Supplementary materials for this article are available at https://doi.org/10.1007/s13253-018-00340-4.

Developing Integer Calibration Weights for Census of Agriculture

Luca SARTORE, Kelly TOPPIN, Linda YOUNG, and Clifford SPIEGELMAN

When conducting a national survey or census, administrative data may be available that can provide reliable values for some of the variables. Survey and census estimates should be consistent with reliable administrative data. Calibration can be used to improve the estimates by further adjusting the survey weights so that estimates of targeted variables honor bounds obtained from administrative data. The commonly used methods of calibration produce non-integer weights. For the Census of Agriculture, estimates of farms are provided as integers so as to insure consistent estimates at all aggregation levels; thus, the calibrated weights are rounded to integers. The calibration and rounding procedure used for the 2012 Census of Agricultural produced final weights that were substantially different from the survey weights that had been adjusted for under-coverage, non-response, and misclassification. A new method that calibrates and rounds as a single process is provided. The new method produces integer, calibrated weights that tend to be consistent with more calibration targets and are more correlated with the modeled census weights. In addition, the new method is more computationally efficient.

Supplementary materials accompanying this paper appear online.

Key Words: Discrete optimization; Coordinate descent; Rounding to integers; Local minimizer; Survey weights estimation; Relative errors.

1. INTRODUCTION

Weights are often adjusted to reduce non-response and coverage errors in a census or a survey. The application of interest here is the United States (US) Census of Agriculture, which use a capture–recapture methodology. However, capture–recapture has been used to evaluate coverage in the US Census of population (see Xi and Tang 2011; Griffin and Mule 2008; Mule 2008; Alho et al. 1993; Hogan 1993). Other methods have been used to adjust weights for coverage or response (see Young et al. 2013; Henry and Valliant 2012; Alho 1990; Tilling and Sterne 1999; Cochran 1978). In all these cases, the weights are

L. Sartore (✉) National Institute of Statistical Sciences, Washington, USA (E-mail: *lsartore@niss.org*). L. Sartore, K.Toppin, L. Young and C. Spiegelman, USDA National Agricultural Statistics Service, Washington, USA (E-mail: *kelly.toppin@nass.usda.gov*), (E-mail: *linda.young@nass.usda.gov*) and (E-mail: *cliff@stat.tamu.edu*). C. Spiegelman Texas A&M University, College Station, USA.

Journal of Agricultural, Biological, and Environmental Statistics
https://doi.org/10.1007/s13253-018-00340-4

Published online: 12 November 2018

generally non-integer and may not provide estimates equal to known population values. Weighting calibration methodology can provide a set of better weights for censuses or any survey for which administrative totals are available. These calibrated weights are assigned to respondent records to account for *under-coverage*, *non-response*, *misclassification*, and fluctuations from known population values. A two-step approach is used to make the adjustment. First, the census or survey design weights are initially adjusted to compensate for *under-coverage*, *non-response*, and/or *misclassification*. Second, calibration is applied to further adjust the weights so that estimates are consistent with population totals obtained from reliable administrative data and other trusted sources.

The US Department of Agriculture's (USDA) National Agricultural Statistics Service (NASS) used the before-mentioned weighting strategy with an additional step of rounding to integers of the post-calibrated weights for its 2012 Census of Agriculture. This rounding step is critical because it ensures that all NASS tables and breakdowns, including those done at the county level, provide counts of farms (not fractional farms) that sum to the grand total. The calibration and rounding phases proved challenging during the 2012 Census of Agriculture. In reviewing the 2012 methodology in preparation for the 2017 Census of Agriculture, it became clear that the process of calibration as well as the subsequent rounding of the final weights to integer values needed improvement. Thus, a new method that produces calibrated, integer-valued survey weights was developed and is the focus of this paper.

In the next section, the needs for new approaches to calibration and rounding are highlighted. Section 1.2 provides a brief overview of the history of calibration. The problem of rounding the calibrated weights to integer values is described in Sect. 1.3. Section 1.4 provides a short introduction about coordinate descent algorithms. A combined calibration and rounding method is presented and developed in Sect. 2. In Sect. 3, the new approach is applied to the data from the 2012 Census of Agriculture, and the results are compared to the calibration and rounding approach used in 2012. Final remarks and conclusions are found in Sect. 4.

1.1. Study Problem: A Review of the 2012 Census of Agriculture

NASS conducts hundreds of surveys each year and prepares reports that cover every aspect of US agriculture, including production of commodities, food supplies, prices paid and received by farms, and farm finances. In the years ending in 2 and 7, NASS also conducts the Census of Agriculture, which provides information on the characteristics of US farms and ranches and the people who operate them. The Census's goal is to account for all farms in the USA. Census estimates are produced at the national, state, and county levels and are used by federal, state, and local governments as well as those who provide services to farms and rural communities. The Census provides a foundation for farm programs and policies, and it impacts community planning, availability of operational loans, and other funding.

Data for the Census of Agriculture are primarily collected from mailed questionnaires. NASS maintains a list frame of agricultural operations in the USA. These potentially meet USDA's farm definition: any place from which $1000 or more of agricultural products were produced and sold, or normally would have been sold, during the year (O'Donoghue et al.

2009). At a specific time, the continuous updates are momentarily interrupted, and the list frame becomes the Census Mail List (CML). All operations on the CML are asked to fill out the Census questionnaire, either online or by mail. Keeping the CML as complete as possible is an ongoing NASS effort, but it does not include all US farms, resulting in list *under-coverage*. Because some farms on the CML do not respond to the Census, *non-response* is present. The operations that do respond are classified as either farms or non-farms based on their response. *Misclassification* occurs when some non-farms are classified as farms, or vice versa. To adjust for *under-coverage*, *non-response*, and *misclassification* NASS uses capture–recapture methodology, also known as dual-system estimation (DSE) (see Young et al. 2017, for details).

NASS obtains information on most commodities from administrative sources, such as USDA Farm Service Agency program data, Agricultural Marketing Services market orders, livestock slaughter data, and cotton ginning data. In 2012, calibration was conducted at the state level for all states except Alaska to ensure that the Census estimates were consistent with state-level administrative data for commodity production. For each state, eight characteristics of farm operations or farm operators were estimated from the DSE weights and used as target values for calibration: 8 categories for the value of agricultural sales, 2 age categories, an indicator for female, 4 categories of race, an indicator for Hispanic origin of the principal farm operator, four sales categories of the 10 major commodities (40 total categories), and 7 categories of farm type. In addition, some commodity estimates were included as targets. Each target was calibrated within a pre-specified tolerance range, which was generally less than 2% of the target. Each state was calibrated separately. The capture–recapture adjusted weights were used as the starting values for the calibration process. However, they were first truncated to be in the interval $[1, 3]$ to enable more of the targeted variables to be within the specified ranges while keeping the calibrated weights within reasonable limits. Then, through calibration, adjusted weights were obtained using an iterative algorithm to solve a constrained least squares problem with the restriction that the calibrated weights were in a specified range. The Census of Agriculture produces estimates at various levels of geography, including the national, state, and county levels. If weights do not have an upper bound, unreasonable estimates may result at the lower levels of geography, especially at the county level. The upper bound varied depending on the size and type of farm. Most farms had weights restricted to the interval $[1, 6]$ (Fetter 2009; Théberge 1999). For some large and/or unique farms, the weights were further restricted to one, two, or three, depending on the farm's influence on the overall production. These are referred to as restricted records.

After calibration, it was rare that all estimates were in the range of their respective targets with all weights being within $[1, 6]$ (Fetter et al. 2005). In these cases, the targets were prioritized. In 2012, the number of farms, total land in farms, and the top cash-receipt commodities accounting for 80% of the state's production were given the highest priority. Within the set of priority targets, the target whose estimate was furthest from the target value was included first. Once that target was hit, the next target with the estimate furthest from the target value was included. If a target could not be hit, it was removed from the list of targets and the next target with the estimate furthest from the target value was included.

Once this process was complete, the step-wise algorithm was again used to calibrate the remaining targets, which had equal priority.

Through calibration, the output weights were set to several decimals; however, Census results are published at the integer level. Rather than rounding estimated totals, weights were rounded to integers. This ensures that all of NASS's published tables and breakdowns, including those at the county level, provide counts of farms (not fractional farms) that sum to the correct grand totals (Scholetzky 2000). Integer weights also guarantee that any user-summary will correctly add up to any level of aggregation.

The 2012 rounding methodology (Kott 2004) used by NASS had a random component. As a consequence of rounding, it was common for some of the calibration equations satisfied during calibration to no longer be satisfied after rounding. Further, the two processes together often produced weights that were quite different from the adjusted (DSE) weights. Since the DSE methodology was used to account for *under-coverage*, *non-response*, and *misclassification* in the 2012 Census of Agriculture, the after-calibration weights (final weights) should be ideally close to the DSE weights. However, this was not always the case at the conclusion of calibration, and it may have caused unnecessary errors for the commodity and demographic variables that were not calibrated.

Given these issues associated with 2012 calibration and rounding methodologies, improvements in the processes of calibration and subsequent rounding of the post-calibrated weights to integer weights were given high priority for the 2017 Census of Agriculture. Thus, a new method that produces calibrated, integer-valued survey weights was developed. The new approach differs from existing ones in two important ways. First, weights are rounded to integers, and then, the optimization is performed by dealing with integer numbers only. Second, the calibration weights are optimized in order to minimize the errors of the totals from the benchmarks while satisfying the range restrictions on the weights instead of minimizing the distance from the DSE weights.

1.2. CALIBRATION: AN OVERVIEW

Lemel (1976) initially introduced calibration to improve the estimates of population totals. Calibration gained importance after Deville (1988), and Deville and Särndal (1992) generalized it. Their methods modify the adjusted weights that appear in the Horvitz–Thompson estimator (Horvitz and Thompson 1952). Suppose that $U = \{1, 2, \ldots, N\}$ is the set of N units in a finite population. Let $t_x = \sum_{i=1}^{N} x_i$ be a total of interest for the variable x. Given a sample $S = \{1, 2, \ldots, n < N\}$ with adjusted weights $d_i = 1/\pi_i$, for any $i \in S$, an estimate for t_x is given by the Horvitz–Thompson estimator $\hat{t}_x^{\text{HT}} = \sum_{i=1}^{n} d_i x_i$. Let $\mathbf{y}_i$ be a vector of variables available for the i-th unit in the sample S for which the population totals are known: $\mathbf{t_y} = \sum_{i=1}^{N} \mathbf{y}_i$. These known totals are the calibration targets. However, the Horvitz–Thompson estimates and the known population totals are often not equal, i.e., $\hat{\mathbf{t}}_{\mathbf{y}}^{\text{HT}} = \sum_{i=1}^{n} d_i \mathbf{y}_i \neq \mathbf{t_y}$, where n denotes the sample size. Calibration resolves this issue by using auxiliary variables to produce the calibration estimator $\hat{t}_x^{\text{Cal}} = \sum_{i=1}^{n} w_i x_i$, where w_i are the calibration weights. Calibration minimizes a distance measure between w_i and the design weights d_i while satisfying the calibration equations

$$\sum_{i=1}^{n} w_i \mathbf{y}_i = \mathbf{t}_{\mathbf{y}}. \tag{1}$$

Deville and Särndal (1992) provided some required properties of the distance measure $G(\mathbf{w}, \mathbf{d})$ and some examples of distance measures. A distance measure of particular importance is the generalized weighted least square given by

$$\frac{1}{2}\sum_{i=1}^{n} \frac{(w_i - d_i)^2}{q_i d_i}, \tag{2}$$

where $1/q_i$ is a known positive weight, which is set to a constant usually equal to 1, or to any other quantity unrelated to d_i (Deville and Särndal 1992). The minimization of Eq. (2) gives $w_i = d_i \left(1 + q_i \mathbf{y}_i^\top \boldsymbol{\lambda}\right)$, where $\boldsymbol{\lambda}$ is a vector of Lagrange multipliers given by $\boldsymbol{\lambda} = \left(\sum_{i=1}^{n} d_i q_i \mathbf{y}_i \mathbf{y}_i^\top\right)^{-1} \left(\mathbf{t}_{\mathbf{y}} - \hat{\mathbf{t}}_{\mathbf{y}}^{\mathrm{HT}}\right)$. The resulting generalized regression estimator of t_x is

$$\hat{t}_x^{\mathrm{GRE}} = \sum_{i=1}^{n} w_i x_i = t_x^{\mathrm{HT}} + \left(\mathbf{t}_{\mathbf{y}} - \hat{\mathbf{t}}_{\mathbf{y}}^{\mathrm{HT}}\right)^\top \mathbf{B}_s,$$

where $\mathbf{B}_s = \left(\sum_{i=1}^{n} d_i q_i \mathbf{y}_i \mathbf{y}_i^\top\right)^{-1} \sum_{i=1}^{n} d_i q_i x_i \mathbf{y}_i$. Although a number of distance measures have been proposed and studied (including those found in Deville and Särndal 1992), $\hat{t}_x^{\mathrm{GRE}}$ is a good point of reference since it has a simple closed-form solution. The calibration estimators derived from many other distance measures are also asymptotically equivalent to $\hat{t}_x^{\mathrm{GRE}}$. Singh and Mohl (1996) heuristically justified the use of calibration estimators, provided some computational algorithms, and compared various methods in terms of weight distribution, estimates, precision, and computational burden. They also distinguished between two types of methods:

1. those that satisfy the calibration equations and iterate until the range restrictions on the weights are met;
2. and those that satisfy the range restrictions and iterate until the calibration equations are met.

They concluded that even though both methods are asymptotically equivalent to the regression method, the interval length of the range restrictions has an impact on the point estimates and associated uncertainty. In fact, neither method may converge to a solution within a reasonable time, if the interval of acceptable weights is short. Théberge (1999) extended the calibration techniques to estimate linear parameters of the population other than totals and means. In addition, when the calibration does not allow for the existence of an optimal solution, he provided a sub-optimal alternative. His method is based on the computation of the calibration weights through a given estimator that exploits linear algebra properties. Duchesne (1999) developed a method that provides robust estimates and satisfies the calibration equations while taking into account the range restrictions on the weights. His procedure performs Newton's iterations to compute the weights according to the approach adopted by Deville and Särndal (1992) by using a restricted mean square metric. The estimator is asymptotically equivalent to that originally provided by Deville and Särndal (1992). Théberge

(2000) further studied the impact of imposing restrictions on the calibration weights and discussed their asymptotic behaviors and the effect of outliers and extreme weights.

Estevao and Särndal (2000) proposed an alternative approach to the calibration problem called the "functional form method." The functional form method removed the limitation that the calibration weights minimize a distance measure. Instead, the calibration weights $\{w_i\}_{i \in S}$ are required to satisfy calibration Eq. (1) and be of the "functional form"

$$w_i = d_i F_i \left(\hat{\boldsymbol{\kappa}}^\top \mathbf{x}_i \right), \tag{3}$$

where $F(\cdot)$ is monotonic function such that $F(0) = 1$ and $F'(0) = 1$ and $\hat{\boldsymbol{\kappa}}$ is a vector of estimated coefficients. $F(\cdot)$ is called the functional equation. This expanded definition of calibration proved helpful for expanding calibration into the realm of adjustments for non-response and coverage errors (Kott 2006).

Some popular functional forms are

1. linear functional
$$F(u) = 1 + u; \tag{4}$$
2. exponential functional
$$F(u) = \exp(u); \tag{5}$$
3. truncated linear functional
$$F(u) = \begin{cases} 1 + u, & \text{if } L - 1 \le u \le U - 1, \\ U, & \text{if } u > U - 1, \\ L, & \text{if } u < L - 1, \end{cases} \tag{6}$$
where U and L are the upper and lower bounds, respectively;
4. logit functional
$$F(u) = \frac{L(U-1) + U(1-L)\exp(Au)}{(U-1) + (1-L)\exp(Au)}, \tag{7}$$
where
$$A = \frac{U - L}{(1-L)(U-1)}.$$

The "linear functional form" can produce negative or extreme weights. The exponential, truncated linear, and logit functionals are often used to counter these problems. However, a lack of closed forms for the exponential, truncated linear, and logit functionals leads to situations where their numerical solvers may not converge.

NASS used a version of the truncated linear functional equation (6) for weight calibration in its 2002, 2007, and 2012 Censuses. However, it was often the case that there was not a set of weights that satisfied both Eqs. (1) and (6). The usual solution is to relax some calibration equations (soft targets) while keeping others strict (hard targets). However, the relaxation of the calibration equations still did not allow the range restrictions to be met. Furthermore, the calibration methodologies reviewed above produce non-integer, calibrated weights.

1.3. ROUNDING

As stated before, given NASS's need for integer weights, a rounding algorithm is required. The literature on rounding in combination with calibration is extremely sparse. The simplest approach to rounding would be to convert the non-integer weights into integers by rounding down if the decimal part is less than 0.5 and rounding up otherwise.

This can be problematic since this method can increase or decrease the weighted estimates enormously. For example, if a single producer accounts for a quarter of a state's production, then rounding their weight up or down changes the total production drastically. The rounding method must be cognizant of that.

In order to minimize the change in the weighted estimates for important census variables, such as weighted number of farms in a county and weighted total land within a county, NASS used probabilistic rounding with the probability of rounding a weight up to its ceiling being proportional to its decimal part (for further details see Kott 2004). While this rounding methodology stabilizes the counties' farm and land counts, it randomizes the other weighted estimated totals.

As a consequence of the probabilistic rounding, it was common for some of the calibration equations satisfied during calibration to no longer be satisfied after rounding. Further, the two processes, calibration and rounding, together often produced weights that were quite different from the input DSE weights. A new algorithm was developed to produce integer calibrated weights. This new approach is called integer calibration (INCA), and it is based on a discrete coordinate descent algorithm.

1.4. COORDINATE DESCENT ALGORITHMS: AN OVERVIEW

The first coordinate descent algorithm was proposed by Cauchy (1847) to solve a system of equations with several variables. Several improvements have been made since then, and it has many applications in machine learning and statistics (Wright 2015). It is especially used to minimize an objective function whose Hessian is singular almost everywhere or too cumbersome to evaluate.

Coordinate descent methods minimize a continuous function $f : \mathbb{R}^n \rightarrow \mathbb{R}$,

$$\min_{\mathbf{w}} f(\mathbf{w}) = f(w_1, \dots, w_n),$$

where $\mathbf{w} \in \mathbb{R}^n$. The minimization is performed iteratively. At each step k, a single component i_k of the gradient ∇f in $\mathbf{w}$ is evaluated, and the value of w_{i_k} is adjusted in the opposite direction of its gradient component such that $f(\mathbf{w}_{k+1}) < f(\mathbf{w}_k)$. These sequential updates are formulated as

$$w_{i_k} \leftarrow w_{i_k} - \alpha \text{ sign}\left([\nabla f]_{i_k}\right),$$

where α represents the length of the adjustment and $[\nabla f]_{i_k}$ denotes the i_k-th component of ∇f (Wright 2015).

In order to reduce the number of iterations, the index i_k can be selected such that the function f is minimized the most. The efficiency of several selection strategies was studied by Nutini et al. (2015), who concluded that the Gauss–Southwell rule (Southwell 1940) is among the most efficient selection criteria. In fact, the index i_k is chosen such that the size of the gradient component $|[\nabla f]_{i_k}|$ is the largest.

2. INTEGER CALIBRATION

In this section, INCA's approach for solving the calibration problem by reformulating it into an optimization problem is established. INCA consists of two sub-algorithms. The first, called rounding, is designed to provide an initial set of integer weights that can satisfy NASS publication requirements. The adjusted weights for under-coverage, non-response, and misclassification are rounded using a computationally efficient procedure that minimizes the error between target values and the weighted totals. The latter sub-algorithm performs a coordinate descent by adding or subtracting integer units to further reduce the distance of the estimates from the targets.

Several computational solutions were implemented to decrease the number of iterations of the algorithms and to perform the minimum number of operations to calculate the quantities needed to provide optimal integer weights. In fact, distinct objective functions are assigned to each sub-algorithm. When rounding, the objective function is designed to force the weighted totals to be closer to the target values and a double penalization is assigned whenever the weighted totals lie outside specified boundaries. Thus, the rounded weights are a sub-optimal initial point for the problem formulated by the calibration objective function, which further forces the weighted totals to be within the boundaries without taking into account the position of calibration targets.

2.1. NOTATION

The following notation is used in this section.

- $\mathbf{d}$ An n-dimensional vector of initial weights
- $\mathbf{w}$ An n-dimensional vector of calibrated weights
- $\mathbf{l_w}$ An n-dimensional vector of lower bounds for the calibrated weights
- $\mathbf{u_w}$ An n-dimensional vector of upper bounds for the calibrated weights
- $\mathbf{A}$ A $p \times n$ matrix of collected data
- $\mathbf{a}_i$ An n-dimensional vector with components of the i-th row of matrix $\mathbf{A}$
- $\mathbf{y}$ A p-dimensional vector of targets (known reliable totals)
- $\mathbf{l_y}$ A p-dimensional vector of lower bounds for the targets
- $\mathbf{u_y}$ A p-dimensional vector of upper bounds for the targets
- ϕ A nonnegative scalar used as tuning parameter

The calibration problem can be reformulated as the following optimization problem:

$$\min_{\mathbf{l_w} \le \mathbf{w} \le \mathbf{u_w}} f(\mathbf{w}) = \min_{\mathbf{l_w} \le \mathbf{w} \le \mathbf{u_w}} [\phi G(\mathbf{w}) + F(\mathbf{w})], \tag{8}$$

subject to the boundary constraints:

$$\mathbf{l_y} \leq \mathbf{Aw} \leq \mathbf{u_y}, \tag{9}$$

where $F(\mathbf{w})$ is an objective function that has a different form in the rounding and calibration sub-algorithms. The function $G(\mathbf{w}) = \|\mathbf{w} - \mathbf{d}\|_1$. Furthermore, to guarantee integer solutions, $\mathbf{w}$ is restricted to integers in the interval of $[\mathbf{l_w}, \mathbf{u_w}]$. Similarly to the LASSO regression methods (Tibshirani 1996), the constant ϕ is used as a trade-off between the objective function and the distance between DSE and calibrated weights. According to Théberge (1999) and Duchesne (1999), it is possible that no solution to the problem in (8) may exist. However, the existence of a sub-optimal solution can be guaranteed by the relaxation of the constraints in (9), which can be reformulated to be included as part of the objective function through $F(\mathbf{w})$.

2.2. INCA Rounding Sub-algorithm

INCA rounds each decimal weight by considering the contributions of the lower and upper integers on the objective function. To restate the problem, the goal is to solve (8) subject to $\mathbf{w} \in \mathcal{N}$, where $\mathcal{N} = \{\mathbf{w} \in \mathbb{N}^p : \mathbf{l_w} \leq \mathbf{w} \leq \mathbf{u_w}\}$.

The objective function $F(\mathbf{w})$ of the rounding sub-algorithm is defined as

$$F(\mathbf{w}) = 2\sum_{i=1}^{p} \frac{|y_i - \hat{y}_i|}{u_i - l_i} + \sum_{i=1}^{p} \begin{cases} (\delta - u_i + \hat{y}_i)/(u_i - \delta), & \text{if } \hat{y}_i > u_i - \delta, \\ (\delta + l_i - \hat{y}_i)/(l_i + \delta), & \text{if } \hat{y}_i < l_i + \delta, \\ 0, & \text{otherwise,} \end{cases} \tag{10}$$

where $\hat{y}_i = \mathbf{a}_i^\top \mathbf{w}$, and if the denominator is zero in any fraction, it is replaced with 1. The scalar δ is a positive constant that prevents false convergence toward unstable solutions, and it is used to shrink the length of the intervals provided around the targets $\mathbf{y}$. The aim of the rounding objective function in (10) is twofold: While forcing the calibrated totals into the target intervals (second addend), it brings the calibrated totals close to the hard targets (first addend). Since the initial vector of real weights is rounded by the minimization of the rounding objective function (10), the integer solution produces totals that might be already an optimal solution for the objective function used for calibration.

The sub-algorithm starts from an initial set of adjusted weights that is forced to be within the feasible set $\mathcal{N}$. This is accomplished by assigning the values of the lower or upper limits to the weights according to the following rule:

$$w_i^* = \begin{cases} l_{w_i}, & \text{if } w_i < l_{w_i}, \\ u_{w_i}, & \text{if } w_i > u_{w_i}, \\ w_i, & \text{otherwise,} \end{cases} \tag{11}$$

for any $i = 1, \ldots, p$.

The rounding sub-algorithm ensures that the integer approximation of the truncated weights (11) reduces the value of the rounding objective function in (10) by optimizing the following problem:

$$\min_{\lfloor \mathbf{w}^* \rfloor \leq \mathbf{w} \leq \lceil \mathbf{w}^* \rceil} [\phi G(\mathbf{w}) + F(\mathbf{w})]. \tag{12}$$

This result is achieved by selecting the upper or lower integer of each non-integer weight such that the minimum formulated in (12) is attained. A sub-optimal solution is obtained by processing a sequence of weights that is formed by assigning the highest priority to the weights having the largest absolute value of the components of the gradient vector. The regular rounding method (i.e., $\lfloor w_i + \frac{1}{2} \rfloor$) is applied to those weights having the gradient equal to zero. The result of these operations is the starting set of integer weights for the coordinate descent sub-algorithm.

2.3. INCA Coordinate Descent Sub-algorithm

Coordinate descent is an optimization technique that successively minimizes an objective function along coordinate directions or coordinate hyperplanes. A coordinate selection rule is used at each iteration k for the determination of the coordinate or block of coordinates to change by maintaining all the others fixed, e.g., the index $i \in \{1, 2, 3, \ldots, n\}$ is chosen for a given initial vector $\mathbf{w}^0 = (w_1^0, w_2^0, w_3^0, \ldots, w_n^0)$, and the solution that minimizes a function f is then obtained as

$$w_i^{k+1} = \min_{w_i^k} f\left(w_1^k, w_2^k, w_3^k, \ldots, w_n^k\right).$$

This step is repeated until a stopping condition is satisfied.

The objective function $F(\mathbf{w})$ of the calibration sub-algorithm is defined as

$$F(\mathbf{w}) = \sum_{i=1}^{p} \begin{cases} (y_i - \hat{y}_i)/(y_i - u_i + \delta), & \text{if } \hat{y}_i > u_i - \delta, \\ (y_i - \hat{y}_i)/(y_i - l_i - \delta), & \text{if } \hat{y}_i < l_i + \delta, \\ 0, & \text{otherwise.} \end{cases} \tag{13}$$

As in the rounding objective function (10), the denominator is replaced with 1 if it is zero in any fraction. This function is designed to lessen the number of iterations during calibration. The sub-algorithm stops when no unit adjustment of the weights leads to any improvements, or when each total is within its boundaries. An optimal value of $\delta > 0$ can be found via cross-validation methods, but these are too expensive in computational terms. Due to the intrinsic nature of this objective function, the constant δ improves the quality of the calibrated weights by reducing the accumulation of totals at the target boundaries. Thus, the minimization of this function enforces the calibration equations within a certain degree of tolerance.

The INCA algorithm performs the discrete search for the optimal integer weights via coordinate descent with the order of possible descent coordinates based on the gradient. The INCA coordinate descent sub-algorithm can be summarized by the following steps:

1. Evaluate the gradient of the calibration objective function at $\mathbf{w}^{(0)} = (w_1^{(0)}, w_2^{(0)}, \cdots, w_n^{(0)})$.

2. Determine the order of the possible descent coordinates based on the largest absolute gradient component.

3. Move the appropriate weight and call the new value $\mathbf{w}^{(1)}$.

4. Repeat steps 1 through 3 with $\mathbf{w}^{(0)}$ by $\mathbf{w}^{(1)}$.

In the INCA coordinate descent, a descent on each coordinate/weight can be regarded as descent in the direction ± 1. For each iteration, the gradient of the calibration objective function is computed and the list of priority indexes is produced using the absolute of the weight's gradient components. INCA processes the first weight w_k on this list, the weight with the largest absolute gradient component, by determining whether the change of the selected weight, $w_k \pm 1$, decreases the calibration objective function. If so, then the weight w_k is updated to this value and INCA moves to the next iteration, recalculating the gradient and the priority list. If not, the next weight on the list is processed and so on. The algorithm stops when the unit shifts do not produce any improvement as measured by the calibration objective function. The selection of coordinates based on the gradient is called the Gauss–Southwell rule (Nutini et al. 2015).

Before discussing the convergence of the algorithm, some definitions are needed. Let $\mathbf{w}^* \in \mathcal{N}$.

Definition 1. The unit neighborhood of $\mathbf{w}^*$, $U(\mathbf{w}^*)$ is defined by

$$U(\mathbf{w}^*) = \left\{\mathbf{w} \in \mathcal{N} \mid w_i \in \{w_i^* - 1, w_i^*, w_i^* + 1\},\ i = 1, \ldots, p\right\}.$$

Definition 2. (*Local minimizer*) A point $\mathbf{w}^* \in \mathcal{N}$ is called an U local minimizer of $f(\cdot)$ over $\mathcal{N}$ if $f(\mathbf{w}^*) \leq f(\mathbf{w})$ for all $\mathbf{w} \in U(\mathbf{w}^*)$ and $\mathbf{w} \neq \mathbf{w}^*$.

Luo and Tseng (1992) provided some results on the convergence of the sequence $\{\mathbf{w}^r\}$ generated by the coordinate descent method based on the Gauss–Southwell rule. These results show that INCA steps converge at least linearly to an element in the set of local minimizers.

3. CALIBRATION OF THE 2012 CENSUS OF AGRICULTURE

In this study, the weights of respondent farms in 49 states were calibrated with INCA. The weights of the farms in Alaska were not considered, since Alaska is excluded from the calibration process. Each state was treated separately from the others to achieve an efficient use of the computational resources both in terms of memory allocation and processing time.

The data collected for the 2012 US Census of Agriculture were used to create a common environment with the same setting for both INCA and the old algorithm. This means that the calibration targets, the target boundaries, and the weight restrictions of each state were set identical to those employed for the production of the official estimates. The DSE weights of the farms in the 49 states were generated with the methodology presented in Sect. 1.1, and used as a common starting point to test and compare the performances of the proposed algorithm.

To show the superiority of INCA compared to the 2012 methodology, a comparison between the numerical achievements of the two algorithms is provided. Three general indexes designed to evaluate the effectiveness of INCA were recorded for each state:

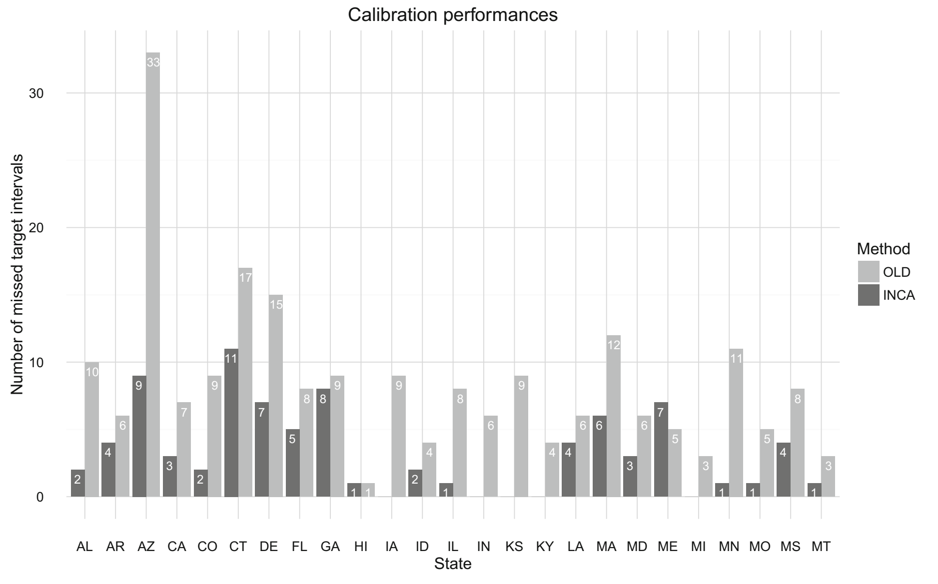

Figure 1. Missed target intervals comparison between INCA and the previous methodology for the 25 states.

- the number of calibrated totals that do not lie within the intervals containing the calibration targets,
- the correlation between the DSE weights and the final weights, and
- the computational running time measured in seconds.

The first index measures how effectively the methods produce the intended results. The second index was provided as an alternative to the mean absolute deviation. In fact, both can quantify how close the final weights are to the initial. The last index determines the computational efficiency of a calibration procedure.

INCA produced less than 5 missed target intervals for the majority of states (38 of the 49 states), and only two states had 10 or more missed target intervals. The 2012 method missed less than 5 target intervals in 14 states and more than 10 in 10 states. In 44 of the 49 states, INCA attained more targets than the 2012 method. Maine, New Hampshire, Nevada, South Carolina, and Vermont are the only states where the final weights from the 2012 method achieved more target intervals than INCA (see Figs. 1 and 2). The inventory for pullets, layers at least 20 weeks old, beef cows, hogs and pigs, other crops (such as grass seed, hay, grass silage, and mint), the number of farms with a total value of production (TVP) between \$1,000,000 and \$2,500,000, and the number of non-equine farms with TVP less than \$2500 were among the most difficult target intervals to achieve. Each of these were missed in at least 10 states. From a computational point of view, Arizona, New Mexico, Nevada, Texas, and the states in New England were the most complex to calibrate since 1% or more of the DSE weights was greater than 6; therefore, smaller weights were forced to increase to meet most of the calibration targets.

The correlations between the calibrated weights from INCA with the DSE weights are higher than those computed with the weights produced by the 2012 method, with the excep-

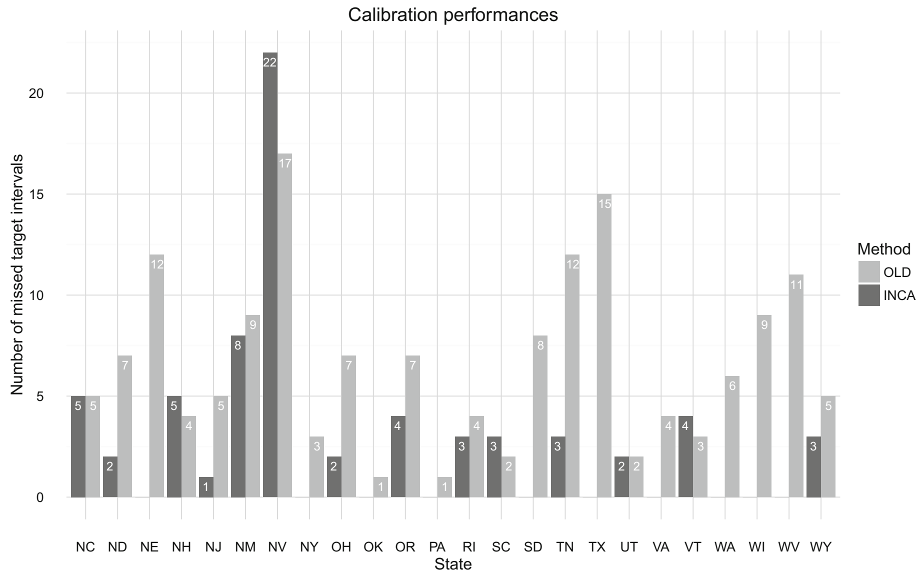

Figure 2. Missed target intervals comparison between INCA and the previous methodology for the 24 states.

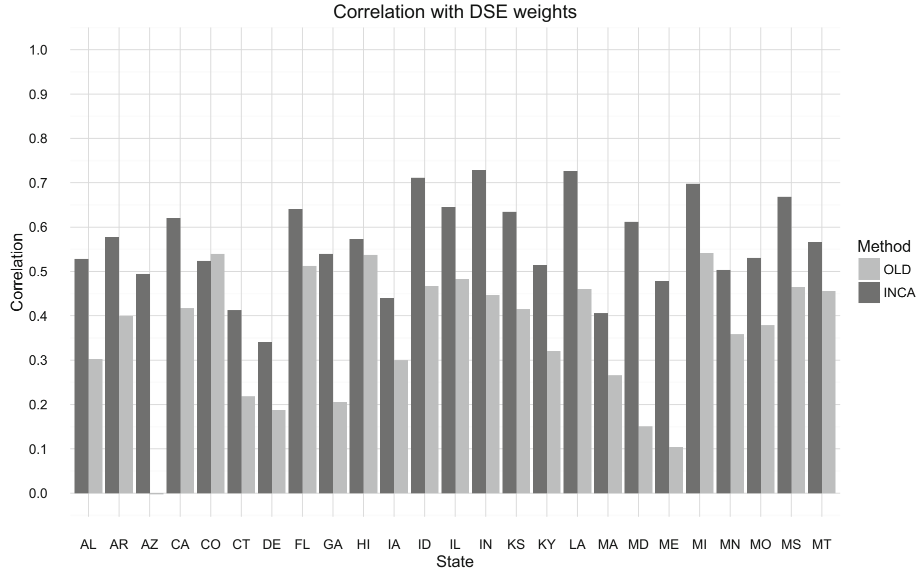

Figure 3. Correlation between final weights and the DSE weight for INCA and the previous methodology for the 25 states.

tion of Colorado (see Figs. 3 and 4). The correlations obtained with the weights produced by INCA for a majority of the 49 states (34 states) have a correlation of at least 0.5 with their capture–recapture weights (DSE). On the other hand, a majority of the correlations obtained with the weights from the old method are mostly below 0.5 (43 states). The old algorithm drastically reduced these correlations in Arizona, Maine, New Hampshire, Nevada, and Vermont.

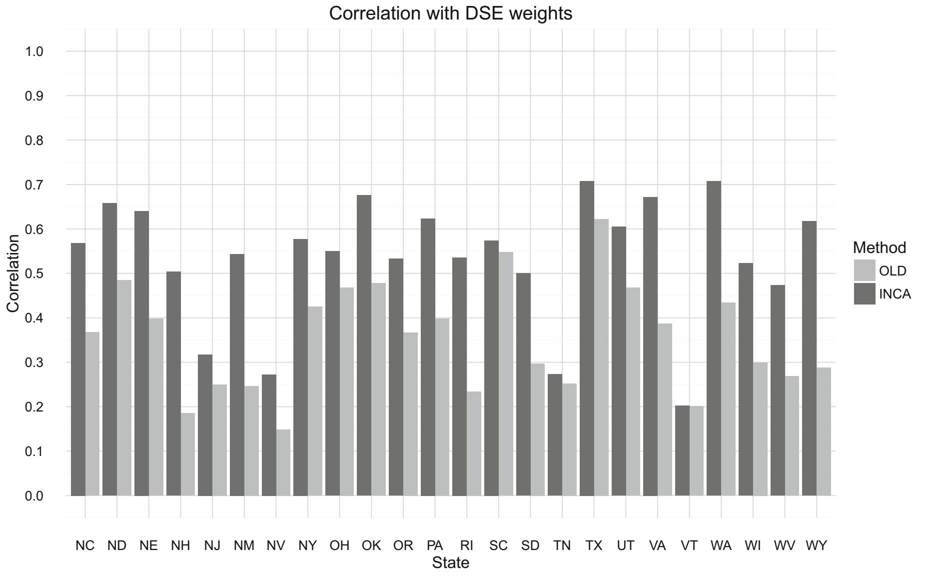

Figure 4. Correlation between final weights and the DSE weight for INCA and the previous methodology for the 24 states.

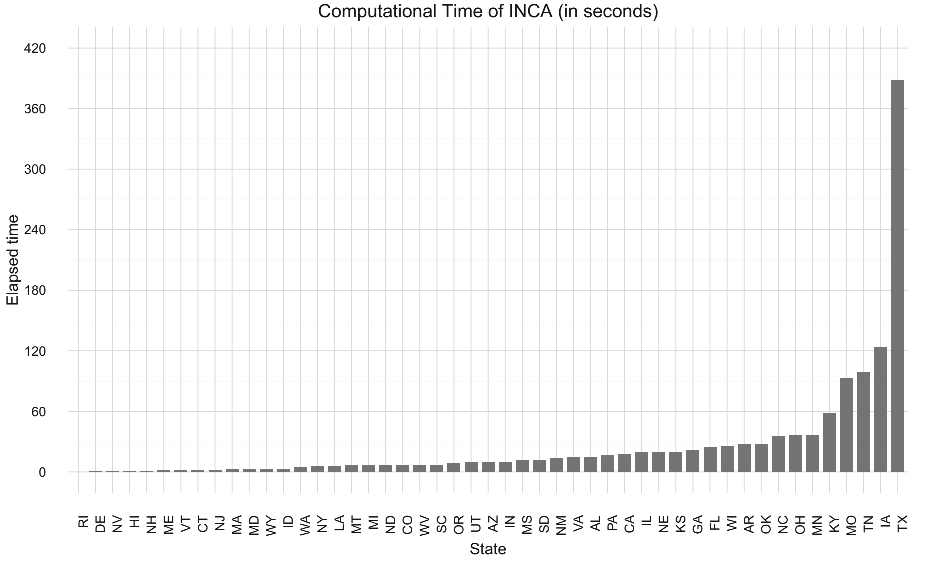

Figure 5. The computation time of INCA for the 49 states included in the Census of Agriculture calibration.

The speed of INCA is remarkable compared to the 2012 method. Using INCA, each state was processed in less than 7 min (see Fig. 5), while the 2012 method needed about 40 min (on average) per state. This radical computational improvement in INCA is also due to the adoption of a sparse representation of data matrices so that redundant information is not processed.

The variances of the point estimates can be calculated using re-sampling methods (Kott 2001; Antal and Tillé 2011; Mashreghi et al. 2016).

4. CONCLUSION

For some applications, calibrated integer weights are needed to develop estimates from survey data. INCA is a new integer calibration algorithm that is a significant upgrade over the current approach of calibration followed by rounding. INCA is the first routine that combines calibration and rounding; in fact, the traditional 2-step process, where calibration takes place before rounding, is reversed in INCA with rounding occurring before calibration. INCA employs a simple yet subtle solution to the calibration problem while producing integer weights. Past calibration methods attempt to minimize the distance of the calibrated weights from the design weights. In contrast, INCA reduces the relative errors of the calibration equations while satisfying the range restrictions of the weights.

INCA relaxes the benchmark constraints and seeks to satisfy them simultaneously. The constraints are relaxed by providing bounds on the population values. This generally makes it more likely that an optimal solution is found compared to considering only exact benchmark constraints. To our knowledge, all other calibration methods consider the benchmark constraints sequentially. The two features of having a range of values associated with each benchmark constraint and processing all benchmark at once allow INCA to meet more targets concurrently and be computationally more efficient.

Moreover, by dealing with integer numbers only, the optimization performed by INCA is more robust and stable. In fact, INCA uses a gradient-based procedure that bypasses the inversion of the Hessian matrix. Conventional methods use pseudo-inverse matrices as stated by Rao and Singh (1997), but the presence of multicollinearity among the variables often leads to unstable estimates. The removal of the calculation of inverse matrices allows INCA to provide accurate results even when the data have variables with perfect collinearity.

The case study with the 2012 Census of Agriculture data highlights how INCA misses fewer targets than the 2012 algorithm while producing final integer weights that are closer to the DSE weights. The computational efficiency of INCA is impressive. Additional numerical tests were performed by simulating an experiment based on a 2^2 factorial design. The source code and the results of these tests are included in the supplemental materials of this paper.

[Received November 2017. Accepted October 2018.]

APPENDICES

A GRADIENTS OF THE OBJECTIVE FUNCTIONS

The gradient of the objective function used for the rounding is

$$\nabla F(\mathbf{w}) = -\mathbf{A}^{\top}\mathbf{v},$$

where the components of $\mathbf{v}$ are given by

$$v_i = 2\,\frac{\operatorname{sign}(\varepsilon_i)}{u_i - l_i} + \begin{cases} 1/(u_i - \delta), & \text{if } \mathbf{a}_i^\top \mathbf{w} > u_i - \delta, \\ -1/(l_i + \delta), & \text{if } \mathbf{a}_i^\top \mathbf{w} < l_i + \delta, \\ 0, & \text{otherwise,} \end{cases} \tag{14}$$

where $\varepsilon_i = y_i - \mathbf{a}_i^\top \mathbf{w}$, for any $i = 1, \ldots, n$.

The gradient of the objective function used for the calibration is

$$\nabla F(\mathbf{w}) = -\mathbf{A}^\top \mathbf{v},$$

where the components of $\mathbf{v}$ are given by

$$v_i = \begin{cases} 1/(y_i - u_i + \delta), & \text{if } \mathbf{a}_i^\top \mathbf{w} > u_i - \delta, \\ 1/(y_i - l_i - \delta), & \text{if } \mathbf{a}_i^\top \mathbf{w} < l_i + \delta, \\ 0, & \text{otherwise.} \end{cases} \tag{15}$$

B EXAMPLE

Consider the following illustration of the INCA methodology. To simplify the computation of the objective function and its gradient, let $\phi = 0$.

- **The setup**

  Bounds on final calibrated weights

$$[1, 6]$$

  The targets

$$\mathbf{y}^\top = \begin{bmatrix} 92 & 61 & 72 \end{bmatrix}$$

  Targets' lower bound

$$\mathbf{l}_\mathbf{y} = \begin{bmatrix} 88 & 58 & 69 \end{bmatrix}$$

  Targets' upper bound

$$\mathbf{u}_\mathbf{y} = \begin{bmatrix} 96 & 64 & 75 \end{bmatrix}$$

  The data matrix

$$\mathbf{A} = \begin{bmatrix} 3 & 0 & 5 & 7 & 9 \\ 1 & 2 & 0 & 8 & 5 \\ 6 & 9 & 5 & 4 & 0 \end{bmatrix}$$

The DSE weights

$$\mathbf{w}_{\text{DSE}}^{\top} = \begin{bmatrix} 15.9 & 0.5 & 1.3 & 3.2 & 1.8 \end{bmatrix}$$

Initial totals

$$\hat{\mathbf{y}}^{\top} = \begin{bmatrix} 63.1 & 42.6 & 64.3 \end{bmatrix}$$

- **Rounding**
  - *Truncation (Pre-rounding adjustments)*
    First, truncate the weights outside the bounds to either 1 or 6.

$$\mathbf{w}^{\top} = \begin{bmatrix} 6 & 1 & 1.3 & 3.2 & 1.8 \end{bmatrix}$$

  - *Initial errors and objective function calculations*
    Initial errors are given by

$$\boldsymbol{\varepsilon} = \mathbf{y} - \mathbf{A}\mathbf{w}$$

$$\begin{bmatrix} 28.9 \\ 18.4 \\ 7.7 \end{bmatrix} = \begin{bmatrix} 92 \\ 61 \\ 72 \end{bmatrix} - \begin{bmatrix} 3 & 0 & 5 & 7 & 9 \\ 1 & 2 & 0 & 8 & 5 \\ 6 & 9 & 5 & 4 & 0 \end{bmatrix} \begin{bmatrix} 6 \\ 1 \\ 1.3 \\ 3.2 \\ 1.8 \end{bmatrix}$$

For example, by setting $\delta = 2$, the initial total loss is given by

$$\begin{aligned} 17.48 = & 2 * (92 - 63.1)/(95 - 89) + (-63.1 + 89 + 2)/91 + \\ & 2 * (61 - 42.6)/(65 - 57) + (-42.6 + 57 + 2)/59 + \\ & 2 * (72 - 64.3)/(75 - 69) + (-64.3 + 69 + 2)/71 \end{aligned}$$

  - *The gradient of the rounding objective function*

$$\nabla F(\mathbf{w}) = -\mathbf{A}^{\top}\mathbf{v}.$$

$$-\mathbf{A}^{\top} = -\begin{bmatrix} 3 & 1 & 6 \\ 0 & 2 & 9 \\ 5 & 0 & 5 \\ 7 & 8 & 4 \\ 9 & 5 & 0 \end{bmatrix}$$

$$\mathbf{v} = \begin{bmatrix} 0.239 \\ 0.317 \\ 0.319 \end{bmatrix}$$

$$\nabla F(\mathbf{w}) = \begin{bmatrix} -2.948 \\ -3.505 \\ -2.790 \\ -5.485 \\ -3.736 \end{bmatrix} = -\begin{bmatrix} 3 & 1 & 6 \\ 0 & 2 & 9 \\ 5 & 0 & 5 \\ 7 & 8 & 4 \\ 9 & 5 & 0 \end{bmatrix} \begin{bmatrix} 0.239 \\ 0.317 \\ 0.319 \end{bmatrix}$$

– *Order of processing*
By taking the absolute value of the gradient

$$|\nabla F(\mathbf{w})| = \begin{bmatrix} 2.948 & 3.505 & 2.790 & 5.485 & 3.736 \end{bmatrix},$$

the following processing order of the weights is obtained:

$$w_4,\ w_5,\ w_2,\ w_1,\ w_3$$

– *Processing the weight in position 4*

$$\mathbf{w}_{lw_4} = \begin{bmatrix} 6 & 1 & 1.3 & 3 & 1.8 \end{bmatrix}$$
$$\mathbf{w}_{uw_4} = \begin{bmatrix} 6 & 1 & 1.3 & 4 & 1.8 \end{bmatrix}$$

The total loss using $\mathbf{w}_{lw_4}$ is given by

$$\begin{aligned} 18.67 = {} & 2 * (92 - 61.7)/(95 - 89) + (-61.7 + 89 + 2)/91 \\ & + 2 * (61 - 41)/(65 - 57) + (-41 + 57 + 2)/59 \\ & + 2 * (72 - 63.5)/(75 - 69) + (-63.5 + 69 + 2)/71 \end{aligned}$$

The total loss using $\mathbf{w}_{uw_4}$ is given by

$$\begin{aligned} 12.73 = {} & 2 * (92 - 68.7)/(95 - 89) + (-68.7 + 89 + 2)/91 \\ & + 2 * (61 - 49)/(65 - 57) + (-49 + 57 + 2)/59 \\ & + 2 * (72 - 67.5)/(75 - 69) + (-67.5 + 69 + 2)/71 \end{aligned}$$

Since the objective function is smaller using $\mathbf{w}_{uw_4}$ than using $\mathbf{w}_{lw_4}$, w_4 is rounded to 4. The new total loss is 12.73.

– *Processing the remaining non-integer weights*
The weight w_5 is similarly rounded, and then, w_3 is processed in the same way. The following output is the resulting vector of weights after the completion of the rounding sub-algorithm:

$$w^\top = \begin{bmatrix} 6 & 1 & 2 & 4 & 2 \end{bmatrix},$$

with a total rounding loss of 9.089.

- **Calibration**
  - *Computing the calibration total loss*

$$\begin{aligned}20 = {} & (92 - 74)/(92 - 88 - 2)\\ & + (61 - 50)/(61 - 58 - 2)\end{aligned}$$

  - *The gradient of the calibration objective function*

$$\nabla F(\mathbf{w}) = -\mathbf{A}^\top \mathbf{v}.$$

$$\mathbf{v} = \begin{bmatrix} 0.5 \\ 1 \\ 0 \end{bmatrix}$$

$$\nabla F(\mathbf{w}) = \begin{bmatrix} -2.5 \\ -2 \\ -2.5 \\ -11.5 \\ -9.5 \end{bmatrix} = -\begin{bmatrix} 3 & 1 & 6 \\ 0 & 2 & 9 \\ 5 & 0 & 5 \\ 7 & 8 & 4 \\ 9 & 5 & 0 \end{bmatrix} \begin{bmatrix} 0.5 \\ 1 \\ 0 \end{bmatrix}$$

  - *Order of processing*
    By taking the absolute value of the gradient

$$|\nabla F(\mathbf{w})| = \begin{bmatrix} 2.5 & 2 & 2.5 & 11.5 & 9.5 \end{bmatrix},$$

the following processing order of the weights is obtained:

$$w_4,\ w_5,\ w_3,\ w_1,\ w_2$$

  - *Iteration 1: processing w_4*
    Compute $F(\mathbf{w})$ by adjusting w_4 in the opposite direction of the gradient. Thus, $w_4 + 1 = 5$, and if $w_4 = 5$, then $F(\mathbf{w}) = 11.5$.

$$11.5 = (92 - 81)/(92 - 88 - 2) + (61 - 58)/(61 - 58 - 2) + (72 - 75)/(72 - 75 + 2)$$

When $w_4 = 5$, then $F(\mathbf{w}) < 20$. Therefore, the updated weights are

$$\mathbf{w}^\top = \begin{bmatrix} 6 & 1 & 2 & 5 & 2 \end{bmatrix}$$

  - *Iteration 2: set priorities for the second step of calibration*

$$\nabla F(\mathbf{w}) = -\mathbf{A}^\top \mathbf{v}.$$

$$\mathbf{v} = \begin{bmatrix} 0.5 \\ 1 \\ -1 \end{bmatrix}$$

$$\nabla F(\mathbf{w}) = \begin{bmatrix} 3.5 \\ -7 \\ 2.5 \\ -7.5 \\ -9.5 \end{bmatrix} = -\begin{bmatrix} 3 & 1 & 6 \\ 0 & 2 & 9 \\ 5 & 0 & 5 \\ 7 & 8 & 4 \\ 9 & 5 & 0 \end{bmatrix} \begin{bmatrix} 0.5 \\ 1 \\ -1 \end{bmatrix}$$

By taking the absolute value of the gradient

$$|\nabla F(\mathbf{w})| = \begin{bmatrix} 3.5 & 7 & 2.5 & 7.5 & 9.5 \end{bmatrix},$$

the following processing order of the weights is obtained:

$$w_5,\ w_4,\ w_2,\ w_1,\ w_3$$

– *Iteration 2: processing the weights*
Compute $F(\mathbf{w})$ by adjusting w_5 in the opposite direction of the gradient. For $w_5 + 1 = 3$, then $F(\mathbf{w}) = 5$.

$$5 = (61 - 63)/(61 - 64 + 2) + (72 - 75)/(72 - 75 + 2)$$

When $w_5 = 3$, then $F(\mathbf{w}) < 11.5$. Thus, the updated weights are

$$\mathbf{w}^\top = \begin{bmatrix} 6\ 1\ 2\ 5\ \mathit{3} \end{bmatrix}$$

– *Iteration 3: set priorities for the second step of calibration*

$$\nabla F(\mathbf{w}) = -\mathbf{A}^\top \mathbf{v}.$$

$$\mathbf{v} = \begin{bmatrix} 0 \\ -1 \\ -1 \end{bmatrix}$$

$$\nabla F(\mathbf{w}) = \begin{bmatrix} 7 \\ 11 \\ 5 \\ 12 \\ 5 \end{bmatrix} = -\begin{bmatrix} 3 & 1 & 6 \\ 0 & 2 & 9 \\ 5 & 0 & 5 \\ 7 & 8 & 4 \\ 9 & 5 & 0 \end{bmatrix} \begin{bmatrix} 0 \\ -1 \\ -1 \end{bmatrix}$$

By taking the absolute value of the gradient

$$|\nabla F(\mathbf{w})| = \begin{bmatrix} 7 & 11 & 5 & 12 & 5 \end{bmatrix},$$

the following processing order of the weights is obtained:

$$w_4,\ w_2,\ w_1,\ w_5,\ w_3$$

– *Iteration 3: processing the weights*

- Compute $F(\mathbf{w})$ by adjusting w_4 in the opposite direction of the gradient. For $w_4 - 1 = 4$: $F(\mathbf{w}) = 10.5$

$$10.5 = (92-83)/(92-88-2) + (61-55)/(61-58-2)$$

  Since if $w_4 = 4$, then $F(\mathbf{w}) > 5$, w_4 is not updated. Therefore, w_2 is consider next.

- Compute $F(\mathbf{w})$ for $w_2 - 1 = 0$: Since w_2 cannot be 0, one cannot decrease w_2. Therefore, one moves to w_1.

- Compute $F(\mathbf{w})$ for $w_1 - 1 = 5$: $F(\mathbf{w}) = 5.5$.

$$5.5 = (92-87)/(92-88-2) + (72-69)/(72-69-2)$$

  Since if $w_1 = 5$, then $F(\mathbf{w}) > 5$, it is not possible to update w_1. Therefore, one moves to w_3

- Compute $F(\mathbf{w})$ for $w_3 - 1 = 1$: $F(\mathbf{w}) = 7.5$.

$$\begin{aligned}&(92-85)/(92-88-2) + (61-63)/(61-64+2)\\&+(72-70)/(72-69-2) = 7.5\end{aligned}$$

  Since if $w_3 = 1$, then $F(\mathbf{w}) > 5$, there is no need to do update w_3. Therefore, one moves to w_5

- Compute $F(\mathbf{w})$ for $w_5 - 1 = 2$: $F(\mathbf{w}) = 11.5$.

$$\begin{aligned}11.5 = {}&(92-81)/(92-88-2) + (61-58)/(61-58-2)\\&+(72-75)/(72-75+2)\end{aligned}$$

  Since if $w_5 = 2$, then $F(\mathbf{w}) > 5$, it is not necessary to update w_5. Therefore, the algorithm stops.

- **Final Weights**

The final calibrated weights are

$$\mathbf{w}^\top = \begin{bmatrix}6 & 1 & 2 & 5 & 3\end{bmatrix}.$$

The final calibrated totals are

$$\hat{\mathbf{y}} = \begin{bmatrix}90 & 63 & 75\end{bmatrix}.$$

By construction of the matrix $\mathbf{A}$, the summation of the weights is not part of the targets in this example. The purpose of this example is to show how the algorithm works rather than showing what type of results are attainable. At the end, the correlation between the initial vector of DSE weights and the final vector of calibrated weights is about 0.8, which is even higher than those obtained in the real case example provided in Sect. 3 (see Figs. 3 and 4).

REFERENCES

Alho, J. M., Mulry, M. H., Wurdeman, K., and Kim, J. (1993). Estimating heterogeneity in the probabilities of enumeration for dual-system estimation. *Journal of the American Statistical Association*, 88(423):1130–1136.

Alho, J. M. (1990). Logistic regression in capture-recapture models. *Biometrics*, 46(3):623–635.

Antal, E. and Tillé, Y. (2011). A direct bootstrap method for complex sampling designs from a finite population. *Journal of the American Statistical Association*, 106(494):534–543.

Cauchy, A. (1847). Méthode générale pour la résolution des systemes déquations simultanées. *Comp. Rend. Sci. Paris*, 25(1847):536–538.

Cochran, W. G. (1978). Laplace's ratio estimator. In *Contributions to survey sampling and applied statistics*, pages 3–10. Elsevier.

Deville, J.-C. (1988). Estimation linaire et redressement sur information auxiliaire d'enqutes par sondage.

Deville, J.-C. and Särndal, C.-E. (1992). Calibration estimators in survey sampling. *Journal of the American Statistical Association*, 87(418):376–382.

Duchesne, P. (1999). Robust calibration estimators. *Survey Survey Methodology*, 25:43–56.

Estevao, V. M. and Särndal, C.-E. (2000). A functional form approach to calibration. *Journal of Official Statistics*, 16(4):379–399.

Fetter, M. (2009). An overview of coverage adjustment for the 2007 census of agriculture. In *Proceeding of the Government Statistics Section, JSM*, pages 3228–3236.

Fetter, M., Gentle, J., and Perry, C. (2005). Calibration adjustment when not all the targets can be met. In *Proceeding of the Survey Research Method section Statistics Section, ASA*, pages 3031–3035.

Griffin, R. and Mule, T. (2008). Spurious Events in Dual System Estimation. Technical Report 2010-E-20, United States Department of Commerce, United States Census Bureau, Washington, DC.

Henry, K. and Valliant, R. (2012). Methods for adjusting survey weights when estimating a total. *Proceedings of the 2012 Federal Committee on Statistical Methodologys Research Conference*.

Hogan, H. (1993). The 1990 post-enumeration survey: Operations and results. *Journal of the American Statistical Association*, 88(423):1047–1060.

Horvitz, D. G. and Thompson, D. J. (1952). A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685.

Kott, P. (2006). Using calibration weights to adjust for nonresponse and coverage errors. *Survey Methodology*, 32:133–142.

Kott, P. S. (2001). The delete-a-group jackknife. *Journal of Official Statistics*, 17(4):521.

Kott, P. S. (2004). Collected techincal notes on weighting and its impact to the 2002 census of agriculture. *United States Department of Agriculture Report*.

Lemel, Y. (1976). Une gnralisation de la mthode du quotient pour le redressement des enqutes par sondage. *Annales de l'ins*, (22/23):273–282.

Luo, Z. Q. and Tseng, P. (1992). On the convergence of the coordinate descent method for convex differentiable minimization. *Journal of Optimization Theory and Applications*, 72(1):7–35.

Mashreghi, Z., Haziza, D., Léger, C., et al. (2016). A survey of bootstrap methods in finite population sampling. *Statistics Surveys*, 10:1–52.

Mule, T. (2008). 2010 Census Coverage Measurement Estimation Methodology. Technical Report 2010-E-18, United States Department of Commerce, United States Census Bureau, Washington, DC.

Nutini, J., Schmidt, M., Laradji, I., Friedlander, M., and Koepke, H. (2015). Coordinate descent converges faster with the gauss-southwell rule than random selection. In Bach, F. and Blei, D., editors, *Proceedings of the 32nd*

International Conference on Machine Learning, volume 37 of *Proceedings of Machine Learning Research*, pages 1632–1641, Lille, France. PMLR.

O'Donoghue, E., Hoppe, R. A., Banker, D., and Korb, P. (2009). Exploring alternative farm definitions: implications for agricultural statistics and program eligibility. *Economic Information Bulletin-USDA Economic Research Service*, 49.

Rao, J. and Singh, A. (1997). A ridge-shrinkage method for range-restricted weight calibration in survey sampling. In *Proceedings of the section on survey research methods*, pages 57–65. American Statistical Association Washington, DC.

Scholetzky, W. (2000). Evaluation of integer weighting for the 1997 Census of Agriculture. Technical Report RD-00-01, United States Department of Agriculture, National Agricultural Statistics Service, Washington, DC.

Singh, A. and Mohl, C. (1996). Understanding calibration estimators in survey sampling. *Survey Methodology*, 22(2):107–115.

Southwell, R. V. (1940). *Relaxation Methods in Engineering Science—a Treatise On Approximate Computation*. Oxford University Press, Oxford.

Théberge, A. (1999). Extension of calibration estimators in survey sampling. *Journal of the American Statistical Association*, 94(446):635–644.

Théberge, A. (2000). Calibration and restircted weights. *Survey Methodology*, 26(1):99–107.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288.

Tilling, K. and Sterne, J. A. (1999). Capture-recapture models including covariate effects. *American journal of epidemiology*, 149(4):392–400.

Wright, S. J. (2015). Coordinate descent algorithms. *Mathematical Programming*, 151(1):3–34.

Xi, C. S. and Tang, C. Y. (2011). Properties of census dual system population size estimators. *International Statistical Review*, 79(3):336–361.

Young, L. J., Lamas, A. C., and Abreu, D. A. (2017). The 2012 Census of Agriculture: a capture–recapture analysis. *Journal of Agricultural, Biological and Environmental Statistics*, 22(4):523–539.

Young, L. J., Lamas, A. C., Abreu, D. A., Wang, S., and Adrian, D. (2013). Statistical methodology for the 2012 us census of agriculture. *In the Proceeding 59th ISI World Statistics Congress*, pages 1063–1068.