

FULL-INFORMATION ESTIMATION FOR HIERARCHICAL DATA

RYAN CUMINGS-MENON

ABSTRACT. The U.S. Census Bureau’s 2020 Disclosure Avoidance System (DAS) bases its output on noisy measurements, which are population tabulations added to realizations of mean-zero random variables. These noisy measurements are observed in a set of hierarchical geographic units, *e.g.*, the U.S. as a whole, states, counties, census tracts, and census blocks. The noisy measurements from the 2020 Redistricting Data File and Demographic and Housing Characteristics File statistical data products are now public. The purpose of this paper is to describe a method to leverage the hierarchical structure within these noisy measurements to compute confidence intervals for arbitrary tabulations and in arbitrary geographic entities composed of census blocks. This method is based on computing a weighted least squares estimator (WLS) and its variance matrix. Due to the high dimension of this estimator, this operation is not feasible using the standard approach, since this would require evaluating products with the inverse of a dense matrix with several billion (or even several trillion) rows and columns. In contrast, the approach we describe in this paper computes the required estimate and its variance with a time complexity and memory requirement that scales linearly in the number of census blocks.

1. INTRODUCTION

As described by Abowd et al. (2022); Cumings-Menon et al. (2023) in more detail, the U.S. Census Bureau’s 2020 Disclosure Avoidance System (DAS) uses formally private methods (Bun and Steinke, 2016; Dwork et al., 2006) to protect the confidentiality of 2020 Census respondents. One important advantage of formally private methods is that they allow for a much greater level of transparency than the methods the Census Bureau used for disclosure limitation in previous decennial censuses. For example, the Census Bureau has released noisy measurements that were used to create the tabulations found in the published 2020 Census statistical data products (U.S. Census Bureau, 2023b), which are called the noisy measurement files (NMFs). These scalar noisy measurements are defined as population tabulations for specific geographic units (*e.g.*, the count of respondents that identify as Black and Asian in Utah, the count of respondents that identify as Hispanic or Latino in the U.S. as a whole, *etc.*) added to realizations of mean-zero random variables. The geographic units used to define these noisy measurements are also defined hierarchically in a rooted tree, *e.g.*, the U.S. as a whole, states, counties, census tracts, and census blocks. The purpose of this paper is to describe a way to use these hierarchical NMFs to compute the best linear unbiased estimator (BLUE) and confidence intervals (CIs) for arbitrary cross tabulations and for arbitrary geographic regions.

After introducing notation in the next subsection, the remainder of this section describes our primary motivating use case, *i.e.*, computing point estimates and CIs using the publicly available NMFs, and the previous literature on methods related to the ones we describe here. To make this paper accessible for other use cases, when possible, we avoid terminology that is specific to the NMFs; one of the main requirements for the application of these methods to other use cases is that the input data is composed of observations for vertices in a rooted tree, as explained in more detail in the next section. Section 2 describes our proposed approach for estimating the full-information weighted least squares (WLS) estimator, which we call the Two-Pass Algorithm, since this algorithm performs operations recursively from the bottom of the tree upward and then, in a second stage, from the root of the tree downward. Section 3 describes a computationally efficient way to compute the covariance between the WLS estimates of arbitrary pairs of vertices of the tree. Section 4 describes how the proposed algorithms can be used to compute confidence intervals for user-defined queries in arbitrary geographic regions. Section 5 provides a numerical experiment of our proposed approach, to explore the accuracy of the resulting confidence interval estimates and the computational requirements of

Date: April, 2024.

The views expressed in this technical paper are those of the author and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data (Project No. 7502798, Disclosure Review Board (DRB) approval number: CBDRB-FY24-CED005-0002).

the proposed algorithms. Our proposed two-pass estimation approach is related to the approach described by Hay et al. (2010), which is described in Section 6 in more detail, and Section 7 concludes.

1.1. Notation. Throughout the paper we denote matrices using uppercase, and vectors using bold lowercase. Given the matrices $A, B \in \mathbb{R}^{N \times N}$ we use $A \leq B$ to denote the property that all eigenvalues of $B - A$ are real and the minimum eigenvalue of $B - A$ is non-negative. Given the random vectors $\mathbf{a} \in \mathbb{R}^M$ and $\mathbf{b} \in \mathbb{R}^N$, each with elements with finite variance, we use $\text{Var}(\mathbf{a}) \in \mathbb{R}^{M \times M}$ to denote the variance matrix of \mathbf{a} and $\text{Cov}(\mathbf{a}, \mathbf{b}) \in \mathbb{R}^{M \times N}$ to denote the covariance matrix between \mathbf{a} and \mathbf{b} . We denote the Kronecker product of real matrices A, B by $A \otimes B$. The length N column vector with each element equal to $c \in \mathbb{R}$ is denoted by \mathbf{c}_N and an $M \times N$ matrix with each element equal to $C \in \mathbb{R}$ by $C_{M \times N}$. Also, when there is little risk of confusion, we omit the subscript and write \mathbf{c} (respectively, C) instead. We denote the $n \times n$ identity matrix by I_n , the i^{th} row of $A \in \mathbb{R}^{m \times n}$ by $A[i, \cdot]$, and the i^{th} column of A by $A[\cdot, i]$. We also denote the set of nonnegative real numbers by \mathbb{R}_+ , the positive reals by \mathbb{R}_{++} , and the nonnegative integers by \mathbb{Z}_+ . We also denote the cardinality of the finite set S by $\text{Card}(S)$.

1.2. Setting. We make use of some terms used by the Census Bureau to describe geographic entities, most of which can alternatively be translated into concepts from graph theory. For example, the Census Bureau’s concept of a *geographic spine*, or *spine*, can alternatively be viewed as a rooted tree. Vertices in this rooted tree are called *geographic units*, or *geounits*, and each geounit is associated with a geographic region. For example, in the tabulation US spine, as defined for each decennial census by the Census Bureau’s Geography Division, the root vertex is the US geounit, which has a geographic extent corresponding to all 50 states and Washington DC. Also, in the tabulation Puerto Rico (PR) spine, the root geounit is the PR geounit, corresponding to Puerto Rico as a whole. The geographic extent of each geounit on the spine corresponds to an element of a partition of the geographic extent of the root geounit. The collection of geounits corresponding to elements within the same partition is called a *geographic level*, or a *geolevel*. For example, in the tabulation US spine, the US geolevel contains only the US geounit, and the state geolevel consists of the children of the US geounit, *i.e.*, 51 geounits, each corresponding to either one of the 50 states or Washington DC. The other geolevels on the tabulation US spine are the county, census tract, block group, and block geolevels. Figure 1 provides a graphical representation of the tabulation US spine, including an example of a path in the 2020 vintage of this spine from a block to the US geounit.

While we use the terminology of the Census Bureau when describing the primary use case we have in mind, such as in the numerical experiment we present in Section 5, most often we use terminology from graph theory instead. Specifically, let the rooted tree(/spine) be denoted by \mathcal{G} . We use $g \in \mathcal{G}$ to denote vertex g of \mathcal{G} . To denote the subtree of \mathcal{G} rooted at node $g \in \mathcal{G}$, *i.e.*, the subset of \mathcal{G} that includes $g \in \mathcal{G}$ and all of its descendants, we use \mathcal{G}_g . We assume that \mathcal{G} is defined so that the shortest path from any given leaf vertex to the root vertex all have a length equal to L , and for any level index $l \in \{0, \dots, L\}$, we use $\text{Level}(\mathcal{G}, l)$ to denote the set of vertices in level l . We also use $\text{Children}(g)$ to denote the set of child vertices of $g \in \mathcal{G}$.

For each vertex $g \in \mathcal{G}$, we will denote the vector of (unknown) independent variables associated with the vertex by $\beta(g) \in \mathbb{R}^n$. Throughout the paper, we assume that parent-child consistency holds; in other words, for any vertex $g \in \mathcal{G}$ that is not a leaf, we assume

$$(1) \quad \sum_{c \in \text{Children}(g)} \beta(c) = \beta(g).$$

We also will associate with each vertex $g \in \mathcal{G}$ a coefficient matrix $S(g) \in \mathbb{R}^{m \times n}$, which we assume has full column rank. Note that we do not assume that the coefficient matrix of different vertices are identical.¹ For each vertex g , the vector of observed noisy measurements of g will be denoted by $\mathbf{y}(g) = S(g)\beta(g) + \mathbf{u}(g)$, where each element of $\mathbf{u}(g) \in \mathbb{R}^m$ is an independently distributed random variable with zero mean and finite variance.

We perform various stacking operations on the attributes of the vertices below, *e.g.*, $\text{stack}(\{S(g)\beta(g)\}_{g \in \mathcal{G}})$, and these operations require that the vertices are ordered. Specifically, we assume this ordering ensures that,

¹One implication of our use of m and n to denote the number of rows and columns of $S(g)$ throughout the paper is that this implicitly requires that these dimensions are the same for each vertex $g \in \mathcal{G}$; however, we actually only require the number of columns of $S(g)$, *i.e.*, $n \in \mathbb{N}$, to be the same for each vertex.

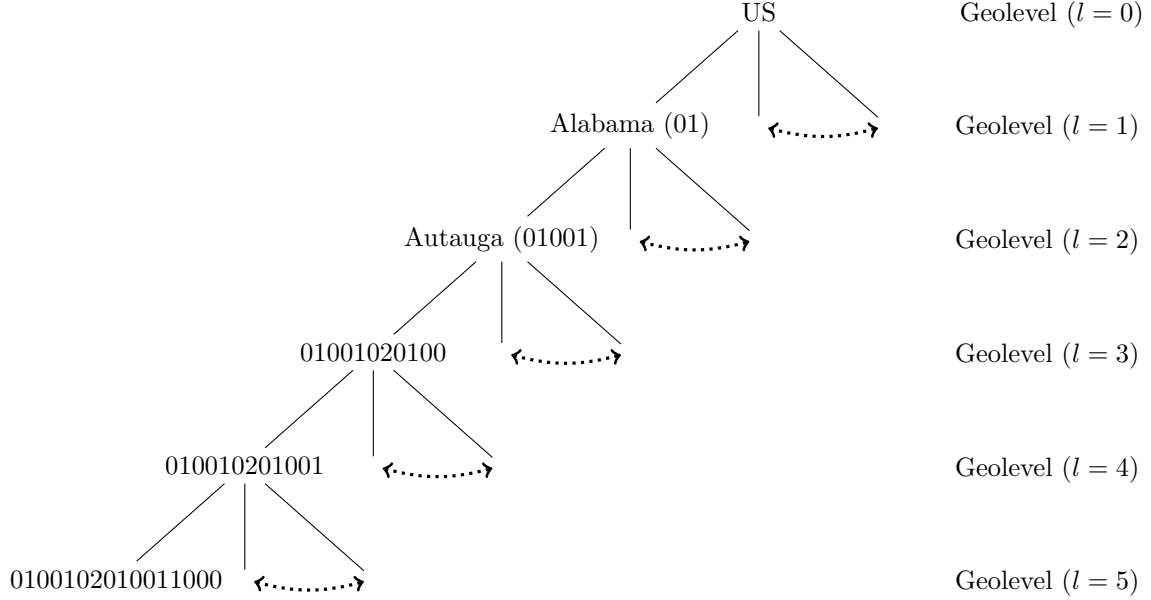


FIGURE 1. This is a graphical depiction of the 2020 tabulation US spine, as defined by the Geography Division of the Census Bureau. The right column provides the geolevel names and indices. The left side of the figure provides an example of a path from the US geounit to a block geounit. Specifically, in our ordering of geounits within each geolevel, this particular path passes through the first geounit in each geolevel.

if $0 \leq l < l' \leq L$, then each vertex in level l precedes vertices in level l' and also that the children of the u^{th} vertex of level l preceded the children of the $(u + 1)^{\text{st}}$ vertex of level l . For example, the path provided in Figure 1 provides the first geounit in this ordering in each geolevel, along with the numerical geounit identifier codes, or *geocodes*, for these geounits in the 2020 tabulation US spine. These geocodes are formatted so that the first part of each geocode is the same as the geocode of the parent. Defining geocodes in this way is a straightforward way to define an ordering with both of our required properties, since this geounit ordering can be defined so that it corresponds to the ordering of the lexicographically-sorted list of these geocodes.²

It may also be helpful to see how these constructs fit into the use case that is the main motivation for this paper, *i.e.*, estimating arbitrary population cross tabulations for arbitrary geographic entities using the public census NMFs. In this setting, for each geounit $g \in \mathcal{G}$, $\beta(g) \in \mathbb{Z}_+^n$ is defined as a vector of histogram counts, *i.e.*, a flattened fully saturated contingency table, $S(g) \in \{0, 1\}^{m \times n}$ is called a strategy matrix, and $\mathbf{y}(g) := S(g)\beta(g) + \mathbf{u}(g)$ is called the noisy measurements of geounit g . Also, in this use case, for each geounit $g \in \mathcal{G}$, $\mathbf{u}(g) \in \mathbb{Z}^m$ is a draw from the mean-zero discrete Gaussian distribution, as described by Canonne et al. (2020). One particular example we use in our numerical experiment in Section 5 is the case of the NMFs used for the persons universe of the Redistricting Data File. For each geounit, these tabulations are composed of cross products of the attributes of the schema of the histogram, which, as described by Abowd et al. (2022) in more detail, consist of the following attributes.

- Household or Group Quarters Type (HHGQ), 8 levels: Provides counts of those living in households, correctional facilities for adults, juvenile facilities, nursing facilities/skilled-nursing facilities, other institutional facilities, college/university student housing, military quarters, and those living in other noninstitutional facilities
- Voting Age (VOTING_AGE), 2 levels: Provides counts of individuals that are age 17 or younger and those that are age 18 or older

²More specifically, this ordering is defined by a sort operation that results in geocodes with a lower width preceding longer-width geocodes and geocodes with the same width, *i.e.*, geocodes of geounits within the same geolevel, being sorted lexicographically.

- Hispanic or Latino Origin (HISPANIC), 2 levels: Provides counts of individuals that are Hispanic/Latino and those that are not Hispanic or Latino
- Census Race (CENRACE), 63 levels: Provides counts of individuals that identify as each combination of Black/African American, American Indian/Native Alaskan, Asian, Native Hawaiian/Pacific Islander, White, and some other race, except “none of the above.”

Note that U.S. Census Bureau (2023b) provides instructions for downloading the Redistricting Data File NMFs, both with unobserved histogram counts defined using 2010 and using 2020 data, as well as the NMFs for Demographic and Housing Characteristics (DHC) files DAS executions. The DHC NMFs use larger schemas than the Redistricting Data File schemas in order to support the estimation of more granular queries. For example, while the persons universe Redistricting Data File schema described above consists of $n = 2,016 = 8 \times 2 \times 2 \times 63$ detailed histogram cells for each geounit, the persons universe DHC schema consists of 1,227,744 detailed cells.

1.3. Motivation: The Weighted Least Squares (WLS) Estimator. In this section we will define the WLS estimator for $\{\beta(g)\}_{g \in \mathcal{G}}$. To do so, note that our parent-child consistency assumption implies that, for any $g \in \mathcal{G}$, $S(g)\beta(g)$ can be expressed as a linear function of $\{\beta(g)\}_{g \in \text{Level}(\mathcal{G}, L)}$. In other words, if we let $\beta \in \mathbb{R}^N$ denote $\text{stack}(\{\beta(g)\}_{g \in \text{Level}(\mathcal{G}, L)})$, then there exists $S \in \mathbb{R}^{M \times N}$ such that the vector $\text{stack}(\{S(g)\beta(g)\}_{g \in \mathcal{G}})$ is given by $S\beta$. Similarly, let $\mathbf{u} \in \mathbb{R}^M$ denote the stacked error terms, *i.e.*, $\mathbf{u} := \text{stack}(\{\mathbf{u}(g)\}_{g \in \mathcal{G}})$, and let \mathbf{y} denote the stacked vector of observed noisy measurements from all vertices, *i.e.*, $\mathbf{y} := \text{stack}(\{\mathbf{y}(g)\}_{g \in \mathcal{G}})$. Using this notation, we have $\mathbf{y} = S\beta + \mathbf{u}$. Note that our assumptions in the previous section imply that, for every $g \in \mathcal{G}$, $S(g)$ has full column rank and that $\mathbf{u}(g)$ is a random variable such that $\text{Var}(\mathbf{u}(g))$ is a diagonal matrix with finite elements. The next example describes $\mathbf{u}, \mathbf{y}, \beta$, and S for one case in which the tree \mathcal{G} is particularly simple.

Example 1. Suppose the tree \mathcal{G} is composed of one root vertex, say g_0 , and this vertex has two child vertices, say g_1 and g_2 , respectively. In this case, we would have $\mathbf{u} = \text{stack}(\{\mathbf{u}(g_0), \mathbf{u}(g_1), \mathbf{u}(g_2)\})$, $\mathbf{y} = \text{stack}(\{\mathbf{y}(g_0), \mathbf{y}(g_1), \mathbf{y}(g_2)\})$, and $\beta := \text{stack}(\{\beta(g_1), \beta(g_2)\})$. Also, S would be defined by,

$$S = \begin{bmatrix} S(g_0) & S(g_0) \\ S(g_1) & 0 \\ 0 & S(g_2) \end{bmatrix}.$$

We can use this notation to define the full-information WLS estimator and its variance matrix as

$$(2) \quad \tilde{\beta} := \arg \min_{\beta} (S\beta - \mathbf{y})^\top \text{Var}(\mathbf{u})^{-1} (S\beta - \mathbf{y}) = (S^\top \text{Var}(\mathbf{u})^{-1} S)^{-1} S^\top \text{Var}(\mathbf{u})^{-1} \mathbf{y}$$

$$(3) \quad \text{Var}(\tilde{\beta}) := (S^\top \text{Var}(\mathbf{u})^{-1} S)^{-1}.$$

The motivation for this paper is to describe a computationally efficient method that takes as input the observed noisy measurements $\mathbf{y} \in \mathbb{R}^M$, an arbitrary vector $\mathbf{q} \in \mathbb{R}^n$, and an arbitrary subset of leaf vertices $H \subset \text{Level}(\mathcal{G}, L)$, and then returns the best linear unbiased estimate of

$$(4) \quad \beta_{H, \mathbf{q}} := \sum_{g \in H} \mathbf{q}^\top \beta(g),$$

along with the variance of this estimate and its CI. Note that for graphs with a particularly low number of leaf vertices, so that $\text{Var}(\tilde{\beta})$ is not a high dimensional matrix, the standard approach to WLS estimation could be carried out in a straightforward manner. Specifically, if we define $\mathbf{h} \in \{0, 1\}^{\text{Card}(\text{Level}(\mathcal{G}, L))}$ so that $\mathbf{h}[i] = 1$ if and only if the i^{th} leaf vertex is in H , then this estimator and its variance would be given by

$$(5) \quad \tilde{\beta}_{H, \mathbf{q}} = (\mathbf{h} \otimes \mathbf{q})^\top \tilde{\beta}$$

$$(6) \quad \text{Var}(\tilde{\beta}_{H, \mathbf{q}}) = (\mathbf{h} \otimes \mathbf{q})^\top \text{Var}(\tilde{\beta}) (\mathbf{h} \otimes \mathbf{q}),$$

These two values could in turn be used to estimate confidence intervals using a Gaussian approximation for $\tilde{\beta}$.

In this paper, we describe how to compute $\tilde{\beta}_{H,q}$, $\text{Var}(\tilde{\beta}_{H,q})$, and the corresponding CI, without explicitly computing and storing the matrix $\text{Var}(\tilde{\beta})$. The use cases we have in mind are ones in which $\text{Var}(\tilde{\beta})$ is too high dimensional to compute or store explicitly, primarily because \mathcal{G} has a high number of leaves. For example, the Redistricting Data File NMFs contain noisy measurements for approximately six million blocks, so the outer inverse in (3) required to define $\text{Var}(\tilde{\beta})$ directly would not be possible to evaluate in practice, even when n , *i.e.*, the number of columns of each $\{S(g)\}_{g \in \mathcal{G}}$, is low. In contrast, our proposed approach does require that n is small enough to allow for the efficient evaluation of inverses of $n \times n$ matrices, since our proposed algorithm does evaluate inverses of matrices of this size explicitly. When using standard methods to evaluate matrix inverses that are used in practice, our proposed algorithm has a time complexity and a memory requirement that is $O(n^3 \sum_l \text{Card}(\text{Level}(\mathcal{G}, l)))$.³ Our proposed approach works by precomputing a set of state variables for each vertex of \mathcal{G} using the two-pass algorithm that we describe in the next section. We also describe algorithms that take these returned state variables as input to compute $\tilde{\beta}_{H,q}$, $\text{Var}(\tilde{\beta}_{H,q})$, and the corresponding WLS-based $1 - \alpha$ confidence level CI.

1.4. Previous Literature. Our proposed estimation algorithm is related to two existing algorithms. First, the approach described here is a generalization of the approach described by Hay et al. (2010), in the sense that these two approaches provide identical estimates, *i.e.*, $\tilde{\beta}_{H,q}$, when $m = n = 1$, $\text{Var}(\mathbf{u}(g))$ is the same for each vertex $g \in \mathcal{G}$, and the number of children of each parent vertex is constant. More detail on this connection is provided in an example in Section 6. Note that, in addition to providing an algorithm for the WLS estimate itself, we also provide a method that computes the exact variance of the estimate, rather than a bound on this variance.

Second, there are also similarities between the approach described here and *forward-backward* algorithms described in the literature on hidden Markov models on trees; for example, Willsky (2002) describes several related approaches for WLS estimation using hierarchical data. At first, these approaches may seem only tangentially related because the most standard model in this literature does not satisfy our parent-child consistency assumption in equation (1), and instead assumes that, for each child $c \in \mathcal{G}$ with parent $g \in \mathcal{G}$, $\beta(c)$ is given by $\beta(c) = A(c)\beta(g) + w(c)$, where $A(c) \in \mathbb{R}^{n \times n}$ and $w(c) \in \mathbb{R}^n$ is a random variable, which is known as the Markov property. In other words, while we define $\{\beta(g)\}_{g \in \mathcal{G}}$ in a non-random and bottom-up manner, the hidden Markov model literature normally defines these vectors using a random process that proceeds in a top-down manner. However, Willsky (2002) describes Markov models on trees that are said to be *degenerate*, *i.e.*, models where the Markov property is not required to hold, and other models that are said to be *internal*, *i.e.*, models in which the histogram cell counts of each parent vertex g depends on $\{\beta(c)\}_{c \in \text{Children}(g)}$. Thus, the algorithm described in this paper can be viewed as an example of a forward-backward estimation approach from the hidden Markov model literature to estimate a model that is both degenerate and internal.

1.5. A Preliminary Result. We make use of the inverse-variance weighted mean for a vector at multiple points in this paper. The next Lemma provides the formula for the resulting vector and its variance estimate.

Lemma 1. *Suppose $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ are realizations of random variables, each with finite variances and with a mean equal to $\psi \in \mathbb{R}^n$. Then, the minimum variance unbiased estimate of ψ that is linear in \mathbf{a}, \mathbf{b} is given by*

$$\hat{\psi} = \text{Var}(\hat{\psi}) \left(\text{Var}(\mathbf{a})^{-1} \mathbf{a} + \text{Var}(\mathbf{b})^{-1} \mathbf{b} \right)$$

where

$$\text{Var}(\hat{\psi}) = \left(\text{Var}(\mathbf{a})^{-1} + \text{Var}(\mathbf{b})^{-1} \right)^{-1}$$

Proof. Let $\mathbf{c} := (\mathbf{a}^\top, \mathbf{b}^\top)^\top$, $X := \text{stack}(\{I_n, I_n\})$, and let the block matrix Ω be defined as

$$\Omega := \begin{bmatrix} \text{Var}(\mathbf{a}) & 0 \\ 0 & \text{Var}(\mathbf{b}) \end{bmatrix}.$$

³Note that we describe a way that our proposed approaches can be used to estimate unbiased point estimates, variances, and confidence intervals of estimates for queries in the much larger DHC schemas, despite the fact that n is prohibitively large in these cases for our proposed methods to be applied directly, in Section 4.1.

Now consider applying a WLS estimator to a dataset with the i^{th} independent and dependent variables given by Y_i and $X_{i,\cdot}^\top$, respectively. The resulting estimator is equal to

$$\check{\psi} = (X^\top \Omega^{-1} X)^{-1} X^\top \Omega^{-1} \mathbf{c},$$

which has variance given by

$$\text{Var}(\check{\psi}) = (X^\top \Omega^{-1} X)^{-1}.$$

Also, the Gauss-Markov theorem implies this estimator is the best linear unbiased estimator for ψ ; see for example, (Davidson et al., 2004). The final result follows from

$$\text{Var}(\check{\psi}) = (\text{Var}(\mathbf{a})^{-1} + \text{Var}(\mathbf{b})^{-1})^{-1} = \text{Var}(\hat{\psi}),$$

and

$$\check{\psi} = (X^\top \Omega^{-1} X)^{-1} X^\top \Omega^{-1} Y = \text{Var}(\hat{\psi}) (\text{Var}(\mathbf{a})^{-1} \mathbf{a} + \text{Var}(\mathbf{b})^{-1} \mathbf{b}) = \hat{\psi}.$$

□

2. THE TWO-PASS ALGORITHM

Before describing the main algorithm in this section, it will be helpful to describe the primary data structures that we use. As described in the following subsections, our main proposed algorithm performs a series of updates on state variables within each vertex $g \in \mathcal{G}$, which includes an estimate of $\beta(g)$ and its variance matrix. To reduce space requirements while running this algorithm, note that it is often possible to delete prior state variables after performing some of the updates described below. Table 1 summarizes some notation we use in this section, some of which is introduced in the following subsections.

Similar to the forward-backward algorithms described by Willsky (2002), the result of the two-pass algorithm described below is the BLUE for $\beta(g)$ for each vertex $g \in \mathcal{G}$. This is because, for each vertex $g \in \mathcal{G}$, each of the three updates to the state variable estimate of $\beta(g)$ is given by the BLUE for a progressively larger information set, *i.e.*, a subset of the noisy measurements. Specifically, the first step of the algorithm is to define this state variable estimate of each vertex g as the WLS estimator based only on the noisy measurements of vertex g . Given our assumptions that $S(g)$ has full column rank and $\mathbf{u}(g)$ has a mean of zero and finite variance, the Gauss-Markov theorem implies that this initial WLS estimate is the BLUE for this particular information set; see for example, (Davidson et al., 2004). In the bottom-up pass, this state variable is updated as the inverse-variance-weighted mean of this initial estimate and an estimate defined as the sum over the bottom-up pass estimates of the children of vertex g . As described in Lemma 1, using the inverse-variance-weighted mean results in the linear unbiased estimator with the lowest possible variance, so the bottom-up pass update defines the estimate of $\beta(g)$ as the BLUE for the information set given by the noisy measurements of the graph \mathcal{G}_g . Similarly, in the top-down pass, this information set is expanded further to also include the noisy measurements of higher levels of \mathcal{G} by projecting the estimate from the bottom-up pass onto the column space of S , using a projection defined as the minimizer of the inverse-variance-weighted least squares objective function. This step also ensures the final estimate satisfies the same parent-child consistency property as the (unobserved) vector of histogram counts, *i.e.*, equation (1). As described in more detail in Theorem 1 below, the variance-minimizing projection is also used to define this final state variable update.

2.1. Bottom-Up Pass: Initialization Step. The first step of the algorithm is to initialize state variables of each vertex $g \in \mathcal{G}$ as

$$(7) \quad \hat{\beta}(g|g) := (S(g)^\top \text{Var}(\mathbf{u}(g))^{-1} S(g))^{-1} S(g)^\top \text{Var}(\mathbf{u}(g))^{-1} \mathbf{y}(g),$$

$$(8) \quad \text{Var}(\hat{\beta}(g|g)) := (S(g)^\top \text{Var}(\mathbf{u}(g))^{-1} S(g))^{-1}.$$

For each of the leaf of the tree, *i.e.*, each vertex in $\text{Level}(\mathcal{G}, L)$, also let

$$(9) \quad \hat{\beta}(g|g-) := \hat{\beta}(g|g)$$

$$(10) \quad \text{Var}(\hat{\beta}(g|g-)) := \text{Var}(\hat{\beta}(g|g))$$

\mathcal{G}	The rooted tree, or, in our main motivating use case, the geographic spine
$\text{Children}(g)$	The set of vertices in \mathcal{G} that are children of g
\mathcal{G}_g	The subtree of \mathcal{G} rooted at vertex g
$S(g) \in \mathbb{R}^{m \times n}$	The coefficient matrix for vertex g
$\mathbf{y}(g) \in \mathbb{R}^m$	The observed noisy measurements of vertex g
$\mathbf{u}(g) \in \mathbb{R}^m$	The error component of the noisy measurements of vertex g
$\text{Var}(\mathbf{z})$	The variance matrix of the random variable \mathbf{z}
$\boldsymbol{\beta}(g) \in \mathbb{R}^n$	The (unobserved) vector of independent variables for vertex g
$\hat{\boldsymbol{\beta}}(g g) \in \mathbb{R}^n$	Estimate for vertex g based on $\mathbf{y}(g)$
$\hat{\boldsymbol{\beta}}(g g-) \in \mathbb{R}^n$	Estimate for vertex g based on noisy measurements in \mathcal{G}_g
$\tilde{\boldsymbol{\beta}}(g) \in \mathbb{R}^n$	Estimate for vertex g based on the observed random variables of all vertices in \mathcal{G}

TABLE 1. Some of the common notational conventions we use throughout the paper

2.2. Bottom-Up Pass: Recursion Step. The updates described in this section are performed first for vertices in $\text{Level}(\mathcal{G}, L-1)$, and afterward in each level moving up the tree. Specifically, given $\{\hat{\boldsymbol{\beta}}(c|c-)\}_{c \in \text{Children}(g)}$ and $\{\text{Var}(\hat{\boldsymbol{\beta}}(c|c-))\}_{c \in \text{Children}(g)}$, let

$$(11) \quad \hat{\boldsymbol{\beta}}(g|\text{Children}(g)-) := \sum_{c \in \text{Children}(g)} \hat{\boldsymbol{\beta}}(c|c-),$$

$$(12) \quad \text{Var}(\hat{\boldsymbol{\beta}}(g|\text{Children}(g)-)) := \sum_{c \in \text{Children}(g)} \text{Var}(\hat{\boldsymbol{\beta}}(c|c-)).$$

Now, since we have two independent estimates of $\boldsymbol{\beta}(g)$, *i.e.*, $\hat{\boldsymbol{\beta}}(g|g)$ and $\hat{\boldsymbol{\beta}}(g|\text{Children}(g)-)$, for each vertex g , we can define a new estimate as the linear combination of these estimates that has the lowest possible variance, which is the invariance-variance weighted mean of these two vector estimates, as described in Lemma 1. In other words, in this step, we update the state variables using

$$(13) \quad \text{Var}(\hat{\boldsymbol{\beta}}(g|g-)) := \left(\text{Var}(\hat{\boldsymbol{\beta}}(g|g))^{-1} + \text{Var}(\hat{\boldsymbol{\beta}}(g|\text{Children}(g)-))^{-1} \right)^{-1}$$

$$(14) \quad \hat{\boldsymbol{\beta}}(g|g-) := \text{Var}(\hat{\boldsymbol{\beta}}(g|g-)) \left(\text{Var}(\hat{\boldsymbol{\beta}}(g|g))^{-1} \hat{\boldsymbol{\beta}}(g|g) + \text{Var}(\hat{\boldsymbol{\beta}}(g|\text{Children}(g)-))^{-1} \hat{\boldsymbol{\beta}}(g|\text{Children}(g)-) \right).$$

2.3. Top-Down Pass: Initialization Step. Prior to starting the top-down pass, we initialize the state variables of the root vertex g as

$$(15) \quad \tilde{\boldsymbol{\beta}}(g) := \hat{\boldsymbol{\beta}}(g|g-)$$

$$(16) \quad \text{Var}(\tilde{\boldsymbol{\beta}}(g)) := \text{Var}(\hat{\boldsymbol{\beta}}(g|g-)).$$

2.4. Top-Down Pass: Recursion Step. The updates in the top-down pass can be viewed as solutions to the optimization problem

$$(17) \quad \{\tilde{\boldsymbol{\beta}}(c)\}_{c \in \text{Children}(g)} = \arg \min_{\{\boldsymbol{\beta}(c)\}_{c \in \text{Children}(g)}} \sum_{c \in \text{Children}(g)} \|\text{Var}(\hat{\boldsymbol{\beta}}(c|c-))^{-1/2} (\hat{\boldsymbol{\beta}}(c|c-) - \boldsymbol{\beta}(c))\|_2^2 \text{ such that:}$$

$$\sum_{c \in \text{Children}(g)} \boldsymbol{\beta}(c) = \tilde{\boldsymbol{\beta}}(g).$$

The Karush–Kuhn–Tucker (KKT) conditions of this optimization problem are

$$\text{Var}(\hat{\boldsymbol{\beta}}(c|c-))^{-1} (\tilde{\boldsymbol{\beta}}(c) - \hat{\boldsymbol{\beta}}(c|c-)) = \boldsymbol{\lambda} \quad \forall c \in \text{Children}(g)$$

$$\sum_{c \in \text{Children}(g)} \tilde{\boldsymbol{\beta}}(c) = \tilde{\boldsymbol{\beta}}(g).$$

These two conditions imply that λ can be found using

$$\begin{aligned} \sum_{c \in \text{Children}(g)} \tilde{\beta}(c) - \hat{\beta}(c|c-) &= \sum_{c \in \text{Children}(g)} \text{Var}(\hat{\beta}(c|c-)) \lambda \\ \Rightarrow \tilde{\beta}(g) - \sum_{c \in \text{Children}(g)} \hat{\beta}(c|c-) &= \sum_{c \in \text{Children}(g)} \text{Var}(\hat{\beta}(c|c-)) \lambda \\ \Rightarrow \lambda &= \left(\sum_{c \in \text{Children}(g)} \text{Var}(\hat{\beta}(c|c-)) \right)^{-1} \left(\tilde{\beta}(g) - \sum_{c \in \text{Children}(g)} \hat{\beta}(c|c-) \right). \end{aligned}$$

After substituting this value of λ into the first KKT condition, we have

$$(18) \quad \tilde{\beta}(c) := \hat{\beta}(c|c-) + A(c) \left(\tilde{\beta}(g) - \sum_{c' \in \text{Children}(g)} \hat{\beta}(c'|c'-) \right),$$

where

$$(19) \quad A(c) := \text{Var}(\hat{\beta}(c|c-)) \left(\sum_{c' \in \text{Children}(g)} \text{Var}(\hat{\beta}(c'|c'-)) \right)^{-1}.$$

Next we will derive the update to the variance matrix state variable, *i.e.*, $\text{Var}(\tilde{\beta}(c))$. To do so, first let \tilde{Q}_g be defined so that $\tilde{\beta}(g) = \tau + \tilde{Q}_g \sum_{c \in \text{Children}(g)} \hat{\beta}(c|c-)$, where τ is a random variable that is independent of $\{\hat{\beta}(c|c-)\}_{c \in \text{Children}(g)}$, *i.e.*, τ is the component of $\tilde{\beta}(g)$ that is a linear function of the noisy measurements of vertex g and the vertices that are ancestors of g . Also, let $B_c := \hat{\beta}(c|c-) - A(c) \sum_{c' \in \text{Children}(g)} \hat{\beta}(c'|c'-)$. To derive $\text{Var}(\tilde{\beta}(c))$, we will first derive a few intermediate properties. First, the definition of $A(c)$ implies

$$(20) \quad A(c) \left(\sum_{c' \in \text{Children}(g)} \text{Var}(\hat{\beta}(c'|c'-)) \right) = \text{Var}(\hat{\beta}(c|c-)).$$

Next, we have

$$\begin{aligned} \text{Cov}(B_c, \tilde{\beta}(g)) &= \text{Cov}(\hat{\beta}(c|c-) - A(c) \sum_{c' \in \text{Children}(g)} \hat{\beta}(c'|c'-), \tau + \tilde{Q}_g \sum_{c' \in \text{Children}(g)} \hat{\beta}(c'|c'-)) \\ &= \text{Var}(\hat{\beta}(c|c-)) \tilde{Q}_g^\top - A(c) \left(\sum_{c' \in \text{Children}(g)} \text{Var}(\hat{\beta}(c'|c'-)) \right) \tilde{Q}_g^\top \\ (21) \quad &= \left(\text{Var}(\hat{\beta}(c|c-)) - \text{Var}(\hat{\beta}(c|c-)) \right) \tilde{Q}_g^\top = 0, \end{aligned}$$

where the penultimate equality follows from (20).

Equation (21) implies that $\text{Var}(\tilde{\beta}(c))$ can be written as

$$\begin{aligned} \text{Var}(\tilde{\beta}(c)) &= \text{Cov}(\tilde{\beta}(c), \tilde{\beta}(c)) = \text{Cov}(B_c + A(c) \tilde{\beta}(g), B_c + A(c) \tilde{\beta}(g)) \\ &= \text{Cov}(B_c, B_c) + \text{Cov}(B_c, A(c) \tilde{\beta}(g)) + \text{Cov}(A(c) \tilde{\beta}(g), B_c) + \text{Cov}(A(c) \tilde{\beta}(g), A(c) \tilde{\beta}(g)) \\ &= \text{Var}(B_c) + A(c) \text{Var}(\tilde{\beta}(g)) A(c)^\top \\ &= \text{Var}(\hat{\beta}(c|c-)) - \text{Var}(\hat{\beta}(c|c-)) A(c)^\top - A(c) \text{Var}(\hat{\beta}(c|c-)) \\ &\quad + A(c) \left(\sum_{c' \in \text{Children}(g)} \text{Var}(\hat{\beta}(c'|c'-)) \right) A(c)^\top + A(c) \text{Var}(\tilde{\beta}(g)) A(c)^\top, \end{aligned}$$

so this and equation (20) imply

$$\begin{aligned}
 \text{Var}(\tilde{\beta}(c)) &= \text{Var}(\hat{\beta}(c|c-)) - \text{Var}(\hat{\beta}(c|c-))A(c)^\top - A(c)\text{Var}(\hat{\beta}(c|c-)) \\
 &\quad + \text{Var}(\hat{\beta}(c|c-))A(c)^\top + A(c)\text{Var}(\tilde{\beta}(g))A(c)^\top \\
 (22) \quad &= \text{Var}(\tilde{\beta}(c)) := \text{Var}(\hat{\beta}(c|c-)) - A(c)\text{Var}(\hat{\beta}(c|c-)) + A(c)\text{Var}(\tilde{\beta}(g))A(c)^\top.
 \end{aligned}$$

The following theorem describes the way in which the projection defining the estimator $\tilde{\beta}(c)$ is optimal. This theorem is proved in Appendix A.

Theorem 1. *Suppose that the following conditions hold.*

- (a) $\text{stack}(\{\tilde{\beta}(c)\}_{c \in \text{Children}(g)}) \in \mathbb{R}^{n \text{Card}(\text{Children}(g))}$ is any unbiased estimator of $\text{stack}(\{\beta(c)\}_{c \in \text{Children}(g)})$
- (b) $\text{stack}(\{\tilde{\beta}(c)\}_{c \in \text{Children}(g)})$ is linear in $\tilde{\beta}(g)$ and each element of $\{\hat{\beta}(c|c-)\}_{c \in \text{Children}(g)}$
- (c) $\sum_{c \in \text{Children}(g)} \tilde{\beta}(c) = \tilde{\beta}(g)$

Out of all such estimators, the estimator $\{\tilde{\beta}(c)\}_{c \in \text{Children}(g)}$ defined in (18), is optimal, in the sense that,

$$\text{Var}(\text{stack}(\{\tilde{\beta}(c)\}_{c \in \text{Children}(g)})) \leq \text{Var}(\text{stack}(\{\tilde{\beta}(c)\}_{c \in \text{Children}(g)})).$$

Algorithm 1 provides pseudocode that summarizes this two-pass algorithm, but a description of how this algorithm fits into the motivation for this paper, as described in Section 1.3, may also be helpful. Specifically, the final state variables $\{\tilde{\beta}(u)\}_{u \in \mathcal{G}}$ and $\{\text{Var}(\tilde{\beta}(u))\}_{u \in \mathcal{G}}$ are directly related to the standard WLS estimate, $\tilde{\beta}$, and its variance matrix, as defined in (2-3), because $\tilde{\beta} = \text{stack}(\{\tilde{\beta}(u)\}_{u \in \text{Level}(\mathcal{G}, L)})$ and the diagonal blocks of $\text{Var}(\tilde{\beta})$ are given by $\{\text{Var}(\tilde{\beta}(u))\}_{u \in \text{Level}(\mathcal{G}, L)}$. This implies that, after running this two-pass algorithm, we can take as input $\mathbf{q} \in \mathbb{R}^n$ and $H \subset \text{Level}(\mathcal{G}, L)$ and output $\tilde{\beta}_{H, \mathbf{q}}$, as defined in (5). In cases in which H is the set of leaves that are descendants of a given vertex $u \in \mathcal{G}$, we could also compute $\text{Var}(\tilde{\beta}_{H, \mathbf{q}})$ using $\mathbf{q}^\top \text{Var}(\tilde{\beta}(u)) \mathbf{q}$, but we would not be able to compute $\text{Var}(\tilde{\beta}_{H, \mathbf{q}})$ in more general cases. The next section describes how to compute $\text{Cov}(\tilde{\beta}(u), \tilde{\beta}(v))$ for arbitrary $u, v \in \mathcal{G}$, which, as described in Section 4 below in more detail, can be used to compute $\text{Var}(\tilde{\beta}_{H, \mathbf{q}})$ in all cases.

Algorithm 1: Two-Pass Estimation: Returns matrices and vectors that can be used in Algorithm 3 to generate point estimates and confidence intervals of user-defined queries for an arbitrary set of leaves

input : \mathcal{G} : The rooted tree, with noisy measurements and coefficient matrices associated with each vertex

input : $\{\mathbf{y}(g), S(g), \text{Var}(\mathbf{u}(g))\}_{g \in \mathcal{G}}$: The set of noisy measurements, coefficient matrices, and variance matrices of the noisy measurements for each vertex $g \in \mathcal{G}$

// For each vertex g in graph \mathcal{G} , use equations (7) and (8) to define $\hat{\beta}(g|g)$ and $\text{Var}(\hat{\beta}(g|g))$, respectively. For each vertex $g \in \text{Level}(\mathcal{G}, L)$, let $\hat{\beta}(g|g-)$ and $\text{Var}(\hat{\beta}(g|g-))$ be defined as $\hat{\beta}(g|g)$ and $\text{Var}(\hat{\beta}(g|g))$, respectively, as described in equations (9) and (10).

// For each l in $L-1, \dots, 0$ and each vertex $g \in \text{Level}(\mathcal{G}, l)$, use equations (14) and (13) to define $\hat{\beta}(g|g-)$ and $\text{Var}(\hat{\beta}(g|g-))$.

// Initialize the top-down pass using (15) and (16), i.e., for g given by the root vertex, let $\tilde{\beta}(g)$ and $\text{Var}(\tilde{\beta}(g))$ be defined as $\hat{\beta}(g|g-)$ and $\text{Var}(\hat{\beta}(g|g-))$, respectively.

// For each level l in $1, \dots, L$ and each vertex $c \in \text{Level}(\mathcal{G}, l)$, define $A(c), \tilde{\beta}(c)$ and $\text{Var}(\tilde{\beta}(c))$ using (19), (18), and (22), respectively.

return: $\{\tilde{\beta}(g), \text{Var}(\tilde{\beta}(g)), A(g), \text{Var}(\tilde{\beta}(g|g-))\}_{g \in \mathcal{G}}$

3. COVARIANCE BETWEEN ESTIMATES OF ARBITRARY VERTICES

Similar derivations that were used to derive $\text{Var}(\tilde{\beta}(c))$ in the preceding section can also be carried out to find the covariance $\text{Cov}(\tilde{\beta}(u), \tilde{\beta}(v))$ for arbitrary $u, v \in \mathcal{G}$. Specifically, this involves recursively substituting instances of $\tilde{\beta}(\cdot)$ in the covariance function inputs for its definition in (18), until the only instance of $\tilde{\beta}(\cdot)$ is the closest common ancestor of u and v in the tree \mathcal{G} . Afterward, these functional forms can be simplified significantly using algebraic simplifications and the property (20). We provide these covariance matrices, rather than the underlying derivations in detail, because the derivations themselves are both straightforward and tedious. It turns out that the resulting covariance matrices are very similar to the ones described in Section IV.4 by Willsky (2002), as we describe in more detail below.

Before describing these covariance matrices, some additional notation will be helpful. Specifically, let $u \wedge v$ denote the vertex that is the closest common ancestor of u and v . Also, let $\omega(u, v)$ denote the shortest path from u to v . We also denote the t^{th} element (respectively, the t^{th} element from the end) of this path by $\omega(u, v)[t-1]$ (respectively, $\omega(u, v)[-t]$). In other words, for $u, v \in \mathcal{G}$, the shortest path from u to v is given by $\omega(u, v)[0](=u), \omega(u, v)[1], \dots, \omega(u, v)[-2], \omega(u, v)[-1](=v)$. Also, we use the same notation defined in the last subsection for B_c , i.e., $B_c := \hat{\beta}(c|c-) - A(c) \sum_{c' \in \text{Children}(u \wedge v)} \hat{\beta}(c'|c-)$.

Next we will derive the covariance matrix $\text{Cov}(\tilde{\beta}(u), \tilde{\beta}(v))$ for two specific cases. First, we consider the case in which $u, v \in \mathcal{G}$ are sibling vertices. Note that, because of property (20), we have

$$\begin{aligned} \text{Cov}(B_u, B_v) &= A(u) \left(\sum_{c \in \text{Children}(u \wedge v)} \text{Var}(\hat{\beta}(c|c-)) \right) A(v)^\top - \text{Var}(\hat{\beta}(u|u-))A(v)^\top - A(u)\text{Var}(\hat{\beta}(v|v-)) \\ &= \text{Var}(\hat{\beta}(u|u-))A(v)^\top - \text{Var}(\hat{\beta}(u|u-))A(v)^\top - A(u)\text{Var}(\hat{\beta}(v|v-)) \\ &= -A(u)\text{Var}(\hat{\beta}(v|v-)), \end{aligned}$$

which implies

$$\begin{aligned} \text{Cov}(\tilde{\beta}(u), \tilde{\beta}(v)) &= \text{Cov}(B_u + A(u)\tilde{\beta}(u \wedge v), B_v + A(v)\tilde{\beta}(u \wedge v)) \\ &= \text{Cov}(B_u, B_v) + \text{Cov}(B_u, A(v)\tilde{\beta}(u \wedge v)) \\ &\quad + \text{Cov}(A(u)\tilde{\beta}(u \wedge v), B_v) + \text{Cov}(A(u)\tilde{\beta}(u \wedge v), A(v)\tilde{\beta}(u \wedge v)) \\ (23) \quad &= \text{Cov}(\tilde{\beta}(u), \tilde{\beta}(v)) = A(u)\text{Var}(\tilde{\beta}(u \wedge v))A(v)^\top - A(u)\text{Var}(\tilde{\beta}(v|v-)). \end{aligned}$$

Next we will consider the case in which $g, c \in \mathcal{G}$ and $c \in \text{Children}(g)$. In this case we have

$$\begin{aligned} \text{Cov}(\tilde{\beta}(c), \tilde{\beta}(g)) &= \text{Cov}(B_c + A(c)\tilde{\beta}(g), \tilde{\beta}(g)) \\ &= \text{Cov}(B_c, \tilde{\beta}(g)) + A(c)\text{Cov}(\tilde{\beta}(g), \tilde{\beta}(g)) \\ (24) \quad &= \text{Cov}(\tilde{\beta}(c), \tilde{\beta}(g)) = A(c)\text{Var}(\tilde{\beta}(g)). \end{aligned}$$

Equations (23-24) can also be generalized to describe the functional form of $\text{Cov}(\tilde{\beta}(u), \tilde{\beta}(v))$ for arbitrary $u, v \in \mathcal{G}$. This functional form closely resembles the one provided in Willsky (2002) for the covariance between the values of a random process at two vertices on a rooted tree. The functional form provided by Willsky (2002) is based on the recursive equation describing the coarse-to-fine dynamics of the process, which in our case is (18), and can be viewed as a series of operations involving matrices associated with vertices along the shortest path from u to v in \mathcal{G} , e.g., in our case $\text{Var}(\tilde{\beta}(\cdot))$, $\text{Var}(\hat{\beta}(\cdot|\cdot-))$, and $A(\cdot)$. The covariance matrix provided by Willsky (2002) can be derived by starting from $u \wedge v$ by initializing the covariance matrix as $\text{Var}(\tilde{\beta}(u \wedge v))$ and then recursively moving outward on the path $\omega(u, v)$ from $u \wedge v$, to account for the relationship between $\tilde{\beta}(u \wedge v)$ and vertices further away from $u \wedge v$ along this path. Our functional forms are slightly different than the ones provided by Willsky (2002) because the inclusion of term $-A(c) \sum_{c' \in \text{Children}(g)} \hat{\beta}(c'|c-)$ in (18) prevents the random process resulting from the two pass-approach described in the previous section from satisfying the Markov property exactly. Specifically, while the Markov property requires that, for any $u, v \in \mathcal{G}$, $\tilde{\beta}(u)$ and $\tilde{\beta}(v)$ are independent after conditioning on only $\tilde{\beta}(u \wedge v)$, we

instead have independence between any $\tilde{\beta}(u)$ and $\tilde{\beta}(v)$ after conditioning on $\{\tilde{\beta}(u \wedge v)\} \cup \{\tilde{\beta}(c)\}_{c \in \text{Children}(u \wedge v)}$. This difference between our setting and the model considered by Willsky (2002) requires some additional care when reasoning about the relationship between elements of $\{\tilde{\beta}(u \wedge v)\} \cup \{\tilde{\beta}(c)\}_{c \in \text{Children}(u \wedge v)}$, so we will use the covariance matrices in (23-24) directly when accounting for these relationships.

First we consider cases in which $v \in \mathcal{G}$ is an ancestor of $u \in \mathcal{G}$. In this case, the resulting covariance matrix is identical to the ones provided in Section IV.4 by Willsky (2002). Specifically,

$$(25) \quad \begin{aligned} \text{Cov}(\tilde{\beta}(u), \tilde{\beta}(v)) &= \left(\prod_{k \in \omega(u, v) / \{v, \omega(u, v)[-2]\}} A(k) \right) \text{Cov}(\omega(u, v)[-2], \tilde{\beta}(v)) \\ &= \left(\prod_{k \in \omega(u, v) / v} A(k) \right) \text{Var}(\tilde{\beta}(v)). \end{aligned}$$

Next we consider the case in which $u, v \in \mathcal{G}$ are not ancestors of one another. In this case, we have

$$(26) \quad \begin{aligned} \text{Cov}(\tilde{\beta}(u), \tilde{\beta}(v)) &= \left(\prod_{k \in \omega(u, u \wedge v) / \{\omega(u, u \wedge v)[-2], u \wedge v\}} A(k) \right) \text{Cov}(\tilde{\beta}(\omega(u, u \wedge v)[-2]), \tilde{\beta}(\omega(u \wedge v, v)[1])) \\ &\quad \left(\prod_{k \in \omega(u \wedge v, v) / \{\omega(u \wedge v, v)[1], u \wedge v\}} A(k)^\top \right). \end{aligned}$$

Algorithm 2: Compute_Covariance: Computes the matrix $\text{Cov}(\tilde{\beta}(u), \tilde{\beta}(v))$ for arbitrary $u, v \in \mathcal{G}$

```

input :  $\mathcal{G}$  : The rooted tree
input :  $\{\tilde{\beta}(g), \text{Var}(\tilde{\beta}(g)), A(g), \text{Var}(\hat{\beta}(g|g-))\}_{g \in \mathcal{G}}$  : the output of Algorithm 1
input :  $u, v \in \mathcal{G}$  : This algorithm returns  $\text{Cov}(\tilde{\beta}(u), \tilde{\beta}(v))$ 
if  $u = v$  then
    return:  $\text{Var}(\tilde{\beta}(u))$ 
end
if  $v = u \wedge v$  then
    // See equation (25) for this case:
    return:  $\left( \prod_{k \in \omega(u, v) / v} A(k) \right) \text{Var}(\tilde{\beta}(v))$ 
end
if  $u = u \wedge v$  then
    // Switch the inputs  $u, v$  and transpose the output of this function:
    return:  $\text{Compute\_Covariance}(\mathcal{G}, \{\tilde{\beta}(g), \text{Var}(\tilde{\beta}(g)), A(g), \text{Var}(\hat{\beta}(g|g-))\}_{g \in \mathcal{G}}, v, u)^\top$ 
end
// The remaining case follows from (26):
 $u' \leftarrow \omega(u, u \wedge v)[-2]$ 
 $v' \leftarrow \omega(u \wedge v, v)[1]$ 
 $\text{Cov}(\tilde{\beta}(u'), \tilde{\beta}(v')) \leftarrow A(u') \text{Var}(\tilde{\beta}(u' \wedge v')) A(v')^\top - A(u') \text{Var}(\hat{\beta}(v'|v'-))$ 
return:  $\left( \prod_{k \in \omega(u, u') / \{u'\}} A(k) \right) \text{Cov}(\tilde{\beta}(u'), \tilde{\beta}(v')) \left( \prod_{k \in \omega(v', v) / \{v'\}} A(k)^\top \right)$ 
```

When $u, v \in \text{Level}(\mathcal{G}, L)$, the derivations in this section can be used to derive the off-diagonal blocks of $\text{Var}(\tilde{\beta})$, as defined in (3), so, combined with the derivation for the diagonal blocks provided in the previous section, these results allow for the computation of arbitrary elements of the variance matrix of the standard WLS estimate.

Algorithm 2 provides pseudocode that summarizes the methods described in this section. The next section describes how Algorithms 1 and 3 can be used to compute $\tilde{\beta}_{H,q}$, as defined in (5), its variance, and confidence intervals for this estimate.

4. CONFIDENCE INTERVAL ESTIMATION

The previous two sections describe each of the state variable updates required for the two-pass estimation approach and the computation of $\text{Cov}(\tilde{\beta}(u), \tilde{\beta}(v))$ for arbitrary $u, v \in \mathcal{G}$, as summarized in Algorithms 1 and 2, respectively. It is straightforward to use these two algorithms to derive confidence intervals centered at $\tilde{\beta}_{H,q}$, which is done below in Algorithm 3.

Algorithm 3: Estimate_Confidence_Interval: Returns the full-information BLUE count estimate and confidence interval of a user-defined query evaluated in an arbitrary subset of the leaf vertices of \mathcal{G}

```

input :  $\mathcal{G}$  : The rooted tree
input :  $\{\tilde{\beta}(g), \text{Var}(\tilde{\beta}(g)), A(g), \text{Var}(\tilde{\beta}(g|g-))\}_{g \in \mathcal{G}}$  : the output of Algorithm 1
input :  $q \in \mathbb{R}^n$  : The linear counting query of interest
input :  $H \subset \text{Level}(\mathcal{G}, L)$  : The subset of the leaves of  $\mathcal{G}$  for which to estimate the linear query
input :  $\alpha$  : This algorithm will return the  $1 - \alpha$  confidence interval of  $\beta_{H,q}$ 
vertices  $\leftarrow \emptyset$ 
// For each level  $l$  in  $L - 1, \dots, 0$  and for each vertex  $g \in \text{Level}(\mathcal{G}, l)$ , if all children
  of  $g$  are in  $H$ , remove Children( $g$ ) from  $H$  and add vertex  $g$  to  $H$ .
Var( $\tilde{\beta}_{H,q}$ )  $\leftarrow 0$ 
 $\tilde{\beta}_{H,q} \leftarrow 0$ 
for  $u \in H$  do
   $\tilde{\beta}_{H,q} \leftarrow \tilde{\beta}_{H,q} + q^\top \tilde{\beta}(u)$ 
  for  $v \in H$  do
    Cov( $\tilde{\beta}(u), \tilde{\beta}(v)$ )  $\leftarrow \text{Compute\_Covariance}(\mathcal{G}, \{\tilde{\beta}(g), \text{Var}(\tilde{\beta}(g)), A(g), \text{Var}(\tilde{\beta}(g|g-))\}_{g \in \mathcal{G}}, u, v)$ 
    Var( $\tilde{\beta}_{H,q}$ )  $\leftarrow \text{Var}(\tilde{\beta}_{H,q}) + q^\top \text{Cov}(\tilde{\beta}(u), \tilde{\beta}(v)) q$ 
  end
end
//  $\Phi(\cdot)$  denotes the cumulative distribution function of the standard normal
  distribution:
 $c \leftarrow \sqrt{\text{Var}(\tilde{\beta}_{H,q})} \Phi^{-1}(1 - \alpha/2)$ 
return:  $\tilde{\beta}_{H,q}, (\tilde{\beta}_{H,q} - c, \tilde{\beta}_{H,q} + c)$ 

```

A number of theoretical properties can also be shown for the estimate $\tilde{\beta}_{H,q}$ and its confidence interval using standard results on the weighted least squares estimator. Recall that in Section 1.3 we assumed that $S \in \mathbb{R}^{M \times N}$ has full column rank and that $\mathbf{u} \in \mathbb{R}^M$ is a vector of independent mean-zero random variables, each with a known and finite variance. As a result of these assumptions, the Gauss-Markov theorem implies that the estimate $\tilde{\beta}_{H,q}$ that is output from Algorithm 3 has the smallest possible variance out of all estimators that are unbiased and linear in the noisy measurements \mathbf{y} ; see for example, (Davidson et al., 2004). In addition, if $\mathbf{u}(g)$ is normally distributed for each $g \in \mathcal{G}$, then $\tilde{\beta}_{H,q}$ is also normally distributed, which can be shown using standard finite sample properties of the WLS estimator (Davidson et al., 2004). As a result, when this additional normality condition holds, this finite sample property implies that the confidence intervals output from Algorithm 2 are statistically valid in this case.

When $\mathbf{u}(g)$ is not normally distributed (but our other assumptions still hold), we could use standard asymptotic theory for regressions to show the same result, *i.e.*, normality of $\tilde{\beta}_{H,q}$ along with its same implication of statistical validity of the confidence intervals output from Algorithm 2, holds in some limiting cases Davidson et al. (2004). However, these limiting cases are not realistic in our primary motivating use case, as they require the number of noisy measurements to diverge, so we will not expand on these further.

After describing our implementation of Algorithm 3 in the next subsection, we will provide a numerical experiment to assess the accuracy of these confidence intervals in practice for our motivating use case. While \mathbf{u} is not distributed as a Gaussian in this motivating use case, there are still reasons to be optimistic about the accuracy of the confidence intervals output from Algorithm 3; for example, as described in Section 1.2, the distribution of \mathbf{u} is closely related to the Gaussian distribution, *i.e.*, it is the discrete Gaussian distribution Canonne et al. (2020).

4.1. Computational Considerations. For many use cases, the limiting factor for the feasibility of the algorithms proposed in this paper will be the need to store many dense $n \times n$ matrices. For example, one possible use case would be to apply these methods to the NMFs for the Redistricting Data File, and, as described above, in this case n is equal to 2,016. Since our choice of output for Algorithm 1, *i.e.*, $\{\tilde{\beta}(g), \text{Var}(\tilde{\beta}(g)), A(g), \text{Var}(\hat{\beta}(g|g-))\}_{g \in \mathcal{G}}$, requires storing three dense $n \times n$ matrices for every vertex and since \mathcal{G} would have approximately 6 million leaf vertices when using the production US DAS execution Redistricting Data File NMFs, at least several hundreds of terabytes would be required to store the output of Algorithm 1 without using methods to reduce storage space in this case. This is problematic because it would be ideal to store the output of Algorithm 1 after it is run only once, so that this output can be used to support the computation of confidence intervals of linear query answers in arbitrary set of leaves over an extended period of time. With this in mind, this section describes multiple ways to reduce the storage requirements of the output of Algorithm 1.

For the primary motivating use case we have in mind, one of the most straightforward and effective ways to keep space requirements of the output of Algorithm 1 low would be to reduce the size of the histogram schema by marginalizing the NMFs over certain levels of attributes in the original schema. For example, in the case of the NMFs for the Redistricting Data File, rather than running Algorithm 1 on the full schema HHGQ \times HISPANIC \times VOTING-AGE \times CENRACE, we could marginalize the NMFs over all attributes other than HHGQ, save the output of Algorithm 1, and then repeat this process for each of the three remaining attributes. Given the number of levels of each of these attributes, as provided in Section 1.2, this alternative would reduce the space requirements for all of the outputs of Algorithm 1 by approximately a factor of 1,005.8 ($\approx 8 \times 2 \times 2 \times 63 / (8 + 2 + 2 + 63)$), and would also substantially reduce the computational requirements of the executions of this algorithm. Likewise, calling Algorithm 1 after marginalizing over each possible pair of these four attributes would reduce the space requirements by a factor of approximately 14.2 relative to using all attributes simultaneously. The downside of this approach is that we could only support the computation of user-defined queries that depend on only a subset of the attributes; however, in the vast majority of use cases for decennial census data, the queries of interest do not depend on all attributes of the full histogram schema used to generate the NMFs. For this reason, in practice, we expect that use cases using decennial census NMFs will most often only require $n \leq 2,016$. Note that, unlike the NMFs for Redistricting Data File DAS executions, the schemas used to generate the NMFs for the persons and units universes of the DHC DAS executions have more than 1.2 million detailed histogram cells, so marginalizing over attributes is a requirement for computational feasibility of our proposed approach in these cases.

There are also several algorithmic modifications that can be considered to reduce the size of the output of Algorithm 1. First, storage of the matrices $\{A(g)\}_g$ can optionally be avoided entirely, since these matrices can be recomputed using $\{\text{Var}(\tilde{\beta}(g)), \text{Var}(\hat{\beta}(g|g-))\}_{g \in \mathcal{G}}$ when they are required. Second, one can also leverage the fact that, in many use cases, some of the matrices output from Algorithm 1 will be the same, to avoid storing multiple instances of duplicate matrices. For example, for our primary use case of interest, vertices in the lowest level of \mathcal{G} have the matrix $\text{Var}(\hat{\beta}(g|g-))$ given by one of only two possible values. Third, one can use the fact that the variance matrices in $\{\text{Var}(\tilde{\beta}(g)), \text{Var}(\hat{\beta}(g|g-))\}_{g \in \mathcal{G}}$ are symmetric to reduce the storage requirement of these matrices by approximately one half. Fourth, the storage requirement of the output of Algorithm 1 can also be reduced using a compressed format. In the numerical experiments provided in the next section, we use these final two methods to reduce the space requirements for the output of Algorithm 1.

5. NUMERICAL EXPERIMENT

In this section we provide a numerical experiment using a Python/PySpark implementation of the algorithms we proposed above. For the purposes of this numerical experiment, this implementation was executed on Amazon Web Services (AWS) using Elastic Map Reduce (EMR) version 6.12 and Apache Spark version

TABLE 2. CI Empirical Coverage: For each of the four CIs considered, and each of the three sets of queries considered, this table provides the CI coverage, which is the proportion of CIs that contain the true CEF-based count out of all queries in the given query set for all block groups in PR.

Query	CI Type			
	0.90 CI	0.95 CI	0.90 CI, Nonnegative	0.95 CI, Nonnegative
All Redistricting Tabulations	0.8994	0.9496	0.9473	0.9737
Total Population	0.8922	0.9363	0.8922	0.9363
CENRACE	0.9002	0.9498	0.9479	0.9739

3.4. The DAS is run on EMR clusters built from AWS r5.24xlarge virtual machines, in which each virtual machine has 96 cores and 768GiB of RAM. We configured the cluster so that it had one primary node and 5 core nodes, and configured Spark so that there were 5 executors, each with 40 executor cores.

Rather than using all noisy measurements generated by a full US DAS execution, in this numerical experiment, we used the publicly available NMFs from the persons universe Redistricting Data File Puerto Rico (PR) DAS execution of the 2010 Demonstration Data Product; instructions for downloading these NMFs are available at U.S. Census Bureau (2023b). This DAS execution used settings that were identical to that of the 2020 production DAS execution that was used to publish the 2020 PR persons tables in the Redistricting Data File, other than the fact that the 2010 Census data was used as input instead of the 2020 Census data. As described previously, unlike the tabulation US spine described in Figure 1, the geographic spine that is used internally within the DAS implementation for PR executions is rooted at the geounit corresponding to PR as a whole, which is included in the state geolevel, and subsequent geolevels include county, tract, optimized block group, and the block geolevels. Optimized block groups are distinct from the census block groups that are included in the tabulation spine described above, so the 2,543 PR tabulation census block groups are not on the optimized spine used internally within DAS; for more detail on the optimized spine, see Cumings-Menon et al. (2022). Since these PR tabulation census block groups are not in the spine used to generate the NMFs, they provide interesting off-spine geographic entities for the purpose of testing our proposed estimation approaches, so we will focus on the estimation of CIs of 2010 tabulations for PR census block groups.

We will consider CIs at the 0.9 and 0.95 confidence levels in particular. In the context of this numerical experiment, we will use *empirical coverage*, or *coverage*, of a set of CIs to refer to the proportion of the CIs in this set that contain the true population count. We will estimate the coverage of several CIs for several sets of queries in this section; note that a $1 - \alpha$ confidence level CI should ideally have a coverage that is at least approximately $1 - \alpha$. Note that one complication in this numerical example is that the point estimates (and therefore the CIs) are not entirely independent with one another. We do not expect the impact of this issue to be large, because the strength of the dependence between typical pairs of PR census block groups becomes weaker as the fanout values, *i.e.*, the number of child geounits of parent geounits, increases, and the fanout values in the spine are fairly high, especially at higher geolevels. For example, the county geolevel is composed of the 78 municipios of PR.

Since we know each true count is nonnegative, we can optionally increase the coverage of each CI, and reduce its width, by rounding the endpoints of the CI up to zero in the event they are negative, which we will refer to as a *nonnegative CI*. Tables 2 and 3 provide the empirical coverage and the average width, respectively, of the 0.9 and 0.95 CIs, both before and after rounding up to zero for three sets of queries. Specifically, we use “All Redistricting Tabulations” to refer to all persons Redistricting Data File tabulations published in 2020, *i.e.*, tables P1, P2, P3, P4, and P5, as described by U.S. Census Bureau (2023a), “Total Population” to refer to the total population, and “CENRACE” to refer to the collection of queries that provide the population counts in each of the 63 race combination categories of this attribute, as described in Section 1.2. For each of these three sets of queries, Tables 2 and 3 consider the CIs for each query in the set, evaluated in each 2010 vintage PR tabulation census block group.

Table 2 shows that coverage of the CIs before rounding up to zero, has a coverage that is fairly close to the confidence level of the CI considered, for each of the three sets of queries. The final two columns demonstrate that rounding these endpoints up to zero improves the coverage of these CIs, particularly for CIs of more granular queries that are more likely to have negative CI endpoints. Table 3 also shows that

TABLE 3. CI Averaged Widths: For each of the four CIs considered, and each of the three sets of queries considered, this table provides the CI width, averaged over all queries in the given query set and all block groups in PR.

Query	CI Type			
	0.90 CI	0.95 CI	0.90 CI, Nonnegative	0.95 CI, Nonnegative
All Redistricting Tabulations	100.8	120.1	54.84	65.74
Total Population	314.4	374.6	313.1	373.2
CENRACE	97.81	116.6	50.66	60.78

rounding up to zero can significantly narrow the widths of these same CIs of more granular queries. This later table also shows that the widths of these CIs are fairly wide on average. This is because census block groups were not on the optimized spine used internally in Redistricting Data file DAS executions. Note that one possibility to increase accuracy further would be to also leverage the information in the DHCP NMFs, after marginalizing over all cells other than the 2,016 detailed histogram cells in the Redistricting Data file schema, as described in Section 1.2. However, as described in Cumings-Menon et al. (2022), the DHC DAS executions used a different spine than the Redistricting Data file DAS executions, so if this is done, Algorithm 1 would need to be run separately on both sets of NMFs. Afterward, for each query and each geographic entity, the WLS-based point estimates and variances could be computed separately using each of the outputs of Algorithm 1. If these point estimates are denoted by $\tilde{\beta}_{H,q,1}, \tilde{\beta}_{H,q,2} \in \mathbb{R}$, then the inverse-variance weighted query estimate would be equal to

$$(27) \quad \check{\beta}_{H,q} = \left(\tilde{\beta}_{H,q,1} / \text{Var}(\tilde{\beta}_{H,q,1}) + \tilde{\beta}_{H,q,2} / \text{Var}(\tilde{\beta}_{H,q,2}) \right) / \left(1 / \text{Var}(\tilde{\beta}_{H,q,1}) + 1 / \text{Var}(\tilde{\beta}_{H,q,2}) \right),$$

which has a variance given by

$$(28) \quad \text{Var}(\check{\beta}_{H,q}) = 1 / \left(1 / \text{Var}(\tilde{\beta}_{H,q,1}) + 1 / \text{Var}(\tilde{\beta}_{H,q,2}) \right),$$

and then the final confidence interval can be computed using the same approach as in Algorithm 3, *i.e.*, with endpoints given by $\check{\beta}_{H,q} \pm \sqrt{\text{Var}(\check{\beta}_{H,q})} \Phi^{-1}(1 - \alpha/2)$.


Note that the analyses above assess statistical validity of these CIs at only two confidence levels. To provide a way to assess statistical validity at all confidence levels, we also constructed the Z -scores of a set of estimates, which for each estimate $\check{\beta}_{H,q}$ is defined as,

$$(29) \quad z_{H,q} = \left(\check{\beta}_{H,q} - \beta_{H,q} \right) / \sqrt{\text{Var}(\check{\beta}_{H,q})}.$$

If $\mathbf{u}(g)$ were normally distributed for each $g \in \mathcal{G}$, then we would have $z_{H,q} \sim N(0, 1)$, so we can assess the error due to our normality modeling assumption by computing the Z -scores for a set of queries to see how closely these scores appear to follow a standard normal distribution. To do this, we computed the Z -scores for the total population query for each of the 2,543 census block groups in PR. Figure 2 provides a Q-Q plot based on these Z -scores. For each blue point, the horizontal axis provides the empirical quantiles of the 2,543 sample Z -scores and the vertical axis provides the corresponding theoretical quantile value based on a standard normal distribution. The red line provides the identity function. Since the blue points are very close to the identity function, this set of Z -scores appear to follow a Gaussian distribution quite closely. This plot also implies that, for any $\alpha \in (0, 1)$, the $1 - \alpha$ CIs for this set of queries would have empirical coverage close to $1 - \alpha$.

To provide an idea of the computational requirements of our implementation of the proposed algorithms for the numerical experiment explored here, the total runtime of our implementation of Algorithm 1 was 8.6 hours, and estimating all of the CIs afterward took approximately 17.1 hours. Our implementation of Algorithm 1 stored its output in AWS S3 storage in a Apache Parquet format, and the size of this final output was approximately 795.6 GB for all geolevels in this compressed format, with the vast majority of this space (699.8 GB) being required to store the block geolevel matrices. The total number of blocks in PR in 2010 with either a housing unit or a GQ, *i.e.*, the universe of blocks we consider in this numerical example, was 51,937, which is about 1/113 of the total number of blocks in the spine of the 2020 US DAS

FIGURE 2. Q-Q Plot: The blue points in this plot can be viewed as points on a parametric curve, parameterized by the quantile $q \in (0, 1)$. The horizontal axis provides the q quantile of the standard normal distribution and the vertical axis provides the sample quantile of the Z -score. The red line provides the identity function, $f(x) = x$.



optimized spine. Given this difference in the scale of the 2020 US and the 2010 PR DAS optimized spines and the computational requirements of this numerical experiment for the 2010 PR NMFs, a similar exercise may also be possible for the 2020 US production NMFs, albeit with a more significant computational cost. Also, the persons universe Redistricting Data File DAS schema likely has a number of detailed cells, *i.e.*, $n = 2,016$, near the highest value that would still allow for estimation in the 2020 US DAS optimized spines without marginalizing prior to estimation. As described previously, estimation using the 2020 US DHCP or DHCH NMFs would require marginalizing the NMFs to smaller(/coarser) histograms prior to estimation, since both of these two schemas have more than $n = 1.2$ million detailed cells for each geounit.

6. COMPARISON WITH HAY ET AL. (2010)

As described above, the approach proposed in this paper can be viewed as a generalization of the approach proposed by Hay et al. (2010), which we will describe in more detail in this section. To do so, we will assume we are in the setting described by Hay et al. (2010). In other words, we will suppose that, for every $g \in \mathcal{G}$, $S(g)$ is equal to the 1×1 matrix $[1]$, each vertex $g \in \mathcal{G}$ has $k \in \mathbb{N}$ children, and that the PLB allocated to each query, and therefore also the variance of each noisy measurement, is the same. Although it is not assumed by Hay et al. (2010), we will also suppose the variance of each query is equal to one specifically, since this variance does not appear in their algorithm and it simplifies the derivations below slightly.

In this setting Hay et al. (2010) show that the OLS estimator can be found using a two-step approach. For each $g \in \mathcal{G}$, before stating this theorem, they first define $z(g) \in \mathbb{R}$, which in our notation is given by

$$(30) \quad z(g) = \begin{cases} y(g) & g \in \text{Level}(\mathcal{G}, L) \\ \left(\frac{k^{L-l+1} - k^{L-l}}{k^{L-l+1} - 1} \right) y(g) + \left(\frac{k^{L-l} - 1}{k^{L-l+1} - 1} \right) \sum_{c \in \text{Children}(g)} z(c) & g \in \text{Level}(\mathcal{G}, l) \text{ and } l \neq L \end{cases}.$$

The relevant result from Hay et al. (2010) is stated below.

Theorem 2. (Theorem 3 from Hay et al. (2010)) *If, for every $g \in \mathcal{G}$, $S(g)$ is equal to the matrix $[1]$, each vertex $g \in \mathcal{G}$ has $k \in \mathbb{N}$ children, and the variance of each noisy measurement is one, then the WLS solution can be found using the recurrence relation*

$$(31) \quad \check{\beta}(v) = \begin{cases} z(v) & v \in \text{Level}(\mathcal{G}, 0) \\ z(v) + \frac{1}{k} \left(\check{\beta}(g) - \sum_{c \in \text{Children}(v)} z(c) \right) & v \notin \text{Level}(\mathcal{G}, 0) \text{ and } v \in \text{Children}(g) \end{cases}.$$

Below we show that our estimator is a special case of the estimator described by Hay et al. (2010).

Theorem 3. *If, for every $g \in \mathcal{G}$, $S(g)$ is equal to the matrix $[1]$, each vertex $g \in \mathcal{G}$ has $k \in \mathbb{N}$ children, and the variance of each noisy measurement is one, then $\tilde{\beta}(g) = \hat{\beta}(g)$.*

Proof. First we will show that $z(g) = \hat{\beta}(g|g-)$ under the assumptions of the theorem. To do so, for $g \in \text{Level}(\mathcal{G}, l)$, we will begin by deriving a convenient functional form for $v(l) := \text{Var}(\hat{\beta}(g|g-))$ under the conditions of the theorem. Using the definition from Section 2.1 and the assumptions of this theorem, for any $g \in \text{Level}(\mathcal{G}, L)$, we have $v(L) = \text{Var}(\mathbf{y}(g)) = 1$. Using notation introduced in Section 2.2, we have $\text{Var}(\hat{\beta}(g|\text{Children}(g)-)) = kv(l)$ for any $g \in \text{Level}(\mathcal{G}, l-1)$, so the variance of the inverse variance weighted mean implies

$$v(l-1) = \frac{1}{(1 + 1/(kv(l)))} = 1 - \frac{1}{1 + kv(l)}$$

After solving this recursive equation with the boundary condition that $v(L) = 1$, we have,

$$v(l) = \frac{(k-1)k^L}{k^{L+1} - k^l}$$

Note that, for $g \in \text{Level}(\mathcal{G}, l)$, equation (13) implies

$$\hat{\beta}(g|g-) = v(l) \left(\mathbf{y}(g) + \frac{1}{kv(l+1)} \sum_{c \in \text{Children}(g)} \hat{\beta}(c|c-) \right).$$

Note that, for any $g \in \text{Level}(\mathcal{G}, L)$, we clearly have $z(g) = \hat{\beta}(g|g-)$. Also, in the case of vertices in levels above level L , the definition of $z(g)$ and $\hat{\beta}(g|g-)$ follows a similar form, in the sense that, for any $g \in \text{Level}(\mathcal{G}, l)$ with $l < L$, both estimators are weighted means of $\mathbf{y}(g)$ and the sum over the estimates of the children of g . In the case of both $z(g)$ and $\hat{\beta}(g|g-)$, the weights of these two terms sum to one, so we will establish that $z(g) = \hat{\beta}(g|g-)$ for all $g \in \text{Level}(\mathcal{G}, l)$ with $l < L$ by showing that the ratio of these weights are the same for both of these estimators. In the case of $\hat{\beta}(g|g-)$, this ratio of these weights is given by $kv(l+1)$, and in the case of $z(g)$ this ratio is given by

$$\frac{k^{L-l+1} - k^{L-l}}{k^{L-l} - 1}.$$

Thus, to establish that $z(g) = \hat{\beta}(g|g-)$ for all $g \in \text{Level}(\mathcal{G}, l)$ with $l < L$, we need to show

$$\begin{aligned} \frac{k^{L-l+1} - k^{L-l}}{k^{L-l} - 1} = kv(l+1) &\iff \frac{k^{L-l+1} - k^{L-l}}{k^{L-l} - 1} = \frac{(k-1)k^{L+1}}{k^{L+1} - k^{l+1}} \\ &\iff \frac{k^L(k-1)}{k^L - k^l} = \frac{k^L(k-1)}{k^L - k^l}, \end{aligned}$$

which implies $z(g) = \hat{\beta}(g|g-)$ for all $g \in \mathcal{G}$.

The final result follows from the fact that, after replacing $z(g)$ with $\hat{\beta}(g|g-)$ in the recurrence relation in (31), this recursive formula is equal to the top-down pass described in Section 2. \square

7. CONCLUSION

This paper describes a two-pass estimation approach for hierarchical data that is capable of providing WLS-based estimates and confidence intervals. We also provide a numerical exercise to demonstrate feasibility of the approach for our motivating use case. As described in the previous section, this two-pass approach can be viewed as a generalization of the approach described by Hay et al. (2010). This section concludes by describing some advantages of our proposed approach relative to another approach in the context of our primary motivating use case, *i.e.*, estimating tabulations and their corresponding confidence intervals using census NMFs, as described in Section 1.3.

In the context of our motivating use case, an alternative to using the WLS estimator described above is to simply derive point estimates from the NMFs directly. For example, one can estimate the total population of a geographic entity that corresponds to a geounit using the total population primitive DP answer of the geounit directly. An alternative estimator for this same query answer can be derived by summing over the total population noisy measurements of the children of this geounit. More generally, in this section, we will refer to all estimates that can be defined by adding and/or subtracting a subset of the scalars in the

NMFs as *naïve estimates* of a tabulation. For use cases that do not require high precision, and only require tabulation estimates for geographic entities that are geounits on the spine used to generate the NMFs, these naïve estimates may be adequate. However, unlike the naïve estimates, the full-information WLS estimator provides a unique estimate, with the added benefit of also providing the lowest possible variance out of all estimators that are unbiased linear functions of the noisy measurements.

These advantages of the two-pass WLS estimator are most obvious when estimating queries for a geographic entity that is far from the spine used to generate the noisy measurements, meaning many geounits would need to be added and subtracted from one another to derive the geographic entity. For example, when deriving a total population estimate for Navajo Nation, it would not be obvious how to choose a relatively small set of noisy measurements that provides an unbiased naïve estimate with a sufficiently low variance. Examples of approaches that are capable of constructing these subsets of the NMFs would include considering adding the NMFs at the block geolevel NMFs, considering the NMFs for the least-granular set of geounits that are entirely within a geounit on the spine, *i.e.*, the approach described in the first comment of Algorithm 3, or considering the NMFs of the smallest set of geounits that can be used to derive the off-spine geographic entity by adding and/or subtracting the geounits in this set.⁴ As a result, data users confronted with this issue when using one of the naïve estimators would often choose different sets of noisy measurements and thus potentially arrive at dissimilar estimates. In contrast, the WLS estimator proposed here allows users to construct the unique minimum-variance point estimates, and their corresponding CIs using a normality approximation, even for geographic entities that are far from the geographic spine. We expect data users will view the combination of accuracy, approximate normality, and uniqueness provided by the WLS-based tabulation estimators to be advantageous.

It is also worth pointing out that the Census Bureau is developing an alternative method to compute confidence intervals, based on simulating the DAS using the 2020 production DAS execution outputs, *i.e.*, the microdata detailed files (MDFs), as the input of each simulation iteration. This simulation-based approach appears to perform well, with CIs that are typically more narrow than the WLS-based CIs of the approach described in this paper, since this simulation-based approach leverages the increased accuracy due to the inequality constraints used within DAS. The main advantage of the WLS-based approach described here over this simulation-based approach is that we avoid the assumption that DAS runs using the private census data as input have similar error distributions as DAS runs that use the 2020 MDFs as input. Since both approaches have advantages, which of these two approaches is best for a given application will typically be dependent on the specific goals of the use case.

REFERENCES

- Abowd, J. M., Ashmead, R., Cumings-Menon, R., Garfinkel, S., Heineck, M., Heiss, C., Johns, R., Kifer, D., Leclerc, P., Machanavajjhala, A., et al. (2022). The 2020 Census disclosure avoidance system topdown algorithm. *Harvard Data Science Review*, (Special Issue 2).
- Bun, M. and Steinke, T. (2016). Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer.
- Canonne, C. L., Kamath, G., and Steinke, T. (2020). The discrete Gaussian for differential privacy. *Advances in Neural Information Processing Systems*, 33:15676–15688.
- Cumings-Menon, R., Abowd, J. M., Ashmead, R., Kifer, D., Leclerc, P., Ocker, J., Ratcliffe, M., and Zhuravlev, P. (2022). Geographic spines in the 2020 census disclosure avoidance system. *arXiv preprint arXiv:2203.16654*.
- Cumings-Menon, R., Ashmead, R., Kifer, D., Leclerc, P., Spence, M., Zhuravlev, P., and Abowd, J. M. (2023). Disclosure avoidance for the 2020 Census Demographic and Housing Characteristics File. *arXiv preprint arXiv:2312.10863*.
- Davidson, R., MacKinnon, J. G., et al. (2004). *Econometric theory and methods*, volume 5. Oxford University Press New York.
- Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer.
- Hay, M., Rastogi, V., Miklau, G., and Suciu, D. (2010). Boosting the accuracy of differentially private histograms through consistency. *Proceedings of the VLDB Endowment*, 3(1).

⁴This latter set of geounits is described in more detail in Cumings-Menon et al. (2022).

U.S. Census Bureau (2023a). Decennial Census P.L. 94-171 Redistricting Data.

U.S. Census Bureau (2023b). Developing the DAS: Demonstration Data and Progress Metrics.

Willsky, A. S. (2002). Multiresolution markov models for signal and image processing. *Proceedings of the IEEE*, 90(8):1396–1458.

APPENDIX A. RESULT ON OPTIMALITY OF TOP-DOWN PASS PROJECTION

This section provides the proof of Theorem 1. Before doing this, we first prove three lemmas that provide properties of the state variables for the parent vertex $g \in \mathcal{G}$. To do this, some additional notation will also be helpful. Specifically, for this parent vertex $g \in \mathcal{G}$, let $\kappa := \text{Card}(\text{Children}(g))$ and $E \in \mathbb{R}^{n\kappa \times n}$ be defined as

$$E := \begin{bmatrix} I_n \\ \vdots \\ I_n \end{bmatrix}.$$

Also, let $\tilde{\beta}_g$ be defined as $\text{stack}(\{\tilde{\beta}(c)\}_{c \in \text{Children}(g)})$, where $\tilde{\beta}(c)$ is defined in (18), and let $\hat{\beta}_g$ be defined as $\text{stack}(\{\hat{\beta}(c|c-)\}_{c \in \text{Children}(g)})$, where $\hat{\beta}(c|c-)$ is defined in (14).

The first lemma in this section provides the functional form for the matrix \tilde{Q}_g , as introduced in Section 2.4.

Lemma 2. *In the top-down pass, for each parent vertex $g \in \mathcal{G}$, \tilde{Q}_g satisfies*

$$(32) \quad \tilde{Q}_g = \text{Var}(\tilde{\beta}(g)) \left(\sum_{c \in \text{Children}(g)} \text{Var}(\hat{\beta}(c|c-)) \right)^{-1}.$$

Proof. The initialization step of the top-down pass implies that for the root vertex g , *i.e.*, equations (15)-(16), we have $\tilde{\beta}(g) = \hat{\beta}(g|g-)$, and

$$\text{Var}(\tilde{\beta}(g)) = \text{Var}(\hat{\beta}(g|g-)) = \left(\text{Var}(\hat{\beta}(g|g))^{-1} + \text{Var}(\hat{\beta}(g|\text{Children}(g)-))^{-1} \right)^{-1}.$$

Thus, equation (14) implies

$$\tilde{Q}_g = \text{Var}(\hat{\beta}(g|g-)) \text{Var}(\hat{\beta}(g|\text{Children}(g)-))^{-1} = \text{Var}(\tilde{\beta}(g)) \left(\sum_{c \in \text{Children}(g)} \text{Var}(\hat{\beta}(c|c-)) \right)^{-1},$$

so the property in the statement of this result, *i.e.*, equality (32), holds for the root vertex.

Now we will suppose equality in (32) holds for an arbitrary parent vertex $g \in \mathcal{G}$ and we show that this implies it also holds for the children of g . To do so, first recall that, for any vertex $p \in \text{Children}(g)$, equation (18) implies

$$\begin{aligned} \tilde{\beta}(p) &= \hat{\beta}(p|p-) + A(p) \left(\tilde{\beta}(g) - \sum_{p' \in \text{Children}(g)} \hat{\beta}(p'|p'-) \right) \\ &= \hat{\beta}(p|p-) + A(p) \left(\tau_g + \tilde{Q}_g \left(\sum_{p' \in \text{Children}(g)} \hat{\beta}(p'|p'-) \right) - \sum_{p' \in \text{Children}(g)} \hat{\beta}(p'|p'-) \right). \end{aligned}$$

By (14), we can substitute each instance of $\hat{\beta}(p|p-)$ with

$$\tau_p + \text{Var}(\hat{\beta}(p|p-)) \left(\sum_{c \in \text{Children}(p)} \text{Var}(\hat{\beta}(c|c-)) \right)^{-1} \left(\sum_{c \in \text{Children}(p)} \hat{\beta}(c|c-) \right),$$

where $\tilde{\tau}_p$ is a random vector that is independent of $\hat{\beta}(c|c-)$, for each $c \in \text{Children}(p)$. Thus, for τ_p defined as another random vector that is independent of $\hat{\beta}(c|c-)$, for each $c \in \text{Children}(p)$, we have

$$\begin{aligned}\tilde{\beta}(p) &= \tau_p + \text{Var}(\hat{\beta}(p|p-)) \left(\sum_{c \in \text{Children}(p)} \text{Var}(\hat{\beta}(c|c-)) \right)^{-1} \left(\sum_{c \in \text{Children}(p)} \hat{\beta}(c|c-) \right) \\ &\quad + A(p) (\tilde{Q}_g - I) \text{Var}(\hat{\beta}(p|p-)) \left(\sum_{c \in \text{Children}(p)} \text{Var}(\hat{\beta}(c|c-)) \right)^{-1} \left(\sum_{c \in \text{Children}(p)} \hat{\beta}(c|c-) \right) \\ &= \tau_p + (I + A(p) (\tilde{Q}_g - I)) \text{Var}(\hat{\beta}(p|p-)) \left(\sum_{c \in \text{Children}(p)} \text{Var}(\hat{\beta}(c|c-)) \right)^{-1} \left(\sum_{c \in \text{Children}(p)} \hat{\beta}(c|c-) \right).\end{aligned}$$

Since we supposed the equality (32) holds for vertex $g \in \mathcal{G}$, this implies

$$\tilde{Q}_p = \left(I + A(p) \left(\text{Var}(\tilde{\beta}(g)) \left(\sum_{p \in \text{Children}(g)} \text{Var}(\hat{\beta}(p|p-)) \right)^{-1} - I \right) \right) \text{Var}(\hat{\beta}(p|p-)) \left(\sum_{c \in \text{Children}(p)} \text{Var}(\hat{\beta}(c|c-)) \right)^{-1},$$

so since this result requires that

$$\tilde{Q}_p = \text{Var}(\tilde{\beta}(p)) \left(\sum_{c \in \text{Children}(p)} \text{Var}(\hat{\beta}(c|c-)) \right)^{-1},$$

we need to show,

$$\text{Var}(\tilde{\beta}(p)) = \left(I + A(p) \left(\text{Var}(\tilde{\beta}(g)) \left(\sum_{p \in \text{Children}(g)} \text{Var}(\hat{\beta}(p|p-)) \right)^{-1} - I \right) \right) \text{Var}(\hat{\beta}(p|p-)).$$

By (22), this equality holds if and only if

$$\begin{aligned}& \text{Var}(\hat{\beta}(p|p-)) - A(p) \text{Var}(\hat{\beta}(p|p-)) + A(p) \text{Var}(\tilde{\beta}(g)) A(p)^\top \\ &= \left(I + A(p) \left(\text{Var}(\tilde{\beta}(g)) \left(\sum_{p \in \text{Children}(g)} \text{Var}(\hat{\beta}(p|p-)) \right)^{-1} - I \right) \right) \text{Var}(\hat{\beta}(p|p-)) \\ &\iff A(p) \text{Var}(\tilde{\beta}(g)) A(p)^\top = A(p) \text{Var}(\tilde{\beta}(g)) \left(\sum_{p \in \text{Children}(g)} \text{Var}(\hat{\beta}(p|p-)) \right)^{-1} \text{Var}(\hat{\beta}(p|p-)) \\ &\iff A(p)^\top = \left(\sum_{p \in \text{Children}(g)} \text{Var}(\hat{\beta}(p|p-)) \right)^{-1} \text{Var}(\hat{\beta}(p|p-)),\end{aligned}$$

which is satisfied by the definition of $A(p)$, i.e., (19). \square

Lemma 3 provides an eigenvalue bound that will be useful in the proof of the main result of this section.

Lemma 3. *For any parent vertex $g \in \mathcal{G}$, we have $\sum_{c \in \text{Children}(g)} \text{Var}(\hat{\beta}(c|c-)) \geq \text{Var}(\tilde{\beta}(g))$.*

Proof. For arbitrary parent vertex $g \in \mathcal{G}$, let $s_g := \sum_{c \in \text{Children}(g)} \text{Var}(\hat{\beta}(c|c-))$.

First, we prove the result for the root vertex $g \in \mathcal{G}$. In this case, (16) and (13) imply

$$\text{Var}(\tilde{\beta}(g)) = \text{Var}(\hat{\beta}(g|g-)) = \left(\text{Var}(\hat{\beta}(g|g))^{-1} + s_g^{-1} \right)^{-1}$$

Thus, the result in this case follows from

$$\begin{aligned} s_g \geq \text{Var}(\tilde{\beta}(g)) &\iff s_g \geq \left(\text{Var}(\hat{\beta}(g|g))^{-1} + s_g^{-1} \right)^{-1} \\ &\iff \left(\text{Var}(\hat{\beta}(g|g))^{-1} + s_g^{-1} \right) s_g \geq I \iff \text{Var}(\hat{\beta}(g|g))^{-1} s_g \geq 0, \end{aligned}$$

which is satisfied because both $\text{Var}(\hat{\beta}(g|g))^{-1}$ and s_g are positive definite by our assumptions on $\{S(g)\}_{g \in \mathcal{G}}$ and the variance of the mean zero errors $\{\mathbf{u}(g)\}_{g \in \mathcal{G}}$.

Second, we prove the result for the arbitrary parent vertex $c \in \mathcal{G}$, which has a parent vertex that we denote by $g \in \mathcal{G}$. Recall equation (22) is

$$\begin{aligned} \text{Var}(\tilde{\beta}(c)) &= \text{Var}(\hat{\beta}(c|c-)) - A(c)\text{Var}(\hat{\beta}(c|c-)) + A(c)\text{Var}(\tilde{\beta}(g))A(c)^\top \\ &= \text{Var}(\hat{\beta}(c|c-)) - \text{Var}(\hat{\beta}(c|c-))s_g^{-1}\text{Var}(\hat{\beta}(c|c-)) + \text{Var}(\hat{\beta}(c|c-))s_g^{-1}\text{Var}(\tilde{\beta}(g))s_g^{-1}\text{Var}(\hat{\beta}(c|c-)) \\ &= \left(\text{Var}(\hat{\beta}(c|c))^{-1} + s_c^{-1} \right)^{-1} - \left(\text{Var}(\hat{\beta}(c|c))^{-1} + s_c^{-1} \right)^{-1} s_g^{-1} \left(\text{Var}(\hat{\beta}(c|c))^{-1} + s_c^{-1} \right)^{-1} \\ &\quad + \left(\text{Var}(\hat{\beta}(c|c))^{-1} + s_c^{-1} \right)^{-1} s_g^{-1} \text{Var}(\tilde{\beta}(g))s_g^{-1} \left(\text{Var}(\hat{\beta}(c|c))^{-1} + s_c^{-1} \right)^{-1}, \end{aligned}$$

where the final equality follows from (13). Thus, $s_c \geq \text{Var}(\tilde{\beta}(c))$ is satisfied if and only if

$$\begin{aligned} s_c &\geq \left(\text{Var}(\hat{\beta}(c|c))^{-1} + s_c^{-1} \right)^{-1} - \left(\text{Var}(\hat{\beta}(c|c))^{-1} + s_c^{-1} \right)^{-1} s_g^{-1} \left(\text{Var}(\hat{\beta}(c|c))^{-1} + s_c^{-1} \right)^{-1} \\ &\quad + \left(\text{Var}(\hat{\beta}(c|c))^{-1} + s_c^{-1} \right)^{-1} s_g^{-1} \text{Var}(\tilde{\beta}(g))s_g^{-1} \left(\text{Var}(\hat{\beta}(c|c))^{-1} + s_c^{-1} \right)^{-1} \\ &\iff \left(\text{Var}(\hat{\beta}(c|c))^{-1} + s_c^{-1} \right) s_c \left(\text{Var}(\hat{\beta}(c|c))^{-1} + s_c^{-1} \right) \geq \text{Var}(\hat{\beta}(c|c))^{-1} \\ &\quad + s_c^{-1} - s_g^{-1} + s_g^{-1} \text{Var}(\tilde{\beta}(g))s_g^{-1} \\ (33) \quad &\iff \text{Var}(\hat{\beta}(c|c))^{-1} s_c \text{Var}(\hat{\beta}(c|c))^{-1} + \text{Var}(\hat{\beta}(c|c))^{-1} + s_g^{-1} - s_g^{-1} \text{Var}(\tilde{\beta}(g))s_g^{-1} \geq 0. \end{aligned}$$

A sufficient condition for (33) is

$$s_g^{-1} - s_g^{-1} \text{Var}(\tilde{\beta}(g))s_g^{-1} \geq 0 \iff s_g \geq \text{Var}(\tilde{\beta}(g)),$$

so, since this latter condition is satisfied by the root vertex, inequality (33), and thus also the final result, follows from induction. \square

Theorem 1 shows that the projection used in the top-down pass, *i.e.*, the solution to the optimization problem in (17), provides the minimum-variance estimator in the family of estimators that are unbiased, linear in the input random variables, and that satisfy the parent-child consistency constraints. The next lemma provides the functional form of all such estimators that are in this family.

Lemma 4. *Suppose that the following conditions hold.*

- (a) $\check{\beta}_g := \text{stack}(\{\hat{\beta}(c)\}_{c \in \text{Children}(g)}) \in \mathbb{R}^{n\kappa}$ is any unbiased estimator of $\text{stack}(\{\beta(c)\}_{c \in \text{Children}(g)})$
- (b) $\check{\beta}_g$ is linear in $\tilde{\beta}(g)$ and $\hat{\beta}_g := \text{stack}(\{\hat{\beta}(c|c-)\}_{c \in \text{Children}(g)})$
- (c) $E^\top \check{\beta}_g = \sum_{c \in \text{Children}(g)} \check{\beta}(c) = \tilde{\beta}(g)$

Then, there exists $R \in \mathbb{R}^{n\kappa \times n}$ such that

$$\check{\beta}_g = (I - RE^\top)\hat{\beta}_g + R\tilde{\beta}(g)$$

and also $E^\top R = I$.

Proof. By the linearity property of $\check{\beta}_g$ that is assumed in the statement of the result, *i.e.*, assumption (b), there exists $Q \in \mathbb{R}^{n\kappa \times n\kappa}$ and $R \in \mathbb{R}^{n\kappa \times n}$ such that

$$(34) \quad \check{\beta}_g = (I - Q)\hat{\beta}_g + R\tilde{\beta}(g).$$

After left multiplying both sides of this equality by E^\top , assumption (c) implies

$$\begin{aligned}
(35) \quad & \tilde{\beta}(g) = (E^\top - E^\top Q)\hat{\beta}_g + E^\top R\tilde{\beta}(g). \\
& \iff E^\top = E^\top Q \\
(36) \quad & \text{and } E^\top R = I
\end{aligned}$$

After taking the expectation of both sides of (34), assumption (a) implies

$$\begin{aligned}
(37) \quad & E(\tilde{\beta}_g) = \text{stack}(\{\beta(c)\}_{c \in \text{Children}(g)}) = ((I - Q) + Re^\top) \text{stack}(\{\beta(c)\}_{c \in \text{Children}(g)}) \\
& \iff Q = Re^\top
\end{aligned}$$

□

The three lemmas above are used in the following proof of Theorem 1, which shows that the projection used in the top-down pass provides the best unbiased, linear estimate that satisfies the parent-child consistency constraints.

Theorem 1. *Suppose that the following conditions hold.*

- (a) $\tilde{\beta}_g := \text{stack}(\{\tilde{\beta}(c)\}_{c \in \text{Children}(g)}) \in \mathbb{R}^{n\kappa}$ is any unbiased estimator of $\text{stack}(\{\beta(c)\}_{c \in \text{Children}(g)})$
- (b) $\tilde{\beta}_g$ is linear in $\tilde{\beta}(g)$ and $\hat{\beta}_g := \text{stack}(\{\hat{\beta}(c|c-)\}_{c \in \text{Children}(g)})$
- (c) $\sum_{c \in \text{Children}(g)} \tilde{\beta}(c) = \tilde{\beta}(g)$

If $\tilde{\beta}(c)$ is defined as in (18) and $\tilde{\beta}_g$ is defined as $\text{stack}(\{\tilde{\beta}(c)\}_{c \in \text{Children}(g)})$, then

$$(38) \quad \text{Var}(\tilde{\beta}) \leq \text{Var}(\tilde{\beta}_g).$$

Proof. Equation (18) implies that $\tilde{\beta}_g = (I - RE^\top)\hat{\beta}_g + R\tilde{\beta}(g)$, where $R := \text{stack}(\{A(c)\}_{c \in \text{Children}(g)})$. Since $E^\top R = I$ by definition of R , Lemma 4 implies that there exists $D \in \mathbb{R}^{n\kappa \times n}$ such that $E^\top D = 0$ and

$$\tilde{\beta}_g = (I - (R + D)E^\top)\hat{\beta}_g + (R + D)\tilde{\beta}(g),$$

which implies

$$\begin{aligned}
(39) \quad & \text{Var}(\tilde{\beta}_g) = (I - (R + D)E^\top)\text{Var}(\hat{\beta}_g)(I - E(R + D)^\top) + (R + D)\text{Var}(\tilde{\beta}(g))(R + D)^\top \\
& + (I - (R + D)E^\top)\text{Cov}(\hat{\beta}_g, \tilde{\beta}(g))(R + D)^\top + (R + D)\text{Cov}(\tilde{\beta}(g), \hat{\beta}_g)(I - E(R + D)^\top).
\end{aligned}$$

After expanding, substituting $\text{Cov}(\hat{\beta}_g, \tilde{\beta}(g))$ for $\text{Cov}(\tilde{\beta}(g), \hat{\beta}_g)^\top$, and then using Lemma 2 to substitute $\text{Cov}(\tilde{\beta}(g), \hat{\beta}_g)$ for $E^\top \text{Var}(\hat{\beta}_g)E\tilde{Q}_g^\top$, (39) becomes

$$\begin{aligned}
(40) \quad & \text{Var}(\tilde{\beta}_g) = \text{Var}(\tilde{\beta}_g) - (I - (R + D)E^\top)\text{Var}(\hat{\beta}_g)ED^\top - DE^\top \text{Var}(\hat{\beta}_g)(I - ER^\top) \\
& + (R + D)\text{Var}(\tilde{\beta}(g))D^\top + D\text{Var}(\tilde{\beta}(g))R^\top \\
& + (I - (R + D)E^\top)\text{Var}(\hat{\beta}_g)E\tilde{Q}_g^\top D^\top - DE^\top \text{Var}(\hat{\beta}_g)E\tilde{Q}_g^\top R^\top \\
& - (R + D)\tilde{Q}_g E^\top \text{Var}(\hat{\beta}_g)ED^\top + D\tilde{Q}_g E^\top \text{Var}(\hat{\beta}_g)(I - ER^\top)
\end{aligned}$$

Since (21) implies $(I - RE^\top)\text{Var}(\hat{\beta}_g)E = 0$, (40) can be written as

$$\begin{aligned}
(41) \quad & \text{Var}(\tilde{\beta}_g) = \text{Var}(\tilde{\beta}_g) + DE^\top \text{Var}(\hat{\beta}_g)ED^\top + (R + D)\text{Var}(\tilde{\beta}(g))D^\top + D\text{Var}(\tilde{\beta}(g))R^\top \\
& - DE^\top \text{Var}(\hat{\beta}_g)E\tilde{Q}_g^\top (R + D)^\top - (R + D)\tilde{Q}_g E^\top \text{Var}(\hat{\beta}_g)ED^\top
\end{aligned}$$

Since Lemma 2 implies $\tilde{Q}_g = \text{Var}(\tilde{\beta}(g)) \left(E^\top \text{Var}(\hat{\beta}_g) E \right)^{-1}$, (41) can be written as

$$\begin{aligned}
 \text{Var}(\check{\beta}_g) &= \text{Var}(\tilde{\beta}_g) + DE^\top \text{Var}(\hat{\beta}_g) ED^\top + (R + D) \text{Var}(\tilde{\beta}(g)) D^\top + D \text{Var}(\tilde{\beta}(g)) R^\top \\
 &\quad - DE^\top \text{Var}(\hat{\beta}_g) E \left(E^\top \text{Var}(\hat{\beta}_g) E \right)^{-1} \text{Var}(\tilde{\beta}(g)) (R + D)^\top \\
 &\quad - (R + D) \text{Var}(\tilde{\beta}(g)) \left(E^\top \text{Var}(\hat{\beta}_g) E \right)^{-1} E^\top \text{Var}(\hat{\beta}_g) ED^\top \\
 &\iff \text{Var}(\check{\beta}_g) = \text{Var}(\tilde{\beta}_g) + DE^\top \text{Var}(\hat{\beta}_g) ED^\top + (R + D) \text{Var}(\tilde{\beta}(g)) D^\top + D \text{Var}(\tilde{\beta}(g)) R^\top \\
 &\quad - D \text{Var}(\tilde{\beta}(g)) (R + D)^\top - (R + D) \text{Var}(\tilde{\beta}(g)) D^\top \\
 (42) \quad &\iff \text{Var}(\check{\beta}_g) = \text{Var}(\tilde{\beta}_g) + D \left(E^\top \text{Var}(\hat{\beta}_g) E - \text{Var}(\tilde{\beta}(g)) \right) D^\top
 \end{aligned}$$

Since Lemma 3 implies $E^\top \text{Var}(\hat{\beta}_g) E - \text{Var}(\tilde{\beta}(g))$ is positive semi-definite, (42) implies $\text{Var}(\check{\beta}_g) \geq \text{Var}(\tilde{\beta}_g)$. \square