Model-based seasonally adjusted estimates and sampling error

Estimating certain CPS series with a model that filters out sampling error may reduce volatility in the time series, facilitating more meaningful trend analysis

Richard Tiller
and
Marisa Di Natale

The Current Population Survey (CPS) is the source of the Nation's official estimates of total employment and unemployment. The CPS is a nationally representative, scientifically selected monthly sample survey of approximately 60,000 households. The survey yields data that are rich in demographic detail, including such characteristics as age, sex, race, and Hispanic or Latino ethnicity. Estimates from the survey are published monthly in the BLS news release *Employment Situation* and in the BLS publication *Employment and Earnings*.

In order to make the time-series data collected from the CPS more useful to analysts and policymakers, the monthly data from the survey are adjusted for seasonal fluctuations. As is well known, the purpose of seasonally adjusting a series is to remove seasonal fluctuations in the data so that users can more easily observe fundamental changes in the level or trend of the series that are associated with business cycle contractions and expansions. Approximately 116 time series from the CPS are directly seasonally adjusted, and many more are indirectly seasonally adjusted, as sums or ratios of the original 116.

There is, however, a source of spurious random fluctuations in the CPS data that arises because the CPS samples only a fraction—1 in 2,200, on average, of the working-age population each month: *sampling error*—the difference between the survey estimates and the values that would be produced by a complete census of the population.

Simultaneously removing both seasonality in the data and noise due to sampling error can prove quite challenging. The monthly estimates produced for the national aggregated series, such as total employment and total unemployment, are highly reliable relative to smaller, disaggregated series. Many of the more detailed demographic series, such as employment and unemployment for blacks, are based on relatively small sample sizes, so that survey error dominates movements in the underlying level of the series. The standard error for a (not seasonally adjusted) month-to-month change based on the CPS can be quite high for some of these series. For example, the standard error for a change in the unemployment rate of adult black males can be as large as 0.8 to 0.9 percentage point, compared with 0.2 percentage point for the unemployment rate for all persons aged 16 years and older. As a result, drawing meaningful conclusions about trends or month-to-month changes is difficult, even after the data have been adjusted for seasonal movements.

As an alternative to conventional seasonal adjustment, the study reported in this article applies an experimental model-based method to selected CPS demographic series. The method is designed to remove the effects of sampling error, as well as those of seasonality, from the series, thereby making it easier to discern underlying trends in the data.

Richard Tiller is a mathematical statistician on the Statistical Methods Staff, and Marisa Di Natale is an economist formerly in the Division of Labor Force Statistics, Office of Current Employment Analysis, Bureau of Labor Statistics.

Approaches to seasonal adjustment

The presence of large survey errors in the detailed CPS series represents a major challenge to conventional methods of seasonal adjustment. Currently, the Bureau of Labor Statistics uses a seasonal adjustment program called X-12-ARIMA to seasonally adjust its CPS series. This program is based on the empirical moving-average approach to seasonal adjustment.

An alternative that is gaining increasing attention is the *model-based* approach to seasonal adjustment. A comparison of the two approaches suggests that the model-based approach provides much-needed flexibility in controlling for the effects of sampling error. Such flexibility is not possible with the conventional moving-average approach.

Conventional approaches to seasonal adjustment are based on the classical decomposition of a time series, which assumes that the series is composed of trend (or trend-cycle), seasonal, and irregular components, in either an additive or a multiplicative relationship. The first two components respectively account for the long- and short-run systematic variation in the series. The irregular component is a residual, usually assumed to be purely random variation with a fixed variance.

Because the three components of the classical decomposition are not directly observable, they must be estimated in order to perform seasonal adjustment. The moving-average method uses weighted moving averages of the original data over a period of many years to produce a smooth trend and a seasonal pattern. The estimated trend and seasonal components are removed from the series, and the residual is the irregular component. This approach makes no attempt to define, in any formal statistical way, what is being estimated, but rather applies a series of moving averages directly to a series. While some of the moving averages are chosen to satisfy a mathematical smoothness criterion, the method was derived largely from empirical work with a wide range of series.

By far the most successful application based on the moving-average approach is the X-11 program,[1] which has gone through several major revisions. The latest, enhanced version is X-12-ARIMA.[2] The original X-11 program, however, remains at the core of X-12.

As an alternative to the moving-average approach, model-based seasonal adjustment has been gaining increased attention. The model-based approach specifies explicit statistical models of the trend, seasonal, and irregular components of the classical decomposition.[3] To seasonally adjust the data, weighted moving averages of the observed data actually are used in the model-based approach, but with the important difference that they are derived directly from the model. An essential characteristic of the approach is its use of standard statistical procedures to estimate the unobserved components of the time series and to provide associated statistical measures such as confidence intervals and significance tests.

There is a large body of literature on the comparative properties of the two approaches.[4] Each has its supporters and critics. The major concerns with the model-based approach are that it may be difficult to develop good models for some series and that the model may fail occasionally when new data become available. These concerns raise issues about the robustness of the adjustment and the associated statistical measures.

One major criticism of the moving-average approach is that it lacks standard statistical measures. The absence of standard errors for published seasonally adjusted data tends to promote the mistaken impression that the final seasonally adjusted values are exact rather than estimates. Moreover, the lack of confidence intervals makes analysis of change in the estimates and the location of turning points more difficult. Still, supporters of the method argue that it is robust and nonparametric; thus, its lack of an explicit statistical model is viewed as an advantage. Even so, the absence of statistical measures of reliability remains a major shortcoming.[5]

The model-based approach makes (testable) assumptions about the underlying probability distribution generating the data. Along with estimates of the model parameters, these assumptions provide the means for constructing confidence intervals and other statistical measures to quantify the uncertainty in the estimates. Non-model-based estimates do not, in general, afford a basis for producing measures of uncertainty in the estimates.

Another criticism of the moving-average approach is that it is not tailored to the specific properties of the series being adjusted.[6] In contrast, the model-based approach develops a model on the basis of goodness-of-fit diagnostics. The resulting seasonal adjustment is based on the properties of the series as represented by the model. In theory, under the assumptions of the model, the seasonal adjustment is "optimal" for the specific series. While the moving-average method can make no such theoretical claim, its moving averages were originally selected because they work well for a very large number of series. This more generic approach continues to work well in practice and may have an advantage over the model-based approach when good models for seasonal adjustment cannot be developed.

Clearly, both approaches to seasonal adjustment have their merits and limitations. Indeed, there have been a number of studies of the relative performance of the two approaches, but no general agreement as to how to interpret the results. Perhaps a more balanced approach is to treat them as complementary tools for performing seasonal adjustment.[7]

Dealing with "noisy" CPS data series

The types of data series that are the focus of the study discussed in this article—survey series with large sampling errors—represent a class of series that presents special problems for the moving-average approach to seasonal adjustment. For these series, that approach (specifically, the X-12 program) performs poorly, not because it has trouble removing seasonality from the series, but because it cannot adequately remove the effects of sampling error. The result is a seasonally adjusted series that often is dominated by sampling error, masking the underlying trend in the series. This occurs because the moving averages are not tailored to the specific properties of survey series with sampling errors. What is implicitly assumed in applying these moving averages to survey data is that sampling error can be adequately treated as part of the conventional irregular component, which, conceptually, is a purely random series with a fixed variance. The design of the survey, however, determines the properties of the sampling errors, which may deviate in important ways from typical irregular behavior. Moreover, for well-designed probability samples, the characteristics of the sampling errors are known, or at least, good estimates of them can be obtained.

For the CPS, the standard errors provided along with the point estimates from the survey routinely yield information on the magnitude of the sampling errors. These standard errors, however, are not constant and vary substantially over time for some series, due to fluctuations in labor force levels, redesigns of the survey, and changes in the sample size. Another important characteristic of the CPS that is relevant to seasonal adjustment is its panel structure, which generates strong correlations of the sampling errors with their past values. The CPS has a rotating panel whereby three-fourths of the households are carried over from the previous sample each month and one-half are carried over from the previous year.[8]

These characteristics represent a serious challenge to conventional approaches to seasonal adjustment. First, the relatively large and changing magnitude of the survey error will directly distort estimates of the trend and seasonal components. In addition, because the survey error is correlated, the moving-average approach treats the induced correlations as if they were related to the trend. Consequently, too much of the sampling error is absorbed into the estimated trend and, to some extent, into the estimated seasonal component, and not enough goes to the irregular component.

The standard model-based approach will do no better than the moving-average approach, because it also is based on the classical time-series decomposition, which ignores survey error as an important source of variation. The model-based approach, however, has flexibility, which is not available with the moving-average approach. As mentioned earlier, the latter, by virtue of its nonparametric structure, is a more generic approach to seasonal adjustment, whereas the model-based approach can be tailored to the idiosyncratic properties of a particular series. From the model-based perspective, survey error is just another unobserved component of the time series for which a separate model can be specified, as is done with the trend, seasonal, and irregular components of the series. Then there is the further advantage, in modeling survey error, of having external information from the survey on the standard errors and correlations. This information can be directly used to specify the parameters of the model. Of course, measures of the survey error characteristics, such as standard errors and autocorrelations, are estimates which are themselves subject to errors that can adversely affect the decomposition. For example, overestimates of the variances could lead to "oversmoothing" the series. Still, given good estimates of the survey error properties, a rather precise and objective identification of the effects of survey error is possible, resulting in a much cleaner separation of the estimated trend and seasonal components from the survey series than can be achieved by the moving-average approach.

As described in the appendix to this article, research has been conducted that uses a time-series model to seasonally adjust some of the more volatile CPS series. For example, model estimates have been produced for blacks 16 years and older, for adult black men and women, and for black male and female teens. The model producing these estimates has two submodels as components: a "signal" model, which models the true value of the specific demographic group; and a "noise" model, which models the survey error associated with that series. This model-based approach filters out the survey error and the seasonal component of the series, allowing the trend to be more cleanly separated from the error. Using estimates from the model instead of from X-12 results in smaller standard errors of the estimates and smoother seasonally adjusted series. This type of model estimation has been widely investigated in the context of "small-area" estimation and has been used in the BLS Local Area Unemployment Statistics program for more than 15 years to estimate employment and unemployment for States and selected metropolitan areas.[9]

Historical model-based estimates

There are two ways of processing time-series data to produce model-based estimates. One, called the "forward filter," processes each observation as it first becomes available in real time. Filtering is similar to the "concurrent" seasonal adjustment method used to seasonally adjust new CPS data series as they become available each month.[10] The most recent value of the time series is "filtered" by using all the available data up to and including the latest month. Prior months' estimates are not revised as new data become available.

The second approach to processing the data, called "smoothing," incorporates data before and after a particular observation. All the data are processed in batch mode, rather than one observation at a time. During the smoothing process, estimates derived from filtering at each point in time are revised to reflect information in all of the data. This approach creates a much smoother seasonally adjusted series than does filtering, although it does require revision of the entire time series. (See the appendix for further discussion of smoothing and filtering.)

Counterparts of smoothing and filtering also exist within the conventional moving-average approach. For X-12, the smoothest estimates are produced by symmetric moving averages, which use an equal number of observations before and after the time point being adjusted. Toward the end of the series, where there are fewer observations after that point, less desirable asymmetric moving averages must be used. When the X-12 program is executed to seasonally adjust the most recent observation, the result is referred to as the *concurrent estimate*, which is the counterpart of the model's filtered estimate.

Charts 1 and 2 compare smoothed model-based and X-12 seasonally adjusted unemployment rates and employment levels from January 1999 through December 2003 for blacks aged 16 and older.[11] Chart 2 also includes the X-12 trend series, discussed later. The X-12 seasonally adjusted series are published each month in the BLS *Employment Situation* news release. The graphs show that the model estimates create smoother seasonally adjusted lines or trend lines for both series than do the corresponding X-12 seasonally adjusted series.

The conventional time-series decomposition represents the observed series as being composed of trend, seasonal, and irregular components. The seasonally adjusted series consists of the trend and irregular components and thus is generally less smooth than the trend. For the series modeled, the distinction between the seasonally adjusted series and the trend series is not important: once the estimated survey error is removed from the series, very little residual irregular component is left. The model-based seasonally adjusted series, therefore, is virtually identical to the trend component. For X-12, the irregular component is relatively large; thus, its seasonally adjusted series is noticeably less smooth than its trend.

One advantage of the modeled estimates is evident from the graphs: peaks and troughs in these series are more easily identifiable, which is not the case with the X-12 seasonally adjusted series. This makes identifying labor market turning points for specific demographic groups easier. For instance, in the black employment- and unemployment-rate series in charts 1 and 2, the turning point clearly occurred in mid-2000.

X-12's trend estimates, shown in chart 2, clearly involve

Chart 1. Unemployment rates, blacks aged 16 years and older, seasonally adjusted, 1999–2003

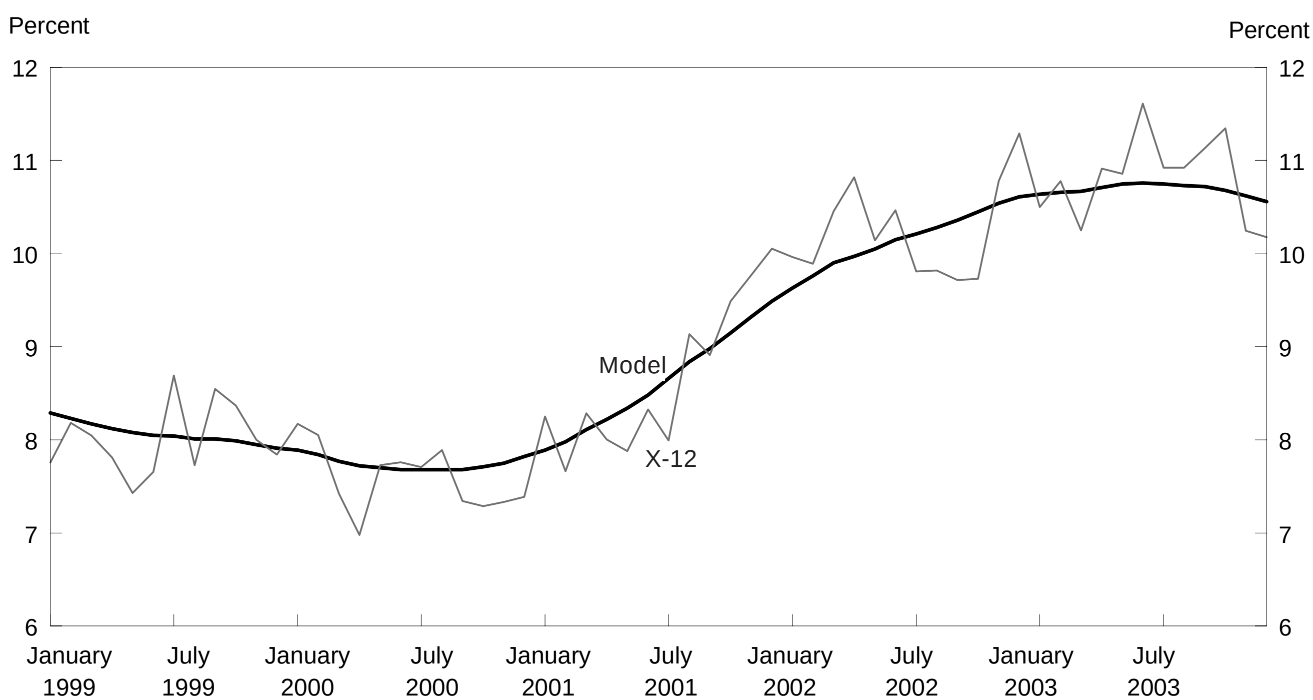

Chart 2. **Employment levels, blacks aged 16 years and older, seasonally adjusted, 1999–2003**

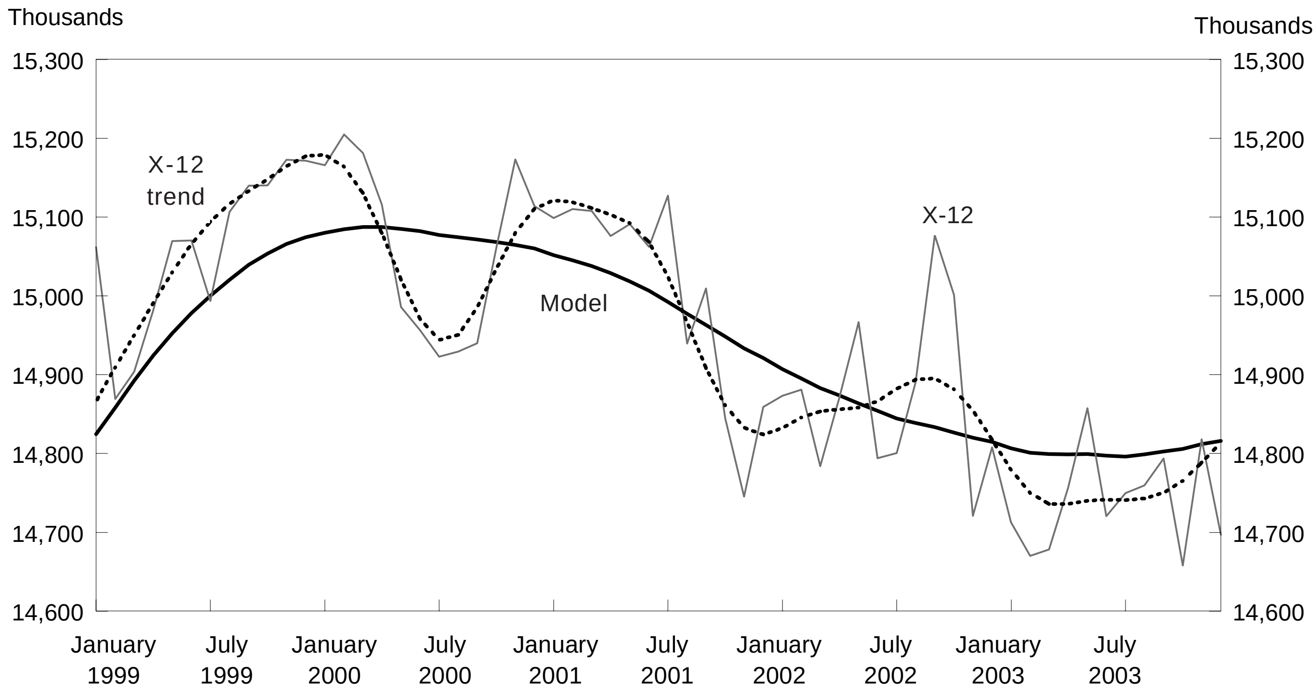

Chart 3. **Unemployment rates, blacks aged 16 years and older, concurrent, seasonally adjusted, 2003–04**

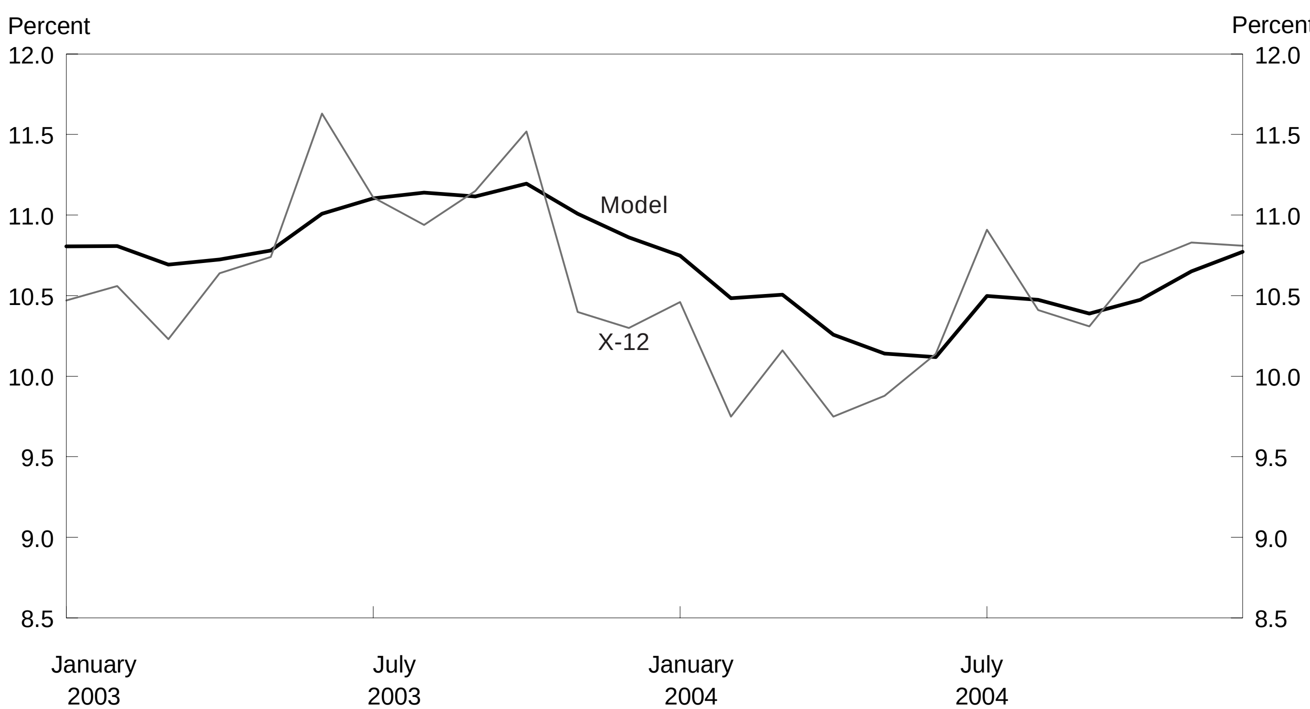

further smoothing of the seasonally adjusted series, but this does not help identify real turning points. Due to the effects of the survey, the X-12 trend series displays spurious oscillations around the model-based trend series. These oscillations could be misinterpreted as real turning points.

Real-time current-year estimates

Because the smoothed estimates use data from the entire sample (January 1976–December 2004), they do not correspond to the estimates that would have been available in 2004 to data users in real time. As mentioned in the previous section, the latest monthly estimate in a modeled series is a filtered estimate that incorporates prior months' data, as well as data for the current month. Filtered estimates are not as smooth as the "smoothed" series, because filtering cannot incorporate data from the future, nor does it incorporate the latest information into past months' estimates. As a result, the month-to-month changes in the filtered estimates are much more volatile than the smoothed estimates, but still are smoother than the X-12 estimates.

Charts 3 through 5 show forward-filter estimates of the black unemployment rate and employment level, and of the unemployment level for adult black women, for January 2003 through December 2004. In chart 3, it is clear why the modeled data may be preferable to the X-12 estimates. The X-12 series shows two rather large spikes in black unemployment in June and October 2003. The June spike appears to be the larger of the two. In contrast, the model series shows, not the June spike, but rather a small, statistically insignificant rise, followed by a flattening out of the rate and a slower decline from October 2003. The model suggests that the rate peaked at around 11.0 percentage points during the summer and early fall of 2003, whereas X-12 gives a more confusing picture. The employment level for blacks is shown in chart 4. The modeled data indicate that the employment level bottomed out around August 2003, whereas the X-12 series depicts the lowest level of employment for this group during the first quarter of the year. In chart 5, the unemployment level of adult black women is shown. Although in both the modeled series and the X-12 series, the unemployment levels wind up being roughly the same at the beginning and end of the period shown, the modeled series is clearly less volatile over the year.

An important feature of the model-based approach, as emphasized in the previous section, is that it provides estimates of the standard errors for the point estimates,

Chart 4. Employment levels, blacks aged 16 years and older, concurrent, seasonally adjusted, 2003–04

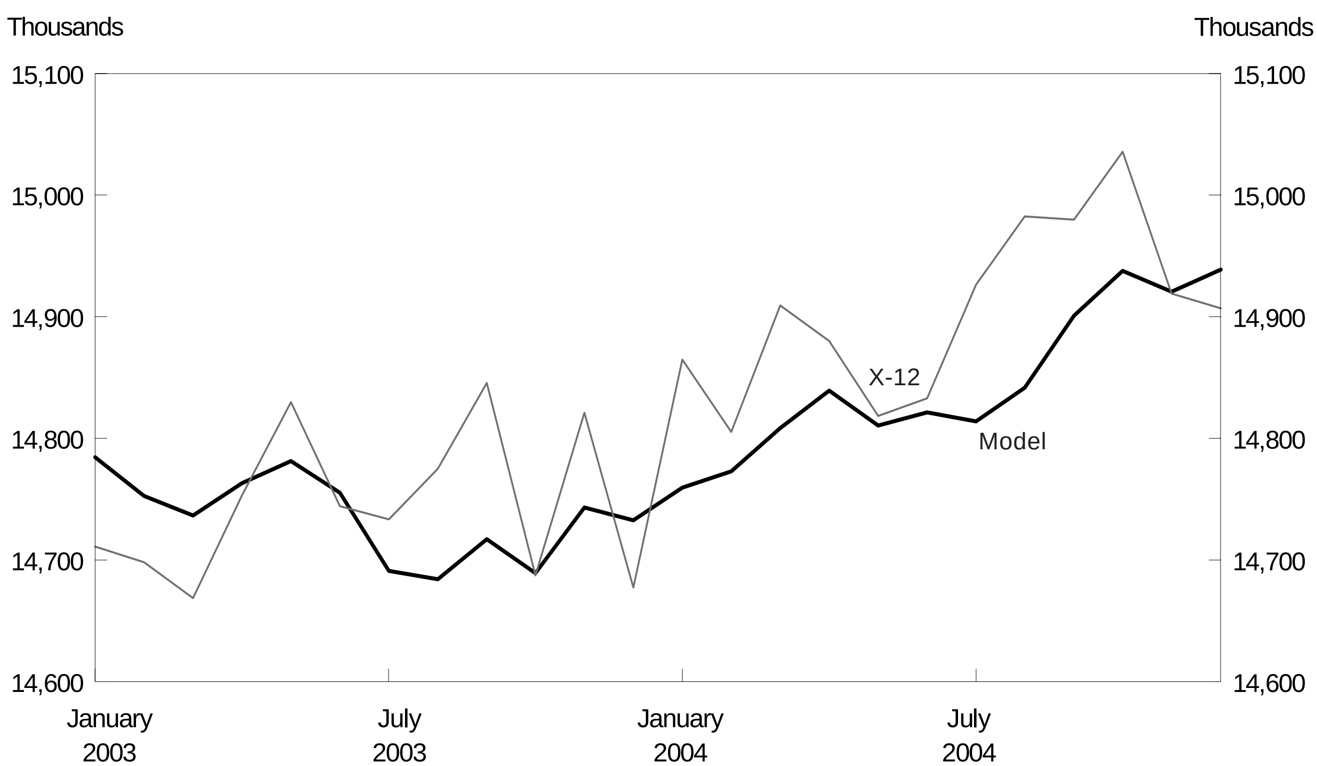

Chart 5. Unemployment levels, adult black women, concurrent, seasonally adjusted, 2003–04

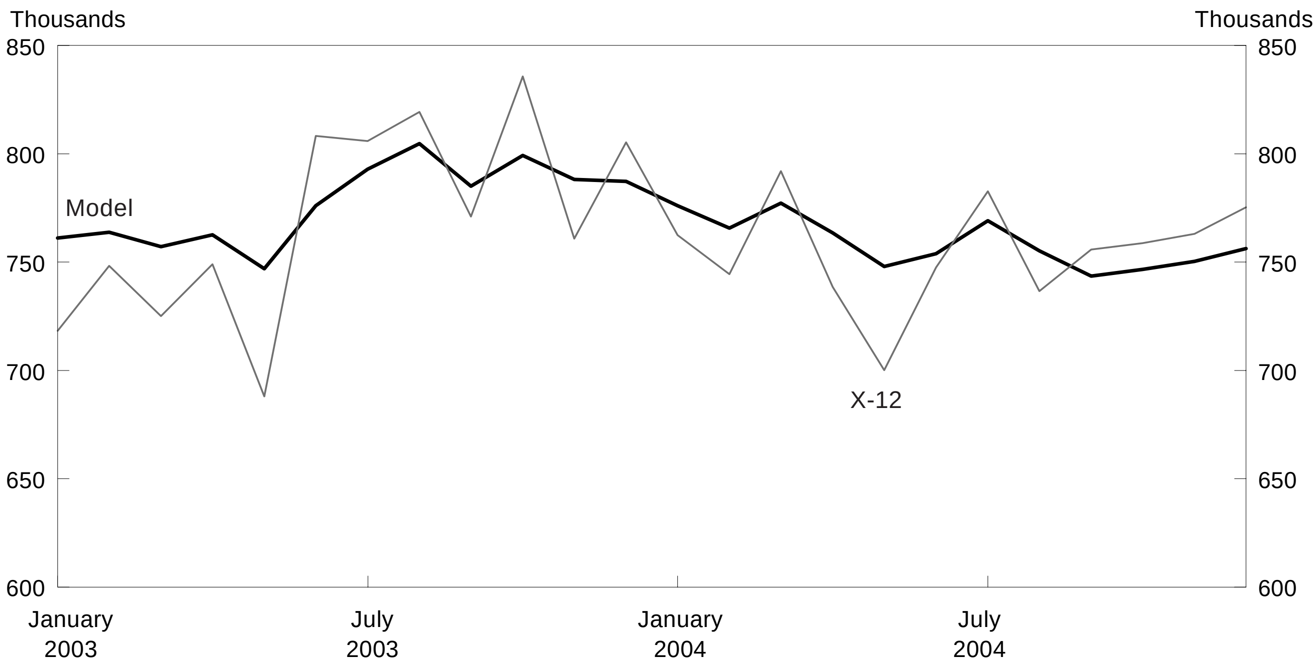

whereas X-12 does not. It is common practice for analysts to use the standard errors for the not-seasonally-adjusted CPS estimates as a proxy measure of error for the X-12 estimates. While this practice is not strictly correct, recent research suggests that these standard errors may be a reasonable approximation (an overestimate of around 10 percent to 15 percent, sometimes less) for detailed series in which sampling error is an important source of variation.[12] The following tabulation compares the model-based standard errors of monthly change from the forward filter with the CPS proxy measures and computes the minimum magnitude of monthly change in the respective estimates required to achieve significance at the 90-percent level for the period from January 2000 to July 2004 (the figures shown are all mean values):

Source and measure	*Unemployment rate, blacks*	*Employment level, blacks*	*Unemployment, black adult women*
Standard error:			
CPS	0.51	101,271	49,261
Model	.21	41,580	14,047
Minimum change required for significance:			
CPS	±.84	±166,591	±81,034
Model	±.35	±68,399	±23,107

For the total black employment level, the minimum magnitude of monthly change required for significance has to be at least 167,000 with the CPS proxy measure—almost two-and-a-half times the magnitude required for the model estimates. For unemployed black adult women, the CPS proxy measure requires a change of at least 81,000 in the X-12 seasonally adjusted estimates, or 3.5 times the change required in the model estimates. For the total black unemployment rate, the CPS proxy requires a change of more than 0.8 percentage point, compared with less than 0.4 percentage point for the model. Thus, even though the CPS proxy may overestimate the required monthly change in X-12 by up to 15 percent in relative terms, that still leaves major gains from using the model estimates.

Research to this point suggests that the model-based approach to seasonal adjustment is a useful tool for helping to discern trends in data series with large variances due to relatively small sample sizes. The approach, however, is not necessarily superior to the moving-average approach currently used by the Bureau of Labor Statistics for adjusting aggregate series such as total employment and unemployment when sampling error is a much less important source of variation than it is for the detailed demographic series. Research into the merits and limitations of the model-based approach to smoothing time series and creating estimates of variance is ongoing in order to evaluate possible extensions of the method to other series. □

Notes

[1] Julius Shiskin, Alan H. Young, and John C. Musgrave, *The X-11 variant of the census method II seasonal adjustment program*, technical paper 15 (Bureau of Economic Analysis, 1967).

[2] David B. Findley, Brian C. Monsell, William R. Bell, Mark C. Otto, and Bor-Chung Chen, "New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program," *Journal of Business and Economic Statistics*, April 1998, pp. 127–77 (with discussion).

[3] Examples are TRAMO-SEATS (see Victor Gomez and Agustin Maravall, "Programs Seats and Tramo: Instructions for the User," Working Paper No. 9628 (Madrid, Bank of Spain, 1996)); and STAMP (see Siem J. Koopman, Andrew C. Harvey, Jurgen A. Doornik, and Neil Shephard, *STAMP : Structural Time Series Analyser Modeller and Predictor* (London, Timberlake Consultants, 2000)).

[4] See, for example, William R. Bell and Steve C. Hillmer, "Issues Involved with the Seasonal Adjustment of Economic Time Series," *Journal of Business and Economic Statistics,* October 1984, pp. 291–320; Agustin Maravall, "Unobserved Components in Economic Time Series," *Handbook of Applied Econometrics*, ed. M. Hashem Pesaran, and Mike Wickens (Oxford, U.K., Basil Blackwell, 1993), pp. 12–72; and Findley, Monsell, Bell, Otto, and Chen, "New Capabilities and Methods."

[5] Recent research suggests that the absence of explicit statistical models in X-11 may not prevent statistical inference from being carried out after all. Stuart Scott and Danny Pfeffermann have experimented with a promising approach to the development of standard error measures for X-11 estimators. (See Stuart Scott and Danny Pfeffermann, "Evaluation of Two Variance Methods for X-11 Seasonally Adjusted Series," ASA *Proceedings of the Joint Statistical Meetings* (Alexandria, VA, American Statistical Association, 2003), pp. 3760–67.)

[6] X-12 allows for some flexibility by automatically varying the length of its filters on the basis of series properties.

[7] This view appears to be gaining momentum with the effort to have X-12 and model-based seasonal adjustment as alternative options in a single software program called X-12-ARIMA/SEATS. (See Brian C. Monsell, John A. D. Aston, and Siem J. Koopman, "Toward X-13?" *ASA Proceedings of the Joint Statistical Meetings* (Alexandria, VA, American Statistical Association, 2003), pp. 1459–66.)

[8] As a result, if the sample, say, overestimated the true value in the previous month, it is likely to continue overestimating in the current and successive months. In other words, the behavior of the survey errors is more similar to a trend with cumulative up-and-down movements than to one with purely random irregular variation.

[9] For a discussion of these models, see R. Tiller, "Time series modeling of sample survey data from the U.S. Current Population Survey," *Journal of Official Statistics*, June 1992, pp. 149–66.

[10] For more on the switch to concurrent seasonal adjustment in the CPS, see Richard B. Tiller and Thomas Evans, "Revision of Seasonally Adjusted Labor Force Series in 2004," *Employment and Earnings*, January 2004, pp. 3–9; on the Internet at **http://www.bls.gov/cps/cpsrs2004.pdf**.

[11] Both the model-based and the X-12 series are based on data extending back to January 1976. In chart 2, the estimates have been adjusted to smooth out the effect of the introduction of new population controls from Census 2000 in January of that year.

[12] Stuart Scott, Michail Sverchkov, and Danny Pfeffermann, "Variance Measures for Seasonally Adjusted Employment and Employment Change," ASA *Proceedings of the Joint Statistical Meetings* (Alexandria, VA, American Statistical Association, 2004), pp. 1328–35.

APPENDIX: Methodology

To account for the special properties of the CPS series, a signal-plus-noise model is developed that combines a time-series model of the "true" values and its unobserved components (trend, seasonal, and irregular) with a model of the sampling errors. The latter is treated as an additional unobserved component of the time series, with the special advantage that its variance-covariance structure is objectively identified by design information.

This appendix discusses, in general terms, how the components of the CPS are modeled, briefly considers the series that is directly modeled in the text of this article, and examines the use of the Kalman filter and smoother algorithms in the text to produce seasonally adjusted estimates in both real and historical time.[1]

Component models of the CPS

For labor force series, seasonality is an important source of variation. The nonseasonal, or trend, part of the series cannot be observed directly; instead, trend values must be inferred solely from observations of the aggregate series. Seasonal adjustment is a special case of the more general "signal-plus-noise" formulation of an observed series as being composed of a number of unobserved components, where the components of interest are referred to as the signal and all other components are considered to be noise.

The goal is to "filter out" the signal from the noise when all that can directly be observed is the data corrupted by noise. The solution to this problem requires a model that specifies, either implicitly or explicitly, the properties of each of the underlying components. In the conventional seasonal adjustment problem, the observed series, Y_t, is assumed to be measured without error and to be composed of trend, T_t, seasonal, S_t, and irregular, I_t, components in additive form:

$$Y_t = T_t + S_t + I_t.$$

Multiplicative seasonality is handled by taking logarithms of the data, fitting the model to the data, and then taking the antilogarithm of the component estimates back to the original scale.

The signal is the nonseasonal component (trend plus irregular component), and the noise includes the seasonal component (when irregular fluctuations are relatively large, it may be more appropriate to define the signal as just the trend and treat the irregular component as part of the noise):

$$\text{Signal: } \mathrm{Sig}_t = T_t + I_t;$$

$$\text{Noise: } N_t = S_t.$$

Following Andrew C. Harvey, the trend, seasonal, and irregular components are specified as structural time-series models.[2] Each component has an associated variance that determines its properties. The trend takes a local linear form, with its level and slope varying randomly. A wide variety of patterns is a possibility, depending on the relative magnitudes of the variances of the level and slope. When the two variances are both zero, we have the smoothest possible trend, which assumes a fixed linear form over the entire observation period. A rapidly fluctuating trend results when the variance of the level shift is much larger than the variance of the slope.

The seasonal component also evolves over time according to its variance. A seasonal pattern that is fixed from year to year implies a zero variance, while patterns that change gradually imply a small positive variance. Finally, the magnitude of the irregular component depends directly on the size of its variance. A zero variance implies the absence of irregular variation in the series. The four variances of the structural model presented in the text of this article are unknown parameters that will be estimated from the data once a complete model is set up as described herein.

In the case of CPS data, there is a fourth component: the sampling error e_t, which arises because only a fraction of the total population is sampled each month. This means that the true values Y_t are no longer directly observed, but instead, the survey estimates,

$$y_t = Y_t + e_t,$$

are what is observed. The definition of the signal is unchanged, but now the noise includes survey error as well as seasonality; that is,

$$N_t = S_t + e_t.$$

Since the presence of sampling error in the data can radically alter the results of seasonal adjustment, a model is required that takes the sampling error into account when it is an important source of variation in the data. Unlike the situation with the other components, the properties of this component do not have to be estimated from just the time series itself. Rather, direct information from the survey microdata may be used to estimate the sampling error model independently of the structural time-series model.

The Census Bureau routinely produces data to assess the reliability of national CPS statistics on an ongoing basis. The process involves drawing a set of random subsamples, or replicates, from the full sample surveyed each month, using the same principles of selection as are used for the full sample, and applying the regular CPS estimation procedures to the replicates that are drawn. Each month, 160 replicates are produced for a large number of characteristics. The variability in the replicates provides the basis for computing empirical variances and autocorrelations for the sampling error. Variances are computed by fitting generalized variance functions to the monthly replicate variances to smooth out volatility in the latter. Lag covariances are computed by averaging each replicate lag covariance over time and normalizing for changes in variance.

The empirical sampling-error variances and autocorrelation estimates are then used to directly derive the parameters of the survey error model, which is specified as an autoregression that relates the current value of the survey error to its past values. Note that the coefficients are determined from the sampling-error autocorrelations, and the variances are adjusted to conform to the survey variances. In this way, both the effect of the rotating panel design and the magnitude of the sampling error are taken into account.

The sampling-error model is combined with the time-series model of the trend, seasonal, and irregular components, with the latter variances estimated directly from the historical CPS series. Thus, the basic properties of the overall model are tailored to the empirical behavior of the CPS series.

Model fitting to CPS series

Black employment- and unemployment-level series for the following four age-gender groups are directly seasonally adjusted, and variance estimates are produced, with the models presented in the text of the article:

- black male youths aged 16–19 years
- black female youths aged 16–19 years
- black men aged 20 years and older
- black women aged 20 years and older.

Seasonally adjusted total black employment and unemployment levels are obtained indirectly by summing the appropriate model estimates, and the unemployment rate is derived from the estimates of these levels. Although seasonal adjustment of the totals is done indirectly, the variances are produced by directly fitting the models to the aggregate data.

The observation period of the study begins in January 1976 and ends in December 2004. For each series directly modeled, a preliminary model consists of the trend level, together with the slope, seasonal, irregular, and survey error components. The model is fine-tuned through diagnostic testing. Evaluating the test results leads to decisions regarding (1) the need for a logarithmic transformation, (2) the presence of a trend slope, an irregular component, and outliers of various forms, and (3) the composition of the seasonal component. If necessary, impulse or step dummies can be incorporated into the model to allow for exceptional temporary shocks or even more permanent shifts in the series.

For all models, the estimated variance of the irregular component was very close to zero. Therefore, the seasonally adjusted series (trend plus irregular component) is virtually identical to the trend. Thus, although there may be irregular variation in the true series, it is empirically difficult to detect in the presence of the large sampling-error variation.[3]

Real-time and historical estimates

Given estimates of the unknown parameters of the model presented in this article, seasonal adjustment of the data may commence. Because CPS data are generated each month, there are two ways to approach estimation. One way is to make an estimate each time new CPS data become available; the other way is to wait until data accumulate and then produce all the estimates at once. The first approach occurs in real time when seasonal adjustment is performed immediately after the latest data are available. (This approach is sometimes referred to as *concurrent seasonal adjustment*.) The second approach corresponds to the usual practice of revising the entire series of seasonally adjusted estimates at the end of the year once the latest data for the entire year become available. In the time-series literature, the first approach is called *filtering*, the second *smoothing*.

To produce seasonally adjusted estimates, the model is first cast into the so-called state-space form, which allows the use of the Kalman filter, a powerful algorithm for computing the estimates in real time.

A state vector includes the values of all the components of the model at a specific point in time. Two equations define the state-space form. First, there is the transition equation that specifies how the state vector, Z_t, behaves over time. This equation takes a simple form, with the current value of the state vector depending only on its value in the previous period, plus V_t, a vector containing the random disturbances representing uncertainties in the dynamics of the system. (The vector V_t is assumed to be normal and independently distributed.) Mathematically,

$$Z_t = F_t Z_{t-1} + V_t,$$

where F_t is a known transition matrix. This simple autoregressive structure may appear to be very limiting, but in fact it encompasses a wide variety of models by introducing artificial state variables.

The state variables, which are the signal and noise components, are not directly observable. Instead, the information on the state of the system is conveyed by the sample data y_t, which is related to the state vector via the observation equation

$$y_t = H_t Z_t = \mathrm{Sig}_t + N_t.$$

This is the second of the two equations. Here, H_t, which has known values, is called the *observation matrix* and sums up the components of the state vector such that the observed datum is equal to the sum of the signal (trend plus irregular component) and the noise (seasonal plus sampling error components).

The task is to find the expected values of the state vector and its variance, given the observed sample values. These are the "best" estimates of the state vector, in the sense of minimizing the mean square error (or, more precisely, the best linear unbiased predictors of the state variables). The expected value is

$$E\left(Z_t \mid y_t, \ldots, y_1\right) = \hat{Z}_{t|t},$$

and the variance is

$$\mathrm{Var}\left(Z_t \mid y_t, \ldots, y_1\right) = \hat{P}_{t|t},$$

where $\hat{Z}_{t|t}$ is the expected value of the state vector at time t, based on all of the observed data up to time t.

The solution to finding a "best" estimate from noisy data as they come in each period is given by the celebrated Kalman filter,[4] which needs only the previous period's estimates $\hat{Z}_{t-1|t-1}$ and $\hat{P}_{t-1|t-1}$ (which are based on all of the data up to time $t-1$) and the current observation y_t to compute the "best" current-period estimates $\hat{Z}_{t|t}$ and $\hat{P}_{t|t}$. This is done in two steps: a prediction step followed by an update step. First, the state vector is predicted for time t by projecting forward its previous period's estimate by means of the aforementioned transition equation, where the disturbance terms, V_t, are set to their expected values of zero.

In fact, however, the disturbances are never exactly zero, and the previous period's estimates contain errors. Accordingly, in the update step, the predictions are improved with the use of information from the current sample values. The observation equation provides the means for getting feedback from the sample by forming a prediction of the new data value as the sum of the signal and noise predictions, $\hat{y}_{t|t-1}$, and comparing it with the actual value y_t:

$$\hat{y}_{t|t-1} = H_t \hat{Z}_{t|t-1} = \hat{\mathrm{Sig}}_{t|t-1} + \hat{N}_{t|t-1}.$$

In this equation, the subscript $t|t-1$ indicates that the prediction is for time t, using data up to time $t-1$ only.

Next, an overall prediction error

$$y_t - \hat{y}_{t|t-1} = \left(\mathrm{Sig}_t - \hat{\mathrm{Sig}}_{t|t-1}\right) + \left(N_t - \hat{N}_{t|t-1}\right)$$

is computed. The error in predicting the latest y_t value using data only up to the previous period is the sum of the unobserved prediction errors for the signal and noise. In the update step, the signal and noise estimates are corrected in proportion to their expected contribution to the total prediction error:

$$\hat{\mathrm{Sig}}_{t|t} = \hat{\mathrm{Sig}}_{t|t-1} + g_t\left[y_t - \hat{y}_{t|t-1}\right];$$
$$\hat{N}_{t|t} = \hat{N}_{t|t-1} + \left(1 - g_t\right)\left[y_t - \hat{y}_{t|t-1}\right].$$

The subscript $t|t$ indicates that information from time t has been added to the estimate.

The weight g_t, which varies between zero and unity, determines how much of the prediction error is allocated to the signal and how much to the noise. It is a function of the ratio of the variance in the signal to the variance in the noise. As the noise variance gets larger, g_t gets smaller, and more weight is given to the previous period's prediction of the signal, because now the observed data at time t are less reliable and therefore the sampling error for this observation is more likely to account for the largest share of the overall prediction error. The signal and noise weights on the prediction error add to unity, ensuring that the updated estimated signal and noise add to the observed sample value.

Each period, the Kalman filter makes a prediction $\hat{y}_{t|t-1}$ of the next observed value y_t, using only the most recent estimates of the signal and noise. The filter calculates corrected estimates of these components from the prediction error, which incorporates information from the latest available data at time t. Because the predicted values are based on all the data available up to time $t-1$, the corrected estimates for time t reflect all the available information from both the historical and current values of the series. After each prediction and update step, the prediction correction process is repeated. This recursive nature is one of the appealing features of the Kalman filter. A simple illustration follows.

Filtering is tailored to real-time processing of one observation at a time as it first becomes available. Conceptually, this is analogous to updating a running average $\overline{x}_t$ of numbers $x_1, x_2, \ldots, x_t$ by the recursive formula

$$\overline{x}_t = \left(\frac{t-1}{t}\right)\overline{x}_{t-1} + \left(\frac{1}{t}\right)x_t,$$

which requires only the previous period's estimate plus the current observation, rather than by starting over each time with the formula

$$\bar{x}_t = \frac{1}{t}\left(x_1 + \cdots + x_t\right),$$

which requires all of the data. The significant point is that the Kalman filter does no more work to process the last observation than it does the first. The net result is an algorithm tailored to real-time applications, whereby data keep coming in and information about the current value of the signal is needed immediately.

The Kalman filter, however, is not well suited to producing historical estimates for a fixed set of data observations, because it is designed to produce an estimate for the current period only and not to revise any earlier estimates. Nonetheless, it is standard practice to revise a seasonally adjusted estimate made at time t by using information that had arisen subsequently, for the succeeding observations $y_{t+1}, y_{t+2}, \ldots, y_n$ up to the last observation at $t = n$ are bound to convey information about the trend and other components that can supplement the information $y_1, \ldots, y_t$ that was available at time t. Clearly, the Kalman filter estimates for the observations at the beginning of the time series will be extremely weak.

The retrospective improvement of the estimates by using ex post information is achieved by a process conveniently described as "smoothing." This is a matter of revising each of the filter estimates for a period running from $t = 1$ to $t = n$ once the full set of observations $y_1, \ldots, y_n$ has become available. These "retrospective" estimates are obtained from the "Kalman smoother," which runs the Kalman filter recursion backwards from $t = n$ to $t = 1$ through the earlier data, revising the estimates produced by filtering at each time point. Smoothing is batch processing in the sense that it operates on all of the data at once, in contrast to the Kalman filter real-time processing of one observation at a time.

Not surprisingly, the estimates from the smoother typically look "smoother" than those from the filter. For historical analysis, the smoothed estimates are superior (they have a smaller error variance) to the filtered estimates. But it is important to note that, because these smoothed estimates use data from the entire sample, they do not correspond to estimates that would have been available to data users in real time. In reality, smoothing would be done once a year, in accordance with the BLS policy of revising official seasonally adjusted estimates. Smoothing could be performed each month, but there is an obvious disadvantage to publishing the results, because all of the previous month's estimates would be revised each month, a practice that could be confusing to data users.

Notes to the appendix

[1] A more technical discussion is given in R. Tiller, "Seasonal Adjustment of CPS Time Series with Large Survey Errors," paper presented at the Federal Economic Statistics Advisory Committee Meeting, Washington, DC, December 13–14, 2001.

[2] Andrew C. Harvey, *Forecasting, Structural Time Series, Models and the Kalman Filter* (Cambridge, U.K., Cambridge University Press, 1990).

[3] Tiller, "Seasonal Adjustment."

[4] Harvey, *Forecasting*.