

## **Assessing Conditions in the Labor Market**

Stephanie Aaronson, Bruce Fallick, Christopher Nekarda, and William Wascher  
Division of Research and Statistics  
November 30, 2012

### **Introduction**

In the FOMC statements following the September and October meetings, the Committee stated that it will be looking for a substantial improvement in the outlook for the labor market in deciding whether to continue or expand its purchases of longer-term assets. In addition, the Committee has been considering including language in the FOMC statement indicating that the lift-off of the funds rate is unlikely to commence until certain conditions prevail, such as the unemployment rate falling below a certain level. Both of these aspects of the FOMC statement elevate the explicit role of labor market conditions in determining the future course of policy and, in doing so, raise some important issues for policymakers to consider.

In particular, in its statement language about possible future longer-term asset purchases, the Committee has not explicitly indicated how it intends to judge whether the outlook for the labor market has improved substantially. Similarly, in the Committee's discussion of thresholds, participants have expressed a range of views as to whether the threshold for economic conditions should refer only to the unemployment rate or to a broader set of indicators. In this context, the Committee may need to consider whether to be more explicit about what indicators it would use to assess whether a substantial improvement in current or prospective labor market conditions has occurred and how best to convey that information to the public.

This memo is intended to provide some background information that might be helpful in making such judgments. We first review the evidence on the merits of the unemployment rate as a summary measure of labor market conditions. If one wants a single indicator of labor market slack, evidence suggests that the unemployment rate is a good choice. However, the unemployment rate does not contain all of the information policymakers might want to consider, and we next discuss how other information can usefully complement the unemployment rate in describing the state of the labor market at any particular time. We then turn to the question of how best to judge the outlook for the labor market and, in particular, whether the Committee could be reasonably confident about projecting an improvement in the labor market based on the past behavior of the unemployment rate and other indicators. Our analysis suggests that a variety of information can be helpful in this regard, but that forecasts of the unemployment rate at a two-quarter horizon have relatively wide confidence bands—with the Tealbook, which has the smallest errors of those considered here, having a 90-percent confidence interval of approximately plus or minus  $\frac{1}{2}$  percentage point.

## **Is the unemployment rate a good summary measure of current labor market conditions?**

The unemployment rate is arguably the best single summary indicator of the state of the labor market, for both conceptual and statistical reasons. From a conceptual standpoint, the unemployment rate attempts to measure the extent of underutilization of available labor resources and thus has a close connection with the full-employment leg of the Federal Reserve's statutory mandate; indeed, the discussion of maximum employment in the Committee's January 2012 statement on longer-run goals and policy strategy includes a reference to the participants' assessments of the longer-run normal rate of unemployment. In addition, the unemployment rate is released on a timely basis and is probably the most widely cited statistical measure of economic activity in the United States. Focusing on such a well-known and timely indicator can help to enhance communication between the Committee and the public.

Aside from these considerations, the unemployment rate also ranks highly as a summary indicator of labor market conditions from a statistical standpoint. It is highly correlated with a broad set of labor market indicators, which, together, might be thought to represent labor market conditions along a number of dimensions. This is demonstrated in exhibit 1, which plots the six-month change in the unemployment rate along with the first principal component of 37 labor market indicators. The two series move fairly closely together; the correlation between them is 0.9.<sup>1</sup> The unemployment rate is also individually highly correlated with particular indicators that the staff and other analysts emphasize in assessing the state of the labor market, most notably payroll employment. Finally, staff research points to the unemployment rate as a reliable indicator of the overall state of the business cycle.<sup>2</sup>

Another advantage of the unemployment rate relative to other possible summary indicators is that historically it fluctuates around a relatively stable average value. Every labor market indicator must be measured against some notion of "normal;" in the case of the unemployment rate, the relevant standard is an estimate of the natural rate of unemployment. In the staff's view, the natural rate varies less than the corresponding concepts of "normal" for other prominent indicators. (For example, the "trend" in payroll employment has varied significantly over time given changing trends in labor force participation and population.) That said, the natural rate does change over time, and its level is difficult to estimate with confidence. Staiger, Stock, and Watson (1997) find that, within a Phillips curve framework, the 95-percent confidence interval for NAIRU estimates typically extends 1½ percentage points on either side of the point estimate, and is often even wider.<sup>3</sup> Moreover, even relatively low-frequency changes in the trend unemployment rate, which may be apparent in retrospect, are difficult to identify in real time. Nonetheless, the problem in identifying normal with regard to the unemployment rate is probably less acute than for other variables.<sup>4</sup>

---

<sup>1</sup> These results are qualitatively robust to a number of specifications and choices of indicators, but many variations are reasonable, and we would caution against placing too much emphasis on this particular model.

<sup>2</sup> Fleischman and Roberts (2011) and Nalewaik (2011).

<sup>3</sup> For a discussion of the Board's staff approach to estimating the natural rate of unemployment and the uncertainty surrounding those estimates, see the memo to the FOMC by Bruce Fallick and Jeremy Rudd, "The Staff's Assessment of Economic Slack," August 28, 2012.

<sup>4</sup> That said, as was noted in the recent memo sent to the FOMC on thresholds, the Committee's assessment of the natural rate would likely be an important consideration in choosing the value of an unemployment threshold, and thus a misperception of its true value could lead policymakers to keep the stance of monetary policy easier or tighter

### *Limitations*

Despite these advantages, there are instances in which the unemployment rate may not accurately characterize labor market slack. Notably, changes in the rate of labor force participation may, at times, distort the degree of labor underutilization implied by the unemployment rate, in the sense that there are more or fewer individuals available for work than the unemployment rate alone would indicate. Ideally, changes in the trend labor force participation rate that influence the unemployment rate, such as those that stemmed from the entry of the baby boom generation in the 1970s, would be captured in our estimate of the natural rate, and so leave the unemployment rate gap as a clean measure of labor market slack.<sup>5</sup> In contrast, demand-driven changes in the participation rate, such as those associated with worker discouragement, may cause the unemployment rate to move in ways that do not accurately reflect the changes in labor demand. For example, an increase in the demand for labor that led to the return of a large number of discouraged workers to the labor force would tend to mute the decline in the unemployment rate and thus mask some of the actual improvement in labor market conditions. Likewise, a decline in the unemployment rate driven by an increase in the number of unemployed workers becoming discouraged and therefore deciding to leave the labor force presumably would not be indicative of an improvement in the labor market.<sup>6</sup>

However, identifying the sources of changes in the labor force participation rate is difficult, especially in real time. As an alternative, some have suggested that the employment-population ratio, which abstracts from labor force participation, may be preferable to the unemployment rate as a measure of labor market conditions. In our view, however, this measure is overly broad in that it treats everyone who is not employed as being equally available for work; in this sense, proper interpretation of the employment-population ratio requires estimates of both the trend in labor force participation and the natural rate of unemployment. Instead, we find it more useful to refer to an intermediate measure of labor underutilization, such as the Bureau of Labor Statistics' (BLS) U-5 rate, which is derived from the same survey used to calculate the unemployment rate (the Current Population Survey, or CPS). This measure adds to the unemployed those persons who are both officially out of the labor force and meet the BLS's criteria for "marginal attachment" to the labor market.<sup>7</sup>

Another potential shortcoming in the unemployment rate is that unemployment, as defined by the BLS, is an all-or-nothing concept. In practice, however, labor market slack may also manifest

---

than intended (Eric Engen, David López-Salido, Jean-Philippe Laforte, Edward Nelson, William Nelson, Dave Reifschneider, and Robert Tetlow, "Using Thresholds to Clarify the Conditionality in the Committee's Forward Guidance for the Federal Funds Rate," October 16, 2012).

<sup>5</sup> Of course, not all changes in trend labor force participation need affect the unemployment rate. For example, the current aging of the population does not appear to be having first-order effects on the unemployment rate at this time.

<sup>6</sup> In a similar vein, persons entering the labor market in response to declines in wealth or to job losses among other members of their household may raise the unemployment rate. In bad times, these "added workers" are usually fewer in number than workers who are discouraged by job prospects (which we believe are probably understated by the BLS's official measure of discouraged workers).

<sup>7</sup> Persons marginally attached to the labor force are defined as those who currently are neither working nor looking for work but indicate that they want and are available for a job and have looked for work sometime in the past 12 months.

itself in persons working fewer hours per week than they would if they were able to find additional work. A measure such as the number of persons working part-time for economic reasons (from the CPS) addresses this additional margin of underutilization of labor, and the BLS's broadest published measure of labor underutilization (U-6) includes these individuals.

Given these limitations in scope, would we do better to concentrate on the U-5 or U-6 measures rather than the official unemployment rate? As a rule, we think not. The three measures are highly correlated, both in levels and in changes, and thus provide a similar signal about the relative cyclical position of the labor market most of the time. On those occasions when they diverge from the unemployment rate, U-5 and U-6 may indeed provide useful information, but in each case it is desirable to evaluate the content of that information before taking it on board.<sup>8</sup> Whereas the unemployment rate is defined with reference to relatively objective behavioral criteria concerning job search, both U-5 and U-6 rely upon more subjective reporting of a person's intentions and availability. For example, marginally attached workers are historically less likely to move into employment than are unemployed workers. And, many of those who contributed to the increase in the number of individuals working part-time for economic reasons during the recent recession had previously reported themselves as part-time for noneconomic reasons; that is, their part-time status did not change, only their reported contentedness with that status changed, possibly because of job losses or hours reductions for others in their households.<sup>9</sup> Therefore, while we view information on marginally-attached workers and part-time work as complementary to the unemployment rate, placing too much emphasis on those measures could be misleading.<sup>10</sup>

Finally, the unemployment rate is subject to various sources of measurement error, reducing the reliability of high-frequency movements as indicators of a change in labor market conditions. Sampling variation alone results in a 90-percent confidence interval for the one-month change in the unemployment rate of about plus or minus 0.2 percentage point. Perhaps more telling, historically about one-third of the months in which the published unemployment rate has changed have been followed by a change in the opposite direction the next month. As a result, we have greater confidence in changes in the unemployment rate that are sustained over several months.

### **Using additional indicators to complement the unemployment rate**

While the unemployment rate is generally the best indicator of labor market slack, this discussion of its limitations suggests that we might want to bring other indicators to bear in evaluating the state of the labor market in any particular situation. As noted above, other data from the CPS can be consulted to help determine whether the limited conceptual scope of the

---

<sup>8</sup> Since the beginning of the recent recession, U-5 has moved fairly closely with the unemployment rate. However, U-6 has diverged somewhat, reflecting a large increase in the number of persons working part-time for economic reasons that has been slow to reverse.

<sup>9</sup> Moreover, for those working part-time for economic reasons, only the additional hours they wish to work should be included in a measure of slack.

<sup>10</sup> Indeed, in contrast to the high level of part-time work in the CPS, the level of the average workweek from the payroll survey has risen over the past couple of years, perhaps suggesting that employers are substituting away from employment and into higher workweeks, and thus raising the unemployment rate above where it might otherwise be. It is not clear how to reconcile these seemingly conflicting signals.

unemployment rate may be reducing its signal value. However, these additional measures have their own limitations and do not address all of the issues of concern, including the limitations inherent in an over-reliance on just one survey. Accordingly, we think that the Committee would be well-served by using other labor market indicators to help filter the signal from the unemployment rate.

Conceptually, the most obvious other indicator to use for this purpose is the change in nonfarm payroll employment, which is derived from a separate survey of establishments with a sample that covers about one-third of all nonfarm payroll employment. Like the unemployment rate, payroll employment measures aggregate labor market activity. Estimates of payroll employment arrive with the same timeliness as the unemployment rate, and the measure is nearly as well-known to the public. Statistically, the change in payroll employment is also highly correlated with the first principal component of a large set of labor market indicators (see exhibit 2).<sup>11</sup> And, in our experience, a decline in the unemployment rate is more likely to accurately signal an improvement in labor market conditions if it occurs together with above-trend increases in payroll employment. For example, the strong pace of employment growth last winter gave us confidence that the decline in the unemployment rate at that time represented a genuine improvement in labor market conditions, whereas the lack of payroll employment growth seen in late 2009 raised doubts about the signal being sent by the decline in the unemployment rate at that time.

Of course, payroll employment is also an imperfect measure, and on some occasions the initial estimates of employment growth have turned out to be unreliable.<sup>12</sup> Moreover, a reference point for “trend” employment depends on several factors beyond the natural rate of unemployment that are also difficult to estimate, including the trend in labor force participation and the discrepancy between employment as measured in the payroll survey and employment as measured in the CPS.

Thus, although we typically give primacy to the unemployment rate and payroll employment for evaluating the state of the labor market, we also look to a variety of other indicators. These include measures of job loss, hiring rates, job vacancies, quit rates, and businesses’ and consumers’ expectations and current assessments of labor market conditions.

Measures of job losses and hiring, being highly cyclical themselves, provide corroborating evidence for changes in unemployment and employment; they also provide information on the gross flows of workers that underlie the net changes reflected in those headline numbers. Thus far, for example, the improvements we have seen during the current recovery have been mainly driven by reductions in job loss rather than increases in hiring—indeed, by one measure (the Job Openings and Labor Turnover Survey, or JOLTS), layoffs returned to pre-recession levels some time ago. This suggests that further sustained declines in the unemployment rate most likely will

---

<sup>11</sup> One might consider using the principal component from a well-chosen factor analysis itself as a primary indicator. However, in addition to being difficult to communicate to the public, such a measure strikes us as too mechanical by not allowing different indicators to be given different emphases at different times. To give just one example, initial claims for unemployment insurance do not currently figure heavily in our thinking because of the influence of Hurricane Sandy on those figures.

<sup>12</sup> To take a particularly notable example, the decline in payroll employment between March 2008 and March 2009 was eventually revised to be 900,000 larger than initially estimated.

be achieved only if there is further improvement in the rate of hiring. In addition to job finding and hiring rates themselves, we look to vacancy rates and hiring plans as potential indicators of the likelihood that such a pick-up is under way.

We also find that surveys of consumers' and businesses' assessments of the state of the labor market are useful indicators of aggregate labor market conditions. What one person sees in his or her individual sphere may not reflect the aggregate situation in the economy, but when averaged together, these perceptions may provide a good independent reading on the state of the aggregate labor market. Similarly, quit rates, which are highly procyclical, presumably reflect workers' assessments of their job prospects, as well as their actual success in finding another desirable job. Consumers' expectations of improvement or deterioration in employment and unemployment also have proven useful as a forward-looking indicator in our near-term models of employment growth.

Among the various labor market indicators that are available on a fairly timely basis, we pay particular attention to initial claims for unemployment insurance; job openings, hiring, and quit rates from the JOLTS; and employment expectations from surveys of consumers. However, a broader view from a wide range of indicators is likely to be the most useful in gauging the state of the labor market.

Indicators of economic activity as a whole, including real GDP and GDI growth, can also be useful in assessing labor market conditions. In particular, from the viewpoint of Okun's law, a decline in the unemployment rate is more likely to be an accurate indicator of improvement in the labor market if it is accompanied by above-trend growth in output. That said, we see output as a less reliable indicator than those that pertain directly to the labor market. The relationship between unemployment and output can be weak over short periods. In addition, real-time estimates of output growth are often subject to considerable revision. And, even after revisions, the labor market need not move in line with output growth. For example, in 2011 the unemployment rate fell considerably despite lackluster output growth, which, in our assessment, partly represented a catch-up for its precipitous rise around the turn of 2009. Finally, the growth rate of potential output depends upon yet more factors that are difficult to estimate.

Overall, then, we see the unemployment rate as the most useful single summary indicator of labor market conditions. However, we also think that it is important to look at additional indicators, foremost among which is nonfarm payroll employment.

### **Forecasting an improvement in the labor market**

The previous discussion points to a range of indicators that can be examined to assess the current state and recent trajectory of the labor market, and so identify a sustained improvement in labor market conditions. In addition, however, the Committee's recent statements have indicated that longer-term asset purchases will continue until there has been a substantial improvement in the *outlook* for the labor market. In this context, we briefly assess, using a variety of approaches, how well one can predict future improvements in the labor market, particularly in the unemployment rate, both in terms of whether an improvement already seen is likely to continue,

and whether an improvement can be reliably foreseen before it has begun.<sup>13</sup> Our analysis suggests that while the unemployment rate itself contains some predictive power for its future path, the confidence intervals around univariate forecasts are wide, and uncertainty remains high even after conditioning on other information about the state of the economy.

### *Univariate approaches*

The unemployment rate is fairly persistent (in a statistical sense), both in levels and in changes. Thus, the recent history of the unemployment rate should provide some information about its future path. A simple exercise illustrates this point. In a sample from 1978 to the present, we calculated the probability that the quarterly average unemployment rate would fall in coming quarters conditional on its behavior over the previous two quarters. To help eliminate measurement error, we considered increases or decreases of 0.15 percentage point or less, the median absolute change over this period, as little changed. In periods when the unemployment rate was little changed over the preceding two quarters, it fell in the subsequent two quarters roughly 40 percent of the time and had about equal chances of being little changed or increasing.<sup>14</sup> In cases in which the unemployment rate declined by more than 0.15 percentage point over two quarters, it subsequently declined further about 60 percent of the time. And when the unemployment rate declined by  $\frac{1}{2}$  percentage point or more, it subsequently fell further about 75 percent of the time.<sup>15</sup> This simple exercise suggests that the unemployment rate itself can provide useful information in determining the outlook for the labor market, but that a great deal of uncertainty remains. In particular, if the unemployment rate has been flat or has fallen only slightly, its history alone provides only a little guidance about whether it will decline in the future.

A more formal way to investigate the predictive content of recent changes in the unemployment rate is by estimating a simple model that uses the unemployment rate as a univariate predictor of itself. We examined one such model preferred by several previous authors, in which the current-month unemployment rate is a function of its previous two lags and the previous month's residual (Montgomery and others, 1998). To reflect the environment within which forecasters

---

<sup>13</sup> Recall that in the October Tealbook, the staff projection assumed that asset purchases would end in mid-2013. This assumption was keyed off of a prospective improvement in the labor market—specifically, the projection anticipated no noticeable decline in the unemployment rate from current levels by that date, but a  $\frac{1}{2}$  percentage point decline over the subsequent year.

<sup>14</sup> We also performed the calculations using real-time data and found similar results. Limiting the sample to periods in which the unemployment rate was above our estimate of its natural rate also yielded similar results, although in this sample, periods in which the unemployment rate was little changed were followed by declines in the unemployment rate about 55 percent of the time, somewhat more frequently than was the case for the sample as a whole.

<sup>15</sup> This means that in 25 percent of cases the unemployment rate did not continue falling. In recent years, there have been two notable episodes when this has occurred. Between 2010:Q4 and 2011:Q1 the unemployment rate fell 0.6 percentage point but was then little changed over the next two quarters as economic growth slowed following the tsunami in Japan, a rise in energy prices, and an intensification of the sovereign debt crisis in Europe. Similarly, the unemployment rate fell 0.9 percentage point between 2011:Q3 and 2012:Q1 but then leveled out over the spring and summer.

**Table A: Comparing Unemployment Rate Forecast Errors, 1997–2012**  
(Root-mean-squared error, percentage points)

<i>Forecast</i>	<i>Current quarter</i>	<i>One quarter ahead</i>	<i>Two quarters ahead</i>
1. Univariate model	0.34	0.51	0.84
2. VAR with payroll employment growth	0.26	0.38	0.62
3. Labor force flows model	0.11	0.34	0.59
4. Greenbook/Tealbook	0.14	0.32	0.51

Notes: Evaluation of all forecasts is based on the timing of 122 historical staff forecasts from January 1997 to September 2012. Models were estimated using real-time data and used only information available to the staff at the time of the projection.

must operate, we estimated our models and generated historical forecasts using real-time data.<sup>16</sup> Over the past 15 years or so, the model’s root-mean-squared error (RMSE) for one-quarter-ahead out-of-sample forecasts was over 0.5 percentage point, and the RMSE for two-quarter-ahead forecasts was over 0.8 percentage point (Table A, line 1).

### *Multivariate approaches*

A natural extension of the univariate approach is to try to improve the forecasts of the unemployment rate by incorporating additional information.<sup>17</sup> One such approach is to add in variables that provide independent readings on the labor market, such as payroll employment growth or initial UI claims. As shown by line 2 in Table A, a VAR that includes the unemployment rate and the change in nonfarm payroll employment performs a bit better than the univariate model described above.<sup>18</sup> While it is possible to improve the forecasting performance using pooling procedures and nonlinear models, the literature suggests that even the best of these

<sup>16</sup> The models we examine in this memo use monthly data and make monthly forecasts. To compare with the Tealbook forecast, we take the quarterly average of the monthly forecasts. Because the publication of the Tealbook is dictated by the date of the upcoming FOMC meeting, forecasts are made at different points in a quarter; as a result, some forecasts have no monthly labor market data for the current quarter whereas others have two months of data. Each historical model forecast was made using the same information available to the staff at the time the forecast was prepared.

<sup>17</sup> A different univariate approach is to directly account for asymmetric movements in the unemployment rate (increases tend to be sharper than decreases) by using a nonlinear model, such as a threshold autoregressive model or a Markov switching model. The staff currently does not use such a framework. Results in the literature on nonlinear methods are mixed. A comprehensive study by Stock and Watson (1999) concluded that linear univariate models generally dominated nonlinear models for out-of-sample forecasts of macroeconomic time series, including the unemployment rate. However, recent papers find modest forecast improvements using nonlinear methods, particularly around recessions (for example, Milas and Rothman, 2008). Note that the staff’s flows model (discussed below) inherently captures the asymmetric movements in unemployment because this model incorporates the underlying sources of the asymmetry: the flows themselves.

<sup>18</sup> The improvement is similar when UI claims were included. A simple VAR that includes the unemployment rate, payroll employment growth, and GDP growth performed about the same as the VAR that did not include GDP. This lack of improvement from including GDP may be partly due to mismatch in timing between the monthly unemployment rate data and the GDP data, which are available only at a quarterly frequency. This mismatch results in the first reading of the GDP data for a given quarter being available only one month after the quarter ends, which is about three months after the first reading on labor market data for that quarter. Within the Federal Reserve System there are more elaborate models that may have better forecasting properties.



models does not significantly reduce the confidence interval around a two-quarter-ahead forecast (Milas and Rothman, 2008).

Another strategy is to exploit the leading relationship between the “steady-state” rate of unemployment implied by the underlying labor force flows from the CPS and the actual unemployment rate (Barnichon and Nekarda, 2012). Because the unemployment rate typically converges toward this steady-state rate with a short lag, the current flows provide information about the future unemployment rate. The forecast horizon of such a model can be extended by projecting the labor force flows themselves using additional labor market variables. In out-of-sample exercises using real-time data, this model outperformed the simple univariate model just described, with RMSEs of 0.3 and 0.6 percentage point at horizons of one quarter and two quarters, respectively (Table A, line 3). Currently, the staff’s flows model projects that the unemployment rate will remain at 7.9 percent for the next six months.

The approaches to forecasting the outlook for the labor market discussed thus far are largely atheoretical and make use of a fairly narrow set of data. Another possibility is to examine the outlook for the labor market within the context of a forecast of the economy as whole, such as the staff projection that is published in the Tealbook. Advantages of the Tealbook approach include its use of a range of information and statistical models, the ability of the staff to estimate current-quarter economic activity from incoming data, and the fact that the staff projection explicitly links forecasts of the unemployment rate to projections of economic activity. A potential disadvantage is that the Tealbook is based on one particular view (in this case the staff’s) of the structure of the economy, while the projection of the unemployment rate itself will depend on the accuracy of the Tealbook’s projections for other variables (such as GDP, potential output, and the natural rate).<sup>19</sup>

Over the one- to two-quarter-ahead horizon under consideration here, the Tealbook forecast for the unemployment rate does better than the other models discussed previously, although it is still surrounded by sizable confidence intervals.<sup>20</sup> Since 1997, the RMSE for forecasts one quarter ahead was about 0.3 percentage point and for two quarters ahead was roughly 0.5 percentage point (Table A, line 4).<sup>21</sup>

To better illustrate the uncertainty in the staff projection, we performed an exercise, similar to the one presented above, in which we assessed the accuracy of Tealbook unemployment rate

---

<sup>19</sup> Of course, many of the pros and cons of the Tealbook projection of the unemployment rate are the same as for other projections that embed the unemployment rate forecast in a forecast of the economy as a whole. Such forecasts—including those from Board models such as FRB/US, EDO, and SIGMA; DSGE models at the Reserve Banks; and those of outside forecasters—also incorporate a broad array of data on the economy and depend on a particular view of the structure of the economy.

<sup>20</sup> At horizons beyond two quarters, the Tealbook forecast of the unemployment rate is substantially better than that of the other models considered here.

<sup>21</sup> Note that while the other models made use of the monthly unemployment rate data, the errors for the Tealbook forecast are based on quarterly data. That said, the evaluation of the models is consistent because the comparison was designed so that the models and the Tealbook forecasts all had the same information on the unemployment rate.

forecasts conditional on observed movements in the unemployment rate. As before, we consider increases or decreases of 0.15 percentage point or less to constitute little change. In cases when the unemployment rate was little changed over the preceding two quarters (44 cases since 1997), the staff was able to correctly forecast the sign of the movement in the unemployment rate over the subsequent 2 quarters about 55 percent of the time; the success rate was the same in the 38 cases in which the unemployment rate had recently declined. However, of the 23 times the staff predicted a decline in the unemployment rate (whether preceded by a flat unemployment rate or a decline), they were right 18 times (78 percent), and the unemployment rate rose only twice.<sup>22</sup>

Overall, these findings suggest that it is possible to improve upon the performance of a univariate model. All of the multivariate approaches improved upon the univariate forecasts, and the difference in the RMSE between the univariate model forecast and the Tealbook forecast is approximately ½ percentage point. Nonetheless, despite this improvement, there is still considerable uncertainty associated with even the best of these forecasts.

## Conclusions

Our main conclusions regarding the most appropriate indicators to consider when judging current conditions in—or the outlook for—the labor market are summarized below:

- In addition to the unemployment rate's advantages in terms of communicating with the public, it is also a reasonably reliable summary measure of contemporaneous labor market conditions.
- However, changes in the unemployment rate—even those sustained over the course of a few months—can, at times, be misleading indicators of underlying labor market conditions. As a result, other indicators of the labor market or the state of the economy can be helpful in judging the extent to which an observed decline in the unemployment rate represents an actual improvement in labor market conditions.
- Above-trend growth in payroll employment appears to be the most relevant supplemental piece of information in this regard, but a variety of other measures covering particular aspects of the labor market (for example, measures of layoffs, hiring, job vacancies, and labor supply changes) can also be useful.
- Movements in the unemployment rate have historically been persistent, implying that recent changes have been useful predictors of future changes in unemployment. However, the accuracy of these predictions is relatively modest.
- Other information from the labor market or about the economy more generally can help to improve forecasts of the unemployment rate as compared with a univariate approach.
- That said, historically even the most accurate of real-time forecasts considered here has an RMSE for two-quarter-ahead forecasts on the order of ½ percentage point.


---

<sup>22</sup> That is, when the staff forecasts a decline in the unemployment rate, they are often correct; however, the staff does not predict a decline often enough, conditional on having observed little change or a decline over the previous two quarters.

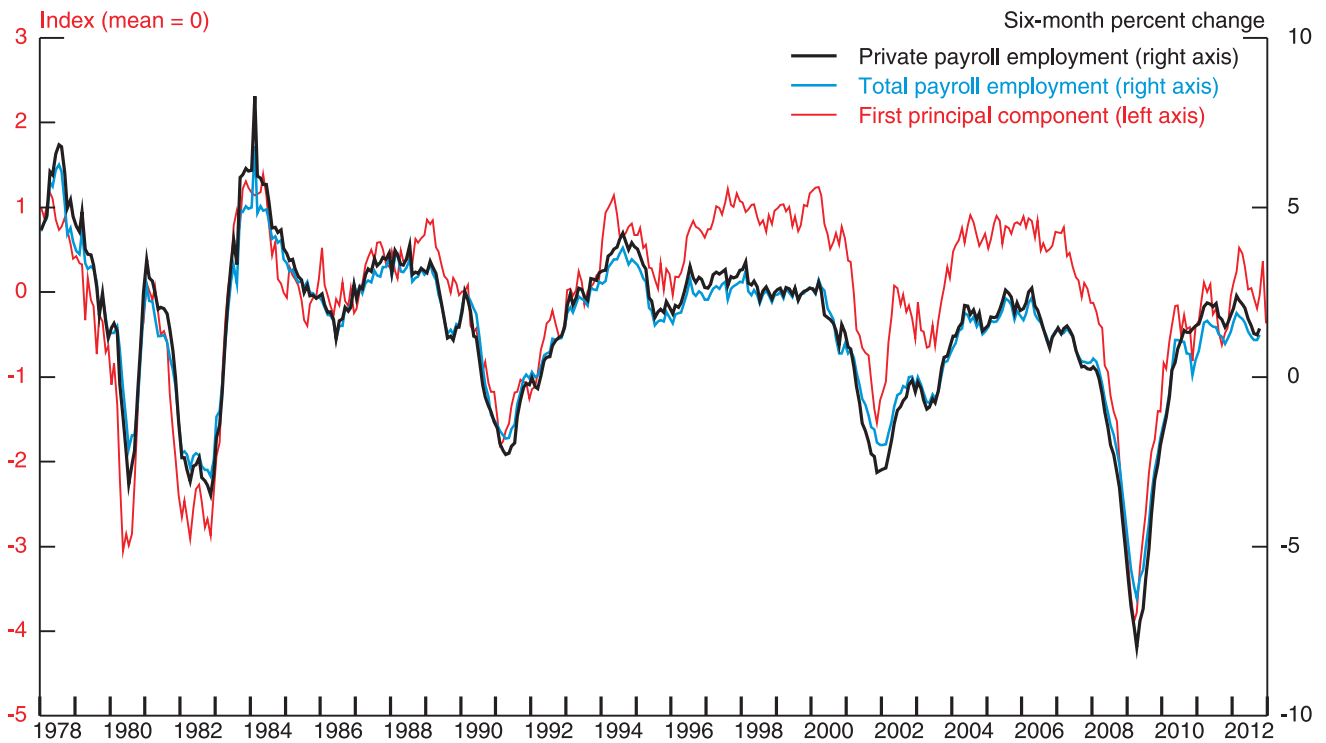
## References

- Barnichon, Regis and Christopher J. Nekarda (2012). “The Ins and Outs of Forecasting Unemployment: Using Labor Force Flows to Forecast the Labor Market,” *Brookings Papers on Economic Activity*, Fall.
- Fleischman, Charles A., and John M. Roberts (2011). “From Many Series, One Cycle: Improved Estimates of the Business Cycle from a Multivariate Unobserved Components Model.” Federal Reserve Board Finance and Economics Discussion Series Paper no. 2011-46, October.
- Milas, Costas, and Philip Rothman (2008). “Out-of-sample Forecasting of Unemployment Rates with Pooled STVECM Forecasts,” *International Journal of Forecasting*, 24: 101–21.
- Montgomery, Alan L., Victor Zarnowitz, Ruey S. Tsay, and George C. Tiao (1998). “Forecasting the U.S. Unemployment Rate,” *Journal of the American Statistical Association*, 93(442): 478–93.
- Nalewaik, Jeremy J. (2011). “Forecasting Recessions Using Stall Speeds.” Federal Reserve Board Finance and Economics Discussion Series Paper no. 2011-24, April.
- Staiger, Douglas, James H. Stock, and Mark W. Watson (1997). “How Precise Are Estimates of the Natural Rate of Unemployment?” In Christina D. Romer and David H. Romer (eds.), *Reducing Inflation: Motivation and Strategy*, 195–242. Chicago, IL: University of Chicago Press.
- Stock, James H., and Mark W. Watson (1999). “A Comparison of Linear and Nonlinear Univariate Models for Forecasting Macroeconomic Time Series.” In Robert Engle and Halbert White (eds.), *Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive W.J. Granger*, 1–44. Oxford: Oxford University Press.

### Exhibit 1: Unemployment Rate and First Principal Component



### Exhibit 2: Measures of Employment and First Principal Component



Note: First principal component of 37 labor market indicators.