

[Description](#)[Remarks and examples](#)[References](#)[Also see](#)

Description

The methods and formulas for the `gsem` command are presented below.

Remarks and examples

stata.com

Remarks are presented under the following headings:

*Introduction**Families of distributions**The Bernoulli family**The binomial family**The ordinal family**The multinomial family**The Poisson family**The negative binomial family**The gamma family**The Gaussian family**Reliability**Link functions**The logit link**The probit link**The complementary log-log link**The log link**The identity link**The likelihood**Gauss–Hermite quadrature**Adaptive quadrature**Laplacian approximation**Postestimation**Empirical Bayes**Other predictions*

Introduction

`gsem` fits generalized linear models with latent variables via maximum likelihood. Here is a table identifying the family/link combinations that `gsem` allows.

	logit	probit	cloglog	log	identity
Bernoulli	x	x	x		
binomial	x	x	x		
ordinal	x	x	x		
multinomial	x				
Poisson				x	
negative binomial				x	
gamma				x	
Gaussian				x	x

Log-likelihood calculations for fitting any model with latent variables require integrating out the latent variables. One widely used modern method is to directly estimate the integral required to calculate the log likelihood by Gauss–Hermite quadrature or some variation thereof. `gsem` implements four different methods for numerically evaluating the integrals.

1. Gauss–Hermite quadrature (GHQ)
2. Mean-variance adaptive quadrature (MVAGH)
3. Mode-curvature adaptive quadrature (MCAGH)
4. Laplacian approximation

The default method is MVAGH. The numerical integration method for MVAGH is based on [Rabe-Hesketh, Skrondal, and Pickles \(2005\)](#), and the other numerical integration methods described in this manual entry are based on [Skrondal and Rabe-Hesketh \(2004, chap. 6.3\)](#).

Families of distributions

`gsem` implements the most commonly used distribution families associated with generalized linear models. `gsem` also implements distributions for ordinal and multinomial outcomes.

In this manual entry, observed endogenous variables are also known as generalized responses or generalized outcomes, but we will simply refer to them as responses or outcomes. The random variable corresponding to a given response will be denoted by Y . An observed value of Y will be denoted by y , and the expected value of Y by μ . For the ordinal and multinomial families, we will refer to a linear prediction, denoted by z , instead of the expected value.

The Bernoulli family

The Bernoulli family is a binary response model. The response Y is assumed to take on the values 0 or 1; however, `gsem` allows any nonzero and nonmissing value to mean 1.

The log of the conditional probability mass function is

$$\log f(y|\mu) = y \log \mu + (1 - y) \log(1 - \mu)$$

where μ is also known as the probability of a success. The default link for the Bernoulli family is the `logit` link.

The binomial family

The binomial family is a count response model and generalizes the Bernoulli family by taking the sum of k independent Bernoulli outcomes. The response Y is assumed to take on the values $0, 1, \dots, k$.

The log of the conditional probability mass function is

$$\begin{aligned} \log f(y|\mu) &= \log\{\Gamma(k+1)\} - \log\{\Gamma(y+1)\} - \log\{\Gamma(k-y+1)\} \\ &\quad + y \log \mu + (1-y) \log(1-\mu) \end{aligned}$$

where μ is the expected value for a single Bernoulli outcome. The default link for the binomial family is the `logit` link.

The ordinal family

The ordinal family is a discrete response model. The response Y is assumed to take on one of k unique values. The actual values are irrelevant except that higher values are assumed to correspond to “higher” outcomes. Without loss of generality, we will assume that Y takes on the values $1, \dots, k$. The ordinal family with k outcomes has cutpoints $\kappa_0, \kappa_1, \dots, \kappa_k$, where $\kappa_0 = -\infty$, $\kappa_y < \kappa_{y+1}$, and $\kappa_k = +\infty$.

Given a linear prediction z , the probability that a random response Y takes the value y is

$$\Pr(Y = y|z) = \Pr(Y^* < \kappa_y - z) - \Pr(Y^* < \kappa_{y-1} - z)$$

where Y^* is the underlying stochastic component for Y . The distribution for Y^* is determined by the link function. `gsem` allows `logit`, `probit`, and `cloglog` for the ordinal family. The `logit` link assigns Y^* the extreme value distribution that is synonymous with the logit link for Bernoulli outcomes. The `probit` link assigns Y^* the standard normal distribution that is synonymous with the probit link for Bernoulli outcomes. The `cloglog` link assigns Y^* the distribution that is synonymous with the complementary log-log link for Bernoulli outcomes. The default link for the ordinal family is the `logit` link.

The multinomial family

The multinomial family is a discrete response model. The response Y is assumed to take on one of k unique values. The actual values are irrelevant and order does not matter; however, `gsem` requires that the values are nonnegative integers. Without loss of generality, we will assume that Y takes on the values $1, \dots, k$. Each of the k outcomes has its own linear prediction. For the model to be identified, one of the outcomes is chosen to be the base or reference. The linear prediction for the base outcome is constrained to be 0 for all observations. Without loss of generality, we will assume the base outcome is the first outcome. Let z_i be the prediction for outcome i , where $z_1 = 0$ for the base outcome.

Given the k linear predictions $\mathbf{z}' = (z_1, z_2, \dots, z_k)$, the log of the conditional probability mass function is

$$\log f(y|\mathbf{z}) = z_y - \log \left\{ \sum_{i=1}^k \exp(z_i) \right\}$$

The only link allowed for the multinomial family is the `logit` link.

The Poisson family

The Poisson family is a count-data response model. The response Y is assumed to take on nonnegative integer values.

The log of the conditional probability mass function is

$$\log f(y|\mu) = -\mu + y \log \mu - \log \Gamma(y + 1)$$

The only link allowed for the Poisson family is the [log](#) link.

The negative binomial family

The negative binomial family is another count-data response model. It is commonly thought of as a Poisson family with overdispersion. `gsem` allows two parameterizations for the dispersion in this family: mean dispersion and constant dispersion.

The log of the conditional probability mass function is

$$\begin{aligned} \log f(y|\mu, \alpha) = & \log\{\Gamma(y + m)\} - \log\{\Gamma(y + 1)\} - \log\{\Gamma(m)\} \\ & + m \log p + y \log(1 - p) \end{aligned}$$

where m and p depend on the form of dispersion.

The only link allowed for the negative binomial family is the [log](#) link.

For mean dispersion, we have

$$\begin{aligned} m &= 1/\alpha \\ p &= \frac{1}{1 + \alpha\mu} \end{aligned}$$

where μ is the expected value of Y and α is the scale parameter. `gsem` fits α in the log scale.

For constant dispersion, we have

$$\begin{aligned} m &= \exp(\log \mu - \log \delta) \\ p &= \frac{1}{1 + \delta} \end{aligned}$$

where μ is the expected value of Y and δ is the scale parameter. `gsem` fits δ in the log scale.

The gamma family

The gamma family is a continuous response model. The response Y is assumed to be a nonnegative real value.

The log of the conditional probability density function is

$$\log f(y|\mu, s) = -\log \Gamma(s^{-2}) - \left(\log \mu + 2 \log s + \frac{y}{\mu} \right) s^{-2} + (s^{-2} - 1) \log y$$

where μ is the expected value of Y and s is the scale parameter. `gsem` fits s in the log scale.

The only link allowed for the gamma family is the [log](#) link.

The Gaussian family

The Gaussian family is a continuous response model and is synonymous with the normal distribution.

When the Gaussian family is specified with the `identity` link but no censoring, `gsem` fits this family by using a single multivariate density function and allows the following two special features:

1. `gsem` can fit covariances between the Gaussian error variables.
2. `gsem` can fit paths between Gaussian responses, including nonrecursive systems.

The log of the conditional probability density function is

$$\log f(\mathbf{y}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) = -\frac{1}{2} \{d \log 2\pi + \log |\boldsymbol{\Sigma}| + (\mathbf{y} - \boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \boldsymbol{\mu})\}$$

where d is the dimension of the observed response vector \mathbf{y} , $\boldsymbol{\mu}$ is the mean of the responses, and $\boldsymbol{\Sigma}$ is the variance matrix of their unexplained errors.

When the Gaussian family is specified with the `log` link or censoring, the two special features described above no longer apply. In addition, the multivariate density function is no longer used. Instead, for each response using the `log` link, the log of the conditional probability density function corresponds to the formula above with $d = 1$. For censored responses, the log likelihood corresponds to the one in the [Methods and formulas](#) for [R] `intreg`.

Reliability

For a given Gaussian response variable with the `identity` link, the reliability Y may be specified as p or $100 \times p\%$. The variance of Y 's associated error variable is then constrained to $(1 - p)$ times the observed variance of Y .

Link functions

Except for the ordinal and multinomial families, the link function defines the transformation between the mean and the linear prediction for a given response. If Y is the random variable corresponding to an observed response variable y , then the link function performs the transformation

$$g(\mu) = z$$

where $\mu = E(Y)$ and z is the linear prediction. In practice, the likelihood evaluator function uses the inverse of the link function to map the linear prediction to the mean.

The logit link

The logit link is

$$g(\mu) = \log \mu - \log(1 - \mu)$$

and its inverse is

$$\mu = g^{-1}(z) = \frac{1}{1 + e^{-z}}$$

The probit link

The probit link is

$$g(\mu) = \Phi^{-1}(\mu)$$

and its inverse is

$$\mu = g^{-1}(z) = \Phi(z)$$

where $\Phi(\cdot)$ is the cumulative distribution function for the standard normal distribution and $\Phi^{-1}(\cdot)$ is its inverse.

The complementary log-log link

The complementary log-log link is

$$g(\mu) = \log\{-\log(1 - \mu)\}$$

and its inverse is

$$\mu = g^{-1}(z) = 1 - \exp\{-\exp(z)\}$$

The log link

The log link is

$$g(\mu) = \log \mu$$

and its inverse is

$$\mu = g^{-1}(z) = e^z$$

The identity link

The identity link is $g(\mu) = \mu$.

The likelihood

`gsem` fits generalized linear models with latent variables via maximum likelihood. The likelihood for the specified model is derived under the assumption that each response variable is independent and identically distributed across the estimation sample. The response variables are also assumed to be independent of each other. These assumptions are conditional on the latent variables and the observed exogenous variables.

The likelihood is computed by integrating out the latent variables. Let $\boldsymbol{\theta}$ be the vector of model parameters, \mathbf{y} be the vector of observed response variables, \mathbf{x} be the vector of observed exogenous variables, and \mathbf{u} be the $r \times 1$ vector of latent variables. Then the marginal likelihood looks something like

$$\mathcal{L}(\boldsymbol{\theta}) = \int_{\mathfrak{R}^r} f(\mathbf{y}|\mathbf{x}, \mathbf{u}, \boldsymbol{\theta})\phi(\mathbf{u}|\boldsymbol{\mu}_u, \boldsymbol{\Sigma}_u)\partial\mathbf{u}$$

where \mathfrak{R} denotes the set of values on the real line, \mathfrak{R}^r is the analog in r -dimensional space, $\boldsymbol{\theta}$ is a vector of the unique model parameters, $f(\cdot)$ is the conditional probability density function for the observed response variables, $\phi(\cdot)$ is the multivariate normal density for \mathbf{u} , $\boldsymbol{\mu}_u$ is the expected value of \mathbf{u} , and $\boldsymbol{\Sigma}_u$ is the covariance matrix for \mathbf{u} . All auxiliary parameters are fit directly without any further parameterization, so we simply acknowledge that the auxiliary parameters are among the elements of $\boldsymbol{\theta}$.

The \mathbf{y} variables are assumed to be independent, conditionally on \mathbf{x} and \mathbf{u} , so $f(\cdot)$ is the product of the individual conditional densities. One exception to this is when \mathbf{y} contains two or more Gaussian response variables with the identity link, in which case the Gaussian responses are actually modeled using a multivariate normal density to allow for correlated errors and nonrecursive systems among Gaussian responses. This one exception does not change how the integral is numerically evaluated, so we make no effort to represent this distinction in the formulas.

For a single-level model with n response variables, the conditional joint density function for a given observation is

$$f(\mathbf{y}|\mathbf{x}, \mathbf{u}, \boldsymbol{\theta}) = \prod_{i=1}^n f_i(y_i|\mathbf{x}, \mathbf{u}, \boldsymbol{\theta})$$

For a two-level model, the likelihood is computed at the cluster level, so the conditional density is also a product of the observation-level density contributions within a given cluster

$$f(\mathbf{y}|\mathbf{x}, \mathbf{u}, \boldsymbol{\theta}) = \prod_{i=1}^n \prod_{j=1}^t f_i(y_{ij}|\mathbf{x}_j, \mathbf{u}, \boldsymbol{\theta})$$

where t is the number of individuals within the cluster. This extends to more levels by expanding the products down to the observations nested within the hierarchical groups. Because the single-level model is a special case of a two-level model where all the groups have a single observation, we will now use the two-level notation and subscripts.

Except for the ordinal and multinomial families, we use the link function to map the conditional mean

$$\mu_{ij} = E(y_{ij}|\mathbf{x}_j, \mathbf{u})$$

to the linear prediction

$$z_{ij} = \mathbf{x}'_j\boldsymbol{\beta}_i + \mathbf{x}'_j\boldsymbol{\Lambda}_i\mathbf{u}$$

where $\boldsymbol{\beta}_i$ is the vector of the fixed-effect coefficients and $\boldsymbol{\Lambda}_i$ is the matrix of the latent loadings for y_{ij} . For notational convenience, we will overload the definitions of $f(\cdot)$ and $f_i(\cdot)$ so that they are functions of the responses and model parameters through the linear predictions $\mathbf{z}' = (z_1, \dots, z_n)$. Thus $f(\mathbf{y}|\mathbf{x}, \mathbf{u}, \boldsymbol{\theta})$ is equivalently specified as $f(\mathbf{y}, \mathbf{z}, \boldsymbol{\theta})$, and $f_i(y_{ij}|\mathbf{x}_j, \mathbf{u}, \boldsymbol{\theta})$ is equivalently specified as $f_i(y_{ij}, z_{ij}, \boldsymbol{\theta})$. In this new notation, the likelihood for a given cluster is

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}) &= \int_{\mathfrak{R}^r} f(\mathbf{y}, \mathbf{z}, \boldsymbol{\theta}) \phi(\mathbf{u} | \boldsymbol{\mu}_u, \boldsymbol{\Sigma}_u) \partial \mathbf{u} \\ &= \frac{1}{(2\pi)^{r/2} \sqrt{|\boldsymbol{\Sigma}_u|}} \int_{\mathfrak{R}^r} \exp \left\{ \log f(\mathbf{y}, \mathbf{z}, \boldsymbol{\theta}) - \frac{1}{2} (\mathbf{u} - \boldsymbol{\mu}_u)' \boldsymbol{\Sigma}_u^{-1} (\mathbf{u} - \boldsymbol{\mu}_u) \right\} \partial \mathbf{u} \end{aligned} \quad (1)$$

`gsem` allows nonrecursive systems between Gaussian response variables with the identity link, but non-Gaussian responses and Gaussian responses with the log link are not allowed to participate in any nonrecursive systems. This means that if a given response y is specified with a family other than Gaussian or a link other than identity, then y cannot have a path that ultimately leads back to itself. Any response may participate in a recursive system because the participating responses may be treated as exogenous variables when predicting other responses in a recursive system.

The latent vector \mathbf{u} consists of stacked collections of the latent variables from each level. Within each level, the latent endogenous variables $\boldsymbol{\eta}$ are stacked over the latent exogenous variables $\boldsymbol{\xi}$. Within a given level, the latent exogenous variables and latent endogenous errors are assumed independent and multivariate normal

$$\begin{aligned} \boldsymbol{\xi} &\sim N(\boldsymbol{\kappa}, \boldsymbol{\Phi}) \\ \boldsymbol{\epsilon} &\sim N(\mathbf{0}, \boldsymbol{\Psi}) \end{aligned}$$

so according to the linear relationship

$$\boldsymbol{\eta} = \mathbf{B}\boldsymbol{\eta} + \boldsymbol{\Gamma}\boldsymbol{\xi} + \mathbf{A}\mathbf{x} + \boldsymbol{\epsilon}$$

we have that the latent variables are jointly multivariate normal. This linear relationship implies that `gsem` allows latent variables to predict each other, but only within level. It also means that `gsem` allows paths from observed variables to latent variables; however, the observed variable must be constant within group if the path is to a group-level latent variable.

For our two-level model, we have

$$\mathbf{u} \sim N(\boldsymbol{\mu}_u, \boldsymbol{\Sigma}_u)$$

where

$$\begin{aligned} \boldsymbol{\mu}_u &= \begin{pmatrix} \boldsymbol{\mu}_\eta \\ \boldsymbol{\kappa} \end{pmatrix} & \boldsymbol{\Sigma}_u &= \begin{pmatrix} \boldsymbol{\Sigma}_{\eta\eta} & \boldsymbol{\Sigma}_{\eta\xi} \\ \boldsymbol{\Sigma}_{\xi\eta} & \boldsymbol{\Phi} \end{pmatrix} \\ \boldsymbol{\mu}_\eta &= (\mathbf{I} - \mathbf{B})^{-1} (\boldsymbol{\Gamma}\boldsymbol{\kappa} + \mathbf{A}\mathbf{x}) \\ \boldsymbol{\Sigma}_{\eta\eta} &= (\mathbf{I} - \mathbf{B})^{-1} (\boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Gamma}' + \boldsymbol{\Psi}) \{(\mathbf{I} - \mathbf{B})^{-1}\}' \\ \boldsymbol{\Sigma}_{\eta\xi} &= (\mathbf{I} - \mathbf{B})^{-1} \boldsymbol{\Gamma}\boldsymbol{\Phi} \end{aligned}$$

The vector $\boldsymbol{\theta}$ is therefore the set of unique model parameters taken from the following:

$\boldsymbol{\beta}_i$ is the vector of fixed-effect coefficients for y_{ij} .

$\boldsymbol{\Lambda}_i$ is the matrix of latent loadings for y_{ij} .

\mathbf{B} is the matrix of latent endogenous coefficients.

$\mathbf{\Gamma}$ is the matrix of latent exogenous coefficients.

\mathbf{A} is the matrix of latent fixed-effect coefficients.

$\boldsymbol{\kappa}$ is the vector of latent exogenous means.

$\mathbf{\Phi}$ is the matrix of latent exogenous variances and covariances.

$\mathbf{\Psi}$ is the matrix of latent endogenous error variances and covariances.

Auxiliary parameters that result from some of the distribution families.

Each level of a multilevel model will have its own set of the following parameters: \mathbf{B} , $\mathbf{\Gamma}$, \mathbf{A} , $\boldsymbol{\kappa}$, $\mathbf{\Phi}$, and $\mathbf{\Psi}$. For multilevel models, $\boldsymbol{\Sigma}_u$ is a block-diagonal matrix with a block for each level.

Gauss–Hermite quadrature

The integral in (1) is generally not tractable, so we must use numerical methods. In the univariate case, the integral of a function multiplied by the kernel of the standard normal distribution can be approximated using Gauss–Hermite quadrature (GHQ). For q -point GHQ, let the abscissa and weight pairs be denoted by (a_k^*, w_k^*) , $k = 1, \dots, q$. The GHQ approximation is then

$$\int_{-\infty}^{\infty} f(x) \exp(-x^2) dx \approx \sum_{k=1}^q w_k^* f(a_k^*)$$

Using the standard normal distribution yields the approximation

$$\int_{-\infty}^{\infty} f(x) \phi(x) dx \approx \sum_{k=1}^q w_k f(a_k)$$

where $a_k = \sqrt{2}a_k^*$ and $w_k = (w_k^*)/\sqrt{\pi}$.

We can use a change-of-variables technique to transform the multivariate integral (1) into a set of nested univariate integrals. Each univariate integral can then be evaluated using GHQ. Let \mathbf{v} be a random vector whose elements are independently standard normal, and let \mathbf{L} be the Cholesky decomposition of $\boldsymbol{\Sigma}_u$, that is, $\boldsymbol{\Sigma}_u = \mathbf{L}\mathbf{L}'$. In the distribution, we have that $\mathbf{u} = \boldsymbol{\mu}_u + \mathbf{L}\mathbf{v}$, and the linear predictions vector as a function of \mathbf{v} is

$$z_{ij} = \mathbf{x}'_j \boldsymbol{\beta}_i + \mathbf{x}'_j \boldsymbol{\Lambda}_i (\boldsymbol{\mu}_u + \mathbf{L}\mathbf{v})$$

so the likelihood for a given cluster is

$$\mathcal{L}(\boldsymbol{\theta}) = (2\pi)^{-r/2} \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} \exp \left\{ \log f(\mathbf{y}, \mathbf{z}, \boldsymbol{\theta}) - \frac{1}{2} \sum_{k=1}^r v_k^2 \right\} dv_1 \dots dv_r \quad (2)$$

where r is the number of latent variables.

Consider an r -dimensional quadrature grid containing q quadrature points in each dimension. Let the vector of abscissas $\mathbf{a}_k = (a_{k1}, \dots, a_{kr})'$ be a point in this grid, and let $\mathbf{w}_k = (w_{k1}, \dots, w_{kr})'$ be the vector of corresponding weights. The GHQ approximation to the likelihood for a given cluster is

$$\mathcal{L}^{\text{GHQ}}(\boldsymbol{\theta}) = \sum_{k_1=1}^q \dots \sum_{k_r=1}^q \left[\exp \left\{ \sum_{i=1}^n \log f_i(y_{ij}, z_{ijk}, \boldsymbol{\theta}) \right\} \prod_{s=1}^r w_{k_s} \right]$$

where

$$z_{ijk} = \mathbf{x}'_j \boldsymbol{\beta} + \mathbf{x}'_j \boldsymbol{\Lambda}_i (\boldsymbol{\mu}_u + \mathbf{L} \boldsymbol{\alpha}_k)$$

Adaptive quadrature

This section sets the stage for mean-variance adaptive Gauss–Hermite quadrature (MVAGH) and mode-curvature adaptive Gauss–Hermite quadrature (MCAGH).

Let's reconsider the likelihood in (2). If we fix the observed variables and the model parameters, we see that the posterior density for \mathbf{v} is proportional to

$$\phi(\mathbf{v}) f(\mathbf{y}, \mathbf{z}, \boldsymbol{\theta})$$

It is reasonable to assume that this posterior density can be approximated by a multivariate normal density with mean vector $\boldsymbol{\mu}_v$ and variance matrix $\boldsymbol{\tau}_v$. Instead of using the prior density of \mathbf{v} as the weighting distribution in the integral, we can use our approximation for the posterior density,

$$\mathcal{L}(\boldsymbol{\theta}) = \int_{\mathbb{R}^r} \frac{f(\mathbf{y}, \mathbf{z}, \boldsymbol{\theta}) \phi(\mathbf{v})}{\phi(\mathbf{v}, \boldsymbol{\mu}_v, \boldsymbol{\tau}_v)} \phi(\mathbf{v}, \boldsymbol{\mu}_v, \boldsymbol{\tau}_v) d\mathbf{v}$$

The likelihood is then approximated with

$$\mathcal{L}^*(\boldsymbol{\theta}) = \sum_{k_1=1}^q \dots \sum_{k_r=1}^q \left[\exp \left\{ \sum_{i=1}^n \log f_i(y_{ij}, z_{ijk}^*, \boldsymbol{\theta}) \right\} \prod_{s=1}^r \omega_{k_s} \right]$$

where

$$z_{ijk}^* = \mathbf{x}'_j \boldsymbol{\beta} + \mathbf{x}'_j \boldsymbol{\Lambda}_i (\boldsymbol{\mu}_u + \mathbf{L} \boldsymbol{\alpha}_k)$$

and $\boldsymbol{\alpha}_k$ and the ω_{k_s} are the adaptive versions of the abscissas and weights after an orthogonalizing transformation, which eliminates posterior covariances between the latent variables. $\boldsymbol{\alpha}_k$ and the ω_{k_s} are functions of $\boldsymbol{\alpha}_k$ and \mathbf{w}_k and the adaptive parameters $\boldsymbol{\mu}_v$ and $\boldsymbol{\tau}_v$.

For MVAGH, $\boldsymbol{\mu}_v$ is the posterior mean and $\boldsymbol{\tau}_v$ is the posterior variance of \mathbf{v} . They are computed iteratively by updating the posterior moments by using the MVAGH approximation, starting with a 0 mean vector and identity variance matrix.

For MCAGH, $\boldsymbol{\mu}_v$ is the posterior mode for \mathbf{v} and $\boldsymbol{\tau}_v$ is the curvature at the mode. They are computed by optimizing the integrand in (2) with respect to \mathbf{v} .

Laplacian approximation

Let's reconsider the likelihood in (1) and denote the argument in the exponential function by

$$\begin{aligned} h(\mathbf{u}) &= \log f(\mathbf{y}, \mathbf{z}, \boldsymbol{\theta}) - \frac{1}{2}(\mathbf{u} - \boldsymbol{\mu}_u)' \boldsymbol{\Sigma}_u^{-1}(\mathbf{u} - \boldsymbol{\mu}_u) \\ &= \sum_{i=1}^n \sum_{j=1}^t \log f_i(y_{ij}, z_{ij}, \boldsymbol{\theta}) - \frac{1}{2}(\mathbf{u} - \boldsymbol{\mu}_u)' \boldsymbol{\Sigma}_u^{-1}(\mathbf{u} - \boldsymbol{\mu}_u) \end{aligned}$$

where

$$z_{ij} = \mathbf{x}'_j \boldsymbol{\beta}_i + \mathbf{x}'_j \boldsymbol{\Lambda}_i \mathbf{u}$$

The Laplacian approximation is based on a second-order Taylor expansion of $h(\mathbf{u})$ about the value of \mathbf{u} that maximizes it. The first and second partials with respect to \mathbf{u} are

$$h'(\mathbf{u}) = \frac{\partial h(\mathbf{u})}{\partial \mathbf{u}} = \sum_{i=1}^n \sum_{j=1}^t \frac{\partial \log f_i(y_{ij}, z_{ij}, \boldsymbol{\theta})}{\partial z_{ij}} \boldsymbol{\Lambda}'_i \mathbf{x}_j - \boldsymbol{\Sigma}^{-1}(\mathbf{u} - \boldsymbol{\mu}_u)$$

$$\mathbf{H}(\mathbf{u}) = \frac{\partial^2 h(\mathbf{u})}{\partial \mathbf{u} \partial \mathbf{u}'} = \sum_{i=1}^n \sum_{j=1}^t \mathbf{x}'_j \boldsymbol{\Lambda}_i \frac{\partial^2 \log f_i(y_{ij}, z_{ij}, \boldsymbol{\theta})}{\partial z_{ij} \partial z_{ij}} \boldsymbol{\Lambda}'_i \mathbf{x}_j - \boldsymbol{\Sigma}^{-1}$$

The maximizer of $h(\mathbf{u})$ is $\hat{\mathbf{u}}$ such that $h'(\hat{\mathbf{u}}) = \mathbf{0}$. The integral in (1) is proportional to the posterior density of \mathbf{u} given the data, so $\hat{\mathbf{u}}$ is also the posterior mode.

The second-order Taylor approximation then takes the form

$$h(\mathbf{u}) \approx h(\hat{\mathbf{u}}) + \frac{1}{2}(\mathbf{u} - \hat{\mathbf{u}})' \mathbf{H}(\hat{\mathbf{u}})(\mathbf{u} - \hat{\mathbf{u}}) \quad (3)$$

because the first order derivative term is 0. The integral is approximated by

$$\int_{\mathfrak{R}^r} \exp\{h(\mathbf{u})\} d\mathbf{u} \approx \exp\{h(\hat{\mathbf{u}})\} (2\pi)^{r/2} |-\mathbf{H}(\hat{\mathbf{u}})|^{-1/2}$$

because the second term in (3) is the kernel of a multivariate normal density once it is exponentiated. The Laplacian approximation for the log likelihood is

$$\log \mathcal{L}^{\text{Lap}}(\boldsymbol{\theta}) = -\frac{1}{2} \log |\boldsymbol{\Sigma}_u| - \frac{1}{2} \log |-\mathbf{H}(\hat{\mathbf{u}})| + h(\hat{\mathbf{u}})$$

Postestimation

We begin by considering the prediction of the latent variables \mathbf{u} for a given cluster in a two-level model. Prediction of latent variables in multilevel generalized linear models involves assigning values to the latent variables, and there are many methods for doing so; see [Skrdonal and Rabe-Hesketh \(2009\)](#) and [Skrdonal and Rabe-Hesketh \(2004, chap. 7\)](#) for a comprehensive review. Stata offers two methods of predicting latent variables: empirical Bayes means (also known as posterior means) and empirical Bayes modes (also known as posterior modes). Below we provide more details about the two methods.

Empirical Bayes

Let $\widehat{\theta}$ denote the estimated model parameters. Empirical Bayes (EB) predictors of the latent variables are the means or modes of the empirical posterior distribution with the parameter estimates θ replaced with their estimates $\widehat{\theta}$. The method is called “empirical” because $\widehat{\theta}$ is treated as known. EB combines the prior information about the latent variables with the likelihood to obtain the conditional posterior distribution of latent variables. Using Bayes’ theorem, the empirical conditional posterior distribution of the latent variables for a given cluster is

$$\begin{aligned}\omega(\mathbf{u}|\mathbf{y}, \mathbf{x}; \widehat{\theta}) &= \frac{f(\mathbf{y}|\mathbf{u}, \mathbf{x}; \widehat{\theta}) \phi(\mathbf{u}; \widehat{\boldsymbol{\mu}}_u, \widehat{\boldsymbol{\Sigma}}_u)}{\int f(\mathbf{y}|\mathbf{u}, \mathbf{x}; \widehat{\theta}) \phi(\mathbf{u}; \widehat{\boldsymbol{\mu}}_u, \widehat{\boldsymbol{\Sigma}}_u) d\mathbf{u}} \\ &= \frac{f(\mathbf{y}|\mathbf{u}, \mathbf{x}; \widehat{\theta}) \phi(\mathbf{u}; \widehat{\boldsymbol{\mu}}_u, \widehat{\boldsymbol{\Sigma}}_u)}{\mathcal{L}(\widehat{\theta})}\end{aligned}$$

The denominator is just the likelihood contribution of the given cluster.

EB mean predictions of latent variables, $\widetilde{\mathbf{u}}$, also known as posterior means, are calculated as

$$\widetilde{\mathbf{u}} = \int_{\mathbb{R}^r} \mathbf{u} \omega(\mathbf{u}|\mathbf{y}, \mathbf{x}; \widehat{\theta}) d\mathbf{u}$$

where we use the notation $\widetilde{\mathbf{u}}$ rather than $\widehat{\mathbf{u}}$ to distinguish predicted values from estimates. This multivariate integral is approximated by MVAGH. If you have multiple latent variables within a level or latent variables across levels, the calculation involves orthogonalizing transformations with the Cholesky transformation because the latent variables are no longer independent under the posterior distribution.

When all the response variables are normal, the posterior density is multivariate normal, and EB means are also best linear unbiased predictors (BLUPs); see [Skrondal and Rabe-Hesketh \(2004, 227\)](#). In generalized mixed-effects models, the posterior density tends to multivariate normal as cluster size increases.

EB modal predictions can be approximated by solving for $\widetilde{\mathbf{u}}$ such that

$$\frac{\partial}{\partial \mathbf{u}} \log \omega(\mathbf{u}|\mathbf{y}, \mathbf{x}; \widehat{\theta}) \Big|_{\mathbf{u}=\widetilde{\mathbf{u}}} = \mathbf{0}$$

Because the denominator in $\omega(\cdot)$ does not depend on \mathbf{u} , we can omit it from the calculation to obtain the EB mode. The calculation of EB modes does not require numerical integration, and for that reason they are often used in place of EB means. As the posterior density gets closer to being multivariate normal, EB modes get closer and closer to EB means.

Just like there are many methods of assigning values to the random effects, there exist many methods of calculating standard errors of the predicted random effects; see [Skrondal and Rabe-Hesketh \(2009\)](#) for a comprehensive review.

Stata uses the posterior standard deviation as the standard error of the posterior means predictor of random effects. For a given level, the EB posterior covariance matrix of the random effects is given by

$$\text{Cov}(\widetilde{\mathbf{u}}|\mathbf{y}, \mathbf{x}; \widehat{\theta}) = \int_{\mathbb{R}^r} (\mathbf{u} - \widetilde{\mathbf{u}})(\mathbf{u} - \widetilde{\mathbf{u}})' \omega(\mathbf{u}|\mathbf{y}, \mathbf{x}; \widehat{\theta}) d\mathbf{u}$$

The posterior covariance matrix and the integrals are approximated by MVAGH.

Conditional standard errors for the estimated posterior modes are derived from standard theory of maximum likelihood, which dictates that the asymptotic variance matrix of $\tilde{\mathbf{u}}$ is the negative inverse of the Hessian matrix.

Other predictions

In what follows, we show formulas with the posterior means estimates of random effects $\tilde{\mathbf{u}}$, which are used by default or if the `means` option is specified. If the `modes` option is specified, $\tilde{\mathbf{u}}$ are simply replaced with the posterior modes $\tilde{\mathbf{u}}$ in these formulas.

For the i th response in the j th observation within a given cluster in a two-level model, the linear predictor is computed as

$$\hat{z}_{ij} = \mathbf{x}'_j \hat{\boldsymbol{\beta}} + \mathbf{x}'_j \hat{\boldsymbol{\Lambda}}_i \tilde{\mathbf{u}}$$

The linear predictor includes the offset or exposure variable if one was specified during estimation, unless the `nooffset` option is specified. If the `fixedonly` option is specified, the linear predictor is computed as

$$\hat{z}_{ij} = \mathbf{x}'_j \hat{\boldsymbol{\beta}}$$

The predicted mean, conditional on the predicted latent variables, is

$$\hat{\mu}_{ij} = g^{-1}(\hat{z}_{ij})$$

where $g^{-1}(\cdot)$ is the inverse link function defined in [Link functions](#) above. For the ordinal and multinomial families, the predicted mean is actually a probability, and `gsem` can produce a probability for each outcome value as described in [The ordinal family](#) and [The multinomial family](#) above.

References

- Rabe-Hesketh, S., A. Skrondal, and A. Pickles. 2005. Maximum likelihood estimation of limited and discrete dependent variable models with nested random effects. *Journal of Econometrics* 128: 301–323.
- Skrondal, A., and S. Rabe-Hesketh. 2004. *Generalized Latent Variable Modeling: Multilevel, Longitudinal, and Structural Equation Models*. Boca Raton, FL: Chapman & Hall/CRC.
- . 2009. Prediction in multilevel generalized linear models. *JRSSA* 172: 659–687.

Also see

[SEM] `gsem` — Generalized structural equation model estimation command