# An Introduction to the R Statistical Computing Environment 

John Fox<br>McMaster University

ICPSR 2021

## Outline

(1) Linear Models in R
(2) Generalized Linear Models in R
(3) Mixed-Effects Models in R

44 Using the Tidyverse for Data Management
(5) R Programming

## Outline

(1) Linear Models in R

- Review of Dummy-Variable Regression
- Type-II Tests
- Arguments of the lm() Function
- Regression Diagnostics: Unusual Cases
- Regression Diagnostics: Added-Variable (AV) Plots
- Regression Diagnostics: Component-Plus-Residuals (C+R) Plots
- The Bulging Rule for Linearizing a Relationship
(2) Generalized Linear Models in R
(3) Mixed-Effects Models in R
(4) Using the Tidyverse for Data Management


## Linear Models in R

Review of Dummy-Variable Regression

- Defining a dummy-variable regressor for a dichotomous explanatory variable - e.g., gender in the regression of income $Y$ on gender and education $X$.


## Linear Models in R

Review of Dummy-Variable Regression

- Defining a dummy-variable regressor for a dichotomous explanatory variable - e.g., gender in the regression of income $Y$ on gender and education $X$.
- Let $D=0$ for women and $D=1$ for men.


## Linear Models in R

Review of Dummy-Variable Regression

- Defining a dummy-variable regressor for a dichotomous explanatory variable - e.g., gender in the regression of income $Y$ on gender and education $X$.
- Let $D=0$ for women and $D=1$ for men.
- Then the additive dummy-regression model is

$$
Y=\alpha+\beta X+\gamma D+\varepsilon
$$

## Linear Models in R

Review of Dummy-Variable Regression

- Defining a dummy-variable regressor for a dichotomous explanatory variable - e.g., gender in the regression of income $Y$ on gender and education $X$.
- Let $D=0$ for women and $D=1$ for men.
- Then the additive dummy-regression model is

$$
Y=\alpha+\beta X+\gamma D+\varepsilon
$$

- So, for women (treating $X$ as
conditionally fixed)

$$
\begin{aligned}
Y & =\alpha+\beta X+\gamma \times 0+\varepsilon \\
E(Y) & =\alpha+\beta X
\end{aligned}
$$

## Linear Models in R

Review of Dummy-Variable Regression

- Defining a dummy-variable regressor for a dichotomous explanatory variable - e.g., gender in the regression of income $Y$ on gender and education $X$.
- Let $D=0$ for women and $D=1$ for men.
- Then the additive dummy-regression model is

$$
Y=\alpha+\beta X+\gamma D+\varepsilon
$$

- So, for women (treating $X$ as conditionally fixed)

$$
\begin{aligned}
Y & =\alpha+\beta X+\gamma \times 0+\varepsilon \\
E(Y) & =\alpha+\beta X
\end{aligned}
$$

## Linear Models in R

Review of Dummy-Variable Regression

- Defining a dummy-variable regressor for a dichotomous explanatory variable - e.g., gender in the regression of income $Y$ on

$$
E(Y)=(\alpha+\gamma)+\beta X
$$ gender and education $X$.

- Let $D=0$ for women and $D=1$ for men.
- Then the additive dummy-regression model is

$$
Y=\alpha+\beta X+\gamma D+\varepsilon
$$

- So, for women (treating $X$ as
- And, for men

$$
Y=\alpha+\beta X+\gamma \times 1+\varepsilon
$$


conditionally fixed)

$$
\begin{aligned}
Y & =\alpha+\beta X+\gamma \times 0+\varepsilon \\
E(Y) & =\alpha+\beta X
\end{aligned}
$$

## Linear Models in R

- Defining a dummy-variable regressor for a dichotomous explanatory variable - e.g., gender in the regression of income $Y$ on gender and education $X$.
- Let $D=0$ for women and $D=1$ for men.
- Then the additive dummy-regression model is

$$
Y=\alpha+\beta X+\gamma D+\varepsilon
$$

- So, for women (treating $X$ as
conditionally fixed)

$$
\begin{aligned}
Y & =\alpha+\beta X+\gamma \times 0+\varepsilon \\
E(Y) & =\alpha+\beta X
\end{aligned}
$$

- And, for men

$$
Y=\alpha+\beta X+\gamma \times 1+\varepsilon
$$

$$
E(Y)=(\alpha+\gamma)+\beta X
$$



- In R notation with data in Data: model <- lm(income ~ education
+ gender, data=Data).


## Linear Models in R

Review of Dummy-Variable Regression

- Different slopes for women and men ("different slopes for different folks") can be modelled by introducing an interaction regressor, the product of $X$ and $D$, into the model:

$$
Y=\alpha+\beta X+\gamma D+\delta(X \times D)+\varepsilon
$$

## Linear Models in R

Review of Dummy-Variable Regression

- Different slopes for women and men ("different slopes for different folks") can be modelled by introducing an interaction regressor, the product of $X$ and $D$, into the model:

$$
Y=\alpha+\beta X+\gamma D+\delta(X \times D)+\varepsilon
$$

- Then, for women

$$
\begin{aligned}
Y & =\alpha+\beta X+\gamma \times 0+\delta(X \times 0)+\varepsilon \\
E(Y) & =\alpha+\beta X
\end{aligned}
$$

## Linear Models in R

Review of Dummy-Variable Regression

- Different slopes for women and men ("different slopes for different folks") can be modelled by introducing an interaction regressor, the product of $X$ and $D$, into the model:

$$
Y=\alpha+\beta X+\gamma D+\delta(X \times D)+\varepsilon
$$

- Then, for women

$$
\begin{aligned}
Y & =\alpha+\beta X+\gamma \times 0+\delta(X \times 0)+\varepsilon \\
E(Y) & =\alpha+\beta X
\end{aligned}
$$

- And, for men

$$
Y=\alpha+\beta X+\gamma \times 1+\delta(X \times 1)+\varepsilon
$$

$E(Y)=(\alpha+\gamma)+(\beta+\delta) X$

## Linear Models in R

Review of Dummy-Variable Regression

- Different slopes for women and men ("different slopes for different folks") can be modelled by introducing an interaction regressor, the product of $X$ and $D$, into the model:

$$
Y=\alpha+\beta X+\gamma D+\delta(X \times D)+\varepsilon
$$

- Then, for women

$$
\begin{aligned}
Y & =\alpha+\beta X+\gamma \times 0+\delta(X \times 0)+\varepsilon \\
E(Y) & =\alpha+\beta X
\end{aligned}
$$

- And, for men

$$
\begin{aligned}
Y & =\alpha+\beta X+\gamma \times 1+\delta(X \times 1)+\varepsilon \\
E(Y) & =(\alpha+\gamma)+(\beta+\delta) X
\end{aligned}
$$

## Linear Models in R

Review of Dummy-Variable Regression

- Different slopes for women and men ("different slopes for different folks") can be modelled by introducing an interaction regressor, the product of $X$ and $D$, into the model:

$$
Y=\alpha+\beta X+\gamma D+\delta(X \times D)+\varepsilon
$$

- Then, for women

$$
\begin{aligned}
Y & =\alpha+\beta X+\gamma \times 0+\delta(X \times 0)+\varepsilon \\
E(Y) & =\alpha+\beta X
\end{aligned}
$$

- And, for men

$$
Y=\alpha+\beta X+\gamma \times 1+\delta(X \times 1)+\varepsilon
$$

$$
E(Y)=(\alpha+\gamma)+(\beta+\delta) X
$$

## Linear Models in R

- Polytomous explanatory variables-i.e., factors with more than two levels—are handled by creating a set of dummy regressors, one fewer than the number of levels.


## Linear Models in R

- Polytomous explanatory variables-i.e., factors with more than two levels—are handled by creating a set of dummy regressors, one fewer than the number of levels.
- For example, for gender with levels female, male, and nonbinary, we can code two dummy regressors:


## Linear Models in R

- Polytomous explanatory variables—i.e., factors with more than two levels—are handled by creating a set of dummy regressors, one fewer than the number of levels.
- For example, for gender with levels female, male, and nonbinary, we can code two dummy regressors:

| Gender | $D_{1}$ | $D_{2}$ |
| :--- | :--- | :--- |
| female | 0 | 0 |
| male | 1 | 0 |
| nonbinary | 0 | 1 |

## Linear Models in R

- Then we can fit the model

$$
Y=\alpha+\beta X+\gamma_{1} D_{1}+\gamma_{2} D_{3}+\delta_{1}\left(X \times D_{1}\right)+\delta_{2}\left(X \times D_{2}\right)+\varepsilon
$$

## Linear Models in R

- Then we can fit the model

$$
Y=\alpha+\beta X+\gamma_{1} D_{1}+\gamma_{2} D_{3}+\delta_{1}\left(X \times D_{1}\right)+\delta_{2}\left(X \times D_{2}\right)+\varepsilon
$$

- and

$$
\begin{aligned}
\text { female : } E(Y) & =\alpha+\beta X+\gamma_{1} \times 0+\gamma_{2} \times 0+\delta_{1}(X \times 0)+\delta_{2}(X \times 0) \\
& =\alpha+\beta X \\
\text { male }: E(Y) & =\alpha+\beta X+\gamma_{1} \times 1+\gamma_{2} \times 0+\delta_{1}(X \times 1)+\delta_{2}(X \times 0) \\
& =\left(\alpha+\gamma_{1}\right)+\left(\beta+\delta_{1}\right) X \\
\text { nonbinary : } E(Y) & =\alpha+\beta X+\gamma_{1} \times 0+\gamma_{2} \times 1+\delta_{1}(X \times 0)+\delta_{2}(X \times 1) \\
& =\left(\alpha+\gamma_{2}\right)+\left(\beta+\delta_{2}\right) X
\end{aligned}
$$

## Linear Models in R

- Type II tests are constructed in conformity to the principle of marginality: Each term in the model is tested assuming that its higher-order relatives are zero (and hence are ignored).


## Linear Models in R

- Type II tests are constructed in conformity to the principle of marginality: Each term in the model is tested assuming that its higher-order relatives are zero (and hence are ignored).
- Thus, a main effect (e.g., $X$ ) is tested assuming that the interaction or interactions to which the main effect is marginal (e.g., $\mathrm{X}: \mathrm{A}, \mathrm{X}: \mathrm{A}: \mathrm{B}$ ) are zero.


## Linear Models in R

- Type II tests are constructed in conformity to the principle of marginality: Each term in the model is tested assuming that its higher-order relatives are zero (and hence are ignored).
- Thus, a main effect (e.g., $X$ ) is tested assuming that the interaction or interactions to which the main effect is marginal (e.g., X:A, X:A:B) are zero.
- For example, consider the model $y \sim a * b * c$ or in longer form $\mathrm{y} \sim 1+\mathrm{a}+\mathrm{b}+\mathrm{c}+\mathrm{a}: \mathrm{b}+\mathrm{a}: \mathrm{c}+\mathrm{b}: \mathrm{c}+\mathrm{a}: \mathrm{b}: \mathrm{c}$.


## Linear Models in R

Type-II Tests for Linear (and Other) Models

- For Type-II tests of all terms, we implicitly fit the following models (all in longer form):

| Model | Formula |
| :--- | :--- |
| 1 | $\mathrm{y} \sim 1+\mathrm{a}+\mathrm{b}+\mathrm{c}+\mathrm{a}: \mathrm{b}+\mathrm{a}: \mathrm{c}+\mathrm{b}: \mathrm{c}+\mathrm{a}: \mathrm{b}: \mathrm{c}$ |
| 2 | $\mathrm{y} \sim 1+\mathrm{a}+\mathrm{b}+\mathrm{c}+\mathrm{a}: \mathrm{b}+\mathrm{a}: \mathrm{c}+\mathrm{b}: \mathrm{c}$ |
| 3 | $\mathrm{y} \sim 1+\mathrm{a}+\mathrm{b}+\mathrm{c}+\mathrm{a}: \mathrm{c}+\mathrm{b}: \mathrm{c}$ |
| 4 | $\mathrm{y} \sim 1+\mathrm{a}+\mathrm{b}+\mathrm{c}+\mathrm{a}: \mathrm{b}+\mathrm{b}: \mathrm{c}$ |
| 5 | $\mathrm{y} \sim 1+\mathrm{a}+\mathrm{b}+\mathrm{c}+\mathrm{a}: \mathrm{b}+\mathrm{a}: \mathrm{c}$ |
| 6 | $\mathrm{y} \sim 1+\mathrm{a}+\mathrm{b}+\mathrm{c}+\mathrm{b}: \mathrm{c}$ |
| 7 | $\mathrm{y} \sim 1+\mathrm{b}+\mathrm{c}+\mathrm{b}: \mathrm{c}$ |
| 8 | $\mathrm{y} \sim 1+\mathrm{a}+\mathrm{b}+\mathrm{c}+\mathrm{a}: \mathrm{c}$ |
| 9 | $\mathrm{y} \sim 1+\mathrm{a}+\mathrm{c}+\mathrm{a}: \mathrm{c}$ |
| 10 | $\mathrm{y} \sim 1+\mathrm{a}+\mathrm{b}+\mathrm{c}+\mathrm{a}: \mathrm{b}$ |
| 11 | $\mathrm{y} \sim 1+\mathrm{a}+\mathrm{b}+\mathrm{a}: \mathrm{b}$ |

## Linear Models in R

Type-II Tests for Linear (and Other) Models

- Contrasting pairs of models by subtracting the regression sum of squares for the smaller model from that for the larger model produces the Type-II ANOVA table:

| Term | Models Contrasted |
| :---: | :---: |
| a | $6-7$ |
| b | $8-9$ |
| c | $10-11$ |
| $\mathrm{a}: \mathrm{b}$ | $2-3$ |
| $\mathrm{a}: \mathrm{c}$ | $2-4$ |
| $\mathrm{~b}: \mathrm{c}$ | $2-5$ |
| $\mathrm{a}: \mathrm{b}: \mathrm{c}$ | $1-2$ |

## Linear Models in R

Type-II Tests for Linear (and Other) Models

- Contrasting pairs of models by subtracting the regression sum of squares for the smaller model from that for the larger model produces the Type-II ANOVA table:

| Term | Models Contrasted |
| :---: | :---: |
| a | $6-7$ |
| b | $8-9$ |
| c | $10-11$ |
| $\mathrm{a}: \mathrm{b}$ | $2-3$ |
| $\mathrm{a}: \mathrm{c}$ | $2-4$ |
| $\mathrm{~b}: \mathrm{c}$ | $2-5$ |
| $\mathrm{a}: \mathrm{b}: \mathrm{c}$ | $1-2$ |

- The degrees of freedom for each term are the number of regressors used for that term.


## Linear Models in R

Type-II Tests for Linear (and Other) Models

- Contrasting pairs of models by subtracting the regression sum of squares for the smaller model from that for the larger model produces the Type-II ANOVA table:

| Term | Models Contrasted |
| :---: | :---: |
| a | $6-7$ |
| b | $8-9$ |
| c | $10-11$ |
| $\mathrm{a}: \mathrm{b}$ | $2-3$ |
| $\mathrm{a}: \mathrm{c}$ | $2-4$ |
| $\mathrm{~b}: \mathrm{c}$ | $2-5$ |
| $\mathrm{a}: \mathrm{b}: \mathrm{c}$ | $1-2$ |

- The degrees of freedom for each term are the number of regressors used for that term.
- The estimated error variance used for the denominator of the $F$-tests comes from the largest model fit to the data, here Model 1, and the denominator degrees of freedom for $F$ are the residual degrees of freedom for this model.


## Linear Models in R

Arguments of the $\operatorname{lm}()$ Function

- lm(formula, data, subset, weights, na.action, method = "qr", model $=$ TRUE, $\mathrm{x}=$ FALSE, $\mathrm{y}=$ FALSE, $\mathrm{qr}=$ TRUE, singular.ok $=$ TRUE, contrasts $=$ NULL, offset, ...)


## Linear Models in R

- lm(formula, data, subset, weights, na.action, method = "qr", model $=$ TRUE, $\mathrm{x}=$ FALSE, $\mathrm{y}=$ FALSE, $\mathrm{qr}=$ TRUE, singular. ok $=$ TRUE, contrasts $=$ NULL, offset, ...)
- Operators for the formula argument:

| Expression | Interpretation | Example |
| :--- | :--- | :--- |
| $\mathrm{A}+\mathrm{B}$ | include both A and B | income + education |
| $\mathrm{A}-\mathrm{B}$ | exclude B from A | $\mathrm{a} * \mathrm{~b} * \mathrm{~d}-\mathrm{a}: \mathrm{b}: \mathrm{d}$ |
| $\mathrm{A}: \mathrm{B}$ | interaction of A and B | type:education |
| $\mathrm{A} * \mathrm{~B}$ | $\mathrm{~A}+\mathrm{B}+\mathrm{A}: \mathrm{B}$ | type*education |
| $\mathrm{B} \%$ in\% A | B nested within A | education \%in\% type |
| $\mathrm{A} / \mathrm{B}$ | $\mathrm{A}+\mathrm{B} \%$ in\% A | type/education |
| $\mathrm{A} \wedge \mathrm{k}$ | effects crossed to order k | $(\mathrm{a}+\mathrm{b}+\mathrm{d})^{\wedge} 2$ |

## Linear Models in R

Arguments of the $\operatorname{lm}()$ Function

- data: A data frame containing the data for the model.


## Linear Models in R

Arguments of the $\operatorname{lm}()$ Function

- data: A data frame containing the data for the model.
- subset:


## Linear Models in R

Arguments of the $\operatorname{lm}()$ Function

- data: A data frame containing the data for the model.
- subset:
- a logical vector: subset $=$ gender $==$ "F"


## Linear Models in R

Arguments of the $\operatorname{lm}()$ Function

- data: A data frame containing the data for the model.
- subset:
- a logical vector: subset $=$ gender $==$ "F"
- a numeric vector of observation indices: subset $=1: 100$


## Linear Models in R

Arguments of the $\operatorname{lm}()$ Function

- data: A data frame containing the data for the model.
- subset:
- a logical vector: subset $=$ gender $==$ "F"
- a numeric vector of observation indices: subset $=1: 100$
- a negative numeric vector with observations to be omitted: subset $=-c(6,16)$


## Linear Models in R

Arguments of the $\operatorname{lm}$ () Function

- data: A data frame containing the data for the model.
- subset:
- a logical vector: subset $=$ gender $==$ "F"
- a numeric vector of observation indices: subset $=1: 100$
- a negative numeric vector with observations to be omitted: subset $=-c(6,16)$
- weights: for weighted-least-squares regression


## Linear Models in R

- data: A data frame containing the data for the model.
- subset:
- a logical vector: subset = gender == "F"
- a numeric vector of observation indices: subset $=1: 100$
- a negative numeric vector with observations to be omitted: subset $=-c(6,16)$
- weights: for weighted-least-squares regression
- na.action: name of a function to handle missing data; default given by the na.action option, initially "na.omit"


## Linear Models in R

- data: A data frame containing the data for the model.
- subset:
- a logical vector: subset = gender == "F"
- a numeric vector of observation indices: subset $=1: 100$
- a negative numeric vector with observations to be omitted: subset $=-c(6,16)$
- weights: for weighted-least-squares regression
- na.action: name of a function to handle missing data; default given by the na.action option, initially "na.omit"
- method, model, x, y, qr, singular.ok: technical arguments


## Linear Models in R

- data: A data frame containing the data for the model.
- subset:
- a logical vector: subset = gender == "F"
- a numeric vector of observation indices: subset $=1: 100$
- a negative numeric vector with observations to be omitted: subset $=-c(6,16)$
- weights: for weighted-least-squares regression
- na.action: name of a function to handle missing data; default given by the na.action option, initially "na.omit"
- method, model, x, y, qr, singular.ok: technical arguments
- contrasts: specify a list of contrasts for factors; e.g., contrasts=list(partner.status=contr.sum, fcategory=contr.poly))


## Linear Models in R

- data: A data frame containing the data for the model.
- subset:
- a logical vector: subset = gender == "F"
- a numeric vector of observation indices: subset $=1: 100$
- a negative numeric vector with observations to be omitted: subset $=-c(6,16)$
- weights: for weighted-least-squares regression
- na.action: name of a function to handle missing data; default given by the na.action option, initially "na.omit"
- method, model, x, y, qr, singular.ok: technical arguments
- contrasts: specify a list of contrasts for factors; e.g., contrasts=list(partner.status=contr.sum, fcategory=contr.poly))
- offset: term added to the right-hand-side of the model with a fixed coefficient of 1.


## Linear Models in R

- Influence on the regression coefficients $=$ leverage $\times$ outlyingness.


## Linear Models in R

Regression Diagnostics: Unusual Cases

- Influence on the regression coefficients $=$ leverage $\times$ outlyingness.
- Hat-values measure leverage.


## Linear Models in R

- Influence on the regression coefficients $=$ leverage $\times$ outlyingness.
- Hat-values measure leverage.
- The fitted linear regression model in matrix form is $\mathrm{y}=\mathrm{Xb}+\mathrm{e}$, where y is the $(n \times 1)$ response vector, X is the $(n \times p)$ model matrix, and $\mathrm{b}=\left(\mathrm{X}^{T} \mathrm{X}\right)^{-1} \mathrm{X}^{T} \mathrm{y}$ is the $(p \times 1)$ vector of least squares coefficients.


## Linear Models in R

- Influence on the regression coefficients $=$ leverage $\times$ outlyingness.
- Hat-values measure leverage.
- The fitted linear regression model in matrix form is $\mathrm{y}=\mathrm{Xb}+\mathrm{e}$, where y is the $(n \times 1)$ response vector, X is the $(n \times p)$ model matrix, and $\mathrm{b}=\left(\mathrm{X}^{T} \mathrm{X}\right)^{-1} \mathrm{X}^{T} \mathrm{y}$ is the $(p \times 1)$ vector of least squares coefficients.
- The fitted values are then $\hat{\mathrm{y}}=\mathrm{Xb}=\mathrm{X}\left(\mathrm{X}^{T} \mathrm{X}\right)^{-1} \mathrm{X}^{T} \mathrm{y}=\mathrm{Hy}$, where the $(n \times n)$ hat-matrix is $H=X\left(X^{T} X\right)^{-1} X^{T}$.


## Linear Models in R

- Influence on the regression coefficients $=$ leverage $\times$ outlyingness.
- Hat-values measure leverage.
- The fitted linear regression model in matrix form is $\mathrm{y}=\mathrm{Xb}+\mathrm{e}$, where y is the $(n \times 1)$ response vector, X is the $(n \times p)$ model matrix, and $\mathrm{b}=\left(\mathrm{X}^{T} \mathrm{X}\right)^{-1} \mathrm{X}^{T} \mathrm{y}$ is the $(p \times 1)$ vector of least squares coefficients.
- The fitted values are then $\hat{y}=\mathrm{Xb}=\mathrm{X}\left(\mathrm{X}^{T} \mathrm{X}\right)^{-1} \mathrm{X}^{T} \mathrm{y}=\mathrm{Hy}$, where the $(n \times n)$ hat-matrix is $H=X\left(X^{T} X\right)^{-1} X^{T}$.
- The $h_{i j}$ element of H gives the weight of $Y_{j}$ in determining $\widehat{Y}_{i}$.


## Linear Models in R

- Influence on the regression coefficients $=$ leverage $\times$ outlyingness.
- Hat-values measure leverage.
- The fitted linear regression model in matrix form is $\mathrm{y}=\mathrm{Xb}+\mathrm{e}$, where y is the $(n \times 1)$ response vector, X is the $(n \times p)$ model matrix, and $\mathrm{b}=\left(\mathrm{X}^{T} \mathrm{X}\right)^{-1} \mathrm{X}^{T} \mathrm{y}$ is the $(p \times 1)$ vector of least squares coefficients.
- The fitted values are then $\hat{y}=\mathrm{Xb}=\mathrm{X}\left(\mathrm{X}^{T} \mathrm{X}\right)^{-1} \mathrm{X}^{T} \mathrm{y}=\mathrm{Hy}$, where the $(n \times n)$ hat-matrix is $H=X\left(X^{T} X\right)^{-1} X^{T}$.
- The $h_{i j}$ element of H gives the weight of $Y_{j}$ in determining $\widehat{Y}_{i}$.
- The H matrix is symmetric $\left(\mathrm{H}=\mathrm{H}^{T}\right)$ and idempotent $\left(\mathrm{H}^{2}=\mathrm{H}\right)$, and it follows that the $j$ th diagonal element of $\mathrm{H}, h_{j}=h_{j j}=\sum_{i=1}^{n} h_{i j}^{2}$ summarizes the size of all of the elements in the $j$ th column of of H and hence the leverage of the $j$ th case in determining the fit.


## Linear Models in R

- Influence on the regression coefficients $=$ leverage $\times$ outlyingness.
- Hat-values measure leverage.
- The fitted linear regression model in matrix form is $\mathrm{y}=\mathrm{Xb}+\mathrm{e}$, where y is the $(n \times 1)$ response vector, X is the $(n \times p)$ model matrix, and $\mathrm{b}=\left(\mathrm{X}^{T} \mathrm{X}\right)^{-1} \mathrm{X}^{T} \mathrm{y}$ is the $(p \times 1)$ vector of least squares coefficients.
- The fitted values are then $\hat{y}=\mathrm{Xb}=\mathrm{X}\left(\mathrm{X}^{T} \mathrm{X}\right)^{-1} \mathrm{X}^{T} \mathrm{y}=\mathrm{Hy}$, where the $(n \times n)$ hat-matrix is $H=X\left(X^{T} X\right)^{-1} X^{T}$.
- The $h_{i j}$ element of H gives the weight of $Y_{j}$ in determining $\widehat{Y}_{i}$.
- The H matrix is symmetric $\left(\mathrm{H}=\mathrm{H}^{T}\right)$ and idempotent $\left(\mathrm{H}^{2}=\mathrm{H}\right)$, and it follows that the $j$ th diagonal element of $\mathrm{H}, h_{j}=h_{j j}=\sum_{i=1}^{n} h_{i j}^{2}$ summarizes the size of all of the elements in the $j$ th column of of H and hence the leverage of the $j$ th case in determining the fit.
- The diagonal entries $h_{j}$ of H are the hat-values.


## Linear Models in R

- Influence on the regression coefficients $=$ leverage $\times$ outlyingness.
- Hat-values measure leverage.
- The fitted linear regression model in matrix form is $\mathrm{y}=\mathrm{Xb}+\mathrm{e}$, where y is the $(n \times 1)$ response vector, X is the $(n \times p)$ model matrix, and $\mathrm{b}=\left(\mathrm{X}^{T} \mathrm{X}\right)^{-1} \mathrm{X}^{T} \mathrm{y}$ is the $(p \times 1)$ vector of least squares coefficients.
- The fitted values are then $\hat{y}=\mathrm{Xb}=\mathrm{X}\left(\mathrm{X}^{T} \mathrm{X}\right)^{-1} \mathrm{X}^{T} \mathrm{y}=\mathrm{Hy}$, where the $(n \times n)$ hat-matrix is $H=X\left(X^{T} X\right)^{-1} X^{T}$.
- The $h_{i j}$ element of H gives the weight of $Y_{j}$ in determining $\widehat{Y}_{i}$.
- The H matrix is symmetric $\left(\mathrm{H}=\mathrm{H}^{T}\right)$ and idempotent $\left(\mathrm{H}^{2}=\mathrm{H}\right)$, and it follows that the $j$ th diagonal element of $\mathrm{H}, h_{j}=h_{j j}=\sum_{i=1}^{n} h_{i j}^{2}$ summarizes the size of all of the elements in the $j$ th column of of H and hence the leverage of the $j$ th case in determining the fit.
- The diagonal entries $h_{j}$ of H are the hat-values.
- The hat-values are bounded between $1 / n$ (if the model has an intercept, otherwise 0 ) and 1 , and the average hat-values is $\bar{h}=p / n$.


## Linear Models in R

- Studentized residuals measure outlyingness.


## Linear Models in R

- Studentized residuals measure outlyingness.
- The studentized residuals are

$$
E_{T i}=\frac{E_{i}}{S_{E(-i)} \sqrt{1-h_{i}}}
$$

where $E_{i}$ is the $i$ th element of the least-squares residuals vector e and $S_{E(-i)}$ is the standard deviation of the residuals when the regression is refit with the $i$ th case removed.

## Linear Models in R

- Studentized residuals measure outlyingness.
- The studentized residuals are

$$
E_{T i}=\frac{E_{i}}{S_{E(-i)} \sqrt{1-h_{i}}}
$$

where $E_{i}$ is the $i$ th element of the least-squares residuals vector e and $S_{E(-i)}$ is the standard deviation of the residuals when the regression is refit with the $i$ th case removed.

- If the model is correct, then each studentized residual is distributed at $t$ with $n-p-1$ degrees of freedom, providing a basis for an outlier test based on the the largest absolute studentized residual.


## Linear Models in R

- Studentized residuals measure outlyingness.
- The studentized residuals are

$$
E_{T i}=\frac{E_{i}}{S_{E(-i)} \sqrt{1-h_{i}}}
$$

where $E_{i}$ is the $i$ th element of the least-squares residuals vector e and $S_{E(-i)}$ is the standard deviation of the residuals when the regression is refit with the $i$ th case removed.

- If the model is correct, then each studentized residual is distributed at $t$ with $n-p-1$ degrees of freedom, providing a basis for an outlier test based on the the largest absolute studentized residual.
- But because there are $n$ studentized residuals, it's necessary to correct for simultaneous statistical inference-e.g., a Bonferroni correction, which multiplies the two-sided $P$-value for the $t$-test by $n$.


## Linear Models in R

- Measuring influence on the regression coefficients with dfbeta and Cook's $D$ :


## Linear Models in R

- Measuring influence on the regression coefficients with dfbeta and Cook's D:
- The most direct measure is to refit the model without the $i$ th case and see how the coefficients change.


## Linear Models in R

- Measuring influence on the regression coefficients with dfbeta and Cook's $D$ :
- The most direct measure is to refit the model without the ith case and see how the coefficients change.
- The answer is dfbeta ${ }_{i}=\mathrm{b}-\mathrm{b}_{(-i)}=\left(\mathrm{X}^{T} \mathrm{X}\right)^{-1} \mathrm{x}_{i} E_{i} /\left(1-h_{i}\right)$, where $\mathrm{b}_{(-i)}$ is the vector of least-squares coefficients computed with the $i$ th case deleted, and $x_{i}$ is the ith row of $X$ (written as a column vector).


## Linear Models in R

- Measuring influence on the regression coefficients with dfbeta and Cook's D:
- The most direct measure is to refit the model without the ith case and see how the coefficients change.
- The answer is dfbeta ${ }_{i}=\mathrm{b}-\mathrm{b}_{(-i)}=\left(\mathrm{X}^{T} \mathrm{X}\right)^{-1} \mathrm{x}_{i} E_{i} /\left(1-h_{i}\right)$, where $\mathrm{b}_{(-i)}$ is the vector of least-squares coefficients computed with the $i$ th case deleted, and $x_{i}$ is the $i$ th row of $X$ (written as a column vector).
- Because there are a lot $(n \times p)$ of dfbeta $_{i j}$, it's useful to summarize the $p$ values for each case $i$. The most common such measure is Cook's distance:

$$
\begin{aligned}
D_{i} & =\frac{\text { dfbeta }_{i}^{T} \mathrm{X}^{T} \mathrm{X} \text { dfbeta }_{i}}{p S_{E}^{2}}=\frac{\left(\widehat{y}-\widehat{y}_{(-i)}\right)^{T}\left(\hat{\mathrm{y}}-\widehat{\mathrm{y}}_{(-i)}\right)}{p S_{E}^{2}} \approx \frac{E_{T i}^{2}}{p} \times \frac{h_{i}}{1-h_{i}} \\
& =\text { outlyingness } \times \text { leverage }
\end{aligned}
$$

where $\widehat{y}_{(-i)}$ is the vector of fitted values computed when the ith case is removed.

## Linear Models in R

- Added-variable plots visualize leverage, outlyingness, and influence on each regression coefficient, reducing the $p$-dimensional scatterplot of the data to a series of $p$ two-dimensional scatterplots, one for each coefficient.


## Linear Models in R

- Added-variable plots visualize leverage, outlyingness, and influence on each regression coefficient, reducing the $p$-dimensional scatterplot of the data to a series of $p$ two-dimensional scatterplots, one for each coefficient.
- For example, focusing on the coefficient $B_{1}$ of $X_{1}$ in the regression

$$
Y=A+B_{1} X_{1}+B_{2} X_{2}+\cdots+B_{k} X_{k}+E(\text { so } p=k+1):
$$

## Linear Models in R

- Added-variable plots visualize leverage, outlyingness, and influence on each regression coefficient, reducing the $p$-dimensional scatterplot of the data to a series of $p$ two-dimensional scatterplots, one for each coefficient.
- For example, focusing on the coefficient $B_{1}$ of $X_{1}$ in the regression

$$
Y=A+B_{1} X_{1}+B_{2} X_{2}+\cdots+B_{k} X_{k}+E(\text { so } p=k+1):
$$

- Regress $Y$ on $X_{2}, \ldots, X_{k}$ (and an intercept), obtaining residuals $E^{\left(Y_{1}\right)}$ (i.e., what remains of $Y$ when the effects of $X_{2}, \ldots, X_{k}$ are removed).


## Linear Models in R

- Added-variable plots visualize leverage, outlyingness, and influence on each regression coefficient, reducing the $p$-dimensional scatterplot of the data to a series of $p$ two-dimensional scatterplots, one for each coefficient.
- For example, focusing on the coefficient $B_{1}$ of $X_{1}$ in the regression

$$
Y=A+B_{1} X_{1}+B_{2} X_{2}+\cdots+B_{k} X_{k}+E(\text { so } p=k+1):
$$

- Regress $Y$ on $X_{2}, \ldots, X_{k}$ (and an intercept), obtaining residuals $E^{\left(Y_{1}\right)}$ (i.e., what remains of $Y$ when the effects of $X_{2}, \ldots, X_{k}$ are removed).
- Regress $X_{1}$ on $X_{2}, \ldots, X_{k}$ (and an intercept), obtaining residuals $E^{\left(X_{1}\right)}$ (i.e., what remains of $X_{1}$ when the effects of $X_{2}, \ldots, X_{k}$ are removed).


## Linear Models in R

- Added-variable plots visualize leverage, outlyingness, and influence on each regression coefficient, reducing the $p$-dimensional scatterplot of the data to a series of $p$ two-dimensional scatterplots, one for each coefficient.
- For example, focusing on the coefficient $B_{1}$ of $X_{1}$ in the regression

$$
Y=A+B_{1} X_{1}+B_{2} X_{2}+\cdots+B_{k} X_{k}+E(\text { so } p=k+1):
$$

- Regress $Y$ on $X_{2}, \ldots, X_{k}$ (and an intercept), obtaining residuals $E^{\left(Y_{1}\right)}$ (i.e., what remains of $Y$ when the effects of $X_{2}, \ldots, X_{k}$ are removed).
- Regress $X_{1}$ on $X_{2}, \ldots, X_{k}$ (and an intercept), obtaining residuals $E^{\left(X_{1}\right)}$ (i.e., what remains of $X_{1}$ when the effects of $X_{2}, \ldots, X_{k}$ are removed).
- plot $E^{\left(Y_{1}\right)}$ versus $E^{\left(X_{1}\right)}$.


## Linear Models in R

- Added-variable plots visualize leverage, outlyingness, and influence on each regression coefficient, reducing the $p$-dimensional scatterplot of the data to a series of $p$ two-dimensional scatterplots, one for each coefficient.
- For example, focusing on the coefficient $B_{1}$ of $X_{1}$ in the regression

$$
Y=A+B_{1} X_{1}+B_{2} X_{2}+\cdots+B_{k} X_{k}+E(\text { so } p=k+1):
$$

- Regress $Y$ on $X_{2}, \ldots, X_{k}$ (and an intercept), obtaining residuals $E^{\left(Y_{1}\right)}$ (i.e., what remains of $Y$ when the effects of $X_{2}, \ldots, X_{k}$ are removed).
- Regress $X_{1}$ on $X_{2}, \ldots, X_{k}$ (and an intercept), obtaining residuals $E^{\left(X_{1}\right)}$ (i.e., what remains of $X_{1}$ when the effects of $X_{2}, \ldots, X_{k}$ are removed).
- plot $E^{\left(Y_{1}\right)}$ versus $E^{\left(X_{1}\right)}$.
- Repeat for each of $X_{2}, \ldots, X_{k}$ (and even, if desired, for the constant regressor, $X_{0}=1$ ).


## Linear Models in R

- The AV plot for $X_{j}$ has the following remarkable properties:


## Linear Models in R

- The AV plot for $X_{j}$ has the following remarkable properties:
- The slope of the least-squares line in the plot is the coefficient $B_{j}$ of $X_{j}$ in the multiple regression.


## Linear Models in R

- The AV plot for $X_{j}$ has the following remarkable properties:
- The slope of the least-squares line in the plot is the coefficient $B_{j}$ of $X_{j}$ in the multiple regression.
- The residuals from this line are the same as the residuals $E_{i}$ in the multiple regression.


## Linear Models in R

- The AV plot for $X_{j}$ has the following remarkable properties:
- The slope of the least-squares line in the plot is the coefficient $B_{j}$ of $X_{j}$ in the multiple regression.
- The residuals from this line are the same as the residuals $E_{i}$ in the multiple regression.
- The horizontal variation of $X_{j}$ in the plot is its conditional variation holding the other $X_{\mathrm{s}}$ constant: $S_{X_{j} \mid \text { other } X_{\mathrm{s}}}^{2}=\sum E^{\left(X_{j}\right)^{2}} /(n-k)$.


## Linear Models in R

- The AV plot for $X_{j}$ has the following remarkable properties:
- The slope of the least-squares line in the plot is the coefficient $B_{j}$ of $X_{j}$ in the multiple regression.
- The residuals from this line are the same as the residuals $E_{i}$ in the multiple regression.
- The horizontal variation of $X_{j}$ in the plot is its conditional variation holding the other $X \mathrm{~s}$ constant: $S_{X_{j} \mid \text { other } X_{\mathrm{s}}}^{2}=\sum E^{\left(X_{j}\right)^{2}} /(n-k)$.
- Consequently, the standard error of $B_{j}$ computed from the simple regression corresponding to the plot, $\mathrm{SE}\left(B_{j}\right)=S_{E} / \sqrt{\sum E^{\left(X_{j}\right)^{2}}}$ is the same as the standard error of $B_{j}$ from the multiple regression.


## Linear Models in R

Regression Diagnostics: Component-Plus-Residuals ( $C+R$ ) Plots

- Component-plus-Residuals plots are even a simpler way of reducing the $p$-dimensional scatterplot to a series of 2D plots:


## Linear Models in R

Regression Diagnostics: Component-Plus-Residuals ( $C+R$ ) Plots

- Component-plus-Residuals plots are even a simpler way of reducing the $p$-dimensional scatterplot to a series of 2D plots:
- Add the residuals from the full regression to the linear component representing $X_{1}$ to form the partial residuals: $E^{(1)}=B_{1} X_{1}+E$.


## Linear Models in R

Regression Diagnostics: Component-Plus-Residuals ( $C+R$ ) Plots

- Component-plus-Residuals plots are even a simpler way of reducing the p-dimensional scatterplot to a series of 2D plots:
- Add the residuals from the full regression to the linear component representing $X_{1}$ to form the partial residuals: $E^{(1)}=B_{1} X_{1}+E$.
- Plot $E^{(1)}$ versus $X_{1}$, enhancing the graph with a scatterplot smoother (nonparametric regression line) to judge nonlinearity.


## Linear Models in R

Regression Diagnostics: Component-Plus-Residuals ( $C+R$ ) Plots

- Component-plus-Residuals plots are even a simpler way of reducing the p-dimensional scatterplot to a series of 2D plots:
- Add the residuals from the full regression to the linear component representing $X_{1}$ to form the partial residuals: $E^{(1)}=B_{1} X_{1}+E$.
- Plot $E^{(1)}$ versus $X_{1}$, enhancing the graph with a scatterplot smoother (nonparametric regression line) to judge nonlinearity.
- By construction, the least-squares slope of the $\mathrm{C}+\mathrm{R}$ plot for $X_{1}$ is $B_{1}$ from the multiple regression, and the residuals in the $\mathrm{C}+\mathrm{R}$ plot are just the Es.


## Linear Models in R

- Component-plus-Residuals plots are even a simpler way of reducing the p-dimensional scatterplot to a series of 2D plots:
- Add the residuals from the full regression to the linear component representing $X_{1}$ to form the partial residuals: $E^{(1)}=B_{1} X_{1}+E$.
- Plot $E^{(1)}$ versus $X_{1}$, enhancing the graph with a scatterplot smoother (nonparametric regression line) to judge nonlinearity.
- By construction, the least-squares slope of the $C+R$ plot for $X_{1}$ is $B_{1}$ from the multiple regression, and the residuals in the $\mathrm{C}+\mathrm{R}$ plot are just the Es .
- Under certain reasonably general (but not bulletproof) circumstances, if the partial relationship between $Y$ and $X_{1}$ is nonlinear but incorrectly modelled as linear, the nature of the nonlinearity will be apparent in the $\mathrm{C}+\mathrm{R}$ plot for $X_{1}$.


## Linear Models in R

- Component-plus-Residuals plots are even a simpler way of reducing the p-dimensional scatterplot to a series of 2D plots:
- Add the residuals from the full regression to the linear component representing $X_{1}$ to form the partial residuals: $E^{(1)}=B_{1} X_{1}+E$.
- Plot $E^{(1)}$ versus $X_{1}$, enhancing the graph with a scatterplot smoother (nonparametric regression line) to judge nonlinearity.
- By construction, the least-squares slope of the $C+R$ plot for $X_{1}$ is $B_{1}$ from the multiple regression, and the residuals in the $\mathrm{C}+\mathrm{R}$ plot are just the Es .
- Under certain reasonably general (but not bulletproof) circumstances, if the partial relationship between $Y$ and $X_{1}$ is nonlinear but incorrectly modelled as linear, the nature of the nonlinearity will be apparent in the $\mathrm{C}+\mathrm{R}$ plot for $X_{1}$.
- Repeat for each of $X_{2}, \ldots, X_{k}$.


## Linear Models in R

- It's often possible to linearize a nonlinear relationship between $Y$ and $X$ by transforming one or the other (or both) by a power transformation.


## Linear Models in R

- It's often possible to linearize a nonlinear relationship between $Y$ and $X$ by transforming one or the other (or both) by a power transformation.
- By power transformations, I mean $X \rightarrow X^{p}$ or similarly for $Y$.


## Linear Models in R

- It's often possible to linearize a nonlinear relationship between $Y$ and $X$ by transforming one or the other (or both) by a power transformation.
- By power transformations, I mean $X \rightarrow X^{p}$ or similarly for $Y$.
- The power $p$ may be positive or negative, and need not be a whole number.


## Linear Models in R

- It's often possible to linearize a nonlinear relationship between $Y$ and $X$ by transforming one or the other (or both) by a power transformation.
- By power transformations, I mean $X \rightarrow X^{p}$ or similarly for $Y$.
- The power $p$ may be positive or negative, and need not be a whole number.
- For example, $X^{1 / 2}=\sqrt{X}$ and $X^{-1}=1 / X$.


## Linear Models in R

- It's often possible to linearize a nonlinear relationship between $Y$ and $X$ by transforming one or the other (or both) by a power transformation.
- By power transformations, I mean $X \rightarrow X^{p}$ or similarly for $Y$.
- The power $p$ may be positive or negative, and need not be a whole number.
- For example, $X^{1 / 2}=\sqrt{X}$ and $X^{-1}=1 / X$.
- $p=1$ is no transformation: $X^{1}=X$.


## Linear Models in R

- It's often possible to linearize a nonlinear relationship between $Y$ and $X$ by transforming one or the other (or both) by a power transformation.
- By power transformations, I mean $X \rightarrow X^{p}$ or similarly for $Y$.
- The power $p$ may be positive or negative, and need not be a whole number.
- For example, $X^{1 / 2}=\sqrt{X}$ and $X^{-1}=1 / X$.
- $p=1$ is no transformation: $X^{1}=X$.
- If $p=0$, we use $\log (X)$.


## Linear Models in R

- It's often possible to linearize a nonlinear relationship between $Y$ and $X$ by transforming one or the other (or both) by a power transformation.
- By power transformations, I mean $X \rightarrow X^{p}$ or similarly for $Y$.
- The power $p$ may be positive or negative, and need not be a whole number.
- For example, $X^{1 / 2}=\sqrt{X}$ and $X^{-1}=1 / X$.
- $p=1$ is no transformation: $X^{1}=X$.
- If $p=0$, we use $\log (X)$.
- Following John Tukey, we say that $p>1$ (e.g., $X^{2}, X^{3}$ ) is a transformation "up the ladder of powers" and $p<1$ (e.g., $\left.X^{1 / 2}, \log (X), 1 / X\right)$ is "down the ladder of powers."


## Linear Models in R

The Bulging Rule for Linearizing a Relationship

- This approach works if


## Linear Models in R

The Bulging Rule for Linearizing a Relationship

- This approach works if
(1) The values of the variable to be transformed are all positive.


## Linear Models in R

The Bulging Rule for Linearizing a Relationship

- This approach works if
(1) The values of the variable to be transformed are all positive.
(2) The relationship between the variables is monotone (strictly increasing or decreasing).


## Linear Models in R

The Bulging Rule for Linearizing a Relationship

- This approach works if
(1) The values of the variable to be transformed are all positive.
(2) The relationship between the variables is monotone (strictly increasing or decreasing).
(3) The relationship is simple, in the sense that the direction of curvature doesn't change.


## Linear Models in R

The Bulging Rule for Linearizing a Relationship

- This approach works if
(1) The values of the variable to be transformed are all positive.
(2) The relationship between the variables is monotone (strictly increasing or decreasing).
(3) The relationship is simple, in the sense that the direction of curvature doesn't change.
(9) There are then only four patterns, summarized by Mosteller and Tukey's bulging rule:


## Linear Models in R

- This approach works if
(1) The values of the variable to be transformed are all positive.
(2) The relationship between the variables is monotone (strictly increasing or decreasing).
(3) The relationship is simple, in the sense that the direction of curvature doesn't change.
(9) There are then only four patterns, summarized by Mosteller and Tukey's bulging rule:



## Outline

(1) Linear Models in R
(2) Generalized Linear Models in R

- Review of the Structure of GLMs
- Implementation of GLMs in R: The glm() Function
- GLMs for Binary/Binomial Data
- GLMs for Count Data and Polytomous Data
(3) Mixed-Effects Models in R

4 Using the Tidyverse for Data Management
(5) R Programming

## Generalized Linear Models in R

Review of the Structure of GLMs

- A generalized linear model consists of three components:


## Generalized Linear Models in R

Review of the Structure of GLMs

- A generalized linear model consists of three components:
(1) A random component, specifying the conditional distribution of the response variable, $Y_{i}$, given the predictors. Traditionally, the random component is an exponential family - the normal (Gaussian), binomial, Poisson, gamma, or inverse-Gaussian.


## Generalized Linear Models in R

- A generalized linear model consists of three components:
(1) A random component, specifying the conditional distribution of the response variable, $Y_{i}$, given the predictors. Traditionally, the random component is an exponential family - the normal (Gaussian), binomial, Poisson, gamma, or inverse-Gaussian.
(2) A linear function of the regressors, called the linear predictor,

$$
\eta_{i}=\alpha+\beta_{1} X_{i 1}+\cdots+\beta_{k} X_{i k}
$$

on which the expected value $\mu_{i}$ of $Y_{i}$ depends.

## Generalized Linear Models in R

Review of the Structure of GLMs

- A generalized linear model consists of three components:
(1) A random component, specifying the conditional distribution of the response variable, $Y_{i}$, given the predictors. Traditionally, the random component is an exponential family - the normal (Gaussian), binomial, Poisson, gamma, or inverse-Gaussian.
(2) A linear function of the regressors, called the linear predictor,

$$
\eta_{i}=\alpha+\beta_{1} X_{i 1}+\cdots+\beta_{k} X_{i k}
$$

on which the expected value $\mu_{i}$ of $Y_{i}$ depends.
(3) A link function $g\left(\mu_{i}\right)=\eta_{i}$, which transforms the expectation of the response to the linear predictor. The inverse of the link function is called the mean function: $g^{-1}\left(\eta_{i}\right)=\mu_{i}$.

## Generalized Linear Models in R

- In the following table, the logit, probit and complementary log-log links are for binomial or binary data:

| Link | $\eta_{i}=g\left(\mu_{i}\right)$ | $\mu_{i}=g^{-1}\left(\eta_{i}\right)$ |
| :--- | :---: | :---: |
| identity | $\mu_{i}$ | $\eta_{i}$ |
| log | $\log _{e} \mu_{i}$ | $e^{\eta_{i}}$ |
| inverse | $\mu_{i}^{-1}$ | $\eta_{i}^{-1}$ |
| inverse-square | $\mu_{i}^{-2}$ | $\eta_{i}^{-1 / 2}$ |
| square-root | $\sqrt{\mu_{i}}$ | $\eta_{i}^{2}$ |
| logit | $\log _{e} \frac{\mu_{i}}{1-\mu_{i}}$ | $\frac{1}{1+e^{-\eta_{i}}}$ |
| probit | $\Phi\left(\mu_{i}\right)$ | $\Phi^{-1}\left(\eta_{i}\right)$ |
| complementary $\operatorname{log-log}$ | $\log _{e}\left[-\log _{e}\left(1-\mu_{i}\right)\right]$ | $1-\exp \left[-\exp \left(\eta_{i}\right)\right]$ |

## Generalized Linear Models in R

Implementation of GLMs in R: The glm () Function

- Generalized linear models are fit with the glm() function. Most of the arguments of $g \operatorname{lm}()$ are similar to those of $\operatorname{lm}()$ :


## Generalized Linear Models in R

Implementation of GLMs in R: The glm() Function

- Generalized linear models are fit with the glm() function. Most of the arguments of $g \operatorname{lm}()$ are similar to those of $\operatorname{lm}()$ :
- The response variable and regressors are given in a model formula.
- Generalized linear models are fit with the glm() function. Most of the arguments of glm() are similar to those of $\operatorname{lm}()$ :
- The response variable and regressors are given in a model formula.
- data, subset, and na.action arguments determine the data on which the model is fit.


## Generalized Linear Models in R

- Generalized linear models are fit with the glm() function. Most of the arguments of $g \operatorname{lm}()$ are similar to those of $\operatorname{lm}()$ :
- The response variable and regressors are given in a model formula.
- data, subset, and na.action arguments determine the data on which the model is fit.
- The additional family argument is used to specify a family-generator function, which may take other arguments, such as a link function.


## Generalized Linear Models in R

Implementation of GLMs in R: The glm() Function

- The following table gives family generators and default links:

| Family | Default Link | Range of $Y_{i}$ | $V\left(Y_{i} \mid \eta_{i}\right)$ |
| :--- | :--- | :---: | :---: |
| gaussian | identity | $(-\infty,+\infty)$ | $\phi$ |
| binomial | logit | $\frac{0,1, \ldots, n_{i}}{n_{i}}$ | $\mu_{i}\left(1-\mu_{i}\right)$ |
| poisson | log | $0,1,2, \ldots$ | $\mu_{i}$ |
| Gamma | inverse | $(0, \infty)$ | $\phi \mu_{i}^{2}$ |
| inverse.gaussian | $1 /$ mu^2 $^{2}$ | $(0, \infty)$ | $\phi \mu_{i}^{3}$ |

## Generalized Linear Models in R

Implementation of GLMs in R: The glm() Function

- The following table gives family generators and default links:

| Family | Default Link | Range of $Y_{i}$ | $V\left(Y_{i} \mid \eta_{i}\right)$ |
| :--- | :--- | :---: | :---: |
| gaussian | identity | $(-\infty,+\infty)$ | $\phi$ |
| binomial | logit | $\frac{0,1, \ldots, n_{i}}{n_{i}}$ | $\mu_{i}\left(1-\mu_{i}\right)$ |
| poisson | log | $0,1,2, \ldots$ | $\mu_{i}$ |
| Gamma | inverse | $(0, \infty)$ | $\phi \mu_{i}^{2}$ |
| inverse.gaussian | $1 /$ mu^2 $^{2}$ | $(0, \infty)$ | $\phi \mu_{i}^{3}$ |

- For distributions in the exponential families, the variance is a function of the mean and a dispersion parameter $\phi$ (fixed to 1 for the binomial and Poisson distributions).


## Generalized Linear Models in R

Implementation of GLMs in R: The glm() Function

- The following table shows the links available $(\checkmark)$ for each family in $R$, with the default link marked by $\star$ :

| family | link |  |  |  |  |  |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | identity | inverse | sqrt | $1 / \mathrm{mu}^{\wedge} 2$ | log | logit | probit | cloglog |
| gaussian | $\star$ | $\checkmark$ |  |  | $\checkmark$ |  |  |  |
| binomial |  |  |  |  | $\checkmark$ | $\star$ | $\checkmark$ | $\checkmark$ |
| poisson | $\checkmark$ |  | $\checkmark$ |  | $\star$ |  |  |  |
| Gamma | $\checkmark$ | $\star$ |  |  | $\checkmark$ |  |  |  |
| inverse.gaussian | $\checkmark$ | $\checkmark$ |  | $\star$ | $\checkmark$ |  |  |  |
| quasi | $\star$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| quasibinomial |  |  |  |  |  | $\star$ | $\checkmark$ | $\checkmark$ |
| quasipoisson | $\checkmark$ |  | $\checkmark$ |  | $\star$ |  |  |  |

## Generalized Linear Models in R

Implementation of GLMs in R: The glm () Function

- The following table shows the links available $(\checkmark)$ for each family in $R$, with the default link marked by $\star$ :

| family | link |  |  |  |  |  |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | identity | inverse | sqrt | $1 / \mathrm{mu}^{\wedge} 2$ | log | logit | probit | cloglog |
| gaussian | $\star$ | $\checkmark$ |  |  | $\checkmark$ |  |  |  |
| binomial |  |  |  |  | $\checkmark$ | $\star$ | $\checkmark$ | $\checkmark$ |
| poisson | $\checkmark$ |  | $\checkmark$ |  | $\star$ |  |  |  |
| Gamma | $\checkmark$ | $\star$ |  |  | $\checkmark$ |  |  |  |
| inverse.gaussian | $\checkmark$ | $\checkmark$ |  | $\star$ | $\checkmark$ |  |  |  |
| quasi | $\star$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ | $\checkmark$ |
| quasibinomial |  |  |  |  |  | $\star$ | $\checkmark$ | $\checkmark$ |
| quasipoisson | $\checkmark$ |  | $\checkmark$ |  | $\star$ |  |  |  |

- The quasi, quasibinomial, and quasipoisson family generators do not correspond to exponential families.


## Generalized Linear Models in R

## GLMs for Binary/Binomial

- The response for a binomial GLM may be specified in several forms:


## Generalized Linear Models in R

## GLMs for Binary/Binomial

- The response for a binomial GLM may be specified in several forms:
- For binary data, the response may be


## Generalized Linear Models in R

## GLMs for Binary/Binomial

- The response for a binomial GLM may be specified in several forms:
- For binary data, the response may be
- a variable or an R expression that evaluates to 0s ('failure') and 1s ('success').


## Generalized Linear Models in R

## GLMs for Binary/Binomial

- The response for a binomial GLM may be specified in several forms:
- For binary data, the response may be
- a variable or an R expression that evaluates to 0s ('failure') and 1s ('success').
- a logical variable or expression, such as voted $==$ "yes" (with TRUE representing success, and FALSE failure).


## Generalized Linear Models in R

## GLMs for Binary/Binomial

- The response for a binomial GLM may be specified in several forms:
- For binary data, the response may be
- a variable or an R expression that evaluates to 0s ('failure') and 1s ('success').
- a logical variable or expression, such as voted $==$ "yes" (with TRUE representing success, and FALSE failure).
- a factor (in which case the first category is taken to represent failure and the others success).


## Generalized Linear Models in R

## GLMs for Binary/Binomial

- The response for a binomial GLM may be specified in several forms:
- For binary data, the response may be
- a variable or an R expression that evaluates to 0s ('failure') and 1s ('success').
- a logical variable or expression, such as voted $==$ "yes" (with TRUE representing success, and FALSE failure).
- a factor (in which case the first category is taken to represent failure and the others success).
- For binomial data, the response may be


## Generalized Linear Models in R

## GLMs for Binary/Binomial

- The response for a binomial GLM may be specified in several forms:
- For binary data, the response may be
- a variable or an R expression that evaluates to 0s ('failure') and 1s ('success').
- a logical variable or expression, such as voted $==$ "yes" (with TRUE representing success, and FALSE failure).
- a factor (in which case the first category is taken to represent failure and the others success).
- For binomial data, the response may be
- a two-column matrix, with the first column giving the count of successes and the second the count of failures for each binomial observation.


## Generalized Linear Models in R

## GLMs for Binary/Binomial

- The response for a binomial GLM may be specified in several forms:
- For binary data, the response may be
- a variable or an R expression that evaluates to 0s ('failure') and 1s ('success').
- a logical variable or expression, such as voted $==$ "yes" (with TRUE representing success, and FALSE failure).
- a factor (in which case the first category is taken to represent failure and the others success).
- For binomial data, the response may be
- a two-column matrix, with the first column giving the count of successes and the second the count of failures for each binomial observation.
- a vector giving the proportion of successes, while the binomial denominators (total counts or numbers of trials) are given by the weights argument to glm().


## Generalized Linear Models in R

GLMs for Count Data and Polytomous Data

- Poisson generalized linear models are commonly used when the response variable is a count (Poisson regression) and for modeling associations in contingency tables (loglinear models). The two applications are formally equivalent.


## Generalized Linear Models in R

GLMs for Count Data and Polytomous Data

- Poisson generalized linear models are commonly used when the response variable is a count (Poisson regression) and for modeling associations in contingency tables (loglinear models). The two applications are formally equivalent.
- Poisson GLMs are fit in R using the poisson family generator with glm().


## Generalized Linear Models in R

GLMs for Count Data and Polytomous Data

- Poisson generalized linear models are commonly used when the response variable is a count (Poisson regression) and for modeling associations in contingency tables (loglinear models). The two applications are formally equivalent.
- Poisson GLMs are fit in R using the poisson family generator with glm().
- Overdispersed binomial and Poisson models may be fit via the quasibinomial and quasipoisson families.


## Generalized Linear Models in R

GLMs for Count Data and Polytomous Data

- Poisson generalized linear models are commonly used when the response variable is a count (Poisson regression) and for modeling associations in contingency tables (loglinear models). The two applications are formally equivalent.
- Poisson GLMs are fit in R using the poisson family generator with glm().
- Overdispersed binomial and Poisson models may be fit via the quasibinomial and quasipoisson families.
- The glm.nb() function in the MASS package fits negative-binomial GLMs to count data.


## Generalized Linear Models in R

GLMs for Count Data and Polytomous Data

- Poisson generalized linear models are commonly used when the response variable is a count (Poisson regression) and for modeling associations in contingency tables (loglinear models). The two applications are formally equivalent.
- Poisson GLMs are fit in R using the poisson family generator with glm().
- Overdispersed binomial and Poisson models may be fit via the quasibinomial and quasipoisson families.
- The glm.nb() function in the MASS package fits negative-binomial GLMs to count data.
- The multinom() function in the nnet package fits multinomial GLMs for nominal polytomous responses.


## Generalized Linear Models in R

GLMs for Count Data and Polytomous Data

- Poisson generalized linear models are commonly used when the response variable is a count (Poisson regression) and for modeling associations in contingency tables (loglinear models). The two applications are formally equivalent.
- Poisson GLMs are fit in R using the poisson family generator with glm().
- Overdispersed binomial and Poisson models may be fit via the quasibinomial and quasipoisson families.
- The glm.nb() function in the MASS package fits negative-binomial GLMs to count data.
- The multinom() function in the nnet package fits multinomial GLMs for nominal polytomous responses.
- The polr() function in the MASS package fits the proportional-odds logit model and the ordered probit model to ordinal polytomous responses.


## Generalized Linear Models in R

GLMs for Count Data and Polytomous Data

- Poisson generalized linear models are commonly used when the response variable is a count (Poisson regression) and for modeling associations in contingency tables (loglinear models). The two applications are formally equivalent.
- Poisson GLMs are fit in R using the poisson family generator with glm().
- Overdispersed binomial and Poisson models may be fit via the quasibinomial and quasipoisson families.
- The glm.nb() function in the MASS package fits negative-binomial GLMs to count data.
- The multinom() function in the nnet package fits multinomial GLMs for nominal polytomous responses.
- The polr() function in the MASS package fits the proportional-odds logit model and the ordered probit model to ordinal polytomous responses.
- The clm() function in the ordinal package fits a variety of models (including the proportional-odds model) to ordinal polytomous responses.


## Outline

(1) Linear Models in R
(2) Generalized Linear Models in R
(3) Mixed-Effects Models in R

- The Linear Mixed-Effects Model
- Fitting Mixed Models in R
- A Mixed Model for the Blackmore Exercise Data

4 Using the Tidyverse for Data Management
(5) R Programming

## The Linear Mixed-Effects Model

- The Laird-Ware form of the linear mixed model:

$$
\begin{aligned}
Y_{i j}= & \beta_{1}+\beta_{2} X_{2 i j}+\cdots+\beta_{p} X_{p i j}+B_{1 i} Z_{1 i j}+\cdots+B_{q i} Z_{q i j}+\varepsilon_{i j} \\
B_{k i} \sim & N\left(0, \psi_{k}^{2}\right), \operatorname{Cov}\left(B_{k i}, B_{k^{\prime} i}\right)=\psi_{k k^{\prime}} \\
& B_{k i}, B_{k^{\prime} i^{\prime}} \text { are independent for } i \neq i^{\prime} \\
\varepsilon_{i j} \sim & N\left(0, \sigma^{2} \lambda_{i j j}\right), \operatorname{Cov}\left(\varepsilon_{i j}, \varepsilon_{i j^{\prime}}\right)=\sigma^{2} \lambda_{i j j^{\prime}} \\
& \varepsilon_{i j}, \varepsilon_{i^{\prime} j^{\prime}} \text { are independent for } i \neq i^{\prime}
\end{aligned}
$$

## The Linear Mixed-Effects Model

- where:


## The Linear Mixed-Effects Model

- where:
- $Y_{i j}$ is the value of the response variable for the $j$ th of $n_{i}$ observations in the $i$ th of $m$ groups or clusters.


## The Linear Mixed-Effects Model

- where:
- $Y_{i j}$ is the value of the response variable for the $j$ th of $n_{i}$ observations in the $i$ th of $m$ groups or clusters.
- $\beta_{1}, \beta_{2}, \ldots, \beta_{p}$ are the fixed-effect coefficients, which are identical for all groups.


## The Linear Mixed-Effects Model

- where:
- $Y_{i j}$ is the value of the response variable for the $j$ th of $n_{i}$ observations in the $i$ th of $m$ groups or clusters.
- $\beta_{1}, \beta_{2}, \ldots, \beta_{p}$ are the fixed-effect coefficients, which are identical for all groups.
- $X_{2 i j}, \ldots, X_{p i j}$ are the fixed-effect regressors for observation $j$ in group $i$; there is also implicitly a constant regressor, $X_{1 i j}=1$.


## The Linear Mixed-Effects Model

- where:
- $Y_{i j}$ is the value of the response variable for the $j$ th of $n_{i}$ observations in the $i$ th of $m$ groups or clusters.
- $\beta_{1}, \beta_{2}, \ldots, \beta_{p}$ are the fixed-effect coefficients, which are identical for all groups.
- $X_{2 i j}, \ldots, X_{p i j}$ are the fixed-effect regressors for observation $j$ in group $i$; there is also implicitly a constant regressor, $X_{1 i j}=1$.
- $B_{1 i}, \ldots, B_{q i}$ are the random-effect coefficients for group $i$, assumed to be multivariately normally distributed, independent of the random effects of other groups. The random effects, therefore, vary by group.


## The Linear Mixed-Effects Model

- where:
- $Y_{i j}$ is the value of the response variable for the $j$ th of $n_{i}$ observations in the $i$ th of $m$ groups or clusters.
- $\beta_{1}, \beta_{2}, \ldots, \beta_{p}$ are the fixed-effect coefficients, which are identical for all groups.
- $X_{2 i j}, \ldots, X_{p i j}$ are the fixed-effect regressors for observation $j$ in group $i$; there is also implicitly a constant regressor, $X_{1 i j}=1$.
- $B_{1 i}, \ldots, B_{q i}$ are the random-effect coefficients for group $i$, assumed to be multivariately normally distributed, independent of the random effects of other groups. The random effects, therefore, vary by group.
- The $B_{i k}$ are thought of as random variables, not as parameters, and are similar in this respect to the errors $\varepsilon_{i j}$.


## The Linear Mixed-Effects Model

- where:
- $Y_{i j}$ is the value of the response variable for the $j$ th of $n_{i}$ observations in the $i$ th of $m$ groups or clusters.
- $\beta_{1}, \beta_{2}, \ldots, \beta_{p}$ are the fixed-effect coefficients, which are identical for all groups.
- $X_{2 i j}, \ldots, X_{p i j}$ are the fixed-effect regressors for observation $j$ in group $i$; there is also implicitly a constant regressor, $X_{1 i j}=1$.
- $B_{1 i}, \ldots, B_{q i}$ are the random-effect coefficients for group $i$, assumed to be multivariately normally distributed, independent of the random effects of other groups. The random effects, therefore, vary by group.
- The $B_{i k}$ are thought of as random variables, not as parameters, and are similar in this respect to the errors $\varepsilon_{i j}$.
- $Z_{1 i j}, \ldots, Z_{q i j}$ are the random-effect regressors.


## The Linear Mixed-Effects Model

- where:
- $Y_{i j}$ is the value of the response variable for the $j$ th of $n_{i}$ observations in the $i$ th of $m$ groups or clusters.
- $\beta_{1}, \beta_{2}, \ldots, \beta_{p}$ are the fixed-effect coefficients, which are identical for all groups.
- $X_{2 i j}, \ldots, X_{p i j}$ are the fixed-effect regressors for observation $j$ in group $i$; there is also implicitly a constant regressor, $X_{1 i j}=1$.
- $B_{1 i}, \ldots, B_{q i}$ are the random-effect coefficients for group $i$, assumed to be multivariately normally distributed, independent of the random effects of other groups. The random effects, therefore, vary by group.
- The $B_{i k}$ are thought of as random variables, not as parameters, and are similar in this respect to the errors $\varepsilon_{i j}$.
- $Z_{1 i j}, \ldots, Z_{q i j}$ are the random-effect regressors.
- The $Z \mathrm{~s}$ are almost always a subset of the $X \mathrm{~s}$ (and may include all of the $X \mathrm{~s}$ ).


## The Linear Mixed-Effects Model

- where:
- $Y_{i j}$ is the value of the response variable for the $j$ th of $n_{i}$ observations in the $i$ th of $m$ groups or clusters.
- $\beta_{1}, \beta_{2}, \ldots, \beta_{p}$ are the fixed-effect coefficients, which are identical for all groups.
- $X_{2 i j}, \ldots, X_{p i j}$ are the fixed-effect regressors for observation $j$ in group $i$; there is also implicitly a constant regressor, $X_{1 i j}=1$.
- $B_{1 i}, \ldots, B_{q i}$ are the random-effect coefficients for group $i$, assumed to be multivariately normally distributed, independent of the random effects of other groups. The random effects, therefore, vary by group.
- The $B_{i k}$ are thought of as random variables, not as parameters, and are similar in this respect to the errors $\varepsilon_{i j}$.
- $Z_{1 i j}, \ldots, Z_{q i j}$ are the random-effect regressors.
- The $Z \mathrm{~s}$ are almost always a subset of the $X \mathrm{~s}$ (and may include all of the $X \mathrm{~s}$ ).
- When there is a random intercept term, $Z_{1 i j}=1$.


## The Linear Mixed-Effects Model

- The remaining parameters specify the variance-covariance components (don't get lost!):


## The Linear Mixed-Effects Model

- The remaining parameters specify the variance-covariance components (don't get lost!):
- $\psi_{k}^{2}$ are the variances and $\psi_{k k^{\prime}}$ the covariances among the random effects, assumed to be constant across groups.


## The Linear Mixed-Effects Model

- The remaining parameters specify the variance-covariance components (don't get lost!):
- $\psi_{k}^{2}$ are the variances and $\psi_{k k^{\prime}}$ the covariances among the random effects, assumed to be constant across groups.
- In some applications, the $\psi$ s are parametrized in terms of a smaller number of fundamental parameters.


## The Linear Mixed-Effects Model

- The remaining parameters specify the variance-covariance components (don't get lost!):
- $\psi_{k}^{2}$ are the variances and $\psi_{k k^{\prime}}$ the covariances among the random effects, assumed to be constant across groups.
- In some applications, the $\psi$ s are parametrized in terms of a smaller number of fundamental parameters.
- $\varepsilon_{i j}$ is the error for observation $j$ in group $i$.


## The Linear Mixed-Effects Model

- The remaining parameters specify the variance-covariance components (don't get lost!):
- $\psi_{k}^{2}$ are the variances and $\psi_{k k^{\prime}}$ the covariances among the random effects, assumed to be constant across groups.
- In some applications, the $\psi$ s are parametrized in terms of a smaller number of fundamental parameters.
- $\varepsilon_{i j}$ is the error for observation $j$ in group $i$.
- The errors for group $i$ are assumed to be multivariately normally distributed, and independent of errors in other groups.


## The Linear Mixed-Effects Model

- The remaining parameters specify the variance-covariance components (don't get lost!):
- $\psi_{k}^{2}$ are the variances and $\psi_{k k^{\prime}}$ the covariances among the random effects, assumed to be constant across groups.
- In some applications, the $\psi$ s are parametrized in terms of a smaller number of fundamental parameters.
- $\varepsilon_{i j}$ is the error for observation $j$ in group $i$.
- The errors for group $i$ are assumed to be multivariately normally distributed, and independent of errors in other groups.
- $\sigma^{2} \lambda_{i j j^{\prime}}$ are the covariances between errors in group $i$.


## The Linear Mixed-Effects Model

- The remaining parameters specify the variance-covariance components (don't get lost!):
- $\psi_{k}^{2}$ are the variances and $\psi_{k k^{\prime}}$ the covariances among the random effects, assumed to be constant across groups.
- In some applications, the $\psi$ s are parametrized in terms of a smaller number of fundamental parameters.
- $\varepsilon_{i j}$ is the error for observation $j$ in group $i$.
- The errors for group $i$ are assumed to be multivariately normally distributed, and independent of errors in other groups.
- $\sigma^{2} \lambda_{i j j^{\prime}}$ are the covariances between errors in group $i$.
- Generally, the $\lambda_{i j j^{\prime}}$ are parametrized in terms of a few basic parameters, and their specific form depends upon context.


## The Linear Mixed-Effects Model

- The remaining parameters specify the variance-covariance components (don't get lost!):
- $\psi_{k}^{2}$ are the variances and $\psi_{k k^{\prime}}$ the covariances among the random effects, assumed to be constant across groups.
- In some applications, the $\psi$ s are parametrized in terms of a smaller number of fundamental parameters.
- $\varepsilon_{i j}$ is the error for observation $j$ in group $i$.
- The errors for group $i$ are assumed to be multivariately normally distributed, and independent of errors in other groups.
- $\sigma^{2} \lambda_{i j j^{\prime}}$ are the covariances between errors in group $i$.
- Generally, the $\lambda_{i j j^{\prime}}$ are parametrized in terms of a few basic parameters, and their specific form depends upon context.
- When observations are sampled independently within groups and are assumed to have constant error variance (as is typical in hierarchical models), $\lambda_{i j}=1, \lambda_{i j j^{\prime}}=0$ (for $j \neq j^{\prime}$ ), and thus the only free parameter to estimate is the common error variance, $\sigma^{2}$.


## The Linear Mixed-Effects Model

- The remaining parameters specify the variance-covariance components (don't get lost!):
- $\psi_{k}^{2}$ are the variances and $\psi_{k k^{\prime}}$ the covariances among the random effects, assumed to be constant across groups.
- In some applications, the $\psi \mathrm{s}$ are parametrized in terms of a smaller number of fundamental parameters.
- $\varepsilon_{i j}$ is the error for observation $j$ in group $i$.
- The errors for group $i$ are assumed to be multivariately normally distributed, and independent of errors in other groups.
- $\sigma^{2} \lambda_{i j j^{\prime}}$ are the covariances between errors in group $i$.
- Generally, the $\lambda_{i j j^{\prime}}$ are parametrized in terms of a few basic parameters, and their specific form depends upon context.
- When observations are sampled independently within groups and are assumed to have constant error variance (as is typical in hierarchical models), $\lambda_{i j}=1, \lambda_{i j j^{\prime}}=0$ (for $j \neq j^{\prime}$ ), and thus the only free parameter to estimate is the common error variance, $\sigma^{2}$.
- If the observations in a "group" represent longitudinal data on a single individual, then the structure of the $\lambda$ s may be specified to capture serial (i.e., over-time) dependencies among the errors.


## Fitting Mixed Models in R

with the nlme and Ime4 packages

- In the nlme package (Pinheiro, Bates, DebRoy, and Sarkar):


## Fitting Mixed Models in R

with the nIme and Ime4 packages

- In the nlme package (Pinheiro, Bates, DebRoy, and Sarkar):
- lme(): linear mixed-effects models with nested random effects; can model serially correlated errors.


## Fitting Mixed Models in R

with the nlme and Ime4 packages

- In the nlme package (Pinheiro, Bates, DebRoy, and Sarkar):
- lme(): linear mixed-effects models with nested random effects; can model serially correlated errors.
- nlme(): nonlinear mixed-effects models.


## Fitting Mixed Models in R

with the nlme and Ime4 packages

- In the nlme package (Pinheiro, Bates, DebRoy, and Sarkar):
- lme(): linear mixed-effects models with nested random effects; can model serially correlated errors.
- nlme(): nonlinear mixed-effects models.
- In the Ime4 package (Bates, Maechler, Bolker, and Walker):


## Fitting Mixed Models in R

with the nlme and Ime4 packages

- In the nlme package (Pinheiro, Bates, DebRoy, and Sarkar):
- lme(): linear mixed-effects models with nested random effects; can model serially correlated errors.
- nlme(): nonlinear mixed-effects models.
- In the Ime4 package (Bates, Maechler, Bolker, and Walker):
- lmer(): linear mixed-effects models with nested or crossed random effects; no facility (yet) for serially correlated errors.


## Fitting Mixed Models in R

with the nlme and Ime4 packages

- In the nlme package (Pinheiro, Bates, DebRoy, and Sarkar):
- lme(): linear mixed-effects models with nested random effects; can model serially correlated errors.
- nlme(): nonlinear mixed-effects models.
- In the Ime4 package (Bates, Maechler, Bolker, and Walker):
- lmer(): linear mixed-effects models with nested or crossed random effects; no facility (yet) for serially correlated errors.
- glmer(): generalized-linear mixed-effects models.


## Fitting Mixed Models in R

- In the nlme package (Pinheiro, Bates, DebRoy, and Sarkar):
- lme(): linear mixed-effects models with nested random effects; can model serially correlated errors.
- nlme(): nonlinear mixed-effects models.
- In the Ime4 package (Bates, Maechler, Bolker, and Walker):
- lmer(): linear mixed-effects models with nested or crossed random effects; no facility (yet) for serially correlated errors.
- glmer(): generalized-linear mixed-effects models.
- There are many other CRAN packages that fit a variety of mixed-effects models, perhaps most notably glmmTMB
(see https://bbolker.github.io/mixedmodels-misc/glmmFAQ.html).


## Fitting Mixed Models in R

- In the nlme package (Pinheiro, Bates, DebRoy, and Sarkar):
- lme(): linear mixed-effects models with nested random effects; can model serially correlated errors.
- nlme(): nonlinear mixed-effects models.
- In the Ime4 package (Bates, Maechler, Bolker, and Walker):
- lmer(): linear mixed-effects models with nested or crossed random effects; no facility (yet) for serially correlated errors.
- glmer(): generalized-linear mixed-effects models.
- There are many other CRAN packages that fit a variety of mixed-effects models, perhaps most notably glmmTMB (see https://bbolker.github.io/mixedmodels-misc/glmmFAQ.html).
- There are also Bayesian approaches to modeling hierarchical and longitudinal data that offer certain advantages; see in particular the rstan, rstanarm, and blme packages.


## A Mixed Model for the Blackmore Exercise Data

- A level-1 model specifying a linear "growth curve" for log exercise for each subject:

$$
\log -\text { exercise }_{i j}=\alpha_{0 i}+\alpha_{1 i}\left(\text { age }_{i j}-8\right)+\varepsilon_{i j}
$$

## A Mixed Model for the Blackmore Exercise Data

- A level-1 model specifying a linear "growth curve" for log exercise for each subject:

$$
\log _{- \text {exercise }_{i j}}=\alpha_{0 i}+\alpha_{1 i}\left(\text { age }_{i j}-8\right)+\varepsilon_{i j}
$$

- Our interest in detecting differences in exercise histories between subjects and controls suggests the level-2 model

$$
\begin{aligned}
& \alpha_{0 i}=\gamma_{00}+\gamma_{01} \text { group }_{i}+\omega_{0 i} \\
& \alpha_{1 i}=\gamma_{10}+\gamma_{11} \text { group }_{i}+\omega_{1 i}
\end{aligned}
$$

where group is a dummy variable coded 1 for subjects and 0 for controls.

## A Mixed Model for the Blackmore Exercise Data

## Laird-Ware form of the Model

- Substituting the level- 2 model into the level- 1 model produces

$$
\left.\left.\begin{array}{rl}
\log _{\text {-exercise }}^{i j}
\end{array}=\left(\gamma_{00}+\gamma_{01} \text { group }_{i}+\omega_{0 i}\right)+\left(\gamma_{10}+\gamma_{11} \text { group }_{i}+\omega_{1 i}\right)\left(\text { age }_{i j}-8\right)+\varepsilon_{i j}\right)=\gamma_{00}+\gamma_{01} \text { group }_{i}+\gamma_{10}\left(\text { age }_{i j}-8\right)+\gamma_{11} \text { group }_{i} \times\left(\text { age }_{i j}-8\right)\right)
$$

## A Mixed Model for the Blackmore Exercise Data

## Laird-Ware form of the Model

- Substituting the level- 2 model into the level- 1 model produces

$$
\left.\left.\begin{array}{rl}
\log _{\text {-exercise }}^{i j}
\end{array}=\left(\gamma_{00}+\gamma_{01} \text { group }_{i}+\omega_{0 i}\right)+\left(\gamma_{10}+\gamma_{11} \text { group }_{i}+\omega_{1 i}\right)\left(\text { age }_{i j}-8\right)+\varepsilon_{i j}\right)=\gamma_{00}+\gamma_{01} \text { group }_{i}+\gamma_{10}\left(\text { age }_{i j}-8\right)+\gamma_{11} \text { group }_{i} \times\left(\text { age }_{i j}-8\right)\right)
$$

- in Laird-Ware form,

$$
Y_{i j}=\beta_{1}+\beta_{2} X_{2 i j}+\beta_{3} X_{3 i j}+\beta_{4} X_{4 i j}+\delta_{1 i}+\delta_{2 i} Z_{2 i j}+\varepsilon_{i j}
$$

## A Mixed Model for the Blackmore Exercise Data

Laird-Ware form of the Model

- Substituting the level-2 model into the level-1 model produces

$$
\left.\left.\begin{array}{rl}
\log _{\text {-exercise }}^{i j}
\end{array}=\left(\gamma_{00}+\gamma_{01} \text { group }_{i}+\omega_{0 i}\right)+\left(\gamma_{10}+\gamma_{11} \text { group }_{i}+\omega_{1 i}\right)\left(\text { age }_{i j}-8\right)+\varepsilon_{i j}\right)=\gamma_{00}+\gamma_{01} \text { group }_{i}+\gamma_{10}\left(\text { age }_{i j}-8\right)+\gamma_{11} \text { group }_{i} \times\left(\text { age }_{i j}-8\right)\right)
$$

- in Laird-Ware form,

$$
Y_{i j}=\beta_{1}+\beta_{2} X_{2 i j}+\beta_{3} X_{3 i j}+\beta_{4} X_{4 i j}+\delta_{1 i}+\delta_{2 i} Z_{2 i j}+\varepsilon_{i j}
$$

- Continuous first-order autoregressive process for the errors:

$$
\operatorname{Cor}\left(\varepsilon_{i t}, \varepsilon_{i, t+s}\right)=\rho(s)=\phi^{|s|}
$$

where the time-interval between observations, $s$, need not be an integer.

## A Mixed Model for the Blackmore Exercise Data

 Specifying the Model in 1 me() and $\operatorname{lmer}()$- Using lme() in the nlme package:

```
lme(log.exercise ~ I(age - 8)*group,
random = ~ I(age - 8) | subject,
correlation = corCAR1(form = ~ age |subject)
data=Blackmoore)
```


## A Mixed Model for the Blackmore Exercise Data

- Using lme() in the nlme package:

$$
\begin{aligned}
& \text { lme(log.exercise } \sim I(\text { age }-8) * \text { group, } \\
& \quad \text { random }=I(\text { age }-8) \mid \text { subject }, \\
& \quad \text { correlation }=\text { corCAR1 (form }=\sim \text { age |subject }) \\
& \text { data=Blackmoore })
\end{aligned}
$$

- Using lmer () in the Ime4 package, but without autocorrelated errors:

$$
\begin{gathered}
\operatorname{lmer}(\text { log.exercise } \sim I(\text { age }-8) * g r o u p+(I(\text { age }-8) \mid \text { subject }), \\
\text { data=Blackmoore })
\end{gathered}
$$

## Outline

(1) Linear Models in R
(2) Generalized Linear Models in R
(3) Mixed-Effects Models in R

4 Using the Tidyverse for Data Management

- Overview of the Tidyverse
- Core Tidyverse Packages
- Other Tidyverse Packages
- Should You Commit to the Tidyverse?
(5) R Programming


## Using the Tidyverse for Data Management

Overview of the Tidyverse

- The "Tidyverse" is an integrated set of R packages developed by Hadley Wickham and his collaborators at RStudio (see https://www.tidyverse.org/).


## Using the Tidyverse for Data Management

Overview of the Tidyverse

- The "Tidyverse" is an integrated set of R packages developed by Hadley Wickham and his collaborators at RStudio (see https://www.tidyverse.org/).
- The packages are meant to provide a straightforward way to import data into R and to manipulate the data.


## Using the Tidyverse for Data Management

- The "Tidyverse" is an integrated set of R packages developed by Hadley Wickham and his collaborators at RStudio (see https://www.tidyverse.org/).
- The packages are meant to provide a straightforward way to import data into R and to manipulate the data.
- There are also Tidyverse tools for R programming and statistical graphics.


## Using the Tidyverse for Data Management

- The "Tidyverse" is an integrated set of R packages developed by Hadley Wickham and his collaborators at RStudio (see https://www.tidyverse.org/).
- The packages are meant to provide a straightforward way to import data into R and to manipulate the data.
- There are also Tidyverse tools for R programming and statistical graphics.
- A central goal of the data-oriented Tidyverse packages is to construct, modify, and maintain "tidy data" -rectangular data sets in which the rows represent cases and the columns represent variables.


## Using the Tidyverse for Data Management

- The "Tidyverse" is an integrated set of R packages developed by Hadley Wickham and his collaborators at RStudio (see https://www.tidyverse.org/).
- The packages are meant to provide a straightforward way to import data into R and to manipulate the data.
- There are also Tidyverse tools for R programming and statistical graphics.
- A central goal of the data-oriented Tidyverse packages is to construct, modify, and maintain "tidy data" -rectangular data sets in which the rows represent cases and the columns represent variables.
- Of course, the idea of a rectangular data set greatly antedates the Tidyverse and is incorporated in the standard R data frame.


## Using the Tidyverse for Data Management

## Core Tidyverse Packages

- There are eight "core" Tidyverse packages, which can be installed and loaded via the master tidyverse package:


## Using the Tidyverse for Data Management

Core Tidyverse Packages

- There are eight "core" Tidyverse packages, which can be installed and loaded via the master tidyverse package:
(1) readr: Imports rectangular data sets from plain-text files.


## Using the Tidyverse for Data Management

## Core Tidyverse Packages

- There are eight "core" Tidyverse packages, which can be installed and loaded via the master tidyverse package:
(1) readr: Imports rectangular data sets from plain-text files.
(2) tibble: The specific implementation of rectangular data sets in the Tidyverse is called a "tibble," and tibble objects inherit from the "data.frame" class.


## Using the Tidyverse for Data Management

Core Tidyverse Packages

- There are eight "core" Tidyverse packages, which can be installed and loaded via the master tidyverse package:
(1) readr: Imports rectangular data sets from plain-text files.
(2) tibble: The specific implementation of rectangular data sets in the Tidyverse is called a "tibble," and tibble objects inherit from the "data.frame" class.
(3) tidyr: Provides functions to create and maintain rectangular data sets (e.g., to transform rectangular data sets between "wide" and "long" form).


## Using the Tidyverse for Data Management

Core Tidyverse Packages

- There are eight "core" Tidyverse packages, which can be installed and loaded via the master tidyverse package:
(1) readr: Imports rectangular data sets from plain-text files.
(2) tibble: The specific implementation of rectangular data sets in the Tidyverse is called a "tibble," and tibble objects inherit from the "data.frame" class.
(3) tidyr: Provides functions to create and maintain rectangular data sets (e.g., to transform rectangular data sets between "wide" and "long" form).
(9) dplyr: Provides functions for data manipulation (e.g., adding variables to an existing data set).


## Using the Tidyverse for Data Management

Core Tidyverse Packages

- There are eight "core" Tidyverse packages, which can be installed and loaded via the master tidyverse package:
(1) readr: Imports rectangular data sets from plain-text files.
(2) tibble: The specific implementation of rectangular data sets in the Tidyverse is called a "tibble," and tibble objects inherit from the "data.frame" class.
(3) tidyr: Provides functions to create and maintain rectangular data sets (e.g., to transform rectangular data sets between "wide" and "long" form).
(1) dplyr: Provides functions for data manipulation (e.g., adding variables to an existing data set).
(3) stringr: Provides functions for manipulating text (character-string) data (e.g., searching for text).


## Using the Tidyverse for Data Management

Core Tidyverse Packages

- There are eight "core" Tidyverse packages, which can be installed and loaded via the master tidyverse package:
(1) readr: Imports rectangular data sets from plain-text files.
(2) tibble: The specific implementation of rectangular data sets in the Tidyverse is called a "tibble," and tibble objects inherit from the "data.frame" class.
(3) tidyr: Provides functions to create and maintain rectangular data sets (e.g., to transform rectangular data sets between "wide" and "long" form).
(1) dplyr: Provides functions for data manipulation (e.g., adding variables to an existing data set).
(3) stringr: Provides functions for manipulating text (character-string) data (e.g., searching for text).
(0) forcats: Provides functions for manipulating R factors (e.g., changing the order of levels of a factor).


## Using the Tidyverse for Data Management

- There are eight "core" Tidyverse packages, which can be installed and loaded via the master tidyverse package:
(1) readr: Imports rectangular data sets from plain-text files.
(2) tibble: The specific implementation of rectangular data sets in the Tidyverse is called a "tibble," and tibble objects inherit from the "data.frame" class.
(3) tidyr: Provides functions to create and maintain rectangular data sets (e.g., to transform rectangular data sets between "wide" and "long" form).
(1) dplyr: Provides functions for data manipulation (e.g., adding variables to an existing data set).
(3) stringr: Provides functions for manipulating text (character-string) data (e.g., searching for text).
(0) forcats: Provides functions for manipulating R factors (e.g., changing the order of levels of a factor).
(1) purrr: Provides R programming tools (e.g., alternatives to iteration).


## Using the Tidyverse for Data Management

- There are eight "core" Tidyverse packages, which can be installed and loaded via the master tidyverse package:
(1) readr: Imports rectangular data sets from plain-text files.
(2) tibble: The specific implementation of rectangular data sets in the Tidyverse is called a "tibble," and tibble objects inherit from the "data.frame" class.
(3) tidyr: Provides functions to create and maintain rectangular data sets (e.g., to transform rectangular data sets between "wide" and "long" form).
(1) dplyr: Provides functions for data manipulation (e.g., adding variables to an existing data set).
(3) stringr: Provides functions for manipulating text (character-string) data (e.g., searching for text).
(0) forcats: Provides functions for manipulating R factors (e.g., changing the order of levels of a factor).
(3) purrr: Provides R programming tools (e.g., alternatives to iteration).
(8) ggplot2: A comprehensive alternative graphics system for R (to be discussed when we take up $R$ graphics, and a package that is slightly out-of-place in the Tidyverse).


## Using the Tidyverse for Data Management

## Other Tidyverse Packages

- There are other Tidyverse packages, which can be installed and loaded separately, most notably:


## Using the Tidyverse for Data Management

## Other Tidyverse Packages

- There are other Tidyverse packages, which can be installed and loaded separately, most notably:
- haven: Imports data from other statistical packages.


## Using the Tidyverse for Data Management

## Other Tidyverse Packages

- There are other Tidyverse packages, which can be installed and loaded separately, most notably:
- haven: Imports data from other statistical packages.
- readxI: Imports data from Excel files.


## Using the Tidyverse for Data Management

## Other Tidyverse Packages

- There are other Tidyverse packages, which can be installed and loaded separately, most notably:
- haven: Imports data from other statistical packages.
- readxl: Imports data from Excel files.
- lubridate: For working with dates.


## Using the Tidyverse for Data Management

Other Tidyverse Packages

- There are other Tidyverse packages, which can be installed and loaded separately, most notably:
- haven: Imports data from other statistical packages.
- readxl: Imports data from Excel files.
- lubridate: For working with dates.
- magrittr: The style of data manipulation encouraged by the developers of the Tidyverse makes extensive use of the "pipe" operator, $\%>\%$, which is provided by the magritr package.


## Using the Tidyverse for Data Management

Other Tidyverse Packages

- There are other Tidyverse packages, which can be installed and loaded separately, most notably:
- haven: Imports data from other statistical packages.
- readxl: Imports data from Excel files.
- lubridate: For working with dates.
- magrittr: The style of data manipulation encouraged by the developers of the Tidyverse makes extensive use of the "pipe" operator, $\%>\%$, which is provided by the magritr package.
- magrittr also includes some other programming-oriented functions.


## Using the Tidyverse for Data Management

Other Tidyverse Packages

- There are other Tidyverse packages, which can be installed and loaded separately, most notably:
- haven: Imports data from other statistical packages.
- readxl: Imports data from Excel files.
- lubridate: For working with dates.
- magrittr: The style of data manipulation encouraged by the developers of the Tidyverse makes extensive use of the "pipe" operator, $\%>\%$, which is provided by the magritr package.
- magrittr also includes some other programming-oriented functions.
- The pipe operator is supplied by several of the core Tidyverse packages.


## Using the Tidyverse for Data Management

- There are other Tidyverse packages, which can be installed and loaded separately, most notably:
- haven: Imports data from other statistical packages.
- readxl: Imports data from Excel files.
- lubridate: For working with dates.
- magrittr: The style of data manipulation encouraged by the developers of the Tidyverse makes extensive use of the "pipe" operator, $\%>\%$, which is provided by the magritr package.
- magrittr also includes some other programming-oriented functions.
- The pipe operator is supplied by several of the core Tidyverse packages.
- Pipes can be used with standard R functions.


## Using the Tidyverse for Data Management

- There are few, if any, Tidyverse functions that don't have close analogs in the standard R distribution, but the Tidyverse functions are more uniform and many people claim that they are easier to use (possibly because they're unfamiliar with standard R).


## Using the Tidyverse for Data Management

- There are few, if any, Tidyverse functions that don't have close analogs in the standard $R$ distribution, but the Tidyverse functions are more uniform and many people claim that they are easier to use (possibly because they're unfamiliar with standard R).
- There are hundreds of functions in the core Tidyverse packages. It isn't obvious that it's easier to learn the Tidyverse than to learn standard R.


## Using the Tidyverse for Data Management

- There are few, if any, Tidyverse functions that don't have close analogs in the standard R distribution, but the Tidyverse functions are more uniform and many people claim that they are easier to use (possibly because they're unfamiliar with standard R).
- There are hundreds of functions in the core Tidyverse packages. It isn't obvious that it's easier to learn the Tidyverse than to learn standard R .
- There are both advantages and disadvantages to Tidyverse implementations of ideas.


## Using the Tidyverse for Data Management

- There are few, if any, Tidyverse functions that don't have close analogs in the standard $R$ distribution, but the Tidyverse functions are more uniform and many people claim that they are easier to use (possibly because they're unfamiliar with standard R).
- There are hundreds of functions in the core Tidyverse packages. It isn't obvious that it's easier to learn the Tidyverse than to learn standard R.
- There are both advantages and disadvantages to Tidyverse implementations of ideas.
- For example, the print () method for tibbles is nicer than that for data frames (cf., the brief() function in the car package), but tibbles don't support row names.


## Using the Tidyverse for Data Management

- There are few, if any, Tidyverse functions that don't have close analogs in the standard $R$ distribution, but the Tidyverse functions are more uniform and many people claim that they are easier to use (possibly because they're unfamiliar with standard R).
- There are hundreds of functions in the core Tidyverse packages. It isn't obvious that it's easier to learn the Tidyverse than to learn standard R.
- There are both advantages and disadvantages to Tidyverse implementations of ideas.
- For example, the print () method for tibbles is nicer than that for data frames (cf., the brief () function in the car package), but tibbles don't support row names.
- Tidyverse tools often don't play well with non-Tidyverse tools.


## Using the Tidyverse for Data Management

- There are few, if any, Tidyverse functions that don't have close analogs in the standard $R$ distribution, but the Tidyverse functions are more uniform and many people claim that they are easier to use (possibly because they're unfamiliar with standard R).
- There are hundreds of functions in the core Tidyverse packages. It isn't obvious that it's easier to learn the Tidyverse than to learn standard R.
- There are both advantages and disadvantages to Tidyverse implementations of ideas.
- For example, the print () method for tibbles is nicer than that for data frames (cf., the brief () function in the car package), but tibbles don't support row names.
- Tidyverse tools often don't play well with non-Tidyverse tools.
- For example, the data.table package implements a data frame alternative that is superior to tibbles for large data sets, but data.tables aren't well supported by Tidyverse functions.


## Using the Tidyverse for Data Management

- R is a programming language, and in many cases the simplest and most direct solution to a problem is to write a program.


## Using the Tidyverse for Data Management

Should You Commit to the Tidyverse?

- R is a programming language, and in many cases the simplest and most direct solution to a problem is to write a program.
- Using the Tidyverse tools effectively requires some programming skills, and a beginner's time might be better spent learning more general basic R programming.


## Using the Tidyverse for Data Management

- R is a programming language, and in many cases the simplest and most direct solution to a problem is to write a program.
- Using the Tidyverse tools effectively requires some programming skills, and a beginner's time might be better spent learning more general basic R programming.
- For an interesting general critique of the Tidyverse (with which I don't entirely agree), see an essay by Norm Matloff at https://github.com/matloff/TidyverseSkeptic.


## Outline

## (1) Linear Models in R

(2) Generalized Linear Models in R
(3) Mixed-Effects Models in R

4 Using the Tidyverse for Data Management
(5) R Programming

- MLE Estimation of the Binary Logit Models by Newton-Raphson
- Object-Oriented Programming


## R Programming

- The binary logit model is

$$
\operatorname{Pr}\left(Y_{i}=1\right)=\phi_{i}=\frac{1}{1+\exp \left(-x_{i}^{\top} \boldsymbol{\beta}\right)}
$$

where

## R Programming

- The binary logit model is

$$
\operatorname{Pr}\left(Y_{i}=1\right)=\phi_{i}=\frac{1}{1+\exp \left(-x_{i}^{\top} \boldsymbol{\beta}\right)}
$$

where

- X is the model matrix, with $\mathrm{x}_{i}^{T}$ as its $i$ th row;


## R Programming

- The binary logit model is

$$
\operatorname{Pr}\left(Y_{i}=1\right)=\phi_{i}=\frac{1}{1+\exp \left(-x_{i}^{T} \beta\right)}
$$

where

- X is the model matrix, with $\mathrm{x}_{i}^{T}$ as its $i$ th row;
- y is the response vector (containing 0 s and 1 s ) with $Y_{i}$ as its $i$ th element;


## R Programming

- The binary logit model is

$$
\operatorname{Pr}\left(Y_{i}=1\right)=\phi_{i}=\frac{1}{1+\exp \left(-x_{i}^{T} \beta\right)}
$$

where

- X is the model matrix, with $\mathrm{x}_{i}^{T}$ as its $i$ th row;
- y is the response vector (containing 0 s and 1 s ) with $Y_{i}$ as its $i$ th element;
- $\beta$ is the vector of logistic-regression parameters.


## R Programming

## MLE Estimation of the Binary Logit Models by Newton-Raphson

- The log-likelihood for the model is

$$
\log _{e} L(\boldsymbol{\beta})=\sum y_{i} \log _{e} \phi_{i}+\left(1-y_{i}\right) \log _{e}\left(1-\phi_{i}\right)
$$

## R Programming

## MLE Estimation of the Binary Logit Models by Newton-Raphson

- The log-likelihood for the model is

$$
\log _{e} L(\boldsymbol{\beta})=\sum y_{i} \log _{e} \phi_{i}+\left(1-y_{i}\right) \log _{e}\left(1-\phi_{i}\right)
$$

- The gradient (the vector of partial derivatives) of the log-likelihood with respect to the parameters is

$$
\frac{\partial \log _{e} L}{\partial \beta}=\sum\left(y_{i}-\phi_{i}\right) x_{i}
$$

## R Programming

- The log-likelihood for the model is

$$
\log _{e} L(\boldsymbol{\beta})=\sum y_{i} \log _{e} \phi_{i}+\left(1-y_{i}\right) \log _{e}\left(1-\phi_{i}\right)
$$

- The gradient (the vector of partial derivatives) of the log-likelihood with respect to the parameters is

$$
\frac{\partial \log _{e} L}{\partial \beta}=\sum\left(y_{i}-\phi_{i}\right) x_{i}
$$

- The Hessian (the matrix of second-order partial derivatives) of the log-likelihood is

$$
\frac{\partial \log _{e} L}{\partial \beta \partial \beta^{T}}=X^{T} \mathrm{VX}
$$

where $\mathrm{V}=\operatorname{diag}\left\{\phi_{i}\left(1-\phi_{i}\right)\right\}$. The variance-covariance matrix of the estimated regression coefficients is the inverse of the Hessian.

## R Programming

- The log-likelihood for the model is

$$
\log _{e} L(\boldsymbol{\beta})=\sum y_{i} \log _{e} \phi_{i}+\left(1-y_{i}\right) \log _{e}\left(1-\phi_{i}\right)
$$

- The gradient (the vector of partial derivatives) of the log-likelihood with respect to the parameters is

$$
\frac{\partial \log _{e} L}{\partial \beta}=\sum\left(y_{i}-\phi_{i}\right) x_{i}
$$

- The Hessian (the matrix of second-order partial derivatives) of the log-likelihood is

$$
\frac{\partial \log _{e} L}{\partial \beta \partial \beta^{T}}=X^{T} \mathrm{VX}
$$

where $\mathrm{V}=\operatorname{diag}\left\{\phi_{i}\left(1-\phi_{i}\right)\right\}$. The variance-covariance matrix of the estimated regression coefficients is the inverse of the Hessian.

- Setting the gradient to 0 produces nonlinear estimating equations for $\beta$, which have to be solved iteratively, possibly using the information in the Hessian.


## R Programming

- Newton-Raphson is a general method for solving nonlinear equations iteratively.


## R Programming

- Newton-Raphson is a general method for solving nonlinear equations iteratively. - Here:


## R Programming

- Newton-Raphson is a general method for solving nonlinear equations iteratively. - Here:
(1) Choose initial estimates of the regression coefficients, such as $b_{0}=0$.


## R Programming

- Newton-Raphson is a general method for solving nonlinear equations iteratively. - Here:
(1) Choose initial estimates of the regression coefficients, such as $b_{0}=0$.
(2) At each iteration $t$, update the coefficients:

$$
\mathrm{b}_{t}=\mathrm{b}_{t-1}+\left(\mathrm{X}^{\top} \mathrm{V}_{t-1} \mathrm{X}\right)^{-1} \mathrm{X}^{T}\left(\mathrm{y}-\mathrm{p}_{t-1}\right)
$$

where

## R Programming

- Newton-Raphson is a general method for solving nonlinear equations iteratively. - Here:
(1) Choose initial estimates of the regression coefficients, such as $\mathrm{b}_{0}=0$.
(2) At each iteration $t$, update the coefficients:

$$
\mathrm{b}_{t}=\mathrm{b}_{t-1}+\left(\mathrm{X}^{T} \mathrm{~V}_{t-1} \mathrm{X}\right)^{-1} \mathrm{X}^{T}\left(\mathrm{y}-\mathrm{p}_{t-1}\right)
$$

where

- $\mathrm{p}_{t-1}=\left\{1 /\left[1+\exp \left(-\mathrm{x}_{i}^{T} \mathrm{~b}_{t-1}\right)\right]\right\}$ is the vector of fitted response probabilities from the previous iteration.


## R Programming

- Newton-Raphson is a general method for solving nonlinear equations iteratively.
- Here:
(1) Choose initial estimates of the regression coefficients, such as $\mathrm{b}_{0}=0$.
(2) At each iteration $t$, update the coefficients:

$$
\mathrm{b}_{t}=\mathrm{b}_{t-1}+\left(\mathrm{X}^{T} \mathrm{~V}_{t-1} \mathrm{X}\right)^{-1} \mathrm{X}^{T}\left(\mathrm{y}-\mathrm{p}_{t-1}\right)
$$

where

- $\mathrm{p}_{t-1}=\left\{1 /\left[1+\exp \left(-\mathrm{x}_{i}^{T} \mathrm{~b}_{t-1}\right)\right]\right\}$ is the vector of fitted response probabilities from the previous iteration.
- $\mathrm{V}_{t-1}=\operatorname{diag}\left\{p_{i, t-1}\left(1-p_{i, t-1}\right)\right\}$.


## R Programming

- Newton-Raphson is a general method for solving nonlinear equations iteratively.
- Here:
(1) Choose initial estimates of the regression coefficients, such as $b_{0}=0$.
(2) At each iteration $t$, update the coefficients:

$$
b_{t}=b_{t-1}+\left(X^{\top} V_{t-1} X\right)^{-1} X^{T}\left(y-p_{t-1}\right)
$$

where

- $\mathrm{p}_{t-1}=\left\{1 /\left[1+\exp \left(-\mathrm{x}_{i}^{T} \mathrm{~b}_{t-1}\right)\right]\right\}$ is the vector of fitted response probabilities from the previous iteration.
- $\mathrm{V}_{t-1}=\operatorname{diag}\left\{p_{i, t-1}\left(1-p_{i, t-1}\right)\right\}$.
(3) Step 2 is repeated until $b_{t}$ is close enough to $b_{t-1}$, at which point the MLE $\widehat{\boldsymbol{\beta}} \approx b_{t}$. The estimated asymptotic covariance matrix of the coefficients is given by $\widehat{V}(\widehat{\boldsymbol{\beta}}) \approx\left(\mathrm{X}^{\top} \mathrm{V}_{t} \mathrm{X}\right)^{-1}$.


## R Programming

Object-Oriented Programming in R: The S3 Object System

- Three standard object-oriented programming systems in R: S3, S4, reference classes. Of these, the S 3 object system is the one most commonly used in applications.


## R Programming

Object-Oriented Programming in R: The S3 Object System

- Three standard object-oriented programming systems in R: S3, S4, reference classes. Of these, the S3 object system is the one most commonly used in applications.
- How the S3 object system works:


## R Programming

- Three standard object-oriented programming systems in R: S3, S4, reference classes. Of these, the S3 object system is the one most commonly used in applications.
- How the S3 object system works:
- Method dispatch of the generic function generic() for the object named object, which is of of class "class" (where $\Rightarrow$ means "the interpreter looks for and dispatches"): generic (object) $\Rightarrow$ generic.class(object) $\Rightarrow$ generic.default (object)


## R Programming

- Three standard object-oriented programming systems in R: S3, S4, reference classes. Of these, the S3 object system is the one most commonly used in applications.
- How the S3 object system works:
- Method dispatch of the generic function generic() for the object named object, which is of of class "class" (where $\Rightarrow$ means "the interpreter looks for and dispatches"): generic (object) $\Rightarrow$ generic.class(object) $\Rightarrow$ generic.default (object)
- For example, summarizing an object mod of class "lm": summary (mod) $\Rightarrow$ summary.lm(mod)


## R Programming

- Three standard object-oriented programming systems in R: S3, S4, reference classes. Of these, the S3 object system is the one most commonly used in applications.
- How the S3 object system works:
- Method dispatch of the generic function generic() for the object named object, which is of of class "class" (where $\Rightarrow$ means "the interpreter looks for and dispatches"):
generic (object) $\Rightarrow$ generic.class(object) $\Rightarrow$ generic.default (object)
- For example, summarizing an object mod of class "lm":
summary (mod) $\Rightarrow$ summary. 1 m (mod)
- Objects can have more than one class, in which case the first applicable method is used.


## R Programming

- Three standard object-oriented programming systems in R: S3, S4, reference classes. Of these, the S3 object system is the one most commonly used in applications.
- How the S3 object system works:
- Method dispatch of the generic function generic() for the object named object, which is of of class "class" (where $\Rightarrow$ means "the interpreter looks for and dispatches"):
generic (object) $\Rightarrow$ generic.class(object) $\Rightarrow$ generic.default (object)
- For example, summarizing an object mod of class "lm":
summary (mod) $\Rightarrow$ summary.lm(mod)
- Objects can have more than one class, in which case the first applicable method is used.
- For example, objects produced by glm() are of class c("glm", "lm") and therefore can inherit methods from class "lm".


## R Programming

- Three standard object-oriented programming systems in R: S3, S4, reference classes. Of these, the S3 object system is the one most commonly used in applications.
- How the S3 object system works:
- Method dispatch of the generic function generic() for the object named object, which is of of class "class" (where $\Rightarrow$ means "the interpreter looks for and dispatches"):
generic (object) $\Rightarrow$ generic.class(object) $\Rightarrow$ generic.default (object)
- For example, summarizing an object mod of class "lm": summary (mod) $\Rightarrow$ summary.lm(mod)
- Objects can have more than one class, in which case the first applicable method is used.
- For example, objects produced by glm() are of class c("glm", "lm") and therefore can inherit methods from class "lm".
- Methods are searched from left to right, so if mod is produced by a call to glm(), and if generic (mod) is called, then methods are invoked in the order
generic(mod) $\Rightarrow$ generic.glm(mod) $\Rightarrow$ generic.lm(mod) $\Rightarrow$ generic.default (mod)
and will fail if none of these three methods are available.


## R Programming

- Generic functions take the form:

```
generic <- function(object, other, named, arguments, ...){
    UseMethod("generic")
}
```

where the ellipses (...) "soak up" additional arguments not named in the generic function that may be passed to specific methods when generic() is called.

## R Programming

- Generic functions take the form:

```
generic <- function(object, other, named, arguments, ...){
    UseMethod("generic")
}
```

where the ellipses (...) "soak up" additional arguments not named in the generic
function that may be passed to specific methods when generic() is called.

- For example, the R summary () function is defined as

```
summary <- function(object, ...){
    UseMethod("summary")
}
and summary.lm() is
summary.lm <- function (object, correlation=FALSE, symbolic.cor=FALSE, ...){
    etc.
}
```

