

BIOS 7345: Advanced Regression for Independent Data

Andrew J. Spieker, Ph.D.

Associate Professor of Biostatistics
Vanderbilt University

Set 16: Further considerations for binary outcomes

Version: 06/11/2025

TABLE OF CONTENTS

- 1 Non-collapsibility of the logit link
- 2 Latent variable formulation logistic regression
- 3 The probit link function for binary outcomes
- 4 The clog-log link function for binary outcomes
- 5 Conditional logistic regression

NON-COLLAPSIBILITY OF THE LOGIT LINK

Logistic regression: Non-collapsibility

- Take, for example, a study to determine the association between cardiovascular disease (CVD) and kidney stone history (yes/no).
- Tables stratified by age group (≤ 50 years and > 50 years):

≤ 50 years	CVD	No CVD
Kidney stones	40	60
No kidney stones	10	90

> 50 years	CVD	No CVD
Kidney stones	90	10
No kidney stones	60	40

- Is age group a confounder?
- What is the within-group (adjusted) odds ratio?

Logistic regression: Non-collapsibility

- Now, “collapse” the tables over age group:

	All	CVD	No CVD
Kidney stones	130	70	
No kidney stones	70	130	

- What is the pooled/unadjusted/crude odds ratio?

NON-COLLAPSIBILITY OF THE LOGIT LINK


Illustration:

```
1 expit <- function(x)
2 {
3   expitx <- exp(x)/(1 + exp(x))
4   return(expitx)
5 }
6
7 ## Very large sample size (answer close to truth)
8 n <- 100000
9
10 ## Set seed
11 set.seed(7345)
12
13 ## Generate data
14 X <- rnorm(n, 0, 1)
15 Z <- rnorm(n, 0, 3)
16 p <- expit(0.4*X + 0.8*Z)
17 Y <- rbinom(n, 1, p)
```

NON-COLLAPSIBILITY OF THE LOGIT LINK

Logistic regression: Heterogeneity of $P(Y = 1|X = x, Z = z)$

Histogram



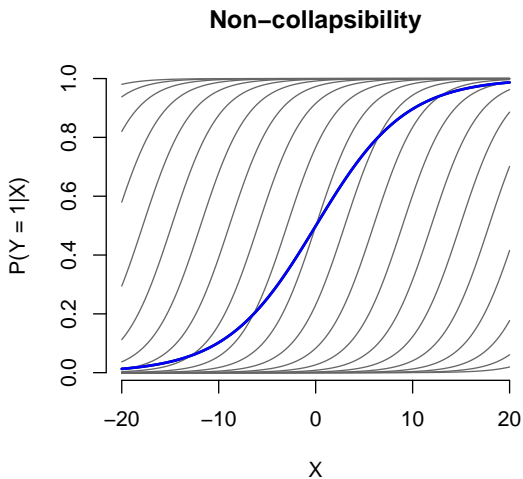
NON-COLLAPSIBILITY OF THE LOGIT LINK

Illustration:

```
1 zz.1 <- glm(Y ~ X, family = binomial(link = "logit"))
2 zz.2 <- glm(Y ~ X + Z, family = binomial(link = "logit"))
3
4 b.1 <- as.numeric(coef(zz.1))
5 b.2 <- as.numeric(coef(zz.2))
6
7 > b.1
8 [1] -0.004947256  0.216420075
9 > b.2
10 [1] -0.009635699  0.401667754  0.796640094
```

NON-COLLAPSIBILITY OF THE LOGIT LINK

Logistic regression: Illustration of non-collapsibility



Logistic regression: Non-collapsibility

- This illustration was a little “unfair” because only the model conditional on Z was actually correctly specified.
- The “fair” approach would be to consider a likelihood based on X that specifically marginalizes over the values of Z (beyond the scope of this course).
- Nevertheless, the estimate of a marginalized likelihood would be consistent for the conditional association (given Z).

TABLE OF CONTENTS

- 1 Non-collapsibility of the logit link
- 2 Latent variable formulation logistic regression
- 3 The probit link function for binary outcomes
- 4 The clog-log link function for binary outcomes
- 5 Conditional logistic regression

Logistic regression: Our typical formulations

- $Y \sim \text{Bernoulli}(p = \text{expit}(\mathbf{x}^T \boldsymbol{\beta}))$.
- $\text{logit}(P(Y = 1|\mathbf{x})) = \mathbf{x}^T \boldsymbol{\beta}$.

Logistic regression: Latent variable formulation

- Imagine that for each Bernoulli trial, there is an underlying but latent continuous variable Y^* distributed as follows:

$$Y^* = \mathbf{x}^T \boldsymbol{\beta} + \epsilon,$$

where $\epsilon \sim \text{Logistic}(0, 1)$, with density:

$$f(\epsilon) = \frac{\exp(-\epsilon)}{(1 + \exp(-\epsilon))^2} = \frac{1}{4} \text{sech}^2(\epsilon/2).$$

- $Y = 1(Y^* > 0)$ marks the observed dichotomous outcome.
- Note: $E[\epsilon] = 0$ and $\text{Var}[\epsilon] = \pi^2/3$.

Logistic regression: Latent variable formulation

- The logistic family is a location-scale family, although there is no reason to allow ϵ to follow a more general form depending upon location and scale parameters:
 - ▶ Location is counteracted by the centering of the intercept, β_0 .
 - ▶ Scale is counteracted by the magnitude of each component of β .
- Latent variable formulation makes it easier to extend to more complicated models with multiple correlated outcomes as well as to compare logistic regression to the closely related probit model.

TABLE OF CONTENTS

- 1 Non-collapsibility of the logit link
- 2 Latent variable formulation logistic regression
- 3 The probit link function for binary outcomes**
- 4 The clog-log link function for binary outcomes
- 5 Conditional logistic regression

Example 16.1: Dose until response


- Consider a study in which you are seeking to determine the relationship between dose and probability of response.
- Imagine each study subject has an underlying dose until response:

$$Y^* \sim \mathcal{N}(\mu, \sigma^2).$$

THE PROBIT LINK FUNCTION FOR BINARY OUTCOMES

Example 16.1: Dose until response

$$\mu=50, \sigma=20$$



Example 16.1: Dose until response

- Each study unit receives one dose, $D = d$, and either responds ($Y = 1$) or not ($Y = 0$).
- Under the normality assumption, it follows that

$$\begin{aligned}P(Y = 1|D = d) &= P(Y^* < D|D = d) \\ &= \Phi\left(\frac{d - \mu}{\sigma}\right) \\ &= \Phi\left(-\frac{\mu}{\sigma} + \frac{1}{\sigma}d\right) \\ &= \Phi(\beta_0 + \beta_1 d).\end{aligned}$$

- By this model, μ and σ are not estimated directly, though you can back-solve for them in this very simple case.
 - ▶ Groups differing by one unit in their received dose differ in their dose-to-response by β_1 standard deviations.

THE PROBIT LINK FUNCTION FOR BINARY OUTCOMES

Example 16.1: Dose until response

```
1 ## Sample size
2 n <- 50
3
4 ## Generate binary outcome accordingly
5 Y <- rbinom(n, 1, pnorm(D, mean = 50, sd = 25))
6
7 ## Fit probit GLM
8 zz <- glm(Y ~ D, family = binomial(link = "probit"))
9
10 ## Salient part of output
11 > summary(zz)
12
13 Coefficients:
14             Estimate Std. Error z value Pr(>|z|)
15 (Intercept) -1.96737    0.63920  -3.078 0.002085 **
16 D           0.04113    0.01099   3.741 0.000183 ***
17
18 ## The true values are  $b_0 = -50/25 = -2$  and  $b_1 = 1/25 = 0.04$ .
19 ## We estimate  $\sigma = 1/0.04133 = 24.2$ .
20 ## We estimate  $\mu = -(-1.96737)/0.04113 = 47.8$ .
```

Latent variable formulation:

- In another latent-variable formulation, we use the linear predictor to form the latent continuous variable, and *then* dichotomize (more common):

$$Y^* = \mathbf{x}^T \boldsymbol{\beta} + \epsilon,$$

where $\epsilon \sim \mathcal{N}(0, \sigma^2)$, and $Y = 1(Y^* > 0)$.

- In this case, $\boldsymbol{\beta}$ is only identifiable *only* up to a scale factor of σ (which is itself not identifiable). For this reason, we say that $\boldsymbol{\beta}$ is *weakly* identifiable.
 - ▶ By convention, we typically set $\sigma = 1$ to estimate $\boldsymbol{\beta}$.

THE PROBIT LINK FUNCTION FOR BINARY OUTCOMES

Example 16.2: Weak identifiability

```
1 ## Sample size
2 n <- 10000
3
4 ## Generate predictor
5 X <- runif(n, 0, 1)
6
7 ## Generate error
8 eps <- rnorm(n, 0, 4)
9
10 ## Set coefficient values
11 b0 <- -2
12 b1 <- 3
13
14 ## Generate latent continuous variable and observed outcome
15 Ystar <- b0 + b1 * X + eps
16 Y <- as.numeric(Ystar > 0)
17
18 ## Fit GLM
19 zz <- glm(Y ~ X, family = binomial(link = "probit"))
20
21 > summary(zz)
22
23 Coefficients:
24             Estimate Std. Error z value Pr(>|z|)
25 (Intercept) -0.51269    0.02546  -20.14  <2e-16 ***
26 X             0.75609    0.04413   17.13  <2e-16 ***
27
28 ## The true values are b0 = -2 and b1 = 3.
29 ## We estimate b0/sigma = -2/4 = -0.5 and b1/sigma = 3/4 = 0.75
```

THE PROBIT LINK FUNCTION FOR BINARY OUTCOMES

Comparison: logit vs. probit




TABLE OF CONTENTS

- 1 Non-collapsibility of the logit link
- 2 Latent variable formulation logistic regression
- 3 The probit link function for binary outcomes
- 4 The clog-log link function for binary outcomes**
- 5 Conditional logistic regression

Another link function:

- Suppose we're interested in drawing inferences about the rates at which a binary outcome occurs.
- Suppose that Y^* denotes the time-to-event during some fixed period of follow-up, τ , that an event occurred (if it did):

$$\lambda(y^*) = \lim_{\delta \rightarrow 0} \frac{P(y^* \leq Y^* < y^* + \delta | Y^* \geq y^*)}{\delta}.$$

- If we assume $\lambda(y^* | X = x) = \lambda_0(y^*) \exp(\beta_1 x)$, then $\exp(\beta_1)$ marks the hazard ratio that compares the instantaneous hazard rate of the event between strata differing in X by one unit.

Deriving the clog-log link:

- Let $Y = 1(Y^* < \tau)$ denote the observed indicator of having observed an event during the fixed period of follow-up. In general, we have:

$$P(Y = 1|X = x) = P(Y^* < \tau|X = x) = 1 - \exp(-\Lambda_x(\tau)).$$

- If the proportional hazards assumption holds, then for any τ ,

$$\begin{aligned}\Lambda_x(\tau) &= \exp(\beta_1 x)\Lambda_0(\tau) \\ \Rightarrow \exp(-\Lambda_x(\tau)) &= \exp(-\exp(\beta_1 x)\Lambda_0(\tau)) = [\exp(-\Lambda_0(\tau))]^{\exp(\beta_1 x)} \\ \Rightarrow 1 - P(Y = 1|X = x) &= (1 - P(Y = 1|X = 0))^{\exp(\beta_1 x)}.\end{aligned}$$

- Because $E[Y|X = x] = P(Y = 1|X = x)$, this is ultimately to say that the proportional hazards assumption lends itself to the link function $g(\mu) = \log(-\log(1 - \mu))$.

The clog-log link: $g(\mu) = \log(-\log(1 - \mu))$


- Bernoulli GLM with the clog-log link:

$$\begin{aligned} \log(-\log(1 - P(Y = 1|X = x))) &= \beta_0 + \beta_1 x \\ \Rightarrow 1 - P(Y = 1|X = x) &= \exp(-\exp(\beta_0 + \beta_1 x)) \\ &= \exp(-\exp(\beta_0))^{\exp(\beta_1 x)} \\ &= (1 - P(Y = 1|X = 0))^{\exp(\beta_1 x)}. \end{aligned}$$

- Note: $P(Y = 1|X = 0) = P(Y^* < \tau|X = 0) = 1 - \exp(-\exp(\beta_0))$.
 - ▶ Note: This CDF is related to a standardized Gumbel distribution (latent variable formulation).
- Note: $\exp(\beta_1)$ is the hazard ratio.
- To what extent can the proportional hazards assumption be assessed in this context?

Example 16.3: Weibull time-to-event with common shape

- Let $X \sim \text{Bernoulli}(p = 0.5)$ and $Y^* \sim \text{Weibull}(k = 2, \lambda = 2 + X)$.
- Consider the cutoff of $\tau = 3$.
 - ▶ $P(Y = 1|X = 0) \approx 0.895$ (`pweibull(3, shape = 2, scale = 2)`).
 - ▶ You can verify that this is consistent with the proportional hazards assumption, with a hazard ratio given by $(2/3)^2 = 4/9 \approx 0.444$, which—by proportional hazards—should be true in general (and not only for this value of τ).

Example 16.3: Weibull time-to-event with common shape

Example 16.3: Weibull time-to-event with common shape

```
1 ## Set seed for reproducibility
2 set.seed(7345)
3
4 ## Set sample size
5 n <- 2000
6
7 ## Generate X
8 x <- rbinom(n, 1, 0.5)
9
10 ## Generate underlying latent continuous variable
11 ystar <- rweibull(n, shape = 2, scale = 2 + x)
12
13 ## Generate dichotomous realization with tau = 3
14 y <- as.numeric(ystar < 3)
15
16 ## Although you could expect this to work for other choices of tau.
```

Example 16.3: Weibull time-to-event with common shape

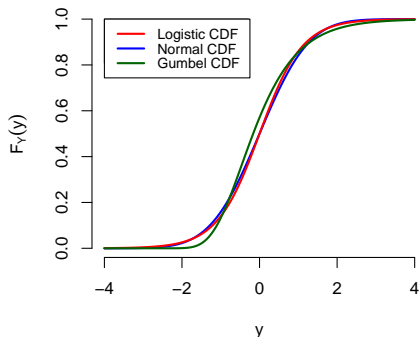
```
1 ## Fit clog-log GLM
2 zz <- glm(y ~ x, family = binomial(link = "cloglog"))
3
4 ## Should be about 0.444 (hazard ratio)
5 > as.numeric(exp(coef(zz)[2]))
6 [1] 0.456471
7
8 ## Should be about 0.895 (baseline prevalence)
9 > as.numeric(1 - exp(-exp(coef(zz)[1])))
10 [1] 0.8849647
```

Link functions: Decisions, decisions

- In this set of notes, we have characterized the logit, probit, and clog-log links for binary outcomes in terms of underlying/latent continuous variables.
- Generally, the choice of a link function for binary outcomes is more heavily driven by the interpretation you seek to get out of the model (and less by trying to fit the data perfectly).

THE CLOG-LOG LINK FUNCTION FOR BINARY OUTCOMES

Comparison: logit vs. probit vs. clog-log



Practical considerations:

- Generally, the choice of a link function for binary outcomes is more heavily driven by the interpretation you seek to get out of the model (and less by trying to fit the data perfectly).

TABLE OF CONTENTS

- 1 Non-collapsibility of the logit link
- 2 Latent variable formulation logistic regression
- 3 The probit link function for binary outcomes
- 4 The clog-log link function for binary outcomes
- 5 Conditional logistic regression

Example 16.4: 1:1 matched case control data (binary exposure)

- Suppose each case ($Y = 1$) is matched to a control ($Y = 0$).
- Let $m = 1, \dots, M$ mark the case-control stratum and let $j = 0, 1$ index the individual in the stratum; consider the model:

$$\text{logit}(P(Y_{jm} = 1 | X_{jm} = x_{jm})) = \alpha_m + x_{jm}\beta$$

- Sample size: $N = 2M$.
- Number of model parameters: $K = M + 1$.
- Maximum likelihood estimate of β is biased (can show $\hat{\beta} \xrightarrow{P} 2\beta$, so that $\widehat{\text{OR}} \xrightarrow{P} \text{OR}^2$).

Example 16.4: Simulation

```
1 ## Set seed for reproducibility
2 set.seed(7345)
3
4 ## Important function
5 expit <- function(x) {exp(x)/(1 + exp(x))}
6
7 ## Set number of simulations
8 nsim <- 1000
9
10 ## Number of strata
11 M <- 500
12
13 ## Set value of parameter of interest
14 beta <- 1
15
16 ## Create place to store results
17 res <- matrix(0, nrow = nsim, ncol = 1)
```

Example 16.4: Simulation

```
1 for (j in 1:nsim) {
2   # Case/control outcome
3   y <- rep(c(0,1), M)
4
5   ## Matched pair number
6   set <- rep(c(1:M), each = 2)
7
8   # Pair-specific parameters
9   alpha <- rep(rnorm(M), each = 2)
10
11  ## Binary exposure
12  x <- rbinom(2*M, 1, expit(alpha + beta*y))
13
14  ## Fit logistic regression model and extract results
15  zz <- glm(y ~ factor(set) + x, family = binomial(link = "logit"))
16  res[j,1] <- as.numeric(coef(zz)[M + 1])
17 }
18
19 ## Result
20 > mean(res)
21 [1] 2.006405
```

Brief background: Conditional likelihood

- Suppose you are interested in estimating β under the likelihood $f(\mathbf{y}; \alpha, \beta)$, where α marks nuisance parameters.
- It is often possible to factor the density in the following useful way:

$$f(\mathbf{w}, \mathbf{v}; \alpha, \beta) = f(\mathbf{v}; \alpha, \beta) f(\mathbf{w}|\mathbf{v}; \beta)$$

where \mathbf{V} and \mathbf{W} denote sufficient statistics for α and β , respectively.

- The density in blue is the conditional likelihood for β under this factorization (**which is not always possible**). When it is possible to factor in this way, the conditional likelihood (which depends upon β but not α) can be used to estimate β .

Brief background: Conditional likelihood in exponential families

- Suppose the density follows the following exponential form:

$$f(\mathbf{y}; \boldsymbol{\alpha}, \boldsymbol{\beta}) = h(\mathbf{y}) \exp \left(\sum_i \beta_i \mathbf{w}_i + \sum_j \alpha_j \mathbf{v}_j - a(\boldsymbol{\alpha}, \boldsymbol{\beta}) \right).$$

- Then, it can be shown that the conditional distribution of \mathbf{w} (given \mathbf{v}) is of the following form:

$$f(\mathbf{w}|\mathbf{v}; \boldsymbol{\beta}) = q(\mathbf{w}, \mathbf{v}) \exp \left(\sum_i \beta_i \mathbf{w}_i - c(\mathbf{v}, \boldsymbol{\beta}) \right).$$

- Note in particular that the conditional likelihood depends only upon $\boldsymbol{\beta}$ and it can be used for inference.

Setup: Matched case control data

- Let $i = 1, \dots, N$ index study units.
- Let $I_{mi} = 1(\text{set}_i = m)$ indicate membership to the m^{th} matched set (of which there are M).
- Let $\mathbf{z}_i = (I_{1i}, \dots, I_{Mi}, \mathbf{x}_i)$.
- Let $\boldsymbol{\theta} = (\boldsymbol{\alpha}, \boldsymbol{\beta}) = (\alpha_1, \dots, \alpha_M, \boldsymbol{\beta})$.
- Then, the likelihood is given by

$$\begin{aligned} \mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\beta}; \mathbf{Z}, \mathbf{y}) &= \prod_{i=1}^N \left[\frac{\exp(\mathbf{z}_i^\top \boldsymbol{\theta})}{1 + \exp(\mathbf{z}_i^\top \boldsymbol{\theta})} \right]^{y_i} \left[1 - \frac{\exp(\mathbf{z}_i^\top \boldsymbol{\theta})}{1 + \exp(\mathbf{z}_i^\top \boldsymbol{\theta})} \right]^{1-y_i} \\ &= \frac{\exp\left(\sum_{i=1}^N \mathbf{z}_i^\top \boldsymbol{\theta} y_i\right)}{\prod_{i=1}^N (1 + \exp(\mathbf{z}_i^\top \boldsymbol{\theta}))}. \end{aligned}$$

Conditional likelihood: Matched case control data

- Suppose there are M matched sets, each with one case matched to K_m controls. Without loss of generality, let the first observation in each matched set (of size $K_m + 1$) be the case, so that $Y_{1,m} = 1$ and $Y_{2,m} = \dots = Y_{K_m+1,m} = 0$.
- It can be shown (details omitted) that the conditional likelihood takes the following form:

$$\mathcal{L}_C(\boldsymbol{\beta}; \mathbf{X}) = \prod_{m=1}^M \frac{\exp(\mathbf{x}_{1m}^\top \boldsymbol{\beta})}{\sum_{k=1}^{K_m+1} \exp(\mathbf{x}_{km}^\top \boldsymbol{\beta})}.$$

Intuition:

- Consider the first stratum of a matched case-control study; say it has one case and two controls. Further, to simplify things, say there is a single dichotomous exposure variable. Label the subjects as $i = 1, 2, 3$.
- Let

$$X_{Di} = \begin{cases} 1 & \text{if subject } i \text{ is case} \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad X_{Ei} = \begin{cases} 1 & \text{if subject } i \text{ is exposed} \\ 0 & \text{otherwise} \end{cases}$$

- If the data consisted of only this stratum, we would maximize the likelihood corresponding to $\text{logit}(p) = \alpha + \beta x_E$ (α is the stratum-specific intercept).

Intuition:

- Make-believe data:

i	x_{Di}	x_{Ei}
1	1	1
2	0	1
3	0	0

- The unconditional likelihood is of the form:

$$P(X_{D1} = 1, X_{D2} = 0, X_{D3} = 0 | X_{E1} = 1, X_{E2} = 1, X_{E3} = 0; \alpha, \beta),$$

and we can find the values of α and β that maximize its value.

- Note that the likelihood depends upon the underlying disease risk (case probability) in this stratum (α) and the association between the exposure and disease incidence (β).

Intuition:

- The unconditional likelihood can be realized as:

$$P(\text{Exactly one is diseased and first is diseased} | X_{E1} = 1, X_{E2} = 1, X_{E3} = 0; \alpha, \beta),$$

- The “exactly one subject is diseased” component of the statement is the stratum-specific risk (probability of being a case).
- The “first subject is diseased” component of the statement is the association between the exposure and risk of disease.

Intuition:

- The conditional likelihood can be realized as:

$$\begin{aligned} P(X_{D1} = 1, X_{D2} = 0, X_{D3} = 0 | X_{E1} = 1, X_{E2} = 1, X_{E3} = 0, X_{D1} + X_{D2} + X_{D3} = 1; \beta) \\ = P(\text{First is diseased} | X_{E1} = 1, X_{E2} = 1, X_{E3} = 0, \text{Exactly one is diseased}; \beta) \end{aligned}$$

- The stratum-specific risk (probability of being a case) is conditioned upon in the conditional likelihood; the association between the risk and exposure does not depend upon the underlying disease risk.
- Since this likelihood does not depend upon α , it does not allow us to estimate α .

Intuitive exploration: Key ideas

- Conditional logistic regression provides estimates of the coefficient of any variable that varies within at least one stratum.
 - ▶ For example, the exposure.
 - ▶ For example, confounders that vary within one more strata.
 - ▶ Interactions of the above two with “stratum determinants” (i.e., matching variables).
- Conditional logistic regression cannot provide estimates for the coefficient of any variable that does not vary within at least one stratum (discordant pairs).
 - ▶ For example, a stratum determinant.
 - ▶ For example, interactions between stratum determinants.

Intuitive exploration: Key ideas

- Model: $\text{logit}(P(Y_{jm} = 1 | X_{jm} = x_{jm})) = \alpha_m + \mathbf{x}_{jm}^T \beta$.
- Even though they don't appear in the computer output of a conditional logistic regression command, the stratum indicators are still "in the model."
- We impose no constraints on how the underlying probabilities vary from stratum to stratum when we perform conditional logistic regression, *including* their interactions with the exposures.
- Therefore, including interactions between \mathbf{x} and stratum determinants does not violate our usual rules about how main effects need to be included in models that contain their interactions.

Example 16.5: Simulation

```
1 ## Important library
2 library(survival)
3
4 ## Important function
5 expit <- function(x) {exp(x)/(1 + exp(x))}
6
7 ## Set seed for reproducibility
8 set.seed(7345)
9
10 ## Set number of simulations
11 nsim <- 1000
12
13 ## Number of strata
14 M <- 500
15
16 ## Set value of parameter of interes
17 beta <- 1
18
19 ## Create place to store results
20 res <- matrix(0, nrow = nsim, ncol = 1)
```

Example 16.5: Simulation

```
1 for (j in 1:nsim) {
2   # Case/control outcome
3   y <- rep(c(0,1), M)
4
5   ## Matched pair number
6   set <- rep(c(1:M), each = 2)
7
8   # Pair-specific parameters
9   alpha <- rep(rnorm(M), each = 2)
10
11  ## Binary exposure
12  x <- rbinom(2*M, 1, expit(alpha + beta*y))
13
14  ## Fit logistic regression model and extract results
15  zz <- clogit(y ~ x + strata(set))
16  res[j,1] <- as.numeric(coef(zz))
17 }
18
19 ## Result (WHEW!!)
20 > mean(res)
21 [1] 1.003203
```

Further thoughts:

- Conditional logistic regression allows you to adjust for confounders even if you didn't match on them (or matched loosely).
- Conditional logistic regression is useful because it easily accounts for matched sets with differing numbers of cases and controls (under an unconditional likelihood, this would be more challenging).
- Conditional logistic regression is useful not *only* for matched data but for setting in which you have a lot of adjustment variables that are not of primary scientific interest.
- I expect you will discuss conditional and marginal likelihoods in much greater depth in BIOS 7346.

This unit:

- Non-collapsibility.
- Link functions for binary outcomes.
- Conditional logistic regression.

SUMMARY: SO FAR

- Random vectors and matrices; multivariate normal theory.
- Ordinary least squares.
- Hypothesis testing and ANOVA.
- Weighted least squares.
- Misspecification.
- Confidence regions and prediction.
- Diagnostics.
- Regularization.
- Bayesian regression.
- Exponential families.
- Generalized linear models.
- Sandwich and bootstrap.
- Quasi-likelihood.
- Hypothesis testing for GLMs.
- Diagnostics for GLMs.
- Further considerations for binary outcomes.

SUMMARY: COMING UP

- Nonlinear least squares.