Addressing the Demand for Time Series and Longitudinal Data on Occupational Employment[*]

Katharine G. Abraham
University of Maryland and NBER

James R. Spletzer
U.S. Bureau of Labor Statistics

October 22, 2009
Revised February 26, 2010

[*] Without implicating them in any way, we are grateful to Laurie Salmon, Dixie Sommers and George Stamas, Office of Employment and Unemployment Statistics, Bureau of Labor Statistics for useful conversations and background information concerning the Occupational Employment Statistics survey.

I. Introduction

Among the important potential effects of increased off-shoring are changes in the occupational composition of U.S. employment. To the extent that firms choose to shift particular tasks to workers located overseas, domestic employment in the occupations that perform those tasks can be expected to decline or to grow less rapidly than would otherwise have been the case. It has been suggested that time series data on occupational employment by industry could be useful for studying these effects in the aggregate. Similarly, longitudinal data on the mix of jobs at individual enterprises could be useful for better understanding the dynamics of outsourcing and offshoring at the level of the individual firm. The development of this sort of information has been recommended as part of a broader set of needed improvements in the data available for the study of off-shoring (see, for example, Sturgeon 2006 and National Academy for Public Administration 2006).

Concerns about sample size and the accuracy with which household respondents report their occupation and industry, together with the fact that household survey data generally cannot be used to study the evolution of employment at individual firms, have led analysts interested in the effects of off-shoring on domestic employment to focus on employer-provided employment data. The Occupational Employment Statistics (OES) survey conducted by the Bureau of Labor Statistics (BLS) is a large employer survey that, each year since 1988, has collected detailed information on employment by occupation.

The OES is designed to produce detailed point-in-time estimates of staffing patterns and wages, not to produce occupational employment time series or to support the analysis of changes in the occupational composition of employment at individual workplaces. For some applications, such as the use of OES employment data to determine weights in the BLS National

Compensation Survey program or the use of OES wage data by the Employment Standard Administration to administer the H-1B visa program, having data that could be compared over time would not be especially valuable. Many users of OES data, however, clearly would benefit from data designed to support cross-year comparisons. For example, annual data designed to track trajectories in staffing patterns would be of great value to the BLS Occupational Employment Projections (OEP) program. Similarly, access to information on trends in occupational staffing patterns and wages would help those who use the data for workforce development, career counseling and career planning purposes. Researchers studying organizational behavior, the sources of productivity growth and other topics could benefit from data that allowed them to track staffing patterns at the level of the individual establishment or individual enterprise.

The purpose of this paper is to evaluate the OES as a source of time series and/or longitudinal data on employment by occupation. This sort of information would be useful not only to identify off-shoring activity and study its impacts, topics that are the focus of the present conference, but for a variety of other purposes. In Section II, we briefly describe the OES survey. Section III discusses the feasibility of using the historical OES data to construct occupational employment time series or for longitudinal analysis at either the establishment or the enterprise level. Unfortunately, the existing survey design and the management approach dictated by current program objectives make the data poorly suited for the analysis of trends, especially over short time intervals. Section IV considers how the OES survey might be re-configured to produce reliable annual time series data and support analysis at the individual establishment or enterprise level. A necessary step would be to collect data from some subset of establishments every year, rather than only once every three years as is the current practice, but

other changes in the survey also would be required. Some concluding thoughts and observations are offered in Section V.

II. The Occupational Employment Statistics (OES) Survey

The Occupational Employment Statistics (OES) survey is an ongoing mail survey conducted by the BLS in collaboration with its state partners. The survey covers all industries exclusive of agriculture. Prior to 1996, industries were surveyed on a three-year rotating cycle. Since 1996, each year's sample has covered establishments across all industries, but except for an annual enumeration of federal and state government establishments, it is still the case that even very large establishments are surveyed only once every three years.

The OES sample is designed to support cross-sectional estimates of staffing patterns developed from data collected over a three-year period. Through 2001, estimates were based on three annual panels, each consisting of approximately 400,000 establishments; within each panel, establishments were assigned an October, November or December reference date. In 2002, the survey transitioned to a design in which estimates are based on six semi-annual panels, with each panel consisting of approximately 200,000 establishments assigned either a May or a November reference date. The May 2008 published estimates, for example, rest on data collected for November 2005, May 2006, November 2006, May 2007, November 2007 and May 2008. Estimates are benchmarked to the average of the most recent May and November employment levels.

Since 1996, the OES has collected information on occupational wages in addition to occupational employment. Establishments selected for the OES are asked to report employment in each cell of a matrix in which the rows refer to different occupations and the columns to wage

intervals. Generally, for firms with 20 or more employees, the survey forms contain between 50 and 225 occupations, depending on the industry of the establishment completing the form. Prior to 2000, employers receiving these forms were asked to list numerically significant or new occupations that could not be reported in a detailed occupation and therefore were reported in an "all other" residual category. This information was used in revising the survey forms for later years. Beginning in 2000, employers have been asked to provide detailed occupational information for workers who cannot be placed in one of the listed occupations.

Since 1999, small employers have received a shorter unstructured form that contains no list of likely occupation titles; rather, the employer is asked to provide a brief description of each occupation represented in the establishment's workforce. The information on these forms is coded into occupational categories by survey staff in the state agencies.[1] Multi-establishment firms may request that their data be collected through the firm's corporate headquarters rather than directly from individual establishments. This is referred to as central office collection (COC). COC reporters provide the OES program with electronic records containing job title and wage information for their employees. The OES staff then builds crosswalks for coding these firms' data into SOC occupations and OES wage intervals.

Approximately 80 percent of establishments sampled for the OES provide usable responses; on an employment-weighted basis, the survey response rate is approximately 75 percent. Nearest neighbor hot deck procedures which take data from another similar establishment are used to impute missing employment information for establishments that do not respond. Missing wage distributions also are imputed using distributions for similar establishments.

[1]Prior to 1999, several states developed their own unstructured short forms that were used to collect data from some small employers, but this was not a part of the formal survey protocol. Beginning in 2004, states were given the discretion to send unstructured forms to establishments with up to 49 employees.

OES data are published by occupation, industry and area. Employment and wage estimates are produced for as many as about 800 occupations, both nationally and for states, metropolitan areas and other geographic areas. In addition, national occupational employment and wage estimates are available for specific industries. BLS does not publish occupation by industry data below the national level and suppression of sparsely-populated cells is common in both the national by industry tables and the cross-industry tables for sub-national areas. The OES program switched from its own survey-specific occupation coding system to the Standard Occupational Classification (SOC) system in 1999 and from the Standard Industrial Classification (SIC) system to the North American Industry Classification System (NAICS) in 2002.

III. Can Existing OES Data Be Used to Construct Annual Time Series?

Several features of the current OES design make it difficult to use the existing data to construct single-year estimates of employment by occupation even at the national level. The current design calls for estimates to be produced using three years of data and the existing survey weights are not suitable for the production of annual estimates. It is possible to construct annual estimation weights, but because, except in federal government, even very large units are surveyed only once every three years, annual estimates tend to be quite variable, especially for more disaggregated employment cells. The significant breaks in both occupation and industry classification caused by the adoption of the SOC and the NAICS are another problem. Finally, other changes in survey operations associated with the adoption of the SOC have affected the comparability of the OES estimates over time. To preview our conclusions, we believe that it is possible to use existing OES data to construct national occupational time series that are suitable

for some analytic purposes, but that these estimates have serious limitations for studying offshoring and its effects.

Lack of Weights for Annual Estimates

The weights used to produce official OES estimates are constructed at the level of cells defined on the basis of industry, establishment size and geography. As noted above, the sample units used to produce each set of estimates are divided into panels spread across three years of data collection. Each sampled establishment is assigned a current weight that reflects its probability of selection into a particular panel.[2] If every cell in a panel contained at least one establishment, the weighted sum of employment calculated for an industry using the current weights would be approximately equal to total national employment in the industry as of the panel reference date(s). There are, however, a very large number of OES sampling cells – as of 2004, the survey was stratified by 343 industries, 7 establishment size classes and 686 metropolitan or balance-of-state geographic areas – and individual panels contain a significant number of empty cells. Because employment in the cells that happen to be empty is not represented, using the current weights to estimate employment in an industry based on the responses to any single panel yields an estimate that lies significantly below the industry's true employment level.[3]

Working with OES data for the private sector over the period 1996 through 2004, Abraham and Spletzer (forthcoming) developed weight adjustment factors to be applied to the OES current weights that were calculated as follows:

[2] The current weights also incorporate adjustments for differences between the way a unit was sampled and the way it was reported (e.g., one establishment at a company sampled but data reported for several establishments together).
[3] In the official estimates, which are based on three years of data, this is not generally a problem because data at the detailed cell level are re-weighted to account for the number of panels in which each cell is represented.

(1) $$ADJFACTOR1_{jt} = \frac{E_{jt}^{CES}}{\sum_i CURRWT_{ijt}^{OES} E_{ijt}^{OES}}$$

where ADJFACTOR1 is the industry weight adjustment factor, E is employment from either the monthly Current Employment Statistics (CES) survey or the OES, CURRWT is the current weight from the OES data file, i indexes individual establishments, j indexes detailed industries and t indexes years. Estimates produced from the OES microdata using weights equal to the product of ADJFACTOR1 and CURRWT reproduce CES national industry employment trends.

In all years, industry weight adjustment factors were calculated at the most detailed industry level for which sample data were available. The SIC classification structure in use through 2001 included 934 detailed industries. In 1996, taking that year as an example, weight adjustment factors were calculated at the 4-digit (most detailed) level for 310 industries, representing 34.0 percent of employment; at the 3-digit level for 383 industries, representing 37.6 percent of employment; at the 2-digit level for 225 industries, representing 27.8 percent of employment; and at the 1-digit level for 16 industries, representing 0.6 percent of employment. The NAICS structure adopted in 2002 includes 1171 detailed industries. In 2004, weight adjustment factors were calculated at the 5-digit (most detailed) level for 424 industries, representing 36.1 percent of employment; at the 4-digit level for 520 industries, representing 47.1 percent of employment; at the 3-digit level for 172 industries, representing 9.8 percent of employment; and at the 2-digit level for 55 industries, representing 7.0 percent of employment.

A further weighting concern is that, although the true distribution of employment by size of establishment appears to have been very stable from 1996 through 2004, the distributions in the data collected for the OES vary considerably from year to year. Factors that appear to have contributed to this variability include the uneven distribution of the largest (certainty) units

across panels; the effects of a 1999 experiment carried out in selected states to determine the feasibility of collecting data from all certainty establishments every year; and the introduction of establishments with 1-4 employees into the survey sample in 1998 (these very small establishments previously had been represented by establishments with 5-9 employees). A second weight adjustment factor was developed to calibrate the share of employment in broad industries accounted for by each of nine establishment size classes to the average share for that size class across the OES benchmark data files for 1998, 2001 and 2004:[4]

$$(2) \qquad ADJFACTOR2_{kst} = \frac{AVESHARE_{ks}^{BMK}}{SHARE_{kst}^{OES}}$$

where ADJFACTOR2 is the size class weight adjustment factor, AVESHARE is the average share of employment accounted for by the designated size class in the benchmark data, SHARE is the current year share in the OES data, k indexes broad industry, s indexes establishment size class and t indexes year. Applying both the industry and the size class adjustment factors yields:

$$(3) \qquad FINALWT_{ijkst} = ADJFACTOR1_{jt}\ x\ ADJFACTOR2_{kst}\ x\ CURRWT_{ijkst}^{OES}$$

Anyone interested in using the historical OES data to construct an annual time series would need to apply some similar procedure to produce weights suitable for annual estimates.

Variability in Annual Estimates

For sufficiently large domains, the adjusted weights just described appear to perform adequately. In cells defined at the national level using broad industries and occupations, employment estimates calculated using these weights seem generally to behave very sensibly (see appendices C and D in Abraham and Spletzer, forthcoming). Because even very large units

[4] The OES data are benchmarked to the quarterly Census of Employment and Wages. Our analysis rests on OES benchmark data files for the years 1998, 2001 and 2004. The size class distributions observed across these three years are very similar.

in the private sector are surveyed only once every three years, however, annual estimates of employment for detailed occupations, estimates for subnational areas and/or estimates for occupation by industry are likely to behave more erratically. To illustrate the potential instability in estimates for small domains, we used the adjusted weights just described to construct annual time series for the ten occupations identified by Jensen and Kletzer (forthcoming) as most off-shorable. These are shown in Figure 1. Each of the ten panels in the figure contains two employment series – one created from year-specific OES microdata with weights adjusted using the method just described and one created from the November OES press releases posted to the BLS website.[5] We should note that using the published OES estimates in this way is not recommended by the BLS, which states on its website that it "does not use or encourage the use of OES data for time series analysis" (http://www.bls.gov/oes/oes_ques.htm#Ques29). Still, we believe that the comparisons shown in the figure are informative.

The estimates in Figure 1 cover the years from 1999, the year that the SOC was adopted in the OES program, through 2004. One difference between the two series is that the microdata estimates cover only employment in the private sector, while the published estimates cover employment in federal, state and local government as well as the private sector. While this difference in coverage has a noticeable effect on the level of some of the series – most obviously, the series for statisticians and medical transcriptionists – it should not much affect their year-to-year variability. More importantly for the series' variability, the microdata estimates are based on establishments that, from 1999 through 2001, represented only about one-third of the full

[5]Through 2001, all OES data were collected with a reference date in October, November or December. In 2002, the OES switched to May and November reference dates. For these later years, in an effort to avoid problems of comparability associated with seasonal differences in staffing patterns at different times of the year, we used only the data from the November panel. The 2003 and 2004 press releases report statistics benchmarked to average employment by industry for the most recent May and November and there may be an issue of comparability between the published estimates for these years as compared to earlier years.

private sector OES sample and, from 2002 through 2004, only about one-sixth of the full private sector sample. The time series created with annual microdata are considerably more volatile than the time series created from the published data. Several of the occupational time series based on annual microdata – those for Mathematical technicians; Credit authorizers, checkers and clerks; Biochemists and biophysicists; Title examiners; Weighers; and Actuaries – show sharp changes from one year to the next that are not apparent in the published estimates based on multiple years of data. For nine of the ten occupations in Figure 1 – excluding only Statisticians – the variance of the series created from the annual microdata is higher than that for the published estimates. Figure 1 suggests that employment time series for detailed occupations that are created from single-year OES microdata are likely to be highly volatile, making them problematic for policy analysis of the effects of offshoring. Increases in the size of the OES sample would be needed to reduce the variance of annual employment estimates.

Breaks in Occupation and Industry Classification Systems

Breaks in both the occupation and industry classification structures are an additional barrier to using the historical OES data to produce detailed annual time series. As already noted, prior to 1999, the OES used its own classification structure; the Standard Occupational Classification (SOC) was introduced in 1999. The North American Industry Classification Structure (NAICS) replaced the older Standard Industrial Classification (SIC) system in 2002. In both cases, the new classification structure was very different from the old. Of the 769 detailed occupations included in the SOC when it was introduced in 1999, only 374 could be cross-walked to occupations that previously existed in the old OES classification structure (Bureau of Labor Statistics, 2001a, p. 24 and p. 175) During the transition to NAICS at the Bureau of

Labor Statistics, only about half of establishments could be assigned NAICS codes based on their SIC classification (Mikkelson, Morisi and Stamas, 2000).

Comparisons at an aggregated level seem more feasible across the classification structure breaks than do more disaggregated comparisons. Matthew Dey of the Bureau of Labor Statistics, for example, has developed a concordance that links cells defined using 19 aggregated occupations and 13 aggregated industries that appear to be reasonably consistent across the breaks in classification system, and Abraham and Spletzer (forthcoming) use a modified version of the Dey concordance. There are detailed occupations within these larger groupings that are defined in the same way in the SOC as in the older OES classification structure, but as already noted there are also many detailed occupations for which no direct linkage is possible. The OES data were never dual coded using the SOC and OES occupational classification structures, so in cases where occupations are not comparable it is difficult to relate the new SOC occupations to the older OES occupations.

The best option for extending the OES annual estimates for detailed occupations back through time likely would be to plot the employment series for each of the occupations and retain those occupations for which there is no evidence of a discontinuity between 1998 and 1999 that might indicate a lack of comparability in how the occupation was defined. Drawing clear conclusions would be difficult, however, because of the substantial underlying variability in annual estimates based on the OES microdata that has already been discussed. Moreover, since industries were surveyed on a three-year rotating cycle prior to 1996, even under the best-case scenario, it would be possible to extend true annual series by only three years, back to 1996.

Other Comparability Issues

In addition to a new coding structure, the SOC also introduced a set of principles intended to guide the classification of workers (Office of Management and Budget, 2000). An important principle is that only individuals who devote at least 80 percent of their time to management activities are to be classified as managers. To implement this SOC guidance, the OES introduced new edit checks to flag establishments that reported employment in a management occupation (e.g., financial manager) without reporting employment in any of the expected subordinate occupations (e.g., financial specialists or clerks). A second set of edit checks was developed to flag establishments with an excessive number of managers. Both sets of edit checks were applied in a limited fashion in 1999 and phased in more fully over the following years.

Implementation of the SOC also included new training designed to explain its structure and coding principles to program staff. Staff who attended SOC training courses in 1999 and subsequent years were instructed that management jobs reported on establishment schedules that did not include an intervening layer of supervision generally would need to be recoded as something else. Management jobs recoded in accord with this advice typically were shifted either to one of the professional occupations or to a first-level supervisor occupation. Because survey program staff code all of the occupations reported on unstructured survey forms submitted by small establishments, the introduction of the unstructured forms as an option in 1999 may have amplified the effects of the SOC training on the OES management employment series. The “rule” that no job should be coded as a management position unless the schedule also includes a first-level supervisor position is easy to apply and seems to have been embraced as a guide to coding the unstructured schedules.

Our best assessment is that the combined effect of these changes was to reduce management employment by very roughly 2 million jobs between 1998 and 2001, with these jobs then assigned instead to other job categories. Although perhaps less important for the analysis of offshoring than for some other purposes, this is nonetheless an additional barrier to comparing OES estimates over time. In our earlier work (Abraham and Spletzer forthcoming), we developed a procedure for "reverse-engineering" the new coding rules that involved reclassifying a sufficient number of non-management jobs in 1999 and later years as management positions to offset the sharp decline in management employment that is evident in the unadjusted data between 1999 and 2001. This effort was unavoidably crude. Further, going forward, insofar as it involves putting new OES data onto the old basis rather than putting old data on the new basis, the application of this procedure would inevitably become less and less appealing.[6]

New SOC training introduced in 2007 makes clear that someone might legitimately be performing management duties without there being an intervening layer of supervision between them and their subordinates. While this change was made to improve coding accuracy, going forward, it too may adversely affect the comparability of the OES data over time.

Problems with Coverage of Units Surveyed

A further limitation of the historical OES data is that the survey sample is not designed to support longitudinal analysis. As has been noted, establishments are asked to report on a three-

[6] The specific jobs that formerly might have been categorized as management jobs were identified by looking for the highest-paid non-management jobs in establishments with too few managers, based on historical patterns. This procedure took jobs from many different occupations and reclassified them as management positions (many to one). We were not able to devise a methodology for putting the old data on the new basis. That would have involved reassigning management positions to many different possible alternate occupations (one to many) and we did not believe we had a sound basis for making such assignments.

year cycle. Large establishments responding to the survey in year t are likely also to have responded in year t-3. Because even large establishments are observed only once every three years, there may be a long lag before important changes in staffing patterns are captured. Further, the establishment is not the obvious unit of analysis for identifying and tracing the effects of offshoring. In a large corporation, sourcing decisions are likely to be made at the corporate level, rather than the establishment level, and a decision to offshore work could take the form of shuttering an entire establishment, rather than transferring portions of the work performed at an individual establishment to another company. In this case, offshoring could not be identified through an establishment-level analysis, but only through an examination of changes in staffing patterns for the company as a whole.

To the extent that the appropriate level of analysis is the company rather than the establishment, the OES suffers from the further limitation that data generally are not collected for all of the establishments at a firm (National Academy of Public Administration 2006). Over the course of a three year period, all establishments large enough to belong to the survey's certainty strata are asked to complete an OES questionnaire, but large firms include many small establishments, and only a fraction of these small establishments would be surveyed even once in a three-year cycle. Among firms in the United States with more than 10,000 employees, for example, in May 2006 there were 377,484 establishments with fewer than 250 employees, accounting for 53.6 percent of these firms' employment; 7,369 establishments with 250 to 499 employees, accounting for an additional 13.9 percent of the firms' employment; and 2,727

establishments with 500-999 employees, representing a further 10.1 percent of the very large firms' employment.[7]

IV. Redesigning the OES for Time Series and Longitudinal Analysis

To produce reliable annual time series, the OES sample would need to be augmented with units that are surveyed every year rather than only once every three years. Being able to track changes in staffing patterns at the establishment level might be of some interest, though as just discussed important developments could be missed with an establishment-level focus. To the extent that analysts are interested in using enterprise-level data to study the effects of offshoring, it would be desirable for at least some firms to provide comprehensive data for all of their establishments. In addition to changes in the design of the survey sample to support time series and/or longitudinal analysis, other changes in the survey's focus also would be necessary.

Sample Redesign Options

We consider three possible options for augmenting the OES sample. The first is intended to be responsive to the demand for a sample that could support annual OES time series; the second suggests an approach to collecting data suitable for longitudinal analysis; and the third is a hybrid of the first two approaches. Many variations on these approaches can be imagined and we do not mean to suggest that the specific options we describe are the only ones that are possible. Further, because it is the information we were able to access given our time constraints, we use data on number of firms and number of establishments from 2006 in our rough cost calculations, though since the size distribution of establishments is very stable over

[7] These figures and others cited in the text are based on tabulations of the Business Employment Dynamics database, which is based on the Quarterly Census of Employment and Wages (QCEW) data file and are designed only to be illustrative.

time and total employment has changed little on net between 2006 and the present, we do not believe our answers would have been much different had we been able to use 2009 data for these calculations. Our intent in any case is to stimulate thinking about sample design alternatives and the rough magnitude of the costs that might be associated with different choices, rather than to recommend a specific plan and attach a specific dollar cost figure to that plan. The rough dollar cost figures we present refer only to the direct costs of additional data collection and do not include the costs of other staffing necessary to edit and process the additional data collected. All of the options considered focus on increases in the size of the sample for the *private sector,* as that is where we would expect the effects of offshoring to be manifested.

Option 1: Survey all large private sector establishments every year. One way to increase the stability of annual OES estimates over time would be to survey all establishments with more than some threshold level of employment every year. Suppose, for example, that the survey were redesigned to collect data from all establishments with 250 or more employees every year. As of May 2006, the universe of private sector establishments eligible for inclusion in the OES included 30,639 establishments with 250-499 employees, 10,894 establishments with 500-999 employees, and 5,470 establishments with 1,000 or more employees. Since even units included in the current certainty strata are surveyed only once every three years, surveying all units with 250 or more employees every year would represent a significant increase in the annual survey sample.

Most of the work of collecting OES data is done by the states. Payments to the states for their work on the OES program totaled $21.5 million in FY2009 (the fiscal year in which the 2008 estimates were published). This figure includes some state overhead expenses and excludes

some modest expenses for data collection work performed by the BLS national office staff, but for present purposes, we treat the payments made to the states as the cost of OES data collection. The money awarded to states is allocated using a formula that takes into account the number of establishments surveyed in the state, the size distribution of those establishments, the number of publication areas in the state and the average wages of state employees. In this allocation formula, establishments with 250-499 employees are treated as equivalent to two units, establishments with 500-999 employees as equivalent to three units and establishments with 1,000 or more employees as equivalent to four units (personal communication with Laurie Salmon, OES program staff member, September 9, 2009). By our calculations, as a very rough approximation, the current allocation formula implies a cost of $50 per establishment of size up to 249 employees, $100 per establishment of size 250-499 employees, $150 per establishment of size 500-999 employees and $200 per establishment of size 1,000 or more employees. Applying these cost estimates, to survey all private sector establishments with 250 or more employees every year would have added roughly $3.9 million dollars to the cost of data collection for the OES program.[8]

Without a more detailed analysis that would be beyond the scope of the present exercise, we cannot say precisely how collecting data from all private sector establishments with 250 or more employees every year would affect the variance of annual estimates, but the number of employees for whom data was collected each year clearly would increase significantly. The current sample includes establishments with total employment of approximately 20.0 million per year over the three years of the survey cycle; the sample augmentation just described would add

[8] We should emphasize that the program does not receive a budget based on the number of units included in the survey sample. Rather, the program receives a total dollar amount of funding that then must be allocated to cover the various costs of program operation. Similar to other surveys, however, the amount of money needed to maintain program operations can be expected to grow over time because of growth in wages, salaries and other expenses and growth in the number of certainty units.

establishments with employment of approximately 19.6 million each year, close to doubling the employment covered.

Option 2: Survey all private sector establishments in large firms every year. Adoption of the preceding option for redesign of the OES sample would reduce the variance of annual OES estimates. Under this plan, however, only large establishments could be followed over time and, in most cases, data collected in a particular year still would cover only a portion of a firm's establishments, making it difficult to use the data for firm-level analysis. If longitudinal analysis at the firm level is a priority, surveying all of the establishments in *firms* with more than a threshold level of employment every year might be an appealing strategy. To the extent that large firms were willing to provide electronic data files containing information on the job classifications of all of their employees, collecting data in this way also could yield significant economies of scale. Indeed, without such economies, this data collection strategy would be prohibitively expensive.

Turning again to the May 2006 universe listing, using EIN as the firm identifier, we found 322,525 establishments that belonged to the 8,295 firms with 1,000 to 4,999 employees, 148,211 establishments that belonged to the 964 firms with 5,000 to 9,999 employees and 429,140 establishments that belonged to the 677 firms with 10,000 or more employees. Using the per-establishment cost figures cited above, even restricting attention to establishments belonging to firms with 10,000 or more employees, the costs of data collection would be projected to rise by $18.7 million dollars per year, an amount that comes close to equaling the program's entire current data collection budget.

It seems likely, however, that economies of scale could be realized if data were collected for entire companies rather than separately for each establishment in the sample. So-called

central office collection (COC) is already in place for some companies who prefer to submit their data in this way. Estimates of the staffing required to complete the current COC workload compiled by BLS national and regional office staff (NCS-OES Data Collection and Processing Cost Team 2008) suggest that a reasonable estimate might be that it costs approximately $6,000 per firm to process these submissions. If COC collection were mandated for all large firms, the per-firm cost likely would be higher. One reason is that firms for which this is not necessarily the preferred reporting method would have to be convinced to report in this way. The need to process all of the firms' establishments, rather than only selected establishments as at present, also would raise per-firm costs. To the extent that firms use common job classification systems across all of their establishments, the effort to assign SOC codes to job titles might not vary a great deal with the number of included establishments, but there could be other complications. For example, past experience suggests that there can be problems with matching establishments listed in the firm records to establishments on the BLS business list and with assigning appropriate industry and geographic identifiers.

For the purpose of producing a rough data collection budget estimate, suppose that it would cost $8,000 per COC firm to process electronically submitted data. If the COC universe were restricted to firms with 10,000 or more employees, this per-firm cost figure implies that total data collection costs would rise by $3.5 million per year, rather than the $18.7 million implied by the establishment-by-establishment collection cost model. If we instead assumed a higher figure of $10,000 per firm, the projected budget increment for added data collection work would be $4.8 million.

Compared to the previous sample redesign option, this option likely would do less to reduce the variance of annual OES estimates. As of May 2006, 14.0 million people were

employed by establishments that would have been added under this strategy, compared to the 19.6 million employed by the establishments added under the previous strategy. In addition, to the extent that establishments of a given firm tend to be similar with respect to their staffing patterns, their addition to the sample will do less to improve the precision of the aggregate estimates than would the addition of a similar number of unaffiliated establishments. A major advantage of this approach is that, because information for all of the establishments at the identified firms would be collected, the data would be well-suited for studying occupational employment trends at the firm level for the covered firms.

Option 3: Survey all private sector establishments in large firms and all other large private sector establishments every year. If a larger amount of money were available, a third option for augmenting the OES sample would be to combine the first two options. Under this option, data would be collected annually for all establishments that belong to firms with 10,000 or more employees and for all other establishments with 250 or more employees. To estimate the cost of this option, we assume that the incremental costs of data collection for establishments belonging to large firms would be $8,000 to $10,000 per firm and that collecting data for any remaining establishments with 250 or more employees would cost $100 to $200 per establishment, depending on establishment size, the same assumptions used to estimate the costs of option 1 and option 2. At $8,000 per firm for COC collections, this option would raise the cost of the OES program by $6.4 million dollars per year; if we instead assume a cost of $10,000 per firm for COC collections, the cost increment would be $7.7 million. These figures are about $1.0 million less than the sum of the incremental costs for option 1 and option 2, due to the overlap between the two covered groups, consisting of large establishments that belong to large firms.

Option 3 would reduce the variance of annual OES estimates by a larger amount than either option 1 or option 2. In total, it would add establishments with 27.8 million jobs to the sample each year. Firm-level analysis also would be supported under this sample design option, at least for the set of large firms for which comprehensive data were collected.

Other options. The options outlined above are of course not an exhaustive set of possibilities. One might, for example, want to modify option 1 by collecting data each year not only for large establishments but also for a subsample of smaller establishments. Option 2 might be modified by focusing collection efforts on firms doing business primarily in sectors that, according to some yet-to-be specified criterion, are likely to be affected by off-shoring activity. Even assuming a cost of no more than $8,000 to $10,000 per firm, reductions in the firm size threshold for data collection that were applied economy-wide would be very expensive. For example, lowering the firm size threshold to include all firms with employment of 5,000 or more would raise the cost of Option 2 to between $10.5 million and $13.8 million. A lower firm size threshold might be more feasible, however, if its application were restricted to certain sectors. Other variants of the sampling options we have outlined also could be devised, but all would have in common the designation of some significant sample of establishments for annual collection.

One question that needs to be asked about the survey redesign options we have suggested is whether businesses would in fact be willing to respond to the survey every year. The results of a test conducted in 2000 are at least somewhat encouraging. In this test, staff in 25 states attempted to collect data from all of the certainty units in their states and 12 states provided data on their collection experience that could be analyzed. These data show roughly comparable response rates for the certainty units originally scheduled to participate in the OES in 2000 and

the added units originally scheduled to report only in another year of the three-year collection cycle. On the negative side, collecting data from both groups was difficult and the combined response rate was lower than would be hoped. An important factor in the response rate obtained appears to have been the amount of staff time available for followup with the surveyed units (Bureau of Labor Statistics, 2001b).

Another question is whether large companies would be willing to submit electronic records for all of their establishments centrally, as is assumed in our cost estimates for option 2 and option 3. Some large companies already do this and, depending on how their records are kept, others might be willing to do so. But if a significant number of large companies cannot or will not agree to central office collection, response rates for options that envision such collection could be adversely affected and the data collection costs for these options could be substantially higher than our rough figures suggest.

Other Changes in Survey Management Practices and Philosophy

In addition to the changes in sample design discussed above, converting the OES from a survey designed only to produce detailed cross-sectional estimates to a survey that also produced usable time series would require a number of other changes to the way in which the survey was managed. These would include changes in editing and imputation procedures, plans for dealing with future changes in the industry and occupation classification structures, and overarching changes in the philosophy governing other management decisions. Making these changes almost certainly would require additional staffing and it should be emphasized that the costs of this added staffing are not reflected in the data collection cost estimates cited above.

Consider first the process for editing the survey responses that are submitted. Under the current program structure, responses submitted by establishments are reviewed in isolation. Edit checks examine whether the combination of occupations reported seems sensible according to a specified set of criteria, but there are no edit checks that examine whether the information reported by an establishment in the current year is consistent with that reported by the same establishment in previous years. Checking the consistency of establishment reports across years would be more important in a program designed to produce time series data. Any new editing procedure would need to allow for the possibility that establishments' staffing patterns might in fact change significantly from one year to the next. One would hope, however, that editing checks could identify large changes in total reported employment that might indicate a change in the unit for which the respondent was reporting or shifts of large blocks of employment from one occupational classification to another that might indicate reporting inconsistencies.

Similarly, under the current survey design, data for establishments that do not respond are imputed using information for other establishments that are deemed to be similar. In a program designed to produce time series data, however, given idiosyncratic variation in staffing patterns across establishments, it would be desirable to develop an imputation methodology that relied more heavily on data reported by the missing establishments in previous time periods.

A recurring issue for a program designed to produce time series estimates of employment by occupation would be dealing with future changes to the SOC and NAICS. The next scheduled NAICS revision is set for 2017 and the next scheduled SOC revision for 2018. From that point forward, current plans call for NAICS revisions approximately once every five years and SOC revisions approximately once every ten years. It seems likely that SOC revisions will pose the most serious challenges for the OES program. Having a set of dual-coded records

containing both the old and the new SOC code could allow the OES staff to reconstruct historical occupational employment series on the new SOC basis, but dual coding is expensive and this would need to be built in to the OES budget plans.

Less tangible, OES management and staff would need to reorient themselves towards a new set of survey objectives, which in turn would drive subsequent decisions. There is always a tension between making changes designed to improve survey estimates and preserving the continuity of historical series. In the current OES program, the program's stated objectives have dictated that, when data improvements are possible, they should be introduced, even when that makes the data less comparable over time. In a program that had as one of its stated objectives the production of annual time series, this balance would need to be set differently.

V. Conclusion

Because the annual sample for the OES survey is large, it is reasonable to think about using the historical OES data to produce annual time series estimates. In practice, however, while useful for certain purposes, the annual times series that the historical data will support have significant limitations for studying the effects of off-shoring. Reflecting the fact that certainty units are surveyed only once every three years, annual estimates for detailed occupations can be expected to have high variance. Breaks in occupation and industry classification systems and other changes in survey practices further complicate any trend analysis based on the OES data. Further, the existing data are not well suited to support longitudinal analysis.

Any redesign of the OES program to produce reliable annual estimates and/or support longitudinal analysis should address each of these factors. We have suggested several options for redesigning the survey sample that would involve the collection of data from all large private

sector establishments each year and/or the collection of data from all of the establishments at selected large firms each year. Ballpark cost estimates for the sample expansions associated with these options range from $3.5 to $7.7 million per year, though we should emphasize both that these estimates are very rough and that they cover only the direct costs of data collection. If the OES survey were to be redesigned along the lines we have suggested, funding also would be required to support new data editing procedures, dual coding of survey records at the time of future changes to the SOC and other survey management activities, though data collection likely would account for the largest share of the total new funding that would be needed.

A significant complication we have not addressed is whether and how annual published estimates would be reconciled with more detailed cross-sectional estimates produced using data from multiple waves of data collection. The existing OES customer base cares a great deal about geographically disaggregated estimates and our guess is that, even with an expanded sample, the needs of this customer base could not be satisfied by purely annual data. This is, however, at least partly an empirical matter that remains to be addressed. Even with the full sample of approximately 1,200,000 establishments currently collected over three years, the BLS does not publish occupational data for industries disaggregated by geographic area and, if we have understood correctly what is being recommended, it seems unrealistic to call for "...the BLS to make the changes to the OES methodology necessary to create time series data on all 820 occupations in the SOC by industry and geographic areas" (Sturgeon 2006). Further work would be required to determine the level of detail in occupational employment estimates, whether by industry or by geographic area but almost certainly not by both simultaneously, that could be supported by different sample redesign options.

Finally, in thinking about a possible OES redesign, it will be important to consider carefully the value of data to support firm-level longitudinal analysis as compared to the value of improved annual time series. Without a doubt, a longitudinal occupational employment database could support interesting research, including useful research on the effects of offshoring. Further, there may be significant economies associated with the collection of data from all establishments of large firms. On the other hand, adding establishments from a small number of firms can be expected to do less to reduce the variability of annual time series estimates than adding a similar number of establishments representing a larger number of firms. In addition, unless new modalities for researcher access to confidential microdata are developed, the number of researchers who would in practice end up working with firm-level data from the OES seems likely to be limited. Those charged with making a decision about the survey's future will need to be clear about the relative importance of different survey objectives as they choose among possible redesign options.

References

Abraham, Katharine G., and James R. Spletzer. Forthcoming. "Are the New Jobs Good Jobs?" In Katharine G. Abraham, James R. Spletzer and Michael Harper, eds., *Labor in the New Economy*, Chicago: University of Chicago Press.

Bureau of Labor Statistics. 2001a. *Occupational Employment and Wages, 1999, Bulletin 2545.* Washington, DC: Government Printing Office. September. Accessed online at http://www.bls.gov/oes/1999/oes_pub99.htm on October 10, 2009.

Bureau of Labor Statistics. 2001b. *Analysis of Certainty Unit Masterfiles, 2000 OES Survey*. Certainty Unit Subgroup. OES Policy Council Meeting, Boston, Massachusetts. October 10-11.

Bureau of Labor Statistics. 2008. *Report on Current Costs, Overhead and Response Measures.* NCS-OES Data Collection and Processing Cost Team. Unpublished report. January 8.

Jensen, J. Bradford and Lori G. Kletzer. Forthcoming. "Measuring Tradeable Services and the Task Content of Offshorable Services Jobs." In Katharine G. Abraham, James R. Spletzer and Michael Harper, eds., *Labor in the New Economy*, Chicago: University of Chicago Press.

Mikkelson, Gordon, Teresa L. Morisi and George Stamas. 2000. "Implementing the NAICS for Business Surveys at BLS." Paper presented at the Second International Conference on Establishment Surveys, Buffalo, New York, June 17-21.

National Academy of Public Administration. 2006. *Off-Shoring: How Big Is It?* Report of a Panel to the U.S. Congress and the Bureau of Economic Analysis. October.

Office of Management and Budget. 2000. Standard Occupational Classification Manual, White Plains, MD: Graphic Systems, Inc.

Sturgeon, Timothy J. 2006. *Why We Can't Measure the Economic Effects of Services Offshoring: The Data Gaps and How to Fill Them*. Final Report of the Services Offshoring Working Group, Industrial Performance Center, Massachusetts Institute of Technology. September 10.

Figure 1: OES Employment of Ten Occupations Identified by Jensen and Kletzer (forthcoming) As Most Off-Shorable, 1999-2004