

Best Linear Unbiased Estimate from Privatized Contingency Tables

Jordan Awan

JAWAN@PURDUE.EDU

Department of Statistics

Purdue University

West Lafayette, IN 47907, USA

Adam Edwards

AMEDWARDS@MITRE.ORG

The MITRE Corporation

McLean, VA 22102, USA

Paul Bartholomew

PBARTHOLOMEW@MITRE.ORG

The MITRE Corporation

McLean, VA 22102, USA

Andrew Sillers

ASILLERS@MITRE.ORG

The MITRE Corporation

McLean, VA 22102, USA

Editor: Pierre Alquier

Abstract

In differential privacy (DP) mechanisms, it can be beneficial to release “redundant” outputs, where some quantities can be estimated in multiple ways by combining different privatized values. Indeed, the DP 2020 Decennial Census products published by the U.S. Census Bureau consist of such redundant noisy counts. When redundancy is present, the DP output can be improved by enforcing self-consistency (i.e., estimators obtained using different noisy counts result in the same value), and we show that the minimum variance processing is a linear projection. However, standard projection algorithms require excessive computation and memory, making them impractical for large-scale applications such as the Decennial Census. We propose the Scalable Efficient Algorithm for Best Linear Unbiased Estimate (SEA BLUE), based on a two-step process of aggregation and differencing that 1) enforces self-consistency through a linear and unbiased procedure, 2) is computationally and memory efficient, 3) achieves the minimum variance solution under certain structural assumptions, and 4) is empirically shown to be robust to violations of these structural assumptions. We propose three methods of calculating confidence intervals from our estimates, under various assumptions. Finally, we apply SEA BLUE to two 2010 Census demonstration products, illustrating its scalability and validity.

Keywords: differential privacy, projection, constrained optimization, confidence interval, U.S. Census Bureau

-
- . Jordan Awan and Adam Edwards are co-first authors and contributed equally to this work.
 - . Approved for Public Release; Distribution Unlimited. Public Release Case Number 24-2176.

1. Introduction

Differential privacy (DP) is a formal privacy framework, proposed by Dwork et al. (2006), which requires introducing randomness into a statistical procedure to provide provable privacy protection. The U.S. Census Bureau has employed differential privacy techniques to protect the 2020 Decennial Census (Abowd et al., 2019), which results in uncertainty about the true counts. The differential privacy mechanism used by the U.S. Census Bureau in 2020 produces millions of noisy counts which are “redundant” in that the same values are counted at various levels of geography (e.g., block, county, state) as well as different levels of detail (e.g., different k -way contingency tables). However, because independent noise is added to each count, these values are not *self-consistent*, meaning that detailed noisy counts do not add up to marginal noisy counts (see Figure 1 for an example). The lack of self-consistency is important because different combinations of the noisy counts can give different estimates of the same quantity, which can be a problem when a single “official” estimate is needed. Furthermore, the existence of multiple independent estimates implies that each individual estimate is sub-optimal, whereas the collection of estimates could be aggregated to produce a single refined estimate.

To process the noisy counts, the U.S. Census Bureau applied the *TopDown Algorithm*, which enforces self-consistency, non-negativity, and integer-valued constraints. While their method substantially improves the average error, it also introduces bias in the DP estimates, and as a non-linear procedure it is difficult to quantify the uncertainty in the resulting estimates (Santos-Lozada et al., 2020; Kenny et al., 2021; Winkler et al., 2021). As an alternative, we propose a linear procedure which enforces self-consistency but allows negative and non-integer-valued estimates. This approach has the following benefits: 1) the procedure results in unbiased estimates, and 2) because it is linear, we can accurately estimate the error and distribution of the final estimates to produce confidence intervals for the true counts.

We first show that the best linear unbiased estimator (BLUE) for the true counts assuming only self-consistency is a linear projection which can be expressed in terms of the self-consistency constraints. This result follows from the Gauss-Markov theorem and agrees with similar prior results in the DP literature (e.g., Gao et al., 2022). However, implementing this projection by standard techniques requires inverting a large matrix, resulting in a high computational and memory cost. This makes this approach inapplicable for large applications, such as the Demographic and Housing Characteristics (DHC) Census product.

To address this limitation, we propose a novel method that we call the Scalable Efficient Algorithm for the Best Linear Unbiased Estimator (SEA BLUE), which consists of a two step process of aggregation and differencing. This procedure is inspired by the up and down passes of Hay et al. (2010); Honaker (2015); Kifer (2021), which process noisy values in a tree structure. However, in our setting detailed counts are constrained to be self-consistent with multiple marginal counts, so we do not have a tree, but a much more complex structure. This makes the identification of the analogous process non-trivial. Our procedure:

1. enforces self-consistency through a linear and unbiased procedure,
2. is computationally and memory efficient,
3. produces the best linear unbiased estimate under certain structural assumptions, and

4. is empirically demonstrated to be robust to mild violations of the assumptions.

Using the linear nature of our procedure, we propose three methods of producing confidence intervals for the original counts. Two of these methods calculate the resulting standard error of the estimates either exactly or via Monte Carlo and produce confidence intervals using a normality assumption. The third method is a distribution-free Monte Carlo method, which gives valid confidence intervals which are valid for all noise distributions.

We illustrate our method in simulation studies, comparing mean-squared error (MSE), confidence interval coverage/width, and computational time/memory. We also apply our methodology to two 2010 Census demonstration products, Redistricting Data (Public Law 94-171, which we abbreviate as PL94) and Demographic and Housing Characteristics (DHC), illustrating the scalability and validity of our methods. Proofs and technical details are found in the appendix.

While our method is motivated by the 2020 Census products, it can be applied in other settings where independent noisy counts are observed at varying levels of detail.

1.1 Related Work

The first processing of redundant DP measurements was due to Hay et al. (2010), which leveraged a binary tree structure of the different true values. Their method used a two stage up and down pass to sequentially update the estimates and finally resulted in minimum variance estimates that are self-consistent. A similar approach has been considered by Honaker (2015) and Kifer (2021), and this method has been applied to derive DP confidence intervals for the median of real-valued data (Drechsler et al., 2022). Recently Cumings-Menon (2024) extended this tree-based approach to process the Census noisy measurement files across geographies, which forms a tree. In contrast, SEA BLUE is designed to optimize estimates within geographies using the structure of different marginal counts, which does not form a tree.

McCartan et al. (2023) and Kenny et al. (2024) proposed a minimum variance weighting of estimates which can be summed up from other tables, which is the same as our collection step, which forms the first step of our SEA BLUE procedure. Kenny et al. (2024) compare these unbiased weighted estimates to the TopDown estimates as well as the estimates from swapping, the privacy protection algorithm used in 2010. While the collection step estimate is optimal “from below,” as we formally prove in Section 4.1, it is not the best linear unbiased estimate as it does not use all information available in the tables. We show in Section 4.2 that our down pass is needed after the collection step to obtain the BLUE.

2. Background and Notation

In this section, we review some basic results in linear algebra and differential privacy, and set the notation for the paper—especially the notation used for contingency tables.

2.1 Linear Algebra

In \mathbb{R}^n the standard inner product is $\langle x, y \rangle = x^\top y$. To account for covariance structure, other inner products also exist: for a square and positive definite matrix Σ , an inner product is $\langle x, y \rangle_\Sigma = x^\top \Sigma y$, which has corresponding norm $\|x\|_\Sigma = \sqrt{\langle x, x \rangle_\Sigma}$.

$\mathbf{Y}_{\mathbf{I}\mathbf{J}}^{\mathbf{A},\mathbf{B}}$	$\mathbf{B}=1$ $\mathbf{B}=2$		$\mathbf{Y}_{\mathbf{I}\cdot}^{\mathbf{A},\mathbf{B}}$	$\mathbf{Y}_{\mathbf{I}\mathbf{J}}^{\mathbf{A},\mathbf{B}}$	$\mathbf{B}=1$ $\mathbf{B}=2$		$\hat{\mathbf{Y}}_{\mathbf{I}\mathbf{J}}^{\mathbf{A},\mathbf{B}}$	$\hat{\mathbf{Y}}_{\mathbf{I}}^{\mathbf{A},\mathbf{B}}$	$\hat{\mathbf{Y}}_{\mathbf{I}\mathbf{J}}^{\mathbf{A},\mathbf{B}}$	$\mathbf{B}=1$ $\mathbf{B}=2$		$\hat{\mathbf{Y}}_{\mathbf{I}\cdot}^{\mathbf{A},\mathbf{B}}$	$\hat{\mathbf{Y}}_{\mathbf{I}}^{\mathbf{A},\mathbf{B}}$
$\mathbf{A}=1$	12	4	16	$\mathbf{A}=1$	13.12	2.28	15.40	17.10	$\mathbf{A}=1$	13.18	2.54	15.72	$\mathbf{A}=1$
$\mathbf{A}=2$	0	0	0	$\mathbf{A}=2$	-1.37	0.94	-0.43	-0.61	$\mathbf{A}=2$	-1.78	1.45	-0.33	$\mathbf{A}=2$
$\mathbf{Y}_{\cdot\mathbf{J}}^{\mathbf{A},\mathbf{B}}$			$\mathbf{Y}_{\cdot\cdot}^{\mathbf{A},\mathbf{B}}$	$\hat{\mathbf{Y}}_{\cdot\mathbf{J}}^{\mathbf{A},\mathbf{B}}$			$\hat{\mathbf{Y}}_{\cdot\cdot}^{\mathbf{A},\mathbf{B}}$	$\hat{\mathbf{Y}}_{\cdot}^{\mathbf{B}}$	$\hat{\mathbf{Y}}_{\cdot\mathbf{J}}^{\mathbf{A},\mathbf{B}}$			$\hat{\mathbf{Y}}_{\cdot\cdot}^{\mathbf{A},\mathbf{B}}$	$\hat{\mathbf{Y}}_{\cdot}^{\mathbf{B}}$
	12	4	16		11.75	3.22	14.97	16.49		11.40	3.99	15.39	
				$\hat{\mathbf{Y}}_{\cdot}^{\mathbf{B}}$			$\hat{\mathbf{Y}}_{\cdot}^{\mathbf{A}}$	$\hat{\mathbf{Y}}_{\cdot}^{\emptyset}$	$\hat{\mathbf{Y}}_{\cdot\mathbf{J}}^{\mathbf{A},\mathbf{B}} = \hat{\mathbf{Y}}_{\mathbf{J}}^{\mathbf{B}}$			$\hat{\mathbf{Y}}_{\cdot\cdot}^{\mathbf{A},\mathbf{B}} = \hat{\mathbf{Y}}_{\cdot}^{\mathbf{A}} = \hat{\mathbf{Y}}_{\cdot}^{\mathbf{B}} = \hat{\mathbf{Y}}^{\emptyset}$	
					11.64	5.63	17.27	16.89					

Figure 1: Left: true sensitive counts. Note that they are self-consistent with the margins. Middle: Dark blue values are raw noisy measurements and light yellow values are computed margins. Note that the values are not self-consistent. Right: Final processed estimates, which are self-consistent with their margins.

For the inner product $\langle \cdot, \cdot \rangle_{\Sigma}$ and a tuple of linearly independent vectors in \mathbb{R}^n , $V = (v_1, \dots, v_k)$, the projection operator onto $S_V = \text{span}(V)$ is $\text{Proj}_{S_V}^{\Sigma} = V(V^{\top}\Sigma V)^{-1}V^{\top}\Sigma$. The projection is the solution to the following “least squares” problem: Let $y \in \mathbb{R}^n$; then $\text{Proj}_{S_V}^{\Sigma} y = \arg \min_{x \in S_V} \|x - y\|_{\Sigma}$. Furthermore, $\text{Proj}_{S_V}^{\Sigma^{-1}} = I - \text{Proj}_{S_V^{\perp}}^{\Sigma}$, where S_V^{\perp} is the orthogonal complement of S_V .

2.2 Contingency Table Notation

In this section we introduce some specialized notation that is necessary to precisely describe our proposed methodology. A summary of the notation is found in Table 7 in Appendix A.

We generally use the variables Y for true (unobserved) counts, \tilde{Y} for the observed noisy counts, \dot{Y} for intermediate estimates of Y from the collection step of the algorithm, and \hat{Y} for final estimators; see Example 1 and Figure 1 for an example. Given a finite set of variables $\mathcal{A} = \{A, B, C, \dots\}$ with an ordering (such as alphabetical), and $S \subset \mathcal{A}$ a subset of these variables, we write \tilde{Y}^S as the table of counts marginalized over the variables not in S to which noise has been applied. For a variable $A \in \mathcal{A}$, $|A|$ denotes the number of levels that A can take on. If $S = \{A_1, \dots, A_k\} \subset \mathcal{A}$, then $V^S = \{1, 2, \dots, |A_1|\} \times \dots \times \{1, 2, \dots, |A_k|\}$ is the set of indices for Y^S ; if $v \in V^S$ is a vector, then \tilde{Y}_v^S represents the single noisy count in \tilde{Y}^S which corresponds to $A_1 = v_1$, $A_2 = v_2$, and so on. We also consider $V_{\bullet}^S = \{\bullet, 1, 2, \dots, |A_1|\} \times \dots \times \{\bullet, 1, 2, \dots, |A_k|\}$, which allows for a “ \bullet ” in some of the entries. If $v \in V_{\bullet}^S$, the entries with a “ \bullet ” indicate that we are summing over all of the values for the corresponding variable. We use the same notation for Y , \dot{Y} , and \hat{Y} as well. For specific examples, we simplify the notation by dropping braces on the sets S and vectors v (e.g., for $S = \{A, B\}$ and $v = (i, j)$, we write $\tilde{Y}_{i,j}^{A,B}$ in place of $\tilde{Y}_{(i,j)}^{\{A,B\}}$).

We use the term *table* to refer to both subsets of variables $S \subset \mathcal{A}$ as well as sets of counts of the form \tilde{Y}^S (similarly for \dot{Y}^S and \hat{Y}^S). It should be clear from context which type of object is being referred to when the term “table” is used. When an index is given, we call objects like \tilde{Y}_v^S and \hat{Y}_v^S either *counts* or *estimates* of Y_v^S .

Example 1. Suppose that we have two binary variables A and B (e.g., Hispanic and Voting Age). Let $Y_{i,j}$ denote the true counts for $A = i$ and $B = j$. We observe noisy counts $\tilde{Y}_{i,j}^{A,B} = Y_{i,j} + \text{noise}$, $\tilde{Y}_i^A = Y_{i,\bullet} + \text{noise}$, $\tilde{Y}_j^B = Y_{\bullet,j} + \text{noise}$ and $\tilde{Y}^\emptyset = Y_{\bullet,\bullet} + \text{noise}$. Note that while $Y_i^A = Y_{i,\bullet}$, this is generally no longer the case for the values in \tilde{Y} : $\tilde{Y}_i^A \neq \tilde{Y}_{i,\bullet}^{A,B}$. However, our final estimates \hat{Y} will be “self-consistent” in this sense: e.g., $\hat{Y}_i^A = \hat{Y}_{i,\bullet}^{A,B}$. See Figure 1 for an illustration.

Some additional notation that will be important to describe and analyze our procedure is as follows: If $S \subset \mathcal{A}$, we write $\underline{S} \subset S$ and $\bar{S} \supset S$ as notation to remind us that \underline{S} is a subset of S and \bar{S} is a superset of S . If $v \in V^S$ and $\bar{S} \supset S$, then we define $v(\bar{S}) \in V_{\bullet}^{\bar{S}}$ as the vector that agrees with v on the entries corresponding to S and has the symbol “ \bullet ” for the entries corresponding to $\bar{S} \setminus S$; the effect of this notation is to marginalize out the variables $\bar{S} \setminus S$ and use v to index the variables from S in the marginalized table. Similarly, for $\underline{S} \subset S$, we define $v[\underline{S}] \in V^{\underline{S}}$ as the vector which extracts the entries of v which correspond to \underline{S} . Combining these two notations, we can also write $v[\underline{S}](S) \in V_{\bullet}^S$ for $\underline{S} \subset S$ and $v \in V^S$, and this vector agrees with v on \underline{S} and has “ \bullet ” on $S \setminus \underline{S}$, which has the effect of marginalizing out $S \setminus \underline{S}$ and using the remaining values in v to index the marginal table on \underline{S} .

2.3 Differential Privacy

Differential privacy (DP) is a probabilistic framework which formally quantifies the amount of privacy protection offered by a mechanism (randomized algorithm) (Dwork et al., 2006). Intuitively, a mechanism satisfies DP if, when it is run on two databases differing in one person’s contribution, the distribution of outputs is similar. There are many different notions of differential privacy, such as pure-DP, approximate-DP, Gaussian-DP (Dong et al., 2022), and zero-concentrated DP (Bun and Steinke, 2016), which differ in how the “similarity” between distributions is measured. In zero-concentrated DP, the version of DP employed by the U.S. Census Bureau (Abowd et al., 2022), the “similarity” is measured in terms of divergences on probability measures.

Definition 1 (Zero-Concentrated DP: Bun and Steinke (2016)). Let $\rho \geq 0$. A mechanism $M : \mathcal{X} \rightarrow \mathcal{Y}$ satisfies ρ -zero concentrated differential privacy (ρ -zCDP) if for any two databases X and X' differing in one entry, we have $D_\alpha(M(X)||M(X')) \leq \rho\alpha$, for all $\alpha \in (1, \infty)$, where $D_\alpha(\cdot||\cdot)$ is the order α Rényi divergence on probability measures:

$$D_\alpha(P||Q) = \frac{1}{\alpha - 1} \log \left(\mathbb{E}_{X \sim P} \left[\frac{p(X)}{q(X)} \right]^{\alpha-1} \right),$$

where p and q are the probability mass/density functions for P and Q , respectively.

In ρ -zCDP, smaller values of ρ ensure that the distributions of $M(X)$ and $M(X')$ are more similar, offering a stronger privacy guarantee.

The simplest and most widely used method of achieving differential privacy is to add independent noise to the entries of a statistic. The noise must be scaled to the *sensitivity* of the statistic. A statistic $T : \mathcal{X} \rightarrow \mathbb{R}^k$ has ℓ_2 -sensitivity Δ if $\Delta \geq \sup_{X, X'} \|T(X) - T(X')\|_2$, where the supremum is over all X and X' which differ in one entry. If T has ℓ_2 -sensitivity

Δ , then $T + N$ satisfies ρ -zCDP, where $N \sim \mathcal{N}(0, \Delta^2/(2\rho)I)$. If T takes values in \mathbb{Z}^n and has ℓ_2 -sensitivity Δ , then $T_i + N_i$, where $N_i \stackrel{\text{iid}}{\sim} N_{\mathbb{Z}}(0, \Delta^2/(2\rho))$ satisfies ρ -zCDP. Here, $N_{\mathbb{Z}}(0, \sigma^2)$ is the *discrete Gaussian* distribution with probability mass function $P(X = t) \propto \exp(-t^2/(2\sigma^2))$ for $t \in \mathbb{Z}$ (Canonne et al., 2020). The U.S. Census Bureau used discrete Gaussian noise to produce noisy measurements of the 2020 Decennial Census tabulations, which satisfy zCDP (Abowd et al., 2022). The methodology developed in this paper can be applied post-process a DP output from any DP framework as long as additive noise is used.

Differential privacy has several important properties including composition, group privacy, and immunity to post-processing (Dwork et al., 2014; Bun and Steinke, 2016). We review immunity to post-processing as it is central to the problem tackled in this paper.

Immunity to Post-Processing: If $M : \mathcal{X} \rightarrow \mathcal{Y}$ satisfies ρ -zCDP, and $g : \mathcal{Y} \rightarrow \mathcal{Z}$ is a (possibly randomized) function that does not depend on the sensitive data, then $g \circ M : \mathcal{X} \rightarrow \mathcal{Z}$ satisfies ρ -zCDP as well. Immunity to post-processing is important because one can first design a privacy mechanism which produces summary statistics, and then process these without compromising the privacy guarantee. Furthermore, by the transparency property of DP, the mechanism M itself can be publicly known and incorporated into the post-processing. For example, the “noisy measurements” (which we denote by \tilde{Y}) published by the U.S. Census Bureau satisfy zCDP; by post-processing, the procedure we propose can give refined estimates \hat{Y} without compromising the privacy guarantee.¹

3. Problem Setup and Theoretical BLUE

For the problem setup, we assume that we observe noisy counts $\tilde{Y}_v^S = Y_v^S + N_v^S$ for $S \in \mathcal{O} \subset \{A, B, C, \dots\}$ and $v \in V^S$, where \mathcal{O} denotes the set of *observed tables*, Y is the original contingency table of sensitive counts, and N_v^S are the noise random variables, which have a known distribution. Our goal is to produce refined estimates \hat{Y}_v^S for $S \in \mathcal{D}$, where \mathcal{D} is the set of *desired tables*, such that the \hat{Y}_v^S are *self-consistent*, meaning that we have $\hat{Y}_{v(S)}^S = \hat{Y}_v^{\underline{S}}$ for all $S, \underline{S} \in \mathcal{D}$, such that $\underline{S} \subset S$ and for all $v \in V^{\underline{S}}$. In other words, marginalizing out the variables of $S \setminus \underline{S}$ in \hat{Y}^S results in the same values as in $\hat{Y}^{\underline{S}}$. We assume that the tables in \mathcal{D} are nested: if $S \in \mathcal{D}$ and $\underline{S} \subset S$, then $\underline{S} \in \mathcal{D}$, but \mathcal{O} need not have this structure. Furthermore, we require our final estimates to be linear and unbiased, with the smallest variance possible.

For the remainder of the paper, the true counts Y will be viewed as fixed, unknown quantities. Thus, statements about the distribution of \tilde{Y} , or other random variables, refer only to the randomness due to the noise mechanism.

We consider a few additional structural assumptions on the problem:

- (A1) [Known noise] The noise random variables N_v^S are mean zero with known covariance $\Sigma = \text{Cov}(N)$. Furthermore, Σ is positive-definite and consists of finite entries.
- (A2) [Uncorrelated noise] The N_v^S are uncorrelated, implying that Σ is a diagonal matrix.

1. Technically the Census noisy measurement files do not strictly satisfy differential privacy because they contain *invariants*, which are values from the data reported without noise. See Gao et al. (2022) and Cho and Awan (2024) for investigations on what the technical guarantee is for these Census products

(A3) [Equal variance within tables] The noise variance is constant within a table: $\text{Var}(\tilde{Y}_v^S) = \text{Var}(\tilde{Y}_w^S)$ for all $v, w \in V^S$ and $S \in \mathcal{O}$.

Note that (A1)-(A3) are all verifiable assumptions on the problem setup. Furthermore, by the transparency property of DP, the privacy mechanism used can be published along with the privatized data \tilde{Y} without weakening the theoretical privacy guarantee, to allow outside users to check these assumptions as well. Assumptions (A1) and (A2) were previously used in the literature (Hay et al., 2010; Honaker, 2015; Kifer, 2021), while (A3) is new to this work and is needed to limit the search of all possible estimators.

Example 2 (Motivating Census Data Products). *The key motivating data sets for this paper are noisy measurement files for the 2020 Decennial Census data products, PL94 and DHC. Restricted to a single geography (e.g., only county-level counts for a particular county), the PL94 product satisfies (A1)-(A3); DHC does not exactly fit our framework as Age variable is binned to different values in different tables. In our application to DHC in Section 6, we use a subset of the noisy counts, which do satisfy (A1)-(A3). More details on these Census products is provided in Section 6.1.*

Remark 2 (Invariants). *Some of the Census tables also have invariant counts, which are values observed without noise. While such counts do not explicitly fit into the assumptions (A1)-(A3), with the present version of our method, our procedure can be applied by setting the variance to be arbitrarily small for these counts. We suspect that a modification to our procedure can formally incorporate these invariants as part of the constraints without such an approximation, but leave this as future work.*

A general solution can be provided to this problem, assuming only property (A1), using the theory of linear models. We show that the best linear unbiased estimator, can be expressed in terms of linear projections, and we give a formula for the projection operator in terms of the constraint matrix. The proof of the result is based on the theory of linearly constrained linear regression.

Theorem 3. *Let $X \in \mathbb{R}^n$ be a random vector with mean μ and covariance Σ , where it is known that $\mu \in S_b = \{x \in \mathbb{R}^n \mid Ax = b\}$. Let v be such that $b = Av$ and call $S_0 = \{x \in \mathbb{R}^n \mid Ax = 0\}$. Call $\hat{\mu} = \text{Proj}_{S_0}^{\Sigma^{-1}}(X) + (\text{Proj}_{S_0^\perp}^\Sigma)^\top v$. Then,*

- $\hat{\mu}$ is the BLUE for μ . Furthermore, $\text{Cov}(\hat{\mu}) = \text{Proj}_{S_0}^{\Sigma^{-1}} \Sigma (\text{Proj}_{S_0}^{\Sigma^{-1}})^\top$.
- If $X \sim \mathcal{N}(\mu, \Sigma)$, then $\hat{\mu}$ is the uniformly minimum variance unbiased estimator for μ ; furthermore, $\hat{\mu} \sim \mathcal{N}(\mu, \text{Proj}_{S_0}^{\Sigma^{-1}} \Sigma (\text{Proj}_{S_0}^{\Sigma^{-1}})^\top)$.

If the rows of A are linearly independent, then

$$\text{Proj}_{S_0}^{\Sigma^{-1}} = \left(I - \Sigma A^\top (A \Sigma A^\top)^{-1} A \right) \quad \text{and} \quad (\text{Proj}_{S_0^\perp}^\Sigma)^\top = \Sigma A^\top (A \Sigma A^\top)^{-1} A.$$

While Theorem 3 derives the BLUE, when n becomes large, computing the projection becomes expensive, especially in terms of memory which scales as $O(pn)$, where p is the dimension of the rowspace or nullspace of A —whichever is smaller. Furthermore, the runtime for this estimator using matrix operations is approximately $O(n^{2.373})$ (Alman and Williams, 2021); see Section 4.3 for a discussion of computational complexity.

Example 3. Applying Theorem 3 to our setting, we have $X = \tilde{Y}$, $\mu = Y$, and $\hat{\mu} = \hat{Y}$. In the constraints, A is the matrix with entries in $\{-1, 0, 1\}$, which encodes the self-consistency constraints, and $b = 0$. Note that the projections are orthogonal in the case that $\Sigma = cI$; in general, assumptions (A1)-(A3) do not ensure that the projections are orthogonal.

Example 4 (Running Toy Example). We introduce a toy example that we will use to illustrate Theorem 3; later we will revisit this example to illustrate the SEA BLUE algorithm. Suppose that the only variable is B , which has 3 levels. The sensitive values are $Y^B = (5, 10, 15)^\top$, with $Y^\emptyset = 30$. After adding independent discrete Gaussian noise with mean zero and variance 1 to each value, we observe the noisy counts $\tilde{Y}^B = (6, 9, 17)^\top$ and $\tilde{Y}^\emptyset = 29$. Note that the noisy values are not self-consistent as $\tilde{Y}_\bullet^B = 32 \neq 29 = \tilde{Y}^\emptyset$. In the notation of Theorem 3, we have $X = \begin{pmatrix} \tilde{Y}^B \\ \tilde{Y}^\emptyset \end{pmatrix}$ with covariance $\Sigma = I$, and constraint $AX = 0$ where $A = (1, 1, 1, -1)$. We calculate that

$$\text{Proj}_{S_0}^I = (I - A^\top(AA^\top)^{-1}A) = \frac{1}{4} \begin{pmatrix} 3 & -1 & -1 & 1 \\ -1 & 3 & -1 & 1 \\ -1 & -1 & 3 & 1 \\ 1 & 1 & 1 & 3 \end{pmatrix}.$$

Thus, by Theorem 3, we have that $\begin{pmatrix} \hat{Y}^B \\ \hat{Y}^\emptyset \end{pmatrix} = \text{Proj}_{S_0}^I \begin{pmatrix} \tilde{Y}^B \\ \tilde{Y}^\emptyset \end{pmatrix} = (5.25, 8.25, 16.25, 29.75)^\top$, which are self-consistent and BLUE.

In the following section, we propose an efficient implementation of BLUE under the additional assumptions of uncorrelated noise (A2) and equal variance within tables (A3).

4. Scalable Efficient Algorithm for BLUE

In this section, we propose a novel two-step estimator that implements BLUE under assumptions (A1)-(A3) and gives a well-motivated heuristic when (A2) and (A3) do not hold. We call our method the Scalable Efficient Algorithm for the Best Linear Unbiased Estimator (SEA BLUE). The two steps of SEA BLUE are a collection step, which aggregates estimators of lower margins from more detailed tables, followed by a down pass, which projects the intermediate estimators resulting from the first step to be self-consistent with the previously calculated margins above. See Figure 2 for an illustration.

In Section 4.1, we describe the collection step of SEA BLUE and prove in Theorem 6 that it results in the BLUE based on all estimators “from below;” in the case of the estimate of the total count, this is simply the BLUE. We also present a pseudo-code implementation of the collection step, which optimizes performance. In Section 4.2 we describe the down pass and prove that, when applied to the output of the collection step, it results in the BLUE for every table. We also develop an efficient implementation of the down pass by optimizing the particular projection operations that are used in this step, and give a pseudo-code implementation. In Section 4.3, we compare both the runtime and memory complexity of SEA BLUE to the matrix projection, demonstrating the increased efficiency of SEA BLUE.

4.1 Collection Step

The first step of SEA BLUE is called the *collection step*, which considers estimators that can be summed up from a single table, and weights these to give preliminary estimators

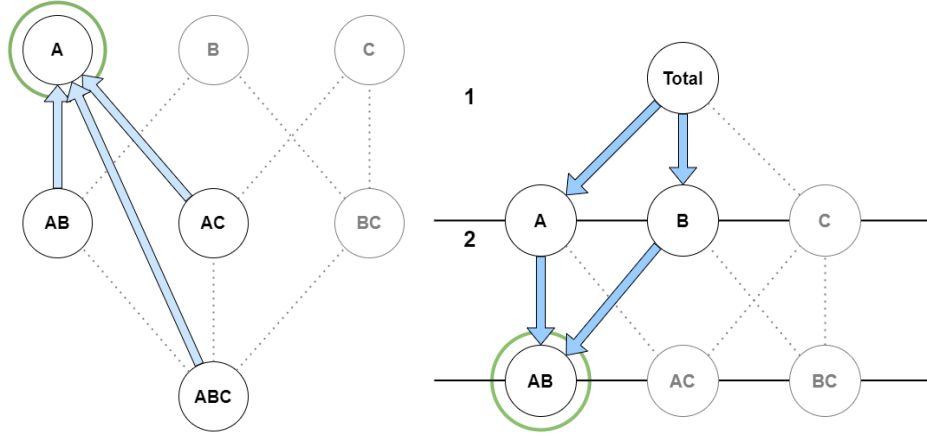


Figure 2: Left: Illustration of the “collection step,” which updates the estimates in the A table using a weighted average of the A margin calculated from A , AB , AC , and ABC tables from the original noisy counts in \tilde{Y} . Right: Illustration of the “down pass,” which updates the AB table to be consistent with the A and B tables via projection; prior to this, the A , B , and C tables were projected to be consistent with the total. The down pass is applied to the intermediate estimates \hat{Y} obtained from the collection step.

of each count, which we denote by \hat{Y}_v^S for a table S and index $v \in V^S$. The collection step was designed with the assumption of uncorrelated noise (A2) in mind, which removes the need to consider the covariance between estimators from different tables. By using inverse-variance weighting at the end of the collection step, the intermediate estimators are optimal “from below.” The collection step does not consider any estimators which leverage information between tables that are not hierarchically related, and the down pass will be needed to incorporate this information.

Definition 4 (Collection Step). *Let $\tilde{Y}_v^S = Y_v^S + N_v^S$ be noisy counts for $S \in \mathcal{O}$ and $v \in V^S$ satisfying (A1). For any $S \subset \mathcal{A}$ and $v \in V^S$, the collection step produces intermediate estimates \hat{Y}_v^S via the following formula:*

$$\hat{Y}_v^S = \frac{\sum_{\bar{S} \in \mathcal{O} | \bar{S} \supset S} \left(\text{Var}(\tilde{Y}_{v(\bar{S})}^{\bar{S}}) \right)^{-1} \tilde{Y}_{v(\bar{S})}^{\bar{S}}}{\sum_{\bar{S} \in \mathcal{O} | \bar{S} \supset S} \left(\text{Var}(\tilde{Y}_{v(\bar{S})}^{\bar{S}}) \right)^{-1}}. \quad (1)$$

Note that because the entries of the noisy counts \tilde{Y} have a known covariance, we can easily calculate $\text{Var}(\tilde{Y}_{v(\bar{S})}^{\bar{S}})$ exactly. If (A2) holds, then the variance of the collection step estimates is given by the following Lemma:

Lemma 5. *If (A1) and (A2) hold, then the variance of the entries of the intermediate estimate \dot{Y}_v^S , denoted by $(\dot{\sigma}^2)_v^S$, is given by*

$$(\dot{\sigma}^2)_v^S := \text{Var}(\dot{Y}_v^S) = \left(\sum_{\bar{S} \in \mathcal{O} | \bar{S} \supset S} \left(\text{Var}(\tilde{Y}_{v(\bar{S})}^{\bar{S}}) \right)^{-1} \right)^{-1}. \quad (2)$$

Under assumptions (A1) and (A2), the collection step produces an optimal estimator, using the estimates that are “from below” and “within-table.” Under the assumption of equal variance within tables (A3), the collection step is simply the optimal estimator “from below.” In particular, the collection step estimate of the total Y^\emptyset is the BLUE.

Theorem 6. *[Optimality for the Total] Assume that (A1)-(A3) hold. Then, \dot{Y}^\emptyset is the BLUE of the total count Y^\emptyset based on the values in \tilde{Y}^S for $S \in \mathcal{O}$, subject to self-consistency.*

While Theorem 6 is explicitly stated in terms of the total Y^\emptyset , it implies that all collection estimates are optimal “from below,” by relabeling a count of interest Y_v^S as Y^\emptyset and only considering tables $\bar{S} \in \mathcal{O}$ which satisfy $\bar{S} \supset S$.

Example 5 (Running Toy Example). *We apply the collection step to the setting in Example 4. Note that assumptions (A1)-(A3) all hold. Since there are no observed tables “below” \tilde{Y}^B , we set $\dot{Y}^B = \tilde{Y}^B$. Using inverse-variance weighting, we have*

$$\dot{Y}^\emptyset = \frac{(1/3)\tilde{Y}_{\bullet}^B + (1)\tilde{Y}^\emptyset}{4/3} = \frac{(1/3)(32) + 29}{4/3} = 29.75.$$

Note that we still do not have self consistency between \dot{Y}^B and \dot{Y}^\emptyset .

Assumption (A2) is needed for Theorem 6 as the collection step only uses inverse-variance weighting, based on the assumption of uncorrelated noise. If noises are correlated, inverse-covariance weighting would be needed (Anderson, 2008), which would be far more computationally and memory expensive. Fortunately many DP statistics use uncorrelated noise, such as those produced by the U.S. Census Bureau, and prior works have also assumed this structure (Hay et al., 2010; Honaker, 2015; Kifer, 2021). The following example demonstrates the suboptimality of \dot{Y}^\emptyset when (A2) does not hold.

Example 6 (Necessity of (A2) for Collection Step). *Suppose that $|A| = |B| = 2$, $\mathcal{O} = \{\{A\}, \{A, B\}\}$, $\tilde{Y}_{i,j}^{A,B}$ are independent with variance 1, and $\tilde{Y}_i^A = \tilde{Y}_{i,\bullet}^{A,B} + N_i$, for $i = 1, 2$, where N_i are independent of all other variables with variance 1. Then the collection step produces $\dot{Y}^\emptyset = ((1/4)\tilde{Y}_{\bullet,\bullet}^{A,B} + (1/6)\tilde{Y}_{\bullet}^A)(1/4 + 1/6) = (3\tilde{Y}_{\bullet,\bullet}^{A,B} + 2\tilde{Y}_{\bullet}^A)/5$, which has variance $\text{Var}(\tilde{Y}_{\bullet,\bullet}^{A,B}) + \frac{4}{25}(\text{Var}(N_1) + \text{Var}(N_2)) = 4 + 8/25$. On the other hand, inverse-covariance weighting tells us that the BLUE for Y^\emptyset is simply $\tilde{Y}_{\bullet,\bullet}^{A,B}$, which has variance 4. Thus, \dot{Y}^\emptyset is not BLUE.*

It may not be immediately obvious why (A3) is necessary to ensure the optimality of the collection step. However, if (A3) does not hold, there may be important between-table estimators that are not considered by the collection step, as demonstrated in the following example.

Example 7 (Necessity of (A3) for Collection Step). *Suppose that $|A| = |B| = 2$, $\mathcal{O} = \{\{A\}, \{A, B\}\}$, $\text{Var}(\tilde{Y}_1^A) = \text{Var}(\tilde{Y}_{2,1}^{A,B}) = \text{Var}(\tilde{Y}_{2,2}^{A,B}) = 1$ and $\text{Var}(\tilde{Y}_2^A) = \text{Var}(\tilde{Y}_{1,1}^{A,B}) = \text{Var}(\tilde{Y}_{1,2}^{A,B}) = 11$. Then (A3) does not hold, since there are unequal variances within both the $\{A, B\}$ table and the $\{A\}$ table. Using (2), we compute*

$$\text{Var}(\dot{Y}^\emptyset) = \left(1/\text{Var}(\tilde{Y}_\bullet^A) + 1/\text{Var}(\tilde{Y}_{\bullet,\bullet}^{A,B})\right)^{-1} = (1/12 + 1/24)^{-1} = 8.$$

However, a between-table estimator, $Z = \tilde{Y}_1^A + \tilde{Y}_{2,1}^{A,B} + \tilde{Y}_{2,2}^{A,B}$ has variance $\text{Var}(Z) = 3$. So, \dot{Y}^\emptyset is not the BLUE.

4.1.1 IMPLEMENTATION OF COLLECTION STEP

The collection step (1) can be easily implemented by first iterating over all desired tables $S \in \mathcal{D}$, and then over all observed tables $R \supset S$, to produce the within-table estimates of Y^S . However, this implementation is inefficient as it fails to leverage the ability to reuse previous computations. In Algorithm 1, we describe an implementation that calculates (1) more efficiently. The idea is to first iterate over the observed tables, which allows us to leverage higher detailed margins when calculating lower detailed margins (e.g., first calculate the one-way margins and then calculate the total from the one-way margins). Algorithm 1 offers the possibility of parallelization by working with different observed tables on different computing nodes and aggregating the summary values C_v^S, D_v^S at the end. Note that the collection step can be further optimized by marginalizing out the variables with the largest number of levels first in order to reduce the number of counts that need to be marginalized over for lower level margins.

4.2 Down Pass

In the down pass, we begin with the intermediate estimates produced by the collection step and iteratively project the i -way margins to be self-consistent with the previously computed $(i - 1)$ -way margin. The total (0-way margin) is simply equal to the end result of the collection step ($\hat{Y}^\emptyset = \dot{Y}^\emptyset$).

Definition 7 (Down Pass). *Let \dot{Y}_v^S be the output of the collection step, and assume that $(\sigma^2)_v^S = \text{Var}(\dot{Y}_v^S)$ is given (such as by Lemma 5 using (A1) and (A2)). Then the down pass produces final estimates \hat{Y}^S that are defined recursively as follows:*

- If $S = \emptyset$, then $\hat{Y}^\emptyset = \dot{Y}^\emptyset$.
- Otherwise, \hat{Y}^S is defined to be the BLUE for Y^S , given the intermediate estimates \dot{Y}^S , subject to the constraint $\hat{Y}_{v(S)}^S = \hat{Y}_v^{\underline{S}}$ for all $\underline{S} \subsetneq S$ and $v \in V^{\underline{S}}$; that is, when marginalizing out the variables of $S \setminus \underline{S}$ in \hat{Y}^S , we obtain the same values as in $\hat{Y}^{\underline{S}}$.

By Theorem 3, we know that \hat{Y}^S can be expressed as a projection, given in Algorithm 2.

From Definition 7, we see that the down pass is “locally optimal” in the sense that if we only had the intermediate estimate \dot{Y}^S , and optimal estimates $\hat{Y}^{\underline{S}}$ are treated as fixed for $\underline{S} \subsetneq S$, then \hat{Y}^S is the optimal estimator. Under assumptions (A1)-(A3), we can in fact say that \hat{Y}^S is BLUE compared to all linear estimators based on the noisy counts \tilde{Y} .

Algorithm 1 Pseudo Code for Collection Step

Require: Noisy counts \tilde{Y}^S for each $S \in \mathcal{O}$, and known variance $\text{Var}(\tilde{Y}_v^S)$ for each $v \in V^S$.

- 1: For each $S \in \mathcal{D}$ and $v \in V^S$, set $C_v^S = D_v^S = 0$.
- 2: **for all** $S \in \mathcal{O}$ (loop through observed tables and create a spanning set of margins) **do**
- 3: Define $\overline{\mathcal{D}}_S \subset 2^S$ such that $\overline{\mathcal{D}}_S \supset \mathcal{D} \cap 2^S$, $\emptyset \in \overline{\mathcal{D}}_S$ and for any $D \in \mathcal{D} \cap 2^S$, there exists a nested sequence $(R_j)_{j=|D|}^{|S|}$ in $\overline{\mathcal{D}}_S$ such that
 1. $R_j \subset R_k$ for all $|D| \leq j \leq k \leq |S|$,
 2. $|R_j| = j$,
 3. $R_{|D|} = D$.

(that is, for any desired table D , involving the variables in S , we have a sequence of tables (R_j) in $\overline{\mathcal{D}}_S$ that increase one variable at a time from D to S ; we refer to $\overline{\mathcal{D}}_S$ as a “spanning set of margins”)
- 4: **for** $i = 0, 1, \dots, |S|$ (number of variables to marginalize) **do**
- 5: **for all** $R \in \overline{\mathcal{D}}_S$ such that $|S| = |R| - i$ (pick a set of variables not marginalized out) **do**
- 6: **for all** $v \in V^R$ (all counts in the table) **do**
- 7: **if** $i = 0$ (implies that $R = S$) **then**
- 8: Define $\tilde{Y}_v^{S \mapsto S} = \tilde{Y}_v^S$
- 9: Define $(\tilde{\sigma}^2)_v^{S \mapsto S} = \text{Var}(\tilde{Y}_v^S)$
- 10: (the notation $S \mapsto R$ represents that we are producing estimates for the table R based on margins calculated from table S)
- 11: **else**
- 12: There exists $\overline{R} \in \overline{\mathcal{D}}_S$ such that $R \subset \overline{R} \subset S$ and $|\overline{R}| = |R| + 1$
- 13: Define $\tilde{Y}_v^{S \mapsto R} = \tilde{Y}_{v(\overline{R})}^{S \mapsto \overline{R}}$ (sum up values from a previously marginalized table)
- 14: Define $(\tilde{\sigma}^2)_v^{S \mapsto R} = (\tilde{\sigma}^2)_{v(\overline{R})}^{S \mapsto \overline{R}}$ (sum up variances from previous margin)
- 15: (we use a previously calculated margin to avoid re-summing the same quantities unnecessarily)
- 16: **end if**
- 17: **if** $R \in \mathcal{D}$ (if the table is one of the desired outputs) **then**
- 18: Update $C_v^R = C_v^R + \tilde{Y}_v^{S \mapsto R} / (\tilde{\sigma}^2)_v^{S \mapsto R}$ (inverse-variance-weighted running sum)
- 19: Update $D_v^R = D_v^R + 1 / (\tilde{\sigma}^2)_v^{S \mapsto R}$ (running sum of the inverse-variances)
- 20: **end if**
- 21: **end for**
- 22: **end for**
- 23: **end for**
- 24: **end for**
- 25: **for all** $S \in \mathcal{D}$ and $v \in V^S$ **do**
- 26: Set $\hat{Y}_v^S = C_v^S / D_v^S$ (by Equation (1))
- 27: Set $(\hat{\sigma}^2)_v^S = 1 / D_v^S$ (by Equation (2))
- 28: **end for**

Theorem 8. *[Optimality of SEA BLUE] Assume (A1)-(A3) hold. Then, \hat{Y}^S resulting from applying the down pass to the output of the collection step is the best linear unbiased estimator of Y^S based on the values in \tilde{Y} , under self-consistency.*

4.2.1 EFFICIENT IMPLEMENTATION OF THE DOWN PASS

Next, we derive an efficient implementation of the projections required in the down pass: projecting a k -way table to be self-consistent with fixed higher margins. Assuming (A1)-(A3) hold, it follows that all of the elements of \hat{Y}^S are independent with equal variance, and so the desired projection is an *orthogonal projection*.

First, we require some technical notation. Let S be a set of variables, and let Z^S be a generic table on these variables. An important subspace is

$$\mathcal{S}_S = \{Z^S \mid \text{for all } \underline{S} \subsetneq S, \text{ and all } v \in V^{\underline{S}}, \text{ we have } Z_{v(S)}^S = 0\}, \quad (3)$$

which is the set of tables with all margins equal to zero.

Given \hat{Y}^S , the output of the collection step, and $\hat{Y}^{\underline{S}}$ for $\underline{S} \subsetneq S$, the desired margin values, we construct two tables: \hat{Z}^S , a table consistent with $\hat{Y}^{\underline{S}}$ which does not use the values in \hat{Y}^S , and \hat{Z}^S , which has the same structure as \hat{Z}^S , but uses the margins of \hat{Y}^S :

$$\hat{Z}_v^S = \sum_{k=1}^{|S|} \sum_{\substack{\{A_1, \dots, A_k\} \subset S \\ \underline{S} = S \setminus \{A_1, \dots, A_k\}}} (-1)^{k+1} \frac{\hat{Y}_{v[\underline{S}]}^R}{|A_1| \cdots |A_k|}, \quad (4)$$

$$\hat{Z}_v^S = \sum_{k=1}^{|S|} \sum_{\substack{\{A_1, \dots, A_k\} \subset S \\ \underline{S} = S \setminus \{A_1, \dots, A_k\}}} (-1)^{k+1} \frac{\hat{Y}_{v[\underline{S}](S)}^S}{|A_1| \cdots |A_k|}. \quad (5)$$

Lemma 9 establishes that \hat{Z}_v^S and \hat{Z}^S have the desired margins and belong to \mathcal{S}_S^\perp .

Lemma 9. *Let $S \in \mathcal{D}$ and assume that the $\hat{Y}^{\underline{S}}$'s are self-consistent for all margins $\underline{S} \subsetneq S$. Let \hat{Z}^S be constructed as in (4) and \hat{Z}^S be from (5). Then we have that*

1. *for all $\underline{S} \subsetneq S$ and all $v \in V^{\underline{S}}$, $\hat{Z}_{v(S)}^S = \hat{Y}_v^{\underline{S}}$ and $\hat{Z}_{v(S)}^S = \hat{Y}_{v(S)}^S$, and*
2. *both $\hat{Z}^S, \hat{Z}^S \in \mathcal{S}_S^\perp$.*

Not only do the constructions \hat{Z}^S and \hat{Z}^S have the desired margins and live in \mathcal{S}_S^\perp , but in fact, \hat{Z}^S is the projection of \hat{Y}^S onto \mathcal{S}_S^\perp . This allows us to write the elegant formula $\hat{Y}^S = \hat{Y}^S - \hat{Z}^S + \hat{Z}^S$ established in the following Theorem:

Theorem 10. *Assume (A1) holds, let \hat{Y}^S be a table achieved after the collection step, and let \hat{Z}^S be the result of (4) using $\hat{Y}^{\underline{S}}$ for $\underline{S} \subsetneq S$. Let \hat{Z}^S be defined as in (5) using the values in \hat{Y}^S . Then,*

1. *$\hat{Z}^S = \text{Proj}_{\mathcal{S}_S^\perp}^I \hat{Y}^S$ (equivalently, $\hat{Y}^S - \hat{Z}^S = \text{Proj}_{\mathcal{S}_S}^I \hat{Y}^S$), and*
2. *if (A2) and (A3) hold, then $\hat{Y}^S = \hat{Y}^S - \hat{Z}^S + \hat{Z}^S$.*

Algorithm 2 Pseudo Code for Down Pass

Require: (A1) holds; \hat{Y}^S and $(\sigma^2)^S$ for each $S \in \mathcal{D}$ as computed from the collection step

- 1: Set $\hat{Y}^\emptyset = \hat{Y}^\emptyset$, which is the total estimate
- 2: **for** $i = 1, 2, \dots, \max_{S \in \mathcal{D}} |S|$ (iteratively go through i -way margins) **do**
- 3: **for all** $S \in \mathcal{D}$ such that $|S| = i$ (pick a set of i -way margins) **do**
- 4: Let \hat{Z}^S be the table constructed in (4)
- 5: **if** (A1)-(A3) hold **then**
- 6: Let \hat{Z}^S be the table constructed in (5)
- 7: Set $\hat{Y}^S = \hat{Y}^S - \hat{Z}^S + \hat{Z}^S$ (by Theorem 10)
- 8: **else if** (A1)-(A2) hold **then**
- 9: Call $\hat{\Sigma} = \text{diag}((\sigma^2)^S)$, which is the covariance matrix of the elements of \hat{Y}^S
- 10: Set $\hat{Y}^S = \text{Proj}_{\hat{\Sigma}^{-1}}^{\hat{\Sigma}^{-1}}(\hat{Y}^S) + \hat{Z}^S - \text{Proj}_{\hat{\Sigma}^{-1}}^{\hat{\Sigma}^{-1}}(\hat{Z}^S)$ (projections are no longer orthogonal)
- 11: **end if**
- 12: **end for**
- 13: **end for**

The intuition behind the formula $\hat{Y}^S = \hat{Y}^S - \hat{Z}^S + \hat{Z}^S$ is that \hat{Y}^S has the correct dependence structure, but the wrong margins; so, we “subtract off the margins” by subtracting \hat{Z}^S and then “add on the correct margins” by adding \hat{Z}^S .

In Example 8, we illustrate how the terms in \hat{Z}^S cancel when calculating margins.

Example 8. Let $S = \{A, B, C\}$, and suppose that the \hat{Y}^R ’s are self-consistent for all $R \subsetneq S$ (e.g., $\hat{Y}_{a,\bullet}^{A,B} = \hat{Y}_a^A$). The construction of \hat{Z}^S in (4) is of the form

$$\hat{Z}_{a,b,c}^S = \left(\frac{\hat{Y}_{a,b}^{A,B}}{|C|} + \frac{\hat{Y}_{a,c}^{A,C}}{|B|} + \frac{\hat{Y}_{b,c}^{B,C}}{|A|} \right) - \left(\frac{\hat{Y}_a^A}{|B||C|} + \frac{\hat{Y}_b^B}{|A||C|} + \frac{\hat{Y}_c^C}{|A||B|} \right) + \frac{\hat{Y}^\emptyset}{|A||B||C|}.$$

If we marginalize over the variable C , we obtain:

$$\hat{Z}_{a,b,\bullet}^S = \left(\hat{Y}_{a,b}^{A,B} + \frac{\hat{Y}_a^A}{|B|} + \frac{\hat{Y}_b^B}{|A|} \right) - \left(\frac{\hat{Y}_a^A}{|B|} + \frac{\hat{Y}_b^B}{|A|} + \frac{\hat{Y}^\emptyset}{|A||B|} \right) + \frac{\hat{Y}^\emptyset}{|A||B|},$$

and we see that all terms cancel except for $\hat{Y}_{a,b}^{AB}$ as claimed in Lemma 9.

Remark 11. For each step of the down pass, the projection only requires calculating the current margins of \hat{Y}^S and the cost to assemble \hat{Z}^S and \hat{Z}^S . In fact, the computational and memory cost for the down pass is of the same rate as the collection step.

Algorithm 2 gives a pseudo-code implementation of the down pass. Note that while each set of margins must be computed in order, margins with the same number of variables can be calculated in parallel.

Example 9. We continue Example 5 by now applying the down pass. First, we calculate $\hat{Z}^S = (32/3, 32/3, 32/3)^\top$, since $32 = \hat{Y}_\bullet^B$ and $|B| = 3$, and $\hat{Z}^S = (29.75/3, 29.75/3, 29.75/3)^\top$, since $\hat{Y}^\emptyset = \hat{Y}^\emptyset = 29.75$ and $|B| = 3$. Thus,

$$\hat{Y}^B = \hat{Y}^B - \hat{Z}^S + \hat{Z}^S = \begin{pmatrix} 6 \\ 9 \\ 17 \end{pmatrix} - \begin{pmatrix} 32/3 \\ 32/3 \\ 32/3 \end{pmatrix} + \begin{pmatrix} 29.75/3 \\ 29.75/3 \\ 29.75/3 \end{pmatrix} = \begin{pmatrix} 5.25 \\ 8.25 \\ 16.25 \end{pmatrix},$$

which agrees with the BLUE calculated in Example 4.

4.3 Complexity of SEA BLUE and Matrix Projection

In this section, we give a detailed analysis of the computation and memory complexity of the matrix projection in Theorem 3 compared to SEA BLUE. For simplicity, we restrict our attention to the case where noisy tables are observed for all possible margins based on $k \geq 1$ variables with $I \geq 2$ levels each and assumptions (A1)-(A3) hold. In this case, we can calculate that there are $n = (I + 1)^k$ noisy counts observed. We consider two asymptotic regimes where either k or I goes to infinity while the other is held fixed.

Matrix Projection: The projection in Theorem 3 requires $O(np + p^{2.373})$ operations to implement (Alman and Williams, 2021) and $O(np)$ units of memory, where p is the smaller dimension of either the column space or nullspace of A (A is defined in Theorem 3).

Lemma 12. *Given all possible tables from k variables, each with I levels,*

1. $p = \min\{I^k, (I + 1)^k - I^k\}$, where p is defined as above, and
2. *matrix projection has runtime $O(I^{2k-1} + I^{2.373(k-1)})$ and memory $O(I^{2k-1})$ as $I \rightarrow \infty$, and runtime $O(I^k(I + 1)^k + I^{2.373k})$ and memory $O(I^k(I + 1)^k)$ as $k \rightarrow \infty$.*

SEA BLUE: In both the collection step and the down pass, it can be seen that the memory requirement never surpasses a constant multiple times n , a large improvement over the rate $O(pn)$ for the matrix projection.

In terms of the runtime, for the collection step, the bottleneck is the time to compute all margins from all of the tables. In Lemma 18, found in the Supplementary Materials, it is established that under (A1)-(A3), the runtime of the down pass is limited by the same quantity. In Lemma 13, we derive the computational cost and memory to compute all margins given either a single table, or all tables from k variables.

Lemma 13. 1. *Let S be a table with k variables, each with I levels. Then it is possible to compute all of its margins in $O(I^k)$ time as $I \rightarrow \infty$ and $O((I + 1)^k)$ time as $k \rightarrow \infty$; the memory requirement is of the same order.*

2. *Given all possible tables from k variables, each with I levels, it is possible to compute all of its margins in $O(I^k) = O(n)$ time as $I \rightarrow \infty$ and $O((2I + 1)^k) = O(n^{\log(2I+1)/\log(I+1)})$ time as $k \rightarrow \infty$; under (A1)-(A3), the memory requirement of SEA BLUE is $O(n)$.*

Remark 14. *In Lemma 13, the quantity $\log(2I + 1)/\log(I + 1)$ is a decreasing function for $I \geq 0$, which approaches 1 as $I \rightarrow \infty$. At $I = 2$ (the smallest number of levels for a variable), it evaluates to $\log(5)/\log(3) \approx 1.46497$. The runtime as $k \rightarrow \infty$ is at worst $o(n^{1.5})$ and approaches $O(n)$ for larger I . On the other hand, as $I \rightarrow \infty$, we have runtime of $O(n)$ no matter the value of k .*

Comparison Between Matrix Projection and SEA BLUE: We see a large improvement in the memory usage of SEA BLUE compared to the matrix projection: $O(n)$ versus $O(np)$, where p is only slightly smaller than n .

In terms of runtime, when $I \rightarrow \infty$, we see a clear improvement to $O(I^k)$ versus $O(I^{2k-1} + I^{2.373(k-1)})$; these runtimes only agree when $k = 1$. As $k \rightarrow \infty$, we can compare $(2I + 1)^k \leq 2^k(I + 1)^k = O(I^k(I + 1)^k)$, where the big- O is tight when $I = 2$ and is little- o for $I > 2$, and $(2I + 1)^k = o(I^{2.373k})$, for $I \geq 2$. Less precisely, we have that SEA BLUE has runtime between $\Omega(n)$ and $o(n^{1.5})$, whereas the matrix projection has a runtime approximately on the order of $O(n^2)$ - $O(n^{2.373})$.

5. Covariance Calculation and Confidence Intervals

When publishing differentially private data, it is important to communicate the error due to privacy in the published estimates. In this section, we show that for any linear procedure (including SEA BLUE), we can either calculate the standard error exactly or approximate it with a Monte Carlo approach. In either case, we show how to construct confidence intervals. We also present a distribution-free method for Monte Carlo confidence intervals.

We introduce additional notation to simplify the presentation of this section: for a table Z^S , we let $\text{vec}(Z^S)$ denote a vectorization of the entries of Z^S , where we assume a fixed ordering of the variables and levels of the variables.

5.1 Exact Covariance Calculation and z -Confidence Intervals

Under assumption (A1)-(A3), we know that SEA BLUE produces the same result as the best linear unbiased estimate described in Theorem 3. In Theorem 3, the covariance of the final estimates \hat{Y} can be obtained by applying the projection procedure to a factorization of the original covariance: $\Sigma^{1/2}$. If the noisy counts \tilde{Y} are normally distributed, we can easily construct confidence intervals for the true values Y using the calculated variances.

Suppose that (A1)-(A3) hold. Let $\Sigma_i^{1/2}$ be the i^{th} column of $\Sigma^{1/2}$. Apply SEA BLUE to $\Sigma_i^{1/2}$ (pretending the vector $\Sigma_i^{1/2}$ contains noisy counts), setting all invariants to zero, and call the result $\hat{\Sigma}_i^{1/2}$. Let $\hat{\Sigma}^{1/2}$ be the matrix whose i^{th} column is $\hat{\Sigma}_i^{1/2}$. Then, $\hat{\Sigma} = \hat{\Sigma}^{1/2}(\hat{\Sigma}^{1/2})^\top$ is the covariance matrix of $\text{vec}(\hat{Y})$. If $\Sigma = \sigma^2 I$ (a stronger version of (A3) where all counts have equal variance), then $\hat{\Sigma}$ is obtained column-wise by simply passing each column of Σ through SEA BLUE and no matrix multiplication is needed.

If \tilde{Y}_v^S are all normally distributed, then a $(1 - \alpha)$ -confidence interval for $\text{vec}(Y)_i$ is $\text{vec}(\hat{Y})_i \pm z_{1-\alpha/2} \sqrt{(\hat{\Sigma})_{i,i}}$.

5.2 Monte Carlo Confidence Intervals

In Proposition 5.1, we saw how to compute the exact covariance of \hat{Y} under assumptions (A1)-(A3). However, if either the assumption of uncorrelated noise (A2) or the assumption of equal variance within tables (A3) are violated, then this result no longer holds. Furthermore, if the variance of every count is desired, running SEA BLUE for n times may be overly expensive. Finally, even if Proposition 5.1 is applicable, if the \tilde{Y} are not normally distributed, producing tight CIs based on the covariance may be challenging, as the use of suboptimal concentration inequalities such as Chebyshev and Gauss can result in confidence intervals with excess width.

To address these limitations, we propose two Monte Carlo confidence interval methods, one under the assumption of normality and the other which holds without any distributional assumptions. These methods simply rely on the fact that the SEA BLUE algorithm is a linear procedure (even if (A2) and (A3) do not hold), and the noise distribution of N is known (as mentioned under “Immunity to Post-Processing” in Section 2.3).

Suppose that (A1) holds and \tilde{Y} are normally distributed. Then, it follows that \hat{Y} are also normally distributed with mean Y . Let $\text{vec}(\tilde{Y})^{(1)}, \dots, \text{vec}(\tilde{Y})^{(R)}$ be i.i.d. samples with the same distribution as $\text{vec}(N)$. Call $(\hat{\sigma}^2)_i^R = \frac{1}{m} \sum_{j=1}^R [\text{vec}(\hat{Y}^{(j)})_i]^2$, where $\hat{Y}^{(j)}$ is the result of applying SEA BLUE to $\tilde{Y}^{(j)}$. Then, $\text{vec}(\hat{Y})_i \pm t_{1-\alpha/2}(R) \sqrt{(\hat{\sigma}^2)_i^R}$ is a $(1 - \alpha)$ -confidence interval for $\text{vec}(Y)_i$.

If the \tilde{Y} ’s are not normally distributed, then we propose a distribution-free Monte Carlo confidence interval using order statistics.

Suppose that (A1) holds. Let $\text{vec}(\tilde{Y})^{(1)}, \dots, \text{vec}(\tilde{Y})^{(R)}$ be i.i.d. samples with the same distribution as $\text{vec}(N)$. Call $\hat{Y}^{(j)}$ the result of applying SEA BLUE to $\tilde{Y}^{(j)}$. Given an index i , let $X_{(1)}^i, \dots, X_{(R)}^i$ be the order statistics of $|\text{vec}(\hat{Y}_i^{(1)})|, \dots, |\text{vec}(\hat{Y}_i^{(R)})|$. If $R \geq (1 - \alpha)/\alpha$, then $\text{vec}(\hat{Y})_i \pm X_{(\lceil (1-\alpha)(R+1) \rceil)}^i$ is a $(1 - \alpha)$ -confidence interval for $\text{vec}(Y)_i$, where $\lceil \cdot \rceil$ is the ceiling function.

In Propositions 5.2 and 5.2, the i.i.d. samples do not use any additional privacy loss budget, as they are simply realizations from the noise distribution and do not access the sensitive data.

Note that in Proposition 5.2, the Monte Carlo samples are used most efficiently when $(1 - \alpha)(R + 1)$ is an integer. For example, if $\alpha = .05$, then $R = 19$ is the smallest value that can be used in Proposition 5.2, and $R = 99$ or $R = 199$ are other sensible choices to reduce the expected width of the confidence interval.

In Table 1, we compare the average relative width of the confidence intervals from Propositions 5.1, 5.2, and 5.2 in the case of normally distributed data. Compared to using the exact covariance, Proposition 5.2 results in about a 10% increase in expected width when using the smallest value of $R = 19$, and this error can be reduced to $< 1\%$ by using $R = 199$. On the other hand, with Proposition 5.2, we can obtain $< 1\%$ error with only $R = 99$. Finally, when comparing the two Monte Carlo methods to each other, the distribution-free approach has 4% increased width compared to Proposition 5.2 with $R = 19$, and less than 1% increased width at $R = 99$. We see that these Monte Carlo methods offer a competitive alternative to Proposition 5.1.

	$R :$	19	99	199
Monte Carlo T / z		1.055	1.010	1.005
Monte Carlo DF / z		1.094	1.017	1.009
Monte Carlo DF / Monte Carlo T		1.038	1.007	1.005

Table 1: Mean relative width of the confidence intervals from Proposition 5.1 (z), Proposition 5.2 (Monte Carlo T), and Proposition 5.2 (Monte Carlo DF). Simulation was conducted with normally distributed random variables, $\alpha = .05$, and 100,000 replicates. In all settings, the Monte Carlo standard errors are $< .001$.

Remark 15. Suppose we want a confidence interval for multiple counts in Y . With either Proposition 5.2 or Proposition 5.2, the same R copies of $\hat{Y}^{(j)}$ can be re-used to make a confidence interval for each count of interest. On the other hand, using Proposition 5.1 requires running SEA BLUE for each count of interest that has a different covariance column. So, while the Monte Carlo methods have an overhead of R runs, the number of runs need not scale with the number of desired confidence intervals.

Remark 16 (Refining Confidence Intervals). While SEA BLUE does not use assumptions of non-negativity and integrality, we can use these constraints to refine any of the confidence intervals presented earlier in this section without affecting the coverage. For example, if it is known that the true values are non-negative integers, then if the initial confidence interval is $[-1.5, 12.7]$ a refined confidence interval would be $[0, 12]$. In general, if $[a(\tilde{Y}), b(\tilde{Y})]$ is a $(1 - \alpha)$ -confidence interval for a count Y_v^S , then so is $[a^*(\tilde{Y}), b^*(\tilde{Y})]$, where $a^*(\tilde{Y}) = \max\{0, \lceil a(\tilde{Y}) \rceil\}$ and $b^*(\tilde{Y}) = \lfloor b(\tilde{Y}) \rfloor$. This method of refining always preserves the coverage and can only reduce the width of the intervals.

6. Simulations and Application to Census Products

In this section, we empirically explore the properties of SEA BLUE in both simulations as well as through an application to two Census products. The code for these experiments is available at <https://github.com/JordanAwan/SeaBlue>.

6.1 Description of Census Products and Simulation Setup

The two U.S. Census Bureau products that we consider are the noisy measurement files for Redistricting Data (PL94) and Demographic and Housing Characteristics (DHC).

PL94 provides tables summarizing population totals as well as counts for race, Hispanic origin, and voting age for each geographic areas as well as housing unit counts by occupancy status and population counts in group quarters. Structurally, the noised PL94 data consists of four variables, sized 2, 2, 8, and 63. The data for each PL94 geography includes contingency tables for 1) the values for each individual variable by itself, 2) values for the full 4-variable cross combination of all variables, 3) three 2-way crosses (both size-2 variables crossed together, and the size-8 variable crossed against each of the size-2 variables), and 4) a single 3-way cross (size-2, 2, and 8 variables crossed together).

The DHC product provides tables on the following variables: age, sex, race, Hispanic or Latino origin, household type, family type, relationship to householder, group quarters population, housing occupancy, and housing tenure. The contingency tables of the DHC that we used as input included 1) a full-cross of five variables, of sizes 2, 2, 42, 63, and 116, 2) a 2-way cross of the size-2 variables, and 3) a 2-way marginal variable by itself. The DHC data set included several contingency tables that are a cross of variables against “bins” of another variable, which weren’t directly usable under our current approach. Instead of crossing sex against each possible age value, the data crosses sex against a modified age variable that expresses the counts within large age ranges, e.g., 0-18, 19-36, etc. Our SEA BLUE approach for DHC summed the counts for these sex-cross-binned-age, eliminate these unusable “binned” counts, in order to get the sex variable by itself.

We use the 2010 demonstration products for PL94 and DHC, rather than the official 2020 releases. These demonstration products were protected using DP in a similar manner as the official 2020 releases and were produced to help data users understand the DP mechanism proposed for the 2020 products. We use these demonstration products in our simulations because we are able to use the Census 2010 Summary File 1 (SF1) counts as the approximately true values for the same geographies. SF1 is the official 2010 tabulation of Census counts and these SF1 values are used to assess mean squared error (MSE) and interval coverage, whereas no analogous product was available for the 2020 release. We also used the 2010 DHC demonstration product to measure timing and memory use for a large contingency table, but we unfortunately did not have the corresponding values to measure the accuracy of DHC results as the 2010 SF1 tables did not match the set of noisy tables in the DHC demonstration product.

We used a sample of 40 Census blocks from Rhode Island for PL94, and for DHC, we used state-level counts from Rhode Island. For PL94, our replications were simulated by taking a sample of geographies, since each geography only has one set of counts. Note that within each product, all geographies have the same set of tables, so timing and memory should be similar for different geographies (e.g., SEA BLUE applied to PL94 for a block in Rhode Island should see similar performance to SEA BLUE applied to a county in California). We also ran SEA BLUE on West Virginia state and its 55 counties for PL94, treating county as another variable. We initially chose Rhode Island because of its small size, however, in the end the runtime for SEA BLUE and matrix projection do not depend on the size. We chose West Virginia for multiple-county execution because it was the state with the largest number of counties with uniform variance in the noisy PL94 product.

For PL94, all tables at a single geography satisfy assumptions (A1)-(A3). However, DHC does not directly fit within our framework as the Age variable is binned at different thresholds in different tables. Due to this, we marginalized out the Age variable whenever binning occurred and used these tables as the input to SEA BLUE and matrix projection. After this pre-processing, assumptions (A1)-(A3) hold for DHC as well. For the West Virginia counties application, each county had the same privacy loss budget, so (A1)-(A3) held for this example as well, including County as an additional variable; however in other states, counties received different privacy loss budgets, which results in unequal variances.

Simulations were conducted in Python, where a custom script was written to implement SEA BLUE, and `numpy` functions were used to perform the matrix operations needed to implement the matrix projection. Tests were run on a Windows laptop with an Intel i7-7820HQ CPU (4 cores / 2.90GHz) and 8 GB RAM. No parallelization was implemented.

In our controlled settings, we generate data for a “ $k \times k$ ” table for $k \in 3, 4, 5, 6$, meaning that there are k variables which each have k levels, and we observe all possible noisy margins. For a single geography we generate the detailed table using a zero inflated Poisson distribution, and aggregate back to the total count. Noise is added using a normal distribution to each count type independently.

Size		Time		Memory	
Dimension	Counts	SEA BLUE	Projection	SEA BLUE	Projection
3×3	64	0.006	0.01	2.52	6.41
4×4	625	0.05	0.024	2.53	12.92
5×5	7,776	0.90	60.45	4.15	947.54
6×6	117,649	18.70	N/A	30.98	N/A
PL94 Block	5,184	0.49	70.55	4.30	421.77
PL94 Multi	290,304	34.36	N/A	116.16	N/A
DHC State	2,897,856	299.07	N/A	1186.35	N/A

Table 2: Comparison of time (in seconds) and memory (in MiB) performance for SEA BLUE versus full projection on tables of various sizes. Time results are averaged over 40 replicates, and memory is averaged over 10 replicates. PL94 block numbers are based on counts from 40 Rhode Island Census blocks, PL94 Multi numbers are based on counts for West Virginia state and its 55 counties. DHC state is based on Rhode Island state counts. Projection did not run for 6×6 , DHC, or multi-geography PL94 due to memory limitations. Only a single instance of SEA BLUE was run on DHC as a test of feasibility.

6.2 Runtime and Memory Usage

We implement both SEA BLUE and the matrix projection and apply them to problems of varying sizes. At each size, we measure the runtime as well as the memory usage as either can potentially be a limiting factor. The results are found in Table 2.

The second main column of Table 2 contains the runtime in seconds, averaged over 40 replicates. For the 3×3 and 4×4 cases, SEA BLUE and projection both run in less than a second. However, for the 5×5 problem, SEA BLUE still runs in under 1s, whereas projection takes 60s. Finally, SEA BLUE takes about 19s to process the 6×6 table, whereas projection was not able to be run for this table size due to memory limitations.

The final column of Table 2 records the maximum memory required for the method, measured in mebibytes (MiB), averaged over 10 replicates (each memory result never varied beyond 3 KiB). We see a substantially improved scaling for memory with SEA BLUE, requiring approximately 4MiB for the 5×5 size compared to nearly 1GiB for the projection. Finally, the 6×6 only required 31MiB for SEA BLUE, whereas the projection was unable to run on our machine due to memory limitations. An error message said that it needed to allocate approximately 103 gibibytes (GiB) for the projection matrix.

We also applied SEA BLUE and the matrix projection to the Decennial Census 2010 demonstration product of PL94, which is protected with zero-concentrated differential privacy. Among the simulations, the PL94 contingency table count is most similar in size to the 5×5 case. Performance-wise, we see SEA BLUE has similar performance on PL94 as in the 5×5 case: the runtime of SEA BLUE for PL94 is about a 95% decrease compared to the matrix projection and the memory is a 99% decrease compared to the projection. When processing many geographies, this difference in performance becomes practically sig-

Experiment	Counts	SEA	Proj	SEA To Proj
One Marginal	4	0.8132	0.8131	0.0000
All Marginals	16	0.8191	0.8191	0.0002
All 2 Way	96	0.7995	0.7880	0.0125
All 3 Way	256	0.7983	0.7606	0.0408
Detailed	256	0.8159	0.8069	0.0109
All Counts 1 Var	500	0.7932	0.7456	0.0451

Table 3: Performance versus true projection in various 4×4 scenarios where (A3) is violated, aggregated over 100 replicates. Three MSE levels are shown: the SEA BLUE estimates to the true counts, projection values to the true counts, and the SEA BLUE estimates to the projection.

nificant. When running SEA BLUE for PL94 for West Virginia state and its 55 counties simultaneously, the projection failed to run while attempting to allocate 624 GiB of memory.

For the larger DHC product, SEA BLUE’s computation took less than 1.2 GiB of memory and less than 30min, while projection failed after attempting to allocate 10TiB of memory. We only ran a single instance of SEA BLUE on DHC as a test of feasibility.

6.3 Unequal Variances

Recall that SEA BLUE is the best linear unbiased estimate only if the assumptions (A1)-(A3) are met. In the PL94 demonstration product, (A1)-(A3) hold for any single geography. However, it may be the case that different counties in the same state have different variances. In this section, we explore through simulations how much error is incurred by using SEA BLUE when the assumption of equal variance within tables (A3) is violated.

We conducted a variety of experiments using a 4×4 simulated contingency tables where a varying number of tables violated (A3). The default variance for each count is set to 2, while in the listed tables the variance was set to be either 1, 2, or 3, uniformly at random. Experiments were conducted where one marginal table, all marginal tables, all 2-way tables, all 3-way tables, the detailed query, or all counts containing a single variable each received non-equal variance. In each case, the MSE was calculated using 100 replicates (because of the small size of the 4×4 tables, we could afford additional replicates here).

The higher proportion of counts that fall within non-uniform tables, the greater the error that SEA BLUE incurs compared to the ideal projection. Nevertheless, these mild violations of (A3) still result in near-optimal MSE (off by about 5% in the worst case). For reference, the MSE of the original counts is approximately 2 in all settings.

6.4 Confidence Interval Validation

In this section, we numerically validate the confidence interval methods of Section 5. In particular, we estimate the coverage and width on both simulated data and the PL94 Census data product. Confidence interval testing against DHC was not possible because reference values were not available for the DHC tables.

Scenario	Exact	Monte Carlo T	Monte Carlo DF
3×3	0.9531	0.9480	0.9582
4×4	0.9520	0.9516	0.9541
5×5	0.9510	0.9500	0.9520
6×6	0.9503	0.9498	0.9523
PL94 Block	0.9505	0.9510	0.9543
PL94 Multi	0.9482	0.9490	0.9733

Table 4: Coverage shown for various scenarios based on 40 replicates using 20 rounds of estimation per trial. $N \times N$ replicates are computed on replicate-generated random noisy data. PL94 runs are based on 40 block-level records from the state of Rhode Island. $\alpha = 0.95$. The Exact method is from Proposition 5.1, Monte Carlo T is from Proposition 5.2, and Monte Carlo DF is from Proposition 5.2.

Scenario	Initial		Exact		T Monte Carlo		DF Monte Carlo	
	Raw	Clip	Raw	Clip	Raw	Clip	Raw	Clip
3×3	5.544	3.413	3.601	2.240	3.772	2.361	3.963	2.546
4×4	5.544	3.506	3.548	2.254	3.728	2.407	3.915	2.574
5×5	5.544	3.549	3.514	2.242	3.694	2.400	3.884	2.568
6×6	5.544	3.568	3.491	2.233	3.669	2.392	3.859	2.5605
PL94 Block	43.802	21.765	19.342	9.319	20.300	9.831	21.349	10.358
PL94 Multi	31.230	16.062	23.740	11.997	23.772	12.013	27.257	13.832
DHC State	13.044	6.421	12.887	6.121	13.544	6.371	14.241	6.727

Table 5: Comparison of confidence interval widths for tables that were supplied as input. Intervals included for initial noise, SEA BLUE variance calculation, T-based, and distribution-free Monte Carlo estimations. Results averaged from 40 replicates using 20 rounds of Monte Carlo estimation per replicate. $N \times N$ replicates are computed on replicate-generated random noisy data. PL94 runs are based on 40 block-level records from the state of Rhode Island. DHC numbers are computed from Rhode Island state-level counts.

In Table 4, we measure the empirical coverage of our three confidence interval methods on both simulated data and on the PL94 product averaged over 40 block-level records from Rhode Island. Note that in the simulated tables, the noise distribution was continuous Gaussian whereas for PL94 the noise distribution was discrete Gaussian. Even with the slight deviation from normality, we still approach the 95% nominal coverage for all of our proposed methods. The over-coverage of the distribution-free approach is to be expected, since this approach makes fewer assumptions on the noise distribution.

In Table 5, we compare the width of our confidence interval procedures. In each case, we consider both the raw and clipped intervals, where the clipping is according to Remark 16. For the simulated examples, the input and output tables are identical; for PL94 and DHC, the output is all tables possible from the marginals, but the input is the subset of tables

Scenario	Exact		T Monte Carlo		DF Monte Carlo	
	Raw	Clip	Raw	Clip	Raw	Clip
PL94 Block	20.9643	10.0505	22.0136	10.5767	23.1337	11.1392
PL94 multi	25.0119	12.8369	25.04334	12.8535	28.7145	14.7892
DHC State	26.5573	13.6419	27.8967	14.3599	29.3217	15.1036

Table 6: Comparison of confidence interval widths for tables that were *not* supplied as input. Intervals included for initial noise, SEA BLUE variance calculation, T-based, and distribution-free Monte Carlo intervals size averaged from 40 replicates using 20 rounds of Monte Carlo estimation per replicate. PL94 runs are based on 40 block-level records from Rhode Island. DHC numbers are computed from Rhode Island state-level counts.

actually supplied in the Census data. In all cases, our clipped intervals are significantly smaller than the initial clipped intervals. In the simulated tables and for PL94, we also see the unclipped confidence interval width is significantly reduced. Interestingly, for DHC the unclipped width is not significantly reduced, but the clipped width is much smaller after SEA BLUE. Clipping has the greatest impact when there are many true counts near zero; we suspect that it must be the case that in DHC, there are many such counts, which are disproportionately affected by using the clipping.

Similarly, Table 6 compares the widths of our confidence intervals but for counts that do not have original noisy estimates. Note that this only applies to PL94 and DHC since the simulated tables have noisy counts for all margins. Compared to Table 5, these widths are larger, which makes sense since no direct measurement was available for these counts. Otherwise, we see similar behavior in Table 6: the clipped intervals are significantly shorter, and the exact variance method is the smallest, followed by the Monte Carlo T-interval, with the distribution-free Monte Carlo interval being the largest.

7. Discussion and Future Work

Under assumptions (A1)-(A3), our SEA BLUE method offers a scalable implementation of the best linear unbiased estimate of true counts, subject to self-consistency constraints. As it is a linear procedure, we can easily produce confidence intervals with provable coverage guarantees. Compared to a naive matrix projection, our method scales significantly better in terms of both runtime and memory requirements and we show that it can be applied to large Census products.

The proposed method was initially designed to process a subset of the Decennial Census products at a single geography (e.g., a single county or a single block), but we were able to also apply our method to incorporate two levels of geography, such as a state with its counties. However, including more than two levels of geography would require an extension of our framework since, for example, different states would have a different set of counties so the “county” variable changes depending on the state. We suspect that a modification of our collection step and down pass approach can accommodate adding in multiple geographies, but we leave this as future work.

While Section 6.3 showed that SEA BLUE is robust to mild violations to (A3), in general it could be arbitrarily far from the projection estimates when there are more egregious violations of the (A3). Performance could be improved by modifying Equations (4) and (5) to handle unequal variances. However, even with this modification, SEA BLUE would not in general produce the projection when (A3) does not hold. As mentioned in the discussion before Theorem 6, when only (A1) and (A2) hold, the collection step produces an optimal estimator using the estimates that are “from below” and “within-table”; thus, this step always improves over the original noisy counts, but could be arbitrarily worse than the optimal estimator (Example 7 offers an example where SEA BLUE produces an estimator with variance 8, where the projection estimator would have variance at most 3). Future work could investigate whether SEA BLUE can be modified to offer more robustness to violations of (A3).

While (A2) is more commonly used in similar literature, and is also satisfied by the 2020 census products, it is also worth considering extensions of SEA BLUE that can handle violations of (A2). As noted in Section 4.1, it seems that without (A2), it may be necessary to track the covariances of the estimates, which would significantly increase the computation and memory requirements. Future work may investigate whether weaker versions of (A2) or modifications to the SEA BLUE procedure can maintain computational efficiency while allowing for dependence in the noise distribution.

As mentioned in Remark 2, this paper does not directly address *invariants*, which are counts observed without noise. We believe that the procedure can be modified to incorporate other linear equality invariants such as these as well.

SEA BLUE only accommodates linear equality constraints in order to remain a linear and unbiased procedure. This has two consequences: 1) SEA BLUE can result in negative and non-integer-valued estimates and 2) SEA BLUE is expected to have higher MSE than the Top Down Algorithm, which incorporates non-negativity and integrality (see McCartan et al. (2023) for a comparison of the Top Down Algorithm against the collection step). As noted in Remark 16, our confidence intervals can be refined to incorporate additional constraints without affecting coverage. Nevertheless, accuracy could be improved by enforcing additional constraints in the estimation procedure. An important direction for future work would be to develop computationally efficient post-processing methods which incorporate all constraints, preserve (at least approximate) unbiasedness, and also give valid confidence regions.

The results of Section 5.1 are focused on coordinate-wise confidence intervals, but one may also be interested in confidence regions that simultaneously cover multiple counts. Extension of Propositions 5.1 and 5.2 to give confidence ellipsoids should be straightforward by leveraging the normality assumption. An extension of Proposition 5.2 is less straightforward, but could potentially be done using the order statistics of a depth statistic (e.g., Yeh and Singh, 1997; Xie and Wang, 2022; Awan and Wang, 2024). However, it is not clear whether these approaches are computationally tractable, especially when making a confidence set for all counts, simultaneously. We leave it to future researchers to investigate the details and feasibility of such confidence regions.

Our simulations did not implement parallelization techniques in the implementation of SEA BLUE. For large databases, the ability to parallelize could enable even larger tables

to be analyzed by SEA BLUE. It would be worth developing R/Python packages that implement the SEA BLUE method with parallel optimizations incorporated.

Acknowledgments

This work was supported by Department of Commerce Contract 1331L523D130S0003 Task Order 1333LB23F00000196. Jordan Awan’s research was also supported in part by NSF grant no. SES-2150615 to Purdue University. The authors are grateful to Nathaniel Cady who contributed early numerical work to the project, and to Ryan Janicki, Philip Leclerc, Mikaela Meyer, as well as the anonymous reviewers for their feedback which improved the presentation of this manuscript.

Appendix A. Summary of Notation

We include in Table 7 a summary of the notation used in this paper, which is defined in Section 2.2.

Notation	Context	Definition/Interpretation
Y		true counts
\tilde{Y}		observed noisy counts
\dot{Y}		intermediate estimate of Y
\hat{Y}		final estimate of Y from SEA BLUE
\mathcal{A}		set of variables in the contingency tables
\mathcal{O}	$\mathcal{O} \subset 2^{\mathcal{A}}$	set of observed tables
\mathcal{D}	$\mathcal{D} \subset 2^{\mathcal{A}}$	set of desired tables
$ A $	$A \in \mathcal{A}$	number of levels that variable A can take on
\tilde{Y}^S	$S \subset \mathcal{A}$	a table of noisy counts for variables in S
V^S	$S \subset \mathcal{A}$	the set of indices for Y^S
\tilde{Y}_v^S	$v \in V^S, S \subset \mathcal{A}$	the single noisy count in \tilde{Y}^S indexed by v
V_{\bullet}^S	$S \subset \mathcal{A}$	same as V^S , but each coordinate also can be “•”
\tilde{Y}_v^S	$v \in V_{\bullet}^S, S \subset \mathcal{A}$	index with v , summing out variables with “•”
$v(\bar{S})$	$\bar{S} \supset S, v \in V^S, S \subset \mathcal{A}$	$v(\bar{S}) \in V_{\bullet}^{\bar{S}}$ agrees with v on S and has “•” on $\bar{S} \setminus S$
$v[\underline{S}]$	$\underline{S} \subset S, v \in V^S, S \subset \mathcal{A}$	$v[\underline{S}] \in V^{\underline{S}}$ extracts entries of v corresponding to \underline{S}
$v[\underline{S}](S)$	$\underline{S} \subset S, v \in V^S, S \subset \mathcal{A}$	$v[\underline{S}](S) \in V_{\bullet}^{\underline{S}}$ agrees with v on \underline{S} and has “•” on $S \setminus \underline{S}$.

Table 7: Summary of notation from Section 2.2.

Appendix B. Proofs and Technical Details

We review some properties of linear transformations of random variables: let $A \in \mathbb{R}^{n \times k}$ and $b \in \mathbb{R}^n$ and let $Y \in \mathbb{R}^k$ be a random variable with mean μ and covariance Σ . Then $\mathbb{E}AY = A\mu$ and $\text{Cov}(AY) = A\Sigma A^\top$. In particular, if $Y \sim N(\mu, \Sigma)$. Then $AY + b \sim N(A\mu + b, A\Sigma A^\top)$.

Theorem 3. Let $X \in \mathbb{R}^n$ be a random vector with mean μ and covariance Σ , where it is known that $\mu \in S_b = \{x \in \mathbb{R}^n \mid Ax = b\}$. Let v be such that $b = Av$ and call $S_0 = \{x \in \mathbb{R}^n \mid Ax = 0\}$. Call $\hat{\mu} = \text{Proj}_{S_0}^{\Sigma^{-1}}(X) + (\text{Proj}_{S_0^\perp}^\Sigma)^\top v$. Then,

- $\hat{\mu}$ is the BLUE for μ . Furthermore, $\text{Cov}(\hat{\mu}) = \text{Proj}_{S_0}^{\Sigma^{-1}} \Sigma (\text{Proj}_{S_0}^{\Sigma^{-1}})^\top$.
- If $X \sim \mathcal{N}(\mu, \Sigma)$, then $\hat{\mu}$ is the uniformly minimum variance unbiased estimator for μ ; furthermore, $\hat{\mu} \sim \mathcal{N}(\mu, \text{Proj}_{S_0}^{\Sigma^{-1}} \Sigma (\text{Proj}_{S_0}^{\Sigma^{-1}})^\top)$.

If the rows of A are linearly independent, then

$$\text{Proj}_{S_0}^{\Sigma^{-1}} = \left(I - \Sigma A^\top (A \Sigma A^\top)^{-1} A \right) \quad \text{and} \quad (\text{Proj}_{S_0^\perp}^\Sigma)^\top = \Sigma A^\top (A \Sigma A^\top)^{-1} A.$$

Proof Since $A\mu = b$, we know that $b \in \text{columnSpace}(A)$. So, there exists v such that $b = Av$. Now, we can say that $\mu - v \in \text{nullSpace}(A)$. We then transform our problem as follows. Let $X' = \Sigma^{-1/2}(X - v)$, $\mu' = \Sigma^{-1/2}(\mu - v)$. Note that $A\Sigma^{1/2}\mu' = 0$, so that $\mu' \in \text{nullSpace}(A\Sigma^{1/2})$. Let M be a matrix whose columns form a basis for $\text{nullSpace}(A\Sigma^{1/2})$. Then there exists a vector β such that $\mu' = M\beta$. So, we have rephrased our problem as $X' = M\beta + e$, where e is a mean-zero vector, with covariance I . By the Gauss-Markov Theorem, the BLUE for β is $\hat{\beta} = (M^\top M)^{-1} M^\top X'$, which implies that the BLUE for μ' is $\text{Proj}_{\text{nullSpace}(A\Sigma^{1/2})}^I X'$. Converting back to our original problem, we have that the BLUE for μ is $\hat{\mu} = \Sigma^{1/2} \text{Proj}_{\text{nullSpace}(A\Sigma^{1/2})}^I (\Sigma^{-1/2}(X - v)) + v$, which can be simplified to $\hat{\mu} = \text{Proj}_{S_0}^{\Sigma^{-1}}(X) + (\text{Proj}_{S_0^\perp}^\Sigma)^\top v$. Similarly, if $X \sim \mathcal{N}(\mu, \Sigma)$, then $\hat{\beta}$ is the uniformly minimum variance unbiased estimator for β , which implies that $\hat{\mu}$ is the uniformly minimum variance unbiased estimator for μ . \blacksquare

Lemma 5. If (A1) and (A2) hold, then the variance of the entries of the intermediate estimate \dot{Y}_v^S , denoted by $(\dot{\sigma}^2)_v^S$, is given by

$$(\dot{\sigma}^2)_v^S := \text{Var}(\dot{Y}_v^S) = \left(\sum_{\bar{S} \in \mathcal{O} \mid \bar{S} \supset S} \left(\text{Var}(\tilde{Y}_{v(\bar{S})}^{\bar{S}}) \right)^{-1} \right)^{-1}. \quad (2)$$

Proof To simplify notation, we write $\sum_{\bar{S} \supset S}$ in place of $\sum_{\bar{S} \in \mathcal{O} \mid \bar{S} \supset S}$. Recall that \dot{Y}_v^S can be expressed as

$$\dot{Y}_v^S = \frac{\sum_{\bar{S} \supset S} \left(\text{Var}(\tilde{Y}_{v(\bar{S})}^{\bar{S}}) \right)^{-1} \tilde{Y}_{v(\bar{S})}^{\bar{S}}}{\sum_{\bar{S} \supset S} \left(\text{Var}(\tilde{Y}_{v(\bar{S})}^{\bar{S}}) \right)^{-1}}.$$

Using the fact that all of the \tilde{Y} values are independent, we calculate the variance as

$$\text{Var}(\dot{Y}_v^S) = \frac{\sum_{\bar{S} \supset S} \left(\text{Var}(\tilde{Y}_{v(\bar{S})}^{\bar{S}}) \right)^{-2} \text{Var}(\tilde{Y}_{v(\bar{S})}^{\bar{S}})}{\left(\sum_{\bar{S} \supset S} \left(\text{Var}(\tilde{Y}_{v(\bar{S})}^{\bar{S}}) \right)^{-1} \right)^2}$$

$$\begin{aligned}
 &= \frac{\sum_{\bar{S} \supset S} \left(\text{Var}(\tilde{Y}_{v(\bar{S})}^{\bar{S}}) \right)^{-1}}{\left(\sum_{\bar{S} \supset S} \left(\text{Var}(\tilde{Y}_{v(\bar{S})}^{\bar{S}}) \right)^{-1} \right)^2} \\
 &= \left(\sum_{\bar{S} \supset S} \left(\text{Var}(\tilde{Y}_{v(\bar{S})}^{\bar{S}}) \right)^{-1} \right)^{-1}.
 \end{aligned}$$

■

Theorem 6. *[Optimality for the Total] Assume that (A1)-(A3) hold. Then, \dot{Y}^\emptyset is the BLUE of the total count Y^\emptyset based on the values in \tilde{Y}^S for $S \in \mathcal{O}$, subject to self-consistency.*

Proof Let \mathcal{T} be the affine space consisting of all possible linear unbiased estimates of Y^\emptyset , based on the counts in \tilde{Y}^S for $S \in \mathcal{O}$. Note that $\{\tilde{Y}_\bullet^S \mid S \in \mathcal{O}\}$ is a set of linearly independent members of \mathcal{T} , where the subscript “ \bullet ” is short hand to represent the vector in V_\bullet^S with “ \bullet ” in every entry. Then, there exists a basis vector b for \mathcal{T} containing \tilde{Y}_\bullet^S as the first entries for all $S \in \mathcal{O}$. Let Λ be the covariance matrix for b .

To see that Λ is positive definite, let ω be a nonzero vector. Then $\text{Var}(\omega^\top b) = \omega^\top \Lambda \omega$. Since \tilde{Y}_v^S are independent with nonzero variance for all $v \in V^S$ and $S \in \mathcal{O}$, and because the entries of b are linearly independent, it follows that there exists a nonzero vector $\tilde{\omega}$ such that $\omega^\top b = \tilde{\omega}^\top \text{vec}(\tilde{Y})$, where $\text{vec}(\tilde{Y}_v^S)$ represents a vectorization of the entries of \tilde{Y}_v^S for all $v \in V^S$ and $S \in \mathcal{O}$. Then, we have that

$$\text{Var}(\tilde{\omega}^\top \text{vec}(\tilde{Y})) = \sum_{S \in \mathcal{O}} \sum_{v \in V^S} \tilde{\omega}_v^S \text{Var}(\tilde{Y}_v^S) > 0,$$

where $\tilde{\omega}_v^S$ represents the entry of $\tilde{\omega}$ that corresponds with \tilde{Y}_v^S . So, we have that $\omega^\top \Lambda \omega = \text{Var}(\tilde{\omega}^\top \text{vec}(\tilde{Y})) > 0$, which implies that Λ is positive definite.

The best linear unbiased estimator for Y^\emptyset uses the inverse covariance weighting of b , which can be expressed in terms of matrix operations as,

$$\frac{\underline{1}^\top \Lambda^{-1} b}{\underline{1}^\top (\Lambda^{-1}) \underline{1}},$$

where $\underline{1}$ is the vector of all 1's of the appropriate dimension, which is the efficient generalized least squares solution (Anderson, 2008).

On the other hand, the collection estimate \dot{Y}^\emptyset can be expressed as

$$\frac{\left(\left(\frac{1}{\text{Var}(\tilde{Y}_\bullet^S)} \right)_{S \in \mathcal{O}}, 0, \dots, 0 \right) b}{\sum_{S \in \mathcal{O}} \frac{1}{\text{Var}(\tilde{Y}_\bullet^S)}},$$

where $\left(1/\text{Var}(\tilde{Y}_\bullet^S) \right)_{S \in \mathcal{O}}$ represents the vector of inverse-variances in the same order as the elements \tilde{Y}_\bullet^S appear in b .

To claim that the collection step results in the BLUE estimate of Y^\emptyset , we need to establish that

$$\frac{\left(\left(\frac{1}{\text{Var}(\tilde{Y}_\bullet^S)}\right)_{S \in \mathcal{O}}, 0, \dots, 0\right)}{\sum_{S \in \mathcal{O}} \frac{1}{\text{Var}(\tilde{Y}_\bullet^S)}} = \frac{\underline{1}^\top \Lambda^{-1}}{\underline{1}^\top \Lambda^{-1} \underline{1}}.$$

It suffices to show that

$$\left(\left(\frac{1}{\text{Var}(\tilde{Y}_\bullet^S)}\right)_{S \in \mathcal{O}}, 0, \dots, 0\right) \Lambda = \underline{1}^\top, \quad (6)$$

as this implies that

$$\left(\left(\frac{1}{\text{Var}(\tilde{Y}_\bullet^S)}\right)_{S \in \mathcal{O}}, 0, \dots, 0\right) = \underline{1}^\top \Lambda^{-1},$$

as well as

$$\sum_{S \in \mathcal{O}} \frac{1}{\text{Var}(\tilde{Y}_\bullet^S)} = \left(\left(\frac{1}{\text{Var}(\tilde{Y}_\bullet^S)}\right)_{S \in \mathcal{O}}, 0, \dots, 0\right) \underline{1} = \underline{1}^\top \Lambda^{-1} \underline{1}.$$

The remainder of the proof aims to establish (6).

Let Z be an element of b . Then, restricting attention to the column of Λ corresponding to Z , we need to show that

$$\sum_{S \in \mathcal{O}} \frac{\text{Cov}(Z, \tilde{Y}_\bullet^S)}{\text{Var}(\tilde{Y}_\bullet^S)} = 1.$$

We can express $Z = \sum_{S \in \mathcal{O}} \sum_{v \in V^S} a_v^S \tilde{Y}_v^S$ for some weights a_v^S . We claim that for Z to be an unbiased estimator for $Y^\emptyset = \sum_{e \in V^{\mathcal{A}}} Y_e$, it must be that $\sum_{S \in \mathcal{O}} \sum_{v \in V^S} a_v^S = 1$, for all $e \in V^{\mathcal{A}}$. To see this, we calculate

$$\begin{aligned} \mathbb{E}Z &= \sum_{S \in \mathcal{O}} \sum_{v \in V^S} a_v^S \mathbb{E}\tilde{Y}_v^S \\ &= \sum_{S \in \mathcal{O}} \sum_{v \in V^S} a_v^S \sum_{e[S]=v} Y_e \\ &= \sum_{e \in V^{\mathcal{A}}} Y_e \sum_{S \in \mathcal{O}} \sum_{v \in V^S, v=e[S]} a_v^S, \end{aligned}$$

where in the second line we used the fact that $\mathbb{E}\tilde{Y}_v^S = Y_v^S = \sum_{e[S]=v} Y_e$. We see that in order for $\mathbb{E}Z = Y^\emptyset = \sum_{e \in V^{\mathcal{A}}} Y_e$, it must be that $\sum_{S \in \mathcal{O}} \sum_{v \in V^S} a_v^S = 1$ as claimed.

Using our expansion of Z , we can now calculate

$$\sum_{S \in \mathcal{O}} \frac{\text{Cov}(Z, \tilde{Y}_\bullet^S)}{\text{Var}(\tilde{Y}_\bullet^S)} = \sum_{S \in \mathcal{O}} \sum_{v \in V^S} \frac{a_v^S \text{Var}(\tilde{Y}_v^S)}{|V^S| \text{Var}(\tilde{Y}_v^S)} \quad (7)$$

$$= \sum_{S \in \mathcal{O}} \sum_{v \in V^S} \frac{a_v^S}{|V^S|} \quad (8)$$

$$= \sum_{S \in \mathcal{O}} \sum_{e \in V^{\mathcal{A}}} \frac{|V^S|}{|V^{\mathcal{A}}|} \sum_{\substack{v \in V^S \\ v=e[S]}} \frac{a_v}{|V^S|} \quad (9)$$

$$= \frac{1}{|V^{\mathcal{A}}|} \sum_{e \in V^{\mathcal{A}}} \sum_{S \in \mathcal{O}} \sum_{\substack{v \in V^S \\ v=e[S]}} a_v^S \quad (10)$$

$$= \frac{1}{|V^{\mathcal{A}}|} \sum_{e \in V^{\mathcal{A}}} 1 \quad (11)$$

$$= 1, \quad (12)$$

where in (7) we use the assumption that the variance is constant within tables (A3); in (9), the factor $|V^{\mathcal{A}}|/|V^S|$ corrects for over-counting, since there are $|V^{\mathcal{A}}|$ e 's and $|V^S|$ v 's, and each corresponds to an equal number. We see that for every column, $\Lambda_{\cdot,i}$, we have $((1/\text{Var}(\tilde{Y}^S)), 0, \dots, 0)\Lambda_{\cdot,i} = 1$ and conclude that \tilde{Y}^\emptyset is the BLUE for Y^\emptyset . \blacksquare

Theorem 8. *[Optimality of SEA BLUE] Assume (A1)-(A3) hold. Then, \hat{Y}^S resulting from applying the down pass to the output of the collection step is the best linear unbiased estimator of Y^S based on the values in \tilde{Y} , under self-consistency.*

Proof For any $S \in 2^{\mathcal{A}} \setminus \mathcal{O}$, artificially set $\tilde{Y}_v^S = 0$ for all $v \in V^S$ and $\sigma^2(\tilde{Y}_v^S) = \infty$. We use the convention that for any non-negative number a , $a/\infty = 0$.

Let $S \in \mathcal{D}$ be chosen, and assume that for all $T_i \subset S$ such that $|T_i| = |S| - 1$, we have the BLUE estimates \hat{Y}^{T_i} (by Proposition 6, we can compute \hat{Y}^\emptyset ; for induction we can assume that we have access to margins “above” S). As noted by Hay et al. (2010, Proof of Theorem 3), when holding \hat{Y}^{T_i} fixed, the values of \hat{Y}^S are independent of \tilde{Y}^R whenever $R \not\supset S$, and the BLUE \hat{Y}^S can be expressed as the solution to the following optimization problem:

$$\begin{aligned} & \text{minimize} \quad \sum_{R \supset S} \sum_{v \in V^R} (\hat{Y}_v^R - \tilde{Y}_v^R)^2 / \sigma^2(\tilde{Y}_v^R). \\ & \text{subject to} \quad \hat{Y}_{v(\bar{R})}^{\bar{R}} = \hat{Y}_v^{\bar{R}}, \quad \text{for all } \bar{R} \supset R \supset S, |\bar{R}| = |R| + 1, \text{ and } v \in V^{\bar{R}}, \\ & \text{and } \hat{Y}_{v(S)}^S = \hat{Y}_v^{T_i}, \quad \text{for all } T_i \text{ and } v \in V^{T_i}. \end{aligned} \quad (13)$$

In the case that $S = \mathcal{A}$ (full detail), this simplifies to

$$\begin{aligned} & \text{minimize} \quad \sum_{v \in V^{\mathcal{A}}} (\hat{Y}_v^{\mathcal{A}} - \tilde{Y}_v^{\mathcal{A}})^2 / \sigma^2(\tilde{Y}_v^{\mathcal{A}}) \\ & \text{subject to} \quad \hat{Y}_{v(\mathcal{A})}^{\mathcal{A}} = \hat{Y}_v^{T_i}, \quad \text{for all } T_i \text{ and } v \in V^{T_i}, \end{aligned}$$

where we used the substitution $\hat{Y}^{\mathcal{A}} = \tilde{Y}^{\mathcal{A}}$, and we see that in this case, the solution to this minimization problem is the projection prescribed in the down pass step.

Now assume that $S \neq \mathcal{A}$. We write out the Lagrangian for (13):

$$\sum_{R \supset S} \sum_{v \in V^R} (\hat{Y}_v^R - \tilde{Y}_v^R)^2 / \sigma^2(\tilde{Y}_v^R) + \sum_{\substack{\bar{R} \supset R \supset S \\ |\bar{R}|=|R|+1}} \sum_{v \in V^R} \lambda_v^{R, \bar{R}} (\hat{Y}_{v(\bar{R})}^{\bar{R}} - \hat{Y}_v^R) + \sum_{T_i} \sum_{v \in V^{T_i}} \lambda_v^{T_i, S} (\hat{Y}_{v(s)}^S - \hat{Y}_v^{T_i}). \quad (14)$$

Note that this is a convex and smooth objective. To understand the solution, we examine the partial derivatives. Taking the derivative with respect to any λ just recovers the original equality constraints. For the other derivatives, we separate them into three cases.

Case 1: Let $v \in V^S$ and take the derivative of (14) with respect to \hat{Y}_v^S and set equal to zero:

$$2(\hat{Y}_v^S - \tilde{Y}_v^S) / \sigma^2(\tilde{Y}_v^S) - \sum_{\substack{\bar{S} \supset S \\ |\bar{S}|=|S|+1}} \lambda_v^{S, \bar{S}} + \sum_{T_i} \lambda_{v[T_i]}^{T_i, S} = 0, \quad (15)$$

where recall that by (A2), $\sigma^2(\tilde{Y}_v^S)$ does not depend on v .

Case 2: For any $R \supsetneq S$, $R \neq \mathcal{A}$ and for any $v \in V^R$, the derivative with respect to \hat{Y}_v^R gives:

$$2(\hat{Y}_v^R - \tilde{Y}_v^R) / \sigma^2(\tilde{Y}_v^R) - \sum_{\substack{\bar{R} \supset R \\ |\bar{R}|=|R|+1}} \lambda_v^{R, \bar{R}} + \sum_{\substack{R \subset \underline{R}, \underline{R} \supset S \\ |\underline{R}|=|R|-1}} \lambda_{v[\underline{R}]}^{R, \underline{R}} = 0. \quad (16)$$

Case 3: In the case of \mathcal{A} and $v \in V^{\mathcal{A}}$, the derivative with respect to $\hat{Y}_v^{\mathcal{A}}$ gives:

$$2(\hat{Y}_v^{\mathcal{A}} - \tilde{Y}_v^{\mathcal{A}}) / \sigma^2(\tilde{Y}_v^{\mathcal{A}}) + \sum_{\substack{|\bar{S}|=|\mathcal{A}|-1 \\ \bar{S} \supset S}} \lambda_{v[\bar{S}]}^{\bar{S}, \mathcal{A}} = 0. \quad (17)$$

Together, all of the instances of (15), (16), and (17), along with the self-consistency constraints describe the solution space. In what follows, we combine these equations to show that the solution space can be decomposed into a down pass to obtain \hat{Y}^S and a projection to compute the other entries.

Let $v \in V^S$ be held fixed. Choose $R \supset S$ such that $|R| = |S| + 1$. Summing up (16) over all $w \in V^R$ which satisfy $w[S] = v$, and dividing through by $|V^R|/|V^S|$ (which is the number of equations being added), gives

$$\frac{|V^S|}{|V^R|} 2(\hat{Y}_v^S - \tilde{Y}_{v(R)}^R) \sigma^2(\tilde{Y}_v^R) - \sum_{\substack{\bar{R} \supset R \\ |\bar{R}|=|R|+1}} \frac{|V^S|}{|V^R|} \sum_{\substack{w \in V^R \\ w[S]=v}} \lambda_w^{\bar{R}, R} + \lambda_v^{S, R} = 0.$$

Using the fact that $\sigma^2(\tilde{Y}_{v(R)}^R) = (|V^R|/|V^S|) \sigma^2(\tilde{Y}_v^R)$ and introducing the notation $\bar{\lambda}_v^{\bar{R}, R} = \frac{|V^S|}{|V^R|} \sum_{\substack{w \in V^R \\ w[S]=v}} \lambda_w^{\bar{R}, R}$, we simplify this as:

$$2(\hat{Y}_v^S - \tilde{Y}_{v(R)}^R) / \sigma^2(\tilde{Y}_{v(R)}^R) - \sum_{\substack{\bar{R} \supset R \\ |\bar{R}|=|R|+1}} \bar{\lambda}_v^{\bar{R}, R} + \lambda_v^{S, R} = 0. \quad (18)$$

Adding (18) to (15) for all $R \supset S$ such that $|R| = |S| + 1$, we see that the negative λ 's in (15) cancel with the positive λ 's in (18). This gives:

$$2 \left(\hat{Y}_v^S \sum_{\substack{R \supset S \\ |R|=|S|+1}} \frac{1}{\sigma^2(\tilde{Y}_{v(R)}^R)} - \sum_{\substack{R \supset S \\ |R|=|S|+1}} \frac{\tilde{Y}_{v(R)}^R}{\sigma^2(\tilde{Y}_{v(R)}^R)} \right) - \sum_{\substack{\bar{R} \supset R \supset S \\ |\bar{R}|=|R|+1=|S|+2}} \bar{\lambda}_v^{R, \bar{R}} + \sum_{T_i} \lambda_{v[T_i]}^{T_i, S} = 0. \quad (19)$$

Similarly, we next choose $R \supset S$ such that $|R| = |S| + 2$, sum up (16) over all $w \in V^R$ such that $w[S] = v$, and divide by $|V^R|/|V^S|$:

$$2(\hat{Y}_v^S - \tilde{Y}_{v(R)}^R)/\sigma^2(\tilde{Y}_{v(R)}^R) - \sum_{\substack{\bar{R} \supset R \\ |\bar{R}|=|R|+1}} \bar{\lambda}_v^{R, \bar{R}} + \sum_{\substack{R \subset R, R' \supset S \\ |\underline{R}|=|R|-1}} \bar{\lambda}_v^{R, R} = 0, \quad (20)$$

where the last term is obtained by simplifying the following:

$$\begin{aligned} \sum_{\substack{R \subset R, R \supset S \\ |\underline{R}|=|R|-1}} \frac{|V^S|}{|V^R|} \sum_{\substack{w \in V^R \\ w[S]=v}} \lambda_{w[\underline{R}]}^{R, R} &= \sum_{\substack{R \subset R, R \supset S \\ |\underline{R}|=|R|-1}} \frac{|V^S|}{|V^R|} \sum_{\substack{v' \in V^R \\ v'[S]=v}} \frac{|V^R|}{|V^R|} \sum_{\substack{w \in V^R \\ w[\underline{R}]=v'}} \lambda_{v'}^{R, R} \\ &= \sum_{\substack{R \subset R, R \supset S \\ |\underline{R}|=|R|-1}} \frac{|V^S|}{|V^R|} \sum_{\substack{v' \in V^R \\ v'[S]=v}} \lambda_{v'}^{R, R} \\ &= \sum_{\substack{R \subset R, R \supset S \\ |\underline{R}|=|R|-1}} \bar{\lambda}_v^{R, R}. \end{aligned} \quad (21)$$

Adding (20) to (19) for all $R \supset S$ such that $|R| = |S| + 2$, we cancel the negative terms of (19) again, obtaining,

$$2 \left(\hat{Y}_v^S \sum_{\substack{R \supset S \\ |R| \leq |S|+2}} \frac{1}{\sigma^2(\tilde{Y}_{v(R)}^R)} - \sum_{\substack{R \supset S \\ |R| \leq |S|+2}} \frac{\tilde{Y}_{v(R)}^R}{\sigma^2(\tilde{Y}_{v(R)}^R)} \right) - \sum_{\substack{\bar{R} \supset R \supset S \\ |\bar{R}|=|R|+1=|S|+3}} \bar{\lambda}_v^{R, \bar{R}} + \sum_{T_i} \lambda_{v[T_i]}^{T_i, S} = 0. \quad (22)$$

Continuing this process, until we have addressed all $R \supset S$, $R \neq \mathcal{A}$, we obtain:

$$2 \left(\hat{Y}_v^S \sum_{\substack{R \supset S \\ R \neq \mathcal{A}}} \frac{1}{\sigma^2(\tilde{Y}_{v(R)}^R)} - \sum_{\substack{R \supset S \\ R \neq \mathcal{A}}} \frac{\tilde{Y}_{v(R)}^R}{\sigma^2(\tilde{Y}_{v(R)}^R)} \right) - \sum_{\substack{R \supset S \\ |R|=|\mathcal{A}|-1}} \bar{\lambda}_v^{R, \mathcal{A}} + \sum_{T_i} \lambda_{v[T_i]}^{T_i, S} = 0. \quad (23)$$

Finally, we sum up (17) over all $w \in V^{\mathcal{A}}$ such that $w[S] = v$ and divide by $|V^{\mathcal{A}}|/|V^S|$:

$$2(\hat{Y}_v^S - \tilde{Y}_{v(\mathcal{A})}^{\mathcal{A}})/\sigma^2(\tilde{Y}_{v(\mathcal{A})}^{\mathcal{A}}) + \sum_{\substack{R \supset S \\ |R|=|\mathcal{A}|-1}} \bar{\lambda}_v^{R, \mathcal{A}} = 0, \quad (24)$$

where the last term is obtained by simplifying an expression similar to (21).

Adding (24) to (23), we have

$$2 \left(\hat{Y}_v^S \sum_{R \supset S} \frac{1}{\sigma^2(\tilde{Y}_{v(R)}^R)} - \sum_{R \supset S} \frac{\tilde{Y}_{v(R)}^R}{\sigma^2(\tilde{Y}_{v(R)}^R)} \right) + \sum_{T_i} \lambda_{v[T_i]}^{T_i, S} = 0. \quad (25)$$

Using the \dot{Y} notation, we can write this as

$$2(\hat{Y}_v^S - \dot{Y}_v^S)/\sigma^2(\dot{Y}_v^S) + \sum_{T_i} \lambda_{v[T_i]}^{T_i, S} = 0. \quad (26)$$

Since all of the steps we took to get (26) are invertible, we see that the original solution space is equivalent to the one implied by (26), (16), (17), and self-consistency. We see that (26) over all $v \in V^S$ along with the constraint $\dot{Y}_{v(s)}^S = \hat{Y}_{v(s)}^{T_i}$ for all T_i and $v \in V^{T_i}$ are equivalent to the following optimization problem:

$$\begin{aligned} & \text{minimize} \quad \sum_{v \in V^S} (\hat{Y}_v^S - \dot{Y}_v^S)^2 / \sigma^2(\dot{Y}_v^S) \\ & \text{subject to} \quad \hat{Y}_{v(s)}^S = \hat{Y}_v^{T_i}, \quad \text{for all } T_i \text{ and } v \in V^{T_i}, \end{aligned}$$

which is an equivalent formulation of the down pass solution \hat{Y}^S . Furthermore holding \hat{Y}^S fixed, (16), (17) and the constraints $\hat{Y}_{v(R')}^{R'} = \hat{Y}_v^R$ for all $R' \supset R \supset S$ such that $|R'| = |R| + 1$ and all $v \in V^R$, are the Lagrange equations for the problem:

$$\begin{aligned} & \text{minimize} \quad \sum_{R \supseteq S} \sum_{v \in V^R} (\hat{Y}_v^R - \tilde{Y}_v^R)^2 / \sigma^2(\tilde{Y}_v^R) \\ & \text{subject to} \quad \hat{Y}_{v(\bar{R})}^{\bar{R}} = \hat{Y}_v^R, \quad \text{for all } \bar{R} \supset R \supset S \text{ such that } |\bar{R}| = |R| + 1 \text{ and all } v \in V^R, \\ & \quad \text{where } \hat{Y}^S \text{ is held fixed,} \end{aligned}$$

which we know has a solution, achieved by a projection. Thus, we have found a solution to the Lagrange equations, which must be equivalent to the corresponding projection. We conclude that under assumptions (A1)-(A3), the SEA BLUE procedure for \hat{Y}^S is the BLUE for Y^S . \blacksquare

Lemma 9. *Let $S \in \mathcal{D}$ and assume that the $\hat{Y}^{\underline{S}}$'s are self-consistent for all margins $\underline{S} \subsetneq S$. Let \hat{Z}^S be constructed as in (4) and \dot{Z}^S be from (5). Then we have that*

1. *for all $\underline{S} \subsetneq S$ and all $v \in V^{\underline{S}}$, $\hat{Z}_{v(\underline{S})}^{\underline{S}} = \hat{Y}_v^{\underline{S}}$ and $\dot{Z}_{v(\underline{S})}^{\underline{S}} = \dot{Y}_{v(\underline{S})}^{\underline{S}}$, and*
2. *both $\hat{Z}^S, \dot{Z}^S \in \mathcal{S}_S^\perp$.*

Proof We only prove the results for \hat{Z}^S as the argument for \dot{Z}^S is the same up to a change of notation. For the reader's convenience, we first recall the definition of \hat{Z}_v^S :

$$\hat{Z}_v^S = \sum_{k=1}^{|S|} \sum_{\substack{\{A_1, \dots, A_k\} \subset S \\ R=S \setminus \{A_1, \dots, A_k\}}} (-1)^{k+1} \frac{\hat{Y}_{v[R]}^R}{|A_1| \cdots |A_k|},$$

1) Choose $A^* \in S$, set $R = S \setminus \{A^*\}$, and let $v \in V^{R^*}$. Then, using the self-consistency of the prior \hat{Y}^R 's, we can write,

$$\begin{aligned} \hat{Z}_{v(S)}^S &= \left(\hat{Y}_v^R + \sum_{\substack{A \in R \\ \underline{R}=S \setminus \{A, A^*\}}} \frac{\hat{Y}_{v[\underline{R}]}^R}{|A|} \right) \\ &\quad - \left(\sum_{\substack{A \in R \\ \underline{R}=S \setminus \{A, A^*\}}} \frac{\hat{Y}_{v[\underline{R}]}^R}{|A|} + \sum_{\substack{A_1, A_2 \in R \\ \underline{R}=S \setminus \{A_1, A_2, A^*\}}} \frac{\hat{Y}_{v[\underline{R}]}^R}{|A_1||A_2|} \right) \\ &\quad + \left(\sum_{\substack{A_1, A_2 \in R \\ \underline{R}=S \setminus \{A_1, A_2, A^*\}}} \frac{\hat{Y}_{v[\underline{R}]}^R}{|A_1||A_2|} + \sum_{\substack{A_1, A_2, A_3 \in R \\ R=S \setminus \{A_1, A_2, A_3, A^*\}}} \frac{\hat{Y}_{v[\underline{R}]}^R}{|A_1||A_2||A_3|} \right) \\ &\quad + \cdots + (-1)^{|S|} \left(\sum_{\substack{A \in R \\ \underline{R}=\{A\}}} \frac{\hat{Y}_{v[\underline{R}]}^R}{\prod_{A_i \notin \{A, A^*\}} |A_i|} + \frac{\hat{Y}^\emptyset}{\prod_{A \neq A^*} |A|} \right) \\ &\quad + (-1)^{|S|+1} \frac{\hat{Y}^\emptyset}{\prod_{A \neq A^*} |A|}, \end{aligned}$$

which we see is a telescoping sum, canceling all of the terms except the first one, giving the desired result: $Z_{v(S)}^* = \hat{Y}_v^R$. In the above expression, each of the quantities in parentheses corresponds to one of the terms in the outer sum over k of (4) with the first term consisting of the margins for which $A^* \notin \underline{R}$ and the second those with $A^* \in \underline{R}$. Note that if for any particular term, $A^* \in \underline{R}$, then we marginalize out the A^* variable with no change to the coefficients. On the other hand, if $A^* \notin \underline{R}$, then we simply get $|A^*|$ copies of the original term, canceling the factor of $|A^*|$ in the denominator. See Example 8 for an illustrative example.

We have established that \hat{Z}^S is self-consistent with the \hat{Y}^R for $R \subset S$ which satisfy $|R| = |S| - 1$. It follows that \hat{Z}^S is self-consistent with the higher margins, as they are already self-consistent with each other.

2) It suffices to show that for all $Z^0 \in \mathcal{S}_S$, $\sum_{v \in V^S} Z_v^0 \hat{Z}_v^S = 0$:

$$\sum_{v \in V^S} Z_v^0 \hat{Z}_v^S = \sum_{v \in V^S} Z_v^0 \left[\sum_{k=1}^{|S|} \sum_{\substack{\{A_1, \dots, A_k\} \subset S \\ \underline{S}=S \setminus \{A_1, \dots, A_k\}}} \frac{(-1)^{k+1} \hat{Y}_{v[\underline{S}]}^S}{\prod_{i=1}^k |A_i|} \right]$$

$$\begin{aligned}
 &= \sum_{k=1}^{|S|} \sum_{\substack{\{A_1, \dots, A_k\} \subset S \\ \underline{S} = S \setminus \{A_1, \dots, A_k\}}} \sum_{v \in V^{\underline{S}}} \frac{(-1)^{k+1} Z_v^0 \hat{Y}_{v[\underline{S}]}^{\underline{S}}}{\prod_{i=1}^k |A_i|} \\
 &= \sum_{k=1}^{|S|} \sum_{\substack{\{A_1, \dots, A_k\} \subset S \\ \underline{S} = S \setminus \{A_1, \dots, A_k\}}} \sum_{w \in V^{\underline{S}}} \sum_{\substack{v \in V^S \\ w=v[\underline{S}]}} \frac{(-1)^{k+1} Z_v^0 \hat{Y}_w^{\underline{S}}}{\prod_{i=1}^k |A_i|} \\
 &= \sum_{k=1}^{|S|} \sum_{\substack{\{A_1, \dots, A_k\} \subset S \\ \underline{S} = S \setminus \{A_1, \dots, A_k\}}} \sum_{w \in V^{\underline{S}}} \frac{(-1)^{k+1} \hat{Y}_w^{\underline{S}}}{\prod_{i=1}^k |A_i|} \sum_{\substack{v \in V^S \\ w=v[\underline{S}]}} Z_v^0.
 \end{aligned}$$

Since $Z^0 \in \mathcal{S}_S$, we have that $\sum_{\substack{v \in V^S \\ w=v[\underline{S}]}} Z_v^0 = Z_{w(S)}^0 = 0$. We conclude that $\sum_{v \in V^S} Z_v^0 \hat{Z}_v^S = 0$ as desired. \blacksquare

Theorem 10. Assume (A1) holds, let \hat{Y}^S be a table achieved after the collection step, and let \hat{Z}^S be the result of (4) using $\hat{Y}^{\underline{S}}$ for $\underline{S} \subsetneq S$. Let \hat{Z}^S be defined as in (5) using the values in \hat{Y}^S . Then,

1. $\hat{Z}^S = \text{Proj}_{\mathcal{S}_S^\perp}^I \hat{Y}^S$ (equivalently, $\hat{Y}^S - \hat{Z}^S = \text{Proj}_{\mathcal{S}_S}^I \hat{Y}^S$), and
2. if (A2) and (A3) hold, then $\hat{Y}^S = \hat{Y}^S - \hat{Z}^S + \hat{Z}^S$.

Proof Call P the operator that maps \hat{Y}^S to \hat{Z}^S . To simplify notation, we write \mathcal{S} in place of \mathcal{S}_S .

1) We will show that $P = \text{Proj}_{\mathcal{S}^\perp}^I$. It suffices to show that a) the elements of \mathcal{S} are eigenvectors of P with eigenvalue 0, b) the elements of \mathcal{S}^\perp are eigenvectors of P with eigenvalue 1, and c) the operator P is self-adjoint with respect to the standard inner product. This suffices because a) and b) establish that P is idempotent, with kernel \mathcal{S} and range \mathcal{S}^\perp , and c) establishes that P is symmetric. Altogether these imply that P is the orthogonal projection onto \mathcal{S}^\perp .

a) Suppose $\hat{Y}^S \in \mathcal{S}$, meaning that all margins of \hat{Y}^S are zero. By the construction of P in (5), we have that $\hat{Z}_v^S = 0$ for all $v \in S$. Thus, elements of \mathcal{S} are eigenvectors of P with eigenvalue 0.

b) Suppose that $\hat{Y}^S \in \mathcal{S}^\perp$. We know that $\hat{Z}^S \in \mathcal{S}^\perp$ by Lemma 9. Since \mathcal{S}^\perp is closed under linear combinations, $\hat{Z}^S - \hat{Y}^S \in \mathcal{S}^\perp$. However, since \hat{Z}^S and \hat{Y}^S have the same margins, by Lemma 9, we have that $\hat{Z}^S - \hat{Y}^S \in \mathcal{S}$. Thus, $\hat{Z}^S - \hat{Y}^S \in \mathcal{S} \cap \mathcal{S}^\perp = \{0\}$, since \mathcal{S} and \mathcal{S}^\perp are orthogonal spaces. Thus, $\hat{Z}^S = \hat{Y}^S$. We see that the elements of \mathcal{S}^\perp are eigenvectors of P with eigenvalue 1.

c) Let X and Y be two tables of the same dimension. We need to show that $(PX)^\top Y = X^\top (PY)$:

$$(PX)^\top Y = (PX)^\top (Y - PY + PY)$$

$$\begin{aligned}
 &= (PX)^\top(Y - PY) + (PX)^\top(PY) \\
 &= (PX)^\top(PY),
 \end{aligned}$$

where we used the fact that $PX \in \mathcal{S}^\perp$, and $Y - PY \in \mathcal{S}$ which can be seen as follows: let $V_1, \dots, V_{\dim \mathcal{S}}$ be a basis for \mathcal{S} and let $W_1, \dots, W_{\dim \mathcal{S}^\perp}$ be a basis for \mathcal{S}^\perp , which together form a basis for all tables on the variables S . Then there exists real-valued coefficients $c_1, \dots, c_{\dim \mathcal{S}}$ and $d_1, \dots, d_{\dim \mathcal{S}^\perp}$ such that

$$Y = \sum_{i=1}^{\dim \mathcal{S}} c_i V_i + \sum_{j=1}^{\dim \mathcal{S}^\perp} d_j W_j.$$

Then,

$$PY = \sum_{j=1}^{\dim \mathcal{S}^\perp} d_j W_j,$$

and we see that $Y - PY = \sum_{i=1}^{\dim \mathcal{S}} c_i V_i \in \mathcal{S}$.

By swapping the roles of X and Y above, we can similarly show that $X^\top(PY) = (PX)^\top(PY)$, which implies that $(PX)^\top Y = X^\top(PY)$, establishing that P is self-adjoint in the standard inner product.

We have established that P is the orthogonal projection with kernel \mathcal{S} and range \mathcal{S}^\perp .

2) By Theorem 3, we know that $\hat{Y}^S = \text{Proj}_{\hat{\Sigma}^S} \dot{Y}^S + \text{Proj}_{\hat{\Sigma}^S} \hat{Z}^S$ since \hat{Z}^S is a table with the desired margins, where $\hat{\Sigma}$ is the covariance of \dot{Y}^S . Under assumptions (A1)-(A3), we have that all elements of $\hat{\Sigma}$ are independent with equal variance. It follows that the projections are orthogonal projections. So, we can simplify this as,

$$\begin{aligned}
 \hat{Y}^S &= \dot{Y}^S - \dot{Z}^S + P\hat{Z}^S \\
 &= \dot{Y}^S - \dot{Z}^S + \hat{Z}^S,
 \end{aligned}$$

since $\dot{Y}^S - \dot{Z}^S = \text{Proj}_{\hat{\Sigma}^S}^I$, $P = \text{Proj}_{\mathcal{S}^\perp}^I$, and $\hat{Z}^S \in \mathcal{S}^\perp$. ■

The lemma below quantifies the dimension of \mathcal{S}_S , which is defined in Equation (3).

Lemma 17. *Let $S \in \mathcal{D}$ consist of k variables, each with n_j levels. Let \mathcal{S}_S be defined as in (3).*

1. *The dimension of \mathcal{S}_S is*

$$\sum_{j=0}^k (-1)^{k-j} \sum_{\{m_1, \dots, m_j\} \subset \{n_1, \dots, n_k\}} m_1 \cdots m_j = \prod_{i=1}^k (n_i - 1).$$

2. *A basis for \mathcal{S}_S is the set of tables $Z^{(i_j)}_{j=1}^k$ for $i_j \in \{1, \dots, n_j - 1\}$, where for $v \in V^S$,*

$$Z_v^{(i_j)_{j=1}^k} = \left[\prod_{j=1}^k (-1)^{I(v_j = i_j + 1)} I(v_j \in \{i_j, i_j + 1\}) \right].$$

Proof Recall that $\mathcal{S}_S = \{Z^S | \text{for all } \underline{S} \subsetneq S, \text{ and all } v \in V^{\underline{S}}, \text{ we have } Z_{v(S)}^S = 0\}$. To establish the dimension of \mathcal{S}_S , we will instead calculate the dimension of its orthogonal complement \mathcal{S}_S^\perp . We know that \mathcal{S}_S^\perp is spanned by the set of self-consistency constraints, which enforce $Z_{v(S)} = 0$ for $\underline{S} \subset S$ such that $|\underline{S}| = |S| - 1$ and all $v \in V^{\underline{S}}$ (note that requiring every margin to be equal to 0 is equivalent to only requiring that the $|S| - 1$ margins are equal to 0). We observe that there are $\sum_{\{m_1, \dots, m_{k-1}\} \subset \{n_1, \dots, n_k\}} m_1 m_2 \cdots m_{k-1}$ such constraints, but they are not linearly independent. Indeed, for any such \underline{S} , the set of constraints $\{Z_{v(S)} = 0 \text{ for } v \in V^{\underline{S}}\}$ also implies that higher margins are equal to zero, giving redundant information for different \underline{S} 's. Using inclusion-exclusion, we can count the rank of these constraints as

$$\sum_{j=0}^{k-1} (-1)^{k-j-1} \sum_{\{m_1, \dots, m_j\} \subset \{n_1, \dots, n_k\}} m_1 \cdots m_j.$$

Then, the dimension of \mathcal{S}_S is

$$\begin{aligned} & (n_1 \cdots n_k) - \sum_{j=0}^{k-1} (-1)^{k-j-1} \sum_{\{m_1, \dots, m_j\} \subset \{n_1, \dots, n_k\}} m_1 \cdots m_j \\ &= \sum_{j=0}^k (-1)^{k-j} \sum_{\{m_1, \dots, m_j\} \subset \{n_1, \dots, n_k\}} m_1 \cdots m_j \\ &= \prod_{i=1}^k (n_i - 1). \end{aligned}$$

For part 2, note that each $Z_v^{(i_j)_{j=1}^k}$ is a member of \mathcal{S}_S since each table has a single “+1” and “−1” in each coordinate. Furthermore, the $Z_v^{(i_j)_{j=1}^k}$ can be seen to be linearly independent since in lexicographic order, they form a lower triangular matrix. Finally, we count that there are $\prod_{i=1}^k (n_i - 1)$ such tables, which is equal to the rank of \mathcal{S}_S , shown in part 1. We conclude that the $Z_v^{(i_j)_{j=1}^k}$ indeed form a basis for \mathcal{S}_S . \blacksquare

Lemma 12. *Given all possible tables from k variables, each with I levels,*

1. $p = \min\{I^k, (I+1)^k - I^k\}$, where p is defined as above, and
2. *matrix projection has runtime $O(I^{2k-1} + I^{2.373(k-1)})$ and memory $O(I^{2k-1})$ as $I \rightarrow \infty$, and runtime $O(I^k(I+1)^k + I^{2.373k})$ and memory $O(I^k(I+1)^k)$ as $k \rightarrow \infty$.*

Proof In Lemma 17, we established that the dimension of the null space when making a table with i variables self-consistent with the margins one level above is $(I-1)^i$. These constraints can be seen to be linearly independent, resulting in

$$\sum_{i=0}^k \binom{k}{i} (I-1)^i = I^k,$$

as the dimension of the complete null space. Thus, $p = \min\{I^k, (I+1)^k - I^k\}$. The runtime and memory results follow by using the formulas $O(np + p^{2.373})$ for runtime and $O(np)$ for memory, along with the simplification $p = O(I^{k-1})$ as $I \rightarrow \infty$ and $p = O(I^k)$ as $k \rightarrow \infty$. ■

Lemma 18. *Let S be a set of variables, with size $|S| = k$. Given self-consistent margins \hat{Y}^R for $R \subsetneq S$, the time to produce \hat{Z}^S is $O(k2^{k-1})$.*

Proof The construction in (4) requires summing over all proper subsets of S , and calculating $\prod_{i=1}^j |A_i|$ for a subset $\{A_1, \dots, A_j\} \subset S$. The product takes $O(j)$ time, so we have a runtime on the order of

$$\sum_{j=1}^k \binom{k}{j} k = k2^{k-1},$$

since there are $\binom{k}{j}$ subsets of S of size j . ■

Lemma 13. 1. *Let S be a table with k variables, each with I levels. Then it is possible to compute all of its margins in $O(I^k)$ time as $I \rightarrow \infty$ and $O((I+1)^k)$ time as $k \rightarrow \infty$; the memory requirement is of the same order.*

2. *Given all possible tables from k variables, each with I levels, it is possible to compute all of its margins in $O(I^k) = O(n)$ time as $I \rightarrow \infty$ and $O((2I+1)^k) = O(n^{\log(2I+1)/\log(I+1)})$ time as $k \rightarrow \infty$; under (A1)-(A3), the memory requirement of SEA BLUE is $O(n)$.*

Proof 1) We will start by marginalizing out one variable at a time and using previously computed margins as the starting point for subsequent margins to avoid marginalizing the same variable more than once. Suppose we consider all tables which have marginalized over i variables. There are $\binom{k}{i}$ such tables, which consist of I^{k-i} counts each; for each count in a margin, it takes I time to compute it from a previous margin. Thus, the runtime is

$$\sum_{i=1}^k \binom{k}{i} I^{k-i+1} = I(I+1)^k - I^{k+1},$$

which is $O(I^k)$ as $I \rightarrow \infty$ and $O((I+1)^k)$ as $k \rightarrow \infty$. The memory required (including the original counts in S) is

$$\sum_{i=0}^k \binom{k}{i} I^{k-i} = (I+1)^k,$$

which is $O(I^k)$ as $I \rightarrow \infty$ and $O((I+1)^k)$ as $k \rightarrow \infty$.

2) Using the above result for a single table and summing over all tables, we have a total runtime of

$$\sum_{j=0}^k \binom{k}{j} \sum_{i=1}^j \binom{j}{i} I^{k-i+1} = I(2I+1)^k - 2^k I^{k+1},$$

which is $O((I+1)^k) = O(n)$ as $I \rightarrow \infty$, and is $O((2I+1)^k) = O(n^{\log(2I+1)/\log(I+1)})$ as $k \rightarrow \infty$.

In the collection step, every table produces at most $O(n)$ estimates, which are aggregated into the running totals C_v^S and D_v^S , which are of order $O(n)$. Thus, at no point does the memory requirement exceed $O(n)$. For the down pass, when (A1)-(A3) hold, we can use the efficient formula for \hat{Y}^S in Theorem 10 which, as Remark 11 points out, has the same computational and memory requirement as the collection step. Note that Lemma 18 establishes that the runtime to assemble \hat{Z}^S and \hat{Z}^S , given the margins, is of a lower order than the time to compute the margins in the first place. \blacksquare

Suppose that (A1)-(A3) hold. Let $\Sigma_i^{1/2}$ be the i^{th} column of $\Sigma^{1/2}$. Apply SEA BLUE to $\Sigma_i^{1/2}$ (pretending the vector $\Sigma_i^{1/2}$ contains noisy counts), setting all invariants to zero, and call the result $\hat{\Sigma}_i^{1/2}$. Let $\hat{\Sigma}^{1/2}$ be the matrix whose i^{th} column is $\hat{\Sigma}_i^{1/2}$. Then, $\hat{\Sigma} = \hat{\Sigma}^{1/2}(\hat{\Sigma}^{1/2})^\top$ is the covariance matrix of $\text{vec}(\hat{Y})$. If $\Sigma = \sigma^2 I$ (a stronger version of (A3) where all counts have equal variance), then $\hat{\Sigma}$ is obtained column-wise by simply passing each column of Σ through SEA BLUE and no matrix multiplication is needed.

If \tilde{Y}_v^S are all normally distributed, then a $(1 - \alpha)$ -confidence interval for $\text{vec}(Y)_i$ is $\text{vec}(\hat{Y})_i \pm z_{1-\alpha/2} \sqrt{(\hat{\Sigma})_{i,i}}$.

Proof For simplicity, we write $P = \text{Proj}_{S_0}^{\Sigma^{-1}}$, where P is the projection matrix defined in Theorem 3 which is implemented by SEA BLUE. We can write

$$\begin{aligned} P\Sigma P^\top &= P\Sigma^{1/2}\Sigma^{1/2}P^\top \\ &= P\Sigma^{1/2}(P\Sigma^{1/2})^\top, \end{aligned}$$

where we use the fact that $\Sigma^{1/2}$ is symmetric. In the case that $\Sigma = \sigma^2 I$, we have that P is an orthogonal projection, so $P^\top = P$ and by the idempotency of projection matrices we have

$$P\sigma^2 I P^\top = \sigma^2 P P = \sigma^2 P.$$

The correctness of the confidence interval follows from Theorem 3. \blacksquare

Suppose that (A1) holds and \tilde{Y} are normally distributed. Then, it follows that \hat{Y} are also normally distributed with mean Y . Let $\text{vec}(\tilde{Y})^{(1)}, \dots, \text{vec}(\tilde{Y})^{(R)}$ be i.i.d. samples with the same distribution as $\text{vec}(N)$. Call $(\hat{\sigma}^2)_i^R = \frac{1}{m} \sum_{j=1}^R [\text{vec}(\hat{Y}^{(j)})_i]^2$, where $\hat{Y}^{(j)}$ is the result of applying SEA BLUE to $\tilde{Y}^{(j)}$. Then, $\text{vec}(\hat{Y})_i \pm t_{1-\alpha/2}(R) \sqrt{(\hat{\sigma}^2)_i^R}$ is a $(1 - \alpha)$ -confidence interval for $\text{vec}(Y)_i$.

Proof Even when (A2) and (A3) may not hold, SEA BLUE is still a linear procedure. It is well known that a linear transformation of a multivariate normal distribution is another multivariate normal. Since the mean of $\hat{Y}^{(j)}$ is known to be zero, $(\sigma^2)_i^m$ is an unbiased

estimator for the variance of $\text{vec}(\hat{Y})_i$ and we have that $\frac{R(\sigma^2)_i^R}{\text{Var}(\text{vec}(\hat{Y})_i)} \sim \chi^2(R)$. Then,

$$\frac{\text{vec}(\hat{Y})_i - \text{vec}(Y)_i}{\sqrt{(\sigma^2)_i^R}} \sim t(R),$$

and the validity of the confidence interval follows. \blacksquare

Suppose that (A1) holds. Let $\text{vec}(\tilde{Y})^{(1)}, \dots, \text{vec}(\tilde{Y})^{(R)}$ be i.i.d. samples with the same distribution as $\text{vec}(N)$. Call $\hat{Y}^{(j)}$ the result of applying SEA BLUE to $\tilde{Y}^{(j)}$. Given an index i , let $X_{(1)}^i, \dots, X_{(R)}^i$ be the order statistics of $|\text{vec}(\hat{Y}_i^{(1)})|, \dots, |\text{vec}(\hat{Y}_i^{(R)})|$. If $R \geq (1 - \alpha)/\alpha$, then $\text{vec}(\hat{Y})_i \pm X_{(\lceil (1 - \alpha)(R + 1) \rceil)}^i$ is a $(1 - \alpha)$ -confidence interval for $\text{vec}(Y)_i$, where $\lceil \cdot \rceil$ is the ceiling function.

Proof Note that $|\text{vec}(\hat{Y})_i - \text{vec}(Y)_i|, |\text{vec}(\hat{Y}_i^{(1)})|, \dots, |\text{vec}(\hat{Y}_i^{(m)})|$ are all i.i.d., implying that any ordering is equally likely. So,

$$P\left(|\text{vec}(\hat{Y})_i - \text{vec}(Y)_i| \leq X_{(\lceil (1 - \alpha)(R + 1) \rceil)}^i\right) \geq \frac{\lceil (1 - \alpha)(R + 1) \rceil}{R + 1} \geq \frac{(1 - \alpha)(R + 1)}{R + 1} = 1 - \alpha,$$

where we used the fact that SEA BLUE is a linear and unbiased procedure to establish that $|\text{vec}(\hat{Y}_i) - \text{vec}(Y)_i|$ is equal in distribution to the $|\text{vec}(\hat{Y}_i^{(j)})|$'s, which is used for the first inequality. Note that the first inequality is actually an equality if the variables are continuous, and the inequality may be strict when ties are possible. Thus, the interval $\text{vec}(\hat{Y})_i \pm X_{(\lceil (1 - \alpha)(R + 1) \rceil)}^i$ is a $(1 - \alpha)$ -confidence interval for $\text{vec}(Y)_i$. \blacksquare

References

- John Abowd, Robert Ashmead, Garfinkel Simson, Daniel Kifer, Philip Leclerc, Ashwin Machanavajjhala, and William Sexton. Census TopDown: Differentially private data, incremental schemas, and consistency with public knowledge. Technical report, US Census Bureau, 2019.
- John M. Abowd, Robert Ashmead, Ryan Cumings-Menon, Simson Garfinkel, Micah Heineck, Christine Heiss, Robert Johns, Daniel Kifer, Philip Leclerc, Ashwin Machanavajjhala, Brett Moran, William Sexton, Matthew Spence, and Pavel Zhuravlev. The 2020 Census Disclosure Avoidance System TopDown Algorithm. *Harvard Data Science Review*, (Special Issue 2), jun 24 2022.
- Josh Alman and Virginia Vassilevska Williams. A refined laser method and faster matrix multiplication. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 522–539. SIAM, 2021.
- Michael L Anderson. Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American statistical Association*, 103(484):1481–1495, 2008.

- Jordan Awan and Zhanyu Wang. Simulation-based, finite-sample inference for privatized data. *Journal of the American Statistical Association*, pages 1–14, 2024.
- Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer, 2016.
- Clément L Canonne, Gautam Kamath, and Thomas Steinke. The discrete Gaussian for differential privacy. *Advances in Neural Information Processing Systems*, 33:15676–15688, 2020.
- Young Hyun Cho and Jordan Awan. Formal privacy guarantees with invariant statistics. *arXiv preprint arXiv:2410.17468*, 2024.
- Ryan Cumings-Menon. Full-information estimation for hierarchical data. *arXiv preprint arXiv:2404.13164*, 2024.
- Jinshuo Dong, Aaron Roth, and Weijie J Su. Gaussian differential privacy. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 84(1):3–37, 2022.
- Jörg Drechsler, Ira Globus-Harris, Audra Mcmillan, Jayshree Sarathy, and Adam Smith. Nonparametric differentially private confidence intervals for the median. *Journal of Survey Statistics and Methodology*, 10(3):804–829, 2022.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pages 265–284. Springer, 2006.
- Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- Jie Gao, Ruobin Gong, and Fang-Yi Yu. Subspace differential privacy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 3986–3995, 2022.
- Michael Hay, Vibhor Rastogi, Gerome Miklau, and Dan Suciu. Boosting the accuracy of differentially private histograms through consistency. *Proceedings of the VLDB Endowment*, 3(1-2):1021–1032, 2010.
- James Honaker. Efficient use of differentially private binary trees. *Theory and Practice of Differential Privacy (TPDP 2015)*, London, UK, 2:26–27, 2015.
- Christopher T Kenny, Shiro Kuriwaki, Cory McCartan, Evan TR Rosenman, Tyler Simko, and Kosuke Imai. The use of differential privacy for census data and its impact on redistricting: The case of the 2020 US Census. *Science Advances*, 7(41):eabk3283, 2021.
- Christopher T Kenny, Cory McCartan, Shiro Kuriwaki, Tyler Simko, and Kosuke Imai. Evaluating bias and noise induced by the US Census Bureau’s privacy protection methods. *Science Advances*, 10(18):eadl2524, 2024.
- Daniel Kifer. Consistency of hierarchical histograms, 2021. private correspondence.

Cory McCartan, Tyler Simko, and Kosuke Imai. Making differential privacy work for census data users. *Harvard Data Science Review*, 5(4), 2023.

Alexis R Santos-Lozada, Jeffrey T Howard, and Ashton M Verdery. How differential privacy will affect our understanding of health disparities in the United States. *Proceedings of the National Academy of Sciences*, 117(24):13405–13412, 2020.

Richelle L Winkler, Jaclyn L Butler, Katherine J Curtis, and David Egan-Robertson. Differential privacy and the accuracy of county-level net migration estimates. *Population Research and Policy Review*, pages 1–19, 2021.

Min-ge Xie and Peng Wang. Repro samples method for finite-and large-sample inferences. *arXiv preprint arXiv:2206.06421*, 2022.

Arthur B Yeh and Kesar Singh. Balanced confidence regions based on Tukey’s depth and the bootstrap. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 59(3):639–652, 1997.