B. Collection of Information Employing Statistical Methods

1. Description of Universe and Sample

a) Universe

Entity	Universe	Sample
Agricultural Region	12	12
Farm Labor Areas	498	90
Farms w/ hired or contract labor	338,373	564
Hired Crop Workers (estimated)	1,400,000	1,500

The universe for the study is the population of field workers active in crop agriculture in the continental United States (U.S.). The National Agricultural Workers Survey (NAWS) will use multi-stage sampling relying on probabilities proportional to size to interview approximately 1,500 randomly selected crop workers in fiscal year (FY) 2011.

b) Response Rate

The sampling design, described below, involves obtaining a random selection of employers. In FY 2009, 66 percent of the randomly selected employers (or their surrogates) who employed workers the day they were contacted by interviewers agreed to cooperate in the survey and interviews were conducted on 59 percent of the eligible establishments. As there are no universe lists of workers, the sampling frame of workers is constructed after contact with the employer.

Once interviewers have a worker frame, a random sample of workers is chosen. The interviewers, who work in pairs, approach workers directly to set up interview appointments in their homes or other agreed-upon locations. In 2008-2009, 92 percent of the approached workers agree to be interviewed.

2. Statistical Methodology

Overview

The goal of the NAWS sampling methodology is to select a nationally representative, random sample of farm workers. The NAWS uses stratified multi-stage sampling to account for seasonal and regional fluctuations in the level of farm employment. The stratification includes three interviewing cycles per year and 12 geographic regions, resulting in 36 time-by-space strata. For each interviewing cycle, NAWS staff draws a random sample of locations within all 12 regions from the universe of 498 Farm Labor Areas (FLAs). FLAs are single or multi-county sampling units which form the primary sampling units (PSUs). Growers within PSUs are the secondary level and workers within growers are the tertiary level of sampling units. The number of interviews allocated to each region is based on regional farm worker employment data (number of agricultural hired and contract workers) from USDA's Farm Labor Survey. Similarly, the number of

interviews allocated to each FLA is proportional to the number of farm workers employed at that time of the year. The FLA size measure (farm labor) is obtained by multiplying a seasonality estimate, derived primarily from the BLS QCEW, by local farm labor expenditure data, from USDA's Census of Agriculture. Interview allocation is thus proportional to stratum size.

In each FLA, a simple random sample of agricultural employers is drawn from a universe list compiled mainly from public agency records. NAWS interviewers then contact the sampled growers or farm labor contractors, arrange access to the work site, and draw a random sample of workers at the work site. Thus, the sample includes only farm workers actively employed in crop agriculture at the time of the interview.

Stratification

Interviewing cycles

To account for the seasonality of the industry, interviews are conducted three times a year in cycles lasting ten to twelve weeks. The cycles start in February, June and October. The number of interviews conducted in each cycle is proportional to the number of agricultural field workers hired at that time of the year. The U.S. Department of Agriculture's (USDA) National Agricultural Statistics Service (NASS) provides the Employment and Training Administration (ETA) the agricultural employment figures, which come from USDA's Farm Labor Survey. In fiscal year 2011, the NAWS will visit a total of 90 interviewing locations. The locations will be similarly apportioned among the cycles using NASS data.

Regions

Regional stratification entails defining 12 distinct agricultural regions that are based on USDA's 17 agricultural regions. At the start of the survey in 1988, the 17 regions were collapsed into 12 by combining those regions that were most similar e.g., Mountain I and Mountain II, based on statistical analysis of cropping patterns. In each cycle, all 12 agricultural regions are included in the sample. The number of interviews per region is proportional to the size of the seasonal farm labor force in that region at that time of the year, as determined by the NASS using information obtained from the Farm Labor Survey.

Sampling within Strata

Farm Labor Areas

Each region is composed of several multi-county sampling units called Farm Labor Areas (FLAs). Originally, the NAWS used USDA Crop Reporting Districts, but experience showed that these units were not homogeneous with respect to farm labor. As a result, using Census of Agriculture data and ETA mappings of seasonal farm labor concentrations, the NAWS staff identified aggregates of counties that had similar farm

labor usage patterns and were roughly similar in size. The resulting FLAs also account for varying county size across the United States. For example, in the East, a Farm Labor Area may include several counties; in the West, a Farm Labor Area may be composed of a single agriculture-intensive county. FLA size is more homogeneous within region than it is across regions. The approximately 3,000 U.S. counties are reduced to 498 farm labor areas. The current list of 90 FLAs contains 425 counties.

While in the past, each cycle's FLAs were drawn from a standing roster of FLA's, for 2011, each cycle's FLAs will be drawn from the universe of 498 FLAs. For each cycle, within each region, a sample of FLAs is drawn using probabilities proportional to size. The size measure used is an estimate of the amount of farm labor in the FLA during a particular cycle. In this case, the measure is based on the hired and contract labor expenses from the most recent Census of Agriculture (CoA), available at the time of drawing the sample. The CoA labor expenses are adjusted using seasonality estimates which identify the percentage of labor expenses that fall into each of the NAWS cycles: fall, spring and summer.

The seasonality estimates are constructed from Quarterly Census of Employment and Wages (QCEW) data. The estimates are made by aggregating the reported monthly employment for each month included in the corresponding NAWS cycle e.g., June, July, August, and September for the summer cycle. The percentage of employment corresponding to each cycle becomes that FLA's seasonality estimate.

Within each region, FLAs are selected systematically from a randomly sorted list of FLAs using probabilities proportional to size. To begin, a selection interval is created by dividing the size measure total for the FLA by the number of FLAs to be selected. The list of FLAs is randomly sorted and a running sum is created of the size measure for each FLA. Thus each FLA covers an interval within the running sum with a range corresponding to the FLA's size measure. (Similar to county example in chart below) FLAs are selected systematically in the following manner. Beginning at a random starting point within the selection interval and proceeding down the randomly sorted list, counties are selected at evenly spaced distances within the running sum that correspond to the selection interval.

Counties

In selecting counties, an iterative sampling procedure is used to ensure an adequate number of counties is selected for each region. In most cases, interviews are completed in the first county and no additional counties are needed. However, because there is tremendous uncertainty about the number of workers in a county, additional counties are occasionally needed to complete the county allocation. Counties are selected one at a time without replacement using probabilities proportional to the size of the farm labor expenditures in that county at that time.

The process of selecting a county within a FLA begins with a randomly sorted list of counties within the FLA. A cumulative sum using the size of the seasonal hired and

contract labor expenditures is constructed for this list similar to the process described above for FLA selection. When selecting a county, the selection number is the product of a random number selected from the uniform distribution multiplied by the cumulative total of the seasonal hired and contract payroll. The county that includes that number in its selection of the cumulative sum is selected.

Example showing selection algorithm for counties within FLAs

County	Seasonal labor expenditures	Cumulative Sum	Selected
A	100,000	100,000	
B	300,000	400,000	
C	800,000	1,200,000	
D	450,000	1,650,000	<=
E	600,000	2,250,000	

Selection

Random number	0.657
Random number * cumulative sum	1,477,205
Selected FLA	D

As mentioned above, often only one county is needed. Should additional counties be needed in a multi-county FLA, additional counties are pulled one at a time without replacement using the method described above. To be prepared, this is done in advance of interviewing but only implemented if needed.

Each county is marked with and ordered by its selection number (e.g., 1,2,3…) Interviews begin in the first selected county and, as a county's work force is depleted, interviewing moves to the second randomly selected county on the list and so forth, until all the allocated interviews in that FLA have been completed. In FLAs where farm work is sparse, interviewers may need to travel to several counties to encounter sufficient workers to complete the FLA's allocation.

Employers

Within each selected county, employers are selected at random from a list of agricultural employers. The list is compiled from administrative lists of employers in crop agriculture. An important component of the list is employer names in selected North American Industrial Classification Codes that the Bureau of Labor Statistics (BLS) provides directly to the contractor per the terms of an interagency agreement between the ETA and the BLS. Because of uncertainty about the conditions of seasonal farm labor in each location, the number of employers to be interviewed is not known in advance.

Workers

Workers are selected using probabilities proportional to the square root of employer size. The process is as follows. Once the randomly selected employer is located, the interviewer determines if he/she is familiar with his/her work force. If not, the interviewer seeks the name of the packinghouse manager, personnel manager, farm labor contractor, or crew leader who can help construct a sampling frame of the workers in the operation. Interviewers document the number of workers employed on the day of worker selection in order to construct worker selection probabilities.

Since the size of the farm labor force varies considerably among agricultural employers, interviews are allocated using the square root of employer size. A chart similar to that below will assist interviewers in determining the number of interviews per farm.

Number of workers	Number of Interviews
1 - 2	1
3 - 6	2
7 -12	3
13-20	4
21-30	5
31-42	6
43-56	7
57-72	8
73-90	9
91-110	10
111-132	11
133 or more	12

On large farms, workers are usually organized into crews consisting of several workers and a supervisor. Crew size can range from a handful to more than 100 but crews of 30 or less are most typical. When the number of crews is large, randomly selecting workers from each crew would be infeasible for interviewers as well as an imposition on the farm employer. For this reason, on farms with more than one crew, the number of crews to be randomly selected will be determined using probabilities proportional to the square root of size. A chart similar to the one below will be provided to interviewers. Interviewers keep track of the number of crews selected and the size of the crew. Interviewers pro-rate the number of interviews across the selected crews based on the size of the crews. Their instructions for this are included in Appendix B.

Number of crews	Number to select randomly
1 to 2	1
3 to 6	2
7 or more	3

Within each crew, the interviewers follow specific sampling instructions that were designed by a sampling statistician to ensure selection of a random sample of workers at each selected employer. Specifically, if n is the number of interviews for that crew and N is the total number of workers within the crew then interviewers place n marked tags and N-n unmarked tags in a pouch and shuffle them. Workers then draw a tag and those with marked tags are included in the sample. If multiple crews are used, the procedure is completed with each selected crew. These instructions are included in Appendix B

Weighting

The NAWS uses a variety of weighting factors to construct weights for calculating unbiased population estimates:

- Sampling weights are used to calculate unbiased population estimates by assigning each sample member a weight corresponding to the inverse of its probability of selection.
- Non-response factors are used to correct sampling weights for deviations from the sampling plan, such as discrepancies in the number of interviews planned and collected in specific locations.
- Post-sampling adjustment factors are used to adjust the weights given to each interview in order to compute unbiased population estimates from the sample data.

As explained below, non-response weights are calculated simultaneously with regional post-sampling adjustment weights.

Sampling weights

Each worker in the sample has a known probability of selection. Information collected at each stage of sampling is used to construct the sampling weights.

Sampling weights are calculated as the inverse of the probability of being selected:

$$WT_k = 1/\text{prob},$$

where prob = workprob*growprob*counprob*flaprob,

with $$workprob = \frac{\text{number of workers interviewed at the farm location}}{\text{total number of workers at that location}},$$

$$\text{growprob} = \frac{\text{number of growers interviewed in the county}}{\text{total number of qualified growers in that county}},$$

Calculating counprob, the county within FLA weight is more complicated. For example, if one of the sampled counties is larger than another, then its probability of selection should be higher than that of the other. If several counties are selected from a particular region, then the selection probability for a particular county is (1) its probability of selection on the first draw, plus (2) the probability of its selection on the second draw, plus (3) the probability of its selection on the third draw, etc.

For the standard method of sampling several items with probabilities proportional to size, without replacement, closed-form formulas for the exact inclusion probabilities do not exist. However, these probabilities can be calculated exactly using multiple summations. This procedure can be implemented in SAS within PROC IML.

Suppose that the population at a particular sampling stage consists of N objects with sizes $s_1, s_2, \ldots, s_N$, having total size $S = \sum_{j=1}^{N} s_j$. Let π_j^i be the probability that the jth item is selected on the ith draw. Then for $j = 1, 2, \ldots, N$,

(1) $$\pi_j^1 = \frac{s_j}{S},$$

(2) $$\pi_j^2 = \sum_{\substack{k=1 \\ k \neq j}}^{N} \frac{s_k}{S} \frac{s_j}{S - s_k},$$

(3) $$\pi_j^3 = \sum_{\substack{k=1 \\ k \neq j}}^{N} \sum_{\substack{l=1 \\ l \neq j \\ l \neq k}}^{N} \frac{s_k}{S} \frac{s_l}{S - s_k} \frac{s_j}{S - s_k - s_l},$$

(4) $$\pi_j^4 = \sum_{\substack{k=1 \\ k \neq j}}^{N} \sum_{\substack{l=1 \\ l \neq j \\ l \neq k}}^{N} \sum_{\substack{m=1 \\ m \neq j \\ m \neq k \\ m \neq l}}^{N} \frac{s_k}{S} \frac{s_l}{S - s_k} \frac{s_m}{S - s_k - s_l} \frac{s_j}{S - s_k - s_l - s_m},$$ and so forth.

These i^{th}-draw probabilities each have the property that $\sum_{j=1}^{N} \pi_j^i = 1$. Finally, the probability that the jth item is included in a sample of size n is $\pi_j = \sum_{i=1}^{n} \pi_j^i$. These inclusion probabilities have the property that $\sum_{j=1}^{N} \pi_j = n$.

Both the FLA and county selection probabilities can be calculated exactly using these formulas.

The formula for FLAprob which is the probability that the FLA was selected within the region derives from the systematic probability proportional to size selection process used.

<u>Consider</u>

N is the number of FLAs in the region

s_1 through s_N are the sizes of the FLAs.

S is the sum of the FLA sizes, so $S = \sum_{i=1}^{N} s_i$.

n is the number of the FLAs to be selected with probabilities proportional to size.

In selecting the FLAs, they were listed in a random order. A column of cumulative FLA sizes was constructed. That is, the cumulative size at the jth FLA will be $S_j = \sum_{i=1}^{j} s_i$.

A random starting point, k_0, was chosen between 1 and $k = \frac{S}{n}$. The integers $k_0, k_0 + k, k_0 + 2k, \ldots, k_0 + (n-1)k$ can be listed. The jth FLA will be selected if one of these integers falls between $S_{j-1} + 1$ and S_j (where S_0 is interpreted to be 0).

Without loss of generality, consider the first FLA on the randomized list. It will be selected if k_0 lies between 1 and s_1. Thus, its probability of selection is $\pi_1 = \frac{s_1}{k} = n\frac{s_1}{S}$.

In general, the probability that the jth FLA is selected is $\pi_j = n\frac{s_j}{S}$.

Non-Response Weighting

Non-response corrections adjust for deviations from the sampling design. If, for example, ten interviews should have been collected at a farm but only two interviews were collected, those two interviews could be given five times the weight they would have received otherwise. Thus, each interview's weight needs to be adjusted to represent a certain value in terms of size. Instead of making this adjustment at the farm level, it could be made at any higher level in the sampling plan. For the NAWS this means at the list-within-county, county, list-within-FLA, FLA, list-within-region, region, or national level.

By raising the level at which adjustments are made, overall size information is, generally, more reliable. This is due to the statistical effect of averaging, greater year-to-year stability over larger geographic areas, and the absence or suppression of data due to confidentiality considerations. On the other hand, lower-level adjustments are more sensitive, if the information used for making the adjustments is reasonably accurate.

For two reasons, the NAWS non-response adjustments are made at the region level. First, the region is the lowest level with enough interview coverage to calculate weights for the size adjustment. All of the 12 NAWS regions are visited in every cycle. If, for some reason, there are too few interviews in a region, the region can be combined with adjacent regions for weighting purposes.

Second, the NAWS uses measures of size provided by the USDA Farm Labor Survey, which are reported by quarter and region. The USDA is the only source of quarterly statistics on levels of farm worker employment. The Census of Agriculture, for instance, reports data annually rather than quarterly and the statistics are published every five years. Thus, by using USDA Farm Labor Survey figures to make the size adjustment, the NAWS can adjust the weights by season and region and construct unbiased population estimates. Non-response adjustments for size, therefore, are made at the region-within-

cycle level to create corrected region weights.

Post-sampling weights

Post-sampling weights are used in the NAWS to adjust the relative value of each interview in order for national estimates to be obtained from the sample. There are five post-sampling weights. Two of the weights adjust for unequal probabilities of selection that can only be determined after the interviews are conducted. These include the unequal probabilities of finding part-time versus full-time workers (day weight) and the unequal probabilities of finding seasonal versus year- round workers (seasonal weight). The next three weights (region, cycle, and year) adjust for the relative importance of a region's data, a sampling cycle, and a sampling year. The measures of size used are obtained from the USDA. The region weight, as discussed below, is calculated simultaneously with non-response weighting. The cycle weight and year weight serve slightly different roles in estimation. They allow different cycles and sampling years to be combined for statistical analysis. These weights are also based on USDA measures of size.

It should be noted that the NAWS sampling plan is based on USDA NASS data collected in the year before the interviews. For example, fiscal year 2009 data is used to plan the NAWS 2010 sample. The weights use NASS data from the year during which the interviews were conducted. This corrects for any discrepancies in allocations due to projecting farm worker distributions based on past year data.

Adjustment for days worked per week: the day weight

The day weight adjusts for the probability of finding part-time versus full-time farm workers. A part-time worker, who works only two or three days per week, has a lower likelihood of being encountered by the interviewing staff than a worker employed six days per week. Therefore, respondents are weighted inversely proportional to the length of their workweek.

A conservative adjustment for the number of days worked is appropriate to avoid excessively large sampling weights. Field reports indicate that relatively few workers are contacted on Sundays, and a review of the interviews indicated that virtually no workers reported Sunday hours without Saturday hours. Accordingly, workers reporting at least six workdays per week nearly always have a full chance of selection. Thus, any workers reporting at least six days of work per week are treated as having a full chance of selection; adjustments are made only for those workers with less than six days of work per week. The day weight (DWTS) is computed as: DWTS = 6 / (length of the workweek) where "length of the workweek" is the number of days per week the respondent reports working at the time of the interview for the current farm task (if two tasks are reported, the one with more days per week is used). Seven-day workweeks are truncated to six, as explained previously. For the few workers not reporting the number of days, DWTS is assigned a default value of 1.

Adjustment for seasons worked per year: the season weight

The calculation of worker-based weights is complicated by the fact that workers could, in

general, be sampled several times a year. Furthermore, the USDA information does not provide a figure for the number of farm workers for the year. The USDA reports the number of farm workers working each season, so the same worker could be reported in multiple seasons. Because of this repetition of workers across seasons, it is impossible to derive the total number of persons working in agriculture during the year.

As employment information is not available for every worker for each quarter of the year, the only way to avoid double-counting of farm workers is to use the 12-month retrospective work history collected in the NAWS. Specifically, predicting future-period employment is achieved by imposing the assumption that workers who report having worked in a previous season would work in the next corresponding season. For example, a worker sampled in spring 2008 who reported working the previous summer (2007) is assumed to work in the following summer (2008). For some purposes, including the calculation of year-to-year work history changes, this assumption cannot not be used. For purposes such as obtaining demographic descriptions of the worker population, however, this assumption provides satisfactory estimates.

Further, it is assumed that a worker has an equal likelihood of being sampled in each season worked. This assumption is dependent on a balance between the amount of farm work done by the worker in each season and the number of interviews obtained in that region for the season. Recall that the NAWS interview allocation is proportional to county-level seasonal agricultural payroll. Thus, the probability of sampling is related to the amount of work performed by individual workers. With these simplifying assumptions, it is possible to calculate a seasonal weight that is simply the inverse of the number of seasons the interviewee did farm work during the previous year.

For the purposes of the NAWS, there are only three seasons per year. An interviewee always performed farm work during the trimester he\she was sampled. From the NAWS interview, it can be determined during which of the two previous trimesters the respondent also did farm work. If the interviewee only worked during the current trimester, the seasonal weight is 1/1 or 1.00. If the interviewee worked during the current trimester and only one of the two prior trimesters, the seasonal weight is 1/2 or 0.50. Finally, if the interviewee worked during the current and both of the prior trimesters, the seasonal weight is 1/3 or 0.33.

This season weight is similar to the day weight in the sense that respondents who spend more time (seasons) working in agriculture have a greater chance of being sampled. Therefore, the weighting has to be inversely proportional to the number of seasons worked in order to account for the unequal sampling probability.

The region weight

The region weight adjusts the relative weight of a region's data in relation to the number of interviews collected in that region. If the number of interviews collected was smaller than the regional allocation in the sampling plan, an adjustment weight greater than one is assigned to each interview in the region, and vice versa. These adjustments ensure that the population estimates are unbiased.

The region weight is based on USDA measures of regional farm employment activity. This is the best source of information available about farm workers. The USDA figures are reported by region and quarter, which allows the weight to be sensitive to seasonal fluctuations.

Correspondence between USDA data and the NAWS sampling cycles

The calculation of the region weight relies on two pieces of information: the USDA regional measures of size and the number of interviews completed in each region. The first step in the process of calculating the region weight is to apportion the USDA quarterly size figures among the three NAWS sampling cycles.

The USDA figures are reported quarterly. NAWS sampling years, however, cover non-overlapping 12-month periods (from September to August), which are divided into three cycles. Accordingly, it is necessary to adjust the USDA figures to fit the NAWS sampling frame by apportioning the four quarters into three cycles.

For example, the number of farm workers in the fall cycle for a region is assumed to be the total number of workers for that region in USDA quarter 1 of the current fiscal year (FY_c) plus one-third the number of workers for that region in USDA quarter 2 of the next fiscal year (FY_p). The formula for the winter, spring and summer cycles is constructed similarly.

Determining the NAWS region grouping according to interview coverage

The region weight (within cycle) is calculated as follows for each region j $(1,\ldots,n_i)$ in cycle i:

$$PWTR_{ij} = \frac{USDA_{ij} \big/ X_{ij}}{\sum_{j=1}^{n_i} (USDA_{ij} / X_{ij})} \times \frac{\sum_{j=1}^{n_i} X_{ij}}{DWTS_{ij}},$$

where $USDA_{ij}$ is the USDA estimate for region j in cycle i, X_{ij} is the sum of the sampling weightss for region j in cycle i, $DWTS_{ij}$ is the sum of farm worker day weights for region j in cycle i. Also, $1 \le DWTS_{ijk} \le 6$ (where k refers to a farm worker), so that

$DWTS_{ij} = X_{ij}$ if all farm workers in region j in cycle i are working full time and

$DWTS_{ij} = 6 \times X_{ij}$ if all farm workers are working 1 day only a week in region j in cycle i.

Combining different sampling cycles: the cycle weight

The NAWS combines data from the different sampling cycles (seasons) within the same sampling year in order to generate more observations for statistical analysis. In order to combine cycles it is necessary to adjust for the number of farm work days represented in each cycle in relation to the number of interviews collected in the cycle. For instance, suppose the NAWS did not do proportional sampling as explained above but rather interviewed the same number of people in all three cycles in the 2007 fiscal year. If the USDA reported more workers for the fall and spring/summer cycles, as compared to the

winter cycle, then the interviews in the fall and spring/summer would be worth relatively more in terms of size than the interviews conducted in the winter cycle. Accordingly, the interviews in the winter would have to be down-weighted in relation to the interviews in the other seasons (cycles) before the cycles could be combined.

The cycle weight is calculated similarly to the region weight, but at the cycle- rather than region- level. The sum of the USDA size for a cycle is divided by the number of interviews in that cycle. The cycle weight (or region weight within year) is calculated as follows for each region j $1,\ldots,n_i$, cycle i in year Y:

$$PWTCR_{ij} = \frac{\dfrac{USDA_{ij}}{X_{ij}} * K_{ij}}{\sum_{i,j\in Y}\left(\dfrac{USDA_{ij}}{X_{ij}} * K_{ij}\right)} * \frac{\sum_{i,j\in Y} X_{ij}}{SSEADWTS_{ij}}$$

where $\text{SSEADWTS}_{\text{ij}} = \sum_{k\in(i,j)} DWTS_k * SEASWTS_k$

and

$$K_{ij} = \frac{\sum_{k\in(i,j)} DWTS_k * SEASWTS_k}{\sum_{k\in(i,j} DWTS_k}$$

$0.33 \le \text{SEASWTS}_{\text{k}} \le 1$ (k refers to a farm worker) and $\text{SEASWTS}_{\text{k}} = 1$ if the farm worker worked only one cycle during the year, so that if all farm workers for region j in cycle i worked full time and only one cycle in the corresponding year $K_{ij} = 1$ and $\text{SSEADWTS}_{\text{ij}} = X_{\text{ij}}$.

Combining different sampling years: the year weight

The year weight allows different sampling years to be combined for statistical analysis. It follows the same rationale as the cycle weight, but at the sampling-year level. If the same number of interviews are collected in each sampling year, those interviews taking place in years with more farm work activity are weighted more heavily in the combined sample.

Sampling years cannot be combined if the interviews are not comparable in terms of agricultural representation. In an extreme case, suppose that the NAWS budget tripled one of the sampling years, consequently tripling the number of interviews. If the two sampling years were joined without adjustment, the larger sampling year would have an unduly large effect on the results.

To avoid this, the year weight is calculated as a ratio of the total number of farm workers in a sampling year to the number of interviews in that sampling year. The year weight (or region weight related to all years of interviews) is calculated as follow for each region j (1..nij), cycle i (the sum over i,j means all farm workers, all cycles all years):

$$PWTYCR_{ij} = \frac{\frac{USDA_{ij}}{X_{ij}} * K_{ij}}{\sum_{i,j}\left(\frac{USDA_{ij}}{X_{ij}} * K_{ij}\right)} * \frac{\sum_{i,j} X_{ij}}{SSEADWTS_{ij}}$$

with the same notations than for the preceding weights.

Obtaining the final weights

Once the individual weight components are calculated, final composite weights are calculated as the product of the day weight, the season weight, and region weight and the sampling weight. The cycle and year are also factored into the composite weights when multiple cycles or sampling years are used. The composite weights are adjusted so the sum of the weights is equal to the total number of interviews at the next level of aggregation. These adjusted composite weights based on farm workers are then used for calculating the estimated proportion of workers with various attributes.

The individual observation weights are obtained at the farm worker level:

$$PWTRD_k = PWTR_{ij,k\in(ij)} * DWTS_k * WT_k$$

This is the weight within cycle; it includes an adjustment for the length of the workweek but no seasonal adjustment.

$$PWTCRD_k = PWTCR_{ij,k\in(ij)} * DWTS_k * SEASWTS_k * WT_k$$

This is the weight within a year; it includes both the length of the workweek and seasonal adjustment. This weight may be used for the analysis of one particular year of interview.

$$PWTYCRD_k = PWTYCR_{ij,k\in(ij)} * DWTS_k * SEASWTS_k * WT_k$$

The composite weight (PWTYCRD) is used for almost all NAWS analysis. This weight allows merging several years of analysis together. It is included in the public access dataset.

3. Statistical Reliability

a) Response

Employer response

To maximize grower response, the contractor sends an advance letter to growers and provides them a brochure explaining the survey. The letter is signed by the survey director and includes the names of the interviewers and their contact information. For further information or questions, the letter and brochure direct growers to contact either the survey contractors (JBS International) at a toll free number or the Department of Labor's (DOL) Contracting Officer's Technical Representative (COTR). Grower calls are returned quickly. In addition, and before the start of every interview cycle, JBS

provides the COTR a list of scheduled interview trips. The list includes the counties and states where interviews will be conducted, the names of the interviewers who will be visiting the selected counties, and the dates the interviewers will be in the selected counties. The COTR refers to the list whenever he receives a grower call to confirm the interviewers' association with the survey.

Both DOL and the contractor make presentations on the survey and provide survey information, e.g., questionnaires, to officials and organizations that work with agricultural employers. The NAWS has received the endorsement of several grower organizations. This improves the response rate since agricultural employers sometimes call their grower organization when considering survey participation.

Intensive and frequent interviewer training is also conducted as a means to increase employer response rates. Interviewers are trained in pitching the survey in various situations and, being well versed in the history, purposes, and use of the survey, are able to easily answer any questions or address any concerns an employer might have. In addition, when explaining the purpose of the survey to employers, interviewers clearly distinguish the survey from enforcement efforts by the Department of Homeland Security, DOL and other Federal agencies, and assure growers that their information is confidential.

Worker response

The survey's methodology has been adapted to maximize response from this hard-to-survey population. Interviewers pitch workers in English or Spanish, as necessary. All interviewers are bilingual and bicultural. In addition, interviewers make sure that potential respondents know that they are not associated with any enforcement agency, e.g., Immigration and Customs Enforcement. Interviewers explain the survey to workers and obtain their informed consent.

b) Non-response

The $20 honorarium to farm workers enables the survey to achieve an estimated worker response rate of 92 percent. This high level of response greatly aids in protecting the survey estimates from non-response bias. To reduce grower non-response, interviewers are instructed to make several contact attempts at different times of the day and on different days of the week. Interviewer contact attempts are logged and the logs are monitored for compliance. When necessary, interviewers are instructed to accommodate a grower's preference for scheduling surveys and, if needed, the interviewer can request an extension of the field period.

To measure the effect of grower non-response on the survey's findings, the survey's statistician, project manger, and COTR are exploring the possibility of using the minimal information known about and/or collected from the non-cooperating growers, e.g., primary crop, county, number of workers employed, and quarterly hired farm payroll to generate proxy grower types. If it is possible to construct such proxies, then the demographic characteristics of workers employed on farms of cooperating growers of a

particular type will be analyzed to determine if there are significant differences in the key demographic and employment characteristics of workers from participating vs. proxy non-participating growers.

c) Reliability

A probability sampling methodology will be used and estimates of the sampling errors will be calculated from the survey data.

Estimation procedure

1. At the highest level of the sampling design, the region/cycle level, stratified sampling was used. Sampling is then carried out at the lower levels, independently within each stratum.

The following description is excerpted from Obenauf[1]:

The stratified sampling technique divides the entire population into relatively homogenous groups that are mutually exclusive and exhaustive. Samples are then drawn from each of these groups (strata) by simple random sampling or an alternate method. The entire sample is a compilation of these independent samples from each of the strata. In stratified sampling, an estimate of the population mean can be made for each of the strata.
Estimate of population mean:

$$\bar{y}_{st} = \frac{\sum_{k=1}^{L} N_k \bar{y}_k}{N},$$

where N_k is the population size of stratum k and L is the number of strata into which the population is divided.
If a simple random sample is taken within each stratum (recall that other schemes can be used to draw a sample from each of the strata), the following represents an unbiased estimate of the variance of $\bar{y}_{st}$:

$$Var(\bar{y}_{st}) \approx \sum_{k=1}^{L} \left(\frac{N_k}{N}\right)^2 \frac{s_k^2}{n_k}(1-f_k).$$

The standard error of the estimator is the square root of this estimated variance, or

$$\text{S.E.}(\bar{y}_{st}) = \sqrt{\sum_{k=1}^{L} \left(\frac{N_k}{N}\right)^2 \frac{s_k^2}{n_k}(1-f_k)}.$$

2. At the second stage of the sampling design, within each stratum, counties (or groups of counties) are treated as clusters.

1 Obenauf, W. (2003), "An Application of Sampling Theory to a Large Federal Survey," Portland State University Department of Mathematics and Statistics.

The following description is another excerpt from Obenauf[1].

The population is again divided into exhaustive, mutually exclusive subgroups and samples are taken according to this grouping. Once the population has been appropriately divided into clusters, one or more clusters are selected … to comprise the sample. There are several methods of estimating the population mean for a cluster sample. The method most pertinent to this study is that involving cluster sampling proportional to size (PPS).

With PPS sampling, the probability (z_j) that a cluster j is chosen on a specific draw is given by $z_j = \frac{M_j}{M}$, where M_j is the size of the jth cluster and M is the population size. An unbiased estimate of the population total is given by

$$\hat{y}_{pps} = \frac{1}{n}\sum_{j=1}^{n}\frac{y_j}{z_j} = \frac{M}{n}\sum_{j=1}^{n}\frac{y_j}{M_j} = M\bar{\bar{y}},$$

where y_j is the sample total for y in the jth cluster, n is the number of clusters in the sample and $\bar{\bar{y}}$ represents the average of the cluster means.

To estimate the population mean, this estimate must be divided by M, the population size.

The variance of the estimator of the population total is given by

$$V\left(\hat{y}_{pps}\right) = \frac{M^2}{n}\left[\sum_{i=1}^{N}\frac{M_i}{M}\left(\frac{y_i}{M_i} - \bar{\bar{y}}\right)^2\right],$$

This is estimated by $V\left(\hat{y}_{pps}\right) \approx \frac{M^2}{n}s^2_{mean}$, where s^2_{mean} is the sample variance of the $\frac{y_j}{m_j}$ values.

For an estimate of the population mean,

$$\bar{y}_{pps} = \bar{\bar{y}} = \frac{1}{n}\sum_{j=1}^{n}\bar{y}_j = \frac{1}{n}\sum_{j=1}^{n}\frac{y_j}{m_j} \text{ and } V\left(\bar{y}_{pps}\right) \approx \frac{s^2_{mean}}{n}.$$

In two-stage cluster sampling, the estimated variance of the estimator is then given by an iterative formula:

$$Var(\bar{y}_{clus}) = E_1\left[Var_2\left(\bar{y}_{clus}\right)\right] + Var_1\left[E_2\left(\bar{y}_{clus}\right)\right].$$

This iterative formula is then generalized to compute the variance of the estimators in multi-stage sampling schemes with three or more levels. Exact formulas become intractable at this point, and the various statistical software packages rely upon either re-sampling methodology or linear approximations in order to estimate the variances and standard errors of the estimators.

The following is an excerpt from the SAS documentation for PROC SURVEYMEANS[2].

> The SURVEYMEANS procedure produces estimates of survey population means and totals from sample survey data. The procedure also produces variance estimates, confidence limits, and other descriptive statistics. When computing these estimates, the procedure takes into account the sample design used to select the survey sample. The sample design can be a complex survey sample design with stratification, clustering, and unequal weighting.
> PROC SURVEYMEANS uses the Taylor expansion method to estimate sampling errors of estimators based on complex sample designs. This method obtains a linear approximation for the estimator and then uses the variance estimate for this approximation to estimate the variance of the estimate itself (Woodruff 1971, Fuller 1975)3,4.

SAS (e.g., Proc Surveymeans), allows the user to specify the details of the first two stages of a complex sampling plan. In the present case, the stratification and clustering at the first two levels are specified in Proc Surveymeans (strata region; cluster FLA). At the lower levels of the sampling scheme, the design attempts to mimic, as closely as is practical, simple random sampling. The software is not able to calculate exact standard errors, since it presumes true simple random sampling beyond the first two levels. The sampling weights will remedy any differences in selection probabilities, so that the estimators will be unbiased. The standard errors, however, are only approximate; the within-cluster variances at stages beyond the first two are assumed to be negligible.

In the "Surveymeans" procedure, the STRATA, CLUSTER, and WEIGHT statements are used to specify the variables containing the stratum identifiers, the cluster identifiers, and the variable containing the individual weights.

For the NAWS, the STRATA are defined as the cycle/region combinations used for the first level of sampling and coded in a variable called dmaregn. The CLUSTER statement contains the primary sampling unit, which is the FLA. The variable for FLA is county_cluster.

The WEIGHT statement references a variable that is for each observation i, the product of both the sampling weight Wt_i and the non-response weight $PWTYCRD_i$. This variable is called pwtycrd for historic reasons.

The Surveymeans procedure also allows for a finite population correction. This option is selected using the TOTAL option on the PROC statement. The total statement allows for the inclusion of the total number of PSUs in each strata. SAS then determines the number of PSUs selected per region from the data and then calculates the sampling rate. In cases such as the NAWS where the sampling rate is different for each strata, the TOTAL option includes a reference to a data set that contains information on all the

---

2 *SAS Institute Inc., SAS/STAT® User's Guide*, Version 8, Cary, NC: SAS Institute Inc., 1999, 61, 3.
3 Woodruff, R. S. (1971), "A Simple Method for Approximating the Variance of a Complicated Estimate," *Journal of the American Statistical Association*, 66, 411–414.
4 Fuller, W. A. (1975), "Regression Analysis for Sample Survey," *Sankhyā*, 37, Series C, Pt. 3, 117–132.

strata and a variable _TOTAL_ that contains the total number of PSUs in that strata.

We include here sample code for Proc Surveymeans to calculate the standard errors for our key estimator WAGET1.

```
      .
proc surveymeans data=naws.crtdvars total=naws.regioninfo;
strata dmaregn;
cluster county_cluster;
var waget1;
weight pwtycrd;
```

Precision of key estimators
Two of the many variables of interest are FWRDAYS, which is the number of days worked per year by a respondent, and WAGET1, which is the average hourly wage of a respondent.

Based on data collected in 2005, and applying the weights that were revised as part of the 2009 independent evaluation (see below under Statistical Consultation), the 2-standard-error confidence interval for the first variable, FWRDAYS, was 183 ± 8.5. That is, with approximately 95% confidence, the average number of days annually worked, per person, lies between 174.5 and 191.5. This constitutes a margin of error of ±4.6% of the estimated value.

For the second variable, the average wage (WAGET1), the interval is 7.95 ± 0.21. With approximately 95% confidence, the average wage lies between $7.74 and $8.16. This yields a margin of error of ±2.6% of the estimated value.

There are numerous other variables of interest, whose standard errors vary greatly. These two are offered as examples that show some of the range of possible precisions obtained.

4. Tests

The questionnaire to be used in the survey was developed by the DOL with input from various Federal agencies. Apart from deleting the Environmental Protection Agency-sponsored questions on pesticide handling, the questionnaire will be unchanged from the version that OMB approved in the last submission. The majority of the questions have been used for over twenty years, are well understood by the sampled respondents, and the data they provide are of high quality.

5. Statistical Consultation

The following individuals have been consulted on statistical aspects of the survey design: Stephen Reder and Robert Fountain, Professors, Portland State University, (503) 725-3999 and 503-725-5204; Phillip Martin, Professor, University of California at Davis

(916) 752-1530; Jeff Perloff, Professor, University of California at Berkeley (510) 642-9574; and John Eltinge, the Bureau of Labor Statistics (BLS) (202) 691-7404.

In 2009, Mathematica Policy Research, Inc. researchers Daniel Kasprzyk, Ph.D., Frank Potter, Ph.D., and Steve Williams ((609) -799-3535) evaluated the equations for the survey's sampling weights and the impact of the weights on key national-level findings.

The data will be collected under contract to the ETA by JBS International, Aguirre Division (650) 373-4900. Analysis of the data will be conducted by Daniel Carroll, ETA (202) 693-2795, and by JBS International, Aguirre Division.

Appendix B: Contacting and Selecting Farm Workers

A. A FARM WORKER QUALIFIES TO PARTICIPATE IN THE NAWS (ELIGIBLE), IF HE/SHE …

1. **WORKS IN** any type of crop agriculture in the United States. This includes "crops" produced in nurseries.
2. **WORKS IN** the production of plants or flowers (including work done in nurseries like planting, cultivating, fertilizing, grafting and seeding).
3. has worked in the last 15 days, at least 4 hours per day, for the contacted employer, and meets any of the criteria mentioned above.

B. A WORKER <u>CANNOT</u> PARTICIPATE IN THE NAWS (INELIGIBLE) IF HE or SHE:

1. Was interviewed by NAWS within the last 12 months in the same location.
2. Is an "H-2A worker." H-2A is a program similar to the "braceros". An H-2A worker is a foreigner who is in the United States on a temporary work visa to work for a specific agricultural employer or association of agricultural employers for a specific period of time (less than a year). At the end of the period, the worker returns to his/her respective country.
3. Works exclusively with livestock (animals: such as bees, horses, fishes, pigs, cows, etc).
4. Hasn't worked for the contacted grower at least one day for 4 hours or more in the last 15 days.
5. Does "non-farm work" for the employer (mechanic, sales, office, etc).
6. Is a family member of the grower or employer and doesn't draw a salary like other farm workers.
7. Is the grower or employer or contractor.
8. Is a sharecropper that makes all operational decisions such as when, where and how to plant, harvest, etc.
9. Works for a packing house or cannery (packing or canning agricultural products) outside of the ranch. **Note:** Workers who are packers or caners <u>can be eligible</u> for the NAWS study <u>if they satisfy the following two requisites</u>:
   - **a)** the canning or packing plant is adjacent or located on the farm, **AND**
   - **b)** at least 50 percent of the produce being packed or canned originated from the ranch of the contacted grower.
10. Works for a landscaping company that just sells, installs, maintains or preserve trees or plants; this includes the planting of ornamental plants and placement of sod.

Whenever a worker doesn't qualify to participate, be gracious and thank him/her for their time and proceed to the next worker.

C. NUMBER OF INTERVIEWS PER GROWER

The Grower Lists indicates the total number of interviews allocated for your assigned county. **NEVER** can the ***total county*** allocation be completed by interviewing workers from ***one single employer***. If this appears likely to happen, call the office for instructions.

Refer to the table below, and find the number of interviews per employer based on the number of workers at the employer on the day visited

Number of workers	Number of Interviews
1 - 2	1
3 - 6	2
7 to 12	3
13-20	4
21-30	5
31-42	6
43-56	7
57-72	8
73-90	9
91-110	10
111-132	11
133 or more	12

Note: Sample the allocated number of workers at the grower (interviewing those that agree to participate) and if the county allocation is not complete, continue onto the next grower. At the last employer complete the number of interviews allocated to that grower on the chart – EVEN IF YOU EXCEED THE COUNTY ALLOCATION.

D. LOCATING THE WORKERS

Once you get permission from the Grower/Employer (and you have documented the number of employed workers) ask the Grower/Employer where you can find the workers. If they are in different locations ask the Grower/Employer: "how many workers are in each location?" Also ask the Grower/Employer (or supervisor assigned by employer) for the best time and location to meet with them.

WORKERS' LOCATIONS

The best time to contact workers
Unless the Grower/Employer gives you permission to speak with his/her employees during working hours, do not make any contacts or appointments or try to interview the workers during their work hours.

Changing work locations
Once the Grower/Employer gives you permission to contact the workers, try to complete your contacts and interviews on the same day the grower gave you permission. You should be aware that from day to day it is common to find that workers in the field change location; and new workers can be in the same field on a different day.

The location of the field is not in the assigned county
If the location of the field or operation of the farm is located outside of the designated county, you **cannot interview** those workers. The farm workers must be physically working in the NAWS assigned county for the particular cycle. That is, it is not unusual that the same Grower/Employer may have farm land and workers in two different counties.

E. HOW TO CHOSE ELIGIBLE WORKERS FOR THE STUDY

Selecting workers located in different areas
If the Grower/Employer informs you that his employees are distributed over more than one fields/crew (in the same county), do the following. Use the table below to identify the number of crews and then randomly select the crews.

Number of crews	Number to select randomly
1 to 2	1
3 to 6	2
7 or more	3

Once you have selected the crews, use the proportional formula, below, to calculate how many from each field/crew you need to interview. The same proportional formula should be used if you locate workers in different residencies. **For example**, if the workers live in two different labor camps or housing then find out how many live in each dwelling and calculate proportionately how many you should interview from each dwelling.

Proportional selection of workers
When you find that workers are divided into different areas, randomly sampling from each group will be necessary to maintain equal likelihood of selection for

everyone. The following formula serves as a guide to calculate the number of workers that should be selected when you find that workers are divided into different areas. In this example, there are 3 sampled fields and you are allowed to conduct 12 interviews for this grower.

a	b	c
Number of workers per location	**Number of workers per location ÷ Total of workers**	**%X# total of interviews = 12**
Field A = 20	20 ÷ 30 = 66.6%	.666 **x** 12 = 08 interviews
Field B = 05	05 ÷ 30 = 16.6%	.166 **x** 12 = 02 interviews
Field C = 05	05 ÷ 30 = 16.6%	.166 **x** 12 = 02 interviews
Workers total = 30		**Total = 12 interviews**

Random Selection
As a sample of workers from a grower/employer is needed, the workers are to be chosen at random. All eligible workers of the grower/employer must have an equal chance of being chosen. Everyone has a chance when selecting crews. Then everyone in the selected crews must have an equal chance of selection. The following are the instructions provided to interviewers:

Random Sampling Instructions for NAWS sampled worksites
Before you go to the site, make sure you have:

- A set of tags with colored stickers on them (at least 12 for each site you expect to visit)
- A set of tags with no stickers (at least 50 for each site you expect to visit)
- A bag (or some other dark container to use to hand out the tags, so that workers can pull the tags without seeing what they're getting)
- Sufficient supplies to carry out surveys with the workers that are selected
- A Sampling Tracking Sheet for each site you expect to visit

Once you have gotten permission from the grower to interview, identify the number of workers on site for that day. Record that number in **Line 1** on the Sampling Tracking Sheet.

NOTE-If the number of workers on the site is less than or equal to the cluster, skip the sampling process and ask all workers to complete the interview. Record the number of workers asked to interview on Line 6 of the Sampling Tracking Sheet and the number completing interviews on Line 7. Leave lines 2-5 blank.

NOTE-for any of these approaches, if any sampled workers refuse the interview- DO NOT REPLACE THEM- move on to the next grower if additional interviews are needed to complete the cluster allocation.

Use the chart above to determine the correct number of interviews to be done; this will be the same number of stickered tags to put into the bag: Record the number of stickered tags you put in the bag on **Line 2** on the Sampling Tracking Sheet.

Next, put enough tags without stickers into the bag so that the total number of tags in the bag equals the number of workers at the site. (For example, if there are 20 workers at the site, and you put 5 stickered tags in the bag, then add another 15 tags.) Record the number of unstickered tags you put in the bag on **Line 3** on the Sampling Tracking Sheet.

One interviewer will go around to each worker and have them pull a tag from the bag, while the other speaks to the group.

At the end of the introduction, the speaker will ask everyone to look at their tags, and ask those who have stickers to come up. Record the number of workers who come up to you with stickered tags, who you ask for an interview on **Line 6** on the Sampling Tracking Sheet.

Carry out the interviews and record the completed number of interviews on **Line 7** of the Sampling Tracking Sheet.

Continue, using the same bag, until you've talked to all workers in the group.

When you have time, count the number of tags left in the bag (if any) and record this number on **Line 4** in the Sampling Tracking Sheet. Count the number of stickered tags left in the bag (if any) and record this number in **Line 5** in the Sampling Tracking Sheet.

Sample Tracking Sheet

County Name_____ Date Visited_____
Grower/Farm name_____ Grower ID_____
Allocation (circle 1) 5 8 10 12

Line	Number of:	
1)	Workers (from employer)	
2)	Stickered tags put in bag(s)	
3)	Unstickered tags put in bag(s) (lines 2+3 should equal line 1)	
4)	Tags left in bag(s) at end (after all groups/after all workers have been offered a tag)	
5)	Stickered tags left	
6)	Workers asked for interview ("contacted" in current system)	
7)	Workers completing interview	

Were there more than one crew: ___ YES ___NO
If yes:
How many crews: _____ How many in each crew (list):
_____
From how many crews did you "randomly" select workers (list):
____

REFERENCES

Fuller, W.A. (1975), "Regression Analysis for Sample Survey," *Sankhyā*, 37, Series C, Pt. 3, 117-132.

Obenauf, W. An Application of Sampling Theory to a Large Federal Survey. Portland State University, Department of Mathematics and Statistics. 2003.

SAS Institute Inc., SAS/STAT® User's Guide, Version 8, Cary, NC: SAS Institute Inc., 1999, 61, 3.

Woodruff, R. S. (1971), "A Simple Method for Approximating the Variance of a Complicated Estimate," *Journal of the American Statistical Association*, 66, 411–414.