Visualizing Labor Market Growth with Many Indicators

HUMAN CAPITAL CURRENTS

The Atlanta Fed's Labor Market Spider Chart, reproduced below, is a visual tool used to gauge the recovery in 13 labor market indicators since the end of the 2007–09 recession. For instance, at 5.7 percent in January 2015, the unemployment rate has reversed just over 80 percent of its increase from its prerecession value (4.8 percent in the fourth quarter of 2007) to its postrecession peak (9.9 percent in the fourth quarter of 2009). The recovery across the other labor market indicators has been somewhat uneven. For example, job openings have recovered more than unemployment perhaps reflecting an outward shift in the Beveridge curve, which depicts the relationship between unemployment and the job vacancy rate.

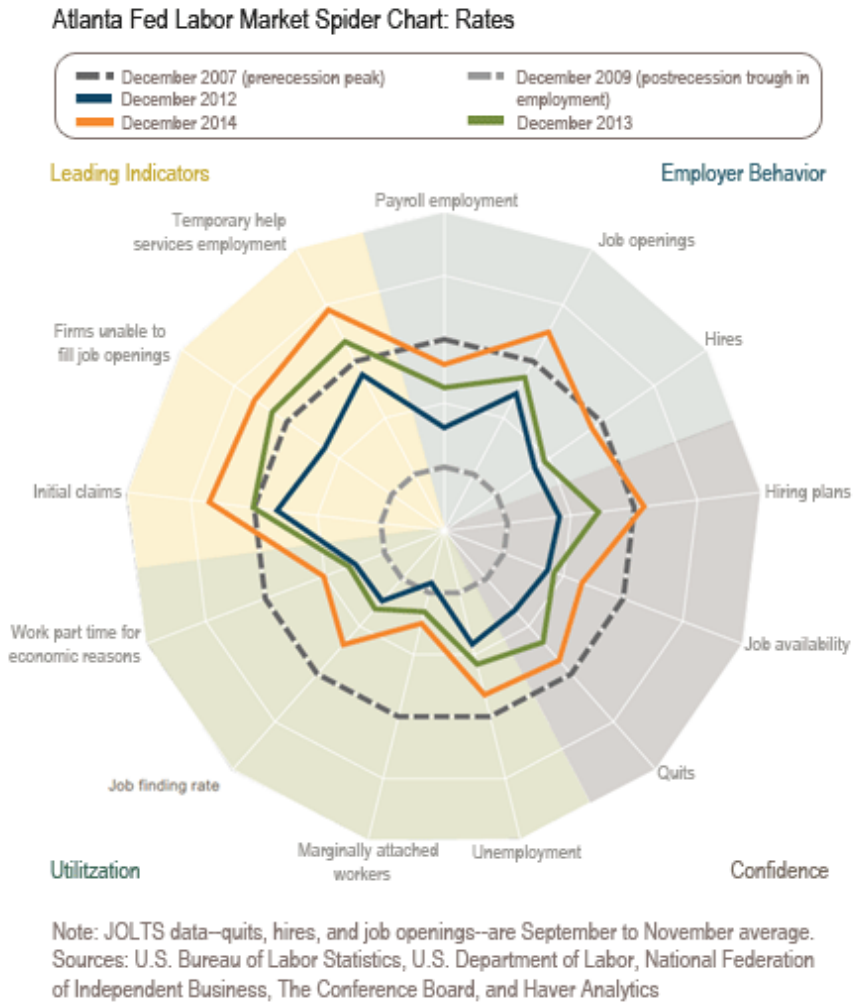

Note: JOLTS data--quits, hires, and job openings--are September to November average.
Sources: U.S. Bureau of Labor Statistics, U.S. Department of Labor, National Federation of Independent Business, The Conference Board, and Haver Analytics

While very useful, the spider chart is probably not the best tool for visualizing changes or growth rates at frequencies longer than six months. The change in nonfarm payroll employment is the most popular measure of short-run labor market growth. Other "combination measures" blend many labor market indicators into a single time series, like the Conference Board's Employment Trends Index (CBET), the "Employment, Unemployment, and Hours" subcomponent of the Chicago Fed National Activity Index (CFNAI-Labor), and the Labor Market Conditions Index (LMCI) from economists at the Federal Reserve Board of Governors. (Economists at the Federal Reserve Bank of Kansas City have also done recent work in this area.) Chart 1 plots seven-month centered moving averages of these series normalized to have mean 0 and standard deviation 1 over the sample period starting in March 1974. The data are truncated below −0.5 to draw attention to growth during expansionary periods. We see the three series are highly, but not perfectly, correlated during expansions.

Chart 1: Centered moving averages of indicators of labor market growth

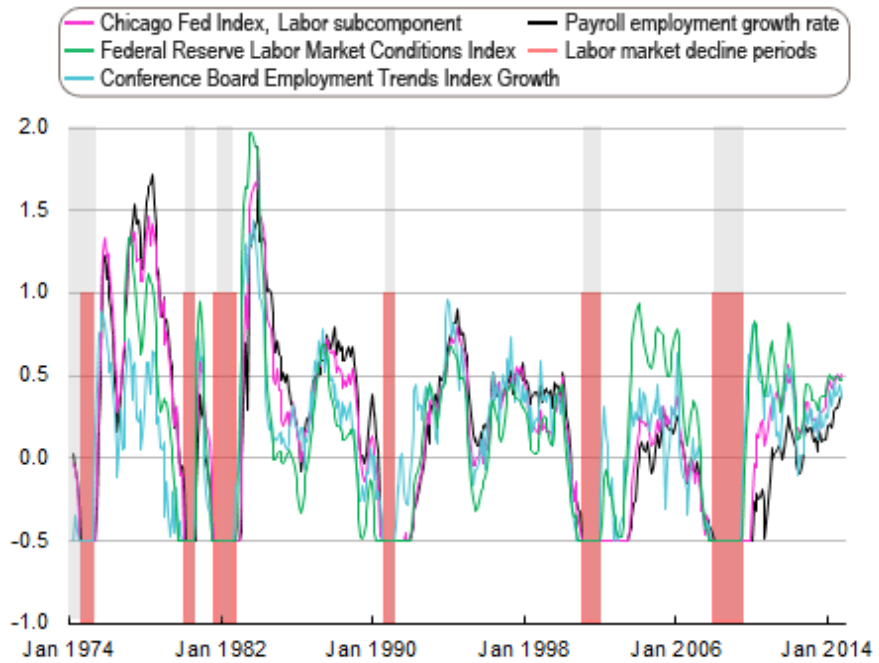

Note: Series have been truncated below at -0.5 to emphasize behavior in expansions. Data represent a standardized seven-month centered moving average.

Sources: Bureau of Labor Statistics, The Conference Board, Federal Reserve Board and Federal Reserve Bank of Chicago

Combination measures of employment trends like those shown in chart 1 are generally based on a statistical model that can appear to be a black box to the uninitiated. An alternative is to incorporate many indicators in one chart, as the spider chart does. This approach has the advantage of transparency, though a number of issues must be confronted as well. Like an episode of *The McLaughlin Group*, we handle each of these in turn as we present our charts.

Issue 1: Some data are noisier than others. Most of our labor market indicators come from surveys of various sizes. For example, the Job Openings and Labor Force Turnover Survey (JOLTS) samples 16,000 establishments each month, and the Establishment (payroll) Survey samples about 144,000 businesses and government agencies each month. (Both of these surveys are conducted by the U.S. Bureau of Labor Statistics.) Not surprisingly, monthly changes in payroll survey data are much less noisy than monthly changes in JOLTS data. To systematically assess the noisiness of a series, we again refer to chart 1. The red bars denote periods where the normalized seven-month centered moving averages of CBET, CFNAI-Labor and LMCI are all below −0.5. (Prior to November 1976, when the centered moving average of LMCI is not available, we identify a labor month decline as a month when the centered moving average of the other two indicators is below −0.5.) These "labor market decline" months correspond quite closely with recession dates established by the National Bureau of Economic Research. In his 2008 working paper, Edward Leamer identified a similar relationship between recession dates and six-month changes in each of payroll employment, the unemployment rate, and industrial production. (We use seven months to make the moving average symmetric.) We identify "labor market expansion" months as periods where none of CBET, CFNAI-Labor, and LMCI are negative in chart 1. To determine the noisiness of the centered *n*-month change of a labor indicator we compute the proportion of times where it is better in "labor market decline" months than in "labor market expansion" months. (If n is even, the centered change assigns a relative weight of 1/2 to the n/2 month lead and lag and a relative weight of 1.0 to the months in between. We provide additional details in the appendix.) If the series was white noise, we would expect this proportion to be 50 percent. The noisiness proportions *p* are converted to signal measures s with the formula $s = (\frac{50\% - p}{50\%})^{\wedge 2}$. In each of charts 2a through 2d, the areas of the points are proportional to the signal measures of the *n*-month change in the series. For the one-month change chart many of the bubbles are small, suggesting, for example, that one should not read too much into a single month's change in the labor force participation rate (LFPR). Each *n*-month change is normalized to have mean 0 and standard deviation 1 over the sample period used (January 1989 to present). We convert the *z*-score units to payroll equivalent (PE) units as described in the appendix. According to chart 2d, if payroll employment's 12-month *z*-score was the same as for the LFPR, total payroll employment would have grown only slightly faster than its trend rate (117,000 jobs per month). But, these charts are based on the assumption that labor trends are constant over time, while in reality...

Chart 2a: Standardized 1-Month Changes

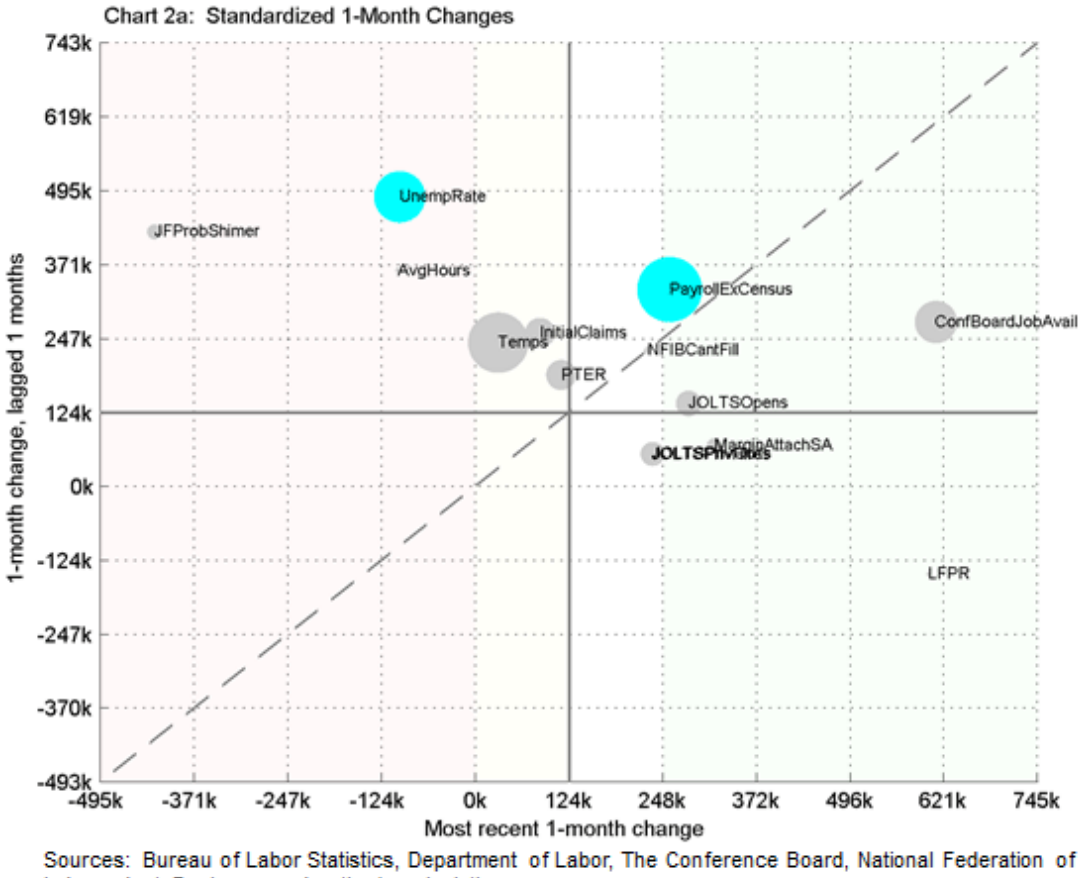

Sources: Bureau of Labor Statistics, Department of Labor, The Conference Board, National Federation of Independent Business, and author's calculations

Chart 2b: Standardized 3-Month Changes

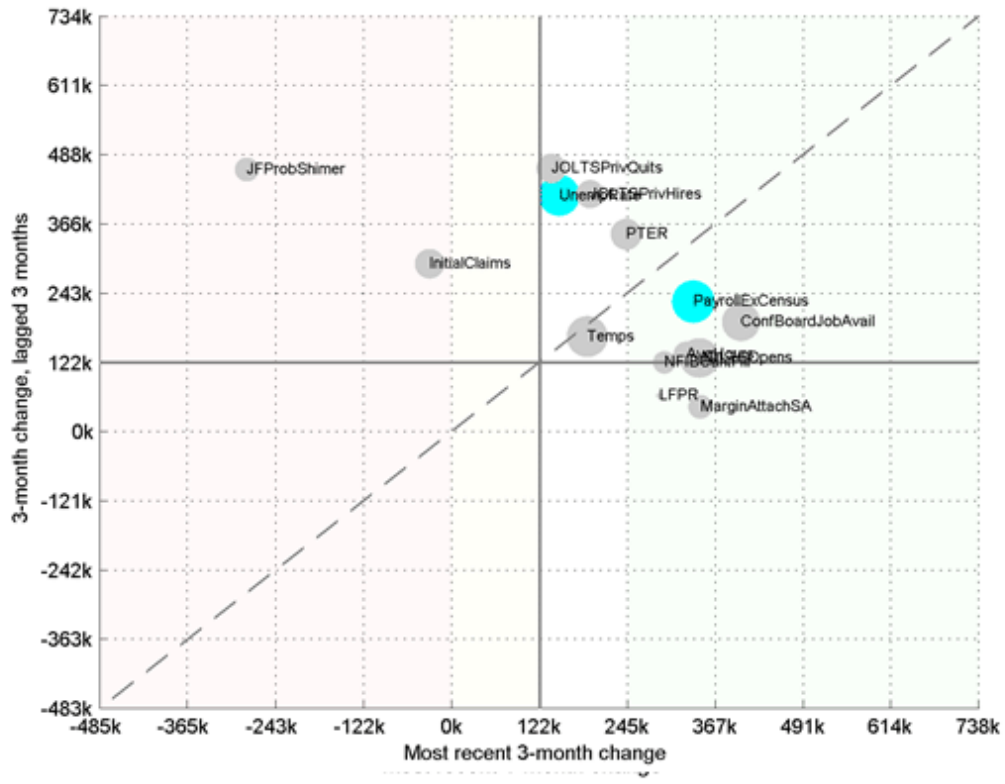

Sources: Bureau of Labor Statistics, Department of Labor, The Conference Board, National Federation of Independent Business, and author's calculations

Chart 2c: Standardized 6-Month Changes

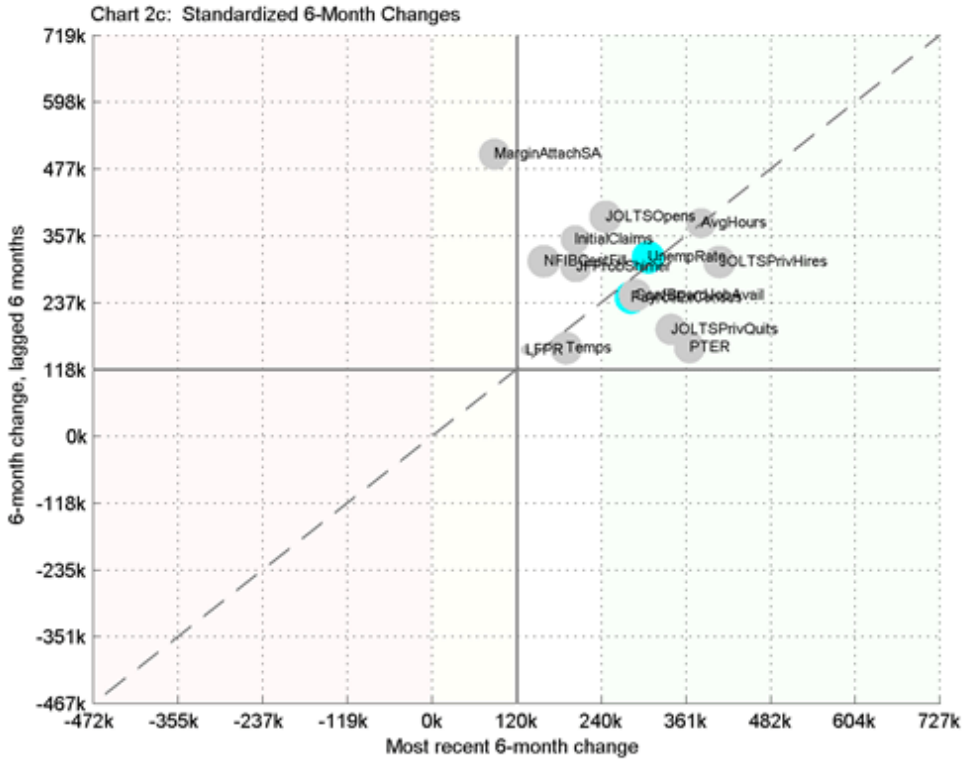

Sources: Bureau of Labor Statistics, Department of Labor, The Conference Board, National Federation of Independent Business, and author's calculations

Chart 2d: Standardized 12-Month Changes

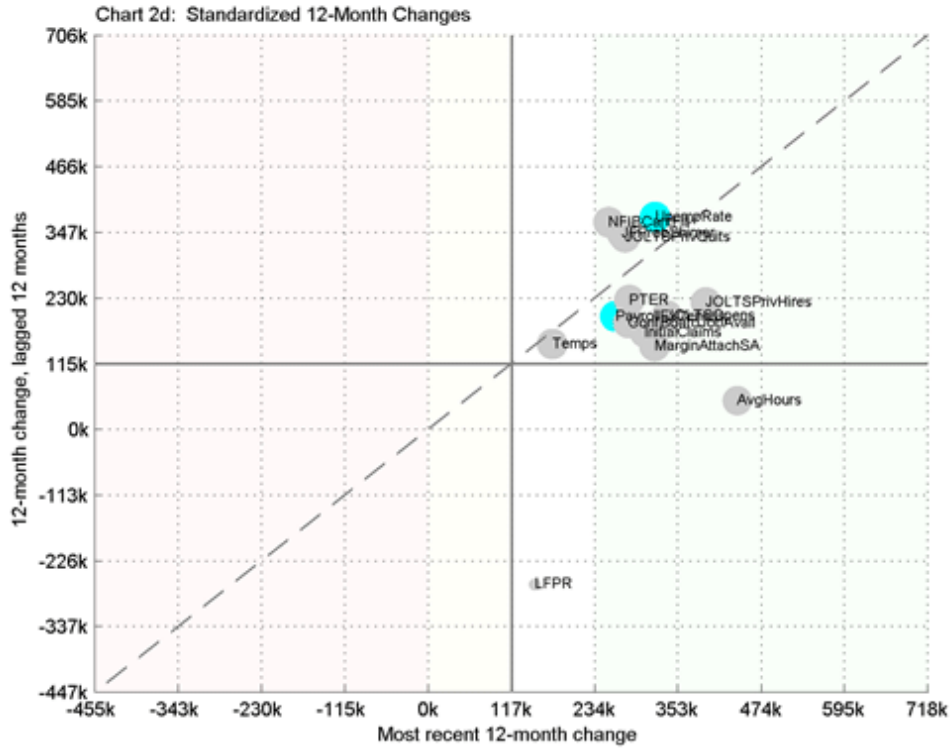

Sources: Bureau of Labor Statistics, Department of Labor, The Conference Board, National Federation of Independent Business, and author's calculations

Issue 2: Labor trends can vary over time. This might be expected with the aging of the baby-boom generation and the slowing growth rate in the working age population. One might be tempted to use an off-the-shelf detrending method like the Hodrick-Prescott (HP) filter. (The HP filter is a standard algorithm for decomposing a time series into trend and cyclical components. The filter has a parameter lambda that determines the smoothness of the trend; we set this to 129,600, a standard value for monthly data.) However, this approach implies that the trend growth rate of payroll employment was −0.7 percent in early 2009 and negative from mid-2007 to mid-2010. We deem this as unrealistically low. The potential labor force grew at least 0.5 percentage point each year during 2007–11, according to the Congressional Budget Office. As an alternative, we assume that the trend unemployment

rate is constant and estimate the time-varying trend in each series as the rate it would have to grow to keep the unemployment rate fixed, using a "constant gain least squares" algorithm described in 2006 research by Branch and Evans (the appendix contains details). Chart 3 shows growth trends for nonfarm payroll employment and LFPR.

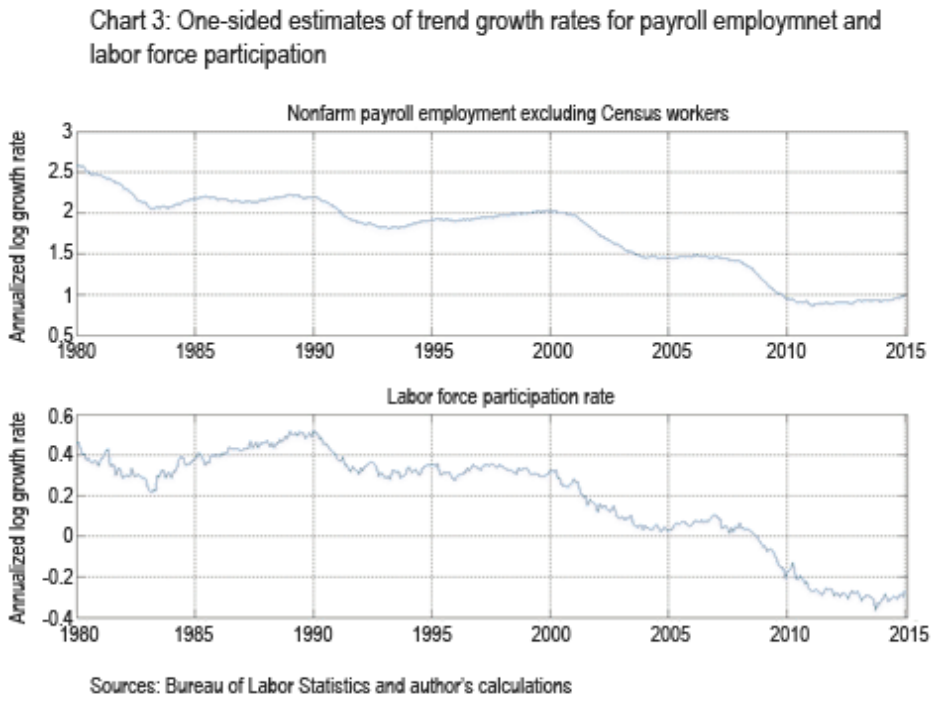

In chart 4A, we modify chart 2d by subtracting the time-varying trends from each series instead of the 1989–2014 sample averages. Chart 2d is reproduced as chart 4b to facilitate comparison. We see that LFPR moves somewhat closer to the other points, but the overall message is not changed. The improvement in the unemployment rate has been stronger than the improvement in many of the other labor market variables, partly because LFPR has not grown much faster than trend.

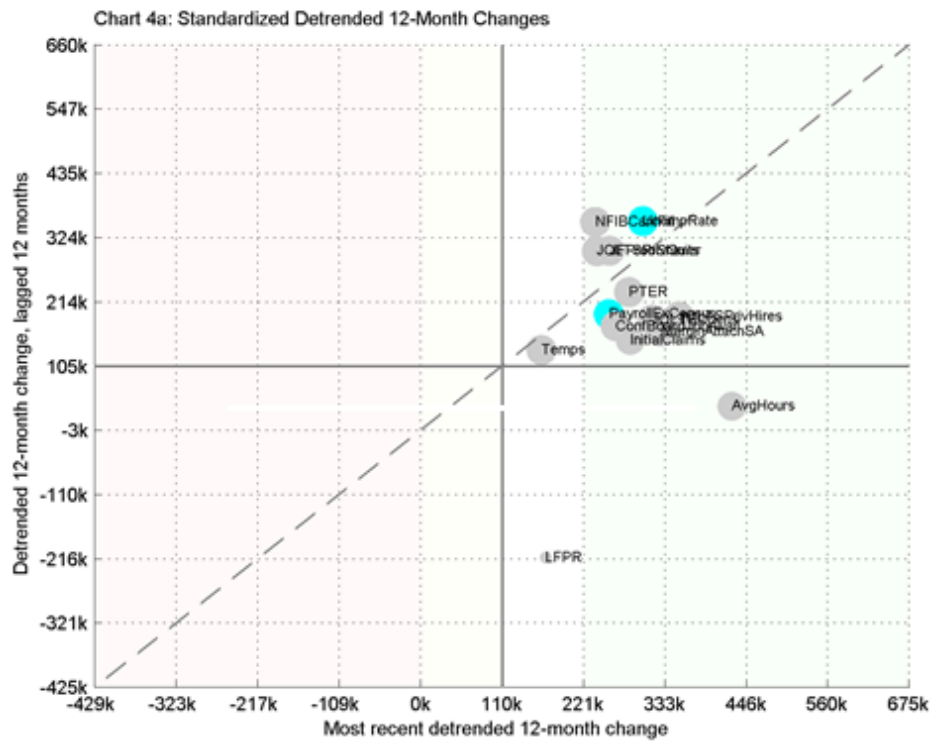

Sources: Bureau of Labor Statistics, Department of Labor, The Conference Board, National Federation of Independent Business, and author's calculations

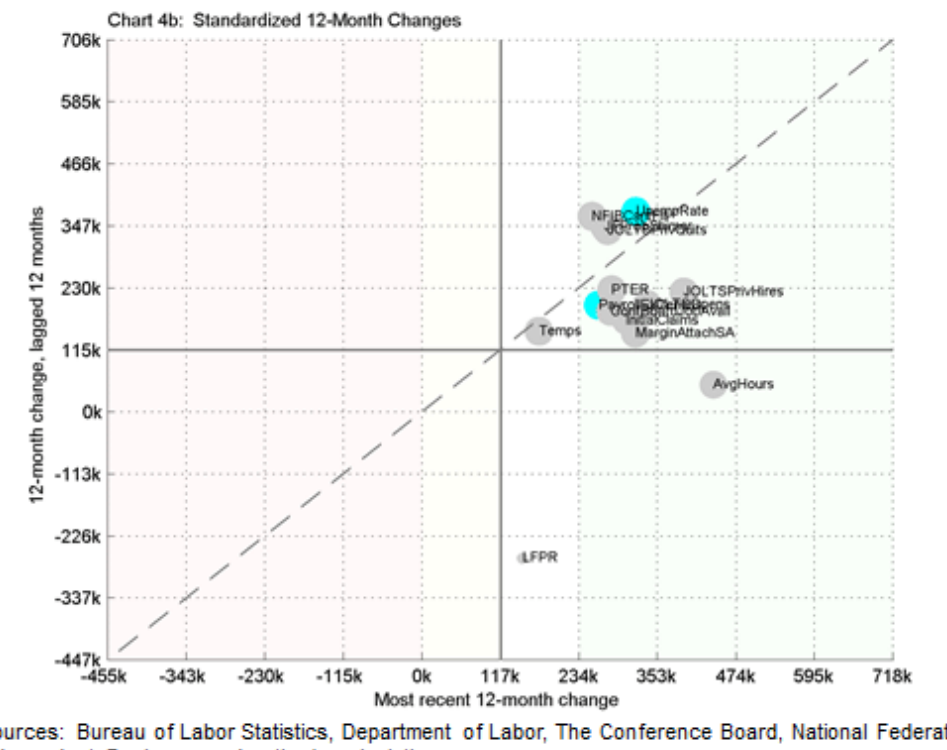

Sources: Bureau of Labor Statistics, Department of Labor, The Conference Board, National Federation of Independent Business, and author's calculations

Issue 3: It's hard to display many time series in one plot. The spider and bubble charts we have shown are essentially snapshots of conditions at particular points in time. Although one can create *movies* over a time interval, these are not well suited for a static document like this. A heat map, as shown in charts 5a through 5d, can be used to show the behavior of many labor market indicators over time. (A Federal Reserve Bank of San Francisco Economic Letter by Daly, Hobijn and Bradshaw (2013) used a somewhat similar approach.) In these charts, white is neutral, green is "good,"

and red is "bad." The intensity of the color is inversely related to how close to neutral the indicator is. (The appendix describes exactly how the colors and intensities are assigned.) In short, for each indicator we ask the question "If the labor market was growing at a rate consistent with no change in the unemployment rate, how likely would the observed n-month change in the indicator be?" For example, the 826,000 nonfarm payroll loss observed in March 2009 would be almost impossible in a stable unemployment rate environment, so this square is colored dark red in chart 5a. On the other hand, relatively few squares for hires in chart 5a are colored dark red because of the large sampling error for the JOLTS survey. Apparently, the only one-month changes that consistently give a clear signal of a recession are the Establishment Survey variables shown in the upper third of the chart. (To a lesser extent, one-month changes in the unemployment rate also signal recessions.) But...

Issue 4: Real-time data differ from revised data. All of the previous charts use the current vintage of data. How do the pictures change with real-time data? For the most part, not too much, except for the JOLTS-based series. This is probably not surprising since the JOLTS data are relatively recent and have gone through a number of important methodological changes.

Where are we now? Overall, these various charts suggest that the labor market is improving. In chart 2D, there are more indicators below the 45-degree line than above, suggesting that the pace of labor market improvement has accelerated over the past 12 months relative to the prior twelve months. Chart 1 suggests that payroll employment growth is a fairly good metric of labor market improvement, and our bubble charts and heat maps are also consistent with that interpretation.

However, payroll growth is not a perfect measure of modest changes in labor market momentum. Research from the BLS shows that the average absolute benchmark revision to the March level of employment for 2004–14 was 314,000 jobs, or roughly 26,000 jobs per month. So distinguishing a labor market truly expanding 150,000 jobs per month from one expanding 175,000 jobs a month is probably impossible in real time. The problem is even more pronounced for frequencies higher than 12-month changes because issues related to seasonal adjustment of payroll employment, and revisions from the first preliminary estimate of payroll employment growth to the "final" (third) will also matter. (Research from the Brookings Institution and the Philadelphia Fed offers good discussions of these issues.)

Finally—and perhaps most importantly—even if we knew for certain how payroll employment would grow, there would still be a fair amount of uncertainty about the associated unemployment rate. In the appendix we show a decomposition of the *n*-month change in the unemployment rate as a sum of five contributions from (1) payroll employment growth, (2) population growth, (3) the discrepancy between payroll and household employment growth, (4) LFPR growth, and (5) a small approximation error. If we perform a thought experiment where a time traveler has each of the five contribution terms above except (3) and (4), and the time traveler forecasts these missing terms to be zero, we get charts 6a and 6b. (These default assumptions, incidentally, are very close to those used in the Atlanta Fed's Jobs Calculator.) The time traveler's average absolute miss on the 12-month change in the unemployment since 1995 would be 0.66 percentage point even with perfect foresight of payroll employment and population. For the 36-month change, the average absolute unemployment rate miss would be 1.78 percentage points. If the time traveler were savvy enough to use a constant, their knowledge of payroll employment growth as well as *ex post* regression coefficients to improve their forecasts of labor force participation and household employment, the 12-month average absolute error would go down to 0.37 percentage point, and the 36-month average absolute error would go down to 0.83 percentage point. These are clearly nontrivial sources of error for someone trying to pinpoint the date the unemployment rate will fall below a particular threshold.

By Pat Higgins, senior economist in the research department of the Atlanta Fed

---

Wage Growth Tracker