New evidence on labor market dynamics over the business cycle

Bhashkar Mazumder

Introduction and summary

Does unemployment rise in a recession mainly because workers lose their jobs at a higher rate or because already unemployed workers are less likely to be hired during a downturn? The answer to this question has important implications for how one thinks about cyclical fluctuations in the economy and policies to address unemployment. For example, one prominent view of the business cycle posits that economic downturns are periods where there has been an adverse shock to productivity that makes the match between employers and workers less profitable. This, in turn, leads firms to increase layoffs. Under this view, we would expect to see much greater cyclicality in the rate of job separation (movements from employment to unemployment) compared with the rate of job finding (movements from unemployment to employment). An alternative view is that there might be reasons why firms prefer to create vacancies during economic upswings, in which case we would expect to see more variability in the job hiring rate over the business cycle. Clearly, documenting the cyclical patterns in job finding and job separation ought to provide important empirical evidence to help distinguish between competing views of unemployment fluctuations and perhaps help guide the development of new theories of the business cycle.

I attempt to shed light on this issue by using a new data source, the U.S. Census Bureau's *Survey of Income and Program Participation* (SIPP), to investigate how much of the cyclicality in unemployment is due to variation in the job finding rate versus the job separation rate. I estimate the job finding rate (or the hiring rate) by calculating the fraction of workers who transition from unemployment to employment in adjacent months. Similarly, the job separation rate is estimated by calculating the fraction of workers who transition from employment to unemployment in adjacent months. The SIPP is particularly well suited for this analysis because it tracks the same individuals over time and asks about their labor market activities during each week of the month. In addition, the SIPP collects information on more than one employer for each worker, enabling the calculation of a long time series on the rate of job-to-job transitions. Job-to-job movements over the business cycle will vary depending on the cyclical properties of the job separation rate and the hiring rate, so new descriptive data on this phenomenon should also help inform our understanding of unemployment dynamics.

The results of this analysis suggest that the hiring rate is highly procyclical and is an important component of the variation in the unemployment rate in recent business cycles. This is in accordance with the findings of some other recent studies in the literature that have utilized other data sources. However, in contrast with these other studies, I also find that there has been a notable degree of countercyclicality in the job separation rate during the two most recent recessions (in 1990–91 and 2001), which has also contributed to the cyclical patterns in unemployment. In other words, the rate of job separation did in fact rise during these recessions. One original contribution of my work is direct evidence that the job-to-job transition rate has been procyclical over the past 20 years. Overall, these results suggest that greater attention should be given to macroeconomic models that emphasize movements in hiring rates as important contributing factors to business cycle fluctuations.

Bhashkar Mazumder is a senior economist in the Economic Research Department and the executive director of the Chicago Research Data Center at the Federal Reserve Bank of Chicago. *The author thanks Phil Doctor for excellent research assistance and Gadi Barlevy, Jeff Campbell, and Eric French for helpful comments. He also thanks seminar participants at the Federal Reserve Bank of Chicago for their comments.*

In addition to exploring the broad macroeconomic patterns, I also utilize the SIPP to assess the extent to which employment dynamics have differed by sex and education level. Exploring heterogeneity in employment dynamics may be useful in understanding whether the cyclical patterns are driven in part by compositional changes. I find that while the transition rates between the labor market states (employment and unemployment) differ by education level, the *cyclical* patterns are actually not very different. Job-to-job transition rates are similar for most groups.

In the following section, I describe some recent research that has reinvigorated the debate concerning labor market dynamics by suggesting that the job finding rate is much more important than the job separation rate in explaining cyclical patterns in unemployment. I then present a simple model for studying labor market transitions. Next, I describe the SIPP data and its relative merits for studying labor market dynamics. I also show that measures of unemployment derived from the SIPP track the official measure of unemployment rather well. Finally, I describe my main results concerning various labor market dynamics.

Recent literature

Until recently, the conventional wisdom held that most of the variation in the unemployment rate over the business cycle was due to fluctuations in job separations. This view was influenced in large part by the empirical work of Davis and Haltiwanger (1990, 1992), who used plant-level data on manufacturing firms and found that job destruction rather than job creation accounted for the lion's share of employment fluctuations over the business cycle. More recently, Shimer (2005a, b) and Hall (2006) have argued that cyclicality in the unemployment rate is primarily explained by cyclicality in the job finding rate. They suggest that the job separation rate has been relatively steady over the course of recent business cycles. These newer findings are based on models of unemployment dynamics utilizing data from surveys of workers in all sectors.

It is useful to review the key findings from Shimer (2005b) in order to provide some background for the original results that I present using the SIPP. Shimer begins by developing a model of the evolution of unemployment and short-term unemployment set in continuous time. This model yields straightforward expressions for the job finding rate and the job separation rate, which simply require *aggregate* data on employment, unemployment, and short-term unemployment. The implied job finding rate and job separation rate for the period from 1951 through 2004 are shown in figure 1. It is immediately evident that the job finding rate is highly procyclical and has been consistently so over the entire period. In contrast, the job separation rate appears to have undergone a steady secular decline and has been virtually acyclical over the two most recent downturns. Shimer demonstrates

FIGURE 1

Job finding and job separation rates, 1951–2004

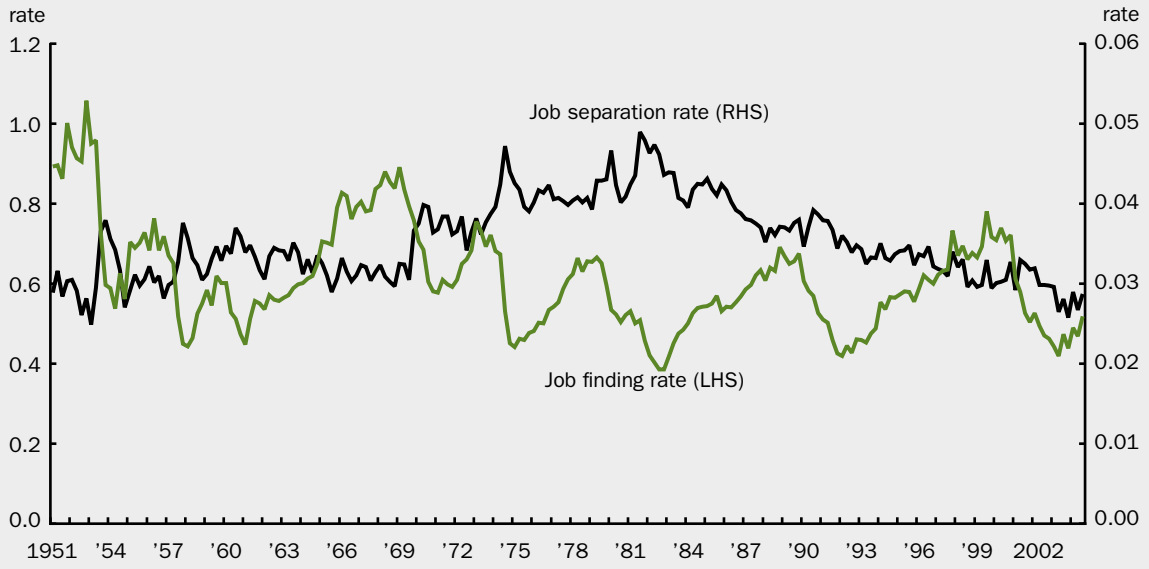

Notes: LHS means left-hand scale. RHS means right-hand scale.
Source: Shimer (2005b).

quantitatively that fluctuations in the job hiring rate are much more strongly associated with unemployment in the two most recent recessions.

Shimer also uses micro-level data on individuals from the matched *Current Population Survey* (CPS) to verify these patterns. He calculates the probability of employment-to-unemployment transition as a second measure of the job separation rate, as well as the probability of unemployment-to-employment transition as a second measure of the job finding rate. He finds virtually the identical cyclical patterns as in figure 1.

Job-to-job transitions

The cyclical properties of job-to-job transitions have also been a focus of attention in the recent literature. Shimer (2005a) calibrates a model of job-to-job transitions using his estimates of the job finding rate and the job separation rate and argues that it implies that the job-to-job transition rate ought to be strongly procyclical. He further argues that the conventional view of a countercyclical job separation rate and an acyclical job finding rate would imply that the job-to-job transition rate should be strongly countercyclical. Therefore, empirical evidence concerning the cyclical properties of job movements potentially provides an independent way to evaluate the competing hypotheses concerning the job separation rate versus the hiring rate.

Until recently, the empirical evidence concerning job-to-job transitions was mainly limited to manufacturing and was inferred by focusing just on workers who quit. Fallick and Fleischman (2004) were the first to exploit the redesign of the CPS in 1994, which enabled researchers to identify month-to-month changes in a worker's employer by matching consecutive surveys. They find that flows between employers are vast and nearly twice the size of flows from employment to unemployment. Using the same data, Nagypal (2004) finds that conditional on staying in the labor force, 70 percent of prime-aged men (25–60 years old) who separate from their employers go to work for a new employer in the subsequent month.

With respect to cyclical patterns, Shimer (2005a), Fallick and Fleischman (2004), and Nagypal (2004) all document that the job-to-job transition rate declined in the 2001 recession, using the matched CPS data.[1] Still, it is unclear based on these limited data whether the most recent recession was typical. In addition, since the rate of job switching appeared to decline sharply from 1994 to 1995 and was fairly flat in the late 1990s despite the booming economy, it is not clear whether this measure exhibits much cyclicality during all phases of the business cycle.

Model of labor market transitions

To organize ideas, it is useful to begin with a simple framework for studying labor market dynamics. For simplicity, I assume that all workers are in the labor force and are either employed or unemployed. I assume that time is discrete and denoted by t, which in practice will be monthly. Let j_{it} denote the job held by individual i at time t. Denote unemployment with a zero so that $j_{it} = 0$ means that individual i was unemployed at time t. Therefore, for individuals who are working at time $t - 1$ (that is, $j_{it-1} \neq 0$), there are three possible labor market transitions: stay on the same job, become unemployed (EU), or switch to a new job (JJ). These transition probabilities are written, respectively, as

1) $\text{Prob}(j_{it} = j_{it-1} | j_{it-1} \neq 0)$,

$\text{Prob}(j_{it} = 0 | j_{it-1} \neq 0)$,

$\text{Prob}(j_{it} \neq j_{it-1}, j_{it} \neq 0 | j_{it-1} \neq 0)$.

Clearly, these must sum to one.

Individuals who are unemployed at time $t - 1$ start a new job (UE) with probability $\text{Prob}(j_{it} \neq 0 | j_{it-1} = 0)$ and fail to start a new job with probability $\text{Prob}(j_{it} = 0 | j_{it-1} = 0)$. Again, these two probabilities must sum to one.

In this article, the job finding rate will be measured by UE, the job separation rate by EU, and the job-to-job transition rate by JJ.

Data

The *Survey of Income and Program Participation* is a series of short panels covering the period from mid-1983 through 2003. These panels are national samples of roughly between 20,000 and 40,000 households where individuals are tracked from two to four years, depending on the particular panel.[2] Households are interviewed every four months (considered one wave) and detailed information is collected on labor market participation and earnings. In addition, up to two distinct employers are identified for each individual in a wave, and this information is used to identify JJ transitions.[3] The sample I use includes all individuals aged 18 through 64.

Comparison with the Current Population Survey

For studying labor market dynamics, there are several key advantages to using the SIPP compared with the CPS. First, since the SIPP is longitudinal, it is designed to follow individuals and track their labor market status over several years. In contrast, the CPS

is a cross-sectional survey that is designed to measure labor market activity at a point in time. The fact that households are potentially reinterviewed in the CPS has been exploited by researchers to try to match some individuals across adjacent months. However, the CPS actually follows street addresses, not *individuals*; so whoever is residing in a given housing unit is interviewed in subsequent meetings. Therefore, when measuring monthly labor market transitions, only the transitions of *stayers* can be measured using the CPS, raising potential selection issues. Thus far there has been little analysis of the potential bias of using the matched CPS to study labor market transitions.[4] Whether there is a potential systematic bias toward finding more or less cyclicality in *UE* or *EU* transitions due to this selection is not immediately obvious and requires greater investigation. With respect to *JJ* transitions, in addition to the potential bias due to people moving, the CPS can only be used to identify employer changes beginning in 1994, when the survey switched to a computerized survey instrument.

Another potential advantage of the SIPP over the CPS is that information about employment status is collected over all weeks during the four-month period, whereas the CPS measures employment status only for one particular week in a month. Therefore, with the SIPP, it is possible to use the underlying weekly data to address measurement problems that may arise if more than one transition occurs between months.

There are also a few drawbacks to using the SIPP. One well-known problem is that respondents tend to underreport transitions within a wave and overreport transitions between interviews. This problem, commonly referred to as "seam bias," suggests that rather than having true monthly data, the researcher actually only has three observations per year. There are also issues relating to consistency across panels. For example, beginning with the 1996 panel, the SIPP was redesigned with the introduction of a computerized interviewing questionnaire.[5] The 1990, 1996, and 2001 panels also oversampled poorer households in order to improve the sample sizes for studying participation in government transfer programs. I have used a statistical procedure to adjust the data for this nonrepresentativeness.[6]

Unemployment in the Survey of Income and Program Participation

Since the CPS is the data set used by the government to calculate the official unemployment rate, it is obviously important that the unemployment rates derived from the SIPP correspond reasonably well with those from the CPS in order for this analysis to be useful. Unemployment as an economic concept is defined by responses to a series of questions in the CPS pertaining to labor market activities for a particular week in the month. The SIPP collects roughly the same information for *each* week in the month and then constructs a monthly variable called *ESR* (employment status recode) summarizing labor market activity for the month.[7] I use *ESR* to construct four monthly indicator variables (0 or 1) for person i in month t. The variable E_{it} stands for being employed the entire month or employed part of the month but with no time spent unemployed during the month. The variable m_{it} stands for being employed part of the month and unemployed (or out of the labor force) part of the month. The variable U_{it} stands for being unemployed for all of the month or being unemployed part of the month and out of the labor force part of the month. The variable O_{it} indicates being out of the labor force the entire month. Being active in the labor force at any point in the month, lf_{it}, is indicated by $1 - O_{it}$. The aggregate unemployment rate in month t is calculated in one of two ways.[8] They are as follows:

$$2) \qquad UR_{narrow,t} = \frac{\sum_i U_{it}}{\sum_i lf_{it}},$$

$$3) \qquad UR_{broad,t} = \frac{\sum_i U_{it} + m_{it}}{\sum_i lf_{it}}.$$

The label "narrow" indicates unemployment that is narrowly defined as only those in the labor force who had no job for the *entire* month. The label "broad" indicates unemployment that is more broadly defined as those who may have been employed part of the month but were also unemployed part of the month. We would expect the narrower measure to understate unemployment relative to the CPS numbers because it omits many of the short-term unemployed that might be captured in the CPS.

Figure 2 plots both unemployment rates calculated using the SIPP, at an annual frequency, along with the official unemployment rate calculated using the CPS. As expected, UR_{narrow} understates the unemployment rate. However, UR_{broad} closely tracks the official unemployment rate based on the CPS with only minor deviations.[9]

Results

Unemployment-to-employment and employment-to-unemployment transition probabilities

Before presenting the results, I describe the exact methodology for calculating the transition probabilities

FIGURE 2

Unemployment rates, 1983–2003

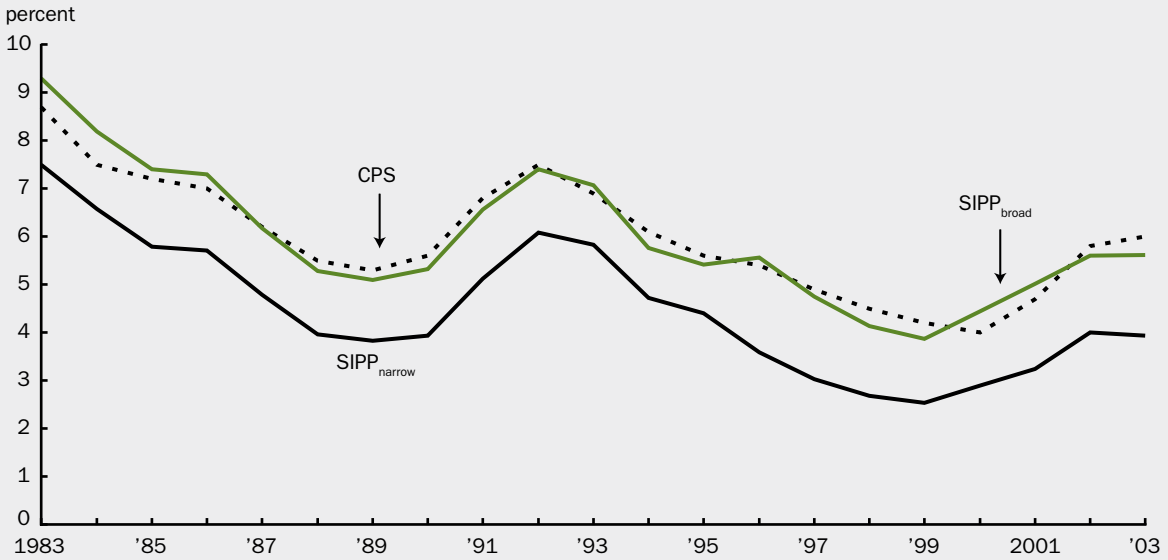

Notes: SIPP means *Survey of Income and Program Participation*. $SIPP_{narrow}$ indicates those in the labor force who are unemployed for the entire month, while $SIPP_{broad}$ indicates those in the labor force who are unemployed for at least part of the month. CPS means *Current Population Survey*. In 1996, there was a survey design change to the SIPP. Also, the SIPP panel for 2000 was discontinued, so the values for 2000 take the average of the values for 1999 and 2001. See the text for further details.

Sources: Author's calculations based on data from the *Survey of Income and Program Participation*; and *Current Population Survey*.

FIGURE 3

Unemployment-to-employment transitions, 1983–2004

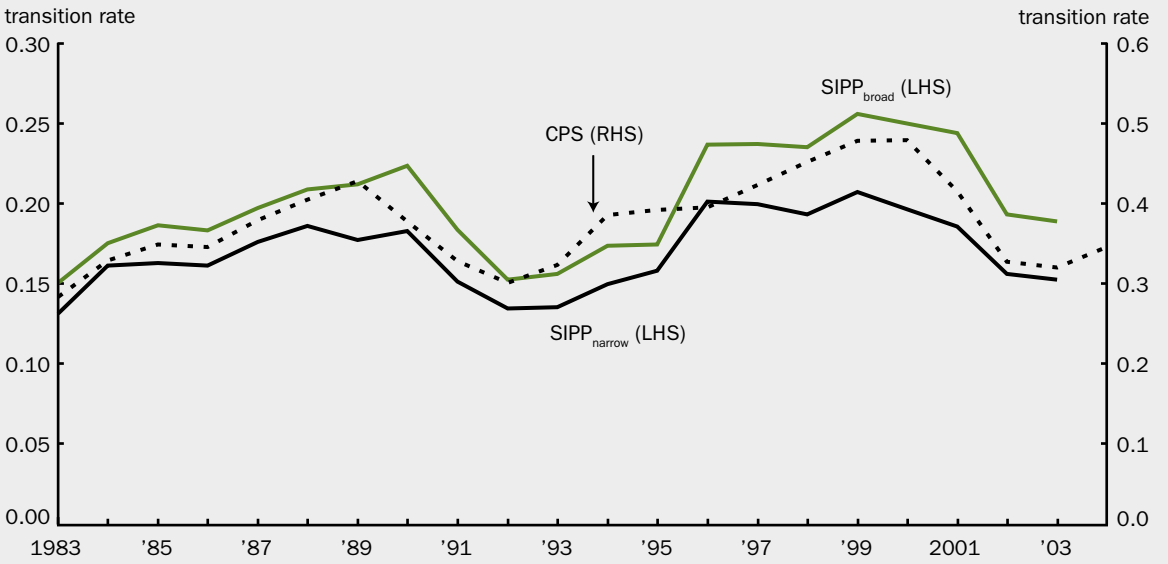

Notes: LHS means left-hand scale. RHS means right-hand scale. SIPP means *Survey of Income and Program Participation*. $SIPP_{narrow}$ indicates those in the labor force who are unemployed for the entire month, while $SIPP_{broad}$ indicates those in the labor who are unemployed for at least part of the month. CPS means *Current Population Survey*. See figure 2 for further details on the SIPP.

Sources: Author's calculations based on data from the *Survey of Income and Program Participation*; and Shimer (2005b).

with the SIPP. Given the previous discussion concerning how to deal with those who are unemployed part of the month, the *UE* and *EU* transitions are each calculated in two ways. One way classifies those who are unemployed part of the month as part of *U* and the other treats them as part of *E*. Formally,

4) $UE_{broad} = \text{Prob}(E_t = 1 | U_{t-1} = 1 \text{ or } m_{t-1} = 1)$,

5) $UE_{narrow} = \text{Prob}(E_t = 1 \text{ or } m_t = 1 | U_{t-1} = 1)$,

6) $EU_{broad} = \text{Prob}(U_t = 1 \text{ or } m_t = 1 | E_{t-1} = 1)$,

7) $EU_{narrow} = \text{Prob}(U_t = 1 | E_{t-1} = 1 \text{ or } m_{t-1} = 1)$.

Figure 3 compares the two measures of the *UE* transition rate calculated from the SIPP with the *UE* transition rate from Shimer (2005b). All series are measured at an annual frequency both to minimize noise and to eliminate seasonal fluctuations. Both of the SIPP measures (plotted on the left-hand scale) produce estimates that are lower than the CPS rate (plotted on the right-hand scale). One reason for this is that Shimer's CPS measures are corrected for what he calls "time aggregation bias." This refers to the fact that some potential *UE* transitions are missed in monthly CPS data when workers who are employed at the interview date in two consecutive months undergo a brief unemployment spell that is missed because of the survey timing. If this is more likely during strong economic times, then the cyclicality of *UE* transitions will tend to be understated and the cyclicality of *EU* transitions may be overstated.

A second reason the SIPP measures are lower could be due to the seam bias problem, which has the likely effect of understating the level of monthly transitions for most months. Interestingly, the cyclical patterns in the data are extremely close. One difference is that the CPS measure peaks in 1991, while the SIPP series peak in 1990. Another difference is in 1996, when a survey design change in the SIPP likely leads to a break in the series. Aside from these minor differences, the cyclical fluctuations are remarkably similar. In that sense, the SIPP results confirm Shimer's findings regarding the importance of hiring in explaining unemployment fluctuations.

The *EU* transition probabilities based on the SIPP are shown in figure 4. It is clear that the survey break in 1996 has a large effect on the EU_{broad} series, but little effect on the EU_{narrow}.[10] Abstracting from the survey break, both series show a similar cyclical pattern and do appear to move with the business cycle.

It is important to point out that the two series in figure 4 still may be susceptible to the time

FIGURE 4

Employment-to-unemployment transitions, 1983–2003

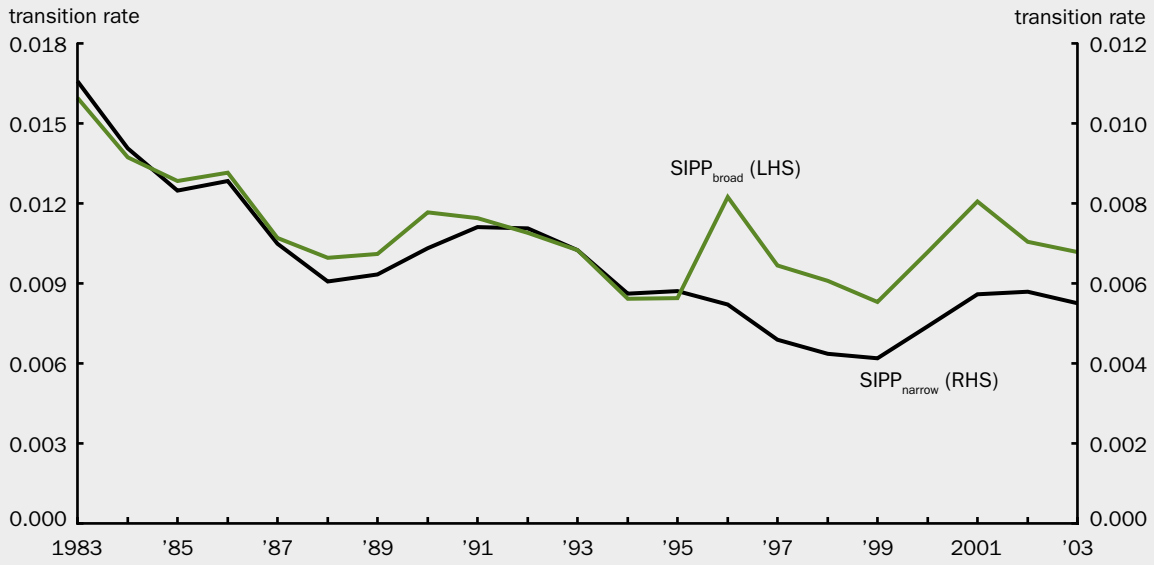

Notes: LHS means left-hand scale. RHS means right-hand scale. SIPP means *Survey of Income and Program Participation*. $SIPP_{narrow}$ indicates those in the labor force who are unemployed for the entire month, while $SIPP_{broad}$ indicates those in the labor who are unemployed for part of the month. See figure 2 for further details on the SIPP.
Source: Author's calculations based on data from the *Survey of Income and Program Participation*.

FIGURE 5

Adjusted employment-to-unemployment transitions and unemployment rates, 1983–2004

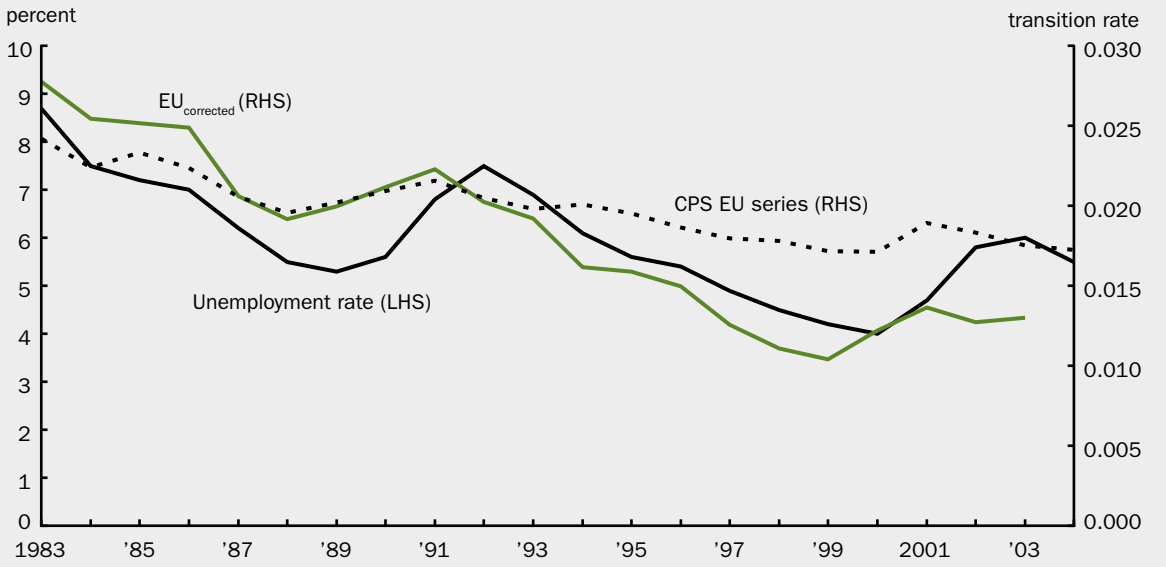

Notes: LHS means left-hand scale. RHS means right-hand scale. CPS means *Current Population Survey*. *EU* means employment-to-unemployment rate. See equation 8 and accompanying text to see how the SIPP measure of *EU* was corrected. See figure 2 for further details on the *Survey of Income and Program Participation*.

Sources: Author's calculations based on data from the *Survey of Income and Program Participation*; *Current Population Survey*; and Shimer (2005b).

FIGURE 6

Job-to-job transition rates, 1983–2005

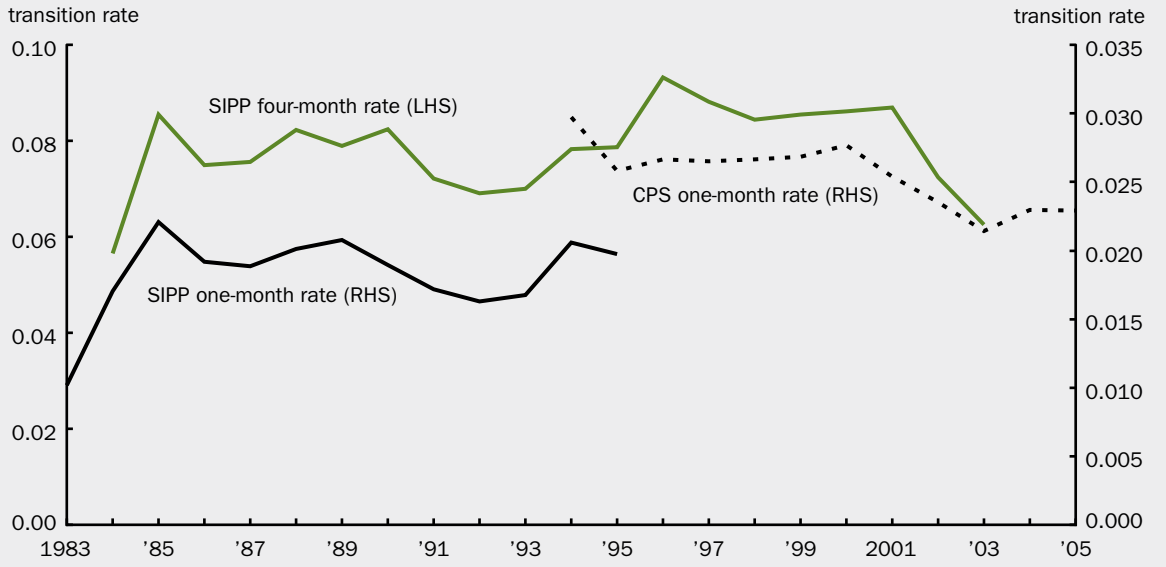

Notes: LHS means left-hand scale. RHS means right-hand scale. SIPP means *Survey of Income and Program Participation*. CPS means *Current Population Survey*. See figure 2 for further details on the SIPP.

Sources: Author's calculations based on data from the *Survey of Income and Program Participation*; and Fallick and Fleischman (2004).

aggregation problem. In order to correct for this, I exploit the weekly information implicit in the monthly *ESR* variable to include employment-to-unemployment-to-employment (*EUE*) transitions that occur within a month. To be specific, I construct a corrected *EU* measured that is defined as follows:

$$8)\quad EU_{corrected} = \text{Prob}(U_t = 1 \mid E_{t-1} = 1 \text{ or } m_{t-1} = 1) + \text{Prob}(E_t = 1 \text{ and } m_{t-1} = 1 \text{ and } E_{t-2} = 1 \mid E_{t-1} = 1 \text{ or } m_{t-1} = 1).$$

In other words, I include workers who were employed for the full month two months ago ($t - 2$) ; were employed part of the month and unemployed part of the month during the previous month ($t - 1$); and are employed for the full month in the current month. This ought to capture individuals who went through an *EUE* transition but were missed by the previous measures. This measure is also adjusted to account for the break in the series due to the survey design change in 1996.[11]

Figure 5 plots the $EU_{corrected}$ series against the unemployment rate and the CPS-based *EU* series from Shimer (2005b). It is worth pointing out that both the SIPP and CPS measures of the *EU* transition rate are corrected (in different ways) for time aggregation. Both measures are plotted on the same scale and show surprisingly similar magnitudes, especially prior to 1996. What is most noteworthy, however, is that the SIPP measure shows more pronounced cyclicality than the CPS measure. Between 1989 and 1991, the *EU* transition rate rose by 11 percent in the SIPP and 7 percent in the CPS. From 1991 to 1999, the SIPP measure (adjusted for the 1996 survey design change) fell by 53 percent, while the CPS measure decreased by just 23 percent. Between 1999 and 2001, the *EU* transition rate rose by 34 percent in the SIPP, but it only increased by 10 percent in the CPS. After further adjusting for the 1996 survey design change, the mean transition rate in the SIPP is 0.018 compared with 0.020 in the CPS. However, the standard deviation for the SIPP is 0.005 compared with just 0.002 for the CPS.

Cyclicality of job-to-job transitions

Perhaps the most important contribution of this article is to provide a longer time series of job-to-job transitions. From 1984 through 1995, it is straightforward to use the SIPP to construct job-to-job transition rates at a monthly frequency, subject to the problem of seam bias discussed earlier. However, since I only examine job switching rates between *waves* (as stated earlier, each wave comprises four months) beginning in 1996, I construct job-to-job transition rates at both the monthly and wave frequency for the earlier period for comparability. The estimates from the SIPP are shown in figure 6 along with similar estimates from the CPS taken from Fallick and Fleischman (2004).[12] It is worth noting that the monthly rate of job-to-job transitions is noticeably lower with the SIPP data than with the CPS data. This is likely due, at least in part, to the seam bias issue that results in understating the transition rate for observations not at the seam. Still, even in the SIPP, the data show that the *JJ* transition rate is vastly larger than the *EU* transition rate, further highlighting the importance of these transitions.

For the period from 1984 through 1995, the one-month and four-month series from the SIPP are quite similar in terms of their cyclical movements. Both series rose sharply in the early 1980s, were relatively flat during the mid to late 1980s, and fell during the early 1990s recession. With the introduction of a new SIPP panel in 1996, the four-month job-to-job transition rate rises sharply from 0.078 in 1995 to 0.094 in 1996 and then declines a bit. There is a pronounced drop in the rate from 2001 to 2003 following the most recent recession. Compared with the more limited data available in the CPS, we now have much clearer evidence that movements between employers are procyclical. These results provide added confirmation to the theoretical findings in Shimer (2005a) that job-to-job transitions would be procyclical if hiring is strongly procyclical and if job separations are acyclical.

Dynamics by demographic subgroups

In this section, I explore how these transition rates and their cyclical properties differ by sex and education level. I now restrict the sample to individuals who are at least 25 years old and who would have most likely completed their schooling. I break down education into four categories of completed schooling: without high school diploma, high school graduate, some college, and college graduate or more. The panels in figure 7 are plotted on a logarithmic scale so that the magnitude of each group's movements are comparable in percentage terms.

Figure 7, panel A shows the EU_{narrow} transition rate for men by level of education. These figures are not corrected for time aggregation bias. The transition rates are highest for those without a high school diploma, averaging 1 percent over the sample period. Those with a college degree or more averaged just 0.3 percent. There does appear to be evidence of convergence in these rates over time. In terms of cyclicality, all groups appeared to follow roughly similar patterns. Indeed, college graduates or those with more advanced education beyond the college degree had a similar percentage increase in the *EU* transition

FIGURE 7

Transitions, 1983–2003

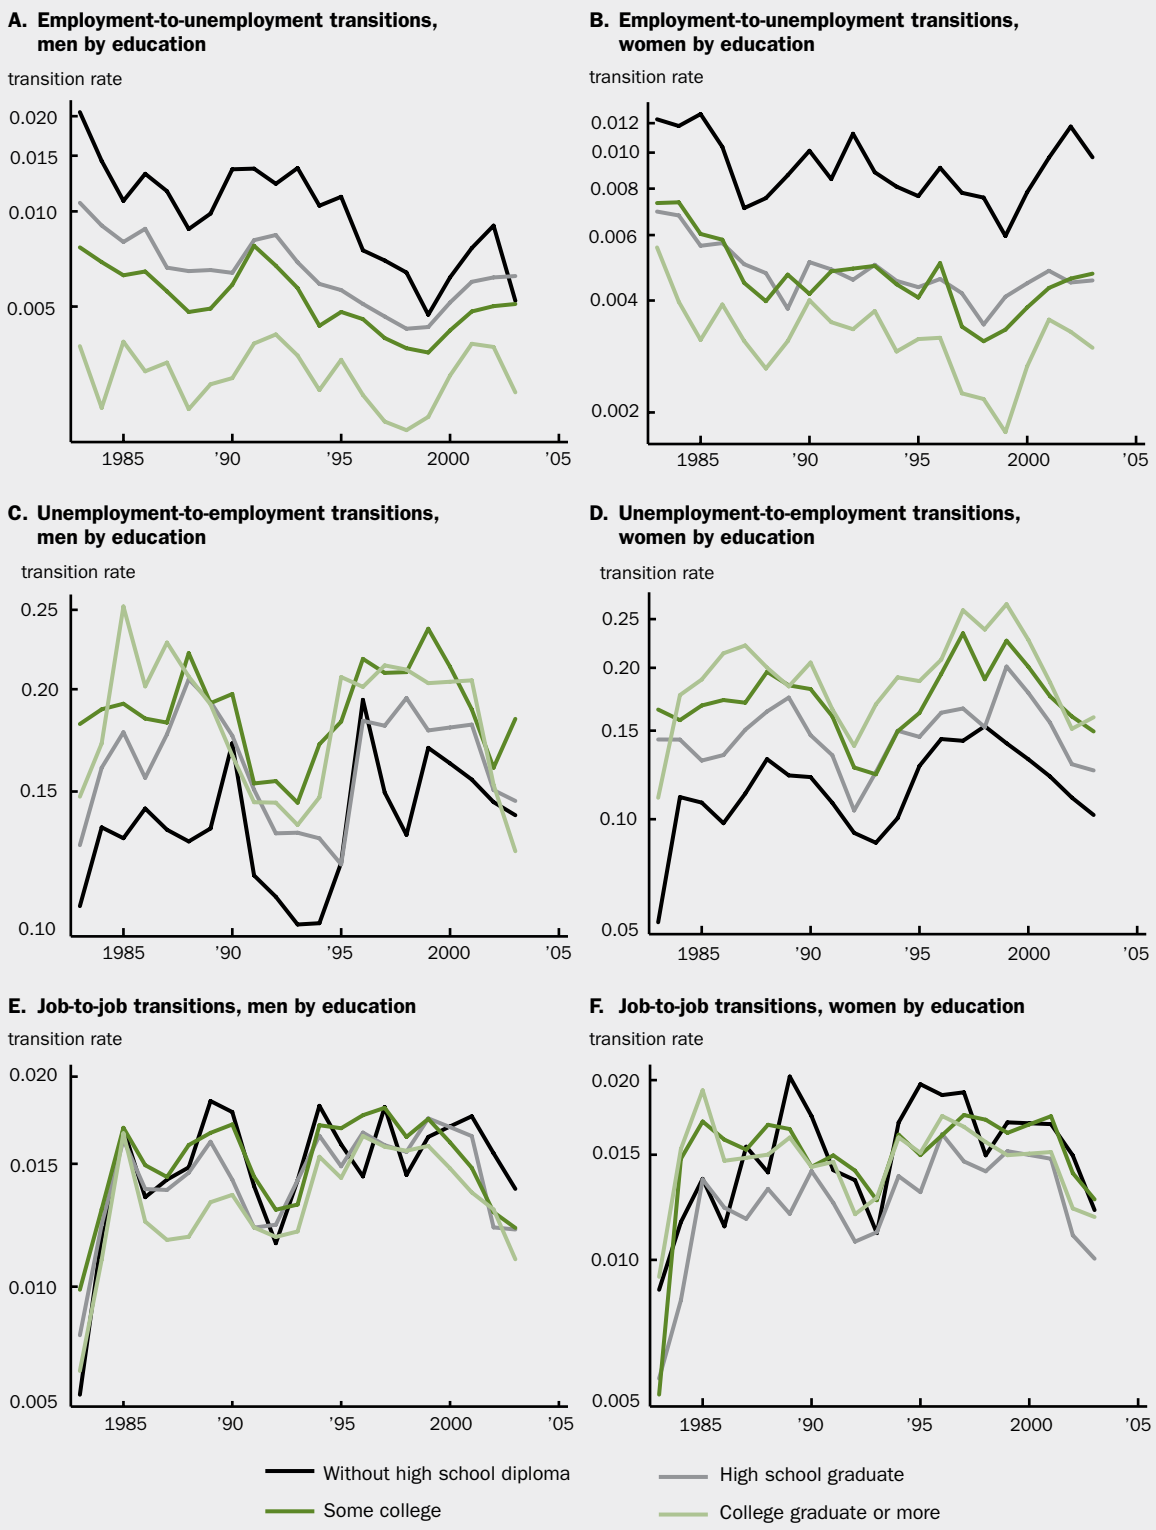

Note: All panels are plotted on a logarithmic scale. See figure 2 for further details on the *Survey of Income and Program Participation*.
Source: Author's calculations based on data from the *Survey of Income and Program Participation*.

probability as those without a high school diploma during the last downturn. Figure 7, panel B shows the same comparison for women. As with men, the groups roughly line up in descending order by education. However, for women, there is no evidence of convergence over time. In addition, the cyclical patterns are really only apparent for those without a high school diploma and those with a college degree or more.

Differences in the *UE* transition for men by education level are shown in figure 7, panel C. Although the differences by education are much less striking than with the *EU* transition, those with some college or more have higher *UE* transition rates on average. However, there have clearly been periods, such as the most recent downturn (2001–03), during which the hiring rate has been similar across all groups. Differences in cyclicality are also not very pronounced across the groups. For women, figure 7, panel D demonstrates that there are clear differences in the level across the groups by education that appear to have been stable over time, but all groups have experienced broadly similar cyclical patterns.

Looking at job-to-job movements (panels E and F of figure 7), there do not appear to be significant differences across education groups in the level of such transitions. It does appear, however, that those without a high school diploma have somewhat more procyclical movements. Given that this series is more volatile in general, it may be prudent not to read too much into this.

Conclusion

An important issue in macroeconomics is the extent to which cyclicality in unemployment is driven by cyclicality in the job separation rate versus the job finding rate. Several recent papers have questioned the prevailing notion that job separations are the primary cause of unemployment fluctuations. These papers present evidence that it is actually the job finding rate that has been the predominant source of cyclicality and that the rate of movement from employment to unemployment has actually been close to acyclical in the most recent recessions. In this article, I use the SIPP, a data source that is especially well suited for studying employment dynamics, to address these questions.

Using the SIPP, I find that the job finding rate as measured by the transition rate from unemployment to employment exhibits virtually the same cyclical patterns as those determined using the CPS. This provides new evidence confirming the importance of cyclical movements in the rate at which firms hire unemployed workers. On the other hand, using the SIPP data, I find a more pronounced countercyclical pattern in the transition rate from employment to unemployment than does Shimer (2005b), who used CPS data. I also present new evidence that covers a longer period than previous studies do; this evidence suggests that the job-to-job transition rate is procyclical. This finding supports the theoretical results in Shimer (2005a)—that is, the job-to-job transition rate is expected to be procyclical if most of unemployment fluctuations are due to a procyclical job finding rate. An examination of these rates by demographic subgroups does not provide any evidence of dramatic differences in the cyclical patterns, suggesting that heterogeneity is not an especially critical concern.

An important area for future research is to explore how wage patterns differ across the business cycle for workers undergoing different types of labor market transitions. Such an analysis would provide a second set of important facts to guide macroeconomic theories of the business cycle.

NOTES

[1]Nagypal (2004) also uses the 1996 and 2001 SIPP panels to show that most of the increase in unemployment in the most recent recession (in 2001) was not due to an increase in the rate at which jobs end, but instead was due to a decline in job-to-job transitions without an intervening unemployment spell.

[2]An SIPP panel for 2000 was discontinued, and so no data is available for the year 2000. In the figures featuring SIPP data throughout this article, values for 2000 take the average of the values for 1999 and 2001.

[3]I use an edited variable produced by the U.S. Census Bureau released in the full panel research files that provides the job identifier of the main job in each month. For the 1990 through 1993 panels, I use a reedited version of the job identifiers produced by the Longitudinal Employer–Household Dynamics (LEHD) program at the U.S. Census Bureau (Stinson, 2003). Unfortunately, full research files for the 1996 and 2001 panels are not currently available. While it is possible to use the start dates and end dates to manually construct monthly job transition rates for these later years, these dates appear to be recorded with considerable error, and so, for this article, I settle for simply using job switches across waves (four-month periods).

[4]Neumark and Kawaguchi (2004) compare the matched CPS data a year apart with the SIPP data to explore the potential effects of selection bias in the matched CPS when estimating the union wage effects and the marriage premiums. Fallick and Fleischman (2004) consider the possible effects of attrition bias on the *level* of the job-to-job transition rate, but not its *cyclicality*. They calculate a lower bound estimate of attrition bias by dividing their sample into cells and reweighting their data based on the probability of a match. However, this makes the strong assumption that conditional on some demographic characteristics, the labor market transition

rates are independent of whether one moves. They also construct an upper bound estimate by assuming that all movers separated from their jobs.

[5]As I will show, for certain time series, this introduces pronounced survey breaks for the post-1996 period. This is similar to the 1994 change in the CPS that affects transition probabilities with that data (Shimer, 2005b).

[6]In 1990, it is easy to identify and remove the oversampled households. However, as far as I am aware, there is no way to identify oversampled households in the 1996 and 2001 surveys. In order to correct for this, I randomly dropped 18 percent of the individuals living in households under the poverty level. SIPP documentation (Westat with Mathematica Policy Research Inc., 2001) suggests that the oversample led to an additional 18 percent of high poverty households being included. I experimented with keeping the full sample and using SIPP-provided weights. I found that when computing the unemployment rate with the full sample, estimates were significantly higher than when using the CPS. Incorporating weights had only a small effect in reducing the estimates.

[7]It is possible to construct a time series of the weekly unemployment rate using the public use files for each wave in each SIPP panel, but this is a fairly large data assembly task. I have settled for using the monthly measure of employment status that is provided with the SIPP's "full panel" public use files for 1984 through 1993.

[8]These calculations also incorporate survey weights provided with the SIPP, which are omitted for simplicity.

[9]One small fluctuation worth noting is that UR_{broad} rises slightly from 1995 to 1996 while UR_{narrow} does not. Indeed, the difference between UR_{broad} and UR_{narrow} is more marked in 1996. This most likely reflects the survey design change that began in 1996. It appears that a higher proportion of individuals are classified as m_t in 1996 than in previous years. This discrepancy begins to revert back in the subsequent years of the 1996 panel, suggesting that some of this might be due to a first-year effect upon introduction of a new survey. In any case, it is clear that the aggregate measures of the labor market from the SIPP describe virtually identical cyclical patterns as those from the CPS.

[10]This is not surprising because, as discussed in note 9, a closer look at the data suggests there was a large change concerning the classification of those who are unemployed part of the month, m_t. Relative to the EU_{narrow} series, the EU_{broad} series removes m_t from the denominator and adds it to the numerator resulting in a sharp and discontinuous increase in the EU_{broad} measure.

[11]I use the percent change in the EU_{narrow} series from 1995 to 1996 and from 1996 to 1997 to produce values for 1996 and 1997. Recall that this series is relatively unaffected by the survey design change. The values for 1998 use the percent change in the original series to update the values from 1997 onward.

[12]Fallick and Fleischman (2004) provide data on the Internet on job-to-job transitions as a share of total population by month; see www.federalreserve.gov/pubs/feds/2004/200434/feds200434.xls. I converted these to shares of the population by using the employment-to-population ratio from U.S. Bureau of Labor Statistics. I then averaged these monthly values over each year.

REFERENCES

Davis, Steven, and John Haltiwanger, 1992, "Gross job creation, gross job destruction, and employment reallocation," *Quarterly Journal of Economics*, Vol. 107, No. 3, August, pp. 818–863.

_____, 1990, "Gross job creation and destruction: Microeconomic evidence and macroeconomic implications," in *NBER Macroeconomics Annual 1990*, Olivier Blanchard and Stanley Fischer (eds.), Cambridge, MA: MIT Press, pp. 123–168.

Fallick, Bruce, and Charles A. Fleischman, 2004, "Employer-to-employer flows in the U.S. labor market: The complete picture of gross worker flows," Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, working paper, No. 2004-34, May.

Hall, Robert, 2006, "Job loss, job finding, and unemployment in the U.S. economy over the past fifty years," *NBER Macroeconomics Annual 2005*, Mark Gertler and Kenneth Rogoff (eds.), Cambridge, MA: MIT Press, pp. 101–137.

Nagypal, Eva, 2004, "Worker reallocation over the business cycle: The importance of job-to-job transitions," Northwestern University, manuscript, April.

Neumark, David, and Daiji Kawaguchi, 2004, "Attrition bias in labor economics research using matched CPS files," *Journal of Economic and Social Measurement*, Vol. 29, pp. 445–472

Shimer, Robert, 2005a. "The cyclicality of hires, separations, and job-to-job transitions," *Federal Reserve Bank of St. Louis Review*, Vol. 87, No. 4, July/August, pp. 493–507.

_____, 2005b. "Reassessing the ins and outs of unemployment," University of Chicago, manuscript, June 21, available at http://home.uchicago.edu/~shimer/.

Stinson, Martha H., 2003, "Technical description of SIPP job identification number editing in the 1990–1993 SIPP panels," U.S. Census Bureau, report, July.

Westat with Mathematica Policy Research Inc., 2001, *Survey of Income and Program Participation Users' Guide*, 3rd ed., report prepared for U.S. Census Bureau, Washington, DC.