

Espuny-Pujol, Ferran; Morrissey, Karyn; Williamson, Paul

Working Paper

A global optimisation approach to range-restricted survey calibration

HEG Working Paper, No. 16-01

Provided in Cooperation with:

Health Economics Group (HEG), University of East Anglia

Suggested Citation: Espuny-Pujol, Ferran; Morrissey, Karyn; Williamson, Paul (2016) : A global optimisation approach to range-restricted survey calibration, HEG Working Paper, No. 16-01, University of East Anglia, Health Economics Group (HEG), Norwich

This Version is available at:

<https://hdl.handle.net/10419/197778>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

HEALTH ECONOMICS GROUP

Faculty of Medicine and Health
Norwich Medical School



University of East Anglia



ESRC Business and Local Government
Data Research Centre

A global optimisation approach to range-restricted survey calibration

Ferran Espuny-Pujol ¹

Karyn Morrissey ²

Paul Williamson ³

¹ Health Economics Group, Norwich Medical School, University of East Anglia

² European Centre for Environment and Human Health, University of Exeter Medical School

³ Department of Geography and Planning, School of Environmental Sciences, University of Liverpool

A global optimisation approach to range-restricted survey calibration*

Ferran Espuny-Pujol¹, Karyn Morrissey², and Paul Williamson³

¹Health Economics Group, Norwich Medical School, U. of East Anglia, UK

²European Centre for Environment and Human Health, U. of Exeter Medical School, UK

³Dept. of Geography & Planning, School of Environmental Sciences, U. of Liverpool, UK

September 7, 2016

Abstract

Survey calibration methods modify minimally unit-level sample weights to fit domain-level benchmark constraints (BC). This allows exploitation of auxiliary information, e.g. census totals, to improve the representativeness of sample data (addressing coverage limitations, non-response) and the quality of estimates of population parameters.

Calibration methods may fail with samples presenting small/zero counts for some benchmark groups or when range restrictions (RR), such as positivity, are imposed to avoid unrealistic or extreme weights. User-defined modifications of BC/RR performed after encountering non-convergence allow little control on the solution, and penalization approaches modelling infeasibility may not guarantee convergence. Paradoxically, this has led to underuse in calibration of highly disaggregated information, when available.

We present an always-convergent flexible two-step Global Optimisation (GO) survey calibration approach. The feasibility of the calibration problem is assessed, and automatically controlled minimum errors in BC or changes in RR are allowed to guarantee convergence in advance, while preserving the good properties of calibration estimators. Modelling alternatives under different scenarios, using various error/change and distance measures are formulated and discussed. The GO approach is validated by calibrating the weights of the 2012 Health Survey for England to a fine age/gender/region cross-tabulation (378 counts) from the 2011 Census in England and Wales.

Keywords Calibration estimation; Calibration weighting; Design-based inference; Generalised regression; Penalized calibration; Raking; Ridge calibration; Range restrictions; Survey weighting

* *Address for correspondence:* Ferran Espuny-Pujol, Norwich Medical School, University of East Anglia, Norwich, NR4 7TJ, United Kingdom. *E-mail:* F.Espuny-Pujol@uea.ac.uk

1 Introduction

Survey calibration incorporates auxiliary information to a sample in two closely related ways: weighting and estimation. *Calibration weights* make a sample consistent with auxiliary information (e.g. census population totals) while in general respecting the initial sample design (Deville and Särndal, 1992). Resulting *calibration estimates* of population parameters (e.g. totals) improve direct sample estimates (e.g. Horvitz-Thompson). Survey calibration methods can be also applied to adjust for non-response or coverage limitations, and to outlier detection (Särndal, 2007). The internal consistency of administrative data can be intrinsically guaranteed with survey calibration, since it may provide a common degree of agreement between estimates from multiple samples of the same population (Wu and Lu, 2016).

Calibration estimates were initially introduced for finite population totals or averages of either categorical or continuous variables. Example methods are the *generalised regression* (GREG) and *raking* estimators (Deville and Särndal, 1992; Singh and Mohl, 1996). Calibration estimates were later developed for variance and bilinear parameters (Théberge, 1999), quantiles and ratios (Särndal, 2007). Given an initial value for frequency tables with no zeroes, if either auxiliary cells or marginal counts are known, the corresponding *post-stratification* problems can also be modelled using survey calibration (Deville and Särndal, 1992). In this context, raking ratio estimates can be obtained recursively using the *Iterational Proportional Fitting* (IPF) algorithm (Wu and Lu, 2016), which dates back to (Deming and Stephan, 1940).

Survey calibration methods are *design-based*: the primary source of randomness is the probability of the sample design (Särndal, 1978). The model-based approach equivalent to calibration estimation is the theory of regression estimation (Fuller, 2002). It requires models of the population to incorporate auxiliary information, and it neither necessarily creates calibration weights nor guarantees to the same extent the internal consistency. Nonetheless, some regression estimators can be obtained as calibration estimators (Särndal, 2007). For example, the GREG calibration estimate of population totals is assisted by a linear effects model, and the raking ratio estimation in frequency tables (using IPF) is assisted by log-linear and logistic regression models (Bishop, 1969; Fuller, 2002). Moreover, all calibration estimators for linear population parameters are asymptotically equivalent to the GREG estimator (Deville and Särndal, 1992).

Survey calibration methods search for (real-valued) calibration weights that: *i*) satisfy a set of *benchmark constraints* (BC) or calibration equations, and, in most cases *ii*) are close to initial weights. Therefore, calibration estimators are: *i*) *design-consistent*, and *ii*) (asymptotically) *design-unbiased* (Deville and Särndal, 1992; Fuller, 2002; Särndal, 2007). In general, the use of auxiliary information in form of BC allows for bias and/or variance reduction in population-level estimates. The bias in calibration estimators is kept

small by staying close to initial (design) weights through the minimisation of a distance measure. Some calibration methods use distances with undesirable effects, that are likely to inflate the bias and/or variance: GREG may produce negative weights, and raking extreme ones. Outliers, small domain estimation or the estimation of non-linear population parameters are also likely to produce extreme and highly variable weights (Théberge, 2000; Wu and Lu, 2016). *Range restrictions* (RR) on weights are imposed in practice in order to avoid weights taking unrealistic or extreme values. These can be imposed directly on weights or through the function measuring the distance to initial weights (Singh and Mohl, 1996).

Failure of survey calibration methods may occur with real data (Sautory, 1991; Tanton et al., 2011). In fact, BC may have no exact solution (zero error), either considered solely or in combination with RR on weights (Singh and Mohl, 1996). This can be due to: the sample not being representative (enough) of every non-void class in the cross-classified BC; the RR being too tight; the BC forming an inconsistent system of equations, due to their derivation from differing data sources or from data with added noise as a result of statistical disclosure control procedures (Tanton et al., 2011). In addition, calibration algorithms may be unstable for too many BC or when multi-collinear survey variables are being benchmarked (Sautory, 1991; Rao and Singh, 1997). The existence of a solution to the range-restricted survey calibration problem was studied theoretically in (Théberge, 2000).

Faced with non-convergence of standard algorithms in practice, alternative approaches have been proposed: non-modelling heuristics and penalized calibration, also known as ridge calibration. Heuristics typically used include: broadening the categorisation of benchmark variables, modifying the benchmarking values, loosening the RR on weights, or even deleting some BC (Sautory, 1991; Bankier et al., 1992). In (Tanton et al., 2011), an upper threshold on the *total absolute error* (TAE) in BC is used to accept non-convergent solutions satisfying RR, with no further control for error in BC. *Penalized calibration* allows a certain degree of relaxation in each BC, being controlled by costs (or tolerances), whilst still providing approximately unbiased and asymptotically design-consistent estimators. Penalized versions of GREG with RR can be found adaptively adjusting the set of tolerances on BC errors in (Rao and Singh, 1997) and addressing its global minimisation in (Wagner, 2013). These methods are asymptotically equivalent to the GREG calibration method (Théberge, 2000), and have regression-based counterparts (Beaumont and Bocci, 2008). However, non-convergence is still reported by penalization methods, being possible even for loose RR on weights.

This paper proposes a Global Optimisation (GO) two-step approach to range-restricted survey calibration. First, the problem feasibility is guaranteed by optimally modifying the BC and (for the first time) RR, if needed and in a controlled manner. Second, the (always-feasible) resulting calibra-

tion problem is solved. The GO method can provide asymptotically design-consistent and realistic solutions, avoiding non-convergence problems and thus overcoming the typical need for user-defined heuristics. Moreover, by keeping close to initial weights, GO is approximately unbiased when design weights are available, otherwise benefiting from the *a priori* information provided by initial weights.

In Section 2, we provide the problem formulation and a technical review of the related methods. In Section 3, it is shown that the feasibility of a range-restricted calibration problem can be checked in advance by solving a (generally sparse) linear program. Moreover, using the ℓ_1 -norm as an error measure, it is shown that feasibility can be achieved if allowing for a minimum (TAE) error in BC and/or a minimum change in RR. In both cases, the formulations are sparse linear programming problems, which have the advantage of being efficiently solvable for many variables; ℓ_1 -norm penalization for errors in addition allows modification of only a small number of benchmark totals or range restrictions. Alternative error functions and modelling options are also discussed.

In Section 4 the generic GO algorithm for an always-convergent globally optimal survey calibration is presented. The Chi-square distance as in GREG is used for demonstration, its minimisation being globally optimal, leading to approximately unbiased estimators. In fact, GREG-based methods allowing for RR on weights are a particular case of GO, provided that the earlier are convergent (to a global optimum).

In Section 5 the performance of the GO method is exemplified with real data by calibrating the weights of the 2012 Health Survey for England to population totals from the 2011 Census in England and Wales. Broad age/gender (20 groups) and region (9 groups) total counts are imposed as exact BC, while optimally controlled errors are allowed to calibrate HSE weights to a fine age/gender/region cross-tabulation (378 counts). The resulting calibration weights are further used to estimate the total counts for a broad age by economic activity cross-tabulation, which is used for validation purposes.

Note. This paper assumes that calibration is performed using within-domain samples. Besides, the approach described herein can be applied also in other fields to perform *indirect survey calibration*, like microsimulation, which may combine surveys and benchmarks corresponding to different periods of time or to non-exactly matching domains or areas (Ballas and Clarke, 2009; Tanton, 2014) or which may calibrate survey-based but otherwise simulated data (Pudney and Sutherland, 1994; Wittenberg et al., 2011).

2 Background

Assume as given survey data corresponding to a sample \mathcal{S} of size n , drawn from a population \mathcal{U} , together with a n -dimensional vector \mathbf{d} of survey

weights (by default, but not limited to, the sample design weights). For \mathbf{x} , a p -dimensional vector of variables, assume as known the $n \times p$ survey matrix \mathbf{X} , which contains the values of \mathbf{x} for all sample units. For simplicity, the totals for variables are considered as the measure of interest; the formulation is similar when using proportions or averages. The variables can be either continuous or categorical, the latter possibly expressed using indicator binary variables (1-0 valued) for each category group in order to exploit known total unit counts for that group. The Horvitz-Thompson direct survey estimate of the population totals for the values of \mathbf{x} is $\mathbf{t}_x^{HT} = \mathbf{X}'\mathbf{d}$ (Horvitz and Thompson, 1952). Also assume as given a more precise estimate \mathbf{t}_x of the totals of \mathbf{x} for the population, this estimate provided for example by administrative census sources.

2.1 Survey calibration

Survey calibration aims at determining new survey weights \mathbf{w} that make the survey compatible with the known auxiliary totals, i.e. satisfying the *benchmark constraints*¹:

$$\mathbf{X}'\mathbf{w} - \mathbf{t}_x = \mathbf{0} . \quad (\text{BC})$$

The weights should realistically represent units: for example, counts of households or persons have to be positive (in general, $\ell_1 \leq \mathbf{w} \leq u_1$, for two constants ℓ_1, u_1). Moreover, a drastic change in any particular weight from its initial value should be avoided (in general, $\ell_2\mathbf{d} \leq \mathbf{w} \leq u_2\mathbf{d}$, for two constants ℓ_2, u_2). Accordingly, the weights can be subject to *range restrictions* in the form

$$\mathbf{l} \leq \mathbf{w} \leq \mathbf{u} , \quad (\text{RR})$$

being \mathbf{l} and \mathbf{u} known vectors². Finally, in order to lead to unbiased estimates, the weights should ideally respect as much as possible the set of initial weights \mathbf{d} , which is achieved by minimising a distance $\mathcal{G}_d(\mathbf{w})$ between \mathbf{w} and \mathbf{d} .

Therefore, the mathematical problem associated to *range-restricted survey calibration* reads:

$$\begin{aligned} \arg \min_{\mathbf{w}} \quad & \mathcal{G}_d(\mathbf{w}) \\ \text{s.t.} \quad & \mathbf{X}'\mathbf{w} - \mathbf{t}_x = \mathbf{0} \quad (\text{BC}) \\ & \mathbf{l} \leq \mathbf{w} \leq \mathbf{u} \quad (\text{RR}) \end{aligned} \quad (1)$$

being \mathbf{w} the *calibration weights* searched for³. The resulting weights are used to make the survey compatible with known auxiliary totals, and in particular can be used to adjust to non-response or coverage errors.

¹The case of a known group total \mathbf{t}_G is a particular case of BC in the form $\mathbf{l}'_G\mathbf{w} - \mathbf{t}_G = \mathbf{0}$.

²The case with no RR is a particular case in which $\mathbf{l} = -\infty, \mathbf{u} = \infty$.

³Alternatively used formulations of the calibration problem computing the relative change in weights $\mathbf{g} = \mathbf{w}/\mathbf{d}$ are equivalent to (1).

Following (Deville and Särndal, 1992), the function $\mathcal{G}_{\mathbf{d}}(\mathbf{w})$ is assumed to be, for every fixed $\mathbf{d} > 0$: non-negative, differentiable, strictly convex, defined on an interval containing \mathbf{d} , such that $\mathcal{G}_{\mathbf{d}}(\mathbf{d}) = 0$, and having a differential continuous and locally invertible at \mathbf{d} . A typical function $\mathcal{G}_{\mathbf{d}}$ for the distance to initial weights is the *modified Chi-square* or generalised least squares distance (Singh and Mohl, 1996)⁴

$$\mathcal{G}_{\mathbf{d}}^{GREG}(\mathbf{w}) = (\mathbf{w} - \mathbf{d})' \mathbf{D}^{-1} (\mathbf{w} - \mathbf{d}) , \quad (2)$$

where $\mathbf{D} = \text{diag}(\mathbf{d})$ is a diagonal matrix with the elements of \mathbf{d} in the diagonal. In that case, the resolution of the survey calibration problem (1) if ignoring any range restrictions (RR) gives the *generalised regression* weights (Deville and Särndal, 1992; Merkouris, 2010):

$$\mathbf{w}^{GREG} = \mathbf{d} + \mathbf{D}\mathbf{X}(\mathbf{X}'\mathbf{D}\mathbf{X})^{-1}(\mathbf{t}_{\mathbf{x}} - \mathbf{X}'\mathbf{d}) , \quad (3)$$

the ratio estimator weights being a particular case for $p = 1$ if replacing \mathbf{D} with $\text{diag}(\mathbf{X})^{-1}\mathbf{D}$ (Deville and Särndal, 1992).

Another commonly used distance function is the *modified discrimination information* associated with the raking estimator (Singh and Mohl, 1996):

$$\mathcal{G}_{\mathbf{d}}^{MDI}(\mathbf{w}) = \sum_{i=1}^n \left(w_i \log \left(\frac{w_i}{d_i} \right) - w_i + d_i \right) . \quad (4)$$

There is no explicit formula to obtain the raking weights, which, when ignoring (RR), have the form $\mathbf{w}^{MDI} = \mathbf{D} \exp(\mathbf{X}\boldsymbol{\lambda})$, for $\boldsymbol{\lambda}$ a p -dimensional vector (Lagrange multiplier) solution of $\mathbf{t}_{\mathbf{x}} = \mathbf{X}'\mathbf{D} \exp(\mathbf{X}\boldsymbol{\lambda})$ (Deville and Särndal, 1992). The raking ratio algorithm in (Deming and Stephan, 1940) provided a solution for the particular case of contingency tables (poststratification), which translates into the benchmark variables being categorical group-membership indicators, some linear combination/s of which is/are unity (i.e. the groups need not to be mutually exclusive) (Deville and Särndal, 1992; Kott, 2009).

It should be understood that formula (3) makes sense only when the matrix $(\mathbf{X}'\mathbf{D}\mathbf{X})$ has full rank. In fact, all the broadly used calibration methods that were proposed in (Singh and Mohl, 1996) involved a term of the form $(\mathbf{X}'\mathbf{D}\mathbf{X})^{-1}$. A well-known source for non-convergence of calibration methods is the lack of representativeness in the survey, which translates into at least one "survey total" being practically zero. In such case, at least one column in \mathbf{X} is close to or equal to zero, and the matrix $(\mathbf{X}'\mathbf{D}\mathbf{X})$ has a row and a column close to or equal zero, meaning that the inverse of that matrix is either not defined or unreliable (in the latter case, the GREG formula has poor performance); this also happens if multiple variables (and so \mathbf{X} columns)

⁴In the case of known totals and initial weights consistent with those totals ($\mathbf{1}'\mathbf{w} = \mathbf{1}'\mathbf{d}$ constant), the minimisation of the Chi-square distance simplifies to minimising $\mathbf{w}'\mathbf{D}^{-1}\mathbf{w}$.

are collinear. Numerical methods based on robust matrix decompositions exist to deal with the resolution of weights without computing the inverse of any matrix, which therefore can address collinearity issues. However, if one column in \mathbf{X} is equal to zero, the equation (BC) can not be satisfied given that the associated benchmark total in \mathbf{t}_x is different from zero. Therefore, in the first case non-convergence is not amendable by any method without changing the original formulation.

2.2 Calibration estimators, small domains and variance

Assume additionally as given a $n \times r$ survey matrix \mathbf{Y} containing the values of \mathbf{y} , a r -dimensional variable of interest, for all n sample units in \mathcal{S} . The population totals \mathbf{t}_y for the variable \mathbf{y} can be estimated using the direct Horvitz-Thompson estimator $\mathbf{Y}'\mathbf{d}$, however the variance of this estimator is high. The calibration estimators make use of calibration weights \mathbf{w}^{Cal} , which account for available auxiliary information, to produce the new estimate

$$\mathbf{t}_y^{Cal} = \mathbf{Y}'\mathbf{w}^{Cal} . \quad (5)$$

In particular, calibration estimators are design-based, not making use of any regression model linking the target variable \mathbf{y} with the auxiliary variables \mathbf{x} . Given that the considered calibration distances $\mathcal{G}_d(\mathbf{w})$ satisfy the properties assumed in Section 2.1, calibration estimators are both asymptotically design unbiased and design-consistent, all of them being asymptotically equivalent (Deville and Särndal, 1992). Moreover, if the auxiliary information is sufficiently related to the variable \mathbf{y} , calibration estimators are more efficient than the Horvitz-Thompson estimator (Fuller, 2002).

When wanting to produce estimators for a small domain \mathcal{D} of the population \mathcal{U} , it is no longer efficient to use the calibration weights (that were computed using X for the whole survey \mathcal{S} and auxiliary totals \mathbf{t}_x for the whole population). Survey calibration at the domain level and/or knowledge on the domain size, or combining information from multiple surveys at the domain level, provides approximately unbiased design-consistent estimators with substantial variance reduction with respect to other estimators (Merkouris, 2010). An intermediate option is adopted in small area estimation (by, for example, but not limited to, spatial microsimulation) when (sufficient) survey data are not available for a small area, consisting in using out-of-area survey data in combination with known area totals (Tanton, 2014). Although not further developed here, the approach to calibration outlined in this paper applies to both of these types of small domain estimation.

The variance of calibration estimators can be approximated asymptotically using the fact that all estimators are asymptotically equivalent to the generalised regression estimator $\mathbf{Y}'\mathbf{w}^{GREG}$ (Deville and Särndal, 1992). A compact form for the asymptotic variance of the generalised regression estimator can be found e.g. in (Merkouris, 2010). Alternative jackknife esti-

mates of variance can be used in a more general context and were shown to outperform Taylor-based techniques for estimating the variance of calibration estimators in (Stukel et al., 1996). This paper will accordingly adopt a non-asymptotic jackknife approach to estimate the variance of calibration estimators, at the expense of a higher computational burden.

2.3 Range-restricted and penalized calibration

Several iterative methods have been used to solve the survey calibration problem (1), when including (RR), for various distance functions (Singh and Mohl, 1996). As explained in the Introduction, there can be many sources for non-convergence, not limited to a lack of representativeness in the survey, which was further discussed in Section 2.1, specially when considering RR in addition to BC. This is usually addressed by using heuristics that either modify the BC and/or RR or allow some error in (BC), but however do not perform any optimal control of that error.

Penalized calibration instead searches for weights satisfying (RR) while having a parametric control on the error in (BC). This is done via the minimisation of

$$\mathcal{G}_d^{GREG}(\mathbf{w}) + (\mathbf{X}'\mathbf{w} - \mathbf{t}_x)' \mathbf{\Lambda}^{-1} (\mathbf{X}'\mathbf{w} - \mathbf{t}_x) , \quad (6)$$

where $\mathbf{\Lambda} = \text{diag}(\boldsymbol{\lambda})$ is a diagonal matrix depending on parameters $\boldsymbol{\lambda} = (\lambda_j)$. The smallest possible values for these parameters are iteratively searched for in (Rao and Singh, 1997), where in fact these are obtained as a function of user-specified tolerances on the errors in (BC). Although this approach is shown to reduce the discrepancy in respecting (BC) for given (RR), its dependence on parameters used to control for errors in (BC) is critical for convergence. See e.g. Théberge (2000) for a closed-form solution to the problem, and (Beaumont and Bocci, 2008) and Section 9 in (Fuller, 2002) for closely related model-based "ridge regression" approaches.

Similarly, in (Wagner, 2013) a vector of unknowns $\boldsymbol{\varepsilon}_{\mathcal{B}}$ was used to model the multiplicative error in part of the benchmark totals so that the corresponding subset \mathcal{B} of BC are satisfied: $\mathbf{X}'_{\mathcal{B}}\mathbf{w} = \text{diag}(\mathbf{t}_{x,\mathcal{B}})\boldsymbol{\varepsilon}_{\mathcal{B}}$. The generalised regression distance \mathcal{G}_d^{GREG} was minimised in combination with the squared discrepancy between $\boldsymbol{\varepsilon}_{\mathcal{B}}$ and $\mathbf{1}$, a vector of 1s, weighted by user-defined parameters $\boldsymbol{\delta}$. Interestingly, a constant named *Gelman bound* (κ_{GB}) was introduced to control for the ratio of the largest to the smallest calibrated weight. Finally, RR on the weights \mathbf{w} and errors $\boldsymbol{\varepsilon}_{\mathcal{B}}$ were allowed:

$$\begin{aligned} \arg \min_{\mathbf{w}, \boldsymbol{\varepsilon}_{\mathcal{B}}, \alpha, \beta} \quad & \mathcal{G}_d^{GREG}(\mathbf{w}) + (\boldsymbol{\varepsilon}_{\mathcal{B}} - \mathbf{1})' \text{diag}(\boldsymbol{\delta}) (\boldsymbol{\varepsilon}_{\mathcal{B}} - \mathbf{1}) \\ \text{s.t.} \quad & \mathbf{X}'_{\mathcal{A}}\mathbf{w} - \mathbf{t}_{x,\mathcal{A}} = \mathbf{0}, \quad \mathbf{X}'_{\mathcal{B}}\mathbf{w} - \text{diag}(\mathbf{t}_{x,\mathcal{B}})\boldsymbol{\varepsilon}_{\mathcal{B}} = \mathbf{0}, \\ & \alpha \leq \mathbf{w} \leq \beta, \quad -\kappa_{GB}\alpha + \beta \leq 0, \\ & \mathbf{l} \leq \mathbf{w} \leq \mathbf{u}, \quad \mathbf{l}_{\mathcal{B}} \leq \boldsymbol{\varepsilon}_{\mathcal{B}} \leq \mathbf{u}_{\mathcal{B}}. \end{aligned} \quad (7)$$

In a simulation with $\kappa_{\text{GB}} = 35$, $\delta \equiv 1000$, convergence problems arose even if allowing any value for the errors $\varepsilon_{\mathcal{B}}$ in some BC, and increased significantly (14% to 23% failure) when imposing actual bounds on those errors.

3 Optimal control for RR and BC to allow successful range-restricted calibration

All existing methods addressing the range-restricted survey calibration problem (1) run into non-convergence issues or lack of control of the errors in BC, even if using penalization formulations like (6) that theoretically allow for the minimum error in BC. The usual approach consists in running a calibration method and at the end (after a long running time), if encountering non-convergence, require an user to adjust the RR and/or BC. We propose instead to assess if the given values for RR and BC allow the existence for a solution (*feasibility*), and allows the computation of optimal alternative values that guarantee the feasibility in case of foreseen non-convergence, given user-specified tolerances on changes in RR and/or errors in BC.

The natural questions that we will address in this section are ⁵, given RR vectors \mathbf{l}, \mathbf{u} : is problem (1) feasible? is it feasible if allowing a certain error ε in BC? in fact, what is the minimum error that needs to be allowed in BC to achieve feasibility? alternatively, what is the smallest change in RR that we need to perform to obtain feasibility (even if possibly allowing for some error ε in BC)?

3.1 An introductory example

Before entering into details, let us inspect the previous questions in a very simple scenario with $n = 100$ individuals in a sample with initial weights $\mathbf{d} \equiv 20$ to be calibrated using one known benchmark constraint BC:

$$w_1 + \dots + w_{100} = 2016 \tag{8}$$

If we impose as RR the positivity of weights $\mathbf{0} \leq \mathbf{w}$, the BC and RR can be satisfied simultaneously, e.g. by setting all weights equal to 20.16, or by setting 99 weights to 20 and just one weight to 36 (the calibration solution will depend on the distance function used to measure changes in initial weights). However, if we impose that $\mathbf{0} \leq w \leq \mathbf{20}$ as RR on weights in order to avoid any survey unit to represent more than 20 total units, the BC (8) cannot be satisfied thus the combination of BC and RR are incompatible. In that case, if we allow a small arbitrary error of 100 in (8), the problem becomes feasible by taking all weights equal to 20. In doing so, we obtain an error in BC equal to 16, which is in fact the minimum needed for compatibility with

⁵If setting to $\pm\infty$ either of the RR vectors, we obtain the same questions for the other RR vector.

the given RR. Alternatively, the upper bound on weights can be set to **20.16** (or any higher value) in order to have feasibility given the original BC.

3.2 Can the problem be solved?

The existence of a solution (*feasibility*) for the range-restricted calibration problem (1) is equivalent to the existence of solutions for the set of constraints:

$$\begin{cases} \mathbf{X}'\mathbf{w} - \mathbf{t}_x = \mathbf{0}, \\ \mathbf{1} \leq \mathbf{w} \leq \mathbf{u}. \end{cases} \quad (9)$$

If this set is non void, then it will be possible to find in it the vector of weights \mathbf{w} at minimum distance $\mathcal{G}_d(\mathbf{w})$ to the set of initial weights \mathbf{d} . By expressing the equality as two inequalities, the system (9) can be seen as a set of affine inequalities in the unknown \mathbf{w} , and therefore its feasibility could be checked using direct methods like system reduction by repeatedly using Fourier-Motzkin elimination (Dantzig and Eaves, 1973). An alternative algorithm for assessing the existence of a solution to the range-restricted calibration problem was developed in Théberge (2000). However, the existence of a solution can also be addressed by computing the minimum change needed in BC or RR for feasibility (following sections): if no change is needed, then the original problem is feasible; otherwise, the minimum needed change has already been computed.

3.3 Feasibility guarantee allowing minimum error in BC

The minimum *total absolute error* (TAE) in BC needed for their compatibility with the given RR is

$$\begin{aligned} TAE^* &= \min_{\mathbf{w}} \|\mathbf{X}'\mathbf{w} - \mathbf{t}_x\|_1 \\ \text{s.t. } &\mathbf{1} \leq \mathbf{w} \leq \mathbf{u}, \end{aligned} \quad (10)$$

being $\|\mathbf{v}\|_1$ the ℓ_1 -norm of a p -dimensional vector \mathbf{v} , defined by $\|\mathbf{v}\|_1 = \sum_{i=1}^p |v_i|$. The total absolute error has a very easy physical interpretation given that its units are the same as those of the population totals \mathbf{t}_x . The minimisation of the TAE error can be written as the minimisation of $\|\tilde{\boldsymbol{\varepsilon}}\|_1 = \mathbf{1}' \tilde{\boldsymbol{\varepsilon}} = \tilde{\varepsilon}_1 + \dots + \tilde{\varepsilon}_p$, for a non-negative vector $\tilde{\boldsymbol{\varepsilon}}$ such that $|\mathbf{X}'\mathbf{w} - \mathbf{t}_x| \leq \tilde{\boldsymbol{\varepsilon}}$ (component-wise). By decomposing the absolute value we obtain two vector inequalities⁶, and therefore:

$$\begin{aligned} TAE^* &= \min_{\mathbf{w}, \tilde{\boldsymbol{\varepsilon}}} \mathbf{1}' \tilde{\boldsymbol{\varepsilon}} \\ \text{s.t. } &\mathbf{X}'\mathbf{w} - \tilde{\boldsymbol{\varepsilon}} \leq \mathbf{t}_x, \quad -\mathbf{X}'\mathbf{w} - \tilde{\boldsymbol{\varepsilon}} \leq -\mathbf{t}_x, \\ &\mathbf{1} \leq \mathbf{w} \leq \mathbf{u}, \quad \tilde{\boldsymbol{\varepsilon}} \geq \mathbf{0}. \end{aligned} \quad (11)$$

Given that both the objective and all constraint functions are affine in the unknowns $(\mathbf{w}, \tilde{\boldsymbol{\varepsilon}})$, we have shown that finding the optimal TAE is a linear

⁶For any x, y real numbers, $|x| \leq y$ is equivalent to $x \leq y$ and $-x \leq y$

programming problem (Boyd and Vandenberghe, 2004). This class of convex optimization problems may be solved quickly and with global optimality convergence guaranteed by exploiting duality relations and optimality theorems (Boyd and Vandenberghe, 2004). If the solution is $TAE^* = 0$ it means that the calibration problem (1) is feasible; non-zero minimum TAE values require the modification of the original problem as proposed in Section 4.

3.4 Feasibility search by allowing minimum change in RR and user-specified error in BC

Assume now that we do not want to have a TAE error in BC greater than a scalar value ε (ideally equal to 0), but that we allow a small change in RR provided by two non-negative vectors $\boldsymbol{\lambda}, \boldsymbol{\mu}$, while keeping the weights inside a limiting range: $\mathbf{L} \leq \mathbf{w} \leq \mathbf{U}$. The smallest possible *total absolute change* (TAC) in RR that guarantees feasibility with TAE error below ε and final weights inside the maximum limiting range, if it exists, can be computed as

$$\begin{aligned} TAC^* = \min_{\boldsymbol{\lambda}, \boldsymbol{\mu}, \mathbf{w}} \quad & \|\boldsymbol{\lambda}\|_1 + \|\boldsymbol{\mu}\|_1 \\ \text{s.t.} \quad & \|\mathbf{X}'\mathbf{w} - \mathbf{t}_x\|_1 \leq \varepsilon, \\ & \mathbf{L} \leq \mathbf{1} - \boldsymbol{\lambda} \leq \mathbf{w} \leq \mathbf{u} + \boldsymbol{\mu} \leq \mathbf{U}, \boldsymbol{\lambda} \geq \mathbf{0}, \boldsymbol{\mu} \geq \mathbf{0}. \end{aligned} \quad (12)$$

The associated linear programming problem reads:

$$\begin{aligned} TAC^* = \min_{\boldsymbol{\lambda}, \boldsymbol{\mu}, \mathbf{w}, \tilde{\boldsymbol{\varepsilon}}} \quad & \mathbf{1}' \boldsymbol{\lambda} + \mathbf{1}' \boldsymbol{\mu} \\ \text{s.t.} \quad & \mathbf{X}'\mathbf{w} - \tilde{\boldsymbol{\varepsilon}} \leq \mathbf{t}_x, \quad -\mathbf{X}'\mathbf{w} - \tilde{\boldsymbol{\varepsilon}} \leq -\mathbf{t}_x, \quad \mathbf{1}' \tilde{\boldsymbol{\varepsilon}} \leq \varepsilon, \\ & -\mathbf{w} - \boldsymbol{\lambda} \leq -\mathbf{1}, \quad \mathbf{w} - \boldsymbol{\mu} \leq \mathbf{u}, \\ & \mathbf{0} \leq \boldsymbol{\lambda} \leq \mathbf{1} - \mathbf{L}, \quad \mathbf{0} \leq \boldsymbol{\mu} \leq \mathbf{U} - \mathbf{u}, \quad \tilde{\boldsymbol{\varepsilon}} \geq \mathbf{0}. \end{aligned} \quad (13)$$

This problem has the trivial solution $TAC^* = 0$ for values ε above or equal to TAE^* , given that TAE^* is the minimum error needed without modifying the RR. It may happen that the problem has no solution for a given ε smaller than TAE^* , for instance for $\varepsilon = 0$ with non-consistent BC, in which case intermediate values between ε and TAE^* could be explored for the allowed error in BC, given that the RR were non-trivial. If $TAC^* > 0$, then a modification of the original calibration problem (1) is required, as proposed in Section 4.

3.5 Alternative possibilities to model feasibility

The global optimality and simplicity of the previous approach are not affected if constant factors are introduced. For example, different relative weights \mathbf{r} may be assigned to different BC, by just replacing in (11) and (13) the expression $\mathbf{1}'\tilde{\boldsymbol{\varepsilon}}$ with $\mathbf{r}'\tilde{\boldsymbol{\varepsilon}}$. In particular, in the case that the benchmark values are provided from administrative totals of cross-tabulated variables, it is possible to normalize the TAE by dividing each benchmark constraint

by the relevant total of the corresponding administrative table, so that the global measure of error is a sum of comparable errors. Similarly, it is possible to divide the benchmark totals \mathbf{t}_x by the relevant benchmark table totals \mathbf{v} so that the calibration weights represent proportions, by just replacing \mathbf{t}_x with $\text{diag}(\mathbf{v})^{-1}\mathbf{t}_x$.

Different precision levels on BC can be also achieved, e.g. having some “exact” BC as in (Wagner, 2013) or some BC with a smaller penalization to errors as in (Rao and Singh, 1997), by setting in (11) and (13) the corresponding components of $\tilde{\boldsymbol{\varepsilon}}$ to the desired precision values. This option will be used in the experimental validation of the paper, where the exact BC will correspond to a broad cross-tabulation and a finer cross-tabulation will be used as inexact BC. If wanting to add a “Gelman” bound control to RR as in (Wagner, 2013), two scalar variables α, β and the same affine constraints as in (7) need to be added to the optimization programs: The modified problems remain linear given that these constraints are linear: $\alpha \leq \mathbf{w} \leq \beta$, $-\kappa_{GB}\alpha + \beta \leq 0$, $\alpha \geq 0$, $\beta \geq 0$. The problem in Section 3.4 can usually be simplified: if the lower bound \mathbf{I} has to be 0 (positivity of weights) then only the upper bound can be varied, which can be done by only estimating $\boldsymbol{\mu}$ while setting $\boldsymbol{\lambda}$ to zero; or if the changes can be the same for all RR, then only one parameter needs to be used for each of the increment vectors $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$.

The ideal choice of a penalization function should be based on distributional assumptions for errors in BC and desired changes in RR⁷, depending on the problem and available computational power. The proposed ℓ_1 -norm allows a linear programming implementation and is a robust penalization that in practice produces many very small residuals, allowing to identify many BC that can be satisfied exactly e.g. by looking at the components of $\tilde{\boldsymbol{\varepsilon}}$ in (11). A simpler approach can use the ℓ_∞ -norm to penalize errors (the maximum component in a vector being penalized), which is equivalent to considering as single-valued the unknown vectors $\tilde{\boldsymbol{\varepsilon}}$ in (11) and $\boldsymbol{\lambda}, \boldsymbol{\mu}$ in (13).

Using a ℓ_2 -norm or a weighted least squares penalization converts the feasibility programs into quadratic and quadratically constrained quadratic, which are solvable for less variables and are more time-consuming than linear programs (Boyd and Vandenberghe, 2004). In fact, Théberge proposed using a quadratic norm to allow minimum errors in BC for GREG without RR in (Théberge, 1999) (a closed-form solution exists). In (Théberge, 2000) he further formulated the problem with RR (Section 2) but adopted a not always convergent penalization-like formulation for its resolution (Section 4).

⁷Assuming independent identically distributed random errors in a linear system, the ℓ_1 -norm penalization gives the MLE for a Laplacian distribution of errors, whereas the ℓ_2 - and ℓ_∞ -norm penalizations give the MLE for Gaussian and uniform error distributions, respectively (Boyd and Vandenberghe, 2004).

4 Optimally modelling and solving range-restricted survey calibration

We have seen in Section 3 that the feasibility of the range-restricted calibration problem (1) has linear complexity, independently of the chosen distance function $\mathcal{G}_{\mathbf{d}}(\mathbf{w})$, and it can be achieved with equal complexity level by allowing optimally controlled errors in BC and/or changes in RR. In this section we propose Algorithm 1 to solve the range-restricted survey calibration problem (1), allowing for optimal modification(s) of the BC tolerance and/or RR bounds, only if needed and controlled by user-defined parameters, ε and δ respectively. We also discuss the choice of a distance function $\mathcal{G}_{\mathbf{d}}(\mathbf{w})$ in Algorithm 1, and focus on the particular Chi-square distance for demonstration.

For sufficiently high values of ε tolerance to error in BC (infinity in the extreme case), the calibration is performed in two steps: first, the TAE^* minimum value is computed; then, the calibration problem allowing TAE error in BC equal to TAE^* is solved optimally. The convergence of this two-step approach is guaranteed by construction while respecting the asymptotic design consistency. For values of ε smaller than TAE^* , a further step is performed trying to achieve error in BC below ε by modifying the RR with a minimum total absolute change TAC^* , given that this value is below the user-specified tolerance δ . In case the modification of RR cannot lead to a feasible problem, Algorithm 1 proposes to use $\varepsilon = TAE^*$, but an exploration for smaller values could be performed as explained before.

In Algorithm 1, the range-restricted calibration problem (1) is solved if it is feasible (line 3), and otherwise it is replaced by a problem of the form

$$\begin{aligned} \arg \min_{\mathbf{w}, \mathbf{y}} \quad & \mathcal{G}_{\mathbf{d}}(\mathbf{w}) \\ \text{s.t.} \quad & \mathbf{A} \begin{pmatrix} \mathbf{w} \\ \mathbf{y} \end{pmatrix} \leq \mathbf{a} , \\ & \mathbf{b} \leq \mathbf{w} \leq \mathbf{c} , \end{aligned} \tag{14}$$

for a matrix \mathbf{A} , vectors \mathbf{a} , \mathbf{b} , \mathbf{c} , and auxiliary variables \mathbf{y} defined by the algorithm. More specifically, the optimisation domain from (1) is modified in line 6 by replacing the BC with $\|\mathbf{X}'\mathbf{w} - \mathbf{t}_{\mathbf{x}}\|_1 \leq TAE^*$, and in line 10 additional constraints of similar nature are added. The resulting domains can be expressed using affine inequalities in the form (14) with the help of auxiliary variables, as done in Section 3. Therefore, the complexity of the resulting problems will be mainly associated to that of the distance function $\mathcal{G}_{\mathbf{d}}(\mathbf{w})$.

Under the assumptions of Section 2.1 for the distance function $\mathcal{G}_{\mathbf{d}}(\mathbf{w})$, the problems (14) are convex with smooth objective and therefore can be solved with global optimality convergence guarantees by exploiting duality relations and optimality theorems like the necessity and sufficiency of Karush–Kuhn–Tucker conditions (Boyd and Vandenberghe, 2004). The resolution can be done with fast convergence in particular cases; a closely related

Algorithm 1 Optimal calibration via optimal control for errors in BC and changes in RR

Require: $\{\mathbf{X}, \mathbf{t}_x, \mathbf{d}\}$ data, $\{\mathbf{l}, \mathbf{u}, \mathbf{L}, \mathbf{U}\}$ RR and bounds, $\mathcal{G}_d(\mathbf{w})$ distance function, ε max allowed error in BC, δ max allowed change in RR

- 1: Compute TAE* min total absolute error achievable in BC, using (11)
- 2: **if** TAE* is 0 (Original problem feasible) **then**
- 3: **return** \mathbf{w} = solution of (1), i.e.
 $\mathbf{w} = \arg \min_{\mathbf{w}} \mathcal{G}_d(\mathbf{w})$
 s.t. $\mathbf{X}'\mathbf{w} - \mathbf{t}_x = \mathbf{0}$ (BC)
 $\mathbf{l} \leq \mathbf{w} \leq \mathbf{u}$ (RR)
- 4: **else**
- 5: **if** TAE* $\leq \varepsilon$ **then**
- 6: **return** \mathbf{w} = sol. of (1) replacing BC with $\|\mathbf{X}'\mathbf{w} - \mathbf{t}_x\|_1 \leq \text{TAE}^*$
- 7: **else**
- 8: Search TAC* min total absolute change needed in RR to have a TAE error in BC below ε , using (13)
- 9: **if** TAC* exists **and** TAC* $\leq \delta$ **then**
- 10: **return** \mathbf{w} = sol. of (1) replacing BC with $\|\mathbf{X}'\mathbf{w} - \mathbf{t}_x\|_1 \leq \varepsilon$ and replacing RR with $\|\boldsymbol{\lambda}\|_1 + \|\boldsymbol{\mu}\|_1 \leq \text{TAC}^*$,
 $\mathbf{L} \leq \mathbf{l} - \boldsymbol{\lambda} \leq \mathbf{w} \leq \mathbf{u} + \boldsymbol{\mu} \leq \mathbf{U}$, $\boldsymbol{\lambda} \geq \mathbf{0}$, $\boldsymbol{\mu} \geq \mathbf{0}$
- 11: **else**
- 12: **error** the problem cannot be solved for the given data and parameters; go to line 6
- 13: **end if**
- 14: **end if**
- 15: **end if**

example is the semismooth Newton method proposed in (Wagner, 2013) for the Chi-square distance (2) and the raking distance (4). Note that despite its possible efficient minimisation, a ℓ_1 distance function is not suitable since it would allow a few weights being very distant from initial ones, which could undesirably cause a high bias in calibration estimators. Rather than developing resolution methods for different distances, this paper has focussed on developing a flexible always-convergent optimal calibration framework, which is exemplified by adapting the range-restricted generalised regression (GREG) estimator.

The range-restricted calibration problem (1) for the Chi-square distance (2) was identified as a quadratic programming problem in (Isaki et al., 2000) and its fast optimal resolution exploiting duality principles was addressed recently (Wagner, 2013), however its feasibility has not yet been guaranteed by any method. If using Algorithm 1 for this purpose, the resulting modified problems (14) are inequality-constrained quadratic convex optimization programming problems. These can be solved in polynomial time, and in practice

relatively quickly, while assessing the global optimality of the solution (Boyd and Vandenberghe, 2004).

5 Evaluation

The proposed methods are demonstrated and validated using real datasets: the Health Survey for England 2012 (HSE) and the 2011 Census in England and Wales (CEW). The HSE is representative of the English population living in private households (Craig and Mindell, 2012), and it is drawn in 2011 using a multi-stage stratified sampling approach. Available survey weights adjust for selection, non-response, and population age/gender and strategic health authority region profiles. CEW tables DC1104EW and DC1602EWLA provide population total counts for non-institutional residents in England.

The initial survey sample for the experiments consisted of 10,308 individuals from HSE 2012⁸. In all experiments, census-based population totals for ten age groups⁹ cross-tabulated with gender and population totals for the nine regions in England (20+9 counts for a population of size 52,059,931) were imposed as exact constraints, and the positivity of calibration weights was imposed as part of range restrictions (RR) on weights. An additional BC was imposed with 378 population totals for 21 age groups¹⁰ cross-tabulated with gender and region. This fine cross-tabulation was not available at the time of release of the HSE data. Population totals for five age groups¹¹ by four economic activity groups (in-employment, ILO unemployed, retired, and other inactive) were used for validation purposes. The experiments did not use any continuous benchmarks for the sake of simplicity, but benchmarks on continuous data, e.g. average age per region, could be incorporated.

Independently of any RR choice (or the initial sample weights), the fine age/gender/region cross-tabulation with 378 group totals cannot be satisfied exactly by calibration weights, given that the sample has a zero count for one group¹². Further motivated by the presence of some small counts, traditional calibration would only impose a broad cross-tabulation on the survey weights like the 29 age/gender and region counts that we will impose exactly in all experiments. However, the fine age/gender/region counts provide a much richer picture of the joint distribution of those variables, and the original weights are distant from correctly representing that picture: the total absolute error (TAE) of the HSE 2012 weights for that BC is of 7,858,083 units (a 15.14% of the total population size). Given that the cross-tabulation

⁸25 adults not having a valid economic activity were discarded.

⁹Ages 0-4, 5-9, 10-15, 16-24, 25-34, 35-44, 45-54, 55-64, 65-74, 75+.

¹⁰Ages 0-4, 5-7, 8-9, 10-14, 15, 16-17, 18-19, 20-24, 25-29, 30-34, 35-39, 40-44, 45-49, 50-54, 55-59, 60-64, 65-69, 70-74, 75-79, 80-84, 85+.

¹¹Ages 16-24,25-34,35-49,50-64,65+

¹²The HSE 2012 sample does not contain females in the North East aged 15.

is categorical, the total absolute error is counting the number of individuals wrongly assigned to each cross-tabulation group. Since we are imposing exact broad age/gender population counts, the estimated population total remains fixed. Thus half of the TAE is the number of individuals being misclassified, which for the HSE estimate is 7.5% of the English population.

Instead of ignoring the fine age/gender/region cross-tabulation, it can be used to calibrate the sample weights if we allow some error in BC (which arises as a natural need for the given data), consequently resulting in better estimates on age/gender/region related variables. In order to allow a direct comparison of the gain in adding this strategy to the traditional approach, we imposed as exact the already described broad age/gender and region cross-tabulations. We performed three validation experiments returning positive weights \mathbf{w} at minimum Chi-square distance (2) to the initial sample weights:

Ex1. Minimum TAE error in BC;

Ex2. Minimum TAE error in BC and $0.5 \leq \mathbf{w} \leq 3.5$ as RR on weights;

Ex3. Minimum TAC change in the previous RR so that $\text{TAE} \leq 0.1\%$.

The corresponding optimisation programs are summarised in Algorithm 1.

For all experiments and for the original HSE data we show in Table 1 the modified Chi-squared distance (2) between the HSE and each considered set of weights and provide descriptive statistics of the latter. Jackknife standard deviation (SD) estimates for the BC and validation count estimates were obtained using 94 groups of primary sampling units (PSUs), built-up by deleting six PSUs at each time as described in (Kott, 1998). The error measures used for estimated counts were: TAE the total absolute error, the TAE as a percentage of the cross-tabulation, TRE the total relative error (sum of relative errors over all group counts, as a %), and RMSE the root mean square error. Table 2 contains the average SD and fitting errors for the fine age/gender/region BC cross-tabulation estimates, and Table 3 the average SD and errors for the age by economic activity validation count estimates.

The original HSE weights (HSE, first row in all tables) do not present very extreme values, the highest ratio between weights being 18.32 (Table 1). As already explained, the HSE weights perform badly in estimating the fine age/gender/region distribution in England: 15% of total absolute error, and 6,829.8% of total relative error, with an average SD equal to 0.95% of the total population (Table 2). The broad age by economic activity cross-tabulation is quite well estimated by HSE with TAE error below 5% (Table 3). This is possibly in part because the broad categorisation of age is similar to the age categorisation which the HSE weights adjust for.

The Ex1 calibration weights (Ex1, second row in all tables) are at average Chi-square distance of 0.5 to the HSE weights, and have slightly more

extreme values, the highest ratio between weights being 23.31 (Table 1). The resulting Ex1 age/gender/region BC count estimates have much smaller SD than the HSE estimates, the proposed method being therefore more efficient, and have by construction a very small bias: the minimum possible TAE error in BC (Table 2). All considered SD and error indicators indicate consistently an improvement in performance when applying the Ex1 calibration weights to estimate the non-fitted validation totals (Table 3).

Experiment Ex2 (third row in all tables) provides an example of a practice commonly followed by practitioners and found in the literature, see e.g. (Singh and Mohl, 1996; Stukel et al., 1996), consisting in the arbitrary selection of range restriction values for the calibration weights and *a posteriori* observation of errors (in case of convergence). The Ex2 selected RR values result in a (user-defined) low dispersion in Ex2 weights (highest ratio between weights being 7), which computation required a minimum TAE error in BC of 0.65% for convergence. As a result the average deviances and all the fitting errors and validation errors are higher for Ex2 than those for the Ex1 weights (Tables 2 and 3). So far we have seen that Ex1 provided an optimal fit of the BC but at the expense of slightly more extreme weights, and also that an arbitrary choice of RR on weights in Ex2 achieved more centred weights at expense of increasing the SD and the (minimum) errors in both the BC and validation estimates.

It would certainly be time consuming to perform an exploration of possible values for RR in order to obtain satisfactory weights with non-extreme values and low SD and low (minimum) fitting errors for the BC. Instead, Ex3 (fourth row in all tables) searches for the minimum change in provided initial values for RR at expense of allowing a (user-specified) 0.1% TAE error in fitting the fine age/gender/region BC counts. Compared with Ex2, both fitting and validation errors were smaller for Ex3. Compared with Ex1, Ex3 resulted in a set of weights with less extreme values, the highest ratio between weights being 14.00, and an small increase in (controlled) bias and SD in fitting the BC. Nonetheless, Ex3 provided (slightly) better estimates of the validation counts, pointing at a possible dangerous over-fitting effect if using the Ex1 approach: fitting too-closely a fine cross-tabulation (having small counts) may well increase the bias and variance in estimation for non-fitted variables. However, in the case considered here, this effect was tiny compared to the efficiency gain and bias reduction with respect to estimates obtained using the initial HSE weights.

Overall, the three experiments Ex1-Ex3 used Algorithm 1 to minimally modify the HSE weights to adjust them to the fine age/gender/region BC cross-tabulation totals, overcoming the fact that the survey sample had small and even zero counts for that cross-tabulation. This was done by allowing for a minimum error in BC, which can be seen as equivalent to clustering some benchmark groups. Thus the experiments optimally improved a practice often followed arbitrarily to avoid non-convergence. In all experiments, not

	Chi-sq	Min	Q1	Median	Q3	Max	Max/Min
HSE	0	0.36	0.77	0.91	1.14	6.67	18.32
Ex1	570.3	0.31	0.73	0.90	1.14	7.26	23.31
Ex2	524.6	0.50	0.73	0.90	1.14	3.50	7.00
Ex3	549.1	0.36	0.73	0.90	1.14	5.04	14.00

Table 1: Chi-square distance and distribution statistics for the HSE and obtained calibration weights.

	SD	TAE	TAE (%)	TRE (%)	RMSE
HSE	205,835.4	7,858,083.0	15.09	6,829.8	27,415.5
Ex1	3,091.3	29,678.0	0.06	121.0	1,079.4
Ex2	11,298.3	339,509.2	0.65	604.1	3,227.8
Ex3	3,525.2	52,059.9	0.10	192.3	1,177.1

Table 2: Age by gender by region cross-tabulation estimates (378 counts): average standard deviation (SD) over all estimates, and fitting errors.

only was the fitted BC age/gender/region distribution much better approximated than with the original HSE weights, but also efficiency and performance improved when estimating age by economic activity validation total counts.

6 Summary and conclusion

This paper has presented a two-step global optimisation (GO) approach to design-based survey calibration with guaranteed convergence, allowing for range-restrictions on weights while controlling for those range-restrictions and the (minimum) error in benchmark constraints.

First, GO assesses the feasibility of the range-restricted calibration problem, with infeasible problems being transformed into feasible ones by allowing minimal errors in the benchmark constraints (BC) and/or minimal changes in the weights' range restrictions (RR). For this purpose, GO identifies the minimum achievable difference between the calibrated (reweighted) survey and the benchmark totals, taking into account any RR specified for the solution weights. It also identifies the minimum needed change in those RR, allowing exploration of an alternative solution, more respectful of the original problem, having zero or below-minimum error in BC (in general, the existence of solution is only guaranteed if allowing the minimum error in BC). All the problems involved in this first step assessing/guaranteeing feasibility are modelled using the robust ℓ_1 -norm penalization (ℓ_∞ - and ℓ_2 -norm alternatives, as well as weighted versions, were discussed in the text) and as a result can be solved efficiently using sparse linear programming. Second, the GO approach applies global optimisation techniques for minimising the change

	SD	TAE	TAE (%)	TRE (%)	RMSE
HSE	1,027,677.0	2,059,120.9	4.89	633.1	165,721.7
Ex1	861,528.8	1,763,839.5	4.19	494.1	130,193.5
Ex2	869,717.6	1,786,060.9	4.24	497.3	130,407.6
Ex3	856,300.1	1,749,775.5	4.16	492.8	129,310.1

Table 3: Age by economic activity estimates (20 counts): average standard deviation (SD) over all estimates, and validation errors.

in weights subject to allowing only the minimum error in BC or change in RR required for feasibility (already computed in the previous step). The approach has been theoretically exemplified with the Chi-square distance being used to measure the change in weights with respect to initial (design) ones. Other distances have been considered, for which modern optimisation techniques will be useful to solve the resulting calibration problems.

The first step to assess/achieve feasibility represents an efficient modelling alternative to the current approaches in which convergence is known only after running a calibration method (time costly) and the reasons for non-convergence are not always clear. Moreover, existing approaches either make use of heuristics after encountering non-convergence, which do not offer enough control on the solution, or require user-defined parameters to model infeasibility, which in practice may not avoid non-convergence.

For survey calibration problems where the BC can be met, GO will provide a solution equivalent to that produced by calibration methods that allow RR on weights (assuming the chosen number of iterations in iterative methods poses no limit to convergence). GO-based estimators preserve the good properties of survey calibration estimators, design consistency and asymptotic design-unbiasedness, while adding guaranteed convergence and global optimality. In a real-data experiment we showed a double-win situation (gain in both bias and variance), achieved through two-level calibration: broad group cross-tabulations were imposed exactly, whereas a small group cross-tabulation (leading to zero counts in the survey) was managed optimally using the proposed approach.

Acknowledgements

The work reported in this paper was funded through the ESRC Secondary Data Analysis Initiative grant ES/K004433/1 and the ESRC Business and Local Government Data Research Centre grant ES/L011859/1.

References

- Ballas, D. and Clarke, G. P. (2009) Spatial microsimulation. In *Handbook of spatial analysis* (eds. A. Fotheringham and P. Rogerson), 277–298. London: Sage.
- Bankier, M. D., Rathwell, S. and Majkowski, M. (1992) Two step generalized least squares estimation in the 1991 Canadian census. In *JSM Proc., Surv. Res. Method. Sect. ASA*.
- Beaumont, J. F. and Bocci, C. (2008) Another look at ridge calibration. *METRON – Int. J. Stat.*, **LXVI**, 5–20.
- Bishop, Y. M. M. (1969) Full contingency tables, logits and split contingency tables. *Biometrics*, **25**, 119–128.
- Boyd, S. and Vandenberghe, L. (2004) *Convex optimisation*. Cambridge University.
- Craig, R. and Mindell, J. (2012) *Health Survey for England. Methods and Documentaiton*. The Health and Social Care Information Centre.
- Dantzig, G. and Eaves, B. (1973) Fourier-Motzkin elimination and its dual. *J. Combin. Theory Ser. A*, **14**, 288–297.
- Deming, W. E. and Stephan, F. F. (1940) On a least square adjustment of a sampled frequency table when the expected marginal totals are known. *Ann. Math. Stat.*, 427–444.
- Deville, J.-C. and Särndal, C.-E. (1992) Calibration estimators in survey sampling. *J. Am. Stat. Assoc.*, **87**, 376–382.
- Fuller, W. A. (2002) Regression estimation for survey samples. *Surv. Methodol.*, **28**, 5–23.
- Horvitz, D. and Thompson, D. (1952) A generalization of sampling without replacement from a finite universe. *J. Am. Stat. Assoc.*, **47**, 663–685.
- Isaki, C. T., Ikeda, M. M., Tsay, J. H. and Fuller, W. A. (2000) An estimation file that incorporates auxiliary information. *J. Off. Stat.*, **16**, 155–172.
- Kott, P. (1998) Using the delete-a-group jackknife variance estimator in practice. In *JSM Proc., Surv. Res. Method. Sect. ASA*.
- (2009) *Calibration weighting: combining probability samples and linear prediction models*, vol. 29B of *Handbook of Statistics – Sample Surveys: Inference and Analysis*, chap. 25, 55–82. Elsevier.

- Merkouris, T. (2010) Combining information from multiple surveys by using regression for efficient small domain estimation. *J. R. Stat. Soc., Ser. B, Stat. Methodol.*, **72**, 27–48.
- Pudney, S. and Sutherland, H. (1994) How reliable are microsimulation results? an analysis of the role of sampling error in a UK tax-benefit model. *J. Public Econ.*, **53**, 327–365.
- Rao, J. N. K. and Singh, A. C. (1997) A ridge-shrinkage method for range restricted weight calibration in survey sampling. In *JSM Proc., Surv. Res. Method. Sect. ASA*.
- Sautory, O. (1991) *La macro CALMAR. Redressement d'échantillons d'enquêtes auprès des ménages par calage sur marges*. INSEE. F9310.
- Singh, A. C. and Mohl, C. A. (1996) Understanding calibration estimators in survey sampling. *Surv. Methodol.*, **22**, 107–115.
- Stukel, D. M., Hidiroglou, M. A. and Särndal, C.-E. (1996) Variance estimation for calibration estimators: A comparison of Jackknifing versus Taylor linearization. *Surv. Methodol.*, **22**, 117–125.
- Särndal, C.-E. (1978) Design-based and model-based inference in survey sampling. *Scan. J. Statist.*, **5**, 27–52.
- (2007) The calibration approach in survey theory and practice. *Surv. Methodol.*, **33**, 99–119.
- Tanton, R. (2014) A review of spatial microsimulation methods. *Int. J. Microsim.*, **7**, 4–25.
- Tanton, R., Vidyattama, Y., Nepal, B. and McNamara, J. (2011) Small area estimation using a reweighting algorithm. *J. R. Stat. Soc., Ser. A, Stat. Soc.*, **174**, 931–951.
- Théberge, A. (1999) Extensions of calibration estimators in survey sampling. *J. Am. Stat. Assoc.*, **94**, 635–644.
- (2000) Calibration and restricted weights. *Surv. Methodol.*, **26**, 99–107.
- Wagner, M. (2013) *Numerical optimisation in survey statistics*. Ph.D. thesis, Universität Trier.
- Wittenberg, R., Hu, B., Hancock, R., Morciano, M., Comas-Herrera, A., Malley, J. and King, D. (2011) Projections of demand for and costs of social care for older people in England, 2010 to 2030, under current and alternative funding systems. PSSRU 2811/2, University of Kent.
- Wu, C. and Lu, W. W. (2016) Calibration weighting methods for complex surveys. *Int. Stat. Rev.*, **84**, 79–98.