


Generalized Models: Part 2

- Topics:
 - **Models for skewed discrete outcomes**
 - Models for skewed continuous outcomes
 - Models for zero-inflated discrete or continuous outcomes
 - Complications for generalized multilevel models
 - Multivariate or multilevel data

3 Parts of Generalized Multilevel Models



1. Link Function (different from general): How the conditional mean of a non-normal outcome is made **unbounded** so that the model fixed and random effects can predict it linearly
2. Linear Predictor (same as in general): How the model predictors linearly relate to the outcome conditional mean
3. Model for Level-1 Residuals (different than general): how the level-1 residuals should be distributed given the sample space (possible values) of the actual outcome

A Taxonomy of Not-Normal Outcomes

- **“Discrete” outcomes**—all responses are **whole** numbers
 - **Categorical variables** in which **values are labels**, not amounts
 - Binomial (2 options) or multinomial (3+ options) distributions
 - Question: Are the values ordered → which link?
 - **Count of things that happened**, so values < 0 cannot exist
 - Sample space goes from 0 to positive infinity
 - Poisson or Negative Binomial distributions (usually)
 - Log link (usually) so predicted outcomes can't go below 0
 - Question: Are there *extra* 0 values? What to do about them?
- **“Continuous” outcomes**—responses can be **any** number
 - Question: What does the residual distribution look like?
 - Normal-ish? Skewed? Cut off? Mixture of different distributions?

A Revised Taxonomy

- Rather than just separating into discrete vs. continuous, think about models based on their shape AND kinds of data they fit
 - Note: You can use continuous models for discrete data (that only have integers), but not discrete models for continuous data (with decimals)
- 1. Skewed-looking distributions
 - Discrete: Poisson, Generalized Poisson, Negative Binomial (NB)
 - Continuous: Log-Normal, Beta, Gamma
- 2. Skewed with a pile of 0's: Becomes **If 0** and **How Much**
 - These models will differ in how they define the "If 0" part
 - Discrete: Zero-Inflated Poisson or NB, Hurdle Poisson or NB
 - Continuous: Two-Part (with normal or lognormal for the how much part)

Discrete Models for Count Outcomes

- Counts: non-negative integer unbounded responses
 - e.g., how many cigarettes did you smoke this week?
 - Traditionally uses natural log link so that predicted outcomes stay ≥ 0
- $g(\bullet)$ $\text{Log}(E(y_i)) = \text{Log}(\mu_i) = \text{model} \rightarrow$ predicts mean of y_i
- $g^{-1}(\bullet)$ $E(y_i) = \exp(\text{model}) \rightarrow$ to un-log it, use $\exp(\text{model})$
 - e.g., if $\text{Log}(\mu_i) = \text{model}$ provides predicted $\text{Log}(\mu_i) = 1.098$, that translates to an actual predicted count of $\exp(1.098) = 3$
 - e.g., if $\text{Log}(\mu_i) = \text{model}$ provides predicted $\text{Log}(\mu_i) = -5$, that translates to an actual predicted count of $\exp(-5) = 0.006738$
 - If people have different lengths of time in which counts were observed, you can include that info as an "offset" in the MODEL statement
- So that's how linear model predicts μ_i , the conditional mean for y_i , but what about residual variance?


Poisson Distribution for Residuals

- Poisson distribution has one parameter, λ , which is **both its mean and its variance** (so $\lambda = \text{mean} = \text{variance}$ in Poisson)

- $f(y_i|\lambda) = \text{Prob}(y_i = y) = \frac{\lambda^y * \exp(-\lambda)}{y!}$

$y!$ is factorial of y

- PDF: $\text{Prob}(y_i = y|\beta_0, \beta_1, \beta_2) = \frac{\mu_i^y * \exp(-\mu_i)}{y!}$



The dots indicate that only integer values are observed.

Distributions with a small expected value (mean or λ) are predicted to have a lot of 0's.

Once $\lambda > 6$ or so, the shape of the distribution is close to a that of a normal distribution.

3 potential problems for Poisson...

- The standard Poisson distribution is rarely sufficient, though
- **Problem #1: When mean \neq variance**
 - If variance $<$ mean, this leads to “under-dispersion” (not that likely)
 - If variance $>$ mean, this leads to “over-dispersion” (happens frequently)
- **Problem #2: When there are *no* 0 values**
 - Some 0 values are expected from count models, but in some contexts $y_i > 0$ always (but subtracting 1 won't fix it; need to adjust the model)
- **Problem #3: When there are *too many* 0 values (stay tuned)**
 - Some 0 values are expected from the Poisson and Negative Binomial models already, but many times there are even more 0 values observed than that
 - To fix it, there are two main options, depending on what you do to the 0's
- Each of these problems requires a model adjustment to fix it...

Problem #1: Variance $>$ mean = over-dispersion


- To fix it, we must add another parameter that allows the variance to exceed the mean... becomes a **Negative Binomial** distribution
 - Says residuals are a mixture of Poisson and gamma distributions, such that λ itself is a random variable with a gamma distribution
 - So expected mean is still given by λ , but the variance will differ from Poisson
- Model: $\text{Log}(y_i) = \text{Log}(\mu_i) = \beta_0 + \beta_1 X_i + \beta_2 Z_i + e_i^G$
- Negative Binomial PDF with a new k dispersion parameter is now:
 - $\text{Prob}(y_i = y | \beta_0, \beta_1, \beta_2) = \frac{\Gamma(y + \frac{1}{k})}{\Gamma(y+1) * \Gamma(\frac{1}{k})} * \frac{(k\mu_i)^y}{(1+k\mu_i)^{y+\frac{1}{k}}}$

DIST = NEGBIN in SAS

So \approx Poisson if $k = 0$

 - k is dispersion, such that $\text{Var}(y_i) = \mu_i + k\mu_i^2$
 - Can test whether $k > 0$ via $-2LL$ test, although LL for $k = 0$ is undefined
- An alternative model with the same idea is the **generalized Poisson**:
 - Mean: $\frac{\lambda}{1-k}$, Variance: $\frac{\mu}{(1-k)^2}$, that way LL is defined for $k = 0$
 - Is in SAS FMM (and in GLIMMIX through user-defined functions)

Negative Binomial (NB) = “Stretchy” Poisson...



Mean = λ
Dispersion = k

$$\text{Var}(y_i) = \lambda + k\lambda^2$$

A Negative Binomial model can be useful for count residuals with extra skewness, but otherwise follow a Poisson distribution.

- Because its k dispersion parameter is fixed to 0, the Poisson model is nested within the Negative Binomial model—to test improvement in fit:
- Is $-2(LL_{Poisson} - LL_{NegBin}) > 3.84$ for $df = 1$? Then $p < .05$, keep NB

Problem #2: There are no 0 values

- **“Zero-Altered”** or **“Zero-Truncated”** Poisson or Negative Binomial: ZAP/ZANB or ZTP/ZTNB (used in hurdle models)
 - Is usual count distribution, just not allowing any 0 values
 - Poisson version is readily available within SAS PROC FMM using DIST=TRUNCPOISSON (next version should have TRUNCNEGBIN, too)
 - Could be fitted in SAS NL MIXED or Mplus, too
- Poisson PDF was: $\text{Prob}(y_i = y | \mu_i) = \frac{\mu_i^y * \exp(-\mu_i)}{y!}$
- Zero-Truncated Poisson PDF is:
 - $\text{Prob}(y_i = y | \mu_i, y_i > 0) = \frac{\mu_i^y * \exp(-\mu_i)}{y! [1 - \exp(-\mu_i)]}$
 - $\text{Prob}(y_i = 0) = \exp(-\mu_i)$, so $\text{Prob}(y_i > 0) = 1 - \exp(-\mu_i)$
 - Divides by probability of non-0 outcomes so probability still sums to 1



Software for Discrete Outcomes

- There are many choices for modeling not-normal **discrete** outcomes (that include integer values only); most use either an identity or log link
- Single-level, univariate generalized models in SAS:
 - GENMOD: DIST= (and default link): Binomial (Logit), Poisson (Log), Zero-Inflated Poisson (Log), Negative Binomial (Log), Zero-Inflated Negative Binomial (Log)
 - FMM: DIST= (and default link): Binomial (Logit), Poisson (Log), Generalized Poisson (Log), Truncated Poisson (Log), Negative Binomial (Log), Uniform
- Multilevel, multivariate generalized models in SAS through GLIMMIX:
 - Binomial (Logit), Poisson (Log), Negative Binomial (Log)
 - BYOBS, which allows multivariate models by which you specify DV-specific link functions and distributions estimated simultaneously
 - User-defined variance functions for special cases (e.g., generalized Poisson)
- NLMIXED can also be used to fit any user-defined model

Generalized Models: Part 2



- Topics:
 - Models for skewed discrete outcomes
 - **Models for skewed continuous outcomes**
 - Models for zero-inflated discrete or continuous outcomes
 - Complications for generalized multilevel models
 - Multivariate or multilevel data

Log-Normal Distribution (Link=Identity)






- $e_i \sim \text{LogNormal}(0, \sigma_e^2) \rightarrow \mathbf{\log}$ of residuals is normal
 - Is same as log-transforming your outcome in this case...
 - The log link keeps the predicted values positive, but slopes then have an exponential (not linear) relation with original outcome

Log-Normal Distribution (Link=Identity)



- GLIMMIX parameterization gives μ (= intercept) and $scale$ = (variance) to convert back into original data as follows:
 - $Mean(Y) = \exp(\mu) * \sqrt{\exp(scale)}$
 - $Variance(Y) = \exp(2\mu) * \exp(scale) * [\exp(scale) - 1]$

Gamma Response Distribution



- GLIMMIX parameterization with LINK=LOG gives μ (= intercept) and $scale$ = (dispersion) to convert into original data as follows:
 - Mean(Y) = $\exp(\mu) \approx (shape*scale)$
 - Variance(Y) = $\exp(\mu)^2 * dispersion \approx (shape * scale^2)$

Software for Continuous Outcomes

- There are many choices for modeling not-normal **continuous** outcomes (that can include non-integer values); most use either an identity or log link
- Single-level, univariate generalized models in SAS:
 - GENMOD: DIST= (and default link): Gamma (Inverse), Geometric (Log), Inverse Gaussian (Inverse²), Normal (Identity)
 - FMM: DIST= (and default link): Beta (Logit), Betabinomial (Logit), Exponential (Log), Gamma (Log), Normal (Identity), Geometric (Log), Inverse Gaussian (Inverse²), LogNormal (Identity), TCentral (Identity), Weibull (Log)
- Multilevel or multivariate generalized models in SAS through GLIMMIX:
 - Beta (Logit), Exponential (Log), Gamma (Log), Geometric (Log), Inverse Gaussian (Inverse²), Normal (Identity), LogNormal (Identity), TCentral (Identity)
 - BYOBS, which allows multivariate models by which you specify DV-specific link functions and distributions estimated simultaneously (e.g., two-part)
- NLMIXED can also be used to fit any user-defined model

Generalized Models: Part 2

- Topics:
 - Models for skewed discrete outcomes
 - Models for skewed continuous outcomes
 - **Models for zero-inflated discrete or continuous outcomes**
 - Complications for generalized multilevel models
 - Multivariate or multilevel data


Modeling Not-Normal Outcomes

- Previously we examined models for skewed distributions
 - Discrete: Poisson, Generalized Poisson, Negative Binomial (NB)
 - Continuous: Log-Normal, Gamma (also Beta from Part 1)
- Now we will see additions to these models when the outcome also has a pile of 0's: Model becomes **If 0** and **How Much**
 - These models will differ in how they define the "If 0" part
 - Discrete: Zero-Inflated Poisson or NB, Hurdle Poisson or NB
 - Continuous: Two-Part (with normal or lognormal for how much)
 - Many of these can be estimated directly in Mplus or SAS GLIMMIX, but some will need to be programmed in SAS GLIMMIX or NL MIXED
 - More options for single-level data in SAS PROC FMM

Problem #3: Too many 0 values, Option #1

- **“Zero-Inflated”** Poisson (DIST=ZIP) or Negative Binomial (DIST=ZINB); available within SAS PROC GENMOD (and Mplus)
 - Distinguishes **two kinds of 0 values: expected** and **inflated** (“structural”) through a mixture of distributions (Bernoulli + Poisson/NB)
 - Creates two submodels to predict “if **extra 0**” and “if not, how much”?
 - Does not readily map onto most hypotheses (in my opinion)
 - But a ZIP example would look like this... (ZINB would add k dispersion, too)
- Submodel 1: $\text{Logit}(y_i = \text{extra } 0) = \beta_{01} + \beta_{11}X_i + \beta_{21}Z_i$
 - Predict being an extra 0 using Link = Logit, Distribution = Bernoulli
 - Don’t have to specify predictors for this part, can simply allow an intercept (but need ZEROMODEL option to include predictors in SAS GENMOD)
- Submodel 2: $\text{Log}(E(y_i)) = \beta_{02} + \beta_{12}X_i + \beta_{22}Z_i$
 - Predict rest of counts (including 0’s) using Link = Log, Distribution = Poisson

Example of Zero-Inflated Outcomes



Zero-inflated distributions have extra “structural zeros” not expected from Poisson or NB (“stretched Poisson”) distributions.

This can be tricky to estimate and interpret because the model distinguishes between *kinds of zeros* rather than zero or not...

Image borrowed from Atkins & Gallop, 2007

Figure 1. Histogram of Marital Status Inventory with predicted probabilities from regressions. NB = negative binomial; ZIP = zero-inflated Poisson; ZINB = zero-inflated negative binomial.

Problem #3: Too many 0 values, Option #1

- The Zero-Inflated models get put back together as follows:

- ω_i is the predicted probability of being an extra 0, from:

$$\omega_i = \frac{\exp[\text{Logit}(y_i = \text{extra } 0)]}{1 + \exp[\text{Logit}(y_i = \text{extra } 0)]}$$

- μ_i is the predicted count for the rest of the distribution, from:

$$\mu_i = \exp[\text{Log}(y_i)]$$

- ZIP: Mean (original y_i) = $(1 - \omega_i)\mu_i$

- ZIP: Variance(original y_i) = $\mu_i + \frac{\omega_i}{(1-\omega_i)} \mu_i^2$

- ZINB: Mean (original y_i) = $(1 - \omega_i)\mu_i$

- ZINB: Variance(original y_i) = $\mu_i + \left[\frac{\omega_i}{(1-\omega_i)} + \frac{k}{1-\omega_i} \right] \mu_i^2$

Problem #3: Too many 0 values, Option #2

- “**Hurdle**” models for Poisson or Negative Binomial
 - PH or NBH: Explicitly **separates 0 from non-0 values** through a mixture of distributions (Bernoulli + Zero-Altered Poisson/NB)
 - Creates two submodels to predict “if any 0” and “if not 0, how much”?
 - Easier to think about in terms of prediction (in my opinion)
- Submodel 1: $\text{Logit}(y_i = 0) = \beta_{01} + \beta_{11}X_i + \beta_{21}Z_i$
 - Predict being any 0 using Link = Logit, Distribution = Bernoulli
 - Don't have to specify predictors for this part, can simply allow it to exist
- Submodel 2: $\text{Log}(E(y_i) > 0) = \beta_{02} + \beta_{12}X_i + \beta_{22}Z_i$
 - Predict rest of positive counts using Link = Log, Distribution = ZAP or ZANB
- These models are not readily available in SAS, but NBH is in Mplus
 - Could be fit in SAS NLMIXED (as could ZIP/ZINB)

Two-Part Models for Continuous Outcomes

- A two-part model is an analog to hurdle models for zero-inflated count outcomes (and could be used with count outcomes, too)
 - Explicitly **separates 0 from non-0 values** through a mixture of distributions (Bernoulli + Normal or LogNormal)
 - Creates two submodels to predict “if not 0” and “if not 0, how much?”
 - Easier to think about in terms of prediction (in my opinion)
- Submodel 1: $\text{Logit}(y_i > 0) = \beta_{01} + \beta_{11}X_i + \beta_{21}Z_i$
 - Predict being any 0 using Link = Logit, Distribution = Bernoulli
 - Usually do specify predictors for this part
- Submodel 2: $(y_i | y_i > 0) = \beta_{02} + \beta_{11}X_i + \beta_{21}Z_i$
 - Predict rest of positive amount using Link = Identity, Distribution = Normal or Log-Normal (often rest of distribution is skewed, so log works better)
- Two-part is not readily available in SAS, but is in Mplus
 - Can be estimated as a multivariate model in SAS GLIMMIX or NL MIXED
 - Is related to “tobit” models for censored outcomes (for floor/ceiling effects)

Pile of 0's Taxonomy

- What kind of **amount** do you want to predict?
 - Discrete: Count → Poisson, Stretchy Count → Negative Binomial
 - Continuous: Normal, Log-Normal, Gamma
- What kind of **If 0** do you want to predict?
 - Discrete: Extra 0 beyond predicted by amount?
→ Zero-inflated Poisson or Zero-inflated Negative Binomial
 - Discrete: Any 0 at all?
→ Hurdle Poisson or Hurdle Negative Binomial
 - Continuous: Any 0 at all?
→ Two-Part with Continuous Amount (see above)
 - Note: Given the same amount distribution, these alternative ways of predicting 0 will result in the same empty model fit

Comparing Generalized Models

- Whether or not a dispersion parameter is needed (to distinguish Poisson and NB) can be answered via a likelihood ratio test
 - For the most fair comparison, keep the linear predictor model the same
- Whether or not a zero-inflation model is needed should, in theory, also be answerable via a likelihood ratio test...
 - But people disagree about this
 - Problem? Zero-inflation probability can't be negative, so is bounded at 0
 - Other tests have been proposed (e.g., Vuong test—see SAS macro online)
 - Can always check AIC and BIC (smaller is better)
- In general, models with the same distribution and different links can be compared via AIC and BIC, but one cannot use AIC and BIC to compare across alternative distributions (e.g., normal or not?)
 - Log-Likelihoods are not on the same scale due to using different PDFs
 - You can compute predicted values under different models to see how reasonably they approximate the data for some unofficial guidance

Generalized Models: Part 2

- Topics:
 - Models for skewed discrete outcomes
 - Models for skewed continuous outcomes
 - Models for zero-inflated discrete or continuous outcomes
 - **Complications for generalized multilevel models**
 - **Multivariate or multilevel data**

Multivariate Data in PROC GLIMMIX

- Multivariate models can be fitted in PROC GLIMMIX using stacked data, same as in MIXED... first, the bad news:
 - There is no **R** matrix in true ML, only **G**, and **V** can't be printed, either, which sometimes makes it hard to tell what structure is being predicted
 - There is no easy way to allow different scale factors given the same link and distribution across multivariate outcomes (as far as I know)
 - This means that a random intercept can be included to create constant covariance across outcomes, but that any differential variance (scale) or covariance must be included via RANDOM statement as well (to go in **G**)
- Now, the good news:
 - It allows different links and distributions across outcomes using LINK=BYOBS and DIST=BYOBS (Save new variables called "link" and "dist" to your data to tell GLIMMIX what to use per outcome)
 - It will do $-2\Delta LL$ tests for you using the COVTEST option! (not in MIXED)

From Single-Level to Multilevel...

- Multilevel generalized models have the same 3 parts as single-level generalized models:
 - Link function transforms bounded DV into unbounded conditional mean
 - Linear model that directly predicts link-transformed conditional mean
 - Alternative distribution of level-1 residuals used (e.g., Bernoulli)
- But in adding random effects (i.e., additional piles of variance) to address dependency in longitudinal or multivariate data:
 - Piles of variance are ADDED TO, not EXTRACTED FROM, the original residual variance pile when it is fixed to a known value (e.g., 3.29), which causes the model coefficients to change scale across models
 - ML estimation is way more difficult because normal random effects + not-normal residuals does not have a known distribution like MVN
 - No such thing as REML for generalized multilevel models

Empty Multilevel Model for Binary Outcomes

- **Level 1:** $\text{Logit}(y_{ti}) = \beta_{0i}$
- **Level 2:** $\beta_{0i} = \gamma_{00} + U_{0i}$
- **Composite:** $\text{Logit}(y_{ti}) = \gamma_{00} + U_{0i}$
- σ_e^2 residual variance is not estimated $\rightarrow \pi^2/3 = 3.29$
 - (Known) residual is in model for actual Y, not prob(Y) or logit(Y)
- Logistic ICC = $\frac{BP}{BP+WP} = \frac{\tau_{U_0}^2}{\tau_{U_0}^2 + \sigma_e^2} = \frac{\tau_{U_0}^2}{\tau_{U_0}^2 + 3.29}$
- Can do $-2\Delta LL$ test to see if $\tau_{U_0}^2 > 0$, although the ICC is somewhat problematic to interpret due to non-constant residual variance
 - Have not seen equivalent ICC formulas for other generalized models besides binary

Note what's NOT in level 1...

Random Linear Time Model for Binary Outcomes

- **Level 1:** $\text{Logit}(y_{ti}) = \beta_{0i} + \beta_{1i}(\text{time}_{ti})$
- **Level 2:** $\beta_{0i} = \gamma_{00} + U_{0i}$
 $\beta_{1i} = \gamma_{10} + U_{1i}$
- **Combined:** $\text{Logit}(y_{ti}) = (\gamma_{00} + U_{0i}) + (\gamma_{10} + U_{1i})(\text{time}_{ti})$
- σ_e^2 residual variance is still not estimated $\rightarrow \pi^2/3 = 3.29$
- Can test new fixed or random effects with $-2\Delta LL$ tests (or Wald test p -values for fixed effects as usual)

Random Linear Time Model for **Ordinal** Outcomes ($C = 3$)

- **L1:** $\text{Logit}(y_{ti1}) = \beta_{0i1} + \beta_{1i1}(\text{time}_{ti})$
 $\text{Logit}(y_{ti2}) = \beta_{0i2} + \beta_{1i2}(\text{time}_{ti})$
- **L2:** $\beta_{0i1} = \gamma_{001} + \mathbf{U}_{0i1}$ $\beta_{1i1} = \gamma_{101} + \mathbf{U}_{1i1}$
 $\beta_{0i2} = \gamma_{002} + \mathbf{U}_{0i2}$ $\beta_{1i2} = \gamma_{102} + \mathbf{U}_{1i2}$
- Assumes proportional odds \rightarrow
 $\gamma_{001} \neq \gamma_{002}$ and $\gamma_{101} = \gamma_{102}$ and $\mathbf{U}_{0i1} = \mathbf{U}_{0i2}$ and $\mathbf{U}_{1i1} = \mathbf{U}_{1i2}$
 - Testable via nominal model (all unequal) or using NLMIXED to write a custom model in which some can be constrained
 - σ_e^2 residual variance is still not estimated $\rightarrow \pi^2/3 = 3.29$

New Interpretation of Fixed Effects

- In general linear mixed models, the fixed effects are interpreted as the “average” effect for the sample
 - γ_{00} is “sample average” intercept
 - u_{0i} is “individual deviation from sample average”
- What “average” means in *generalized* linear mixed models is different, because the natural log is a nonlinear function:
 - So the mean of the logs \neq log of the means
 - Therefore, the fixed effects are not the “sample average” effect, they are the effect for ***specifically for $U_i = 0$***
 - Fixed effects are *conditional* on the random effects
 - This gets called a “unit-specific” or “subject-specific” model
 - This distinction does not exist for normally distributed outcomes

Comparing Results across Models

- NEW RULE: Coefficients cannot be compared across models, because they are not on the same scale! (see Bauer, 2009)
- e.g., if residual variance = 3.29 in binary models:
 - When adding a random intercept variance to an empty model, the **total variation in the outcome has increased** → the fixed effects will increase in size because they are *unstandardized* slopes

$$\gamma_{\text{mixed}} \approx \sqrt{\frac{\tau_{U_0}^2 + 3.29}{3.29}} (\beta_{\text{fixed}})$$

- **Level-1 predictors cannot decrease the residual variance** like usual, so all other models estimates have to go up to compensate
 - If X_{ti} is uncorrelated with other X 's and is a pure level-1 variable ($\text{ICC} \approx 0$), then fixed and $\text{SD}(U_{0i})$ will increase by same factor
- **Random effects variances can decrease**, though, so level-2 effects should be on the same scale across models if level-1 is the same

A Little Bit about Estimation

- Goal: End up with maximum likelihood estimates for all model parameters (because they are consistent, efficient)
 - When we have a \mathbf{V} matrix based on multivariate **normally** distributed \mathbf{e}_{ti} residuals at level-1 and multivariate normally distributed \mathbf{U}_i terms at level 2, ML is easy
 - When we have a \mathbf{V} matrix based on multivariate **Bernoulli** distributed \mathbf{e}_{ti} residuals at level-1 and multivariate normally distributed \mathbf{U}_i terms at level 2, ML is much harder
 - Same with any other kind model for “not normal” level 1 residual
 - **ML does not assume normality unless you fit a “normal” model!**
- 3 main families of estimation approaches:
 - Quasi-Likelihood methods (“marginal/penalized quasi ML”)
 - Numerical Integration (“adaptive Gaussian quadrature”)
 - Also Bayesian methods (MCMC, newly available in SAS or Mplus)

2 Main Types of Estimation


- **Quasi-Likelihood methods** → older methods
 - “Marginal QL” → approximation around fixed part of model
 - “Penalized QL” → approximation around fixed + random parts
 - These both underestimate variances (MQL more so than PQL)
 - 2nd-order PQL is supposed to be better than 1st-order MQL
 - QL methods DO NOT PERMIT MODEL $-2\Delta LL$ TESTS
 - HLM program adds Laplace approximation to QL, which then does permit $-2\Delta LL$ tests (also in SAS GLIMMIX and STATA xtmelogit)
- **ML via Numerical Integration** → gold standard
 - Much better estimates and $-2\Delta LL$ tests, but can take for-freaking-ever (can use PQL methods to get good start values)
 - Will blow up with many random effects (which make the model exponentially more complex, especially in these models)
 - Relies on assumptions of local independence, like usual → all level-1 dependency has been modeled; level-2 units are independent

ML via Numerical Integration

- **Step 1:** Select **starting values** for all fixed effects
- **Step 2:** Compute the **likelihood** of each observation given by the *current* parameter values using chosen distribution of residuals
 - Model gives link-predicted outcome given parameter estimates, but the U 's themselves are not parameters—their variance and covariance are instead
 - But so long as we can assume the U 's are MVN, we can still proceed
 - Computing the likelihood for each set of possible parameters requires *removing* the individual U values from the model equation—by **integrating** across possible U values for each Level-2 unit
 - Integration is accomplished by “Gaussian Quadrature” → summing up rectangles that approximate the integral (area under the curve) for each Level-2 unit
- **Step 3:** Decide if you have the right answers, which occurs when the log-likelihood changes very little across iterations (i.e., it converges)
- **Step 4:** If you aren't converged, choose new parameters values
 - Newton-Rhapon or Fisher Scoring (calculus), EM algorithm (U 's = missing data)

ML via Numerical Integration

- More on Step 2: Divide the U distribution into rectangles
 - Called “Gaussian Quadrature” (# rectangles = # “quadrature points”)
 - Start by dividing the whole distribution into rectangles, then repeat, taking the most likely section for each level-2 unit and retriangulating that
 - This is “adaptive quadrature” and is computationally more demanding, but gives more accurate results with fewer rectangles (SAS will pick how many)



The likelihood of each level-2 unit's outcomes at each **U** rectangle is then weighted by that rectangle's probability of being observed (from the multivariate normal distribution). The weighted likelihoods are then summed across all rectangles...
→ ta da! “**numerical integration**”

Example of Numeric Integration: Binary DV, Fixed Linear Time, Random Intercept Model

1. Start with values for fixed effects: intercept: $\gamma_{00} = 0.5$, time: $\gamma_{10} = 2.0$,
2. Compute likelihood for real data based on fixed effects and plausible U_{0i} (-2,0,2) using model: $\text{Logit}(y_{ti}=1) = \gamma_{00} + \gamma_{10}(\text{time}_{ti}) + U_{0i}$
 - Here for one person at two occasions with $y_{ti}=1$ at both occasions

	$U_{0i} = -2$	$\text{Logit}(y_{ti})$	IF $y_{ti}=1$ Prob	IF $y_{ti}=0$ 1-Prob	Likelihood if both $y=1$	Theta prob	Theta width	Product per Theta
Time 0	$0.5 + 1.5(0) - 2$	-1.5	0.18	0.82	0.091213	0.05	2	0.00912
Time 1	$0.5 + 1.5(1) - 2$	0.0	0.50	0.50				
	$U_{0i} = 0$	$\text{Logit}(y_{ti})$	Prob	1-Prob				
Time 0	$0.5 + 1.5(0) + 0$	0.5	0.62	0.38	0.54826	0.40	2	0.43861
Time 1	$0.5 + 1.5(1) + 0$	2.0	0.88	0.12				
	$U_{0i} = 2$	$\text{Logit}(y_{ti})$	Prob	1-Prob				
Time 0	$0.5 + 1.5(0) + 2$	2.5	0.92	0.08	0.90752	0.05	2	0.09075
Time 1	$0.5 + 1.5(1) + 2$	4.0	0.98	0.02				
Overall Likelihood (Sum of Products over All Thetas):								0.53848

(do this for each occasion, then multiply this whole thing over all people)

(repeat with new values of fixed effects until find highest overall likelihood)

Summary: Generalized Multilevel Models

- Analyze link-transformed conditional mean (e.g., via logit, log, log-log...)
 - **Linear** relationship between X's and **transformed** conditional mean of Y
 - **Nonlinear** relationship between X's and **original** Y
 - Original e_{ij} residuals are assumed to follow some non-normal distribution
- In models for binary or categorical data, level-1 residual variance is set
 - So it can't go down after adding level-1 predictors, which means that the scale of everything else has to go UP to compensate
 - Scale of model will also be different after adding random effects for the same reason—the total variation in the model is now bigger
 - Fixed effects may not be comparable across models as a result
- Estimation is trickier and takes longer
 - Numerical integration is best but may blow up in complex models
 - Start values are often essential (can get those with pseudo-likelihood estimators)

Generalized Models Part 2: Summary

- There are many options for “amount” variables whose residuals may not be normally distributed
 - Discrete: Poisson, Negative Binomial
 - Continuous: Lognormal, Gamma, Beta
 - Too many 0’s: Zero-inflated or hurdle for discrete; two-part
- Multivariate and multilevel versions of all the generalized models we covered *can* be estimated...
 - But it’s harder to do and takes longer due to numeric integration (trying on random effects at each iteration)
 - But there are fewer ready-made options for modeling differential variance/covariance across DVs (no easy R matrix structures in true ML)
- Program documentation will always be your friend to determine exactly what a given model is doing!