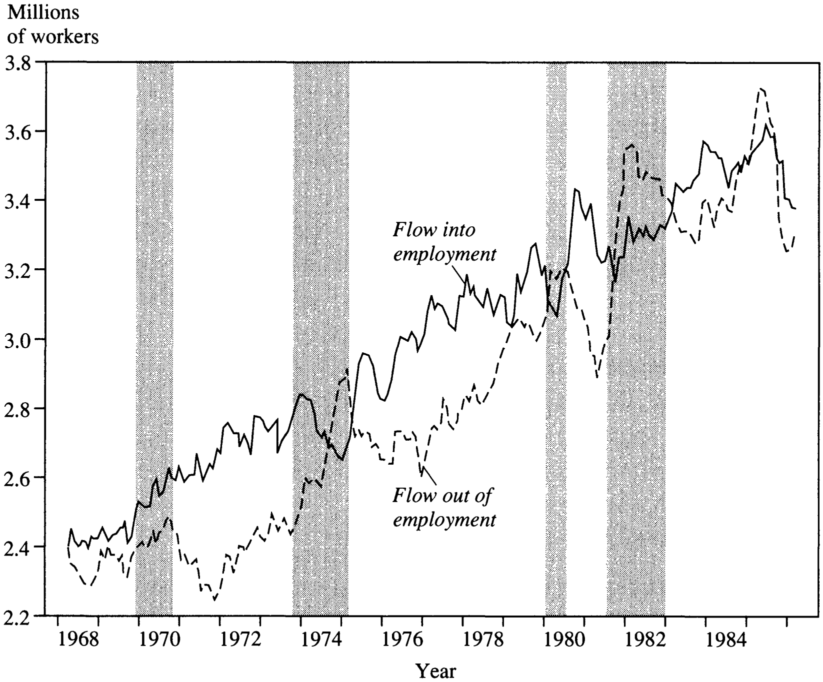

Source: Authors' own calculations using Abowd-Zellner data that have been seasonally adjusted and filtered using a seven-month moving average. Shaded areas represent recessions.

We wondered whether the asymmetry found in figure 5 was not hiding some nonlinearity in the data that may involve the way in which recessions differ from expansions. Thus, we have explored whether our results are sensitive to the inclusion of the two main recessions. We found that they are not sensitive to the exclusion of the 1973–75 recession, suggesting that oil prices do not explain much. The exclusion of the twin 1980–82 recessions, however, removes most of the asymmetry. Also, the results are not strongly affected by mechanically dividing the sample according to the level of unemployment.

For those skeptical of econometrics, figure 6 simply plots the flows into and out of employment, filtered by a seven-month moving average. In both recessions, it is clear to the naked eye that the increase in the outward flow accounts for more of the decrease in employment than the decrease in the inward flow. For the flow into employment, the standard deviation of the difference between it and a quadratic trend equals 202, with differences from trend ranging from -620 to 690; the standard deviation of the difference between the outward flow and a quadratic trend is 212, with differences from trend ranging from -730 to 880. The fact that the ranges differ much while the standard deviations differ little is consistent with our findings on the importance of removing the 1980–82 recession.

Should the finding of asymmetry, or even of symmetry, when the 1980–82 recession is omitted, surprise anyone? It surprised us. Suppose that we take the stylized fact as one of symmetry. To the extent that quits to nonemployment (which are included in the flow out of employment) and replacement of those quits (which are included in the inward flow) are procyclical, symmetry in the two flows implies much larger movements in job destruction than job creation, something we did not expect. We explore this theme when we look at the other data sets.[20]

Flows into and out of Manufacturing

Flows of workers to and from manufacturing firms were measured until 1981, when the collection was discontinued. These flows differ from those measured by the CPS in four major ways. First, they refer to the manufacturing sector only, so that a comparison with the CPS data sheds light on the difference between that sector and the entire economy. Second, the manufacturing data provide a different type of disaggregation, by quits versus layoffs, and new hires versus rehires. Third, as these data are collected by firms, the movement of an employed worker directly to another job is recorded, which is not the case with the CPS data. We refer to such movements as employment-to-employment movements, or *EE* flows. Fourth, they record cumulative changes rather than point-in-time status.

Thus, for comparability of the two sets of flows, we subtract movements from manufacturing employment to other employment from the flow out of employment, and subtract employment-to-manufacturing-employment movements from the flow into employment.[21] No time

20. These other data sets differ in many ways—coverage, stocks versus accumulated flows, and so on. We discuss the main differences in the text. A detailed comparison is available in Appendix B.

21. The evidence on employment-to-employment (*EE*) quits was the subject of a recent paper by Akerlof, Rose, and Yellen (1988).

series exists for *EE* quits. In our previous paper we assume that a constant fraction, 0.4, of quits were *EE* quits. We have examined this assumption using new data. Based on tabulations provided by Kevin Murphy of annual CPS data on individuals with different combinations of numbers of employers, stays out of the labor force, and stretches of unemployment, we have constructed upper- and lower-bound series for *EE* movements for males between 1968–88. These movements are surprisingly consistent with our earlier assumption that the proportion of quits that are *EE* quits is roughly constant, but the fraction's value is closer to 0.6.[22]

We thus construct two series, one for the flow out of manufacturing employment, constructed as separations minus 50 percent of quits, and one for the flow into manufacturing employment, constructed as the sum of accessions, minus 50 percent of quits.[23] Figure 7, akin to figure 6, gives the constructed flows into and out of employment for manufacturing; the flows are smoother than the CPS flows and do not require filtering to show their cyclical properties.

Comparison of figures 6 and 7 suggests that manufacturing shows greater symmetry than the aggregate flows. To get a more precise feel, we use the same methodology as earlier, running auxiliary regressions of the flows of layoffs, quits, new hires, and rehires for the period 1968–81 on the three aggregate stock variables, and tracing the effects of a cyclical shock on the accumulated flows. Because the manufacturing flows exhibit more serial correlation than their CPS counterparts, we allow for four lagged values of each of the four flows in each regression. Table 3 gives the responses in manufacturing employment and in manufacturing flows, disaggregated into layoffs, quits, new hires, and rehires, that result from a negative shock to aggregate activity. Table 4 gives the accumulated flows for both the economy as a whole, which are repeated from figure 5, and for the constructed flows into and out of manufacturing, which are the flows adjusted for 50 percent of quits.

For the period 1968–81, manufacturing employment accounts for 22 percent of employment, and the mean inward and outward manufacturing

22. The details of that computation, as well as the time series so constructed, are given in Appendix C.

23. The series constructed from the CPS tabulations cannot be used because they are annual. We need monthly series.

Figure 7. Flows into and out of Manufacturing Employment, January 1968–December 1981

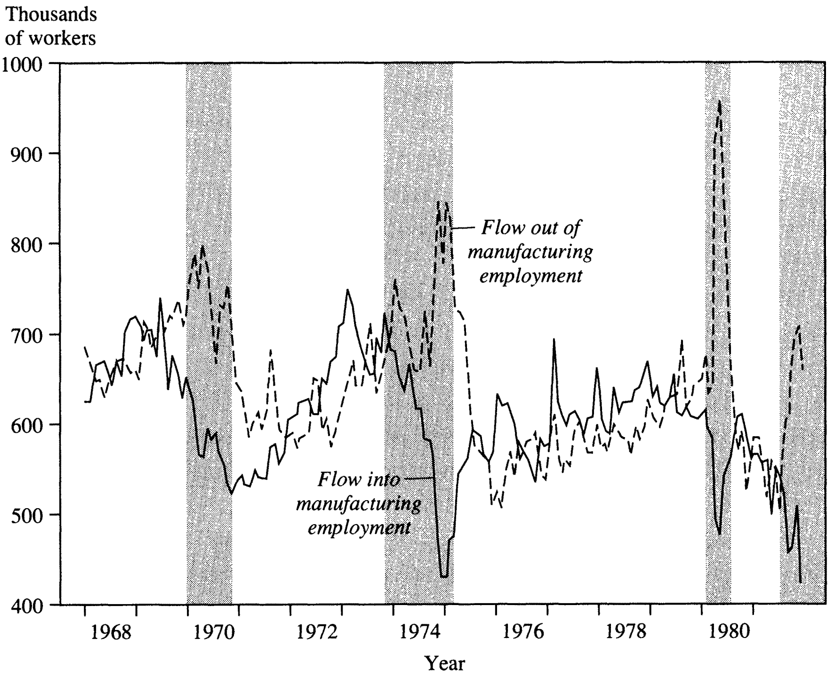

Source: Bureau of Labor Statistics, *Employment and Earnings*, monthly reports. The flow into manufacturing equals accessions less one-half of quits. The flow from manufacturing equals separations less one-half of quits. Shaded areas represent recessions.

flows are 21 and 25 percent of their aggregate counterparts respectively. However, the responses of the manufacturing flows and of manufacturing employment to an aggregate shock represent a much larger fraction of the response of the aggregate flows. The ratio of the response of manufacturing to total employment, using directly estimated responses, equals 26 percent in the current month, increasing to 56 percent over the year ahead. The ratio is even higher using the response of the stock that is implied by accumulation of the flows.[24]

24. After six months, the implied stock response for manufacturing employment exceeds the directly estimated response and becomes implausible. To us, this indicates misspecification of the dynamics of the flows in the auxiliary regressions used for those simulations. Therefore, one should probably focus on the response over the first six months alone.

Table 3. Response of Manufacturing Employment Stocks and Flows to a Negative Shock, January 1968–December 1981
Thousands of workers

	Manufacturing					
	Layoffs	*Quits*	*New hires*	*Rehires*	*Implied E*[a]	*Direct E*[b]
Means	269	401	573	232	19,739	. . .
Accumulated response						
Month 1	18 (2)	−5 (1)	−12 (1)	−1 (1)	−26	−26
3	69 (5)	−27 (4)	−52 (5)	5 (3)	−89	−88
6	135 (12)	−77 (11)	−101 (17)	28 (4)	−131	−153
12	203 (31)	−191 (35)	−273 (50)	80 (10)	−205	−163

Source: Authors' own calculations. Standard deviations are given in parentheses. The shock is defined as a one-standard-deviation negative innovation to the aggregate activity disturbance.
a. Implied E equals the response of manufacturing employment implied by cumulation of the estimated flow responses.
b. Direct E equals the measured response of manufacturing employment.

In response to a negative shock, layoffs go up but gradually the increase is offset in larger and larger proportion by a decrease in quits. Recalls increase, to reach 21 percent of layoffs after six months and 39 percent after a year. New hires decrease strongly.

Finally, turning to the constructed flows, we find rough symmetry between their cyclical responses, rather than the asymmetry found in the aggregate data. Thus, given procyclical quits and replacement tendencies, manufacturing data also imply a larger amplitude of fluctuations in job destruction than in job creation.

Table 4. Response of Employment Stocks and Flows to a Negative Shock, Manufacturing and Aggregate Data, January 1968–December 1981
Thousands of workers

	Manufacturing flows		*Aggregate flows*		*Aggregate*
	into E[a]	*out of E*[b]	*into E*	*out of E*	*direct E*
Means	605	641	2,852	2,609	89,543
Accumulated response					
Month 1	−11	16	−25	55	−101
3	−33	56	−72	90	−197
6	−61	96	−96	140	−277
12	−97	107	−92	160	−292

Source: Authors' own calculations. Aggregate flow data are repeated from figure 5. The shock is defined as a one-standard-deviation negative innovation to the aggregate activity disturbance.
a. The manufacturing flow into employment equals the hire rate plus the rehire rate less one-half of the quit rate.
b. The manufacturing flow out of employment equals the layoff rate plus one-half of the quit rate.

Job Creation and Job Destruction

To focus on job creation and destruction, we look next at the data set put together by Davis and Haltiwanger, which uses quarterly data on manufacturing firms' employment levels. They have constructed two series which they call "POS" and "NEG." The first is the sum of the positive changes in employment over all manufacturing firms, divided by manufacturing employment; the second is the sum of the negative changes in employment over all firms, again divided by manufacturing employment. Davis and Haltiwanger have analyzed these series at length; thus, we limit ourselves to the question of whether their series behave like the series we examined above. To find the answer, we construct two sets of series for job creation and destruction, one using the manufacturing flows above and one using the Davis-Haltiwanger series.

The Davis-Haltiwanger POS and NEG series differ from true job creation and destruction in that they ignore three phenomena: gross job creation and destruction within firms (they take a net number for each establishment); job creations offset by job destructions within a quarter for a given firm (they use quarterly point-in-time stocks); and the fact that firms may not be able to find workers to fill newly created jobs. It is impossible to adjust for the first two. To adjust roughly for the third, we add the change in the vacancy rate, Δv, to POS, assuming the vacancy rate in manufacturing moves with the aggregate. Thus, we construct a job creation and a job destruction series from the Davis-Haltiwanger data:

$$\text{DHJC} = \text{POS} + \Delta v,$$
$$\text{DHJD} = \text{NEG}.$$

The manufacturing flows into employment (hires and rehires) and out of employment (layoffs and quits) differ from true job creation and destruction in that the outward flow includes quits not associated with job destruction, and the inward flow includes replacements of those quits. The empirical difficulty is clearly that of constructing a series for quits that are not replaced. From a 1975 Department of Labor survey of jobseeking methods, we can obtain a rough estimate of "quits not replaced" at a particular point in time, namely January 1973. These data

Figure 8. Job Destruction and Job Creation, January 1972–December 1985

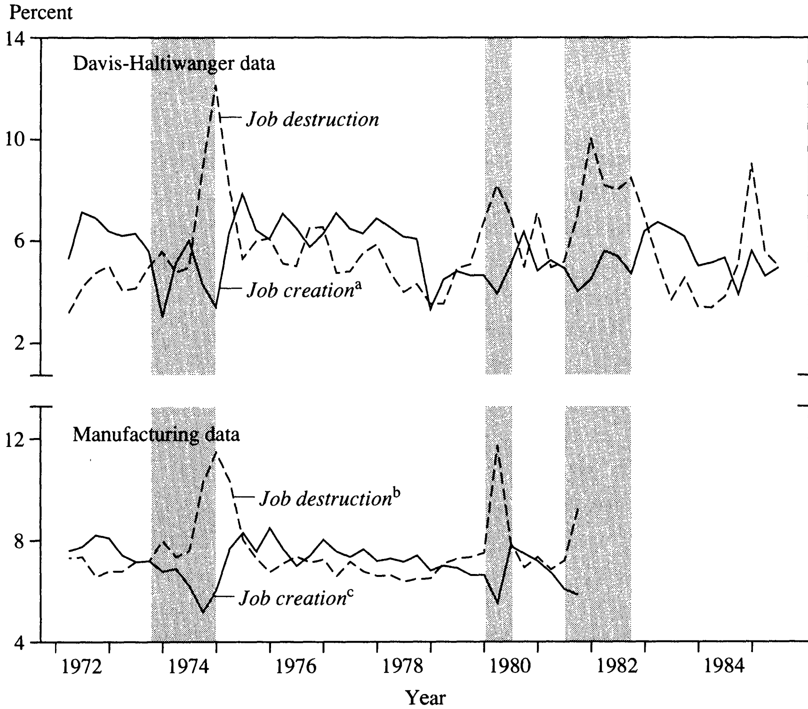

Sources: Davis and Haltiwanger (1990) and BLS, *Employment and Earnings*, monthly reports. Shaded areas represent recessions.

a. Job creation equals POS plus change in vacancy rate.

b. Job destruction equals the layoff rate plus 15 percent of quits.

c. Job creation equals hire rate plus rehire rate less 85 percent of quits plus the change in the vacancy rate.

suggest a lower bound of about 10 percent.[25] Given this lower bound, we proceed on the assumption that 85 percent of quits are replaced, and construct two series for manufacturing job creation and destruction:

$$\text{MJC} = h + r - 0.85q + \Delta v,$$
$$\text{MJD} = l + 0.15q,$$

where h, r, q, and l refer to the hire, rehire, quit, and layoff rates respectively.

In figure 8, we plot both sets of series. The series MJC and MJD have means of 7.0 and 7.5 percent respectively, compared to a mean of 5.5

25. We derive this number in Appendix D.

percent for DHJC and DHJD. This is what we would expect given that the Davis-Haltiwanger series are net of intrafirm job creation and destruction, and monthly movements in job creation and destruction reversed over the quarter. The cyclical behavior of both sets of series is similar. The standard deviations of DHJC and DHJD equal 1.1 and 1.6 percent respectively; the standard deviations of MJC and MJD equal 0.7 and 1.3 percent respectively. The correlation between MJC and DHJC is 0.71; between MJD and DHJD it is 0.82. In the two recessions covered by both sets of series, job destruction increases substantially more than job creation decreases.

It is natural to hypothesize that the asymmetry seen in the data can be explained by the way in which firms choose to adjust their employment levels. If, for example, firms adjust employment more through firing than hiring, this would generate flows into and out of employment that are consistent with our aggregate and manufacturing results. This explanation, however, is not sufficient to explain the pattern of flows observed in the Davis-Haltiwanger data. This is because the Davis-Haltiwanger data are based on net changes in employment at the plant level, and are, therefore, invariant to whether hiring or firing varies the most. To explain the Davis-Haltiwanger asymmetry, we would have to consider additional factors. We do so below.

The Davis-Haltiwanger data set allows disaggregation by sector or by size of firm. Among the results they derive, two will be important when we turn to interpretations. First, the larger amplitude of job destruction than of job creation is not an artifact of distribution effects across sectors: the same asymmetry holds at the four-digit level.[26] Second, while bankruptcies and plant closures are countercyclical, the proportion of job destruction due to closures is no higher in recessions.

Tentative Interpretations

The three data sets yield a consistent picture. Movements in employment appear to be associated with much larger fluctuations in job destruction than in job creation. Recessions are associated with large increases in job destruction and only small decreases in job creation. And, while direct evidence exists only for manufacturing, the indirect

26. Davis and Haltiwanger (1989).

evidence suggests that, if anything, the asymmetry is even stronger for the economy as a whole.

Any interpretation must first deal with the question of why, abstracting from cyclical movements, we simultaneously observe high rates of job creation and destruction. Clearly, much of the turnover comes from varying efficiencies, over time and across firms, in producing the same goods. The turnover also results from changes in tastes and incomes, leading to temporary or permanent changes in demand for specific goods in specific places. The high rates of turnover in small businesses come to mind as obvious examples.[27] Consider the demand source of turnover first; thinking of restaurants helps the intuition here. In a recession, we would expect fewer new restaurants to open and more to close. But are there reasons to expect that the adjustment will happen mainly through an increase in closings rather than a decrease in openings, as would be needed to explain the aggregate data? We can think of no good reason why this should be so. Indeed, economic theory gives a simple argument why the opposite is more likely true. To enter the market, new restaurants must anticipate covering average cost while existing restaurants, which already have their capital in place, only have to cover marginal cost. Thus, we would expect entry to be affected more than exit, or in the terminology used in this paper, we would expect job creation to vary more than job destruction.

The other source of turnover, technological progress, does suggest one possible explanation for the observed asymmetry. Suppose that the process of growth is one in which new, more productive processes replace old ones. Under that assumption, new jobs that produce at low marginal cost are largely immune to variations in demand. Old, high marginal cost jobs are not, and fluctuations in aggregate demand primarily affect the rate of job destruction. This tentative explanation, however, still has not confronted the marginal cost–average cost distinction we discussed with regard to restaurants. Moreover, we cannot think of direct evidence that new products are indeed more immune to aggregate demand fluctuations.

Turnover is not simply a mechanical process triggered by taste changes or technological progress. Firms in fact have the choice of whether and when to introduce new technologies or phase out obsolete ones, as well

27. See Brown, Hamilton, and Medoff (1990), for example.

as whether and when to hire or lay off workers. Thus, one potential line of explanation for the asymmetry is that the timing of job destruction is endogenous and concentrated in recessions.

If recessions are times when firms decide, or are forced, to close down unprofitable product lines, then job destruction will indeed be concentrated in recessions, leading to large variations in job destruction compared to job creation.[28] Davis and Haltiwanger, working within the framework of equilibrium cycles, suggest an explanation based on the idea that aggregate fluctuations are the result of transitory productivity shocks. If reallocating labor takes time, it is efficient to do so in recessions, which are periods of transitorily low productivity. We find their particular story unconvincing and believe there may be other reasons why additional job destruction takes place during recessions. Bankruptcies would have been a good candidate for explaining the asymmetry; the data by Davis and Haltiwanger show, however, that the proportion of job destruction due to plant closings actually decreases slightly in recessions. The argument may be rescued by invoking the effects of the fear of bankruptcy, and the presence of x-inefficiency. In good times considerable slack enters the operation of firms, and in bad times, when bankruptcy looms, the slack is squeezed out. This is a common but, as far as we know, little documented view of the value of recessions.[29]

28. For evidence at the plant level, see Cooper and Haltiwanger (1990). Lawrence Summers has suggested to us an analogy with heat waves. Heat waves are associated with higher death rates. This may be because heat waves precipitate the deaths of some people who would have died very soon anyway. This phenomenon is also known as intertemporal bunching.

29. This idea has also been suggested by Lasky (1990), who shows that high profits are followed by lower productivity growth. We want to mention two explanations that, while possible in the abstract, are ruled out by aspects of the data. Suppose that in recessions, firms rotate workers on layoffs. This would give rise both to additional movements into and out of employment in both the CPS and the manufacturing flows and to apparent additional job creation and destruction in recessions, leading to the observed asymmetry. But, by the nature of the Davis-Haltiwanger data, it would not appear in their measure of job creation and destruction.

The other explanation uses a similar argument, but does so across plants. Suppose that, because of fixed costs of operation, firms operate only a subset of plants in recessions and rotate among them. This would also lead to additional job creation and destruction, which would appear in all three data sets. Such rotations within firms, however, are likely to be present mostly in manufacturing, while the asymmetry is as at least as large outside of manufacturing.

Another line of explanation is based on asymmetric hiring and firing costs. Consider, for example, a firm that faces no cost of hiring but fixed (that is, independent of the number of workers fired) costs of firing, and that faces variations in demand.[30] Such a firm will hire workers often, and more or less as it needs them (not quite myopically, as it will take into account the potential costs of firing them if the need arose), but will fire workers infrequently and in large batches.[31] Such behavior appears to roughly fit our pattern of job creation and destruction. But we have learned to question whether such micro-asymmetries carry over to macroeconomic variables. If all firms were identical and faced identical shocks, the argument would trivially carry to the aggregate. But the very coexistence of job creation and destruction is itself proof of heterogeneity. Ricardo Caballero tackles precisely the issue of aggregation.[32] He concludes that the micro-asymmetry is unlikely to carry over to macro data. There are, however, two cases in which it may. The first arises when recessions are much more abrupt than expansions. In this case, at the beginning of a recession, many firms will quickly reach the firing threshold, leading to a large burst of job destruction. Our finding that leaving out the 1980–82 twin recessions reduces the asymmetry suggests that this may be part of the answer. The second arises if the costs of hiring and firing vary systematically over the cycle in a similar way for many firms. In a depressed labor market, for example, the risk of losing workers on layoffs to other firms is lower; thus, the cost of placing workers on temporary layoff is smaller.

Our main conclusion from this section is that changes in employment are dominated by movements in job destruction rather than in job creation. This finding rules out several explanations of the business cycle. One theory that fares poorly is the idea, along Schumpeterian lines, that explains booms as times of bunching of new product introductions.[33] However, we cannot be sure what actually explains the asym-

30. William Brainard has suggested reasons why this may be so. In contrast to hiring, firing may affect morale and productivity, as well as create problems with unions. While the effect is probably related to the number of workers fired, it may still be closer to a fixed cost than a proportional one.

31. We have learned a lot recently about the behavior of employment under such rules. See, in particular, Bertola and Caballero (1990).

32. Caballero (1990).

33. See Shleifer (1986).

metry in our result. The explanation that strikes us as most plausible is the one that views recessions as times of cleaning up. That argument, however, needs to be examined using other sources of data.

Flows between Employment, Unemployment, and Not in the Labor Force

The method of estimation and simulation outlined earlier allows us to trace the effects of a cyclical shock on each of the six CPS flows between *E, U,* and *N*. In an extension of figure 5, figure 9 gives the cumulated response of each flow to a one-standard-deviation negative innovation to the aggregate activity disturbance. Again, all numbers are in thousands of workers and refer to differences with the no-innovation case. The numbers alongside the arrows are the cumulated changes in the flows from month one (when the innovation takes place) to month *n*. Two numbers are given for the stocks: the first obtained by cumulation of the estimated response of the flows; the second (in brackets) obtained directly from equation 1 for *E* and *U*, and from an auxiliary regression of the "want-a-job" group on stocks for *N*. Finally, for reasons that will be clear later, we give in brackets under the *UE* flow the estimated response of recalls in manufacturing from table 3. Again, to allow readers to judge whether our econometric results capture the basic aspects of the data, figures 10–12 plot the flows, filtered with a centered seven-month moving average, between *E* and *U, E* and *N*, and *N* and *U* respectively.

The Sharp Difference between U and N in Aggregate Data

The main point made by the figures, and visible to the eye in figures 10 and 11, is the sharply different cyclical behavior of the flows between *E* and *U* on one hand, and between *E* and *N* on the other. In a recession, both the *EU* flow and, to a lesser degree, the *UE* flow increase. In contrast, the *EN* flow and the *NE* flow both decrease. Put another way, looking at the composition of the flows out of *E* in a recession, note that the flow into *U* increases sharply while the flow into *N* decreases sharply; looking at the flows into *E*, note that the flow out of *U* increases while the flow out of *N* decreases.

Figure 9. Response of Gross Flows between Employment, Unemployment, and Not in the Labor Force to an Aggregate Activity Shock, over Selected Intervals[a]

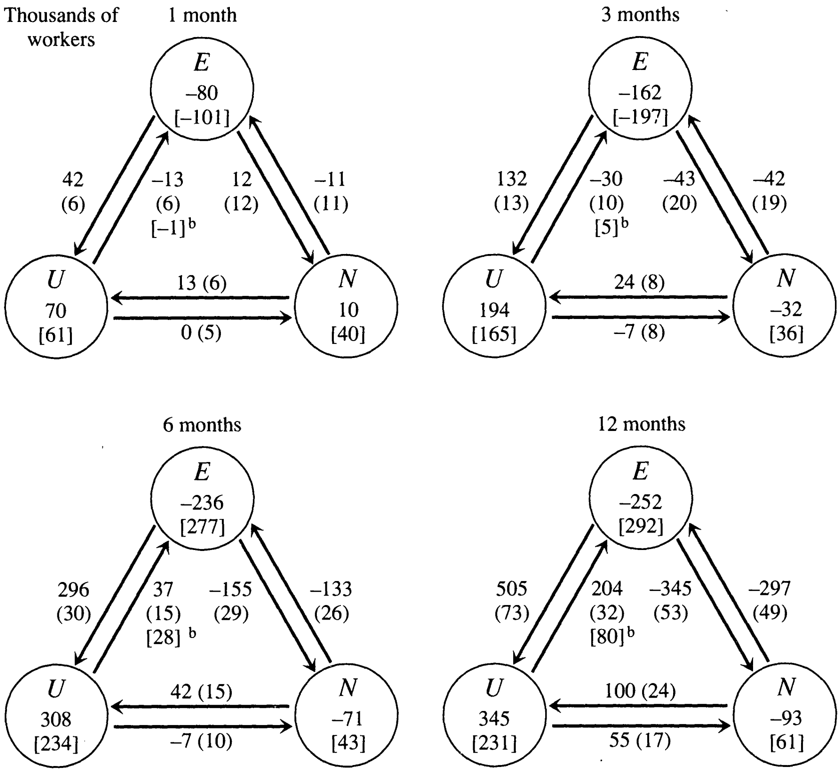

Source: Authors' own calculations using stock numbers from the Current Population Survey and the Abowd-Zellner adjusted gross flow series. The variables E, U, and N represent employment, unemployment, and not in the labor force respectively.

a. The figure shows the effects of a one-standard-deviation negative innovation in the aggregate activity disturbance on stocks and flows of employment, unemployment, and not in the labor force. Numbers refer to differences with the no-innovation case. Numbers alongside the arrows are the cumulated flow changes from the month of the shock to month n. Two numbers are given for stocks—the first gives the cumulated response of the flows, the second (in brackets) gives the response of employment obtained by direct estimation of equation 1 for E and U, and from an auxiliary regression of "want a job" on stocks for N. Standard deviations appear in parentheses.

b. Numbers in brackets under the UE flow equal the estimated response of recalls in manufacturing, from table 3.

The fact that the *UE* flow increases in a recession raises an obvious question: whether the probability for an unemployed worker becoming employed actually rises in a recession! To answer this question, we compute the probabilities, also known as the hazard rates, of going from one of the three pools into one of the other two. The probabilities are implied by the dynamic effects of the aggregate activity disturbance

Figure 10. Flows between Employment and Unemployment, January 1968–December 1981

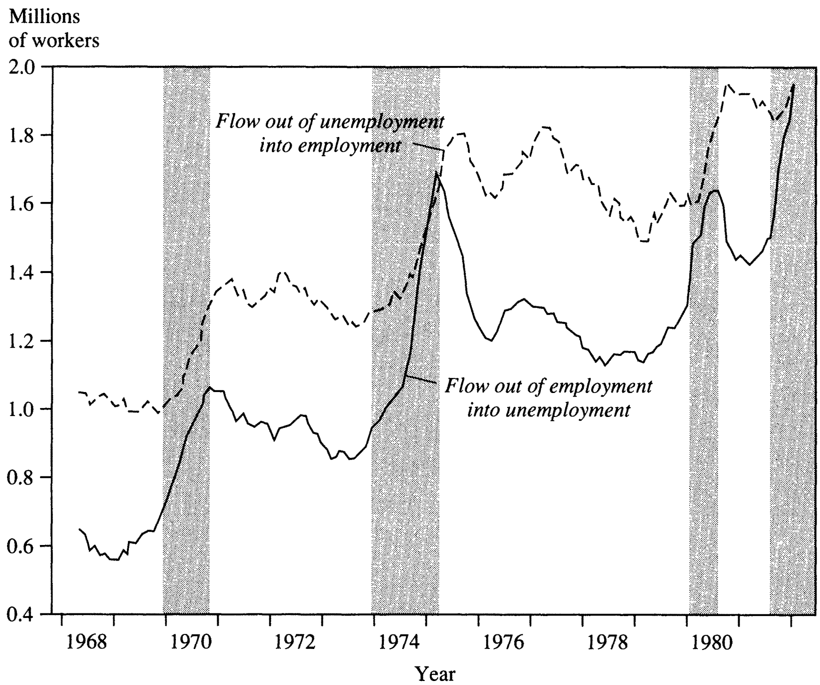

Source: Authors' own calculations using Abowd-Zellner data that have been filtered using seven-month moving average. Shaded areas represent recessions.

on the flows and the stocks. For each month, we compute the hazard rate from one stock to another, say from X to Y, to use a general example, as the ratio of the flow from X to Y in that month divided by the stock of X in the previous month implied by cumulation of the flows. We use as initial values for E and U the mean values of employment and unemployment. For the initial value of N, we use the mean value of the pool of people not in the labor force but who want a job. The changes in hazard rates from their preshock values are given in figure 13 (average hazard rates themselves are easily derived from the mean flows and stocks in figure 1).

While the flow from unemployment to employment increases in a recession, the hazard rate decreases as the pool of unemployed increases proportionately more than the flow. The hazard rate from U to E goes

Figure 11. Flows between Employment and Not in the Labor Force, January 1968–December 1981

Source: Authors' own calculations using Abowd-Zellner data that have been filtered using seven-month moving average. Shaded areas represent recessions.

from a preshock value of 24.0 percent to an average value of 23.6 percent over the six months following the shock.[34]

Some of the flows between E and U are the result of temporary layoffs and recalls. A natural question is how much these two flows contribute to the changes in the total flows between E and U. However, neither the flow data for layoffs nor those for recalls are available; indeed, firms often do not know whether a layoff will be temporary or permanent. Recalls are available for manufacturing, and we derived their response to a cyclical shock in table 3.[35] Also, one may recall that figure 9 offers the estimated responses of recalls in manufacturing in the brackets under

34. These are small changes; remember that the increase in the unemployment rate as a result of a one-standard-deviation innovation is itself small, about 0.3 percent.

35. Also see Lilien (1980).

Figure 12. Flows between Unemployment and Not in the Labor Force, January 1968–December 1981

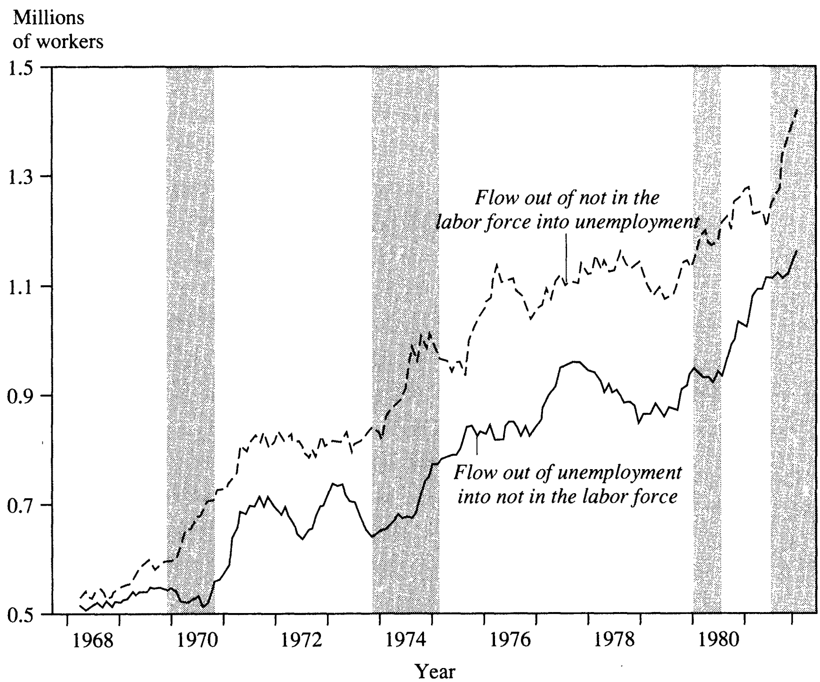

Source: Authors' own calculations using Abowd-Zellner data that have been filtered using seven-month moving average. Shaded areas represent recessions.

the *UE* flows. To the extent that some recalls also take place outside of manufacturing, this number is an underestimate of the effect of recalls. The proportion of the increase in the *UE* flow that is the result of the increase in recalls varies over time; initially, recalls increase while the *UE* flow decreases. A year after the shock, the increase in recalls accounts for 38 percent of the increase in the accumulated *UE* flow.

Two other aspects of figure 9 also bear mention. First, neither of the flows between *U* and *N* is strongly cyclical. The *NU* flow increases slightly, while the *UN* flow first decreases and then increases. (The large change in hazard rates comes from the denominators, not the numerators. Thus, if the relevant stock of would-be job takers in *N* is considerably larger than the measured want-a-job series, the change in the hazard rates would be less.) Second, while the responses in *E* (constructed by

Figure 13. Response of Hazard Rates between Employment, Unemployment, and Not in the Labor Force to Aggregate Activity Shock[a]

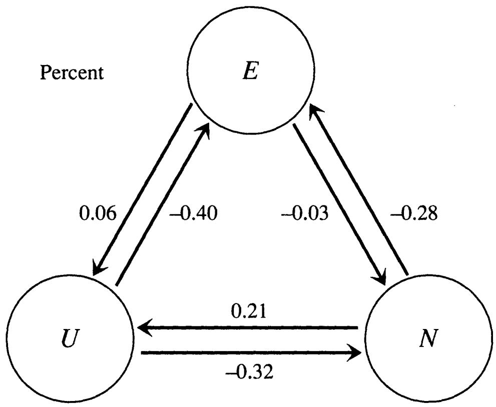

Source: Authors' own calculations using stock numbers from the Current Population Survey and the Abowd-Zellner adjusted gross flow series. The variables *E*, *U*, and *N* represent employment, unemployment, and not in the labor force respectively.

a. Numbers refer to differences in hazard rates with their preshock values, averaged over the first six months following a shock. Again, the shock is defined as a one-standard-deviation innovation in the aggregate activity disturbance.

accumulation and estimated directly) closely coincide, the same is not true of the responses in *N*, obtained by accumulation of the flows or by direct estimation of the response of the pool of people not in the labor force but who want a job. The directly estimated response shows an increase in *N* in a recession, which conforms with our beliefs. The implied response, however, shows a decrease. This rough check suggests that something is wrong with either the flow series or the assumptions underlying our approach.[36] We have explored at length whether this result is sensitive to variations in sample, lag length, and so on. We regret to say that it is an extremely robust feature of our results.

Disaggregation by Age and Sex

We now turn to the disaggregated evidence. The question we have in mind is whether the cyclical characteristics of the aggregate flows are

36. Our approach to identification, however, cannot be blamed for this result. No matter how we identify our aggregate activity shock, the two sets of responses should be identical if the other assumptions are satisfied.

Table 5. Response of Adjusted and Unadjusted Gross Flows to a Negative Shock, January 1968–May 1986

	Flows					
	EU	*EN*	*UE*	*NE*	*UN*	*NU*
Abowd-Zellner adjusted flows						
Month 1	42	12	−11	−11	0	13
3	132	−43	−30	−42	−7	24
6	296	−155	38	−133	−7	42
12	505	−345	204	−297	55	100
Unadjusted flows						
Month 1	37	13	−17	−11	1	13
3	110	−16	−29	−33	12	34
6	243	−81	27	−86	61	95
12	413	−233	171	−202	226	248

Source: Authors' own calculations using the Current Population Survey unadjusted flow data and the Abowd-Zellner adjusted flow series. The shock is defined as a one-standard-deviation negative innovation to the aggregate activity disturbance.

reflected in most age-sex groups or whether there are instead sharp differences across groups.

In looking at the disaggregated gross flows, we have to use unadjusted data because adjusted data do not exist since adjustments become more problematic as the sample size dwindles. To judge whether it makes sense to look at the cyclical properties of the unadjusted data, we compute the responses to a cyclical shock using both adjusted and unadjusted *aggregate* flows. Table 5 presents the results. Impulse responses derived from unadjusted flows tend to underpredict the responses of the flows into and out of E and to overpredict considerably the responses of flows between N and U. These under- and overpredictions should be kept in mind when interpreting the results below.

The disaggregated results for males are presented in table 6, and for females in table 7. In each case, results are given for eight age groups, for each of which we use the same approach as for the aggregate flows. The set of six flows for each group is regressed, as in equation 4, on current and lagged values of the three aggregate stocks. The simulation traces the effects of a one-standard-deviation negative innovation on each of the six flows. Each table gives the cumulated flows after six months, the mean hazard rates, and the average change in those hazard rates over the first six months following the shock.

A pattern of clear cyclical differences across groups emerges. Starting

Table 6. Response of the Flows and Hazard Rates for Male Workers to a Negative Shock, Disaggregated by Age Group, January 1968–May 1986

Age group	Flows: EU	EN	UE	NE[a]	UN	NU[a]	Mean E
	Accumulated flow responses after six months (in thousands)						
16–19	12*	−14*	−9*	−18*	16*	20*	3,638
20–24	33*	−5	3	−7*	12*	10*	6,273
25–34	56*	−1	20*	−2	4*	7*	13,714
35–44	33*	−1	11*	−1	2	1	11,193
45–54	24*	0	8*	1	0	1	10,084
55–59	8*	2	1	−1	1	1	4,239
60–64	4*	−1	0	−4*	0	0	2,805
65+	4*	−4*	−1	−4*	−2	0	1,921
Total	174	−24	33	−36	33	40	. . .
	Mean hazard rates (in percent)						
16–19	4.5	11.4	27.9	12.9	30.2	8.0	. . .
20–24	3.2	3.2	31.0	14.9	13.8	9.3	. . .
25–34	1.6	0.7	29.2	13.0	8.9	10.6	. . .
35–44	1.1	0.5	29.2	10.5	8.5	7.3	. . .
45–54	1.0	0.7	27.7	6.9	10.1	4.0	. . .
55–59	0.9	1.2	24.7	4.8	12.7	2.2	. . .
60–64	0.8	2.9	20.4	3.4	19.9	1.1	. . .
65+	0.6	10.4	15.9	2.0	36.8	0.3	. . .
	Change in hazard rates (average over first six months)						
16–19	0.08	−0.02	−0.65	−0.11	−0.14	0.08	. . .
20–24	0.08	−0.01	−0.54	−0.13	0.00	0.10	. . .
25–34	0.07	0.00	−0.25	−0.04	−0.12	0.17	. . .
35–44	0.04	0.00	−0.30	−0.10	−0.18	0.02	. . .
45–54	0.04	0.00	−0.33	−0.03	−0.25	0.05	. . .
55–59	0.03	0.00	−0.47	−0.02	−0.15	0.00	. . .
60–64	0.02	0.00	−0.17	−0.05	−0.26	0.00	. . .
65+	0.02	−0.01	−0.87	−0.02	. . .[b]	0.00	. . .

Source: Authors' own calculations using data from Current Population Survey. The shock is defined as a one-standard-deviation negative innovation to the aggregate activity disturbance. One asterisk represents significance at the 5 percent level.

a. Because data series on individuals not in the labor force but who want a job are not available at this level of disaggregation, we computed the hazard rates out of *N* using the total number of people not in the labor force for *N*. This tends to make the absolute levels of the *NE* and *NU* hazard rates meaningless for both the young and older workers.

b. Estimate is unreliable.

with males, and without doing too much violence to the facts, one can distinguish among three groups: the young workers (16–19 year olds), who are a small proportion of the total labor force and, on average, exhibit different behavior as reflected in hazard rates; the mature workers (20–59); and the older workers.

Young workers account for half of the decrease in the *EN* and *NE* flows. Interestingly, except at the very beginning, the decrease in the *EN* flow results more from a decrease in the number of employed workers

Table 7. Response of the Flows and Hazard Rates for Female Workers to a Negative Shock, Disaggregated by Age Group, January 1968–May 1986

Age group	*Flows* *EU*	*EN*	*UE*	*NE*[a]	*UN*	*NU*[a]	*Mean E*
	Accumulated flow responses after six months (in thousands)						
16–19	7*	−3	−6*	−11*	12*	17*	3,093
20–24	14*	−10*	−2	−4	5	11	5,366
25–34	18*	−9*	5	−8*	8	10*	8,922
35–44	10*	−7	−1	−6*	3	6*	7,462
45–54	11*	−3	2	−1	3	5*	6,635
55–59	4*	−3	1	−4*	0	1	2,679
60–64	2*	−2	1	−2	0	1	1,682
65+	0	−3	−1	−2	0	0	1,072
Total	66	−40	−1	−38	31	52	...
	Mean hazard rates (in percent)						
16–19	3.5	12.6	26.4	9.9	36.4	6.8	...
20–24	2.1	5.3	27.2	7.5	28.2	5.5	...
25–34	1.4	4.3	23.4	5.1	28.4	3.0	...
35–44	1.1	3.9	23.7	5.4	27.8	2.2	...
45–54	1.0	3.7	22.0	4.4	27.1	1.6	...
55–59	0.8	3.9	20.6	3.1	25.6	0.9	...
60–64	0.7	5.6	19.9	2.2	27.6	0.4	...
65+	0.6	11.7	16.4	0.8	37.7	0.1	...
	Change in hazard rates (average over first six months)						
16–19	0.05	0.03	−0.65	−0.04	−0.36	0.08	...
20–24	0.05	−0.02	−0.62	−0.01	−0.41	0.07	...
25–34	0.04	−0.01	−0.15	−0.02	−0.13	0.03	...
35–44	0.03	−0.01	−0.45	−0.02	−0.32	0.02	...
45–54	0.03	−0.00	−0.29	−0.00	−0.32	0.02	...
55–59	0.02	−0.02	−0.14	−0.03	−0.34	0.01	...
60–64	0.02	−0.02	−0.10	−0.01	−0.56	0.00	...
65+	0.00	−0.03	−0.56	−0.00	−0.70	0.00	...

Source: Authors' own calculations using data from Current Population Survey. The shock is defined as a one-standard-deviation negative innovation to the aggregate activity disturbance. One asterisk represents significance at 5 percent level.

a. Because data series on individuals not in the labor force but who want a job are not available at this level of disaggregation, we computed the hazard rates out of *N* using the total number of people not in the labor force for *N*. This tends to make the absolute levels of the *NE* and *NU* hazard rates meaningless for both the young and older workers.

than from a sharply declining hazard rate. The decrease in the *NE* flow is associated with a sharp decline in the *NE* hazard rate. Young workers also have the highest increase in the *EU* hazard rate and the second highest decrease in the *UE* hazard rate. They experience the largest decrease in the *UE* flow. They also experience abnormally large increases in the flows between *N* and *U*, but unadjusted flows tend to overstate such movements.

Cyclical movements in the flows of *mature* workers are concentrated in the flows between *E* and *U*. For all age groups, both the *EU* and *UE* flows increase in response to a negative shock to aggregate activity. For all age groups also, the *EU* hazard rate increases, and, because the pool of *U* increases, the *UE* hazard rate decreases. There is a fairly clear effect of age on the cyclical change in hazard rates: the older the workers, the more insulated they are from cyclical fluctuations. The decrease in the *EN* flow is small and associated with unchanged hazard rates. The decrease in the *NE* flow is also small and associated with decreasing hazard rates.

The cyclical behavior of *older* workers reflects elements of both mature and younger workers. They share with mature workers a very small increase in the *EU* hazard rate, and with younger workers the very large decrease in the *UE* hazard rate. They share with younger workers a decrease in the *EN* hazard rate, and with mature workers the absence of any cyclical effects on the flows between *N* and *U*.

The cyclical behavior of the flows of females resembles that of males, and one can again distinguish among the three age groups. However, there are differences. The heterogeneity across female groups is less striking than for males. All age groups, except for the youngest group, experience a decrease in the *EN* hazard rate. Females account for much more of the movement between *E* and *N* than males.

Heterogeneity in labor market transitions has been a popular theme in the labor–macroeconomics literature.[37] Our results show that this heterogeneity extends to the cyclical behavior of flows as well.

Labor Market Dynamics with Two Types of Workers

In this section, we examine features of the labor market that can explain the empirical findings previously discussed. The evidence leads us to consider a class of models that has two types of workers, "primary" and "secondary." We present a simple formalization and show how it can explain the data. Extending our earlier work on matching, we give empirical evidence for a crucial aspect of our model: the preference of firms for hiring primary over secondary workers. Finally, we discuss several extensions and implications of our model.

37. In particular, see Hall (1970) and Clark and Summers (1979).

Primary and Secondary Workers

We see the evidence as consistent with the existence of two groups of workers who differ in their attachments to the labor market. We see the various age-sex groups as composed of different proportions of these two underlying groups.[38]

"Primary" worker conjures the image of a head of household, with long job tenure, infrequent movements into and out of the labor force, and brief spells of unemployment. "Secondary" worker conjures the image of a teenager hopping from job to job, with intervening periods both of unemployment and time out of the labor force. There are many relevant aspects to the labor market attachments of these two groups. Secondary workers may quit their jobs more often (in what follows, we shall use quits to signify quits to nonemployment). Secondary workers may drop out of the labor force more often, while if primary workers leave employment, they go mainly to U. Search behavior may differ. And the two types may be perceived differently by firms, leading firms to prefer hiring primary over secondary workers and to prefer firing secondary workers first.

A Model of Primary and Secondary Workers

Consider the following assumptions. The economy is subject to continual job creation and destruction, which leads firms to continually lay off workers and advertise new vacancies. Of the two types of workers, primary workers do not quit; they leave employment only when laid off and then go only into unemployment. Secondary workers leave employment both through layoffs and quits; at that time, they drop out of the labor force, going into N.[39] Both primary and secondary workers are acceptable to firms; however, when firms have a choice, they prefer hiring a primary rather than a secondary worker.[40]

38. A similar two-types model was proposed by Darby, Haltiwanger, and Plant (1985) to explain the behavior of hazard rates from U to E as a function of duration.

39. The model could easily be extended by assuming that some fraction of the secondary workers were in U. This could then accommodate flows between U and N.

40. In Blanchard and Diamond (1990a) we impose a similar assumption and derive the form of the matching functions from an explicit description of a meeting process between firms and workers. The resulting functional forms differ from the simple Cobb-Douglas forms used in both the simulation and estimation here. In Blanchard and Diamond (1990b) we study what happens when ranking follows from differences in training costs between otherwise identical workers who are free to bargain over wages.

These assumptions leave out differences in layoff behavior, search behavior, and cyclical quit rates. Nevertheless, they are sufficient to generate the qualitative cyclical behavior of the gross flows and even do a decent quantitative job of fitting the impulse responses.

More specifically, we assume the economy is composed of one type of job and two types of workers—primary (denoted by 1) and secondary (denoted by 2). Primary workers can be either employed or unemployed. Let L_1 denote the primary labor force. Thus, $L_1 = E_1 + U$, where E_1 denotes primary worker employment and U denotes unemployment (because there are only primary workers in unemployment, there is no need for a subscript). Secondary workers can be either employed or not in the labor force. Thus, $L_2 = E_2 + N$, where N is the number of secondary workers not in the labor force. The values for L_1 and L_2 are given.

Jobs take three forms: filled, unfilled with a vacancy posted, or unfilled with no vacancy posted. Each job requires one worker. Let K be the total number of jobs, F the number of filled jobs, V the number of vacancies, and I the number of unfilled jobs with no vacancy posted, or idle capacity. Thus, $K = F + V + I$. Obviously $F = E = E_1 + E_2$. K is given. Each of the K jobs produces, if filled, a gross (of wages) revenue of either 1 or 0. The 0–1 productivity for each job follows a Markov process in continuous time. A productive job becomes unproductive with flow probability π_0. An unproductive job becomes productive with flow probability π_1. At any point in time, some jobs become productive and others unproductive. This is the black box mechanism we use to capture the large gross flows of job creation and destruction that exist in the economy.

In addition to the movement of workers that results from job creation and destruction, there is also movement because of quits. Primary workers do not quit. Secondary workers quit at the constant rate q.

The process of matching workers and jobs is captured by a matching function, where hires, h, is a function of the pool of nonemployed workers and of vacancies. We assume that primary and secondary workers have identical search behavior, and that both types are acceptable to all firms. Thus, the aggregate matching function is given by

$$h = m[(U + N), V],$$

where m_U, $m_V \geq 0$. We also assume that when firms have the choice, they prefer primary over secondary workers; in other words, firms *rank* primary workers above secondary ones. This assumption implies the existence of a matching function that gives hires of primary workers the form

$$h_1 = m_1(U, V),$$

where we assume again $m_{1,U}, m_{1,V} \geq 0$. The important characteristic of this function is that N does not appear. Given vacancies, a larger number of secondary workers in N does not affect the chances of the unemployed primary workers. The hiring function for secondary workers is given by $h_2 = h - h_1$.

These assumptions lead to the following equations of motion:

$$dV/dt = -h - \pi_0 V + \pi_1 I + qE_2,$$

$$dE_1/dt = -\pi_0 E_1 + h_1,$$

$$dE_2/dt = -(\pi_0 + q)E_2 + h_2.$$

For a job to produce one unit of output, it must not only be productive but must also be matched with a worker. To do so, a vacancy must be posted and a worker must be recruited. New vacancies come from quits, qE_2, and from jobs that were previously unproductive becoming productive, $\pi_1 I$. Vacancies decrease for two reasons: some are filled by new hires, h; some become unproductive before they are filled, $\pi_0 V$. The change in primary worker employment equals hires, h_1, minus layoffs, $\pi_0 E_1$. The change in secondary worker employment equals hires, h_2, minus layoffs and quits, $(\pi_0 + q)E_2$.

Using the various identities, these three equations can be written as a differential equation system in U, N, and V. We can then suppose that the economy is subject to an adverse cyclical shock, leading to an increase in the rate of job destruction, π_0, and a decrease in the rate of job creation, π_1. We then ask what happens to the flows between E and U and between E and N.

Intuition suggests the answer. As layoffs increase, the flows of primary and secondary workers out of employment increase. As the pool of employed secondary workers decreases, however, the number of quits decreases, even at a constant quit rate. Thus, while the *EU* flow unambiguously increases, what happens to the *EN* flow is ambiguous. On the hiring side, decreases in job creation and in quits lead to a decrease in job vacancies. This combined with ranking, and with the increase in the pool of unemployed primary workers, sharply decreases the chances of secondary workers finding work. Thus, the *NE* flow decreases. What happens to the *UE* flow is ambiguous. The larger pool of unemployed may offset the effect of decreased vacancies and lead to an increase in the number of hires from *U*; that is, to an increase in the *UE* flow. Thus, the model has the potential to explain the signs of all four flows between *E* and *U*, and *E* and *N*. It obviously cannot explain movements between *U* and *N*; they have been ruled out for simplicity's sake.

To see how well such a model could do, we have performed simple simulations. The simulation first requires the specification of the two hiring functions. We take h and h_1 to be constant-returns Cobb-Douglas functions in the relevant pool and in vacancies:

$$h = a(U + N)^{1-\alpha} V^{\alpha};$$

$$h_1 = a_1 U^{1-\alpha_1} V^{\alpha_1}.$$

We then calibrate the model so as to replicate steady-state values of the flows and stocks. We take those to be, in millions, $L_1 = 80$, $L_2 = 20$, $U = 5$, $N = 4$, $V = 3$, $I = 11$, and $EU = UE = 1.4$, $EN = NE = 1.4$.

This steady-state calibration ties down all the coefficients, except for two pairs—the combinations of a and α, and a_1 and α_1 in the two matching functions. Different combinations lead to the same steady-state flows of hiring, but to different strengths of the ranking effect. Estimation of the aggregate matching function corresponding to h in our previous paper suggests a value for α of around 0.6. Our estimates of the matching function h_1 suggest a value of α_1 of 0.2. These two estimated coefficients imply strong ranking, which means a strong deterioration of hiring prospects for secondary workers when unemployment increases. The resulting set of parameters is $\pi_0 = 0.019$; $\pi_1 = 0.16$; $q = 0.07$; $a = 0.6$; $\alpha = 0.6$; $a_1 = 0.3$; $\alpha_1 = 0.2$.

The last step is to specify the cyclical shock, which we define as an increase in job destruction and a decrease in job creation. We calibrate the shock so as to replicate the rate and level of the increase in unemployment shown in our impulse responses, roughly 0.3 percent over the first twelve months. Also, based on results from the first part of the paper, we choose changes in π_0 and π_1 that imply an increase in the rate of job destruction that is larger than the decrease in the rate of job creation.

Table 8. Responses of the Stock of Unemployment and the Flows into and out of Employment to a Negative Shock

Thousands of workers

						Flows[a]	
	U	*EN*	*EU*	*NE*	*UE*	*into E*	*out of E*
Estimated response							
Month							
1	70	12	42	−11	−13	−25	55
6	308	−155	296	−133	37	−96	140
12	345	−345	505	−297	204	−92	160
Simulation response							
Month							
1	23	4	22	−25	−1	−26	26
6	197	−11	280	−374	83	−290	269
12	305	−197	672	−866	367	−499	474

Source: Authors' own calculations. Estimated responses are repeated from figure 9. The shock is defined as a one-standard-deviation negative innovation to the aggregate activity disturbance.

a. Numbers may not add because of rounding.

More specifically, we assume that π_0 and π_1 adjust at a rate of 30 percent per month from their initial values to new steady-state values of 0.020 and 0.145 respectively. The results of the simulation are given in table 8, which also repeats their empirical counterparts from figure 9.

The simulation results indicate that the model replicates the signs of the responses of the four flows, confirming intuition. The model even does a decent job of fitting the quantitative responses of the flows between E and U. However, it generates a much larger decrease in the NE flow and a much smaller decrease in the EN flow. If we allow the quit rate to depend on labor market conditions, that would help explain the quantitative features of the NE flow but would further lower the EN flow. It may not be desirable to fit all aspects of the VAR results. In particular, they imply a decrease in N in a recession, which we do not find plausible.

Evidence on Ranking of Primary and Secondary Workers

In our previous paper, we looked at hires from the unemployment and not in the labor force pools and concluded that a model that assumed perfect substitutability was strongly rejected by the data. We now explore

Table 9. Hiring from Unemployment and Not in the Labor Force, February 1970–December 1981

			Logarithms of lagged stock variables			*Summary statistics*		
	Constant	*Time* ($\times 10^{-2}$)	$\ln(V_{-1})$	$\ln(U_{-1})$	$\ln(N_{-1})$	*RTS*	*DW*	R^2
Dependent variable								
$\ln(H_U)$	−1.24	0.01	0.13	0.78	0.00	0.91	2.0	0.77
		(0.1)	(1.3)	(3.6)	(0.0)			
	−1.30	−0.01	0.15	0.80	0.04	1.00*	2.0	0.77
		(0.3)	(1.6)	(3.7)	(0.1)	(SF = 0.7)		
$\ln(H_N)$	0.12	−0.06	0.35	−0.40	0.26	0.21	1.9	0.42
		(0.7)	(2.7)	(0.7)	(1.0)			
	−0.43	−0.10	0.55	−0.26	0.71	1.00*	1.8	0.39
		(3.9)	(4.6)	(0.9)	(1.9)	(SF = 0.003)		

Source: Authors' own calculations. See Blanchard and Diamond (1989b) for a more detailed analysis. The t-statistics are in parentheses. The variables H_U and H_N stand for hires out of unemployment and hires out of the want-a-job category of not in the labor force respectively. The estimated degree of returns to scale is designated by RTS. The asterisk means a constant returns to scale restriction was imposed on the regression. The SF notation stands for the significance level of a test that RTS were constant.

whether the matching–ranking assumptions made in this paper are more consistent with the data.[41]

We report estimated matching functions for workers in the U and N pools in table 9.[42] We take hires from unemployment, H_U, to equal the flow from unemployment to employment minus 1.5 times the number of recalls in manufacturing. The number 1.5 is a rough estimate of the importance of recalls outside manufacturing. Recalls do not require the posting of vacancies. We use the number of workers not in the labor force but who want a job for N, and take hires from N, H_N, to equal the flow of workers from not in the labor force into employment. Data for N exist only in quarterly form and are, therefore, linearly interpolated. We use adjusted vacancies for V, and use the unemployment pool net of "job losers on layoff"—that is, workers who consider themselves to have a job—for U. The specification of the matching functions under

41. See Blanchard and Diamond (1989b, table 4) for a comparison of past and present results. The minor differences in the results come from a different method of seasonal adjustment (we used a band filter in 1989 and use the X11 program here).

At this time, we shall not attempt an estimation of the structural model described in the text; instead, we restrict ourselves to a test of the matching–ranking process. In future work with Kevin Murphy, however, we do plan to examine the cyclical behavior of the number of workers who have spells of both U and N, as opposed to the number who have spells of just E and U.

42. See Blanchard and Diamond (1989b) for definition of variables.

ranking heavily depends on the exact nature of ranking. In a centralized market, the vacancies would go to primary workers first, and to secondary workers only when no unemployed remained. In a decentralized market, however, secondary workers may be the only applicants for a job, and therefore get the job even though there are unemployed primary workers elsewhere. While we have derived functional forms from first principles in another paper, here we look only at log-log specifications, which should be thought of as approximations.[43] The assumption that U is composed only of primary workers and N only of secondary workers is clearly an oversimplification. If the assumption is not too far from the truth, ranking has two main implications for the coefficients of such regressions: N, the pool of secondary workers, should not appear in H_U, but U should appear negatively in H_N. In contrast, the alternative assumption of lower search intensity of the secondary workers should lead to negative effects of U on H_N and negative effects of N on H_U.

Table 9 gives the estimation results, with and without imposing constant returns to scale. The period of estimation, February 1970 to December 1981, is determined by the availability of recall and "want-a-job" series. We see the effects of ranking in both sets of regressions. The evidence appears very consistent with ranking, although the amount of data construction involved and the implicit assumption that the pools of workers in U and N correspond exactly to the pools of primary and secondary workers prevent us from claiming too much.[44]

Extensions and Implications

Our model has provided a simple interpretation of the stylized aggregate and disaggregated facts about the flows of workers, using an approach with two types of workers, primary and secondary. In a recession, job destruction increases and both types experience layoffs. But because of ranking in hiring, nonemployed secondary workers experience a much larger decline in their chance of getting a job. Because of high quit rates and lower accession rates, the pool of employed secondary workers shrinks relative to the pool of primary workers.

43. Blanchard and Diamond (1990a).

44. The same approach could be used to look for more general patterns of ranking, or for differences in search intensity. To do so, stocks and flows could be decomposed by age and sex.

In many ways, however, this model is too simple. By equating nonemployed primary workers with unemployment, and nonemployed secondary workers with those not in the labor force, our model allows for a simple mapping between types of workers and observable pools. But as we know from the mean gross flows, there is considerable movement between unemployment and not in the labor force. Thus, an obvious extension is to assume that secondary workers move stochastically back and forth between the two nonemployment states. This stochastic movement can explain not only the level but also the cyclical behavior of the flows of workers between U and N—the larger the pool of nonemployed secondary workers, as is the case during a recession, the larger the flows between U and N. Also, as some of the unemployed are secondary workers, the difference between the cyclical behavior of the flows of primary and secondary workers into employment is more pronounced than the difference between the flows from U and N into E.

Our model also suffers in its sharp distinction between primary and secondary workers. Many secondary workers are young workers who eventually become primary workers. Thus, a recession, with lower movement of secondary workers into and out of jobs, is likely to slow this process of transition.

The model also has implications for wages.[45] It is clear that the assumption of a single type of job is far less satisfactory when thinking about wages than when thinking about flows. If wages are set either through worker–firm bargaining, or unilaterally by the firms based on efficiency considerations, hazard rates both into and out of employment are likely to be important. Under Nash bargaining, for example, the wage will divide the surplus from filling the job represented by the gap between the marginal product of the job and the value of time if not employed. How close the wage comes to the marginal product or the value of time will depend on how easily firms can replace workers and how easily workers can find another job. This is where hazard rates matter.

Suppose that the relative gap were the same for secondary and primary workers. Then the larger variation in labor market conditions for second-

45. As in our previous paper, we do not feel ready to go to the data, but we want to sketch what these implications may be. We took a stab at explaining wages empirically in Blanchard and Diamond (1989b). That stab convinced us to proceed more slowly.

ary workers would imply more relative variation in their wage. But the assumption that the gap is the same for both types is surely wrong. The higher quit rate of secondary workers indicates that their reservation wage is indeed close to their wage, or that the gap is small. Teenagers may value time not employed more than older workers, relative to their marginal products. Moreover, firms faced with two types of workers may well offer two types of jobs, with lower marginal products for secondary jobs.[46] If the gap is smaller, what gets divided is smaller and so there is little scope for wage variation for secondary workers. An offsetting effect follows from the fact that the hiring rate of primary workers includes their taking low-wage secondary jobs. This effect implies that primary wages will fall by more than if primary workers found primary jobs. Thus, while the stylized facts about flows suggest that secondary workers are more strongly affected by cyclical fluctuations, it does not follow that their wage will vary more.[47] Indeed it may be that the wage of primary workers does not move much because these workers have more job security, and the wage of secondary workers does not move much because the surplus to be divided is small. Consideration of average wages must also reflect both the drop in the fraction of the employed who are secondary workers and the extent to which primary workers take secondary jobs.

APPENDIX A

Identification

OUR TWO Brookings papers use identical approaches, though different notations.[48] The results from the papers differ slightly because here we use different time periods, levels of variables rather than logarithms, and different seasonal adjustment methods.

46. The existence of two types of jobs is an old theme. Much recent debate has focused on trends in the proportions of "good" versus "bad" jobs. Interactions between the two types of jobs and two types of workers raise additional issues. For example, at the end of a recession, many primary workers may find themselves in bad jobs, and will move to good jobs as the economy improves.

47. Other considerations, such as the presence of minimum wages, or constraints governing the relative wages of primary and secondary workers, may well be involved.

48. For more detailed discussion of the assumptions, see Blanchard and Diamond (1989a) and Yellen's accompanying comment.

From the VAR to the Structural Model

As indicated in the text, we assume the existence of an underlying structural model of the form:

$$\text{(A1)} \qquad \mathbf{X} = \mathbf{A(L)}\boldsymbol{\epsilon};$$

$$\mathbf{X} = [E\ U\ V]'; \ \boldsymbol{\epsilon} = [\epsilon_c\ \epsilon_s\ \epsilon_f]',$$

$$\mathbf{A(L)} = \mathbf{A}_0 + \mathbf{A}_1\mathbf{L} + \mathbf{A}_2\mathbf{L}^2 + \ldots; \ \mathcal{E}(\boldsymbol{\epsilon\epsilon}') = \mathbf{D}.$$

We assume that movements in employment, unemployment, and vacancies result from the dynamic effects of underlying, unobservable shocks. The three shocks are assumed to be aggregate activity shocks, reallocation shocks, and labor force shocks, denoted by ϵ_c, ϵ_s, and ϵ_f. The vector of white noise innovations to those shocks is $\boldsymbol{\epsilon}$. The matrix polynomial $\mathbf{A(L)}$ gives the dynamic effects of each of the innovations in the shocks on each of the three stock variables—employment, unemployment, and vacancies. The matrix $\mathbf{D}$ is the variance–covariance matrix of the innovations.

We estimate a VAR for $\mathbf{X}$, by OLS, with twelve lags on each of the three variables, a constant, and time, for the period January 1968 to May 1986. The estimated VAR can be written in moving average form:

$$\text{(A2)} \qquad \mathbf{X} = \mathbf{B(L)}\boldsymbol{\eta};$$

$$\boldsymbol{\eta} = [e\ u\ v]';$$

$$\mathbf{B(L)} = \mathbf{B}_0 + \mathbf{B}_1\mathbf{L} + \mathbf{B}_2\mathbf{L}^2 + \ldots; \ \mathcal{E}(\boldsymbol{\eta\eta}') = \boldsymbol{\Sigma}.$$

The vector $\boldsymbol{\eta}$ is the vector of reduced-form disturbances, or, put more simply, the vector of residuals to the three VAR equations. The constant $\mathbf{B}_0$ is the identity matrix.

Comparison of equations A1 and A2 shows that

$$\text{(A3)} \qquad \boldsymbol{\eta} = \mathbf{A}_0\boldsymbol{\epsilon}; \mathbf{A(L)} = \mathbf{B(L)}\mathbf{A}_0.$$

The residuals of the VAR are linear functions of the underlying innovations to the shocks. The linear relation is given by the matrix $\mathbf{A}_0$. The matrix lag polynomial of the structural model is equal to the matrix lag polynomial of the moving-average representation of the VAR, postmultiplied by $\mathbf{A}_0$. Thus, to go from the VAR to the structural model, one needs to obtain the matrix $\mathbf{A}_0$.

The relation between $\boldsymbol{\eta}$ and $\boldsymbol{\epsilon}$ imposes the following restrictions on $\mathbf{A}_0$, $\mathbf{D}$, and $\boldsymbol{\Sigma}$,

$$\boldsymbol{\Sigma} = \mathbf{A}_0\mathbf{D}\mathbf{A}_0'. \tag{A4}$$

An estimate of $\boldsymbol{\Sigma}$ can be obtained from the data, but neither $\mathbf{A}_0$ nor $\mathbf{D}$ is observable. Further identification assumptions are needed to obtain $\mathbf{A}_0$.

Identification of $\mathbf{A}_0$

Write the relation between $\boldsymbol{\eta}$ and $\boldsymbol{\epsilon}$ fully out as

$$\begin{aligned} e &= a_{11}\epsilon_c + a_{12}\epsilon_s + a_{13}\epsilon_f \,, \\ u &= a_{21}\epsilon_c + a_{22}\epsilon_s + a_{23}\epsilon_f \,, \\ v &= a_{31}\epsilon_c + a_{32}\epsilon_s + a_{33}\epsilon_f \,. \end{aligned} \tag{A5}$$

To identify the nine elements of $\mathbf{A}_0$, $((a_{ij})$ for $i, j = 1, 2, 3)$, we make several assumptions. First, we normalize $\mathbf{A}_0$ so that an innovation to the cyclical shock has an effect of 1 on vacancies, that an innovation to the reallocation shock has an effect of -1 on employment, and finally that a labor force innovation increases the labor force by 1. The normalizations imply $a_{31} = 1$, $a_{12} = -1$, and $a_{13} + a_{23} = 1$. Second, we assume the three structural innovations are contemporaneously uncorrelated, which implies that $\mathbf{D}$ is a diagonal matrix.

Counting equations and unknowns, note that in equation A4, $\boldsymbol{\Sigma}$ has six distinct moments. Under the assumptions we have made so far, the matrix $\mathbf{A}_0$ has six unknown coefficients, and the matrix $\mathbf{D}$ has three unknown diagonal elements. We therefore need three further restrictions for identification.

One restriction is that the effects of employment on labor force participation rates are the same whether the movements are caused by cyclical or reallocation shocks. This implies that $a_{11}/a_{21} = a_{12}/a_{22}$. Another restriction is that, in the current month, an innovation to the labor force, ϵ_f, shows up equally in employment and unemployment. That is, of those people who decide to join the labor force for reasons unrelated to changes in labor market conditions, half of them do so by going directly into employment, the other half by going into unemployment. This implies that $a_{13} = 0.5$, and thus, given the normalization, $a_{23} = 0.5$.

These two restrictions narrow underidentification to the pair a_{11} and a_{32}, which are the effects of a cyclical innovation on employment and of a reallocation innovation on vacancies respectively. Given one, the other is identified. Our restriction here is based on simple theoretical considerations. For some period of time, aggregate activity shocks should move unemployment and vacancies in opposite directions, and reallocation shocks should move unemployment and vacancies in the same direction. Thus, we examine values of a_{11}; for each value, we identify $\mathbf{A}_0$ and trace the effects of each type of innovation. Then we search for values such that the signs of the effects of the two innovations on unemployment and vacancies are as predicted for the first ten months. This gives us a range of values for a_{11} between 1.6 and 2.6. We choose a value of 2.2. Then, we can use the method of moments to solve for equation A4 and obtain $\mathbf{A}_0$. Equation A5 can then be written,

$$e = 2.2\epsilon_c - 1.0\epsilon_s + 0.50\epsilon_f\ ,$$

$$u = -0.4(2.2\epsilon_c - 1.0\epsilon_s) + 0.50\epsilon_f\ ,$$

$$v = \epsilon_c + 0.08\epsilon_s - 0.05\epsilon_f\ .$$

Given $\mathbf{A}_0$, we can trace the effects of innovations to all three structural shocks on E, U, and V. In our previous paper, we concentrated on the dynamic effects of all three. Here, we focus on the effects of ϵ_c, which dominate medium-term movements in E, U, and V.

Identification assumptions are by definition controversial. We have therefore examined the robustness of our results with respect to variations in those assumptions. We have considered values of a_{13} between 0.4 and 0.6, and values of a_{11} between 1.6 and 2.6. While these variations affect the dynamic responses to ϵ_s and ϵ_c, they leave the responses of the stocks to ϵ_c nearly unaffected. The flavor of the results is given in table A1, which replicates the results in figure 5 for alternative identification assumptions.

Table A1. Response of Flows into and out of Employment to a Negative Shock[a]
Thousands of workers

Experiment	Months after shock	Total flows: out of E	Total flows: into E
Benchmark case			
($a_{11}=2.2$, $a_{13}=0.5$)	1	55	−25
	6	140	−96
Varying a_{11}			
($a_{11}=1.6$, $a_{13}=0.5$)	1	47	−15
	6	135	−91
($a_{11}=2.6$, $a_{13}=0.5$)	1	59	−30
	6	142	−97
Varying a_{13}			
($a_{11}=2.2$, $a_{13}=0.6$)	1	55	−25
	6	141	−96
($a_{11}=2.2$, $a_{13}=0.4$)	1	55	−24
	6	138	−94

Source: Authors' own calculations.
a. The table shows the estimated response of cumulative flows under alternative identification assumptions to a negative one-standard-deviation innovation to ϵ_c.

APPENDIX B

Relations among the Three Data Sets

IF WE DEFINE the flow into employment as the number of workers who go from nonemployment to employment within a month, and the flow out of employment as the number of workers who go from employment to nonemployment, how do the flows from the CPS and the manufacturing survey relate to these true series?

The CPS data are point-in-time data, which record the status of workers at two points in time. Therefore, the data *exclude* sequences of changes between survey dates that do not result in a change in status, such as movement from employment to unemployment to employment. They also *misclassify* sequences of changes between survey dates that show up as a change in status, such as a movement from employment to unemployment to not in the labor force. The short duration of unemployment and its variation over the cycle imply that the exclusions and misclassifications may be substantial and cyclical. We do not correct for them.

The manufacturing flow data, collected from firms, are integrals of the flows of separations and accessions. They too differ from the true series: they refer only to the manufacturing sector; they *include* in the flow out of employment those workers who go directly to another job, and therefore remain in employment; and they *include* in the flow into employment those workers who come directly from another job, and therefore remain in employment. Thus, we attempt to adjust the flows into and out of employment for manufacturing-employment-to-employment quits, and employment-to-manufacturing-employment quits respectively.

If we define the flow of job creation as the sum of jobs that did not previously exist and for which the firm needs an additional worker, and if we define the flow of job destruction as the sum of jobs that previously existed and for which the firm no longer needs a worker, then how do the manufacturing and the Davis-Haltiwanger series relate to those true series?

The two series differ in several ways. The flow of separations *includes* separations that are replaced, and thus are not associated with job destruction. Also, the flow of accessions *includes* accessions that are replacements of those workers who have left and need to be replaced, and are therefore not associated with job creation. These replacements are thus a distributed lag of the previous flow. The flow of accessions also *includes* accessions that fill previously created vacancies. Further, the flow of accessions *excludes* those job creations for which a worker has not yet been hired. The flow of separations *excludes* those job destructions that have not yet led to worker separations, due for example to notice periods or firing costs. It also *excludes*, although this is likely to be small, those vacancies that are canceled before they are filled.

We attempt to correct for the first two facts by estimating the number of nonreplaced quits (see text and Appendix D) and removing them from separations and accessions, ignoring the distributed lag effect for accessions. We adjust for the third and fourth by adding the change in vacancies (scaled for manufacturing) to the flow of accessions. We do not adjust for the fifth.

The Davis-Haltiwanger data are first differences of point-in-time employment stock data for establishments. They define job creation as the sum of changes in employment over all firms with positive changes in employment between the two quarterly survey dates. Job destruction is defined as the sum of changes in employment over all firms with negative changes in employment. Therefore, these data do not *include* job creation and destruction within establishments, only the net amount. Nor do they *include* sequences of job creation and destruction that do not show up as net changes between survey dates. Like manufacturing data, the data *exclude* those changes in job creation and job destruction that have not yet led to either firing or hiring of workers. We do not correct for the first two. We correct for the third, by adding the change in vacancies (scaled for manufacturing) to job creation.

APPENDIX C

Construction of EE Time Series

EVERY YEAR, BLS carries out a retrospective survey of individuals in the CPS. As a result, we can characterize individuals by three numbers, x, y, and z. The x is a 0–1 variable—equal to 1 if the individual was not in the labor force at any point during the year, and 0 otherwise. The y runs from 0 to 3, and stands for the number of employers—where 3 stands for three employers or more. The z runs from 0 to 3 and stands for stretches of unemployment—where 3 stands for three stretches or more. The number of stretches of unemployment for those who did not work during the year and were not in the labor force at least once is not known; they are classified separately.

Kevin Murphy has provided us with tabulations of males in each category for the period 1975–87. From those tabulations, we construct three series that we consider lower bounds (EE_{small}), upper bounds (EE_{big}), and best guesses (EE_{med}) on EE movements. The series are constructed as

$$EE_{\text{small}} = X_{020} + 2X_{030} + X_{031} + X_{130} ,$$

$$\begin{aligned} EE_{\text{big}} = {} & X_{020} + X_{021} + X_{022} + X_{023} + 2(X_{030} + X_{031} \\ & + X_{032} + X_{033}) + X_{120} + X_{121} + X_{122} + X_{123} \\ & + 2(X_{130} + X_{131} + X_{132} + X_{133}), \end{aligned}$$

$$\begin{aligned} EE_{\text{med}} = {} & X_{020} + 0.5X_{021} + 2X_{030} + 1.5X_{031} + 0.5X_{120} \\ & + 1.5X_{130} + X_{131} . \end{aligned}$$

Figure C1. Manufacturing Quits versus Employment-to-Employment Movements, 1975–85

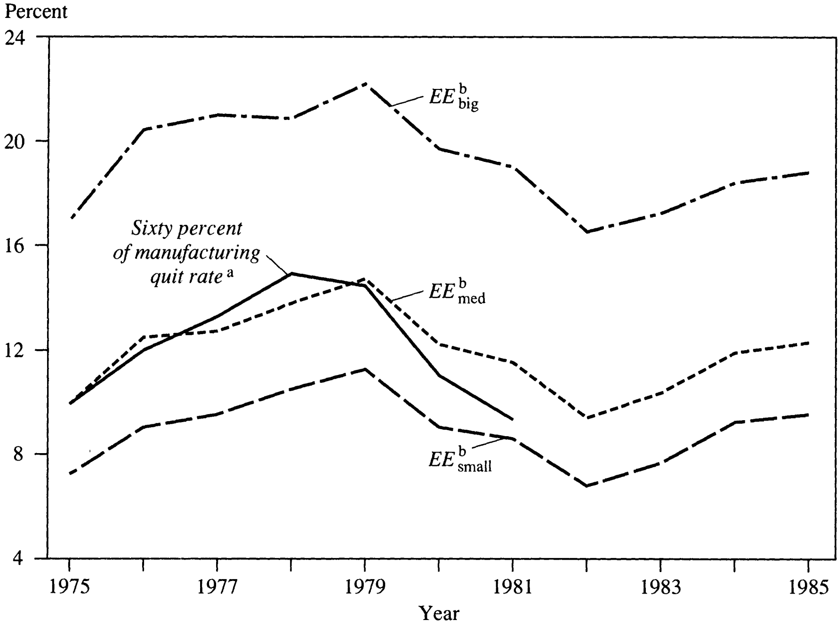

Source: Current Population Survey and authors' own calculations using data provided by Kevin Murphy.

a. The manufacturing quit rate, taken as the sum of monthly quit rates, is multiplied by a constant of proportionality of 0.6.

b. EE_{small}, EE_{big}, and EE_{med} refer to series constructed by the authors giving lower bounds, upper bounds, and best guesses on employment-to-employment movements respectively. See the text for equations using individuals' employment status, number of employers, and stretches of unemployment to construct these series. Each series is divided by total employment (E) to give a quit rate.

where X_{xyz} is the proportion of individuals in each category. The three series, each divided by mean employment for the year, are plotted in figure C1 for the period 1975–85. We also plot in the figure a series equal to 60 percent of the quit rate in manufacturing, taken as the sum of the monthly quit rates. This series is available only until 1981. The figure suggests that constructed *EE* quits move closely with total quits and that the constant of proportionality is around 0.6.

The assumption that *EE* quits are a constant proportion of quits thus appears empirically reasonable. We also think it is theoretically reasonable. There are three types of quits. The first, called an *EM* quit, is a quit to a generalized alternative because the current position does not seem attractive any longer. These quits decline in a recession since the value of being in *M* goes down. The second, called an *EE* quit, is a quit because a better job has been found. These also decline in a recession because there are fewer alternative jobs and they are more actively pursued by others. The third, and probably less quantitatively important, type is a quit in anticipation of a layoff (see Appendix D for an estimate). Many people receive advance notice of a layoff, which gives them time to search for a new job and the possibility of finding one before the layoff occurs. While there are more layoffs in a recession, the likelihood of finding a job is down. Thus it is not clear how these anticipatory quits behave over the cycle.

APPENDIX D

Proportion of Quits That Are Not Replaced

THERE ARE two types of nonreplaced quits. First, when a worker leaves a position, firms may eliminate that job, something they would not have done otherwise. Second, a worker may quit because he learns that his position will be terminated and successfully locates another job. We have no information on the first type. We have some rough information on the second, from a BLS survey on jobseeking methods used by workers in 1972.[49]

Let X be the pool of workers whose jobs are to be terminated. Let a be the fraction of those who search for another job before termination and b be the fraction of those searchers who are successful. Thus, abX workers quit to another job in anticipation of job termination, and $(1 - ab)X$ stay until job termination. Let Q be the total number of quits (quits in anticipation of termination and for other reasons) and S the number of separations by job termination, so that $S = (1 - ab)X$. Then the ratio of quits in anticipation of job termination to total quits—that is, the proportion of nonreplaced quits—is equal to $[ab/(1 - ab)](S/Q)$.

Table 7 in the BLS survey gives the proportion of those who began their search while employed, disaggregated by reason for separation. It suggests a value for a of 0.20–0.25. Table G6 in the BLS survey gives a cross tabulation of the number of weeks of on-the-job search and the

49. Bureau of Labor Statistics (1975).

number of weeks of total search. The cross tabulation is too coarse to give an exact number for b, but suggests a range of 0.60–0.70. Table 7 also gives a ratio for S/Q of $(2304/3087) = 0.75$. These values imply a range of 9 to 13 percent for the proportion of nonreplaced quits. However, this is a lower bound, as it only captures the second type of nonreplaced quits.

Comments and Discussion

Robert E. Hall: I see this paper as part of the resurgence of interest in labor market dynamics that has taken place in the past few years, a resurgence in which Peter Diamond and Olivier Blanchard have been leaders. Their work has focused primarily on the worker's point of view. In particular, they have been looking at the flows recorded in the CPS, which I think is an important enterprise.

An attractive feature of this paper is its comparison of the gross flows with other ways of looking at the same issues, particularly research by Steve Davis and John Haltiwanger, which is the other major prong of the resurgent interest in labor market dynamics. Davis and Haltiwanger focus on the establishment point of view. A reconciliation of the two perspectives and a demonstration that the two views reflect the same underlying features of the labor market is one of the most important contributions of this paper. This paper increased my understanding of the literature on the worker's point of view and the establishment-based research.

The model of labor market flows has much in common with Blanchard and Diamond's earlier work. I especially like the clever ideas used to identify the model within a Beveridge curve framework, as I did in 1989. Because of the limited space, however, I say little else about the model and instead concentrate on the comparison of establishment and worker data.

This paper responds to the problem in the CPS flow data raised by the dominance of the short end of the duration distribution for jobs. High turnover and low job duration dominate gross movements. The median duration of a job in the United States is one day—more than half of each day's job placements that occur are for day work.

You can alleviate that problem if you look from one month to the next. Day work becomes much less important, but short-duration employment still dominates flows. There is an awful lot of one-month employment, which doesn't contribute much to total labor input but does add an enormous amount of turnover. This is one of the big obstacles to doing research with gross flows data, and, in particular, is what led me to look at tenure data as a way of getting around the problem. One point I would make is that there is another body of data not mentioned here, namely, the distribution of job tenure. I don't think the tenure data contradicts the CPS in any important way, but such a confirmation could be valuable.

As I said before, I find the comparison with the Davis-Haltiwanger research especially interesting and valuable and want to focus on that along with the theme that recessions are times of higher job destruction more so than times of lower job creation.

The striking finding of this paper is the small response of the *UE* hazard rate in recessions. To put it differently, unemployment does not seem to last much longer in recessions. We tend to think of recessions as times when it is harder to find work, and yet the job-finding rate, the hazard rate from *U* to *E*, is relatively insensitive to the cycle. Figure 5 shows the model's inference about the response to a negative cyclical shock. At all horizons, the flow out of employment is the dominant source of the decline in employment. The job-finding rate is much less important.

Figure 6 reports the raw data. The flow into employment is the solid line and the flow out of employment is the broken line. The 1973–75 recession is typical; it is the strongest generic example of these results. The rise in the broken line is much larger than the decline in the solid line. In 1981–82, which the authors point out dominates the data, the evidence is even sharper. The raw data leave little doubt about their finding.

Figure 7 shows related data for manufacturing. Although the responses are essentially symmetrical, even when quits are accounted for, the variations in the flows out of employment still dominate employment changes. Davis and Haltiwanger have done a highly complementary study of employment changes by firm. They look at the shift in the cross-sectional distribution of employment change at the establishment level that occurs in a recession.

They compile the probability distribution of the size of the employment change. A horizontal axis is the amount of the employment change and a vertical axis is the density. They examine the change in the distribution between a recession period and a neutral or boom period. With respect to positive changes, the distribution looks just the same. The distribution for small negative changes is also mostly unchanged. The important changes in the cross-sectional distribution of employment changes that occur between a recession year and a normal year are in the left tail of the distribution. There is an increased incidence of large negative employment changes that characterize a recession.

By and large, the labor market functions normally in a recession period; for most firms, it is business as usual. The fraction of firms that actually hire more is the same in a recession *and* in a boom. This is a very remarkable finding.

The story that emerges from all the empirical work, both the authors' and Davis and Haltiwanger's, is that, for the great majority of establishments, employment growth proceeds normally in a recession. It is only a few establishments that contract sharply. Blanchard and Diamond observe that this finding contradicts the standard view of entry and exit, where sunk costs would make exit insensitive to economic conditions. In the standard view, all the action occurs on entry. The employment-growth side of the Davis-Haltiwanger distribution would vary, not the employment-contraction side. The sunk cost entry-exit model, which is the model that most students start with, is wrong.

Blanchard and Diamond point to an alternative model that would explain the left side of the cross-sectional distribution. In that model, there are established units, embodying old technology, where the capital is approaching a market value of zero and so the units are about to be shut down. Shutdown occurs when a recession strikes.

Here the authors address the ideas that are ready to take off in this literature. Valerie Ramey has described one idea as the economics of pitstops. She observes that in an auto race, when a driver sees a yellow flag—meaning everyone has to go slow—that is the time to take a pitstop. If one thinks of pitstops as corresponding to those times when the economy slows down and workers are laid off, then one would expect that the actions of firms—to shut down and discharge their workers—would cluster at particular times. These times are called recessions, and can be thought of as yellow flags. Russ Cooper and John Haltiwanger's paper on replacement led Ramey to make this analogy.

Pitstop theory is likely to develop in different versions. There is the Davis-Haltiwanger analytical model already cited. A recession is a time of technical regress. It is also a time when a shutdown could occur, which could explain the shift in the cross-sectional distribution.

Also, nonconvexity of some aspect of the technology, perhaps layoff costs, will be crucial to an ultimate explanation of this finding. Whenever there is bunching, nonconvexity comes to mind. But, as Blanchard and Diamond point out, there is a Caplin-Spulberg problem here: supersensitive units in the economy that get pushed past a certain point will distribute themselves randomly with respect to that crossover point; over time they will not concentrate in the aggregate, even though their firm-level behavior is concentrated. Giuseppe Bertola and Ricardo Caballero have shown that some effects of nonconvexity survive aggregation. I think that pitstop economics is going to be an important area of further research into the nature of recession.

Another promising branch of theory is the economics of thick-and-thin markets, an area pioneered by Peter Diamond. A relevant version of this model explains that some pieces of the economy move from their normal high-level thick-market equilibrium to a thin-market equilibrium. These pieces are scattered, both geographically and across the product space, but they shut down as they move to their thin-market equilibrium in unison. When they do, that time is a recession.

The last idea concerns contagion in employment policies. Economists have agonized for years over the fact that contagion may be a factor in the stock market. Bob Shiller's results, I think, show that the reason everyone tried to sell their stocks in October 1987 was that everyone else was doing it. The same factor may apply in many other arenas of economic behavior, one of which may be that when other businesses are clearing house, shutting down, or dramatically reducing employment, others do too. There is contagion in employment policy, just as there is contagion in portfolio management.

Blanchard and Diamond's paper will be seen as an important step in the development of a new kind of fluctuations theory. Their unification of CPS data and establishment findings is a particularly significant contribution.

Kevin Murphy: I would like to begin by saying that Olivier Blanchard and Peter Diamond should be praised for drawing together data from a variety of sources and using these data to identify some key aspects of the employment–unemployment picture. Their use of gross flows data along with other data from the CPS and the Davis-Haltiwanger manufacturing data represents a step forward in the effort to understand labor market dynamics from a broad empirical perspective. In general I found their approach relatively attractive. My comments on this paper are first about the data and second about their possible interpretation.

First, the data used in this paper come from a survey (the Current Population Survey); as a result they are subject to sampling error. Even though you may have, say, 100,000 people in the basic CPS sample, the authors are looking at relatively rare events (such as an employment-to-unemployment transition) that happen to less than 1 percent of the population in a typical month. With a sample of 100,000 individuals, the fact that they look at transitions across months cuts the sample to 75,000, since one quarter of the sample is not interviewed the next month. This implies that for an employment-to-unemployment transition one might expect to see about 500 transitions. Because this is simply a binomial with a very low probability, we can use the Poisson approximation and say that the variance of this number is also about 500. This implies a standard deviation of about 25 so that getting either 450 or 550 transitions could occur because of sampling error. This would then imply a standard error from month to month of about 35 (=(1.44)(25)), which is about 7 percent of the level. By this reasoning, a change in the gross flows of even 15 percent from one month to the next could simply be the result of normal sampling error. Undoubtedly this must account for much of what Blanchard and Diamond refer to as "high-frequency" movements in the gross flows data.

Second, Blanchard and Diamond use both "adjusted" and "unadjusted" gross flows data, with the adjusted data coming from the work of John Abowd and Arnold Zellner. The need for the adjusted series arises because the gross flows data are derived from point-in-time questions for the same individuals in two consecutive months. Hence, while classification error may represent an unimportant component of the levels of employment, unemployment, and not in the labor force, it will be very important in determining the number of individuals classified in different states in the two months (since most individuals do not move between states in a typical month). While this measurement error means that one cannot use the levels of these transition rates without adjustment, it is less clear to what extent it biases estimates of changes in gross flows through time. In fact, in the one specification test in the paper, where Blanchard and Diamond compare the changes in stocks of employed, unemployed, and not in the labor force with the accumulated changes predicted by the gross flows data, the unadjusted flow series appears to do a better job of approximating the change in the actual stocks than does the adjusted series. Hence I am pleased that they present both sets of numbers for the majority of their calculations.

A third point is that at a more basic level the compilers of the CPS in essence do some recoding of their own when they generate the employment status variable used to generate the gross flows data. People are asked to list their major activity and if they respond that they are "in school" or "keeping house" they are then asked if they did any work or did anything to find work. In fact many of the people listed as unemployed do not respond that they are unemployed when asked their major activity. As with the debate about adjusted versus unadjusted data, the question here is, does this just affect the level or is it also important for the cyclical and secular changes? I don't really know the answer to this question, but since both the standard employment question and the major activity question are reported in the data, one could ask how the stories told by the two measures differ. In particular, when unemployment rates rise in a recession, are more people likely to be recoded into unemployment from not in the labor force? Again, this may be important for measuring the details of what happens over the cycle as Blanchard and Diamond attempt to do.

One final data point has to do with measurement. In general Blanchard and Diamond look at gross flows—for example, the number of people that move from unemployment to employment in a given month. Alternatively, one can express flows using hazard rates—the number of individuals that change from, say, unemployment to employment as a percentage of the number of unemployed persons. In addition, one can look at either the arithmetic value or the log of either the flow or the hazard rate. For employment, where the base for the hazard rates moves very slowly through time, the distinction between these measures is trivial. However, the distinctions for unemployment are more crucial.

For example, while Blanchard and Diamond point out that the absolute size of the flow from unemployment to employment rises in a recession, the hazard rate (the fraction of people that become employed) actually falls. As Blanchard and Diamond point out, the primary reason why the flows out of unemployment rise during a recession is that more people are unemployed and hence more people exit unemployment.

In addition to their basic results for the levels of the gross flows, Blanchard and Diamond also present calculations of the implied hazard rates. In my opinion these calculations are somewhat less than satisfactory, given that their model is estimated on the gross levels. Because the model is estimated to fit changes in the levels of the gross flows data, it is somewhat unclear how well the model fits these hazard rates and whether the hazard rates implied by their model correspond to actual features of the empirical hazards or are sensitive to the indirect estimation strategy employed. This brings me to the log versus levels specification issue. If they were to specify the empirical model in terms of the logs of the flows rather than the arithmetic levels of the flows, it would be straightforward to move from the flow data to the hazard rate data since the flows in logs are simply the log of the hazard plus the log of the base population. Personally I would have preferred they specify the model in logs rather than levels since this would have allowed us to infer the hazard rate results directly and examine how well the model can explain changes in the empirical hazard rates. It should be noted, however, that once the model is changed to a log specification, the aggregation across unemployment and not in the labor force required to move from their two-state model to their three-state model becomes much more complicated. However, I still prefer the log specification and would like to see the estimates for the resulting hazards.

I now turn to the interpretation of the authors' results. In some sense we already knew much of the message of this paper. Business cycle fluctuations are caused by increases in the number of individuals entering unemployment more than they are by decreases in the number of individuals that leave unemployment. Many papers in the labor literature have pointed out that during a recession the entry rate into unemployment (that is, the exit rate from employment) rises sharply while the hazard rates for leaving unemployment fall modestly (hence durations rise somewhat). Thus, in the labor literature, the business cycle is much more incidence than it is duration. For gross flows, this implies that gross exits from unemployment will rise as long as the exit hazard rate falls less (in percentage terms) than the unemployment rate rises. Given the sharp rise in entry rates and the relatively modest fall in exit rates from unemployment, it is not surprising that gross entry (job destruction) rises much more than gross exits (job creation) fall.

Given that the findings of Blanchard and Diamond confirm what I would contend is conventional wisdom, what is added by their observations? I believe that there are at least two major contributions. First, their analysis recasts the results from the individual level, which has been the focus of all of the previous literature, to the implications these data have for the aggregate picture. While the aggregate fluctuations generated follow directly from the individual data, it is nevertheless important to draw out these implications. In this way Blanchard and Diamond translate the results on entry and exit into results about changes in the rates at which employee–employer matches are being created and destroyed. Second, all of what I have said to this point refers to the flows of individuals between unemployment, employment, and not in the labor force. Blanchard and Diamond extend these results to statements about jobs by showing that the individual-based evidence corresponds closely with the job-based evidence produced by Davis and Haltiwanger. At both the individual level and the job level, fluctuations in the breakup of employment matches are far more important than fluctuations in the rate of creation of employment matches in generating cyclical fluctuations in employment and unemployment.

My final comments have to do with the difference between the cyclical and secular changes in unemployment. In their analysis Blanchard and Diamond use a VAR together with some explicit normalizations to identify what they term the cyclical component of unemployment (shown in figure 4). In my work it is becoming increasingly clear that it is important to distinguish this level of unemployment from the trend increase in unemployment that has occurred over the period of their data. In particular there are at least two key differences. First, the trend increase in unemployment is due much more to a decrease in exit rates from nonemployment (a rise in durations) than the cyclical fluctuations. Secondly, the trend increases in unemployment are concentrated among low-wage workers to a much greater extent than are cyclical increases in unemployment. For example, in joint work with Chinhui Juhn and Robert Topel we find that the bottom decile of the male wage distribution accounts for about 30 percent of the trend rise in unemployment but only about 20 percent of cyclical unemployment. Hence in trying to understand the long-term rise in unemployment we must largely explain why low-wage men are working much less, while over the cycle we must also ask why a much broader spectrum of the population spends more time unemployed. In interpreting Blanchard and Diamond's results, the reader must keep in mind that they are talking about what happens over the cycle and not about why unemployment is higher in the late 1980s than during the late 1960s.

General Discussion

The panelists made a number of observations about the characterization of recessions as times of "cleaning up." Christopher Sims suggested that a putty-clay model of capital goods could explain the larger fluctuations in job destruction than in job creation. He reasoned that the asymmetry could be explained by having either a greater substitutability between labor and newer capital equipment or a decreasing returns technology for production of new capital goods. William Brainard noted, however, that this result relied upon the inability to restart old equipment. In standard models, fluctuations in employment involve fluctuations in the utilization of the oldest capital. Kevin Murphy suggested that to understand this job creation–job destruction pattern one ought to explain why firms that are doing badly are more cyclically sensitive than firms doing well. John Haltiwanger warned against a vintage-model explanation since his research with Steve Davis suggests that only 15 percent of the time-series variance in gross job reallocation could be explained by allowing differential responses across age-groups of firms. He also noted that while the rates of job creation and destruction decline sharply with plant age, 65 percent of the work force is employed in plants fifteen or more years old, so most of the job creation and destruction comes from older plants.

There were two remarks on how the data did not accord well with a "cleaning up" or "pitstop" model of recessions. Martin Baily observed that productivity fell during a recession, contrary to what would be expected if managers were streamlining their production process. Matthew Shapiro argued that lack of a spike up in employment after a recession hurts the pitstop metaphor since after a pitstop people get going right away.

George Perry pointed out that during recessions many people switch from long-duration jobs to temporary jobs, resulting in a greater proportion of the work force in short-duration jobs. If the amount of job creation and destruction is relatively constant in the temporary jobs, then the destruction is taking place in the long-duration jobs. This view provides a harsher picture of what happens during a recession than one would get if the change in job composition were ignored. Murphy observed, however, that the average number of unemployment spells does not change much over the business cycle. George Akerlof felt that the large magnitude of gross job flows (the rate of the flow out of manufacturing is on the order of 5 percent per month) might mask the effect that Perry is describing. Haltiwanger reported some information on the persistence of job destruction: on average 80 percent of jobs destroyed at a given plant have not been restored at that plant two years after they were originally destroyed. This, along with similar findings for job creation, suggests that the time series of gross job creation and destruction reported in the paper primarily reflect permanent reallocation of jobs. Perry observed that the differences the authors attribute to gender might instead be due to industrial structure since the industrial sector of the economy has predominantly male employees and the service sector predominantly female.

References

Abowd, John M., and Arnold Zellner. 1985. "Estimating Gross Labor Force Flows." *Journal of Business and Economic Statistics* 3: 254–83.

Akerlof, George A., Andrew K. Rose, and Janet L. Yellen. 1988. "Job Switching and Job Satisfaction in the U.S. Labor Market." *BPEA, 2:1988*, 495–582.

Bertola, Giuseppe, and Ricardo Caballero. 1990. "Macroeconomic Implications of Kinked Adjustment Costs." *NBER Macroeconomics Annual* 5: 237–87.

Blanchard, Olivier Jean, and Peter Diamond. 1989a. "The Beveridge Curve." *BPEA, 1:1989*, 1–60.

———. 1989b. "The Aggregate Matching Function." Working Paper 3175. Cambridge, Mass.: National Bureau of Economic Research (November). Forthcoming in *Growth-Productivity-Unemployment*, edited by Peter Diamond. Cambridge, Mass.: MIT Press.

———. 1990a. "Ranking, Unemployment Duration, and Wages." MIT Working Paper 546.

———. 1990b. "Ranking, Unemployment Duration, and Training Costs." Massachusetts Institute of Technology (January).

Brown, Charles, James Hamilton, and James Medoff. 1990. *Employers Large and Small.* Cambridge, Mass.: Harvard University Press.

Burda, Michael, and Charles Wyplosz. 1990. "Gross Labour Market Flows in Europe: Some Stylized Facts." INSEAD working paper (June).

Bureau of Labor Statistics. 1975. *Jobseeking Methods Used by American Workers.* Bulletin 1886. Washington: Department of Labor.

———. 1977. *Employment in Perspective: A Cyclical Analysis of Gross Flows in the Labor Force.* Washington: Department of Labor.

———. 1982. *Gross Flow Data from the Current Population Survey, 1970–1980.* Washington: National Technical Information Service.

Caballero, Ricardo. 1990. "A Fallacy of Composition." Columbia University (April).

Clark, Kim B., and Lawrence H. Summers. 1978. "Labor Force Transitions and Unemployment." Working Paper 277. Cambridge, Mass.: National Bureau of Economic Research (April).

———. 1979. "Labor Market Dynamics and Unemployment: A Reconsideration." *BPEA, 1:1979*, 13–60.

Cooper, Russell, and John C. Haltiwanger. 1990. "The Macroeconomic Implications of Machine Replacement: Theory and Evidence." Boston University.

Darby, Michael R., John C. Haltiwanger, and Mark W. Plant. 1985. "Unemployment Rate Dynamics and Persistent Unemployment under Rational Expectations." *American Economic Review* 75: 614–37.

Davis, Steve, and John C. Haltiwanger. 1989. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation." Working paper, University of Chicago and University of Maryland.

———. 1990. "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications." *NBER Macroeconomics Annual* 5: 123–68.

Hall, Robert E. 1970. "Why Is the Unemployment Rate So High at Full Employment?" *BPEA, 3:1970*, 369–402.

Lasky, Mark J. 1990. "Do Low Profits Stimulate Productivity Growth?" Chapter one, Ph.D. dissertation, Massachusetts Institute of Technology.

Lilien, David M. 1980. "The Cyclical Pattern of Temporary Layoffs in United States Manufacturing." *Review of Economics and Statistics* 62: 24–31.

Perry, George L. 1972. "Unemployment Flows in the U.S. Labor Market." *BPEA, 2:1972*, 245–78.

Poterba, James M., and Lawrence H. Summers. 1986. "Reporting Errors and Labor Market Dynamics." *Econometrica* 54: 1319–38.

Shleifer, Andrei. 1986. "Implementation Cycles." *Journal of Political Economy* 94: 1163–90.

Smith, Ralph E., Jean E. Vanski, and Charles C. Holt. 1974. "Recession and the Employment of Demographic Groups." *BPEA, 3:1974*, 737–58.