

# Inference in generalized bilinear models

Jeffrey W. Miller\*

Department of Biostatistics, Harvard University

and

Scott L. Carter

Division of Computational Biology, Dana–Farber Cancer Institute

October 9, 2020

## Abstract

Latent factor models are widely used to discover and adjust for hidden variation in modern applications. However, most methods do not fully account for uncertainty in the latent factors, which can lead to miscalibrated inferences such as overconfident p-values. In this article, we develop a fast and accurate method of uncertainty quantification in generalized bilinear models, which are a flexible extension of generalized linear models to include latent factors as well as row covariates, column covariates, and interactions. In particular, we introduce delta propagation, a general technique for propagating uncertainty among model components using the delta method. Further, we provide a rapidly converging algorithm for *maximum a posteriori* GBM estimation that extends earlier methods by estimating row and column dispersions. In simulation studies, we find that our method provides approximately correct frequentist coverage of most parameters of interest. We demonstrate on RNA-seq gene expression analysis and copy ratio estimation in cancer genomics.

*Keywords:* batch effects, factor analysis, Fisher information, negative-binomial regression, uncertainty quantification.

---

\*J.W.M. was supported by the Department of Defense grant W81XWH-18-1-0357, the National Institutes of Health grant R01GM083084, and the Zhu Family Center for Global Cancer Prevention. S.L.C. was supported by National Cancer Institute grant 1R01CA227156-01, the Department of Defense grant W81XWH-18-1-0357, the American Brain Tumor Association, the Wong Family Awards in Translational Oncology, and the LUNGstrong Fund. S.L.C. is also affiliated with the Department of Biostatistics at Harvard University and the Broad Institute of MIT and Harvard.

# 1 Introduction

Latent factor models have become an essential tool for analyzing and adjusting for hidden sources of variation in complex data. Building on generalized linear model theory, generalized bilinear models (GBMs) provide a flexible framework incorporating latent factors along with row covariates, column covariates, and interactions to analyze matrix data (Choulakian, 1996; Gabriel, 1998; de Falguerolles, 2000; Perry and Pillai, 2013; Hoff, 2015; Buettner et al., 2017). However, uncertainty quantification is a persistent problem in large complex models, and GBMs are particularly challenging due to the non-linearity of multiplicative terms, constraints and strong dependencies among parameters, inapplicability of normal model theory, and the fact that the number of parameters grows with the data.

Most latent factor methods do not fully account for uncertainty in the latent factors. For example, to remove batch effects in gene expression analysis, several methods first estimate a factorization  $UV^T$  and then treat  $V$  as a known matrix of covariates, accounting for uncertainty only in  $U$  (Leek and Storey, 2007, 2008; Sun et al., 2012; Risso et al., 2014). In copy number variation detection, it is common to treat the estimated  $UV^T$  as known and subtract it off (Fromer et al., 2012; Krumm et al., 2012; Jiang et al., 2015). In principle, Bayesian inference provides full uncertainty quantification (Carvalho et al., 2008), however, Markov chain Monte Carlo tends to be slow in large parameter spaces with strong dependencies, as in the case of GBMs. Variational Bayes approaches are faster (Stegle et al., 2010; Buettner et al., 2017; Babadi et al., 2018), but rely on factorized approximations that tend to underestimate uncertainty. Meanwhile, the classical method of inverting the Fisher information matrix is computationally prohibitive in large GBMs.

In this article, we introduce a novel method for uncertainty quantification in GBMs, focusing on the case of count data with negative binomial outcomes. The basic idea is to propagate uncertainty between model components using the delta method, which can be done analytically using closed-form expressions involving the gradient and the Fisher information; we refer to this as *delta propagation*. The method facilitates computation of p-

values and confidence intervals with approximately correct frequentist properties in GBMs. Further, we provide an algorithm for *maximum a posteriori* GBM estimation that extends previous work by estimating row- and column-specific dispersion parameters, improving numerical stability, and explicitly handling identifiability constraints.

In a suite of simulation studies, we find that our methods perform favorably in terms of consistency, frequentist coverage, computation time, algorithm convergence, and robustness to the outcome distribution. We then apply our methods to gene expression analysis, (a) comparing performance with DESeq2 (Love et al., 2014) on a benchmark dataset of RNA-seq samples from lymphoblastoid cell lines, and (b) testing for age-related genes using RNA-seq data from the Genotype-Tissue Expression (GTEx) project (Melé et al., 2015). Finally, we apply our methods to copy ratio estimation in cancer genomics, comparing performance with the Genome Analysis Toolkit (GATK) (Broad Institute, 2020) on whole-exome sequencing data from the Cancer Cell Line Encyclopedia (Ghandi et al., 2019).

The article is organized as follows. In Section 2, we define the GBM model and we address identifiability, interpretability, and residuals. In Sections 3 and 4, we describe our estimation and inference methods, respectively. In Section 5, we establish theoretical results, and Section 6 contains simulation studies. In Sections 7 and 8, we apply our methods to gene expression analysis and copy ratio estimation in cancer genomics. The supplementary material contains a discussion of previous work and challenges, additional empirical results, mathematical derivations and proofs, and step-by-step algorithms.

## 2 Model

In this section, we define the class of models considered in this paper and we provide conditions guaranteeing identifiability and interpretability of the model parameters. For  $i \in \{1, \dots, I\}$  and  $j \in \{1, \dots, J\}$ , suppose  $Y_{ij} \in \mathbb{R}$  is a random variable such that

$$g(\mathbb{E}(Y_{ij})) = \sum_{k=1}^K x_{ik} a_{jk} + \sum_{\ell=1}^L b_{i\ell} z_{j\ell} + \sum_{k=1}^K \sum_{\ell=1}^L x_{ik} c_{k\ell} z_{j\ell} + \sum_{m=1}^M u_{im} d_{mm} v_{jm} \quad (2.1)$$




Figure 1: Schematic of the generalized bilinear model (GBM) structure.

where  $x_{ik}$  and  $z_{jl}$  are observed covariates,  $a_{jk}$ ,  $b_{il}$ ,  $c_{kl}$ ,  $u_{im}$ ,  $d_{mm}$ , and  $v_{jm}$  are parameters to be estimated, and  $g(\cdot)$  is a smooth function such that  $g'$  is positive, referred to as the link function. In matrix form, denoting  $\mathbf{Y} = (Y_{ij}) \in \mathbb{R}^{I \times J}$ , Equation 2.1 is equivalent to

$$g(\mathbb{E}(\mathbf{Y})) = \mathbf{X}\mathbf{A}^T + \mathbf{B}\mathbf{Z}^T + \mathbf{X}\mathbf{C}\mathbf{Z}^T + \mathbf{U}\mathbf{D}\mathbf{V}^T \quad (2.2)$$

where  $g$  is applied element-wise to the matrix  $\mathbb{E}(\mathbf{Y})$ . To be able to use capital  $Y$  to denote scalar random variables, we use bold  $\mathbf{Y}$  to denote the data matrix. We refer to this as a *generalized bilinear model* (GBM), following the terminology of Choulakian (1996).

In the genomics applications in Sections 7 and 8, we use negative binomial outcomes with log link  $g(\mu) = \log(\mu)$ , and the role of each piece is as follows (see Figure 1):  $Y_{ij}$  is the read count for feature  $i$  in sample  $j$ ,  $\mathbf{X} \in \mathbb{R}^{I \times K}$  contains feature covariates and  $\mathbf{A} \in \mathbb{R}^{J \times K}$  contains the corresponding coefficients,  $\mathbf{Z} \in \mathbb{R}^{J \times L}$  contains sample covariates and  $\mathbf{B} \in \mathbb{R}^{I \times L}$  contains the corresponding coefficients,  $\mathbf{C} \in \mathbb{R}^{K \times L}$  contains intercepts and coefficients for interactions between the  $x$ 's and  $z$ 's, and  $\mathbf{U}\mathbf{D}\mathbf{V}^T$  is a low-rank matrix that captures latent effects due, for example, to unobserved covariates such as batch.

## 2.1 Identifiability and interpretation

For identifiability and interpretability, we impose certain constraints (Conditions 2.1 and 2.2). The only constraints on the covariates are that  $\mathbf{X}$  and  $\mathbf{Z}$  are full-rank, centered, and include a column of ones for intercepts; see below. The rest of the constraints are enforced on the parameters during estimation. We write  $\mathbf{I}$  to denote the identity matrix to distinguish it from the number of features,  $I$ .



**Condition 2.1** (Identifiability constraints). *Assume the following constraints:*

- (a)  $X^T X$  and  $Z^T Z$  are invertible,
- (b)  $X^T B = 0$ ,  $Z^T A = 0$ ,  $X^T U = 0$ , and  $Z^T V = 0$ ,
- (c)  $U^T U = I$  and  $V^T V = I$ ,
- (d)  $D$  is a diagonal matrix such that  $d_{11} > d_{22} > \dots > d_{MM} > 0$ , and
- (e) the first nonzero entry of each column of  $U$  is positive,

where  $A \in \mathbb{R}^{J \times K}$ ,  $B \in \mathbb{R}^{I \times L}$ ,  $C \in \mathbb{R}^{K \times L}$ ,  $D \in \mathbb{R}^{M \times M}$ ,  $U \in \mathbb{R}^{I \times M}$ ,  $V \in \mathbb{R}^{J \times M}$ ,  $X \in \mathbb{R}^{I \times K}$ ,  $Z \in \mathbb{R}^{J \times L}$ , and  $M < \min\{I, J\}$ .

In Theorem 5.1, we show that Condition 2.1 is sufficient to guarantee identifiability of  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ , and  $V$  in any GBM satisfying Equation 2.2. More precisely, letting

$$\eta(A, B, C, D, U, V) := XA^T + BZ^T + XCZ^T + UDV^T \quad (2.3)$$

for some fixed full-rank  $X$  and  $Z$ , Theorem 5.1 shows that  $\eta(\cdot)$  is a one-to-one function on the set of parameters satisfying Condition 2.1.

**Condition 2.2** (Interpretability constraints). *Assume that (a)  $x_{i1} = 1$  and  $z_{j1} = 1$  for all  $i, j$ , and (b)  $\sum_{i=1}^I x_{ik} = 0$  and  $\sum_{j=1}^J z_{j\ell} = 0$  for all  $k \geq 2$ ,  $\ell \geq 2$ .*

When Condition 2.2(a) holds, we can rearrange the right-hand side of Equation 2.1 as

$$c_{11} + a_{j1} + b_{i1} + \sum_{k=2}^K (c_{k1} + a_{jk})x_{ik} + \sum_{\ell=2}^L (c_{1\ell} + b_{i\ell})z_{j\ell} + \sum_{k=2}^K \sum_{\ell=2}^L c_{k\ell}x_{ik}z_{j\ell} + \sum_{m=1}^M u_{im}d_{mm}v_{jm}.$$

Using this and assuming Conditions 2.1 and 2.2, we show in Theorem 5.2 that the interpretation of each parameter is: (1)  $c_{11}$  is the overall intercept,  $c_{11} = \frac{1}{IJ} \sum_{i=1}^I \sum_{j=1}^J g(\mathbb{E}(Y_{ij}))$ , (2)  $a_{j1}$  is a sample-specific offset and  $b_{i1}$  is a feature-specific offset, (3)  $c_{k1}$  is the mean effect of the  $k$ th feature covariate and  $a_{jk}$  is the sample-specific offset of this effect, (4)  $c_{1\ell}$  is the

mean effect of the  $\ell$ th sample covariate and  $b_{i\ell}$  is the feature-specific offset of this effect, and (5)  $c_{k\ell}$  is the effect of the interaction  $x_{ik}z_{j\ell}$  for  $k \geq 2$  and  $\ell \geq 2$ .

GBMs can be decomposed in terms of the sum-of-squares of each component's contribution to the overall model fit, enabling one to interpret the proportion of variation explained by each component. Specifically, in Theorem 5.3, we show that

$$\text{SS}(XA^T + BZ^T + X CZ^T + UDV^T) = \text{SS}(XA^T) + \text{SS}(BZ^T) + \text{SS}(X CZ^T) + \text{SS}(UDV^T)$$

whenever Condition 2.1(b) holds, where  $\text{SS}(Q) := \sum_{i,j} q_{ij}^2$  denotes the sum of squares of the entries of a matrix. This extends a similar result by Takane and Shibayama (1991).

## 2.2 Outcome distributions

For the distribution of  $Y_{ij}$ , we focus on discrete exponential dispersion families (Jørgensen, 1987; Agresti, 2015). Specifically, suppose  $Y_{ij} \sim f(y \mid \theta_{ij}, r_{ij})$  where for  $y \in \mathbb{Z}$ ,

$$f(y \mid \theta, r) = \exp(\theta y - r\kappa(\theta))h(y, r) \tag{2.4}$$

is a probability mass function for all  $\theta \in \Theta$  and  $r \in \mathcal{R}$ ; this is referred to as a *discrete EDF*. Here,  $\Theta \subseteq \mathbb{R}$ ,  $\mathcal{R} \subseteq (0, \infty)$ , and  $\mathbb{Z}$  denotes the integers. For any discrete EDF, the mean and variance are  $\mathbb{E}(Y) = r\kappa'(\theta)$  and  $\text{Var}(Y) = r\kappa''(\theta)$  (Jørgensen, 1987); also see Section S4. We refer to  $r$  as the inverse dispersion parameter, and  $1/r$  is the dispersion. Discrete EDFs can be translated into standard EDFs of the form  $\exp(r[\theta y - \kappa(\theta)])h(y, r)$  via the transformation  $y \mapsto ry$ . We refer to a GBM with EDF outcomes as an EDF-GBM.

**Negative binomial (NB) outcomes.** In the applications in Sections 7 and 8, we use negative binomial outcomes:  $Y_{ij} \sim \text{NegBin}(\mu_{ij}, r_{ij})$  where  $\mu_{ij}$  is the mean and  $1/r_{ij}$  is the dispersion. This is a discrete EDF as in Equation 2.4 with  $\theta = \log(\mu/(\mu + r))$  and  $\kappa(\theta) = -\log(1 - \exp(\theta))$ ; thus,  $\mathbb{E}(Y) = \mu$  and  $\text{Var}(Y) = \mu + \mu^2/r$ . The NB distribution is an overdispersed Poisson since if  $Y \mid \lambda \sim \text{Poisson}(\lambda)$  and  $\lambda \sim \text{Gamma}(r, r/\mu)$ , then integrating

out  $\lambda$ , we have  $Y \sim \text{NegBin}(\mu, r)$ . We refer to a GBM with NB outcomes as an NB-GBM.

We parametrize the dispersions as  $1/r_{ij} = \exp(s_i + t_j + \omega)$  and work in terms of  $S = (s_1, \dots, s_I)^\top \in \mathbb{R}^I$ ,  $T = (t_1, \dots, t_J)^\top \in \mathbb{R}^J$ , and  $\omega \in \mathbb{R}$ , subject to the identifiability constraints  $\frac{1}{I} \sum_i e^{s_i} = 1$  and  $\frac{1}{J} \sum_j e^{t_j} = 1$ . Note that this makes  $\frac{1}{IJ} \sum_{i,j} 1/r_{ij} = \exp(\omega)$ .

## 2.3 Residuals and adjusting out selected effects

Residuals are useful for many purposes, such as visualization, model criticism, and downstream analyses. We define GBM residuals as  $\varepsilon_{ij} := g(Y_{ij} + \epsilon) - \eta_{ij}$  where  $\eta = \eta(A, B, C, D, U, V) \in \mathbb{R}^{I \times J}$  as in Equation 2.3 and  $\epsilon$  is a small constant to make  $\varepsilon_{ij}$  well-defined; for NB-GBMs, we use  $\epsilon = 1/8$  as a default. A model-based estimate of the variance of  $\varepsilon_{ij}$  is given by  $\sigma_{ij}^2 g'(\mu_{ij})^2$  where  $\mu_{ij}$  and  $\sigma_{ij}^2$  are the mean and variance of  $Y_{ij}$  under the model. This formula can be derived either from the Fisher information or from a first-order Taylor approximation to  $g$ . It turns out that the corresponding precisions  $w_{ij} := 1/(\sigma_{ij}^2 g'(\mu_{ij})^2)$  play a key role in our GBM estimation and inference algorithms. In the NB case with  $g(\mu) = \log(\mu)$ , these residual precisions are  $w_{ij} = r_{ij} \mu_{ij} / (r_{ij} + \mu_{ij})$ .

Often, it is useful to adjust out some effects but not others. Let  $\mathcal{R}_x$ ,  $\mathcal{R}_z$ , and  $\mathcal{R}_u$  be the indices of the columns of  $X$ ,  $Z$ , and  $U$  (or  $V$ ) that one does not wish to adjust out. We define the partial residuals  $\varepsilon_{ij}^{\mathcal{R}} := \eta_{ij}^{\mathcal{R}} + \varepsilon_{ij}$  where  $\eta^{\mathcal{R}} \in \mathbb{R}^{I \times J}$  is defined as in Equation 2.3 but with  $x_{ik}$ ,  $z_{j\ell}$ , and  $u_{im}$  replaced by  $x_{ik} \mathbf{1}(k \in \mathcal{R}_x)$ ,  $z_{j\ell} \mathbf{1}(\ell \in \mathcal{R}_z)$ , and  $u_{im} \mathbf{1}(m \in \mathcal{R}_u)$ .

## 3 Estimation

We provide a general algorithm for estimating the parameters of a discrete EDF-GBM (Equations 2.2 and 2.4), and we augment the algorithm to estimate the NB-GBM dispersion parameters as well. Here, we only give an outline; see Section S7 for step-by-step details.

**Inputs.** The required inputs are  $\mathbf{Y} \in \mathbb{Z}^{I \times J}$ ,  $X \in \mathbb{R}^{I \times K}$ ,  $Z \in \mathbb{R}^{J \times L}$ , and  $M \in \{0, 1, 2, \dots\}$ . Optional inputs are the maximum step size  $\rho > 0$ , prior precisions  $\lambda_a, \lambda_b, \lambda_c, \lambda_d, \lambda_u, \lambda_v > 0$ , prior means and precisions  $m_s, m_t, \lambda_s, \lambda_t$  for the log-dispersions in the NB-GBM case, and

the convergence criterion. As defaults, we use  $\rho = 5$ ,  $\lambda_a = \lambda_b = \lambda_c = \lambda_d = \lambda_u = \lambda_v = 1$ ,  $m_s = m_t = 0$ ,  $\lambda_s = \lambda_t = 1$ , convergence tolerance  $\tau = 10^{-6}$  for the relative change in log-likelihood+log-prior, and a maximum of 50 iterations.

**Preprocessing.**

- (1) Ensure that  $X$  and  $Z$  satisfy Condition 2.1(a) and Condition 2.2.
- (2) Unless the covariates are already on a common scale in terms of units, standardize  $X$  and  $Z$  such that  $\frac{1}{I} \sum_{i=1}^I x_{ik}^2 = 1$  and  $\frac{1}{J} \sum_{j=1}^J z_{j\ell}^2 = 1$  for all  $k \geq 2$  and  $\ell \geq 2$ .
- (3) Precompute the pseudoinverses  $X^+ = (X^T X)^{-1} X^T$  and  $Z^+ = (Z^T Z)^{-1} Z^T$ .

**Initialization.**

- (1) Solve for  $A$ ,  $B$ , and  $C$  to minimize the sum-of-squares of the GBM residuals  $\varepsilon_{ij}$ .
- (2) Randomly initialize  $D$ ,  $U$ , and  $V$  by computing the truncated singular value decomposition (of rank  $M$ ) of a random matrix with i.i.d.  $\mathcal{N}(0, 10^{-16})$  entries.
- (3) In the case of NB-GBMs, iteratively update  $S$ ,  $T$ , and  $\omega$  for a few iterations.

**Iteration.** In each iteration, we cycle through the components of the model, updating each in turn using an optimization-projection step, consisting of an unconstrained optimization step and a likelihood-preserving projection onto the constrained parameter space. We use a bounded, regularized version of Fisher scoring to perform the unconstrained optimization step for each of  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $UD$ , and  $VD$ , separately, holding all the other parameters fixed. For a generic parameter vector  $\beta$ , the (unbounded) regularized Fisher scoring step is  $\beta \leftarrow \beta + (\mathbb{E}(-\nabla_{\beta}^2 \mathcal{L}) + \lambda \mathbf{I})^{-1} (\nabla_{\beta} \mathcal{L} - \lambda \beta)$  where  $\mathcal{L}$  is the log-likelihood and  $\lambda > 0$  is a regularization parameter. This arises from optimizing the log-likelihood plus the log-prior, where the prior on  $\beta$  is  $\pi(\beta) = \mathcal{N}(\beta \mid 0, \lambda^{-1} \mathbf{I})$ , since then the gradient and Fisher information are  $\nabla_{\beta}(\mathcal{L} + \log \pi) = \nabla_{\beta} \mathcal{L} - \lambda \beta$  and  $\mathbb{E}(-\nabla_{\beta}^2(\mathcal{L} + \log \pi)) = \mathbb{E}(-\nabla_{\beta}^2 \mathcal{L}) + \lambda \mathbf{I}$ . Since these Fisher scoring steps occasionally diverge, for numerical stability we bound them using

$$\begin{aligned} \xi &\leftarrow (\mathbb{E}(-\nabla_{\beta}^2 \mathcal{L}) + \lambda \mathbf{I})^{-1} (\nabla_{\beta} \mathcal{L} - \lambda \beta) \\ \beta &\leftarrow \beta + \xi \min\{1, \rho \sqrt{\dim(\xi)} / \|\xi\|\} \end{aligned} \tag{3.1}$$

where  $\|\xi\| = (\sum_i |\xi_i|^2)^{1/2}$  is the Euclidean norm. The idea is that  $\xi \min\{1, \rho\sqrt{\dim(\xi)}/\|\xi\|\}$  points in the same direction as  $\xi$ , but its root-mean-square is capped at  $\rho$ . Similarly, for  $S$  and  $T$  in the NB-GBM, we use bounded regularized Newton steps for the unconstrained optimizations, since the (expected) Fisher information for  $S$  and  $T$  is not closed form.

See Section S7 for full step-by-step details. See Section 5 for time complexity analysis.

## 4 Inference

In this section, we introduce our methodology for computing approximate standard errors for the parameters of a GBM. Since the standard technique of inverting the Fisher information matrix does not work well on GBMs, we develop a novel technique for propagating uncertainty from one part of the model to another; see Section 4.1. We provide an outline of our inference algorithm in Section 4.2, and step-by-step details are in Section S8.

### 4.1 Delta propagation method

In fixed-dimension parametric models, the asymptotic covariance of the maximum likelihood estimator is equal to the inverse of the Fisher information matrix. Thus, classically, approximate standard errors are given by the square roots of the diagonal entries of the inverse Fisher information. However, in GBMs, inverting the full (constraint-augmented) Fisher information does not work well for two reasons: (1) it is computationally intractable for large data matrices, and (2) it does not yield well-calibrated standard errors in terms of coverage, presumably because the number of parameters grows with the amount of data. Meanwhile, the inverse Fisher information for each component individually (for instance,  $F_a^{-1}$  for  $A$ ) is computationally efficient, but severely underestimates uncertainty since it treats the other components as known, and thus, can be thought of as representing the conditional uncertainty in each component given the other components.

We propose a general technique for approximating, for each model component, the additional variance due to uncertainty in the other components. By adding this to the

conditional variance (that is,  $\text{diag}(F_a^{-1})$  in the case of  $A$ ), we obtain approximate variances that are better calibrated, empirically. The basic idea is to write the estimator for each component as a function of the other components, and propagate the variance of the other components through this function using the same idea as the delta method.

In general, suppose we have a model with parameters  $\theta \in \mathbb{R}^d$  and  $\nu \in \mathbb{R}^k$ , and we wish to quantify the uncertainty in  $\theta$  due to uncertainty in  $\nu$ . Suppose the true values are  $\theta_0$  and  $\nu_0$ , and let  $\hat{\theta}$  and  $\hat{\nu}$  be the maximum likelihood estimators. Define

$$h(\nu) := \theta_0 + F(\theta_0, \nu)^{-1}g(\theta_0, \nu)$$

where  $g(\theta, \nu) := \nabla_{\theta}\mathcal{L}$  is the gradient of the log-likelihood and  $F(\theta, \nu) := \mathbb{E}(-\nabla_{\theta}^2\mathcal{L})$  is the Fisher information matrix for  $\theta$ , evaluated at  $(\theta, \nu)$ . The interpretation is that  $\hat{\theta} \approx h(\hat{\nu})$ , since  $h(\hat{\nu})$  is a Fisher scoring step on  $\theta$  starting at  $\theta_0$ , when the current estimate of  $\nu$  is  $\hat{\nu}$ .

By a Taylor approximation to  $h$  at  $\nu_0$ , we have  $h(\hat{\nu}) \approx h(\nu_0) + h'(\nu_0)^{\text{T}}(\hat{\nu} - \nu_0)$  where  $h'(\nu) \in \mathbb{R}^{k \times d}$  such that  $h'(\nu)_{ij} = \partial h_j / \partial \nu_i$ . Define the random element  $\varphi = (h(\nu_0), h'(\nu_0))$ , where the randomness comes from the data. Assuming  $\hat{\nu}|\varphi \approx \mathcal{N}(\nu_0, \Sigma_{\nu})$  for some  $\Sigma_{\nu}$ , we have  $h(\hat{\nu})|\varphi \approx \mathcal{N}(h(\nu_0), h'(\nu_0)^{\text{T}}\Sigma_{\nu}h'(\nu_0))$ . Meanwhile, under standard regularity conditions,  $\mathbb{E}(h(\nu_0)) = \theta_0$  and  $\text{Cov}(h(\nu_0)) = F(\theta_0, \nu_0)^{-1}$ . Then, by the law of total covariance,

$$\text{Cov}(h(\hat{\nu})) = \text{Cov}(\mathbb{E}(h(\hat{\nu})|\varphi)) + \mathbb{E}(\text{Cov}(h(\hat{\nu})|\varphi)) \approx F(\theta_0, \nu_0)^{-1} + \mathbb{E}(h'(\nu_0)^{\text{T}}\Sigma_{\nu}h'(\nu_0)).$$

The interpretation of this decomposition is that  $F(\theta_0, \nu_0)^{-1}$  represents the uncertainty in  $\theta$  given  $\nu$ , and  $\mathbb{E}(h'(\nu_0)^{\text{T}}\Sigma_{\nu}h'(\nu_0))$  represents the uncertainty in  $\theta$  due to uncertainty in  $\nu$ . Since  $\hat{\theta} \approx h(\hat{\nu})$ , plugging in empirical estimates leads to the approximation

$$\text{Cov}(\hat{\theta}) \approx F(\hat{\theta}, \hat{\nu})^{-1} + \hat{h}'(\hat{\nu})^{\text{T}}\hat{\Sigma}_{\nu}\hat{h}'(\hat{\nu}), \tag{4.1}$$

where  $\hat{h}(\nu) = \hat{\theta} + F(\hat{\theta}, \nu)^{-1}g(\hat{\theta}, \nu)$ .




Figure 2: Diagram of uncertainty propagation scheme for GBM inference.

To compute the Jacobian matrix  $h'(\nu)^T$ , observe that the  $i$ th column of  $h'(\nu)^T$  is

$$\frac{\partial h}{\partial \nu_i} = \frac{\partial F^{-1}}{\partial \nu_i} g + F^{-1} \frac{\partial g}{\partial \nu_i} = -F^{-1} \frac{\partial F}{\partial \nu_i} F^{-1} g + F^{-1} \frac{\partial g}{\partial \nu_i} \quad (4.2)$$

by Bishop (2006, Eqn C.21), where  $F = F(\theta_0, \nu)$ ,  $g = g(\theta_0, \nu)$ , and  $\partial/\partial \nu_i$  is applied element-wise. Equation 4.2 also holds for  $\partial \hat{h}/\partial \nu_i$ , but with  $\hat{\theta}$  in place of  $\theta_0$ ; this facilitates computing  $\hat{h}'(\hat{\nu})$  in Equation 4.1. To obtain standard errors for each element of  $\theta$ , computation is simplified by the fact that we only need the diagonal of  $\text{Cov}(\hat{\theta})$ , and to simplify computation even further we use a diagonal matrix for  $\hat{\Sigma}_\nu$  when we apply this technique to GBMs. Additionally, since we use *maximum a posteriori* estimates, we use the regularized Fisher information and the gradient of the log-posterior in the formulas above.

## 4.2 Outline of inference algorithm

Here we outline our procedure for computing GBM standard errors; see Section S8 for step-by-step details. The strategy is to handle  $U$  and  $V$  jointly by inverting the regularized, constraint-augmented Fisher information matrix for  $(U, V)$  and then employ delta propagation for  $A$ ,  $B$ ,  $C$ ,  $S$ , and  $T$ ; see Figure 2. Empirically, we find that some of the delta propagation terms are negligible, thus, these have been excluded.

For notational convenience, we vectorize the parameter matrices as follows. For  $Q \in \mathbb{R}^{m \times n}$ , define  $\text{vec}(Q) := (q_{11}, q_{21}, \dots, q_{m1}, q_{12}, q_{22}, \dots, q_{m2}, \dots, q_{mn}) \in \mathbb{R}^{mn}$ . Denote  $\vec{a} := \text{vec}(A^T)$ ,  $\vec{b} := \text{vec}(B^T)$ ,  $\vec{c} := \text{vec}(C)$ ,  $\vec{d} := \text{diag}(D)$ ,  $\vec{u} := \text{vec}(U^T)$ , and  $\vec{v} := \text{vec}(V^T)$ . It is

easier to work with the vectorized transposes of  $A$ ,  $B$ ,  $U$ , and  $V$  (that is,  $\text{vec}(A^T)$  rather than  $\text{vec}(A)$ ), since then the Fisher information matrices have a block diagonal structure. We write  $F_a$ ,  $F_b$ ,  $F_c$ ,  $F_d$ ,  $F_u$ , and  $F_v$  to denote the regularized Fisher information for  $\vec{a}$ ,  $\vec{b}$ ,  $\vec{c}$ ,  $\vec{d}$ ,  $\vec{u}$ , and  $\vec{v}$ , respectively, for instance,  $F_a = \mathbb{E}(-\nabla_{\vec{a}}^2 (\mathcal{L} + \log \pi))$ . We write  $F_{s,\text{obs}}$  and  $F_{t,\text{obs}}$  for the regularized, observed Fisher information for  $S$  and  $T$ , that is,  $F_{s,\text{obs}} = -\nabla_S^2 (\mathcal{L} + \log \pi)$ .

**Inputs.** The required inputs are  $\mathbf{Y}$ ,  $X$ ,  $Z$ , and the estimates of  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ ,  $V$ ,  $S$ ,  $T$ , and  $\omega$ . Optional inputs are the prior parameters  $(\lambda_a, \lambda_b, \lambda_c, \lambda_d, \lambda_u, \lambda_v, \lambda_s, \lambda_t, m_s, m_t)$ .

**Compute conditional uncertainty for each component.**

- (1) Compute  $F_c^{-1}$  and the diagonal blocks of  $F_a^{-1}$ ,  $F_b^{-1}$ ,  $F_u^{-1}$ , and  $F_v^{-1}$ .
- (2) Compute the diagonals of  $F_{s,\text{obs}}^{-1}$  and  $F_{t,\text{obs}}^{-1}$ .

**Compute joint uncertainty in  $(U, V)$  accounting for constraints.**

- (1) Compute  $\tilde{F}_{(u,v)}$ , the regularized constraint-augmented Fisher information for  $\begin{bmatrix} \vec{u} \\ \vec{v} \end{bmatrix}$ .
- (2) Compute  $\text{diag}(\tilde{F}_{(u,v)}^{-1})$ . (It is key to do this in a computationally efficient way.)
- (3) Define  $\widehat{\text{var}}_u$  and  $\widehat{\text{var}}_v$  to be the entries of  $\text{diag}(\tilde{F}_{(u,v)}^{-1})$  corresponding to  $U$  and  $V$ .

**Propagate uncertainty between components using delta propagation.**

- (1) Propagate uncertainty in  $(U, V)$  to  $A$  and  $B$ , to obtain  $\widehat{\text{var}}_{(u,v) \rightarrow a}$  and  $\widehat{\text{var}}_{(u,v) \rightarrow b}$ , the additional variance of the estimators of  $A$  and  $B$  due to uncertainty in  $(U, V)$ .
- (2) Propagate uncertainty in  $A$  and  $B$  through to  $C$ , to obtain  $\widehat{\text{var}}_{(a,b) \rightarrow c}$ .
- (3) Propagate uncertainty in  $A$ ,  $B$ ,  $U$ ,  $V$  to  $S$  and  $T$ , to get  $\widehat{\text{var}}_{(a,b,u,v) \rightarrow s}$  and  $\widehat{\text{var}}_{(a,b,u,v) \rightarrow t}$ .

**Compute approximate standard errors.**

- (1)  $\hat{s}e_a \leftarrow \text{sqrt}(\text{diag}(F_a^{-1}) + \widehat{\text{var}}_{(u,v) \rightarrow a})$  and  $\hat{s}e_b \leftarrow \text{sqrt}(\text{diag}(F_b^{-1}) + \widehat{\text{var}}_{(u,v) \rightarrow b})$
- (2)  $\hat{s}e_c \leftarrow \text{sqrt}(\text{diag}(F_c^{-1}) + \widehat{\text{var}}_{(a,b) \rightarrow c})$
- (3)  $\hat{s}e_u \leftarrow \text{sqrt}(\widehat{\text{var}}_u)$  and  $\hat{s}e_v \leftarrow \text{sqrt}(\widehat{\text{var}}_v)$
- (4)  $\hat{s}e_s \leftarrow \text{sqrt}(\text{diag}(F_{s,\text{obs}}^{-1}) + \widehat{\text{var}}_{(a,b,u,v) \rightarrow s})$  and  $\hat{s}e_t \leftarrow \text{sqrt}(\text{diag}(F_{t,\text{obs}}^{-1}) + \widehat{\text{var}}_{(a,b,u,v) \rightarrow t})$



Here,  $\text{sqrt}(\cdot)$  is the element-wise square root. We do not provide standard errors for  $D$  and  $\omega$ , since it seems difficult to estimate them without non-negligible bias. See Section S8 for the complete step-by-step algorithm. See Section 5 for computation time complexity.

## 5 Theory

In this section, we provide theoretical results on GBMs. The proofs are in Section S10.

**Theorem 5.1** (Identifiability). *If  $(A_1, B_1, C_1, D_1, U_1, V_1, X, Z)$  and  $(A_2, B_2, C_2, D_2, U_2, V_2, X, Z)$  satisfy Condition 2.1 and*

$$XA_1^T + B_1Z^T + XC_1Z^T + U_1D_1V_1^T = XA_2^T + B_2Z^T + XC_2Z^T + U_2D_2V_2^T, \quad (5.1)$$

*then  $A_1 = A_2$ ,  $B_1 = B_2$ ,  $C_1 = C_2$ ,  $D_1 = D_2$ ,  $U_1 = U_2$ , and  $V_1 = V_2$ . In particular, for any GBM satisfying Equation 2.2 for some  $X$ ,  $Z$ , and  $M$ , if Condition 2.1 holds, then  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ , and  $V$  are identifiable in the sense that they are uniquely determined by the distribution of  $\mathbf{Y}$ ; in fact, they are uniquely determined by  $\mathbb{E}(\mathbf{Y})$ .*

**Theorem 5.2** (Interpretation of parameters). *If Conditions 2.1 and 2.2 hold and  $\mu_{ij} := \mathbb{E}(Y_{ij})$  satisfies Equation 2.1, then:*

- (a)  $\sum_{j=1}^J a_{jk} = 0$  and  $\sum_{i=1}^I b_{i\ell} = 0$  for all  $k \in \{1, \dots, K\}, \ell \in \{1, \dots, L\}$ ,
- (b)  $\sum_{i=1}^I u_{im} = 0$  and  $\sum_{j=1}^J v_{jm} = 0$  for all  $m \in \{1, \dots, M\}$ ,
- (c)  $\frac{1}{I} \sum_{i=1}^I g(\mu_{ij}) = c_{11} + a_{j1} + \sum_{\ell=2}^L c_{1\ell} z_{j\ell}$  for all  $j \in \{1, \dots, J\}$ ,
- (d)  $\frac{1}{J} \sum_{j=1}^J g(\mu_{ij}) = c_{11} + b_{i1} + \sum_{k=2}^K c_{k1} x_{ik}$  for all  $i \in \{1, \dots, I\}$ , and
- (e)  $\frac{1}{IJ} \sum_{i=1}^I \sum_{j=1}^J g(\mu_{ij}) = c_{11}$ .

For a matrix  $Q \in \mathbb{R}^{m \times n}$ , we write  $\text{SS}(Q) := \sum_{i=1}^m \sum_{j=1}^n q_{ij}^2$  for the sum of squares.

**Theorem 5.3** (Sum-of-squares decomposition). *If Condition 2.1(b) holds, then*

$$\text{SS}(XA^T + BZ^T + XCZ^T + UDV^T) = \text{SS}(XA^T) + \text{SS}(BZ^T) + \text{SS}(XCZ^T) + \text{SS}(UDV^T).$$

Theorem 5.4 shows that the projections we use in the estimation algorithm are likelihood-preserving. The idea is that, for example,  $\tilde{A}$  is the result of an unconstrained optimization step on  $A$ , and we project  $(\tilde{A}, C)$  to  $(A_1, C_1)$  to enforce the constraints in Condition 2.1 without affecting the likelihood. To interpret items 3 and 4 of Theorem 5.4, note that in the algorithm, we optimize with respect to  $G := UD$  and  $H := VD$ , rather than  $U$  and  $V$ . We write  $Q^+$  to denote the pseudoinverse. When  $(Q^T Q)^{-1}$  exists,  $Q^+ = (Q^T Q)^{-1} Q^T$ .

**Theorem 5.4** (Likelihood-preserving projections). *Suppose  $(A, B, C, D, U, V, X, Z)$  satisfies Condition 2.1. Fix  $X$  and  $Z$  and define  $\eta(\cdot)$  as in Equation 2.3.*

1. *Let  $\tilde{A} \in \mathbb{R}^{J \times K}$ . Define  $A_1 = \tilde{A} - Z(Z^+ \tilde{A})$  and  $C_1 = C + (Z^+ \tilde{A})^T$ . Then  $\eta(A_1, B, C_1, D, U, V) = \eta(\tilde{A}, B, C, D, U, V)$  and  $(A_1, B, C_1, D, U, V)$  satisfies Condition 2.1.*
2. *Let  $\tilde{B} \in \mathbb{R}^{I \times L}$ . Define  $B_1 = \tilde{B} - X(X^+ \tilde{B})$  and  $C_1 = C + (X^+ \tilde{B})$ . Then  $\eta(A, B_1, C_1, D, U, V) = \eta(A, \tilde{B}, C, D, U, V)$  and  $(A, B_1, C_1, D, U, V)$  satisfies Condition 2.1.*
3. *Let  $\tilde{G} \in \mathbb{R}^{I \times M}$ . Define  $G_0 = \tilde{G} - X(X^+ \tilde{G})$  and let  $U_1 D_1 V_1^T$  be the compact SVD (of rank  $M$ ) of  $G_0 V^T$ . Assume the singular values are distinct and positive, and choose the SVD in such a way that Conditions 2.1(d) and 2.1(e) are satisfied. Define  $A_0 = A + V(X^+ \tilde{G})^T$ ,  $A_1 = A_0 - Z(Z^+ A_0)$ , and  $C_1 = C + (Z^+ A_0)^T$ . Then  $\eta(A_1, B, C_1, D_1, U_1, V_1) = \eta(A, B, C, I, \tilde{G}, V)$  and  $(A_1, B, C_1, D_1, U_1, V_1)$  satisfies Condition 2.1.*
4. *Let  $\tilde{H} \in \mathbb{R}^{J \times M}$ . Define  $H_0 = \tilde{H} - Z(Z^+ \tilde{H})$  and let  $U_1 D_1 V_1^T$  be the compact SVD (of rank  $M$ ) of  $U H_0^T$ . Assume the singular values are distinct and positive, and choose the SVD in such a way that Conditions 2.1(d) and 2.1(e) are satisfied. Define  $B_0 = B + U(Z^+ \tilde{H})^T$ ,  $B_1 = B_0 - X(X^+ B_0)$ , and  $C_1 = C + X^+ B_0$ . Then  $\eta(A, B_1, C_1, D_1, U_1, V_1) = \eta(A, B, C, I, U, \tilde{H})$  and  $(A, B_1, C_1, D_1, U_1, V_1)$  satisfies Condition 2.1.*

Table 1: Computation time complexity of each update in the estimation algorithm.

Operation	Time complexity
Computing $\eta$	$O(IJ \max\{K, L, M\})$
Updating $A$	$O(IJK^2)$
Updating $B$	$O(IJL^2)$
Updating $C$	$O(IJ \max\{K^2, L^2\})$
Updating $D, U,$ and $V$	$O(IJM^2)$
Updating $S$ and $T$	$O(IJ)$
Total per iteration	$O(IJ \max\{K^2, L^2, M^2\})$

Table 2: Computation time complexity of the inference algorithm.

Operation	Time complexity
Preprocessing	$O(IJ \max\{K, L, M\})$
Conditional uncertainty for each component	$O(IJ \max\{K^2, L^2, M^2\})$
Joint uncertainty in $(U, V)$ accounting for constraints	$O(IJ^2M^3)$
Propagate uncertainty between components	$O(IJ \max\{K^3, L^3, M^3\})$
Compute approximate standard errors	$O(IJ)$
Total	$O(IJ \max\{K^3, L^3, JM^3\})$

Next, we provide the computation time complexity of our estimation and inference algorithms in Sections 3 and 4. To simplify the expressions, here we assume

$$\max\{K^2, L^2, M\} \leq \min\{I, J\}. \quad (5.2)$$

For the estimation algorithm, preprocessing and initialization take  $O(IJ \max\{K, L, M\})$  time, and Table 1 summarizes the computation time for updating each model component. The table first breaks out the time required to compute  $\eta$  (Equation 2.3), which is a prerequisite within each other update, and then lists the time required for each update given  $\eta$ . In total, it takes  $O(IJ \max\{K^2, L^2, M^2\})$  time to perform each overall iteration.

For the inference algorithm, Table 2 shows the computation time assuming Equation 5.2 and also assuming  $I \geq J$ . These are one-time costs since there are not repeated iterations. When  $M > 0$ , the most expensive operation tends to be computing the joint uncertainty in  $(U, V)$ , and as  $J$  grows this dominates the cost. We have experimented extensively but have not found a faster alternative that provides well-calibrated standard errors.

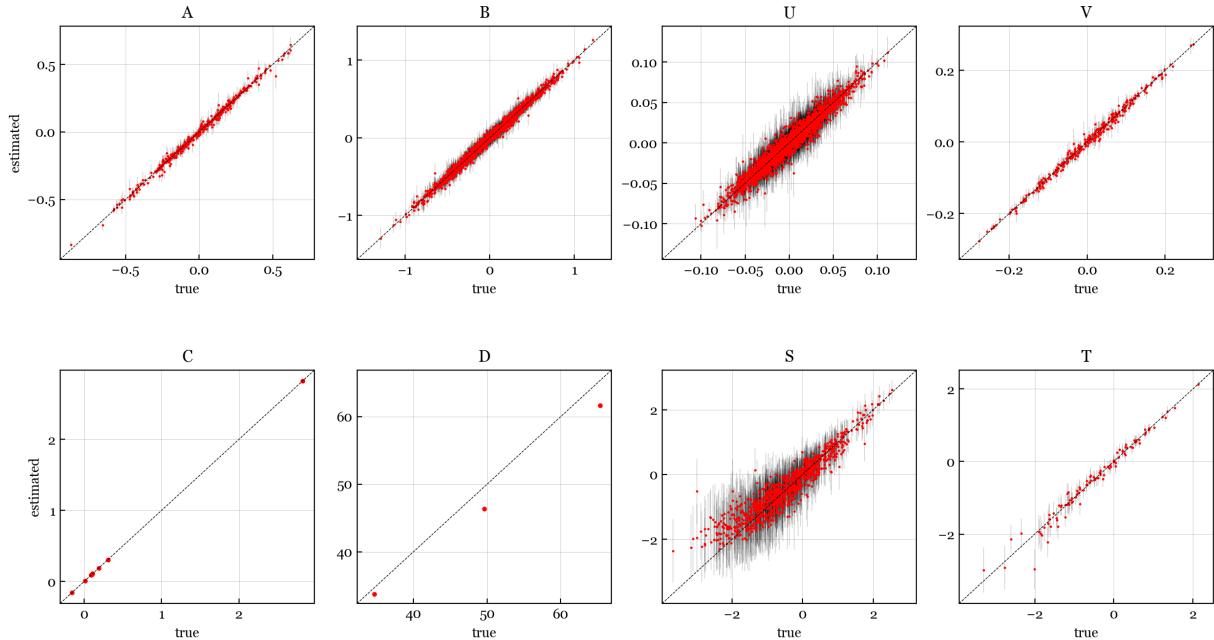


Figure 3: Scatterplots of estimated versus true parameters for a typical simulated data matrix.

## 6 Simulations

In this section, we present simulation studies assessing (a) consistency and statistical efficiency, (b) accuracy of standard errors, (c) computation time and algorithm convergence, and (d) robustness to the outcome distribution. See Section S2 for more simulation results.

In each simulation run, the data are generated as follows; see Section S2 for full details. We generate the covariates using one of three schemes, `Normal`, `Gamma`, or `Binary`, then we generate the true parameters using either a `Normal` or `Gamma` scheme, and finally we generate the outcome data using the log link and a `NB` (negative binomial), `LNP` (log-normal Poisson), `Poisson`, or `Geometric` distribution. For brevity, we refer to each combination of choices by the triplet of outcomes/covariates/parameters, for instance, `NB/Binary/Normal`.

**Typical example.** Figure 3 shows scatterplots of the estimated versus true parameters for an `NB/Normal/Normal` simulation with  $I = 1000$ ,  $J = 100$ ,  $K = 4$ ,  $L = 2$ , and  $M = 3$ . Each dot represents a single univariate parameter, for example, the plot for  $A$  contains  $JK$  dots, one for each entry  $a_{jk}$ . The error bars are  $\pm 2\hat{s}e$ . Visually, the estimates are close to true values, and the standard errors look appropriate. Since the likelihood and prior are

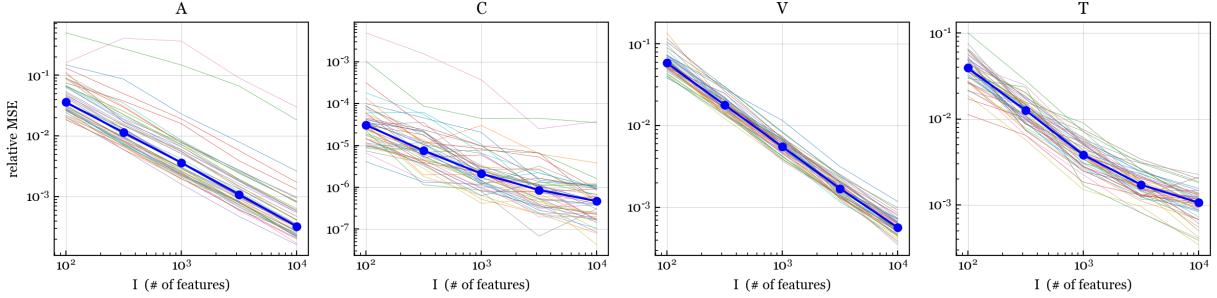


Figure 4: Relative mean-squared error between estimated and true parameter values. Each plot contains 50 thin lines, one for each run, along with the median over the 50 runs (thick blue line).

invariant to permutations and sign changes of the latent factors, in this section we permute and flip signs to find the correct assignment to the true latent factors. Note that the  $s_i$  estimates are biased upward when the true value of  $s_i$  is very low; this is because very low values of  $s_i$  make row  $i$  roughly Poisson, and in this case any value of  $s_i$  from  $-\infty$  to  $\approx -2$  could yield a reasonable fit. The prior on  $s_i$  prevents the estimate from diverging, but also leads to an upward bias when the true value is very low.

## 6.1 Consistency and statistical efficiency

In many applications,  $I$  is much larger than  $J$ . One would hope that for any  $J$ , the estimates of  $A$ ,  $C$ ,  $V$ , and  $T$  would be consistent as  $I \rightarrow \infty$  since then these parameters have fixed dimension. Meanwhile, one cannot hope for consistency in  $B$ ,  $U$ , and  $S$  as  $I \rightarrow \infty$ .

To assess consistency and efficiency, for each  $I \in \{100, 316, 1000, 3162, 10000\}$  we use the NB/Normal/Normal simulation scheme to generate 50 data matrices with  $J = 100$ ,  $K = 4$ ,  $L = 2$ , and  $M = 3$ , each with a different set of covariates and parameters. For each data matrix, we run our NB-GBM estimation algorithm with convergence tolerance  $\tau = 10^{-8}$ . Figure 4 shows the relative mean-squared error (MSE) between the estimates and the true values for  $A$ ,  $C$ ,  $V$ , and  $T$ . For  $T$ , we measure the relative MSE in the dispersion parametrization (rather than log-dispersion) since there is little difference between, say,  $t_j = -5$  and  $t_j = -100$ ; both make column  $j$  approximately Poisson distributed.

We see that for  $A$ ,  $C$ ,  $V$ , and  $T$ , the relative MSE is decreasing to zero, suggesting

that the estimates of these parameters are consistent as  $I \rightarrow \infty$ . Further, for  $A$  and  $V$ , the relative MSE appears to be  $O(1/I)$ , which is the optimal rate of convergence even for fixed-dimension parametric models. For  $B$ ,  $U$ , and  $S$  (Figure S1), the relative MSE hovers around a small nonzero value, but does not appear to be trending to zero, as expected. For  $D$  and  $\omega$  (Figure S1), the relative MSE is small and the trend is suggestive but less clear.

## 6.2 Accuracy of standard errors

Next, we assess the accuracy of the standard errors produced by our inference algorithm, in terms of the coverage. Ideally, a 95% confidence interval would contain the true parameter 95% of the time, but even when the model is correct, this is not guaranteed since intervals are usually based on an approximation to the distribution of an estimator. To assess coverage, for each  $I \in \{100, 1000, 10000\}$ , we use the `NB/Normal/Normal` scheme to generate 50 data matrices with  $J = 100$ ,  $K = 4$ ,  $L = 2$ , and  $M = 3$ , each with a different set of covariates and parameters. For each data matrix, we run our NB-GBM estimation algorithm (with tolerance  $\tau = 10^{-8}$ ) and then we run our NB-GBM inference algorithm to obtain approximate standard errors. We construct Wald-type confidence intervals for each univariate parameter, for example, the 95% confidence interval for  $a_{jk}$  is  $\hat{a}_{jk} \pm 1.96 \hat{s}e$  where  $\hat{s}e$  is the approximate standard error for  $\hat{a}_{jk}$ .

Figure 5 shows actual coverage versus target coverage, estimated by combining across all 50 runs and across all entries of each parameter matrix/vector. Perfect coverage would be a straight line on the diagonal. We exclude  $c_{11}$ ,  $D$ , and  $\omega$  in these coverage results since it seems challenging to estimate them without non-negligible bias, skewing the results. We see in Figure 5 that the actual coverage for  $A$ ,  $B$ ,  $C$ ,  $U$ , and  $S$  is excellent at every target coverage level from 0% to 100%. For  $V$  and  $T$ , the coverage is reasonably good for the smaller values of  $I$  but appears to degrade when  $I$  increases.

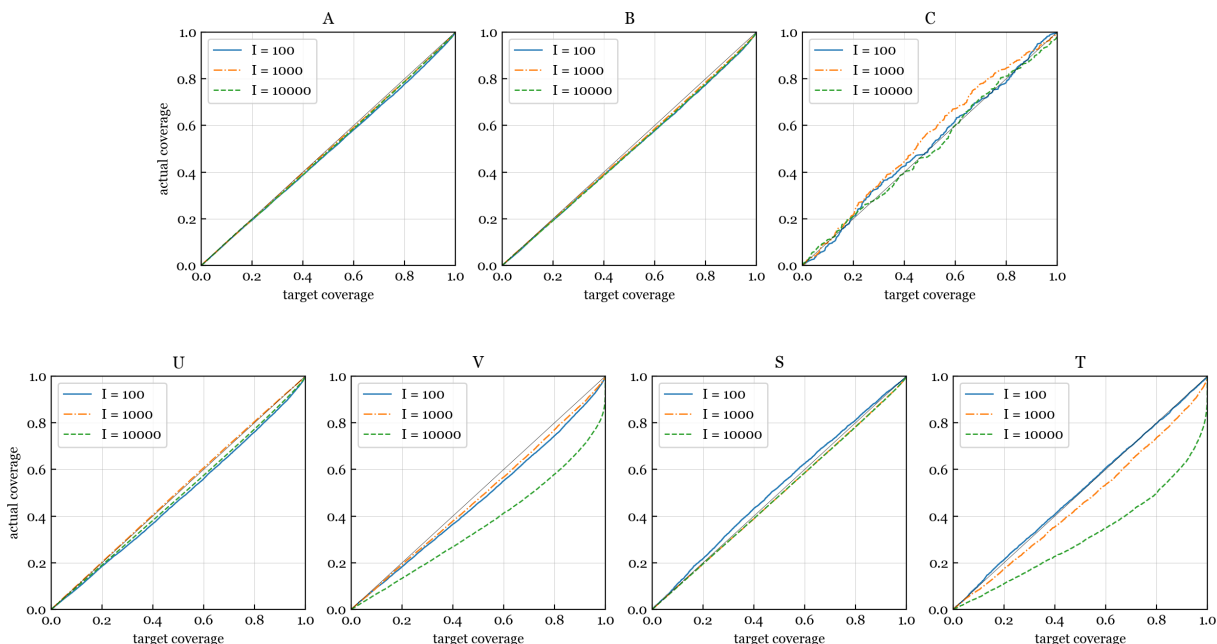


Figure 5: Coverage of confidence intervals for the entries of each parameter matrix/vector.

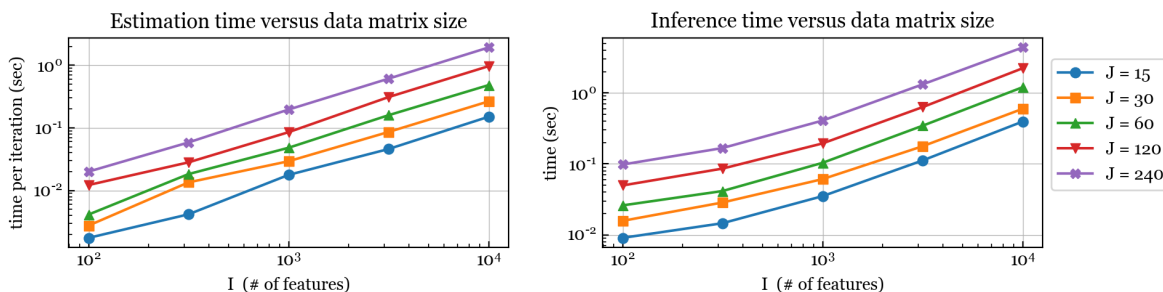


Figure 6: Computation time of our GBM algorithms as a function of  $I$  and  $J$ .

### 6.3 Computation time and algorithm convergence

**Computation time.** For each combination of  $I \in \{100, 316, 1000, 3162, 10000\}$  and  $J \in \{15, 30, 60, 120, 240\}$ , we generate 10 data matrices using the NB/Normal/Normal simulation scheme with  $K = 4$ ,  $L = 2$ , and  $M = 3$ . For each  $I$  and  $J$ , Figure 6 shows the average computation time per iteration of the estimation algorithm, along with the average computation time for the inference algorithm. These empirical results agree with our theory in Section 5 showing that the time per iteration scales like  $IJ$  (that is, linearly with the size of the data matrix) and the time for inference scales like  $IJ^2$ .

**Algorithm convergence.** Next, we evaluate the number of iterations required for

the estimation algorithm to converge. Similarly to before, for  $I \in \{100, 1000, 10000\}$ , we run the NB/Normal/Normal scheme 25 times with  $J = 100$ ,  $K = 4$ ,  $L = 2$ , and  $M = 3$ . Figure S2 shows the log-likelihood+log-prior (plus a constant) versus iteration number for each simulation run. In these simulations, the log-likelihood+log-prior levels off after around 5 or fewer iterations, indicating that the algorithm converges rapidly.

## 6.4 Robustness to the outcome distribution

To assess the robustness of the NB-GBM to the assumption that the outcome distribution is negative binomial, we rerun the experiments in Sections 6.1 and 6.2 using the following data simulation schemes: (a) LNP/Normal/Normal (b) Poisson/Normal/Normal, and (c) Geometric/Normal/Normal. The results in Figures S3, S4, and S5 show that the algorithms are quite robust to misspecification of the outcome distribution.

## 7 Application to gene expression analysis

In this section, we evaluate our GBM algorithms on RNA-seq gene expression data. An RNA-seq dataset consists of a matrix of counts in which entry  $(i, j)$  is the number of high-throughput sequencing reads that were mapped to gene  $i$  for sample  $j$ . These read counts are related to gene expression level, but there are many biases – both sample-specific and gene-specific. More generally, there are often significant sources of unwanted variation—both biological and technical—that obscure the signal of interest. Most methods use pipelines that adjust for each bias sequentially, rather than in an integrated way. GBMs enable one to use a single coherent model that adjusts for gene covariates and sample covariates as well as unobserved factors such as batch effects.



## 7.1 Comparing to DESeq2 on lymphoblastoid cell lines

As a test of our GBM methods, we compare with DESeq2 (Love et al., 2014), a leading method for RNA-seq differential expression analysis. We first consider a benchmark dataset used by Love et al. (2014) consisting of 161 samples from lymphoblastoid cell lines (Pickrell et al., 2010). We use the subset of 20,815 genes with nonzero median count across samples.

In both DESeq2 and the GBM, we adjust for two sample covariates: sequencing center (Argonne or Yale) and cDNA concentration. To adjust for GC content, which is often the most important gene covariate, in the GBM we construct the  $X$  matrix using a natural cubic spline basis with knots at the 2.5%, 25%, 50%, 75%, and 97.5% quantiles of GC content. DESeq2 does not have a built-in capacity to adjust for gene covariates, so for DESeq2, we adjust for GC using their recommended approach of pre-computing normalization factors using CQN (Hansen et al., 2012), which uses the same spline basis. Since DESeq2 does not adjust for latent factors, we first set  $M = 0$  for direct comparison; later, we set  $M = 2$ .

It is natural to use negative binomial (NB) outcomes for sequencing data since the technical variability is close to Poisson (Marioni et al., 2008), and biological variability introduces overdispersion (Robinson et al., 2010). These modeling choices yield an NB-GBM with  $I = 20,815$ ,  $J = 161$ ,  $K = 7$ , and  $L = 3$ . DESeq2 also uses an NB model, so the main difference between DESeq2 and this particular GBM is the way that the parameters and standard errors are estimated. Using a 1.8GHz processor, GBM estimation and inference took 42 seconds, whereas DESeq2+CQN took 105 seconds.

**Correctness of p-values under mock null comparisons.** First, we assess the calibration of p-values for testing for differential expression between two conditions. Under the null hypothesis of no difference, the p-values would ideally be uniformly distributed on  $[0, 1]$ . Since the Pickrell samples appear to be relatively homogeneous (when adjusting for sequencing center and cDNA concentration), we can assess how well this ideal is attained by randomly splitting the samples into two groups and testing for differential expression.

To this end, we add a sample covariate  $z_{j4}$  consisting of a dummy variable for the




Figure 7: Comparison of p-value distributions on Pickrell data using DESeq2 and the GBM. (Left) CDFs of p-values for mock null comparisons over 50 random splits. (Middle) Same, but zoomed in to region shaded in left-hand figure. (Right) CDFs of p-values for testing for a difference between sequencing centers. The x-axis ends at the significance threshold for controlling FWER at 0.05.

assignment of samples to the two random groups. Thus, the null hypothesis of no difference for gene  $i$  is  $b_{i4} = 0$  and the alternative is  $b_{i4} \neq 0$ . The (two-sided) p-value for gene  $i$  is  $p_i := 2(1 - \Phi(|\hat{b}_{i4}/\hat{se}(b_{i4})|))$  where  $\Phi(x)$  is the standard normal CDF (cumulative distribution function). Figure 7 shows the p-value CDF over all genes, aggregating over 50 random splits into two groups containing 80 and 81 samples, respectively. Both DESeq2 and the GBM yield p-values that are very close to the ideal uniform distribution. This indicates that both methods are accurately controlling the false positive rate.

**Sensitivity to detect actual differences.** To compare sensitivity, we test for differential expression between sequencing centers by computing p-values  $p_i = 2(1 - \Phi(|\hat{b}_{i\ell}/\hat{se}(b_{i\ell})|))$  where  $\ell$  is the index of the sequencing center covariate. Using Bonferroni to control the family-wise error rate (FWER) at 0.05, the number of genes detected as differentially expressed by the GBM and DESeq2 are 1038 and 892, respectively. Figure 7 shows the lower tail of the p-value CDFs. In these results, the GBM yields equal or greater sensitivity.

**Visualization using GBM latent factors.** The latent factors of the GBM provide a model-based approach to visualizing high-dimensional count data, while adjusting for covariates. To illustrate, we modify the model to use  $M = 2$ . Figure 8 shows a scatterplot of  $v_{j2}$  versus  $v_{j1}$  for the estimated  $V$  matrix. Observe that this yields very tightly grouped clusters of samples from the same subject. This is analogous to plotting the first two




Figure 8: Visualization of Pickrell data using GBM latent factors, adjusting for covariates. Each dot represents one of the 161 samples, and each subject ID is indicated by a different combination of color and shape (by cycling through 32 colors and 3 shapes).

scores in principal components analysis (PCA). Thus, for comparison, Figure S15 shows the PCA plots based on (a) log-transformed TPMs (Transcripts per Million), specifically,  $\log(\text{TPM}_{ij} + 1)$ , and (b) the variance stabilizing transform (VST) in the DESeq2 package, using the GC adjustment from CQN. The DESeq2 model does not estimate latent factors, which is why PCA is used in DESeq2. The TPM plot is very noisy in terms of subject ID clusters. The VST plot is better than TPMs, but still not quite as clean as the GBM plot.

Overall, in terms of sensitivity, controlling false positives, computation time, and visualization, these results suggest that the GBM performs very well. When  $J$  is very small, it may be beneficial to augment the GBM to use DESeq2-like shrinkage estimates for  $s_i$ .

## 7.2 Analyzing GTEx data for aging-related genes

Next, we test our methods on an application of scientific interest, using RNA-seq data from the Genotype-Tissue Expression (GTEx) project (Melé et al., 2015), consisting of 8,551 samples from 30 tissues in the human body, obtained from 544 subjects. We apply




Figure 9: Visualization of GTEx data using NB-GBM latent factors, adjusting for covariates. Each dot represents one of the 8,551 samples, and the color indicates the tissue type.

the GBM to find genes whose expression changes with age, adjusting for technical biases. See Jia et al. (2018) and Zeng et al. (2020) for studies of age-related genes using GTEx.

We use the GTEx RNA-seq data from recount2 (Collado-Torres et al., 2017),<sup>1</sup> normalized using the `scale_counts` function in the `recount` R library. We use the subset of 8,551 samples that passed GTEx quality control, and the subset of genes in chromosomes 1–22 that have an HGNC-approved gene symbol and have nonzero median across all samples.

To visualize the samples, we take a random subset of 5,000 genes and estimate an NB-GBM with two latent factors and no sample covariates. For the gene covariates, we use  $\log(\text{length}_i)$ ,  $\text{gc}_i$ , and  $(\text{gc}_i - \overline{\text{gc}})^2$  where  $\text{length}_i$  is the sum of the exon lengths and  $\text{gc}_i$  is the GC content of gene  $i$ . Thus, in this initial model for visualization,  $I = 5,000$ ,  $J = 8,551$ ,  $K = 4$ ,  $L = 1$ , and  $M = 2$ . Figure 9 shows the latent factors ( $v_{j2}$  versus

<sup>1</sup>Downloaded from <https://jhubiostatistics.shinyapps.io/recount> on 8/7/2020.

$v_{j1}$ ), similarly to Figure 8. The samples tend to fall into clusters according to the tissue from which they were taken. Some tissues, such as brain and blood, clearly contain two or more subclusters which turn out to correspond to subtissue types (Figure S16). Meanwhile, when more latent factors are used (that is,  $M > 2$ ), some clusters that overlap in Figure 9 become well-separated in higher latent dimensions. For comparison, running PCA on the log TPMs is not nearly as clear in terms of tissue/subtissue clusters (Figure S17).

**Testing for age-related genes.** To find genes that are related to aging, we add subject age as a sample covariate. Each gene then has a coefficient describing how its expression changes with age, and we compute a p-value for each gene to test whether its coefficient is nonzero. Due to the heterogeneity of tissue/subtissue types, we analyze each subtissue type separately. To perform both exploratory analysis and valid hypothesis testing, we used a random subset of 108 subjects during an exploratory model-building phase and then used the remaining 436 subjects during a testing phase with the selected model.

In the exploratory phase, we considered adjusting for various technical sample covariates and gene covariates, and varied  $M$  from 0 to 10. For each model and each subtissue type, we used the GBM to find the set of genes exhibiting a significant association with age, controlling FWER at 0.05 using the Bonferroni correction. To score the relevance of each of these gene sets in terms of aging biology, we computed its F1 score for overlap with the set of aging-related genes identified by De Magalhães et al. (2009).<sup>2</sup> Based on this exploratory analysis, we chose to keep  $\log(\text{length}_i)$ ,  $\text{gc}_i$ , and  $(\text{gc}_i - \overline{\text{gc}})^2$  as gene covariates, and use `smexncrt` (exonic rate, the fraction of reads that map within exons) as well as `age` (subject age, coded as a numerical value in  $\{25, 35, \dots, 75\}$ ) as sample covariates. For each subtissue, we chose the  $M$  that yielded the highest F1 score on the exploratory data.

In the testing phase, we apply the selected model for each subtissue to test for age-associated genes. For illustration, we present results for the “Heart - Left Ventricle” subtissue (Heart-LV), which had the highest F1 score across all subtissues on the exploratory data. We ran the GBM on the 176 Heart-LV samples in the test set, using the 19,853 genes

---

<sup>2</sup>From [https://genomics.senescence.info/gene\\_expression/signatures.html](https://genomics.senescence.info/gene_expression/signatures.html) on 8/11/2020.




Figure 10: Estimated PCMT1 expression based on log counts, log TPMs, and GBM residuals.

with nonzero median across these samples, with  $M = 3$  based on the exploratory phase. Thus, in this model,  $I = 19,853$ ,  $J = 176$ ,  $K = 4$ ,  $L = 3$ , and  $M = 3$ .

**Results.** We found 2,444 genes to be significantly associated with age in Heart-LV, controlling FWER at 0.05. For comparison, simple linear regression on the log TPMs yields only 1 significant gene; thus, the GBM has much greater power than a simple standard approach. To validate the biological relevance of the GBM hits, we compare with what is known from the aging literature. First, the top GBM hit for Heart-LV is PCMT1 (Entrez gene ID 5110) with a p-value of  $1.1 \times 10^{-47}$ . PCMT1 is involved in the repair and degradation of damaged proteins, and is a well-known aging gene, being one of 307 human genes in the GenAge database (build version 20) from Tacutu et al. (2018). Figure 10 shows the estimated expression of PCMT1 versus age for the Heart-LV samples. The GBM-estimated expression exhibits a clear downward linear trend with age. For comparison, Figure 10 shows that the log TPMs are considerably noisier and the trend is much less clear. Simple linear regression on the log TPMs yields a p-value of  $1.1 \times 10^{-3}$  for PCMT1, which does not reach the Bonferroni significance threshold of  $0.05/I = 2.5 \times 10^{-6}$ . Here, we define the GBM-estimated expression as the partial residual  $\hat{c}_{11} + \hat{a}_{j1} + \hat{b}_{i1} + (\hat{c}_{1\ell} + \hat{b}_{i\ell})z_{j\ell} + \varepsilon_{ij}$  where  $\ell$  is the index of the `age` column in  $Z$ , and  $\varepsilon_{ij}$  is the GBM residual (Section 2.3).

To evaluate the GBM hits altogether for biological relevance to aging, we test for enrichment of Gene Ontology (GO) term gene sets using DAVID v6.8 (Huang et al., 2009a,b).

We run DAVID on the top 1000 GBM hits for Heart-LV, using all 19,853 tested genes as the background list. (DAVID allows at most 1000 genes.) Tables S2 and S3 show the top 20 enriched GO terms in the Biological Process and Cellular Component categories. These results are highly consistent with known aging biology (López-Otín et al., 2013).

## 8 Application to cancer genomics

Next, we apply the GBM to estimate copy ratios for sequencing data from cancer cell lines. Copy ratio estimation is an essential step in detecting somatic copy number alterations (SCNAs), that is, duplications or deletions of segments of the genome. The input data is a matrix of counts where entry  $(i, j)$  is the number of reads from sample  $j$  that map to target region  $i$  of the genome. The goal is to estimate the copy ratio of each region, that is, the relative concentration of copies of that region in the original DNA sample.

Simple estimates based on row and column normalization are very noisy and are contaminated by significant technical biases. State-of-the-art methods employ a panel of normals (that is, sequencing samples from non-cancer tissues) to estimate technical biases using principal components analysis (PCA), and then use linear regression to remove these biases from the cancer samples of interest. We take an analogous approach, first running a GBM on a panel of normals, and then running a GBM on the cancer samples using a feature covariate matrix  $X$  that includes the  $U$  matrix estimated from the panel of normals.

To assess performance, we compare with the state-of-the-art method provided by the Broad Institute’s Genome Analysis Toolkit (GATK) (Broad Institute, 2020) on the 326 whole-exome sequencing samples from the Cancer Cell Line Encyclopedia (CCLE) (Ghandi et al., 2019). These samples are from a wide range of cancer types, including lung, breast, colon, prostate, brain, and many others. We use the subset of 180,495 target regions that are in chromosomes 1–22 and have nonzero median count across the 326 samples.

Since there are essentially no normal samples in the CCLE dataset, we create a panel of pseudo-normals by taking a random subset of 163 samples as a training set and de-

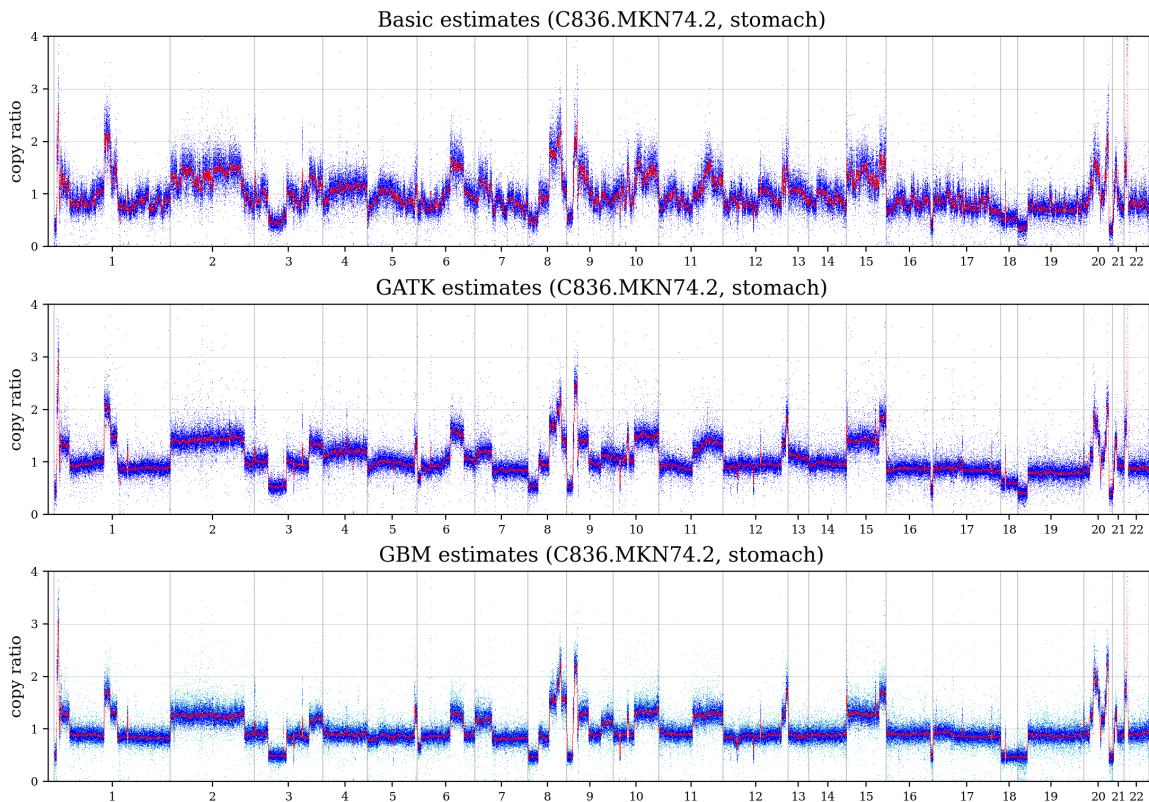


Figure 11: Copy ratio estimates for an illustrative sample from the CCLE data. The x-axis is genomic position, and each blue dot is the estimate for one region; moving averages are in red. For the GBM, regions with high and low precision estimates are plotted in blue and cyan, respectively.

segmenting them to adjust for copy number alterations; see Section S3.2 for details. The remaining 163 samples are used as a test set. For the GBM, we use  $\log(\text{length}_i)$ ,  $\text{gc}_i$ , and  $(\text{gc}_i - \overline{\text{gc}})^2$  as region covariates, no sample covariates, and 5 latent factors. Thus,  $I = 180,495$ ,  $J = 163$ ,  $K = 4$ ,  $L = 1$ , and  $M = 5$  on the training set, while  $I = 180,495$ ,  $J = 163$ ,  $K = 9$ ,  $L = 1$ , and  $M = 0$  on the test set. We define the GBM copy ratio estimates as the exponentiated residuals  $\tilde{Y}_{ij}/\hat{\mu}_{ij}$  where  $\tilde{Y}_{ij} := Y_{ij} + 1/8$ ; see Section 2.3. The GBM took 10 minutes and 4.3 minutes to run on the training and test sets, respectively, while GATK took 3.3 minutes and 28 minutes on training and test, respectively. The slowness of GATK on the test set is likely due to having to run it separately on every test sample.

Figure 11 shows the GBM and GATK copy ratio estimates for an illustrative sample from the test set. As a baseline, we also show the simple normalization-based estimates defined as  $\tilde{Y}_{ij}/(\alpha_i\beta_j)$  where  $\alpha_i = \frac{1}{J} \sum_{j=1}^J \tilde{Y}_{ij}$  and  $\beta_j = \frac{1}{I} \sum_{i=1}^I \tilde{Y}_{ij}/\alpha_i$ . A major advantage of






Figure 12: Performance of the GBM versus GATK on the 163 test samples from the CCLE whole-exome sequencing dataset. The GBM exhibits better performance in terms of both metrics.

the GBM is that it provides uncertainty quantification. Here, the estimated precision (that is, the inverse variance) of each log copy ratio estimate is  $w_{ij}$  (Section 2.3). In the GBM plot in Figure 11, this is illustrated by using cyan for the regions with low relative precision; see Section S3.2. By downweighting regions with low estimated precision, downstream analyses such as SCNA detection can be made more accurate.

To quantify performance, Figure 12 compares the GBM and GATK in terms of two performance metrics. The local relative standard error quantifies the variability of the log copy ratio estimates around a weighted moving average, accounting for the precision of each estimate. Meanwhile, the weighted median absolute difference quantifies the typical magnitude of the slope of a weighted moving average. On these data, the GBM exhibits better performance in terms of both metrics; see Section S3.2 for details. Figures S18 and S19 show the GBM and GATK copy ratio estimates for all 163 test samples. The GBM estimates are visibly less noisy than the GATK estimates.

Overall, the GBM appears to perform very well in terms of removing technical biases and denoising, particularly when using uncertainty quantification to downweight low precision regions. The improved performance appears to be due to (a) model-based uncertainty quantification and (b) using a robust probabilistic model for count data.

## 9 Conclusion

Generalized bilinear models provide a flexible framework for the analysis of matrix data, and the delta propagation method enables accurate GBM uncertainty quantification in modern applications. In future work, it would be interesting to extend to the more general model of Gabriel (1998), provide theoretical guarantees for delta propagation, and try applying delta propagation to other models.

## Acknowledgments

We would like to thank Jonathan Huggins, Will Townes, Mehrtash Babadi, Samuel Lee, David Benjamin, Robert Klein, Samuel Markson, Philipp Hähnel, and Rafael Irizarry for many helpful conversations.

## References

- Agresti, A. *Foundations of Linear and Generalized Linear Models*. John Wiley & Sons, 2015.
- Babadi, M., Lee, S. K., and Smirnov, A. N. GATK gCNV: accurate germline copy-number variant discovery from sequencing read-depth data. The International Conference on Probabilistic Programming (PROBPROG), Oct 2018.
- Bishop, C. M. *Pattern Recognition and Machine Learning*. springer, 2006.
- Broad Institute. Genome Analysis Toolkit (GATK) v4.1.8.1, 2020. URL <https://gatk.broadinstitute.org/>.
- Buettner, F., Pratanwanich, N., McCarthy, D. J., Marioni, J. C., and Stegle, O. f-scLVM: scalable and versatile factor analysis for single-cell RNA-seq. *Genome Biology*, 18(1):212, 2017.
- Carvalho, C. M., Chang, J., Lucas, J. E., Nevins, J. R., Wang, Q., and West, M. High-dimensional sparse factor modeling: applications in gene expression genomics. *Journal of the American Statistical Association*, 103(484):1438–1456, 2008.
- Choulakian, V. Generalized bilinear models. *Psychometrika*, 61(2):271–283, 1996.
- Collado-Torres, L., Nellore, A., Kammers, K., Ellis, S. E., Taub, M. A., Hansen, K. D., Jaffe, A. E., Langmead, B., and Leek, J. T. Reproducible RNA-seq analysis using recount2. *Nature Biotechnology*, 35(4):319–321, 2017.
- de Falguerolles, A. GBMs: GLMs with bilinear terms. In *COMPSTAT*, pages 53–64. Springer, 2000.

- De Magalhães, J. P., Curado, J., and Church, G. M. Meta-analysis of age-related gene expression profiles identifies common signatures of aging. *Bioinformatics*, 25(7):875–881, 2009.
- Fromer, M., Moran, J. L., Chambert, K., Banks, E., Bergen, S. E., Ruderfer, D. M., Handsaker, R. E., McCarroll, S. A., O’Donovan, M. C., Owen, M. J., et al. Discovery and statistical genotyping of copy-number variation from whole-exome sequencing depth. *The American Journal of Human Genetics*, 91(4):597–607, 2012.
- Gabriel, K. R. Generalised bilinear regression. *Biometrika*, 85(3):689–700, 1998.
- Ghandi, M., Huang, F. W., Jané-Valbuena, J., Kryukov, G. V., Lo, C. C., McDonald, E. R., Barretina, J., Gelfand, E. T., Bielski, C. M., Li, H., et al. Next-generation characterization of the cancer cell line encyclopedia. *Nature*, 569(7757):503–508, 2019.
- Hansen, K. D., Irizarry, R. A., and Wu, Z. Removing technical variability in RNA-seq data using conditional quantile normalization. *Biostatistics*, 13(2):204–216, 2012.
- Hoff, P. D. Multilinear tensor regression for longitudinal relational data. *The Annals of Applied Statistics*, 9(3):1169, 2015.
- Huang, D. W., Sherman, B. T., and Lempicki, R. A. Bioinformatics enrichment tools: paths toward the comprehensive functional analysis of large gene lists. *Nucleic Acids Research*, 37(1):1–13, 2009a.
- Huang, D. W., Sherman, B. T., and Lempicki, R. A. Systematic and integrative analysis of large gene lists using DAVID bioinformatics resources. *Nature Protocols*, 4(1):44, 2009b.
- Jia, K., Cui, C., Gao, Y., Zhou, Y., and Cui, Q. An analysis of aging-related genes derived from the Genotype-Tissue Expression project (GTEx). *Cell Death Discovery*, 4(1):1–14, 2018.
- Jiang, Y., Oldridge, D. A., Diskin, S. J., and Zhang, N. R. CODEX: A normalization and copy number variation detection method for whole exome sequencing. *Nucleic Acids Research*, 43(6):e39–e39, 2015.
- Jørgensen, B. Exponential dispersion models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 49(2):127–145, 1987.
- Krumm, N., Sudmant, P. H., Ko, A., O’Roak, B. J., Malig, M., Coe, B. P., Quinlan, A. R., Nickerson, D. A., and Eichler, E. E. Copy number variation detection and genotyping from exome sequence data. *Genome Research*, 22(8):1525–1532, 2012.
- Leek, J. T. and Storey, J. D. Capturing heterogeneity in gene expression studies by surrogate variable analysis. *PLoS Genetics*, 3(9):e161, 2007.
- Leek, J. T. and Storey, J. D. A general framework for multiple testing dependence. *Proceedings of the National Academy of Sciences*, 105(48):18718–18723, 2008.
- López-Otín, C., Blasco, M. A., Partridge, L., Serrano, M., and Kroemer, G. The hallmarks of aging. *Cell*, 153(6):1194–1217, 2013.
- Love, M. I., Huber, W., and Anders, S. Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome Biology*, 15(12):550, 2014.

- Marioni, J. C., Mason, C. E., Mane, S. M., Stephens, M., and Gilad, Y. RNA-seq: An assessment of technical reproducibility and comparison with gene expression arrays. *Genome Research*, 18(9):1509–1517, 2008.
- Melé, M., Ferreira, P. G., Reverter, F., DeLuca, D. S., Monlong, J., Sammeth, M., Young, T. R., Goldmann, J. M., Pervouchine, D. D., Sullivan, T. J., et al. The human transcriptome across tissues and individuals. *Science*, 348(6235):660–665, 2015.
- Perry, P. O. and Pillai, N. S. Degrees of freedom for combining regression with factor analysis. *arXiv preprint arXiv:1310.7269*, 2013.
- Pickrell, J. K., Marioni, J. C., Pai, A. A., Degner, J. F., Engelhardt, B. E., Nkadori, E., Veyrieras, J.-B., Stephens, M., Gilad, Y., and Pritchard, J. K. Understanding mechanisms underlying human gene expression variation with RNA sequencing. *Nature*, 464(7289):768–772, 2010.
- Risso, D., Ngai, J., Speed, T. P., and Dudoit, S. Normalization of RNA-seq data using factor analysis of control genes or samples. *Nature Biotechnology*, 32(9):896–902, 2014.
- Robinson, M. D., McCarthy, D. J., and Smyth, G. K. edgeR: A Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics*, 26(1):139–140, 2010.
- Stegle, O., Parts, L., Durbin, R., and Winn, J. A Bayesian framework to account for complex non-genetic factors in gene expression levels greatly increases power in eQTL studies. *PLoS Comput Biol*, 6(5):e1000770, 2010.
- Sun, Y., Zhang, N. R., and Owen, A. B. Multiple hypothesis testing adjusted for latent variables, with an application to the AGEMAP gene expression data. *The Annals of Applied Statistics*, 6(4):1664–1688, 2012.
- Tacutu, R., Thornton, D., Johnson, E., Budovsky, A., Barardo, D., Craig, T., Diana, E., Lehmann, G., Toren, D., Wang, J., et al. Human ageing genomic resources: New and updated databases. *Nucleic Acids Research*, 46(D1):D1083–D1090, 2018.
- Takane, Y. and Shibayama, T. Principal component analysis with external information on both subjects and variables. *Psychometrika*, 56(1):97–120, 1991.
- Zeng, L., Yang, J., Peng, S., Zhu, J., Zhang, B., Suh, Y., and Tu, Z. Transcriptome analysis reveals the difference between “healthy” and “common” aging and their connection with age-related diseases. *Aging Cell*, 19(3):e13121, 2020.

# Supplementary material for “Inference in generalized bilinear models”

## S1 Discussion

### S1.1 Previous work

There is an extensive literature on models involving an unknown low-rank matrix, going by a variety of names including latent factor models, factor analysis models, multiplicative models, bi-additive models, and bilinear models. In particular, a large number of models can be viewed as special cases of generalized bilinear models (GBMs). Since a full review is beyond the scope of this article, we settle for covering the main threads in the literature.

#### S1.1.1 Normal bilinear models without covariates.

Principal components analysis (PCA) is equivalent to maximum likelihood estimation in a GBM with only the  $UDV^T$  term ( $K = 0$ ,  $L = 0$ ,  $M > 0$ ), assuming normally distributed outcomes with common variance  $\sigma^2$ . PCA (or equivalently, the SVD) is often performed after centering the rows and columns of the data matrix, and from a model-based perspective, this is equivalent to including intercepts ( $K = 1$ ,  $L = 1$ ,  $M > 0$ ):

$$Y_{ij} = c + a_i + b_j + \sum_{m=1}^M u_{im}d_m v_{jm} + \varepsilon_{ij} \quad (\text{S1.1})$$

where  $\varepsilon_{ij}$  is a normal residual. Similarly, scaling the rows and columns is analogous to using a rank-one factorization of the variance, that is,  $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma_i^2 \sigma_j^2)$ .

**Estimation.** Equation S1.1 is often called the AMMI (additive main effects and multiplicative interaction) model, and a range of techniques for using it have been developed (Gauch, 1988). The least squares fit of an AMMI model can be obtained by first fitting the linear terms  $c + a_i + b_j$  ignoring the non-linear term, and then estimating the non-linear

term  $\sum_{m=1}^M u_{im}d_mv_{jm}$  using PCA on the residuals (Gilbert, 1963; Gollob, 1968; Mandel, 1969; Gabriel, 1978). Estimation is more difficult when each entry is allowed to have a different variance, that is, when  $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma_{ij}^2)$  with  $\sigma_{ij}^2$  known; this is sometimes called a weighted AMMI model (Van Eeuwijk, 1995). To handle this, Gabriel and Zamir (1979) develop the criss-cross method of estimation, which successively fits the non-linear terms  $m = 1, \dots, M$ , one by one, using weighted least squares.

**Hypothesis testing for model selection.** While most applications of PCA only use the estimates, without any uncertainty quantification, statistical research on the AMMI model has largely focused on hypothesis testing for which factors  $m$  to include in the model. Early contributions on testing in this model were made by Fisher and Mackenzie (1923), Cochran (1943), Tukey (1949), Williams (1952), Mandel (1961), Gollob (1968), and Mandel (1969). Methods of this type are very widely used, particularly in the study of genotype-environment interactions in agronomy; see reviews by Freeman (1973), Gauch (2006), and Gauch et al. (2008).

**Confidence regions for parameters.** Uncertainty quantification for the AMMI model parameters has also been studied. Asymptotic covariance formulas for the least squares estimates have been given by Goodman and Haberman (1990), Chadoeuf and Denis (1991), Dorkenoo and Mathieu (1993) and Denis and Gower (1996) for the AMMI model and various special cases. These results are based on inverting the constraint-augmented Fisher information matrix (Aitchison and Silvey, 1958; Silvey, 1959); we use the same technique in Section S6 to estimate standard errors for  $U$  and  $V$  in our more general GBM model and we extend it using delta propagation.

### S1.1.2 Normal bilinear models with covariates.

**Estimation.** In a wide-ranging article, Tukey (1962) discussed the possibility of combining regression with factor analysis, by factoring the matrix of residuals after adjusting for covariates. Indeed, for the case of normal outcomes with common variance, Gabriel (1978, Cor 3.1) showed that when using a model of the form  $\mathbf{Y} = XA^T + BZ^T + UDV^T + \boldsymbol{\varepsilon}$ , the

least squares fit can be obtained simply by first fitting  $A$  and  $B$  using regression (ignoring  $UDV^T$ ), then fitting  $UDV^T$  to the residuals. This can be viewed as a generalization of the AMMI estimation procedure. Takane and Shibayama (1991) extend the results of Gabriel (1978) by first fitting  $\mathbf{Y} = XA^T + BZ^T + XCZ^T + \boldsymbol{\varepsilon}$  using least squares, and then using PCA to analyze the residuals as well as each fitted component of the model, that is,  $X\hat{A}^T$ ,  $\hat{B}Z^T$ ,  $X\hat{C}Z^T$ ,  $\hat{\boldsymbol{\varepsilon}}$ , and combinations thereof.

In a complementary direction, reduced-rank regression (Davies and Tso, 1982) and CANDELINC (Carroll et al., 1980) use least squares to fit models of the form  $\mathbf{Y} = XA^T + \boldsymbol{\varepsilon}$  and  $\mathbf{Y} = XCZ^T + \boldsymbol{\varepsilon}$ , respectively, where  $A$  and  $C$  are constrained to be low-rank.

**Hypothesis testing and confidence regions.** In the case of normal outcomes with common variance  $\sigma^2$ , for the model with  $\mathbf{Y} = XA^T + BZ^T + UDV^T + \boldsymbol{\varepsilon}$ , Perry and Pillai (2013) show how to perform inference for univariate entries of  $A$  and  $B$  (and univariate linear projections, more generally) accounting for the uncertainty in  $UDV^T$  via an estimate of the degrees of freedom associated with the latent factors. Further, Perry and Pillai (2013) show that the problem can be reduced to the covariate-free case,  $\mathbf{Y} = UDV^T + \boldsymbol{\varepsilon}$ , enabling one to use results on hypothesis testing in the AMMI model (Gollob, 1968; Mandel, 1969) which provide estimates of the degrees of freedom. However, this approach appears to rely on the assumption of normal outcomes with common variance.

### S1.1.3 Generalized bilinear models without covariates.

In many applications, it is unreasonable to use a normal outcome model. A classical approach is to transform the data and then apply a normal outcome model, however, as discussed by Van Eeuwijk (1995), there is unlikely to be a transformation that simultaneously achieves (a) approximate normality, (b) common variance, and (c) additive effects.

A more principled approach is to extend the generalized linear model (GLM) framework to handle latent factors, as suggested by Gower (1989). Goodman’s RC models are early contributions in this direction (Goodman, 1979, 1981, 1986, 1991), consisting of count models with multinomial or Poisson outcomes where  $\log(\mathbb{E}(Y_{ij})) = c + a_i + b_j + \sum_{m=1}^M u_{im}d_m v_{jm}$ .

More generally, Van Eeuwijk (1995) develops the generalized AMMI (GAMMI) model, which is a GLM version of the AMMI model in Equation S1.1, specifically,  $g(\mathbb{E}(Y_{ij})) = c + a_i + b_j + \sum_{m=1}^M u_{im}d_mv_{jm}$ . Van Eeuwijk (1995) introduces a coordinate descent algorithm and discusses approaches for choosing  $M$ , however, he does not consider uncertainty quantification for parameters, does not estimate dispersion parameters, and only demonstrates the method on very small datasets ( $11 \times 5$  and  $17 \times 12$ ).

Correspondence analysis (Benzécri, 1973; Greenacre, 1984) is an SVD-based exploratory analysis method for matrices of categorical data, and has been reinvented under many names, such as reciprocal averaging and dual scaling (de Falguerolles, 2000). Correspondence analysis bears resemblance to estimation methods for the GAMMI model, however, it is primarily descriptive in perspective, and thus typically does not involve quantification of uncertainty.

#### S1.1.4 Generalized bilinear models with covariates.

Choulakian (1996) defines a class of GBMs of the same form as in this article, where  $g(\mathbb{E}(\mathbf{Y})) = XA^T + BZ^T + XCZ^T + UDV^T$  and  $g$  is the (a) canonical, (b) identical, or (c) logarithmic link function. For the case of no covariates (that is, the GAMMI model), Choulakian (1996) proposes an estimation algorithm that involves univariate Fisher scoring updates, which is attractive for its simplicity, but may exhibit slow convergence or failure to converge due to strong dependencies among parameters. While the defined model class is general, some limitations of the paper by Choulakian (1996) are that uncertainty quantification is not addressed, the estimation algorithm is for the special case of GAMMI, a single common dispersion is assumed and estimation of dispersion is not addressed, identifiability constraints are not enforced, no initialization procedure is provided, and only very small datasets are considered ( $10 \times 7$  and  $11 \times 11 \times 2$ ).

Gabriel (1998) considers a very general class of models of the form  $g(\mathbb{E}(\mathbf{Y})) = \sum_{k=1}^K X_k \Theta_k Z_k$ , where  $X_k$  and  $Z_k$  are observed matrices (for instance, covariates) and  $\Theta_k$  is a low-rank matrix of parameters for each  $k = 1, \dots, K$ . He extends the criss-cross estimation



algorithm of Gabriel and Zamir (1979) to this model. While the model of Gabriel (1998) is very elegant, some limitations are that estimation is performed using a vectorization approach that is computationally prohibitive on large matrices, uncertainty quantification is not addressed, a common dispersion parameter is assumed for all entries, and only very small datasets are considered ( $10 \times 9$  and  $17 \times 2$ ). Also, it is not clear what identifiability constraints are assumed on the  $\Theta_k$  matrices.

In recent work, Townes (2019) considers a model of the form  $g(\mathbb{E}(\mathbf{Y})) = XA^T + BZ^T + UDV^T + \mathbf{1}\delta^T$  where  $\mathbf{1}$  is a vector of ones and  $\delta \in \mathbb{R}^J$  is a vector of fixed column-specific offsets. Townes (2019) derives diagonal approximations to Fisher scoring updates for  $\ell_2$ -penalized maximum likelihood estimation, and in a postprocessing stage, enforces orthogonality constraints to aid interpretability. Other differences compared to the present work are that only estimation is considered (uncertainty quantification is not addressed) and overdispersion parameters are not estimated.

The overview by de Falguerolles (2000) provides an interesting and insightful discussion of several threads in the literature.

### **S1.1.5 Recent applications of bilinear models.**

Several authors have used bilinear models or GBMs in genetics and genomics, usually to remove unwanted variation such as batch effects. However, most of these methods do not fully account for uncertainty in the latent factors, which may lead to miscalibrated inferences such as overconfident p-values. For example, to remove batch effects in gene expression analysis, several approaches involve first estimating  $UDV^T$  and then treating  $V$  as a known matrix of covariates, accounting for uncertainty only in  $UD$  using standard regression (Leek and Storey, 2007, 2008; Sun et al., 2012; Risso et al., 2014); this is also done to adjust for population structure in genetic association studies (Price et al., 2006). In copy number variation detection, it is common to simply treat the estimated  $UDV^T$  as known and subtract it off (Fromer et al., 2012; Krumm et al., 2012; Jiang et al., 2015).

Carvalho et al. (2008) use a Bayesian sparse factor analysis model with covariates, employing evolutionary stochastic search for model selection and Markov chain Monte Carlo (MCMC) for posterior inference within models. Stegle et al. (2010), Buettner et al. (2017), and Babadi et al. (2018) use complex hierarchical models that can be viewed as Bayesian GBMs with additional prior structure, and they employ variational methods for approximate posterior inference. Another application in which bilinear models have seen recent use is longitudinal relational data such as networks, for which Hoff (2015) employs an interesting Bayesian model with  $\mathbf{Y}_t = AX_tB^T + \boldsymbol{\varepsilon}$ , where  $X_t$  is a matrix of observed covariates that depend on time  $t$ .

## S1.2 Challenges and solutions

Estimation and inference in large GBMs is complicated by a number of nontrivial challenges. In this section, we discuss several issues and how we resolve them.

**Estimating the dispersion parameters.** There are several issues with estimating the negative binomial (NB) dispersions  $1/r_{ij}$ . First, since there is insufficient information to estimate all  $IJ$  dispersions individually, we use the rank-one parametrization  $1/r_{ij} = \exp(s_i + t_j + \omega)$ . Second, the choice of identifiability constraints matters — the natural choice of constraints,  $\sum_i s_i = 0$  and  $\sum_j t_j = 0$ , leads to noticeably biased estimates of  $s_i$  and  $t_j$ , particularly for higher values; see Figure S6. Instead, we constrain  $\frac{1}{I} \sum_i e^{s_i} = 1$  and  $\frac{1}{J} \sum_j e^{t_j} = 1$ , which effectively mitigates this bias, empirically. Third, the maximum likelihood estimates sometimes exhibit a severe downward bias, particularly for low values of log-dispersion; we use a simple heuristic bias correction to deal with this. Fourth, to avoid arithmetic underflow/overflow in the log-dispersion update steps, we develop carefully constructed expressions for the gradient and Hessian. Finally, to prevent occasional lack of convergence due to oscillating estimates, we employ an adaptive maximum step size.

**Inapplicability of standard GLM methods.** Since  $XA^T + BZ^T + XCZ^T$  is linear in the parameters, one could vectorize and write it as  $\text{vec}(XA^T + BZ^T + XCZ^T) = \tilde{X}\beta$  where

$\beta = (\text{vec}(A)^T, \text{vec}(B)^T, \text{vec}(C)^T)^T \in \mathbb{R}^{JK+IL+KL}$  and  $\tilde{X} \in \mathbb{R}^{IJ \times (JK+IL+KL)}$  is a function of  $X$  and  $Z$ . In principle, one could then apply standard GLM estimation methods for estimating  $\beta$  to construct a joint update to  $(A, B, C)$ . However, this vectorization approach is only computationally feasible for small data matrices since computing the matrix inverse  $(\tilde{X}^T \tilde{X})^{-1}$  takes on the order of  $(JK + IL + KL)^3$  time. Further, this update would need to be done repeatedly since  $D, U, V, S, T$ , and  $\omega$  also need to be simultaneously estimated, and the vectorization approach does not help estimate these parameters.

**Inapplicability of the singular value decomposition.** At first glance, it might appear that the singular value decomposition (SVD) would make it straightforward to estimate  $U, D$ , and  $V$  given the other parameters. However, when used for estimation, the SVD implicitly assumes that every entry has the same variance. This is far from true in GBMs, and consequently, naively using the SVD to update  $UDV^T$  leads to poor estimation accuracy. The criss-cross algorithm of Gabriel and Zamir (1979) yields a low-rank matrix factorization that accounts for entry-specific variances, and our algorithm provides another way of doing this while adjusting for covariates in a GBM. In our algorithm, we only directly use the SVD for enforcing the identifiability constraints, not for estimation of  $UDV^T$ .

**Computational efficiency.** The genomics applications in Sections 7 and 8 involve large count matrices  $\mathbf{Y} \in \mathbb{R}^{I \times J}$  where the number of features  $I$  is on the order of  $10^4$  to  $10^6$  and the number of samples  $J$  can be as large as  $10^4$  or more. Consequently, computational efficiency is essential for practical usage of the method. For estimation, we exploit the special structure of the GBM to derive computationally efficient Fisher scoring updates to each component of the model. For inference, we develop a novel method for efficiently propagating uncertainty between components of the model. Assuming  $\max\{K^2, L^2, M\} \leq J \leq I$ , our estimation algorithm takes  $O(IJ \max\{K^2, L^2, M^2\})$  time per iteration, and our inference algorithm requires  $O(IJ \max\{K^3, L^3, JM^3\})$  time, making them computationally feasible on large data matrices.

**Numerical stability.** Using a good choice of initialization is crucial for numerical stability. To initialize the estimation algorithm, we analytically solve for values of  $A, B$ ,

and  $C$  to approximate the data matrix and then, for NB-GBMs, we iteratively update  $S$  and  $T$  for a few iterations. Even with a good initialization, optimization methods occasionally diverge. In a large GBM, there are so many parameters that even occasional divergences cause the algorithm to fail with high probability. We reduce the frequency of divergences to be negligible by enforcing a bound on the norm of the optimization steps; see Section 3.

**Enforcing identifiability constraints.** Rather than performing constrained optimization steps, we use a combination of unconstrained optimization steps and likelihood-preserving projections onto the constrained parameter space. Although the construction of likelihood-preserving projections in a GBM is not obvious, we show that they can be efficiently computed using simple linear algebra operations. This optimization-projection approach has a number of advantages; see Section S1.3 for further discussion.

**Dependencies in latent factors.** Optimizing the latent factor term  $UDV^T$  is challenging due to the dependencies among  $U$ ,  $D$ , and  $V$  as well as the orthonormality constraints  $U^T U = I$  and  $V^T V = I$ . Consequently, updating  $U$ ,  $D$ , and  $V$  individually does not seem to work well. To resolve this issue, we relax the dependencies and constraints by defining  $G := UD$  and  $H := VD$ , and updating  $G$ ,  $H$ , and  $D$  separately.

**Prior / regularization.** To improve estimation accuracy in a high-dimensional setting, we place independent normal priors on the entries of  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ ,  $V$ ,  $S$ , and  $T$ , and use *maximum a posteriori* (MAP) estimation, which is equivalent to  $\ell_2$ -penalization/shrinkage for this choice of prior. An additional benefit of using priors is that it improves the numerical stability of the estimation algorithm. See Section S9 for prior details.

### S1.3 Enforcing the GBM identifiability constraints

It might seem preferable to perform unconstrained optimization throughout the estimation algorithm until convergence, and then enforce the identifiability constraints as a postprocessing step. However, in general, this would not converge to a local optimum in the constrained space because the prior does not have the same invariance properties as the

likelihood. Thus, we maintain the constraints throughout the algorithm by applying a projection at each step.

When updating each component ( $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ ,  $V$ ,  $S$ , and  $T$ ), rather than using a constrained optimization step such as equality-constrained Newton’s method (Boyd and Vandenberghe, 2004), we use an unconstrained optimization step followed by a likelihood-preserving projection onto the constrained space.

It is crucial to preserve the likelihood when projecting onto the constrained space, since otherwise the projection might undo all the gains obtained by the unconstrained optimization step — in short, otherwise we might end up “taking one step forward and two steps back.” To this end, we employ likelihood-preserving projections for each component of the GBM. By Theorem 5.4, the likelihood is invariant under these operations and the projected values satisfy the identifiability constraints. The optimization-projection approach has several major advantages.

1. In the likelihood surface, there can be strong dependencies among the parameters within each row of  $A$ ,  $B$ ,  $U$ , and  $V$ , whereas the between-row dependencies are much weaker (specifically, they have zero Fisher cross-information). Thus, it is desirable to optimize each row jointly, however, this is complicated by the fact that the constraints create dependencies between rows. Consequently, using equality-constrained Newton appears to be computationally infeasible since it would require a joint update of each parameter matrix in entirety.
2. Since each optimization-projection step modifies multiple components of the GBM, it effectively performs a joint update on multiple components. For instance, the likelihood-preserving projection for  $A$  also modifies  $C$ , so the optimization-projection step on  $A$  is effectively a joint update to  $A$  and  $C$ . This has the effect of enlarging the constrained space within which each update takes place, improving convergence.
3. For  $U$  and  $V$ , the constrained space is particularly difficult to optimize over since it involves not only within-column linear dependencies ( $X^T U = 0$  and  $Z^T V = 0$ ), but

also quadratic dependencies within and between columns ( $U^T U = I$  and  $V^T V = I$ ). The optimization-projection approach makes it easy to handle these constraints.

4. It is straightforward to perform unconstrained optimization for each component separately, and the projections that we derive turn out to be very easy to apply.

## S2 Additional simulation results and details

We present additional simulation results and details supplementing Section 6.

### Simulating covariates, parameters, and data

Here, we provide the details of how the simulation data are generated. First, the covariates are generated using a copula model as follows. We describe the procedure for the feature covariate matrix  $X \in \mathbb{R}^{I \times K}$ ; the sample covariate matrix  $Z \in \mathbb{R}^{J \times L}$  is generated in the same way but with  $J$  and  $L$  in place of  $I$  and  $K$ . We generate a random covariance matrix  $\Sigma = Q^T Q$  where the entries of  $Q \in \mathbb{R}^{K \times K}$  are  $q_{kk'} \sim \mathcal{N}(0, 1)$  i.i.d., and then we compute the resulting correlation matrix  $\tilde{\Sigma} \in \mathbb{R}^{K \times K}$  by setting  $\tilde{\Sigma}_{kk'} = \Sigma_{kk'} / \sqrt{\Sigma_{kk} \Sigma_{k'k'}}$ . We generate  $(\tilde{x}_{i1}, \dots, \tilde{x}_{iK})^T \sim \mathcal{N}(0, \tilde{\Sigma})$  i.i.d. for  $i = 1, \dots, I$ , and define  $X \in \mathbb{R}^{I \times K}$  by setting  $x_{ik} = h(F^{-1}(\Phi(\tilde{x}_{ik})))$  where  $h(x) = \text{sign}(x) \min\{100, |x|\}$ ,  $\Phi(x)$  is the standard normal CDF, and  $F^{-1}$  is the generalized inverse CDF for the desired marginal distribution, which we take to be  $\mathcal{N}(0, 1)$  for the **Normal** scheme,  $\text{Gamma}(2, \sqrt{2})$  for the **Gamma** scheme, and  $\text{Bernoulli}(1/2)$  for the **Binary** scheme. Finally, we standardize  $X$  by setting  $x_{i1} = 1$  for all  $i$  and centering/scaling so that  $\sum_{i=1}^I x_{ik} = 0$  and  $\frac{1}{I} \sum_{i=1}^I x_{ik}^2 = 1$  for  $k \geq 2$ .

The true parameters  $A_0, B_0, C_0, D_0, U_0, V_0, S_0, T_0$ , and  $\omega_0$  are then generated as follows. First, we generate matrices  $\tilde{A}, \tilde{B}$ , and  $\tilde{C}$  with i.i.d. entries as follows: (**Normal** scheme)  $\tilde{a}_{jk} \sim \mathcal{N}(0, 1/(4K))$ ,  $\tilde{b}_{il} \sim \mathcal{N}(0, 1/(4L))$ , and  $\tilde{c}_{k\ell} \sim \mathcal{N}(0, 1/(KL)) + 3 \mathbf{1}(k = 1, \ell = 1)$ , or (**Gamma** scheme)  $\tilde{a}_{jk} \sim \text{Gamma}(2, 2\sqrt{2K})$ ,  $\tilde{b}_{il} \sim \text{Gamma}(2, 2\sqrt{2L})$ , and  $\tilde{c}_{k\ell} \sim \text{Gamma}(2, \sqrt{2KL}) + 3 \mathbf{1}(k = 1, \ell = 1)$ . These distributions are defined so that the scale of the entries of  $X\tilde{A}^T, \tilde{B}Z^T$ , and  $X\tilde{C}Z^T$  is not affected by  $K$  and  $L$ .

Then we set  $A_0 = \tilde{A} - Z(Z^+ \tilde{A})$ ,  $B_0 = \tilde{B} - X(X^+ \tilde{B})$ , and  $C_0 = \tilde{C}$ . Next, we set  $U_0 = \tilde{U} - X(X^+ \tilde{U})$  and  $V_0 = \tilde{V} - Z(Z^+ \tilde{V})$  where  $\tilde{U} \in \mathbb{R}^{I \times M}$  and  $\tilde{V} \in \mathbb{R}^{J \times M}$  are sampled uniformly from their respective Stiefel manifolds, that is, uniformly subject to  $\tilde{U}^T \tilde{U} = I$  and  $\tilde{V}^T \tilde{V} = I$ . The diagonal entries of  $D_0$  are evenly spaced from  $\sqrt{I} + \sqrt{J}$  to  $2(\sqrt{I} + \sqrt{J})$ ; this scaling is motivated by the Marchenko–Pastur law for the distribution of singular values of  $I \times J$  random matrices (Marchenko and Pastur, 1967). For the log-dispersion parameters ( $S_0$ ,  $T_0$ , and  $\omega_0$ ), we generate  $\tilde{s}_i, \tilde{t}_j \sim \mathcal{N}(0, 1)$  i.i.d., and set  $\omega_0 = -2.3$ ,  $s_{0i} = \tilde{s}_i - \log(\frac{1}{I} \sum_{i=1}^I \exp(\tilde{s}_i))$ , and  $t_{0j} = \tilde{t}_j - \log(\frac{1}{J} \sum_{j=1}^J \exp(\tilde{t}_j))$ ,

Given the true parameters and covariates, the data matrix  $\mathbf{Y} \in \{0, 1, 2, \dots\}^{I \times J}$  is generated as follows. We compute the mean matrix  $\mu_0 := g^{-1}(XA_0^T + B_0Z^T + XC_0Z^T + U_0D_0V_0^T)$ , where the inverse link function  $g^{-1}(x) = e^x$  is applied element-wise, and we compute the inverse dispersions  $r_{0ij} := \exp(-s_{0i} - t_{0j} - \omega_0)$ . Then we sample  $Y_{ij} \sim \mathcal{D}(\mu_{0ij}, r_{0ij})$  where  $\mathcal{D}(\mu, r) = \text{NegBin}(\mu, r)$  in the NB scheme,  $\mathcal{D}(\mu, r) = \text{LNP}(\mu, \log(1/r + 1))$  in the LNP scheme,  $\mathcal{D}(\mu, r) = \text{Poisson}(\mu)$  in the Poisson scheme, or  $\mathcal{D}(\mu, r) = \text{Geometric}(1/(\mu + 1))$  in the Geometric scheme, so that  $\mathbb{E}(Y_{ij}) = \mu_{0ij}$  in each case. Here, for  $y \in \{0, 1, 2, \dots\}$ , Geometric( $p$ ) has p.m.f.  $f(y) = (1 - p)^y p$ , whereas LNP( $\mu, \sigma^2$ ) has p.m.f.

$$f(y) = \int \text{Poisson}(y|\lambda) \text{LogNormal}(\lambda \mid \log(\mu) - \frac{1}{2}\sigma^2, \sigma^2) d\lambda \quad (\text{S2.1})$$

for  $\mu > 0$  and  $\sigma^2 > 0$ . These outcome distributions are defined so that in each case, if  $Y \sim \mathcal{D}(\mu, r)$  then  $\mathbb{E}(Y) = \mu$ . Further, in the LNP case,  $\text{Var}(Y) = \mu + \mu^2/r$ , so the interpretation of  $r$  is the same as in the NB case.

### Consistency and statistical efficiency – Details on Section 6.1

In these simulations, to accurately measure the trend with increasing  $I$ , we generate the covariates, true parameters, and data with  $I = 10000$  and project them onto the lower-dimensional spaces for smaller  $I$  values; for  $\mathbf{Y}$ ,  $X$ , and  $S_0$  this projection simply consists of taking the first  $I$  rows/entries,  $Z$  and  $T_0$  are unaffected by the projection, and  $A_0$ ,  $B_0$ ,

$C_0$ ,  $D_0$ ,  $U_0$ , and  $V_0$  are projected by matching the first  $I$  rows of the mean matrix  $\mu_0$ .

We use the relative MSE rather than the MSE to facilitate interpretability, since this puts the errors on a common scale that does not depend on the magnitude of the parameters. For instance, the relative MSE for  $A$  is defined as

$$\text{MSE}_{\text{rel}}(A, A_0) = \frac{\sum_{j=1}^J \sum_{k=1}^K |a_{jk} - a_{0jk}|^2}{\sum_{j=1}^J \sum_{k=1}^K |a_{0jk}|^2}$$

where  $A$  is the estimate and  $A_0$  is the true value.

Figure S1 shows the relative MSE plots for all GBM parameter components. For  $A$ ,  $C$ ,  $V$ , and  $T$ , the relative MSE appears to be decreasing to zero. The trend for  $D$  and  $\omega$  is suggestive but not as clear, making it difficult to gauge whether  $D$  and  $\omega$  are likely to be consistent based on these experiments. The relative MSEs for  $B$ ,  $U$ , and  $S$  are small but do not appear to be going to zero as  $I \rightarrow \infty$  with  $J$  fixed; this is expected since for these parameters the amount of data informing each univariate entry is fixed.

## Accuracy of standard errors – Details on Section 6.2

To estimate the actual coverage at every target coverage level from 0% to 100%, we use the fact that the actual coverage of a  $100 \times (1 - \alpha)\%$  interval for some parameter  $\theta$  can be written as

$$\mathbb{P}(|\hat{\theta} - \theta_0| < z_{\alpha/2} \hat{\text{se}}) = \mathbb{P}\left(1 - 2(1 - \Phi(|\hat{\theta} - \theta_0|/\hat{\text{se}})) < 1 - \alpha\right)$$

where  $z_{\alpha/2} = \Phi^{-1}(1 - \alpha/2)$  and  $\Phi(x)$  is the  $\mathcal{N}(0, 1)$  CDF. Thus, since  $1 - \alpha$  is the target coverage, the curve of actual coverage versus target coverage is simply the CDF of the random variable  $1 - 2(1 - \Phi(|\hat{\theta} - \theta_0|/\hat{\text{se}}))$ . The plots in Figure 5 are empirical CDFs of this random variable, aggregating across all entries of each parameter matrix/vector.



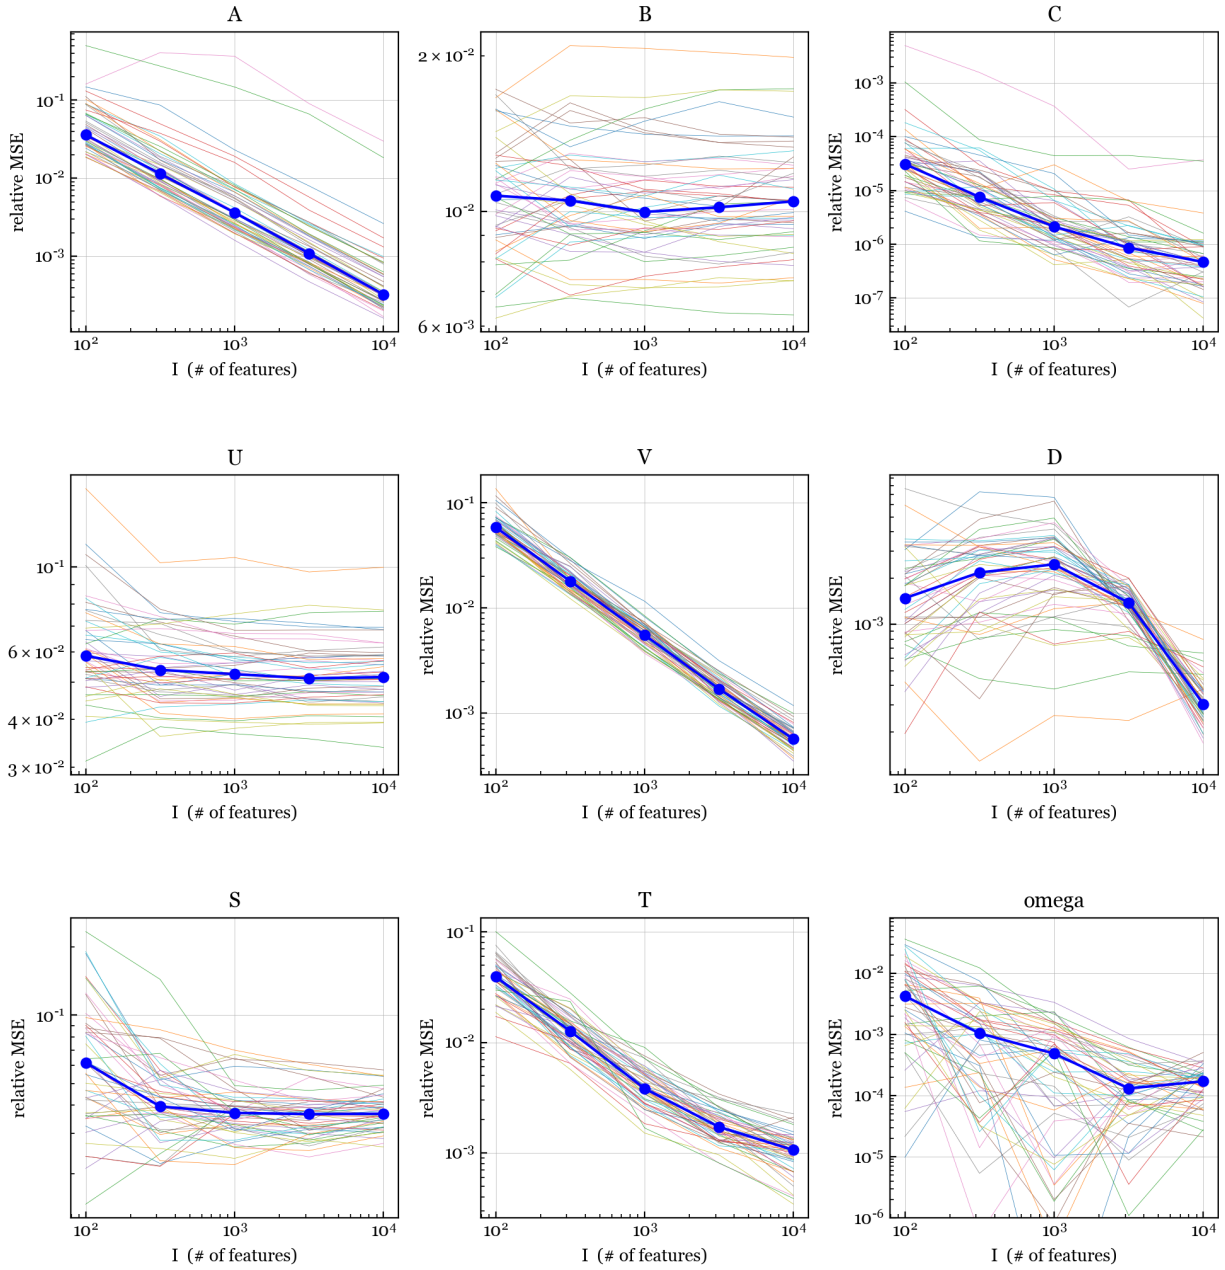


Figure S1: Relative mean-squared error between estimated and true parameter values. Each plot contains 50 thin lines, one for each run, along with the median over the 50 runs (thick blue line).




Figure S2: Log posterior density (plus constant) versus iteration for the estimation algorithm.

### Algorithm convergence – Details on Section 6.3

Figure S2 shows the results of the algorithm convergence experiments described in Section 6.3. For each  $I$ , the plot shows 25 curves of the log-likelihood+log-prior (plus a constant) versus iteration number, one for each of the 25 simulation runs. For visual interpretability, we add a constant to each curve such that the final value after iteration 50 is equal to zero. Based on our experiments, the estimation algorithm converges rapidly.

### Robustness to the outcome distribution – Details on Section 6.4

Figures S3, S4, and S5 show the results of the robustness experiments in Section 6.4. With LNP outcomes (Figure S3), the results are very similar to when the true outcome distribution is actually NB (Figures S1 and 5). With Poisson outcomes (Figure S4), the estimation accuracy is even better than with NB outcomes, presumably because there is less variability and the Poisson distribution is a limiting case of NB when the dispersion goes to zero. The coverage with Poisson outcomes is essentially the same as NB, slightly better for some parameters and slightly worse for others. With Geometric outcomes (Figure S5), the same parameters appear to be consistently estimated, however, the estimation accuracy in terms of relative MSE is worse by roughly a factor of 10. Compared to NB outcomes, the coverage with Geometric outcomes is similar for some parameters and slightly worse for others. Overall, it appears that the estimation and inference algorithms are quite robust to misspecification of the outcome distribution.

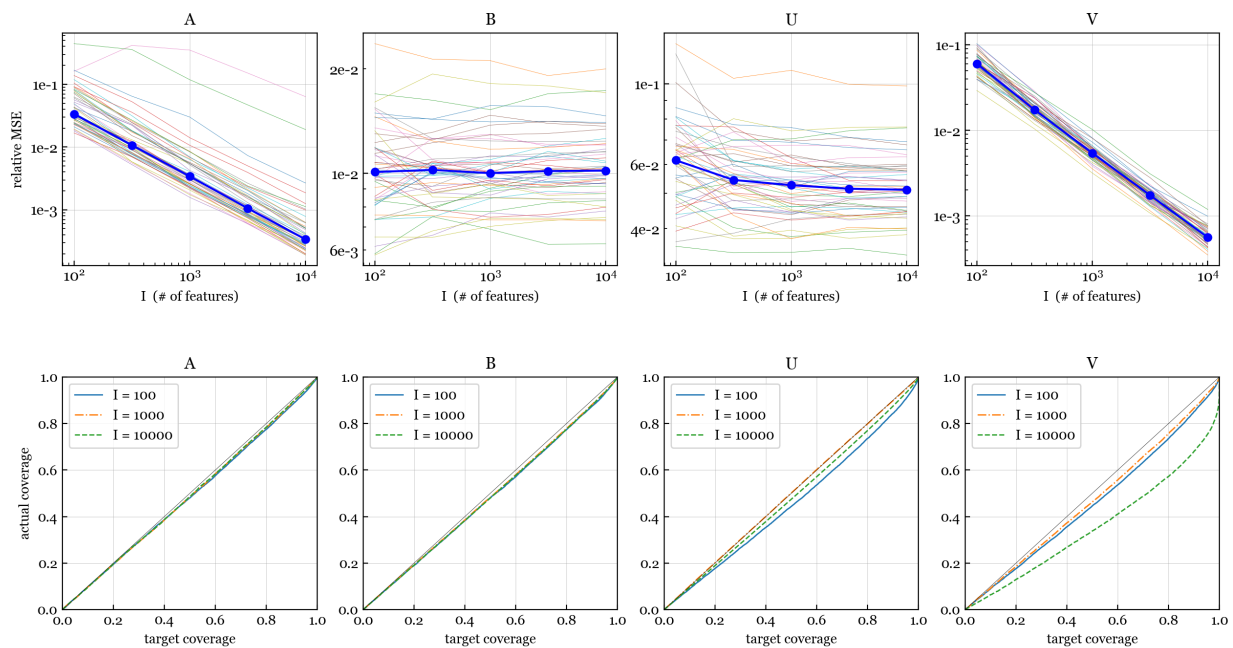


Figure S3: Robustness to outcome:  $K = 4, L = 2, M = 3, \text{LNP/Normal/Normal}$ .

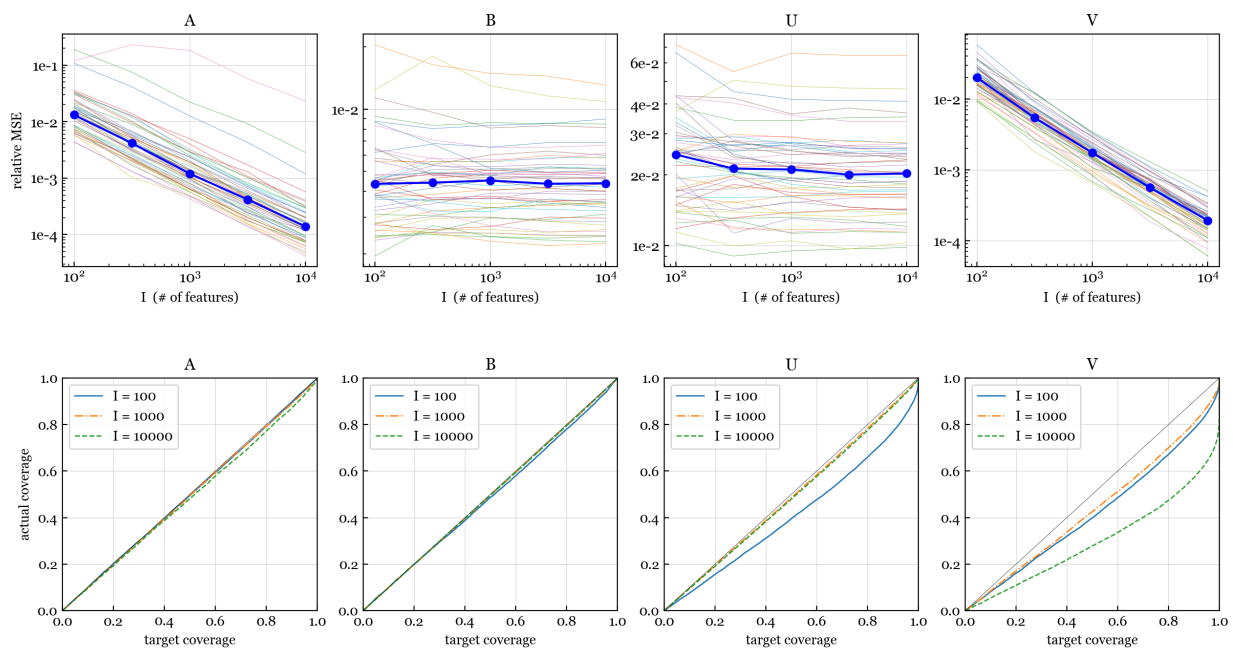


Figure S4: Robustness to outcome:  $K = 4, L = 2, M = 3, \text{Poisson/Normal/Normal}$ .

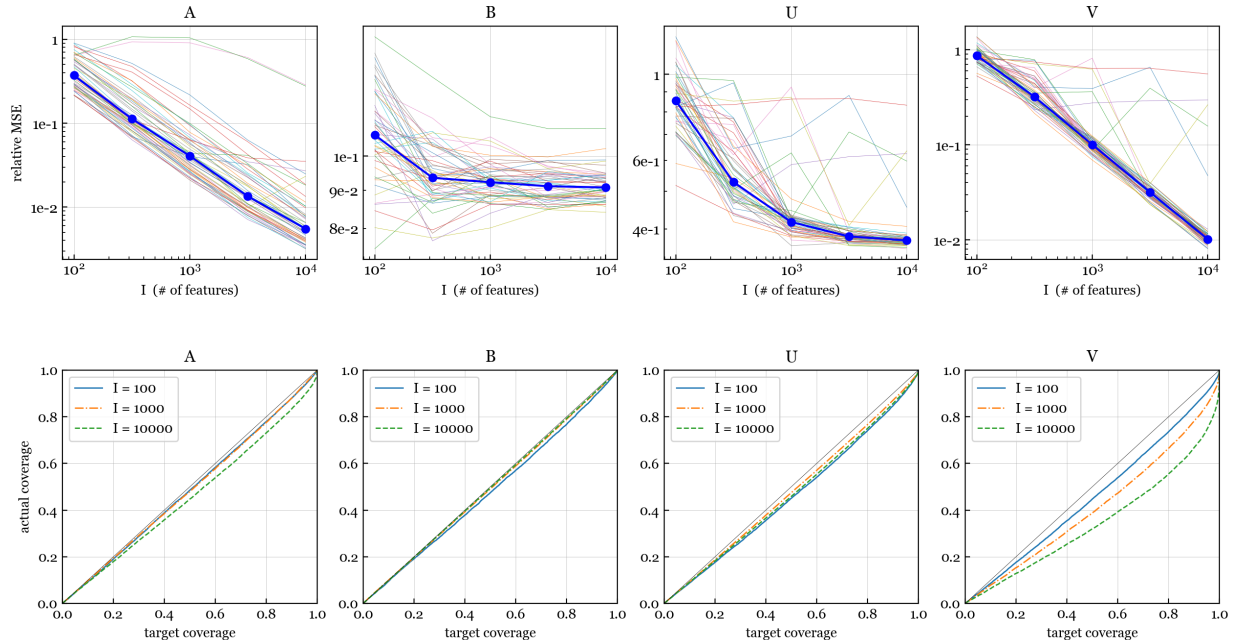


Figure S5: Robustness to outcome:  $K = 4$ ,  $L = 2$ ,  $M = 3$ , Geometric/Normal/Normal.

### Convergence to global optimum

Based on the consistency experiments reported in Figure S1, it seems unlikely that the algorithm is getting stuck in a suboptimal local mode. To more directly assess whether the algorithm is converging to a global optimum, we compare the estimates obtained when (a) initializing the estimation algorithm to the true parameter values versus (b) initializing using our proposed approach. Table S1 shows that the difference between these estimates is negligible. In detail, we generate 50 data matrices using the NB/Normal/Normal scheme with  $I = 1000$ ,  $J = 100$ ,  $K = 4$ ,  $L = 2$ , and  $M = 3$ . For each data matrix, we run the estimation algorithm with initialization approaches (a) and (b) and compute the relative MSE between the two resulting estimates. In Table S1, for each parameter we report the largest observed value of the relative MSE over these 50 simulation runs. The maximum relative MSE is extremely small in each case, meaning that initializing at the true parameter values yields nearly identical estimates as initializing with our proposed approach. This suggests that our estimation algorithm is not getting stuck in a suboptimal local mode.

Table S1: Initialization error. Maximum relative MSE between estimates when initializing at the true values versus the proposed initialization approach, over 50 runs with  $I = 1000$  and  $J = 100$ .

$A$	$B$	$C$	$D$	$U$	$V$	$S$	$T$	$\omega$
$2 \times 10^{-7}$	$9 \times 10^{-7}$	$7 \times 10^{-9}$	$1 \times 10^{-8}$	$4 \times 10^{-6}$	$3 \times 10^{-7}$	$3 \times 10^{-7}$	$4 \times 10^{-8}$	$2 \times 10^{-9}$

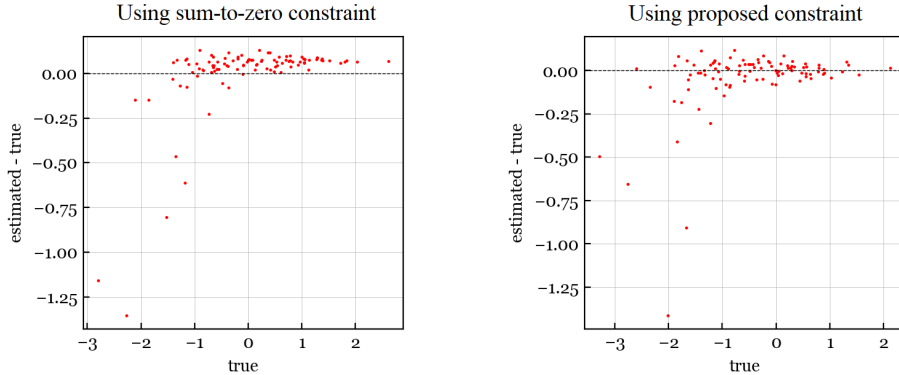


Figure S6: Scatterplots of  $\hat{t}_j - t_{0j}$  versus  $t_{0j}$  for a typical simulated data matrix.

### Choice of identifiability constraint on log-dispersions

The choice of identifiability constraint on the log-dispersions  $S$  and  $T$  has a significant effect on estimation accuracy. Perhaps the most obvious choice would be sum-to-zero constraints:  $\sum_i s_i = 0$  and  $\sum_j t_j = 0$ . However, it turns out that this leads to poor performance in terms of estimation accuracy. The reason is that low log-dispersions (for instance, around  $t_j \approx -2$  or less) are difficult to estimate accurately, and noisy estimates of a few low values of  $t_j$  can have an undue effect on  $\sum_j t_j$ , causing significant error in the rest of the  $t_j$ 's. Figure S6 (left panel) illustrates the issue in a simulated example using the NB/Normal/Normal scheme with  $I = 10000$ ,  $J = 100$ ,  $K = 4$ ,  $L = 2$ , and  $M = 3$ . Here, the sum-to-zero constraints are used on both the true parameters and the estimates.

In our proposed algorithm, we instead constrain  $\frac{1}{I} \sum_i e^{s_i} = 1$  and  $\frac{1}{J} \sum_j e^{t_j} = 1$ , which has the effect of downweighting the low values (which are harder to estimate) and upweighting the large values (which are easier to estimate). As illustrated in Figure S6 (right panel), this effectively resolves the issue. Here, we run the same algorithm on the same data matrix, but instead use our proposed constraints. We can see that the undesirable shift is effectively removed.

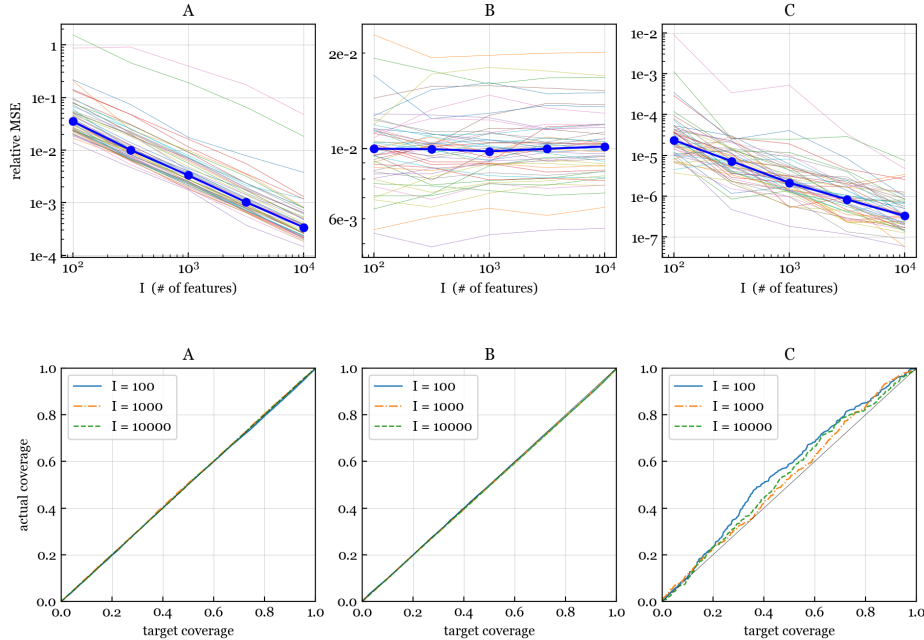


Figure S7: Special case of no latent factors:  $K = 4$ ,  $L = 2$ ,  $M = 0$ , NB/Normal/Normal.

### Special case of no latent factors

If the number of latent factors is zero, that is,  $M = 0$ , then the model is no longer “bilinear” since the  $UDV^T$  term is no longer present. The resulting model with  $g(\mathbb{E}(\mathbf{Y})) = XA^T + BZ^T + XCZ^T$  can be viewed as a standard GLM via the vectorization approach discussed in Section S1.2, however, computation using standard GLM methods becomes intractable when  $I$  or  $J$  is large. Thus, our methods provide a computationally tractable approach to estimation and inference in this quite general class of GLMs for matrix data. In Figure S7, we present a subset of simulation results using the NB/Normal/Normal simulation scheme with  $J = 100$ ,  $K = 4$ ,  $L = 2$ , and  $M = 0$ . As expected, the estimates for  $A$  and  $C$  appear to be consistent and rapidly converging to the true values. Further, the coverage is nearly perfect for  $A$  and  $B$ , and possibly slightly conservative for  $C$ .

### Effect of small sample size

In applications such as gene expression analysis, sometimes a very small number of samples are available. To explore this regime, we run simulations using the NB/Normal/Normal

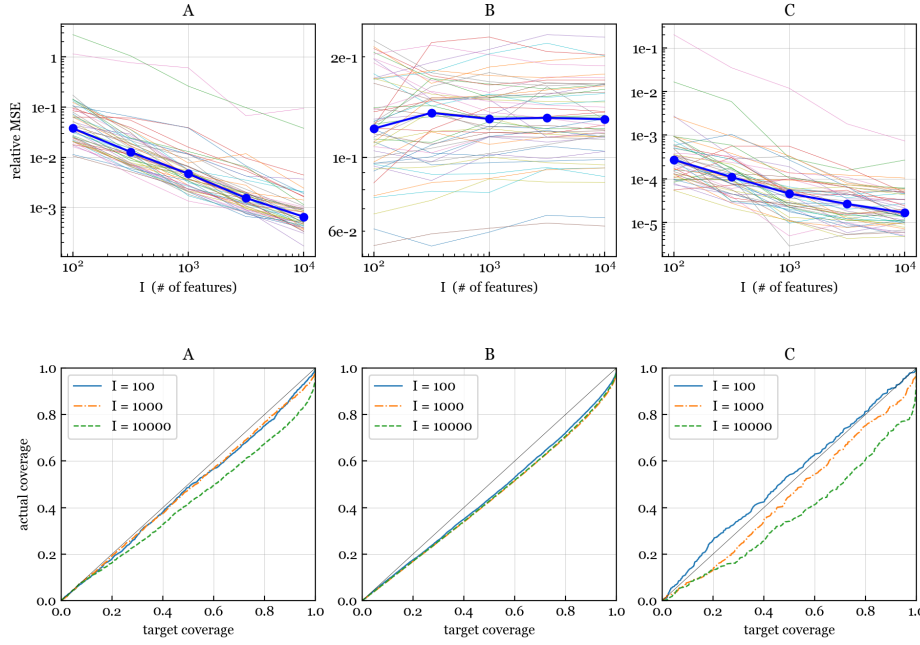


Figure S8: Effect of small sample size:  $J = 10$ ,  $K = 4$ ,  $L = 2$ ,  $M = 0$ , NB/Normal/Normal.

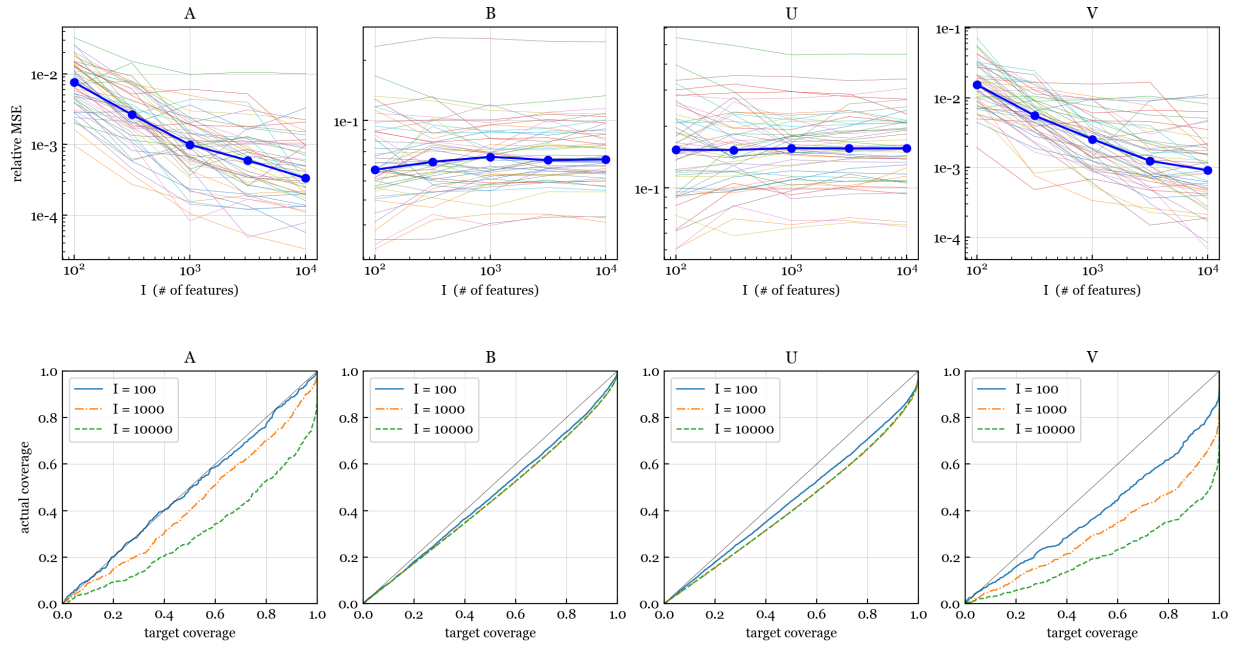


Figure S9: Effect of small sample size:  $J = 10$ ,  $K = 1$ ,  $L = 1$ ,  $M = 1$ , NB/Normal/Normal.

scheme with  $J = 10$  and (a)  $K = 4, L = 2, M = 0$ , and (b)  $K = 1, L = 1, M = 1$ . As Figures S8 and S9 show, consistency for the same parameters appears to hold, although in some cases the rate is less clear than when  $J = 100$  (Figures S7 and S10). Meanwhile, the coverage performance is noticeably worse when  $J = 10$  than when  $J = 100$ .

### **Effect of varying latent dimension**

The “bilinear” latent factorization term  $UDV^T$  is more challenging in terms of estimation and inference compared to the linear terms  $XA^T, BZ^T$ , and  $XCZ^T$ . To assess the effect of the latent dimension  $M$  on consistency, efficiency, and accuracy of standard errors, we run the simulations using: (a)  $K = 1, L = 1, M = 1$ , and (b)  $K = 1, L = 1$ , and  $M = 5$ . This is similar to the GAMMI model (Van Eeuwijk, 1995), but also includes  $S, T$ , and  $\omega$ . We see in Figure S10 that when  $M = 1$ , the estimates of  $A$  and  $V$  are converging to the true values as expected, and the coverage for  $A, B, U$  and  $V$  is excellent, except that the coverage for  $V$  degrades slightly when  $I = 10000$ . When the latent dimension is increased to  $M = 5$ , Figure S11 shows that the performance is similar, except that the coverage for  $V$  degrades more severely when  $I = 10000$ .

Based on these and other experiments, we find that as long as the entries of  $D$  are not too small and are sufficiently well-spaced, the performance remains good. The magnitude of each entry of  $D$  controls the strength of the signal for the corresponding latent factor; thus, a small entry in  $D$  makes it difficult to estimate the corresponding columns of  $U$  and  $V$ . Meanwhile, if two entries of  $D$  are similar in magnitude, then there is near non-identifiability of  $U$  and  $V$  with respect to those two latent factors due to rotational invariance, leading to a loss of performance in terms of statistical efficiency and coverage.

### **Varying the distribution of covariates and parameters**

As usual in regression, since the model is defined conditionally on the covariates, we would expect the results to be robust to the distribution of the covariates. To verify that this is indeed the case, we run the simulations with the following simulation schemes



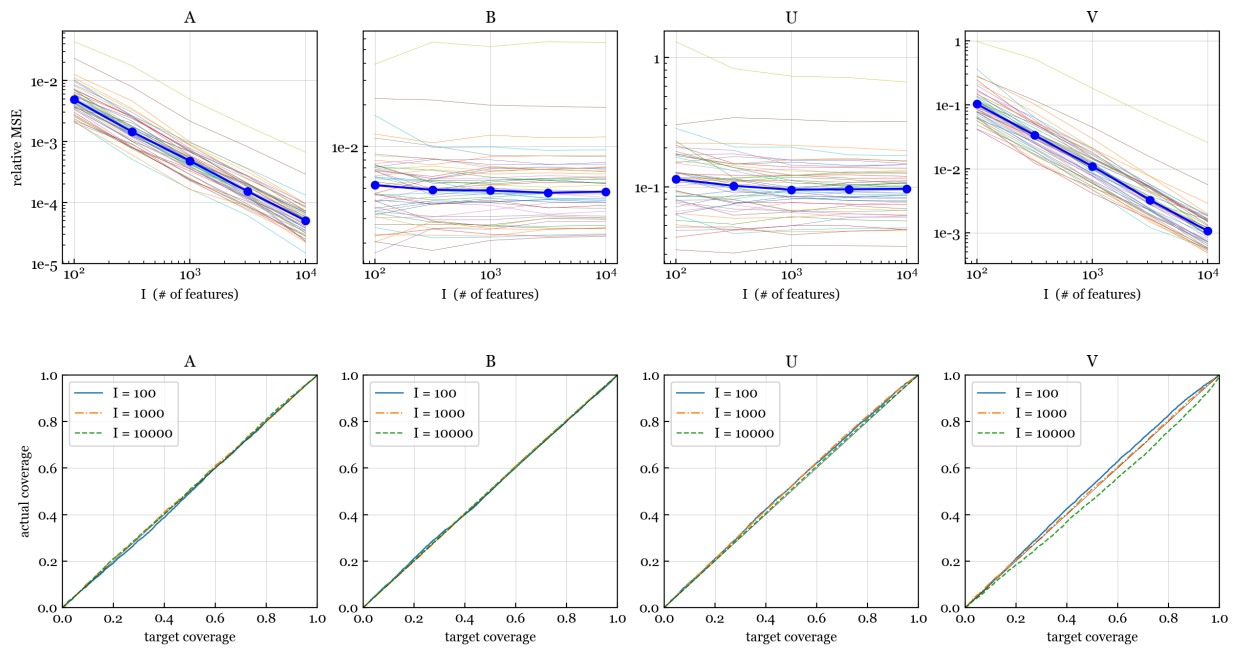


Figure S10: Effect of varying latent dimension:  $K = 1, L = 1, M = 1, \text{NB/Normal/Normal}$ .

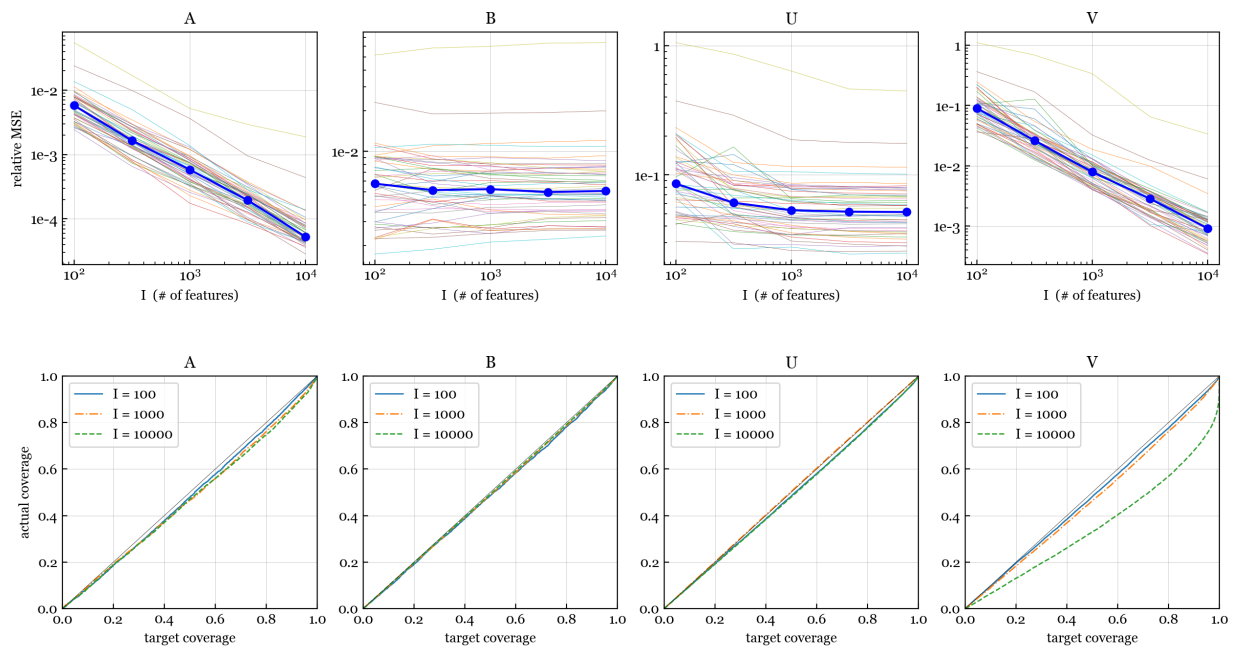


Figure S11: Effect of varying latent dimension:  $K = 1, L = 1, M = 5, \text{NB/Normal/Normal}$ .

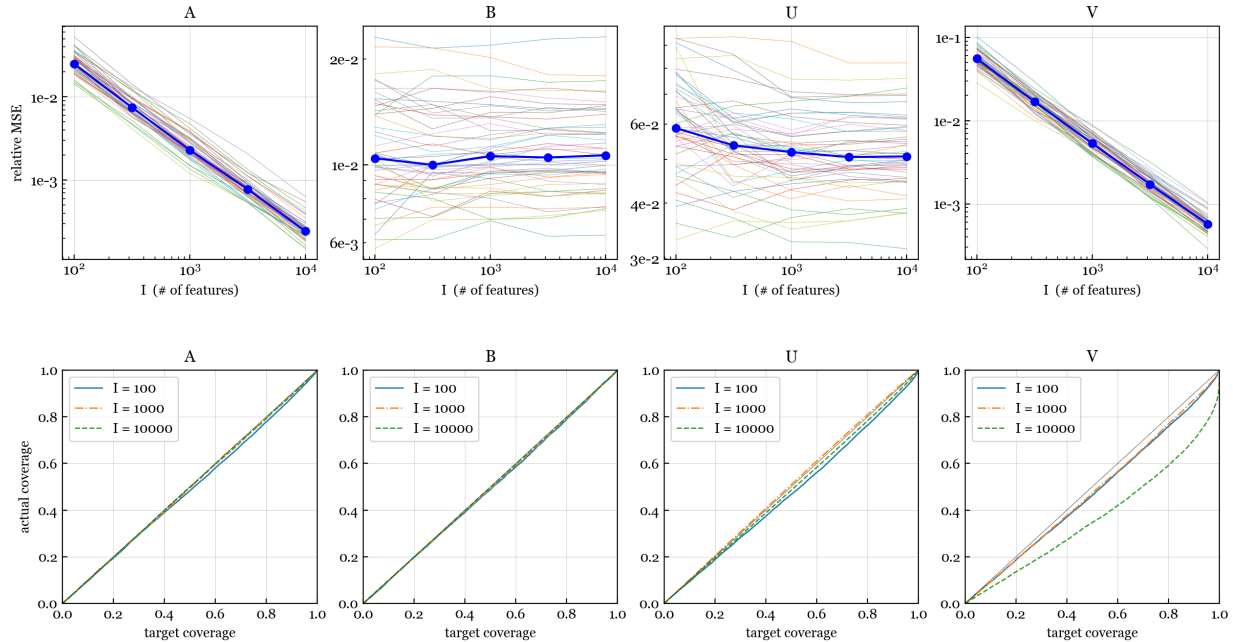


Figure S12: Varying the covariate distribution:  $K = 4$ ,  $L = 2$ ,  $M = 3$ , NB/Bernoulli/Normal.

as defined above: (a) NB/Bernoulli/Normal (Bernoulli-distributed covariates) and (b) NB/Gamma/Normal (Gamma-distributed covariates), with  $J = 100$ ,  $K = 4$ ,  $L = 2$ , and  $M = 3$ ; see Figures S12 and S13. Similarly, although we have a prior on the parameters, we would not expect the results to be highly sensitive to the distribution of the true parameters. To check this, we also run the simulations with the NB/Normal/Gamma scheme (Gamma-distributed true parameters) with  $J = 100$ ,  $K = 4$ ,  $L = 2$ , and  $M = 3$ ; see Figure S14. The results are nearly identical to the case of normally distributed covariates and true parameters.

## S3 Additional application results and details

### S3.1 Gene expression application – Details on Section 7

Here we provide additional results and details on the gene expression applications in Section 7. For the Pickrell data, Figure S15 shows the PCA plots based on (a) log-transformed TPMs, specifically,  $\log(\text{TPM}_{ij} + 1)$ , and (b) the variance stabilizing transform (VST) method

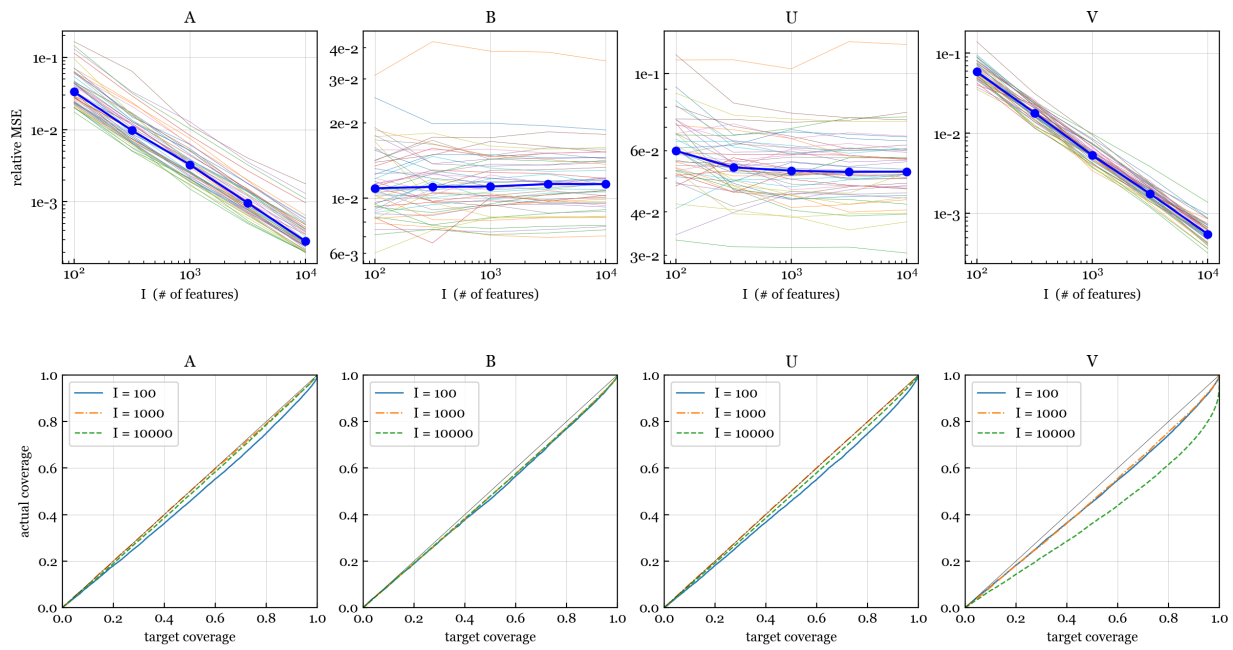


Figure S13: Varying the covariate distribution:  $K = 4$ ,  $L = 2$ ,  $M = 3$ , NB/Gamma/Normal.

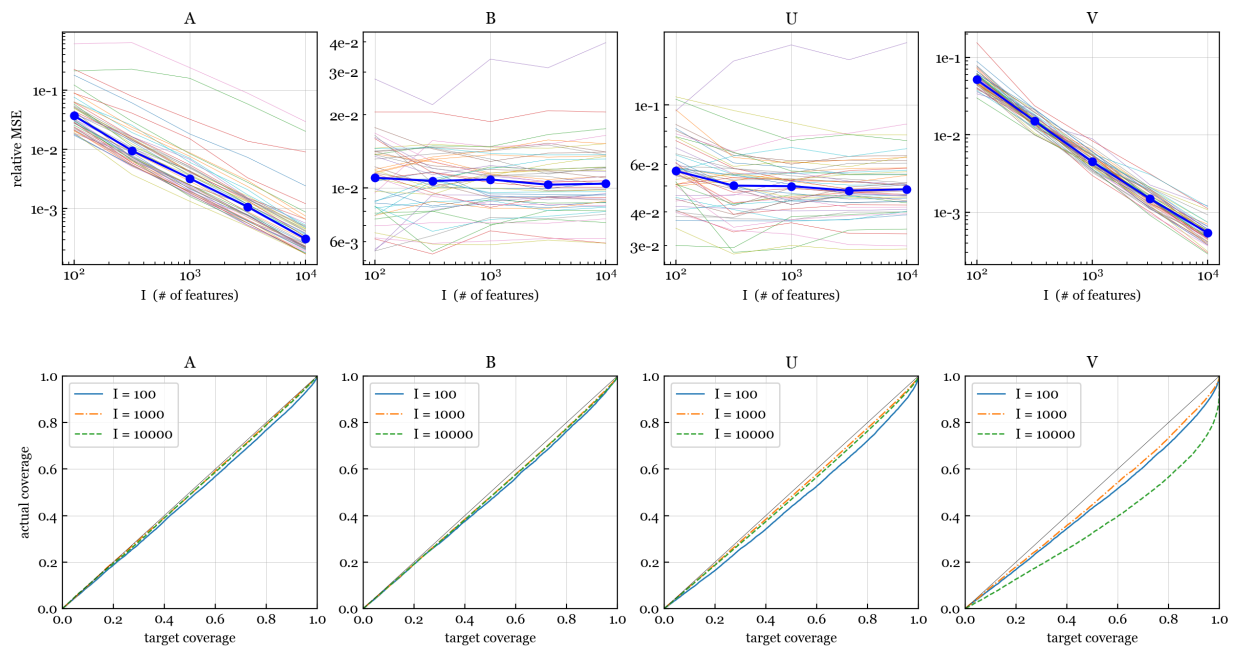


Figure S14: Varying the parameter distribution:  $K = 4$ ,  $L = 2$ ,  $M = 3$ , NB/Normal/Gamma.

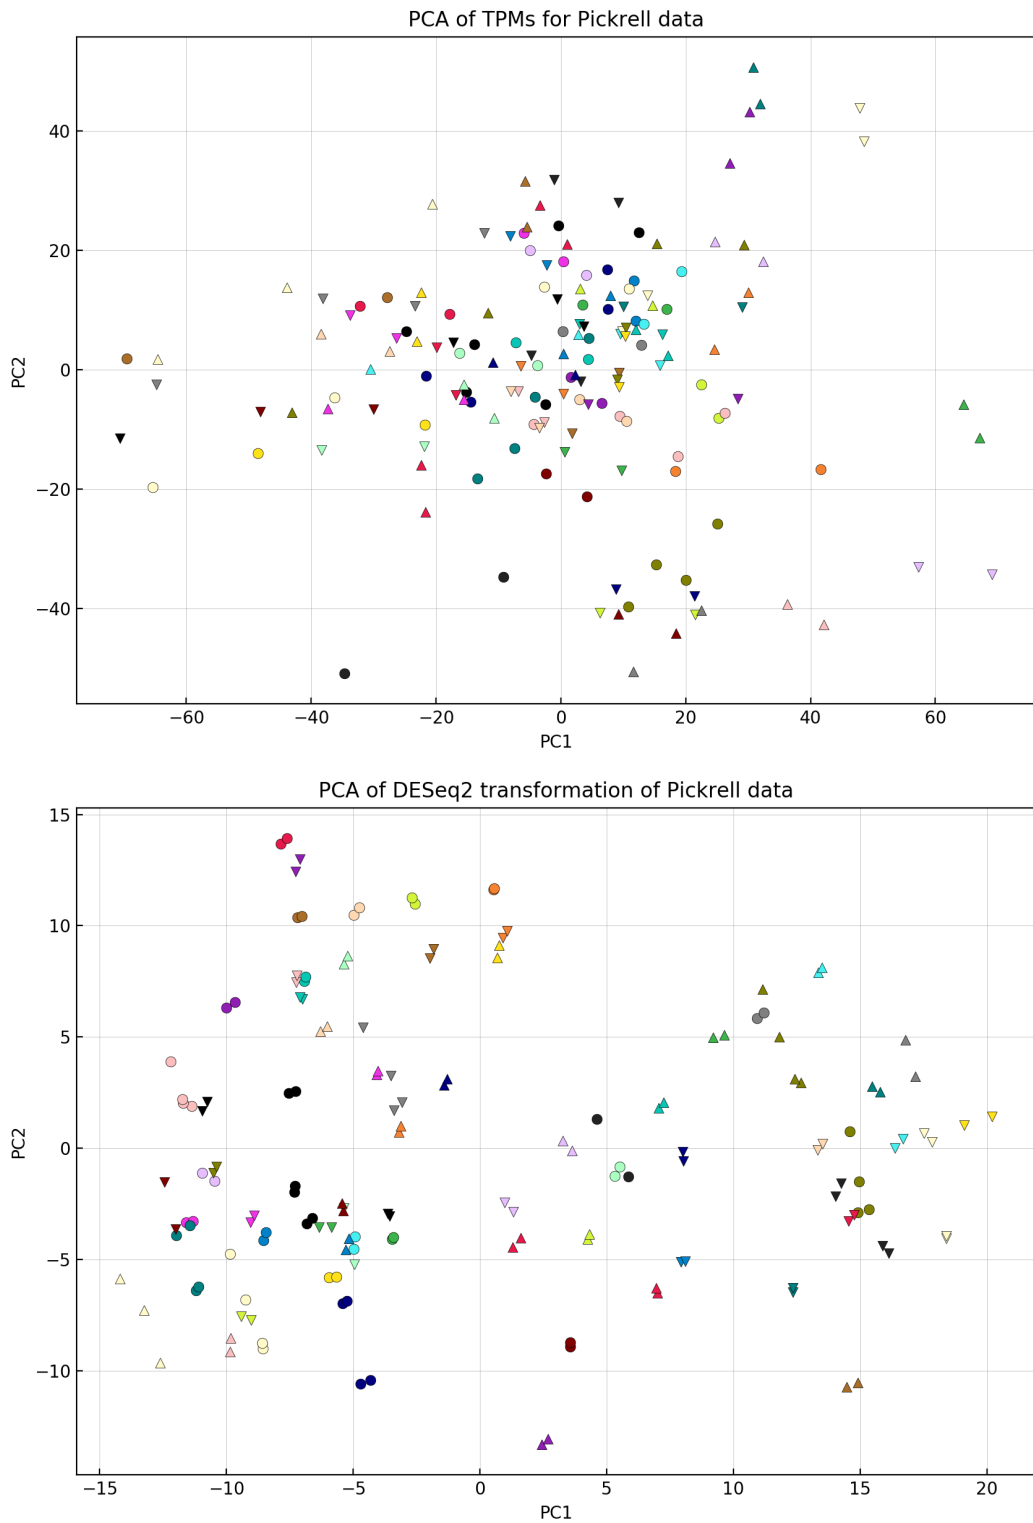


Figure S15: PCA of the Pickrell data using: (top) log-transformed TPMs and (bottom) the VST method in the DESeq2 software package, including the GC bias adjustment from CQN. Compare with the GBM approach in Figure 8.

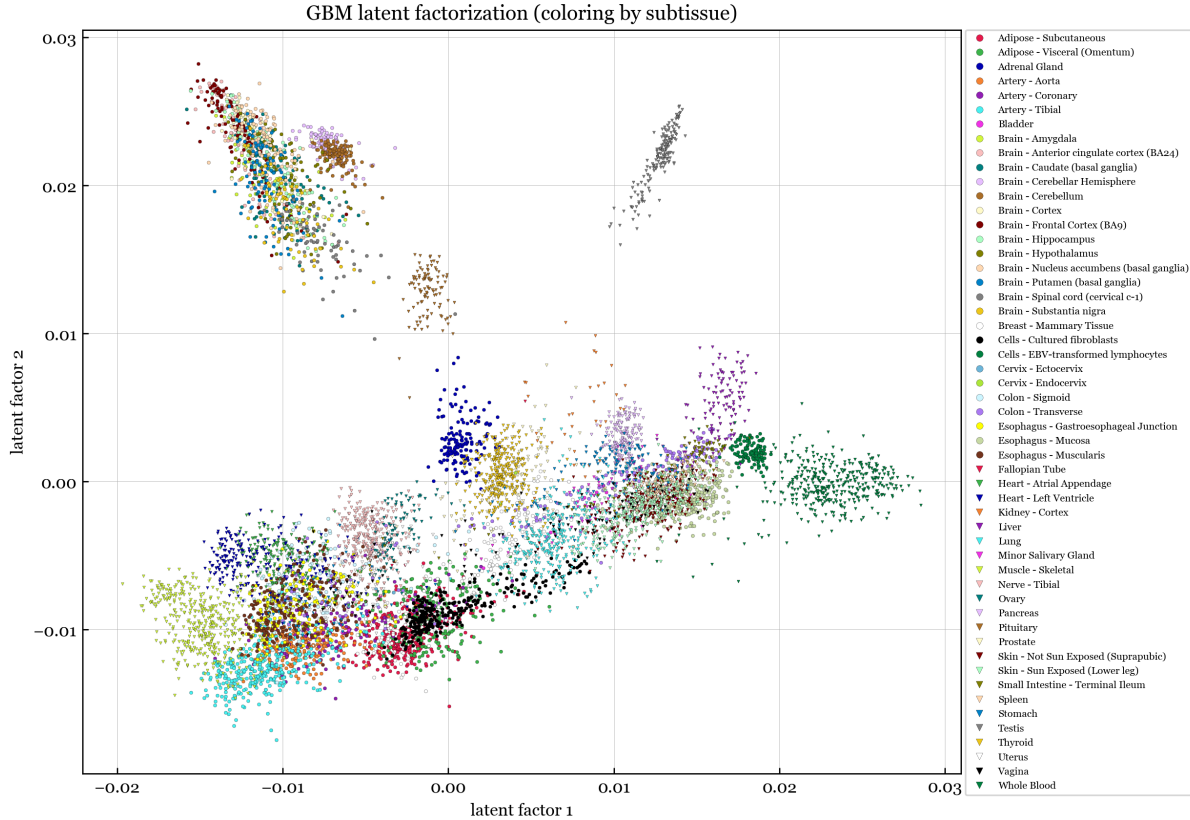


Figure S16: Visualization of GTEx data using NB-GBM latent factors, adjusting for covariates. Each dot represents one of the 8,551 samples, and the color indicates the subtissue type.

in the DESeq2 software package, using the GC adjustment from CQN. For the GTEx data, Figure S16 shows the latent factors ( $v_{j2}$  versus  $v_{j1}$ ) as in Figure 9, but coloring the points according to tissue subtype instead tissue type. We see that the samples tend to fall into clusters according to tissue subtype, further resolving the clustering in Figure 9. Figure S17 shows a PCA plot of the log TPMs of the GTEx data, which is not nearly as clear as Figure 9 in terms of tissue type clustering. For the analysis of aging-related genes in the Heart-LV subtissue using the GTEx data (Section 7.2), Tables S2 and S3 show the top 20 enriched GO terms in the Biological Process and Cellular Component categories.




Figure S17: PCA of the GTEx data using log-transformed TPMs, specifically,  $\log(\text{TPM}_{ij} + 1)$ . Each dot represents one of the 8,551 samples, and the color indicates the tissue type.

Table S2: Top GO terms (Biological Process) for age-related expression in Heart-LV.

GO term ID	Description	Count	p-value	Benjamini
GO:0098609	cell-cell adhesion	48	5.1e-12	1.5e-08
GO:0006418	tRNA aminoacylation for protein translation	16	1.4e-09	2.0e-06
GO:0006099	tricarboxylic acid cycle	12	3.7e-07	3.6e-04
GO:1904871	positive regulation of protein localization to Cajal body	7	1.1e-06	6.1e-04
GO:1904851	positive regulation of establishment of protein localization to telomere	7	1.1e-06	6.1e-04
GO:0006607	NLS-bearing protein import into nucleus	10	1.3e-06	6.2e-04
GO:0006914	autophagy	22	1.8e-05	7.6e-03
GO:0016192	vesicle-mediated transport	24	2.6e-05	8.3e-03
GO:0006511	ubiquitin-dependent protein catabolic process	24	2.6e-05	8.3e-03
GO:0006888	ER to Golgi vesicle-mediated transport	24	3.5e-05	1.0e-02
GO:0006886	intracellular protein transport	31	4.3e-05	1.1e-02
GO:1904874	positive regulation of telomerase RNA localization to Cajal body	7	8.3e-05	2.0e-02
GO:0006090	pyruvate metabolic process	8	9.6e-05	2.1e-02
GO:0070125	mitochondrial translational elongation	16	1.1e-04	2.2e-02
GO:0006446	regulation of translational initiation	10	1.5e-04	2.8e-02
GO:0043039	tRNA aminoacylation	5	1.6e-04	3.0e-02
GO:0018107	peptidyl-threonine phosphorylation	10	2.9e-04	4.9e-02
GO:0000462	maturation of SSU-rRNA from tricistronic rRNA transcript	9	3.3e-04	5.4e-02
GO:0006610	ribosomal protein import into nucleus	5	3.7e-04	5.6e-02
GO:0016236	macroautophagy	14	4.0e-04	5.9e-02

Table S3: Top GO terms (Cellular Component) for age-related expression in Heart-LV.

GO term ID	Description	Count	p-value	Benjamini
GO:0016020	membrane	220	9.8e-21	3.7e-18
GO:0005739	mitochondrion	157	1.2e-20	3.7e-18
GO:0070062	extracellular exosome	242	4.3e-16	9.1e-14
GO:0005829	cytosol	282	1.0e-15	1.6e-13
GO:0005913	cell-cell adherens junction	57	9.5e-15	1.2e-12
GO:0005737	cytoplasm	380	2.3e-13	2.4e-11
GO:0043209	myelin sheath	36	4.7e-13	4.2e-11
GO:0005759	mitochondrial matrix	47	5.7e-09	4.5e-07
GO:0005654	nucleoplasm	217	1.1e-08	7.8e-07
GO:0000502	proteasome complex	18	1.4e-08	8.0e-07
GO:0005743	mitochondrial inner membrane	56	1.4e-08	8.0e-07
GO:0042645	mitochondrial nucleoid	14	3.5e-07	1.8e-05
GO:0014704	intercalated disc	14	8.5e-07	4.2e-05
GO:0005832	chaperonin-containing T-complex	7	2.5e-06	1.1e-04
GO:0005643	nuclear pore	16	5.2e-06	2.2e-04
GO:0043231	intracellular membrane-bounded organelle	55	2.7e-05	1.1e-03
GO:0002199	zona pellucida receptor complex	6	2.9e-05	1.1e-03
GO:0043034	costamere	8	5.4e-05	1.9e-03
GO:0043234	protein complex	42	7.8e-05	2.6e-03
GO:0045254	pyruvate dehydrogenase complex	5	1.5e-04	4.6e-03

### S3.2 Cancer genomics application – Details on Section 8

**Data acquisition and preprocessing.** We downloaded BAM files for the 326 CCLE whole-exome samples from the Genomic Data Commons (GDC) Legacy Archive of the National Cancer Institute (<https://portal.gdc.cancer.gov/legacy-archive/>), using the GDC Data Transfer Tool v1.6.0. Using the PreprocessIntervals tool from the Genome Analysis Toolkit (GATK) v4.1.8.1 (<https://github.com/broadinstitute/gatk/>) running on Java v1.8, we preprocessed the CCLE exome target region interval list to pad the intervals by 250 base pairs on either side (options: `padding 250`, `bin-length 0`, `interval-merging-rule OVERLAPPING_ONLY`). Then, to convert each BAM file to a vector of counts, we counted the number of reads in each target region using the CollectReadCounts tool from GATK (options: `interval-merging-rule OVERLAPPING_ONLY`). For analysis, we included all target regions in chromosomes 1–22 that have nonzero median across the 326 samples.

**De-segmenting the training samples.** We randomly selected half of the samples to use as a training set, and the rest were used as a test set to evaluate performance.

To be able to treat the training samples as “pseudo-normal” (non-cancer) samples, we de-segment them as follows. We first compute a rough estimate of copy ratio, defined as  $\rho_{ij} := \tilde{Y}_{ij}/(\alpha_i\beta_j)$  where  $\alpha_i = \frac{1}{J} \sum_{j=1}^J \tilde{Y}_{ij}$ ,  $\beta_j = \frac{1}{I} \sum_{i=1}^I \tilde{Y}_{ij}/\alpha_i$ , and  $\tilde{Y}_{ij} = Y_{ij} + 1/8$ ; here,  $1/8$  is a pseudocount that avoids issues when taking logs. For each sample  $j$ , we then run a standard binary segmentation algorithm (Killick et al., 2012, Eqn 2) on  $\log \rho_{ij}$  to detect changepoints. For binary segmentation, we use cost function  $\mathcal{C}(x_{1:n}) = -(\frac{1}{\sqrt{n}} \sum_{i=1}^n x_i/\sigma_j)^2$  and penalty  $\beta = 1000$  where  $\sigma_j^2$  is the sample variance of  $(\log \rho_{ij} - \log \rho_{i+1,j})/\sqrt{2}$ . Define  $o_{ij}$  to be the average of  $\log \rho_{ij}$  over the segment containing region  $i$ .

We then compute the de-segmented counts  $Y_{ij}^{\text{deseg}} := \text{round}(\alpha_i\beta_j \exp(\log(\rho_{ij}) - o_{ij}))$ . The idea is that this adjusts out the departures from copy neutral (that is, from normal diploid) as inferred by the segmentation algorithm, to create a panel of pseudo-normals.

**Running the GBM to estimate copy ratios.** We run the GBM estimation algorithm on the de-segmented training data using  $M = 5$  latent factors. To avoid overfitting of the latent factors on the training data, we fix the sample-specific log-dispersion  $T$  after running the initialization procedure – that is, we do not update  $T$  at each iteration of the algorithm. We use the defaults for all other algorithm settings. For covariates, we construct  $X$  to include  $\log(\text{length}_i)$ ,  $\text{gc}_i$ , and  $(\text{gc}_i - \overline{\text{gc}})^2$ , and we use no sample covariates.

On the test data, we update  $T$  at each iteration as usual. We construct  $X$  to include the same covariates as before, along with 5 additional covariates equal to the columns of the  $U$  matrix that was estimated on the training data. We use no sample covariates, and we set  $M = 0$  on the test data. We define the GBM copy ratio estimates as the exponentiated residuals  $\rho_{ij}^{\text{GBM}} = \tilde{Y}_{ij}/\hat{\mu}_{ij}$  where  $\tilde{Y}_{ij} = Y_{ij} + 1/8$ ; also see Section 2.3.

**Running GATK to estimate copy ratios.** We run the CreateReadCount-PanelOfNormals GATK tool on the de-segmented training data, and to enable comparison with the GBM, we set the options to use 5 principal components and include all regions (`number-of-eigensamples 5`, `minimum-interval-median-percentile 0`, `maximum-zeros-in-interval-percentile 100`). To estimate copy ratios for the test data, we run the DenoiseReadCounts GATK tool using the pon file from CreateReadCount-



PanelOfNormals. On the test data, we use the original counts (not de-segmented).

**Performance metrics.** Suppose  $x, w \in \mathbb{R}^n$  and  $k \in \{0, 2, 4, \dots\}$ , where  $x_1, \dots, x_n$  represent noisy measurements of a signal of interest,  $w_i$  is a weight for point  $x_i$ , and  $k$  is a smoothing bandwidth. We define the local relative standard error as

$$\text{LRSE}(x, w, k) = \sqrt{\frac{1}{n} \sum_{i=1}^n (w_i/\bar{w}_i)^2 (x_i - \bar{x}_i)^2}$$

where  $\bar{w}_i$  is the moving average of  $w$  using bandwidth  $k$ , and  $\bar{x}_i$  is the weighted moving average of  $x$  using bandwidth  $k$  and weights  $w$ . More precisely,  $\bar{w}_i = \frac{1}{|A_i|} \sum_{j \in A_i} w_j$  and

$$\bar{x}_i = \frac{\sum_{j \in A_i} w_j x_j}{\sum_{j \in A_i} w_j} \tag{S3.1}$$

where  $A_i = \{\max(1, i - k/2), \dots, \min(n, i + k/2)\}$ . The idea is that if one is trying to estimate the mean of the signal, then a natural approach would be to use a weighted moving average, and the LRSE approximates the standard deviation (times  $\sqrt{k}$ ) of this estimator.

We define the weighted median absolute difference as

$$\text{WMAD}(x, w, k) = \text{median}\{k |\bar{x}_{i+1} - \bar{x}_i| : i = 1, \dots, n - 1\}$$

where  $\bar{x}_i$  is the same as above. The WMAD is similar to the median absolute deviation metric that is frequently used in this application, but it allows one to account for weights.

To assess copy ratio estimation performance for sample  $j$ , we take  $x_i$  to be the log copy ratio estimate for region  $i$  and we use a bandwidth of  $k = 100$  for both metrics. We take  $w_i = W_{ij}$  for the GBM (following Section 2.3), and for GATK we take  $w_i = 1$  since GATK does not provide weights/precisions.

**Estimated copy ratios for all test samples.** Figures S18 and S19 show heatmaps of the GATK and GBM copy ratio estimates on the  $\log_2$  scale for all 163 samples in the test set. For visualization, these heatmaps are smoothed using a moving weighted average

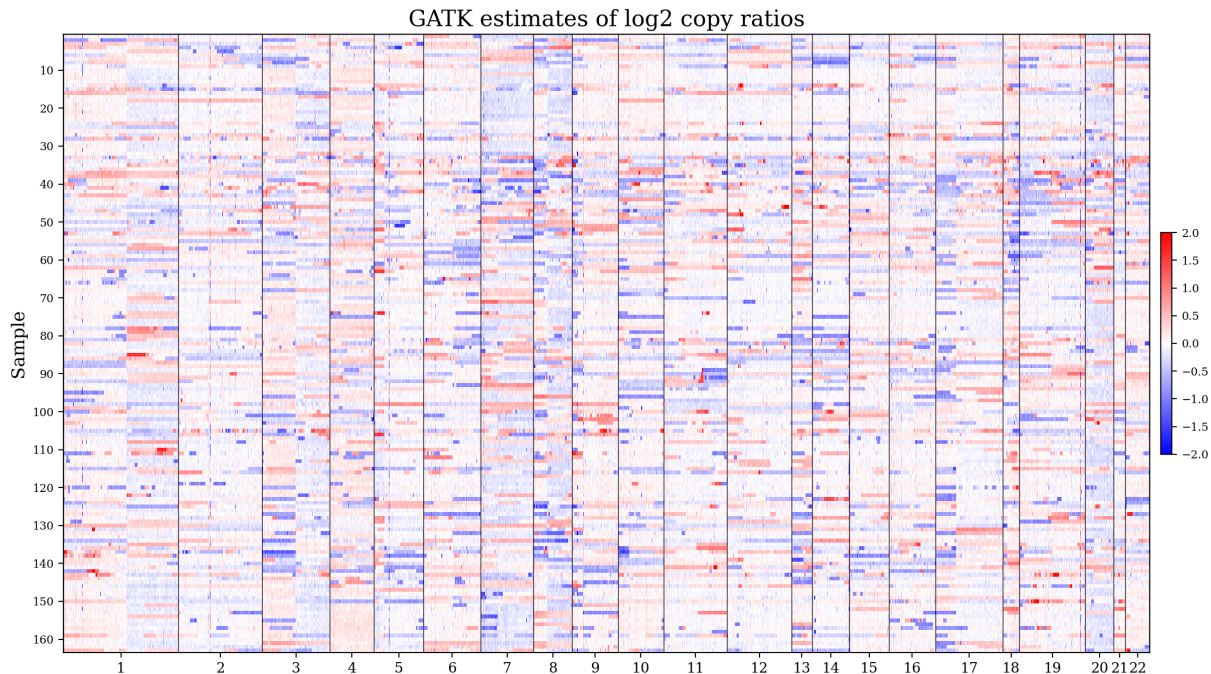


Figure S18: GATK copy ratio estimates (on the log<sub>2</sub> scale) for all 163 test samples from the CCLE whole-exome sequencing dataset. The x-axis represents genomic position. Red and blue indicate copy gains and losses, respectively.

as in Equation S3.1 with  $k = 3$ . The GBM estimates are visibly less noisy and appear to infer copy neutral regions (white portions of the heatmap) more accurately than GATK.

## S4 Exponential dispersion families

For completeness, we state the basic results for discrete EDFs that we use in this paper. See Jørgensen (1987) or Agresti (2015) for reference on this material. Suppose

$$f(y \mid \theta, r) = \exp(\theta y - r\kappa(\theta))h(y, r) \tag{S4.1}$$

is a p.m.f. on  $y \in \mathbb{Z}$  for  $\theta \in \Theta$  and  $r \in \mathcal{R}$ , where  $\Theta \subseteq \mathbb{R}$  and  $\mathcal{R} \subseteq (0, \infty)$  are convex open sets. If  $Y \sim f(y \mid \theta, r)$ , then

$$\mathbb{E}(Y \mid \theta, r) = r\kappa'(\theta) \quad \text{and} \quad \text{Var}(Y \mid \theta, r) = r\kappa''(\theta). \tag{S4.2}$$




Figure S19: GBM copy ratio estimates (on the log2 scale) for all 163 test samples from the CCLE whole-exome sequencing dataset. The x-axis represents genomic position. Red and blue indicate copy gains and losses, respectively.

These properties are straightforward to derive from the fact that  $\log \sum_{y \in \mathbb{Z}} \exp(\theta y) h(y, r) = r\kappa(\theta)$ . For  $y, \theta$ , and  $r$  such that  $h(y, r) > 0$ , let

$$\begin{aligned} \mathcal{L} &= \mathcal{L}(y, \theta, r) := \log f(y \mid \theta, r) = \theta y - r\kappa(\theta) + \log h(y, r), \\ \mu &= \mu(\theta, r) := \mathbb{E}(Y \mid \theta, r) = r\kappa'(\theta), \text{ and} \\ \sigma^2 &= \sigma^2(\theta, r) := \text{Var}(Y \mid \theta, r) = r\kappa''(\theta). \end{aligned}$$

Assume  $r$  is not functionally dependent on  $\theta$ . Then we have

$$\frac{\partial \mathcal{L}}{\partial \theta} = y - r\kappa'(\theta) = y - \mu \quad \text{and} \quad \frac{\partial^2 \mathcal{L}}{\partial \theta^2} = -r\kappa''(\theta) = -\sigma^2.$$

Assume  $\kappa'$  is invertible; this holds as long as  $\text{Var}(Y \mid \theta, r) > 0$  for all  $\theta$  and  $r$ . Then since  $\mu = r\kappa'(\theta)$ , we have  $\theta = \kappa'^{-1}(\mu/r)$ , and it follows that  $\partial\theta/\partial\mu = 1/(r\kappa''(\theta)) = 1/\sigma^2$  by the inverse function theorem. Similarly, defining  $\eta := g(\mu)$ , where  $g(\cdot)$  is a smooth function

such that  $g'$  is positive, we have  $\partial\mu/\partial\eta = 1/g'(\mu)$ . Therefore, by the chain rule,

$$\frac{\partial\mathcal{L}}{\partial\eta} = \frac{\partial\mathcal{L}}{\partial\theta} \frac{\partial\theta}{\partial\mu} \frac{\partial\mu}{\partial\eta} = (y - r\kappa'(\theta)) \frac{1}{r\kappa''(\theta)} \frac{1}{g'(\mu)} = \frac{y - \mu}{\sigma^2 g'(\mu)},$$

and

$$\frac{\partial^2\mathcal{L}}{\partial\eta^2} = \left( -\frac{\partial\mu}{\partial\eta} \right) \frac{1}{\sigma^2 g'(\mu)} + (y - \mu) \frac{\partial}{\partial\eta} \left( \frac{1}{\sigma^2 g'(\mu)} \right).$$

Thus, if  $Y \sim f(y | \theta, r)$  and  $\mathcal{L} = \log f(Y | \theta, r)$ , then

$$\frac{\partial\mathcal{L}}{\partial\eta} = \frac{Y - \mu}{\sigma^2 g'(\mu)} \quad \text{and} \quad \mathbb{E} \left( -\frac{\partial^2\mathcal{L}}{\partial\eta^2} \right) = \frac{1}{\sigma^2 g'(\mu)^2}. \quad (\text{S4.3})$$

If  $\eta$  depends on some parameters  $\alpha$  and  $\beta$  (and  $r$  does not), then

$$\frac{\partial\mathcal{L}}{\partial\alpha} = \frac{\partial\mathcal{L}}{\partial\eta} \frac{\partial\eta}{\partial\alpha} \quad \text{and} \quad \frac{\partial^2\mathcal{L}}{\partial\alpha\partial\beta} = \frac{\partial^2\mathcal{L}}{\partial\eta^2} \frac{\partial\eta}{\partial\alpha} \frac{\partial\eta}{\partial\beta} + \frac{\partial\mathcal{L}}{\partial\eta} \frac{\partial^2\eta}{\partial\alpha\partial\beta}.$$

Therefore, by Equation S4.3,

$$\frac{\partial\mathcal{L}}{\partial\alpha} = \frac{Y - \mu}{\sigma^2 g'(\mu)} \frac{\partial\eta}{\partial\alpha} \quad (\text{S4.4})$$

$$\mathbb{E} \left( -\frac{\partial^2\mathcal{L}}{\partial\alpha\partial\beta} \right) = \frac{1}{\sigma^2 g'(\mu)^2} \frac{\partial\eta}{\partial\alpha} \frac{\partial\eta}{\partial\beta} \quad (\text{S4.5})$$

since  $\partial^2\eta/\partial\alpha\partial\beta$  does not depend on  $Y$  and  $\mathbb{E}(\partial\mathcal{L}/\partial\eta) = 0$ .

## S5 Gradient and Fisher information for EDF-GBMs

In Section S5.1, we derive the gradient and Fisher information matrix with respect to each of  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ , and  $V$  individually, for a discrete EDF-GBM. In Section S5.2, we specialize these formulas to the case of NB-GBMs, and we derive the gradient and observed Fisher information for  $S$  and  $T$  in this case. For completeness, in Section S5.3 we also provide the cross-terms of the Fisher information between all pairs of components, although not

all of these are needed for our approach. The formulas contain factors reminiscent of the gradient and Fisher information for a standard GLM, which take a standard form based on the link function and the mean-variance relationship (Agresti, 2015).

For our estimation algorithm, we only require the Fisher information matrix with respect to each parameter matrix/vector individually, rather than jointly. Meanwhile, for inference, we also use the constraint-augmented Fisher information matrix for  $U$  and  $V$  jointly; see Section S6.2. To enable comparison with the standard approach of using the full Fisher information matrix, in Section S6.3 we provide the constraint-augmented Fisher information matrix for all parameters jointly.

Consider a GBM with discrete EDF outcomes, that is, suppose  $Y_{ij} \sim f(y \mid \theta_{ij}, r_{ij})$  where  $\theta_{ij} := \kappa'^{-1}(\mu_{ij}/r_{ij})$ ,  $\mu_{ij} := g^{-1}(\eta_{ij})$ , and  $\eta = \eta(A, B, C, D, U, V) \in \mathbb{R}^{I \times J}$  as in Equation 2.3. This makes  $\mu_{ij} = \mathbb{E}(Y_{ij} \mid \theta_{ij}, r_{ij})$ ; also, define  $\sigma_{ij}^2 := \text{Var}(Y_{ij} \mid \theta_{ij}, r_{ij})$ . Let

$$\mathcal{L} := \sum_{i=1}^I \sum_{j=1}^J \mathcal{L}_{ij} = \sum_{i=1}^I \sum_{j=1}^J \log f(Y_{ij} \mid \theta_{ij}, r_{ij}) \quad (\text{S5.1})$$

denote the overall log-likelihood, where  $\mathcal{L}_{ij} := \log f(Y_{ij} \mid \theta_{ij}, r_{ij})$ . By Equations S4.4 and S4.5, for any two univariate entries of  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ , and  $V$ , say  $\alpha$  and  $\beta$ , we have

$$\frac{\partial \mathcal{L}_{ij}}{\partial \alpha} = e_{ij} \frac{\partial \eta_{ij}}{\partial \alpha} \quad \text{and} \quad \mathbb{E} \left( - \frac{\partial^2 \mathcal{L}_{ij}}{\partial \alpha \partial \beta} \right) = w_{ij} \frac{\partial \eta_{ij}}{\partial \alpha} \frac{\partial \eta_{ij}}{\partial \beta} \quad (\text{S5.2})$$

where we define the matrices  $E \in \mathbb{R}^{I \times J}$  and  $W \in \mathbb{R}^{I \times J}$  such that

$$e_{ij} := \frac{\partial \mathcal{L}_{ij}}{\partial \eta_{ij}} = \frac{Y_{ij} - \mu_{ij}}{\sigma_{ij}^2 g'(\mu_{ij})} \quad \text{and} \quad w_{ij} := \mathbb{E} \left( - \frac{\partial^2 \mathcal{L}_{ij}}{\partial \eta_{ij}^2} \right) = \frac{1}{\sigma_{ij}^2 g'(\mu_{ij})^2}. \quad (\text{S5.3})$$

The partial derivatives  $\partial \eta_{ij} / \partial \alpha$  with respect to each entry of  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ , and  $V$  are:

$$\begin{aligned} \frac{\partial \eta_{ij}}{\partial a_{j'k}} &= x_{ik} \mathbf{1}(j = j') & \frac{\partial \eta_{ij}}{\partial v_{j'm}} &= u_{im} d_{mm} \mathbf{1}(j = j') \\ \frac{\partial \eta_{ij}}{\partial b_{i'\ell}} &= z_{j\ell} \mathbf{1}(i = i') & \frac{\partial \eta_{ij}}{\partial u_{i'm}} &= v_{jm} d_{mm} \mathbf{1}(i = i') \end{aligned} \quad (\text{S5.4})$$

$$\frac{\partial \eta_{ij}}{\partial c_{kl}} = x_{ik} z_{jl} \qquad \frac{\partial \eta_{ij}}{\partial d_{mm}} = u_{im} v_{jm}.$$

## S5.1 Gradient and Fisher information for each component

By Equations S5.2 and S5.4, we have

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial a_{jk}} &= \sum_{i=1}^I x_{ik} e_{ij} & \frac{\partial \mathcal{L}}{\partial v_{jm}} &= \sum_{i=1}^I u_{im} d_{mm} e_{ij} \\ \frac{\partial \mathcal{L}}{\partial b_{il}} &= \sum_{j=1}^J z_{jl} e_{ij} & \frac{\partial \mathcal{L}}{\partial u_{im}} &= \sum_{j=1}^J v_{jm} d_{mm} e_{ij} \\ \frac{\partial \mathcal{L}}{\partial c_{kl}} &= \sum_{i=1}^I \sum_{j=1}^J x_{ik} e_{ij} z_{jl} & \frac{\partial \mathcal{L}}{\partial d_{mm}} &= \sum_{i=1}^I \sum_{j=1}^J u_{im} e_{ij} v_{jm}. \end{aligned} \tag{S5.5}$$

To express these equations in matrix notation, we vectorize the parameter matrices as in Section 4.2:  $\vec{a} := \text{vec}(A^\text{T})$ ,  $\vec{b} := \text{vec}(B^\text{T})$ ,  $\vec{c} := \text{vec}(C)$ ,  $\vec{d} := \text{diag}(D)$ ,  $\vec{u} := \text{vec}(U^\text{T})$ , and  $\vec{v} := \text{vec}(V^\text{T})$ . By Equation S5.5, the gradient of  $\mathcal{L}$  with respect to each vectorized component is then:

$$\begin{aligned} \nabla_{\vec{a}} \mathcal{L} &= \text{vec}(X^\text{T} E) & \nabla_{\vec{v}} \mathcal{L} &= \text{vec}((UD)^\text{T} E) \\ \nabla_{\vec{b}} \mathcal{L} &= \text{vec}(Z^\text{T} E^\text{T}) & \nabla_{\vec{u}} \mathcal{L} &= \text{vec}((VD)^\text{T} E^\text{T}) \\ \nabla_{\vec{c}} \mathcal{L} &= \text{vec}(X^\text{T} E Z) & \nabla_{\vec{d}} \mathcal{L} &= \text{diag}(U^\text{T} E V). \end{aligned} \tag{S5.6}$$

For each component of the model, the entries of the Fisher information matrices are as follows, by Equations S5.2 and S5.4:

$$\begin{aligned} \mathbb{E} \left( - \frac{\partial^2 \mathcal{L}}{\partial a_{jk} \partial a_{j'k'}} \right) &= \mathbb{1}(j = j') \sum_{i=1}^I w_{ij} x_{ik} x_{ik'} \\ \mathbb{E} \left( - \frac{\partial^2 \mathcal{L}}{\partial b_{il} \partial b_{i'l'}} \right) &= \mathbb{1}(i = i') \sum_{j=1}^J w_{ij} z_{jl} z_{j'l'} \\ \mathbb{E} \left( - \frac{\partial^2 \mathcal{L}}{\partial c_{kl} \partial c_{k'l'}} \right) &= \sum_{i=1}^I \sum_{j=1}^J w_{ij} x_{ik} x_{ik'} z_{jl} z_{j'l'} \end{aligned} \tag{S5.7}$$

$$\begin{aligned}\mathbb{E}\left(-\frac{\partial^2 \mathcal{L}}{\partial v_{jm} \partial v_{j'm'}}\right) &= \mathbb{1}(j = j') \sum_{i=1}^I w_{ij} u_{im} d_{mm} u_{im'} d_{m'm'} \\ \mathbb{E}\left(-\frac{\partial^2 \mathcal{L}}{\partial u_{im} \partial u_{i'm'}}\right) &= \mathbb{1}(i = i') \sum_{j=1}^J w_{ij} v_{jm} d_{mm} v_{j'm'} d_{m'm'} \\ \mathbb{E}\left(-\frac{\partial^2 \mathcal{L}}{\partial d_{mm} \partial d_{m'm'}}\right) &= \sum_{i=1}^I \sum_{j=1}^J w_{ij} u_{im} v_{jm} u_{im'} v_{j'm'}.\end{aligned}$$

To express these equations in matrix form, we introduce the following notation. We write  $\text{Diag}(Q_1, \dots, Q_n)$  to denote the block diagonal matrix with blocks  $Q_1, \dots, Q_n$ ,

$$\text{Diag}(Q_1, \dots, Q_n) := \begin{bmatrix} Q_1 & 0 & \cdots & 0 \\ 0 & Q_2 & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & Q_n \end{bmatrix},$$

and we write  $\text{block}(Q_{ij} : i, j \in \{1, \dots, n\})$  to denote the block matrix with block  $i, j$  equal to  $Q_{ij}$ . For a matrix  $Q \in \mathbb{R}^{m \times n}$ , we write  $Q_{i*}$  and  $Q_{*j}$  to denote the diagonal matrices constructed from the  $i$ th row and  $j$ th column, respectively, that is,  $Q_{i*} := \text{Diag}(q_{i1}, \dots, q_{in})$  and  $Q_{*j} := \text{Diag}(q_{1j}, \dots, q_{mj})$ . Then, by Equation S5.7, the Fisher information matrices for each component of the model are:

$$\begin{aligned}\mathbb{E}(-\nabla_{\mathbf{a}}^2 \mathcal{L}) &= \text{Diag}(X^T W_{*1} X, \dots, X^T W_{*J} X) & (\text{S5.8}) \\ \mathbb{E}(-\nabla_{\mathbf{b}}^2 \mathcal{L}) &= \text{Diag}(Z^T W_{1*} Z, \dots, Z^T W_{I*} Z) \\ \mathbb{E}(-\nabla_{\mathbf{c}}^2 \mathcal{L}) &= \text{block}\left(\sum_{j=1}^J z_{j\ell} z_{j\ell'} (X^T W_{*j} X) : \ell, \ell' \in \{1, \dots, L\}\right) \\ \mathbb{E}(-\nabla_{\mathbf{v}}^2 \mathcal{L}) &= \text{Diag}((UD)^T W_{*1} (UD), \dots, (UD)^T W_{*J} (UD)) \\ \mathbb{E}(-\nabla_{\mathbf{u}}^2 \mathcal{L}) &= \text{Diag}((VD)^T W_{1*} (VD), \dots, (VD)^T W_{I*} (VD)) \\ \mathbb{E}(-\nabla_{\mathbf{d}}^2 \mathcal{L}) &= \sum_{j=1}^J (UV_{j*})^T W_{*j} (UV_{j*}).\end{aligned}$$

## S5.2 NB-GBM with log link: Gradient and Fisher information

The negative binomial distribution has probability mass function

$$\text{NegBin}(y \mid \mu, r) = \frac{\Gamma(y+r)}{\Gamma(y+1)\Gamma(r)} \left(\frac{\mu}{\mu+r}\right)^y \left(\frac{r}{\mu+r}\right)^r$$

for  $y \in \{0, 1, 2, \dots\}$ , given  $\mu > 0$  and  $r > 0$ . This is a discrete EDF of the form in Equation S4.1 with  $\theta = \log(\mu/(\mu+r))$  and  $\kappa(\theta) = -\log(1 - \exp(\theta))$ . Observe that

$$\begin{aligned} \kappa'(\theta) &= \frac{e^\theta}{1 - e^\theta} = \frac{\mu}{r} \\ \kappa''(\theta) &= \frac{e^\theta}{(1 - e^\theta)^2} = \frac{\mu + \mu^2/r}{r}. \end{aligned}$$

Thus, letting  $Y \sim \text{NegBin}(\mu, r)$ , we have  $\mathbb{E}(Y) = \mu$  and  $\text{Var}(Y) = \mu + \mu^2/r$  by Equation S4.2. For an NB-GBM with log link  $g(\mu) = \log(\mu)$ , the gradients and Fisher information matrices for  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ , and  $V$  are given by Equations S5.6 and S5.8 where, by Equation S5.3,

$$\mu_{ij} = \exp(\eta_{ij}), \quad w_{ij} = \frac{r_{ij}\mu_{ij}}{r_{ij} + \mu_{ij}}, \quad \text{and} \quad e_{ij} = (Y_{ij} - \mu_{ij}) \frac{w_{ij}}{\mu_{ij}}. \quad (\text{S5.9})$$

Next, we derive the gradient and observed information matrix for the log-dispersions. First, consider a single entry. Letting  $\psi(x)$  denote the digamma function, we have

$$\begin{aligned} \frac{\partial}{\partial r} \log \text{NegBin}(Y \mid \mu, r) &= \psi(Y+r) - \psi(r) + \log(r) + 1 - \log(\mu+r) - \frac{Y+r}{\mu+r} \\ \frac{\partial^2}{\partial r^2} \log \text{NegBin}(Y \mid \mu, r) &= \psi'(Y+r) - \psi'(r) + \frac{1}{r} - \frac{1}{\mu+r} - \frac{\mu-Y}{(\mu+r)^2}. \end{aligned} \quad (\text{S5.10})$$

We work with the observed information rather than the expected information since  $\mathbb{E}(\psi'(r+Y))$  does not seem to have a simple expression.

It turns out that Equation S5.10 tends to lead to arithmetic overflow/underflow in extreme cases. Although these extreme cases occur only occasionally for individual en-



tries, large GBMs have so many entries that these failures occur persistently and must be addressed. To avoid arithmetic overflow/underflow, we rewrite Equation S5.10 as follows:

$$\begin{aligned}\frac{\partial}{\partial r} \log \text{NegBin}(Y \mid \mu, r) &= \psi_{\Delta}(Y, r) - \log 1p(\mu/r) - \frac{Y - \mu}{r + \mu} + \text{err}(Y, r) \\ \frac{\partial^2}{\partial r^2} \log \text{NegBin}(Y \mid \mu, r) &= \psi'_{\Delta}(Y, r) + \frac{Y + \mu^2/r}{(r + \mu)^2} + \text{err}'(Y, r)\end{aligned}\tag{S5.11}$$

where

$$\begin{aligned}\log 1p(x) &= \begin{cases} x + \log((1+x)/e^x) & \text{if } x < 1 \\ \log(1+x) & \text{if } x \geq 1 \end{cases} \\ \psi_{\Delta}(y, r) &= \begin{cases} \psi(y+r) - \psi(r) & \text{if } r < 10^8 \\ \log 1p(y/r) & \text{if } r \geq 10^8 \end{cases} \\ \psi'_{\Delta}(y, r) &= \begin{cases} \psi'(y+r) - \psi'(r) & \text{if } r < 10^8 \\ -(y/r)/(y+r) & \text{if } r \geq 10^8 \end{cases}\end{aligned}\tag{S5.12}$$

and the error terms  $\text{err}(y, r)$  and  $\text{err}'(y, r)$  are typically exceedingly small and can be safely ignored. Mathematically,  $\log 1p(x) = \log(1+x)$ , however, numerically, the expression in Equation S5.12 computes this value in a way that helps avoid arithmetic overflow and underflow. Similarly,  $\psi_{\Delta}(y, r)$  and  $\psi'_{\Delta}(y, r)$  compute  $\psi(y+r) - \psi(r)$  and  $\psi'(y+r) - \psi'(r)$ , respectively, to very high accuracy while avoiding overflow/underflow; the errors in these approximations are  $\text{err}(y, r)$  and  $\text{err}'(y, r)$ , respectively. To derive Equation S5.11, we group terms of similar magnitude and we use the asymptotics of  $\psi(x)$  and  $\psi'(x)$ .

Now, we derive the gradient and observed information for  $S$  and  $T$  in the NB-GBM. Recall that we parametrize  $r_{ij} = \exp(-s_i - t_j - \omega)$  and we work in terms of the vector of feature-specific log-dispersion offsets  $S \in \mathbb{R}^I$ , the vector of sample-specific log-dispersion offsets  $T \in \mathbb{R}^J$ , and the overall log-dispersion  $\omega$ , subject to the identifiability constraints  $\frac{1}{I} \sum_i e^{s_i} = 1$  and  $\frac{1}{J} \sum_j e^{t_j} = 1$ . Let  $\mathcal{L}_{ij} = \log \text{NegBin}(Y_{ij} \mid \mu_{ij}, r_{ij})$  and  $\mathcal{L} = \sum_{i=1}^I \sum_{j=1}^J \mathcal{L}_{ij}$

as before. The derivatives with respect to  $s_i$  and  $t_j$  are then

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial s_i} &= \sum_{j=1}^J \frac{\partial \mathcal{L}_{ij}}{\partial r_{ij}} (-r_{ij}) & \frac{\partial^2 \mathcal{L}}{\partial s_i^2} &= \sum_{j=1}^J \left( \frac{\partial^2 \mathcal{L}_{ij}}{\partial r_{ij}^2} r_{ij}^2 + \frac{\partial \mathcal{L}_{ij}}{\partial r_{ij}} r_{ij} \right) \\ \frac{\partial \mathcal{L}}{\partial t_j} &= \sum_{i=1}^I \frac{\partial \mathcal{L}_{ij}}{\partial r_{ij}} (-r_{ij}) & \frac{\partial^2 \mathcal{L}}{\partial t_j^2} &= \sum_{i=1}^I \left( \frac{\partial^2 \mathcal{L}_{ij}}{\partial r_{ij}^2} r_{ij}^2 + \frac{\partial \mathcal{L}_{ij}}{\partial r_{ij}} r_{ij} \right)\end{aligned}\tag{S5.13}$$

since

$$\frac{\partial r_{ij}}{\partial s_i} = \frac{\partial r_{ij}}{\partial t_j} = -r_{ij} \quad \text{and} \quad \frac{\partial^2 r_{ij}}{\partial s_i^2} = \frac{\partial^2 r_{ij}}{\partial t_j^2} = r_{ij}.$$

By Equation S5.11,

$$\begin{aligned}\frac{\partial \mathcal{L}_{ij}}{\partial r_{ij}} &= \psi_{\Delta}(Y_{ij}, r_{ij}) - \log 1p(\mu_{ij}/r_{ij}) - \frac{Y_{ij} - \mu_{ij}}{r_{ij} + \mu_{ij}} + \text{err}(Y_{ij}, r_{ij}) \\ \frac{\partial^2 \mathcal{L}_{ij}}{\partial r_{ij}^2} &= \psi'_{\Delta}(Y_{ij}, r_{ij}) + \frac{Y_{ij} + \mu_{ij}^2/r_{ij}}{(r_{ij} + \mu_{ij})^2} + \text{err}'(Y_{ij}, r_{ij})\end{aligned}\tag{S5.14}$$

where the error terms are negligible. Note that  $\omega$  is implicitly optimized in the optimization-projection steps for  $S$  and  $T$  due to the likelihood-preserving projections, making it unnecessary to optimize with respect to  $\omega$  directly.

In practice, we do not use the off-diagonal terms of the observed information matrix for the log-dispersion parameters, since they seem to have a very small effect on the results. However, for completeness, we note here that

$$\frac{\partial^2 \mathcal{L}}{\partial s_i \partial t_j} = \frac{\partial^2 \mathcal{L}_{ij}}{\partial r_{ij}^2} r_{ij}^2 + \frac{\partial \mathcal{L}_{ij}}{\partial r_{ij}} r_{ij}$$

for all  $i, j$ , and

$$\frac{\partial^2 \mathcal{L}}{\partial s_i \partial s_{i'}} = 0 \quad \frac{\partial^2 \mathcal{L}}{\partial t_j \partial t_{j'}} = 0$$

for all  $i \neq i'$  and  $j \neq j'$ .

### S5.3 Fisher information between components of the model

In this section, we provide formulas for the off-block-diagonal entries of the full Fisher information matrix for a discrete EDF-GBM, that is, the entries that involve more than one of  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ ,  $V$ ,  $S$ ,  $T$ , and  $\omega$ . With the exception of the entries involving  $U$  and  $V$ , we do not use these formulas in our proposed algorithms, however, we provide the formulas here to enable comparison with the standard inference approach based on the full Fisher information matrix. As in Equation S5.3, we define

$$e_{ij} = \frac{\partial \mathcal{L}_{ij}}{\partial \eta_{ij}} \quad \text{and} \quad w_{ij} = \mathbb{E} \left( - \frac{\partial^2 \mathcal{L}_{ij}}{\partial \eta_{ij}^2} \right).$$

For any univariate entry of  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ , or  $V$ , say  $\beta$ , we have

$$\mathbb{E} \left( - \frac{\partial e_{ij}}{\partial \beta} \right) = \mathbb{E} \left( - \frac{\partial^2 \mathcal{L}_{ij}}{\partial \eta_{ij}^2} \right) \frac{\partial \eta_{ij}}{\partial \beta} = w_{ij} \frac{\partial \eta_{ij}}{\partial \beta}. \quad (\text{S5.15})$$

Using Equation S5.15 along with Equations S5.4 and S5.5, the cross terms among  $A$ ,  $B$ , and  $C$  are:

$$\mathbb{E} \left( - \frac{\partial^2 \mathcal{L}}{\partial a_{jk} \partial b_{il}} \right) = w_{ij} x_{ik} z_{jl} \quad (\text{S5.16})$$

$$\mathbb{E} \left( - \frac{\partial^2 \mathcal{L}}{\partial a_{jk} \partial c_{k'\ell}} \right) = \sum_{i=1}^I w_{ij} x_{ik} x_{ik'} z_{j\ell} \quad (\text{S5.17})$$

$$\mathbb{E} \left( - \frac{\partial^2 \mathcal{L}}{\partial b_{il} \partial c_{k'\ell}} \right) = \sum_{j=1}^J w_{ij} x_{ik} z_{j\ell} z_{j\ell'}. \quad (\text{S5.18})$$

Similarly, the cross terms among  $U$ ,  $V$ , and  $D$  are:

$$\mathbb{E} \left( - \frac{\partial^2 \mathcal{L}}{\partial u_{im} \partial v_{jm'}} \right) = w_{ij} u_{im'} d_{m'm'} d_{mm} v_{jm} \quad (\text{S5.19})$$

$$\mathbb{E} \left( - \frac{\partial^2 \mathcal{L}}{\partial u_{im} \partial d_{m'm'}} \right) = \sum_{j=1}^J w_{ij} u_{im'} v_{jm'} v_{jm} d_{mm} \quad (\text{S5.20})$$

$$\mathbb{E}\left(-\frac{\partial^2 \mathcal{L}}{\partial v_{jm} \partial d_{m'm'}}\right) = \sum_{i=1}^I w_{ij} u_{im} d_{mm} u_{im'} v_{jm'}. \quad (\text{S5.21})$$

The cross terms between  $A, B, C$  and  $U, V$ , and  $D$  are:

$$\mathbb{E}\left(-\frac{\partial^2 \mathcal{L}}{\partial a_{jk} \partial u_{im}}\right) = w_{ij} x_{ik} v_{jm} d_{mm} \quad (\text{S5.22})$$

$$\mathbb{E}\left(-\frac{\partial^2 \mathcal{L}}{\partial a_{jk} \partial v_{j'm}}\right) = \mathbf{1}(j = j') \sum_{i=1}^I w_{ij} x_{ik} u_{im} d_{mm} \quad (\text{S5.23})$$

$$\mathbb{E}\left(-\frac{\partial^2 \mathcal{L}}{\partial a_{jk} \partial d_{mm}}\right) = \sum_{i=1}^I w_{ij} x_{ik} u_{im} v_{jm} \quad (\text{S5.24})$$

$$\mathbb{E}\left(-\frac{\partial^2 \mathcal{L}}{\partial b_{i\ell} \partial u_{i'm}}\right) = \mathbf{1}(i = i') \sum_{j=1}^J w_{ij} z_{j\ell} v_{jm} d_{mm} \quad (\text{S5.25})$$

$$\mathbb{E}\left(-\frac{\partial^2 \mathcal{L}}{\partial b_{i\ell} \partial v_{jm}}\right) = w_{ij} z_{j\ell} u_{im} d_{mm} \quad (\text{S5.26})$$

$$\mathbb{E}\left(-\frac{\partial^2 \mathcal{L}}{\partial b_{i\ell} \partial d_{mm}}\right) = \sum_{j=1}^J w_{ij} z_{j\ell} u_{im} v_{jm} \quad (\text{S5.27})$$

$$\mathbb{E}\left(-\frac{\partial^2 \mathcal{L}}{\partial c_{k\ell} \partial u_{im}}\right) = \sum_{j=1}^J w_{ij} x_{ik} z_{j\ell} v_{jm} d_{mm} \quad (\text{S5.28})$$

$$\mathbb{E}\left(-\frac{\partial^2 \mathcal{L}}{\partial c_{k\ell} \partial v_{jm}}\right) = \sum_{i=1}^I w_{ij} x_{ik} z_{j\ell} u_{im} d_{mm} \quad (\text{S5.29})$$

$$\mathbb{E}\left(-\frac{\partial^2 \mathcal{L}}{\partial c_{k\ell} \partial d_{mm}}\right) = \sum_{i=1}^I \sum_{j=1}^J w_{ij} x_{ik} z_{j\ell} u_{im} v_{jm}. \quad (\text{S5.30})$$

All of the cross terms between  $(S, T, \omega)$  and  $(A, B, C, D, U, V)$  are zero. To see this, first note that by differentiating Equation S5.10 with respect to  $\mu$ ,

$$\frac{\partial^2}{\partial \mu \partial r} \log \text{NegBin}(Y \mid \mu, r) = \frac{Y - \mu}{(r + \mu)^2}.$$

Hence, for any entry of  $A, B, C, D, U$ , or  $V$ , say  $\beta$ , we have  $\mathbb{E}(-\partial^2 \mathcal{L} / \partial \beta \partial s_i) = 0$  since

$$\mathbb{E}\left(-\frac{\partial^2 \mathcal{L}_{ij}}{\partial \beta \partial s_i}\right) = \mathbb{E}\left(-\frac{\partial^2 \mathcal{L}_{ij}}{\partial \mu_{ij} \partial r_{ij}} \frac{\partial \mu_{ij}}{\partial \beta} \frac{\partial r_{ij}}{\partial s_i}\right) = -\mathbb{E}\left(\frac{Y_{ij} - \mu_{ij}}{(r_{ij} + \mu_{ij})^2}\right) \frac{\partial \mu_{ij}}{\partial \beta} \frac{\partial r_{ij}}{\partial s_i} = 0.$$

Similarly,  $\mathbb{E}(-\partial^2 \mathcal{L} / \partial \beta \partial t_j) = 0$  and  $\mathbb{E}(-\partial^2 \mathcal{L} / \partial \beta \partial \omega) = 0$ .

## S6 Constraint-augmented Fisher information

During estimation, we enforce constraints that ensure identifiability of the parameters, however, these constraints are not accounted for in the Fisher information matrix. Consequently, it turns out that to appropriately quantify uncertainty in the parameters, it is necessary to augment the Fisher information matrix to account for the identifiability constraints; see Section S6.1.

In our proposed algorithm, we use this constraint-augmentation technique only for  $U$  and  $V$  jointly (see Section S6.2), and for the remaining components we use our delta propagation technique, which does not involve special handling of the constraints. We also compare our proposed method to the standard approach of inverting the full constraint-augmented Fisher information matrix. In Section S6.1, we review the constraint-augmentation technique in general, and in Section S6.3, we provide formulas for the full constraint-augmented Fisher information for GBMs.

### S6.1 Constraint-augmentation technique

In this section, we review the constraint-augmentation technique of Aitchison and Silvey (1958) and Silvey (1959) in the setting of a finite-dimensional parametric model and we extend it to our setting; also see Silvey (1975, Section 4.7). Consider an i.i.d. model with parameter  $\theta \in \mathbb{R}^d$ , and suppose we want to perform estimation and inference subject to the constraint that  $g(\theta) = 0$ , where  $g(\theta) = (g_1(\theta), \dots, g_k(\theta))^T \in \mathbb{R}^k$  is a differentiable function. Suppose  $\hat{\theta}$  is the maximum likelihood estimator subject to  $g(\theta) = 0$ ; this is referred to as a “restricted maximum likelihood estimator” (Silvey, 1975).

Let  $J_\theta \in \mathbb{R}^{k \times d}$  be the Jacobian matrix of  $g$ , that is,  $J_{\theta,ij} = \partial g_i / \partial \theta_j$ . Let  $\mathcal{I}_\theta \in \mathbb{R}^{d \times d}$  be the Fisher information matrix for  $\theta$ , that is,  $\mathcal{I}_\theta = \mathbb{E}(-\nabla_\theta^2 \mathcal{L})$  where  $\mathcal{L}$  is the log-likelihood. Suppose  $\theta_0$  is the true value of the parameter. Aitchison and Silvey (1958) show that under

regularity conditions, when  $\mathcal{I}_{\theta_0}$  is invertible, the covariance matrix of  $\hat{\theta}$  is approximately equal to the leading  $d \times d$  submatrix of  $\begin{bmatrix} \mathcal{I}_{\theta_0} & J_{\theta_0}^\top \\ J_{\theta_0} & 0 \end{bmatrix}^{-1}$ . However, in GBMs, we need to consider situations in which  $\mathcal{I}_{\theta_0}$  is not invertible since the model is overparametrized unless the identifiability constraints are imposed. When  $\mathcal{I}_{\theta_0}$  is not invertible, Silvey (1959) extends the technique by showing that the covariance matrix of  $\hat{\theta}$  is approximately equal to the leading  $d \times d$  submatrix of  $\begin{bmatrix} \mathcal{I}_{\theta_0} + J_{\theta_0}^\top J_{\theta_0} & J_{\theta_0}^\top \\ J_{\theta_0} & 0 \end{bmatrix}^{-1}$ .

Since we employ *maximum a posteriori* estimates, we modify the technique to use the regularized Fisher information matrix  $F_\theta = \mathbb{E}(-\nabla_\theta^2(\mathcal{L} + \log \pi))$  in place of  $\mathcal{I}_\theta$ . For our choice of prior,  $F_{\theta_0}$  is invertible even when  $\mathcal{I}_{\theta_0}$  is not invertible, and consequently it turns out that the leading  $d \times d$  submatrices of  $\begin{bmatrix} F_{\theta_0} & J_{\theta_0}^\top \\ J_{\theta_0} & 0 \end{bmatrix}^{-1}$  and  $\begin{bmatrix} F_{\theta_0} + J_{\theta_0}^\top J_{\theta_0} & J_{\theta_0}^\top \\ J_{\theta_0} & 0 \end{bmatrix}^{-1}$  coincide; see Proposition S6.1. Thus, when using  $F_\theta$  instead of  $\mathcal{I}_\theta$ , it is not necessary to employ the extended version provided by Silvey (1959); this provides a big advantage, computationally, when we apply the method to perform inference for  $(U, V)$  in GBMs. Although Aitchison and Silvey (1958) and Silvey (1959) justify the technique in the i.i.d. setting, empirically we find that it works well in our non-i.i.d. setting also.

**Proposition S6.1.** *Let  $F \in \mathbb{R}^{d \times d}$  and  $J \in \mathbb{R}^{k \times d}$ . If  $F$  and  $JF^{-1}J^\top$  are invertible, then the leading  $d \times d$  submatrices of  $\begin{bmatrix} F & J^\top \\ J & 0 \end{bmatrix}^{-1}$  and  $\begin{bmatrix} F + J^\top J & J^\top \\ J & 0 \end{bmatrix}^{-1}$  are equal.*

*Proof.* Let  $A := F + J^\top J$ . By the formula for inverting a  $2 \times 2$  block matrix,

$$\begin{bmatrix} A & J^\top \\ J & 0 \end{bmatrix}^{-1} = \begin{bmatrix} A^{-1} - A^{-1}J^\top(JA^{-1}J^\top)^{-1}JA^{-1} & * \\ * & * \end{bmatrix}, \quad (\text{S6.1})$$

and the same formula applies for  $\begin{bmatrix} F & J^\top \\ J & 0 \end{bmatrix}^{-1}$ , but with  $F$  in place of  $A$ . Here,  $*$  denotes unneeded entries. By the Woodbury matrix inversion formula,

$$A^{-1} = F^{-1} - F^{-1}J^\top(I + JF^{-1}J^\top)^{-1}JF^{-1}. \quad (\text{S6.2})$$

Defining  $C := JF^{-1}J^T$ , this implies  $JA^{-1}J^T = C - C(I + C)^{-1}C = C(I + C)^{-1}$ , hence

$$(JA^{-1}J^T)^{-1} = (I + C)C^{-1}. \quad (\text{S6.3})$$

Using Equation S6.2 again, we see that  $JA^{-1} = JF^{-1} - C(I + C)^{-1}JF^{-1}$ , and thus

$$(JA^{-1}J^T)^{-1}JA^{-1} = (I + C)C^{-1}JF^{-1} - JF^{-1} = C^{-1}JF^{-1}. \quad (\text{S6.4})$$

Likewise,  $A^{-1}J^T = F^{-1}J^T - F^{-1}J^T(I + C)^{-1}C$ , so combining this with Equations S6.4 and S6.2 and canceling, we have

$$A^{-1} - A^{-1}J^T(JA^{-1}J^T)^{-1}JA^{-1} = F^{-1} - F^{-1}J^T(JF^{-1}J^T)^{-1}JF^{-1}.$$

The result follows by Equation S6.1. □

## S6.2 Constraint-augmented Fisher information for $(U, V)$

We have found that it is important to quantify uncertainty in  $U$  and  $V$  jointly and account for the identifiability constraints. In this section, we derive a computationally efficient method for computing the diagonal of the inverse of the constraint-augmented Fisher information matrix (Section S6.1) to obtain approximate standard errors for  $U$  and  $V$ .

Let  $J_u$  and  $J_v$  denote the constraint Jacobian matrices for  $U$  and  $V$ ; see Section S6.3.

The regularized, constraint-augmented Fisher information matrix for  $(U, V)$  is then

$$\tilde{F}_{(u,v)} := \begin{bmatrix} F_u & F_{uv} & J_u^T & 0 \\ F_{uv}^T & F_v & 0 & J_v^T \\ J_u & 0 & 0 & 0 \\ 0 & J_v & 0 & 0 \end{bmatrix}$$

where  $F_u = \mathbb{E}(-\nabla_u^2 \mathcal{L}) + \lambda_u \mathbf{I}$ ,  $F_v = \mathbb{E}(-\nabla_v^2 \mathcal{L}) + \lambda_v \mathbf{I}$ , and  $F_{uv} = \mathbb{E}(-\nabla_u \nabla_v^T \mathcal{L})$ ; formulas for each of these expectations are given in Equations S5.8 and S5.19.

We need the first  $IM + JM$  entries of  $\text{diag}(\tilde{F}_{(u,v)}^{-1})$  in order to obtain approximate standard errors for the entries of  $U$  and  $V$ , however, naively performing this matrix inversion is computationally intractable when  $I$  (or  $J$ ) is large. To compute this efficiently, we structure the calculation as follows. First, let  $P$  be the permutation matrix such that

$$P\tilde{F}_{(u,v)}P^T = \begin{bmatrix} F_u & J_u^T & F_{uv} & 0 \\ J_u & 0 & 0 & 0 \\ F_{uv}^T & 0 & F_v & J_v^T \\ 0 & 0 & J_v & 0 \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^T & \mathbf{C} \end{bmatrix}$$

where we define

$$\mathbf{A} = \begin{bmatrix} F_u & J_u^T \\ J_u & 0 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} F_{uv} & 0 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{C} = \begin{bmatrix} F_v & J_v^T \\ J_v & 0 \end{bmatrix}.$$

(We use sans serif font for these block matrices, such as  $\mathbf{A}$ , to distinguish them from parameter matrices such as  $A$ .) Since  $\text{diag}((P\tilde{F}_{(u,v)}P^T)^{-1}) = \text{diag}(P\tilde{F}_{(u,v)}^{-1}P^T) = P\text{diag}(\tilde{F}_{(u,v)}^{-1})$ , we can compute the diagonal of the inverse of  $P\tilde{F}_{(u,v)}P^T$  and then permute back to get  $\text{diag}(\tilde{F}_{(u,v)}^{-1})$ . By the formula for inversion of a  $2 \times 2$  block matrix,

$$(P\tilde{F}_{(u,v)}P^T)^{-1} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^T & \mathbf{C} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{A}^{-1} + \mathbf{A}^{-1}\mathbf{B}(\mathbf{C} - \mathbf{B}^T\mathbf{A}^{-1}\mathbf{B})^{-1}\mathbf{B}^T\mathbf{A}^{-1} & * \\ * & (\mathbf{C} - \mathbf{B}^T\mathbf{A}^{-1}\mathbf{B})^{-1} \end{bmatrix}$$

where  $*$  denotes entries that are not needed. Similarly, by the same formula,

$$\mathbf{A}^{-1} = \begin{bmatrix} F_u & J_u^T \\ J_u & 0 \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{D} & \mathbf{E} \\ \mathbf{E}^T & * \end{bmatrix}$$

where  $\mathbf{D} = F_u^{-1} - F_u^{-1}J_u^T(J_uF_u^{-1}J_u^T)^{-1}J_uF_u^{-1}$  and  $\mathbf{E} = F_u^{-1}J_u^T(J_uF_u^{-1}J_u^T)^{-1}$ . Thus,

$$(\mathbf{C} - \mathbf{B}^T\mathbf{A}^{-1}\mathbf{B})^{-1} = \begin{bmatrix} F_v - F_{uv}^T\mathbf{D}F_{uv} & J_v^T \\ J_v & 0 \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{G} & * \\ * & * \end{bmatrix},$$



where  $\mathbf{G}$  is defined to be the leading  $JM \times JM$  block of this matrix. Putting these pieces together justifies defining the approximate variance of each entry of  $\vec{v} = \text{vec}(V^T)$  as

$$\widehat{\text{var}}(\vec{v}) := \text{diag}(G)$$

since these are the entries of  $\text{diag}((P\tilde{F}_{(u,v)}P^T)^{-1})$  corresponding to  $\vec{v}$ . To approximate the variance of the entries of  $\vec{u} = \text{vec}(U^T)$ , observe that

$$\begin{aligned} \mathbf{A}^{-1}\mathbf{B}(\mathbf{C} - \mathbf{B}^T\mathbf{A}^{-1}\mathbf{B})^{-1}\mathbf{B}^T\mathbf{A}^{-1} &= \begin{bmatrix} \mathbf{D} & \mathbf{E} \\ \mathbf{E}^T & * \end{bmatrix} \begin{bmatrix} F_{uv} & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{G} & * \\ * & * \end{bmatrix} \begin{bmatrix} F_{uv}^T & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{D} & \mathbf{E} \\ \mathbf{E}^T & * \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{D}F_{uv}\mathbf{G}F_{uv}^T\mathbf{D} & * \\ * & * \end{bmatrix}. \end{aligned}$$

Therefore, we define

$$\widehat{\text{var}}(\vec{u}) := \text{diag}(\mathbf{D} + \mathbf{D}F_{uv}\mathbf{G}F_{uv}^T\mathbf{D})$$

since these are the entries of  $\text{diag}((P\tilde{F}_{(u,v)}P^T)^{-1})$  corresponding to  $\vec{u}$ .

### S6.3 Full constraint-augmented Fisher information for GBMs

To facilitate comparison with the classical approach, we derive the constraint-augmented Fisher information matrix for all of the components  $(A, B, C, D, U, V)$  jointly, even though our approach only requires the matrix for  $(U, V)$ . It is not necessary to include the log-dispersion parameters  $(S, T, \omega)$  since the Fisher information (as well as the constraint Jacobian) between these parameters and  $(A, B, C, D, U, V)$  is zero (see Section S5.3); thus, in the classical approach, inference for  $(S, T, \omega)$  and  $(A, B, C, D, U, V)$  can be performed independently since the constraint-augmented Fisher information decomposes.

For each of  $A, B, C, D, U$ , and  $V$ , we vectorize both the parameter matrix and the corresponding constraints, as follows. For  $A$ , the constraint  $Z^T A = 0$  can be written as  $g(A) = 0$  where  $g(A) = \text{vec}(A^T Z) \in \mathbb{R}^{KL}$ . The constraint Jacobian for  $\vec{a} = \text{vec}(A^T)$  is then  $J_a := Z^T \otimes \mathbf{I}_K$ , where  $\otimes$  denotes the Kronecker product and  $\mathbf{I}_K$  is the  $K \times K$  identity

matrix. Likewise, vectorizing the constraint on  $B$  as  $\text{vec}(B^T X) = 0$ , the constraint Jacobian for  $\vec{b} = \text{vec}(B^T)$  is  $J_b := X^T \otimes I_L$ .

For uncertainty quantification, the key constraints on  $U$  and  $V$  are  $X^T U = 0$ ,  $Z^T V = 0$ ,  $U^T U = I$ , and  $V^T V = I$ . Vectorizing, the constraints on  $U$  can be written as  $\text{vec}(U^T X) = 0$  and  $\text{vec}(U^T U - I) = 0$ . Thus, the constraint Jacobian for  $\vec{u}$  is

$$J_u := \begin{bmatrix} X_{1:} \otimes I_M & \cdots & X_{I:} \otimes I_M \\ (U_{1:} \otimes I_M) + (I_M \otimes U_{1:}) & \cdots & (U_{I:} \otimes I_M) + (I_M \otimes U_{I:}) \end{bmatrix} \in \mathbb{R}^{(MK+M^2) \times IM}.$$

Here, for a matrix  $Q \in \mathbb{R}^{m \times n}$ , we write  $Q_{i:}$  and  $Q_{:j}$  to denote the  $i$ th row and  $j$ th column as column vectors, respectively, that is,  $Q_{i:} := (q_{i1}, \dots, q_{in})^T \in \mathbb{R}^n$  and  $Q_{:j} := (q_{1j}, \dots, q_{mj})^T \in \mathbb{R}^m$ . The constraint Jacobian for  $\vec{v}$ , namely  $J_v$ , is computed the same way as  $J_u$  but with  $V$ ,  $Z$ ,  $J$ , and  $L$  in place of  $U$ ,  $X$ ,  $I$ , and  $K$ , respectively. There are no constraints on  $C$ , and the remaining constraints on  $D$ ,  $U$ , and  $V$  do not reduce the dimensionality of the parameter space. Thus, altogether, the constraint Jacobian for  $(A, B, C, D, U, V)$  is

$$\mathcal{J} := \begin{bmatrix} J_a & 0 & 0 & 0 & 0 \\ 0 & J_b & 0 & 0 & 0 \\ 0 & 0 & 0 & J_u & 0 \\ 0 & 0 & 0 & 0 & J_v \end{bmatrix}$$

where the column of zero blocks corresponding to  $(C, D)$  has width  $KL + M$  and height  $JK + IL + IM + JM$ .

Let  $\mathcal{I} = \mathbb{E}(-\nabla^2 \mathcal{L})$  denote the full Fisher information matrix for  $(A, B, C, D, U, V)$ . Formulas for all of the entries of  $\mathcal{I}$  are given in Equation S5.8 and in Section S5.3. The full constraint-augmented Fisher information matrix for all of the parameters is

$$\tilde{\mathcal{I}} := \begin{bmatrix} \mathcal{I} & \mathcal{J}^T \\ \mathcal{J} & 0 \end{bmatrix}.$$

The classical approach is then to define approximate standard errors as the square roots of

the entries of  $\text{diag}(\tilde{\mathcal{I}}^{-1})$  corresponding to each component (Section S6.1 and Silvey, 1975); for instance, the first  $JK$  entries of  $\text{diag}(\tilde{\mathcal{I}}^{-1})$  are the variances of the entries of  $\text{vec}(A^T)$ .

## S7 Step-by-step estimation algorithm

Given the inputs and preprocessing as described in Section 3, the algorithm is as follows.

### S7.1 Initialization procedure

We initialize the GBM estimation algorithm by (a) solving for values of  $A$ ,  $B$ , and  $C$  to minimize the sum-of-squares of the GBM residuals  $\varepsilon_{ij}$ , (b) randomly initializing  $D$ ,  $U$ , and  $V$ , and in the NB-GBM case, (c) iteratively updating  $S$ ,  $T$ , and  $\omega$  for a few iterations. This approach has the advantage of being simple, fast, and effective; see below for details.

It is somewhat tricky to initialize the algorithm well due to a chicken-and-egg problem. The issue is that having decent estimates of  $S$  and  $T$  is important to avoid overfitting to outliers, but we need a reasonable estimate of the mean matrix in order to estimate  $S$  and  $T$ . Our solution to this problem is to exclude  $D$ ,  $U$ , and  $V$  from the initial fitting of  $A$ ,  $B$ , and  $C$  in step (a) above. This prevents the latent factors from overfitting to outlier samples or outlier features, thus helping avoid getting stuck at a suboptimal point. In detail, we initialize as follows.

- (1) Compute  $\check{Y}_{ij} = g(Y_{ij} + \epsilon)$  where  $\epsilon = 1/8$ , and set  $\check{Y} = (\check{Y}_{ij}) \in \mathbb{R}^{I \times J}$ .
- (2)  $C \leftarrow X^+ \check{Y} (Z^+)^T$
- (3)  $A \leftarrow (X^+ \check{Y} - CZ^T)^T$
- (4)  $B \leftarrow \check{Y} (Z^+)^T - XC$
- (5) Sample  $q_{ij} \sim \mathcal{N}(0, 10^{-16})$  i.i.d. for all  $i, j$ , and set  $Q = (q_{ij}) \in \mathbb{R}^{I \times J}$ .
- (6) Compute  $U$ ,  $D$ , and  $V$  such that  $UDV^T = Q$  is the compact SVD of rank  $M$ .
- (7) Initialize  $S = 0$ ,  $T = 0$ , and  $\omega = 0$ , and then run the updates to  $S$  and  $T$  as defined in Section S7.2 for 4 iterations, using the current values of  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ , and  $V$ . Note that this also modifies  $\omega$ .

## S7.2 Updates to each component of the model

In this section, we provide step-by-step algorithms for updating each component of the model using the optimization-projection approach. The optimization part of each update is based on the bounded regularized Fisher scoring step in Equation 3.1, using the formulas for the gradients and Fisher information matrices derived in Section S5. The projection part of each update is based on Theorem 5.4. In the updates to  $G = UD$  and  $H = VD$ , we use the priors on  $G$  and  $H$  induced by the priors on  $U$  and  $V$ , given  $D$ ; see Section S9.

For a matrix  $Q \in \mathbb{R}^{m \times n}$ , we denote  $Q_{i:} := (q_{i1}, \dots, q_{in})^T \in \mathbb{R}^n$ ,  $Q_{:j} := (q_{1j}, \dots, q_{mj})^T \in \mathbb{R}^m$ ,  $Q_{i*} = \text{Diag}(Q_{i:})$ ,  $Q_{*j} = \text{Diag}(Q_{:j})$ ,  $\text{vec}(Q)$  is the column-wise vectorization of  $Q$ , and  $\text{block}(Q_{ij} : i, j \in \{1, \dots, n\})$  is the block matrix with blocks  $Q_{ij}$ . When multiplying by a diagonal matrix such as  $Q_{i*}$  or  $Q_{*j}$ , we do not allocate the full diagonal matrix and perform matrix multiplication. Instead, it is much more efficient to simply multiply each row (when left-multiplying) or column (when right-multiplying) by the corresponding diagonal entry.

### Computing $\eta$ , $\mu$ , $W$ , and $E$

- (1)  $\eta \leftarrow XA^T + BZ^T + XCZ^T + UDV^T$
- (2) For  $i = 1, \dots, I$  and  $j = 1, \dots, J$ ,
  - (a)  $\mu_{ij} \leftarrow \exp(\eta_{ij})$
  - (b)  $w_{ij} \leftarrow r_{ij}\mu_{ij}/(r_{ij} + \mu_{ij})$
  - (c)  $e_{ij} \leftarrow (Y_{ij} - \mu_{ij})w_{ij}/\mu_{ij}$

### Updating $A$

- (1) Recompute  $W$  and  $E$  using the current parameter estimates.
- (2) For  $j = 1, \dots, J$ 
  - (a)  $\xi \leftarrow (X^T W_{*j} X + \lambda_a \mathbf{I})^{-1} (X^T E_{:j} - \lambda_a A_{:j})$  (compute Fisher scoring step)
  - (b)  $A_{:j} \leftarrow A_{:j} + \xi \min\{1, \rho\sqrt{K}/\|\xi\|\}$  (apply modified step to  $A_{:j}$ )
- (3)  $Q \leftarrow Z^+ A$  (efficiently structure computation of projection)

- (4)  $A \leftarrow A - ZQ$  (enforce  $Z^T A = 0$  by projecting onto nullspace of  $Z^T$ )
- (5)  $C \leftarrow C + Q^T$  (compensate to preserve likelihood)

### Updating $B$

- (1) Recompute  $W$  and  $E$  using the current parameter estimates.
- (2) For  $i = 1, \dots, I$ 
  - (a)  $\xi \leftarrow (Z^T W_{i*} Z + \lambda_b \mathbf{I})^{-1} (Z^T E_{i*} - \lambda_b B_{i*})$  (compute Fisher scoring step)
  - (b)  $B_{i*} \leftarrow B_{i*} + \xi \min\{1, \rho\sqrt{L}/\|\xi\|\}$  (apply modified step to  $B_{i*}$ )
- (3)  $Q \leftarrow X^+ B$  (efficiently structure computation of projection)
- (4)  $B \leftarrow B - XQ$  (enforce  $X^T B = 0$  by projecting onto nullspace of  $X^T$ )
- (5)  $C \leftarrow C + Q$  (compensate to preserve likelihood)

### Updating $C$

- (1) Recompute  $W$  and  $E$  using the current parameter estimates.
- (2)  $F \leftarrow \text{block}\left(\sum_{j=1}^J z_{j\ell} z_{j\ell'} (X^T W_{*j} X) : \ell, \ell' \in \{1, \dots, L\}\right)$  (compute Fisher info)
- (3)  $\xi \leftarrow (F + \lambda_c \mathbf{I})^{-1} (\text{vec}(X^T E Z) - \lambda_c \text{vec}(C))$  (compute Fisher scoring step)
- (4)  $\text{vec}(C) \leftarrow \text{vec}(C) + \xi \min\{1, \rho\sqrt{KL}/\|\xi\|\}$  (apply modified step to  $C$ )

### Updating $D$

- (1) Recompute  $W$  and  $E$  using the current parameter estimates.
- (2)  $F \leftarrow \sum_{j=1}^J (UV_{j*})^T W_{*j} (UV_{j*})$  (compute Fisher information)
- (3)  $\xi \leftarrow (F + \lambda_d \mathbf{I})^{-1} (\text{diag}(U^T E V) - \lambda_d \text{diag}(D))$  (compute Fisher scoring step)
- (4)  $\text{diag}(D) \leftarrow \text{diag}(D) + \xi \min\{1, \rho\sqrt{M}/\|\xi\|\}$  (apply modified step to  $D$ )

### Updating $G = UD$

- (1) Recompute  $W$  and  $E$  using the current parameter estimates.
- (2)  $G \leftarrow UD$
- (3)  $\Lambda \leftarrow (D^2/\lambda_u)^{-1}$  (precision matrix for prior on each row of  $G$ )

- (4) For  $i = 1, \dots, I$ 
  - (a)  $\xi \leftarrow (V^T W_{i*} V + \Lambda)^{-1} (V^T E_{i:} - \Lambda G_{i:})$  (compute Fisher scoring step)
  - (b)  $G_{i:} \leftarrow G_{i:} + \xi \min\{1, \rho\sqrt{M}/\|\xi\|\}$  (apply modified step to  $G_{i:}$ )
- (5)  $Q \leftarrow X^+ G$  (efficiently structure computation of projection)
- (6)  $G \leftarrow G - XQ$  (enforce  $X^T G = 0$  by projecting onto nullspace of  $X^T$ )
- (7)  $A \leftarrow A + VQ^T$  (compensate to preserve likelihood)
- (8)  $Q \leftarrow Z^+ A$  (efficiently structure computation of projection)
- (9)  $A \leftarrow A - ZQ$  (enforce  $Z^T A = 0$  by projecting onto nullspace of  $Z^T$ )
- (10)  $C \leftarrow C + Q^T$  (compensate to preserve likelihood)
- (11) Run compact SVD of rank  $M$  on  $GV^T$ , yielding  $U, D, V$  such that  $UDV^T = GV^T$ .

### Updating $H = VD$

- (1) Recompute  $W$  and  $E$  using the current parameter estimates.
- (2)  $H \leftarrow VD$
- (3)  $\Lambda \leftarrow (D^2/\lambda_v)^{-1}$  (precision matrix for prior on each row of  $H$ )
- (4) For  $j = 1, \dots, J$ 
  - (a)  $\xi \leftarrow (U^T W_{*j} U + \Lambda)^{-1} (U^T E_{:j} - \Lambda H_{j:})$  (compute Fisher scoring step)
  - (b)  $H_{j:} \leftarrow H_{j:} + \xi \min\{1, \rho\sqrt{M}/\|\xi\|\}$  (apply modified step to  $H_{j:}$ )
- (5)  $Q \leftarrow Z^+ H$  (efficiently structure computation of projection)
- (6)  $H \leftarrow H - ZQ$  (enforce  $Z^T H = 0$  by projecting onto nullspace of  $Z^T$ )
- (7)  $B \leftarrow B + UQ^T$  (compensate to preserve likelihood)
- (8)  $Q \leftarrow X^+ B$  (efficiently structure computation of projection)
- (9)  $B \leftarrow B - XQ$  (enforce  $X^T B = 0$  by projecting onto nullspace of  $X^T$ )
- (10)  $C \leftarrow C + Q$  (compensate to preserve likelihood)
- (11) Run compact SVD of rank  $M$  on  $UH^T$ , yielding  $U, D, V$  such that  $UDV^T = UH^T$ .

## Updating $S$ in an NB-GBM

For the updates to  $S$  and  $T$ , we employ adaptive maximum step sizes  $\rho_{s_i}$  and  $\rho_{t_j}$  for  $s_i$  and  $t_j$ , respectively. This helps prevent occasional lack of convergence due to oscillating estimates. At the start of the algorithm, we initialize  $\rho_{s_i} \leftarrow \rho$  and  $\rho_{t_j} \leftarrow \rho$ . Define  $\log1p(x)$ ,  $\psi_{\Delta}(y, r)$ , and  $\psi'_{\Delta}(y, r)$  as in Equation S5.12. Note that we do not explicitly update  $\omega$ , since  $\omega$  is implicitly updated in the projection part of the updates to  $S$  and  $T$ .

- (1) Compute  $\mu$  using the current parameter estimates.
- (2) For  $i = 1, \dots, I$  and  $j = 1, \dots, J$ , (differentiate each term in the log-likelihood)
  - (a)  $\delta_{ij} \leftarrow -r_{ij} \left( \psi_{\Delta}(Y_{ij}, r_{ij}) - \log1p(\mu_{ij}/r_{ij}) - (Y_{ij} - \mu_{ij})/(r_{ij} + \mu_{ij}) \right)$
  - (b)  $\delta'_{ij} \leftarrow -\delta_{ij} + r_{ij}^2 \psi'_{\Delta}(Y_{ij}, r_{ij}) + (Y_{ij} + \mu_{ij}^2/r_{ij})/(1 + \mu_{ij}/r_{ij})^2$
- (3) For  $i = 1, \dots, I$ 
  - (a)  $g \leftarrow -\lambda_s(s_i - m_s) + \sum_{j=1}^J \delta_{ij}$  (derivative of log-posterior with respect to  $s_i$ )
  - (b)  $h \leftarrow -\lambda_s + \sum_{j=1}^J \delta'_{ij}$  (second derivative of log-posterior with respect to  $s_i$ )
  - (c) If  $h < 0$  then  $\xi \leftarrow -g/h$ , otherwise,  $\xi \leftarrow g$ . (Newton if valid, otherwise gradient)
  - (d)  $s_i \leftarrow s_i + \xi \min\{1, \rho_{s_i}/|\xi|\}$  (apply modified optimization step to  $s_i$ )
  - (e) If  $|\xi| > \rho_{s_i}$  then  $\rho_{s_i} \leftarrow \rho_{s_i}/2$ , otherwise,  $\rho_{s_i} \leftarrow \rho$ . (adapt maximum step size)
- (4)  $c \leftarrow \log(\frac{1}{I} \sum_{i=1}^I e^{s_i})$
- (5)  $S \leftarrow S - c$  (enforce constraint by projecting)
- (6)  $\omega \leftarrow \omega + c$  (compensate to preserve likelihood)
- (7)  $r_{ij} \leftarrow \exp(-s_i - t_j - \omega)$  for  $i = 1, \dots, I$  and  $j = 1, \dots, J$  (update inverse dispersions)

## Updating $T$ in an NB-GBM

Steps (1)-(2) and (7) are the same as in the update to  $S$ . Steps (3)-(6) become:

- (3) For  $j = 1, \dots, J$ 
  - (a)  $g \leftarrow -\lambda_t(t_j - m_t) + \sum_{i=1}^I \delta_{ij}$  (derivative of log-posterior with respect to  $t_j$ )
  - (b)  $h \leftarrow -\lambda_t + \sum_{i=1}^I \delta'_{ij}$  (second derivative of log-posterior with respect to  $t_j$ )
  - (c) If  $h < 0$  then  $\xi \leftarrow -g/h$ , otherwise,  $\xi \leftarrow g$ . (Newton if valid, otherwise gradient)

- (d)  $t_j \leftarrow t_j + \xi \min\{1, \rho_{t_j}/|\xi|\}$  (apply modified optimization step to  $s_i$ )
- (e) If  $|\xi| > \rho_{t_j}$  then  $\rho_{t_j} \leftarrow \rho_{t_j}/2$ , otherwise,  $\rho_{t_j} \leftarrow \rho$ . (adapt maximum step size)
- (4)  $c \leftarrow \log(\frac{1}{J} \sum_{j=1}^J e^{t_j})$
- (5)  $T \leftarrow T - c$  (enforce constraint by projecting)
- (6)  $\omega \leftarrow \omega + c$  (compensate to preserve likelihood)

### Bias correction for $S$ and $T$ in an NB-GBM

Empirically, when the true values of  $S$  and  $T$  are low, the maximum likelihood estimates tend to exhibit a downward bias. Occasionally, this leads to massive underestimation of some of the log-dispersion values. This issue is mitigated somewhat by using a prior to shrink the estimates toward zero, however, it seems difficult to tune the prior to appropriately balance the bias. Thus, we employ the following simple bias correction procedure, applied after the final iteration of the estimation algorithm. Choose lower bounds  $s_*$  and  $t_*$  on  $s_i$  and  $t_j$ , respectively; we use  $s_* = t_* = -4$  as defaults.

- (1)  $s_i \leftarrow s_* + \log(\exp(s_i - s_*) + 1)$  for  $i = 1, \dots, I$  (apply bias correction to  $S$ )
- (2)  $c \leftarrow \log(\frac{1}{I} \sum_{i=1}^I e^{s_i})$
- (3)  $S \leftarrow S - c$  (enforce constraint by projecting)
- (4)  $\omega \leftarrow \omega + c$  (compensate to preserve likelihood)

The same procedure is applied to  $T$ , with  $t_*$  in place of  $s_*$ . We find that this improves the accuracy of the log-dispersion estimates when the true values are at the low end.

### S7.3 Remarks on the estimation algorithm

We continue iterating until either (a) the relative change in log-likelihood+log-prior (Equations S5.1 and S9.2) from one iteration to the next is less than the convergence tolerance  $\tau$  or (b) the maximum number of iterations has been reached.

In the updates to  $G = UD$  and  $H = VD$ , we use the compact SVD to enforce the constraints on  $D$ ,  $U$ , and  $V$ . Fast computation of the compact SVD can be done using



procedures for the truncated SVD, which allows one to specify the rank (that is, the number of latent factors  $M$ ). Procedures for the truncated SVD are available in many programming languages. It is not necessary to enforce the ordering and sign constraints on  $D$ ,  $U$ , and  $V$  (Conditions 2.1(d) and 2.1(e)) during the iterative updates since both the likelihood and the prior are invariant to the order and sign of the latent factors.

Note that, due to the symmetry of the model, the updates for  $A$  and  $B$  are similar enough that a single function can be used to compute both of them, with an option to handle the transpose for  $C$ . Likewise, a single function can be used to compute both the  $UD$  and  $VD$  updates, with an option to handle transposes appropriately.

## S8 Step-by-step inference algorithm

**Notation.** For matrices  $A$  and  $B$ , we use  $A \otimes B$  to denote the Kronecker product and  $A \odot B$  for the element-wise product. Normally, we write  $AB$  for matrix multiplication, but for improved clarity we use  $A \times B$  to denote matrix multiplication when multi-letter variables such as `invFc` are involved. We write `hcat`( $A_1, \dots, A_n$ ) for the horizontal concatenation of matrices  $A_1, \dots, A_n$ , that is, `hcat`( $A_1, \dots, A_n$ ) := [ $A_1 \ \dots \ A_n$ ]. Likewise, `vcat` denotes vertical concatenation. We define `block`( $j, K$ ) :=  $((j-1)K+1, (j-1)K+2, \dots, jK)$ . For a matrix  $A \in \mathbb{R}^{m \times n}$ , `colsums`( $A$ ) denotes the vector of column sums, that is, `colsums`( $A$ ) =  $(\sum_i a_{i1}, \dots, \sum_i a_{in})^T \in \mathbb{R}^n$ . Likewise, `rowsums` denotes the row sums. For a vector  $x \in \mathbb{R}^{mn}$ , we define `reshape`( $x, m, n$ ) to be the matrix  $A \in \mathbb{R}^{m \times n}$  such that  $x = \text{vec}(A)$ . For a vector  $x \in \mathbb{R}^n$ , we write `Diag`( $x$ ) to denote the  $n \times n$  diagonal matrix with  $x$  on the diagonal. For a matrix  $A \in \mathbb{R}^{m \times n}$  and vectors  $x \in \mathbb{R}^m$ ,  $y \in \mathbb{R}^n$ , with  $m \neq n$ , we extend the  $\odot$  operator as follows:  $A \odot x = x \odot A := \text{Diag}(x)A$  and  $A \odot y = y \odot A := A \text{Diag}(y)$ . We write `sqrt`( $\cdot$ ) to denote the element-wise square root. We use  $I_n$  to denote the  $n \times n$  identity matrix.

### Preprocessing.

- (1) Compute the inverse dispersions  $r_{ij} \leftarrow \exp(-s_i - t_j - \omega)$  for all  $i, j$ .
- (2) Compute  $\mu$ ,  $W$ , and  $E$  as in Section S7.

- (3) Compute  $\mathbf{dWM} \in \mathbb{R}^{I \times J}$  where  $\mathbf{dWM}_{ij} \leftarrow \mu_{ij} r_{ij}^2 / (r_{ij} + \mu_{ij})^2$ .
- (4) Compute  $\mathbf{dEM} \in \mathbb{R}^{I \times J}$  where  $\mathbf{dEM}_{ij} \leftarrow -\mu_{ij} r_{ij} (r_{ij} + Y_{ij}) / (r_{ij} + \mu_{ij})^2$ .
- (5)  $\mathbf{gradA} \leftarrow E^T X$
- (6)  $\mathbf{gradB} \leftarrow EZ$
- (7)  $\mathbf{gradC} \leftarrow X^T EZ$
- (8) Compute  $\delta_{ij}$  and  $\delta'_{ij}$  for all  $i, j$  using the formula from the  $S$  update during estimation.
- (9)  $\mathbf{gradS}_i \leftarrow -\lambda_s (s_i - m_s) + \sum_{j=1}^J \delta_{ij}$  for  $i = 1, \dots, I$ .
- (10)  $\mathbf{gradT}_j \leftarrow -\lambda_t (t_j - m_t) + \sum_{i=1}^I \delta_{ij}$  for  $j = 1, \dots, J$ .

**Compute conditional uncertainty for each component.**

- (1)  $\mathbf{invFa}_j \leftarrow (X^T W_{*j} X + \lambda_a \mathbf{I})^{-1}$  for  $j = 1, \dots, J$ .
- (2)  $\mathbf{invFb}_i \leftarrow (Z^T W_{i*} Z + \lambda_b \mathbf{I})^{-1}$  for  $i = 1, \dots, I$ .
- (3)  $\mathbf{invFc} \leftarrow (\mathbb{E}(-\nabla_c^2 \mathcal{L}) + \lambda_c \mathbf{I})^{-1}$  where  $\mathbb{E}(-\nabla_c^2 \mathcal{L})$  is given in Equation S5.8.
- (4)  $\mathbf{invFu}_i \leftarrow ((VD)^T W_{i*} (VD) + \lambda_u \mathbf{I})^{-1}$  for  $i = 1, \dots, I$ .
- (5)  $\mathbf{invFv}_j \leftarrow ((UD)^T W_{*j} (UD) + \lambda_v \mathbf{I})^{-1}$  for  $j = 1, \dots, J$ .
- (6)  $\mathbf{invFs}_i \leftarrow 1 / (\lambda_s - \sum_{j=1}^J \delta'_{ij})$  for all  $i = 1, \dots, I$ .
- (7)  $\mathbf{invFt}_j \leftarrow 1 / (\lambda_t - \sum_{i=1}^I \delta'_{ij})$  for all  $j = 1, \dots, J$ .
- (8)  $\mathbf{invFs} \leftarrow (\mathbf{invFs}_1, \dots, \mathbf{invFs}_I)^T$
- (9)  $\mathbf{invFt} \leftarrow (\mathbf{invFt}_1, \dots, \mathbf{invFt}_J)^T$

**Compute constraint Jacobians for  $U$  and  $V$ .**

- (1)  $\mathbf{Ju}_i \leftarrow \mathbf{vcat}(X_{i:} \otimes \mathbf{I}_M, (U_{i:} \otimes \mathbf{I}_M) + (\mathbf{I}_M \otimes U_{i:}))$  for  $i = 1, \dots, I$ .
- (2)  $\mathbf{Ju} \leftarrow \mathbf{hcat}(\mathbf{Ju}_1, \dots, \mathbf{Ju}_I)$
- (3)  $\mathbf{Jv}_j \leftarrow \mathbf{vcat}(Z_{j:} \otimes \mathbf{I}_M, (V_{j:} \otimes \mathbf{I}_M) + (\mathbf{I}_M \otimes V_{j:}))$  for  $j = 1, \dots, J$ .
- (4)  $\mathbf{Jv} \leftarrow \mathbf{hcat}(\mathbf{Jv}_1, \dots, \mathbf{Jv}_J)$

**Compute joint uncertainty in  $(U, V)$  accounting for constraints.**

- (1)  $\mathbf{Fuv}_i \leftarrow \mathbf{hcat}(w_{i1}(DV_{1:})(DU_{i:})^T, \dots, w_{iJ}(DV_{J:})(DU_{i:})^T)$  for  $i = 1, \dots, I$
- (2)  $\mathbf{Fuv} \leftarrow \mathbf{vcat}(\mathbf{Fuv}_1, \dots, \mathbf{Fuv}_I)$
- (3)  $\mathbf{FJ} \leftarrow \mathbf{vcat}(\mathbf{invFu}_1 \times \mathbf{Ju}_1^T, \dots, \mathbf{invFu}_I \times \mathbf{Ju}_I^T)$

- (4)  $\text{FuvFJ} \leftarrow \text{Fuv}^T \times \text{FJ}$
- (5)  $\text{invJFJ} \leftarrow (\text{Ju} \times \text{FJ})^{-1}$
- (6)  $\text{FFuv} \leftarrow \text{vcat}(\text{invFu}_1 \times \text{Fuv}_1, \dots, \text{invFu}_I \times \text{Fuv}_I)$
- (7)  $\text{FuvFFuv} \leftarrow \text{Fuv}^T \times \text{FFuv}$
- (8)  $\text{Fv} \leftarrow$  block diagonal matrix with  $j$ th block equal to  $(UD)^T W_{*j}(UD) + \lambda_v \text{I}$ .
- (9)  $\text{A} \leftarrow \text{Fv} - \text{FuvFFuv} + \text{FuvFJ} \times \text{invJFJ} \times \text{FuvFJ}^T$
- (10)  $\text{B} \leftarrow \begin{bmatrix} \text{A} & \text{Jv}^T \\ \text{Jv} & 0 \end{bmatrix}^{-1}$
- (11)  $\text{C} \leftarrow \text{B}[1:JM, 1:JM]$  (that is,  $\text{C}$  is the leading  $JM \times JM$  block of  $\text{B}$ )
- (12)  $\text{FuvD} \leftarrow \text{FFuv}^T - \text{FuvFJ} \times \text{invJFJ} \times \text{FJ}^T$
- (13)  $\text{d} \leftarrow \text{vcat}(\text{diag}(\text{invFu}_1), \dots, \text{diag}(\text{invFu}_I))$
- (14)  $\text{f} \leftarrow \text{colsums}(\text{FJ}^T \odot (\text{invJFJ} \times \text{FJ}^T))$
- (15)  $\text{g} \leftarrow \text{colsums}(\text{FuvD} \odot (\text{C} \times \text{FuvD}))$
- (16)  $\text{varU} \leftarrow \text{d} - \text{f} + \text{g}$
- (17)  $\text{varV} \leftarrow \text{diag}(\text{C})$

**Propagate uncertainty from  $U$  to  $A$ .**

- (1) For  $j = 1, \dots, J$ ,
  - (a)  $\text{Q}_j \leftarrow (-\text{invFa}_j \times (X \odot \text{dWM}_{:,j})^T) \odot (X \times \text{invFa}_j \times \text{gradA}_{j,:})^T + \text{invFa}_j \times (X \odot \text{dEM}_{:,j})^T$
- (2)  $\text{dA} \leftarrow \text{vcat}(\text{Q}_1 \otimes (DV_{1,:})^T, \dots, \text{Q}_J \otimes (DV_{J,:})^T)$
- (3)  $\text{varAfromU} \leftarrow \text{colsums}(\text{dA}^T \odot (\text{varU} \odot \text{dA}^T))$

**Propagate uncertainty from  $V$  to  $A$ .**

- (1) For  $j = 1, \dots, J$ ,
  - (a) Initialize  $\text{dA} \in \mathbb{R}^{K \times M}$  to all zeros.
  - (b) For  $m = 1, \dots, M$ ,
    - (i)  $\text{XdE} \leftarrow X^T(\text{dEM}_{:,j} \odot (D_m U_{:,m}))$
    - (ii)  $\text{XdWX} \leftarrow (X^T((\text{dWM}_{:,j} \odot (D_m U_{:,m})) \odot X))$
    - (iii)  $\text{dA}_{:,m} \leftarrow (-\text{invFa}_j \times \text{XdWX}) \times (\text{invFa}_j \times \text{gradA}_{j,:}) + \text{invFa}_j \times \text{XdE}$
  - (c)  $\text{varAfromV}_j \leftarrow \text{colsums}(\text{dA}^T \odot (\text{varV}[\text{block}(j, M)] \odot \text{dA}^T))$

$$(2) \text{ varAfromV} \leftarrow \text{vcats}(\text{varAfromV}_1, \dots, \text{varAfromV}_J)$$

**Propagate uncertainty from  $U$  and  $V$  to  $B$ .**

- (1) Computing  $\text{varBfromU}$  is identical to calculating  $\text{varAfromV}$ , but with  $Z$ ,  $V$ ,  $\text{varU}$ ,  $\text{invFb}$ ,  $\text{gradB}$ ,  $\text{dWM}^T$ , and  $\text{dEM}^T$  in place of  $X$ ,  $U$ ,  $\text{varV}$ ,  $\text{invFa}$ ,  $\text{gradA}$ ,  $\text{dWM}$ , and  $\text{dEM}$ , respectively.
- (2) Computing  $\text{varBfromV}$  is identical to calculating  $\text{varAfromU}$ , but with  $Z$ ,  $U$ ,  $\text{varV}$ ,  $\text{invFb}$ ,  $\text{gradB}$ ,  $\text{dWM}^T$ , and  $\text{dEM}^T$  in place of  $X$ ,  $V$ ,  $\text{varU}$ ,  $\text{invFa}$ ,  $\text{gradA}$ ,  $\text{dWM}$ , and  $\text{dEM}$ , respectively.

**Propagate uncertainty from  $A$  to  $C$ .**

- (1) Initialize  $\text{dC} \in \mathbb{R}^{KL \times JK}$  to all zeros.
- (2) For  $j = 1, \dots, J$  and  $k = 1, \dots, K$ ,
  - (a)  $\text{dF} \leftarrow (Z_{j,:} Z_{j,:}^T) \otimes (X^T((\text{dWM}_{:,j} \odot X_{:,k}) \odot X))$
  - (b)  $\text{dgradC} \leftarrow (X^T \times (\text{dEM}_{:,j} \odot X_{:,k})) Z_{j,:}^T$
  - (c)  $\text{dC}[:, (j-1)K + k] \leftarrow \text{invFc} \times (-\text{dF} \times (\text{invFc} \times \text{vec}(\text{gradC})) + \text{vec}(\text{dgradC}))$
- (3)  $\text{invFdC}_j \leftarrow \text{invFa}_j \times \text{dC}[:, \text{block}(j, K)]^T$  for  $j = 1, \dots, J$
- (4)  $\text{invFdC} \leftarrow \text{vcats}(\text{invFdC}_1, \dots, \text{invFdC}_J)$
- (5)  $\text{varCfromA} \leftarrow \text{colsums}(\text{dC}^T \odot \text{invFdC})$

**Propagate uncertainty from  $B$  to  $C$ .**

- (1) Initialize  $\text{dC} \in \mathbb{R}^{KL \times IL}$  to all zeros.
- (2) For  $i = 1, \dots, I$  and  $\ell = 1, \dots, L$ ,
  - (a)  $\text{dF} \leftarrow (Z^T((\text{dWM}_{i,:} \odot Z_{:, \ell}) \odot Z)) \otimes (X_{i,:} X_{i,:}^T)$
  - (b)  $\text{dgradC} \leftarrow X_{i,:}((\text{dEM}_{i,:} \odot Z_{:, \ell})^T \times Z)$
  - (c)  $\text{dC}[:, (i-1)L + \ell] \leftarrow \text{invFc} \times (-\text{dF} \times (\text{invFc} \times \text{vec}(\text{gradC})) + \text{vec}(\text{dgradC}))$
- (3)  $\text{invFdC}_i \leftarrow \text{invFb}_i \times \text{dC}[:, \text{block}(i, L)]^T$  for  $i = 1, \dots, I$
- (4)  $\text{invFdC} \leftarrow \text{vcats}(\text{invFdC}_1, \dots, \text{invFdC}_I)$
- (5)  $\text{varCfromB} \leftarrow \text{colsums}(\text{dC}^T \odot \text{invFdC})$

**Compute approximate variances for  $A$ ,  $B$ , and  $C$ .**

- (1)  $\text{varA} \leftarrow \text{vcat}(\text{diag}(\text{invFa}_1), \dots, \text{diag}(\text{invFa}_J)) + \text{varAfromU} + \text{varAfromV}$
- (2)  $\text{varB} \leftarrow \text{vcat}(\text{diag}(\text{invFb}_1), \dots, \text{diag}(\text{invFb}_I)) + \text{varBfromU} + \text{varBfromV}$
- (3)  $\text{varC} \leftarrow \text{diag}(\text{invFc}) + \text{varCfromA} + \text{varCfromB}$

**Propagate uncertainty from  $(A, B, U, V)$  to  $S$ .**

First, we describe how to compute  $\text{varSfromU}$  and  $\text{varSfromV}$ .

- (1) Compute  $Q \in \mathbb{R}^{I \times J}$  where  $q_{ij} \leftarrow -w_{ij}e_{ij}/r_{ij}$ .
- (2) Compute  $P \in \mathbb{R}^{I \times J}$  where  $p_{ij} \leftarrow 2w_{ij}q_{ij}/\mu_{ij}$ .
- (3)  $\text{dgradS} \leftarrow QVD$
- (4)  $\text{dF} \leftarrow \text{dgradS} - PVD$
- (5)  $\text{dS} \leftarrow (-\text{invFs} \odot \text{dF} \odot \text{invFs} \odot \text{gradS}) + (\text{invFs} \odot \text{dgradS})$
- (6)  $\text{varSfromU} \leftarrow \text{rowsums}(\text{dS} \odot \text{reshape}(\text{varU}, M, I)^T \odot \text{dS})$
- (7) For  $i = 1, \dots, I$ ,
  - (a)  $\text{dgradS} \leftarrow Q_{i:}(DU_{i:})^T$
  - (b)  $\text{dF} \leftarrow \text{dgradS} - P_{i:}(DU_{i:})^T$
  - (c)  $\text{dS} \leftarrow -\text{invFs}_i \cdot \text{dF} \cdot \text{invFs}_i \cdot \text{gradS}_i + \text{invFs}_i \cdot \text{dgradS}$
  - (d)  $\text{varSfromV}_i \leftarrow \text{varV}^T \times \text{vec}((\text{dS} \odot \text{dS})^T)$
- (8)  $\text{varSfromV} \leftarrow (\text{varSfromV}_1, \dots, \text{varSfromV}_I)^T$

Next,  $\text{varSfromB}$  and  $\text{varSfromA}$  are computed in exactly the same way as  $\text{varSfromU}$  and  $\text{varSfromV}$ , respectively, but with  $X$ ,  $Z$ ,  $I$ ,  $\text{varB}$ , and  $\text{varA}$  in place of  $U$ ,  $V$ ,  $D$ ,  $\text{varU}$ , and  $\text{varV}$ , respectively.

**Propagate uncertainty from  $(A, B, U, V)$  to  $T$ .**

- (1) We compute  $\text{varTfromV}$  and  $\text{varTfromU}$  in exactly the same way as  $\text{varSfromU}$  and  $\text{varSfromV}$ , respectively, but with  $Y^T$ ,  $\mu^T$ ,  $W^T$ ,  $E^T$ ,  $r^T$ ,  $V$ ,  $U$ ,  $\text{gradT}$ ,  $\text{invFt}$ ,  $\text{varV}$ , and  $\text{varU}$  in place of  $Y$ ,  $\mu$ ,  $W$ ,  $E$ ,  $r$ ,  $U$ ,  $V$ ,  $\text{gradS}$ ,  $\text{invFs}$ ,  $\text{varU}$ , and  $\text{varV}$ , respectively.
- (2) We compute  $\text{varTfromA}$  and  $\text{varTfromB}$  in exactly the same way as  $\text{varSfromU}$  and  $\text{varSfromV}$ , respectively, but with  $Y^T$ ,  $\mu^T$ ,  $W^T$ ,  $E^T$ ,  $r^T$ ,  $Z$ ,  $X$ ,  $I$ ,  $\text{gradT}$ ,  $\text{invFt}$ ,  $\text{varA}$ ,

and `varB` in place of  $Y$ ,  $\mu$ ,  $W$ ,  $E$ ,  $r$ ,  $U$ ,  $V$ ,  $D$ , `gradS`, `invFs`, `varU`, and `varV`, respectively.

**Compute approximate standard errors.**

- (1)  $\hat{s}_A \leftarrow \text{reshape}(\text{sqrt}(\text{varA}), K, J)^T$
- (2)  $\hat{s}_B \leftarrow \text{reshape}(\text{sqrt}(\text{varB}), L, I)^T$
- (3)  $\hat{s}_C \leftarrow \text{reshape}(\text{sqrt}(\text{varC}), K, L)$
- (4)  $\hat{s}_U \leftarrow \text{reshape}(\text{sqrt}(\text{varU}), M, I)^T$
- (5)  $\hat{s}_V \leftarrow \text{reshape}(\text{sqrt}(\text{varV}), M, J)^T$
- (6)  $\hat{s}_S \leftarrow \text{sqrt}(\text{invFs} + \text{varSfromA} + \text{varSfromB} + \text{varSfromU} + \text{varSfromV})$
- (7)  $\hat{s}_T \leftarrow \text{sqrt}(\text{invFt} + \text{varTfromA} + \text{varTfromB} + \text{varTfromU} + \text{varTfromV})$

We do not attempt to provide standard errors for  $D$ , since it seems difficult to estimate  $D$  without significant bias. Note that here, we reshape the vectorized standard errors to matrices having the same dimensions as the corresponding components, for instance,  $\hat{s}_A$  has the same dimensions as  $A$ , namely  $J \times K$ .

## S9 Priors for regularization

We place independent normal priors on all the entries of  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ , and  $V$ , and in the NB-GBM, on the entries of  $S$  and  $T$  as well. Specifically, the prior is  $\pi(A, B, C, D, U, V, S, T) = \pi_a(A)\pi_b(B)\pi_c(C)\pi_d(D)\pi_u(U)\pi_v(V)\pi_s(S)\pi_t(T)$  where

$$\begin{aligned}
 \pi_a(A) &= \prod_{j,k} \mathcal{N}(a_{jk} \mid 0, \lambda_a^{-1}) & \pi_u(U) &= \prod_{i,m} \mathcal{N}(u_{im} \mid 0, \lambda_u^{-1}) \\
 \pi_b(B) &= \prod_{i,l} \mathcal{N}(b_{il} \mid 0, \lambda_b^{-1}) & \pi_v(V) &= \prod_{j,m} \mathcal{N}(v_{jm} \mid 0, \lambda_v^{-1}) \\
 \pi_c(C) &= \prod_{k,\ell} \mathcal{N}(c_{k\ell} \mid 0, \lambda_c^{-1}) & \pi_s(S) &= \prod_i \mathcal{N}(s_i \mid m_s, \lambda_s^{-1}) \\
 \pi_d(D) &= \prod_m \mathcal{N}(d_{mm} \mid 0, \lambda_d^{-1}) & \pi_t(T) &= \prod_j \mathcal{N}(t_j \mid m_t, \lambda_t^{-1}).
 \end{aligned} \tag{S9.1}$$

Thus, the log-prior is

$$\begin{aligned} \log \pi = \text{const} &- \frac{1}{2}\lambda_a \sum_{j,k} a_{jk}^2 - \frac{1}{2}\lambda_b \sum_{i,\ell} b_{i\ell}^2 - \frac{1}{2}\lambda_c \sum_{k,\ell} c_{k\ell}^2 - \frac{1}{2}\lambda_d \sum_m d_{mm}^2 \\ &- \frac{1}{2}\lambda_u \sum_{i,m} u_{im}^2 - \frac{1}{2}\lambda_v \sum_{j,m} v_{jm}^2 - \frac{1}{2}\lambda_s \sum_i (s_i - m_s)^2 - \frac{1}{2}\lambda_t \sum_j (t_j - m_t)^2. \end{aligned} \quad (\text{S9.2})$$

For the prior parameters, we use the following default settings:  $\lambda_a = \lambda_b = \lambda_c = \lambda_d = \lambda_u = \lambda_v = 1$ ,  $m_s = m_t = 0$ , and  $\lambda_s = \lambda_t = 1$ . These defaults are fairly generally applicable since they are acting on coefficients that are on the same scale in terms of units, due to the fact that we standardize the covariates to have zero mean and unit variance, that is,  $\frac{1}{I} \sum_{i=1}^I x_{ik} = 0$  and  $\frac{1}{I} \sum_{i=1}^I x_{ik}^2 = 1$  for all  $k \geq 2$  and  $\frac{1}{J} \sum_{j=1}^J z_{j\ell} = 0$  and  $\frac{1}{J} \sum_{j=1}^J z_{j\ell}^2 = 1$  for all  $\ell \geq 2$ . However, specific applications may call for departures from these defaults.

For the updates to  $G = UD$  and  $H = VD$  in the GBM estimation algorithm (Section S7), we use the priors on  $G$  and  $H$  induced by the priors on  $U$  and  $V$ , given  $D$ . First consider  $G$ . For any fixed  $D$ , the induced prior on  $g_{im} = u_{im}d_{mm}$  is  $\pi(g_{im}) = \mathcal{N}(g_{im} \mid 0, d_{mm}^2/\lambda_u)$ . Thus, given  $D$ , the prior on each row of  $G$  is  $G_{i:} \sim \mathcal{N}(0, \Lambda^{-1})$  where  $\Lambda = (D^2/\lambda_u)^{-1}$ . The gradient and Hessian of the log-prior on  $G_{i:}$  are therefore  $-\Lambda G_{i:}$  and  $-\Lambda$ , respectively. Hence, with this prior, the regularized Fisher scoring approach (Equation 3.1) yields the  $G$  update formulas used in the algorithm (Section S7). The  $H$  update is similar, except that the induced prior is  $H_{j:} \sim \mathcal{N}(0, \Lambda^{-1})$  where  $\Lambda = (D^2/\lambda_v)^{-1}$ . It seems reasonable to hold  $D$  fixed when computing the induced priors on  $G$  and  $H$ , rather than integrating it out, since  $D$  tends to be more accurately estimated than  $U$  or  $V$  (in terms of relative MSE), presumably due to the fact that  $D$  has only  $M$  nonzero entries, each of which is informed by all of the data; see Figure S1 for an empirical example.

## S10 Proofs

### S10.1 Identifiability and interpretation

*Proof of Theorem 5.1.* Left-multiplying both sides of Equation 5.1 by  $X^T$ , we have

$$X^T X A_1^T + X^T X C_1 Z^T = X^T X A_2^T + X^T X C_2 Z^T \quad (\text{S10.1})$$

by Condition 2.1(b). Since  $X^T X$  is invertible, this implies

$$A_1^T + C_1 Z^T = A_2^T + C_2 Z^T. \quad (\text{S10.2})$$

Right-multiplying Equation S10.2 by  $Z$ , we have  $C_1 Z^T Z = C_2 Z^T Z$  by Condition 2.1(b). Since  $Z^T Z$  is invertible, this implies that  $C_1 = C_2$ . Plugging  $C_1 = C_2$  into Equation S10.2 yields  $A_1 = A_2$ . Plugging  $A_1 = A_2$  and  $C_1 = C_2$  into Equation 5.1, we have

$$B_1 Z^T + U_1 D_1 V_1^T = B_2 Z^T + U_2 D_2 V_2^T. \quad (\text{S10.3})$$

Right-multiplying Equation S10.3 by  $Z$ , using Condition 2.1(b), and using the fact that  $Z^T Z$  is invertible, we obtain  $B_1 = B_2$ . This implies that  $U_1 D_1 V_1^T = U_2 D_2 V_2^T$ . By the uniqueness properties of the singular value decomposition, Conditions 2.1(c) and 2.1(d) imply that  $D_1 = D_2$ ,  $U_1 = U_2 S$ , and  $V_1^T = S V_2^T$  for a diagonal matrix  $S$  of the form  $S = \text{Diag}(\pm 1, \dots, \pm 1)$  (Blum et al., 2020). By Condition 2.1(e),  $S = I$ . Therefore,  $U_1 = U_2$  and  $V_1 = V_2$ . This proves that  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $U$ , and  $V$  are uniquely determined by  $\mathbb{E}(\mathbf{Y})$  for any given  $X$ ,  $Z$ ,  $M$ .  $\square$

*Proof of Theorem 5.2.* First,  $\sum_{j=1}^J a_{jk} = 0$  follows from the fact that  $z_{j1} = 1$  for all  $j$  by Condition 2.2(a) and  $Z^T A = 0$  by Condition 2.1(b). Likewise,  $\sum_{i=1}^I b_{i\ell} = 0$  follows from  $x_{i1} = 1$  and  $X^T B = 0$ . In the same way,  $\sum_{i=1}^I u_{im} = 0$  and  $\sum_{j=1}^J v_{jm} = 0$  since  $x_{i1} = 1$ ,  $z_{j1} = 1$ ,  $X^T U = 0$ , and  $Z^T V = 0$ .



When Condition 2.2(a) holds, we can rearrange Equation 2.1 as

$$\begin{aligned}
g(\mu_{ij}) &= c_{11} + a_{j1} + b_{i1} + \sum_{k=2}^K (c_{k1} + a_{jk})x_{ik} + \sum_{\ell=2}^L (c_{1\ell} + b_{i\ell})z_{j\ell} \\
&\quad + \sum_{k=2}^K \sum_{\ell=2}^L c_{k\ell}x_{ik}z_{j\ell} + \sum_{m=1}^M u_{im}d_{mm}v_{jm}.
\end{aligned} \tag{S10.4}$$

Averaging Equation S10.4 over all  $i$ , and using these sum-to-zero properties (specifically, using that  $\sum_{i=1}^I x_{ik} = 0$  for  $k \geq 2$ ,  $\sum_{i=1}^I b_{i\ell} = 0$ , and  $\sum_{i=1}^I u_{im} = 0$ ),

$$\begin{aligned}
\frac{1}{I} \sum_{i=1}^I g(\mu_{ij}) &= c_{11} + a_{j1} + \left(\frac{1}{I} \sum_i b_{i1}\right) + \sum_{k=2}^K (c_{k1} + a_{jk}) \left(\frac{1}{I} \sum_i x_{ik}\right) + \sum_{\ell=2}^L (c_{1\ell} + \left(\frac{1}{I} \sum_i b_{i\ell}\right)) z_{j\ell} \\
&\quad + \sum_{k=2}^K \sum_{\ell=2}^L c_{k\ell} \left(\frac{1}{I} \sum_i x_{ik}\right) z_{j\ell} + \sum_{m=1}^M \left(\frac{1}{I} \sum_i u_{im}\right) d_{mm} v_{jm} \\
&= c_{11} + a_{j1} + \sum_{\ell=2}^L c_{1\ell} z_{j\ell}.
\end{aligned} \tag{S10.5}$$

In the same way, averaging Equation S10.4 over all  $j$  (and using that  $\sum_{j=1}^J z_{j\ell} = 0$  for  $\ell \geq 2$ ,  $\sum_{j=1}^J a_{jk} = 0$ , and  $\sum_{j=1}^J v_{jm} = 0$ ), we have

$$\frac{1}{J} \sum_{j=1}^J g(\mu_{ij}) = c_{11} + b_{i1} + \sum_{k=2}^K c_{k1} x_{ik}.$$

Finally, averaging Equation S10.5 over all  $j$ , we have  $\frac{1}{IJ} \sum_{i=1}^I \sum_{j=1}^J g(\mu_{ij}) = c_{11}$ .  $\square$

*Proof of Theorem 5.3.* For any  $Q \in \mathbb{R}^{m \times n}$ , we have  $\text{SS}(Q) = \text{tr}(Q^T Q)$ , where  $\text{tr}(\cdot)$  denotes the trace. Define  $Q = XA^T + BZ^T + XCZ^T + UDV^T$ . By using  $X^T B = 0$ ,  $Z^T A = 0$ ,  $X^T U = 0$ , and  $Z^T V = 0$ , we have that

$$\begin{aligned}
Q^T Q &= (AX^T + ZB^T + ZC^T X^T + VDU^T)(XA^T + BZ^T + XCZ^T + UDV^T) \\
&= AX^T XA + AX^T XCZ^T + ZB^T BZ^T + ZB^T UDV^T + ZC^T X^T XA^T \\
&\quad + ZC^T X^T XCZ^T + VDU^T BZ^T + VDU^T UDV^T.
\end{aligned}$$

By the cyclic property of the trace,

$$\begin{aligned}\operatorname{tr}(AX^T X CZ^T) &= \operatorname{tr}(XCZ^T AX^T) = 0, \\ \operatorname{tr}(ZB^T UDV^T) &= \operatorname{tr}(B^T UDV^T Z) = 0, \\ \operatorname{tr}(ZC^T X^T X A^T) &= \operatorname{tr}(X A^T ZC^T X^T) = 0, \\ \operatorname{tr}(VDU^T BZ^T) &= \operatorname{tr}(BZ^T VDU^T) = 0.\end{aligned}$$

Therefore, by the linearity of the trace,

$$\begin{aligned}\operatorname{SS}(Q) &= \operatorname{tr}(Q^T Q) = \operatorname{tr}(AX^T X A^T) + \operatorname{tr}(ZB^T BZ^T) + \operatorname{tr}(ZC^T X^T X CZ^T) + \operatorname{tr}(VDU^T UDV^T) \\ &= \operatorname{SS}(X A^T) + \operatorname{SS}(BZ^T) + \operatorname{SS}(XCZ^T) + \operatorname{SS}(UDV^T).\end{aligned}$$

□

## S10.2 Likelihood-preserving projections

*Proof of Theorem 5.4.* (1.) For the projection of  $\tilde{A}$ , plugging in the definitions of  $A_1$  and  $C_1$ , we have

$$X A_1^T + X C_1 Z^T = X \tilde{A}^T - X(Z^+ \tilde{A})^T Z^T + X C Z^T + X(Z^+ \tilde{A})^T Z^T = X \tilde{A}^T + X C Z^T,$$

and therefore,  $\eta(A_1, B, C_1, D, U, V) = \eta(\tilde{A}, B, C, D, U, V)$ . To see that Condition 2.1 is satisfied, first note that  $Z^T(I - ZZ^+) = Z^T - Z^T Z(Z^T Z)^{-1} Z^T = 0$ , and therefore

$$Z^T A_1 = Z^T(\tilde{A} - ZZ^+ \tilde{A}) = Z^T(I - ZZ^+) \tilde{A} = 0.$$

(2.) Similarly, for the projection of  $\tilde{B}$ , we have  $B_1Z^T + XC_1Z^T = \tilde{B}Z^T + XCZ^T$  and  $X^TB_1 = 0$ . (3.) For the projection of  $\tilde{G}$ , we have

$$\begin{aligned} XA_1^T + XC_1Z^T + U_1D_1V_1^T &= XA_0^T - X(Z^+A_0)^T Z^T + XCZ^T + X(Z^+A_0)^T Z^T + G_0V^T \\ &= XA^T + X(X^+\tilde{G})V^T + XCZ^T + \tilde{G}V^T - X(X^+\tilde{G})V^T \\ &= XA^T + XCZ^T + \tilde{G}IV^T, \end{aligned}$$

and thus,  $\eta(A_1, B, C_1, D_1, U_1, V_1) = \eta(A, B, C, I, \tilde{G}, V)$ . To check that Condition 2.1 is satisfied, first observe that  $Z^TA_1 = Z^T(I - ZZ^+)A_0 = 0$  and  $X^TG_0 = X^T(I - XX^+)\tilde{G} = 0$ . Hence,

$$0 \stackrel{(a)}{=} X^TG_0V^TV_1D_1^{-1} \stackrel{(b)}{=} X^TU_1D_1V_1^TV_1D_1^{-1} \stackrel{(c)}{=} X^TU_1$$

where we have used (a)  $X^TG_0 = 0$ , (b)  $G_0V^T = U_1D_1V_1^T$ , and (c)  $V_1^TV_1 = I$  and  $D_1D_1^{-1} = I$ . Likewise, since  $V^TZ = 0$  by assumption,

$$0 = D_1^{-1}U_1^TG_0V^TZ = D_1^{-1}U_1^TU_1D_1V_1^TZ = V_1^TZ$$

since  $U_1^TU_1 = I$  and  $D_1^{-1}D_1 = I$ .

(4.) For the projection of  $\tilde{H}$ , in an altogether similar way, we have

$$B_1Z^T + XC_1Z^T + U_1D_1V_1^T = BZ^T + XCZ^T + UI\tilde{H}^T.$$

Further,  $X^TB_1 = X^T(I - XX^+)B_0 = 0$  and  $Z^TH_0 = 0$ , thus

$$\begin{aligned} 0 &= Z^TH_0U^TU_1D_1^{-1} = Z^TV_1D_1U_1^TU_1D_1^{-1} = Z^TV_1 \\ 0 &= X^TUH_0^TV_1D_1^{-1} = X^TU_1D_1V_1^TV_1D_1^{-1} = X^TU_1. \end{aligned}$$

□

Computationally, it is highly advantageous to structure the calculation of the projec-

tions in Theorem 5.4 as follows. First, one can precompute the pseudoinverses  $X^+$  and  $Z^+$  since  $X$  and  $Z$  are fixed throughout the algorithm. In the updates to  $A$  (or  $B$ ), it is advantageous to first compute  $Z^+\tilde{A}$  (or  $X^+\tilde{B}$ , respectively) in order to avoid explicitly computing and storing  $XX^+ \in \mathbb{R}^{I \times I}$  and  $ZZ^+ \in \mathbb{R}^{J \times J}$ . Likewise, in the projection of  $\tilde{G}$  (or  $\tilde{H}$ ), first compute  $X^+\tilde{G}$  and  $Z^+A_0$  (or  $Z^+\tilde{H}$  and  $X^+B_0$ , respectively). We use this approach in the step-by-step algorithm provided in Section S7.

The interpretation of the operations in Theorem 5.4 is as follows. For  $\tilde{A}$ , the idea is that  $A_1 = (I - ZZ^+)\tilde{A}$  is an orthogonal projection of the columns of  $\tilde{A}$  onto the nullspace of  $Z^T$ , and  $C_1$  is a shifted version of  $C$  to compensate for the shift from  $\tilde{A}$  to  $A_1$ . Likewise for  $\tilde{B}$ , but with  $X$  instead of  $Z$ . For  $\tilde{G}$ , the idea is that (a)  $G_0$  is a projection of  $\tilde{G}$  onto the nullspace of  $X^T$ , (b) the SVD enforces the orthonormality, ordering, and sign constraints on  $U_1$ ,  $D_1$ , and  $V_1$  while maintaining  $X^TU_1 = 0$  and  $Z^TV_1 = 0$ , (c)  $A_0$  compensates for the shift from  $\tilde{G}$  to  $G_0$ , (d)  $A_1$  projects  $A_0$  onto the nullspace of  $Z^T$ , and (e)  $C_1$  compensates for the shift from  $A_0$  to  $A_1$ . For  $\tilde{H}$ , the interpretation is similar.

### S10.3 Estimation time complexity

In this section, we justify the expressions in Section 5 giving the time complexity of the estimation algorithm as a function of  $I$ ,  $J$ ,  $K$ ,  $L$ , and  $M$ , assuming  $\max\{K^2, L^2, M\} \leq \min\{I, J\}$  (Equation 5.2). The outline of the estimation algorithm is in Section 3, and the step-by-step details are in Section S7. Denote  $x \wedge y = \min\{x, y\}$  and  $x \vee y = \max\{x, y\}$ . For the updates to each of  $A$ ,  $B$ ,  $C$ ,  $D$ ,  $UD$ ,  $VD$ ,  $S$ , and  $T$ , we report the computation time complexity after  $\eta$ ,  $\mu$ ,  $W$ , and  $E$  have been recomputed.

**Cost of preprocessing and initialization.** Precomputing the pseudoinverses  $X^+$  and  $Z^+$  takes  $O(IK^2)$  and  $O(JL^2)$  time, respectively; thus, both are  $O(IJ)$  by Equation 5.2. Computing  $C \leftarrow X^+\check{Y}(Z^+)^T$  takes  $O(IJ(K \wedge L))$  time,  $A \leftarrow (X^+\check{Y} - CZ^+)^T$  takes  $O(IJK)$  time, and  $B \leftarrow \check{Y}(Z^+)^T - XC$  takes  $O(IJL)$  time. Computing  $D$ ,  $U$ , and  $V$  takes  $O(IJM)$  time, since the truncated SVD of rank  $M$  for an  $I \times J$  matrix can be done in  $O(IJM)$  time

(Halko et al., 2011). Finally, each update to  $S$  and  $T$  takes  $O(IJ)$  time (see below). Thus, overall, preprocessing and initialization takes  $O(IJ(K \vee L \vee M))$  time.

**Cost of computing  $\eta$ ,  $\mu$ ,  $W$ , and  $E$ .** Computing  $\eta = XA^T + BZ^T + XCZ^T + UDV^T$  takes  $O(IJ(K \vee L \vee M))$  time, since  $XA^T$ ,  $BZ^T$ ,  $XCZ^T$ , and  $UDV^T$  take  $O(IJK)$ ,  $O(IJL)$ ,  $O(IJ(K \wedge L))$ , and  $O(IJM)$  time, respectively. Computing  $\mu$ ,  $W$ , and  $E$  takes  $O(IJ)$  time given  $\eta$ .

**Cost of updating  $A$ .** For each  $j$ , computing the Fisher scoring step takes  $O(IK^2)$  time, so altogether the  $J$  steps take  $O(IJK^2)$  time. For the projection, we compute  $Q \leftarrow Z^+A$  and  $ZQ$ , which takes  $O(JKL)$  time. By Equation 5.2, we have  $L \leq I$ , so the cost of computing the projection can be absorbed into the cost of the Fisher scoring steps.

**Cost of updating  $B$ .** By symmetry, this takes  $O(IJL^2)$  time.

**Cost of updating  $C$ .** Computing the Fisher information matrix  $F$  takes  $O(IJK^2 + JK^2L^2)$  time, which is  $O(IJK^2)$  since  $L^2 \leq I$  by Equation 5.2. Inverting  $F + \lambda_c I$  takes  $O((KL)^3) = O(IJKL)$  time (using Equation 5.2), and computing  $X^T E Z$  takes  $O(IJK)$  time, so the update to  $C$  can be done in  $O(IJ(K^2 \vee L^2))$  time.

**Cost of updating  $D$ .** Computing the Fisher information matrix  $F$  takes  $O(IJM^2)$  time, inverting  $F + \lambda_d I$  takes  $O(M^3)$  time, and computing  $\text{diag}(U^T E V)$  takes  $O(IJM)$  time. Thus, the update to  $D$  takes  $O(IJM^2)$  time.

**Cost of updating  $G = UD$ .** By comparison with the  $B$  update, the Fisher scoring steps cost  $O(IJM^2)$  time. The projection steps (except for the SVD) take  $O(IKM + JKM + JKL)$  time, and by Equation 5.2, this is  $O(IJM)$ . Computing the truncated SVD of rank  $M$  for an  $I \times J$  matrix can be done in  $O(IJM)$  time (Halko et al., 2011). Therefore, the cost of the projection can be absorbed into the Fisher scoring steps.

**Cost of updating  $H = VD$ .** By symmetry, this takes  $O(IJM^2)$  time.

**Cost of updating  $S$  and  $T$ .** Given  $\mu$ , updating  $S$  and  $T$  takes  $O(IJ)$  time, since computing  $\delta$  and  $\delta'$  involves a loop over all  $i$  and  $j$ .

## S10.4 Inference time complexity

Here, we justify the expressions in Section 5 giving the time complexity of the inference algorithm, assuming  $\max\{K^2, L^2, M\} \leq \min\{I, J\}$  (Equation 5.2) and also assuming  $I \geq J$ . The outline of the inference algorithm is in Section 4.2, and the step-by-step details are in Section S8. Denote  $x \wedge y = \min\{x, y\}$  and  $x \vee y = \max\{x, y\}$ .

**Cost of preprocessing.** Computing  $\eta$ ,  $\mu$ ,  $W$ , and  $E$  takes  $O(IJ(K \vee L \vee M))$  time. Computing  $\text{gradA}$ ,  $\text{gradB}$ , and  $\text{gradC}$  takes  $O(IJK)$ ,  $O(IJL)$ , and  $O(IJK)$  time, respectively. All of the other preprocessing steps take  $O(IJ)$  time.

**Cost of computing conditional uncertainty for each component.** Computing  $\text{invFa}$ ,  $\text{invFb}$ , and  $\text{invFc}$  take  $O(IJK^2)$ ,  $O(IJL^2)$ , and  $O(IJ(K^2 \vee L^2))$  time, respectively. Both  $\text{invFu}$  and  $\text{invFv}$  take  $O(IJM^2)$  time, and  $\text{invFs}$  and  $\text{invFt}$  take  $O(IJ)$  time. Thus, overall, this part is  $O(IJ(K^2 \vee L^2 \vee M^2))$ .

**Cost of computing constraint Jacobians for  $U$  and  $V$ .** Computing  $J_u$  and  $J_v$  take  $O(I(KM^2 + M^3))$  and  $O(J(LM^2 + M^3))$  time, respectively.

**Cost of computing joint uncertainty in  $(U, V)$  accounting for constraints.** The most expensive steps are computing  $\text{FuvFFuv} \leftarrow \text{Fuv}^T \times \text{FFuv}$ ,

$$\mathbf{B} \leftarrow \begin{bmatrix} \mathbf{A} & \mathbf{Jv}^T \\ \mathbf{Jv} & 0 \end{bmatrix}^{-1},$$

and  $\mathbf{g} \leftarrow \text{colsums}(\text{FuvD} \odot (\mathbf{C} \times \text{FuvD}))$ , which take  $O(IJ^2M^3)$ ,  $O(J^3M^3)$ , and  $O(IJ^2M^3)$  time, respectively. Since we assume  $I \geq J$ , these are all  $O(IJ^2M^3)$ . It is tedious but straightforward to check that all of the other steps in this part take less than  $O(IJ^2M^3)$  time, assuming  $I \geq J$  and  $\max\{K^2, L^2, M\} \leq \min\{I, J\}$  (Equation 5.2).

**Cost of propagating uncertainty from  $U$  and  $V$  to  $A$  and  $B$ .** Computing  $\text{varAfromU}$  and  $\text{varAfromV}$  take  $O(IJK(K \vee M))$  and  $O(IJK^2M)$  time, respectively; combined, this is  $O(IJK^2M)$ . By symmetry,  $\text{varBfromU}$  and  $\text{varBfromV}$  take  $O(IJL^2M)$  time.

**Cost of propagating uncertainty from  $A$  and  $B$  to  $C$ .** First, consider comput-

ing `varCfromA`. Each step in the loop over  $j$  and  $k$  takes  $O(IK^2)$  time (since  $L^2 \leq I$ ), thus, computing `dC` takes  $O(IJK^3)$  time altogether. Computing `invFdC` takes  $O(JK^3L)$  time, and the last step is  $O(JK^2L)$ . Thus, overall, `varCfromA` takes  $O(IJK^3)$  time. By symmetry, `varCfromB` takes  $O(IJL^3)$  time.

**Cost of propagating uncertainty from  $(A, B, U, V)$  to  $S$ .** Computing `varSfromA`, `varSfromB`, `varSfromU`, and `varSfromV` take  $O(IJK)$ ,  $O(IJL)$ ,  $O(IJM)$ , and  $O(IJM)$  time, respectively. Thus, overall this takes  $O(IJ(K \vee L \vee M))$  time.

**Cost of propagating uncertainty from  $(A, B, U, V)$  to  $T$ .** By symmetry, `varTfromA`, `varTfromB`, `varTfromU`, and `varTfromV` takes  $O(IJ(K \vee L \vee M))$  time overall.

**Cost of computing approximate standard errors.** Given the approximate element-wise variances, this is only takes as much time as the dimension of the each of the parameter matrices/vectors; namely,  $O(JK)$ ,  $O(IL)$ ,  $O(KL)$ ,  $O(IM)$ ,  $O(JM)$ ,  $O(I)$ , and  $O(J)$  for each of  $A$ ,  $B$ ,  $C$ ,  $U$ ,  $V$ ,  $S$ , and  $T$ , respectively. Using Equation 5.2 to easily upper bound each of these shows that, overall, this is  $O(IJ)$ .

## References

- Agresti, A. *Foundations of Linear and Generalized Linear Models*. John Wiley & Sons, 2015.
- Aitchison, J. and Silvey, S. Maximum-likelihood estimation of parameters subject to restraints. *The Annals of Mathematical Statistics*, pages 813–828, 1958.
- Babadi, M., Lee, S. K., and Smirnov, A. N. GATK gCNV: accurate germline copy-number variant discovery from sequencing read-depth data. The International Conference on Probabilistic Programming (PROBPROG), Oct 2018.
- Benzécri, J. L’analyse des Correspondances. L’analyse des Données, Vol. 2. *Dunod. Paris*, 1973.
- Blum, A., Hopcroft, J., and Kannan, R. *Foundations of Data Science*. Cambridge University Press, 2020.
- Boyd, S. and Vandenberghe, L. *Convex Optimization*. Cambridge University Press, 2004.
- Buettner, F., Pratanwanich, N., McCarthy, D. J., Marioni, J. C., and Stegle, O. f-scLVM: scalable and versatile factor analysis for single-cell RNA-seq. *Genome Biology*, 18(1):212, 2017.
- Carroll, J. D., Pruzansky, S., and Kruskal, J. B. CANDELINC: A general approach to multidimensional analysis of many-way arrays with linear constraints on parameters. *Psychometrika*, 45(1):3–24, 1980.

- Carvalho, C. M., Chang, J., Lucas, J. E., Nevins, J. R., Wang, Q., and West, M. High-dimensional sparse factor modeling: applications in gene expression genomics. *Journal of the American Statistical Association*, 103(484):1438–1456, 2008.
- Chadoeuf, J. and Denis, J. B. Asymptotic variances for the multiplicative interaction model. *Journal of Applied Statistics*, 18(3):331–353, 1991.
- Choulakian, V. Generalized bilinear models. *Psychometrika*, 61(2):271–283, 1996.
- Cochran, W. The comparison of different scales of measurement for experimental results. *The Annals of Mathematical Statistics*, 14(3):205–216, 1943.
- Davies, P. and Tso, M. K.-S. Procedures for reduced-rank regression. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 31(3):244–255, 1982.
- de Falguerolles, A. GBMs: GLMs with bilinear terms. In *COMPSTAT*, pages 53–64. Springer, 2000.
- Denis, J.-B. and Gower, J. C. Asymptotic confidence regions for biadditive models: Interpreting genotype-environment interactions. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 45(4):479–493, 1996.
- Dorkenoo, K. and Mathieu, J.-R. Etude d’un modele factoriel d’analyse de la variance comme modele lineaire generalise. *Revue de Statistique Appliquée*, 41(2):43–57, 1993.
- Fisher, R. and Mackenzie, W. Studies in Crop Variation: The Manurial Response of Different Potato Varieties. *Journal of Agricultural Sciences*, 13:311–320, 1923.
- Freeman, G. Statistical methods for the analysis of genotype-environment interactions. *Heredity*, 31(3):339–354, 1973.
- Fromer, M., Moran, J. L., Chambert, K., Banks, E., Bergen, S. E., Ruderfer, D. M., Handsaker, R. E., McCarroll, S. A., O’Donovan, M. C., Owen, M. J., et al. Discovery and statistical genotyping of copy-number variation from whole-exome sequencing depth. *The American Journal of Human Genetics*, 91(4):597–607, 2012.
- Gabriel, K. R. Least squares approximation of matrices by additive and multiplicative models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 40(2):186–196, 1978.
- Gabriel, K. R. Generalised bilinear regression. *Biometrika*, 85(3):689–700, 1998.
- Gabriel, K. R. and Zamir, S. Lower rank approximation of matrices by least squares with any choice of weights. *Technometrics*, 21(4):489–498, 1979.
- Gauch, H. G. Jr. Model selection and validation for yield trials with interaction. *Biometrics*, pages 705–715, 1988.
- Gauch, H. G. Jr. Statistical analysis of yield trials by AMMI and GGE. *Crop Science*, 46(4):1488–1500, 2006.
- Gauch, H. G. Jr., Piepho, H.-P., and Annicchiarico, P. Statistical analysis of yield trials by AMMI and GGE: Further considerations. *Crop Science*, 48(3):866–889, 2008.



- Gilbert, N. Non-additive combining abilities. *Genetics Research*, 4(1):65–73, 1963.
- Gollob, H. F. A statistical model which combines features of factor analytic and analysis of variance techniques. *Psychometrika*, 33(1):73–115, 1968.
- Goodman, L. A. Simple models for the analysis of association in cross-classifications having ordered categories. *Journal of the American Statistical Association*, 74(367):537–552, 1979.
- Goodman, L. A. Association models and canonical correlation in the analysis of cross-classifications having ordered categories. *Journal of the American Statistical Association*, 76(374):320–334, 1981.
- Goodman, L. A. Some useful extensions of the usual correspondence analysis approach and the usual log-linear models approach in the analysis of contingency tables. *International Statistical Review/Revue Internationale de Statistique*, pages 243–270, 1986.
- Goodman, L. A. Measures, models, and graphical displays in the analysis of cross-classified data. *Journal of the American Statistical Association*, 86(416):1085–1111, 1991.
- Goodman, L. A. and Haberman, S. J. The analysis of nonadditivity in two-way analysis of variance. *Journal of the American Statistical Association*, 85(409):139–145, 1990.
- Gower, J. Discussion of the paper by van der Heijden, de Falguerolles and de Leeuw. *Applied Statistics*, 38:273–276, 1989.
- Greenacre, M. J. *Theory and applications of correspondence analysis*. London (UK) Academic Press, 1984.
- Halko, N., Martinsson, P.-G., and Tropp, J. A. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- Hoff, P. D. Multilinear tensor regression for longitudinal relational data. *The Annals of Applied Statistics*, 9(3):1169, 2015.
- Jiang, Y., Oldridge, D. A., Diskin, S. J., and Zhang, N. R. CODEX: A normalization and copy number variation detection method for whole exome sequencing. *Nucleic Acids Research*, 43(6):e39–e39, 2015.
- Jørgensen, B. Exponential dispersion models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 49(2):127–145, 1987.
- Killick, R., Fearnhead, P., and Eckley, I. A. Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*, 107(500):1590–1598, 2012.
- Krumm, N., Sudmant, P. H., Ko, A., O’Roak, B. J., Malig, M., Coe, B. P., Quinlan, A. R., Nickerson, D. A., and Eichler, E. E. Copy number variation detection and genotyping from exome sequence data. *Genome Research*, 22(8):1525–1532, 2012.
- Leek, J. T. and Storey, J. D. Capturing heterogeneity in gene expression studies by surrogate variable analysis. *PLoS Genetics*, 3(9):e161, 2007.

- Leek, J. T. and Storey, J. D. A general framework for multiple testing dependence. *Proceedings of the National Academy of Sciences*, 105(48):18718–18723, 2008.
- Mandel, J. Non-additivity in two-way analysis of variance. *Journal of the American Statistical Association*, 56(296):878–888, 1961.
- Mandel, J. The partitioning of interaction in analysis of variance. *Journal of Research of the National Bureau of Standards, Series B*, 73:309–328, 1969.
- Marchenko, V. A. and Pastur, L. A. Distribution of eigenvalues for some sets of random matrices. *Matematicheskii Sbornik*, 114(4):507–536, 1967.
- Perry, P. O. and Pillai, N. S. Degrees of freedom for combining regression with factor analysis. *arXiv preprint arXiv:1310.7269*, 2013.
- Price, A. L., Patterson, N. J., Plenge, R. M., Weinblatt, M. E., Shadick, N. A., and Reich, D. Principal components analysis corrects for stratification in genome-wide association studies. *Nature Genetics*, 38(8):904–909, 2006.
- Risso, D., Ngai, J., Speed, T. P., and Dudoit, S. Normalization of RNA-seq data using factor analysis of control genes or samples. *Nature Biotechnology*, 32(9):896–902, 2014.
- Silvey, S. D. The Lagrangian multiplier test. *The Annals of Mathematical Statistics*, 30(2):389–407, 1959.
- Silvey, S. D. *Statistical Inference*. CRC Press, 1975.
- Stegle, O., Parts, L., Durbin, R., and Winn, J. A Bayesian framework to account for complex non-genetic factors in gene expression levels greatly increases power in eQTL studies. *PLoS Comput Biol*, 6(5):e1000770, 2010.
- Sun, Y., Zhang, N. R., and Owen, A. B. Multiple hypothesis testing adjusted for latent variables, with an application to the AGEMAP gene expression data. *The Annals of Applied Statistics*, 6(4):1664–1688, 2012.
- Takane, Y. and Shibayama, T. Principal component analysis with external information on both subjects and variables. *Psychometrika*, 56(1):97–120, 1991.
- Townes, F. W. Generalized principal component analysis. *arXiv preprint arXiv:1907.02647*, 2019.
- Tukey, J. W. One degree of freedom for non-additivity. *Biometrics*, 5(3):232–242, 1949.
- Tukey, J. W. The future of data analysis. *The Annals of Mathematical Statistics*, 33(1):1–67, 1962.
- Van Eeuwijk, F. A. Multiplicative interaction in generalized linear models. *Biometrics*, pages 1017–1032, 1995.
- Williams, E. J. The interpretation of interactions in factorial experiments. *Biometrika*, 39(1-2):65–81, 1952.