

Assessing the jobless recovery

Daniel Aaronson, Ellen R. Rissman, and Daniel G. Sullivan

Introduction and summary

By all measures, employment growth since the recession that ended in November 2001 has been surprisingly weak. But exactly how weak and why have been the subjects of much discussion. In this article, we review the evidence on recent employment trends, provide some new evidence on the role of self-employment, and offer some thoughts on the strengths and weaknesses of several proposed explanations of the causes of what many have called a jobless recovery.

Differing estimates of employment growth from the Bureau of Labor Statistics' (BLS) two monthly surveys have led to some controversy over exactly how weak the labor market has been over the last few years, with the survey of households suggesting somewhat more upbeat labor market conditions than the survey of business establishments. In order to understand the differences in the estimates of employment growth, it is helpful to first adjust for the fact that the two surveys differ in the employment concept they are attempting to estimate. The biggest such difference is that the household survey attempts to measure the number of unincorporated self-employed workers, while the payroll survey does not. Because self-employment has been growing in recent years, this difference in coverage accounts for a portion of the difference in employment growth estimates.

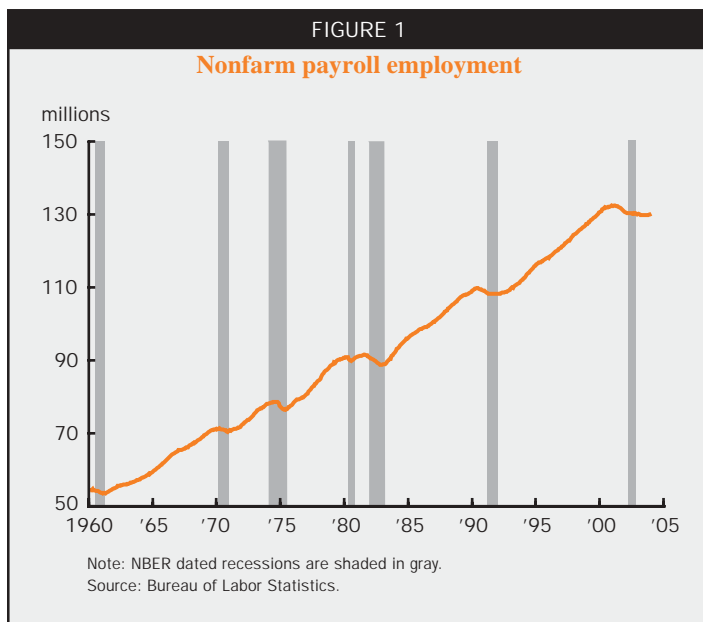
Some analysts have pointed to the growth of self-employment as a hopeful sign of increased entrepreneurial activity. Undoubtedly, many of the newly self-employed are entrepreneurs who will be significant employers in the future. But, we suggest that for some, self-employment is likely a temporary status, which they will quickly relinquish when labor market conditions improve. This seems especially likely in the case of those self-employed individuals who have not incorporated their businesses.

To better understand the nature of the increase in self-employment, we study its relationship with unemployment. Looking across states, as well as over time, we find that higher unemployment rates tend to imply higher rates of unincorporated self-employment. Moreover, when we use our estimates to predict by how much self-employment should have risen given trends in unemployment over the last few years, the results explain a good part of the observed increase since the end of the last expansion. Therefore, this share of the increase in self-employment is likely a reflection of the weak labor market conditions of the last three years.

Even after subtracting the number of unincorporated self-employed workers and making other adjustments to align the household survey estimate of employment with the employment concept of the payroll survey, significant differences remain in recent employment *growth* estimates. Much of this gap is due to a sizable difference in employment counts that developed between 1998 and 2002 that has since receded. Therefore, interpreting recent growth rates, say during the recovery, depends very much on which survey we believe is more accurate. As we discuss in detail below, each survey measure has some potential weaknesses.

In the case of the payroll data, the most significant potential source of error is likely in estimates of the difference between employment in newly opened and newly closed establishments. For such estimates, the BLS is forced to rely on a statistical model rather than hard data. However, the BLS has benchmarked the payroll survey estimates to a nearly full count of paid

Daniel Aaronson is a senior economist and economic advisor, Ellen R. Rissman is an economist, and Daniel G. Sullivan is a vice president and senior economist at the Federal Reserve Bank of Chicago. The authors would like to thank Kate Anderson for research assistance and Spencer Krane and Craig Furfine for their helpful comments and suggestions.



employment through March 2003. Thus, if there is a pickup in employment that has been missed by the payroll survey, it has occurred in the last year.

In the case of the household data, the biggest potential source of error is likely in the estimates of the population by which the survey's estimates of the employment–population ratio are multiplied. Unmeasured flows of immigrants can seriously bias such estimates of the population, and lead to corresponding biases in estimates of total employment. As we discuss, this latter problem is probably more serious. For this and other reasons discussed below, we put greater faith in the payroll survey.


This article concludes by reviewing a list of explanations that analysts have given for the jobless recovery. These include the possibility of an increased need for sectoral labor reallocation, the emergence of just-in-time hiring practices, the rising cost of health care benefits, a fall in labor supply, and the failure of aggregate demand to keep pace with more rapid productivity growth. The evidence in support of the various theories is quite meager. We offer some highly speculative thoughts on their strengths and weaknesses, arguing that sectoral reallocation, health insurance costs, and falling labor supply are unlikely to be major culprits, but that just-in-time hiring and inadequate aggregate demand likely have played a larger role in the weak hiring trends.

Ultimately, however, we are forced to conclude that it is not yet clear what forces have kept employment from growing more robustly in the last two years.

Employment trends

Figure 1 plots the level of payroll employment since 1960. The shaded bars identify recession periods as defined by the National Bureau of Economic Research (NBER). For example, the NBER believes the most recent recession began (the peak) in March 2001 and ended (the trough) in November 2001.¹ For the vast majority of these 43 years, payroll employment is a coincident indicator of economic activity, falling when the economy is contracting (shaded areas) and rising when the economy is expanding (nonshaded areas).

To show this cyclical feature more clearly, figure 2 charts employment during the 12 months before and the 28 months after cyclical troughs. The T refers to the month of the cycle trough. Months prior to the end of the recession (including the recession itself) are represented by the negative numbers to the left of the T, and months after the end of the recession (“the recovery”) are represented by the positive numbers to the right of the T. For example, month 12 is the twelfth month into the recovery whereas month –12 is the twelfth month before the end of the recession. Employment levels are given relative to the value at the trough.



The light orange line in figure 2 gives the average path of employment during the five recessions and recoveries of the 1960s, 1970s, and 1980s. Contrary to recent experience, employment began growing almost immediately after the end of these recessions.² Twenty-six months into the recovery, where we stand as of this writing, employment was, on average, 5.4 percent higher than the trough month and 3.6 percent higher than the previous expansion's employment peak.

However, following the 1990–91 recession (represented by the dark orange line), employment continued to fall. It was 14 months into the recovery before employment returned to the level of the trough and an additional nine months before it exceeded the previous expansion's employment peak. As of January 2004, 26 months into this recovery (the black line in figure 2), the economy has yet to reach the employment level of the November 2001 trough and is almost 2 percent or 2.4 million jobs below the March 2001 employment peak.

An even gloomier picture emerges for total hours worked. The average workweek, at least for the 80 percent of the work force that the BLS labels production or nonsupervisory workers, has been slow to recover relative to previous recoveries.³ The index of total production worker hours is 1.7 percent below the level reached at the cycle trough and 4.9 percent below the previous expansion's production-hours peak. By comparison, in the cycles during the 1960s, 1970s, and 1980s, after 26 months, production worker hours averaged a gain of 6.0 percent relative to the cycle trough and 0.8 percent relative to the previous expansion's peak.

Over the last year, a number of analysts have questioned whether the recent employment performance has actually been as bad as is indicated in figure 2.⁴ The employment numbers in that figure are from the *Current Employment Statistics* data (often referred to as the establishment or payroll survey), a large, nationally representative monthly survey of roughly 160,000 businesses and government agencies covering 400,000 establishments. Two important and related measurement issues in the payroll survey may be of particular importance around economic turning points.

First, the survey may be slow picking up job growth due to unusual levels of firm entry and exit. This is not to say that the BLS ignores firm births and deaths. However, at least until final data revisions are complete, the BLS uses a statistical model to estimate net job changes due to firm entry and exit. It is certainly possible that these statistical relationships are inaccurate over short periods, particularly when births or deaths deviate from historical averages. Second, because the survey only counts paid employees, it does


not count unpaid workers, unpaid family workers, and proprietors who own unincorporated businesses. We discuss this measurement issue in much more detail below. For now, we note that so long as this is a problem, the payroll survey could understate employment growth during recoveries by missing the acceleration of business openings, particularly unincorporated concerns, as business conditions improve.

Those skeptical of the accuracy of the payroll survey note that there is an alternative data source: the monthly survey of households (the household or *Current Population Survey*), a nationally representative sample of 60,000 households that is the basis for the monthly unemployment rate. The household survey may be more timely in accounting for business births and deaths because it simply asks household members whether they are employed in a given month. There is no need to find and measure new employers. Furthermore, since the household survey counts all non-institutionalized persons, there is less concern that it will miss subgroups of workers that are not considered paid employees.

However, these two surveys are fundamentally different instruments and consequently many survey concepts, including the very definitions of *who* and *what* are being counted, differ. The household survey is a more inclusive count of employment, including a number of categories of workers—agricultural, private household workers, owners of unincorporated businesses, unpaid family business employees, or those on unpaid leave of absence (for example, maternity leave)—that are excluded from the payroll counts.⁵ Furthermore, what is actually being tallied differs. The household survey counts the number of people employed, while the payroll survey counts the number of jobs occupied. Thus, in the payroll survey, multiple jobholders are counted for each job they hold.

The quantitative importance of these distinctions is shown in figure 3, which plots the level of employment reported by the household (dark orange line) and payroll survey (black line).⁶ The reported level of employment in the surveys has deviated from 4.5 million to 10.5 million over the last ten years, with the gap varying to some degree over the business cycle.

Fortunately, it is relatively straightforward to adjust for the majority of the surveys' conceptual differences. The light orange line in figure 3 is one such attempt by the BLS.⁷ This series adjusts the household survey count to be consistent with the coverage and concepts (jobs rather than people at work) in the payroll survey. Once such adjustments are made, the surveys match quite closely from 1994 to 1998, at which point the payroll survey began to grow substantially



faster than the household survey. This process reversed around 2002 to the point where, as of January 2004, the adjusted household survey's jobs count is higher, albeit by less than half a million workers or 0.4 percent of total employment.

Figure 4 shows the growth in jobs recorded by the two surveys since the NBER trough in November 2001. At that time, a fairly substantial gap of roughly 1.9 million jobs, favoring the payroll survey, was already in place. Since then, the adjusted household survey has made up all of this ground and more, recording almost 2.4 million or 1.8 percent more jobs than the payroll survey through January 2004. Household employment, unadjusted for payroll survey coverage and concepts, surpassed the November 2001 NBER trough level within six months; the adjusted series reached this goal even quicker. Accordingly, some have argued that the worst of the labor market news is a figment of the payroll data.

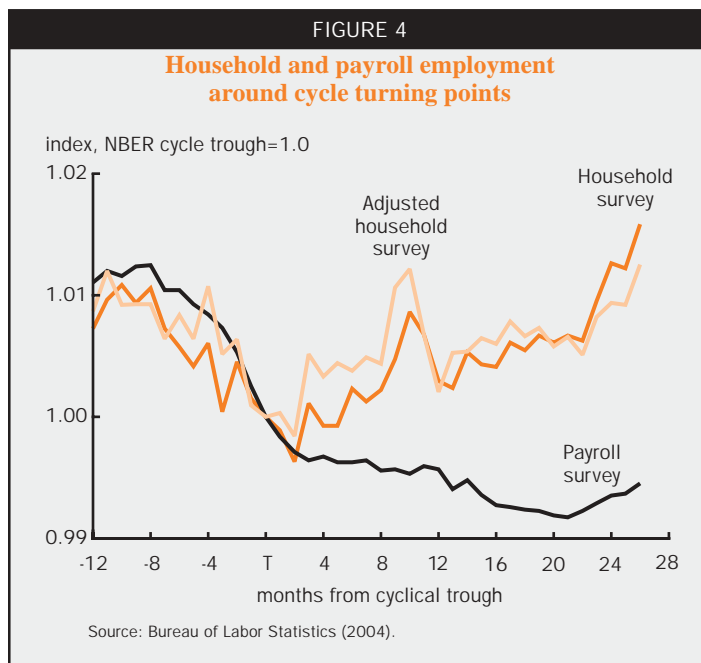
That is not to say that growth in the household survey has been especially strong either. Through January 2004, household employment had grown 1.6 percent (unadjusted, 1.3 percent) during this recovery, well below the typically robust growth rates of the 1960–80s recoveries. Indeed, it is even below the substandard performance recorded in the early 1990s jobless recovery, when the payroll and unadjusted household surveys

recorded roughly 1.8 percent and 2.5 percent gains, respectively, through an equivalent period.

Although both surveys are useful indicators of current labor market conditions, there are reasons to de-emphasize, although not completely discount, the household survey estimate of employment. The payroll survey is much larger,⁸ covers a far higher fraction of employment, and is benchmarked to a universe count of jobs from the unemployment insurance (UI) records once a year, albeit with a lag.⁹ This benchmarking implies that the payroll survey represents a full population of paid employees. The last such benchmarking, reported in January 2004 and covering data through March 2003, showed little adjustment to the jobs picture was necessary, suggesting

that the BLS's statistical model of firm births and deaths was fairly accurate.

Furthermore, the household survey has its own measurement issues. In particular, as detailed in Nardone et al. (2003), household employment growth may have been overstated during the early 2000s. Aggregate household employment growth is derived from two statistics—the fraction of people employed (the employment–population ratio) and the population level. The former is estimated directly from the household survey. Population is enumerated only every ten



years in the U.S. population census and, therefore, must be estimated in intervening years. These estimates can be subject to substantial measurement error. Moreover, the importance of the population estimates is particularly acute for a discussion of recent survey accuracy because the household employment–population ratio has been falling since 2001. Therefore, the entire increase in the household survey’s employment count is due to estimated population growth.

So why might the population estimates be off? The methodology used to produce population estimates for years between the decennial censuses does not account for the state of the economy. However, it is likely that when U.S. labor market conditions are weak (relative to conditions in home countries), fewer immigrants will enter or remain in the country.¹⁰ Consequently, estimated population growth may be too high when U.S. labor market conditions are tepid. This causes the household survey’s estimate of employment growth to be too high as well. Furthermore, overestimation of the population may have been exacerbated by the failure to account for reductions in immigration due to the restrictions imposed after September 11, 2001.¹¹ In fact, as displayed in figure 5, the January 2004 employment release reports a downward adjustment to population of 348,000, by far the largest of its kind in post World War II data, based on revised *estimates* of net international migration. Until the next full census count in 2010, it is hard to gauge whether the problem has been completely fixed.

Interestingly, the large spike in January 2000 represents the opposite dynamic resulting from the strong economy in the late 1990s. Then, low unemployment rates likely led to there being more workers in the country than was expected based on population projections. Consequently, household employment estimates grew significantly slower than payroll, a bias that was only corrected by the full 2000 Census population count.¹²

The recent increase in self-employment

The adjusted household series in figures 3 and 4 are computed in a way that mimics the payroll survey’s coverage. Consequently, changes in self-employment are overlooked. However, self-employment, and in particular its recent rise, is a feature of the household survey that has received considerable attention recently. In particular, a growing number of commentators have pointed to the

spread of self-employment as a sign of surging entrepreneurship and, consequently, an indication of a healthy labor market and economy.¹³

Table 1 gives some indication of the size of this recent increase, and figure 6 charts self-employment rates in a longer historical context. As the recession began at the end the first quarter of 2001, roughly 9.2 million people, or 6.8 percent of nonfarm workers, reported themselves self-employed in nonfarm sector businesses.¹⁴ This figure fell to 8.9 million, or just under 6.7 percent of the nonfarm work force, by the end of the recession and continued to fall to 8.7 million or 6.5 percent of the nonfarm work force during the first quarter of 2002, before rising to 9.5 million, or roughly 7 percent of the nonfarm work force, by the end of 2003. Since the beginning of the recession, the increase in self-employment has been a rather modest one-quarter of a million workers or 0.14 percentage points on the self-employment rate. Measured from different points, however, the increase looks more significant. For example, since the end of the recession, over one-half million workers have, on net, become self-employed, increasing the self-employment rate by 0.31 percentage points. And from its recent quarterly low in early 2002, over 800,000 more workers have become self-employed, raising the self-employment rate by about 0.49 percentage points.

To evaluate whether this increase is a sign of labor market strength, we compare the actual increases with what might be predicted based on a simple statistical model of past relationships between the self-employment rate and a measure of labor market conditions. This approach is described more formally in Rissman (2003).




TABLE 1		
Self-employment during this recovery		
Selected dates	Number of self-employed (thousands)	Self-employment as share of nonfarm payroll employment
2001:Q1 (end of last expansion)	9,245	6.84
2001:Q4 (end of recession)	8,926	6.67
2002:Q1 (self-employment reached low)	8,665	6.49
2003:Q4	9,493	6.98

In that paper, the decision to become self-employed is modeled when self-employment is a low-paying alternative to wage work. Workers are either employed in wage work, unemployed, or self-employed and looking for wage work. Workers shift from unemployment into self-employment if the expected return to job search from self-employment exceeds the expected return to job search while unemployed. Rissman argues that self-employment is countercyclical. During a cyclical downturn, the likelihood of being laid off rises and the probability of generating a job offer falls. In contrast, the return to searching from self-employment remains relatively unaffected. Consequently, in evaluating the two alternatives of 1) searching from unemployment, and 2) establishing a business while continuing to search, the latter option of self-employment becomes relatively more attractive when the economy is weak. Consistent with this hypothesis, for males aged 21 years or older, Rissman finds that increases in local unemployment rates are associated with increases in self-employment.¹⁵

As we show below, the household survey itself shows a similar general pattern. Historically, many of these businesses form during weaker labor markets, when wage and salary jobs are scarce, and subsequently disappear as labor market conditions improve. Therefore, if this historical pattern plays out again, a good part of the recent rise in self-employment could be reversed when labor market conditions improve.

The cyclical nature of self-employment

This section provides evidence on the extent to which self-employment is cyclical. We derive these estimates from the household survey, which as previously


noted is a representative sample of approximately 60,000 households. Participating households are surveyed for four months, left out of the sample survey for eight months, and finally surveyed again for four additional months. Those households in the fourth and eighth months of their participation are known as the outgoing rotation groups (ORGs), and we use them in these calculations. A major advantage of the ORGs files, relative to other household-based samples like the *Panel Study of Income Dynamics* or the *National Longitudinal Surveys*, is the large sample sizes, comprising roughly

180,000 households per year.¹⁶ However, the data only go back to 1979, a relatively short period to evaluate. We end our analysis at the first quarter of 2001, the end of the last expansion.

In order to evaluate the relationship between self-employment and labor market conditions, we compute quarterly local area measures of self-employment and unemployment rates back to 1979. Our geographic boundary is the state, although we check the robustness of our results to using cities, or more specifically metropolitan statistical areas (MSAs), as well. Employing local measures of the key variables introduces significantly more degrees of freedom to our estimation, relative to a simple time series, by taking advantage of different cyclical conditions across the country.

The actual regression that we run is:

$$1) S_{it} = \alpha + \beta U_{it} + \delta T + \lambda T^2 + \varepsilon_i + \varepsilon_{it}$$



where S_{it} and U_{it} are the log self-employment and unemployment rates for state i at time t , T references the quarters elapsed since 1979, ε_{it} is a state-specific error term, and ε_{it} is a normal disturbance term that is independent of the other control variables.¹⁷ We include quarterly dummies to account for any seasonal patterns in the self-employment rate. We eliminate the state-specific error term, ε_{it} , by including state fixed effects. State fixed effects account for any time-invariant unobserved characteristic of the area, in this case including laws or customs related to state differentials in self-employment incentives or ease of business incorporation. We use the remaining variation, how changes in the state unemployment rates covary with changes in state self-employment rates, to identify the cyclicity of self-employment.

Finally, we perform a number of specification checks on equation 1. Two are particularly relevant. First, in this particular example, secular trends in self-employment are captured by a national time trend.¹⁸ However, it is reasonable to imagine that this secular trend is heterogeneous, related to changes in state laws, economic conditions, or local customs. Therefore, we also estimate a version of equation 1, in which T and T^2 are interacted with the state dummies, allowing each state to have its own time trend.

Second, obviously not all small businesses are alike. This heterogeneity prompts one to wonder whether it is high or low quality firms that tend to open during booms and busts and whether this matters in thinking about how to analyze the recent run-up in self-employment. Unfortunately, we do not have financial measures of firm quality in the household survey. Instead, we use an alternative proxy of quality: whether the business is incorporated. This stratification has a second, important advantage. Since the payroll survey includes incorporated but not unincorporated firms, it is useful to know any distinctions between how the two groups react to the business cycle.

Why do we link incorporation status with firm quality? The main reason is that the cost, both in terms of the entrepreneur's time and the direct outlays required to form and dissolve a corporation, is likely to discourage businesses with an expected short window. Granted, the direct outlays, which vary from state to state and depend upon whether the newly formed entity is a corporation or limited liability corporation, are not large. For example, in Illinois, various websites offer to incorporate an Illinois business for under \$500.¹⁹ But combined with the potentially substantial cost of the entrepreneur's time in the process, this may be large enough to discourage incorporation for

those with lower expected success.²⁰ One would also expect that the main advantages of incorporation—personal financial protection to investors, officers, and directors through limited liability—are likely to be more important among larger, high asset businesses.²¹

Table 2 provides estimates of the cyclicity of the self-employment rate, broken down by incorporation status. Results from two general specifications are reported in the first and second row. The first uses a national time trend and the second allows each state to follow its own time trend. Each cell contains two numbers, the impact of a 10 percent increase in the unemployment rate (the elasticity $\beta \times 10$) and, in parentheses, the standard error attached to that point estimate.

Under the two columns labeled “Total self-employed,” we report results when all self-employed workers, regardless of incorporation status, are evaluated. The first column uses weighted least squares (weighted by the size of the state's labor force) with Huber–White and state cluster-corrected standard errors. The second column uses a biweight robust regression technique that we prefer for its high degree of efficiency in the face of the kind of heavy-tailed data that we employ here. For the most part, the results are robust to different specifications and estimation techniques: A 10 percent increase in the local unemployment rate increases the local self-employment rate by about 0.1 percent to 0.2 percent, although none of the estimates are statistically significant at standard significance levels.

However, as we see in the next four columns, the legal type of business matters a great deal. These columns stratify the self-employed into those that own incorporated and unincorporated businesses. In general, we find that there is a statistically and economically important cyclical effect on unincorporated businesses but not on incorporated ones. A 10 percent increase in the state unemployment rate increases the state unincorporated self-employment rate by 0.2 percent to 0.3 percent. However, there is no discernable effect on the state's incorporated self-employment rate.²²

What does this imply about the aggregate increase in self-employment over the last two years? The last two columns of table 2 provide an answer. From the beginning of the last recession until the end of 2003, the unemployment rate increased from 4.2 percentage points to 5.9 percentage points or just over 40 percent.²³ Based on the estimates reported in table 2, we would expect that such an increase in the unemployment rate would increase the unincorporated self-employment rate but not the incorporated rate, causing the total self-employment rate to rise by roughly 0.05 percentage points from its early 2001 level.²⁴ This encompasses over one-third of the 0.14 percentage point increase

TABLE 2								
Cyclicality of self-employment								
	Impact of 10% increase in the unemployment rate						Predicted percentage point increase in total self-employment rate, 2001:Q1 to 2003:Q4 (actual = 0.14)	
	Total self-employed		Incorporated		Unincorporated		WLS	Robust
	WLS	Robust	WLS	Robust	WLS	Robust		
Time trend								
State-specific	0.24 (0.14)	0.12 (0.08)	0.18 (0.20)	0.13 (0.16)	0.28 (0.15)	0.17 (0.09)	0.05	0.03
National	0.14 (0.17)	0.12 (0.08)	-0.45 (0.24)	-0.19 (0.16)	0.37 (0.21)	0.29 (0.09)	0.07	0.05

Notes: These estimates are β (or standard error) $\times 10$. WLS regressions are weighted by state labor force size. Standard errors (in parentheses) are Huber-White and corrected for state clustering. Robust regressions are based on the rreg algorithm described in StataCorp (2001). Sample used is nonfarm workers aged 16 and over in the *Current Population Survey*, 1979 to first quarter of 2001. Observations are aggregated to the state level.

that we have seen since the first quarter of 2001. To be clear, these predicted effects evaluate the growth of self-employment since the beginning of the recession. The model has much less success in forecasting the large decline and offsetting increase (see table 1) in the intervening period.

Furthermore, given the time it takes to open a new business, there is reason to believe that the relationship between labor market conditions and the self-employment decision is not contemporaneous, as we have assumed thus far. Rather, it is some combination of past labor market conditions that matter. Consequently, we reestimated equation 1 but included four quarters of lagged unemployment rates on the right hand side. Table 3 reports the sum of these coefficients (including the contemporaneous estimates) and the resulting prediction for unincorporated self-employment.

Here, we find the magnitude of the cyclicality of self-employment to be about double that reported in table 2. For example, a 10 percent increase in the state unemployment rate increases the state unincorporated self-employment rate by 0.4 percent to 0.7 percent, compared with 0.2 percent to 0.3 percent when lagged effects are not included. Again, assuming no effect on incorporated firms, a 40 percent increase in the unemployment rate would imply a 0.08 percentage point to 0.13 percentage point increase in the self-employment rate, explaining almost all of the small 0.14 percentage point gain that has occurred since early 2001.

These results suggest that a sizable portion of the rise in the self-employment rate since the beginning of the recession is likely related to unincorporated firms surfacing during weak economic times. Many of these new businesses are likely to disappear when the wage and salary sector improves.

What caused the jobless recovery?

Why has employment failed to grow more vigorously the last two years? Analysts have suggested many theories. These include the possibility of an increased need for sectoral labor reallocation, the emergence of just-in-time hiring practices, the rising cost of health care benefits, a fall in labor supply, and the failure of aggregate demand to keep pace with more rapid productivity growth. In this section, we briefly review these theories and offer a few, very speculative observations on their strengths and weaknesses.

TABLE 3				
Cyclicality of unincorporated self-employment (includes four-quarter lag in state unemployment rate)				
	Impact of 10% increase in the unemployment rate: Sum of coefficients		Predicted percentage point increase in total self-employment rate, 2001:Q1 to 2003:Q4 (actual = 0.14)	
	WLS	Robust	WLS	Robust
Time trend				
State-specific	0.47 (0.30)	0.42 (0.25)	0.09	0.08
National	0.61 (0.48)	0.68 (0.26)	0.11	0.13

Notes: Standard errors in parenthesis. See table 2 for more detail.

Sectoral labor reallocation

One frequently mentioned explanation for the jobless recovery is that there is currently an unusually great need to reallocate labor resources across sectors of the economy. The movement of large numbers of workers from one sector to another can be necessitated by changes in trade patterns, shifts in product demand, productivity growth, and other factors. Such movement is a normal feature of a dynamic economy in which some firms contract or close, while others expand or open. Indeed, the reallocation of labor from less to more productive uses is an important source of overall productivity gains in the economy and, thus, of rising living standards. However, because it frequently takes substantial time for displaced workers to find new employers, these long-term benefits often impose substantial short-term costs. This is especially true if jobless workers need to be retrained to acquire the skills that growing employers seek. Thus, if the pace of sectoral reallocation had recently risen, the result would be a temporary increase in the natural rate of unemployment and a temporary fall in employment growth.²⁵

Many analysts claim a link between the disappointing employment growth of the last two years and international trade. Especially great attention has been focused on the new possibilities for “offshoring” service jobs to countries such as China and India, facilitated by the Internet. We interpret such claims as contending that offshoring has forced an abnormally large number of workers to make major career changes and that the increased need for employment transitions has temporarily reduced employment growth.

Further, some commentators suggest that offshoring permanently lowers U.S. employment. However, this view underestimates the ability of our economy to adjust to changing circumstances. New jobs are constantly being created to replace those that disappear. If necessary, wages adjust over time to ensure that the growth of employment closely matches the growth of the labor force. History makes it clear that increased foreign trade is no permanent barrier to employment growth. Indeed, over the last several decades there have been continuing concerns about job losses from the U.S., first to Japan, then to Korea, Taiwan, and South East Asia, and more recently to Mexico. All the while, however, U.S. employment has continued to grow.

To say that the effects of increased trade are temporary, however, is not to say that they are unimportant. New trade patterns could force many workers to make significant career changes, increasing the need for costly job search and retraining.²⁶ The same is true of reallocation caused by shifts in relative product demand or differences across industries in rates of productivity

growth. A large enough increase in sectoral labor reallocation, whether due to increased trade or other reasons, could explain the weak employment growth of the last two years.

But, has there been a major increase in reallocation? We are skeptical. The most frequently cited evidence of such an increase is contained in a recent article by Groshen and Potter (2003). On the basis of an analysis of industry-level employment data, they conclude that the need for workers to shift to new jobs “largely explains why the payroll numbers have been so slow to rise.” However in Aaronson, Rissman, and Sullivan (2004), our other article in this issue, we argue that the statistic proposed by Groshen and Potter is unlikely to be a good proxy for the extent of reallocation. Moreover, when we compute a measure based on Rissman (1997), which better captures the need for reallocation across industries, we find that reallocation rose less during the last two recessions than in previous downturns and that once the recent recessions ended, reallocation returned relatively quickly to low levels. In essence, we find that employment shares by industry are relatively stable and that what shifts we do see are typical of the patterns usually observed when overall labor market conditions have been weak. Thus, we do not think that the need to reallocate workers across industries is a likely explanation of the jobless recovery.

Our findings in Aaronson, Rissman, and Sullivan (2004) indicate that the need to reallocate workers across *industries* has not been unusually great. However, our results do not necessarily rule out an increase in *other* forms of sectoral reallocation, such as those in which workers are forced to change occupations or geographic regions. However, as we also discuss in Aaronson, Rissman, and Sullivan (2004), there is some evidence that overall rates of job destruction and job creation are both at low levels, which seems inconsistent with a major role for any form of labor reallocation. Thus, we are skeptical that any form of sectoral reallocation provides the explanation for the jobless recovery.

Just-in-time hiring

We are more sympathetic to the theory that the employment practices some are calling just-in-time hiring have played a role in restraining employment growth since the end of the recession. To understand this theory, consider the staffing decisions of a firm that expects, but is not sure, that demand for its product is about to increase. Initially, bringing on new workers takes significant time and expense to find, screen, and deploy. In addition, once new workers are hired, they are expensive to let go, perhaps because of the possibility of lawsuits. The firm faces a tough choice.

On the one hand, if it increases its employment and perhaps builds inventories, but demand turns out to be weak, then high payroll costs will mean poor profits. On the other hand, if the firm maintains its current level of employment but demand picks up, then it may initially be unable to increase production in line with orders. By the time it has increased employment, it may have missed making some profitable sales. Firms will deal with this tradeoff in various ways, but at least some are likely to hire workers in *anticipation* of increasing demand.

Now suppose that new institutions allow the firm to bring on additional workers on very short notice. Moreover, the workers it brings on can be dismissed on equally short notice. This new institution makes the firm's choice much easier. It can wait to see whether demand picks up and, if it does, summon additional workers to increase production. It doesn't have to worry about having unneeded workers on its payrolls or being unable to fill new orders. Given the new option, fewer firms are likely to hire workers and begin to build inventories before they can verify the increase in demand.

To a significant extent, such new hiring institutions have emerged over the period of the two jobless recoveries. The most obvious example is the temporary services industry. When the economy was recovering from the recessions of the early 1980s, this industry was still very small and mainly dedicated to providing clerical workers to fill vacancies created by temporary absences of clients' regular employees.²⁷ Since then, the industry has grown very rapidly as can be seen in panel A of figure 7 overleaf. At its peak in 2000, it accounted for 2 percent of nonfarm payroll employment.²⁸ Furthermore, the nature of the workers the industry provides to client firms has changed. It now provides many more light industrial, call center, and technical workers than in the mid-1980s.²⁹ In addition, temporary service workers are now sometimes the majority of workers at a client's establishment. Thus, this industry's role has expanded from one of helping clients deal with short employee absences to one of allowing them to vary the scale of their operation.

The temporary services industry is not the only one adding flexibility to firms' staffing levels. Another example is management and technical consulting. Consultants can quickly be deployed to a firm needing increased access to specialized skills. As panel B of figure 7 shows, this industry has also grown rapidly in recent years.

In addition to the growth of industries that directly provide labor services, a number of other developments have likely made it easier for firms to hire new

workers quickly. These include fax machines and the Internet, which can speed up the process of screening potential hires, as well as new firms that match employers and employees and maintain enormous files of potential job candidates.³⁰ All together, these new developments may reduce hiring costs and make it more feasible for firms to hire workers "just in time" to perform needed tasks.

The increase in the flexibility of firms' hiring decisions is at least consistent with more sluggish employment growth at the beginning of business expansions. Such periods are characterized by much more than the usual degree of uncertainty. Indeed, it is often not clear how strong and broad a pickup in economic activity is. It may not even be clear that a pickup actually has started or if there is a significant chance of a "double dip" recession. So, it is natural for firms to be cautious. In the absence of the new hiring institutions, the fear of getting caught with a large work force and few orders often may be less than the fear of not having enough workers to meet rising demand. So firms may tend to hire and build inventories in anticipation of demand. However, with confidence that they can increase employment rapidly when needed or with the temporary services industry as a backup, firms may be more likely to wait to hire until demand is definitely strong. We cannot claim that the evidence linking the growth of flexible employment options to the jobless recoveries is extremely strong, but the timing is right and it agrees with at least some anecdotal evidence.

This argument for just-in-time hiring institutions having contributed to sluggish employment growth at the beginning of expansions does not necessarily imply that there must be a major increase in temporary services employment during these periods. It is the option value that temps provide to firms that restrains hiring that might otherwise have taken place in anticipation of demand growth. However, under this theory, if temporary services hiring picks up, it is likely to signal a pickup in overall hiring some time later. And, in fact, Segal and Sullivan (1997a) and others have shown that temporary services employment is a leading indicator for overall employment growth.

Panels C and D of figure 7 show how employment in the temporary and management and technical consulting services industries has grown over the last two recovery periods. In 1991, temp employment was flat for approximately half a year after the trough, but then began to grow very steadily. This takeoff was six to nine months before that of overall employment. After the most recent recession, temporary employment started to grow after a few months, but then fell back, and did not start increasing in a significant way until

the summer of 2003. If the lag between the start of its growth and that of overall employment growth matched that of the previous cyclical episode, the jobs figures should have begun improving around the beginning of 2004. The March employment report indicates improvement and the relatively robust growth of temporary services jobs is one reason for at least modest optimism going forward.

Health care costs

Another prominent theory is that increases in health care costs have played a significant role in restraining employment growth.³¹ There seem to be two versions of this theory. One emphasizes that rapid growth in health care premium costs may have temporarily raised total employment costs above equilibrium levels, thus reducing labor demand. The other emphasizes

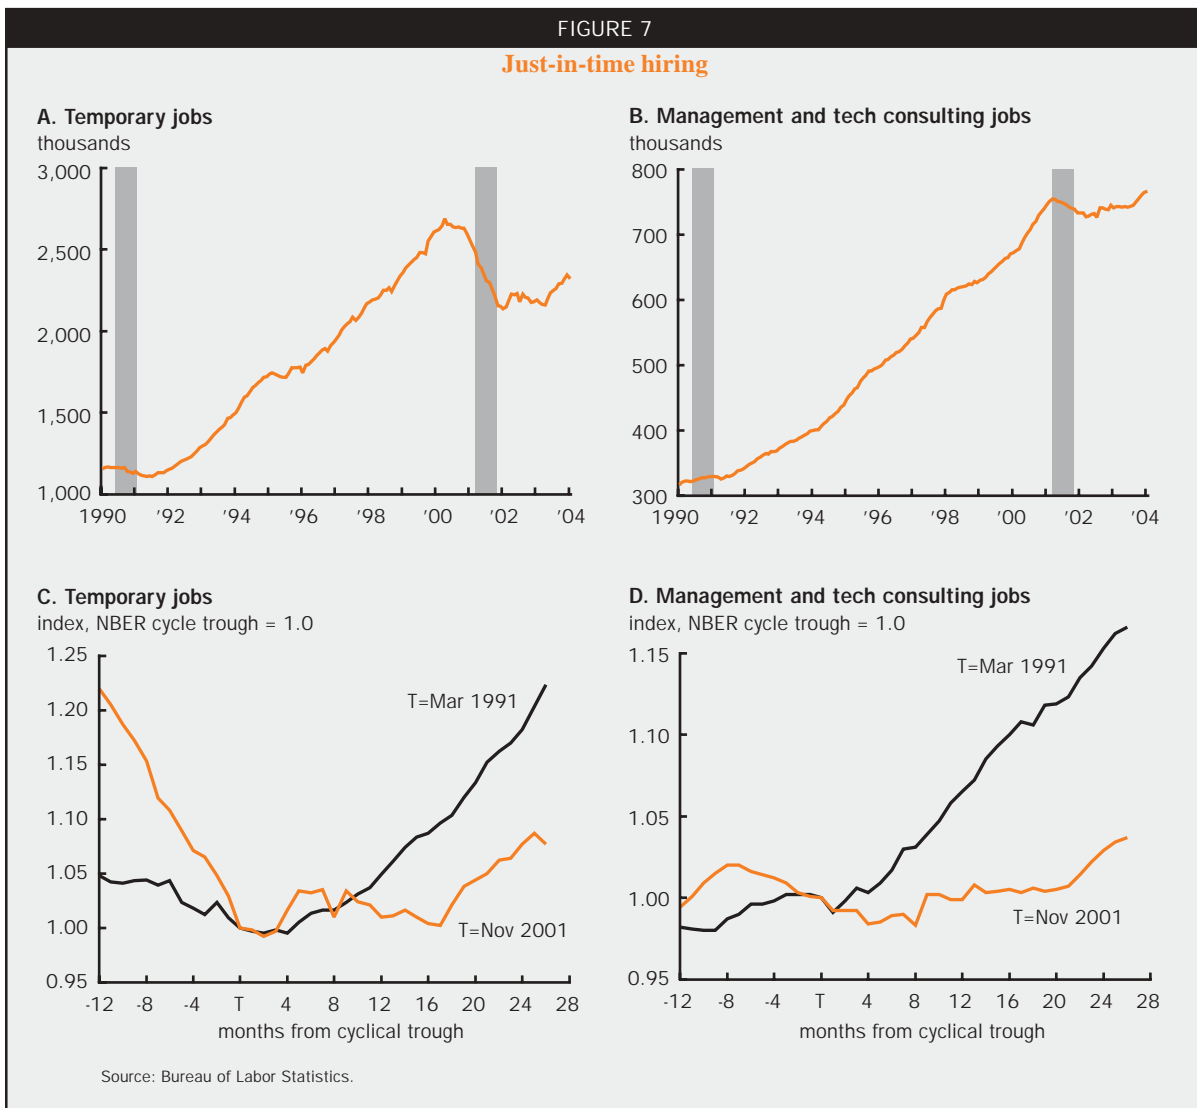
the increase in the importance of costs that are fixed per worker, independent of hours worked or salary level. Such costs can lead employers to reduce employment by requiring existing employees to work longer hours or by employing fewer, more highly skilled workers.

Figure 8 shows the rate of increase in health care costs relative to the overall personal consumption expenditures (PCE) deflator. Clearly the growth in health care costs has significantly outpaced overall consumer price inflation in recent years. Such increases in health care prices raise the cost of employers' existing medical benefit plans and could make hiring less attractive. This view gains some additional support from the observation that health care costs also rose very rapidly during early 1990s, the time of the first jobless recovery.

Two major challenges can be offered to this theory. The first is that employers can respond to such cost

FIGURE 7

Just-in-time hiring



increases by reducing the generosity of health insurance benefits or by altering other aspects of compensation such as the level of wages and salaries. Thus, an increase in premiums does not have to lead to higher total employment costs. A possible response to this challenge is that the adjustment of other aspects of the employers' compensation packages can take time. In the short run, firms might not be able to fully offset medical cost increases with cuts elsewhere.³²

Another challenge to the theory that rapidly rising health insurance costs are behind the jobless recovery is that the recent growth of *total* labor compensation costs has not been particularly rapid, especially when judged relative to productivity growth. Growth in the employment cost index has fallen only modestly since the onset of the recession, going from 4.1 percent for the 12 months ending in December of 2000 to 3.8 percent for the 12 months ending in December of 2003. However, other measures of labor compensation growth have dropped by substantially more. In particular, over the same period, the four-quarter growth of hourly compensation in the nonfarm business sector as reported in the productivity and costs data fell from 7.0 percent to 3.2 percent. Moreover, this fall occurred while productivity growth was increasing from 2.8 percent to 4.4 percent. Thus, the growth in unit labor costs, which measure the nominal cost of producing a unit of output, fell from 4.2 percent to -1.2 percent. Unit labor costs have been falling since the first quarter of 2002, accumulating to a 4.1 percent decline, the largest such drop since at least the early 1960s. As a result, the share of labor in total costs has fallen sharply since the first quarter of 2001. This drop is also the largest since the early 1960s. Given that total labor compensation levels are not rising particularly quickly, we do not place much weight on the possibility that rapidly rising medical costs are boosting overall compensation costs above levels consistent with employment growth.

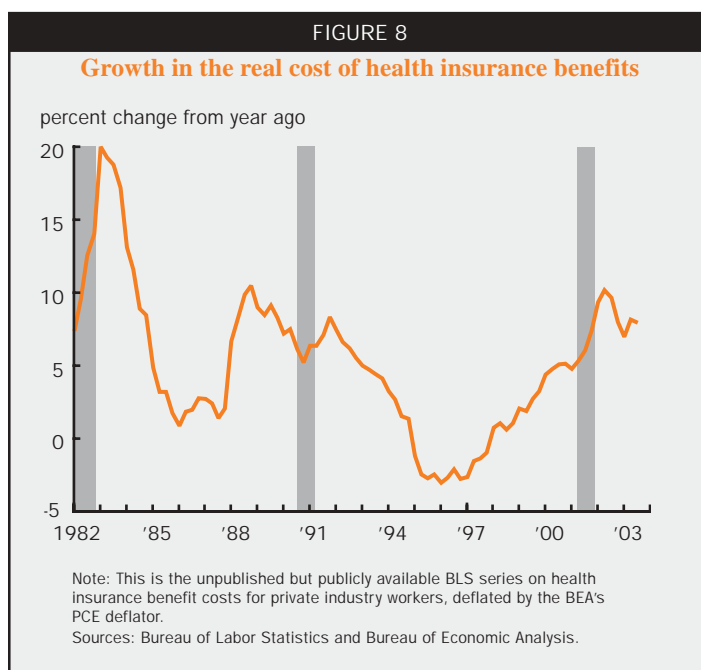
What about the version of the health costs theory that points to higher fixed costs per worker? As figure 9 shows, health insurance costs rose from 5.4 percent of compensation in 1999 to 6.5 percent in 2003. This fraction rose to similar levels during the jobless recovery of the early 1990s before declining in the mid-1990s. So, high health costs per worker have tended to correlate with sluggish employment growth.³³

An increase in fixed costs per worker tends to give firms an incentive to hire fewer workers, while increasing their average work week.³⁴ However, the average work week has not been rising. In fact, average weekly hours of production workers are at the same level as they were at the trough of the recession in November 2001. It is possible that work hours are, in fact, increasing, but that this is somehow not being recorded in the data. But, we know of no evidence that this is the case. Moreover, one would expect that if actual hours per worker were rising, there would be at least some increase in measured hours as well. Possibly, the reason average hours per worker are not rising is that some firms have responded to higher insurance premiums by keeping more workers' hours below the threshold they use to qualify employees for benefits. However, if this were true then fixed benefits costs would not be restraining employment growth. Thus, we do not find the fixed-cost version of the employment costs theory very persuasive either.³⁵

Labor supply

Just-in-time hiring and health care costs are reasons why labor *demand* may be unusually low. Logically, another possibility is that labor *supply* may have declined. Under this theory, employment growth may be lagging in part because relatively few people want to work.

As support for this theory, one might point to the unusually large decline in labor force participation over the course of the recession and the following two years. Historically, labor force participation has been only



modestly procyclical. But since the most recent recession began, labor force participation has fallen substantially, from 67.1 percent (just off the all-time high) at the March 2001 peak to 65.9 percent in February 2004, the lowest value since 1988. Because of the steep drop in labor force participation, the unemployment rate has fallen considerably since the end of the recession even while the growth in employment (by either the payroll or household survey measures) has been slower than the growth in population.

An obvious objection to interpreting the decline in labor force participation as a fall in labor supply is that it may reflect people's discouragement over their chances of finding a job rather than a decline in their desire to work. A corollary of this view would be that the standard unemployment rate offers a misleading view of the degree of slack in the labor market. However, as Barrow (2004) discusses, relatively little of the increase in nonparticipation is attributable to greater numbers of discouraged workers; the increase in those out of the labor force who report wanting a job has been quite small. Therefore, the majority of the decline in the participation rate appears to be a genuine fall in labor supply.

Given that the drop in participation appears to a substantial extent to represent a fall in labor supply, a quick bounce back to the levels of late 2000 does not seem particularly likely. Rather, those pre-recession levels may have been the result of some overshooting in the labor market, much like the unemployment rates of 4 percent were likely lower than the long-run

equilibrium unemployment rate. Moreover, the standard unemployment rate may not be a particularly good indicator of labor market conditions.

Table 4 breaks down the changes in labor force participation by age and gender. These figures suggest two groups are primarily driving the decline. First, there has been a remarkably large fall in the labor force participation rate of teenagers and young adults in their early 20s. The decline in teenage labor force participation since January 2001 accounts for over 0.5 percentage points, or about half, of the decline in the *economy-wide* labor force participation rate. The second critical group is women between the ages of 25 and 44. The fall in this group's labor force participation since January 2001 accounts for about 0.4 percentage points, or roughly 40 percent, of the aggregate decline.


The drop in teenage and young adult labor force participation seems especially large, recently falling to post-World War II lows among males and 30-year lows among females. Is this a normal cyclical response to an extended period of disappointing employment growth or a permanent shift in the school-work decision? Table 5 offers some clues.

Here, we display the *annualized growth* in the share of the teen and young adult (20 to 24) population that are in school (column 2) and not in school (column 5).³⁶ Between 2000 and 2002 (latest available data), the share of teenagers in school grew by 1.2 percent per year, double the rate of the mid- to late 1990s, when labor markets were strengthening. However, the share of students in the teen population was growing even faster—1.7 percent per year—during the recession and jobless recovery of the early 1990s. The growth in the share of young adults 20 to 24 that are in school was roughly 2.5 percent per year during both jobless recoveries, slightly faster than the 1990s and 1980s expansions. So there has been a noticeable shift in school going activity but it does not seem to be especially large, at least relative to the previous recession.

However, there is at least one notable distinction during the most recent period. In the remaining columns, we look at the joint decision to be in school and in the labor force. Whereas the gain in schooling during the early 1990s was fairly evenly split between those who remained in the labor force and those who did not, this has not been the case in the early 2000s. Recently, there has been a sizable increase in the share of young people in school

FIGURE 9

**Health insurance benefits as a share of total compensation
(private industries, all workers)**



Notes: Employer costs for employee compensation (ECEC) survey. For 1986 to 1990, only total insurance is reported. We assumed that health insurance is 92% of total insurance in these years, as it is in 1991.
Source: Bureau of Labor Statistics.

TABLE 4				
Labor force participation by age and gender				
	January 2001	January 2004	% change	Share of the noninstitutionalized 16+ population January 2004
Females				
16–19	50.5	44.3	–12.3	3.6
20–24	73.6	70.1	–4.8	4.5
25–34	75.9	73.7	–2.9	8.8
35–44	77.8	75.1	–3.5	9.9
45–54	76.6	76.7	0.1	9.4
55+	26.4	30.1	13.9	15.6
Males				
16–19	53.0	44.5	–16.0	3.7
20–24	83.4	79.8	–4.3	4.5
25–34	93.4	92.0	–1.5	8.7
35–44	92.8	91.9	–1.0	9.6
45–54	88.6	88.3	–0.4	9.0
55+	40.8	43.2	6.0	12.7

Source: Bureau of Labor Statistics.

but not in the labor force—4.6 percent per annum for teens and 8.9 percent per annum for 20–24 year olds—but a drop in the share both in the labor force and in school. Thus, not only has there been a fall in labor force participation of those not in school but, in contrast to previous patterns, of those in school as well.

It is hard to say why this has happened. One highly speculative explanation for the school group’s behavior is that the strong productivity growth of late 1990s might have seemed transitory as it was occurring. Now, after almost a decade of stronger growth, it seems more permanent. Consequently, workers expect wages to be higher in the future. A classic “intertemporal” response to such expectations is to work

less now and more in the future. Furthermore, if productivity growth is expected to be higher in higher skilled occupations, we would expect an increase in school enrollments as well. This story might also be consistent with the labor supply behavior of secondary earners in general and, thus, help to explain some of the fall in the labor force participation of women in the traditional child-rearing age groups.

However, in the end, we doubt that the decline in labor supply is the primary reason for the sluggish employment growth of the last two years. If declining labor supply were the cause, we would expect to see relatively strong wage growth. However, as we noted in the discussion of health care costs, growth in labor

TABLE 5						
Growth in share of 16 to 24 year olds in school and labor force (annualized percentage growth rates)						
	In school			Not in school		
	Total	In labor force	Not in labor force	Total	In labor force	Not in labor force
16 to 19 year olds						
2000–2002	1.2	–3.5	4.6	–5.0	–6.3	–0.7
1993–2000	0.6	1.0	0.4	1.5	2.0	0.3
1990–1993	1.7	1.7	1.7	–5.2	–6.3	–1.7
1984–1990	0.0	0.5	–0.3	–3.0	–2.9	–3.5
20 to 24 year olds						
2000–2002	2.5	–1.2	8.9	–1.9	–2.6	1.6
1993–2000	1.7	2.4	0.6	0.6	0.9	–1.0
1990–1993	2.6	3.5	1.1	–0.9	–1.5	1.9
1984–1990	2.1	2.8	1.2	–2.1	–2.0	–2.6

Source: Bureau of Labor Statistics.

compensation has not been particularly strong recently, especially relative to productivity growth. Nevertheless, even if a decline in labor supply is not the primary cause of the jobless recovery, it may be a contributing factor. If fewer workers had withdrawn from the labor force, the unemployment rate would have been higher and wage growth might have fallen further. Lower wage growth could, in turn, have led to an increase in firms' employment of workers.

Inadequate aggregate demand growth


Probably the most frequently mentioned explanation for the weak employment growth of the last two years is rapid productivity growth. As can be seen in figure 10, productivity growth has been very impressive of late. In the two years since the trough of the recession, output per hour worked in the nonfarm business sector has risen by more than 9 percent, approximately 3.5 percentage points more than in the typical post-recession period. Moreover, relative to the *peak* of the last business cycle, productivity growth has been even more impressive, with output per hour having risen by more than 12 percent or 5.5 percentage points more than the typical cyclical pattern.

Holding constant growth in output, greater growth in productivity implies less growth in hours worked and likely, therefore, less growth in employment. In this sense, rapid productivity growth explains weak employment growth. However, there are no fixed speed limits to economic growth. So, there is no reason to expect output growth to remain constant in the face of more rapid productivity growth. Rather, economic theory suggests that a pickup in the growth of aggregate supply induced by more rapid productivity growth would ordinarily lead to expectations of higher incomes and a comparable pickup in aggregate demand and output growth. Thus, those who attribute the jobless recovery to rapid productivity growth must implicitly be attributing it to some break in the normal link between aggregate supply and aggregate demand.

Why should aggregate demand have failed to increase sufficiently in response to the pickup in productivity growth? Quite possibly, demand growth has been held back over the last two years by a series of shocks—concerns over terrorism, the buildup to the Iraq war, revelations of poor corporate governance, and the hangover from the fall in the stock

market. These shocks likely made firms unusually cautious and, therefore, reluctant to invest and hire.³⁷ Similarly, in the early 1990s, the savings and loan crisis may have created a credit crunch that restrained aggregate demand and hiring. Such an interpretation suggests that absent such post-recession shocks, future recoveries may see employment expansions more like those before the 1990s.

Another possible explanation for the failure of demand to keep up with the strong productivity growth of the last two years is that recent productivity gains may owe more to the ebbing of adjustment costs associated with previously high investment levels than to new breakthroughs in technology that increase expectations for future income. Under this interpretation, rapid technological advance led to very high levels of investment in new equipment in the late 1990s. The need to make all this new equipment work may have led in turn to what now looks like excessive hiring. This story suggests that firms recently have been eliminating those excesses. As a result, measured productivity growth has been strong. But because the driver of productivity growth is not new breakthrough technology, there is no support for additional increases in expectations of future incomes. Thus, increased measured productivity growth is not associated with increased aggregate demand and hiring. Rather, employment levels might be interpreted as returning to normal after a period in which they were temporarily elevated in order to implement new technology.



Conclusion

This article reviews evidence on the jobless recovery. We began by summarizing recent employment trends, showing, as has the BLS, that the major surveys of employment can be reconciled for much of the last ten years. But a gap developed between 1998 and 2002 that has led to somewhat different interpretations of what has happened to the labor market during the last three years. While we have sketched reasons to put more weight on the survey of employers rather than households, regardless of which survey we use employment growth has been historically weak. Furthermore, we have provided new evidence that the small increase in self-employment since the beginning of the recession, which has been interpreted by some as a hopeful sign for labor markets, is likely a reflection of the weak labor market conditions of the last three years.

Finally, we have offered some speculative thoughts on the strengths and weaknesses of the various explanations analysts have given for the jobless recovery. These include the possibility of an increased need for sectoral labor reallocation, the emergence of just-in-time hiring practices, the rising cost of health care benefits, a fall in labor supply, and the failure of aggregate demand to keep pace with more rapid productivity growth. We argue that the most plausible culprits are just-in-time hiring and inadequate aggregate demand in the face of rapid productivity growth. Ultimately, however, we must conclude that it is not yet clear what forces have kept employment from growing more robustly the last two years.

NOTES

¹Recent newspaper reports suggest that the NBER dating committee is considering the unusual step of changing the date of the cycle peak. See, for example, Henderson (2004).

²In fact, employment was growing within four months, and in two cases within one month, of the trough in all five of the 1960–80s recoveries. To a degree, this is no accident; the NBER uses payroll employment, along with real gross domestic product (GDP), real income, industrial production, and real manufacturing and trade sales, as one of its key indicators for dating cycle turning points. There is no fixed rule as to how these variables are weighted, although real GDP and other output measures get “considerable weight.” See the discussion at www.nber.org/cycles/recessions.html.

³The survey on which these numbers are based, the payroll survey, reports hours data for production or nonsupervisory workers only. By industry, this would include production workers in mining and manufacturing, construction workers in construction, and nonsupervisory workers in service-providing industries.

⁴For example, see Meltzer (2003). Articles and op-eds questioning the payroll survey’s accuracy have been published in numerous magazines and newspapers over the last year.

⁵On the other hand, 15-year-olds are included in the payroll survey but not the household survey.

⁶The household employment series are smoothed to account for population control corrections that are incorporated discretely into the reported figures.

⁷The adjusted series is from BLS (2004). See Nardone et al. (2003) for a detailed discussion and attempted reconciliation of the gap that developed during the 1990s and BLS (2004) for a discussion of the recent data.

⁸To quantify the importance of sample size, any monthly change above 108,000 jobs is considered statistically different from no change in the payroll survey. To claim statistical significance in the household survey, the monthly change must be greater than 290,000. See BLS (2004).

⁹The UI records provide a count of the number of employees covered by unemployment insurance laws. They cover over eight million establishments and nearly 97 percent of total nonfarm employment. The BLS uses alternative sources for the population not covered by UI. This full employment count is benchmarked to the sample-based counts in the payroll survey once a year. See www.bls.gov/ces/cesmetho.htm#10 for more details about the UI records, benchmarking, as well as the models used to estimate firm births and deaths prior to the UI benchmarking.

¹⁰The empirical evidence on this point is quite thin though. Hanson and Spilimbergo (1999) use U.S.–Mexico border apprehensions to show that illegal immigration is quite responsive to changes in U.S. and Mexican real wages.

¹¹One bit of corroborating evidence is that only five million immigrants applied for green card visas in 2003, roughly 60 percent below a typical year’s applicant pool.

¹²Population controls adjustments provided by the 2000 Census increased the level of household employment by 1.7 million, narrowing but not fully eliminating the cumulative 1990s gap between the household and payroll employment counts.

¹³For example, Meltzer (2003), Hilsenrath (2003), and Kudlow (2004).

¹⁴These calculations are based on the raw household employment figures and therefore do not smooth out the population adjustments. This adjustment would have little impact on the results reported below. We report quarterly figures because of concern about the variability of the monthly figures.

¹⁵Rissman uses the *National Longitudinal Survey of Youth*, a large panel data set that follows individuals over time. The individuals are interviewed annually from 1979–94 and again in 1996 and 1998. The participants range in age from 14 to 22 in 1979.

¹⁶We use the 16 and over population and exclude agricultural workers. The results are robust to reasonable changes to both of these groups.

¹⁷By specifying the regression equation in logs, the coefficient estimates have an intuitive interpretation as an elasticity. A 1 percent increase in the unemployment rate is associated with a β percent increase in self-employment. Furthermore, by framing the regression in logs we avoid the problem of having predicted values of self-employment less than 0 or larger than 1.

¹⁸An alternative approach to control for secular trends is including year dummies. Unfortunately, this sops up too much of the variation and results in highly imprecise parameter estimates. In future work, we plan to explore these specification issues further.

¹⁹Without a service, the current filing fee is \$150, plus an additional franchise fee of at least \$25 and a \$25 fee for reserving the company name. Incorporation in Delaware costs about half as much. Additionally, owners may wish to consult with an attorney or CPA and will likely need to open a bank checking account at some cost. Unfortunately, we do not have information about the cost of dissolving a corporation.

²⁰As a barometer consider the average start-up cost for a small business. In the 1992 *Characteristics of Business Owners Survey*, 27 percent of respondents report no start-up capital and another 34 percent report start-up capital of less than \$5,000. Only 10 percent had over \$50,000 in capital at start-up. So the fees and opportunity costs associated with incorporation would be substantial relative to the start-up cost of many businesses.

²¹Limited liability means that if the corporation suffers losses, the corporation itself must bear those losses. The personal assets of the individual shareholders are not at risk. This is in contrast to unincorporated businesses, which do not offer personal liability protection to their owners or employees. Other advantages of incorporation include the ease of ownership transfer. Corporate shares may be transferred without dissolving the corporation. Additionally, incorporation offers different tax options that encourage pension, profit sharing, and stock option plans. The main disadvantage of a corporation is that the dividends paid by corporations to their shareholders are taxed twice—once as income to the corporation and again as dividend income to the shareholder. Some corporate structures, such as a Subchapter S Corporation, eliminate this double taxation. The S Corporation allows certain income, deductions, and losses to be passed through the corporation to the individual tax return of each shareholder. For unincorporated businesses, net income flows to the owner and the owner is taxed.

²²Similar results arise using a different definition of the local labor market—metropolitan statistical area (MSA). Because the signal to noise ratio is low in smaller MSAs with few observations, we include only the largest 100 MSAs and do the analysis at an annual frequency.

²³Technically, this unemployment rate includes agricultural workers. However, adjusting the unemployment rate in a rough way to be more consistent with a nonfarm measure would have a small impact on the change in the unemployment rate over the last two years.

²⁴This assumes the effect on the incorporated firms is zero, which cannot be rejected by our estimates. The total self-employment predictions are derived by multiplying the relevant elasticity by the percentage change in the unemployment rate, the share of workers in that category, and the initial self-employment rate. For example, the 0.05 percentage point figure for the WLS estimate reported in table 2 is calculated as 0.28 (the unincorporated elasticity when the state-specific time trend is used) \times 0.405 (the percentage change in the unemployment rate) \times 0.67 (the share of self-employed who own unincorporated firms) \times 0.068 (the self-employment rate in the first quarter of 2001).

²⁵Lilien (1982) argues that variation in the pace of sectoral reallocation may contribute significantly to unemployment fluctuations. Davis and Haltiwanger (1990) and Davis, Haltiwanger, and Schuh (1996) discuss the role of job creation and destruction over the business cycle.

²⁶The costs to the affected workers can also be quite significant as is discussed, for example, in Jacobson, LaLonde, and Sullivan (1993a and 1993b).

²⁷See Segal and Sullivan (1997a and 1997b) for more on the temporary services industry. Schreft and Singh (2003) also discuss the role of the industry over the business cycle.

²⁸Moreover, as Segal and Sullivan (1997b) show, because the industry has very high turnover, a significantly higher fraction of workers have a temporary services job over a period of a year or two.

²⁹Many temp firms also provide general computer training, a service that Autor (2001a) argues allows these firms to screen potential employees for their clients.

³⁰See Autor (2001b) for a discussion of the role of the Internet in the labor market.

³¹See, for example, Wessel (2004).

³²However, Gruber and Krueger (1991) and Anderson and Meyer (2000) find that the additional costs of workers compensation and unemployment insurance are largely passed onto employees in the form of lower wages. A related objection to the health costs theory is that medical insurance costs also rose rapidly in the early 1980s, but employment growth was very rapid following the recession trough of 1982. A possible response to this challenge is to argue that the lower inflation environment that has prevailed since the late 1980s makes it more difficult for firms to offset medical insurance cost increases with lower wage and salary growth. This would be the case if it were easier to reduce nominal wage growth from, say, 8 percent to 6 percent in a high inflation environment than to reduce wage growth from, say, 3 percent to 1 percent in a low inflation environment. This argument relies on some form of “money illusion,” which we do not find very appealing. However, researchers have found evidence of nominal wage rigidities that may be consistent with such effects—for example, Altonji and Devereux (2000) and Lebow et al. (1999).

³³Unfortunately, these data only go back to 1986. Thus, we don’t know whether health care was a high or low percentage of employment costs in the periods of rapid employment growth following earlier recessions.

³⁴Alternatively, they might tilt their hiring towards fewer more highly skilled workers.

³⁵A somewhat similar theory is that increased requirements to fund pensions for retired workers impose an increased burden on firms that is restraining hiring. However, retiree pension obligations are fixed costs that do not depend on current employment or production levels. Thus, basic economic theory predicts they should have no effect on firms' decisions, except perhaps their decisions on whether to stay in business. Some researchers have found evidence that accounting changes can affect aspects of firm behavior such as investment. However, we are skeptical that such costs have played a significant role in the jobless recoveries.

³⁶The data are based on special education supplements to the household survey given every October.

³⁷The emergence of just-in-time hiring practices might also have damped demand growth if it led firms to delay hiring and rebuilding inventories. Alternatively, a lack of employment growth caused by the growth of just-in-time hiring practices could have had a deleterious effect on consumer or business confidence.

REFERENCES

- Aaronson, Daniel, Ellen Rissman, and Daniel Sullivan**, 2004, "Can sectoral reallocation explain the jobless recovery?," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 28, Second Quarter.
- Altonji, Joseph, and Paul Devereux**, 2000, "The extent and consequences of downward nominal wage rigidity," *Research in Labor Economics*, Vol. 19, Amsterdam: Elsevier Science, pp. 383–431.
- Anderson, Patricia, and Bruce Meyer**, 2000, "The effects of the unemployment insurance payroll tax on wages, employment, claims, and denials," *Journal of Public Economics*, Vol. 78, No. 1-2, pp. 81–106.
- Autor, David**, 2001a, "Why do temporary help firms provide free general skills training?," *Quarterly Journal of Economics*, Vol. 116, No. 4, pp. 1409–1448.
- _____, 2001b, "Wiring the labor market," *Journal of Economic Perspectives*, Vol. 15, No. 1, pp. 25–40.
- Barrow, Lisa**, 2004, "Is the official unemployment rate misleading? A look at labor market statistics over the business cycle," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 28, Second Quarter.
- Bureau of Labor Statistics**, 2004, "Employment from the BLS household and payroll surveys: Summary of recent trends," Washington, DC, available at www.bls.gov/cps/ces_cps_trends.pdf, posted March 5, 2004.
- Davis, Steven, and John Haltiwanger**, 1990, "Gross job creation and destruction: Microeconomic evidence and macroeconomic implications," *NBER Macroeconomics Annual 1990*, Cambridge, MA, and London: MIT Press, pp. 123–168.
- Davis, Steven, John Haltiwanger, and Scott Schuh**, 1996, *Job Creation and Destruction*, Cambridge, MA, and London: MIT Press.
- Groshen, Erica, and Simon Potter**, 2003, "Has structural change contributed to a jobless recovery?," *Current Issues in Economics and Finance*, Federal Reserve Bank of New York, Vol. 9, No. 8, August.
- Gruber, Jonathan, and Alan Krueger**, 1991, "The incidence of mandated employer-provided insurance: Lessons from workers' compensation insurance," in *Tax Policy and the Economy*, Vol. 5, David Bradford (ed.), Cambridge, MA: MIT Press, pp. 111–144.
- Hanson, Gordon, and Antonio Spilimbergo**, 1999, "Illegal immigration border enforcement and relative wages: Evidence from apprehensions at the U.S.–Mexico border," *American Economic Review*, Vol. 89, No. 5, pp. 1337–1357.
- Henderson, Nell**, 2004, "Economists say recession started in 2000," *Washington Post*, January 22.
- Hilsenrath, Jon**, 2003, "Self-employed boost the economic recovery," *Wall Street Journal*, December 1.
- Jacobson, Louis, Robert LaLonde, and Daniel Sullivan**, 1993a, "Earnings losses of displaced workers," *American Economic Review*, Vol. 83, No. 4, pp. 685–709.
- _____, 1993b, "Long-term earnings losses of high seniority displaced workers," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 17, No. 6, pp. 2–20.
- Kudlow, Lawrence**, 2004, "This recovery's not broken," *National Review On-line*, March 8.

Lebow, David, Raven Saks, and Beth Anne Wilson, 1999, "Downward nominal wage rigidity: Evidence from the Employment Cost Index," Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, working paper, No. 31.

Lilien, David, 1982, "Sectoral shifts and cyclical unemployment," *Journal of Political Economy*, Vol. 90, No. 4, pp. 777–793.

Meltzer, Alan, 2003, "A jobless recovery?," *Wall Street Journal*, October 1.

Nardone, Thomas, Mary Bowler, Jurgen Kropf, Katie Kirkland, and Signe Wetrogan, 2003, "Examining the discrepancy in employment growth between the CPS and CES," Bureau of Labor Statistics, working paper.

Rissman, Ellen R., 2003, "Self-employment as an alternative to unemployment," Federal Reserve Bank of Chicago, working paper.

_____, 1997, "Measuring labor market turbulence," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 21, No. 3, pp. 2–14.

Schreft, Stacey, and Aarti Singh, 2003, "A closer look at jobless recoveries," *Economic Review*, Federal Reserve Bank of Kansas City, pp. 45–73.

Segal, Lewis, and Daniel Sullivan, 1997a, "The growth of temporary services work," *Journal of Economic Perspectives*, Vol. 11, No. 2, pp. 117–136.

_____, 1997b, "Temporary services employment durations: Evidence from state UI data," Federal Reserve Bank of Chicago, working paper.

StataCorp, 2001, *Stata Statistical Software: Release 7.0 Manual*, College Station, TX: Stata Corporation.

Wessel, David, 2004, "Health care costs blamed for hiring gap," *Wall Street Journal*, March 11, p. A2.