Estimating gross flows consistent with stocks in the CPS

The basic gross-flow table formerly used by the Bureau of Labor Statistics in examining labor market flows was expanded and the resulting tables were raked iteratively in order to produce labor market flow statistics compatible with CPS stocks

Harley J. Frazis,
Edwin L. Robison,
Thomas D. Evans,
and
Martha A. Duff

The Current Population Survey (CPS) is primarily a cross-sectional survey designed to estimate the distribution of labor force states—employed (E), unemployed (U), or not in the labor force (N)—among the population[1] for a given month. However, the CPS also can be used to examine the number of persons who change their labor force state between months.

Gross-flow estimates describe the month-to-month transitions from one labor force state to another. The following 3 × 3 matrix gives an example in which EU represents the number of persons who were employed in the previous month (May) and are unemployed in the current month (June), and similarly for the other entries:

		Current month (June)		
		E	U	N
Previous month (May)	E	EE	EU	EN
	U	UE	UU	UN
	N	NE	NU	NN

Gross-flow estimation is possible in the CPS because households are interviewed for 4 consecutive months, are then rotated out of the survey for 8 months, and are then interviewed for another 4 consecutive months. About three-fourths of the sample households are in common across 2 consecutive months. Household records can be linked, and month-to-month labor force transitions determined. for most persons in those households.

Harley J. Frazis is an economist, and Edwin L. Robison, Thomas D. Evans, and Martha A. Duff are matematical statisticians, in the Office of Employment and Unemployment Statistics, Bureau of Labor Statistics.

Gross-flow statistics from the CPS were published from 1948 until 1952. Publication was stopped because there were clear discrepancies between labor force changes derived from the flows and labor force changes derived from the monthly stock estimates. (The sources of these differences are explained later.) Over the years, many analysts have called for the Bureau of Labor Statistics (BLS) to resume publishing gross flows. This article describes a new method of obtaining flow statistics that are compatible with the monthly stock numbers. Seasonal adjustment of gross-flow series also is discussed.

Existing gross-flow data problems

The Census Bureau generates unpublished gross-flow estimates as part of its monthly production of CPS data. The current procedure used by the Census Bureau to generate the tabulations each month starts by matching respondents in the current month to respondents in the previous month; about 72 percent are matched. Next, the sampling weights of the matched respondents are adjusted so that weighted sample totals, by sex, match known population totals. The adjusted weights are then used to compute weighted estimates of labor force transition flows. The analysis that follows focuses on two types of error inherent in this procedure: classification error and margin error.

Classification error. For a variety of reasons, some CPS respondents may be classified into the wrong labor force state. Errors in classifying the respondent can have large effects on gross-flow calculations. In stock data, classification errors tend to offset each other, whereas in flow data, errors tend to be additive. For example, if equal numbers of respondents are erroneously classified as employed when they are unemployed and as unemployed when they are employed, stock data will be unaffected, but both EU and UE flows will be increased.

Although research indicates that classification error may have large effects on gross flows, the Bureau of Labor Statistics has no current plans to publish classification-error-corrected flows. While measurement error probabilities could be derived from reinterview data, it is not entirely clear how such data should be used. In their attempts to correct for classification error, John Abowd and Arnold Zellner,[2] and James Poterba and Lawrence Summers,[3] used "reconciled" reinterview data, whereby the interviewer attempted to establish a true labor force state in the case of contradiction between the original survey and the reinterview. Because of data quality problems, however, reconciled reinterview data are no longer being produced. Tin Chiu Chua and Wayne Fuller[4] used unreconciled reinterview data, but doing so requires additional statistical assumptions. Moreover, the reinterview sample may not be representative of the CPS sample as a whole, because response rates are lower than they are for the CPS. (Currently, reinterview response rates are approximately 80 percent, compared with 90–95 percent in the CPS.)

Margin error. CPS stock estimates for a given month use responses from all eight panels. Gross-flow estimates, by contrast, are restricted to the six panels in months-in-sample (MIS) groups that continue from one month to the next. These groups are MIS1–MIS3 and MIS5–MIS7 in the previous month and become MIS2–MIS4 and MIS6–MIS8, respectively, in the current month. Moreover, because the CPS does not track persons who change their residence, and because respondents may be absent or refuse to complete the survey in a given month, not all persons in the survey can be matched, even in the continuing months-in-sample.

It is natural to attempt to derive implied changes in stocks by adding up gross flows. For example, one could derive an estimate of the change in employment by adding the flows into employment (UE and NE) and subtracting the flows out of employment (EU and EN). Consistently, the implied changes in stocks do not match the changes in stocks estimated from the CPS as a whole. The gross flows tend to show net flows out of the labor force, in contrast to the stock numbers. For example, the following tabulation shows both the average (not seasonally adjusted) monthly changes in stocks directly from the CPS between December 1994 and December 2004 and the changes in stocks implied by adding up the flows calculated by the current procedure over the same period (numbers are in thousands):

	Change in—		
Method	*Employed*	*Unemployed*	*Not in labor force*
CPS	130	8	87
Current method	–203	–231	434

Whereas the stock numbers show large increases in employment and a smaller increase in the number not in the labor force, adding up the flows implies large decreases in both employment and unemployment and a large increase in those not in the labor force.

There are three sources of the discrepancy between published estimates of changes in stocks and estimates from gross-flow data. First, there is the problem of "nonidenticals": respondents who match from month to month may be systematically different from those in the relevant MIS who do not match due to nonresponse, changing addresses, and so forth. Unmatched respondents are part of the stocks, but not the flows. Second is the problem of rotation group effects: it is known that respondents' answers to questions about their labor force status systematically differ by MIS. Third, the current gross-flow method does not take into account changes in the population aged 16 years and older.

The research preparatory to this article showed that rotation group effects are the most important reason for the discrepancy between stocks and flows. It has long been known that the labor force state a respondent reports in the CPS is affected by MIS. For example, in 2003, the average weighted percentage of respondents who reported that they were not in the labor force ranged from 32.8 percent in MIS1, to 33.4 percent in MIS2, to 34.3 percent in MIS8, with corresponding decreases in both the percentage employed and the percentage unemployed.[5]

These rotation group effects have clear implications for gross flows. Because respondents in MIS1 and MIS5 cannot be matched to the previous month, the increase in respondents reporting that they were not in the labor force for other MIS's implies that matched samples will show flows out of the labor force even if there is no change in the stocks.

As mentioned, another source of margin discrepancies is that the current gross-flow method does not account for population growth and, more broadly, does not account for flows into and out of the scope of the CPS. Abowd and Zellner pointed out that the in-scope population for the CPS is not static.[6] Thus, a complete table of flows would include not only flows between labor force states, but also flows from the labor force out of the scope of the CPS (because of death, entry into the Armed Forces, emigration, and the like) and flows from out-of-scope states to a labor force state (for example, due to turning 16 and immigration). It is easy to deal with entering the scope of the survey due to turning age 16, because 15-year-olds are in the survey (and their birth dates are collected). Other flows into and out of scope are more difficult to deal with, because complete data do not exist within the CPS: persons exiting households are imperfectly tracked, and persons entering households are not asked retrospective questions.[7]

A method for correcting margin error

The basic approach set forth in this article is to adjust the flows in the matched CPS data by a method known as *raking*, so that they correspond to the labor force stocks in both the current and the previous month. The method accounts for flows into and out of the scope of the CPS and thus corrects for all sources of margin discrepancies, so that implied changes in stocks derived from the flows match changes in CPS stock estimates. As a basis for discussing this method, table 1 expands upon the basic 3 × 3 gross-flow matrix shown earlier, in order to deal with flows into and out of the scope of the CPS.

Deaths refers to those individuals who were in scope in the previous month, but had died by the current month. *Just 16* refers to those individuals who just turned 16 in the current month. The portions of the table set in roman type correspond to the estimates of flows and stocks presented here, with P denoting the previous month and C denoting the current month. The remaining parts of the table consist of indirectly estimated residual totals that are used to make the table "add up."

As noted, the four boldface estimates JE, JU, JN, and JP in the "Just 16" row can be computed directly from the CPS by using the known ages of respondents in the previous and current months, as well as their labor force status in the current month. (Note that JP is simply the stock estimate of those who just turned of age as of the current month.) Deaths are reported in the CPS, but for various reasons are undercounted by nearly half, so to estimate flows that are out of scope due to death, a less direct approach needs to be adopted. To get a more accurate estimate, average death rates for each gender are derived from mortality tables published annually by the National Center for Health Statistics.[8] These death rates are then applied to the CPS data to estimate total deaths each month (DC in table 1). Finally, the deaths are allocated among labor force states on the basis of those states' average allocation of deaths from historical CPS data, generating the boldface estimates ED, UD, and ND in the "Deaths" column of the table.

Three cells in table 1 are defined to be zero: those who would be classified as inflows, but were immediate outflows due to death, and those who would be defined simultaneously as *other inflows* and *other outflows*.

Details

This section discusses in detail the computation of raked tables. All cases presented refer to the table structure defined in table 1. Gross-flow tables are computed for men and women separately.

Margin adjustment step. Construct the stock labor force estimates EP, UP, NP, and JP for the previous month, using the previous month's sampling weight, and construct the stock labor force estimates EC, UC, and NC for the current month, using that month's sampling weight. Now construct the death estimates ED, UD, and ND by first estimating the number of deaths (the sum of ED, UD, and ND) by taking the previous month's total population and multiplying that by a death rate (obtained from records on vital statistics) appropriate to each sex. Next, distribute the total death estimate among the three labor force estimates (ED, UD, and ND) on the basis of 3-year

Table 1. Structure of basic gross flows

Row and column category			Current month: Labor force status E	Current month: Labor force status U	Current month: Labor force status N	Current month: Other outflow: Deaths	Current month: Other outflow: Other outflow	Current month: Row total
Previous month	Labor force status	E	EE	EU	EN	ED	EO	EP
		U	UE	UU	UN	UD	UO	UP
		N	NE	NU	NN	ND	NO	NP
	Inflow	Just 16	JE	JU	JN	0	JO	JP
		Other inflow	IE	IU	IN	0	0	IP
	Column total		EC	UC	NC	DC	OC	Total

NOTE: Shading indicates original 3 × 3 gross-flow matrix used as example early in text.

average estimates of the proportion of deaths by labor force classification. These estimates are obtained from the CPS.

The adjustments so far yield a total of P0 = (EP + UP + NP + JP) – (ED + UD + ND) as a potential population in the scope of the CPS during the current month. The actual in-scope population is P1 = EC + UC + NC. Usually, due to immigration, P1 will be greater than P0. If so, then *other inflow* estimates (IE, IU, and IN) are set by allocating the discrepancy (P1 – P0) to labor force states in proportion to their shares of the current month's population, and *other outflow* is set to zero. Some of the time when CPS population controls are adjusted, P0 will be greater than P1. In this case, *other outflow* estimates (EO, UO, NO, and JO) are set by allocating (P0 – P1) in proportion to EC, UC, NC, and JP, respectively (the totals by labor force state in the previous month, plus JP).

Matching step. Construct weighted counts of the 12 flow cells (EE, EU, EN, UE, UU, UN, NE, NU, NN, JE, JU, and JN), using the sampling weight for the current month for those individuals who had a labor force status in both the previous and the current month. These totals are constructed for each sex. The weighted flow counts will be too small by approximately 25 percent, because about 75 percent of the sample overlaps from month to month. The initial iteration step presented next will correct this undercount.

Iteration step. The table constructed in the matching step is not entirely consistent, because not all the cells are guaranteed to add up to the appropriate row and column totals. In order to obtain consistency, iterative raking is performed. The death estimates (ED, UD, ND, and DC) and the inflows or outflows are held fixed, while the remaining interior cells of the tables are raked by iterative proportional fitting to ensure additivity to the stock estimates in the row and column totals.

Final factor calculation. The procedure presented next produces gross-flow microweights. Compute factors for each of the 12 raked cells (EE, EU,..., JN) for each table (for men and for women) by dividing the final estimate obtained in the iteration step by the initial cell value (the weighted sample count, using the current-month sampling weight, for those individuals who are in the sample both months). When applied to the current-month weights of matched individuals, the factors yield gross-flow weights that enable the re-creation of the flows for both men 16 years and older and women 16 years and older. The computation of flows for demographic groups other than those two will be facilitated by these weights, but will not be perfectly consistent with CPS stock numbers.

Results

The foregoing procedure was implemented with data from 1990 through 2004. Table 2 shows a comparison of row percentages—flows as a percentage of the population in a labor force state the previous month—for the current procedure and for the procedure just set forth.[9] As can be seen, the average differences are fairly small, 7 percent or less in magnitude. This is to be expected, because the aim of the raking procedure is to revise the existing flow data as little as possible while forcing compatibility with the margins. Earlier, it was noted that rotation group bias caused the current procedure to show more substantial flows out of the labor force than could be reconciled with the stock data. Table 2 reflects this movement, because the new procedure slightly reduces flows out of the labor force. In some months, there are more substantial differences. For example, in June 1992, NU flows were 3.7 percent of the previous month's not-in-the-labor-force stocks under the current procedure, but 4.2 percent under the new procedure, a 12-percent increase. This discrepancy is associated with an unusually high increase in unemployment, from 9.4 million to 10.4 million. Chart 1 shows the time series of NU flows. The June 1992 point in the adjusted series is a multiyear high, whereas the point in the current series is not much different from other Junes.

Table 2. Gross-flow row percentages, 1990-2004 CPS, different methods of estimation

Flow	Average row percentage, current method	Average row percentage, adjusted method divided by current method	Average percent increase, adjusted method divided by current method	Average absolute percent change, adjusted method divided by current method	Minimum percent increase	Maximum percent increase
EE	95.8	95.9	0	0	0	0
EU	1.4	1.4	4	4	–2	13
EN	2.8	2.7	–4	4	–11	6
UE	27.5	27.3	0	2	–5	4
UU	49.7	50.9	2	2	–2	6
UN	22.8	21.8	–4	4	–10	0
NE	4.7	4.9	4	4	–6	15
NU	2.5	2.7	7	7	1	16
NN	92.8	92.4	0	0	–1	0

Chart 1. NU flows, 1990–2004, row percent

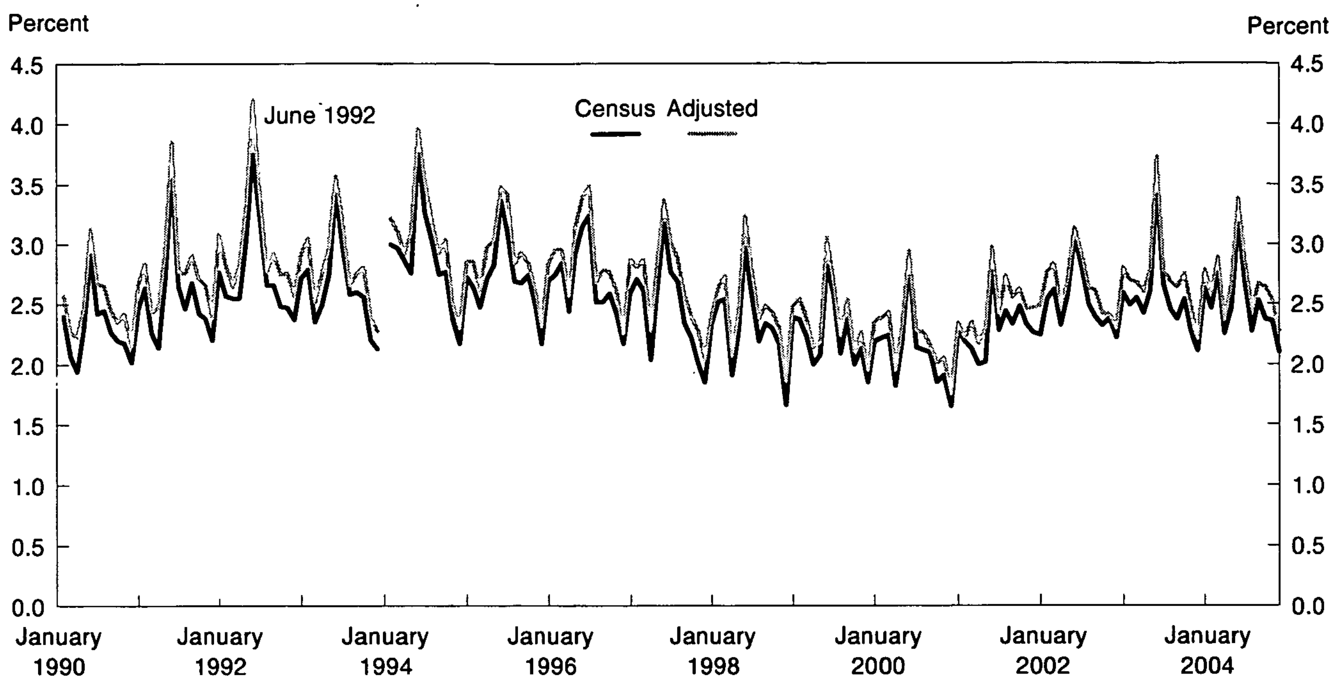

Seasonal adjustment

Seasonal adjustment is a natural way to further improve the utility of gross flows, and it is also highly useful as an analysis tool to evaluate data quality and to uncover hidden or hard-to-see characteristics. One major wrinkle in seasonally adjusting gross-flow data is that the adjusted flow data need to be made compatible with seasonally adjusted stocks. Accordingly, once the raked gross flows are seasonally adjusted, they are raked a second time, this time to the total seasonally adjusted stock numbers from the CPS. Final raked and seasonally adjusted flows were examined as part of the research for this article, because it was not known exactly how reraking would affect the flows. It is important that extraneous seasonality or any other systematic effects not be introduced and that the seasonal factors remain reasonably stable from year to year.

Seasonal adjustment is performed with the Census Bureau's X-12-ARIMA program, the current standard at the Bureau of Labor Statistics.[10] Data from February 1990 to March 2004 were examined, with a missing data point in January 1994 due to a survey redesign. Initially, the research analyzed flows for a variety of demographic breaks by age, sex, and race. However, because these breaks resulted in smaller sample sizes and flow estimates that were subject to greater sampling error, it was determined that only flows for men 16 years and older and women 16 years and older could be adequately seasonally adjusted.

The 3 × 3 flows between the labor force states were seasonally adjusted. The X-12-ARIMA procedure indicated that all of the flows showed seasonality, although it was relatively weak for UN flows for both men and women and UU flows for women. These flows may be even more seasonal than is indicated here, but sampling error is likely masking some seasonal properties, because the diagnostic statistics indicating the presence of seasonality increase when flows for men and women are added. Flows with the largest seasonal factors are EU, EN, and NE, while EE and NN factors are by far the smallest. In general, the seasonally adjusted flows appear to have explainable seasonal patterns.

A few of the flows show signs of seasonal factors that are unstable from year to year. This instability can be expected, because, clearly, some of the smaller flows are fairly noisy. Plotting the seasonal factors for each month over the years of the sample period showed no factors unstable enough to cause serious concern.

An examination of the seasonally adjusted flow data revealed another pattern. The CPS is fielded on the week containing the 19th of each month[11] and asks questions referring to the week of the 12th. Thus, CPS reference weeks are separated by either 4 or 5 weeks. The adjusted series revealed systematic differences in the flows, depending on the distance between the reference periods. Note that such differences are quite plausible: respondents have more opportunities to change their labor force status, and thus end up in the off-diagonal cells, when reference periods are 5 weeks apart. Population growth also implies that reference-week intervals can affect levels of all the cells. Modeling these calendar effects often results in smoother seasonally adjusted series.[12] In this case, modeling the calendar effects with dummy variables resulted in seasonally adjusted series with no apparent residual seasonality.

After seasonal adjustment, the flows for men and women are raked to their respective published CPS seasonally adjusted controls. The complete matrix for the flows, as shown in table 1, has a row for 15-year-olds turning 16 (JE, JU, and JN). These series are too small to adjust reliably, so, as reasonable approximations, CPS seasonal factors for 16-to-19-year-old employed, unemployed, and not-in-the-labor-force men and women are applied to the cells corresponding to those categories in the "Just 16" row of the matrix. Other cells for flows into and out of the scope of the survey are assumed to be nonseasonal.

Once the seasonally adjusted flows were raked, they were reexamined for differences with the unraked series, residual seasonality, and stability of the seasonal factors. Basic statistics for differences between the unraked seasonally adjusted flows and the raked seasonally adjusted flows are presented as row percentages (as in table 2) in table 3. The overall effect of raking seasonally adjusted flows to seasonally adjusted stocks is less than the effect of raking the flows that are not seasonally adjusted, as is shown in table 2. The mean percentage difference between raked and unraked seasonally adjusted flows is in all cases less than 0.5 percent, and the mean absolute percentage difference is in all cases less than 2 percent. None of the percentage differences is large; the largest differences between raked and unraked flows were approximately 5 percent. Not surprisingly, the two largest flows have by far the smallest percentage differences.

The final check on the raked flows was to see whether seasonal or calendar patterns were reintroduced by raking. A comparison of the seasonally adjusted flows and the raked seasonally adjusted flows showed some differences along this dimension. Unfortunately, it appears that some of the calendar effects are reintroduced into the final raked seasonally adjusted series, but not enough to cause serious concern.

MORE THAN 50 YEARS AGO, the Bureau of Labor Statistics stopped publishing series of labor market flows when it was discovered that the published flows were incompatible with the monthly labor force stock numbers. This article describes efforts to produce labor market flow statistics from the CPS without the incompatibilities that led to their curtailment. The basic gross-flow table was expanded to estimate flows into and out of the scope of the CPS. The resulting gross-flow tables were then raked in an iterative process to match both the previous month's and the current month's stock estimates. As part of this project, a method was developed to seasonally adjust the flow series while maintaining the flows' compatibility with the seasonally adjusted stocks. The Bureau of Labor Statistics is continuing to work on this research, with the expectation of eventually resuming publication of the gross flows.[13] □

Table 3. Seasonally adjusted gross-flow row percentages, 1990–2004 CPS, unraked and raked

Flow	Average row percentage, seasonally adjusted	Average row percentage, seasonally adjusted and raked	Average percent increase; adjusted divided by raked	Average absolute percent change, adjusted divided by raked	Minimum percent increase	Maximum percent increase
EE	95.9	95.9	0	0	0	0
EU	1.4	1.4	0	1	–5	5
EN	2.7	2.7	0	1	–4	3
UE	27.4	27.4	0	1	–2	3
UU	50.8	50.8	0	1	–3	3
UN	21.8	21.8	0	1	–3	3
NE	4.9	4.9	0	1	–3	4
NU	2.7	2.7	0	1	–5	5
NN	92.4	92.5	0	0	0	0

Notes

ACKNOWLEDGMENT: The authors thank Fran Horvath of the Office of Employment and Unemployment Statistics for his contributions to the research preparatory to this article.

[1] "Population" refers to the U.S. civilian noninstitutional population aged 16 and older.

[2] John M. Abowd and Arnold Zellner, "Estimating Gross Labor-Force Flows," *Journal of Business and Economic Statistics*, July 1985, pp. 254–83.

[3] James M. Poterba and Lawrence Summers, "Reporting Errors and Labor Market Dynamics," *Econometrica*, November 1986, pp. 1319–38.

[4] Tin Chiu Chua and Wayne A. Fuller, "A Model for Multinomial

Response Error Applied to Labor Flows," *Journal of the American Statistical Association*, March 1987, pp. 46–51.

[5] Here is an example that shows how little difference matching makes relative to the rotation group effect. In 1996, the average monthly change in the percentage of the population that was employed was 0.16 percent, the average change for the percentage unemployed was –0.07 percent, and the average change for the percentage not in the labor force was 0.09 percent. If, instead, one computes the same figures, but this time using the changes in percentage from rotation groups MIS1-3 and MIS5-7 in the first month to rotation groups MIS2-4 and MIS6-8, respectively, in the next month, whether the observations do or do not match between months, the numbers are 0.02 percent for the employed, –0.16 percent for the unemployed, and 0.13 percent for those not in the labor force. (Due to rounding, the preceding numbers may not add to zero.) If one now computes the same figures, but using only matched observations from rotation groups MIS1-3 and MIS5-7 in the first month and rotation groups MIS2-4 and MIS6-8 in the next month, the numbers are 0.03 percent for the employed, –0.17 percent for the unemployed, and 0.14 percent for those not in the labor force—very close to the unmatched figures, showing that failure to match makes little additional difference after accounting for the rotation group.

[6] Abowd and Zellner, "Estimating Gross Labor-Force Flows."

[7] The CPS uses an address-based sample and does not attempt to interview persons who moved out of the sampled address. If a household member exits the household between interviews, a remaining household member is asked the reason for the exit. If the entire household is replaced, no such question is asked.

[8] National Center for Health Statistics, *Vital Statistics of the United States* (Atlanta, Centers for Disease Control and Prevention, published annually).

[9] Flows from April 1994 are missing in the current series; January 1994 also is excluded, due to linking problems associated with a redesign of the survey.

[10] *X-12-ARIMA Reference Manual*, version 0.2.10 (Census Bureau, 2002).

[11] For those Decembers when this date would conflict with holidays, the survey is fielded during the week that includes the 12th of the month.

[12] Stephanie Cano, Patricia M. Getz, Jurgen Kropf, Stuart Scott, and George Stamas, "Adjusting for a Calendar Effect in Employment Time Series," *Proceedings of the Survey Research Methods Section, 1996* (Alexandria, VA, American Statistical Association, 1996), explain how this approach is similarly implemented in seasonal adjustment for the BLS Current Employment Statistics survey. For an in-depth discussion of calendar effects in the present case, see Thomas D. Evans, "Analysis of Raking on Seasonally Adjusted Household Gross Flows Data," *ASA Proceedings of the Joint Statistical Meetings, 2004*, (Alexandria, VA, American Statistical Association, 2004), pp. 1166-73.

[13] The utility of the new gross-flow series to analysts is discussed in the companion article to this one, "Analyzing CPS data using gross flows," by Randy Ilg, this issue, pp. 10–18.