Model selection

& some other Miscelaneous

Mabel Carabali

EBOH, McGill University

updated: 2024-11-25

What do you think?

Racial discrimination in mortgage lending has declined sharply in America

Endogeneity tweet

Expected Competencies

- Understand the difference between predictive and etiological epidemiology principles.
- Knows statistical assumptions for different regression models.
- Knows basic model fit or "goodness-of-fit" statistics.

Objectives

- To revise the principles and differences between predictive and etiological research.
- To revise the use of main "goodness-of-fit" statistics.
- To revise the overall framework for model parameterization and specification for prediction models.
- To revise variable parameterization and selection.

What is model selection?

- There are two important considerations when building a regression model:
  - The kind of outcome data you want to model (continuous, count, binary, etc.)
  - **The variables and interactions you include in your model**

- **Model selection** is a process that attempts to find the best model for a given purpose.

- The two main purposes of regression models are:
  - **Causal Inference:** To estimate the effect of one or more variables while adjusting for the possible confounding effects of other variables
  - **Prediction:** To predict outcomes for a set of similar individuals

Causal Inference and Prediction: Different Goals!

- **Causal Inference**: estimating the effect of a variable on an outcome
  - Usually adjusted for confounding
  - Want a model with good confounder selection, determined via a DAG and substantive knowledge

- **Prediction**: predicting a future outcome using a set of covariates
  - We want predictions that are close to the actual value, but avoid overfitting
  - May prioritize measurable predictors over strong predictors
- **The two goals do NOT require the same model selection approach!**

Prediction??

- **Prediction** aims to anticipate some future outcome using a set of covariates
- Prediction models are typically built using data from an existing group of individuals, with the goal of being able to predict the outcome for future individuals.
- Here, we are concerned with getting the best model that gives "optimal" predictions for future subjects.
- A good predictive model doesn't necessarily tell you anything useful about how to intervene to change the outcome!
    - Age, sex, and prior hospitalization may be good predictors of heart failure, but we wouldn't stop admitting people to the hospital to prevent heart failure

Prediction: Examples

Not only weather...

- Credit scores
- Netflix recommendations
- Cardiovascular risk scores *
- Kidney function estimates *
- Predicting mortality

* Some, including the misstep "race correction", different form "race-adjustment".

Prediction Examples

Framingham cardiovascular disease (10-year risk) calculator

FRS Calculator

Age	42	years
Sex	Female	
HDL-C	1.52	mmol/L
Total-C	4.33	mmol/L
Systolic BP	110	mmHg
	No — BP Treated	
	No — Smoker	
	No — Diabetes	
	Yes — Fam Hx. of premature CVD	

Calculate

10-Year CVD Risk **3.0 % (Low-Risk)**

Heart Age **31**

Or AHA's CV Risk Calculator

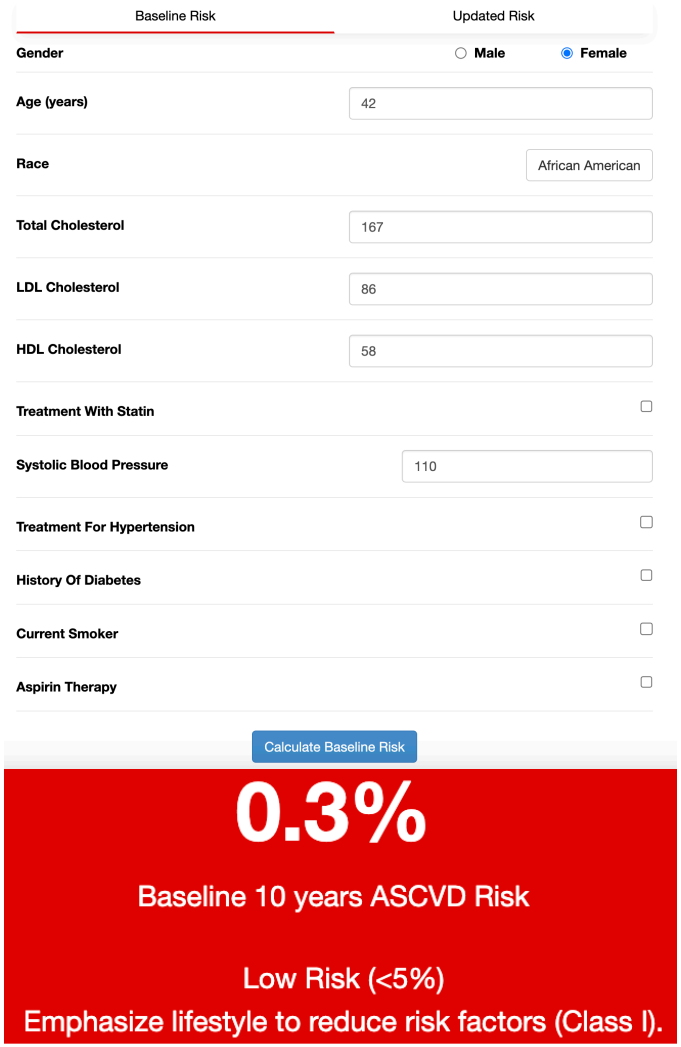

Causal Inference vs. Prediction

Important: We can't interpret coefficients from prediction models as causal parameters

- Not appropriately controlled for confounding (because we aren't worried about this)
- This is often a pitfall that people fall into!

American Journal of Epidemiology

Vol. 177, No. 4
DOI: 10.1093/aje/kws412
Advance Access publication:
January 30, 2013

Commentary

The Table 2 Fallacy: Presenting and Interpreting Confounder and Modifier Coefficients

Daniel Westreich* and Sander Greenland

* Correspondence to Dr. Daniel Westreich, Department of Obstetrics and Gynecology, Duke Global Health Institute, Duke University, DUMC 3967, Durham, NC 27710 (e-mail: daniel.westreich@duke.edu).

OK, but how do I build a model?

What is required to build a "Good" model?

Criteria:

- Is our research question causal or predictive?
  - This will determine the "philosophy" we use to build our model
- Sample size (Power and Precision)
- Choice of statistical model/distribution/link (Poisson, logistic, linear, etc.)
- Consideration of different sources of bias (for causal inference).
- What data do we have available to us?

Overall framework for model building / specification:

1) Variable specification (including assessment of number and parameterization)

2) Interaction assessment (including assessment of heterogeneity)*

3) Confounding assessment (including consideration of precision)*

*Steps 1 and 2 can occur iteratively and depending on the research question investigated differently.

When do we say our model is "good"? - Variable specification

Especially when we have many predictor variables (with many possible interactions), it can be difficult to find a good model.

- Which main effects do we include?
- Which interactions do we include?
- With 6 variables, there are 64 potential models with just main effects!

When do we say our model is "good"? - Variable specification

- An active research problem in statistics
- Model selection procedures try to simplify this task.

Evaluating Model Fit

To implement a model selection procedure, we first need a criterion to compare models.

The goal is to select the model with the optimal value of the criterion.

To that end, we assess the: **Goodness-of-Fit Statistics and use some statistical model metrics**

Evaluating Model Fit: Recall

R-squared R^2

- R^2 represents the proportion of variance in the outcome explained by all the predictors in the model.
- **Criterion:** Choose the model with the largest R^2
- **Problem:** R^2 always increases with model size
- If only use R^2 means simply choosing the largest model: not very useful.
- Helpful when comparing models with the same number of parameters.

Adj. R-squared R^2

- Adjusted R^2 represents the proportion of variance in the outcome explained by all the predictors in the model, **penalized for the number of variables included in the model**.
- **Criterion:** Choose the model with the largest adjusted R^2.
- Here, the largest model is not necessarily the best model.

AIC = Akaike Information Criterion

$$AIC = -2ln(Likelihood) + 2p$$

- **The best model is the one with the smallest AIC**.
- The AIC is formed by two terms:
  - The likelihood: measure of fit
  - The penalty term: $2p$, accounts for adding more terms to the model.

The first term **always** decreases as more terms are added to the model, so $2p$ is needed for "balance".

- AIC can be used whenever we have a likelihood, so this generalizes to many statistical models.

BIC = Bayesian Information Criterion

$$BIC = -2ln(Likelihood) + pln(n)$$

- **The best model is the one with the smallest BIC**.
- AIC and BIC are very similar - only the last term changes
- BIC will always choose a model as small or smaller than the AIC (if using the same search strategy).

Selection Strategies - Prediction

- Now that we know how to evaluate model fit, we need to figure out how to find the model with the **best** fit.

- **Best subset:** Search all possible models and take the one with the highest R^2, or lowest MSE/AIC/BIC, etc.
    - Such searches are typically only feasible when you have less than 30 potential predictor variables.

- **Stepwise (forward, backward, or both) searches:** Useful when the number of potential predictor variables is large.

Illustration with the COVID-19 in Kenya data

```
library(epib.704.data)
data("covidkenya")
glimpse(covidkenya)
```

```
## Rows: 355
## Columns: 15
## $ adm_sex            <chr> "Male", "Male", "Male", "Male", "Male", "Male", "Fe…
## $ sex                <int> 1, 1, 1, 1, 1, 1, 2, 1, 1, 2, 1, 1, 2, 1, 2, 2, 1, …
## $ adm_agemons        <int> 20, 58, 5, 44, 39, 56, 15, 105, 6, 44, 77, 44, 7, 3…
## $ age_cat_new        <chr> "12-23 months", "2-5 years", "<6 months", "2-5 year…
## $ covid_status       <int> 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, …
## $ height             <dbl> 72.60, 108.70, 67.40, 107.20, 112.15, NA, 78.50, 12…
## $ weight             <dbl> 7.16, 17.40, 8.20, 15.30, 16.07, 14.40, 9.30, 24.30…
## $ hfaz               <dbl> -4.35, -0.06, 0.54, 1.43, 3.56, NA, 0.01, -0.62, -3…
## $ stunting_cat       <chr> "Severe stunting", "No stunting", "No stunting", "N…
## $ muac_average       <dbl> 11.15, 15.50, 14.90, 14.30, 13.95, 14.55, 15.40, 16…
## $ nutritional_status <chr> "SAM", "No Malnutrition", "No Malnutrition", "No Ma…
## $ muac_cat           <int> 1, 3, 3, 3, 3, 3, 3, NA, 2, 3, NA, 3, 3, 3, 3, 3, N…
## $ hh_covid_pos       <chr> "No", "No", "No", "No", "No", "No", "No", "No", "No…
## $ birth_order        <chr> "First", "Second", "Third and above", "Second", "Se…
## $ breast_feeding1    <chr> "nobreastfeed", "nobreastfeed", "nobreastfeed", "no…
```

Regression subset selection: Illustration with the COVID-19 in Kenya data

```
library(leaps) #Regression subset selection package

regfit_full <- regsubsets(covid_status ~ .,
                      data = covidkenya,
                      method = "exhaustive")
```

```
## Reordering variables and trying again:
```

```
reg_summary <- summary(regfit_full)
summary(regfit_full)
```

```
## Subset selection object
## Call: regsubsets.formula(covid_status ~ ., data = covidkenya, method = "exhaustive")
## 19 Variables  (and intercept)
##                                  Forced in Forced out
## adm_sexMale                          FALSE      FALSE
## adm_agemons                          FALSE      FALSE
## age_cat_new12-23 months              FALSE      FALSE
## age_cat_new2-5 years                 FALSE      FALSE
## age_cat_new6-11 months               FALSE      FALSE
## height                               FALSE      FALSE
## weight                               FALSE      FALSE
## hfaz                                 FALSE      FALSE
## stunting_catNo stunting              FALSE      FALSE
## stunting_catSevere stunting          FALSE      FALSE
## muac_average                         FALSE      FALSE
## nutritional_statusNo Malnutrition    FALSE      FALSE
## nutritional_statusSAM                FALSE      FALSE
## muac_cat                             FALSE      FALSE
```

Illustration with the COVID-19 in Kenya data

```
par(mfrow = c(1,2))

plot(reg_summary$adjr2, xlab = "Number of Variables", ylab = "Adjusted RSq", type = "l")
adj_r2_max = which.max(reg_summary$adjr2)
points(adj_r2_max, reg_summary$adjr2[adj_r2_max], col ="red", cex = 2, pch = 20)

plot(reg_summary$bic, xlab = "Number of Variables", ylab = "BIC", type = "l")
bic_min = which.min(reg_summary$bic)
points(bic_min, reg_summary$bic[bic_min], col = "red", cex = 2, pch = 20)
```

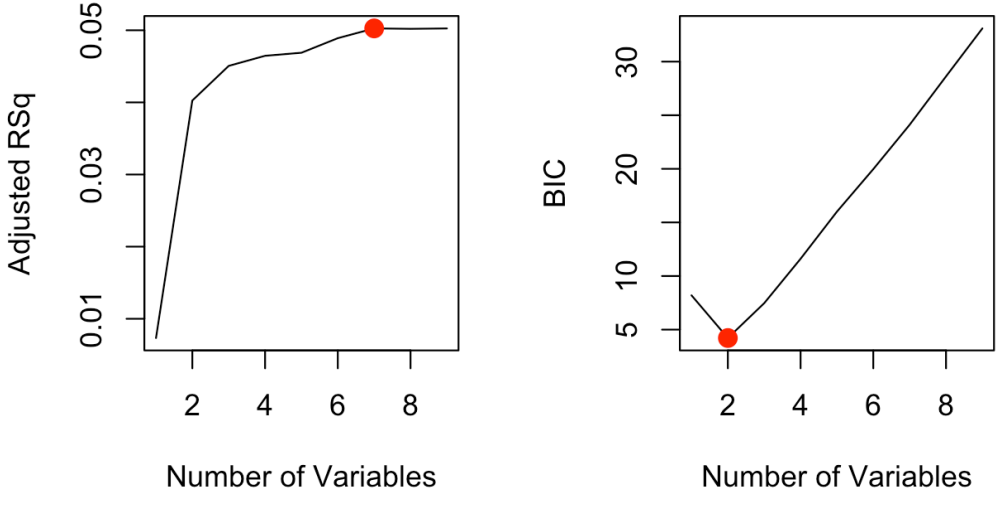

```
plot(regfit_full, scale = "adjr2",
     main= "'adjr2' for COVID-19 in Kenya data")
```

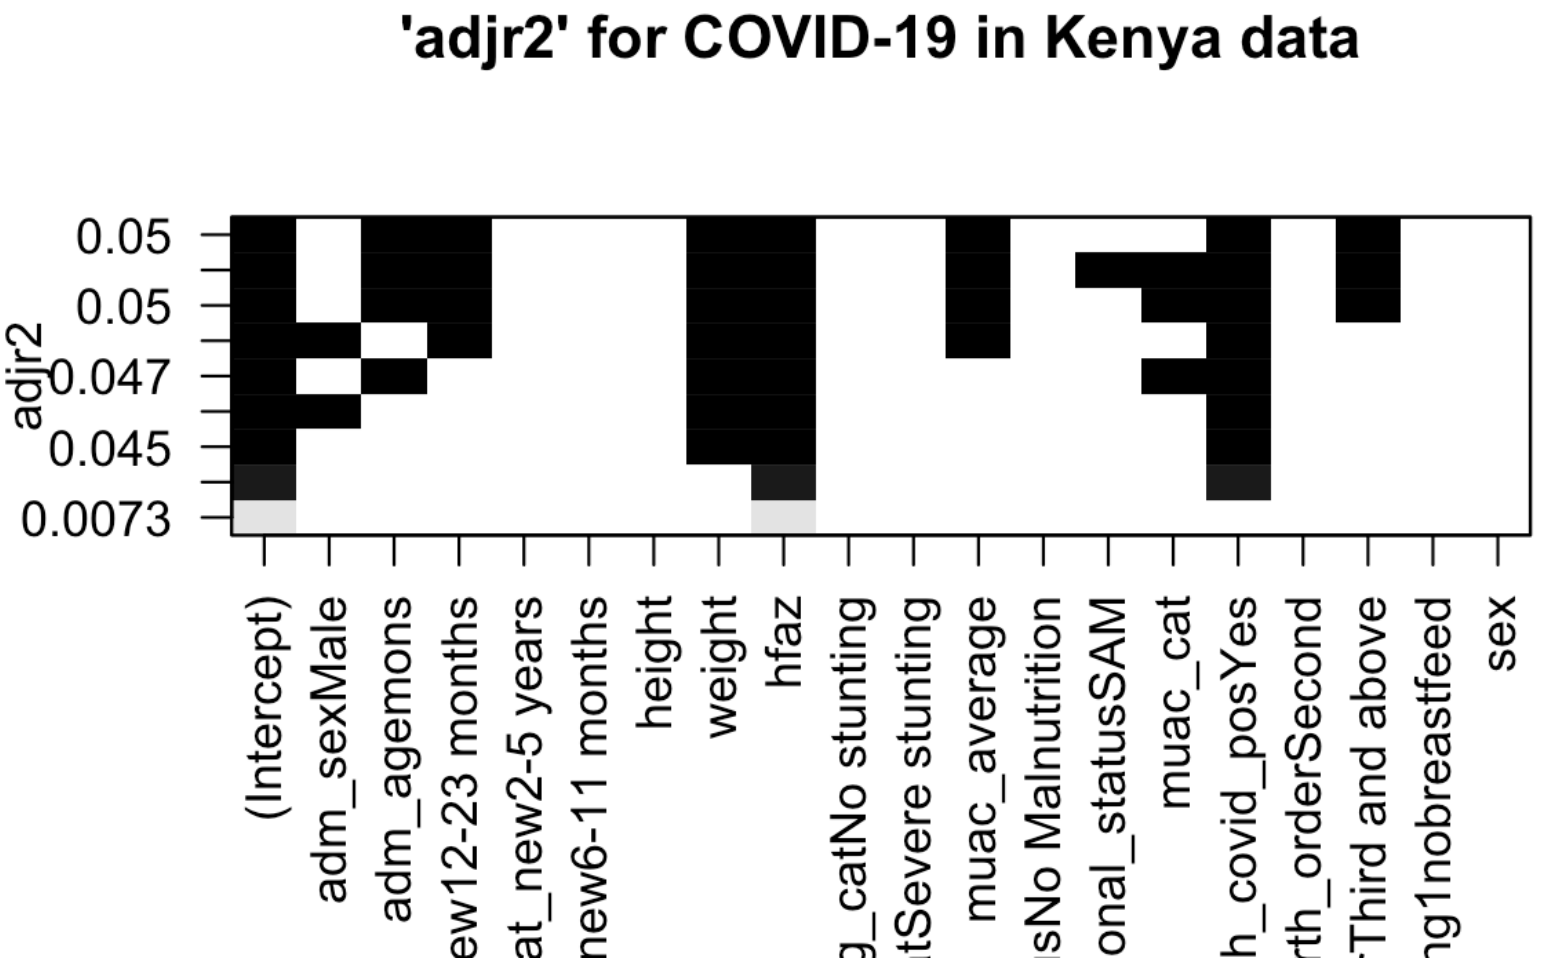

```
plot(regfit_full, scale = "bic",
     main= "'BIC' for COVID-19 in Kenya data")
```

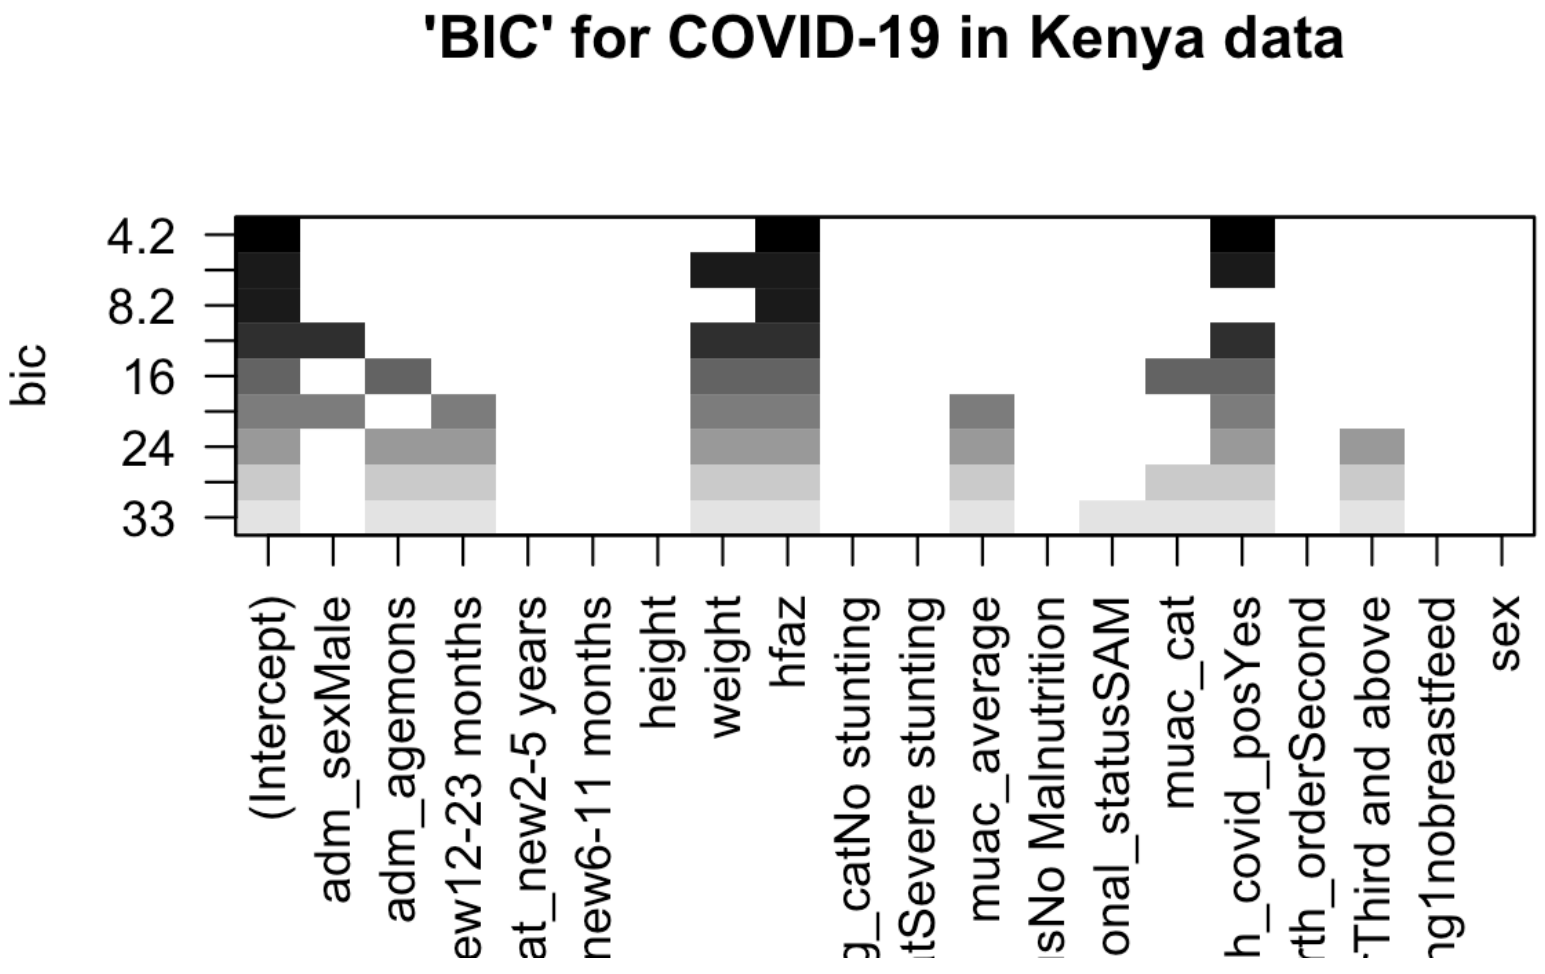

Illustration of selection with the COVID-19 in Kenya data

Subset selection object $\rightarrow$ automatically selects the "best" sets:
regsubsets.formula(covid_status ~ ., data = covidkenya, method = "exhaustive")

- 19 Variables/ parameter (and intercept) with 1 subsets of each size up to 9 variables

- Based on this, we would probably choose to include a model with 2-8 variables:

  - Likely age, weight, hfaz (height for age), MUAC (continuous), hh_covid, depending on the choice of metrics...

Do you trust this approach? Any concerns??

Even the *"smartest"* software, program or technology will need, until now, a bit of guidance...

Selecting a pool of variables to choose from, may help

```
dat<- covidkenya %>% select(adm_sex, age_cat_new, stunting_cat,
                            nutritional_status, hh_covid_pos,
                            birth_order, breast_feeding1,
                            covid_status, hfaz, muac_average)
```

Then re-formulating the selection

```
regfit_full1 <- regsubsets(covid_status ~ .,
                      data = dat,
                      method = "exhaustive")
```

```
reg_summary1 <- summary(regfit_full1)
```

Illustration of selection with the COVID-19 in Kenya data

From the selected covariates, there are 9 variables and 15 parameters to be estimated.

- The package proposes 8 models with 8 combinations of the parameters/covariates.

```
## Subset selection object
## Call: regsubsets.formula(covid_status ~ ., data = dat, method = "exhaustive")
## 15 Variables  (and intercept)
##                                  Forced in Forced out
## adm_sexMale                          FALSE      FALSE
## age_cat_new>=5 years                 FALSE      FALSE
## age_cat_new12-23 months              FALSE      FALSE
## age_cat_new2-5 years                 FALSE      FALSE
## age_cat_new6-11 months               FALSE      FALSE
## stunting_catNo stunting              FALSE      FALSE
## stunting_catSevere stunting          FALSE      FALSE
## nutritional_statusNo Malnutrition    FALSE      FALSE
## nutritional_statusSAM                FALSE      FALSE
## hh_covid_posYes                      FALSE      FALSE
## birth_orderSecond                    FALSE      FALSE
## birth_orderThird and above           FALSE      FALSE
## breast_feeding1nobreastfeed          FALSE      FALSE
## hfaz                                 FALSE      FALSE
## muac_average                         FALSE      FALSE
## 1 subsets of each size up to 8
## Selection Algorithm: exhaustive
##          adm_sexMale age_cat_new>=5 years age_cat_new12-23 months
```

Illustration of selection with the COVID-19 in Kenya data

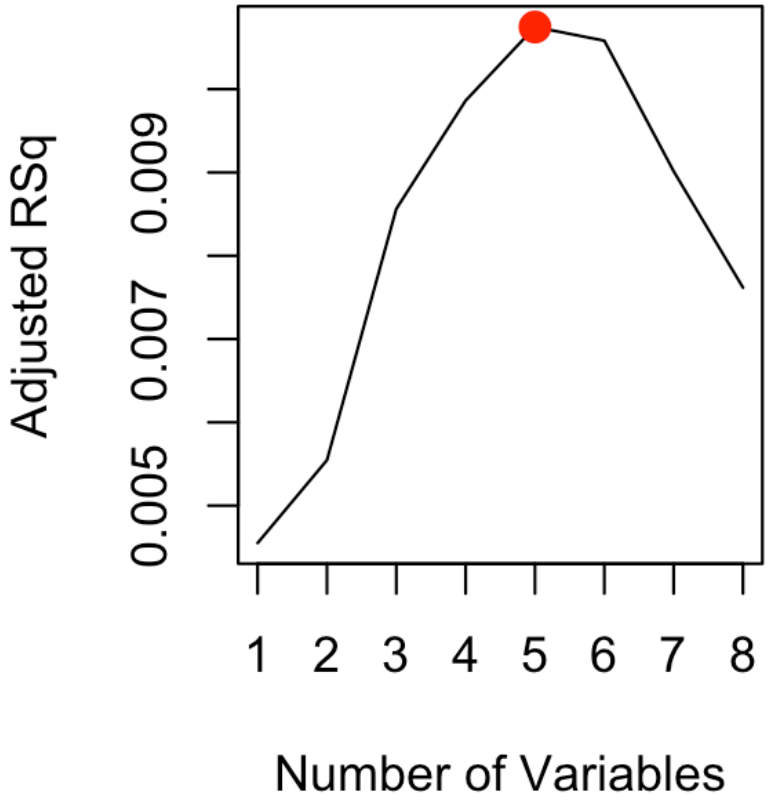

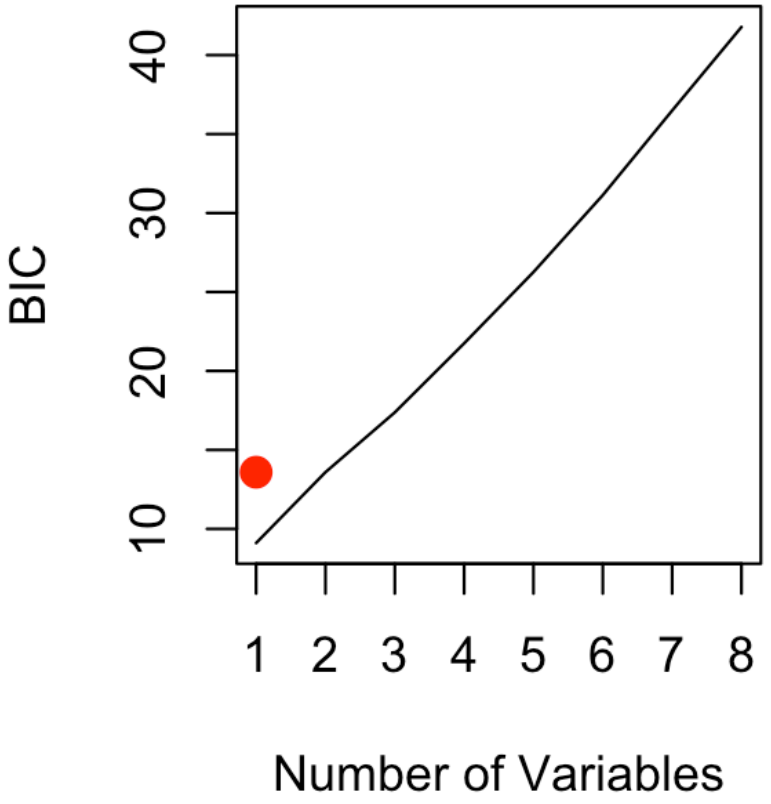

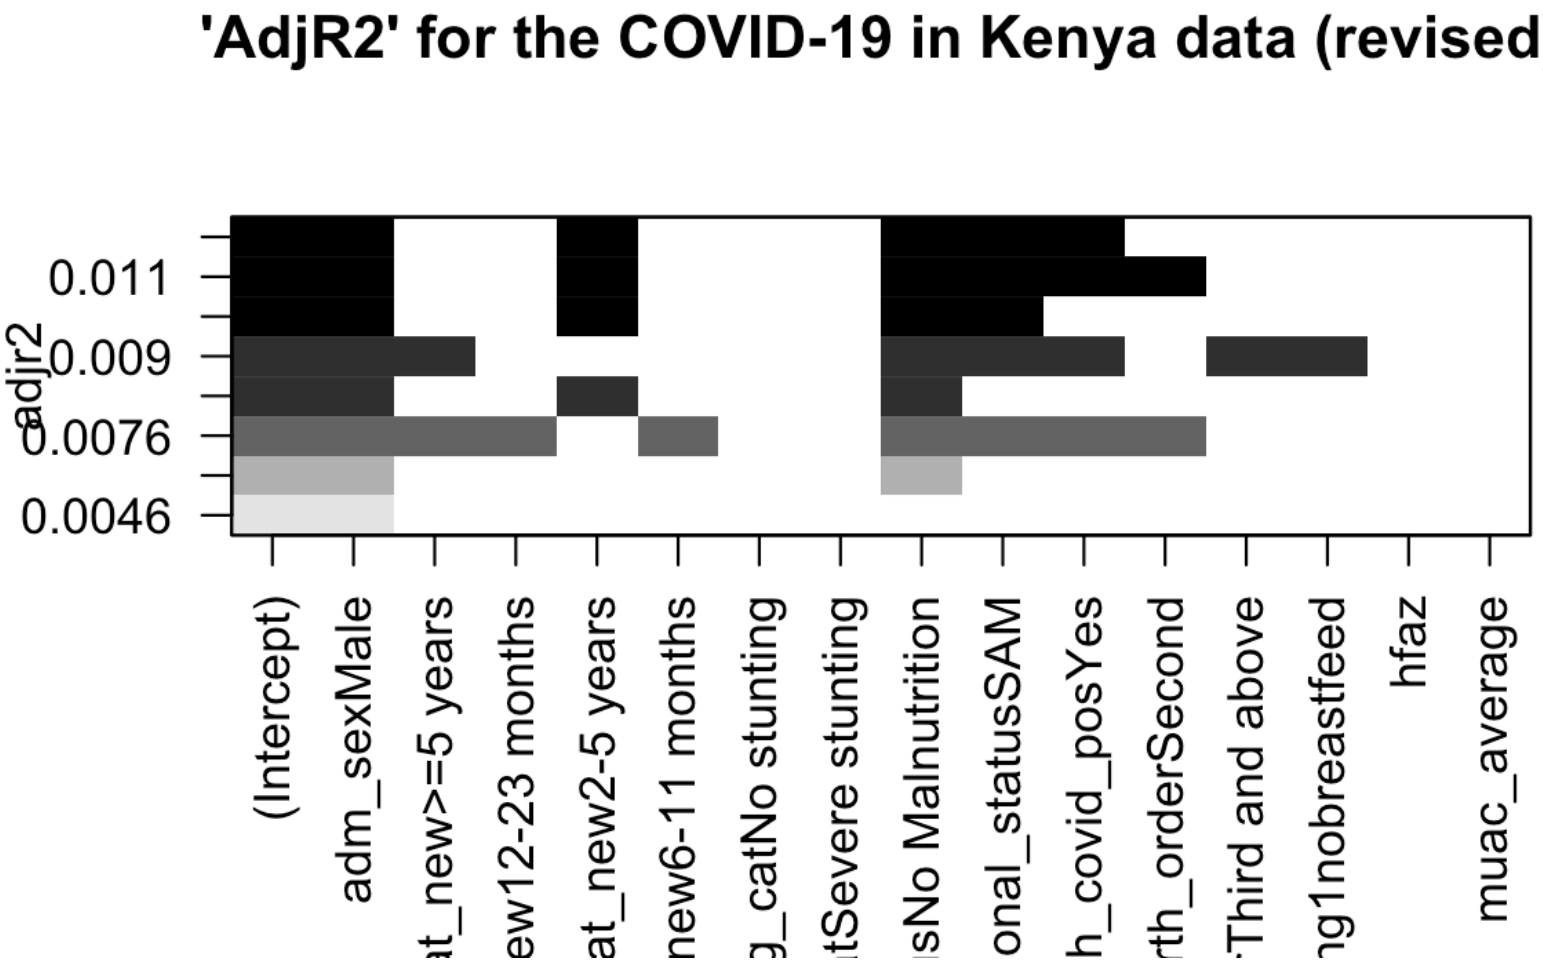
'AdjR2' for the COVID-19 in Kenya data (revised)
adjr2
0.011
0.009
0.0076
0.0046
(Intercept)
adm_sexMale
at_new>=5 years
ew12-23 months
at_new2-5 years
new6-11 months
g_catNo stunting
tSevere stunting
sNo Malnutrition
onal_statusSAM
h_covid_posYes
rth_orderSecond
Third and above
ng1nobreastfeed
hfaz
muac_average

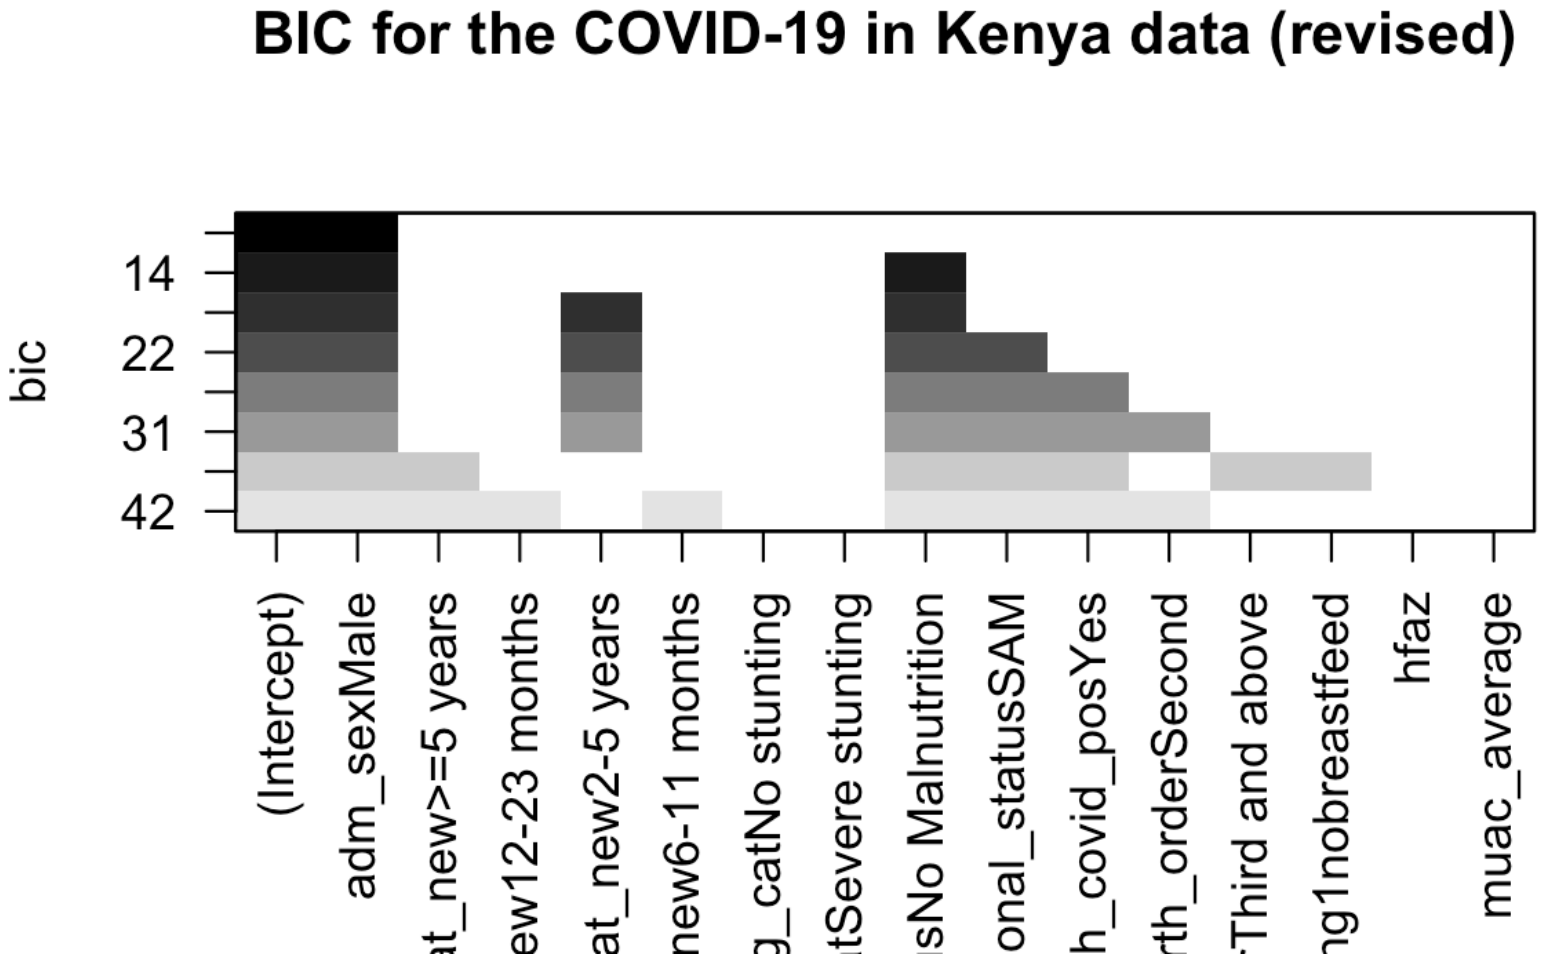
BIC for the COVID-19 in Kenya data (revised)
bic
14
22
31
42
(Intercept)
adm_sexMale
ıt_new>=5 years
ew12-23 months
at_new2-5 years
new6-11 months
g_catNo stunting
tSevere stunting
sNo Malnutrition
onal_statusSAM
h_covid_posYes
rth_orderSecond
Third and above
ng1nobreastfeed
hfaz
muac_average

Other Subset Selection

```
#Backwards selection
regsubsets(outcome ~ .,
           data = covidkenya,
           method = "backward")

#Forwards selection
regsubsets(outcome ~ .,
           data = covidkenya,
           method = "forward")

#Stepwise selection
regsubsets(outcome ~ .,
           data = covidkenya,
           method = "seqrep")
```

Still any decision should be made with caution!

Selection Strategies: Caution!

- Different model fit metrics (e.g., R^2, AIC, Mallow's Cp) can show different models as the "best".

- Can be a recipe for a type I error (chance significant findings) recall Ioannidis paper?
  - Particularly when the number of candidate predictors is large relative to the sample size.

- Automated selection procedures make it easy for researchers to ignore good practices for causal inference, like choosing variables based on a DAG.

- High risk of overfitting to your data set
  - Selected variables may strongly discriminate among individuals in your data set, but may have less ability to do so on other data

Overall framework for model specification:

1) Variable specification:

> - What is the universe of variables I would consider?
>   - There is a limited number available, some of the ones we would need, will not be.
>   - We need to think this through and choose.
>   - Leaving variables out is a strong assumption about that not having an effect ($\beta = 0$)

2) Interaction assessment (including assessment of heterogeneity)

3) Confounding assessment (including consideration of precision)

Selection and Specification

> *"The approach is guided by several principles, including the adherence to a hierarchically defined initial (full) model, and a backward elimination strategy.*

Corollaries

- Collinearity, correlation (other hierarchical structures)
- Moderate or correct for your own Degrees of Freedom
- Multiple testing and Bonferroni corrections

Summary

- **Prediction and inference are different types of problems**.
- If the goal is causal inference, we must think about confounding (and other biases!) and model interpretation is important.
- If the goal is prediction, model interpretation can be less important, and the choice of predictor variables is driven more by the data we have available and model evaluation metrics.
- Model selection strategies can be used for both, but much more care must be taken if you choose to use them for an inference question

“All models are wrong, but some are useful”

— George Box (1919-2013)

QUESTIONS?

COMMENTS?

RECOMMENDATIONS?

Resources:

- Greenland et al. 2016. "Outcome modelling strategies in epidemiology: traditional methods and basic alternatives." Int. J. Epidemiol.

- Steyerberg, Ewout W. 2019. Clinical Prediction Models: A Practical Approach to Development, Validation, and Updating. Springer, Cham.

- Harrell, Frank E., Jr. 2015. Regression Modeling Strategies: With Applications to Linear Models, Logistic and Ordinal Regression, and Survival Analysis. Springer Series in Statistics. Cham: Springer International Publishing.

Miscelaneous

Goodness-of-Fit Statistics and model metrics

- **R-squared** R^2 : % of variation in `Y` explained by `"predictors"` variables. The higher R^2, the better the model.
- **Root Mean Squared Error (RMSE)**: the average error performed by the model in predicting the outcome for an observation. RMSE = sqrt(MSE); The $\uparrow$ the RMSE, the better the model.
- **Residual Standard Error (RSE):** the model sigma, a variant of RMSE adjusted # predictors in the model. The $\downarrow$ RSE, the better the model.
- **Mean Absolute Error (MAE)**, measures the prediction error. MAE = mean(abs(observed-predicted)). Less sensitive to outliers compared to RMSE.

Goodness-of-Fit Statistics and model metrics

- **AIC:** (Akaike's Information Criteria) penalizes the inclusion of additional variables to a model.
- **AICc:** is a version of AIC corrected for small sample sizes.
- **BIC:** (or Bayesian information criteria) is a variant of AIC with a stronger penalty.
- **Mallows Cp:** A variant of AIC developed by Colin Mallows.
- **WAIC:** Widely (Watanabe) Application Criterion (Bayesian)
- **PSIS:** Pareto-Smoothed Importance Sampling (Bayesian)
- **DIC:** Deviance Information Criterion (Bayesian)

Evaluating Prediction Models

- The model fit in the **training data** (i.e., the data used to develop a predictive model) is not our primary interest
- What we are primarily interested in is the accuracy of our model predictions **when our model is applied to new data** that was not used as part of the model development process (i.e., **test data**)
- **A good model fit in the training data doesn't necessarily ensure a good ability to predict using other data**.

Paramaterization, Trends, Dose-Response?

Recall this?

Table 1. Summary of Study Characteristics

Characteristic	1">N = 355[1]
covid_status	55 (16%)
adm_agemons	
Median (Q1, Q3)	30 (15, 63)
muac_average	
Median (Q1, Q3)	14.70 (13.40, 15.70)
caregiver_educ1	
None	21 (5.9%)
Primary	216 (61%)
Secondary	87 (25%)
Above secondary	30 (8.5%)
nutrition	
malnutrition	123 (35%)

Paramaterization, Trends, Dose-Response?

Let's consider Education of the caregiver and nutritional status measured by Mid Upper Arm Circumference (MUAC) in cm, as the continuous outcome.

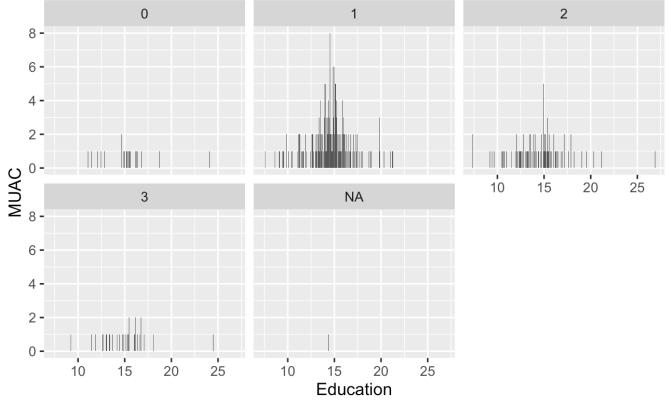


```
## # A tibble: 5 × 2
##   caregiver_educ_n m.muac
##              <dbl>  <dbl>
## 1                0   15.2
## 2                1   14.6
## 3                2   14.7
## 4                3   15.2
## 5               NA   14.4
```

Paramaterization, Trends, Dose-Response?

How can we assess their relationship? **Using a continuous variable assumptions?**

```
muac1<- glm(muac_average ~ caregiver_educ_n, data = L25data)
round(j_summ(muac1, confint = T)$coeftable, 2)
```

```
##                    Est.  2.5% 97.5% t val.    p
## (Intercept)       14.67 14.11 15.22  52.15 0.00
## caregiver_educ_n -0.03 -0.39  0.33  -0.18 0.86
```

Using a categorical version of the variable assumptions?

```
muac2<- glm(muac_average ~ factor(caregiver_educ_n), data = L25data)
round(j_summ(muac2, confint = T)$coeftable, 2)
```

```
##                                Est.  2.5% 97.5% t val.    p
## (Intercept)                   15.23 14.18 16.29  28.23 0.00
## factor(caregiver_educ_n)1 -0.65 -1.76  0.46  -1.15 0.25
## factor(caregiver_educ_n)2 -0.79 -1.97  0.39  -1.31 0.19
## factor(caregiver_educ_n)3 -0.24 -1.62  1.14  -0.35 0.73
```

Paramaterization, Trends, Dose-Response?

Let's consider COVID-19 status (disease yes/no) and Education, can we assess their relationship?

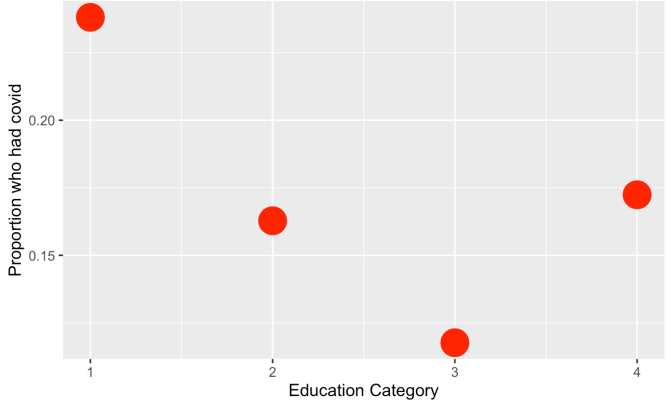

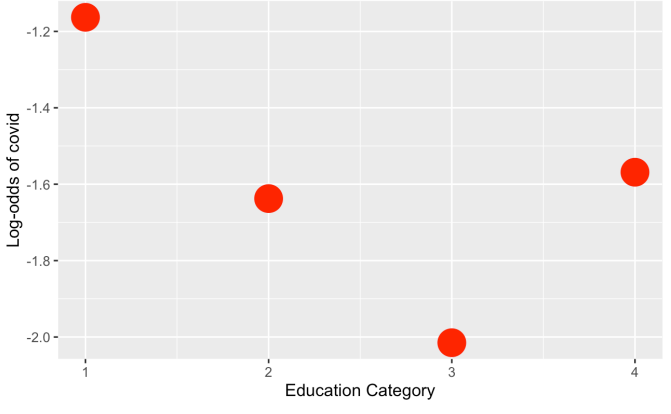

Clinical epidemiology of COVID-19 among hospitalized children in rural western Kenya

Paramaterization, Trends, Dose-Response?

Dichotomous Outcome

```
muac3<- glm(covid_status ~ caregiver_educ_n
            data = L25data, family= binomia
round(j_summ(muac3, confint = T)$coeftable,
```

```
##                    Est.  2.5% 97.5% z val.    p
## (Intercept)      -1.44 -2.05 -0.83  -4.64 0.00
## caregiver_educ_n -0.18 -0.60  0.24  -0.85 0.39
```

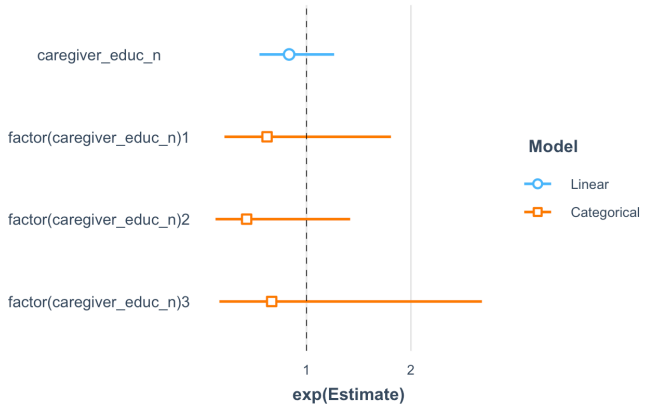


```
muac4<- glm(covid_status ~ factor(caregiver
            data = L25data, family= binomia
round(j_summ(muac4, confint = T)$coeftable,
```

```
##                              Est.  2.5% 97.5% z val.    p
## (Intercept)                -1.16 -2.17 -0.16  -2.27 0.02
## factor(caregiver_educ_n)1 -0.47 -1.54  0.59  -0.87 0.38
## factor(caregiver_educ_n)2 -0.85 -2.05  0.35  -1.39 0.16
## factor(caregiver_educ_n)3 -0.41 -1.80  0.99  -0.57 0.57
```

Paramaterization, Trends, Dose-Response?

What happens when the referent group has a very small sample size?

- OR estimates are all imprecise (since the imprecision for the referent group is propogated)
  - One way are this is to use incremental coding of the dummy variables

```
L25data$edu1<- L25data$edu2<-L25data$edu3<-0

L25data$edu1[L25data$caregiver_educ_n>=1]<- 1
L25data$edu2[L25data$caregiver_educ_n>=2]<- 1
L25data$edu3[L25data$caregiver_educ_n>=3]<- 1

L25data$edu1[is.na(L25data$caregiver_educ_n)==T]<- NA
L25data$edu2[is.na(L25data$caregiver_educ_n)==T]<- NA
L25data$edu3[is.na(L25data$caregiver_educ_n)==T]<- NA
```

Clinical epidemiology of COVID-19 among hospitalized children in rural western Kenya

Paramaterization, Trends, Dose-Response?

Incremental categories*

```
muac5<- glm(covid_status ~ edu1+edu2+edu3,
            data = L25data,
            family= binomial(link = "logit"
round(j_summ(muac5, confint = T)$coeftable,
```

```
##                Est.  2.5% 97.5% z val.    p
## (Intercept) -1.16 -2.17 -0.16  -2.27 0.02
## edu1        -0.47 -1.54  0.59  -0.87 0.38
## edu2        -0.38 -1.13  0.38  -0.98 0.33
## edu3         0.45 -0.72  1.61   0.75 0.45
```

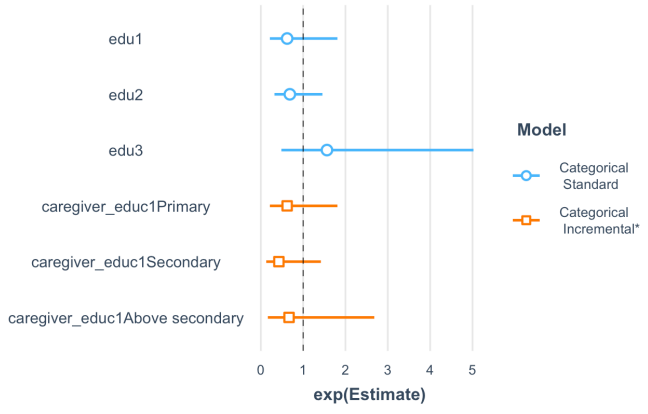

OR what we usually do...

```
muac6<- glm(covid_status ~ caregiver_educ1,
            data = L25data, family= binomia
round(j_summ(muac6, confint = T)$coeftable,
```

```
##                                  Est.  2.5% 97.5% z val.    p
## (Intercept)                     -1.16 -2.17 -0.16  -2.27 0.02
## caregiver_educ1Primary          -0.47 -1.54  0.59  -0.87 0.38
## caregiver_educ1Secondary        -0.85 -2.05  0.35  -1.39 0.16
## caregiver_educ1Above secondary -0.41 -1.80  0.99  -0.57 0.57
```

*Incremental means, that the reference is the immediately previous category/group

Splines examples!

- Y= muac_average (MUAC in cms);
- X= hfaz (Height for age, Z-score)

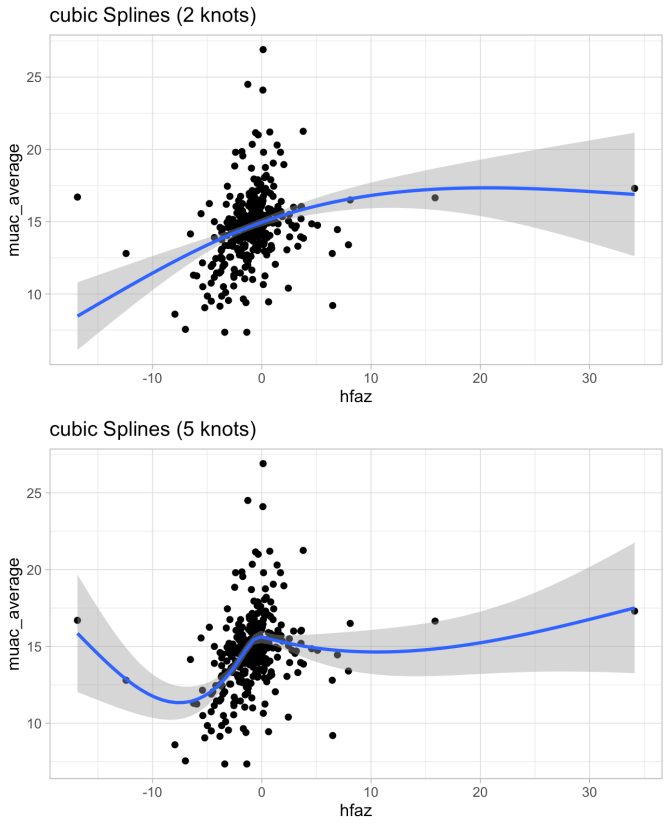


```
require(splines)
lmmod6.11a<- glm(muac_average~ breast_feedi
                 ns(hfaz, df=2), data = c
#round(j_summ(lmmod6.11a, confint = T)$coef
lmmod6.11b<- glm(muac_average~ breast_feedi
                 ns(hfaz, df=5), data = c
round(j_summ(lmmod6.11b, confint = T)$coeft
```

```
##                                Est.  2.5% 97.5% t
## (Intercept)                   12.95  9.55 16.35
## breast_feeding1nobreastfeed -0.08 -0.74  0.59  -
## adm_sexMale                    0.08 -0.34  0.49
## age_cat_new>=5 years           3.99  3.01  4.96
## age_cat_new12-23 months        1.58  0.66  2.50
## age_cat_new2-5 years           2.53  1.59  3.46
## age_cat_new6-11 months         1.24  0.26  2.22
## ns(hfaz, df = 5)1             -0.64 -3.91  2.63  -
## ns(hfaz, df = 5)2              0.03 -3.32  3.39
## ns(hfaz, df = 5)3              3.32  0.24  6.40
## ns(hfaz, df = 5)4             -2.67 -9.94  4.60  -
## ns(hfaz, df = 5)5              7.34  3.22 11.46
```

```
#attr(terms(lmmod6.11b), "predvars")
```

Last Splines examples!

- Y= muac_average (MUAC in cms);
- X= adm_agemons (Age in months)

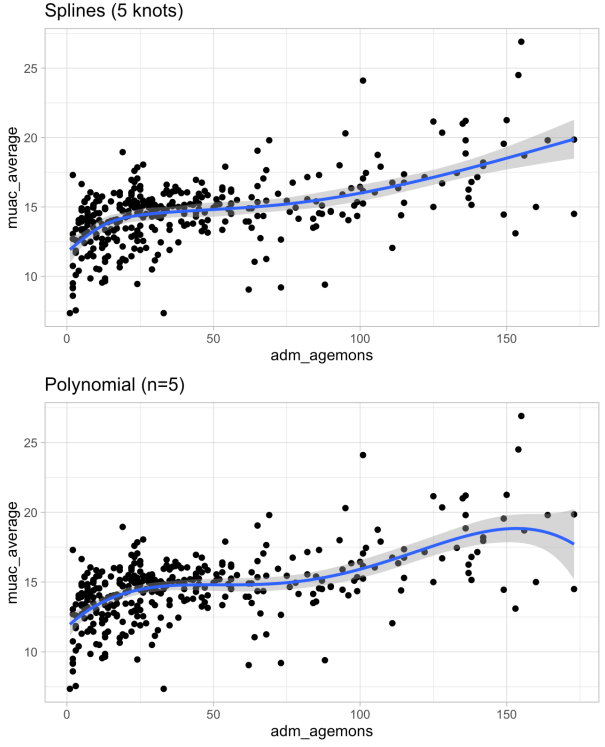


```
lmmod6.12a<- glm(muac_average~ breast_feedi
                        ns(adm_agemons, df=5), d
#round(j_summ(lmmod6.12a, confint = T)$coef
lmmod6.12b<- glm(muac_average~ breast_feedi
                        poly(adm_agemons, 5, raw
round(j_summ(lmmod6.12b, confint = T)$coeft
```

```
##                                    Est.  2.5% 97.5%
## (Intercept)                       11.70 10.75 12.66
## breast_feeding1nobreastfeed       -0.29 -0.97  0.38
## adm_sexMale                        0.13 -0.30  0.56
## poly(adm_agemons, 5, raw = T)1     0.21  0.09  0.34
## poly(adm_agemons, 5, raw = T)2     0.00 -0.01  0.00
## poly(adm_agemons, 5, raw = T)3     0.00  0.00  0.00
## poly(adm_agemons, 5, raw = T)4     0.00  0.00  0.00
## poly(adm_agemons, 5, raw = T)5     0.00  0.00  0.00
```

```
#attr(terms(lmmod6.12b), "predvars")
```